The Impacts of AI at Scale: Evidence from Research Scientists*

Zhengyi Yu†

September 7, 2025

Abstract

This paper studies the impact of AI on productivity and inequality. While prior research has examined the effects of AI adoption in small experiments, my paper studies the impacts of AI at scale by leveraging the introduction of AlphaFold2. This AI algorithm is capable of accurately predicting the structure of proteins, a task traditionally performed by structural biologists. To capture the net impact of AI on structural biologists, I implement a difference-in-differences strategy comparing them to life scientists in other fields. While structural biologists did not change their overall number of publications with the availability of AlphaFold2, they experienced a 10% increase in citations to their new projects, a 4% rise in publications in high-impact journals, and a 3% shift from their original research trajectory. However, the emergence of AI intensifies citation polarization between highly cited and less-cited researchers. To explain this growing inequality, I find that highly cited scientists are twice as likely to incorporate AlphaFold2 into their research as their less-cited peers. Moreover, AI increases the demand for multidisciplinary expertise and raises the average age of major contributing scientists in protein structure research team. Overall, these findings suggest that while AI boosts scientific productivity on average, it also exacerbates inequality.

*This paper was written using data obtained in May, 2025, from Digital Science's Dimensions platform, available at https://app.dimensions.ai. The author did not receive any external research support for this work and declare that they have no relevant financial relationships or potential conflicts of interest.

†Department of Economics, University of Southern California, yuzhengy@usc.edu

1 Introduction

The rapid diffusion of artificial intelligence (AI) has sparked growing attention to their effects on labor productivity and inequality. Recent studies tend to find that AI can enhance the productivity of adopters and help level the playing field for less-skilled workers in small-scale experimental settings (Brynjolfsson et al., 2025; Noy and Zhang, 2023; Dell'Acqua et al., 2023). Nevertheless, questions persist about the net benefit of AI at scale. In experiments, workers with and without AI are typically randomly assigned and prevented from competing directly, allowing for a clean estimate of AI's partial-equilibrium treatment effect. In the real world, by contrast, AI adoption is usually voluntary, with adopters and non-adopters competing in the same labor or product markets. Thus, gains for AI adopters may come at the expense of non-adopters, particularly in the short run, when demand adjusts slowly. When evaluating the impact of AI at scale, we must also account for its potential externalities. Moreover, while AI tools may offer greater potential benefits to less-skilled workers, individuals with lower education or income are less likely to adopt them in practice (Humlum and Vestergaard, 2024; Bick et al., 2024), leaving it unclear whether AI is, on net, democratizing.

To evaluate the impact of AI on labor productivity and inequality at scale, this paper exploits the introduction of AlphaFold2, a Nobel Prize–winning AI algorithm that accurately predicts protein three-dimensional (3D) structures. This breakthrough most directly affected structural biologists, for whom determining protein structures is a core task. AlphaFold2 can provide valuable hypotheses, accelerating the process and improving the accuracy of protein structure determination. I employ a Difference-in-Differences (DiD) approach with coarsened exact matching (CEM) to identify the causal impact of AlphaFold2, comparing the academic records of structural biologists to those of comparable life scientists who do not solve protein structure, before and after the introduction of AI. Rather than comparing AI users to non-users within structural biology, I use cross-field variation in exposure to capture the net impact of AlphaFold2 on all structural biologists.

Overall, the release of AlphaFold2 did not increase the quantity of research by structural biologists but enhanced its quality and shifted the direction of their scientific inquiry. However, these productivity gains are concentrated among already highly cited and well-resourced researchers, raising concerns about widening inequality. Consistent with the growing polarization, correlation results show that highly cited scientists are more likely to use AlphaFold2 in their research than their less-cited peers. Furthermore, following the emergence of AI, team leaders of protein structure determination increasingly come from interdisciplinary backgrounds, and the average academic age of key contributors rises, suggesting that AI also reshapes human capital dynamics within research organizations.

Several features of AlphaFold2 and structural biology make it ideal for examining the net impacts of AI on the knowledge workers. First, AlphaFold2 is a free, well-established, and popular tool in real-world research. Since its release in 2021, structural biologists have widely integrated AlphaFold2 into their protein structure research (Kovalevskiy et al., 2024). Therefore, the popularity of AlphaFold2 in structural biology community allows me to estimate the impact of AI at scale. Second, structural biology researchers compete vigorously to determine the structure of important proteins, aiming for priority and publication in academic journals (Hill and Stein, 2025a,b). This setting provides an opportunity to assess the net impact of AI in a highly competitive market. Third, academic research often tackles questions with no prior answers, providing an opportunity to test AI's role in helping knowledge workers solve open-ended problems and shape the direction of innovation.

Three major bibliometric databases support the delineation of treatment and control groups in my DiD design and enable the construction of a panel dataset of academic outputs. First, I use the PubMed database to locate publications in the field of life sciences. Second, I rely on the Protein Data Bank (PDB), a repository containing structural data for biological macromolecules, primarily proteins. Each PDB entry is linked to the corresponding publication via its PubMed ID, allowing me to identify PubMed papers that focus on protein structure determination. Third, I use the Dimensions database, which provides identifiers that enable tracking the academic history of authors of specific PubMed papers and PDB-linked publications over time.

The analysis sample consists of Principal Investigators (PIs) in the life sciences, selected based on their publication records from 2013 to 2016. Structural biologists (the treatment group) are defined as PIs with at least two publications linked to the PDB. The control group includes life scientists who are PIs with at least two PubMed-indexed papers published in journals that also feature PDB-linked papers, allowing for comparability in research field. I further exclude control PIs with any PDB-linked publications to rule out direct involvement in protein structure determination. Lastly, I match control PIs to treated PIs using coarsened exact matching (CEM) based on key characteristics such as institution and citation level to improve similarity across observed attributes. After defining the treatment and control groups, I query Dimensions for the complete academic histories of the researchers in my sample. The results are robust to alternative definitions of treatment.

I document three main findings in this study. First, to assess the impact of AI on the productivity of human researchers, I examine how AlphaFold2 affected the quantity, academic influence and quality of structural biologists' research. The results show that AlphaFold2 did not change the overall number of projects led by structural biology PIs but significantly improved their academic impact and quality. Specifically, structural biologists

experienced a 10% increase in citations to their new publications after the introduction of AlphaFold2. Furthermore, structural biologists increase the number of publications in high-impact-factor (IF) journals by 4%. The introduction of AI also shifts treated scientists' research trajectory, making it 3% more distant from their original work, as measured by text analysis of titles, abstracts, and keywords. At the same time, AlphaFold2 had no detectable effect on the extensive margin of publication, funding, or retention in academic institutions among treated scientists. Suggestive evidence indicates that the positive impact of AI on the structural biology labor force is driven by researchers who integrate AI into their work. Although we see a decline in productivity associated with researchers not using AI, the overall impact of AI on the productivity remains positive. These findings alleviate concerns about the scalability of AI's benefits, showing that AI can enhance the overall quality of research output and steer scientific inquiry in new directions.

The second set of results examines the heterogeneous impact of AI. I use baseline citation levels as a proxy for researcher skill. The findings reveal that the overall improvement in research quality driven by AI masks a growing productivity divide: highly cited researchers experience a 12% increase in total citations, while less-cited researchers see no significant change. This pattern is consistent with evidence that highly cited researchers are more likely to integrate AI into their work. I argue that variation in AI adoption is not solely driven by the information channel, as AlphaFold2 is still more frequently used in studies of complex proteins, even after accounting for researcher and year fixed effects. In addition to the skill-biased effect of AI, I find further heterogeneity by resources and geographic location. Researchers with greater funding, and those based in the top three countries with the largest structural biology communities experience larger productivity gains.

The third set of results examines how AI technology affects the organization and the dynamics of the field. I use a repeated cross-section sample to compare the team compositions of protein structure papers and other life science papers published in similar journals. The emergence of AI does not affect the number of PIs or non-PIs involved in protein structure projects. However, the number of institutional affiliations associated with PIs in these papers increases by 10% following the introduction of AI, suggesting a growing reliance on diverse expertise. In addition, the average age of first and corresponding authors increases by 4%. These increases in institutional diversity and major contributors' age are consistent with the findings of Jones (2009), which documents that when the frontier of knowledge advances, the innovation requires more interdisciplinary inputs and delay the age of innovators. Therefore, the emergence of AI may complement incumbent workers while increasing competition within the field.

Overall, the findings suggest that AI improves the quality of researchers' output without

reducing its quantity. I do not claim that the conclusions of this paper apply universally to all types of AI. Broadly, AI can be categorized into two types: those designed to replace human labor, such as fully autonomous vehicles, and those that complement human, such as AI in healthcare (Acemoglu and Restrepo, 2020; Agrawal et al., 2019). AlphaFold2 belongs to the latter one. Although AlphaFold2 is a domain-specific tool in biology, it embodies two central functions of augmentative AI for knowledge workers: (1) generating valuable hypotheses for testing, and (2) streamlining labor-intensive processes. As such, the findings of this paper offer broader insights into the application of AI in tasks that require the rapid generation, testing, and refinement of complex solutions—such as drug discovery and the delivery of highly customized services.

This paper contributes to the broad literature on the impact of technology on human labor. Many theoretical and empirical evidence suggest that the development of robotic automation and information technology have reduced employment—at least in the short run—and disproportionately benefited high-skill workers, leading to increased wage inequality (Rosen, 1981; Katz and Murphy, 1992; Berman et al., 1994; Autor et al., 1998; Garicano and Rossi-Hansberg, 2006; Bartel et al., 2007; Michaels et al., 2014; Akerman et al., 2015; Autor, 2015; Acemoglu and Restrepo, 2018; Feigenbaum and Gross, 2024). Unlike earlier technologies primarily automating routine tasks, AI has the capacity for cognitive functions such as generating hypotheses and ideas. Therefore, it remains unclear whether the labor market impact of AI technologies will mirror that of previous technological revolutions. Although Acemoglu et al. (2022) found that AI had no significant effect on employment or wages, they acknowledged that it may be too early to observe measurable impacts. Overall, research on the labor market effects of AI remains nascent (Restrepo, 2023).

My work is most closely related to the emerging literature on the impact of AI on labor productivity. Recent studies show that AI may affect the labor market differently from previous technologies by serving as a complementary tool and narrows the gap between high- and low-skill workers (Brynjolfsson et al., 2025; Noy and Zhang, 2023; Caplin et al., 2024; Peng et al., 2023; Choi et al., 2024; Dell'Acqua et al., 2023; Kanazawa et al., 2025).[1] While these studies offer micro-level evidence on the partial-equilibrium effects of AI on its adopters, they leave open questions about AI's net impact on labor productivity and inequality for two reasons. First, productivity gains among AI adopters may come at the expense of non-adopters, potentially resulting in no net increase in total productivity at scale.[2] Second, even if AI appears to benefit low-skill workers among the AI adopters, its

[1] Otis et al. (2023) and Wang et al. (2023) show that the "low-skill bias" of AI is context-dependent, and in some settings AI may widen rather than narrow productivity differences, especially in non-routine work.

[2] For example, Crépon et al. (2013) find that a program can be effective when applied to a small group of participants, but offers little net benefit at scale due to displacement effects.

aggregate effect on productivity inequality remains ambiguous if high-skill individuals are more likely to adopt AI. This is particularly likely when the adoption of AI is voluntary and requires expertise, as recent survey evidence suggests (Humlum and Vestergaard, 2024; Bick et al., 2024; Handa et al., 2025). Leveraging a quasi-experimental setting and rich bibliometric data, this paper provides the first empirical evidence on the effects of AI on productivity and inequality among scientists at scale.

This paper also contributes to the growing literature on the impact of technology on scientific research. Agrawal and Goldfarb (2008), Teodoridis (2018), Furman and Stern (2011), and Ding et al. (2010) show that advances in technology can enhance follow-on scientific innovation and collaboration. Furthermore, Furman and Teodoridis (2020) and Kim (2023) show that the emergence of automation technologies can shift the direction of scientific inquiry. Previous technologies primarily focused on automating specific tasks. In contrast, a key innovation of AI lies in its unprecedented ability to reason and generate hypotheses for scientific testing. As a result, AI has the potential to influence scientific research in ways that differ fundamentally from prior technologies. Doshi et al. (2024), for instance, examine how generative AI affects the self-evaluation of research proposal novelty in an experimental setting. This paper contributes to the existing literature by examining how AI affects real-world research quality, democratization, direction and the competitive dynamics of the field over time.

The remainder of the paper is organized as follows. Section 2 introduces the empirical setting, discussing the development of AlphaFold2 and its role in protein structure determination. Section 3 describes the data, sample construction, and empirical strategy. Section 4 presents the main results on the impact of AI on researcher productivity. Section 5 explores the distributional effects of AI. Section 6 examines the influence of AI on the dynamics of research teams in the field. Section 7 concludes.

2 Background

This section provides an overview of structural biology, the development of AI within the field, and how AI affects the research of protein structure.

2.1 Structural Biology

Structural biology is a branch of molecular biology, biochemistry, and biophysics that focuses on the molecular structure and dynamics of biological macromolecules, particularly proteins. Proteins are the "workhorses" of biological life, which are composed of a linear sequence of

amino acids that fold into a 3D structure. Proteins support virtually all biological processes and their 3D structures are key to understanding how proteins interact, function biologically, and cause diseases when they malfunction. Scientists often understand the amino acid sequence of proteins but they cannot easily deduce the 3D structure from the sequence. Experimentally solving a protein's 3D structure often requires structural biologists to spend months or even years with advanced experimental equipments.

2.2 Protein Structure Prediction

The computational prediction of the protein structure has been a pursuit in structural biology for over five decades. Various software were developed prior to AlphaFold2 but they all lacked the accuracy needed for practical applications. In 2018, by employing deep learning, AlphaFold1 was able to predict protein structures with enhanced accuracy and won first place in the 13th Critical Assessment of Structure Prediction (CASP13)[3]. However, AlphaFold1's accuracy only exceeded 60% in 2018, which was still far from the practical usage threshold of 90%.

The 2020 release of AlphaFold2 exceeded public expectations with a dramatic improvement in accuracy, marking a paradigm shift in structural biology[4]. On November 30, 2020, AlphaFold2 was declared the winner of CASP14, demonstrating an astounding accuracy rate of over 90%. Panel A in Appendix Figure A1 shows the comparison between AlphaFold2's predictions and experimentally determined structures in CASP14.

AlphaFold2's achievement surpassed most scientists' expectations and was hailed as solving the 50-year-old "protein-folding problem." The AlphaFold team published their AI algorithm and released a database of predicted protein structures in July 2021, which I define as the first treatment year in my analysis. In July 2022, this database expanded over 200 million protein structures, encompassing nearly all known proteins. The AlphaFold2 prediction database and the corresponding algorithm has been freely accessible from the outset, and researchers are only required to cite the relevant AlphaFold2 papers if they use its data or methodology (Jumper et al., 2021; Varadi et al., 2022). In 2024, two leaders of the AlphaFold team (Demis Hassabis and John Jumper) were awarded the Nobel Prize in Chemistry.

Despite the high accuracy of structural predictions, AI cannot fully replace human structural biologists (Terwilliger et al., 2023). The number of protein structure papers remains stable after the debut of AlphaFold2 in top 3 natural science journals: *Cell*, *Nature*, *Science*

[3]CASP is a worldwide competition held every two years since 1994 to foster the study of computational prediction of protein 3D structures.

[4]Appendix Section B outlines the evolution of AlphaFold2 from AlphaFold1 and addresses the potential concern of anticipatory effects related to the arrival of AlphaFold2.

(CNS) and journals publishing protein structure research (see Panel B of Appendix Figure A1). Two key reasons underscore the continued need for human characterization of protein structures. First, like any predictive AI, AlphaFold2 is not 100% accurate for all protein structures, and experiments conducted by structural biologists remain essential to validate its predictions. For example, since AlphaFold2 is trained by the existing protein structure, its accuracy can be diminished for proteins that have little similarity to previously resolved structures, such as certain membrane proteins (Azzaz et al., 2022). Second, AlphaFold2's role is confined to predicting the shape of 3D structures, while structural biologists' work extends beyond this to explore the functions and bio-mechanics explanation underlying protein structures.

2.3 Structural Biology with AI

The debut of AlphaFold2 generated significant excitement among structural biologists by enabling research on large and novel proteins that were previously difficult to resolve (Subramaniam and Kleywegt, 2022). AlphaFold2 has rapidly become a widely adopted tool in protein studies. Panel C of Appendix Figure A1 shows that over 40% of protein structure papers in the top three natural science journals cited AlphaFold2 in 2023.

Prior to the invention of AlphaFold2, protein structure determination involved four major steps (see Panel A in Figure 1). First, structural biologists need to identify functional and stable regions of the protein to produce recombinant protein domains. This step often relies on expert judgment and trial-and-error, which can take a significant amount of time. Second, structural biologists use expensive techniques such as X-ray crystallography or cryo-electron microscopy (Cryo-EM) to gather information about the protein architecture. X-ray crystallographer first crystallize the protein, then subject it to X-ray diffraction. Cryo-EM researchers observe protein structures by rapidly freezing samples and using electron beams to capture high-resolution images. Third, structural biologists construct a 3D protein model to match the experimental results from step 2. This step is time-intensive and requires iterative refinement to ensure accuracy. Finally, structural biologists analyze the structural data to identify active sites, interaction interfaces, or conformational dynamics that explain the protein's biological role.

Panel B of Figure 1 illustrates the typical workflow for protein structure determination with the help of AI. There are two steps in the workflow where AlphaFold2 can involve. First, AlphaFold2's structural predictions play a crucial role as hypotheses, offering researchers a starting point to guide experimental validation and refinement (Goulet and Cambillau, 2022). AlphaFold2's predictions offer structural hypotheses for target proteins, enabling structural

biologists to pinpoint potential functional areas and avoid unstable domains. By leveraging these hypotheses, structural biologists can design and produce recombinant proteins more efficiently, reducing the time and resources needed.

