

Networks in economic development

Emily Breza,^{*} Arun Chandrasekhar,^{**} Benjamin Golub,^{***}
and Aneesh Parvathaneni^{****}

Abstract: This chapter surveys the implications of studies in network economics for economic development. We focus on information flow and risk-sharing—two topics where work in theory, empirics, and policy analysis have been especially intensive and complementary. In analysing information, we distinguish models of information diffusion and aggregation, and highlight how different models imply very different guidance regarding the right way to seed information. In discussing risk-sharing, we look at the key frictions that impede efficient informal insurance, and some potential unintended consequences when policymakers intervene to help. Throughout, we stress practical insights that can be used with limited measurement of the details of networks.

Keywords: social networks, learning, information diffusion, risk sharing, econometrics of networks

JEL classification: D83, D85, L14, O20

I. Introduction

A large and growing literature in the social sciences documents an important role of social and economic *networks* in developing economies. Network interactions are ones mediated by interpersonal and local relationships.¹ Examples of network-based processes include learning about a new government scheme (say, an education subsidy) from a neighbour who is using it, borrowing kerosene from a nearby cousin, or participating in a ROSCA (Rotating Savings and Credit Association) with friends in the neighbourhood. Networks are often crucial in facilitating economic activities when formal institutions and markets are missing or unsuited to people's needs. Particular

^{*} Harvard University; e-mail: ebreza@fas.harvard.edu

^{**} Stanford University; e-mail: arungc@stanford.edu

^{***} Harvard University; e-mail: ben.golub@gmail.com

^{****} Yale University; e-mail: aneesha.parvathaneni@yale.edu

We thank the Foundations of Human Behavior Initiative at Harvard and the National Science Foundation (SES-1658940, SES-1629446). We are grateful to Matthew O. Jackson for useful conversations, as well as to Ria Granzier-Nakajima, Yixi Jiang, and Vasu Chaudhary for excellent research assistance.

¹ By *local* we simply mean that each person has a distinct 'neighborhood', and the situation is not well-approximated by some uniform or homogeneous interaction within a large group. For book-length introductions to network analysis in economics, see Jackson (2010), Easley *et al.* (2010), and Goyal (2012). For a survey of field experiments in development, see Breza (2016).

doi:10.1093/oxrep/grz025

© The Author(s) 2019. Published by Oxford University Press.

For permissions please e-mail: journals.permissions@oup.com

attention has been paid, both theoretically and empirically, to the two examples we have mentioned: the *provision of information* (and the closely related topic of technology adoption) and *financial contracting*, especially risk sharing.

In this article, we survey the implications of network economics research for the two above-mentioned areas—where work in theory, empirics, and policy analysis have been especially intensive and complementary. We focus on the lessons a policy-maker might draw from this literature. For example, if a policy-maker is organizing an information campaign, who should be seeded with information to effect the largest aggregate change in behaviour? If a policy-maker is overseeing an expansion of the formal financial system, how can she assess the crowding out of informal networks of financial relationships and mitigate the harms of that?

Networks are intrinsically very rich and intricate, with a huge amount of potential structure: among 24 people, there are 276 possible bilateral connections, and 2^{276} possible undirected networks.² Researchers who focus on networks find this richness interesting. A network theorist might study how each conceivable realization of the structure will affect some outcome (e.g. information flow), under various models of behaviour. An empirical researcher might then examine the resulting theories: measuring networks rather finely, testing which models of behaviour hold at a detailed level, estimating their parameters, and testing the validity of predictions in the aggregate.

A policy-maker—ultimately, someone we hope would be an important ‘client’ of this enterprise—could understandably see this activity and worry that it is fantastically resource-intensive, especially insofar as it involves collecting detailed *network data*. And the policy-maker might wonder whether and how this enterprise can realistically play a role in policy.

Throughout the article, we return several times to a simple but important theme. Though advancing our understanding of networks can, indeed, be very demanding theoretically and in terms of data, that understanding can yield simple and practical policy guidance. In some cases, network studies suggest policies that involve modest data demands—for example, a policy may require a manageable number of sufficient statistics. (Examples of relevant statistics include who is reported as being influential by others, or a segregation measure—e.g. how much correlation there is between people’s caste and that of their social contacts.) The detailed studies are often essential for devising and justifying the policies, but micro-level network measurement is not ultimately a part of using the insights. We discuss several applications of this principle, building up to some novel developments in the statistics of networks required for the ‘sufficient statistics’ approach. In addition to their practical value, these advances are intellectually exciting in their own right. At the same time, there are domains in which we are not far enough along in our understanding of either the detailed workings of networks or in the statistical approaches that are necessary for practical purposes, and we highlight when this is the case, hoping to stimulate further research.

To summarize, the scientific study of network processes and the application of the results for practical policy advice are extremely different enterprises, but also highly complementary to one another. Our aim here is to elucidate aspects of both; to highlight the

² More than the number of atoms in the universe, as [Jackson \(2019\)](#) points out.

complementarity and some of what it delivers; and to emphasize areas where we believe it can and should deliver more.

Network theory and empirics constitute a wide-ranging set of methods in economics. So, beyond the main application areas we've emphasized—information-delivery and risk-sharing—we also touch, more briefly, on other, related areas: migration, targeting recipients of redistributive transfers and related interventions, and take-up of immunization. In all these topics, our focus is on the aspects of network theory and empirics that, in our view, are likeliest to be of special interest to policy analysts.

II. Information delivery

(i) Motivation and outline

Information campaigns. Substantial resources are spent on information campaigns in the developing world.³ These campaigns include agricultural extension, financial extension, and health services, among others. The goal is quite simple: to diffuse useful information widely. This information can be about technologies, government programmes, or relevant conditions, such as prices.

Simple interventions providing better information can have considerable welfare benefits. For example, [Jensen \(2010\)](#) finds that providing accurate information about returns to secondary school increases schooling by 0.20 to 0.35 years. [Duflo and Saez \(2003\)](#) show that, in a US workplace, monetarily incentivizing employee attendance for a session providing benefits information increases retirement savings enrolment for the treated group considerably.⁴

[Dupas \(2011\)](#) shows that providing information about how the risk of HIV depends on the age of one's partner leads to a large reduction in risky sexual behaviour, as proxied for by teenage pregnancy.⁵ In short, information-delivery interventions that amount to acquiring at most a few hours of someone's attention can have literally life-changing economic consequences.⁶ How to deliver information well is therefore an important question.

Leveraging social networks. A crucial potential lever in information campaigns—a way to scale up the benefit per unit of resources invested—is the activation of a social network to disseminate and amplify messages. At least potentially, social channels can be a viable alternative or complement to media campaigns, personal outreach by government

³ To take just a few examples, in 2014–15, India allocated around \$250m for agricultural extension and training ([Gulati et al., 2018](#)).

⁴ By 1.25 percentage points on a base of 1 percentage point after 11 months.

⁵ In particular, teenage pregnancy falls by 28 per cent in the treated group. Self-reported data indicate that teenage girls substituted away from older partners (who are more likely to have HIV) and towards same-age partners.

⁶ For instance, [Dupas \(2011\)](#) reports that about 20 teenage pregnancies were averted at a total cost of just under US\$2,000. The paper also estimates that the cost of the programme per averted HIV infection in the targeted population is about \$700, not counting the avoidance of infections (outside the target population).

extension programme workers, and the like. For instance, [BenYishay and Mobarak \(2018\)](#) find that if peer farmers (those who have similar agricultural set-ups as target farmers) are incentivized to spread information about a new technology, adoption of that technology goes up by 10–14 per cent relative to a control. Outside of a development context, well-identified evidence exists that viral campaigns can considerably scale up take-up of a product or behaviour ([Aral and Walker, 2011](#)). Indeed, in advanced economies, marketers often choose ‘viral’ marketing strategies over traditional ones. Facilitating viral diffusion can be considerably more efficient than attempting to inform everyone directly. But whether networks can be leveraged to spread information is far from obvious. Indeed, in other contexts, information-provision interventions that attempt to exploit the network have little effect. In the context of agricultural extension in Mali, carefully choosing individuals to spread information about a new farming technology resulted in little take-up and performed similarly to informing randomly-selected farmers ([Beaman *et al.*, 2018](#)). Beyond the risk that network-based diffusion may prove ineffective, there is the additional possibility that it may cause harm. For instance, in the case of the 2009 avian flu pandemic, news spread through informal channels, as did false rumours suggesting that the best way to respond was to take one’s chickens indoors to avoid their confiscation by the government ([Banerjee, 2008](#)). During the 2016 demonetization of large banknotes in India, informal discussions of the policy caused many people to believe the false rumour that certain small coins had also ceased to be legal tender ([Banerjee *et al.*, 2018a](#)).

Two questions and a role for research. These cases highlight the obvious question of *whether* social networks will work to diffuse information. There is also a natural next question: if network-mediated processes are expected to be important, how should a policy-maker take this into account? How should the set of initially seeded nodes be chosen? How should the message be designed to optimize its spread? [Banerjee *et al.* \(2019b\)](#) find that the take-up of a microfinance scheme is potentially quite sensitive, in practice, to the choice of who is initially contacted in the village. [Alatas *et al.* \(2018\)](#) similarly find that in an online social network, seed identity plays an essential role in the reach and effectiveness of a message. Thus, the reality that theories must contend with from the start is that *when* network-mediated processes matter, the details of how they are managed may also matter a great deal.

If policy-makers aim to leverage social networks to distribute messages, they need guidance from social science: principles to decide when social networks will work to diffuse information and how the outcomes depend on policy choices. These questions motivate a large theoretical literature on social learning in networks, as well as a growing empirical one, and the insights of this work are our focus in this section.

A toolbox of models

It is worth saying at the outset that there is *no canonical model of learning or behaviour change in networks*. This is, at a minimum, a fact about current practice among both theoretical and empirical researchers: there is no single model that most social scientists working on these topics would consider the default starting point for questions about information campaigns.

Given the current state of the art, the best we can offer is a medium-sized menagerie of models that have proved enlightening in various settings. These models elucidate very different forces. Different features of behaviours and networks are pivotal in each of these. The greatest successes of network-based analysis in development contexts have come not from taking some single model literally but from using a variety of model-based insights and complementary tools to organize our thinking about messy real-life processes. The models here provide the building blocks for such ‘organizing’ work.

For each type of model, we present the key structure, discuss when that model is a natural tool, and present some of the basic analysis. As we develop the tools, we discuss some practical case studies and real-world considerations illustrating both the insights delivered by each model and its limitations.

Before diving into the toolbox, it is useful to take a high-level look at it. All the models we discuss have something in common: they consider a process happening on a network—usually a social network, where links represent communication or observation opportunities. Given a network and an initial state—who is initially informed, what opinions individuals hold, etc.—the process gives us a prediction of dynamics and long-run outcomes. The policy-maker can use such models to reason about various approaches to diffusing information.⁷

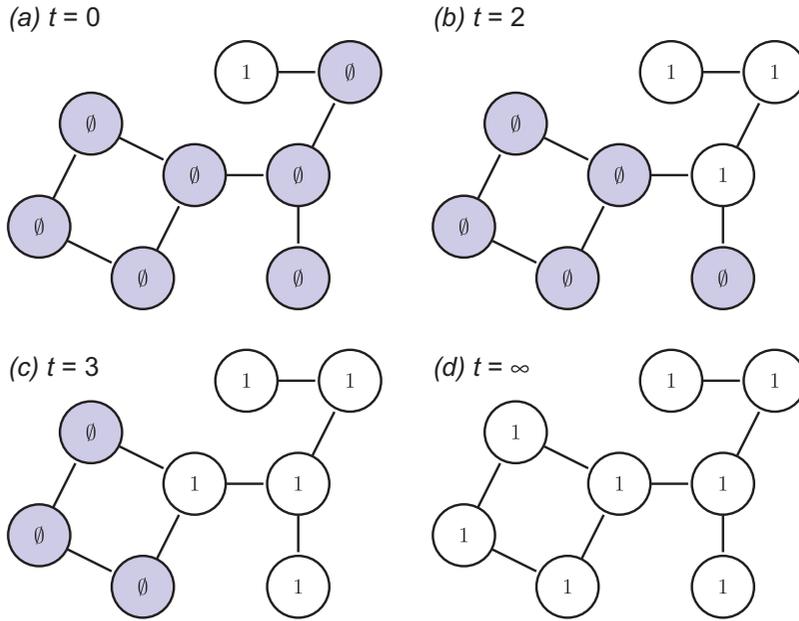
Within this framework, the models are very diverse. We give a non-exhaustive list of dimensions on which they differ.

- **Diffusion vs aggregation.** Does the model focus on diffusion, aggregation, or a combination? *Diffusion* refers to the spread of a state akin to a virus, where social transmission can activate neighbours but not deactivate them. Often the state is some form of awareness—e.g. of a new product or a piece of news. In contrast, *aggregation* refers to the synthesis and reconciliation of different, and possibly competing, information. For example, if it is unclear which mobile phone is better, aggregation refers to the process whereby people learn about this and update their views and behaviours by observing others.
- **Discrete vs continuous.** Does the model focus on discrete states or continuous levels? Discrete models will track, for example, whether each person uses product A or B. Continuous models stand in contrast to these and track a fine level of intensity for each individual: how strongly she believes that product A is better, or how much of her land she devotes to crop A as opposed to B.
- **Simple vs sophisticated.** How sophisticated, rational, or subtle is individual behaviour? Some models posit very simple rules of individual behaviour: e.g. someone becomes aware as soon as at least one neighbour is. Other models involve agents making optimal inferences from all information observed to date, which can entail quite complex, and often counterintuitive, reasoning.

Which choices an analyst makes on each dimension will be driven by substantive facts about the setting as well as pragmatic considerations—which models can be effectively solved and estimated.

⁷ As we have mentioned already, this analysis will not generally require omniscience about the network or other aspects of the model; for more on this, see section IV.

Figure 1: An example of a simple contagion process



(ii) Viral information transmission

The first and perhaps simplest model is that of a classic viral process. For a concrete example, consider a situation in which people learn of an obviously interesting technology or opportunity, see a funny meme, or report a piece of urgent news. In all these cases, awareness of or familiarity with the item spreads from person to person like a virus when people repeat the information or otherwise make it known to others. Crucially, the process of spreading, which we call *diffusion*, is *irreversible*: awareness can be activated by social transmission, but it cannot be contagiously deactivated.

Initially informed individuals in a diffusion process are called *seeds*. Figure 1 illustrates a few steps of a diffusion process starting from a single seed.

In this example, transmission is deterministic: a node that is informed then informs, with certainty, everyone connected to it.⁸ In practice, analysts use a generalization of this model: a node remains ‘contagious’ during a *window* of length τ after being infected; during each of these periods, the node infects each neighbour independently with a certain probability p (typically much less than 1).⁹ Then the node does not infect anyone any more—an unrealistic assumption we return to. This process is called the SIR model: each node is susceptible, infected (and infectious), and then recovered (no longer infectious). To take a practical example, when A learns a piece of news, it does

⁸ A network or graph is defined by specifying its nodes (people, locations, etc.) and links or edges between those nodes, which in this case are undirected. The *neighbours* of a node are all those nodes that have edges to it.

