Introduction: The Economics of Artificial Intelligence: Health Care Challenges
By Ajay Agrawal, Joshua Gans, Avi Goldfarb, and Catherine Tucker

Healthcare applications of artificial intelligence (AI) and machine learning (ML) have received a great deal of attention in academia. Over the past decade, a number of conferences have brought computer scientists and medical scholars together to develop AI healthcare tools. There are thousands of papers describing new ways to use AI in healthcare, and thousands more detailing what might go wrong as AI diffuses.

Despite this scholarly attention, AI adoption in healthcare has lagged other industries (Goldfarb, Taska, and Teodoridis 2020). To understand the potential for, and the barriers to, AI in healthcare, the second day of the 2022 National Bureau of Economic Research AI conference focused on AI, economics, and healthcare. Specifically, as set out in the invitation, “The goal of the day is to set the research agenda for economists, emphasizing how AI might enable a reimagining of the healthcare system and the economics of healthcare delivery.”

Leading health economists were asked to write up and present ideas. Each presentation had two discussants, typically one with expertise in healthcare and one with expertise in machine learning. Other scholars attended from economics, machine learning, medicine, law, and public health, along with industry experts and clinicians.

This volume contains the four invited articles, along with commentary by a number of other attendees. As such it includes a multi-disciplinary commentary rooted in the perspective of the health economists who wrote the articles. The papers and comments highlight the most important open questions for economists to address. Three themes emerge. First, each paper recognizes that AI has potential to improve healthcare, whether measured by better clinical outcomes or by reduced costs. Second, the papers identify different barriers to the successful deployment of AI in healthcare that may explain the current slow rate of adoption. Third, each paper offers at least a glimmer of hope for overcoming these barriers.

In terms of potential, Dranove and Garthwaite (chapter 1) note that advances in AI “offer new and unprecedented opportunities to improve medical decision-making”. Sahni, Stein, Zemmel, and Cutler (chapter 2) estimate that AI could reduce healthcare costs by 5-10%, yielding hundreds of billions of dollars of savings in the United States each year. Mullainathan and Obermeyer (chapter 3) describe how AI can improve clinical outcomes, emphasizing an example of using AI to test for heart attack in the emergency department. Stern (chapter 4) cites the potential of contemporary deep learning to transform healthcare.

After recognizing this potential, the invited articles in this volume focus more on the barriers to successful adoption of AI. Each paper emphasizes a different headwind to AI’s diffusion in healthcare: incentives (Dranove and Garthwaite), management (Sahni, Stein, Zemmel, and Cutler), data (Mullainathan and Obermeyer), and regulation (Stern).

Dranove and Garthwaite highlight the central role of the physician in the healthcare system. Therefore “the success of AI may depend on buy-in from the very individuals whose success it threatens–physicians.” This threat to physicians comes from the potential for AI in diagnosis. Automated diagnosis “could result in physicians ceding much of their practices to lower cost allied medical professionals". Given the central role of physicians in decision-making, this represents an important barrier to adoption. If "Development and adoption of AI for medical decision making will require the active participation of physicians and other medical decision makers before it is adopted" then broader changes to medical systems are needed, including in payment structures, decision rights, and malpractice risks. An AI that simply replaces physician diagnosis with machine diagnosis will face strong resistance from physician decision-makers. They conclude with a message for economists and machine learning experts interested in healthcare applications of AI: "A particularly important point is for actors outside healthcare to understand how the incentives of existing medical providers can influence the future of AI".

Bell, an executive at pharmaceutical company Novartis and former pharmacy professor, also emphasizes the central role of humans in the loop of medical decision-making. Her emphasis is on the physician, and the recognition of the consequences of physician errors. AI has potential to reduce these errors and improve patient outcomes. She concludes "We don't need to concern ourselves with replacing physicians just yet, let's just work on getting them all to play at the top of their game."

Sahni, Stein, Zemmel, and Cutler emphasize that "management barriers, both at the organizational level and industry level, have been the challenge in healthcare". They provide several specific examples of AI use cases, and argue against the barrier being related to physician payment schemes. Much of the paper is dedicated to measuring the potential impact of AI on health costs if these management barriers could be overcome, emphasizing not just improved clinical outcomes but also how AI could improve administrative efficiency.

Chan's comment on Sahni et al seeks to understand these barriers by asking, "What makes technology adoption different in healthcare relative to other industries?" He notes that IT adoption in healthcare was slow, and that a rich data-driven literature developed to understand how economic incentives affected health IT. He concludes by asking for analysis that leverages heterogeneity in adoption of AI in healthcare, arguing that a closer look at "the effects of adoption on spending and outcomes would likely yield significant insights into the intended and unintended consequences of AI on the healthcare industry as a whole.".

