

# Benjamin S. Manning

---

|   |   |   |      |
|---|---|---|------|
| CONTACT                                   | bmanning (at) mit.edu ◊ <a href="http://benjaminmanning.io">benjaminmanning.io</a>  |   |      |
| EDUCATION                                 | <b>Massachusetts Institute of Technology</b>  |   |      |
|   | Ph.D. Management Science, Information Technology  | 2022 - Present  |      |
|   | S.M., Management Research   | 2024  |      |
|   | <b>Harvard University</b>   |   |      |
| M.P.P. Economic & Social Policy           | 2021  |   |      |
| <b>Washington University in St. Louis</b> |   |   |      |
|   | B.A. Mathematics (College Honors)   | 2017  |      |
| REFERENCES                                | John Horton (Chair)<br>MIT Sloan School of Management<br>jjhorton (at) mit.edu  | Christopher Norio Avery<br>Harvard Kennedy School<br>christopher_avery (at) hks.harvard.edu |      |
|   | Colin F. Camerer<br>California Institute of Technology<br>camerer (at) caltech.edu  | Siddharth Suri<br>Microsoft Research<br>siddharth.suri (at) microsoft.com                   |      |
|   | GRANTS<br>(\$483,325)   | Schmidt Sciences AI at Work Award (\$193,325; PI)   | 2025 |
|   |   | Emergent Ventures Fast Grant (\$50,000; PI)   | 2025 |
|   | MIT GenAI Impact Consortium & OpenAI (\$180,000; Co-PI w/ John Horton)  | 2025  |      |
|   | Schmidt Sciences Grant Proposal Award (\$10,000; PI)  | 2025  |      |
|   | Emergent Ventures Fast Grant (\$50,000; PI)   | 2024  |      |
| AWARDS &<br>SCHOLARSHIPS                  | WISE Best Paper Award   | 2025  |      |
|   | Harvard Distinction in Student Teaching Award   | 2025  |      |
|   | Becker Friedman Institute Price Theory Summer Camp  | 2025  |      |
|   | Kosciuszko Foundation Tuition Scholarship (\$6,500)   | 2019 & 2022   |      |
| WORKING<br>PAPERS                         | <b>General Social Agents</b><br>(Benjamin S. Manning <sup>1</sup> & John Horton)<br><i>Under Review</i>   |   |      |
|   | <ul style="list-style-type: none"><li>• WISE 2025 Best Paper Award</li><li>• Media: <a href="#">Marginal Revolution</a></li></ul>   |   |      |
|   | <b>Abstract:</b> Useful social science theories predict behavior across settings. However, applying a theory to make predictions in new settings is challenging; rarely can it be done without ad hoc modifications to account for setting-specific factors. We argue that AI agents put in simulations of those novel settings offer an alternative for applying theory, requiring minimal or no modifications. We present an approach for building such “general” agents that use theory-grounded natural language instructions, existing empirical data, and knowledge acquired by the underlying AI during training. To demonstrate the approach in settings where no data from that data-generating process exists—as is often the case in applied prediction problems—we design a heterogeneous population of 883,320 novel games. AI agents are constructed using human data from a small set of conceptually related but structurally distinct “seed” games. In preregistered experiments, on average, agents predict initial human play in a random sample of 1,500 games from the population better than (i) a cognitive hierarchy model, (ii) game-theoretic equilibria, and (iii) out-of-the-box agents. For a small set of separate novel games, these simulations predict responses from a new sample of human subjects <i>better</i> even than the most plausibly relevant published human data. |   |      |
|   | <b>Automated Social Science: Language Models as Scientist and Subjects</b><br>(Benjamin S. Manning <sup>1</sup> , Kehang Zhu <sup>1</sup> , & John Horton)<br>Reject & Resubmit at <i>The Quarterly Journal of Economics</i>  |   |      |
|   | <ul style="list-style-type: none"><li>• Media: <a href="#">Marginal Revolution</a></li></ul>  |   |      |
|   | <b>Abstract:</b> We present an approach for automatically generating and testing, <i>in silico</i> , social scientific hypotheses. This automation is made possible by recent advances in large language models (LLM), but the key feature of the approach is the use of structural causal models. Structural causal models provide a language to state hypotheses,   |   |      |