⁹ The example illustrated in the figure is a special case where the infection probability is 1 and the contagious period is one unit of time long.

not necessarily follow that B immediately learns it or begins sharing it. Instead, B learns with some probability. If B does not learn during the ‘window’ when A is talking, then this conduit of transmission goes unused.

Where does this process fit on each of the dimensions that we highlighted in our rough taxonomy of network processes? For concreteness, we continue with the ‘piece of news’ example in illustrating this discussion. First, the model is, by design, concerned exclusively with *diffusion*, as opposed to aggregation. Nodes are not weighing or reconciling opposing views: they are simply switched on irreversibly into a state of awareness. Second, the model has discrete states: it focuses on a zero–one binary outcome. We are not modelling, for example, some measure of how strongly people respond to the news; those are, for some pieces of news, important dimensions, but they are not the ones this model is focused on. Finally, agents’ behaviour here is extremely simple and mechanical: there are no complicated calculations underlying their updating—e.g. Bayesian inference about some underlying state, or anticipation of neighbours’ behaviour.

A key concept: The basic reproductive number

It is helpful to develop a sense of how a viral process works and what we can say about its outcomes. This provides a useful reference point as we discuss the dimensions of the model that are crucial to consider in policy applications. We discuss this somewhat informally, with the goal of developing a heuristic understanding, and make references as appropriate to formalizations.

Let us start with a seemingly very extreme thought experiment. Let us suppose society is shaped like a tree, with each individual having one ‘superior’ above and d ‘subordinates’ below. Each node’s window is $\tau = 1$ period after becoming infected itself. Suppose we start the process at a node in this tree and ask how many individuals below her will, in expectation, be infected in T rounds of transmission. Recalling that each link transmits the state with probability p , a simple calculation shows that the answer is $\sum_{t=0}^T (pd)^t$. Thus we see that the viral process can infect ‘arbitrarily many’ people if $pd > 1$ —which is to say, if the expected number of new infections following a single initial infection exceeds 1. The number pd is called the ‘basic reproductive number’ (BRN) of the process. To summarize, if the BRN exceeds 1, then the process is viral. If it is much less than 1, the process is severely bounded in how large it is likely to grow.

A striking fact about viral processes of the SIR form is that in circumstances much less specific than our extreme thought experiment, this intuition survives and the ‘tree’ approach (with suitable adjustment of the parameters) gives a useful sense of how the viral process behaves. Indeed, a subfield of applied mathematics examines the study of epidemic processes rigorously. To roughly summarize a deep set of results, one can analogously define the ‘average number of new infections a random node produces’ and think of this number, the generalized BRN, as the ‘scaling factor’ of the viral process as it proceeds. (The subtleties are in defining what we mean by ‘average’, ‘random’, etc.) Thus, if the number exceeds 1, the

process is ‘expansive’ and a large number of people will be infected with positive probability starting from any random seed in the network, once T is reasonably large. Strikingly, this works even if the network does not resemble a tree (instead containing many cycles¹⁰) and even if it is very far from regular.¹¹ Thus, policy-makers can potentially take advantage of exponential growth in network diffusion. In this sense the theory fleshes out a rather dramatic instance of the promise of leveraging the network.

Practically important dimensions of a viral process

At this point it is useful to step back and ask how a policy-maker would decide whether and how to apply the basic viral model. We highlight *duration* as an important threshold consideration in this regard, and then turn to measuring and maximizing diffusion.

Duration. A crucial dimension is how long information spreads through the community—that is, how many rounds of communication occur.

One can posit that there is an unlimited number of periods, and the process runs until no changes occur. In the example we gave in [Figure 1](#), where the process operates deterministically, a viral process with unlimited time saturates the connected component in which it starts. More generally, when the BRN is above 1, with positive probability a single node gives rise to a cascade of awareness of unbounded size as time proceeds. Indeed, if we modify the ‘recovered’ assumption of the SIR model to something arguably more realistic, where already knowledgeable nodes who have ‘recovered’ have some probability of mentioning the information and initiating a renewed contagion, then with a BRN above 1, it is easy to see that everyone who physically can become infected eventually will. We return to this, and other observations about virality, in the next section.

A more realistic set-up for many situations is that the analyst is concerned with the progress of the viral process not at an indefinite time in the future but at a particular finite time. This might, for instance, be a deadline to make the right decision about a government policy (e.g. applying to a programme, or exchanging old currency) or a time at which people have to make a decision about which technology to use (e.g. planting time). Alternatively, community interest in a topic may naturally decrease over time, so the effective duration of the process is the time during which it holds people’s attention. Another concern is that if the message is reasonably easy to get wrong in retelling, then the information will only be useful for a small number of ‘steps’ of indirect transmission, rather than arbitrarily many.

¹⁰ By a cycle, we mean a path of links that returns to where it started. Cycles are absent in a tree. For more background on terms from network theory, such as degree, path, cycle, etc., several standard references are helpful: [Jackson \(2010\)](#) is tailored to social scientists, whereas [Diestel \(2005\)](#) and [Bollobás \(1998\)](#) explore more of the mathematical theory.

¹¹ A network is called regular if each node has the same number of neighbours. The ideas mentioned here take us far beyond our scope in this article, but readers can start on the fascinating journey into the general theory of viral processes, percolation theory, etc., with [Jackson \(2010\)](#), [Karlin \(2014\)](#), and [Durrett \(2007\)](#), to name a few entry points.

In such cases of finite duration, a reasonable estimate of the eventual diffusion is obtained by summing the T -step ‘reach’ of each seed, across seeds. That is, if $T = 1$, we can sum up the degrees (numbers of direct connections) of all seeds. If $T = 2$, we look at the sizes of their second-order neighbourhoods: ‘friends and friends of friends’. To do better than this approximation, the analyst would have to account for the overlaps between different seeds’ extended neighbourhoods.

From this we can see that the duration of a viral process matters greatly for who should be initially seeded for maximum diffusion. To see this in an example, if individuals only speak for one round and then stop ($T = 1$), a policy-maker who can inform one agent maximizes diffusion by simply informing the individual with the most friends. In contrast, imagine there are multiple rounds of communication. If individual A has five friends and none of them has any others, but individual B has 2 friends, but each of them has 3 friends, with two rounds of communication ($T = 2$) the node B is a much more valuable seed than A .

Examining the impact of the duration T is a good illustration of the subtleties arising even in a simple model of viral transmission. Duration can be pivotal to what aspects of the network an analyst will focus on. For short duration, highly local aspects of the network will be relevant. When duration is longer, the virality factor (i.e. basic reproductive number) becomes more important.

Virality. Many analyses of viral processes abstract from duration, taking it to be effectively infinite. This is an interesting polar case to examine. Under this assumption, the degree of virality (determined by d and p in our simple example) becomes crucial. We now flesh out this point and describe practical measurements of virality.

Let us first consider a large population, such as users on an online network (WhatsApp or Facebook) and a relatively small number of seeds. Recalling that our estimated eventual reach of eventual diffusion is

$$\sum_{t=0}^T (pd)^t$$

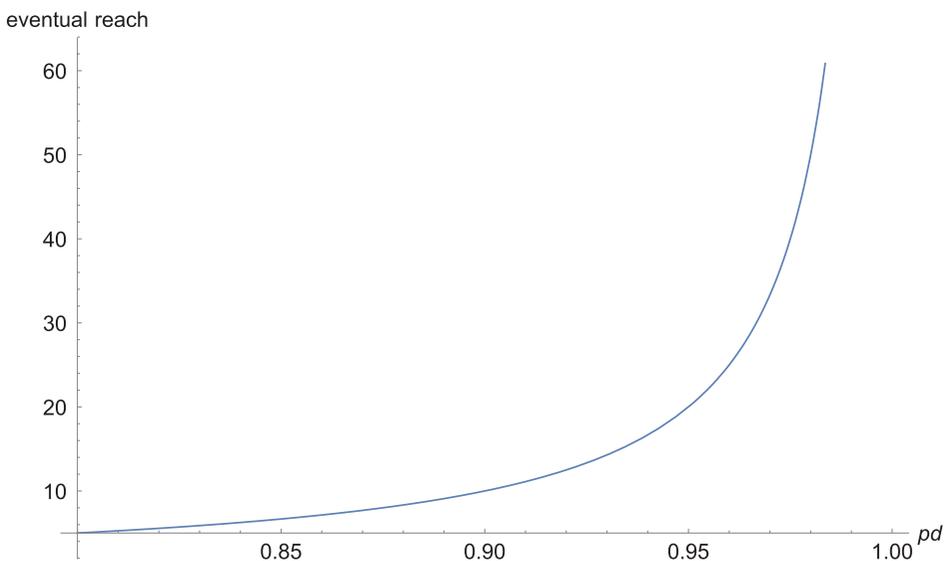
(with suitably calibrated p and d), we can observe that this sum diverges if $pd > 1$. Moreover, it is approximated by $1/(1 - pd)$ for pd less than 1.¹² This calculation is for one seed; if we start with S seeds whose extended neighbourhoods have negligible overlap, then we can expect a reach of

$$S \sum_{t=0}^T (pd)^t$$

or $S/(1 - pd)$, when $pd < 1$, or an unbounded expected reach (with a positive probability of infecting an arbitrarily large network) if $pd > 1$. The reach in the former case is depicted in [Figure 2](#).

¹² [Golub and Jackson \(2010b\)](#) discuss some challenges of estimating the BRN from observed diffusion data.

Figure 2: An illustration of eventual reach as a function of the basic reproductive number, pd , in our example, with $S = 1$



Notes: We have plotted the function for $pd < 1$ only. When $pd > 1$, the reach of diffusion is unbounded in principle, but bounded in practice by the size of the ambient population.

As we can see, for cascades of any substantial reach (which require the BRN to be near 1 or above 1) virality as measured by pd matters a great deal. Around the critical threshold of 1, increasing virality matters much more than most other details, since returns are so convex in the BRN. This helps explain why Internet platforms invest a great deal in fine-tuning the virality of a message. Increasing pd from 0.99 to 0.995, for instance, doubles the expected reach of an item, so this seemingly tiny improvement can be worth a lot. Thus, it can make sense to do extensive testing on how small features of a stimulus or its recipients affect its propagation.

On the other hand, in development contexts—e.g. the work by [BenYishay and Mobarak \(2018\)](#) mentioned earlier or [Banerjee et al. \(2019b\)](#)—the reach we measure in practice does not seem to be consistent with truly exponential growth. Indeed, since almost all important technologies are known to someone in the community or around it, information that is ‘viral’ in nature has often already spread by natural means. The kinds of messages that a policy-maker wants to transmit are likely to be less contagious. In short, both on theoretical grounds and based on some measurements, viral contagion models with large T do not seem like a natural fit for many relevant applications.

Beyond viral processes

Diffusion can be more subtle and more complex as well. The simple models that we have presented leave out whether, for instance, there are global complementarities or rivalries in the diffusion opportunities. Perhaps an individual is less likely to tell others about a health service if there are only a limited number of slots to take up the service: the virality then is endogenously affected. If being the first to tell others about the topic

is important, then virality itself can die out as the topic becomes more widely known. The basic ‘epidemic’ structure of the process is also very simplistic, and rules out more complex forms of interdependence between nodes.

Even abstracting from these complications, pure viral models do not seem well-suited to tractable analysis of optimal targeting or seeding. When targeting is known to matter empirically, researchers have turned to other models, such as a theory in which the number of times a message ‘goes around’ matters for its effectiveness (Banerjee *et al.*, 2013). These theories are interesting in their own right, and seem to be more useful tools in studies of seeding. We mention these in discussing our next set of models.

(iii) Aggregation models

The study of aggregation focuses on the details of how something of an intermediate intensity—e.g. an opinion or belief—adjusts in response to neighbours. For example, consider a case in which uncertainty is of first-order importance, e.g. people are trying to decide how much of their land to devote to a potentially promising new crop. Models of aggregation posit that individuals somehow combine current opinions in their community, and possibly private information, to update their views.

For concreteness, imagine that individuals have access to a new technology, such as a new irrigation system or fertilizer, but it is not known whether it is worth paying the cost to start using it. There are some informative ‘signals’ held by some people about the new technology—e.g. via the experiences of their extended family members in other villages. The ideal thing to do would be to have everyone in the village aggregate all the signals—all the information at their village’s disposal—to fully learn and estimate the technology’s value and then make decisions accordingly. In practice, this aggregation happens in a decentralized manner through word-of-mouth communication in the social network.

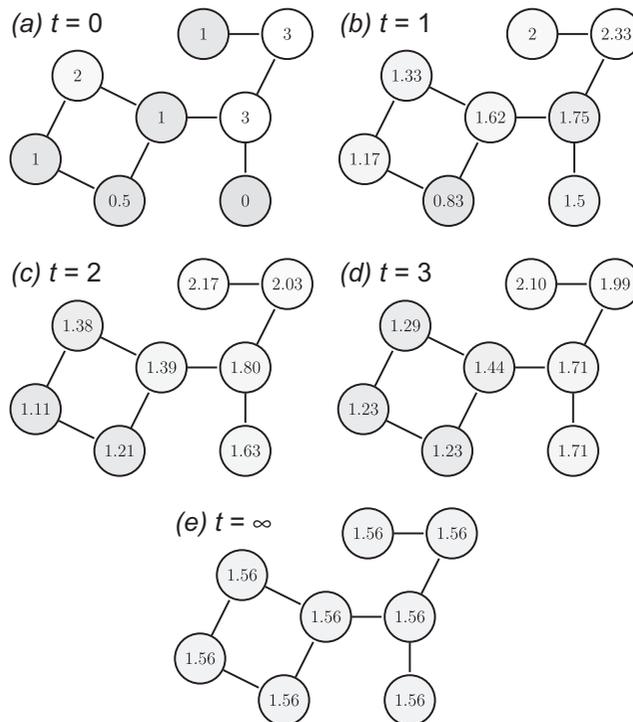
Having a way to study aggregation is crucial in investigating social learning. When a new technology that is potentially profitable is introduced to a community, information diffuses about its existence or basic description but, in the process, individuals also often aggregate views about how good it is, etc. This applies not only to new technologies, but also to health programmes, job opportunities, and so on. For a policy-maker, the reason to consider the role of social networks in this context is that the quality of information aggregation (and therefore the probability of individuals making correct decisions in adopting useful technology) may crucially depend on their structure.

Average-based updating: the DeGroot model

The so-called DeGroot model relies on a very simple ‘rule of thumb’ model of aggregation, taking a perspective informed by behavioural economics (DeMarzo *et al.*, 2003). It has proven to be a useful and intuitive perspective for thinking about aggregation questions in networks.