Sendak, Gulamali, and Balu discuss their experience developing and implementing over 15 AI solutions within Duke Health and their interviews of leaders across US health systems. The comment discusses the barriers to AI use in each of the delivery domains for patient care highlighted in Sahni et al. After specifying examples of AI implementation that addresses each of these domains, Sendak, Gulamali, and Balu highlight that "most health system and provider practice AI use cases do not generate financial value", especially under a fee-for-service reimbursement model. They highlight the need for several policy and organizational changes if AI solutions are going to be effective at scale.

Mullainathan and Obermeyer focus on "the lack of accessible clinical data". They describe the various barriers researchers and practitioners face in getting access to health data for building AI tools. First, identifying the necessary data requires deep knowledge of medicine, applied econometrics, and software engineering. Second, data access typically requires the researcher have an appointment at a given hospital. They note that "even faculty members at universities affiliated with the hospital are typically ineligible." Any analyst working with the data is likely to need an appointment in the hospital, limiting the ability of academic researchers and AI companies to build clinical AI tools. They note that open data are a classic public good. “No single actor has a strong incentive to act”. The paper argues that research is fundamental to the advancement of scientific fields, and that the challenges of accessing data for research mean that progress for AI in healthcare is slow.

Stern brings novel data on the regulation of medical AI to describe the ways in which the state of regulation has limited adoption of AI and other medical technologies. The barrier is not regulation per se. She notes that effective regulation can encourage adoption. Instead, the regulatory barriers relate to an uncertain regulatory environment and the challenges that software poses for medical regulation because of frequent updates. She emphasizes, “The value of regulatory innovation and regulatory clarity may be particularly important in the context of AI devices because such a large share of innovations to date have emerged from smaller firms and those from other countries” who may lack US regulatory expertise.

Each paper then offers some hope for overcoming these barriers. Dranove and Garthwaite emphasize the incentives within the medical system, particularly as they relate to the role of the physician. The last two sentences of the chapter provide a hint of how the barriers might be overcome. “A particularly important point is for actors from outside of healthcare to understand how the incentives of existing medical providers can influence the future of AI. This could highlight areas where a greater degree of intervention from outside of the sector may be warranted.” In other words, changes in incentives that come from outside healthcare represent a way for healthcare to eventually benefit from AI’s promise.

For Sahni, Stein, Zemmel, and Cutler, the hope comes through case studies, randomized control trials, and improved data that could be developed to prove the impact of AI in clinical domains. They note long timelines but that the overall promise of the technology suggests that the payoff may be worth the investment. Implicitly, they argue that demonstrated clinical and operational value can overcome incentive-related barriers.

After Mullainathan and Obermeyer’s discussion of the data barriers, they describe steps they have taken to overcome these barriers. They are building two organizations, one non-profit and one for-profit, to make data available to clinical researchers and practitioners. Their non-profit, Nightingale Open Science, aims to catalyze research by supporting the creation of previously unseen datasets and making them accessible to a global community of researchers in ways that preserve patient privacy. Their for-profit, Dandelion Health, focuses on AI product development by building “the largest and highest-quality AI-ready training dataset in the world”. This is perhaps an unusual role for academic economists. After recognizing that “it is perhaps surprising that market forces have not solved the problem of data access”, they are trying to provide their own solutions.

Three comments focus on the access of these two organizations, and how to increase the chance that the data makes a meaningful difference to healthcare globally. Eloundou and Mishkin, both of OpenAI, call for more data and discuss their organization’s approach to developing models which requires high quality data and deep expertise in the model training phase.

Gichoya, a radiologist and informatician, emphasizes incentives for data sharing and legal and reputational risks as limiting the sharing of healthcare data. She highlights several of the competencies that Nightingale and Dandelion help deliver, concluding that they help with organization-level barriers related to “compliance, data science, finances, intellectual property, and legal expertise of data use agreements”..

The comment by Papyan, Donoho, and Donoho (a mathematician-computer scientist, pediatric neurosurgeon, and statistician) provides a detailed discussion of the role of data platforms and how Mullainathan and Obermeyer’s Nightingale initiative could lead “to the unleashing of a great deal of research energy as traditional barriers to research are shattered”. They provide evidence that “the data platform concept–join researchers with data–has been proven to work in field after field across decades”. Citing DARPA, IMAGENET, and MNIST, the make the bold claim that “From this viewpoint, [Turing Award winner] Yann LeCun made a bigger impact by developing the MNIST dataset and publishing it, than by the specifics of any actual ML models he constructed for use with MNIST. Those early neural net models have been superseded, but MNIST is still powering research papers today.” They also note that the success of data platforms typically hinges on rewards, and so economists should not ignore our expertise in incentives in developing and promoting health data platforms. The specific nature of such rewards and the data available through the platform should recognize advances in AI technology, recognizing that the goal is to incentivize the next generation of healthcare AI, the supply of data, and the development of models. They summarize, “By incorporating elements of the successful Common Task Framework, such as rewards and leaderboards, Nightingale can encourage participation and drive progress in the field. .

