

This PDF is a selection from a published volume from the National Bureau of Economic Research

Volume Title: The Economics of Artificial Intelligence: Health Care Challenges

Volume Authors/Editors: Ajay Agrawal, Joshua Gans, Avi Goldfarb, and Catherine Tucker, editors

Volume Publisher: University of Chicago Press

Volume ISBNs: 978-0-226-83311-8 (cloth); 978-0-226-83312-5 (electronic)

Volume URL: <https://www.nber.org/books-and-chapters/economics-artificial-intelligence-health-care-challenges>

Conference Date: September 22–23, 2022

Publication Date: March 2024

Chapter Title: Comment on Chapters 1 and 3: Artificial Intelligence and Decision Making in Health Care: Prediction or Preferences?

Chapter Author(s): M. Kate Bundorf, Maria Polyakova

Chapter URL: <https://www.nber.org/books-and-chapters/economics-artificial-intelligence-health-care-challenges/comment-chapters-1-and-3-artificial-intelligence-and-decision-making-health-care-prediction-or>

Chapter pages in book: p. 144 – 147

which directions? Under which conditions will lower-cost, more accessible, AI-powered services prosper and when will the human touch present in traditional healthcare prevail? The answer to these questions is not clear *ex ante*. In the context of mental health (a context where privacy and human interaction are key), AI-powered technology platforms for providing mental health services have been highly successful.<sup>1</sup>

Overall, the outlook for research at the intersection of AI and healthcare is bright, and there are several exciting areas for researchers to explore. I encourage myself and other researchers to be thoughtful about their pursuits in this space and look forward to research that helps healthcare integrate AI in ways that maximize welfare across the board.

## Reference

Gaynor, M. 2007. "Competition and Quality in Health Care Markets." *Foundations and Trends® in Microeconomics* 2 (6): 441–508.

## Comment on Chapters 1 and 3

M. Kate Bundorf  
and Maria Polyakova

### **Artificial Intelligence and Decision Making in Health Care: *Prediction or Preferences?***

People make most of their decisions in the context of uncertainty. Because each of our decisions can have many possible outcomes, we can never know with certainty how the choices we make today will affect us tomorrow. In the canonical economic model, people make decisions under uncertainty by assessing both the likelihood of potential outcomes associated with alternative courses of action and their utility of each outcome and then

M. Kate Bundorf is the J. Alexander McMahon Distinguished Professor of Health Policy and Management and a professor in the Sanford School of Public Policy and the Department of Population Health Sciences at Duke University, a core faculty member at the Duke-Margolis Center for Health Policy, and a research associate of the National Bureau of Economic Research.

Maria Polyakova is an assistant professor in the Department of Health Policy at Stanford University School of Medicine, and a faculty research fellow of the National Bureau of Economic Research.

For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see <https://www.nber.org/books-and-chapters/economics-artificial-intelligence-health-care-challenges/artificial-intelligence-and-decision-making-healthcare-prediction-or-preferences>.

1. See <https://www.brookings.edu/techstream/the-wellness-industrys-risky-embrace-of-ai-driven-mental-health-care/>.

choosing the alternative with the greatest expected utility (Morgenstern and Von Neumann 1953; Savage 1954). In this model, people need two types of information to make a choice among alternatives: the distribution of future outcomes associated with each alternative, and how they value each of the possible future outcomes. Here we refer to the first as “prediction” and the second as “preferences.” This distinction is consistent with the theoretical model developed by Agrawal, Gans, and Goldfarb (2018a), which differentiates between prediction and “judgment” in understanding the impact of artificial intelligence on decision making.

The distinction between prediction and preferences is central for understanding how artificial intelligence will impact health care treatment decisions. Uncertainty is at the heart of clinical decision making; both diagnosis and treatment are probabilistic. A well-working AI tool can be enormously beneficial for understanding the distribution of potential diagnoses and treatment outcomes. As Obermeyer and Mullainathan discuss in this volume, there are important challenges in generating predictions and deploying them in clinical settings including incomplete or biased data, challenges of communicating algorithmic predictions to human experts, and barriers to implementation. Overall, however, prediction is essential for clinical medicine and prediction technologies can reduce the cost or increase the quality of this important input into decision making (Agrawal et al. 2018b).

The ability of artificial intelligence to incorporate patient preference into decision making, however, is less straightforward. Clinical decision making requires interpreting information about probabilistic outcomes for patients. In many cases, this requires making trade-offs among uncertain, even if well-predicted outcomes. Patient preferences are critical for making these trade-offs. A challenge, which we explored empirically in the context of expert advice for the choice of health insurance plans (Bundorf, Polyakova, and Tai-Seale 2019), is that patients may not have full information about their own preferences. Analogously, in many clinical settings, patients may not have well-formed preferences over treatment outcomes. Indeed a key component of decision aids for health care treatment decisions is helping patients understand “either explicitly or implicitly, the value they place on potential benefits and harms” (Stacey et al. 2017). Physicians and other types of clinicians, due to their experience and expertise, naturally play a role in helping patients formulate their preferences. Can AI do this instead?