Basic mechanics. Consider the following set of rules describing individual communication and opinion updating. (See Figure 3 for an example of DeGroot updating.) Everyone initially receives some signals, where the signals can be represented as real numbers (e.g. an estimate of the net gain that can be expected from adopting the technology). These

Figure 3: DeGroot updating in a sample social network



signals are individuals' $t = 0$ estimates. In every period, they communicate their estimates to others and form new estimates by averaging what they hear in their neighbourhood. So perhaps Ann received a signal of 2, while her friends Bob and Carol received signals of 1. Then Ann simply averages 2, 1, and 1 and now updates to an opinion of 1.33. Meanwhile Bob and Carol have done the same with their own friends; they have opinions of 1.17 and 1.62, respectively. Then in the next period, Ann would hear these and update her opinion to an average of 1.33 (her $t = 1$ opinion), 1, and 1.63 to yield an estimate of 1.39. This repeated averaging process continues iteratively, for an indefinite duration, or until some given stopping time. The example as we have presented it has all links being bilateral, so if i pays attention to j then j pays attention to i as well; this need not be the case, and one can define a version of the model with directed links.¹³

In the taxonomy we presented earlier, the DeGroot model is, like the viral contagion process, simple and mechanical. However, this model is designed to study aggregation rather than diffusion, and unlike the viral model, it features continuous opinions rather than discrete states.

¹³ In the example, individuals take a simple arithmetic mean of all opinions in their neighbourhood, including their own past opinion. In a more general version of the model, we could have individuals taking weighted averages, with weights reflecting how much they trust or respect the opinions of others. See [Golub and Sadler \(2017\)](#) for a more formal and detailed presentation.

Convergence and consensus. Some basic facts about the DeGroot model are as follows. First, assuming the network is connected, individuals will converge to a consensus opinion: everyone will tend to agree about the gains to the innovation. Second, this consensus will be a weighted average of initial opinions, with the weight of an individual being equal to her *eigenvector centrality*.¹⁴ That is, we can write

$$c = \sum_i \pi_i x_i(0), \quad (1)$$

where c is the consensus, π_i is the centrality of i , and $x_i(0)$ is the $t = 0$ estimate or opinion of i . Thus, a one-unit change in i 's initial estimate makes for a π_i unit change in the consensus.

The defining property of eigenvector centrality (which turns out to be essentially uniquely defined for connected networks) is that i 's centrality is equal to a weighted sum of the centralities of i 's neighbours. The weight that neighbour j 's centrality gets in this calculation is the amount of attention j pays to i in the updating process. Thus, someone is influential if she is listened to by influential network neighbours; the more influential those neighbours are, and the more they listen to her, the more she is influential. For example, the centrality of the agent in the middle of the [Figure 3](#) network who starts with opinion 1 is equal to 0.25, so she has a weight of 25 per cent in determining the network's opinion. The simple relation between consensus opinions and initial estimates (equation 1), and the fact that network centrality provides the coefficients in the formula, gives a sense of why the DeGroot model is such a nice 'pocket calculator' for social learning questions.

The model also provides a key insight about which individuals a policymaker may want to seed or influence. If the policy-maker can approach several people and persuade them that the innovation is of high quality, she is best-off doing this for the high-centrality individuals. This is related to the fact that their opinions 'go around' more than those of the less central.

If the right answer is known, then it is best to give it to the most central, and it is helpful to have central people in the population. On the other hand, when information is widely dispersed throughout the population, we will see in the next section that more egalitarian communities are better able to arrive at the truth.

Wisdom of crowds. We now use the DeGroot model to study the aggregation of dispersed information. Suppose individuals have initial estimates $x_i(0) = \theta + \varepsilon_i$, whose expectation is the true value, θ , of the parameter of interest (e.g. the true value of the technology), and ε_i are uncorrelated errors. We can give some conditions such that there is a 'wisdom of crowds', with the consensus estimate being close to θ . Indeed, we can show that the expectation of the consensus c in equation (1) is θ . The variance of the consensus estimate is bounded by a constant times the maximum of anyone's centrality in the network. In a reasonably egalitarian network, the maximum centrality will tend to 0 as the network

¹⁴ See [Golub and Jackson \(2010a\)](#) for an exposition in the context of this model, and proofs of the statements made in this paragraph; see [Jackson \(2010\)](#) for a more general introduction to eigenvector centrality and other related measures. We normalize eigenvector centrality so that the entries of the vector π sum to 1.

grows large. Thus, as long as no individual is too central, the consensus is much closer to the truth than the initial opinion of a typical individual.¹⁵ On the other hand, if some individual is highly central, then that person's idiosyncratic error affects everyone, and biases the group's beliefs in a sort of 'echo chamber' effect.

Speed of convergence. Another interesting issue that the DeGroot model lets us study concerns the speed of convergence and the nature of disagreement. Networks that exhibit *homophily*, with inward-looking groups that do not pay much attention outside themselves, can be very slow to converge to the group consensus (Golub and Jackson, 2012). This is an issue of primary importance in the context of the developing world: ethnic, caste-based, or religious segregation is frequent (McPherson and Smith-Lovin, 1987; Banerjee *et al.*, 2013) and in these settings social learning may operate very slowly. Because people in a tight-knit sub-group tend to listen to each other, loosely speaking they will converge to an internal group consensus first, and only then slowly inch toward the consensus view of the larger community. These issues are explored in detail in Golub and Jackson (2012).

Rational aggregation

We have, thus far, presented the DeGroot rule as a simple updating dynamic, appealing for its simplicity and tractability, without discussing incentives or optimization. It is natural for economists to ask in what settings this rule is in some sense rational. In this section we consider what is known about rational (Bayesian) aggregation, explain what challenges modellers face in taking the full Bayesian approach to networks, and some current work in meeting those challenges.

First, we give a basic sense of what we mean by Bayesian updating. Let us consider a version of our running example. A group of individuals commonly observe a set of signals—for instance, reports (external to the network) about the quality of a new fertilizer. There are two possible states: the new fertilizer is worth adopting or it is not. A Bayesian model assumes they use Bayes' rule to update their beliefs about the state. Suppose there are six sources whose information is independent conditional on the true state. Each source issues a recommendation of whether to use the technology; it is 2/3 likely to be correct. Four of the sources report that the technology is good, and two report that it is bad. If *ex ante* the probability of the technology being good is 1/2, then after observing their information, agents would place probability 80 per cent on the technology being worthwhile to adopt.

A large literature, going back at least to the jury theorem of Condorcet (1785), considers Bayesian learning.¹⁶ Bayesian models that are simple and tractable in networks typically assume information flows rather frictionlessly. For example, Acemoglu *et al.* (2014) and Mobius *et al.* (2015) both study communication models where all agents pass around signals 'tagged' with the names of the nodes where they originated, so that everyone learns the signals without needing to do any complex inference. In this case, Bayesian learning is arguably the right benchmark: once signals are effectively public,

¹⁵ These 'wisdom of crowds' issues are studied in detail in Golub and Jackson (2010a).

¹⁶ For development applications, papers in this spirit include Foster and Rosenzweig (1995) and Conley and Udry (2010).

the situation is reduced to the example we gave above. Strikingly, [Mobius *et al.* \(2015\)](#) presents some evidence that when there are not ‘too many’ signals to process, college students in a field experiment are capable of learning in a way that is not too far from the predictions of a tagged information model!

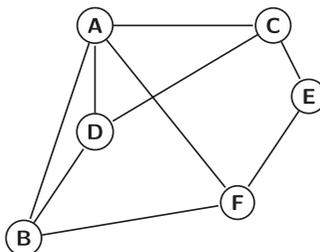
Beyond tagging: Inferring the meaning of neighbours’ actions. However, in some cases it is clear that individuals do not, and cannot, communicate ‘everything they know’, tagged with its origin. There may simply be too much information to keep track of. For instance, to continue with our running example, suppose that information about the new fertilizer initially reaches different nodes in the network, who must then communicate it to ensure its diffusion. Instead of there being set probabilities of each signal’s correctness, each signal requires, for its interpretation, auxiliary information about the source’s credibility, its potential bias, and the relevance of its experience to the present decision. In realistic circumstances this is a vast amount of information, and all agents have to keep track of it to do Bayesian updating based on all the available data.

A tempting alternative to the burdensome tagged protocol is for nodes to report a much less detailed summary than ‘everything they have heard’—for example, their best guess of what the right decision is between the two options (use the technology or not). In the spirit of a Bayesian model, it is also tempting to assume that agents learn from others rationally, understanding their true information content. In terms of our taxonomy, we have come to a type of model that is focused on aggregation, has discrete actions, and where behaviour is very sophisticated—held to the high standard of Bayesian rationality.

Sequential Bayesian updating. We start with one of the few models in this class that is tractable—a classic perspective due to [Bikhchandani *et al.* \(1992\)](#) and [Banerjee \(1992\)](#). These models show that an arbitrarily large group of decision-makers can systematically make the wrong choice (e.g. not choosing the right technology). This point is made in a stylized setting, where each individual only learns and acts once, rather than repeatedly. Imagine a long line of villagers, each deciding whether he should adopt the new irrigation system or maintain the old one. Every individual receives a signal and observes all past decisions. The first individual, 1, receives a signal and uses it to choose one of the two options. The second individual, 2, has his own signal and sees the decision of 1. Crucially, 2 does not get to see 1’s signal, only the choice made. Then 2 makes his decision. Next, suppose that 3 comes along and draws a signal but also sees 1’s and 2’s choices, but again does not see their signal realizations. The process continues this way.

The first result of the herding models is that even if the line has an arbitrarily large number of individuals, it is entirely possible that the vast majority of them make the wrong choice. This is so even though there are arbitrarily many signals, which—if they could all be seen by the community—would lead to a very strong belief in favour of the right decision.

Why does aggregation fail here? The key observation is that *past decisions do not convey all the underlying information*. A simple illustration is as follows. Consider the signal structure we introduced before, where each individual receives a signal that is correct with probability $2/3$ and incorrect with probability $1/3$. If 1 and 2 make the same choice and 3 sees it, he will never follow his own signal, instead deferring to the majority.

Figure 4: A sample network which engages in Bayesian learning

In essence, two signals are always more powerful than one's own private signal, so 3 will find it optimal to ignore his private information and thereby stop adding to the pool of public information. But this causes a problem for 4, who now sees all prior individuals making the same decision, and will, *a fortiori*, ignore her private information.¹⁷

This simple model makes a few key points. First, even if individuals are rational, as opposed to being governed by mechanical rules, there is no guarantee of good aggregation. Learning can inherently involve externalities, even without any bounds on rationality. Second, at a more methodological level, in some cases we can analyse aggregation even when agents have to infer what others' actions mean. However, as we will now see, these features of the Bayesian learning model do not extend easily to more realistic environments.

The difficulty of Bayesian updating in general. In practice, agents do not make decisions in this neat sequential way; they aggregate information and report opinions repeatedly. As it turns out, doing Bayesian extraction of information from one's neighbours' reports can become extremely burdensome in that case. To see this, consider a case (see [Figure 4](#)) in which individual *A* speaks to *B*, *C*, and *D*. *B* also speaks to *D*. *C* speaks to *D* and *E*. Finally *F* receives information from *B*, *E*, and *A*. The connections are bilateral: conversations happen in both directions along any link. This is an environment typical of a social network in the real world with overlapping friend sets. If *F* makes a decision about which technology to use or recommend based on conversations with *B*, *E*, and *A*, it is a difficult problem for *F* to figure out how much of the information is actually independent, and how much is communicated from a common, upstream source. This is a very complex computation and there is a question of whether individuals are likely to be well-described by a model that has them doing these computations with ease.¹⁸ In fact, evidence from a variety of experiments suggests that this may not be a good description of behaviour (see [Grimm and Mengel \(2018\)](#) and [Chandrasekhar et al. \(2019\)](#)). In addition, in Bayesian learning models, there is often the assumption that the agents know the network structure (and know that everyone knows it—in fact, have *common knowledge* of it), which is quite unrealistic ([Kumbasar et al., 1994](#); [Casciaro, 1998](#); [Chandrasekhar et al., 2018](#)). In the case where the network is unknown,

¹⁷ See [Golub and Sadler \(2017\)](#) for a much more detailed discussion of this type of model and its generalizations.

¹⁸ For example, [Eyster and Rabin \(2014\)](#) argue that any rational learning model requires agents to 'anti-imitate' at least some of their network neighbours to correct for the redundancy of signals originating at a common source, and that this behaviour is unrealistic in practice. See [Häzla et al. \(2017\)](#) for some formal theoretical hardness results.

the problem becomes even more outlandishly difficult. Thus, while the basic idea of Bayesian updating models is appealing on a normative basis, we see that even in the apparently simple case of a smallish network and a fixed state, the calculations that must at least implicitly occur for Bayesian updating are arguably not plausible.

Is Bayesian updating the right descriptive model? We have seen that it is hard, even for an omniscient analyst, to give explicit formulas for how agents would optimally solve the problems we have been discussing in general settings. Some theorists have heroically managed to prove some things about what *would* happen if agents were able to solve these problems. Mossel *et al.* (2015), show that rather generally, for a wide class of large networks, a group of Bayesian learners will indeed come to the right conclusion about the state of the world—in our example, whether adopting the technology is the right decision.¹⁹ This prediction shows that, even if we are willing to assume individuals are very powerful in terms of their computational abilities, fully Bayesian models with repeated communication may run into difficulties matching the observed fact that individuals often do not learn well.

(iv) The frontier in learning models

Now that we have surveyed some canonical concepts and building blocks in learning models, we turn to the frontier of research. This subsection examines some current perspectives that combine features of the models we have discussed with the aim of bringing the insights closer to policy applications.

Naive learning with uninformed agents: A hybrid model

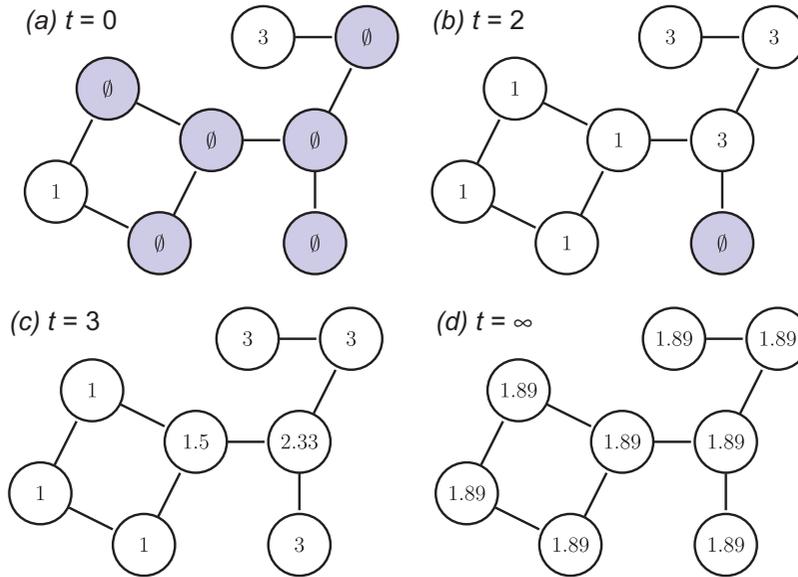
Thus far, we have thought of diffusion and information aggregation as separate concepts, but this conceptual separation, helpful for theoretical purposes, is unrealistic for policy analysis. A simple generalization of the DeGroot framework, due to Banerjee *et al.* (2019a), allows us to consider both in one model.

