Application: NBER Digital Economics and AI Tutorial

Spring 2027, Stanford University

Contact information

Màxim Ventura-Bolet
PhD Student, Department of Economics, Columbia University
Email: mv3010@columbia.edu

Graduate study

I am a third-year PhD student in Economics at Columbia University with interests in game theory and the economics of AI. I spent the Summer of 2026 as a research intern in the Economics and Computation group at Microsoft Research (working with Nicole Immorlica and Brendan Lucier). I do not have a dissertation title nor committee at this stage of my PhD but you can contact Navin Kartik or Joseph Stiglitz for recommendation letters.

Research projects

Published Work

The Impact of AI and Digital Platforms on the Information Ecosystem (with Joseph E. Stiglitz). Published in *The Economics of Transformative AI*, NBER / University of Chicago Press, 2026, pp. 367–389. We show how AI and digital platforms can weaken incentives to gather and process new information, the very inputs AI depends on, so that the information ecosystem may deteriorate, especially in the short to medium run. A companion piece on policy responses appears in *The Digitalist Papers, Vol. 2* (Stanford Digital Economy Lab, 2025).

Working Papers

Ask a Buddy or a Bot? Unbundling Answers from Relationships (with Nicole Immorlica, Brendan Lucier, and Joseph E. Stiglitz). *Submitted to the NBER Digital Economics and AI Meeting, Spring 2027.* People routinely turn to other people to answer questions they cannot answer themselves. These consultations deliver answers, but they also build social capital, which raises the quality of future human answers. We develop a dynamic model in which each question can be asked to a human, asked to AI, or left unanswered, and in which the value of asking a human increases with a common stock of social capital that depreciates over time and is replenished by human consultations. AI generates immediate welfare gains: it improves answers to some questions and makes previously unanswered questions worth asking. When it displaces human consultations, however, social capital erodes monotonically and flow welfare declines toward the highest steady state below the pre-AI level, which may involve the collapse of human consultations. The total number of questions asked rises on impact but can eventually fall below its pre-AI level. If long-run welfare falls below the pre-AI benchmark, AI reduces lifetime welfare for sufficiently patient

societies. Because individuals do not internalize their contribution to social capital, decentralized users undertake fewer human consultations than a social planner.

Belief Protocols for AI Agents (with Nicole Immorlica and Brendan Lucier; in progress). A developer selling AI agents must choose how those agents form beliefs about strategic opponents. When the protocol is common knowledge, it affects both opponents' strategies and customers' decisions to deploy. We study rules mapping an agent's observable environment (the continuation game, available history, and its own activation) into beliefs. With one AI delegate, we compare two protocols. *Deployment maximization* chooses the belief that maximizes the probability of being deployed, jointly maximizing the probability that a deployment opportunity arises and the probability of activation given that opportunity. *Rationalizability maximization* chooses the belief supporting equilibria that rationalize activation for the largest set of the principal's possible outside options, a forward-induction argument (that is, maximizes probability of deployment once an opportunity arises). The protocols coincide in a static benchmark and can differ in finite and infinite dynamic games, although equivalence holds under additional conditions.

Other working papers: *The Indirect Cost Function of LLMs: A Microfoundation* (with Cole Wittbrodt), which models the cost to a user of getting an LLM to carry out a desired search, and *Straightening Single Crossing* (with Navin Kartik and Daniel Rappoport).