---

<sup>1</sup>Indicates first author or co-first authors

a blueprint for constructing LLM-based agents, an experimental design, and a plan for data analysis. The fitted structural causal model becomes an object available for prediction or the planning of follow-on experiments. We demonstrate the approach with several scenarios: a negotiation, a bail hearing, a job interview, and an auction. In each case, causal relationships are both proposed and tested by the system, finding evidence for some and not others. We provide evidence that the insights from these simulations of social interactions are not available to the LLM purely through direct elicitation. When given its proposed structural causal model for each scenario, the LLM is good at predicting the signs of estimated effects, but it cannot reliably predict the magnitudes of those estimates. In the auction experiment, the *in silico* simulation results closely match the predictions of auction theory, but elicited predictions of the clearing prices from the LLM are inaccurate. However, the LLM's predictions are dramatically improved if the model can condition on the fitted structural causal model. In short, the LLM knows more than it can (immediately) tell.

### Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Silicus?

(John Horton<sup>1</sup>, Apostolos Filippas<sup>1</sup>, Benjamin S. Manning<sup>1</sup>)

Revise & Resubmit at *The Review of Economics and Statistics*

- Extended abstract at *ACM Conference on Economics & Computation (EC '24)*

**Abstract:** We argue that newly-developed large language models (LLMs), because of how they are trained and designed, are implicit computational models of humans—a *Homo silicus*. LLMs can be used like economists use *Homo economicus*: they can be given endowments, information, preferences, and so on, and then their behavior can be explored in scenarios via simulation. Experiments using this approach, derived from Charness and Rabin (2002), Kahneman et al. (1986), Samuelson and Zeckhauser (1988), Oprea (2024b), and Horton (2025), show qualitatively similar results to the original, and when they differ, it is often generative for future research. We discuss potential applications, conceptual issues, and why this approach can inform the study of humans.

### Prompt Adaptation as a Dynamic Complement in Generative AI Systems

(Eaman Jahani<sup>1</sup>, Benjamin S. Manning, Joe Zhang, Hong-Yi TuYe, Mohammed Alsobay, Christos Nicolaides, Siddharth Suri, & David Holtz)

Revise & Resubmit at *Information Systems Research*

- Media: [Marginal Revolution](#), [MIT Ideas Made to Matter](#)

**Abstract:** As generative AI systems rapidly improve, a key question emerges: how do users adapt to these changes, and when does such adaptation matter for realizing performance gains? Drawing on theories of dynamic capabilities and IT complements, we study *prompt adaptation*—how users adjust their inputs in response to evolving model behavior—using a common experimental design applied to two preregistered tasks with 3,750 total participants who submitted nearly 37,000 prompts. We show that the importance of prompt adaptation depends critically on task structure. In a task with fixed evaluation criteria and an unambiguous goal, user prompt adaptation accounts for roughly half of the performance gains from a model upgrade. In contrast, in an open-ended creative task where the space of acceptable outputs is effectively unbounded and quality is subjective, performance improvements are driven primarily by model capability; prompt adaptation plays a limited role. We further show that automated prompt rewriting cannot generally substitute for human adaptation: when aligned with task objectives, it can modestly improve performance, but when misaligned, it can actively undermine the gains from model improvements. Together, these findings position prompt adaptation as a dynamic complement whose importance depends on task structure and system design, and suggest that without it, a substantial share of the economic value created by advances in generative models may go unrealized.