Typically, policy-makers introduce a new technology to a select set of individuals, and in effect the seeds get signals about the technology. Others, however, are at first not aware it exists and have no basis on which to form beliefs about it. To consider our example again, a new irrigation system may be demonstrated to a set of seed farmers. Then the information about the new technology diffuses (others become aware of it); at the same time, there is also aggregation going on: once aware, people update their opinions about it based on others' views.

This type of model therefore merges both diffusion and aggregation. If an individual has never heard of the topic, he updates his belief about the topic according to his neighbour's belief the first time he hears about it. But those who have heard of it learn from each other according to the DeGroot model. So, if *A* newly hears about the technology from *B*, then he simply adopts *B*'s belief. But if *A* has heard about it and has belief *x*, while his neighbour *B* has belief *y*, his neighbour *C* has belief *z*, and his last neighbour *D* has never heard of it, then in the next period *A* has belief $(x + y + z)/3$.

¹⁹ The conditions have the flavour of network balance: there cannot be segments of the network that are only very distantly connected to those who observe them.

Figure 5: DeGroot updating with uninformed agents in a sample social network



The process is illustrated in [Figure 5](#).

One can think of the network as consisting of domains of influence by each of the seeds: for every seed the collection of individuals who are closer to him than other seeds. Individuals' first beliefs are inherited from their closest seeds. If the domains of influence of various seeds are very unequal in size, then that seed can effectively exert an undue level of influence, disrupting good aggregation, for the same reasons as we discussed above in analysing the DeGroot model. On the other hand, if the domain of influence of each seed is roughly of the same order of magnitude, then this obstruction to good learning does not arise.

Suppose that an authority delivers information to a set of seeds, and some of these seeds happen to form a connected subnetwork of the community. Then only the friends who have many connections to the network outside that community will have meaningfully large domains of influence, but those who are predominantly connected within this group will have negligible domains and therefore the signals of these individuals in the group will have very limited reach in seeding the aggregation process.

A policy insight coming out of this model is that giving information to a cluster of connected individuals could actually be detrimental to learning compared to a policy where the seeding is less concentrated. This may seem counterintuitive, but is a direct implication of the hybrid of seeding and diffusion we have presented. A broader implication is that diffusion vs aggregation is not an 'either/or' dichotomy: these processes coexist in reality, and we can see that when they are considered in tandem, results can be quite different from studying each on its own.

Strategic interactions

Up to this point we have mainly considered examples in which take-up has only an individual pay-off. In reality, pay-offs may depend on others' take-up. For instance, many products such as communication applications for a phone have the property of *strategic complements*, meaning that the technology creates more value when one's contacts in an appropriately defined social network are using it.²⁰ Even technologies that are apparently independent in their use, such as fertilizer for one's farm, may have a strategic complementarity arising from the fact that if friends use it, one can receive advice and support from them if difficulties arise.

The presence of strategic complementarities makes diffusion more complex, at least insofar as agents behave strategically in view of them. When an individual hears about the new product, he needs to assess how many people likely will eventually hear and take it up themselves. This is explored in detail in [Sadler \(2018\)](#).

The results are quite subtle. On the one hand, complementarities can accelerate diffusion. When a product is going viral, adoption is helped by the expectation that one's friends will soon be using it, too. On the other hand, strategic considerations can reverse standard intuitions. For instance, a typical result from viral diffusion models is that when there are more highly central individuals, there should be more widespread diffusion (holding network density—i.e. the number of links divided by the population size—fixed). However, [Sadler \(2018\)](#) shows that once complementarities in take-up are present, the opposite can be true in equilibrium. Though the strategic considerations leading to this result are beyond our scope here, this shows that basic forces can be entirely reversed by taking strategic behaviour into account in a richer model.

An alternative to strategic complementarities is strategic substitutes, meaning that it is *less* appealing to adopt when one's friends are adopting. For example, a learning opportunity (say, an educational programme focusing on health) may be effectively a public good, so that when one's friends take it up, there is less benefit to doing so because one can simply ask them for help if the need arises. This will, again, change intuitions of standard targeting analysis and lead the analyst to focus on very different centrality statistics from the basic ones we introduced above in the sub-section, *Average-based updating: The DeGroot model* ([Galeotti et al., 2019](#)).

'Negative feedback' effects in development interventions can be serious and pernicious, as shown by [Kremer and Miguel \(2007\)](#). Prior to this study, some individuals had randomly received an incentive to take up a drug to fight helminths (e.g. tapeworms, hookworms). This follow-up study found that individuals in areas with high treatment believed less in the effectiveness or value of the deworming drug. The fact that the drug protected everyone in the community actually led to lowered incentives to adopt.

Thus, the policy implication here is that strategic complements or substitutes can change the basic insights of optimal seeding for diffusion. Generally, policy-makers should be aware of the limitations of mechanical viral models, even if the direction of a strategic effect seems intuitive at first glance. A full exploration of these forces is just beginning, but policy-makers should at the very least assess whether strategic consideration of others' adoption decisions is playing a role in the process before applying models that focus purely on learning.

²⁰ This may or may not be the network that mediates learning about the innovation.

(v) Closing reflections

In this section, we have surveyed models that are helpful for analysing the flow of information in networks. The broader question at hand is how to take these insights to a policy setting. We now offer some closing reflections on this issue.

Should we be using networks?

A first question is whether to bother with a network approach at all. In the scenarios discussed above, the policy-maker might maximize take-up by broadcasting a message and effectively ‘seeding’ every individual, or at least attempting to do so. A natural question is why use network-based policies as opposed to such broadcasts. To sharpen this question, recall that a key assumption in the model of a simple viral contagion is that a single exposure can be enough to inform a node, assuming the information ‘gets across’. This entails, loosely speaking, that information can be immediately trustworthy. But when such a situation arises in practice, there are often inexpensive means to broadcast that information widely—e.g. through television, radio, SMS campaigns, letters, and so on.

Our view is that broadcasting often *is* better, but by the time a policy-maker is thinking seriously about a network approach, there is some reason that broadcasting is not effective enough. This may be because information is more interesting, more trustworthy, or more salient if conveyed by social contacts. It may also be that the target population is simply hard to reach: they may not engage with the media the policy-maker can use for broadcasting.

In some examples, there is good reason to believe that interest in leveraging network learning is justified, and that a network-focused policy is better than a broadcast. [Banerjee *et al.* \(2019b\)](#) show that, in advertising an immunization campaign, delivering information to 30 random seeds does not have a meaningful benefit whereas contacting the ‘right’ seeds suggested by network theory leads to a marked increase in vaccination take-up. [Schilbach \(2015\)](#) finds that an agricultural technology spreads much more effectively if farmers have a means of teaching their contacts how best to use it. [Banerjee *et al.* \(2018a\)](#) show that individuals’ willingness to engage in social learning may actually be decreased by broadcasting information.

Thus, a blanket statement that broadcasting is superior is unlikely to hold. Nevertheless, a simple but important observation is that the policy-maker ought to consider *why* broadcasting is not a better policy: this will illuminate which of the types of models we have mentioned is likely to be most relevant. Since the assumptions that make for effective broadcasting are often the same as those that make for effective viral diffusion, this suggests that *when* networks are useful and whom to seed is an important question, it is a live possibility that models of the kind we have focused on from section II(iii) onward will be the more relevant ones.

What to measure?

Where a network approach *is* necessary, an obvious question is what a policy-maker should measure to implement the policy. A first observation, and a crucial one, is that network data are not necessary at all to implement some of the insights we have surveyed.

For example, recall the basic model of viral information transmission from section II(ii). There, we saw that the basic reproductive number was of central importance. The

fact that low-dimensional structural parameters of the diffusion model, rather than details of a network, are most important has practical implications. If diffusion is, indeed, viral, then a policy-maker is best served by fine-tuning the message to make it spread as fast as possible. [Cheng et al. \(2014\)](#) (one lead-in to a sizable literature) shows that assessing virality is realistic in practice based on observing messages spread in a network.

Other insights are robust to details of network structure as well. In section II(iv), Naive learning with uninformed agents: A hybrid model, we saw that clustering informed nodes can actually limit the reach of their messages, and there is a case for ‘spreading out’ the initially informed across a large network. Detailed network data are not needed to achieve a high probability of such spreading out. Random seeding will work fine, and the techniques of section IV can be used to refine this policy.

In some cases, network data are relevant. We have seen that eigenvector centrality predicts the influence of a node in a DeGroot aggregation model. In related models, such as [Banerjee et al. \(2019b\)](#), both the theory and evidence show that seeding central nodes can be critical, and superior to broader seeding of random nodes. Fortunately, [Banerjee et al. \(2019b\)](#) also describe strategies for cost-effective identification of the seeds likeliest to work well. In essence, the policy-maker can simply ask people who is best at spreading information—the paper’s term for such a person is a *gossip*. As the paper argues, the very qualities that make a node able to spread a message well also make that node’s prominence legible to the other inhabitants of the network: a good spreader of information is someone that people are likely to hear about indirectly. That fact can be used to elicit the identities of gossips without the need for collecting network data. Indeed, [Banerjee et al. \(2019b\)](#) show that a very simple elicitation in the spirit of ‘name the best person to spread information’ does a remarkably good job of identifying the optimal seeds that would be identified by a much more elaborate network census.

III. Financial contracting

(i) Motivation and outline

Expanding financial access. Expanding financial market access to underserved households is a major concern of policy-makers in developing countries. Most financial market expansions aim to either (i) improve household risk management or (ii) relax credit (or savings) constraints that inhibit profitable investments from being made. Financial inclusion policies take many forms, including subsidies or guarantees for savings accounts, loans, payment systems, and insurance products. To take just one example, 225m bank accounts were opened between 2014 and 2016 through an initiative of the Indian government ([Agarwal et al., 2017](#)). Despite these efforts, 1.7 billion adults globally remain unbanked, and many households with formal accounts use them only rarely, if ever ([Demirgüç-Kunt et al., 2018](#); [Badarinza et al., 2019](#)).

Informal financial transactions and social networks. While many households in developing countries live outside of the formal financial system, most nonetheless engage in frequent financial transactions. Financial diaries and survey data from numerous countries

and sources paint a picture of complex (informal) financial portfolios (Rosenzweig, 1988; Banerjee and Duflo, 2007; Collins *et al.*, 2010). Many of these transactions take place between households embedded in the same social network and may facilitate risk management or investments. In fact, consistent with an insurance motive, some informal loans exhibit state-contingent repayments, and transfers between households respond to income shocks (Udry, 1994; Fafchamps and Lund, 2003).

Key questions and a role for research. What implications, if any, does the prevalence of informal financial transactions have for the policy-maker trying to design financial inclusion policies? In this section, we look to the networks literature for guidance on this question.

We first consider risk sharing. In the absence of formal insurance mechanisms, how well can the community self-insure on its own? If full insurance is not possible, then what determines the boundaries of the risk-sharing group? After all, the potential value of formal insurance mechanisms is likely to be greatest exactly where informal mechanisms break down—so it is important to understand ‘where’ that occurs.

We next turn to the impacts of credit expansions in the presence of informal financial transactions. First, how does the formal product interact with informal sources? In other words, is there an impact on total credit supply, or does the formal product simply crowd out the informal insurance arrangements? Second, does the new product change incentives for network formation, affecting the financial relationships and perhaps some others as well?

Risk-sharing arrangements do not exist in a vacuum: they are a special case of enforcing contracts between parties connected by personal relationships of observation and habitual exchange. Risk-sharing interacts with other decisions that shape people’s economic lives and social networks. We consider two case studies that zoom out from risk sharing to inseparable broader issues of practical importance. Specifically, we discuss recent research on the role that information flows in a network play in enforcing contracts. We then examine some evidence on how risk-sharing considerations affect migration, and thus shape the very composition of a local risk-sharing community.

It is worth noting at the outset that the discussion here often centres more on theoretical models and, in many places, highlights where empirical studies of networks’ role in risk-sharing are limited. This opens up exciting avenues for research and policy experiments, but makes the orientation of this section necessarily somewhat more tentative than cases where, as in the information-sharing literature, there has been more practical application of network theory.

(ii) Risk sharing

Households in developing countries face numerous sources of income risk with limited scope for formal insurance. The promise of informal insurance is that risk-averse agents with uncertain incomes can arrange for state-contingent income transfers among themselves that are Pareto improving.

A benchmark model of risk sharing

We can consider a venerable simple model that provides a useful benchmark: conditions such that the community can fully insure agents against all diversifiable risks (Wilson, 1968).²¹ This basic benchmark model does not include a meaningful role for networks, but it is a useful point of departure. Individuals have uncertain incomes, which may be affected by their effort. Information flows perfectly, so all individuals have common knowledge of others' income, effort, savings (if saving is possible), and so on. Furthermore, agents can all commit to contracts: full specifications of how much is transferred to each individual in each state of the world. Given income realizations, then, individuals can effectively divide the resources—the total income of the village or total amount to be consumed given a savings decision—in a way that is Pareto efficient given available resources. This entails that they equate marginal utility ratios across states for every pair of people. As it turns out, this is equivalent to consumption being allocated to maximize expected utilitarian welfare with some Pareto weights stipulated by the social planner (taking into account heterogeneities across agents in risk-tolerance, etc.).

Endogenizing network structure

One natural question unresolved by the benchmark model is how a risk-sharing community is determined. Relatedly, how much of an agent's total income risk is informally insured?

Underconnected equilibrium networks. The benchmark model can be expanded to explore these questions by explicitly modelling network formation (Bramoullé and Kranton, 2007). Rather than taking the risk-sharing group as given, we can ask what it will endogenously emerge to be in a model where its composition is the result of local choices.

Say that i and j are directly connected if they can engage in direct bilateral transfers of resources. Next, we can define a notion of indirect linkage: i and j are in the same network component if there is a series of links making a path from one to the other. In this case, we will say they are in the same risk-sharing group. This is a big assumption: it entails that any physically feasible transfers and communication can be costlessly carried out in this network. We return to weakening it later, but this is an interesting starting point to consider which network will be formed.

To this end, suppose that maintaining any link has a social cost. Recalling that the full benefits of risk-sharing can be achieved by any indirectly connected individuals, ideally (from the perspective of social efficiency) there are no redundant links whatsoever: every individual is connected to everyone else in the cost-minimizing way. Thus, the efficient network consists of a single connected component that (i) satisfies full risk sharing, and (ii) is such that removing any link would disconnect the network.

The efficient network need not be the one formed in equilibrium. To formalize this notion, network theorists have introduced notions of *stable networks* (see Jackson (2010) for a discussion). In general terms, a network is stable if some given set of deviations is not in agents' interests. At a minimum, for a network to be stable, no agent who

²¹ See also Debreu (1959); Arrow (1964); Diamond (1967).

is paying for a link should want to stop supporting that link and lose it. Notions of stability also consider the possibility of *adding* links: in the [Bramoullé and Kranton \(2007\)](#) case, it should be that no pair of agents profits from creating and paying for a link.