Stern notes that many of the regulatory barriers are already being overcome. She notes “a dramatic uptick in the commercialization of AI products over recent years” and that “regulators have begun a germane and important discussion of how such devices could be regulated constructively in the future. Risk classification and the regulation of software updates are important areas for regulatory innovation. Compared to the other chapters, Stern’s is relatively optimistic due to ongoing regulatory innovation. The right kind of regulation creates incentives to invest in new technology.

Babic, a philosophy professor, highlights Stern’s data and analysis, noting the increase in approvals of AI devices and the preliminary evidence on safety is encouraging. He provides some more detail on the regulatory innovation, building on Stern’s chapter and Babic et al (2019) to note that “the FDA has required software based medical devices to undergo a new round of review every time the underlying code is changed”, and a new proposed framework, still in its infancy, “could made for a much more productive partnership between the FDA and medical AI manufacturers”. He concludes that an alternative regulatory approach, with a separate agency for governing algorithms across all domains, would allow for a more AI-focused regulatory environment.

Several commenters touch on themes that appear in multiple papers. These comments mainly focused on strategies for overcoming the various barriers.

Operations professor Lu uses Sahni et al to highlight the potential and then suggests ways that the physician resistance emphasized by Dranove and Garthwaite might be overcome. Specifically, with respect to physician resistance, she argues that “allowing physicians to freely choose whether to seek opinions from AI would greatly decrease the tension between physicians and AI and promote physician trust in AI.”

Legal scholar Price recognizes “the key role of humans in the loop” while worrying that “not all health-care providers will be able, adequately trained, or well-resourced enough to catch errors in the system or to ensure it works as intended”. He highlights how the anticipated use and the anticipated user affect design and regulation. The specific person in the loop matters. He notes that if “the value of AI systems in health-care settings may in fact be greatest in situations where human experts are least available” then systems designed for an expert in the loop will eliminate much of the value of AI in healthcare.

Economists Bundorf and Polyakova also explore the changing role of the physician, building on ideas on Dranove and Garthwaite’s chapter and using a framework familiar to economists: decision-making under uncertainty. AI generates predictions, but clinical decision-making is still likely to benefit from incorporating patient preferences. Patients may not have well-formed preferences over treatment outcomes and so physicians and other medical professionals play a role in helping patients formulate preferences. Bundorf and Polyakova note that AI remains limited in its ability to predict such preferences and is likely to assume away preference heterogeneity across patients. This, in turn, suggests a heterogeneous impact on clinicians. Those who are skilled at helping “patients translate prediction into decisions by incorporating patient preferences will have skills which are complementary to AI.” Others will not. Ultimately, they note that an important aspect of the opportunity for AI to improve healthcare is to “incentivize physicians to focus on what medical students often say motivated them to choose medicine–listening to the patient.”

Information systems professor Adjerid’s comment discusses incentives and how technology can change the practice of healthcare. His insights come from research on the diffusion of electronic medical records, noting some parallels and a handful of differences. Epidemiologist Rosella also highlights lessons from electronic medical records, noting unintended negative consequences. Her comment discusses incentives, trust, and interface design.

We finish the volume with Rosella's conclusion as it effectively summarizes the four barriers highlighted in the papers, but reframes them as “the building blocks needed for AI to have a meaningful impact in healthcare”: (1) designing AI to support the transparency needed in healthcare decisions (from Dranove and Garthwaite), (2) understanding of the complex healthcare environment (from Sahni et al), (3) enrich the data used and ensure it is made available in a responsible way (from Mullainathan and Obermeyer), and (4) innovative models of regulation (from Stern).

References

Agrawal, Ajay, Joshua Gans, and Avi Goldfarb, Eds. 2019. The Economics of Artificial Intelligence: An Agenda. University of Chicago Press, Chicago IL.

Goldfarb, Avi, Bledi Taska, and Florenta Teodoridis. 2020. Artificial Intelligence in Healthcare? Evidence from Online Job Postings. *American Economic Association Papers & Proceedings,* 110: 400-404.