

How Does AI Accelerate Drug Development: A Global Perspective

Panle Jia Barwick Hongyuan Xia Tianli Xia Jing Yi*

June 2026

1 Introduction

Artificial intelligence is increasingly embedded in biomedical research, from target discovery and molecular design to clinical trial design and real-world evidence generation (for example, Zhang et al. [2025](#)). Yet it remains unclear whether this technological shift has translated into measurable improvements in the productivity of drug development. This distinction matters because AI may expand the volume and sophistication of biomedical research without necessarily accelerating the translation of scientific ideas into patents, clinical trials, approved therapies, or patient benefits. Quantifying this link is therefore essential for understanding whether AI represents a genuine improvement in the biomedical innovation process or primarily a new set of tools layered onto existing research practices.

Moreover, the global geography of pharmaceutical innovation has been characterized by a clear hierarchy: The United States, alongside Europe and Japan, has led drug innovation, while developing countries have mainly consumed existing medicines or produced generics (Garthwaite [2025](#); Lakdawalla [2018](#)). Whether AI narrows or widens this global gap is theoretically ambiguous. On the one hand, AI may package existing scientific knowledge and lower the cost of accessing frontier research, potentially offering latecomer countries a shortcut to innovation (Agrawal, McHale, and Oettl [2026](#)). On the other hand, the downstream

*Barwick: Department of Economics, UW-Madison, NBER, and CEPR, pbarwick@wisc.edu. H. Xia: Department of Economics, Cornell University, hx276@cornell.edu. T. Xia: Simon Business School, University of Rochester, tianli.xia@rochester.edu. Yi: Department of Agricultural & Applied Economics, UW-Madison, jing.yi@wisc.edu

impact of AI depends on complementary capabilities, including research institutions that can absorb new methods, firms that can translate discoveries into development programs, clinical trial infrastructure, and regulatory systems that support innovation (Babina et al. 2024; Brynjolfsson, Rock, and Syverson 2021).

This paper therefore asks two questions. First, does AI adoption in upstream biomedical science translate into downstream drug development outcomes, such as patents, clinical trials, and approved therapies? Second, how does this relationship vary across countries, and does AI enable emerging innovation systems to catch up with established pharmaceutical leaders? We focus in particular on China as a central case. Over the past two decades, China has rapidly expanded its pharmaceutical innovation capacity, clinical trial activity, biomedical research output, and regulatory infrastructure (Barwick, Xia, and Xia 2026). At the same time, it has become one of the world’s major producers of AI research and applications (Fang et al. 2026; Maslej et al. 2025). This setting allows us to examine whether AI serves as an equalizing force in global pharmaceutical innovation or instead reinforces existing advantages held by established leaders.

Answering these questions is empirically challenging for several reasons. First, AI adoption in upstream biomedical science is difficult to measure systematically. AI includes a broad set of methods used across biomedical domains, and standard adoption measures are not well suited to this setting. Job-posting or skill-based measures capture firm-level labor demand but are difficult to map to disease-specific scientific activity, while patent-based AI measures capture later-stage invention and may misclassify biological concepts such as signal processing and molecular computation as artificial intelligence. Second, quantifying the translation from science to drug development requires linking publications to diseases, countries, patents, clinical trials, drug projects, and approvals. Third, AI adoption is likely endogenous, as disease areas with stronger commercial potential, scientific opportunities, or research infrastructure may attract both AI research and downstream development. Finally, the welfare implications of AI remain ambiguous because the social value of additional in-

novation depends on the diseases targeted, approval probabilities, disease burden, and the distribution of benefits across countries.

To address these challenges, we combine large-scale publication, patent, clinical trial, drug project, approval, and disease burden data. We construct a publication-based measure of AI adoption in upstream biomedical science using PubMed publications from 2010 to 2025. Rather than relying only on keywords, we build a training sample and use manual review and LLM-assisted labeling to identify substantive AI use in biomedical research. We then map AI-related publications to diseases and countries using MeSH terms and author affiliations, and link this disease-country-year measure to PATSTAT patents, Citeline clinical trials, drug project and approval data, and disease burden measures. Our main analysis examines whether greater upstream AI exposure is followed by more downstream innovation. To address endogenous adoption, we construct shift-share instruments that combine pre-period exposure to computer science with subsequent global growth in AI and related computer science fields. Finally, we use the estimates to construct counterfactual innovation trajectories and evaluate the welfare implications of AI-driven pharmaceutical innovation across diseases and countries.

