

The Past and Future of Network Analysis in Economics *

Matthew O. Jackson[†]

November, 2014

To appear in *The Oxford Handbook on the Economics of Networks*.

Abstract

This is a short essay offering some perspective on the literature for the *The Oxford Handbook on the Economics of Networks*.

Keywords: Social Networks, Network Games, Graphical Games, Games on Networks, Diffusion, Learning, Network Formation, Community Detection,

JEL Classification Codes: D85, C72, L14, Z13

1 An Explosion of Research on Networks in Economics

Over the past two decades research on networks in economics has grown exponentially - from a handful of papers by the late 1990s to thousands today. This explosive growth is due to a number of factors. First, and foremost, to understand many economic behaviors, from the dynamics of product adoption to financial contagions, it is necessary to account for the patterns of interactions. Failing to include network structure can lead one to a deficient understanding of an observed behavior and can lead to poor policy design. Second, there are increasingly well-understood features of networks (e.g., how densely connected a population is, how segregated it is, among others) that have specific and important implications for economic behaviors. This improved understanding has widened the collection of settings

*Financial support from the NSF under grant SES-1155302, from the AFOSR and DARPA under grant FA9550-12-1-0411, and from ARO MURI under award No. W911NF-12-1-0509 is gratefully acknowledged. I thank the editors, Yann Bramoullé, Andrea Galeotti, Brian W. Rogers, for the opportunity to offer thoughts on the literature. I also thank Matt Elliott, Stephen Nei, and Yiqing Xing for helpful comments on an earlier draft. Remark for working paper version: This is written to appear in a handbook and so I only reference chapters in the book rather than original materials - except for references to a couple of recent essays that articulate specific views of the literature not appearing in the handbook.

[†]Department of Economics, Stanford University, Stanford CA 94305, USA, Santa Fe Institute and CIFAR.
<http://www.stanford.edu/~jacksonm> e-mail: jacksonm@stanford.edu

in which networks have been analyzed in conjunction with economic consequences, as it is increasingly clear how to account for network structure and relate it to behavior. Third, increasingly available data and improvements in computational capabilities are enabling the testing and application of models that could not be analyzed even a few decades ago. Moreover, data on networks of interactions is often available in conjunction with behaviors - which is essential for understanding the economic implications of network structure, and allows us to test theories and evaluate policies, measure social learning, diffusion, and peer effects.

Thus, this handbook thus comes at an opportune time. We have learned an immense amount about social and economic networks and developed important new tools, and the set of applications is rapidly expanding. The literature has grown so much that taking stock of it will help researchers new to the area to access the large toolbox for analyzing networked interactions. In addition, there is much to yet to be studied, and so taking stock of the current knowledge base also highlights the research frontier and illuminates important open problems.

Here, I highlight a few of the areas in which the literature has made substantial progress, discuss what enabled this progress, as well as where current tools are poised to make further contributions, and point out some of the most important open problems.

2 Successes: What Economics has Brought to the Table

Economists' forays into the study of networks over the past couple of decades have helped advance the literature in several ways.

The neoclassical economics paradigm of ('rational') individual choice has driven new modeling of network formation and of behavior on networks based on strategic interactions. While arguably narrow, that paradigm is nevertheless powerful, and it drove us to understand how individual decisions to form relationships impact the networks that emerge [see the chapters by Mauleon and Vannetelbosch; and Pin and Rogers]. In particular, the early theme in the economics literature on networks was the tension between individual incentives to form relationships and the broader societal welfare from the resulting network, as individuals do not internalize the indirect impact that their relationships have on others through a multitude of factors: increased information diffusion, contagion, changes in bargaining strengths, exposure to risks, opportunities, etc. The understanding of network externalities was greatly enhanced by those studies, and an economist's inherent welfarist perspective was essential in asking new questions and generating answers that were absent from the

literature, despite the fact that there was a rich and healthy social networks literature in sociology.