Second, AlphaFold2 accelerates the process and improves the quality of 3D modeling for both X-ray crystallography and Cyro EM. For example, X-ray crystallography often encounters the "phase problem," where information loss during diffraction complicates reconstruction. While molecular replacement (MR) can address this issue, successful MR typically requires a previously solved structure of a similar protein with at least 30% amino acid sequence similarity (Kim, 2023). With AlphaFold2's ability to predict structures for all proteins, structural biologists can now reconstruct phases for novel proteins that lack similar solved structures (McCoy et al., 2022). Furthermore, AlphaFold2 aids Cryo-EM experimentalists in interpreting lower-resolution electron density maps. Therefore, AlphaFold2 allows scientists to refine experimental data with greater precision, enhancing the accuracy and completeness of protein structures. In general, AlphaFold2 exemplifies two central roles that AI can play for knowledge workers: (1) generating valuable hypotheses for testing, and (2) streamlining labor-intensive processes.

In addition to its benefits for structure determination, AlphaFold2's predictions have also aided structural biologists in other areas of protein structure research and even opened new avenues, including studies of protein mechanisms, protein design, protein interactions, and drug discovery. Furthermore, researchers can develop new AI models in various domains based on AlphaFold2's framework. In mid 2024, the AlphaFold team introduced AlphaFold3, a new AI algorithm designed to predict interactions among all biological molecules, with the goal of accelerating drug discovery (Abramson et al., 2024).

Beyond its complementary uses, the emergence of AlphaFold also challenges the traditional role of structural biologists in academic research. Reliance on experimental protein structure determination may decline as AlphaFold2 provides direct predictions of molecular structures. In the past, studies of biological molecules depended on structures resolved through experimental methods. Now, for many applications, the precision achieved by AlphaFold2 is sufficient, reducing the demand for experimentally solved structures. This shift may diminish the academic influence of structural biologists who focus on relatively simple proteins that AlphaFold2 can predict with high confidence, and who do not pivot toward the new AI-augmented research.

3 Data, Sample, and Empirical Strategy

3.1 Data Sources

This paper focuses on the impact of AI on scientists' academic performance and the composition of research teams. To measure these outcomes, I combine multiple data sources to construct two datasets: a panel of biologists, where academic outcomes serve as the dependent variable, and a repeated cross-sectional sample of life science papers, where team composition characteristics are the outcomes of interest. Further details on the data sources are provided in Appendix D.

3.1.1 PubMed

PubMed is a freely accessible search engine that interfaces with the MEDLINE database (Medical Literature Analysis and Retrieval System Online), encompassing the bibliographic information of a comprehensive collection of life science journals. Within PubMed, each scientific paper is assigned a unique PubMed ID (PMID). Therefore, I can use the availability of PMID to identify if a paper is related to life science.

Additionally, PubMed categorizes publications through the use of MeSH (Medical Subject Headings) terms, which are standardized keywords representing the scientific fields to which the papers pertain. It is noteworthy that these MeSH terms are assigned by database indexers rather than the authors of the papers, ensuring a consistent and controlled vocabulary across the literature. However, PubMed does not systematically offer an ID to identify authors unambiguously.

3.1.2 Protein Data Bank

Established in 1971, the Protein Data Bank (PDB) serves as a central repository for protein structures, housing over 200,000 structures as of 2025. The standard protocol for publishing a protein structure paper entails that authors deposit their characterization of the 3D structure in the PDB shortly before or concurrent with the submission to the academic journals. As a result, the PDB comprises the comprehensive collection of all publicly available protein structures. Once a paper is published in a journal, the PDB links its PMID to the corresponding deposited structure. For the remainder of the paper, I refer to PubMed papers linked to the PDB as "PDB papers." Since these papers focus on protein structure characterization, I define structural biologists as researchers who have published PDB papers. The PDB also provides detailed characteristics and quality metrics of protein structures,

including length, weight, amino acid sequence, experimental tools, and resolution. These details help in understanding the research focus of structural biologists.

3.1.3 Dimensions

In order to construct a panel of researchers' publication history, I leverage the Dimensions database (https://www.dimensions.ai/). This extensive database provides metadata of 150 million publications including author affiliation, publication, citation counts, journal, abstracts, patents and funding. Importantly, Dimensions also offers the author identifier (author ID), which enables the unambiguous tracking of an author's publication history, thus facilitating an accurate examination. Furthermore, Dimensions extracts the bibliographic information from various existing databases including PubMed. Therefore, we can link an author's Dimensions ID to the PMID of their publications to determine whether their papers appear in the PDB. Similarly, using Dimensions, we can link the PMID of PDB papers to the journals where they were published. Previous works, such as Jia et al. (2022), have demonstrated the efficacy and reliability of the Dimensions database.

3.1.4 Journal Ranking

To measure the quality of journals, I connect journal with their rankings in the *Journal Citation Reports* (JCR), the most widely used journal ranking system. JCR calculates the *Journal Impact Factor* (JIF) of a journal in year y as follows:

$$\text{JIF}_y = \frac{\text{Citations in year } y \text{ published in years } y-1 \text{ and } y-2}{\text{Number of articles published in years } y-1 \text{ and } y-2}$$

JCR then ranks journals within each field by JIF and divides them into four quartiles, with Q1 containing the highest-impact journals and Q4 the lowest.

3.1.5 ORCID

To trace researchers' employment histories, I leverage the Open Researcher and Contributor ID (ORCID), a non-profit organization that provides unique, persistent digital identifiers (ORCID iDs) for researchers. Researchers can apply for an ORCID iD, which helps disambiguate individuals with similar names and ensures accurate attribution of scholarly work across systems and institutions. Researchers can also voluntarily upload their education, employment, and research histories to ORCID, which can be accessed through ORCID's public data. A recent survey of researchers from 31 research universities shows that ORCID adop-

tion is most common in the biological and biomedical sciences, with rates reaching as high as 93% (Porter, 2025). Life scientists also have more motivation to actively use and maintain their ORCID profiles, as major grant applications increasingly require them, including the National Institutes of Health (NIH).

Dimensions provides ORCID iDs for researchers, allowing me to link individuals to their ORCID profiles and fetch their employment histories. Around 80% of the authors in my sample is linked to an ORCID iD by Dimensions.

3.2 Research Trajectory

To measure the trajectory of research, I first query Dimensions for the abstracts and MeSH terms from PubMed papers for the authors in the sample. Then, I process the text by removing common English stop words and punctuation. Finally, I apply text analysis methods to convert the cleaned text into vector representations and compute the distance between research projects using the approaches described below.

3.2.1 TF-IDF

I use Term Frequency-Inverse Document Frequency (TF-IDF) to represent the text as sparse vectors. Specifically, I apply one-gram TF-IDF, which divides the text into single terms (unigrams) and evaluates their importance within a document relative to a corpus. Term frequency (TF) quantifies how often a term appears in a document, while inverse document frequency (IDF) downweights terms that are frequent across the entire corpus. The product of TF and IDF assigns higher weights to terms that are more informative for distinguishing between documents. The output is a high-dimensional vector for each document, where each dimension corresponds to the weighted importance of a term.

3.2.2 SciBERT

In addition to the traditional vector representation provided by TF-IDF, I leverage a machine learning model for natural language processing (NLP). Specifically, I use SciBERT, a specialized Bidirectional Encoder Representations from Transformers (BERT) model trained on scientific text, to generate dense vector representations. SciBERT utilizes BERT's transformer architecture, capturing contextual and semantic relationships through a self-attention mechanism. Unlike TF-IDF, SciBERT does not rely on term frequency but instead encodes the meaning of words within the context of the entire document. This approach produces dense vector representations that allow for nuanced comparisons of abstracts and keywords. For instance, SciBERT can distinguish the term "elasticity" in an economics paper versus a

physics paper based on the surrounding context in the sentence.

3.2.3 Research Distance Measure

To measure the similarity between two papers, I calculate cosine similarity using the vector representations generated by TF-IDF or SciBERT. Cosine similarity measures the cosine of the angle between two vectors, which captures the degree of similarity between documents while remaining invariant to differences in vector length. The similarity scores range from 0 to 1, where 1 indicates identical vectors and 0 indicates orthogonality. For ease of interpretation, I define the distance as $1 - similarity$ for each pair of papers.

Next, for each author, I compute the average pairwise distance between all papers in each year t after 2016 and the baseline pool of papers from 2013 to 2016. This measure captures how distinct papers in a given year after 2016 are from those in the baseline period. A higher value indicates a bigger shift in research trajectory. If an author does not publish any papers in a given year, the distance for that year is set to zero, reflecting the absence of new research trajectory.

3.3 Sample Construction

3.3.1 A Panel of Incumbent Biologists

To evaluate AlphaFold2's influence on the scientific productivity of structural biologists, I construct an individual-level balanced panel comprising of structural biologists and comparable control life scientists based on two main criteria. First, control researchers do not engage in protein structure determination but conduct research in other life science fields [5]. Second, control and treated biologists must have comparable key baseline characteristics, including citations, publication age, coronavirus research, and affiliations. The investigation period spans from 2017 to 2023, with 2021 serving as the initial treatment year, following the emergence of AlphaFold2 in late 2020.

To meet the first criterion, I identify journals that published at least two PDB papers between 2013 and 2016. I then query Dimensions to retrieve all PubMed-indexed articles published in these journals during 2013-2016. From this set, I extract the Dimensions researcher ID of the corresponding authors for each publication[6]. I focus on corresponding authors for two reasons. First, they typically serve as the principal investigators (PIs) of the

[5] Although AI applications extend beyond structural biology in life sciences, this may result in a conservative estimate of AlphaFold2's impact. Moreover, AlphaFold2 is widely regarded as the most revolutionary AI advancement in life sciences to date.

[6] If the corresponding author is missing, I extract the last author.

research projects. Second, PIs represent the research labs, allowing me to capture lab-level impacts without double-counting individual members. Based on this list, I define structural biologists as PIs with at least two PDB papers, and non-structural biologists as PIs with at least two PubMed publications but no PDB papers.

This approach ensures that both groups are drawn from similar academic venues and excludes new entrants during the investigation period (2017–2023). Although both the treatment and control groups publish in similar journals, direct competition between them is likely minimal, as structural biology comprises only less than 2% of all publications in these outlets. This limited overlap between the treatment and control groups enables an assessment of the net impact of AI availability on the productivity of the entire structural biology labor force.

I successfully identify 3,227 structural biologists and 130,362 non-structural biologists in the author list. Using Dimensions, I collect their complete publication and funding records from 2017 to 2023. I downloaded the data in May 2025.

To meet the second criterion, I implement a "coarsened exact matching" (CEM) procedure (Iacus et al., 2012) to match structural biologists with non-structural biologists based on several key ex ante variables. Establishing comparable treated and control groups strengthens the parallel trends assumption that in the absence of treatment, the outcomes for both groups would have followed similar trajectories over time. Before the CEM process, I perform several data trimmings. First, I exclude the coauthors of structural biologists to minimize potential spillover effects. Second, I drop control authors who published PDB papers during the pre-treatment period (2017–2020). Third, I remove authors with zero publications during the pre-treatment period to focus on active researchers.

For the CEM, I begin by performing an exact match on authors' baseline affiliations during 2017–2020, ensuring comparisons between structural and non-structural biologists within the same institution. Additionally, I balance the following three key aspects using coarsened bins: 2017–2020 citations, academic age[7], and citations of COVID-related papers during 2017–2020. I successfully matched around 70% of the structural biologists and details of the matching process are provided in the Data Appendix. The final panel dataset consists of 11,398 control and 2,119 treated biologists.

The overlapping timeline of the COVID-19 pandemic with my investigation period poses a potential threat to the validity of my estimation. However, the CEM procedure described above mitigates some of the potential confounding effects of COVID-19. First, I use exact matching between treated and control biologists based on their affiliations, addressing potential biases stemming from variations in COVID-19 policies across institutions. Second, I

[7]Academic age is defined as the difference between 2020 and the year of a researcher's first publication.

perform coarsened exact matching on baseline coronavirus-related citations to account for potential bias from researchers' pre-existing focus on COVID-19, which may have attracted greater academic attention during the pandemic (Hill et al., 2025). Moreover, my study period extends through 2023, by which time COVID-19-related research had substantially declined, according to PubMed data.

3.3.2 A Repeated Cross-section Sample of Papers

To assess the impact of AI on the composition of research teams in new protein structure studies, I focus on papers published between 2017 and 2023. I collect author information from 271 journals that published at least two PDB papers during 2013–2016—the same journal list used to construct the panel of biologists. The treatment group consists of PDB papers, while the control group includes PubMed-indexed papers from the same journals that are not listed in the PDB. To examine how AI affects the age dynamics of researchers, I extract the Dimensions researcher IDs of first and corresponding authors and retrieve their publication age using data from Dimensions.

3.4 Summary Statistics

Table 1, Panel A, reports the means of baseline characteristics from 2017 to 2020, separately for treated and control biologists in the panel sample. I measure researchers' productivity only based on the projects where they serve as principal investigators (PIs). Unlike in fields such as economics or management, where coauthors often contribute equally, authorship in the natural sciences tends to be more hierarchical. PIs in the natural sciences are typically responsible for the research design and the overall direction of the project. As a result, the projects that researchers lead provide a more accurate reflection of their research productivity and academic influence.

I use several measures to capture the quantity, academic influence, and quality of work by the researchers. First, PubMed publications represents the number of publications indexed by PubMed. For example, PubMed publications for an researcher i in year t refer to the number of PubMed papers published in year t, where researcher i served as principal investigator. Second, PubMed citations represents the citation augmented publication. For example, PubMed citations in year t refer to the number of citations received by PubMed papers published in year t, measured as of May 2025 (the time of data download). While PubMed publications reflect the quantity of a researcher's output, PubMed citations serve as a proxy for its academic impact and quality. Third, the number of PubMed Q1 journals captures research quality by counting how many publications appeared in journals classified

as Q1 in the JCR. Finally, The number of grants represents the total number of research grants a researcher receives in a given year.

The research trajectory is measured using two approaches: one based on TF-IDF and the other on SciBERT embeddings. The majority of researchers in both treatment and control groups are affiliated with universities.

Although direct evidence of AlphaFold2 usage is unavailable, citations to its publications can serve as a promising proxy, as the AlphaFold team requires users of their prediction and algorithm to cite their work (Jumper et al., 2021; Varadi et al., 2022). To assess whether citations to the AlphaFold2 paper serve as a reliable proxy for actual AI usage, I randomly selected 30 papers led by structural biologists and manually examined how AlphaFold2 was used in each (see Appendix Section C). In 90% of the cases, the authors directly employed AlphaFold2's algorithm or prediction results in their research, rather than citing it solely for background, general acknowledgment, or literature review purposes. Accordingly, I classify structural biologists as AlphaFold2 users if they cite their publications. Approximately 25% of structural biologists have cited AlphaFold2 in their research, compared to around 1% of control biologists. This stark disparity suggests that AlphaFold2 is substantially more relevant to the work of structural biologists.

Panel B, reports the means of baseline characteristics from 2017 to 2020, separately for treated and control publications in the repeated cross-section sample. To assess the impact of AI on the organization and personnel of research teams in protein structure studies, I measure the number of PIs and non-PIs in the publications. To capture the effect of AI on interdisciplinary inputs and collaboration, I use the number of affiliations and the number of countries represented among the PIs on each paper.

To study the impact of AI on the age dynamics of research teams, I use the publication age of the first author and the average publication age of the PIs. Publication age is defined as the difference between the paper's publication year and the first publication year of its first author or PI. The publication age of first authors is significantly lower than that of PIs in both the treatment and control groups. On average, 10% of structure papers cited AlphaFold2 in their research after 2020, compared to around 0.5 % among control papers. To account for the potential influence of COVID-related research on these estimates, I exclude COVID-related papers from both the treatment and control groups.

3.5 Empirical Strategy

Since the main outcome variables used to measure the impact of AlphaFold2 on structural biologists—such as publications and citations—are count data, I employ a Poisson difference-

in-differences (DiD) regression for the main analysis. This approach follows the economics of science literature, such as Azoulay et al. (2019), and the recommendation of Chen and Roth (2023) for estimating treatment effects as percentage changes. Specifically, I used the following regression to estimate the impacts of AI on the incumbent structural biologists:

$$E[Y_{it}|X_{it}] = \exp\left[\sum_{\substack{t=2017 \\ t\neq 2020}}^{2023} \beta_t \cdot \text{Structural_Biologist}_i \times D_t + \alpha_i + \gamma_t \times CEM_i\right] \tag{1}$$

where Y_{it} represents the outcomes of interest, such as the number of PubMed publications and citations in year t for author i. $\text{Structural_Biologist}_i = 1$ if author i belongs to the treatment group, and $\text{Structural_Biologist}_i = 0$ otherwise. D_t is a dummy variable equal to 1 for each year t. α_i denote individual fixed effects, capturing time-invariant characteristics of each author. $\gamma_t \times \text{CEM}_i$ represents the interaction between year fixed effects and the CEM stratum fixed effect for researcher i. This interaction ensures I compared the treatment and control group within the same CEM stratum.

To capture the impact of AI on the evolution of research teams in new protein structure studies, I use the following equation in the repeated cross-section sample of publications:

$$E[Y_i|X_i] = \exp\left[\sum_{\substack{t=2017 \\ t\neq 2020}}^{2023} \beta_t \cdot \text{PDB_Paper}_i \times \mathbf{1}[D_i = t] + \text{PDB_Paper}_i + \text{Year}_i\right] \tag{2}$$

where Y_i represents the outcome of interest for paper i, such as the number of PIs and non-PIs, the number of affiliations of the PIs, and the age of the first author and PIs. $\text{PDB_Paper}_i = 1$ if paper i is listed in the PDB, and $\text{PDB_Paper}_i = 0$ otherwise. $\mathbf{1}[D_i = t]$ is an indicator that paper i was published in year t. Year_i denotes the publication year fixed effect of paper i. I do not include journal fixed effects in the regression, as they could themselves be influenced by the treatment and thus constitute a bad control.

4 Main Results

This section examines the net impact of AlphaFold2 on the academic productivity, research direction, and employment for structural biologists.

4.1 Publications and Citations

Figure 2 presents event-study DiD estimates for two key measures of productivity and academic influence: the number of PubMed publications and citations. Panel A shows that the annual number of PubMed publications by structural biologists remained stable before and after the introduction of AlphaFold2. However, Panel B reveals a different pattern for the number of citation augmented publication: structural biologists received more citations for the projects published after the availability of AlphaFold2, with no evidence of significant pre-trends. Overall, the evidence indicates that after the debut of AlphaFold2, structural biologists continued to publish the same quantity of research, but their new projects have inspired more follow-on research as measured by citations.

