NBER WORKING PAPER SERIES

BUILDING A WALL AROUND SCIENCE:
THE EFFECT OF U.S.-CHINA TENSIONS ON INTERNATIONAL SCIENTIFIC RESEARCH

Robert Flynn
Britta Glennon
Raviv Murciano-Goroff
Jiusi Xiao

Working Paper 32622
http://www.nber.org/papers/w32622

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2024, Revised May 2025

We are grateful to Fabian Waldinger, Megan MacGarvie, Ina Ganguli, Lee Branstetter, Emilie Feldman, Mae McDonnell, Natalie Carlson, and Heather Berry, as well as seminar participants at the Geography of Innovation Conference, Wharton Emerging Markets Conference, Sussex University Department of Economics, University of Richmond, University of Hawaii, University of Connecticut, Academy of Management, and the Chinese Economists Society Conference. We are especially appreciative of Dimensions from Digital Science, which provided us with the data used in this analysis. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Building a Wall Around Science: The Effect of U.S.-China Tensions on International Scientific Research
Robert Flynn, Britta Glennon, Raviv Murciano-Goroff, and Jiusi Xiao
NBER Working Paper No. 32622
June 2024, Revised May 2025
JEL No. F22, F6, O3

ABSTRACT

This paper quantifies how rising U.S.–China geopolitical tensions reshaped international science along three dimensions: trainee mobility, cross-border knowledge flows, and researcher productivity. Using a difference-in-differences design alongside detailed CV and publications data, we find ethnically Chinese students became 15% less likely to enter U.S. doctoral programs and 4% less likely to remain in the U.S. after graduation. China-based scientists cite U.S. research less frequently, while U.S. citations of China-produced work remain unchanged on average. Productivity among U.S.-based ethnically Chinese researchers declined by 8–11%, with a 10% higher rate of exit, but we do not see the same declines in China-based researcher productivity. These patterns cannot be explained by formal policy changes alone and instead point to a chilling effect driven by perceived hostility. Our results highlight how geopolitical tensions, independent of war or policy, can reshape the global geography of scientific talent and knowledge.

Robert Flynn
Boston University
robflynn@bu.edu

Britta Glennon
University of Pennsylvania
and NBER
bglennon@wharton.upenn.edu

Raviv Murciano-Goroff
Boston University
Questrom School of Business
ravivmg@bu.edu

Jiusi Xiao
Claremont Graduate University
jiusi.xiao@cgu.edu

1 Introduction

The production of science has become increasingly geographically distributed and globally interconnected. For example, the percentage of science and engineering Ph.D. degrees granted by U.S. universities earned by foreign-born individuals has nearly doubled since the 1980s (NSF NCSES). Diaspora networks have driven the diffusion of knowledge and ideas around the globe (Kerr, 2008; Oettl and Agrawal, 2008; Miguelez, 2018). And publications with authors from multiple countries now account for 23% of global publications and 40% of U.S.-based publications (NSF NCSES). Furthermore, there is growing evidence that international collaboration and talent flows lead to higher-impact science (e.g. Hsiehchen, Espinoza and Hsieh, 2015; Freeman and Huang, 2015), while access to frontier knowledge—regardless of its geographic location—is critical for scientific progress (Iaria, Schwarz and Waldinger, 2018).

As science becomes more international, however, it also becomes more susceptible to international conflict and geopolitical tensions with potentially significant, and yet difficult to measure, effects on innovation and the scientific workforce. In this paper, we analyze the impact of rising geopolitical tensions between the United States and China on trainee mobility and retention, cross-border knowledge flows, and scientist productivity. Using difference-in-difference empirical frameworks carefully tailored to each outcome, we examine each dimension from both a "U.S." perspective and from a "China" perspective. Because many of the policies contributing to rising U.S.–China tensions were rooted in nationalistic goals to reduce reliance on the other country and to strengthen domestic scientific capabilities, we examine how such tensions affect science in both countries. While our analysis is structured to capture potential asymmetries, we focus in particular on understanding why the observed effects appear especially pronounced for ethnically Chinese researchers and trainees in the U.S.

Understanding the impact of U.S.-China tensions on science is particularly important given that the relationship between the U.S. and China, until recently, was arguably one of the most important scientific relationships in the world. About a third of visa holders enrolled in U.S. Science and Engineering Ph.D. programs are from China (NSF NCSES) while a quarter of U.S.-based Science and Engineering publications have at least one author based in China (NSF NCSES). But around 2016, that relationship began to come undone, as we document in more detail in Section 2.

To quantify the impact of U.S.-China tensions on the dimensions of mobility, knowledge flows, and productivity, we rely on two rich data sources. For our analysis of the effect on mobility, we analyze a collection of publicly posted curricula vitae (CVs) on ORCID (Open Research and Contributor ID), a website where academics can create and share a digital CV, containing information on employment and educational histories. Critically, these data allow us to examine mobility for scientists at early stages of their careers, even before they have produced any publications. In addition, for tracking changes in knowledge flows and scientists' publication productivity, we utilize bibliometric data from Dimensions from Digital Science, a database of the metadata from scientific publications. In all analyses, we focus on STEM research and trainees given the particular focus of the American and Chinese governments on STEM.

We employ a difference-in-differences empirical design to quantify the impact of these growing U.S.-China tensions on trainee mobility and retention, knowledge flows, and scientist productivity. Such an approach is critical given the concurrent development of Chinese science during this period. Simply estimating the correlation between U.S.-China geopolitical tensions and, for example, the propensity of Chinese students to study or work in the U.S., could lead to bias; as Chinese science continues to advance, prospective Chinese Ph.D. students may become more likely to stay in or return to China. Hence, for each major component of our analysis, we are careful to both select an appropriate control group and to show event-study plots documenting clear trend breaks. We consider the "treatment" to begin in 2016, but we are careful to include dynamic treatment effects to account for the fact that tensions gradually escalated over several years, since rising U.S.-China tensions encompassed more than a single discrete policy change.

Our analysis proceeds in four parts. First, we examine STEM trainee mobility and retention. We find that ethnically Chinese students became 15% less likely to enroll in U.S. doctoral programs after 2016 and, if they did enroll, 4% less likely to stay in the U.S. after graduation. Many shifted toward other anglophone destinations instead. Second, we examine cross-border knowledge flows in both directions. After 2016, China-based scientists significantly reduced citations to U.S. research—particularly to recent and frontier work—while U.S.-based scientists' citation of China-produced publications remained unchanged on average. Third, we study scientific productivity. China-based researchers who previously relied on U.S. science saw little to no decline in output. In contrast, ethnically Chinese researchers in the U.S. became 8-11% less productive and were 10% more likely to stop publishing entirely.

Finally, we probe the underlying mechanism behind these changes in mobility, citation patterns, and productivity. While formal U.S. government actions—such as visa restrictions and the China Initiative—may have played a role, our results suggest a broader mechanism at work. First, we show that many of the observed effects begin around 2016, predating the most concrete policy interventions. Second, the effects extend to individuals unlikely to be directly targeted by U.S. enforcement actions, including ethnically Chinese scientists from third countries, those without professional ties to China, and likely Chinese-Americans (identified as having Chinese surnames but non-Chinese given names). Third, the magnitude of the effects varies by local political climate, with stronger declines in both Chinese graduate applications and enrollments at U.S. institutions in Republican-leaning states, consistent with the interpretation that perceived hostility plays a key role. Fourth, prior research finds no evidence of increased F-1 visa denial rates for Chinese students during this period that might explain the decline in enrollment. Finally, we document a sharp divergence in perception and behavior among ethnically Chinese scientists in the U.S., including reduced citation of China-produced research and widespread self-reported fear. Taken together, these patterns are most consistent with a chilling effect driven not by formal exclusion but by perceived hostility and reputational risk, which operate through anticipatory self-selection rather than policy enforcement.

This paper contributes to the broader literature on the effect of war, conflict, and geopolitics on science. Prior work shows that international conflict can significantly disrupt science. For example, World War I reduced international knowledge flows and scientific cooperation, and lowered the productivity of scientists who depended on frontier knowledge from abroad (Iaria, Schwarz and Waldinger, 2018). In the lead-up to World War II, widespread academic emigration reshaped global science (Waldinger, 2012; Becker et al., 2021): the expulsion of professors from Germany damaged German Ph.D. student outcomes (Waldinger, 2010) and changed the trajectory of U.S. science (Moser, Voena and Waldinger, 2014). The collapse of the Soviet Union triggered an outflow of scientists from the USSR (Abramitzky and Sin, 2014; Borjas and Doran, 2012; Ganguli, 2017), and more recently, the Russian invasion of Ukraine has reduced productivity among Ukrainian scientists and constrained the exchange of scientific knowledge and ideas (Ganguli and Waldinger, 2023).

Our results build on this literature by showing that even geopolitical tensions at a much lower level than the formal expulsion of academics or violent warfare can lead to a significant shift in scientist mobility, knowledge flows, and productivity. The effect of tensions on these dimensions is likely to be driven by a mix of explicit government policy that targets particular foreign groups, where it exists, and nationalist or anti-foreign sentiment. The degree to which such sentiment may affect the production of science–and whether there are important asymmetries in terms of how seriously science in different countries is affected–is not well understood. Our results suggest that bilateral geopolitical tensions, along with growing nationalist and anti-foreign sentiment, significantly impact scientific mobility and productivity. Given the literature linking immigrants to innovation (Bernstein et al., 2022; Hunt and Gauthier-Loiselle, 2010; Moser and San, 2020) and the evidence indicating disastrous long-term effects on universities in the sending countries (Waldinger, 2016) and positive effects on universities and science in the receiving countries (Agrawal, McHale and Oettl, 2017), a change in where scientists migrate has major implications for the global geography of science.