This project offers several contributions. First, the project contributes to the AI-and-productivity literature by studying innovation rather than routine output. While existing studies focus on task efficiency, we examine whether AI adoption changes the drug discovery pipeline itself, including the generation and advancement of drug candidates (Bick, Blandin, and Deming 2026; Brynjolfsson, Li, and Raymond 2025; Cui et al. 2026; Lou and Wu 2021). This connects to a broader view of AI as a technology that may reshape scientific discovery and technology commercialization rather than simply augment or replace existing tasks.

Second, this paper bridges the gap between AI adoption in upstream scientific research and downstream innovation outcomes. Existing work has emphasized the potential of AI to transform biomedical discovery, but there is limited systematic evidence on whether AI-intensive research translates into patents, clinical trials, drug development progress, and

approved therapies (Hill and Stein 2026). By linking AI-related biomedical publications to multiple downstream stages of the drug development process, we provide evidence on whether AI changes not only the production of scientific knowledge but also its translation into commercially and clinically meaningful innovation. Moreover, by translating these effects into welfare estimates, we connect technological progress to human health outcomes in terms that are directly relevant for policy and resource allocation decisions.

Third, the paper contributes to the emerging literature on heterogeneity in the effects of AI by studying whether AI narrows or widens global innovation gaps. Recent studies show that the productivity effects of generative AI vary across workers and contexts, often benefiting less experienced workers more strongly (Brynjolfsson, Li, and Raymond 2025; Cui et al. 2026). Our setting takes a global perspective by asking whether a similar equalizing force operates across countries. In particular, we examine whether AI enables emerging innovation systems, especially China, to translate scientific advances into downstream pharmaceutical innovation, or whether the benefits of AI remain concentrated in established pharmaceutical leaders with stronger complementary capabilities. This contributes to the broader literature on global innovation, catch-up, and the geography of pharmaceutical development.

2 Data and Methodology

2.1 Data Sources

We combine several data sources to measure the diffusion of artificial intelligence in biomedical science and its downstream relationship with drug development. First, we construct a corpus of biomedical publications from PubMed covering 2010–2025. Using publication titles, abstracts, keywords, references, and author affiliations, we identify AI-related biomedical research, classify papers by disease area, and measure AI research intensity at the disease-country-year level. This allows us to trace how AI capabilities diffuse across biomedical fields and geographic regions over time.

Second, we link this publication-based measure to patenting and clinical trial activity. Patent data come from PATSTAT, while clinical trial records come from Citeline. The trial data provide information on disease indications, trial phases, start dates, completion dates, sponsors, study locations, and trial status. We harmonize publications, patents, and clinical trials using disease and indication taxonomies, including MeSH terms, ICD codes, and our own disease classification. This linkage enables us to examine whether AI-intensive biomedical research is followed by greater inventive activity and more active clinical development in the same disease areas.

Third, we use drug approval and disease burden data to study downstream innovation and welfare outcomes. Drug approval information comes from the Citeline Pharmaprojects database, while disease burden measures come from the Global Burden of Disease database. These sources allow us to examine whether AI-related biomedical research is associated not only with more scientific output, patents, and clinical trials, but also with faster development, new drug approvals, and potential health gains. Together, these data provide a comprehensive view of the role of AI in the biomedical innovation pipeline, from scientific discovery to clinical translation and patient-relevant outcomes.

2.2 Measuring AI Adoption in Upstream Biomedical Science

Existing measures of AI adoption are not well suited for our setting. A common approach in the literature measures AI adoption using job postings or worker skills (Goldfarb, Taska, and Teodoridis 2023). These measures are useful for studying firm-level adoption and labor-market adjustment, but they are less informative about the use of AI in upstream biomedical science, where many relevant applications occur within university labs, hospitals, research institutes, and scientific collaborations rather than through explicit AI hiring by firms. They are also difficult to map to specific disease areas, which is essential for linking upstream scientific activity to downstream drug development. Patent-based AI measures face related limitations (Fang et al. 2026). Although they are useful for measuring AI-related invention,

they are less directly suited to capturing the diffusion of AI methods across biomedical research. More importantly, in biomedical and pharmaceutical contexts, broad AI classifiers can mistakenly classify biological concepts such as logic gates, circuits, signal processing, and molecular computation as artificial intelligence.