Table 2 demonstrates the robustness of the DiD estimates to alternative treatment definitions. Since structural biologists may work across multiple research areas — not just on structure determination — I measure treatment intensity using the proportion of PDB papers in structural biologist i's total publication profile during the baseline period (2017–2020). Table A1 presents the distribution of treatment intensity across structural biologists. Columns (1) and (5) report estimates using the preferred treatment definition: a dummy variable for all structural biologists. The results show that the debut of AlphaFold2 had no significant effect on the total number of publications, but it increased citations to post-treatment projects by $e^{0.0977} - 1 = 10.3\%$.

Columns (2) and (6) report results for those classified as highly impacted structural biologists, defined as having an above-median share of PDB papers. The estimated effects for this group are similar to those using the preferred binary treatment definition. Columns (3) and (7) focus on lowly impacted structural biologists, those with a below-median PDB share. For this group, the decrease in publication output is statistically insignificant at the 1% level, and the increase in citations is insignificant at the 6% level.

Finally, Columns (4) and (8) use the percentage of PDB papers as a continuous treatment variable. This specification yields the largest effects, with a marginally significant 5% increase in the number of publications and a 12% increase in citations. However, recent econometric studies have highlighted the difficulties in interpreting coefficients for continuous treatments, even in non-staggered difference-in-differences settings like the one in this paper (Callaway et al., 2024). Therefore, I prefer a binary treatment indicator for all structural biologists instead. Together, the estimates suggest that citations of structural biologists increased, and this effect was stronger among individuals whose research was more exposed to the introduction of AlphaFold.

As noted above, structural biologists conduct both research listed in the PDB and re-

search that is not. It is important to note that papers not listed in the PDB can still relate to protein structure studies; they just do not directly involve solving protein structures. To test whether AlphaFold2 improves the productivity of structural biologists through protein structure determination or other related research, Panels A and B of Appendix Figure A2 decompose PubMed citations of treated biologists into two categories: citations to PDB papers and citations to non-PDB papers. Since control biologists do not have PDB papers by design, I use their total citations as the comparison group for both outcomes. The results show that AlphaFold2 increased citations for structural biologists in both PDB and non-PDB papers to a similar extent. This is consistent with the fact that among the projects by structural biologists citing AlphaFold2, roughly 50% are PDB papers and 50% are non-PDB papers.

Since researchers in the control group are defined by the absence of PDB papers, another concern with my DiD identification strategy is that the control group may include researchers who do not directly solve protein structures but conduct related research that could also benefit from AlphaFold2. However, the results from Figure A2 suggest that this potential contamination would bias the estimated positive effects downward.

Although Covid-19 affected all research in the life sciences, in theory some biologists may have found it easier to pivot to Covid-related topics in order to attract greater academic attention (Hill et al., 2025). To account for the overlap between Covid-19 research and the post-treatment period, I control for Covid-related research prior to 2021 in CEM for the baseline regression model 1. To further assess the robustness of the results to Covid-19, Appendix Figure A3 reports estimates that exclude all Covid-19 related research both before and after the treatment. The positive post-treatment trend in citations remains evident and large. This estimate is more conservative, as it also removes the potential effects of AlphaFold2 in facilitating structural biologists' work on Covid-19 related studies.

4.1.1 Suggestive Evidence: Comparing AI Adopters and Non-Adopters

The DiD estimate from Equation 1 captures the average treatment effect for the entire population of structural biologists, which include both the partial-equilibrium effect of AI usage and any potential externalities. The productivity gains from the emergence of AlphaFold2 are likely to come primarily from structural biologists who adopt the new tool. To test this mechanism, Appendix Figure A4 presents DiD event-study estimates from Equation 1 separately for AlphaFold2 adopters and non-adopters among structural biologists. Prior to AlphaFold2's release, both adopters and non-adopters exhibit similar and stable trends in publication and citation counts. Following its introduction, however, their trajectories diverge: adopters experience increases in both publications and citations, while non-adopters

see a decline. Since the adoption of AlphaFold2 is voluntary for the researchers, these results should not be interpreted as causal effects of AI usage or non-usage. Nonetheless, they reveal a widening performance gap between adopters and non-adopters in the post-AI period.

One possible explanation for the zero equilibrium effect on research quantity, despite the rise in quality, is that structural biologists compete intensively for a limited number of publication slots for protein structures. This is plausible, as Panel B of Appendix Figure A1 shows that the number of protein structure publications in the relevant journals has remained relatively stable. Consequently, the increased output of AI adopters may have been offset by reduced output from non-adopters. By contrast, the "citation space" is less demand-constrained and less competitive than journal publication space. As a result, the increase in citations for AI users does not need to come entirely at the expense of non-users, allowing for a positive equilibrium effect on research citations.

Since the adoption of AlphaFold2 is endogenous, it is possible that AI adopters would have pursued more promising research agendas even absent the technology. To attenuate this concern, Appendix Table A2 reports the correlation between a PDB paper's characteristics and its likelihood of citing AlphaFold2 while controlling for author and year fixed effects. Columns (1) and (2) use a binary outcome indicating whether a PDB paper solves a novel protein structure.[8] Papers citing AlphaFold2 are 29% more likely to report novel structures without author fixed effects, and the premium remains 20% even within the same researcher's publications. Columns (3) to (6) examine the natural log of the longest and heaviest protein structure reported in each PDB paper.[9] The results show that proteins solved in AlphaFold2-citing papers are significantly longer and heavier, both with and without author fixed effects. While there are various ways to assess the importance and complexity of proteins, measures such as novelty, sequence length, and molecular weight are among the most direct indicators. Overall, the results suggest that AI is disproportionately applied to research on more novel and complex proteins, even within the same author's body of work.

Columns (7) and (8) report Poisson regression results using citation counts as the outcome. Column (7) shows that papers citing AlphaFold2 receive, on average, 75% more citations. Column (8) includes author fixed effects to account for researcher-specific factors and finds that, even within the same author's publications, AI-assisted papers receive 57% more citations. Individual fixed effects cannot fully address the endogeneity of AI adoption, since authors may selectively apply AI to more promising projects. Still, the fact that the

[8] A protein structure is classified as novel if no previously solved structure exists with at least 30% amino acid sequence similarity.

[9] Protein length reflects the number of amino acids in a polymer entity, while weight is measured in kilodaltons (kDa). Because a single PDB paper may contain multiple protein structures, I use the longest and heaviest as representative values.

citation premium remains large and statistically significant indicates that the benefits of AI extend beyond just researcher-specific attributes.

Previous literature has examined the conditions under which first-mover advantages arise (Lieberman and Montgomery, 1988; Arora et al., 2023). To test whether such an advantage exists for the adoption of AI among structural biologists, Appendix Table A5 presents an event study comparing early adopters of AlphaFold2 (those citing it before 2023) with late adopters (those first citing it in 2023). The timing of publication and citation spikes aligns with the adoption of AlphaFold2. However, the productivity spike for late adopters in 2023 is even larger than that of early adopters, although the difference is not statistically significant. This pattern suggests that early adopters do not enjoy a distinct advantage over later adopters—at least so far. One alternative explanation could be that we are still in the early stage of AI diffusion in structural biology, so adoption in 2023 may not yet constitute "late" adoption in a meaningful sense.

4.2 Journal Quality, Research Direction, and Employment

Although citations are the most direct measure of research quality and influence on follow-on work, they can be sometimes noisy and affected by non-quality factors. To further assess the impact of AI on research quality and academic recognition for structural biologists, Panel A of Figure 3 presents the effects of AlphaFold2 on their publications in high-impact-factor (IF) journals. High-IF journals are defined as those in Quartile 1 (Q1) of the Journal Citation Reports (JCR). We observe a noticeable increase in the number of high-IF publications by treated researchers in 2022 and 2023. Column (1) of Table 3 confirms this visual pattern and shows that structural biologists published 4% more papers in high-IF journals following the availability of AlphaFold2. These results suggest that editors and referees at top journals also recognized the increasing quality of structural biologists' work in the post-AI period.

A key determinant of whether a technology will augment labor is whether it facilitates a new direction of new products and services (Acemoglu and Restrepo, 2020; Hirvonen et al., 2022). If AlphaFold2 merely replace the certain tasks by structural biologists and does not induce new direction of research, it will likely to reduce their academic influence. Thus, one possible reason structural biologists experienced increased citations after the debut of AI is that AI enabled them to explore new research areas. To test this mechanism, Panels B and C of Figure 3 present event-study estimates of how AI affected the direction of scientific output. The outcome variables measure the distance between post-2016 research and baseline projects from 2013–2016, using TF-IDF and SciBERT applied to publication abstracts, titles, and MeSH terms. Higher values indicate greater divergence from earlier work, reflecting larger

shifts in research trajectory. Because the measures of research distance are continuous and include zeros, I estimate their effects using OLS regressions on the raw values.

Panel B shows that, following the introduction of AlphaFold2, structural biologists' research became more distinct from their baseline work, as measured by TF-IDF scores, although the estimates are accompanied by relatively wide standard errors. Column (2) of Table 3 indicates that, post-treatment, structural biologists pursued research that was marginally more distant—by approximately 1.5% (0.011/0.718)—from their baseline work. Panel C of Figure 3 shows a statistically more precise shift in research direction post-treatment when measured using SciBERT. Column (3) of Table 3 confirms the visual pattern, showing a significant 3.4% (0.004/0.116) shift in research direction following the introduction of AlphaFold2. SciBERT, a deep learning model pre-trained on scientific text, captures semantic and contextual meaning, making it more effective for identifying nuanced innovation in scientific language than TF-IDF, which relies solely on word frequency. Taken together, the results show that the emergence of AlphaFold2 has shifted the research trajectories of structural biologists.

Panel D of Figure 3 examines the impact of AlphaFold2 on the number of grants awarded to structural biologists. In the natural sciences, grants play a critical role in enabling research, advancing academic careers, and supporting retention in academia. If AlphaFold2 were replacing research by structural biologists, we would expect a decline in grant awards. However, the event-study estimates show no observable change in grant activity, and the DiD estimates reported in Column (4) of Table 3 confirm that the effect is small and statistically insignificant.

Panel E of Figure 3 investigates whether the effect of AI on research productivity translates into changes in career trajectories for structural biologists. Given the potential of AlphaFold2 to transform the drug development industry, there is a possibility of brain drain from academia to industry. For instance, Gofman and Jin (2024) documents significant movement of AI faculty from universities to industry during periods of rapid AI advancement. To test for this, I examine a subsample of researchers who have uploaded employment histories to ORCID and restrict the sample to those affiliated with academic institutions as of 2020. The results show no significant change in the likelihood of remaining in academia following the release of AlphaFold2. The pooled DiD estimates reported in Column (5) of Table 3 also confirm that AI has had no measurable effect on academic retention.

Finally, Panel F of Figure 3 examines the impact of AI on the extensive margin of research. While earlier analyses of publication counts reflect both extensive and intensive margins, it remains possible that AlphaFold2 could have reduced participation by some researchers while increasing the productivity of others. However, the event-study plot shows

a flat trend, and the small, statistically insignificant estimate in Column (6) of Table 3 confirms that AlphaFold2 had no discernible effect on the extensive margin. These findings suggest that AI did not lead structural biologists to exit academic research.

5 Distributional Effects of AI

5.1 Heterogeneous Impacts by Structural Biologist Skill

Understanding whether the emergence of AI democratizes skill is critical, as it carries important implications for inequality. The existing literature presents mixed evidence on the skill-biased effects of AI. Many studies find that AI can narrow the productivity gap between high- and lower-skill workers in tasks such as writing, customer support, and driving (Brynjolfsson et al., 2025; Noy and Zhang, 2023; Peng et al., 2023; Kanazawa et al., 2025; Choi et al., 2024). However, Otis et al. (2023) found that AI may widen productivity disparities in non-routine tasks, such as managing small businesses. In this context, we examine whether AI amplifies performance differences between high- and low-skill scientists in the field of research.

To measure researcher skill, I use baseline citation counts from 2017 to 2020, adjusting for academic age. Specifically, structural biologists are grouped into mutually exclusive five-year age bins, and those with above-median citations within their age group are classified as highly-cited, while those below the median are classified as less-cited. Figure 4 presents event-study estimates separately for highly cited and less cited structural biologists. Panel A shows that highly cited researchers gradually increased their number of PubMed publications following the debut of AlphaFold2, while less cited researchers experienced a slight decline. Panel B reveals that AlphaFold2 significantly increased PubMed citations for highly cited researchers, with no evidence of pre-trends, whereas less cited researchers exhibited a relatively flat citation trajectory post-treatment. The results are similar when we classify structural biologists according to their unconditional raw baseline citations (see Appendix Figure A6).

Table 4 summarizes the heterogeneous impacts of AI on various labor outcomes by researcher skill level. Compared to their highly cited counterparts, less cited structural biologists exhibit significantly less publication counts, citation rates, shift of research direction measured by TF-IDF, and likelihood of publishing. Although high-skill structural biologists produce significantly more research in high-impact-factor journals and demonstrate greater shift in research trajectory measured by SciBERT, we lack sufficient statistical power to detect significant differences in AI's impact between the two groups. For grants and academic

retention, there is no evidence of significant effects for highly cited researchers, nor meaningful differences between them and their less cited peers. Overall, the results indicate that the introduction of AlphaFold increased inequality among structural biologists, primarily benefiting researchers who were already highly cited.

5.1.1 AI Adoption and Increasing Productivity Gap

This section investigates why high-skill researchers benefit more from the emergence of AI than their less cited counterparts. Since the productivity gain is likely driven by the adoption of AI, as suggested by Appendix Figure A4 and Appendix Table A2, I examine how a researcher's skill level and other background characteristics predict their likelihood of adopting the new technology. Specifically, I estimate the following specification separately for each ex ante characteristic:

$$\text{Cite_AI}_i = \alpha + \beta\Omega_i + \epsilon_i \tag{3}$$

where $Cite_AI_i$ is a binary indicator equal to 1 if author i cited AlphaFold2 publications after its release. Ω_i denotes a dummy variable indicating whether author i belongs to one of the following predefined types: *Highly Cited*, *High Novelty*, *High Length*, *High Weight*, *Cryo-EM*, *High Age*, *University*, *High Rank*, *Top-3 Countries* or , *High Fund*.

Highly Cited equals 1 if the author is classified as highly cited, as defined earlier. *High Novelty* equals 1 if the share of novel proteins solved by the researcher is above the median. Solving a novel protein is typically more complex because there are no existing structural references to guide the analysis. *High Length* and *High Weight* equal 1 if the average length or molecular weight of proteins solved by the author exceeds the respective sample medians. *Cryo-EM* equals 1 if the author has used cryo-electron microscopy in their protein structure research. Cryo-EM is typically used to study complex protein structures.

High Age indicates whether the author's academic age is above the sample median. *University* equals 1 if the author is affiliated with a university. For university-affiliated researchers, I use the 2020 *QS World University Rankings in Biology*[10] to rank their institutions. *High Rank* equals 1 if the university rank is above the sample median (220). *Top-3 Countries* includes authors affiliated with institutions in the U.S., U.K., or China—the three countries with the largest number of structural biologists in the sample. *High Fund* equals 1 if the author's total funding amount (in US dollars) exceeds the sample median.

[10]The ranking includes the top 600 universities globally. Authors affiliated with institutions outside the top 600 are coded as having a rank below 600.

Column (1) of Table 5 shows that highly cited structural biologists are approximately $0.1980/0.1518 \approx 124\%$ more likely to cite AlphaFold2 than their less-cited counterparts. Columns (2)–(4) indicate that researchers with prior experience working on novel, longer, and heavier proteins are significantly more likely to adopt AlphaFold2. Column (5) shows that structural biologists with experience using advanced experimental techniques, such as cryo-EM, are more likely to cite AlphaFold2. Cryo-EM is often used to solve large and complex proteins. Thus, these results are consistent with Appendix Table A2 and indicate that researchers with experience in solving more novel, longer and larger proteins are more likely to adopt AI in their work.

Column (6) suggests that senior researchers are less likely to adopt the AI technology. Column (7) finds no significant difference in adoption between industry researchers and those at lower-ranked universities. However, researchers affiliated with higher-ranked institutions are more likely to adopt AI. Column (8) indicates that researchers from the top three countries (U.S., U.K., and China) are also more likely to cite AlphaFold2. Column (9) indicates that researchers with more funding amount are also more likely to cite AlphaFold2.

In the full model presented in Column (10), which includes all *ex ante* characteristics, being highly cited emerges as the strongest predictor of AI adoption. This indicates that skill level itself plays a distinct role in determining AI adoption. Meanwhile, prior experience with complex proteins, access to advanced equipment, and ample research resources also play a significant role in AI adoption. These findings are consistent with the previous studies that AI adoption rate is higher among more privileged workers (Humlum and Vestergaard, 2024; Bick et al., 2024; Handa et al., 2025).

To examine whether the AI adoption gap narrows over time and thereby shed light on the dynamics of citation polarization, Appendix Figure A7 plots the adoption rates of AlphaFold2 separately for highly cited and less-cited researchers. The figure shows that highly cited researchers adopted the technology more rapidly in 2022, and the adoption gap between the two groups persists through 2023. Combined with the diverging academic performance between AlphaFold2 adopters and non-adopters shown in Appendix Figure A4, the difference in adoption rates aligns with the observed increase in citation polarization.

5.1.2 Why is not AI Adoption Universal among Researchers?

The previous section shows that highly cited researchers are more likely to adopt AI. Yet, given that AlphaFold2 is freely accessible, a natural question arises: why are highly cited researchers more likely to adopt it? This section explores two possible mechanisms. First, highly cited researchers may face lower costs of integrating the new AI tool into their research. Appendix Table A2 shows that AI tools are more often used in complex protein research.

Thus, prior experience with complex proteins and advanced research equipment could make researchers better positioned to adopt AI. Second, highly cited researchers may derive greater benefits from using AI. If adoption yields higher returns for highly-cited researchers, they are correspondingly more likely to adopt the technology.

To test the first mechanism, Appendix Table A3 examines the correlation between a researcher's citation level and the complexity of protein research prior to the release of AlphaFold2. Columns (1) through (5) show that highly cited researchers were more likely to work with complex proteins, employ advanced techniques such as cryo-EM, and receive greater research funding at baseline. Shifting research directions is costly for innovators, involving both intellectual and logistical challenges (Arts and Fleming, 2018; Hill et al., 2025; Myers, 2020). Since less-cited researchers have less experience solving proteins that necessitate the AI tool, the higher costs they face in pivoting to AI-augmented protein research may help explain their lower likelihood of adopting the new technology.