Our paper also contributes to the emerging literature on the effects of contemporary U.S.–China tensions. Recent work has examined the impact of the China Initiative—an important policy launched in 2018 that disproportionately targeted scientists of Chinese origin in the U.S.—on collaboration patterns and research productivity (Aghion et al., 2023; Jia et al., 2024). We extend this work in two key ways.

First, we take a more comprehensive view of how tensions affect science by analyzing effects on both the U.S. and China and across multiple dimensions: mobility, knowledge flows, and productivity. Second, we document a novel mechanism: a chilling effect driven not by formal policy, but by perceived hostility. Separately, Sim and Hong (2025) show that U.S.–China tensions increased workplace effort among high-skilled Chinese immigrants in the U.S. private sector, pointing to the broader reach of geopolitical conflict beyond academia.

More broadly, our findings provide new evidence that negative sentiment and perceived risk can reshape where scientists choose to study, work, and publish—even in the absence of direct restrictions. This chilling effect operates through anticipatory behavior and reputational concerns, rather than through formal exclusion. It affects not just those with direct ties to China but also individuals who are ethnically Chinese regardless of nationality or affiliation. While prior work has focused on war or explicit legal barriers, ours is among the first to show that ambient geopolitical hostility alone can meaningfully disrupt global science.

2 Empirical Context

Our study focuses on the years surrounding a sharp downturn in U.S.–China relations. Formal scientific cooperation between the two nations dates back to the 1979 U.S.–China Science and Technology Cooperation Agreement. By the 2010s, each country was the other's largest scientific research partner, and China was the largest source of international students in U.S. STEM programs (USCET, 2023). Starting around 2016, however, the relationship began to deteriorate.

First, anti-China rhetoric rose sharply during the 2016 U.S. presidential campaign. Donald Trump repeatedly described China as an economic and strategic threat, citing trade imbalances, IP theft, and espionage. His book Great Again: How to Fix Our Crippled America, which outlined his political agenda, included quotes such as, "there are people who wish I wouldn't refer to China as our enemy. But that's exactly what they are." This narrative translated into policy. Beginning in 2018, the U.S. imposed tariffs on Chinese goods, kicking off the U.S.-China trade war, and pursued a broader confrontational stance.

More relevant to this paper, the U.S. Department of Justice launched the "China

Initiative" in 2018 to investigate (typically) Chinese and Chinese-American scientists under suspicion of IP theft on behalf of the Chinese government. Figure A1 tallies the cases charged under the initiative, according to the MIT Technology Review.[1] Though only 77 cases were brought, the initiative fueled widespread concern about racial profiling and was ultimately shut down in 2022. During this period, the U.S. also revoked student visas[2] and barred entry for some Chinese scholars[3]. In 2018, the State Department imposed new restrictions limiting the duration of F-1 visas for Chinese students in sensitive STEM fields.[4] These actions contributed to an increasingly hostile climate, especially for ethnically Chinese scientists in the U.S.

Consistent with this series of events, anti-Chinese sentiment in the U.S. ticked up from around 55% in 2015 to 81% in 2024, according to the Pew Research Center.[5] The changing policy environment and sentiment had a particularly striking impact on the experience of ethnically Chinese scientists working in the U.S. Faculty protests against investigations of Chinese scholars cited concerns regarding a growing "chilling effect" on academic research by creating a hostile environment for Chinese and Chinese-American researchers in the U.S.[6]

Importantly, these changes have not been entirely one-sided. Many of the U.S. policies were a response to, for example, forced technology transfer in China and significant corporate espionage by China in the United States.[7] The Chinese government has also implemented a wide range of nationalist policies intended to improve self-reliance. For example, the Thousand Talents Program, which began to develop in the early 2000s, and the Junior Thousand Talents Program, implemented in 2013, have long aimed to encourage Chinese (senior and junior, respectively) academics to return to China from abroad. As a more recent example, in

---

[1] https://www.technologyreview.com/2021/12/02/1040656/china-initative-us-justice-department/

[2] https://www.nytimes.com/2018/07/25/us/politics/visa-restrictions-chinese-students.html

[3] https://www.nytimes.com/2019/04/14/world/asia/china-academics-fbi-visa-bans.html

[4] https://www.science.org/content/article/more-restrictive-us-policy-chinese-graduate-student-visas-raises-alarm

[5] https://www.pewresearch.org/global/2024/05/01/americans-remain-critical-of-china/

[6] See, for example, letters by faculty at Stanford, Yale, University of Pennsylvania, and Princeton.

[7] See for example https://www.nytimes.com/2017/08/01/business/trump-china-trade-intellectual-property-section-301.html

March 2020, China's government changed its incentive scheme for academics in China to no longer encourage publication in international journals. But these policies have typically not targeted the U.S. explicitly until recently, and they have been of a more gradual nature. For this reason, we view the primary treatment in our analysis as originating on the U.S. side.

Given this backdrop of escalating rhetoric, policy shifts, and rising hostility—–particularly from the U.S. side—–we treat the 2016 onset of U.S.–China geopolitical tension as a bundled "treatment" comprising changes in policy, political rhetoric, and public sentiment. Because this treatment evolved gradually over time, we employ dynamic treatment effects and then later investigate which specific mechanisms best explain the observed effects.

3 Data

Our analysis investigates how the rise of U.S.-China tensions starting in 2016 impacted STEM trainee mobility (graduate student enrollment and retention), knowledge flows (the usage of scientific works), and scientific productivity in STEM fields. For analyzing trainee mobility, we utilize data from curricula vitae (CVs) posted on Open Research and Contributor ID (ORCID), a website where academics can publicly post their CV in a standardized format. For the latter two outcomes, we utilize scientific publications data from Dimensions, a database of metadata on the near-universe of scientific works published in academic journals. From those two raw data sources (CV data and Publications data), we construct five datasets for our analysis. Figure A1 provides an overview of these constructed datasets while Table 1 provides basic summary statistics for each. Details about their construction are provided below.

3.1 CV Data

We construct two datasets from the curriculum vitaes (CVs) available in ORCID. Each CV in ORCID includes an individual's name as well as their self-reported educational background and employment history. Based on that information, we constructed additional variables for our analysis. We marked if and when the individual enrolled in a degree program, if this program was in the U.S., if this program was at the doctoral level, and whether the individual held a job in the U.S. immediately after studying at a U.S. institution. In addition, we assigned them an academic field based on their listed department and inferred based on their name if they were ethnically Chinese. Specifically, we imputed each individual's ethnicity using an algorithmic name classifier and created a flag for if this ethnicity was Chinese. Additional details of the procedures for determining field and inferring ethnicity are provided in Appendix A.4.

The ORCID website contains over 15 million CVs. We collected the publicly available CVs posted on the site as of 2022. We restricted to individuals reporting complete educational backgrounds, which amounts to 1.8 million CVs. We further restrict to those who graduated from STEM programs, a total of 836,495 CVs.[8]

Our analysis of trainee mobility draws on two datasets constructed from ORCID CVs. The first is the *Doctoral Student dataset* dataset, which includes individuals who entered doctoral programs between 2008 and 2019. To be included, individuals must report at least one prior degree pursued within 10 years of their Ph.D. enrollment; we use the location of this prior degree as a proxy for nationality. This measure of nationality is conceptually and empirically distinct from our measure of ethnicity, which we infer algorithmically based on name. The dataset includes 129,223 individuals.

The second dataset, the *U.S. Graduates dataset*, includes individuals who completed a degree at a U.S. institution between 2008 and 2019 and took a job within three years of graduation. That dataset contains information about 50,890 individuals. Tables 1(a) and 1(b) present simple descriptions of both datasets.

Because individuals must actively create an ORCID profile and fill out a digital CV, our Doctoral Student and U.S. Graduates datasets are not representative of the full population. Appendix Table A10(a) shows that ORCID users tend to have more publications and a higher likelihood of grant funding than non-users. While our sample is selected, this group aligns with our focus on the individuals most likely to contribute to future science.

[8]The following 11 fields are considered STEM: Agriculture, Biological Sciences, Biomedical and Clinical Sciences, Chemical Sciences, Earth Sciences, Engineering, Environmental Sciences, Health Sciences, Information and Computing Sciences, Mathematical Sciences, and Physical Sciences.

Table 1: Basic Summary Statistics Across Datasets

Panel A: Doctoral Students	Mean	SD
Ph.D. first year	2,012.79	2.88
Enrolls in U.S. university	0.24	0.43
Enrolls in U.K. university	0.09	0.29
Enrolls in non-U.S. anglo. university	0.17	0.37
Ethnically CN	0.16	0.37
Observations	129,223	

Panel B: U.S. Graduates	Mean	SD
Job first year	2,014.92	3.16
Job in U.S.	0.85	0.35
Job in U.K.	0.01	0.10
Job in non-U.S. anglo. country	0.03	0.16
Ethnically CN	0.18	0.38
Observations	50,890	

Panel C: Publication-Citation Shares	Mean	SD
Citing U.S.	0.50	0.50
Share of raw references	0.13	0.15
Share of recent references	0.09	0.15
Share of frontier references (1%)	0.19	0.28
Share of recent frontier references (1%)	0.16	0.29
Observations	4,237,614	

Panel D: U.S.-U.K. Publications	Mean	SD
U.S. publication	0.83	0.38
Cite share to CN	0.02	0.04
Recent cite share to CN	0.03	0.07
Frontier cite share to CN (1%)	0.01	0.06
Recent frontier cite share to CN (1%)	0.02	0.09
Observations	2,833,730	

Panel E: China-based Researcher Panel	Mean	SD
Num Pubs	2.15	2.78
Num Pubs in US-based journals	0.65	1.25
Num Impact Factor wt Pubs	4.47	6.53
Impact Factor wt US-based Pubs	1.54	3.50
Predom. Cite US	0.50	0.50
Observations	83,839	

Panel F: U.S.-based Researcher Panel	Mean	SD
Num Pubs	0.92	1.89
Num Pubs in US-based journals	0.54	1.30
Impact Factor wt Pubs	2.64	6.68
Impact Factor wt US-based Pubs	1.70	5.08
Ethnic CN	0.17	0.37
Observations	1,913,212	

Notes: This table provides basic descriptive statistics for the data described in Section 3. The unit of analysis for each dataset is in the panel title. Panels (a) and (b) summarize the data for analyzing trainee mobility (student or graduate level), Panels (c) and (d) for scientific knowledge flows (publication-share and publication level), and Panels (e) and (f) for researcher productivity (researcher by year level).