Similarly, this game-theoretic approach founded another wave of the literature systematically exploring how an individual's behavior is driven by that of the individual's neighbors and hence indirectly by others in the network [see the chapters by Bramoullé and Kranton; Blume; Nava; and Zenou]. This game theoretic approach has added substantially to peer effects studies, and the understanding of delinquency, educational achievement, vaccination decisions, product adoption, participation in various programs, and other interdependent behaviors. Viewing such interactions not just as contagion processes, but ones with complementarities and substitution pressures in behaviors (e.g., choosing compatible technologies, benefiting from spillovers in knowledge, etc.) leads to a much richer mosaic of how behavior interacts with network structure.

The interest in understanding how peoples' decisions interact also requires a deep understanding of social learning and diffusion. Thus, there has been renewed interest in social learning and diffusion, and understanding how network structure impacts outcomes, and research in the past decade has substantially advanced that frontier. The advances include new understandings of when processes converge or reach consensus, speeds of convergence, how dynamics relate to homophily and other network features, as well as how dynamics differ according to the complexity of the learning or diffusion process and the network structure [see the chapters by Golub; Kearns; and Lamberson]. Most importantly, this has resulted in a fairly systematic understanding of how dynamic processes are dictated by network structure, and how basic network characteristics such as the degree distribution, homophily, and local patterns such as clustering, determine the outcome of social learning and diffusion processes. Similar understandings have been found in the study of networked markets [see the chapters by Galeotti; and Elliott and Manea]. This knowledge has enabled new empirical analyses and tests of the theory, and new areas of application [see the chapters by Breza; Mobius; and Munshi].

Ultimately, having such a systematic understanding that relates network structure to behavior and welfare is essential to network science having a lasting impact in economics. This is a point that I expand upon elsewhere (Journal of Economic Perspectives, Fall 2014), and also with Brian Rogers and Yves Zenou (in "The Economic Consequences of Social Networks") we provide an up-to-date taxonomy of network features and their economic consequences, as well as discussion of how accounting for networks can be essential for understanding behavior and shaping policies. The longevity of network analysis in economics will derive from seeing significant network effects in a variety of applications.

In addition to bringing the neoclassical economic paradigm to network analyses, working to develop a systematic understanding of how networks impact economic behaviors, and

exploring new applications (more on that below), another important push from the economics literature has been on cleanly identifying network effects and testing many of the new theories. Economists' preoccupation with the inference of causation and getting clean identification of hypothesized effects has increased the use of lab and field experiments in analyzing social and economic networks [e.g., see the chapters by Choi and Kariv, by Kearns, by Aral, and by Breza]. Such techniques are certainly not new to network analysis, but the emphasis has changed. This is in part due to the progress in modeling, which has resulted in numerous new hypotheses to be tested regarding network formation, peer effects, social learning, and diffusion.¹ This is resulting in some important cross-fertilization with anthropology and sociology, which have strong traditions and expertise in field work. This is also pushing statistics and econometrics researchers to produce new techniques for analyzing data and testing network theories (more on this below).

3 A Bucket List

Despite the richness of the literature, there is so much yet to be studied. The frontiers can be thought of as relating to the recent progress in the literature in several ways. First, existing theory and techniques are enabling new applications. Next, recent theory and empirical work have made evident some particular areas where new models and techniques are needed. Finally, there is a need for some systematic meta-analyses of the tools and findings to date. Let me treat each of these in turn.

I begin with application areas that are ripe for further investigation. As mentioned above, the areas in which studies are incorporating networks are growing rapidly. There is a very natural and unusually strong interplay between theory and application in network studies. Applying the theory is helping to identify the impact of network structure on economic behaviors, and as a result is reshaping some policies, as well as testing and refining theory. One important facet of this is that incorporating network structure into the analysis gives us more precision in distinguishing various social forces, since things like diffusion of information, norms, and pressure from peers, all depend in different ways on network structure; and so including networks can help us to distinguish various forms of peer effects - which again can have substantial policy implications. For example, if people forego education because they are learning about its benefits through their social network that suggests a very

¹As a caution, the pressure to have clean identification should not preclude studies that uncover important correlations without establishing causation. Although one obviously has to be careful regarding what to conclude from correlations, the recent trend in natural and field experiments should not completely crowd out observational studies that unearth important relationships in the data that pave the way for further study.

different policy than if they are aware but forego it due to pressures to match the behavior of their friends. For all of these reasons, a key feature of emerging, and future, studies is that they not only involve observations of network structures, but also include behavioral outcomes of the individuals or organizations involved in the networks.