It is difficult to directly test the second mechanism — whether AlphaFold2 is skill-biased, conditional on its use — in an observational setting, because adoption is an endogenous choice. However, given existing evidence that AI lowers the learning curve and may generate larger gains for lower-skill workers, I develop a simple framework in Appendix Section E, motivated by the Roy model (Roy, 1951; Chiquiar and Hanson, 2005). In this framework, lower-skill researchers may still be less likely to adopt AI, even when their productivity gains are larger conditional on adoption. This occurs when the cost of AI adoption declines with researcher skill. Thus, selection into adoption is a key determinant of whether inequality decreases or increases.

The model also helps to reconcile a common empirical pattern: while generative AI often provides greater benefits to lower-skill workers in experimental settings, surveys consistently show that more educated and higher-wage individuals are more likely to adopt these tools. In experimental settings where both high- and low-skill workers receive instructions on how to use AI, the cost of adoption tends to be similar across skill levels. However, in voluntary settings, higher-skill workers are typically better resourced, more informed, and more familiar with technology advancements, which can lower their cost of AI adoption. As a result, productivity disparities can widen — even when AI offers greater potential benefits to less-skill workers, if they face higher barriers to adoption. The adoption gap between high- and low-skill researchers could be larger when the implementation of AI for certain tasks requires greater expertise. These findings underscore the importance of targeted education and training programs to ensure more equitable AI uptake and to help mitigate emerging inequalities.

5.2 Heterogeneous Impacts by Background Characteristics

Aligned with the broader debate on whether AI will exacerbate inequality among workers, there is also growing concern about the democratization of AI—that is, whether its benefits are concentrated among an elite group or accessible to a wider pool of individuals. Structural biology provides an ideal setting to study this question, as it is a field that demands substantial investments in both financial and human capital.

Panels A and B of Figure 5 examine the heterogeneous effects of AlphaFold2 on the number of publications and citations, respectively. First, senior structural biologists experience a greater increase in both publications and citations compared to their junior peers, although the difference is not statistically significant for citation gains. Second, structural biologists with more baseline funding show significantly higher increases in both publications and citations relative to their less-funded counterparts. Researchers with less funding show no significant improvement in either the quantity or quality of publications, as measured by citation counts.

Third, we observe a significant increase in both publications and citations for structural biologists from academia. In contrast, those from non-academic institutions experience a significant decline in publications and no notable change in citations. One possible explanation is that academic structural biologists are more likely to study novel and complex proteins to advance fundamental knowledge, while non-academic researchers often focus on applied objectives such as drug development. However, we do not find any significant differences in outcomes between researchers from high- and low-ranked institutions.

Fourth, structural biologists based in the top-3 countries (the US, the UK and China) with the largest structural biology communities experience a significant increase in both the number of publications and citations. In contrast, structural biologists from other countries, on average, show no significant improvement in publications or citations. Since structural biology requires substantial investment in both financial and human capital, the distribution of structural biologists is highly uneven across the world, with the top 3 countries accounting for approximately 50% of the total. Although AlphaFold2 is freely accessible worldwide, researchers based in regions with larger scientific communities appear to benefit more.

Finally, structural biologists with greater experience in solving novel proteins appeared to experience larger gains in publications and citations following the introduction of AlphaFold2, relative to peers with less such experience. As we discussed earlier, because AlphaFold2 is more beneficial in novel research, it is consistent that researchers with greater experience solving novel proteins benefit more from AI. Furthermore, novel proteins, by definition lacking previously solved structural analogs, are generally more difficult for AlphaFold2 to predict

accurately, because the AI training relies heavily on existing structural templates. Therefore, the value of jobs involving protein structures that cannot be accurately predicted by AI may increase in the AI era.

Overall, the heterogeneity results are consistent with the skill-biased effect of AlphaFold2. The evidence suggests that the benefits of AlphaFold2 are not democratized but instead concentrated among elite groups of researchers—those with higher funding levels, those based in the top three countries, and more experience with complex proteins. The results are also consistent with Table 5, which shows that more privileged structural biologists are more likely to adopt AI. As the theoretical framework from Garicano and Rossi-Hansberg (2006); Ide and Talamas (2024) suggests, when the cost of acquiring general knowledge falls, the marginal value of specialized or rare knowledge rises.

6 Research Team Structure

This section examines how AlphaFold2 has affected the composition and size of research teams in protein structure studies, using Equation 2. While we do not observe a net reduction in publications by principal investigators, AlphaFold2's ability to accelerate and simplify parts of the protein structure determination process may reduce the need for certain roles within research teams, such as research staff, PhD students, and postdocs. At the same time, AlphaFold2 enables researchers to pursue larger and more complex structures, which may require larger and more skilled teams. As such, the net impact of AI on research team size is theoretically ambiguous.[11]

Panel A of Figure 6 shows no significant impact of AlphaFold2 on the number of non-PI authors involved in protein structure research during the post-treatment period. While the pooled DiD estimate in Column (1) of Table 6 is statistically significant and positive, this appears to be driven by a dip in 2018 rather than a sustained post-treatment trend. Therefore, we cannot confidently conclude that AI increases team size. Panel B indicates that the number of PIs per protein structure publication remains unchanged following the introduction of AlphaFold2, a finding also confirmed by the insignificant estimate in Column (2) of Table 6. Together, the results suggest that AlphaFold2 did not alter the ratio between PIs and non-PI team members in protein structure research.

Next, I analyze the impact of AI on the composition of researchers' institutional affiliations. Panel C of Figure 6 shows that the number of affiliations associated with PIs in

[11] Unlike Cavalli (2022), who examined how AlphaFold affected the organization of other peer computational biology labs developing similar protein prediction methods, this paper investigates how AI influences the research team of structural biology who use the tool as an input in their own scientific discovery.

protein structure research increases following the emergence of AlphaFold2. Column (3) of Table 6 indicates a 10% rise in the number of PI affiliations, suggesting that protein structure research increasingly draws on expertise from more diverse disciplinary backgrounds. Importantly, the rise in the number of affiliations is not driven by an increase in the number of PIs per team, implying that protein structure research increasingly requires individual PIs to draw on a broader set of disciplinary backgrounds, rather than relying on additive expertise through more collaboration. In contrast, Panel D shows no significant change in the number of countries represented among authors, suggesting that AlphaFold2 has not led to an increase in international collaboration.

The increased affiliation among PIs implies that the protein structure studies require more diverse knowledge inputs. According to Jones (2009), when knowledge within a field becomes more complex and competitive, the age of its contributors tends to rise. To test this prediction, I collect the publication age of first and corresponding authors with Dimensions researcher IDs from the Dimensions database. First authors are typically junior researchers who contribute significantly to the research and represent the next generation of structural researchers. Panel E of Figure 6 shows that the publication age of first authors in structural biology increases following the emergence of AlphaFold2. Panel F similarly indicates an increase in the average publication age of first authors. Columns (5) and (6) of Table 6 confirm that both first-author and PI ages rise by approximately 4% after the introduction of AI. Overall, the increased ages of major contributors also signals a more competitive field.

While the previous section examined the impact of AI on incumbent researchers, this section explores how AI influence the new research in the field. Existing evidence suggests that although the emergence of AI has not significantly changed the number of researchers involved in new protein structure studies, it has increased the demand for more interdisciplinary expertise. An older age profile among first authors and PIs of protein structure studies also suggests a more competitive environment for future researchers entering the field.

7 Conclusion

AI technology is rapidly becoming an integral part of society, transforming the way we work and live. One key distinction of AI compared to previous technologies is its encroachment on knowledge workers, who were previously considered "safe" from the automation that primarily affected manual labor. This paper is the first to analyze the causal impact of a widely adopted AI application on scientific productivity and employment. Leveraging a dataset that combines publication and employment records of individual structural biologists

and matched control researchers, I evaluate the impact of AlphaFold2 at scale, which is a Nobel Prize-winning AI tool for protein structure prediction.

The main results show that the debut of AlphaFold2 had no equilibrium effect on the quantity of research by structural biologists, but led to a positive increase in their academic influence and quality, as measured by citations and the ranking of the published journals. The introduction of AI not only enhances research quality but also shift the research trajectory for the researchers. Importantly, the productivity gains from AI have not, so far, displaced human structural biologists.

However, the emergence of AlphaFold2 has widened the citation gap between highly cited and less-cited structural biologists. Consistent with the growing polarization, correlational evidence indicates that highly cited researchers adopt the new AI tool more rapidly. These high-skill researchers have more experience working on complex proteins, where AlphaFold2 provides the greatest advantages. As a result, they face lower costs when integrating AI into their ongoing research, enabling faster adoption.

In line with the skill-biased nature of AI, I also find that structural biologists with greater resources are more likely to benefit from this technology. Finally, this paper sheds light on how AlphaFold2 is reshaping the field of structural biology. The emergence of AI demands more interdisciplinary expertise of the team leaders, and increases the average age of major contributors in the research, reflecting the field's increasing competitiveness.

My paper focuses exclusively on the impacts of AlphaFold2 on structural biologists, as they are the most directly affected by this technology. However, AlphaFold2's influence extends far beyond protein structure determination. Future research could examine its impact on other areas of biology that have historically relied on human protein characterization but now benefit from computational predictions. Broadly, the impact of AI on the knowledge workforce remains unsettled, as the technology continues to evolve rapidly.

References

Abramson, J., Adler, J., Dunger, J., Evans, R., Green, T., Pritzel, A., Ronneberger, O., Willmore, L., Ballard, A. J., Bambrick, J., et al. (2024). Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, pages 1–3.

Acemoglu, D., Autor, D., Hazell, J., and Restrepo, P. (2022). Artificial intelligence and jobs: evidence from online vacancies. *Journal of Labor Economics*, 40(S1):S293–S340.

Acemoglu, D. and Restrepo, P. (2018). The race between man and machine: Implications of technology for growth, factor shares, and employment. *American economic review*, 108(6):1488–1542.

Acemoglu, D. and Restrepo, P. (2020). The wrong kind of ai? artificial intelligence and the future of labour demand. *Cambridge Journal of Regions, Economy and Society*, 13(1):25–35.

Agrawal, A., Gans, J. S., and Goldfarb, A. (2019). Artificial intelligence: the ambiguous labor market impact of automating prediction. *Journal of Economic Perspectives*, 33(2):31–50.

Agrawal, A. and Goldfarb, A. (2008). Restructuring research: Communication costs and the democratization of university innovation. *American Economic Review*, 98(4):1578–1590.

Akerman, A., Gaarder, I., and Mogstad, M. (2015). The skill complementarity of broadband internet. *The Quarterly Journal of Economics*, 130(4):1781–1824.

Arora, A., Belenzon, S., and Dionisi, B. (2023). First-mover advantage and the private value of public science. *Research Policy*, 52(9):104867.

Arts, S. and Fleming, L. (2018). Paradise of novelty—or loss of human capital? exploring new fields and inventive output. *Organization Science*, 29(6):1074–1092.

Autor, D. H. (2015). Why are there still so many jobs? the history and future of workplace automation. *Journal of economic perspectives*, 29(3):3–30.

Autor, D. H., Katz, L. F., and Krueger, A. B. (1998). Computing inequality: have computers changed the labor market? *The Quarterly journal of economics*, 113(4):1169–1213.

Azoulay, P., Fons-Rosen, C., and Zivin, J. S. G. (2019). Does science advance one funeral at a time? *American Economic Review*, 109(8):2889–2920.

Azzaz, F., Yahi, N., Chahinian, H., and Fantini, J. (2022). The epigenetic dimension of protein structure is an intrinsic weakness of the alphafold program. *Biomolecules*, 12(10):1527.

Baek, M., DiMaio, F., Anishchenko, I., Dauparas, J., Ovchinnikov, S., Lee, G. R., Wang, J., Cong, Q., Kinch, L. N., Schaeffer, R. D., et al. (2021). Accurate prediction of protein structures and interactions using a three-track neural network. *Science*, 373(6557):871–876.

Bartel, A., Ichniowski, C., and Shaw, K. (2007). How does information technology affect productivity? plant-level comparisons of product innovation, process improvement, and worker skills. *The quarterly journal of Economics*, 122(4):1721–1758.

Berman, E., Bound, J., and Griliches, Z. (1994). Changes in the demand for skilled labor within us manufacturing: evidence from the annual survey of manufactures. *The quarterly journal of economics*, 109(2):367–397.

Bick, A., Blandin, A., and Deming, D. J. (2024). The rapid adoption of generative ai. Technical report, National Bureau of Economic Research.

Brynjolfsson, E., Li, D., and Raymond, L. (2025). Generative ai at work. *The Quarterly Journal of Economics*, page qjae044.

Callaway, B., Goodman-Bacon, A., and Sant'Anna, P. H. (2024). Difference-in-differences with a continuous treatment. Technical report, National Bureau of Economic Research.

Caplin, A., Deming, D. J., Li, S., Martin, D. J., Marx, P., Weidmann, B., and Ye, K. J. (2024). The abc's of who benefits from working with ai: Ability, beliefs, and calibration. Technical report, National Bureau of Economic Research.

Cavalli, G. (2022). How scientific organizations react to novel methodological advances: the impact of alphafold v1. *Available at SSRN 4286478*.

Chen, J. and Roth, J. (2023). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics*, page qjad054.

Chiquiar, D. and Hanson, G. H. (2005). International migration, self-selection, and the distribution of wages: Evidence from mexico and the united states. *Journal of political Economy*, 113(2):239–281.

Choi, J. H., Monahan, A. B., and Schwarcz, D. (2024). Lawyering in the age of artificial intelligence. *Minn. L. Rev.*, 109:147.

Crépon, B., Duflo, E., Gurgand, M., Rathelot, R., and Zamora, P. (2013). Do labor market policies have displacement effects? evidence from a clustered randomized experiment. *The quarterly journal of economics*, 128(2):531–580.

Dell'Acqua, F., McFowland III, E., Mollick, E. R., Lifshitz-Assaf, H., Kellogg, K., Rajendran, S., Krayer, L., Candelon, F., and Lakhani, K. R. (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper*, (24-013).

Ding, W. W., Levin, S. G., Stephan, P. E., and Winkler, A. E. (2010). The impact of information technology on academic scientists' productivity and collaboration patterns. *Management Science*, 56(9):1439–1461.

Doshi, A. R., Chai, S., and Troebinger, M. (2024). How experience moderates the impact of generative ai ideas on the research process. *Available at SSRN 5013086*.

Feigenbaum, J. and Gross, D. P. (2024). Answering the call of automation: How the labor market adjusted to mechanizing telephone operation. *The Quarterly Journal of Economics*, 139(3):1879–1939.

Furman, J. L. and Stern, S. (2011). Climbing atop the shoulders of giants: The impact of institutions on cumulative research. *American Economic Review*, 101(5):1933–1963.

Furman, J. L. and Teodoridis, F. (2020). Automation, research technology, and researchers' trajectories: Evidence from computer science and electrical engineering. *Organization Science*, 31(2):330–354.

Garicano, L. and Rossi-Hansberg, E. (2006). Organization and inequality in a knowledge economy. *The Quarterly journal of economics*, 121(4):1383–1435.

Gofman, M. and Jin, Z. (2024). Artificial intelligence, education, and entrepreneurship. *The Journal of Finance*, 79(1):631–667.

Goulet, A. and Cambillau, C. (2022). Present impact of alphafold2 revolution on structural biology, and an illustration with the structure prediction of the bacteriophage j-1 host adhesion device. *Frontiers in Molecular Biosciences*, 9:907452.

Handa, K., Tamkin, A., McCain, M., Huang, S., Durmus, E., Heck, S., Mueller, J., Hong, J., Ritchie, S., Belonax, T., et al. (2025). Which economic tasks are performed with ai? evidence from millions of claude conversations. *arXiv preprint arXiv:2503.04761*.

Hill, R. and Stein, C. (2025a). Race to the bottom: Competition and quality in science. *The Quarterly Journal of Economics*, page qjaf010.

Hill, R. and Stein, C. (2025b). Scooped! estimating rewards for priority in science. *Journal of Political Economy*, 133(3):000–000.

Hill, R., Yin, Y., Stein, C., Wang, X., Wang, D., and Jones, B. F. (2025). The pivot penalty in research. *Nature*, pages 1–8.

Hirvonen, J., Stenhammar, A., and Tuhkuri, J. (2022). *New evidence on the effect of technology on employment and skill demand*. Taloustieto Oy.

Humlum, A. and Vestergaard, E. (2024). The adoption of chatgpt. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2024-50).

Iacus, S. M., King, G., and Porro, G. (2012). Causal inference without balance checking: Coarsened exact matching. *Political analysis*, 20(1):1–24.

Ide, E. and Talamas, E. (2024). Artificial intelligence in the knowledge economy. In *Proceedings of the 25th ACM Conference on Economics and Computation*, pages 834–836.

Jia, R., Roberts, M. E., Wang, Y., and Yang, E. (2022). The impact of us-china tensions on us science. Technical report, National Bureau of Economic Research.

Jones, B. F. (2009). The burden of knowledge and the "death of the renaissance man": Is innovation getting harder? *The Review of Economic Studies*, 76(1):283–317.

Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., et al. (2021). Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589.

Kanazawa, K., Kawaguchi, D., Shigeoka, H., and Watanabe, Y. (2025). Ai, skill, and productivity: The case of taxi drivers. *Management Science*.

Katz, L. F. and Murphy, K. M. (1992). Changes in relative wages, 1963–1987: supply and demand factors. *The quarterly journal of economics*, 107(1):35–78.

Kim, S. (2023). Shortcuts to innovation: The use of analogies in knowledge production.

Kovalevskiy, O., Mateos-Garcia, J., and Tunyasuvunakool, K. (2024). Alphafold two years on: Validation and impact. *Proceedings of the National Academy of Sciences*, 121(34):e2315002121.

Lieberman, M. B. and Montgomery, D. B. (1988). First-mover advantages. *Strategic management journal*, 9(S1):41–58.

McCoy, A. J., Sammito, M. D., and Read, R. J. (2022). Implications of alphafold2 for crystallographic phasing by molecular replacement. *Acta Crystallographica Section D: Structural Biology*, 78(1):1–13.