We propose that existing AI approaches to advising consumers based on their preferences are less well suited to health care decision making than in other contexts. A common data-driven approach for capturing preferences is to link the decisions of millions of consumers with objective measures of choice satisfaction to predict whether someone with your characteristics is likely to be satisfied with a particular option. This “consumers-like-you” approach, however, has limitations for decision making in health care. It

assumes consumers are making good decisions, when we have lots of evidence that they aren't (Bernheim, Della Vigna, and Laibson 2019). Indeed, if consumers "like you" made great decisions, then they wouldn't need anyone's advice.

An alternative approach is to provide machine-based expert advice. This type of advice, however, imposes expert preferences on individual choices—in essence, assuming away preference heterogeneity across patients, the type of information we wanted to incorporate in the first place. Whether such computerized experts can replace today's physicians depends on how much we think physicians are able to tailor their expert advice to each patient. Physicians are under increasing pressure to provide that type of personalized advice due to a movement toward "shared decision making," in which patients are full partners with clinicians in health care decision making as opposed to more passive recipients of expert advice (Resnicow et al. 2022). The objective of shared decision making is to ensure that treatment decisions more fully reflect patient values. Patients may prefer tailored even if imperfect advice to a recommendation that is only right for an average patient.

As long as AI remains limited in its ability to "predict" patient preferences, it will remain a complement rather than a substitute for many physicians. As Dranove and Garthwaite in this volume recognize, medical specialties that focus less on patient relationships may be more substitutable by AI. Indeed, as AI tools become better at predicting the probabilities of different potential outcomes, the role of physicians is likely to shift more to the domain of helping patients formulate their preferences in response to AI-generated information.

To summarize, what are the implications of the distinction between prediction and preferences for the role of AI in health care decision making? AI's current state as a data prediction exercise may limit its ability to inform decisions that are highly dependent on preferences or decisions for which patient preferences vary substantially. Clinicians who help patients translate prediction into decisions by incorporating patient preferences will have skills that are complementary to the strengths of AI. Paradoxically, this implies that in contrast to the fear of AI replacing the medical profession promulgated by some, increased integration of AI into a clinician's daily routine may not replace physicians but rather incentivize physicians to focus on what medical students often say motivated them to choose medicine—listening to the patient.

## References

- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018a. "Prediction, Judgment, and Complexity: A Theory of Decision-Making and Artificial Intelligence." In *The Economics of Artificial Intelligence: An Agenda*, edited by Ajay Agrawal, Joshua Gans, and Avi Goldfarb, 89–110. Chicago: University of Chicago Press.

- Agrawal, Ajay, Joshua Gans, and Avi Goldfarb. 2018b. *Prediction Machines: The Simple Economics of Artificial Intelligence*. Boston: Harvard Business Press.
- Bernheim, B. Douglas, Stefano DellaVigna, and David Laibson. 2019. *Handbook of Behavioral Economics-Foundations and Applications 2*. Amsterdam: North-Holland.
- Bundorf, M. Kate, Maria Polyakova, and Ming Tai-Seale. 2019. "How Do Humans Interact with Algorithms? Experimental Evidence from Health Insurance." NBER Working Paper No. 25976. Cambridge, MA: National Bureau of Economic Research.
- Morgenstern, Oskar, and John Von Neumann. 1953. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.
- Resnicow, Ken, Delwyn Catley, Kathy Goggin, Sarah Hawley, and Geoffrey C. Williams. 2022. "Shared Decision Making in Health Care: Theoretical Perspectives for Why It Works and for Whom." *Medical Decision Making* 42 (6): 755–64.
- Savage, Leonard J. 1954. *The Foundations of Statistics*. New York: John Wiley & Sons.
- Stacey, Dawn, France Légaré, Krystina Lewis, Michael J. Barry, Carol L. Bennett, Karen B. Eden, Margaret Holmes-Rovner, Hilary Llewellyn-Thomas, Anne Lydiatt, Richard Thomson, and Lyndal Trevena. 2017. "Decision Aids for People Facing Health Treatment or Screening Decisions." *Cochrane Database of Systematic Reviews* 4 (4): CD001431.

## Comment on Chapters 1 and 4      W. Nicholson Price II

### Health AI, System Performance, and Physicians in the Loop

Accounts of artificial intelligence (AI) in medicine must grapple, in one way or another, with the interaction between AI systems and the humans involved in delivering healthcare. Humans are, of course, involved throughout the process of developing, deploying, and evaluating AI systems, but a particular role stands out: the human in the loop of an algorithmic decision. In medicine, when an algorithm is involved in a decision, a typical view of the system envisions a human healthcare professional mediating that algorithm—deciding whether and how to implement or react to any recommendation, prediction, or other algorithmic output. This person is the human in the loop, and their role is often central, complicated, and contested.

Principal contributions to this volume recognize the key role of humans in the loop. Dranove and Garthwaite (2023) focus on the value chain of healthcare and the physician's role within it, recapitulating the notion of the

W. Nicholson Price II is a professor at University of Michigan Law School and a senior fellow at the University of Copenhagen Centre for Advanced Studies in Biomedical Innovation Law.

For acknowledgments, sources of research support, and disclosure of the author's material financial relationships, if any, please see <https://www.nber.org/books-and-chapters/economics-artificial-intelligence-health-care-challenges/health-ai-system-performance-and-physicians-loop>.