Perhaps the most notable area of such exploration and growth is development economics. This is a very natural area in which to study social networks as most transactions and relations in developing countries are informal: not relying on formal contracts or exogenous institutions, but relying heavily on social interactions [see the chapters by Breza; Mobius; and Munshi]. This arena is important beyond its immediate welfare impact, as not only can network theory help us to understand behaviors and welfare, but small communities in developing countries tend to be relatively closed, thus giving a researcher a holistic view of the patterns of interaction and an unusual degree of control in field experiments. This helps in testing existing theories and generating new models and theories, and is helping test theories of social learning and diffusion. Developing countries also present fairly stark instances in which to study how culture and social norms are shaped by and transmitted through networks - and can help us better understand things like corruption, collective action problems, revolutions, inequality, and growth, all of which have strong network components that are yet to be understood but are well within our current grasp.

Although labor economics is a longer-standing area of network study, it is also an arena in which many fascinating questions remain open [see the chapters by Beaman; Burt; Contractor, and Dessein and Prat]. It has long been clear that networks play important roles in who has access to which opportunities, but richer data sets and new models are opening new questions and avenues for research. For example, how do individuals make decisions on how hard to study, whether to attend university, whether to engage in crime, etc.? Such decisions are shaped by the social setting in which the person is embedded, and heavily influenced by family and friends. We are just beginning to unpack the many different types of interaction that are lumped into the broad category of “peer effects”. Looking carefully at network patterns of interaction should help us to disentangle information spillovers, from norms, and complementarities, from opportunity. This continues to be a promising area for research.

Beyond these two prominent areas of application, we are also seeing increased interest in financial networks and economic fluctuations [see the chapters by Cabrales, Gale, Gottardi; and Acemoglu, Ozdaglar, and Tahbaz-Salehi], as it has become increasingly obvious that a proper understanding of risk and interdependencies in an economy has to account for indirect effects and transmissions of shocks, and these are inherently network phenomena. Here tools are (rapidly) emerging, but the gap between theory and application still needs to be closed. How can one measure whether an institution is not just ‘too big to fail’ but also

‘too connected to fail’? How can we properly measure counter-party risk, accounting for further connections and indirect effects? What are the implications of an increasingly global economy for economic fluctuations within industries and regions? This is an area in which network theory can have immediate as well as a lasting impact, and so it is encouraging to see new work emerging and to see interest from regulators and practitioners in network tools. It is also a particularly interesting area for future development, since the nodes in the network are generally financial institutions whose behavior is sophisticated and often highly strategic, which affects not only investments and production, but the evolution of the network and the resulting transmission of shocks and crises.

There are two other areas in which networks of relationships and externalities are prevalent, and yet are still quite understudied: international trade [see the chapter by Chaney] and international relations [see the chapter by Goyal] and in fact the two areas are intertwined (a point I explore in a recent paper with Stephen Nei). Given the inherent complexity of networks, there is no reason to undertake network analyses unless there are important externalities and interdependencies across relationships. It is thus natural that many economic settings are first studied from a perspective that is either been market-based with many relatively anonymous participants or bilateral with two participants. However, in international trade and international relations, there are fundamental unanswered questions that require modeling beyond a two-at-a-time or a market approach, as the actors involved are inherently networked and face large externalities. In international relations, decisions by countries of which alliances to undertake and which conflicts to enter depend on which allies other countries have. Studying international conflict without a network approach eliminates one’s ability to study more than a third of conflicts, and to understand very basic inter-state history. The same can be said of international trade, where at both a country level and a firm level, terms of trade depend heavily on the opportunities of partners. Each of these applications require developing new models that capture the specific incentives at play, and also obtaining richer data sets that allow us to track networks of trade and alliances over time and see the resulting consequences.