PEER-  
REVIEWED  
PUBLICATIONS

### National Megastudy Shows that Email Nudges to Elementary School Teachers Boost Student Math Achievement, Particularly When Personalized

(Angela Duckworth<sup>1</sup>, Katherine Milkman, ... Benjamin S. Manning, ..., & 26 others)

*Proceedings of the National Academy of Sciences*, 2025

### Effect Size Magnification: No Variable is as Important as the One You're Thinking About—While You're Thinking About It

(Linnea Gandhi<sup>1</sup>, Benjamin S. Manning, & Angela Duckworth)

*Current Directions in Psychological Science*, 2024

- Top 10 most downloaded Association for Psychological Science (APS) journal articles in 2024
- Media: [Under the Cortex](#)

PUBLISHED  
CONFERENCE  
PROCEEDINGS

### Strategic Tradeoffs Between Humans and AI in Multi-Agent Bargaining

(Crystal Qian<sup>1</sup>, Kehang Zhu<sup>1</sup>, John Horton, Benjamin S. Manning, Vivian Tsai, James Wexler, & Nithum Thain )

*ACM Conference on Intelligent User Interfaces (IUI)*, 2026

**The Coasean Singularity? Demand, Supply, and Market Design with AI Agents**  
(Peyman Shahidi<sup>1</sup>, Gili Rusak<sup>1</sup>, **Benjamin S. Manning**<sup>1</sup>, Andrey Fradkin<sup>1</sup>, & John Horton<sup>1</sup>)  
Forthcoming in *The Economics of Transformative AI*

- Media: **Marginal Revolution**

**Abstract:** AI agents—autonomous systems that perceive, reason, and act on behalf of human principals—are poised to transform digital markets by dramatically reducing transaction costs. This chapter evaluates the economic implications of this transition, adopting a consumer-oriented view of agents as market participants that can search, negotiate, and transact directly. From the demand side, agent adoption reflects derived demand: users trade off decision quality against effort reduction, with outcomes mediated by agent capability and task context. On the supply side, firms will design, integrate, and monetize agents, with outcomes hinging on whether agents operate within or across platforms. At the market level, agents create efficiency gains from lower search, communication, and contracting costs, but also introduce frictions such as congestion and price obfuscation. By lowering the costs of preference elicitation, contract enforcement, and identity verification, agents expand the feasible set of market designs but also raise novel regulatory challenges. While the net welfare effects remain an empirical question, the rapid onset of AI-mediated transactions presents a unique opportunity for economic research to inform real-world policy and market design.

**Taste-as-Program: Scalable Preference Elicitation with Artificial Intelligence**  
(Gili Rusak<sup>1</sup>, **Benjamin S. Manning**<sup>1</sup>, & John Horton)

**Abstract:** Many economic mechanisms require detailed information about individuals' preferences. Preference elicitation, however, is costly: individuals may under-report when evaluating many alternatives. We document this pattern in an online labor market experiment, showing that workers' self-reported rankings over job opportunities fail to capture true preferences as the option set grows. While ranking costs grow with the number of jobs, the costs to those workers of expressing *how* they choose jobs—what we call “tastes”—do not. Taste descriptions, which specify what to prioritize, which tradeoffs matter, and so on, can function as “programs” for generating preferences over arbitrarily many options. We show that an AI can execute these programs at scale. The same workers write descriptions of their tastes, and a language model ranks the same job opportunities on their behalf. These AI-inferred preferences remain accurate as choice sets grow, including for jobs workers have never seen. Indeed, we estimate that AI is more accurate than the workers when ranking as few as 17 jobs. In market simulations, participants prefer matches from mechanisms facilitated by AI-inferred preferences over those using their own rankings. These findings suggest that AI can enable market designs in settings where direct preference elicitation is prohibitively costly.

**Recapitulating all of Experimental Economics with AI Agents**  
(**Benjamin S. Manning**<sup>1</sup> & John Horton<sup>1</sup>)

**What Happens When Your News Feed Listens? LLMs, Stated Preferences, and Engagement**  
(**Benjamin S. Manning**<sup>1</sup> & Alex Moehring<sup>1</sup>)

**A Data-Driven Approach to Designing Better Nudges: A Case Study in Math Education**  
(Marcos Gallo<sup>1</sup>, Sean S. Hu<sup>1</sup>, **Benjamin S. Manning**, Angela Duckworth, & Colin Camerer)

**Will The Real Psychohistory Please Stand Up?**  
(**Benjamin S. Manning**<sup>1</sup> & John Horton)

**2026:** MIT Center for Information Systems Research (CISR) Hot Topic Webinar; The Business Implications of Generative AI (BIG.AI); Digital Economy Workshop Athens; Chapman University Brownbag Seminar; AEA Annual Meeting;

