Jimmy Lin (MIT Economics)
jimmylin@mit.edu

I am a joint Economics and Statistics PhD student at MIT. My research interests are in public and behavioral economics, with a focus on (1) using machine learning (ML) as a new tool to study human behavior and (2) understanding how markets interact with algorithms and how they should be regulated. I am excited to apply for the Digital Economics and AI Tutorial, where I would be eager to learn from and discuss research ideas with leading researchers and fellow graduate students in these fields.

My research is advised by Sendhil Mullainathan, Ashesh Rambachan, Amy Finkelstein, and Jesse Shapiro. I describe my two ongoing research projects and my academic background below.

RESEARCH PROJECTS

"Cognitive Mistakes in Repeated Human Decisions: An Algorithmic Assessment of Government Case Workers" (with Steven Ma)

In many high-stakes economic contexts, experts repeatedly make similar but independent decisions based on probabilistic judgments. Doctors diagnose many patients, loan officers review many loan applications, judges rule on many cases, and government administrators interview many welfare applicants. In a rational model, if cases are randomly ordered and independent, the expert decisions should be serially independent. However, positive and negative autocorrelations have been documented in empirical studies of loan officers, asylum interviewers, baseball umpires, doctors, and hiring managers (Chen et al. 2016, Jin et al. 2023, Radbruch & Schiprowski 2024), suggesting that past cases often bias current decisions.

What cognitive biases explain these behavioral phenomena in sequential decision-making, and how do they impact decision quality and welfare? Many theories have been posed, including ex-post rationalization, anchoring, sequential contrast effects, the gambler's fallacy, or implicit quotas; Bordalo et al. (2023) tries to unify these theories in a model of selective memory recall. The empirical challenge for testing these models is that there is a large space of potential mechanisms for selective recall: memories that are recalled could depend on various case characteristics, how long ago they occurred, how quickly cases were handled in succession, or any combination of the above. Rather than individually specifying the vast number of behavioral hypotheses, our approach is to train a black-box algorithm on sequences of human decisions to then identify, interpret, and validate the most explanatory behavioral theories.

We demonstrate this approach in the context of Child Protective Services (CPS), where hotline workers receive referrals for suspected child maltreatment and decide whether to send an investigator to the household (i.e. "screening in" the referral) or not. CPS investigations are pervasive—37% of *all* American children experience a CPS investigation before age 18 (Kim et al. 2017)—and exhibit large racial disparities (Baron et al. 2024). False positives are costly because investigations are invasive to families and expensive for CPS workers; false negatives threaten the child's safety and wellbeing. As such, understanding the reasons for screening mistakes is important not just to refine behavioral models of memory but also to inform potential interventions (e.g. reassigning the order of referrals, developing algorithmic decision aides).

Using rich administrative data on over 100,000 CPS referrals and screening decisions in Allegheny County, we first show that these behavioral effects are large. We are training a supervised learning

Jimmy Lin (MIT Economics)
jimmylin@mit.edu

algorithm (gradient-boosted trees) to predict human screening decisions from the characteristics of the current referral and screener fixed effects. We are also training a transformer model (a custom variant of GPT) to predict human decisions from the full sequence of past case characteristics and decisions. We hope to show that the latter model is significantly more predictive of human decisions, indicating that past cases influence current decisions. Second, we will use interpretability methods (e.g. Ludwig & Mullainathan 2024, Templeton et al. 2024) to describe the behavioral mechanisms implied by the fitted transformer model. We will causally test these theories out of sample and compare the results to a calibrated version of Bordalo et al. (2023). Third, we will quantify how these behavioral biases impact welfare (through decision mistakes) and racial disparities (through inaccurate statistical discrimination).

We hope these findings contribute to the growing literatures of (1) machine learning for scientific discovery (e.g. Ludwig & Mullainathan 2024, Mullainathan & Rambachan 2024), (2) cognitive models of memory (e.g. Bordalo et al. 2023, Enke 2024), and (3) racial discrimination in expert decisions (e.g. Arnold et al. 2022, Baron et al. 2024).

"Tenant Screening Algorithms, Landlord Risk, and Housing Inequality"

Private market transactions often require assessments of risk: landlords evaluate a tenant's risk of late payments or property damage, employers evaluate a candidate's risk of being underqualified, and developers evaluate the risks of natural disasters when deciding where to build. Algorithmic risk scores have been developed and widely adopted in such settings to help inform market decisions by estimating risk from historical data. How do these risk scores affect market transactions in equilibrium, and how should they be regulated?

In the housing market, most landlords use tenant scoring algorithms when considering rental applications. These algorithms are offered by dozens of loosely regulated companies that combine credit reports, eviction filings, and criminal records into an opaque tenant score. Proponents of these scoring tools argue that they could reduce financial risks for landlords (particularly for small landlords or landlords in low-income neighborhoods), decrease security deposits and evictions, and help low-income or minority households that have low risk but would typically be discriminated against through stereotypes. Critics of these scoring tools argue that they could systematically bar opportunities for households with criminal or eviction histories, implicitly learn to discriminate against protected classes like race or Housing Choice Voucher (HCV) status, and increase neighborhood segregation. For example, SafeRent settled a class-action lawsuit from HCV holders alleging discrimination in Massachusetts. They agreed to not display tenant scores for low-income tenants, which could potentially lead to a ban-the-box-style effect that hurts low-income tenants (Agan and Starr 2017).