These areas of application are clearly not the only ones in which network analyses can and should be successful over the next years, as we are also continuing to see growing numbers of network studies in political economy, marketing, patenting, among others. Areas like development economics, international trade, and international relations are areas where the study of networks is relatively late and the questions are quite obvious and important, and so these are areas in which network analyses are likely to produce dramatic advances in the short run. In addition, there are over-arching questions about the ultimate impact of technology that continues to make it easier to communicate quickly and cheaply throughout the world [see the chapters by Economides and Watts].

The growing use of networked data in testing theory is also putting new pressure on the development of statistical and econometric models for studying network formation. The stems from the fact that most network studies of networked behaviors are challenged by the endogeneity of the network, which could end up correlating with behavior and unobserved factors that influence behavior [see the chapters by Aral; Boucher and Fortin, and Chandrasekhar]. A major challenge in developing tractable statistical methods for analyzing network formation is that the formation of relationships are generally correlated. That is, the choice of whether to form a relationship between two parties – whether it be friendship, favor exchange, a financial transaction, sharing of risk or information, etc. – is usually substantially influenced by whom else the two parties are connected to. This leads all of the relationships in a network to be inter-dependent. Coupling this with the fact that many data sets consist of observations of a single network means that the data do not consist of many independent observations, but instead one large observation consisting of many dependent objects (e.g., links). Although there are some existing models that admit interdependencies (e.g., exponential random graph models), they suffer from proven computational problems, and models that are both robustly computable and admit link interdependencies are just emerging. This is an area that should experience breakthroughs and enormous advancements over the next decades.

Beyond expanding areas of application and the development of new statistical methods, there are also exciting frontiers that should be explored in network theory over the coming decades. Here is a partial list of important areas in which significant advances are likely in the near future.

First, although we now have models that are helping us to understand diffusion, games on networks, and social learning, these processes are often influenced and manipulated from outside of the networks in ways that are rapidly changing with technologies and of which we have little understanding. For instance, new products have prices and features that are designed to influence their diffusion, and they come together with marketing campaigns that are increasingly taking advantage of social media [see the chapters by Bakstrom, Kleinberg, and Ugander; Bloch; and Mayzlin]. This interaction of media and word of mouth should have important implications for social learning. More generally, there are many contexts in which outside actors attempt to influence networks and the processes operating on them, such as regulators imposing restrictions on which investments financial institutions can make; and there should be some general insights that will prove useful in understanding such interactions. In addition, such processes are also influenced from within, as people may withhold or distort the information that they transmit which can have a profound impact on the diffusion of information and social learning, in ways that are only beginning to be explored [see the chapter by Golub].

Second, although there has been substantial progress in modeling network formation and games on networks in static settings, most interactions are dynamic by nature. Although there are some dynamic models of network formation, and a few dynamic strategic analyses of favor exchange, risk-sharing, and cooperation on networks [e.g., see the chapter by Nava], this is another area where more breakthroughs loom. ‘Cooperative’ behavior can be induced and enforced through potential changes in network structures over time: people abide by certain social norms because their social standing and relationships could deteriorate if they break with the norms.² While the basic concepts are fairly straightforward and we see some insights in analyses of games with various random matching technologies, we still know little about how the evolution of behavior depends on, and influences, social structure, and why some norms are robust and difficult to change and others are fragile. Developing rich dynamic models for these questions can help shed important light on growth and development, as well as persistent inequality.