We therefore construct a publication-based measure of AI adoption tailored to upstream biomedical science. We begin with the universe of PubMed publications from 2010 to 2025. For each publication, we use titles, abstracts, keywords, MeSH terms, references, journal information, and author affiliations to identify whether the paper applies or develops AI methods in a biomedical context. Rather than relying only on keyword searches, we first use AI-related terms and biomedical text features to construct a candidate set, and then build a training sample labeled manually and/or with the assistance of large language models. We use this labeled sample to train and validate a classifier that distinguishes explicit AI applications, such as machine learning, deep learning, natural language processing, and computer vision, from broader computational biology or bioinformatics papers that use terms such as logic, circuit, signal processing, or computation in a primarily biological rather than AI sense.

We then map each publication to disease areas using MeSH terms, disease keywords, and indication taxonomies, and assign publications to countries using author affiliations. This procedure allows us to construct a disease-country-year measure of upstream AI adoption. Our baseline measure is the share of biomedical publications in disease d , country c , and year t that are AI-related; we also construct alternative measures based on the count and log count of AI-related publications, as well as AI publication intensity relative to total scientific output in the same disease-country cell. Because this measure captures AI adoption at the scientific discovery stage and is constructed at the disease-country-year level, it can be linked directly to downstream outcomes in the same disease and country, including patents, clinical trials, drug development progress, and approvals.

2.3 Empirical Strategy

Our main analysis estimates whether upstream AI adoption in biomedical science translates into downstream drug development. The baseline unit of analysis is the disease-country-year level. We estimate the following specification:

$$Y_{dct} = \beta AIExposure_{dc,t-j} + \lambda_{dc} + \mu_{dt} + \nu_{ct} + \varepsilon_{dct}, \quad (1)$$

where Y_{dct} measures downstream drug development activity for disease d in country c and year t . The outcome can include the number of patents, number of clinical trials, trial initiation rates, phase-specific trial activity, phase transition probabilities, or drug approvals. The key explanatory variable, $AIExposure_{dc,t-j}$, measures lagged AI-related biomedical research in the same disease-country cell.

The specification includes disease-country fixed effects, λ_{dc} , which absorb time-invariant differences across disease-country pairs; disease-year fixed effects, μ_{dt} , which control for global disease-specific shocks; and country-year fixed effects, ν_{ct} , which control for country-level trends in research capacity, regulation, and drug development. The coefficient β captures whether disease-country areas with greater upstream AI exposure subsequently experience more downstream drug development, relative to other diseases in the same country and the same disease in other countries.

To examine whether AI has a differential effect in China, we augment the baseline specification with an interaction between lagged upstream AI exposure and an indicator for China:

$$Y_{dct} = \beta AIExposure_{dc,t-j} + \theta (AIExposure_{dc,t-j} \times China_c) + \lambda_{dc} + \mu_{dt} + \nu_{ct} + \varepsilon_{dct} \quad (2)$$

Here, $China_c$ is an indicator equal to one for China. Because the country-year fixed effects absorb the level effect of China in each year, the interaction term identifies whether the relationship between upstream AI exposure and downstream drug development is stronger or

weaker in China relative to other countries, conditional on disease-country, disease-year, and country-year fixed effects. The coefficient θ therefore captures whether China is especially effective at translating AI-related biomedical research into downstream pharmaceutical innovation. A positive θ would be consistent with AI accelerating China’s catch-up in drug development, whereas a zero or negative θ would suggest that the benefits of AI remain similar across countries or depend more heavily on complementary capabilities outside China.

2.4 Instrumental Variable

A key empirical concern is that AI adoption may be endogenous. For example, diseases with stronger commercial potential or faster scientific progress may attract both more AI research and more clinical trials. To address this issue, we construct shift-share instruments that combine pre-period exposure to computer science with subsequent global growth in AI or computer science fields.

Our first instrument interacts a disease-country unit’s pre-period connection to computer science with global growth in AI publications:

$$Z_{dc,t-j} = CSShare_{dc,0} \times \Delta \log \left(AIPublications_{-c,t-j}^{Global} \right), \quad (3)$$

where $CSShare_{dc,0}$ is the pre-period share of biomedical papers in disease d and country c that are connected to computer science, and $\Delta \log \left(AIPublications_{-c,t-j}^{Global} \right)$ captures the growth of global AI publications excluding country c . The intuition is that disease-country research areas historically closer to computer science should be more exposed to the global expansion of AI.

