Contact Information:
David Almog, 5th Year MECS PhD student at Kellogg, Northwestern University.
david.almog@kellogg.northwestern.edu
https://www.david-almog.com/

I am a 5th year Ph.D. student in the Managerial Economics and Strategy program at the Northwestern University, Kellogg School of Management. Because of my strong interest in behavioral phenomena and their implications, I also enrolled as an exchange student at University of Chicago Booth School of Business to take behavioral classes. During one of my classes at Booth, Professor Mullainathan introduced me to some of the latest research on machine learning and how this approach can outperform humans in relevant contexts. While I was thrilled to see the progress made in machine learning for prediction problems, I was struck by the dearth of evidence on how humans respond when algorithms are incorporated into the decision-making process. Ultimately, understanding how individuals react will determine when and to what degree human decisions should be replaced by machines ones.

Currently, I am working on two papers on human-AI interaction, the beginning of a research agenda on *AI Oversight*. In the most advanced project, a joint working paper with Daniel Martin, Romain Gauriot and Lionel Page, we study the implications of human behavior when AI is granted final decision rights. We do so by leveraging the introduction of Hawk-Eye in professional tennis. In this paper we find that umpires lowered their overall mistake rate after the introduction of Hawk-Eye review, in line with rational inattention given psychological costs of being overruled by AI (which overpowers the incentives to free ride on the new technology). We also find that umpires increased the rate at which they called balls in, which produced a shift from making Type II errors (calling a ball out when in) to Type I errors (calling a ball in when out). This paper shows how the particular way we introduce AI into the decision loop can have meaningful impact on the final outcome.

In more preliminary joint work with Daniel Martin, as AI oversight is a growing part of workplace culture we study whether worker performance under AI oversight is fundamentally different from worker performance under human oversight. In this paper, we study the impact of AI oversight on job performance using an online experiment in which subjects were paid to provide captions for images. We find that subjects perform better at this task when they know that they will be evaluated by humans, both in terms of human-measured performance and AI-measured performance. Additionally, under human oversight, participants exhibit distinct writing behaviors: they paste more text, write shorter captions, and use the delete key more frequently.

For my Job Market paper, I am studying the impact of labor market perceptions on AI utilization. While the project is still in its early stages, I am making steady progress and expect to be well-prepared by February. Attending the NBER Digital Economics and AI Tutorial will provide valuable feedback from experts, helping to refine my paper as I prepare for the job market. On the next page, I list my committee members and provide a brief description of my project.

Dissertation Information

Title: AI Signaling in the Labor Market.

Committee members: Yuval Salant, Alex Imas, Jörg Spenkuch, Daniel Martin, and Alvaro Sandroni.

Short description:

Recent evidence shows that workers override AI recommendations at higher rates than they should. Examples of AI under-utilization have been documented in hiring and promotion decisions (Hoffman, Kahn, and Li 2018), medical diagnoses (Agarwal, Moehring, Rajpurkar, and Salz 2023), and judge pretrial bail detention decisions (Angelova, Dobbie, and Yang 2023). While most of these studies focus on utilization through the lens of difficulties in adopting new information sources and behavioral biases, such as overconfidence, this study examines a less-explored mechanism: the concern for how AI use is perceived by others.

In this study, I aim to investigate whether third-party beliefs about AI adoption can lead to underutilization of the technology. The experiment will involve a tournament structure where participants compete to be selected from a group. There will be two very similar tournaments every participant will be part of. I will examine whether when participants are in the block where the selector is aware of their AI utilization rates are more likely to underuse AI recommendations due to concerns about negatively affecting their selection chances. As a follow-up, if concerns about others' perceptions do indeed reduce AI utilization, I am interested in understanding whether this impacts performance, which I hypothesize would decline.

This experiment offers the controlled environment that is difficult to achieve in the field, enabling us to vary the observability of AI utilization across blocks. If our hypothesis is correct, and workers do indeed care about how they are perceived when their AI usage is visible, we should observe a reduction in AI utilization in the block where utilization summary statistics (i.e., average follow-up recommendation rates) are provided to selectors.