Given the importance of ethnically Chinese scientists in this paper, we might be concerned that ORCID users who are ethnically Chinese differ from those who are not ethnically Chinese, and specifically, that they might have noticeably different research outcomes prior to the onset of rising tensions in 2016. Appendix Table A10(b) shows the mean attributes of these two groups. Reassuringly, while ethnically Chinese researchers also tend to have stronger research outcomes, we do not observe interactions between ORCID self-selection and being ethnically Chinese.

As an additional check on the validity of our constructed data, we compare our insights to those provided by publicly available established sources. Reassuringly, doctoral stay rates inferred from our U.S. graduates dataset are similar to those reported by NSF. Specifically, for the cohort of nationally Chinese STEM doctoral students graduating from U.S. universities in 2012,[9] the NSF reported a stay rate of 83%; we similarly observe that 83% of post-graduation jobs taken in 2012 by Chinese doctoral graduates were in the U.S. For nationally Indian graduates, the rates are 83% and 79.1% respectively. That these rates are comparable validates our retention measure and underlying data.

3.2 Publications Data

We construct three datasets using the bibliometric information about published scientific works available in Dimensions.[10] For each publication in their database, Dimensions provides the names of the authors, the publishing journal, the scientific field(s) of the publication, the publishing year, the addresses of the authors, and the list of papers cited in the references and bibliography. In addition, Dimensions provides algorithmically disambiguated author identifiers enabling the tracking of authors across publications.

As of 2023, the Dimensions data contains over 1.8 billion citations across 140 million publications. Of those, we focus on the 51 million published between 2008 and 2019. For each publication, we construct additional variables based on the publication's metadata. We create a flag for whether all of the authors list first affiliation addresses in China,[11] the U.S., or the U.K. (which is our primary control country). We refer to these publications as being written by China-based, U.S., or U.K research teams respectively. For each author on each publication, we also flag if that author's modal address country between 2008 and 2012 was either China or the U.S. If the modal country for an author during that time was China, we call the researcher "China-based." We similarly define "U.S.-based" researchers. Note that our labels such as 'China-based researcher' or 'China-based research team' refer

[9] 2012 is the last year before our treatment year (2016) for which complete data on this outcome is available from NSF.

[10] Dimensions is similar to other bibliometric databases such as Web of Science and Scopus, but it has been shown to have a wider coverage of scientific journals represented in their data (Singh et al., 2021).

[11] In all of our analyses, when we refer to "China" as a location, we are referring to mainland China.

to institutional affiliation, not ethnicity or nationality. For example, a researcher working at a Chinese institution is labeled 'China-based' regardless of their ethnic background or country of origin.

Using this data, we create three datasets. First, for analyzing if China-based research teams changed their usage of U.S.-produced scientific works, we create a dataset which we call the *Publication-Citation Shares dataset*. Following the methodology of Iaria, Schwarz and Waldinger (2018), for each publication in Dimensions, we create two observations. The first observation contains measures of how much the focal publication references scientific papers produced by research in the U.S., and the second observation contains measures of how much the focal publication references scientific papers produced by research in the U.K. Publications that do not reference any previous works are removed. Further, we focus only on the observations produced from the references of STEM publications authored by China-based research teams between 2011 and 2019. This amounts to 4,237,614 observations from 2,118,807 publications.

We compute four measures of the usage of science from these countries: raw, recent, frontier, and recent frontier. The raw measure captures the share of the publication's total references that cited papers produced in the U.S. and the U.K. Recent citations refer to those published within the preceding five years. The frontier and recent frontier measures restrict these same calculations to citations of high-impact research, defined following Iaria, Schwarz and Waldinger (2018) as research in the top percentile of its field's citation distribution. More details on these usage measures are provided in Appendix A.2.

Second, to examine if U.S.-China tensions impacted the usage of Chinese science by U.S. research teams, we create a dataset which we call the *U.S.-U.K. Publications dataset*. This dataset contains STEM publications authored by U.S. and U.K. research teams between 2011 and 2019. For each publication, we again compute multiple measures of the usage of science (raw, recent, frontier, and recent frontier) but in reverse: measuring usage of China-produced science among U.S. and U.K. publications.[12] This dataset includes 2,833,730 publication observations and

[12] We define a publication as being produced in China primarily according to the publication's corresponding author. See Appendix A.2 for details.

enables us to track how researchers in the U.S. and U.K. changed their usage of China-produced scientific knowledge.[13]

Finally, to examine if U.S.-China tensions have impacted the productivity of China-based and U.S.-based scientists, we create a panel dataset which we call the *Researcher Panel*. The observations in this dataset are created by constructing a strongly balanced panel of the authors listed in publications in the Dimensions data in the years between 2008 and 2019.[14] For each researcher-year observation, we include the number of publications by that author in that year, as well as quality-adjusted measures, such as the number of publications weighted by the impact factors of the publishing journals.

To analyze the productivity impact of U.S.–China tensions on China-based researchers, we restrict the Researcher Panel to China-based STEM[15] scientists who published at least five papers between 2008 and 2012[16] and at least one publication between 2013 and 2019.

For the analysis of U.S.-based researchers, we apply a different restriction. Because treatment status in this case is based on ethnicity rather than citation patterns, we do not require pre-2013 publication activity. We include all U.S.-based STEM researchers who published at least once between 2013 and 2019.

For both panels of China-based and U.S.-based STEM researchers, we apply the Coarsened Exact Matching (CEM) method (Iacus, King and Porro, 2012) to match on pre-analysis observables, as will be described in more detail later, to ensure that the treatment and control researchers are comparable groups. After the CEM procedure, the final China-based STEM researcher panel includes 11,977 unique individuals with 83,839 researcher-year observations. The final U.S.-based STEM researcher panel includes 273,316 unique individuals with 1,913,212 researcher-year observations.

[13] Unlike the Publication-Citation Shares dataset, the U.S.-U.K. Publications dataset is not disaggregated into shares. This is because we are interested in citations to only one country—China—among these publications.

[14] Our analysis focuses on the 2013–2019 period. We use the years 2008–2012 for computing metrics, such as how active researchers were and where they were located.

[15] We define a researcher's field as the modal field of their publications.

[16] The restriction on pre-analysis period publications ensures a sufficient number of observations to identify treatment status, as further detailed in Section 4.3.

Table 1 Panels (c)-(f) present basic summary statistics of these data. Appendix Tables A5 and Appendix A7 report the balance between the treated and control group after the CEM procedure. More detailed explanations and summary statistics for all five datasets can be found in Appendix A.

4 Analysis and Results

We now present the results of our difference-in-differences analyses, structured around the three key outcomes: trainee mobility, cross-border knowledge flows, and researcher productivity. For each, we pair difference-in-differences estimates with event-study plots to assess pre-trends and trace dynamic effects. In the mobility and U.S.-based productivity analyses, treatment is defined by ethnicity (inferred from names), while in the China-based productivity analysis, treatment is based on researchers' pre-period reliance on U.S. versus U.K. science. In the knowledge flows analysis, treatment is assigned at the publication or publication-share level based on the citing or cited country respectively. A comprehensive set of robustness checks, including placebo tests, sensitivity analyses, and alternative sample restrictions, is presented in Section 4.4.

In the following sections, we explain our approach to analyzing each outcome in detail. A summary of these approaches can also be found in Table A2.

4.1 STEM Trainee Mobility

4.1.1 Enrollment in U.S. Doctoral Programs

Attracting talented trainees–and retaining them post-graduation–has long been seen as economically and competitively important for the U.S. to retain its edge in scientific research (Black and Stephan, 2010). But the U.S.-China tensions described in Section 2 suggest that the U.S. may have become a less attractive destination for Chinese STEM trainees. Therefore, we examine how growing U.S.-China tensions starting in 2016 affected the enrollment of ethnically Chinese students in doctoral programs at U.S. universities relative to their non-ethnically Chinese counterparts with the following difference-in-differences model estimated on observations from the Doctoral Student dataset:

$$Y_i = \beta_1 Treat_i + \beta_2 Post_{it} + \beta_3 (Treat_i^* Post_{it}) + \gamma \mathbf{X}_{it} + \epsilon_i \quad (1)$$

In this equation, i is a doctoral student and $Post_{it}$ is defined as an indicator of whether the student began their doctoral studies in 2016 or later. The outcome of interest is Y_i, which is an indicator for whether the student enrolled in a U.S. doctoral program. The treatment group for this analysis are students who are ethnically Chinese, and the control group are students who are not ethnically Chinese. $\mathbf{X}_{it}$ contains fixed effects for the year of enrollment, the scientific field of a student's doctoral studies, and the country where the student received their prior academic degree. Note that, as described in Section 3.1, we define treatment based on ethnicity (inferred from names), while separately controlling for nationality—–proxied by the location of individuals' prior degrees. This approach isolates the effect of rising U.S.–China tensions from broader shifts in the appeal of U.S. doctoral programs.

