My research interest is in coordination of strategic agents under uncertainty, especially as it applies to how markets work: how buyers and sellers find each other, decide whether to trade, and determine a price.

My job market paper looks at a simple model of a dynamic market with fixed preferences and finds that, even if supply is constant, shortages may develop. I construct and investigate a model in which a (partially) nonperishable good with a sticky price can be stockpiled, and agents who fear a shortage may decide to buy more than they wish to immediately consume. If the good has a low income-elasticity of consumption, those fears can become self-fulfilling; such shortages fizzle out quickly for perishable commodities, but may last longer for nonperishable items. For goods that have high income-elasticity, self-fulfilling "runs" on the commodity won't occur with common knowledge of rationality and common knowledge that there is no supply disruption. This ties in with empirical results from a 2014 paper in which the authors study shortages after earthquakes in Japan and Chile, finding that (after an actual disruption to supply) shortages are most common for "emergency" goods and are less likely and shorter-lived for perishable goods. If prices are sufficiently responsive to shortages or stockpiles, the rationing outcome is again self-defeating; agents who know that, if a run took place, they would still be able to buy (albeit at a higher price) in the next period, have less incentive to run. This hints at a family of policy responses.

Other papers I have worked on look more specifically at coordination in more abstract settings. I performed a human subject experiment in which subjects were randomly and anonymously paired with partners to play a stochastic game, in which the game proceeds through discrete-time stages. The game that they played has two states; one state generally has higher payoffs than the other state. In either state, the players would play a 2×2

game in which one action gave a higher stage payoff and the other action resulted in a higher probability that the next stage would be in the higher-payoff state. These options, accordingly, presented subjects with a choice between what could be framed as a safer, more immediate payoff versus a riskier, deferred payoff. In all cases, the players would end up with the same payoff; there was no ex-post conflict in the players' interests. However, if the players had different risk preferences, there might be an ex ante conflict. If the players made different choices, they got both the worse immediate payoff and the worse transition probability; accordingly, they had a motive to coordinate. When the players had similar levels of risk-aversion, this posed little difficulty; in every stage, regardless of state, the two players played the same action. In some pairings, however, the players chose different actions in the early stages; in almost all cases, by the fifth stage in which players were in the same state, they would have settled on an action that they would both play in any subsequent occurence of that state. More often than not in these cases, they settled on the safe choice; in a few cases in which subjects had coordinated on the risky choice, they shifted to the safe choice after some unfavorable realizations of the random transition. Even though subjects were informed of the transition probabilities as a function of state and action profile, the experienced realizations of those transitions played a significant role in their behavior.

My other paper in progress reviews characterizations of "risk dominance" to select between equilibria, and extends the ideas to non-equilibrium situations. One equilibrium is said to "risk-dominate" another if agents unsure which is likely to obtain view it as "safer". This concept has been around for a long time, and was formalized in a narrow category of games in 1988. It can be categorized in a wide variety of ways that coincide for this narrow

category, but diverge for more general games. Futhermore, the logic of risk-dominance can be applied to situations in which players fail to coordinate altogether — that is to say, in a non-equilibrium context. I gather together a number of such characterizations, add some of my own, and motivate their extension to non-equilibrium behavior in finite two-player games of complete information. In particular, I introduce an information-theory based measure of the extent to which an action is the best-response to an agent who is best-responding to almost-correct beliefs. A (strict) Nash equilibrium action may be fragile to second-order beliefs that aren't close enough, but an action that is not part of a Nash equilibrium may be sustained by some second-order beliefs that are almost right.

In the near future, I would like to investigate other information-related modes of market failure, such as what Roth calls "congestion", which I understand to be a situation in heterogeneous markets in which there are so many potential buyers and sellers that agents on one side or the other of the market feel "rushed", and have to agree to transactions before they have gathered information related to their decision. I would also like to build a framework for analyzing measures of financial liquidity; there is a paper by Chordia, Huh, and Subrahmanyam (2009) that provides inspiration for this.