DISSERTATION PROPOSAL

Arnav Sood

"Essays on Learning in Economics"

Tuesday, December 10, 2024
1:00pm
Tepper 4242

Chapter 1: Frequentist Persuasion (joint with James Best)

People don't always know what other people mean. And even when we do, it's not always clear that one should believe what they are saying. Past experiences determine how we interpret communication.

We study a one-shot sender-receiver game where sender's signal is private. Receiver instead has a finite sample of (state, message) pairs, which she uses to estimate it.

The large-sample limit of frequentist persuasion is Bayesian persuasion. But because sender conditions on the sample size, more data is not always better — there will be a finite K where sender provides the most informative signal and maximizes receiver welfare. Conversely, even when preferences are perfectly aligned, full honesty is not always optimal for finite K.

Chapter 2: Dynamic Experimentation and Reputation

We build a model of social learning with myopic and heterogeneous experimenters who are concerned about their reputations. (The focus is on variation in search costs, but other forms — such as the probability of successful search — are also considered.) Experimentation is both risky (in that it could lead to bad outcomes) and costly, but it is also the only way for agents to demonstrate their types.

We plan to characterize the optimal dynamic allocation of searchers, and to show that there generically exist PBEs where reputation motives can overturn the stagnation results in similar models (Callendar 2011, Garfargini and Strulovici 2016.) Currently, we are solving the allocation problem for simpler cases (e.g., with uncorrelated alternatives, with finite but expanding consideration sets.)

Chapter 3: Contracting with Unobserved State Changes: The Case of Quiet Quitting

Employers in the COVID era were forced to confront "quiet quitters" — workers who *silently* convert from rational types who work if incentivized, to behavioral types who always shirk. We consider a principal-agent model where the agent exogenously changes type, and principal must account for this in designing the optimal contract.

First we characterize the Bayes-optimal contract in the case where the quitting distribution is common-knowledge and correctly specified. We then compare this to non-Bayesian sequential hypothesis tests, and study how these simpler heuristics hold up in more complicated cases (e.g., various forms of misspecification by the principal on the change distribution, strategic behavior by the agent on how and when to change type.)

Proposed Committee: Ali Shourideh (Chair), James Best, Maryam Saeedi, and Aislinn Bohren (Penn, External Reader)

Proposal Documents: Dropbox with PDF abstracts can be found here