Michaels, G., Natraj, A., and Van Reenen, J. (2014). Has ict polarized skill demand? evidence from eleven countries over twenty-five years. *Review of Economics and Statistics*, 96(1):60–77.

Myers, K. (2020). The elasticity of science. *American Economic Journal: Applied Economics*, 12(4):103–134.

Noy, S. and Zhang, W. (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Available at SSRN 4375283*.

Otis, N., Clarke, R. P., Delecourt, S., Holtz, D., and Koning, R. (2023). The uneven impact of generative ai on entrepreneurial performance. *Available at SSRN 4671369*.

Peng, S., Kalliamvakou, E., Cihon, P., and Demirer, M. (2023). The impact of ai on developer productivity: Evidence from github copilot. *arXiv preprint arXiv:2302.06590*.

Porter, S. R. (2025). Understanding orcid adoption among academic researchers. *Scientometrics*, 130(5):2783–2797.

Restrepo, P. (2023). Automation: Theory, evidence, and outlook.

Rosen, S. (1981). The economics of superstars. *The American economic review*, 71(5):845–858.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford economic papers*, 3(2):135–146.

Subramaniam, S. and Kleywegt, G. J. (2022). A paradigm shift in structural biology. *Nature Methods*, 19(1):20–23.

Teodoridis, F. (2018). Understanding team knowledge production: The interrelated roles of technology and expertise. *Management Science*, 64(8):3625–3648.

Terwilliger, T. C., Liebschner, D., Croll, T. I., Williams, C. J., McCoy, A. J., Poon, B. K., Afonine, P. V., Oeffner, R. D., Richardson, J. S., Read, R. J., et al. (2023). Alphafold pre-

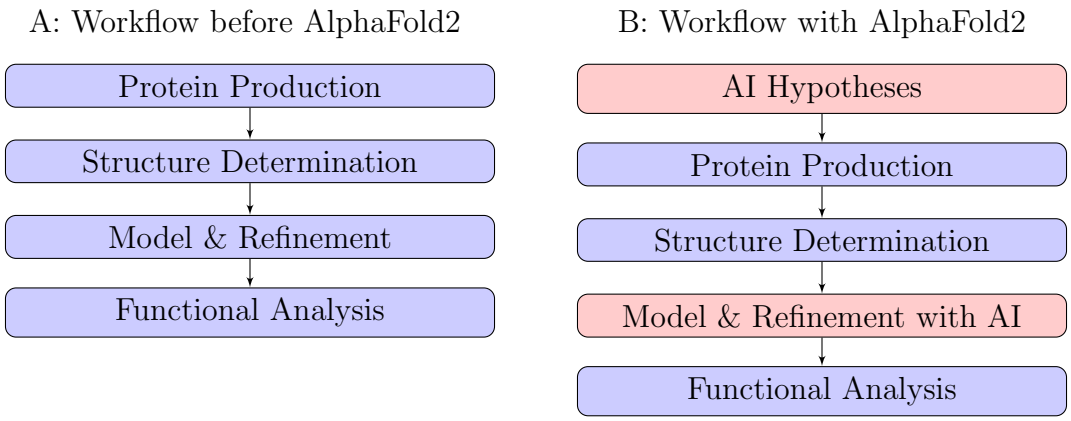
dictions are valuable hypotheses and accelerate but do not replace experimental structure determination. *Nature Methods*, pages 1–7.

Varadi, M., Anyango, S., Deshpande, M., Nair, S., Natassia, C., Yordanova, G., Yuan, D., Stroe, O., Wood, G., Laydon, A., et al. (2022). Alphafold protein structure database: massively expanding the structural coverage of protein-sequence space with high-accuracy models. *Nucleic acids research*, 50(D1):D439–D444.

Wang, W., Gao, G., and Agarwal, R. (2023). Friend or foe? teaming between artificial intelligence and workers with variation in experience. *Management Science*.

Figures & Tables

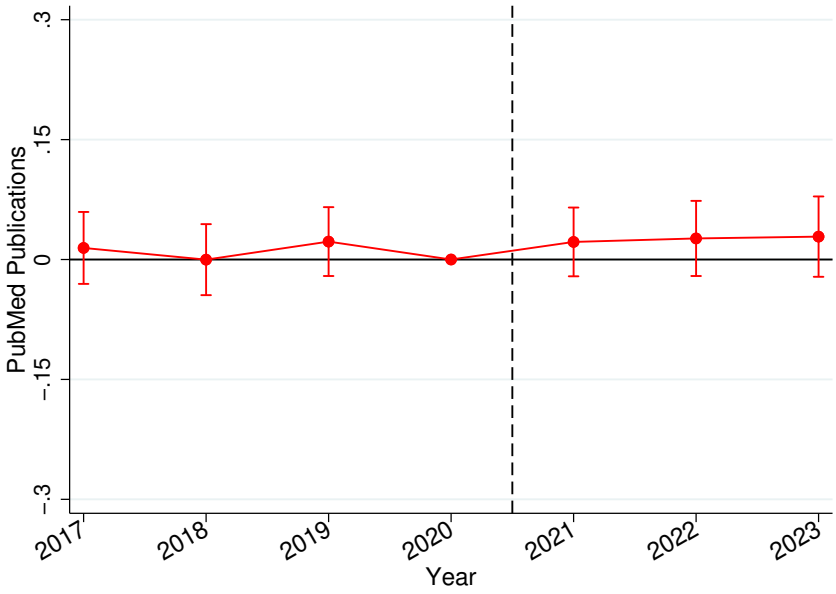
Figure 1: Protein Structure Determination Workflow



Note: Panel A illustrates the workflow of protein structure determination prior to the advent of AI. Panel B depicts the workflow of protein structure determination with AI assistance.

Figure 2: Impacts on PubMed Publications and Citations

A: PubMed Publications

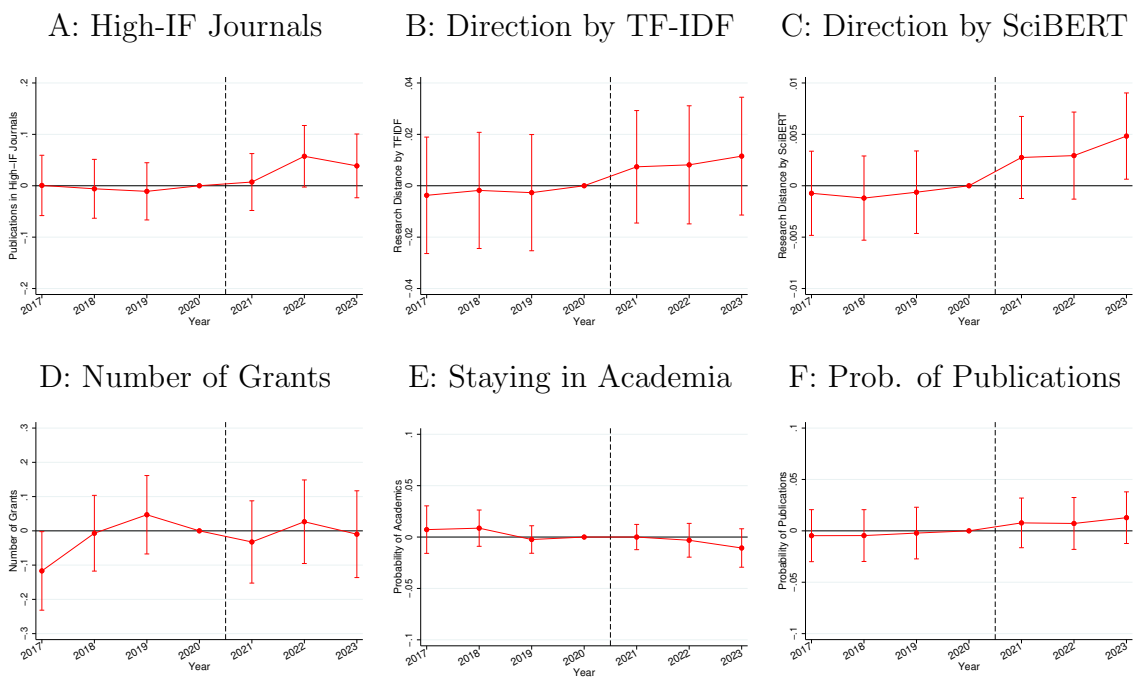


B: PubMed Citations



Note: These figures plot the DiD event-study estimates of AI's impact on structural biologists' academic performance. Panel A reports the impacts on the number of PubMed publications. Panel B focuses on the impacts on the citations. All models are Poisson regressions that include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Each vertical bar represents the 95% confidence interval of the estimate.

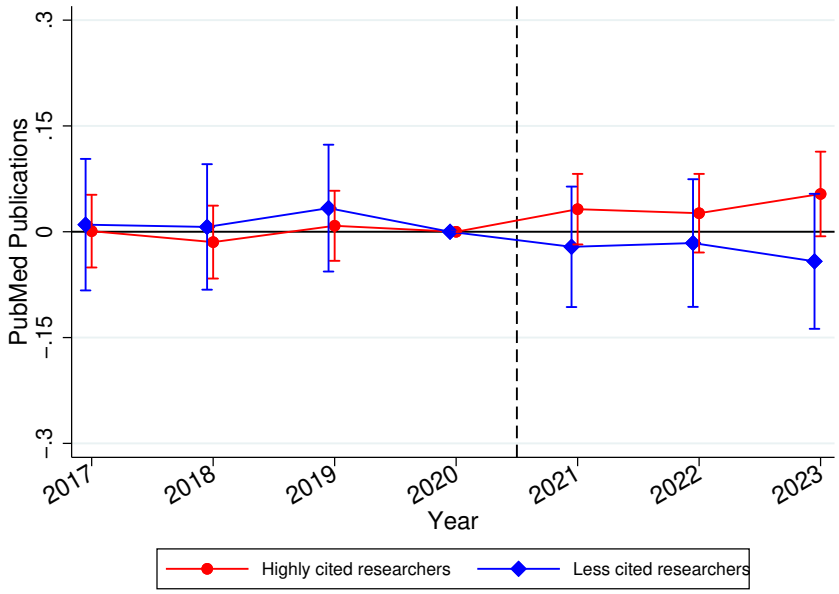
Figure 3: Impacts on Productivity, Research Direction and Employment



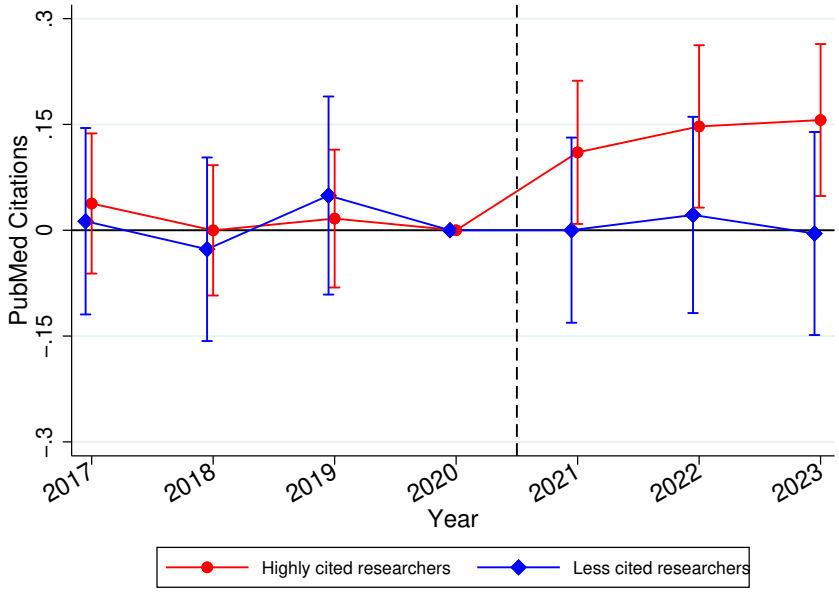
Note: These figures plot the DiD event-study estimates of AI's impact on structural biologists' academic performance and employment. Panel A reports estimates on the number of publication in high IF journals. Panel B shows the estimates on the shift of research direction measured by TF-IDF. Panel C shows the impacts on the shift of research direction measured by SciBERT. Panel D shows the impacts on the number of grants. Panel E shows the probability of staying in academia. Panel F examines the impacts on the probability of publication. All models are Poisson regressions that include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Each vertical bar represents the 95% confidence interval of the estimate.