Stable networks are typically smaller with less risk sharing than the efficient network. This is because individuals benefit if their risk-sharing partners are also linked to others; when individuals make their linking decisions, they do not internalize the benefits received by others connected to the potential partner. [Bramoullé and Kranton \(2007\)](#) make this point in a canonical model of network formation, and also examine how the endogenous formation of the risk-sharing network can lead *ex ante* identical individuals to have unequal outcomes.

Bargaining over network formation: overconnection and implications for inequality.

Picking up this theme, [Ambrus and Elliott \(2018\)](#) examine in more detail how the distribution of risk-sharing rents affects jockeying for network positions, and the implications for efficiency and equity. Specifically, they study a set-up in which there are two phases. In the first phase, individuals invest in relationships among themselves. This comes at a cost; it takes time and effort to build a relationship. In phase 2, the individuals receive random incomes and make transfers to fully insure each other. Here there is full commitment, complete information, and so on. Thus insurance is perfect among all agents in the same component of the network.

The focus therefore is on two issues. First, do all individuals in the community end up in the same network component once the risk-sharing links are formed? Second, and more interestingly, how is the surplus from insurance divided? Specifically, do some individuals receive more of the surplus than others systematically, and if so, why?

[Ambrus and Elliott \(2018\)](#) study a bargaining model to shed light on this issue. Consider a thought experiment in between phases 1 and 2, wherein an individual i can generate a renegotiation, threatening to drop any relationship with another agent j . If that relationship were to disappear, then both would remain in an altered network with only that link removed. This may or may not change the pay-offs for both i and j . If i stands to lose more from the removal than j , then j will, all else equal, be in a stronger position with respect to i in bargaining over surplus.

To see this in an example, consider the case of three players and a line $i - j - k$, with no link between i and k but all other links present. There is some surplus $2V$ that is generated if these three are all mutually linked: this comes from the gain of V when two agents are linked and then another gain of V when the third agent is added. (Conveniently, in the model used in this paper, the risk-sharing surplus in certainty equivalent terms is simply proportional to the number of people in that component minus one.) How is this total surplus divided? It is intuitive that the middle agent j should be able to secure a higher pay-off through bargaining. If j threatens to cut the link with i , then i is alone in autarky, whereas j and k are connected. In that case j receives surplus $V/2$, the value from sharing with k , whereas i receives no surplus from risk sharing. Thus (and we are doing the bargaining theory loosely here!), j can extract a payment of $V/2$ from i using this threat. By symmetry, j will receive V units of surplus and each of i and k only receive $V/2$.

The practical takeaway here is that when individuals invest in relationships in phase 1, they should be mindful of the fact that where they end up in the network—that is, how central they are—can affect the surplus they capture in the subsequent risk-sharing

process. This can lead to *overinvestment* in links, in contrast to the underinvestment we saw before, because individuals want to secure a larger share of the surplus. Individuals create more links than are necessary to connect the network component. (Recall the fact that they can achieve full risk sharing for n people with only $n - 1$ links—for instance, on a line.)

Of particular importance to the developing world is the fact that this model can also be extended to the case of multiple types. In many contexts, individuals belong to different groups such as caste, ethnicity, or nationality. Naturally, people may find it less costly to form links with members of their own group. But, if incomes are less correlated *across* groups than within, the existence of bridging links may substantially improve overall insurance. The central individuals in each group will have the strongest incentives to form the bridging links, as they can extract more of the surplus from members of the same group. This gives rise to further inequality in the level of consumption.

This model also allows us to ask which network structures arise as a function of how costly it is to form links. In a situation where building links is very costly, the only type of network that is stable is a star: all agents only connected to a single central agent and not to each other.²² Practically this means that if relationships are hard to build, we expect high inequality in relationship patterns and therefore in average consumption. But if relationships are cheap to build, then even the most equitable minimally connected network structure—a line—is also possible. This has policy implications if, for example, the policy-maker can invest in technologies that lower the cost of forming links.

Risk sharing under limited commitment

As demonstrated above, endogenous formation of risk-sharing networks can lead to inefficiencies and asymmetries in link formation and inequality in consumption allocations. Crucially, however, full risk sharing still occurs conditional on the network that is formed. However, the assumptions that give rise to full insurance are strong and unlikely to hold in practice. One well-studied relaxation is limited-commitment risk sharing, where individuals may renege on promised contracts (Kocherlakota, 1996). If an individual happens to have a very high income in a given period, he may choose not to pitch in and make transfers to others, preferring to lose the relational contract and keep the money. Knowing this, the insurance demands made from the community will endogenously be lower than in the case with perfect commitment, and the community will not achieve full insurance.²³

The standard limited commitment models assume that if an agent (person or household) defects on the insurance arrangement, the community can punish that agent with full exclusion from all future risk sharing. But such trigger strategies require full knowledge and coordination by all members of the network, and are thus more of a theoretical device than a practical description of social life. Bloch *et al.* (2008) examine in much more detail how a community might come to know of deviations, and thus that study gets into the foundations of the informational assumptions in risk-sharing models. In a model where the network is exogenous, they posit that the same network

²² The centre of the star pays large link costs but enjoys high bargaining power from his links and therefore can seize much of the surplus.

²³ By full insurance, we mean what a utilitarian planner would implement with full information about income realizations and no constraints on transfers.

structure that enables risk sharing also plays a role in diffusing information about deviations by network members. After all, individuals must first learn of the deviation in order to punish the behaviour. This means that the network structure itself mediates the degree to which limited commitment erodes risk sharing.²⁴

This suggests that we should better understand which network shapes are more conducive to cooperation—i.e. which allow for better enforcement of risk-sharing agreements. This can help policy-makers identify those communities that might benefit most from outside intervention. In order to make progress towards this question, Bloch *et al.* (2008) make the simplifying assumption that a social norm governs the size of transfers between any two individuals. Agents must simply decide, for each exogenous network link, whether or not to engage in bilateral transfers according to that norm.²⁵ One key finding is that stability is non-monotonic in the density of the network, holding the number of communication rounds fixed. In sparse networks, the deviator does not have many risk-sharing relationships, so every link is valuable. Therefore, deviation is costly even if only the aggrieved party punishes the deviator. When the network is dense, an individual has many risk-sharing partners. However, news of the deviation is likely to spread to many of those individuals, again making deviations costly. In the middle, because there may be both limited benefits to an incremental link and limited diffusion, risk sharing may actually be harder to sustain.

Ambrus *et al.* (2014) give a different approach to the question of which network structures give rise to better or worse risk sharing by bringing into the picture other (non-risk-sharing) benefits that agents derive from their network. For tractability, they work under the assumptions of exogenous network structure and limited commitment and assume that all agents can observe the full realization of incomes and transfers. However, they allow the extent and structure of risk-sharing agreements to be endogenously determined. To capture non-risk-sharing benefits, they posit that an agent i experiences a direct consumption value from being linked to j , $c(i, j)$, independent of any financial transfer. (This can be thought of as friendship, a reduced-form way of capturing favour-trading, etc.) Thus, agents can leverage ‘social collateral’ to enforce transfers, deleting links in the case of a deviation.²⁶ The potential extent of insurance in a given network is predicted by the ‘expansiveness’ of the network.²⁷ Networks are expansive, loosely speaking, if, no matter how we select a set S of nodes, the number of links this set has to nodes outside itself is not too low on a *per capita* basis (i.e. per member of S). For an allocation of consumption and transfers to be feasible, no group can have an incentive to deviate from that allocation in any state. The maximal punishment that can be imposed on a group of deviators severs all outside links to members of the group. Therefore, the larger the number of outside links for any arbitrary group, relative to group size, the easier it is to enforce *ex post* transfers.

In addition to providing a metric relating network structure to the capacity for risk sharing, the model can also give predictions about the nature of risk-sharing

²⁴ See also Ali and Miller (2013) who explore how a network structure impacts the scope for cooperation when information about deviations must diffuse through the graph.

²⁵ Agents are fully aware of the net obligations of each partner to others.

²⁶ This idea was used by Karlan *et al.* (2009) as a means of enforcing informal loan repayments.

²⁷ To show this, they first demonstrate that the model of risk sharing enforced through individual link deletions is isomorphic to a model of social ostracism (i.e. link deletions by groups of individuals) and group deviations.

relationships in a constrained efficient, second-best world. Ambrus *et al.* (2014) show that the network organizes itself into ‘risk-sharing islands’, which are groups of nodes that share risk fully with one another. The islands are connected to one another, but the enforcement constraints bind on each cross-island link, leading to only partial risk-sharing between the islands.

The logic of Ambrus *et al.* (2014) can help to give a policy-maker guidance when thinking about where to prioritize the introduction of external risk-sharing tools. Less expansive, fractionalized networks will exhibit less insurance and therefore may benefit most from the introduction of formal insurance. Such tools might be especially beneficial if they target the types of risks that are insufficiently insured across groups, rather than crowding out links within-group. We return to the issue of crowding out later.

Endogenizing network structure with contracting frictions

So far, we have considered network formation and contracting frictions separately. But do the implications for the policy-maker change when these features are combined? Jackson *et al.* (2012) introduces endogenous network formation to a model of favour exchange, which can be viewed as a metaphor for risk sharing.

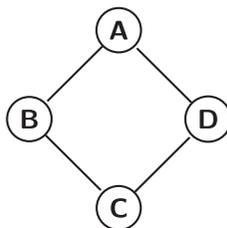
In their model, the need for a favour on a link has a Poisson arrival rate. Doing a favour has cost which is strictly lower than the value of receiving a favour. The authors ask which network structures can sustain favour change in equilibrium, if agents can punish any nodes who withhold favours with permanent link deletion. They argue that, while the grim trigger strategy of deleting all links in the whole network following any deviation is subgame-perfect, it is also ‘drastic and unrealistic’. Therefore, they search for *renegotiation-proof* networks and punishments, which involve punishments that agents would not want to revise.

Let m denote the smallest number of links such that the threat of losing them is enough to incentivize good behaviour.²⁸ Moreover, suppose that if any agent ever fails to perform a favour, the link involved in the deviation is automatically dropped. (This is a reduced form for bilateral punishment strategies.)

The authors give an example using the graph in Figure 6, with four nodes and four edges $g = AB, BC, CD, DA$, and where $m = 2$. Suppose that A fails to perform a favour for B . Then, edge AB is dropped automatically. However, in the resulting network, B only has one neighbour. By the assumption that $m = 2$, one link is not valuable enough to enforce the provision of favours by B . That is, B cannot be incentivized to perform favours if the threat of punishment is the deletion of only 1 link. Therefore edge BC must drop as well. But by this logic, no remaining edges are stable, and the whole network must collapse. This punishment threat is renegotiation-proof, because cooperation cannot be sustained once AB is deleted.

This example also highlights one unattractive feature of some renegotiation-proof equilibria. Some nodes that are not directly connected to the deviator may nonetheless experience a loss of links following a deviation, in this case node C . This leads the

²⁸ In more detail: let c be the cost of performing a favour, v be the benefit, p be the probability of needing or being called upon to provide a favour, and $\delta < 1$ the discount factor. Then m is the smallest integer such that $m \frac{\delta p(v-c)}{1-\delta} > c$: the expected future net benefit from remaining in the network is greater than the cost of providing a favour.

Figure 6: A sample network that is not robust

authors to additionally focus on networks and punishment strategies that are *robust* to social contagion, meaning that they involve only link deletions local to the violation. Jackson *et al.* (2012) show that stable, renegotiation-proof, robust networks are ‘social quilts’: m -cliques (complete sub-graphs of $m + 1$ nodes) attached to one another such that no cycle in the network has length longer than $m + 1$. The example above is not a robust network: note that it fails the cycle length requirement—the longest cycle is $4 > m + 1 = 3$. It also clearly fails the 2-clique condition.

Another implication of the model is that in any robust network, every link must be *supported*.²⁹ That is, for every ij pair, there must be a node k that is linked to both i and j .

This model suggests that networks with structures that look closer to social quilts will be able to sustain more informal financial transactions. It also suggests that one (unsupported) link across two groups with uncorrelated sources of risk might not suffice in a limited commitment environment.

Empirical evidence on the extent of risk sharing

Above, we have described the workhorse models of risk sharing. Incorporating the network, either through endogenous network formation or as a parameter in the punishment technology gives rise to predictions on the extent of possible risk sharing, the limits to membership of risk-sharing groups, and on the network structures that are more or less conducive to informal insurance, each of which is useful for a policy-maker. It should be noted that the models of risk sharing are in general much more complex and subtle than models of information flow alone. While they give conceptual guidance, a diagnosis of how well communities share risk requires empirical evidence. We next take these models to the data.

It is useful to begin with tests of the benchmark, full insurance model. Even for this simplest case, constructing a clean empirical test is very difficult. Most tests take a stand on the unit of risk sharing (e.g. the village) and then look for empirical support of two testable predictions of benchmark risk-sharing models. First, consumption should be uncorrelated with idiosyncratic shocks to income or wealth (Cochrane, 1991). However, this prediction may hold for many reasons besides informal insurance, such as access to savings technologies or formal insurance instruments. Second, consumption should be correlated with the aggregate endowment of the risk-sharing group, and therefore with average group consumption and average group income (Mace, 1991; Townsend,

²⁹ This remains true even when the authors allow heterogeneity in the cost and value of favours across individuals.

1994; Deaton, 1997). This second test is data intensive, requiring panel measurements of income and consumption for a large number of individuals within a risk-sharing group.³⁰ These papers reject full insurance, but nonetheless find considerable risk smoothing within a village.³¹ Note that these tests abstract away from how consumption smoothing is achieved or whether the village is the correct representation of the risk-sharing group. After all, the benchmark theory is largely silent on these questions. Moreover, consumption may co-vary in a group for reasons unrelated to risk sharing, such as preferences and shared consumption experiences such as festivals.

Most of the tests of full insurance posit that the village is the operative unit of risk sharing. However, this need not be the case. Moreover, the constrained efficient risk sharing of Ambrus *et al.* (2014) looks like sub-groups with full insurance within-group and partial insurance across groups. The models suggest factors that influence subgroup composition—namely the costs of linking, the correlation of incomes, and the ability to sanction deviations either through the spread of information on the information network or the deletion of links in the social network. While the costs of forming links might be lower within a village versus outside, other characteristics such as group identities, religion, kinship, or caste might also be conducive to risk sharing.

One direction that the empirical literature has taken is to use observable group memberships as a stand-in for the fully-insured ‘islands’ that we introduced in discussing Ambrus *et al.* (2014). In the context of India, Rosenzweig and Stark (1989), Munshi and Rosenzweig (2016), and Munshi (2017) show that caste groups are a natural place to look for active informal insurance.³² Mazzocco and Saini (2012) find that after incorporating heterogeneity in risk aversion, they can reject full insurance at the village level, but they cannot reject full insurance at the caste level within a village. Extended family is also a natural proxy for risk-sharing groups. Angelucci and De Giorgi (2009) and Angelucci *et al.* (2017) show that the benefits of Progresa, a large, randomized cash-transfer programme in Mexico, spill over onto non-beneficiaries through transfers and loans from extended family network members.