Table 2 presents the results from estimating Equation 1. Column (1) reports the estimated impact on the likelihood of enrolling in a U.S. doctoral program to be -3.5 percentage points (SE = 0.78). This amounts to a 15% decline relative to the sample mean. Table 2 Columns (2) and (3) report the estimated effects on the likelihood of enrolling in a university located in the U.K. or non-U.S. anglophone country respectively.

The estimated effect on the likelihood of enrolling in a U.K. university (Column 2) is 0.9 percentage points (SE = 0.4), while the effect for non-U.S. anglophone universities (Column 3) is 2.2 percentage points (SE = 0.67). The latter represents a 13% increase over the sample mean. Taken together, these findings are consistent with the explanation that other anglophone universities substituted for U.S. universities among ethnically Chinese Ph.D. students after 2016. They also suggest that the results are not due to an increase in the appeal of Chinese universities; the Chinese STEM trainees not going to the U.S. appear to attend other anglophone universities rather than staying in China. An additional robustness check showing no increase in the propensity to enroll in China after 2016 further supports this view (Appendix Figure A3).

Table 2: Main Treatment Effects on Trainee Mobility Among Ethnically Chinese Doctoral Students

	(1) Enrolls in U.S.	(2) Enrolls in U.K.	(3) Enrolls in Anglo.	(4) Enrolls in U.S.	(5) Enrolls in U.S.
Treatment = Ethnically CN	0.0323 (0.00466)	-0.0198 (0.00376)	-0.0238 (0.00483)	0.0249 (0.00510)	
Treatment = Ethnically CN × Post-2016	-0.0348 (0.00782)	0.00924 (0.00404)	0.0218 (0.00670)	-0.0309 (0.00804)	
Treatment = Top Quality Ethnically CN					0.132 (0.0105)
Treatment = Top Quality Ethnically CN × Post-2016					-0.0565 (0.0173)
Field & Cohort FE	Y	Y	Y	Y	Y
Prior Country FE	Y	Y	Y	Y	Y
Prior University Rank FE	N	N	N	Y	N
Model	OLS	OLS	OLS	OLS	OLS
Mean DV	0.239	0.0910	0.169	0.265	0.265
Observations	129204	129204	129204	92069	92069

Notes: This table presents difference-in-differences estimates of the impact of Chinese ethnicity on the probability of enrolling in a U.S. university post-2016. Standard errors are clustered at the field-year level. Dependent variables are listed in the column headers.The sample includes all STEM doctoral students from 2008-2019.

We further investigate whether changes in the quality of students enrolling in U.S. doctoral programs drive these results in Table 2. In Column (4), we control for students' prior degree university rank by decile and find that the estimated impact of U.S.-China tensions on the likelihood of enrolling in a U.S. doctoral program is not significantly different: -3.1 percentage points (SE = 0.80). In Column (5), we redefine the treatment group to be those ethnically Chinese doctoral students whose prior degree was from a top decile university, while keeping the control group unchanged. Among these students, the estimated impact on the likelihood of enrolling in a U.S. doctoral program grows to -5.7 percentage points (SE = 1.7). This suggests that the impact was strongest among students from top-quality universities.[17]

Both to examine pre-trends and to trace the dynamic effects, we estimate the following event-study model specification:

[17] In additional analyses performed, available upon request, we show that the fraction of students enrolling in doctoral programs from prominent Chinese universities is falling over time.

$$Y_i = \beta Treat_i + \sum_{k=2008}^{2014} \delta_k \mathbb{1}(t=k) Treat_i + \sum_{k=2016}^{2019} \tau_k \mathbb{1}(t=k) Treat_i + \gamma \mathbf{X}_{it} + \epsilon_i \quad (2)$$

The variables in this equation are the same as in Equation 1. In this specification, the δ_k coefficients capture any baseline differences in U.S. enrollment rates between ethnically Chinese and non-ethnically Chinese students in the pre-2016 period, while the τ_k coefficients trace out the dynamic effects of U.S.–China tensions in the years after 2016.

Figure 1(a) plots the estimated δ_k and τ_k coefficients from Equation 2. Prior to 2016, the rates at which ethnically Chinese and non-ethnically Chinese students enrolled in U.S. doctoral programs followed similar trends. Beginning in 2016, however, the rate that ethnically Chinese students enrolled in U.S. programs began declining (relative to that of non-ethnically Chinese students) and continued declining through at least 2019. For example, in 2018, the probability that an ethnically Chinese doctoral student enrolled in a U.S. program had decreased by five percentage points relative to the rate in 2015.

The results in this section document that U.S.-China tensions emerging in 2016 decreased the enrollment of ethnically Chinese students in U.S. doctoral programs, but the economic significance of our results can be difficult to interpret. A back-of-the-envelope estimate of the number of nationally Chinese doctoral students displaced during our treatment period leverages the coefficient from Table 2 Column (1) combined with NSF NCSES data[18] data to reason that 5,760 nationally Chinese STEM students did not attend U.S. doctoral programs between 2016 and 2019 as a result of U.S.-China tensions.[19]

We also observed that ethnically Chinese students enrolled in non-U.S. anglo-

[18]https://ncses.nsf.gov/pubs/nsb202332/figure/HED-23

[19]We restrict our estimation to nationally Chinese doctoral students given the absence of annual data on ethnically Chinese matriculation patterns, understanding the former to be an approximate subset of the latter. We assume that the 34,510 nationally Chinese students enrolled in U.S. doctoral programs in 2015, per NSF, were at that point evenly split across cohorts in five-year doctoral programs. Since nationally Chinese students are a smaller group than ethnically Chinese students, we should consider this projection as a lower bound. At the same time, given the selected nature of the ORCID sample and the data we draw from it, we generalize our results to broader populations with caution.

Figure 1: Event Studies for Student Enrollment and Retention

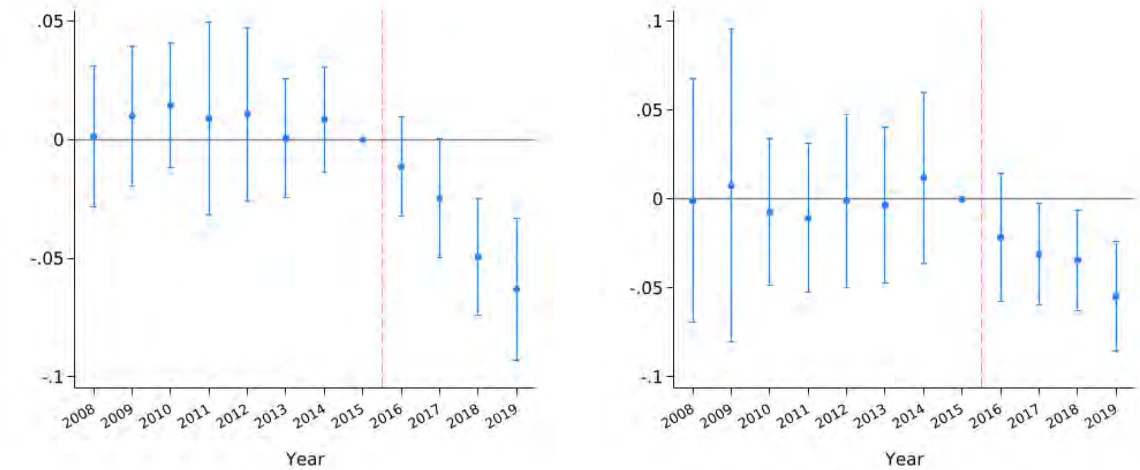

(a) DV: Enrollment **(b)** DV: Retention

Notes: Panel (a) reports event-study coefficients from a regression predicting enrollment in a U.S. university among doctoral students (Equation 2). The treated group is ethnically Chinese doctoral students, and the control group is non-ethnically Chinese doctoral students. The regression includes fixed effects for research field, doctoral cohort, and prior country. Panel (b) reports event-study coefficients from a regression predicting whether a U.S. graduate's first post-graduation job remains in the U.S. The treated group is ethnically Chinese U.S. graduates, and the control group is non-ethnically Chinese U.S. graduates. The regression includes fixed effects for research field and job start year. Standard errors are clustered at the field-year level in both regressions.

phone universities at increased rates in the years following 2016. These patterns are consistent with a shift in Chinese STEM talent from the U.S. to other anglophone countries following the rise in U.S.–China tensions.

4.1.2 Retention of U.S. Graduates

Retaining trained and talented scientists may be equally as important for a nation's economic competitiveness as attracting such talent. Table 3 displays the results from estimating Equation 1 on the U.S. Graduates dataset. The outcome here is whether a graduate's first post-degree job is in the U.S., and each observation is an individual graduating from a U.S. institution. As in the previous section, we define the treated group as ethnically Chinese individuals and the control group as non-ethnically Chinese individuals. We include fixed effects for the year of the

Table 3: Main Treatment Effects on U.S. Retention for Post-U.S. Graduation Jobs

	(1) Job in U.S.	(2) Job in U.K.	(3) Job in Anglo.
Treatment = Ethnically CN	-0.00119	-0.00613	-0.0131
	(0.00767)	(0.00188)	(0.00308)
Treatment = Ethnically CN × Post-2016	-0.0360	0.00330	0.00845
	(0.00946)	(0.00222)	(0.00363)
Field FE	Y	Y	Y
Job Year FE	Y	Y	Y
Model	OLS	OLS	OLS
Mean DV	0.853	0.00949	0.0255
Observations	50890	50890	50890

Notes: This table presents difference-in-differences estimates of the impact of U.S.-China tensions on the probability that an ethnically Chinese student's post-graduation job is in the U.S post-2016. Standard errors are clustered at the field-year level. Dependent variables are listed in the column headers.The sample includes all U.S. STEM graduates from 2008-2019.

graduate's first post-graduation job and the scientific field of their earned degree. Column (1) reports the estimated impact of treatment on the likelihood that a U.S. graduate's first job remains in the U.S. as -3.6 percentage points (SE = 0.95). This amounts to a 4% decline from the sample mean.