Figure 4: Heterogeneous Impacts: Highly Cited vs Less Cited

A: PubMed Publications

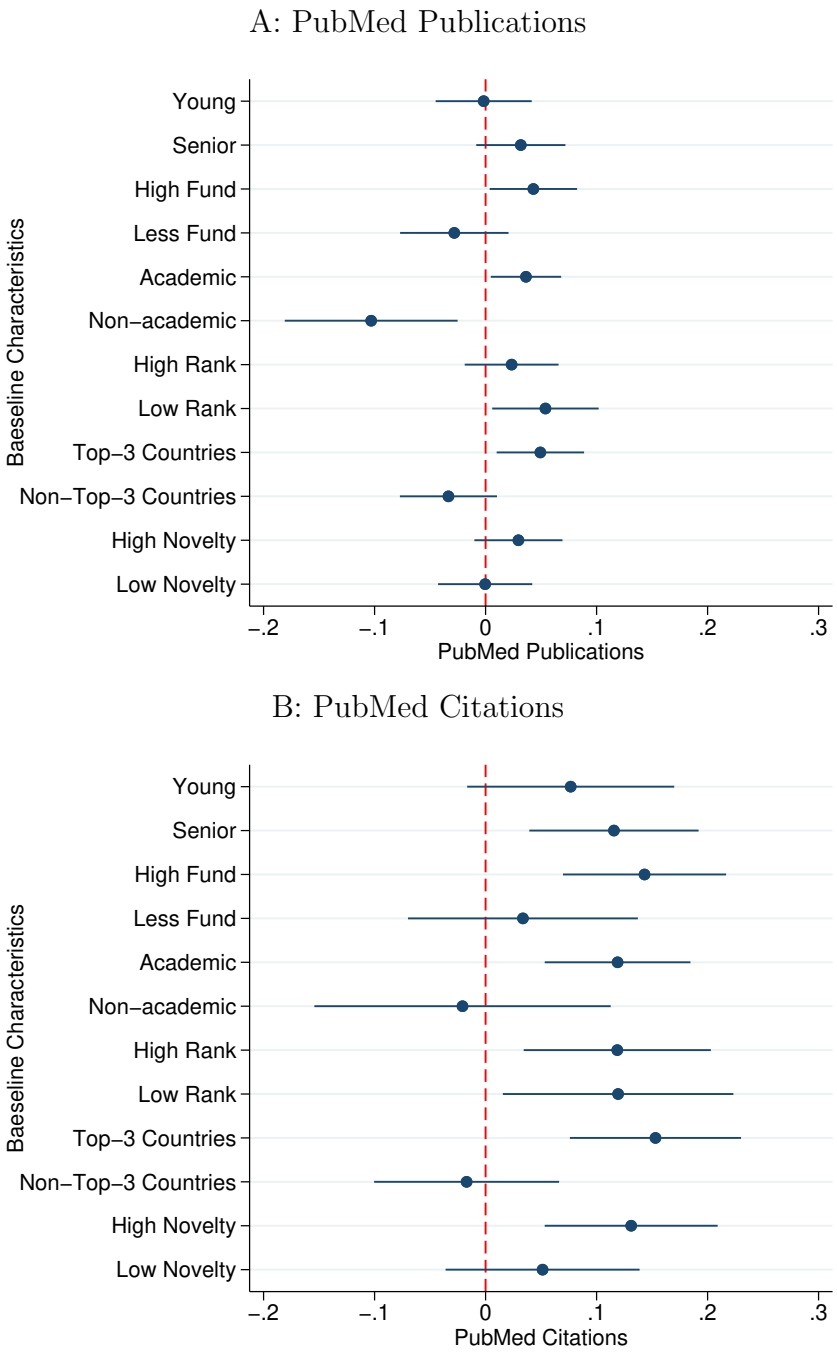


B: PubMed Citations



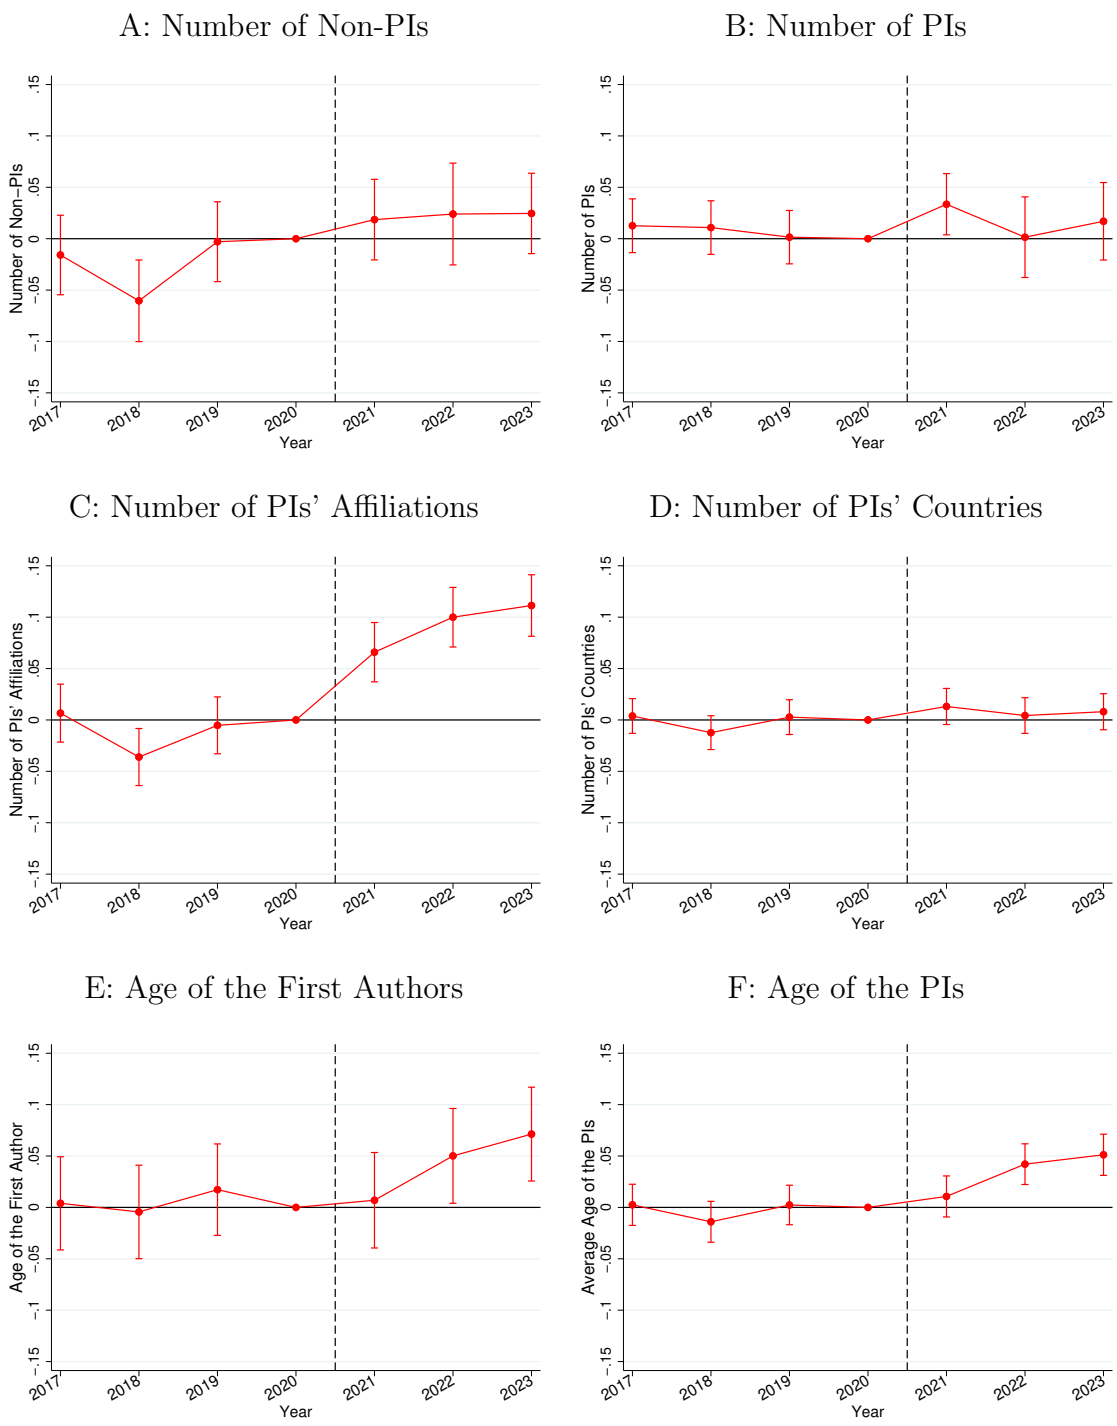
Note: These figures compare the impacts of AI on academic performance between highly cited and less cited researchers. Panels A and B present the outcomes for PubMed publications and PubMed citations. All models include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Each vertical bar represents the 95% confidence interval of the estimate.

Figure 5: Heterogeneous Impacts on Research Productivity



Note: These figures plot the heterogeneous impacts of AI on labor productivity: young vs senior, highly funded vs less funded, academic vs non-academic; high rank vs low rank, top-3 countries vs other countries, and more experience with novel proteins vs less. All models are Poisson regressions that include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Each vertical bar represents the 95% confidence interval of the estimate.

Figure 6: Impacts on the Research Team



Note: These figures plot the DiD event-study estimates of AI's impact on the team of structural biology. Panel A shows the impact on the number of non-PI authors, Panel B shows the impact on the number of PIs, Panel C shows the impacts on the number of affiliations of PIs, Panel D shows the number of countries of PIs, Panel E shows the impacts on the age of the first authors and Panel F shows the impacts on the age of the corresponding authors. All models are Poisson regressions that include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Each vertical bar represents the 95% confidence interval of the estimate.

Table 1: Summary Statistics

	Treatment		Control	
Baseline (2017-2020)	Mean	St. Dev.	Mean	St. Dev.
Panel A: Incumbent Biologists				
PubMed Publications	2.4745	2.2465	2.1311	2.1152
PubMed Citations	109.3754	162.5976	98.6775	158.8471
PubMed Q1 Journals	1.4853	1.6323	1.2753	1.4851
Number of Grants	0.4315	0.5658	0.3747	0.5379
PubMed TF-IDF	0.7184	0.2357	0.6700	0.2458
PubMed SciBERT	0.1157	0.0428	0.1193	0.0495
Academic Age	27.2690	10.1418	25.0405	10.4765
University Scientist	0.7782	0.4156	0.8924	0.3098
AlphaFold2 User (Ex Post)	0.2501	0.4332	0.0140	0.1177
Number of obs	2119		11398	
Panel B: PubMed Publications				
Number of Non-PI	8.1160	6.5774	5.9163	8.7657
Number of PI	1.4968	0.6849	1.2639	2.3632
Number of Affiliations	1.6098	0.9304	1.5528	1.4664
Number of Countries	1.1513	0.4052	1.0982	0.3762
Pub Age of First Author	7.1707	6.8034	7.6141	8.1925
Pub Age of PIs	22.2492	9.2785	18.4637	10.2316
Cite AlphaFold2 (Ex Post)	0.1057	0.3075	0.0055	0.0741
Number of obs	23824		1353451	

Note: This table summarizes the baseline (2017–2020) characteristics. Panel A shows the summary statistics for the researchers from the panel sample. Panel B shows the summary statistics for the papers from the repeated cross-section sample. All values represent the averages of the raw data. Columns (1) and (2) provide the annual averages and standard deviations for treatment group, while Columns (3) and (4) present the corresponding values for control group.

Table 2: Impacts on PubMed Publications and Citations

	PubMed Publication				PubMed Citations			
	(1) Total	(2) Highly Impacted	(3) Lowly Impacted	(4) Continuous	(5) Total	(6) Highly Impacted	(7) Lowly Impacted	(8) Continuous
Treatment × Post	0.0162	0.0245	−0.0139	0.0517*	0.0977***	0.1068***	0.0578	0.1211**
	(0.0150)	(0.0169)	(0.0302)	(0.0297)	(0.0304)	(0.0347)	(0.0567)	(0.0574)
Log likelihood	−135208.91	−109626.88	−116114.27	−135208.04	−2292317.61	−1885034.29	−2007488.54	−2292728.01
Number of obs	93419	74953	80574	93419	93126	74786	80303	93126
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CEM Matching	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: This table presents the DiD regression estimates on the impact of AI on structural biologists' academic performance. Column (1) to (4) report the estimates for PubMed publications using various treatment definitions: a dummy variable for all structural biologists, a dummy for highly impacted researchers, a dummy for lowly impacted researchers, and a continuous treatment based on the percentage of PDB publications at baseline. Columns (5) to (8) present the estimates for PubMed citations using the same treatment definitions as in Columns (1) to (4). All models are Poisson regressions that include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Statistical significance: $^{***}p < 0.01,^{**}p < 0.05,^{*}p < 0.1$.

Table 3: Impacts on Research Productivity, Direction, and Employment

	(1) High-IF Journals	(2) TF-IDF	(3) SciBERT	(4) Grant	(5) Academic	(6) Prob of Pub
Treatment × Post	0.0375**	0.0111*	0.0041***	0.0135	−0.0080	0.0121*
	(0.0185)	(0.0066)	(0.0012)	(0.0344)	(0.0094)	(0.0073)
Log likelihood/ R^2	−103762.76	0.4678	0.4768	−43176.71	0.8666	0.4539
Number of obs	85084	94619	94619	51195	27524	94619
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
CEM Matching	Yes	Yes	Yes	Yes	Yes	Yes

Note: This table presents the DiD regression estimates of the impact of AI on structural biologists' productivity, employment and research direction. Columns (1) to (6) show the impacts on the number of publication in high-IF journals, the research direction measure by TF-IDF, the research direction measure by SciBERT, the number of grants, retention in academia and the probability of publishing. All models are Poisson regressions that include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Statistical significance: $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.

Table 4: Heterogeneous Impacts: Highly Cited vs. less Cited Researchers

	(1) Publications	(2) Citations	(3) High-IF Journals	(4) TF-IDF	(5) SciBERT	(6) Grant	(7) Academic	(8) Prob of Pub
Treatment × Post	0.0366**	0.1161***	0.0515**	0.0278***	0.0060***	−0.0189	−0.0028	0.0296***
	(0.0179)	(0.0346)	(0.0212)	(0.0077)	(0.0014)	(0.0426)	(0.0122)	(0.0083)
Treatment × Post × Less Cited	−0.0799**	−0.1280**	−0.0691	−0.0326**	−0.0035	0.1002	−0.0118	−0.0341**
	(0.0327)	(0.0556)	(0.0430)	(0.0132)	(0.0024)	(0.0730)	(0.0193)	(0.0144)
Log likelihood/ R^2	−135202.76	−2291921.38	−103758.22	0.4679	0.4768	−43175.56	0.8667	0.4540
Number of obs	93419	93126	85084	94619	94619	51195	27524	94619
Individual FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
CEM Matching	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Note: This table presents the heterogeneous impact of AI by skill level on structural biologists' academic performance, measured by publications and citations. Columns (1) to (8) show the impacts on the number publications, citation counts, the number of publication in high-IF journals, the research direction measured by TF-IDF and SciBERT, the number of grants, retention in academia and the probability of publishing. All models are Poisson regressions that include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Statistical significance: $^{***}p < 0.01, ^{**}p < 0.05, ^{*}p < 0.1$.

Table 5: Characteristics of AlphaFold2 Users

	AI User									
Baseline Characteristics	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Highly Cited	0.1980***									0.1417***
	(0.0183)									(0.0193)
High Novelty		0.1816***								0.1266***
		(0.0188)								(0.0195)
High Length			0.1135***							0.0259
			(0.0187)							(0.0687)
High Weight				0.1121***						0.0051
				(0.0187)						(0.0685)
Cyro EM					0.1866***					0.0705**
					(0.0258)					(0.0276)
High Age						−0.0493***				−0.0461**
						(0.0188)				(0.0181)
University							0.0320			0.0261
							(0.0243)			(0.0232)
High Rank							0.0439**			−0.0113
							(0.0205)			(0.0203)
Top 3 Country								0.0421**		−0.0109
								(0.0189)		(0.0185)
High Fund									0.1475***	0.0861***
									(0.0186)	(0.0194)
Constant	0.1518***	0.1785***	0.1934***	0.1942***	0.2214***	0.2753***	0.2075***	0.2273***	0.1764***	0.0742***
	(0.0129)	(0.0118)	(0.0132)	(0.0132)	(0.0101)	(0.0134)	(0.0200)	(0.0139)	(0.0131)	(0.0254)
R^2	0.0522	0.0420	0.0172	0.0167	0.0242	0.0032	0.0045	0.0023	0.0290	0.1002
Number of obs	2119	2119	2119	2119	2119	2119	2119	2119	2119	2119

This table presents the OLS correlation between the usage of AlphaFold2 and the baseline characteristics of structural biologists. The outcome variable equals 1 if the structural biologist has cited AlphaFold2. Statistical significance: $^{***}p < 0.01,^{**}p < 0.05,^{*}p < 0.1$.

Table 6: Impacts on the Research Team

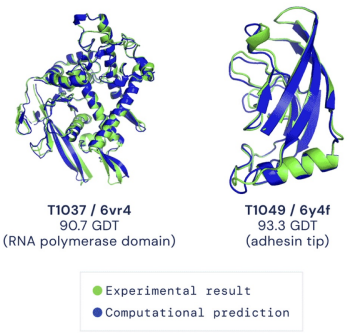
	(1) Number of Non-PI	(2) Number of PI	(3) Number of Affiliations	(4) Number of Countries	(5) Age of First Author	(6) Age of PI
Treatment × Post	0.0416***	0.0112	0.1010***	0.0098**	0.0380***	0.0366***
	(0.0119)	(0.0097)	(0.0080)	(0.0048)	(0.0129)	(0.0055)
Log likelihood	−5251632.52	−1784301.56	−1840608.64	−1436642.52	−6847578.04	−7056375.86
Number of obs	1377275	1377275	1347392	1340694	1296446	1348781
Structure FE	Yes	Yes	Yes	Yes	Yes	Yes
Year × Journal FE	Yes	Yes	Yes	Yes	Yes	Yes

Note: This table presents the impact of AI on structural biologists' team size and age. Column (1) reports the impacts on the number of non-PI members; Column (2) focuses on the number of PI; Column (3) examines the the number of affiliations of PIs; Column (4) looks at the number of countries of PIs; Column (5) examines the age of the first author; and Column (6) focuses on the average age of the PIs. All models are Poisson regressions that include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Statistical significance: $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.

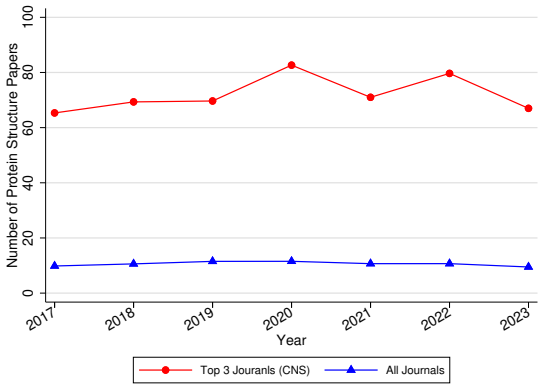
Appendix A Additional Figures and Tables

Figure A1: AlphaFold2 and Protein Structure Papers

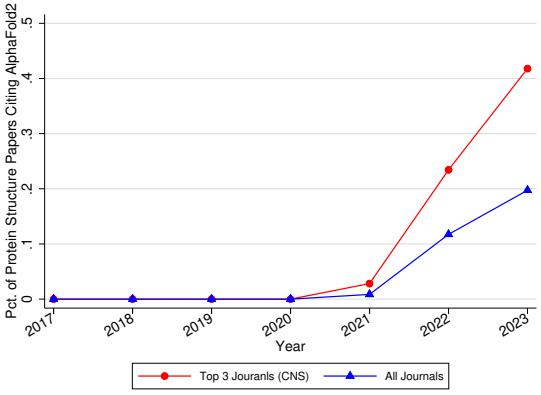
A: AlphaFold2's Prediction of CASP14 Target Protein



B: Protein Structure Papers



C: Proteins Structure Papers Using AI



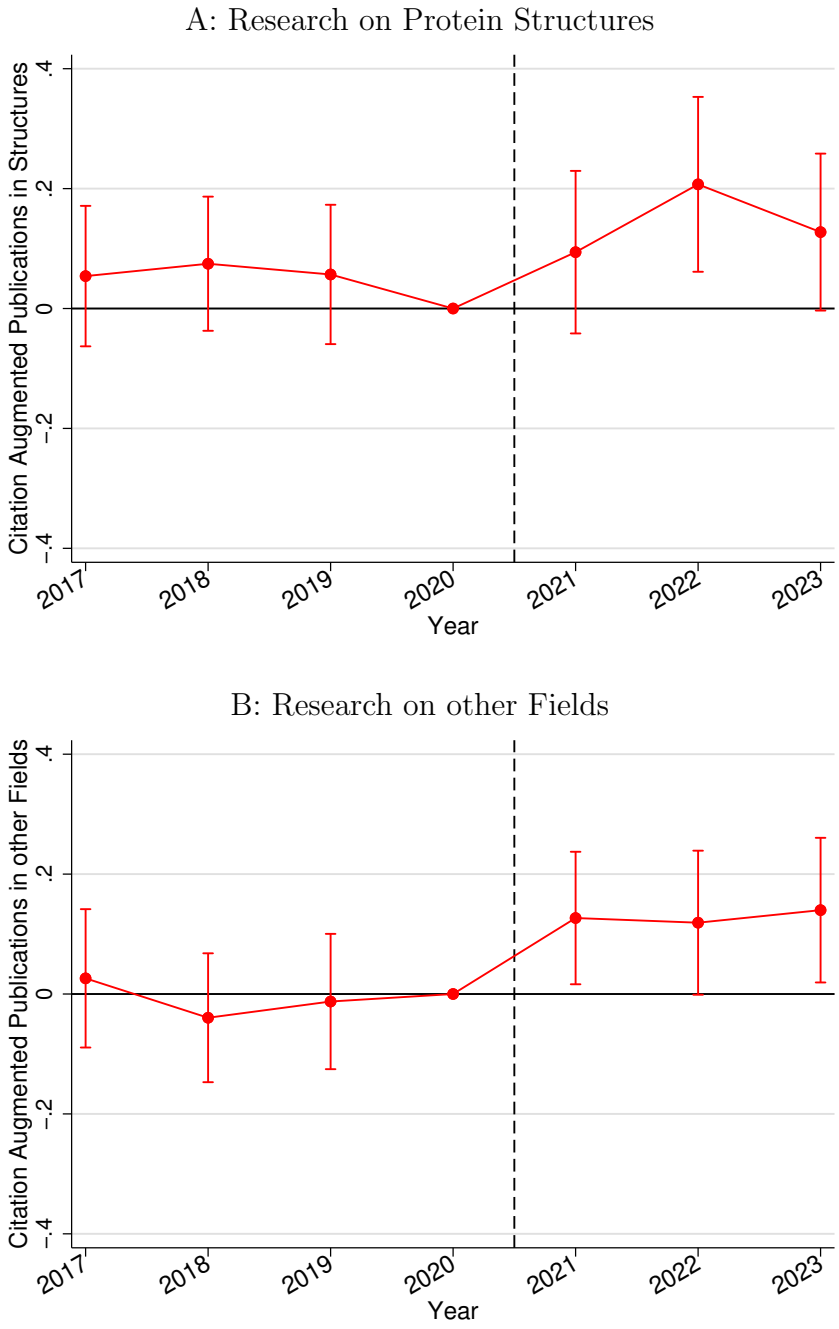
Note: Panel A shows the comparison between the AlphaFold2's prediction and the structure characterized by experiment in CASP14. The left one is CASP14 target T1037 (PDBID: 6VR4). The right one is CASP14 target T1049 (PDBID: 6Y4F). (source: https://deepmind.google/discover/blog/alphafold-a-solution-to-a-50-year-old-grand-challenge-in-biology/). Panel B presents the annual average number of protein structure papers from 2017 to 2023 in *Cell*, *Nature*, *Science*, and all journals that published at least two protein structure papers between 2013 and 2016. Panel C shows the percentage of protein structure papers citing AlphaFold2 from 2017 to 2023 in *Cell*, *Nature*, *Science*, and the same set of journals. The list of protein structure papers are downloaded from Protein Data Bank.

