Markets and other institutions—both formal and informal—function through interactions among their members. Patterns of interaction constitute networks, networks that govern information flows, facilitate exchanges, and shape people's behavior. In my research, I aim to understand how behaviors and information spread, how network structures influence economic outcomes, and how our understanding of these phenomena can help firms, governments, and other organizations achieve their objectives.

The complexity of network models is hard to overstate. Even if the individuals in a network behave very simply, subtle changes in the way people interact can have a profound effect on outcomes. For instance, in a social learning model, whether people can successfully aggregate their information depends on whether they communicate their beliefs to each other or just observe each others' choices. Likewise in a diffusion model, whether links are fixed over time or are periodically reshuffled affects whether viral cascades are possible. Real people are not so simple. What people know about the network, and what they infer about others' behavior, influences how they behave, which in turn affects the network's role. To answer the most interesting questions about social networks, we must examine how networks influence individual incentives, information, and strategies.

In a pair of papers with Ilan Lobel, I develop a detailed understanding of rational social learning in networks. Both study a sequential model in which each player acts after receiving private information and observing the choices of some who act earlier. In the first paper, now published in *Theoretical Economics*, I explore how correlations between players' neighborhoods affect the learning process. Prior work finds that successful information aggregation turns on network connectivity and properties of players' private signals. In the more general setting of my paper, I show that uncertainty about the network structure can cause learning to fail, even when the network is highly connected, and this mechanism produces more severe inefficiencies than those found in earlier work. A second paper in *Operations Research* extends the model in a different direction, allowing players to have heterogeneous preferences over actions. This allows us to study how the network structure impacts the value of different types of connections and how homophily—the tendency to link with those similar to oneself—affects learning. Homophily improves learning outcomes when individuals are sparsely connected, but it becomes harmful when the network is dense.

The intricacies of networks challenge our intuition, making theoretical work particularly vital to our understanding. In some cases, the most straightforward approach to a problem can be misleading. In my job market paper, I show how common simplifying assumptions in diffusion models drive key findings. Instead of making a mean-field approximation, I develop new mathematical tools to study person-to-person diffusion in a fixed network. This allows us to capture viral phenomena and reverses standard comparative statics results. Moreover, by modeling non-myopic players, I uncover a novel strategic effect: when large adoption cascades are possible, exposure to a contagion conveys information about a player's network position. When adopting a behavior entails positive externalities, this effect helps players coordinate their decisions. As a follow up to this project, I am working on a model of the market for "fake news." Much information consumption now occurs through social media, and the nuanced features of diffusion in this environment may have important policy implications.

Another project, joint work with Elliot Lipnowski, reveals a subtle relationship between network structure and coordination, and it shows that network models can yield new insights in epistemic game theory. We introduce a solution concept, peer-confirming equilibrium, in which a network describes players' strategic knowledge: players best-respond to conjectures about others' behavior, conjectures about neighbors are correct, and these two facts are common knowledge. Central players can sometimes help us get close to equilibrium play, but this requires that a central player's best response is highly sensitive to others' choices. In a simple protest game, we show why leaders of different groups might successfully coordinate with each other, even as they miscoordinate with most of the population. In a dynamic game, we show that peer-confirming equilibrium can refine the predictions of both subgame perfect Nash equilibrium and extensive form rationalizability. Due to a novel form of forward induction reasoning, a player can use her action to signal information about others' strategies

Going forward, I am excited about two distinct research programs. In our recent survey for the *Oxford Handbook of the Economics of Networks*, Ben Golub and I point to the need for better models of persistent disagreement. Even when much information is available and widely shared, we often observe strong disagreement on matters of fact—consider the role of greenhouse gases in climate change, or the (lack of) relationship between vaccines and autism. Economic models of social learning have trouble generating disagreement. From Bayesian approaches to ones using heuristic update rules, consensus is a robust outcome in a variety of settings. To productively comment on what network and information structures best promote information aggregation and consensus, we first need a model that can generate realistic disagreement.

To that end, I am working on a model of belief updating when agents can encounter false information. An agent sequentially hears propositions, which may be false, and must choose which to believe. I ask whether a "good" update rule exists and how false information affects social learning. Taking an axiomatic approach to the first question, I find that "willingness-to-learn" axioms are incompatible with "non-manipulability" axioms. In a social learning model, long-run disagreement is generic. Moreover, we can characterize the influence of each individual agent on steady state beliefs and analyze the effect of interventions on learning outcomes.

In a second ongoing project, joint with Ben Golub, I seek to understand how the relationships people choose to form affect spillovers from those relationships. Consider peer effects in education. A large body of research shows that having high achieving classmates helps students elevate their own performance. Turning this insight into policy prescriptions has proved less than straightforward. As shown in a recent study at the U.S. Air Force Academy (Carrell et al., ECMA, 2013), manipulating peer groups to achieve a desired outcome can backfire. Putting people in the same class is no guarantee they will become friends, and we may only get beneficial peer effects if students form friendships. To fill a gap in the peer effects literature, we are developing a model in which individuals exert two types of effort, "productive effort" and "social effort." They choose which relationships to build, and these ties lead to complementarities in achievement. Our model offers a tractable way to estimate peer effects and design interventions, accounting for endogenous network formation.

As this research matures, I plan to study how group composition interacts with performance incentives. Team-based work depends not just on individual productivity but the ability to leverage complementary skills, and members of a team must invest in relationships with one another to do this. Certain types of individuals have trouble working together, but sometimes there are large benefits if they do so successfully—for instance, cross-disciplinary work has led to many important innovations. How should we build teams? How should we reward individual members when we only observe group-level output? I hope that my work in the coming years can improve our answers to these questions.