Figure 1(b) plots the corresponding event-study coefficients from estimating Equation 2 on the U.S. Graduates dataset. The plot shows estimates that are not significantly different from zero until after 2016, highlighting that the rate that ethnically Chinese graduates' jobs remained in the U.S. tracked with that of non-ethnically Chinese graduates prior to 2016. Following 2016, however, a statistically significant divergence emerges, with relative stay rates for ethnically Chinese graduates declining through 2019.

Since the relative rate at which ethnically Chinese U.S. graduates remained in the country decreased, where did they take jobs instead? Table 3 Columns (2) and (3) report the estimated effects of U.S.-China tensions on the likelihood that a U.S. graduate's first job is in the U.K. or in a non-U.S. anglophone country, respectively. Only the latter is statistically significant: Column (3) estimates a 0.85 percentage point increase (SE = 0.36) in the probability of taking a job in a non-U.S. anglophone country. This amounts to an increase over the sample mean of nearly 33%.

These results, while somewhat noisier than those for doctoral enrollment, are consistent with substitution: after 2016, ethnically Chinese U.S. graduates appear more likely to take positions in other anglophone rather than remain in the U.S.

4.2 Knowledge Flows and the Usage of U.S. and Chinese Research

Prior literature indicates that the mobility of scientists correlates with the diffusion of scientific knowledge (Bernstein et al., 2022; Hunt and Gauthier-Loiselle, 2010; Moser and San, 2020; Agrawal, McHale and Oettl, 2017). In particular, graduates of foreign doctoral programs who return home may bring back new ideas, scientific techniques, and personal connections, which can influence their future work. We investigate if the rise in U.S.-China tensions, and the subsequent decline in the mobility of graduate students and trained scientists, also influenced knowledge flows between researchers in the U.S. and China.

We examine how China-based research teams' usage of scientific works produced in the U.S., as well as how U.S.-based research teams' usage of works produced in China, changed because of worsening U.S.-China tensions. In these analyses, we make comparisons using U.K.-authored research and researchers as controls for their counterparts from the U.S. The U.K. serves as a suitable control for the U.S. in this analysis for a number of reasons. First, the U.S. and the U.K. enjoy similar levels of government support and national preference for research as reflected in their R&D workers per capita (Appendix Figure A2(a)). Second, both pursue similar types of research as evidenced in the field composition of their publications (Figure A2(b)). Third, the U.S. and the U.K. are the top two destinations for nationally Chinese researchers studying abroad (Figure A2(c)). Beyond these quantitative similarities, the U.S. and the U.K. share a common language and cultural lineage.

4.2.1 Knowledge Flows from U.S.-Produced Science to Researchers in China

Following the approach of Iaria, Schwarz and Waldinger (2018), we estimate a difference-in-differences model using observations from the Publication-Citation Shares dataset in which the focal publications were written by research teams in China. In this analysis, we compare the references on these publications by China-based teams to U.S. papers versus to U.K. papers. Specifically, we estimate the following specification:

$$Y_{ij} = \beta_1 Treat_{ij} + \beta_2 Post_{it} + \beta_3 (Treat^{*}_{ij} Post_{it}) + \gamma \mathbf{X}_{ijt} + \epsilon_{ij} \tag{3}$$

In this equation, i is a scientific publication authored by a research team in China, and j represents either the U.S. or U.K. The outcome of interest is Y_{ij}, which is the share of publication i's references that cite papers produced in country j. The treatment group for this analysis contains the observations for which j = U.S., and the control group contains observations where j = U.K. The variable $Post_{it}$ is defined as an indicator for if publication i came out in 2016 or later. $\mathbf{X}_{it}$ contains fixed effects for the citing publication. The parameter of interest in this equation is β_3, which can be interpreted as the effect of U.S.-China tensions on the share of references among publications by China-based research teams that cite U.S.-produced papers (versus U.K.-produced papers).

This difference-in-differences approach, as well as setting up the analysis to examine the relative share of references to U.S. papers versus U.K. papers, addresses two potential identification concerns. First, norms regarding citing previous work change over time. By including fixed effects for the year of publication, we negate any concern that these changes are impacting our estimates of the treatment effect. Second, Chinese research has been increasing in quality over time. Therefore, regardless of any geopolitical changes, researchers in China may be increasingly relying on China-produced research rather than research produced elsewhere. By comparing the relative share of U.S.-produced papers to U.K.-produced papers in the reference lists of publications by China-based research teams, we isolate the impact of U.S.-China tensions on the usage of U.S. research from the general trend in Chinese researchers relying less on foreign research.

Table 4 shows estimated coefficients from Equation 3 using publications authored by research teams in China. Column (1) estimates the impact of growing U.S.-China tensions on the share of references to papers published in the U.S. to be -0.014 (SE = 0.004). This corresponds to an 11% decline relative to the overall sample mean and a 6.5% decline relative to the average citation share to U.S. papers

Table 4: Main Treatment Effects on Knowledge Flows among Publications by China-based Teams

	DV: Share Size			
	(1) Raw	(2) Recent	(3) Frontier	(4) Recent Frontier
Treated = Citing U.S.	0.182	0.125	0.281	0.247
	(0.0137)	(0.0150)	(0.0186)	(0.0231)
Treated = Citing U.S.	-0.0138	-0.0140	-0.0142	-0.0321
× Post-2016	(0.00438)	(0.00256)	(0.00514)	(0.00553)
Citing Paper FE	Y	Y	Y	Y
Model	OLS	OLS	OLS	OLS
Mean DV	0.126	0.0853	0.193	0.162
Observations	4237614	4042500	3332908	2303014

Notes: This table presents difference-in-differences coefficients describing how the share of references citing U.S.-produced papers on publications by China-based research teams changed after 2016. Standard errors are clustered at the field level. Dependent variables are listed in the column headers. The sample includes citation shares to U.S. and U.K. papers on China-based publications. The analysis period is 2011–2019. The treated group includes citation shares to U.S. papers, and the control group includes citation shares to U.K. papers.

among publications by China-based research teams.[20] Column (2) reports the estimated impact of U.S.-China tensions on the share of recent references, defined as references to papers published in the previous five years, to be -0.014 (SE = 0.003) or a 16% decline from the sample mean and a 10% decline from the China-U.S. average. That this decline is greater (in percentage terms) suggests that increasing U.S.-China tensions may have had greater influence on the dissemination of recent research. To test if the results are driven by a reduction in citing low-quality research, we repeat the analyses from Columns (1) and (2) using dependent variables capturing citations to papers in the top 1% of their field's citation distribution, which we call "frontier research" and "recent frontier research." Columns (3) and (4) report the estimated impacts on these quality-adjusted shares, revealing a significant reduction in citations to frontier U.S.-produced research and suggesting that this concern is unfounded.

We also estimate an event-study model in order to trace the dynamic effects of U.S.-China tensions and to examine the evidence in support of the parallel trends assumption underlying the previous difference-in-differences estimates. The event-study model specification is as follows:

[20] We provide the percentage change relative to the China-U.S. average given the sizable level difference between the average China-U.S. and China-U.K. citation rates.

$$Y_{ij} = \beta Treat_{ij} + \sum_{k=2008}^{2014} \delta_k \mathbb{1}(t=k) Treat_{ij} + \sum_{k=2016}^{2019} \tau_k \mathbb{1}(t=k) Treat_{ij} + \gamma \mathbf{X}_{ijt} + \epsilon_{ij} \quad (4)$$

The variables in this equation are the same as Equation 3.

The event-study coefficients confirm that the change in citations to U.S.-produced papers by publications written by China-based research teams came about as an abrupt change, starting in 2016, and has continued to decline in the years since. Figure 2(a) shows the estimated coefficients when the dependent variable is the raw share of each publication's reference list. These shares move in parallel in the years prior to 2016, but following that year there is a dramatic decline in the rate of citing U.S.-produced papers. By 2018—–just two years after tensions began to rise—–the share of references to U.S. papers in publications by China-based research teams had already declined by 6.3% relative to 2015 (coefficient = -0.014, SE = 0.004). Appendix D includes event studies for citations to recent, frontier, and recent frontier research. The results remain negative and statistically significant for all versions, but the magnitude of the impact for recent frontier research more than doubles. This implies that the impact of U.S.-China tensions on publications by China-based research teams was particularly pronounced for citations to works on the scientific cutting edge. This pattern supports the interpretation that rising tensions disrupted actual knowledge flows, rather than simply altering citation practices. If the shift were purely behavioral, we would expect a uniform decline across all references. Instead, the drop is concentrated in citations to recent and frontier U.S. research—precisely the kinds of knowledge most essential for keeping up with the scientific frontier.