Table A1: Distribution of Treatment Intensity

Percentile	Pct. of PDB Papers (2017-2020)
1	0.00
25	0.09
50	0.33
75	0.67
99	1
Number of obs	2119

Note: This table presents the distribution of the percentage of PDB publications for structural biologists during 2017–2020.

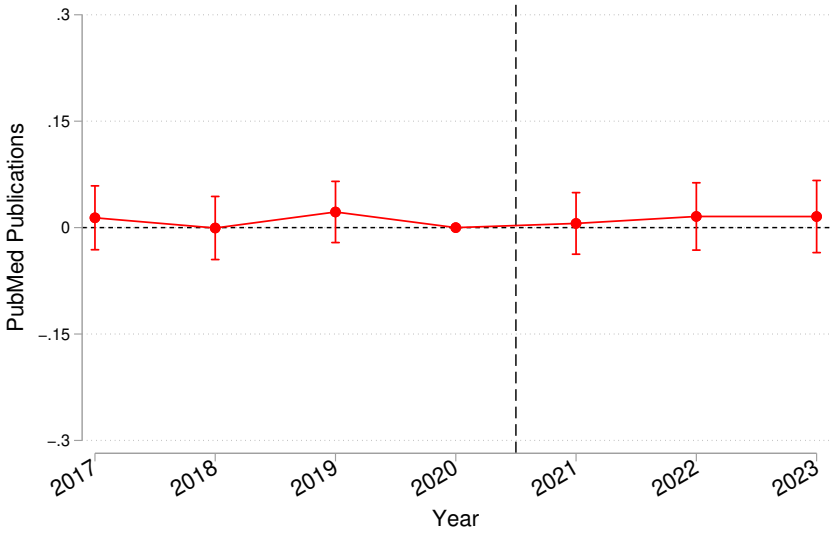
Figure A2: Impacts of AI on the Citations of Different Projects



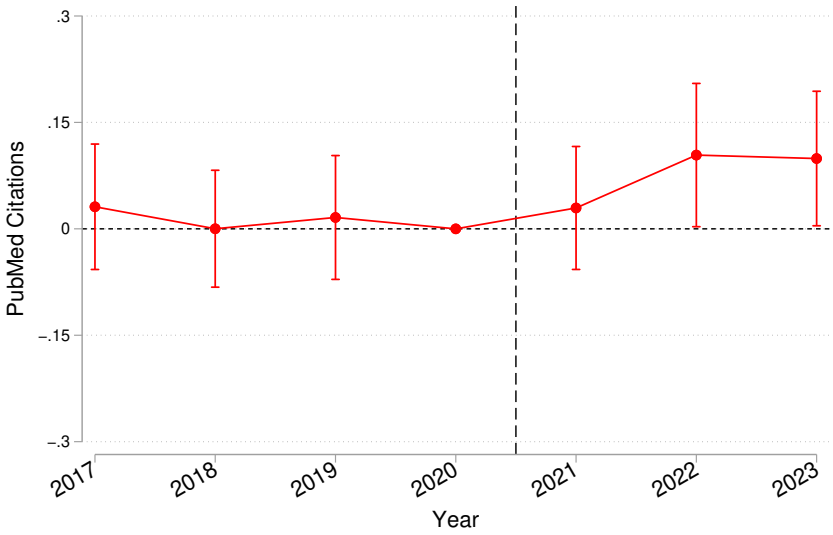
Note: These figures plot the DiD event-study estimates of AI's impact on citations by different fields. Panel A presents the the impacts on the citations of protein structure studies and Panel B shows the results for citations of other publications. All models are Poisson regressions that include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Each vertical bar represents the 95% confidence interval of the estimate.

Figure A3: Impacts on PubMed Publications and Citations (No Covid Research)

A: PubMed Publications



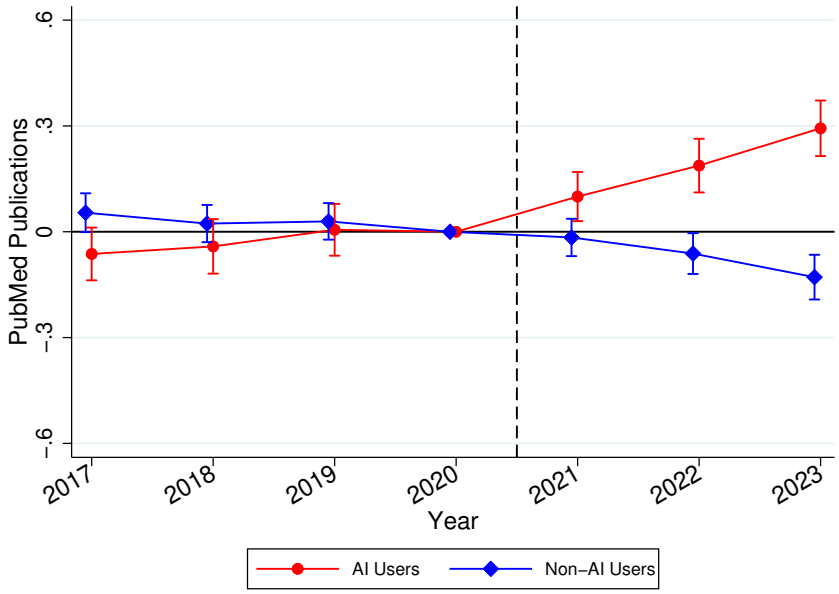
B: PubMed Citations



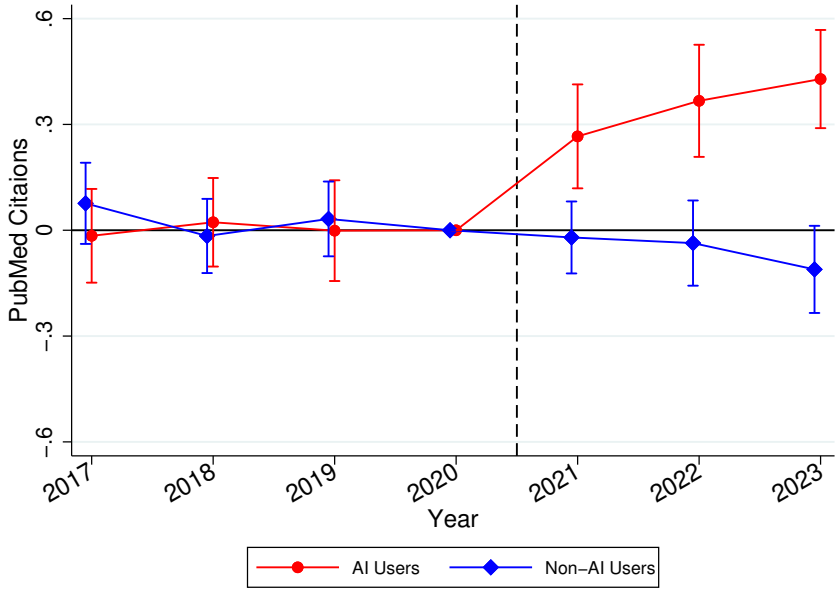
Note: These figures plot the DiD event-study estimates of AI's impact on structural biologists' academic performance, excluding Covid-related research. Panel A reports the impacts on the number of PubMed publications. Panel B focuses on the impacts on the citations. All models are Poisson regressions that include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Each vertical bar represents the 95% confidence interval of the estimate.

Figure A4: Heterogeneous Impacts: AI users vs non-AI users

A: PubMed Publications



B: PubMed Citations



Note: These figures compare the impacts of AI on academic performance between AI users and non-AI users. Panels A and B present the outcomes for PubMed publications and PubMed citations. All models include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Each vertical bar represents the 95% confidence interval of the estimate.

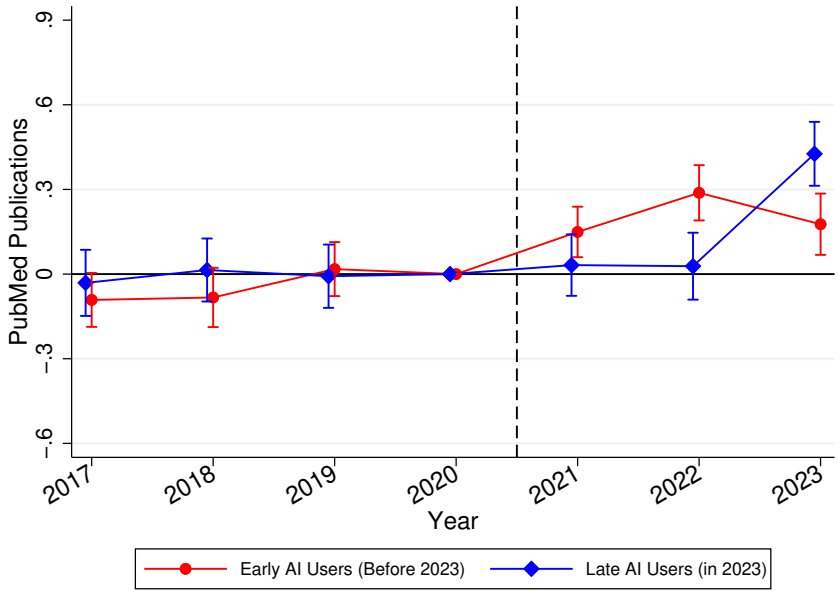
Table A2: Correlations between Protein Characteristics and AI Usage

	Novelty		Length		Weight		Citation	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
AlphaFold2	0.2867***	0.1933***	0.1925***	0.1204*	0.1781**	0.1208*	0.7482***	0.5731***
	(0.0350)	(0.0490)	(0.0635)	(0.0677)	(0.0710)	(0.0710)	(0.1849)	(0.1918)
R^2	0.0372	0.4256	0.0058	0.6616	0.0044	0.7093		
Log likelihood							-25262.73	-9369.90
Individual FE	No	Yes	No	Yes	No	Yes	No	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Number of obs	1855	1392	1855	1392	1855	1392	1855	1390

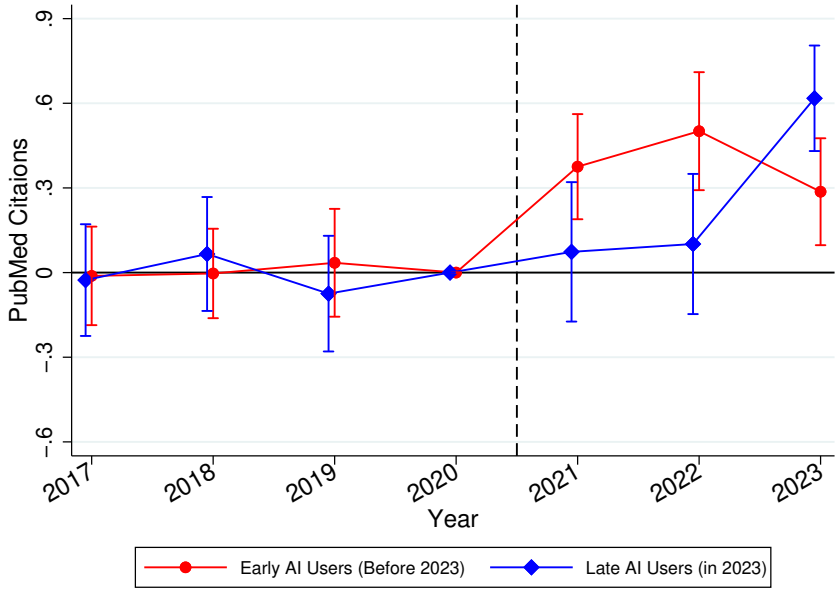
Note: This table presents the correlation between protein characteristics/quality and the usage of AlphaFold2. Columns (1)–(2) report correlation results for a dummy variable indicating whether the protein is novel. Columns (3)–(4) show results for the log-transformed length. Columns (5)–(6) present results for the log-transformed weight. Columns (7)–(8) examine citation counts. Columns (1)–(6) use OLS regressions, while Columns (7)–(8) employ Poisson regression. All regressions include year fixed effects, and Columns (2), (4), (6), and (8) also control for individual fixed effects. Statistical significance: $^{***}p < 0.01$,$^{**}p < 0.05$,$^{*}p < 0.1$.

Figure A5: Heterogeneous Impacts: Early vs Late AI Adopters

A: PubMed Publications



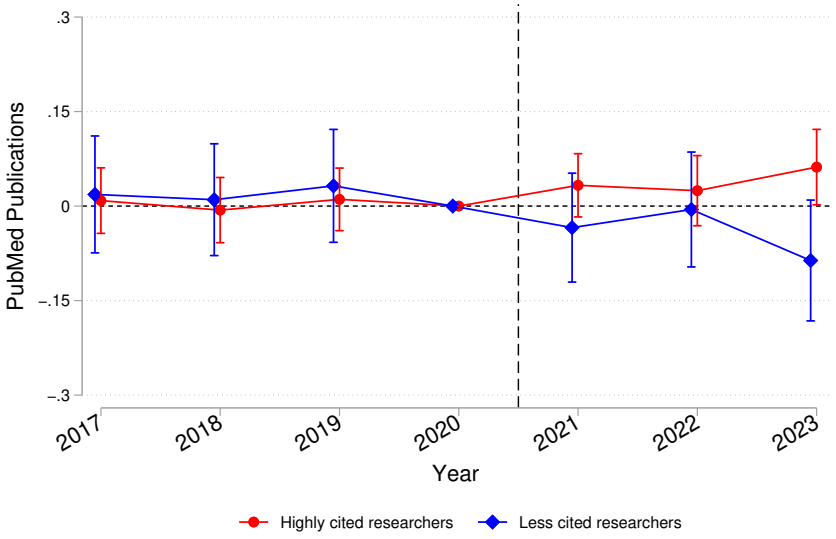
B: PubMed Citations



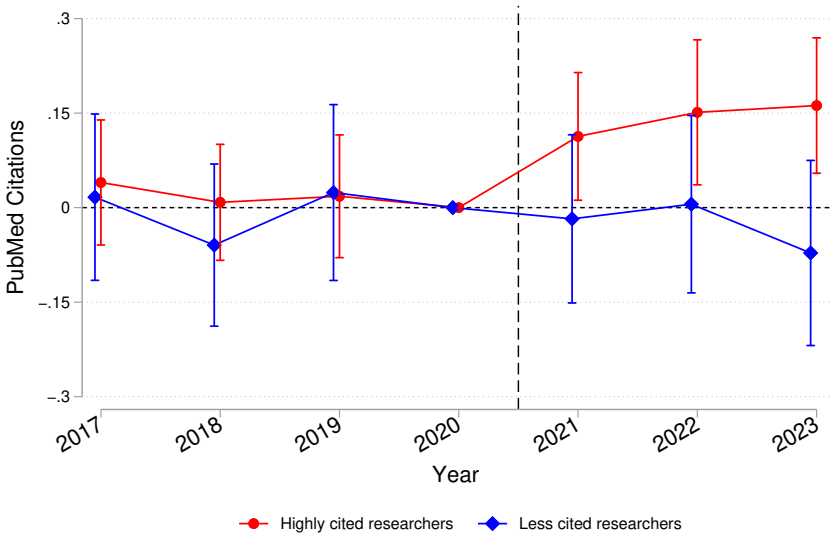
Note: These figures compare the impacts of AI on academic performance between AI users and non-AI users. Panels A and B present the outcomes for PubMed publications and PubMed citations. All models include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Each vertical bar represents the 95% confidence interval of the estimate.

Figure A6: Heterogeneous Impacts: Highly Cited vs Less Cited (Unconditional)

A: PubMed Publications

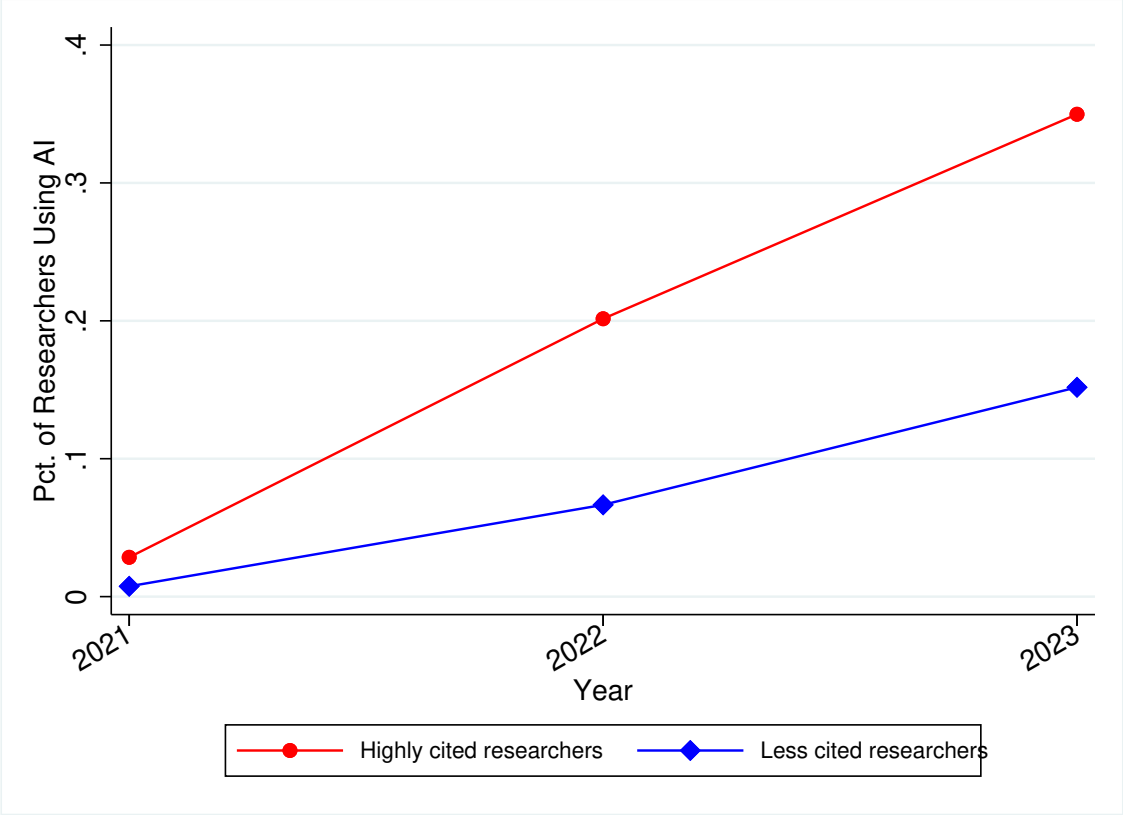


B: PubMed Citations



Note: These figures compare the impacts of AI on academic performance between highly cited and less cited researchers. Panels A and B present the outcomes for PubMed publications and PubMed citations. All models include individual fixed effects, year fixed effects, and CEM. Robust standard errors are clustered at the author level. Each vertical bar represents the 95% confidence interval of the estimate.