Related to this point, networks are often non-stationary and highly dynamic entities, and yet standard approaches still involve either static or stationary processes, not only in modeling but also in representations. This has been driven by tractability as well as a lack of a clear picture of the extent to which the nonstationary nature of networks is a major issue. For example, social media data show highly nonstationary patterns of interaction, with some information relayed for long periods of time, and other information only for short periods, and flu contagions depend on non-stationary school and travel patterns, just to mention two obvious examples. Although this is a subject has received some attention, we still know little about the ultimate impact of such nonstationarities. This is fruitful area not only for model development, but also for empirical investigations.

Third, as the set of network tools and models have expanded, we have little understanding of which tools are appropriate for which circumstances. Which are the right measures of homophily and segregation and how does that depend on the particulars of the application? How can we decide which of the many measures of power or centrality is appropriate for a given analysis? Which algorithms for detecting underlying community structures are appropriate as a function of the setting? This complicates empirical work, as trying many different measures to see which ones yield results can lead to spurious findings, and correcting for this requires complicated statistical corrections for running multiple models on the same data - which most researchers ignore.³ Distinguishing models and techniques requires meta-theory

²This becomes important in settings in which repeated interactions between any two given people are insufficient to enforce behavior via simple folk-theorem arguments, and so the broader structure of interactions becomes important.

³This is also an issue if one splits data and runs models on one part and then checks that they still make good predictions out of sample (on the second part of the data) - as even there, considering *multiple* models can end up selecting ones that happen to spuriously provide good fits on the second data set.

that provides us with an understanding of which properties characterize which tools and help us pre-select among theories and form sharper hypotheses for testing. An axiomatic approach may be quite useful to dealing with these issues, as will meta-analyses of empirical projects where similar questions are explored across various case studies.

Fourth, it is clear that networks and behaviors ‘co-evolve’: friendships influence behaviors and behaviors influence who becomes friends with whom. Although there are some models and studies of this phenomenon [see the chapter by Vega-Redondo], it is another area that is not so extensively studied, either empirically or theoretically. This phenomenon also applies beyond social interactions, having important potential implications for things such as investments by banks and other inter-linked financial entities, whose incentives regarding how to invest and whether to monitor those investments are network dependent, and influence the resulting relationships. Such co-dependencies could have far-reaching policy implications.

Fifth, a related but quite distinct observation is that people interact in many different ways at once. One might exchange favors with co-workers, share information with trading partners, and so forth. The layering of different types of relationships among individuals has not gone unnoticed, as there are many studies of multi-relational or multiplexed networks. Nonetheless, we know little about when and why multiple relationships interact with each other and what the broader consequences of interactions between different types of networks might be. The literature to date has been largely observational, and this is an area where there are enormous potential gains from bringing some simple economic modeling to bear.

4 Closing Thoughts

It has become clear that network modeling and analysis in economics is more than a fad. The ubiquity of networked interactions in economic settings means that, if anything, network analyses should continue to grow in economics. The bucket list above provides some areas where the potential gains from new research are self-evident, and undoubtedly, new areas will emerge over time as the theory and data continue to advance. The multitude of important open questions makes this a most fertile area for research, and makes this volume indispensable.

Finally, as I stated at the outset, network science is inherently inter-disciplinary, drawing on tools from a variety of disciplines ranging from mathematics to sociology, and including economics, statistical physics, statistics, and computer science [see the discussion by Kirman]. Its breadth derives from the presence of networks throughout the social and physical world; and from the wide variety of perspectives and techniques that are useful in representing and analyzing networks as well as collecting and analyzing data. Although researchers are much more aware of each other across disciplines than ever before, and tools and techniques

are crossing borders, substantial homophily remains among the research communities. This is partly driven by the historical silos and departments in which sub-communities reside, which leads to incentives and cultures that are heavily influenced by home-disciplines. This leads to prejudices within economics, just as in other disciplines, that are suspicious of research employing paradigms and methods that are unusual within the discipline. It is thus heartening to see that this handbook involves researchers from outside of economics, and that many of the chapters incorporate substantial amounts of material from outside of economics. A hope is that the trend continues, and research in network science eventually becomes seamless across historical silos and disciplines.

Matthew O. Jackson, Stanford, November 2014