Taking the network models of risk sharing to the data poses several additional complications. Several models are descriptive in nature, and only a few offer sharp, testable predictions about network structure and the extent of risk sharing as a function of network position. It also imposes even more onerous data requirements. First, the ideal data would include detailed social network information in addition to panels of consumption and income information. Very few data sets meet this requirement.³³ Second, as our previous discussion implies, the network data need to be elicited at the right level. In the case of information diffusion, the village is likely the appropriate unit of analysis

³⁰ Townsend (1994) and many subsequent papers perform their tests using data from three Indian villages surveyed by the International Crops Research Institute of the Semi-Arid Tropics (ICRISAT).

³¹ Given the rejection of the full insurance benchmark, others have tested risk-sharing models with incomplete information, typically using the models to generate testable restrictions on the co-movement of contemporaneous and lagged consumption with income. See, for example, Ligon *et al.* (2002) and Kinnan (2019).

³² Caste linkages may also be relevant both within and across village boundaries.

³³ In fact, we are only aware of one, the Townsend Thai Monthly Rural survey (Samphantharak and Townsend, 2010), which includes detailed income and consumption data as well as network edges for a subsample of household nodes. Kinnan *et al.* (2019) use this data to show how health shocks are smoothed by and propagate through the network, considering both risk-sharing and network impacts from effects on the businesses operated by shocked households.

for locally relevant information. In contrast, the risk-sharing network may be much broader, making the collection of network data a much more difficult proposition.

With only network data, it is still possible to examine the predicted risk-sharing properties under the models described above. [Ambrus *et al.* \(2014\)](#), for example, measure the expansiveness properties of network data from Peru. They conclude that the observed network structure is not compatible with full insurance, though the network structure is conducive to ‘very-good risk sharing’. Recall that the model of [Jackson *et al.* \(2012\)](#) has the testable prediction that all risk-sharing links are supported; that is for every connected pair, there exists a third individual with links to each of the pair’s members. The authors show high levels of support in the network data of [Banerjee *et al.* \(2019b\)](#), with rates reaching over 85 per cent for some kinds of favour-exchange relationships. In principle, a policy-maker could look for differences in these measures across populations to help measure the level of informal insurance available to different communities.

Policy interventions in insurance markets

Policy-makers wishing to promote better risk mitigation have several options at their disposal. One is to try to fill in what is missing directly: that is, expand the formal provision of insurance (or credit or savings) through regulations, subsidies, or guarantees. The other is to provide resources that might be complementary with pre-existing risk-sharing, such as decreasing the costs of forming links or of enforcing informal contracts.

Many attempts at the former have been disappointing across a wide set of contexts ([Valdés *et al.*, 1986](#); [Mobarak and Rosenzweig, 2012](#); [Cole *et al.*, 2013](#); [Banerjee *et al.*, 2014](#)). We return to this idea below, with the specific example of microcredit. Our discussion of the risk-sharing literature highlights that it is very hard to measure how broadly risk is being shared, what types of risk are adequately mitigated through informal insurance, and the extent of insurance, even if we observe state-contingent transfers.³⁴ In an ideal world, policy-makers would know the answers to all of these questions before deciding where to allocate scarce resources. This provides a future research opportunity, but currently limits the prospects for efficient targeting of ‘remedial’ formal products.

In contrast, finding ways for the informal market to work better may have more promise. Mobile payment systems such as M-Pesa in Kenya or Bkash in Bangladesh have revolutionized how individuals transact. One interpretation is that these payment technologies decrease the cost of forming a risk-sharing link, especially for geographically distant individuals. Informal insurance in localized regions such as villages is limited by the correlated nature of incomes. Thus, forming new, distant links might be especially valuable by diversifying the incomes in the insurance pool. [Jack and Suri \(2014\)](#) show that the expansion of M-Pesa in Kenya has improved the ability for network members located in other parts of the country to participate in the insurance network.³⁵ Finding other ways to amplify the insurance value of the network may also have policy promise.

³⁴ Clearly informal risk sharing is not the only reason that many of these attempts have failed. Issues such as trust, basis risk, financial literacy, and behavioural biases might also limit demand for formal insurance.

³⁵ See also [Blumenstock \(2012\)](#), [Gagnon and Goyal \(2017\)](#).

(iii) The interaction between formal credit and informal financial transactions

The models of risk sharing discussed above can help us to think about the impacts of introducing new financial products to an existing risk-sharing network. One implication might simply be a lack of demand for the formal product, if the network is already providing the same service at a lower cost. However, access to savings and credit might also crowd out informal insurance under limited commitment. Two individuals who share risk sustain this by punishing each other in the event that one refuses to help the other in a time of need. But with access to savings or credit, the individual is better equipped to smooth consumption in the absence of informal transfers, and therefore the punishment of withholding future informal transfers is less severe. This can reduce the scope for informal risk-sharing and can possibly even reduce welfare.

This is not simply a theoretical possibility. Recent research (Comola and Prina, 2015; Banerjee et al., 2017, 2018b) that looks at the introduction of financial products such as microcredit has found that a consequence may be the loss of insurance links, particularly among the very poor who may not have access to microcredit loans. Specifically, Banerjee et al. (2018b) look at two settings, 75 villages in Karnataka with full network data and 102 slums in Hyderabad with aggregated relational network data (described in section IV) where microfinance was introduced. In the first setting, the study was observational, while the second involved a randomization. In both cases, they used machine learning to classify which households *ex ante* were likely to take up microfinance using eligibility data along with various demographics. They find that in fact those considered highly unlikely to join microcredit suffer the greatest losses in links and in particular this consists of links lost between groups of households all of whom are unlikely to take up the new credit opportunity.³⁶

The loss of informal links can be devastating for vulnerable populations. Practically speaking, a policy-maker must take into account the externalities on the network when partially completing a market and ensure that the design accounts for such spillovers. Moreover, given that network relationships play many roles, the atrophy of links in one domain might impose costs in others.

The expansion of financial products has both beneficial and negative effects, and so their design requires great care. In the case of the microfinance expansion of Banerjee et al. (2017), access to formal credit for individuals with the desire to expand a pre-existing business actually *crowds in* informal financial transactions. Thus, a more targeted approach to microfinance expansion might generate better results and would also require fewer subsidies.

(iv) Financial contracting in a broader context: case studies

Financial contracting and migration

Development economists are interested in rural to urban migration because the returns to such migration appear to be quite high. However, there seems to be not enough

³⁶ Heß et al. (2018) find related results when they look at the introduction of a community driven development programme in the Gambia.

migration given such returns (Gollin *et al.*, 2013; Bryan *et al.*, 2014; Morten, 2016; Baseler, 2018).

Networks play a major role in the migration decision for at least two reasons. The first is that a rural villager, as discussed above, has a risk-sharing network at home. When he chooses to migrate, he is doing so because he has the potential to increase his income (and potentially draw independent income which is uncorrelated with his parents' income and so serves as insurance in part). Notice then that the extent of insurance at home—the risk-sharing network structure—affects the incentives to migrate to the city (Munshi and Rosenzweig, 2016). Relatedly, when one moves to the city, destination incomes are not known to the entire village, and this introduces a 'hidden income' problem potentially affecting the quality of risk-sharing back home. Thus, local risk-sharing both shapes and is shaped by migration, and thus it is crucial to consider who will migrate and how the gains will be allocated.

Further, policies that *ex ante* may not have seemed directly related to migration can also affect migration incentives and therefore the at-home risk-sharing networks. For instance, a nationwide minimum wage scheme, such as NREGA in India, can be an external source of insurance, thereby changing the valuation of migration in the first place (Morten, 2016).

The second reason networks play a major role in migration comes from information. Motivating this fact is the finding in Bryan *et al.* (2014) that simply buying a bus ticket for villagers led to increased migration and subsequent income thereafter. While this study does not conclude a specific mechanism, subsequent work has looked at the channel of information. Baseler (2018) looks at (i) whether poor beliefs about destination wages depress migration in a meaningful way and (ii) how it could be the case that there are persistent poor beliefs about destination wages in equilibrium. To examine this, in a first experiment, he simply provides information about wages in Nairobi to a random set of individuals and shows that there is a substantial uptick in migration. He then considers the at-home network and uses this as a lens to understand why there can be persistently poor beliefs about destination returns. The idea is that in a risk-sharing network (e.g. even with one's parents as he looks at a remittance relationship here), an individual has an incentive to lie downwards about the extent of the destination income. The only people to whom he would report accurately are those he would want to entice to come, but ideally not those who would essentially pollute the beliefs of his risk-sharing network who may request higher transfers. He finds exactly this behaviour in the data. What this means for a policy-maker is that there are strong incentives for individuals to persistently reduce beliefs about destination returns to migration. Convincing campaigns to relieve these information frictions could have very powerful effects and increase migration and perhaps welfare.

Using network information flow to improve contracts

Papers such as Bloch *et al.* (2008) highlight ways in which the information and financial contracting roles of the network are closely linked. Specifically, the (threat of) transmission of information about deviations from a promised contract aids in contract enforcement. Consider the case of a village wherein a very central individual comes to believe that some individual *A* is deficient in some way (e.g. is not financially responsible). When the central individual gossips about *A*, this travels far and wide, and more

importantly is likely to reach even those with whom A may need to cooperate (e.g. start an entrepreneurial endeavour) and as a consequence A may be more likely to behave in ways that signal that he is of high quality (e.g. responsible). This sort of reputational effect is thought to be the basis of numerous informal institutions.

Breza and Chandrasekhar (2019) study this *reputational effect* in the context of a financial product. Individuals set savings goals over a half-year period and know that they will be randomly assigned to have no monitor, a central monitor (an individual in their village who is notified about their savings account amount every two weeks), or a non-central monitor. The experiment demonstrates that savings increase considerably both when there is monitoring, but more importantly when there is random assignment to highly central monitors. Individuals are cognizant of the reputational effects and further the study shows that there are reputational changes in the minds of randomly chosen villagers even a year after the end of the experiment. The reputational effect is stronger when a central monitor has been randomly assigned.

The results speak to an intuition that has been in the literature for a long time, that reputation drives cooperation in these contexts:

the contributing member may admonish his partner for causing him or her discomfort and material loss. He might also report this behavior to others in the village, thus augmenting the admonishment felt. Such behavior is typical of the close-knit communities in some LDCs. (Besley and Coate, 1995)

Finally, the results have clear policy relevance for the design of financial instruments. Individuals are comfortable in many settings (e.g. in places where they participate in rotating savings credit associations) publicizing their goals and progress towards improving their own financial health. Adding these sorts of self-selected monitoring features allows them to use their reputation as collateral and in fact improve their livelihood as a consequence.

IV. Network data for interventions

As networks are very rich and intricate, an empirical researcher faces a challenge. Carefully studying how networks influence economic environments requires obtaining precise, appropriate data for the context. The direct approach is to elicit the network data at the highest level of detail by enumerating all agents and mapping out all links between them. However, in certain contexts such data may be costly or impractical to obtain. In these cases, policy-makers may need to appeal to other techniques to obtain the information about the network needed for their application.

In this section we describe four strategies that empirical researchers and policy-makers could use to obtain information about the network relevant for their work. First, we discuss how one could map out the entire network within a community. This would provide the researcher with very granular data about the nature of links between units. Second, we touch upon network data available through social media, phone records, and other digital sources. These provide the policy-maker with a pre-made map of interactions, though the interactions mapped here must be carefully interpreted.

Third, we discuss a common case that researchers or policy-makers may face: they do not need to know specifically if certain agents are linked, but rather are interested

in general features of the network (such as whether it exhibits fractionalization) or features of a node (such as how central an agent is). If one is willing to assume a statistical model of network formation, then with very cheap and limited data, one can estimate this model and then get a rough image of it: that is, one can estimate the probability that each agent is linked to each other agent, and this is enough for many practical applications.

Finally, we note that in particular contexts, the policy-maker may be interested in a specific feature of the network that can be elicited in a simple manner. For instance, if we wanted to know a respondent's link count, we could simply ask him to enumerate his friends. This is a trivial example, but we note several more subtle generalizations which may be of policy relevance.

(i) Network maps

One approach to studying networks empirically is to first conduct a full, survey-based elicitation of the relevant graph (see, for example, [Banerjee *et al.* \(2019b\)](#)). Typical economic network mapping requires (i) obtaining a census of all population units in an area, (ii) eliciting the names of all network contacts for each individual, (iii) matching the list of social contacts to the census, and (iv) repeating (i)–(iii) across many networks. This process also necessitates grappling with many important conceptual issues inherent to any network data collection exercise. First, what is the 'right' level of granularity at which to study the network, and what are the network features we should focus on? It might be sufficient to analyse village-level networks when studying the spread of information about local jobs or immunization camps. On the other hand, a typical risk-sharing network is likely much larger in scope, encompassing people who live far away. Second, what constitutes a network link? Network interactions take many forms, including information transfer and advice, socialization, and financial exchange. The most detailed network surveys elicit information on many such categories, but judgement calls remain to be made when constructing a network relevant for some process or behaviour. While isolating the financial component of the network may be relevant for studying risk sharing, the advice-only network may *not* be appropriate for studying information diffusion. After all, individuals may spread information while socializing or attending religious activities together. They might also gossip while repaying an informal loan. Third and relatedly, what does a link encode? In the model of the network-based process, can individuals only interact if they share an edge? Does a network link represent a probability of interaction, so that friends of friends can interact, albeit with lower probability?

Setting aside the conceptual issues that arise in interpreting network data, even if we take for granted that the researcher wants full network data, they are often costly and impractical for policy-makers to collect. We next discuss three strategies for obtaining network data at a lower cost.

(ii) Social media and other digital data

A first naturally occurring and potentially very rich type of network data that can be collected comes from social media. For instance, access to Facebook or Twitter data

gives a policy-maker a very rich description of (online) interaction patterns: essentially the entire network, subject to privacy restrictions.³⁷ Social media are widely used throughout the developing world and so may be an important tool in applying insights from network economics to policy. This is directly useful, for instance, if the policy-maker is attempting to conduct a diffusion campaign, as in the immunization campaign throughout Indonesia described above (Alatas *et al.* 2018).³⁸

A second data source consists of phone records. For instance, policy-makers may have access to (anonymized) call record networks, or simply cell-tower records (Blumenstock, 2012). In contexts with robust mobile money networks, data on financial transfers across space may also help policy-makers identify risk-sharing responses to aggregate shocks (Jack and Suri, 2014; Blumenstock *et al.*, 2016; Björkegren, 2018). Clearly, phone and financial transfer data are likely to be less useful when studying interactions at a more local, geographically proximate level. For instance, phone data do not often provide meaningful information when studying word-of-mouth communication in a village. Moreover, each of these data collection strategies, particularly the phone and financial transaction records, must be used with caution owing to privacy concerns.