4.2.2 Knowledge Flows from China-Produced Science to Researchers in the U.S.

Did U.S. research teams similarly change their usage of China-produced scientific knowledge? To investigate that question, we compare the share of references in the

Figure 2: Event Studies for Knowledge Flows

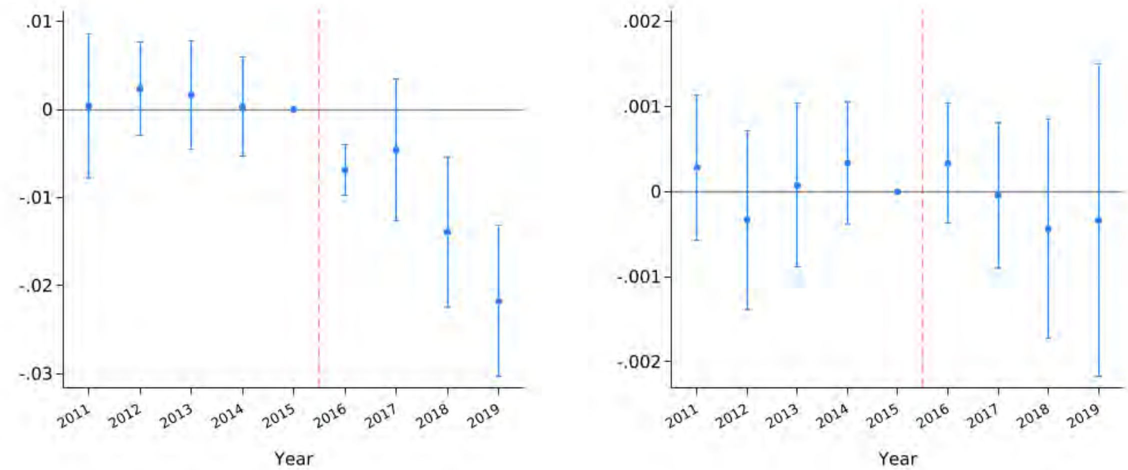

(a) China-based Teams Building on U.S. Science

(b) U.S.-based Teams Building on Chinese Science

Notes: Panel (a) reports event-study coefficients from a regression predicting the citation share to U.S. papers in publications by China-based research teams. The treated group is citation shares to the U.S., and the control group is citation shares to the U.K. The regression includes fixed effects for the citing publication. Panel (b) reports event-study coefficients from a regression predicting the share of references to China-produced papers in publications by U.S. and U.K. research teams. The treated group is U.S. publications, and the control group is U.K. publications. The regression includes fixed effects for publication year and research field. Standard errors are clustered at the field level in both regressions.

publications of U.S. research teams made to China-produced papers with the share of references in the publications of U.K. research teams made to China-produced papers.

This approach is different than the one that we used for analyzing if the usage of U.S.-produced research by China-based research teams had changed. In that analysis, we examined the share of references to U.S.-produced papers versus the share of references to U.K.-produced papers on publications by China-based research teams, which allowed us to account for secular trends—such as the rising quality and quantity of scientific works produced in China—that might decrease the share of references to foreign works more generally. In analyzing the U.S.-side, a similar approach would not be appropriate since no other country could plausibly model a counterfactual case given the unique changes occurring in China's scientific production over the past two decades. Therefore, instead, we examine the share of references to China-produced papers in the publications of U.S. and U.K. research teams. In doing so, we can control for the changes in China-produced science, while isolating the effect of U.S.-China tensions on the usage of China-produced scientific works in U.S. research.

We estimate the difference-in-differences model for this analysis adapting the specification in Equation 3 with observations from the U.S.-U.K. Publications dataset. Specifically, we estimate the following:

$$Y_i = \beta_1 Treat_i + \beta_2 Post_{it} + \beta_3 (Treat_i^{*} Post_{it}) + \gamma \mathbf{X}_{it} + \epsilon_i \quad (5)$$

In this equation, i is a scientific publication authored by a U.S.- or U.K.-based research team. The outcome of interest is Y_i, which is the share of publication i's references that go to papers produced in China. The treatment group contains the set of observations in which publication i is authored by a U.S.-based research team, and the control group contains observations where publication i is authored by a U.K.-based research team. The variable $Post_{it}$ is defined as an indicator for if the publication i came out in 2016 or later. $\mathbf{X}_{it}$ contains fixed effects for the year and the scientific field of the publication. The parameter of interest from this equation is β_3, which can be interpreted as the effect of U.S.-China tensions on the share of references to China-produced papers among publications by U.S. research teams, relative to U.K. research teams.

Table 5 shows the estimates when the dependent variable is the raw share of references, the share of recent references, the share of frontier references, and the share of recent frontier references. The coefficients are all small in magnitude, vary in direction, and none reaches statistical significance. Figure 2(b) and Panels (d)-(f) in Appendix Figure A6, which display the leads and lags from event-study models of Chinese citation shares among U.S. and U.K. publications, show a similar pattern. These plots reveal a mild and statistically insignificant decrease in the rate that U.S. publications (relative to U.K. publications) cite China in 2018 and 2019.

Both Figure 2(b) and the associated event study plots demonstrate that following the rise in U.S.-China tensions, U.S. researchers did not meaningfully change

Table 5: Main Treatment Effects on Knowledge Flows among U.S. and U.K. Publications

	DV: Share Size			
	(1) Raw	(2) Recent	(3) Frontier	(4) Recent Frontier
Treated = U.S. Publication	-0.000472	-0.000363	-0.000552	-0.000451
	(0.000601)	(0.000831)	(0.000344)	(0.000577)
Treated = U.S. Publication	-0.000194	0.000190	-0.000132	-0.000248
× Post-2016	(0.000682)	(0.00101)	(0.000529)	(0.000855)
Field & Year FE	Y	Y	Y	Y
Model	OLS	OLS	OLS	OLS
Mean DV	0.0166	0.0264	0.0109	0.0166
Observations	2833136	2755769	2348683	1826175

Notes: This table presents difference-in-differences coefficients describing how the share of references citing China-produced papers on U.S. publications changed after 2016. Standard errors are clustered at the field level. Dependent variables are listed in the column headers. The sample includes all publications by U.S. or U.K. research teams between 2011 and 2019.

their citation habits with respect to China-based scientific sources. The small and statistically insignificant coefficients both before and after 2016 demonstrate that the propensities of the U.S. and the U.K. to cite Chinese research moved in parallel throughout this period.

Ultimately, the results from this section reveal a shift in knowledge flows, demonstrated by a change in the works that China-based researchers used in their publications. Since 2016, China-based researchers have decreased their usage of U.S.-produced science in the citations of their publications. As this decrease can be seen in the relative usage of U.S.-produced research versus U.K.-produced research, the shift goes beyond any contemporaneous increase in the quality of Chinese science that may be causing Chinese researchers to rely less on science from non-Chinese sources more generally. In contrast, U.S. researchers did not change their usage of China-produced research in a statistically significant manner. This implies that the majority of the effect of U.S.-China tensions after 2016 were felt in the knowledge flows from the U.S. to China and much less so on the knowledge flows from China to the U.S.

4.3 Productivity Impact on STEM Researchers

Scientific productivity is a core driver of technological progress and national competitiveness. In this section, we examine how rising U.S.–China tensions affected the research output of scientists in both countries. There are multiple mechanisms by which productivity could have been affected: reduced international collaboration, fewer cross-border trainees, visa barriers, a decline in cross-border knowledge flows, and a broader chilling effect shaped by fear, discrimination, or reputational risk.

Our empirical strategy focuses on the groups most plausibly exposed to the mechanisms above: (1) China-based researchers who were highly reliant on U.S.-produced science—–a group shown in the previous section to experience a sharp decline in access to U.S. research, and (2) ethnically Chinese researchers based in the U.S., who appear especially vulnerable to perceived hostility and institutional barriers. For each, we estimate difference-in-differences models using matched control groups to isolate the effect of geopolitical tensions on publication productivity. We describe the approach in detail below.

4.3.1 Productivity Impact on China-based Researchers

For assessing the impact of rising U.S.-China tensions on the productivity of China-based researchers, we analyze both the extensive margin, which we define as scientists stopping research altogether, as well as the intensive margin, which we define as changes in the number of publications produced by scientists each year. The data used for these exercises are the sub-sample of observations from the Researcher Panel associated with research-active China-based scientists: researchers who are China-based STEM scientists, who published five or more publications between 2008 and 2012 and at least one publication between 2013 and 2019.[21]

Based on our findings in the previous section, we define the treated group as "U.S.-reliant" researchers, defined as those who are above the 75$^{\text{th}}$ percentile within their field for the portion of their citations that go to papers from the U.S.

[21] We define "research active" scholars as those with at least one publication between 2008 and 2019. To ensure sufficient publications to identify reliance, we restrict the sample to researchers with at least 5 publications during 2008–2012.

and are below the 25th percentile for their field for their citation share to papers from the U.K. In turn, we define the control group as "U.K.-reliant" researchers, defined as researchers above the 75th percentile in citation share within the field to the U.K. and below the 25th percentile in citation share within the field to the U.S.

To isolate the effect of the rise in tensions and boost the precision of our estimates, we match researchers from the treated and control observations by employing the CEM procedure from Iacus, King and Porro (2012). Specifically, we match researchers from the treated and control groups based on the following observables: number of publications produced between 2008 and 2012 (binned into deciles), career age as of 2012 (in four bins),[22] the number of actively publishing years between 2008 and 2012,[23] if the researcher is affiliated with a university, if the researcher is located in a Tier 1 city,[24] and if in a New Tier 1 city.[25] In addition, we included the level and the growth rates for the number of publications and impact-factor-weighted publications between 2013 and 2015 as covariates in the match. A comparison of the treatment and control groups across these covariates can be found in Appendix Table A5. Ultimately, the sample on which we analyze the productivity impact on matched China-based researchers contains 11,977 unique individuals. Of those, 5,983 are in the treated group ("U.S.-reliant" researchers) and 5,994 unique individuals in the control group ("U.K.-reliant" researchers).

