

Agents' Mental Models

Three essays investigating the construction and implications of economic agents' internal representations of problems they face.

Chapter 1: Causal Coherence

Agents with the same information and same preferences can make different choices. Agents differ not only with respect to their preferences and information, but their causal interpretations of that information. This can lead to what agents with the correct causal model would perceive as "irrational mistakes" committed by others; even though the agents with the wrong causal model have rational preferences and are well-informed.

I introduce the *causally coherent* agent, who has a causal model about a causally ambiguous phenomenon that is consistent with data, makes choices rationally, and is unaware of alternative explanations for these data. In essence, her model is not identified so she hazards a guess. I use Causal Bayesian Networks to represent her internal model of the phenomena she encounters, so she is like agents in artificial intelligence who use such a framework.

In this framework, I show how agents with the same information and the same preferences will make different choices. Moreover, with this framework, I can construct a set of reasonable theories that emerge from data the agents see. This provides a framework for constructing agents' conjectures in an arbitrary setting. I apply this framework to an auction to show that agents with wrong models suffer a 'causal curse' similar in kind to the winner's curse.

Chapter 2: Causal Coherence and Non-Public Social Beliefs

This essay investigates how the market learns by drawing an analogy between a population distribution of agents, who believe different causal models, and a belief distribution over models. The winner's curse due to incorrect causal models leads to losses, so agents with incorrect causal models will be driven from the market. Compare this to a single, well-informed observer who has access to all models and sees all market transactions. How does the market population distribution track the observer's belief distribution? This paper investigates the circumstances under which they will be the same, and what must be true of the market to be as logically rational as the well-informed observer.

Chapter 3: Optimal Auctions with Ambiguity

With Subir Bose and Emre Ozdenoren.

A crucial assumption in the optimal auction literature is that each bidder's valuation is known to be drawn from a unique distribution. In this paper we study the optimal auction problem allowing for ambiguity about the distribution of valuations. Agents may be ambiguity averse (modeled using the maxmin expected utility model of Gilboa and Schmeidler (1989).) When the bidders face more ambiguity than the seller we show that (i) an auction that provides full insurance to all types of bidders is always in the set of optimal auctions, and in certain cases the seller can strictly increase his revenue by switching to a full insurance mechanism; (ii) if the seller is ambiguity neutral and any prior that is close enough to the seller's prior is included in the bidders' set of priors then the optimal auction must be a full insurance auction; (iii) in general, neither the first nor the second price auction is optimal (even with suitably chosen reserve prices). When the seller is ambiguity averse and the bidders are ambiguity neutral an auction that fully insures the seller must in the set of optimal mechanisms.