(iii) Aggregate images of networks

Though networks are extremely rich objects, which features of a network the policy-maker would like to know will depend on context. Often a particular network theory or application will suggest the importance of specific details such as how fractionalized a network is (which turns out to be important in DeGroot-type network learning processes), or what its connected components are (which we have seen are important for risk sharing). This may be all the information needed. For instance, to assess whether an information scheme is likely to work, the policy-maker may simply need to assess whether the network is connected enough. To assess the effect of affirmative action policies on take-up of subsidies for a particular sub-group, the policy-maker may have to know how likely it is that information in one sub-community travels to another sub-community.

In this way, the practical objective of a policy-maker may be to estimate specific features of the network or understand how they affect some policy outcome. This raises the question of whether a full network map is needed. Could it be possible, instead, to obtain some data that permit estimation of these key features (sufficient statistics identified by a theory) without paying the cost of collecting the full data? We present such an approach below using the methodology of *aggregated relational data*.

Aggregated relational data

One option is to use a technique called aggregated relational data (ARD) (McCormick and Zheng, 2015). ARD consist of responses to questions of the form:

‘Think of all of the households in your village with whom you «INSERT ACTIVITY». How many of these have trait k ?’

³⁷ While the data track online interactions, there is often enough information in order to construct offline social networks as well.

³⁸ Also see Leider *et al.* (2009) and Mobius *et al.* (2015), who conduct experiments about information diffusion and the motives for financial transfers in the Facebook network.

Here the activity represents the relevant notion of a link for the policy or research question at hand (e.g. friendship or risk-sharing). An example of a trait would be ‘whether an adult in the household had typhoid, malaria, or cholera over the past six months’.

ARD, when collected over many activities, provides the policy-maker with enough information to be able to recover a reasonably granular image of the network. In many cases, this might be sufficient for the policy-maker’s purposes.

For every respondent, ARD records the number of links the person has to people of each trait. This helps to locate a person in relation to a sub-group having each trait. Individual *A* may know many households from his friendship list that have experienced cholera, but very few containing a government employee. Individual *B* may know many households who have experienced cholera, as well as many that include government employees. Individual *C* may have many friends who are government employees but no friends that have had cholera. This indicates that, wherever they are in the network, individual *A* is in a region with more connections to cholera-types (call this the left), individual *B* is somewhere in the ‘middle’ (with links to both cholera and government employee types), and individual *C* is somewhere on the ‘right’ (with only links to government employee types).

Breza *et al.* (2019) lay out the following procedure for collecting ARD, using ARD to estimate the parameters of a network formation model, and finally sampling from a distribution over any node- or graph-level statistic of interest to the policy-maker or researcher.³⁹

- (a) **Conduct ARD survey:** For a share of the population (e.g. 30 per cent), conduct a network survey to elicit the list of network links. Ask 5–8 ARD questions about the ARD traits, such as: ‘How many households in your network do you know where an adult had cholera in the past 6 months?’

Let g_{ij} be a binary indicator for whether households *i* and *j* are linked. The ARD response of interest for a household *i* would be:

$$y_{ik} = \sum_j g_{ij} \cdot 1(j \text{ has had cholera in the past 6 months})$$

where trait *k* denotes the cholera question and $g_{ij} = 1$ implies *i* and *j* have a link. Hence y_{ik} sums up all the network members of *i* who had cholera in the past 6 months.

- (b) **Collect census data:** Collect census data with required demographic indicators, denoted X_i (e.g. GPS coordinates, caste/subcaste) and ARD traits (e.g. whether an adult member in the household has cholera).
- (c) **Estimate network formation model with ARD using pre-packaged software:** The software uses the information from the ARD survey and the counts of the people with each ARD trait from the census to estimate parameters of the network formation model. In this model, the probability that two households *i* and *j* have a *link* depends on a household fixed effect v_i (e.g. gregariousness) and distance in some latent space. The main idea is the distance between the locations of the households on the latent space (denoted by z_i and z_j) is inversely proportional to the probability of them being linked.

³⁹ All of the estimation and simulation codes are available on the authors’ webpages.

$P(g_{ij} = 1 | v_i, v_j, \zeta, z_i, z_j) \propto \exp(v_i + v_j - \zeta \cdot \text{distance}(z_i, z_j))$, where $\zeta \geq 0$ is some coefficient relating distance to the odds of linking.

- For the ARD sample, this step produces estimates of the individual traits, v_i, z_i . The software then estimates a model for predicting these individual-specific traits based on X_i .
- For the households not in the ARD sample, the software predicts v_i, z_i (out of sample) using X_i .

With estimated fixed effects and latent locations for all the n households in the network, the probability of any network graph being drawn is computed.

(d) Compute network statistics of interest using estimated probability model: Use the estimates of ζ , fixed effects v_i , and locations in the latent space z_i to compute the expectation of the statistic $S(g)$ (e.g. average path length of the network): $E[S(g) | Y]$ given ARD $Y = \{y_{ik}\}$.

(e) Estimate economic parameters of interest: Let k index each independent network. Given an economic outcome W_k and estimates $E[S(g_k) | Y_k]$ of interest, one can run regressions such as

$$W_k = \alpha + \beta E[S(g_k) | Y_k] + e_k \text{ or } E[S(g_k) | Y_k] = \alpha + \beta \text{Treatment}_k + e_k.$$

If required, one can also do more complex exercises using the estimates of the network statistics.

It is worth noting that, of course, ARD is not a magic solution in order to estimate all features of a network. First, there are some features that simply cannot be estimated using ARD with any meaningful amount of precision. The simplest example is that we can never know whether a link between i and j exists without asking directly. No manner of statistical modelling can give us that beyond a simple probability. Second, the ARD approach uses a statistical model. The quality of the results is only as good as the extent to which the model describes the real-world network data. [Breza et al. \(2019\)](#) document in detail for which network features ARD performs well in real-world network data.

Finally, [Breza et al. \(2019\)](#) show that ARD offers cost savings of the order of 80 per cent relative to collecting full network data. To take just one example, measurements of the effect of financial products on risk sharing networks ([Banerjee et al., 2017](#)), among other things, could have been just as well estimated with ARD rather than the full network data. Therefore, ARD provides a promising avenue for policy-makers to make use of insights from network economics.

(iv) Specific network features

A final and very practical option available for some types of network statistics is to collect data specific to the question that the policy-maker is interested in. Network economics has taught us that, in certain specific cases, concepts can be captured by features or statistics quite easily, and with that in hand a policy-maker can readily make his design decisions. We caution that this is not the norm: for most network attributes there

is no shortcut and the above techniques must be used. In some specific cases, however, clever insights allow us to bypass typical network data collection.

A first example returns to the idea of seeding individuals with information to increase diffusion. As discussed in section II, under the right conditions, identifying network central individuals is important and can have high returns. However, without network data, how might one go about doing this? The answer comes from an observation that under certain assumptions, even if individuals do not know the network structure nominations of somebody as a good gossip end up highly correlated with that person's network centrality. Therefore, simply by asking about who is a good gossip, a policy-maker can identify better individuals with whom to seed information (Banerjee *et al.*, 2019b).

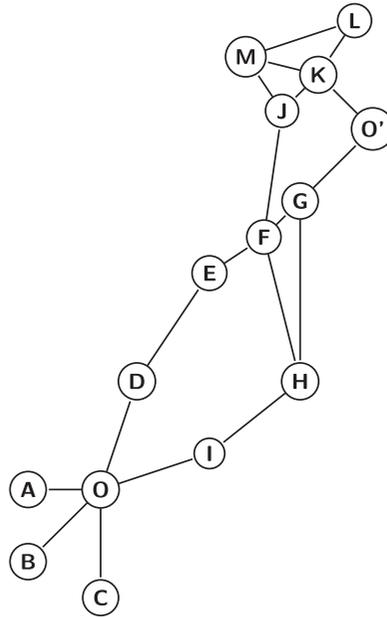
Consider the case (see Figure 7) where an individual H hears about gossip. For simplicity, there are two originators of different pieces of gossip: O and O' . Information spreads through the network and we consider what happens when O' is the originator versus O . H hears from his friends I , F , and G . If O' is very central, then over time H repeatedly hears the gossip through many channels. For instance, H hears from G who heard from O' directly. But also H may hear from F who heard from G who heard from O' . Later on H may hear from I who heard through another path, and so on. And later on F may report the gossip again because he heard from J who heard from M who heard from K who heard from O' , and so on. If O is not very central then there will not be as many such paths for her, so H may only hear O -originated gossip a few times. This means if H ranks these two, O' will be ranked above O , simply because of this phenomenon—just by counting, which is a simple operation.

The practical implication is that one can simply ask a question of the form: 'If one was to try to spread information about X most widely in your village, to whom should the information be given?' and the resulting answer would, under the above logic, be a reasonable approximation of central seed options.

Banerjee *et al.* (2019b) demonstrate in two randomized controlled trials—one a cellphone marketing campaign and the second a large-scale vaccination campaign in hundreds of villages—that there are high returns to seeding so-called gossip nominees. Both take-up of the phone product as well as take-up of vaccines increases dramatically relative to random seeding and even seeding a larger set of non-gossip seeds, as well as relative to seeding village leaders.

A second example can be seen in the case where a policy-maker may want to assess how much a given policy may crowd out the capacity of a network to support certain kinds of cooperative relationships. This ability to sustain cooperation is well captured by *support*, defined in Jackson *et al.* (2012). The idea here is that A and B can sustain a financial link only if there is a C to support the relationship in the sense of having relationships with both A and B . In such a case, poor behaviour by A will be punished by C in addition to B . To compute whether the sample has supported links, the policy-maker needs to simply survey a base sample of N households and then ask for their contacts and then ask how many of those contacts are mutually linked. This is easy enough to do.

A third example relates to social fissures. Policy can affect how individuals interact in potentially socially volatile situations. For instance, policy giving benefits to certain groups (e.g. through affirmative action) may have rather complex effects on the distribution of relationships both within and across social groups (e.g. different ethnicities, races, or castes). The homophily—the relative number of links within versus across

Figure 7: A sample network

groups—is an essential quantity in such a context and is rather trivial to obtain: it suffices to simply sample some individuals and assess the demographic compositions of their friendship networks.

These examples are meant to be illustrative, not exhaustive. After all, there are myriad features the policy-maker could be interested in, given the context. At the same time, the given examples provide a starting point for some policy-relevant elicitations.

V. Discussion

Network economics has provided a useful language to capture why and how some core features of informal economies work in the developing world. Given that the absence of strong institutions is a defining characteristic of developing economies, the relationships between individuals are essential to economic exchange. The models and methods of network economics offer tools to think about information diffusion and aggregation, insurance, migration, and norms, among other important phenomena.

Through the course of this chapter, we have discussed models of diffusion and aggregation of information. We have looked at the trade-offs between different types of information delivery under different models of communication. We have also looked at theoretical work in risk sharing pertaining to (i) how (exogenous) networks can be leveraged for enforcement in the presence of contracting frictions, and (ii) what kind of investments are made into building links that can then be employed to share risk. In addition, we have examined empirical tests of full insurance models in rural settings. We have discussed how while there is considerable informal insurance provided by caste

and family networks, it is not complete. In addition, we have shown that entry of formal institutions into rural and poorer urban areas has an effect on the network structure and even affects those who do not participate in these formal institutions.

These insights can inform policy design. For instance, assessing the nature and incentives of the information to be diffused (e.g. whether it is likely that an individual requires multiple exposures to start using a new technology, whether it is rivalrous) is essential to determine a useful strategy. Finding out whether the information is an unknown unknown or a known unknown will also affect the seeding strategy, as will incentives beyond the information campaign. A similarly broad perspective is useful in assessing financial access expansion schemes—e.g. looking at effects on those who participate, as well as others in the community.

Fundamental research in the development space has benefited from high-quality network data. Access to these network data has allowed the research community to test insights from theory and consider policy implications. The goal from a policy perspective is, insofar as possible, to apply these insights without arduous data requirements. We have offered, in section IV, several options in service of this goal.

References

- Acemoglu, D., Bimpikis, K., and Ozdaglar, A. (2014), 'Dynamics of Information Exchange in Endogenous Social Networks', *Theoretical Economics*, **9**(1), 41–97.
- Agarwal, S., Alok, S., Ghosh, P., Ghosh, S., Piskorski, T., and Seru, A. (2017), 'Banking the Unbanked: What do 255 million New Bank Accounts Reveal about Financial Access?', Columbia Business School Research Paper, 17-12.
- Alatas, V., Chandrasekhar, A., Mobius, M., Olken, B., and Paladines, C. (2018), 'When Celebrities Speak: A Nationwide Twitter Experiment Promoting Vaccination in Indonesia', National Bureau of Economic Research, No. w25589.
- Ali, S. N., and Miller, D. A. (2013), 'Enforcing Cooperation in Networked Societies', Working Paper, available at <http://www-personal.umich.edu/~econdm/resources/AliMiller-Networked.pdf>
- Ambrus, A., and Elliott, M. (2018), 'Investments in Social Ties, Risk Sharing and Inequality', Economic Research Initiatives at Duke (ERID) Working Paper 179.
- Mobius, M., and Szeidl, A. (2014), 'Consumption Risk-sharing in Social Networks', *The American Economic Review*, **104**, 149–82.
- Angelucci, M., and De Giorgi, G. (2009), 'Indirect Effects of an Aid Program: How do Cash Transfers Affect Ineligibles' Consumption?', *American Economic Review*, **99**, 486–508.
- Rasul, I. (2017), 'Consumption and Investment in Resource Pooling Family Networks', *The Economic Journal*, **128**, 2613–51.
- Aral, S., and Walker, D. (2011), 'Creating Social Contagion through Viral Product Design: A Randomized Trial of Peer Influence in Networks', *Management Science*, **57**, 1623–39.
- Arrow, K. J. (1964), 'The Role of Securities in the Optimal Allocation of Risk-bearing', *The Review of Economic Studies*, **31**, 91–6.
- Badarınza, C., Balasubramaniam, V., and Ramadorai, T. (2019), 'The Household Finance Landscape in Emerging Economies', *Annual Review of Financial Economics* available at SSRN 3280962 (2018).
- Banerjee, A. (1992), 'A Simple Model of Herd Behavior', *The Quarterly Journal of Economics*, **107**, 797–817.
- (2008), 'A Bird Flu over the Commie Nest', *Down to Earth: Science and Environment Online*, **16**.
- Dufo, E. (2007), 'The Economic Lives of the Poor', *Journal of Economic Perspectives*, **21**, 141–68.