Figure A7: AlphaFold2 Adoption Over Time



Note: This figure depicts the cumulative percentage of structural biologists adopting AI from 2021 to 2023. The red line represents highly cited researchers, while the blue line represents less cited researchers.

Table A3: Characteristics of Highly Cited Researchers

	(1) High Novelty	(2) High Length	(3) High Weight	(4) Cyro EM	(5) High Fund
Highly Cited	0.1358***	0.1704***	0.1694***	0.1683***	0.2742***
	(0.0210)	(0.0214)	(0.0214)	(0.0153)	(0.0209)
R^2	0.0193	0.0290	0.0287	0.0544	0.0752
Number of obs	2119	2119	2119	2119	2119

Note: This table presents the correlation between a structural biologist's baseline citation level and their expertise in novel, long, and heavy proteins, their use of Cryo-EM in protein research, as well as the amount of fund. Statistical significance: $^{***}p < 0.01,^{**}p < 0.05,^{*}p < 0.1$.

Appendix B Timeline of AlphaFold2 Development

This section examines the relationship between AlphaFold1 and AlphaFold2 and tries to address the concern about the potential anticipatory effects before AlphaFold2's emergence.In 2018, DeepMind's AlphaFold1 made headlines by winning the international protein structure prediction challenge known as the Critical Assessment of Structure Prediction (CASP13), achieving a median Global Distance Test (GDT) score of approximately 60 out of 100. For comparison, the top-performing models in previous CASP competitions had typically scored around 40, representing only moderate accuracy in predicting protein 3D structures. AlphaFold1 demonstrated that deep learning could outperform traditional physics-based and template-based methods, marking a turning point for the field.

Yet the leap from AlphaFold1 to AlphaFold2 was not merely incremental and it represented a wholesale change in technological paradigm. AlphaFold1 relied on a two-stage process: it first used neural networks to predict pairwise distances between amino acids and then applied separate algorithms to construct 3D structures from these predictions. In contrast, AlphaFold2 introduced an end-to-end architecture that directly maps protein sequences to 3D structures. This model integrates multiple sources of information—such as evolutionary relationships from multiple sequence alignments and physical constraints—within a unified, iterative transformer-based framework that refines predictions in a self-consistent manner.

This redesign dramatically improved predictive accuracy. At CASP14 in 2020, AlphaFold2 achieved near-experimental accuracy on most targets, with a median GDT score exceeding 90—an improvement so striking that it effectively solved the long-standing protein folding problem for many practical purposes. As Nobel laureate and structural biologist Venki Ramakrishnan remarked, AlphaFold2's achievement "occurred decades before many people in the field would have predicted." This comment suggests that, although structural biologists may have anticipated the eventual development of a powerful tool for protein structure prediction, most did not expect such rapid progress.

Furthermore, although the DeepMind team published their methods for AlphaFold1 in *Nature* in 2020, they did not release a database of predicted protein structures, as they later did with AlphaFold2. This is largely because AlphaFold1's accuracy was insufficient for most practical research applications. As a result, the AlphaFold1 paper has received fewer than 3,000 citations in PubMed as of May 2025, whereas the AlphaFold2 paper has already garnered over 33,000 citations—a reflection of its far greater utility and impact on the scientific community.

Shortly after the debut of AlphaFold2, RoseTTAFold, a similarly accurate protein prediction software, was introduced by Baker Lab. RoseTTAFold's paper (Baek et al., 2021) was

published on the same day as AlphaFold2's in Science. Given the near-simultaneous timelines and comparable accuracy of both AlphaFold2 and RoseTTAFold, the introduction of the latter does not skew our analysis of AlphaFold2's impact. Notably, since the RoseTTAFold team did not release a comprehensive database akin to AlphaFold's, the latter's influence remains significantly larger. This is the primary reason for this paper's focus on AlphaFold2's impact on structural biologists. In fact, in 2024, David Baker, Demis Hassabis, and John Jumper jointly received the Nobel Prize in Chemistry.

A common identifying assumption in Difference-in-Differences designs is the absence of anticipatory effects—that is, agents do not change their behavior in advance of the treatment or policy shock. In my setting, this concern is mitigated by the nature of AlphaFold2 as a research tool, rather than a policy or incentive shock. Even if structural biologists expected an improved version of AlphaFold to emerge exactly in 2020, they could not access or use AlphaFold2 prior to its release, nor could they incorporate its predictions into their research workflows. In contrast to policies that change expectations or incentives in advance, AlphaFold2 had no productive impact until it was actually released as a functioning and accessible tool. As such, the standard no-anticipation assumption is unlikely to be violated in this context: the treatment effect of AlphaFold2 began only after researchers could use it.

Appendix C Usage of AlphaFold2

Researchers cite prior work for various reasons—including direct use of a paper's results, inspiration from its ideas, or general acknowledgment of related literature. This section investigates whether structural biologists citing AlphaFold2 actively employ its algorithm or prediction outputs in their research. I randomly selected 30 papers led by structural biology principal investigators that cite AlphaFold2 and manually examined how the method was used in each case (see Appendix Table C1).

Among these 30 papers, 17 are directly involved in solving protein structures, as indexed in the Protein Data Bank (PDB). Of these 17 PDB papers, only 1 cites AlphaFold2 solely for general acknowledgment in the conclusion. 9 use AlphaFold2-predicted structures to extend structural analysis or as hypotheses to guide experiments. The remaining 7 not only use AlphaFold2 predictions for structural guidance or experimental design but also incorporate them into model refinement during structure determination.

Among the 13 papers not directly involved in structure determination, only 2 cite AlphaFold2 solely for general acknowledgment in the introduction or conclusion. In contrast, 9 of the 13 non-PDB papers directly use AlphaFold2 predictions in their research. 1 paper uses AlphaFold2's algorithm to validate a synthetic structure. 1 paper evaluates AlphaFold2's

predictive performance, offering insights into its broader research applications. Overall, 90% of the sampled papers citing AlphaFold2 actively incorporate it into their research, rather than citing it merely due to its prominence. This suggests that citations to the AlphaFold2 paper serve as a plausible proxy for actual AI usage by the structural biologists.

Table C1: AI Usage

Paper ID	PDB	Purpose	Actual Usage
1	0	General Acknowledgment	0
2	0	Direct use of AlphaFold2 predictions to gain structural insights	1
3	0	Validation for designed proteins	1
4	0	General Acknowledgment	0
5	0	Direct use of AlphaFold2 predictions to gain structural insights	1
6	1	Generate model hypotheses and refine the model building process	1
7	0	Direct use of AlphaFold2 predictions to gain structural insights	1
8	1	Generate model hypotheses and refine the model building process	1
9	1	Direct use of AlphaFold2 predictions to gain structural insights	1
10	0	Direct use of AlphaFold2 predictions to gain structural insights	1
11	1	Direct use of AlphaFold2 predictions to gain structural insights	1
12	0	Direct use of AlphaFold2 predictions to gain structural insights	1
13	1	Direct use of AlphaFold2 predictions to gain structural insights	1
14	1	Generate model hypotheses and refine the model building process	1
15	1	Generate model hypotheses and refine the model building process	1
16	1	Direct use of AlphaFold2 predictions to gain structural insights	1
17	1	Generate model hypotheses and refine the model building process	1
18	1	Direct use of AlphaFold2 predictions to gain structural insights	1
19	0	Direct use of AlphaFold2 predictions to gain structural insights	1
20	1	Direct use of AlphaFold2 predictions to gain structural insights	1
21	1	Generate model hypotheses and refine the model building process	1
22	1	General Acknowledgment	0
23	0	Direct use of AlphaFold2 predictions to gain structural insights	1
24	0	Test the reliability of AlphaFold2 for further research.	1
25	1	Generate model hypotheses and refine the model building process	1
26	1	Direct use of AlphaFold2 predictions to gain structural insights	1
27	0	Direct use of AlphaFold2 predictions to gain structural insights	1
28	1	Direct use of AlphaFold2 predictions to gain structural insights	1
29	1	Direct use of AlphaFold2 predictions to gain structural insights	1
30	0	Direct use of AlphaFold2 predictions to gain structural insights	1
Total	17		27

Note: This table summarizes how a random sample of papers citing AlphaFold2 make use of it. 17 papers are directly involved in solving protein structures. 27 papers actually utilize AlphaFold2 into their research and only 3 cite AlphaFold2 for general acknowledgment.

Appendix D Data Appendix

D.1 PubMed

PubMed is a free search engine for biomedical and life sciences literature, maintained by the U.S. National Library of Medicine (NLM) at the National Institutes of Health (NIH). It primarily provides access to MEDLINE, a vast database of peer-reviewed biomedical journal articles, along with citations from other sources such as PubMed Central (PMC). Each paper in PubMed is assigned a unique PubMed ID (PMID). Furthermore, PubMed analyzes article content to assign relevant Medical Subject Headings (MeSH) terms which is a set of standardized vocabulary used by the NLM to index and organize biomedical literature. Additionally, PubMed provides a list of PMIDs that cite a particular PubMed papers. The list of PMIDs for papers that cite the AlphaFold2 team's two publications (Jumper et al., 2021; Varadi et al., 2022) was downloaded from PubMed in July 2024.

D.2 Protein Data Bank

The Protein Data Bank (PDB) is a global repository that stores and provides access to three-dimensional structural data of biological molecules, such as proteins and nucleic acids. It contains the universe of solved protein structures. PDB groups polymer entities based on sequence similarity thresholds (e.g., 30%, 50%, 70%, 90%, 100%). The groups are mutually exclusive, and I define a paper as solving a novel protein if its publication date is the earliest among all proteins within the 30% sequence similarity threshold group. I downloaded the data for all structures stored in PDB as of July 2024.

D.2.1 PubMed Paper and Protein Structure Characteristics

Although PDB links each PMID to the corresponding protein structure, there is a misalignment between the outcome level (publication) and the protein structure characteristics. Specifically, characteristics such as resolution, R Free, sequence length, molecular weight, and sequence clusters are defined at the protein level, whereas a single publication may report multiple protein structures. For large proteins, researchers would deposit different structures for its different domains, meaning that some structure characteristics are defined at an entity level, which is even smaller than the protein level.

To address these issues, I assign the weight and length of the longest entity in a paper as the representative protein characteristics for that paper. In general, the longest entity is also the heaviest. Similarly, the resolution and R Free for a paper correspond to those of the

longest entity reported in that publication. A protein entity structure is considered novel if its publication is the earliest among all entities within the same 30% sequence similarity cluster. If a paper contains at least one novel structure, I classify it as solving a novel structure. With these adjustments, I can now regress paper-level outcomes, such as whether a paper cites AlphaFold2, on the protein characteristics of the paper.

D.3 Dimensions

Dimensions database (https://www.dimensions.ai/) is a research analytics platform developed by Digital Science that provides access to a vast database of scholarly content, including journal articles, preprints, grants, and patents. Dimensions encompasses over 140 million publications and more than 7 million grants. It assigns unique researcher IDs to track scholars unambiguously. This paper was written using data obtained on July 15, 2024, from Digital Science's Dimensions platform, available at https://app.dimensions.ai. For publication histories, I use Dimensions' API service to retrieve the following publication details for the sample authors:

D.3.1 Affiliation

Dimensions does not report author affiliation information on a year-by-year basis. Therefore, I rely on the affiliation reported in the "authors" variable, which reflects the affiliation listed in the publication. To determine an author's baseline affiliation, I select the affiliation from their most recent publication prior to the first treatment year (2021). If an author has multiple affiliations, I use the first-listed affiliation as the representative one. Dimensions also provides country information for affiliations in the "authors" variable. I classify an affiliation as an academic institution if its name contains "college," "university," or "institute of technology." To determine the global ranking of academic affiliations, I perform a name match between the reported affiliation and the 2020 rankings from the *QS World University Rankings*.

D.3.2 Covid Literature

To identify whether a paper is related to coronavirus, I rely on its MeSH terms. There are three MeSH terms directly associated with coronavirus research: *Coronavirus*, *SARS-CoV-2*, and *COVID-19*. If a paper contains any of these terms, I classify it as a coronavirus related paper. Since the COVID-19 outbreak occurred at the end of 2019, the MeSH term *COVID-19* only appear in publications from 2020 onward. However, studies related to *Coronavirus* or *SARS-CoV-2* may have been conducted before the COVID-19 outbreak.

D.4 ORCID

ORCID provides employment histories for researchers who voluntarily upload and publicly share their information. Therefore, the sample for the employment analysis is limited to individuals who have made their ORCID employment records publicly available. The data include role title, affiliation, start date, and end date of each employment entry. To identify whether an employing organization is an academic institution, I classify affiliations as academic if their names contain keywords such as "college," "university," or "institute of technology." The employment data were obtained from the ORCID Public Data File 2024, posted on October 4, 2024.

D.5 Coarsened Exact Matching

To ensure that control biologists are comparable to structural biologists in academic background and exhibit similar publication trends prior to the treatment year (i.e., the debut of AlphaFold2), I employ the Coarsened Exact Matching (CEM) procedure. First, I perform an exact match based on authors' affiliations as of 2020, using the affiliation listed in their most recent paper by that year. For authors with multiple affiliations in a single paper, I select the first listed affiliation. Affiliations supporting experiments in protein structure determination can differ significantly from those without structural biologists, and this step helps account for potential confounding effects related to institutional differences.

Next, I select 3 baseline covariates for coarsened matching: publication age, coronavirus citations, and baseline citations. I divide publication age into four bins based on the quartiles of its distribution. For COVID-related citations, I first classify researchers into two groups: those with COVID-related research and those without. Among researchers with COVID-related work, I further divide them into four bins based on the quartiles of their COVID citation distribution. For baseline citations, I group researchers into four bins according to the quartiles of their overall citation distribution. The CEM procedure relies heavily on human judgment, especially when choosing relevant covariates and deciding the bin size. To ensure robustness, I tested various cutoff points for the covariate bins and found that minor adjustments did not significantly alter the main results.

Appendix E A Model of AI Adoption

Structural biologists decide whether to study a protein that requires cutting-edge AI technology. If a researcher continues working on a protein that does not require AlphaFold2, she receives y academic rewards, measured in publications or citations. If the researcher chooses

to solve a complex protein by leveraging AI, she receives $(1+\delta)\cdot y$ academic rewards, where $\delta > 0$ since AI-augmented research tends to attract greater academic attention. Let $\pi \cdot y$ denote the cost of switching to a protein that requires AI, where $\pi \sim U(0, c)$. This cost may arise from learning AlphaFold2 or directing from the current research agenda to proteins where AlphaFold2's predictions are essential. Researchers will adopt AI if the net benefit is positive:

$$((1+\delta)y - y) - \pi \cdot y \geq 0, \quad \text{or simply} \quad \delta - \pi \geq 0.$$

Suppose there are two types of researchers: H and L. H-type structural biologists are high-skill researchers who have more experience with complex proteins, while L-type researchers, on average, focus on less complex proteins. Assume that AI benefits L-type researchers more than H-type researchers: $\delta^L > \delta^H$. Additionally, assume that the cost of adopting AI is, on average, higher or equal for L-type researchers compared to H-type researchers: $c^L \geq c^H$. This assumption follows from the idea that high-skill researchers can more easily integrate the new AI technology into their research. The notion that low-skill researchers benefit more from adopting AI while facing a higher adjustment cost parallels the setting in (Chiquiar and Hanson, 2005), where low-skill migrants can earn more in the destination country but incur a higher migration cost.

Proposition 1. *If the cost of integrating AI is equal for both L- and H-type researchers on average, L-type researchers more likely use AI in their work:*

$$P(\delta^L - \pi^L \geq 0) > P(\delta^H - \pi^H \geq 0).$$

The intuition is that if there is no difference in the cost of applying AI, L-type researchers are more likely to use AI because they benefit more. This situation more likely happens when AI is used for simple tasks, such as using *ChatGPT* for writing emails requires almost no costs regardless of the skill levels.

Proposition 2. *If the cost of integrating AI for L-type researchers is higher than for H-type researchers on average but the net benefit of using AI is still higher for L-type researchers, i.e., $\delta^L - c^L \geq \delta^H - c^H$, then L-type researchers more likely use AI in their work:*

$$P(\delta^L - \pi^L \geq 0) \geq P(\delta^H - \pi^H \geq 0).$$

The intuition is that even though the cost of using AI is higher for lower-skill researchers, they are still more likely to adopt AI if the productivity gains are sufficiently large.

Proposition 3. *If the cost of integrating AI for L-type researchers is higher than for H-type researchers on average and the net benefit of using AI is lower for L-type researchers, i.e.,* $\delta^L - c^L < \delta^H - c^H$*, then L-type researchers less likely use AI in their work:*

$$P(\delta^L - \pi^L \geq 0) < P(\delta^H - \pi^H \geq 0).$$

The intuition is that even if the potential benefit is higher for low-skill researchers, the cost may be too high for them to shift from their current work to AI-augmented research. This situation is more likely when AI is required for specialized tasks.

The difference of productivity growth between high- and low-skill researchers after the availability of AI can be expressed as follows:

$$\frac{p^H(1+\delta^H)y^H + (1-p^H)y^H - y^H}{y^H} - \frac{p^L(1+\delta^L)y^L + (1-p^L)y^L - y^L}{y^L} = p^H\delta^H - p^L\delta^L,$$

where p^H and p^L represent the probability of AI adoption for high- and low-skill researchers, respectively. This expression shows that the productivity gap is jointly determined by both the benefit of AI (δ^L, δ^H) and the probability of adoption (p^L, p^H). Even though AI provides greater benefits to L-type researchers than H-type researchers ($\delta^L > \delta^H$), the inequality between high- and low-skill researchers can widen if low-skill researchers are significantly less likely to adopt AI than their higher-skilled counterparts. Specifically, if $p^H\delta^H > p^L\delta^L$, then the productivity gap becomes larger after the availability of AI.