In this early-stage project, I plan to investigate these claims using data on algorithm adoption and household location outcomes. I leverage the staggered entry of algorithmic scoring companies into housing markets: scoring companies typically entered markets when local eviction court proceedings and criminal records became digitized, which offers a plausibly exogenous instrument for algorithm adoption (c.f. the digitization of house transaction records in Raymond 2024). Data on landlord adoption of tenant scores come from thousands of individual lawsuits against these scoring companies and FOIA data from local public housing authorities (Kirchner 2020), and household characteristics and moving patterns come from the Census. I will estimate how the adoption of

Jimmy Lin (MIT Economics)
jimmylin@mit.edu

tenant scoring algorithms affect rent prices, eviction rates, neighborhood segregation, and location outcomes of low-income, racial minority, or HCV-holding households.

COURSEWORK, TEACHING & HONORS

Coursework. I am currently in the second year of the PhD program, but I have already completed my core and field class requirements (in part by taking PhD core classes as a Harvard undergraduate). In addition to the standard core sequence in microeconomics, macroeconomics, and econometrics, I have taken the following graduate-level field classes:

- **Public** (3 classes). Redistribution and social insurance (Amy Finkelstein and Nathan Hendren), Optimal tax theory and empirics (Jim Poterba and Iván Werning), Public finance and behavioral public (Raj Chetty and Hunt Allcott).
- **Industrial organization** (2 classes). IO theory (Mike Whinston and Tobias Salz), Structural methods and advanced topics (Nikhil Agarwal and Tobias Salz).
- **Theory** (2 classes). Market design (Scott Kominers), Auction theory (Shengwu Li).
- **Statistics** (3 classes). Probability theory (Joseph Blitzstein), Individual risk (Joseph Blitzstein and Cynthia Dwork), Nonlinear econometric methods (Whitney Newey and Alberto Abadie)
- **Machine learning** (2 classes). Machine learning (Konstantinos Daskalakis and Tommi Jaakkola), Algorithms & behavioral science (Sendhil Mullainathan and Ashesh Rambachan).
- **Other**. Labor (Abi Adams and Daron Acemoglu), Advanced research and communication (Amy Finkelstein and Abhijit Banerjee).

Teaching. I was a teaching assistant for the following undergraduate classes at Harvard:

- **Intermediate microeconomics** (Econ 1011a; Edward Glaeser). 3 years as a TA, including 1 as head TA. Co-wrote the course textbook.
- **Economics and computation** (CS 136; David Parkes). Won the Derek Bok Center Distinction for Teaching. Received all 5/5 teaching ratings.
- **Mathematics in the world** (Math 157; Joseph Harris).

I am also co-writing a textbook for Sendhil and Ashesh's graduate class (Algorithms and Behavioral Science), and I taught competitive programming at the AlphaStar Academy.

Honors. I am graciously supported by the Paul and Daisy Soros Fellowship for New Americans, the NSF Graduate Research Fellowship, and the MIT Economics Departmental Fellowship. I was awarded the Thomas Hoopes Prize for Outstanding Research for my undergraduate thesis, "Labor Supply During the COVID Pandemic: The Who, What, Where, and Why of the Great Resignation," which was advised by Edward Glaeser.

Conclusion. At this tutorial, I am eager to learn more about the interaction between algorithms and markets, discuss uses of machine learning to study traditional and new economic questions, and meet faculty and fellow graduate students excited about these fields.

Jimmy Lin (MIT Economics)
jimmylin@mit.edu

References

Agan, Amanda, and Sonja Starr. 2018. "Ban the Box, Criminal Records, and Racial Discrimination: A Field Experiment*." *The Quarterly Journal of Economics* 133 (1): 191–235.

Arnold, David, Will Dobbie, and Peter Hull. 2022. "Measuring Racial Discrimination in Bail Decisions." *American Economic Review* 112 (9): 2992–3038.

Baron, E Jason, Joseph J Doyle Jr, Natalia Emanuel, Peter Hull, and Joseph Ryan. 2024. "Discrimination in Multiphase Systems: Evidence from Child Protection*." *The Quarterly Journal of Economics* 139 (3): 1611–64.

Bordalo, Pedro, John J Conlon, Nicola Gennaioli, Spencer Y Kwon, and Andrei Shleifer. 2023. "Memory and Probability*." *The Quarterly Journal of Economics* 138 (1): 265–311.

Chen, Daniel L., Tobias J. Moskowitz, and Kelly Shue. 2016. "Decision Making Under the Gambler's Fallacy: Evidence from Asylum Judges, Loan Officers, and Baseball Umpires *." *The Quarterly Journal of Economics* 131 (3): 1181–1242.

Enke, Benjamin. 2024. "The Cognitive Turn in Behavioral Economics." *Working Paper*.

Jin, Lawrence, Rui Tang, Han Ye, Junjian Yi, and Songfa Zhong. 2023. "Path Dependency in Physician Decisions." *The Review of Economic Studies*.

Kirchner, Lauren. 2020. "How We Investigated the Tenant Screening Industry – The Markup." May 28, 2020. https://themarkup.org/show-your-work/2020/05/28/how-we-investigated-the-tenant-screening-industry.

Ludwig, Jens, and Sendhil Mullainathan. 2024. "Machine Learning as a Tool for Hypothesis Generation*." *The Quarterly Journal of Economics* 139 (2): 751–827.

Mullainathan, Sendhil, and Ashesh Rambachan. 2024. "From Predictive Algorithms to Automatic Generation of Anomalies." *National Bureau of Economic Research Working Paper Series*.

Radbruch, Jonas, and Amelie Schiprowski. 2024. "Interview Sequences and the Formation of Subjective Assessments." *The Review of Economic Studies*.

Raymond, Lindsey. n.d. "The Market Effects of Algorithms." *Working Paper*.