**2025:** Workshop on Information Systems and Economics (WISE); Conference on AI, Machine Learning, and Business Analytics; George Mason University Center for Public Choice Seminar<sup>3</sup>; Dartmouth Economics Seminar; North American Meeting of the Economic Science Association (ESA); Instacart Economics Team Seminar; NYU Workshop on AI and the Future of Collaborative Innovation; KOF ETH Zurich and IZA Workshop: Matching Workers and Jobs Online; Machine Learning in Economics Summer Conference (MLESC25); Academy of Management (AOM) Annual Meeting; International Conference on Computational Social Science (IC2S2); ZEW Conference on the Economics of ICT<sup>3</sup>; Wharton AI and the Future of Work Conference; AI, Mechanism Design and Human Behavior: Experiments and Theory NBER/CEME Decentralization Conference; Khipu Latin American Meeting in AI; Artificially Intelligent Social Science Workshop (Oxford); AEA Annual Meeting;

**2024:** Workshop on Information Systems and Economics (WISE)<sup>3</sup>; Conference on AI, Machine

---

<sup>2</sup>Includes scheduled talks

<sup>3</sup>Co-author presenting joint work.

Learning, and Business Analytics; NABE Tech Economics Conference; Informs Annual Meeting<sup>3</sup>; Machine Learning in Science Conference; Econometric Society Interdisciplinary Frontiers Conference on Economics and AI+ML; International Conference on Computational Social Science (IC2S2)<sup>3</sup>; NBER Summer Institute—Digital Economics and AI<sup>3</sup>; ZEW Conference on the Economics of ICT; International Conference of the French Association of Experimental Economics; Instacart Economics Team Seminar; Statistical Conference in E-Commerce Research (SCECR); Wharton AI and the Future of Work Conference; International Meeting on Experimental and Behavioral Social Sciences; MIT IDE Annual Conference; Microsoft Research AI, Cognition, and the Economy Workshop; Measuring Development: AI, the Next Generation at the World Bank;

**2023:** Interactive Causal Learning Conference; MIT CODE; MIT CSAIL’s FutureTech Seminar; Talking to Machines AI Workshop at Oxford University

## TEACHING

*15.575*: Economics of Information and Information Technology, **TA for John Horton & Alessandro Acquisti**—MIT, Spring 2026

*QSTBA 830*: Business Experimentation and Causal Methods, **Guest Lecturer for Andrey Fradkin**—Boston University, Spring 2025

*BUSN 38802*: Managerial Decision Making; **TA for Richard Thaler**, University of Chicago—Booth London Campus, Fall 2024 (Rating: 4.7/5.0)

*API-101Z*: Markets and Failures; **TA for Christopher Avery**—Harvard University, Fall 2024 (Rating: 4.6/5.0)

*API-101Z*: Markets and Failures; **TA for Christopher Avery**—Harvard University, Fall 2023 (Rating: 4.3/5.0)

## AFFILIATIONS

MIT Initiative on the Digital Economy, Heterodox Academy

## PROFESSIONAL SERVICE

**Ad-hoc Reviewer:** *Management Science*, *Proceedings of the National Academy of Sciences*, *Southern Economic Journal*, *Conference on Neural Information Processing Systems (NeurIPS)*, *ACM-Collective Intelligence*, *International Conference on Information Systems (ICIS)*, *Sociological Science*, *International Conference on Computational Social Science (IC2S2)*

**Program Committees:** *ACM Conference on Economics and Computation (2026)*, *The Business Implications of Generative AI (BIG.AI) (2026)*, *Conference on Digital Experimentation (CODE) (2023)*

## EMPLOYMENT

**University of Pennsylvania** 2021 - 2022  
Research Assistant for Angela Duckworth

**University of Chicago** 2020 - 2022  
Research Assistant for Jon Rogowski

**Dartmouth College** 2018 - 2019  
Research Assistant for Sydney Finkelstein

**Phetpittayakom School** 2017 - 2018  
High School Math Teacher in Thailand

**Optiver US LLC** 2016  
Execution Trading Intern