For the extensive margin, we estimate the following Cox Proportional Hazard model:

$$h(t \mid Treat_i) = h_0(t) \exp(\beta \cdot Treat_i) \tag{6}$$

In Equation 6, $h(t \mid Treat_i)$ is the the hazard function of failure at time t, conditional on the treatment status. The failure event is defined by the researcher stopping publishing altogether. $h_0(t)$ is the baseline hazard function, $Treat_i$ is a binary indicator of treatment status. In the China-based researcher analysis, $Treat_i$ equals one

[22] Defined as the number of years since their first publication.
[23] The number of years between 2008 and 2012 that a person published at least one publication.
[24] Tier 1 cities include Beijing, Shanghai, Guangzhou, Shenzhen.
[25] New Tier 1 Cities include Chengdu, Chongqing, Hangzhou, Wuhan, Nanjing, Tianjin, Suzhou, Xi'an, Changsha, Shenyang, Qingdao, Zhengzhou, Dalian, Dongguan, Ningbo.

if the China-based researcher is "U.S.-reliant", zero if the China-based researcher is "U.K.-reliant". β is the coefficient of interest that measures the hazard of stopping publication associated with treatment status.

Columns (1) and (2) of Table 6 show the results. U.S.-reliant China-based researchers faced an increased, but not statistically significant, hazard of stopping publication altogether (β = 0.079, SE = 0.049). When the outcome is restricted to publishing in U.S.-based journals (Column 2), the hazard rises to 12% (β = 0.115, SE = 0.062), a statistically marginal effect indicating that U.S.-reliant researchers were more likely to disengage from U.S.-based journals following the rise in tensions. Appendix Figures A9(a) and A9(b) report the Kaplan-Meier Curves that show survival step function curves for the treated and control groups.

For the intensive margin, we estimate the following difference-in-differences specification predicting publication counts for active researchers:

$$Y_{it} = \beta_1 Treat_i + \beta_2 Post_{it} + \beta_3 (Treat^*_{it} Post_{it}) + \gamma \mathbf{X}_{it} + \epsilon_i \tag{7}$$

In Equation 7, i is a researcher and t is the year. The outcome of interest is Y_{it}, which is the number of publications researcher i published in year t in the baseline specification. $Post_{it}$ is defined as an indicator for 2016 or later. $Treat_{it}$ refers to the U.S.-reliant researchers. $\mathbf{X}_{it}$ contains individual fixed effects and year fixed effects. Because the outcome of interest is a count variable, we estimate this specification using a Poisson (PPML) model.

Columns (3)-(6) of Table 6 show the intensive margin results from difference-in-differences models, where the dependent variable varies across each column. Column (3) reports a small and statistically insignificant effect on overall publication counts (β = -0.024, SE = 0.016), suggesting that U.S.-China tensions had little impact on the productivity of U.S.-reliant China-based researchers relative to their U.K.-reliant counterparts.

While the overall effect on the treated group's productivity is small and not statistically significant, it is possible that China-based researchers changed *where* they published their works. For example, these researchers may have found it more challenging to publish in U.S.-based journals after 2016. To test this, Column

(4) restricts the outcome to publications in U.S. journals,[26] yielding an even smaller and again insignificant effect (β = -0.003, SE = 0.027). Finally, we investigate the impact on the quality of scientific production by weighting publications according to the impact factors of the journals in which they appeared.[27] Results are similarly null when using impact-factor-weighted outcomes (Columns 5–6).

To document the dynamic effects and assess the pre-trends, we estimate the following event-study specification with variables defined similarly to Equation 7:

$$Y_i = \beta Treat_i + \sum_{k=2013}^{2015} \delta_k \mathbb{1}(t = k) Treat_i + \sum_{k=2016}^{2019} \tau_k \mathbb{1}(t = k) Treat_i + \gamma \mathbf{X}_{it} + \epsilon_i \tag{8}$$

Figures 3(a) and (b) plot the coefficients from Equation 8. While the average treatment effect is not statistically significant, the event study shows a slight decline with noisy estimates, suggesting that the productivity impact among China-based researchers could become more negative if the trend continues.

Ultimately, these results indicate that the rise in tensions—and specifically the decline in access to U.S. research—did not, or at least had not yet by 2019, significantly affected the productivity of U.S.-reliant China-based researchers.

4.3.2 Productivity Impact on U.S.-based Researchers

We next examine whether rising U.S.–China tensions affected the productivity of ethnically Chinese STEM researchers in the U.S. As with the analysis of China-based researchers, we focus on a sub-sample of the Researcher Panel: U.S.-based STEM researchers with at least one publication between 2008 and 2019. For this analysis, we define the treated group as ethnically Chinese researchers and the control group as non-ethnically Chinese researchers.

Again, we match each treated researcher with a non-ethnically Chinese control

[26] The country of the journal is determined by the location of its publishing house.

[27] Given that we are analyzing recent years, we do not use forward citations to the publications themselves for quality weighting as the forward citations of recent publications would be truncated.

Table 6: Main Treatment Effects on Productivity among China-based Researchers

	Extensive Margin		Intensive Margin			
	(1) Pubs	(2) U.S. Pubs	(3) Pubs	(4) U.S. Pubs	(5) IF wted Pubs	(6) IF wted U.S. Pubs
US-Reliant=1	0.079	0.115				
	(0.049)	(0.062)				
US-Reliant=1 × Post-2016=1			-0.024	-0.003	-0.005	-0.005
			(0.016)	(0.027)	(0.018)	(0.032)
Indiv FE	N	N	Y	Y	Y	Y
Year FE	N	N	Y	Y	Y	Y
Model	Cox	Cox	Poisson	Poisson	Poisson	Poisson
CEM	Y	Y	Y	Y	Y	Y
Mean DV	2,018.594	2,018.200	3.298	1.120	6.891	2.667
Observations	9,566	9,566	54,276	48,408	54,276	48,408

Notes: Robust standard errors in parentheses and clustered at the person level. The dependent variable is in column header. The analysis sample is China-based researcher panel. The analysis period is 2013–2019, where the post treatment period is 2016–2019. "U.S.-Reliant" refers to the China-based researchers whose fraction of pre-2013 raw citation share is above the within-field 75th percentile to U.S. papers and below the 25th percentile to U.K. papers. The control group is China-based researchers with above 75th percentile within field U.K. citation share and below 25th percentile within field U.S. citation share. Column (1) and (2) report the estimated coefficients from a Cox proportional hazards model. The failure event in this analysis is defined as stopping publication. The analytical sample is the matched CEM sample. Column (3)-(6) reports the estimated coefficients from Poisson model. The regression is weighted by the CEM matching weight. Intensive margin specifications also include the treatment indicator, the post dummy, year fixed effects, and individual fixed effects.

researcher using Coarsened Exact Matching based on characteristics measured between 2008 and 2012. These include research output (number and growth of publications and impact-factor-weighted publications), collaboration patterns (number of coathors, number of foreign coauthors, number of China-based coauthors), career age, and institutional factors (university affiliation, ever having a foreign address, ever receiving foreign funding).

Table A7 in the Appendix reports the summary statistics of the covariates to show balance across matched treatment and control groups. In total, the sample on which we analyze the productivity impact on U.S.-based researchers contains 1,913,212 observations, with 45,177 unique ethnically Chinese individuals and 228,139 unique non-ethnically Chinese individuals.

Table 7 Columns (1)-(2) report extensive margin effects (Equation 6) using a Cox hazard model. Column (1) shows a 10% (β = 0.095, SE = 0.011) increase in the hazard of stopping publishing altogether for ethnically Chinese scientists in the U.S., while Column (2) shows a 7% increase in the hazard of stopping publishing in U.S.-based journals (β = 0.072, SE = 0.011, equivalent to a hazard ratio of 1.0747). Both estimates are statistically significant at the 1% level. Appendix Figures A9(c)

Table 7: Main Treatment Effects on Productivity among U.S. based Researchers

	Extensive Margin		Intensive Margin			
	(1)	(2)	(3)	(4)	(5)	(6)
	Pubs	U.S. Pubs	Pubs	U.S. Pubs	IF wted Pubs	IF wted U.S. Pubs
Ethnic CN=1	0.023	0.022				
	(0.006)	(0.006)				
Ethnic CN=1 × Post-2016=1	0.095	0.072	-0.082	-0.112	-0.087	-0.117
	(0.011)	(0.011)	(0.009)	(0.010)	(0.011)	(0.013)
Indiv FE	N	N	Y	Y	Y	Y
Year FE	N	N	Y	Y	Y	Y
Model	Cox	Cox	Poisson	Poisson	Poisson	Poisson
CEM	Y	Y	Y	Y	Y	Y
Mean DV	2,017.898	2,017.670	1.075	0.766	3.176	2.488
Observations	273,316	273,316	1,889,902	1,509,991	1,889,902	1,509,991

Notes: The dependent variable for each model is listed in each column header. The sample is the U.S.-based researcher panel from 2013-2019. The treatment group ("Ethnic CN") comprises researchers identified as ethnically Chinese by name; the control group is non-ethnically Chinese researchers. Column (1) and (2) report the estimated coefficients from Cox proportional hazards models, where the failure event is defined as stopping publication. The analytical sample is the matched CEM sample. Columns (3)-(6) report the estimated coefficients from a Poisson model. The regressions are weighted using CEM. Intensive margin specifications also include year fixed effects, and individual fixed effects. Robust standard errors are in parentheses and clustered at the person level.

and A9(b) report the Kaplan-Meier Curves that show survival step function curves for the treated and control groups.

