Doctoral Student Workshop on Economics of Artificial Intelligence

Contact Information

Lindsey Raymond
lraymond@mit.edu
978-621-7554 (cell)
100 Main Street, E62-489, Cambridge, MA 02142

Course of Study
PhD Candidate
Technological Innovation, Entrepreneurship and Strategic Management Group (TIES) at MIT Sloan
Years of enrollment: 2017 – 2022 (expected)
Dissertation Committee: Scott Stern (chair), Erik Brynjolfsson and Danielle Li
Major Coursework: Economics PhD sequences in Labor, Industrial Organization

Minor coursework: Computer Science PhD sequence in Machine Learning and Statistical Learning Theory

General Research Interests
My research agenda focuses on the economic consequences of AI and ML on workers and firms and how incentives to adopt and develop these technologies affect their impact. In one branch of work, I explore the impact of these technologies for workers and firms. In the second, I explore incentives to produce, adopt and diffuse this type of technology. I am particularly interested in understanding the implications of distinct algorithm design choices and why these choices may have different economic effects. My work relies heavily on a detailed understanding of machine learning and analysis of textual data, and therefore machine learning based methods are also an integral component of my research. My interest in this area extends to policy - I will be spending the next academic year as a Staff Economist at the Council of Economic Advisors, and as part of that role, plan to be working on understanding the existing and potential impacts of AI on the American economy.

Current Projects

- **Ex Ante Worker Productivity and Artificial Intelligence (with Erik Brynjolfsson)**

The potential for technological change to raise demand for higher skill workers, or those who are more educated, able or experienced, occupies a central role in understanding increases in wage inequality and employment trends over the past fifty years. By increasing the potential set of tasks that can be done by capital instead of labor, there has been concern and some early evidence that AI might lead to greater automation and wage polarization. While much of the existing empirical and theoretical work focuses on potential for AI to automate tasks formerly performed by labor, we document that AI, because it is very good at learning optimal ways to perform tasks, has the potential to affect labor productivity by efficiently coaching and training workers and thus disproportionately benefiting the ex-ante less productive workers.

We conceptualize worker productivity is a learning problem - in our context, call centers, team managers spend their time reading customer/agent transcripts, giving feedback to the workers and identifying and disseminating best practices. We study the introduction of an AI that provides real time coaching to call center agents, by suggesting what to say during conversations with customers and providing explanations of why to say it. Introduction of the AI disproportionately benefits lower and middle skill workers, because the AI software picks up best practices from the most productive workers, who are training examples, and disseminates them to the ex-ante less productive ones. While this reduces the value to having a talented team manager, because the AI also provides more real-time feedback to the manager on each agent's performance, this also affects attrition rates, allowing managers to quickly identify low performing workers who do not improve from on-the-job training. We also study how these results vary by a worker's contracting relationship to the firm (outsourced or internal to the firm) and location (US based or international). Finally, we focus on the value of discretion, or allowing the humans to overrule the AI suggestions, to better understand the tradeoffs between full automation and the value of an AI decision support tool plus a human worker.

- **Hiring as Exploration (with Danielle Li and Peter Bergman)**, NBER Working Paper 27736

This paper views hiring as a contextual bandit problem: to find the best workers over time, firms must balance "exploitation" (selecting from groups with proven track records) with "exploration" (selecting from under-represented groups to learn about quality). Yet modern hiring algorithms, based on "supervised learning" approaches, are designed solely for exploitation. Instead, we build a resume screening algorithm that values exploration by evaluating candidates according to their statistical upside potential. Using data from professional services recruiting within a Fortune 500 firm, we show that this approach improves the quality (as measured by eventual hiring rates) of candidates selected for an interview, while also increasing demographic diversity, relative to the firm's existing practices. The same is not true for traditional supervised learning based algorithms, which improve hiring rates but select far fewer Black and Hispanic applicants. In an extension, we show that exploration-based algorithms are also able to learn more effectively about simulated changes in applicant hiring potential over time. Together, our results highlight the importance of incorporating exploration in developing decision-making algorithms that are potentially both more efficient and equitable.