We also consider two related instruments. The first uses pre-period exposure to different computer science fields, such as machine learning, computer vision, natural language processing, bioinformatics, and data systems:

$$Z_{dct}^{Field} = \sum_{s \in \mathcal{S}} Share_{dcs,0} \times \Delta \log \left(CS_{st}^{Global,-c} Publications_{st} \right). \quad (4)$$

The second uses historical connections to frontier computer science institutions:

$$Z_{dct}^{School} = \sum_{u \in \mathcal{U}} Share_{dcu,0} \times \Delta \log (AI_{st} Publications_{ut}), \quad (5)$$

where u indexes top computer science schools. These instruments exploit variation in how different disease-country research portfolios were historically positioned to benefit from later advances in AI. The identifying assumption is that, conditional on the fixed effects, these pre-period connections affect downstream drug development only through their effect on AI-related biomedical research.

2.5 Welfare and Counterfactual Analysis

Finally, we use the estimated relationship between AI exposure and downstream drug development to evaluate the broader welfare implications of AI in biomedical innovation. The first step is to construct counterfactual drug development trajectories under different AI adoption scenarios. For example, we can compare observed trial activity and drug approvals with a counterfactual no-AI or lower-AI baseline over a 10- or 20-year horizon.

We then translate the estimated additional drug development activity into welfare measures. Specifically, we combine disease-level estimates of additional clinical trials, phase transitions, and drug approvals with external estimates of health benefits, including QALYs gained, DALYs averted, and disease-specific burden measures from Global Burden of Disease and WHO data. Where appropriate, we also use value-of-statistical-life and willingness-to-pay approaches to convert health improvements into economic value.

This analysis allows us to assess not only whether AI increases biomedical innovation, but also where its benefits are likely to be largest. We will examine welfare gains by disease category, country, and income region, which helps distinguish whether AI primarily bene-

fits established pharmaceutical markets or also improves access to innovation in lower- and middle-income settings.

References

- Agrawal, Ajay K, John McHale, and Alexander Oettl. 2026. *AI in Science*. Technical report. National Bureau of Economic Research.
- Babina, Tania, Anastassia Fedyk, Alex He, and James Hodson. 2024. “Artificial intelligence, firm growth, and product innovation.” *Journal of Financial Economics* 151:103745.
- Barwick, Panle Jia, Hongyuan Xia, and Tianli Xia. 2026. *From Free Rider to Innovator: The Rise of China’s Drug Development*. Technical report. National Bureau of Economic Research.
- Bick, Alexander, Adam Blandin, and David Deming. 2026. “The Rapid Adoption of Generative AI.” *Management Science*.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. “Generative AI at Work.” *The Quarterly Journal of Economics* 140 (2): 889–942.
- Brynjolfsson, Erik, Daniel Rock, and Chad Syverson. 2021. “The productivity J-curve: How intangibles complement general purpose technologies.” *American Economic Journal: Macroeconomics* 13 (1): 333–372.
- Cui, Kevin Z., Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz. 2026. “The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers.” *Management Science*.
- Fang, Hanming, Xian Gu, Hanyin Yan, and Wu Zhu. 2026. *AI Patents in the United States and China: Measurement, Organization, and Knowledge Flows*. Technical report. National Bureau of Economic Research.
- Garthwaite, Craig. 2025. “Economic markets and pharmaceutical innovation.” *Journal of Economic Perspectives* 39 (2): 3–26.
- Goldfarb, Avi, Bledi Taska, and Florenta Teodoridis. 2023. “Could machine learning be a general purpose technology? A comparison of emerging technologies using data from online job postings.” *Research Policy* 52 (1): 104653.
- Hill, Ryan R, and Carolyn Stein. 2026. *How Artificial Intelligence Shapes Science: Evidence from AlphaFold*. Technical report. National Bureau of Economic Research.
- Lakdawalla, Darius N. 2018. “Economics of the pharmaceutical industry.” *Journal of Economic Literature* 56 (2): 397–449.
- Lou, Bowen, and Lynn Wu. 2021. “AI on drugs: Can artificial intelligence accelerate drug development? Evidence from a large-scale examination of bio-pharma firms.” *MIS quarterly* 45 (3): 1451–1482.
- Maslej, Nestor, Loredana Fattorini, Raymond Perrault, Yolanda Gil, Vanessa Parli, Njenga Kariuki, Emily Capstick, Anka Reuel, Erik Brynjolfsson, John Etchemendy, et al. 2025. “Artificial intelligence index report 2025.” *arXiv preprint arXiv:2504.07139*.
- Zhang, Kang, Xin Yang, Yifei Wang, Yunfang Yu, Niu Huang, Gen Li, Xiaokun Li, Joseph C Wu, and Shengyong Yang. 2025. “Artificial intelligence in drug development.” *Nature medicine* 31 (1): 45–59.