- Banerjee, A., Duflo, E., and Hornbeck, R. (2014), 'Bundling Health Insurance and Microfinance in India: There Cannot be Adverse Selection if There is no Demand', *American Economic Review*, **104**, 291–7.
- Breza, E., Chandrasekhar, A. G., and Golub, B. (2018a), 'When Less is More: Experimental Evidence on Information Delivery during India's Demonetization', NBER Working Paper No. 24679, available at SSRN 3163930.
- — — Mobius, M. (2019a), 'Naive Learning with Uninformed Agents', NBER Working Paper No. 25497.
- — — Duflo, E., and Kinnan, C. (2017), 'Do Credit Constraints Limit Entrepreneurship? Heterogeneity in the Returns to Microfinance', Buffett Institute Global Poverty Research Lab Working Paper No. 17–104.
- Chandrasekhar, A. G., Duflo, E., and Jackson, M. O. (2019b), 'Using Gossips to Spread Information: Theory and Evidence from Two Randomized Controlled Trials', *Review of Economic Studies*, forthcoming.
- — — — (2013), 'The Diffusion of Microfinance', *Science*, **341**, 1236498.
- — — — (2018b), 'Changes in Social Network Structure in Response to Exposure to Formal Credit Markets', available at SSRN 3245656.
- Baseler, T. (2018), 'Hidden Income and the Perceived Returns to Migration: Experimental Evidence from Kenya', available at https://economics.stanford.edu/sites/g/files/sbiybj9386/f/hidden_income_and_migration.pdf
- Beaman, L., BenYishay, A., Magruder, J., and Mobarak, A. M. (2018), 'Can Network Theory-based Targeting Increase Technology Adoption?', NBER Working Paper No. 24912.
- BenYishay, A., and Mobarak, A. M. (2018), 'Social Learning and Incentives for Experimentation and Communication', *The Review of Economic Studies*, **86**, 976–1009.
- Besley, T., and Coate, S. (1995), 'Group Lending, Repayment Incentives and Social Collateral', *Journal of Development Economics*, **46**, 1–18.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992), 'A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades', *Journal of Political Economy*, **100**, 992–1026.
- Björkegren, D. (2018), 'The Adoption of Network Goods: Evidence from the Spread of Mobile Phones in Rwanda', *The Review of Economic Studies*, **86**, 1033–60.
- Bloch, F., Genicot, G., and Ray, D. (2008), 'Informal Insurance in Social Networks', *Journal of Economic Theory*, **143**, 36–58.
- Blumenstock, J. E. (2012), 'Inferring Patterns of Internal Migration From Mobile Phone Call Records: Evidence from Rwanda', *Information Technology for Development*, **18**, 107–25.
- Eagle, N., and Fafchamps, M. (2016), 'Airtime Transfers and Mobile Communications: Evidence in the Aftermath of Natural Disasters', *Journal of Development Economics*, **120**, 157–81.
- Bollobás, B. (1998), *Modern Graph Theory*, 1st edn, Graduate Texts in Mathematics **184**, New York, Springer-Verlag.
- Bramoullé, Y., and Kranton, R. (2007), 'Risk-sharing Networks', *Journal of Economic Behavior & Organization*, **64**, 275–94.
- Breza, E. (2016), 'Field Experiments, Social Networks, and Development', in *The Oxford Handbook of the Economics of Networks*, Oxford, Oxford University Press.
- Chandrasekhar, A. G. (2019), 'Social Networks, Reputation and Commitment: Evidence from a Savings Monitors Experiment', *Econometrica*, **87**(1), 175–216.
- — McCormick, T. H., and Pan, M. (2019), 'Using Aggregated Relational Data to Feasibly Identify Network Structure Without Network Data', *The American Economic Review*, forthcoming
- Bryan, G., Chowdhury, S., and Mobarak, A. M. (2014), 'Underinvestment in a Profitable Technology: The Case of Seasonal Migration in Bangladesh', *Econometrica*, **82**, 1671–748.
- Casciaro, T. (1998), 'Seeing Things Clearly: Social Structure, Personality, and Accuracy in Social Network Perception', *Social Networks*, **20**, 331–51.
- Chandrasekhar, A. G., Breza, E., and Tahbaz-Salehi, A. (2018), 'Seeing the Forest for the Trees? An Investigation of Network Knowledge', National Bureau of Economic Research Working Paper 24359.

- Chandrasekhar, A. G., Larreguy, H., and Xandri, J. (2019), 'Testing Models of Social Learning on Networks: Evidence from Two Experiments', *Econometrica*, forthcoming.
- Cheng, J., Adamic, L., Dow, P. A., Kleinberg, J. M., and Leskovec, J. (2014), 'Can Cascades be Predicted?', in Proceedings of the 23rd International Conference on Worldwide Web, ACM, 925–36.
- Cochrane, J. H. (1991), 'A Simple Test of Consumption Insurance', *Journal of Political Economy*, **99**, 957–76.
- Cole, S., Giné, X., Tobacman, J., Topalova, P., Townsend, R., and Vickery, J. (2013), 'Barriers to Household Risk Management: Evidence from India', *American Economic Journal: Applied Economics*, **5**, 104–35.
- Collins, D., Morduch, J., Rutherford, S., and Ruthven, O. (2010), *Portfolios of the Poor: How the World's Poor Live on 2 Dollars a Day*, Princeton, NJ, Princeton University Press.
- Comola, M., and Prina, S. (2015), 'Treatment Effect Accounting for Network Changes: Evidence from a Randomized Intervention', available at SSRN 2250748.
- Condorcet, N. C. (1785), *Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix*, Paris, Imprimerie Royale.
- Udry, C. (2010), 'Learning About a New Technology: Pineapple in Ghana', *American Economic Review*, **100**, 35–69.
- Deaton, A. (1997), *The Analysis of Household Surveys: A Microeconomic Approach to Development Policy*, Washington, DC, World Bank Publications.
- Debreu, G. (1959), *Theory of Value: An Axiomatic Analysis of Economic Equilibrium*, New Haven, CT, Yale University Press.
- DeMarzo, P. M., Vayanos, D., and Zwiebel, J. (2003), 'Persuasion Bias, Social Influence, and Unidimensional Opinions', *The Quarterly Journal of Economics*, **118**, 909–68.
- Demirgüç-Kunt, A., Klapper, L., Singer, D., Ansar, S., and Hess, J. (2018), *The Global Findex Database 2017: Measuring Financial Inclusion and the Fintech Revolution*, Washington, DC, The World Bank.
- Diamond, P. A. (1967), 'The Role of a Stock Market in a General Equilibrium Model with Technological Uncertainty', *The American Economic Review*, **57**, 759–76.
- Diestel, R. (2005), *Graph Theory (Graduate Texts in Mathematics)*, New York, Springer.
- Duflo, E., and Saez, E. (2003), 'The Role of Information and Social Interactions in Retirement Plan Decisions: Evidence from a Randomized Experiment', *The Quarterly Journal of Economics*, **118**, 815–42.
- Dupas, P. (2011), 'Do Teenagers Respond to HIV Risk Information? Evidence from a Field Experiment in Kenya', *American Economic Journal: Applied Economics*, **3**, 1–34.
- Durrett, R. (2007), *Random Graph Dynamics*, vol. **200**, Cambridge, Cambridge University Press.
- Easley, D., Kleinberg, J., et al. (2010), *Networks, Crowds, and Markets*, vol. **8**, Cambridge, Cambridge University Press.
- Eyster, E., and Rabin, M. (2014), 'Extensive Imitation is Irrational and Harmful', *The Quarterly Journal of Economics*, **129**, 1861–98.
- Fafchamps, M., and Lund, S. (2003), 'Risk-sharing Networks in Rural Philippines', *Journal of Development Economics*, **71**, 261–87.
- Foster, A., and Rosenzweig, M. (1995), 'Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture', *The Journal of Political Economy*, **103**, 1176–209.
- Gagnon, J., and Goyal, S. (2017), 'Networks, Markets, and Inequality', *American Economic Review*, **107**, 1–30.
- Galeotti, A., Golub, B., and Goyal, S. (2019), 'Targeting Interventions in Networks', available at SSRN: <https://ssrn.com/abstract=3054353>
- Gollin, D., Lagakos, D., and Waugh, M. E. (2013), 'The Agricultural Productivity Gap', *The Quarterly Journal of Economics*, **129**, 939–93.
- Golub, B., and Jackson, M. O. (2010a), 'Naïve Learning in Social Networks and the Wisdom of Crowds', *American Economic Journal: Microeconomics*, **2**, 112–49.
- (2010b), 'Using Selection Bias to Explain the Observed Structure of Internet Diffusions', *Proceedings of the National Academy of Sciences*, **107**, 10833–6.

- Golub, B., and Jackson, M. O. (2012), 'How Homophily Affects the Speed of Learning and Best-response Dynamics', *The Quarterly Journal of Economics*, **127**, 1287–338.
- Sadler, E. (2017), 'Learning in Social Networks', in *The Oxford Handbook of the Economics of Networks*, available at SSRN 2919146.
- Goyal, S. (2012), *Connections: An Introduction to the Economics of Networks*, Princeton, NJ, Princeton University Press.
- Grimm, V., and Mengel, F. (2018), 'Experiments on Belief Formation in Networks', *Journal of the European Economic Association*, doi.10.1093/jeea/jvy038.
- Gulati, A., Sharma, P., Samantara, A., and Terway, P. (2018), *Agriculture Extension System in India*, Indian Council For Research On International Economic Relations (ICRIER).
- Hazła, J., Jadbabaie, A., Mossel, E., and Rahimian, M. A. (2017), 'Bayesian Decision Making in Groups is Hard', arXiv preprint arXiv:1705.04770.
- Heß, S. H., Jaimovich, D., and Schündeln, M. (2018), 'Development Projects And Economic Networks: Lessons from Rural Gambia', available at http://hess.org/cdd_gambia.pdf
- Jack, W., and Suri, T. (2014), 'Risk Sharing and Transactions Costs: Evidence from Kenya's Mobile Money Revolution', *American Economic Review*, **104**, 183–223.
- Jackson, M. O. (2010), *Social and Economic Networks*, Princeton, NJ, Princeton University Press.
- (2019), *The Human Network: How Your Social Position Determines Your Power, Beliefs, and Behaviors*, New York, Pantheon.
- Rodriguez-Barraquer, T., and Tan, X. (2012), 'Social Capital and Social Quilts: Network Patterns of Favor Exchange', *American Economic Review*, **102**, 1857–97.
- Jensen, R. (2010), 'The (Perceived) Returns to Education and the Demand for Schooling', *The Quarterly Journal of Economics*, **125**, 515–48.
- Karlan, D., Mobius, M., Rosenblat, T., and Szeidl, A. (2009), 'Trust and Social Collateral', *The Quarterly Journal of Economics*, **124**, 1307–61.
- Karlin, S. (2014), *A First Course in Stochastic Processes*, Cambridge, MA, Academic Press.
- Kinnan, C. (2019), 'Distinguishing Barriers to Insurance in Thai Villages', available at <http://faculty.wcas.northwestern.edu/~cgk281/BtI.pdf>
- Townsend, R., Samphantharak, K., and Vera-Cossio, D. (2019), 'Insurance and Propagation in Village Networks', available at https://www.pier.or.th/wp-content/uploads/2019/09/pier_dp_115.pdf
- Kocherlakota, N. R. (1996), 'Implications of Efficient Risk Sharing without Commitment', *Review of Economic Studies*, **63**, 595–609.
- Kremer, M., and Miguel, E. (2007), 'The Illusion of Sustainability', *The Quarterly Journal of Economics*, **122**, 1007–65.
- Kumbasar, E., Rommey, A. K., and Batchelder, W. H. (1994), 'Systematic Biases in Social Perception', *American Journal of Sociology*, **100**, 477–505.
- Leider, S., Möbius, M. M., Rosenblat, T., and Do, Q.-A. (2009), 'Directed Altruism and Enforced Reciprocity in Social Networks', *The Quarterly Journal of Economics*, **124**, 1815–51.
- Ligon, E., Thomas, J. P., and Worrall, T. (2002), 'Informal Insurance Arrangements with Limited Commitment: Theory and Evidence from Village Economies', *The Review of Economic Studies*, **69**, 209–44.
- McCormick, T. H., and Zheng, T. (2015), 'Latent Surface Models for Networks Using Aggregated Relational Data', *Journal of the American Statistical Association*, **110**, 1684–95.
- Mace, B. J. (1991), 'Full Insurance in the Presence of Aggregate Uncertainty', *Journal of Political Economy*, **99**, 928–56.
- McPherson, M., and Smith-Lovin, L. (1987), 'Homophily in Voluntary Organizations: Status Distance and the Composition of Face-to-Face Groups', *American Sociological Review*, **52**, 370–9.
- Mazzocco, M., and Saini, S. (2012), 'Testing Efficient Risk Sharing with Heterogeneous Risk Preferences', *American Economic Review*, **102**, 428–68.
- Mobarak, A. M., and Rosenzweig, M. R. (2012), 'Selling Formal Insurance to the Informally Insured', available at <http://faculty.som.yale.edu/mushfiqmobarak/papers/insurance.pdf>
- Mobius, M., Phan, T., and Szeidl, A. (2015), 'Treasure Hunt: Social Learning in the Field', NBER Working Paper No. 21014.

- Morten, M. (2016), 'Temporary Migration and Endogenous Risk Sharing in Village India', *Journal of Political Economy*, **127**(1), 1–46.
- Mossel, E., Sly, A., and Tamuz, O. (2015), 'Strategic Learning and the Topology of Social Networks', *Econometrica*, **83**, 1755–94.
- Munshi, K. (2017), 'Caste and the Indian Economy', available at https://www.histecon.magd.cam.ac.uk/km/Munshi_JEL1.pdf.
- Rosenzweig, M. (2016), 'Networks and Misallocation: Insurance, Migration, and the Rural–Urban Wage Gap', *American Economic Review*, **106**, 46–98.
- Rosenzweig, M. R. (1988), 'Risk, Implicit Contracts and the Family in Rural Areas of Low-income Countries', *The Economic Journal*, **98**, 1148–70.
- Stark, O. (1989), 'Consumption Smoothing, Migration, and Marriage: Evidence from Rural India', *Journal of political Economy*, **97**, 905–26.
- Sadler, E. (2018), 'Diffusion Games', forthcoming in *American Economic Review*, available at SSRN 2624865.
- Samphantharak, K., and Townsend, R. M. (2010), *Households as Corporate Firms: An Analysis of Household Finance Using Integrated Household Surveys and Corporate Financial Accounting*, **46**, Cambridge, Cambridge University Press.
- Schilbach, F. N. (2015), 'Essays in Development and Behavioral Economics', Ph.D. thesis, Harvard.
- Townsend, R. M. (1994), 'Risk and Insurance in Village India', *Econometrica: Journal of the Econometric Society*, 539–91.
- Udry, C. (1994), 'Risk and Insurance in a Rural Credit Market: An Empirical Investigation in Northern Nigeria', *The Review of Economic Studies*, **61**, 495–526.
- Valdés, A., Hazell, P., and Pomareda, C. (1986), *Crop Insurance for Agricultural Development: Issues and Experience*, IICA Biblioteca Venezuela.
- Wilson, R. (1968), 'The Theory of Syndicates', *Econometrica: Journal of the Econometric Society*, **36**, 119–32.