A back-of-the-envelope calculation suggests that this decline corresponds to 1,858 additional ethnically Chinese researchers stopping publishing. To get this estimate, we multiply the number of ethnically Chinese researchers who would have dropped out in the absence of rising tensions by the 10% hazard.[28]

Columns (3)-(6) of Table 7 show the results of estimating Equation 7 on the intensive margin using the difference-in-differences approach. Across all four measures of productivity, shown in Columns (3)-(6), U.S.-China tensions had a statistically and economically significant negative effect on the productivity of U.S.-based ethnically Chinese researchers. The estimated declines range from -0.082 (SE = 0.009) for overall publication count to -0.117 (SE = 0.013) for impact-factor-weighted U.S.-based publications. These effects correspond to an 8.2% decrease in overall productivity and an 8.7% decrease in the production of impact-factor-weighted publications, relative to matched non-ethnically-Chinese counterparts.

[28]The number of ethnically Chinese who would have dropped out is calculated by $N \times r \times p$, where N is the sample size, the baseline dropout rate $r = 1 - \frac{\text{No. researchers in 2013}}{\text{No. of researchers in 2019}}$, and percent Chinese $p = \frac{\text{No.ethnically Chinese researchers in 2013}}{\text{No. researchers in 2013}}$.

Figure 3: Event-Study Plots for Productivity Impacts

China-based Researchers

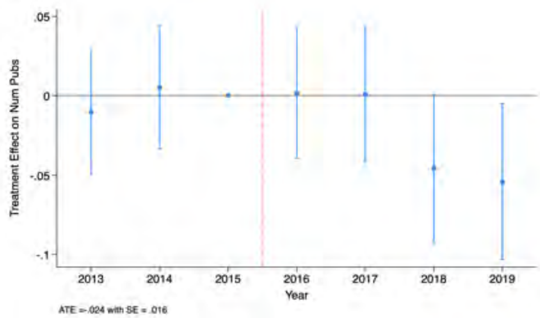

(a) DV: Pubs

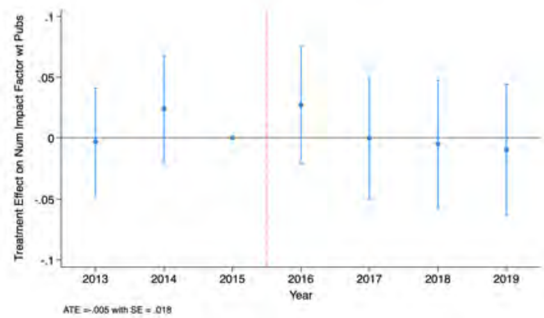

(b) DV: IF wt Pubs

U.S.-based Researchers

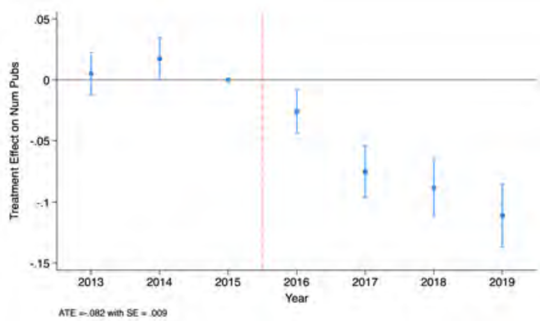

(c) DV: Pubs

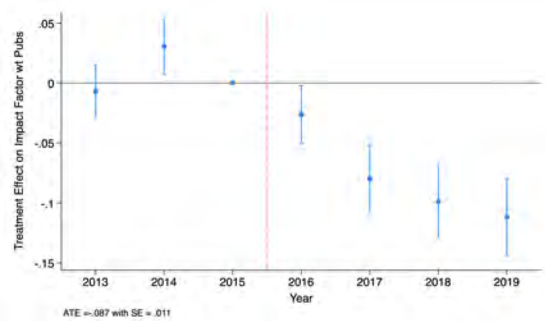

(d) DV: IF wt Pubs

Notes: This plot reports event-study coefficients from Poisson regressions. The dependent variable is indicated in the subfigure titles. Panels (a) and (b) use the China-based researcher panel; the treated group is U.S.-reliant researchers and the control group is U.K.-reliant researchers. Panels (c) and (d) use the U.S.-based researcher panel; the treated group is the U.S.-based ethnically Chinese researchers and the control group is the matched non-ethnically Chinese researchers. Regressions include individual and year fixed effects. Standard errors are clustered at the individual level.

The corresponding event study graphs (Figures 3(c) and (d)) show a clear trend break beginning in 2016, with declines in productivity persisting through 2019.[29]

We also perform a back-of-the-envelope calculation for estimating the decline in the number of publications among active ethnically Chinese researchers. Specifically, we multiply the estimated coefficient by each post-treatment year's total

[29] Event study graphs for other outcomes of interest are available upon request.

publications produced by ethnically Chinese researchers. Summing these products, an eight percent decrease in productivity equates to 92,197 publications that could have been published by ethnically Chinese researchers after 2015.[30]

Both Table 7 and Figures 3(c) and (d) demonstrate that, following the rise in U.S.-China tensions, both productivity and publication quality of ethnically Chinese STEM researchers in the U.S. decreased relative to that of non-ethnically Chinese STEM researchers.

These results underscore a striking asymmetry in the effect of U.S.-China tensions on scientific productivity. Despite a documented decline in access to U.S. research, China-based researchers who had previously relied heavily on U.S. science show no significant drop in output (yet). In contrast, U.S.-based ethnically Chinese researchers experienced meaningful declines in both publication volume and quality.

4.4 Robustness Checks

Alternative Explanations. We test whether factors unrelated to U.S.-China tensions—such as domestic improvements in China, crowd-out from other countries, or changing research quality—could explain our findings. For the mobility results, one possibility is that students chose to remain in China due to better domestic opportunities or policies like the Thousand Talents Program. However, Appendix Figure A3 shows no increase in the propensity to enroll in Chinese universities after 2016, suggesting that they are not simply choosing to stay home. Another possibility is increased competition from India, which could have displaced Chinese applicants from U.S. doctoral programs due to fixed program sizes. Yet Figure A4(a) shows no uptick in the share of Indian enrollment after 2016, suggesting that this kind of crowd-out is unlikely to explain the observed relative decline in Chinese enrollment. A third possibility is that retention trends among Chinese graduates simply reflect a broader shift away from U.S. employers among all international students, perhaps due to worsening labor market conditions or increased opportunities elsewhere. But Figure A4(b) shows that post-graduation stay rates for Indian students are stable or rising. This suggests the effect is specific to Chinese graduates rather

[30]Note that this back-of-the-envelope calculation implicitly that all papers are solo-authored. Thus, this figure could be higher than the true impact.

than a general pattern among international STEM degree recipients.

We also consider whether changes in Chinese research quality explain our knowledge flows results. If Chinese science improved dramatically during this period, Chinese researchers might simply cite less foreign work overall. However, our design directly addresses this by comparing U.S. citation shares to those of the U.K., which serves as a benchmark. But results are also unchanged when we include field-year fixed effects, separate field and year fixed effects, or field-specific time trends (Appendix Table A14).

Sensitivity to Definitions and Sample Construction. We test robustness to alternate measurement definitions and sample restrictions. First, in our mobility analyses, we allowed up to ten years between a student's prior degree and doctoral study, and up to three years between U.S. degree completion and post-graduation employment. To test whether these windows affect our results, we restrict to samples allowing only a one-year lag at most. The results are unchanged in magnitude and significance (Appendix Table A11 Columns (1)-(2)).

Second, when analyzing the retention of ethnically Chinese U.S. graduates, we included students earning degrees at any level. One might wonder if our results are primarily driven by earners of bachelor's and master's degrees rather than research-oriented doctorates. Appendix Table A12 repeats the main analysis using only U.S. graduates from doctoral programs, finding similar declines in trainee retention alongside increases in mobility to other anglophone countries.

A third concern might be that comparisons with students from countries that rarely send doctoral students to the U.S. might distort our estimates. While our baseline models include country fixed effects to address this, we further test robustness by restricting the sample to students from regions with strong U.S. STEM enrollment—Asia, Africa, and Latin America. Table A11 (Column 3)) shows no significant change in the estimated post-2016 enrollment decline for ethnically Chinese students.

Next, in our knowledge flows analysis, we test sensitivity to our definition of frontier research using thresholds of 3% and 5% rather than 1% and find similar declines in U.S. citation shares (Appendix Table A13).

Finally, in the productivity analysis of China-based researchers, we define U.S.-reliant scientists as those with U.S. citation shares above the 75th percentile (and U.K. shares below the 25th percentile). We repeat the analysis using alternative thresholds—50/50 and 90/10—and find similar results: no statistically significant changes in productivity under any definition (Appendix Table A16). Likewise, when we use recent citation shares instead of cumulative shares to define U.S. reliance, results remain unchanged (Appendix Table A16).

Parallel Trends and Pre-Treatment Tests. A key identifying assumption behind our difference-in-differences framework is that, in the absence of treatment, the treated and control groups would have followed parallel trends. To assess the plausibility of this assumption, we implement several diagnostic tests. First, we estimate event study models for each outcome, plotting pre-treatment coefficients to test for baseline trend similarity. Across analyses of mobility, knowledge flows, and productivity, we find that pre-2016 coefficients are small and statistically indistinguishable from zero.

Second, we apply the "pretrends" test developed by Roth (2022) to assess the statistical power of our design to detect violations of the parallel trends assumption. This test does not test for the presence of pre-trends per se, but rather helps evaluate how much power our design has to detect them, and how likely it is that conditioning on passing a pre-trends test could induce spurious results. We report this diagnostic in Appendix Figures A5, A7, and A10 Panels (a) and (d).

Third, we apply the "HonestDiD" sensitivity test developed by Rambachan and Roth (2023), which evaluates how robust our estimates are to potential violations of the parallel trends assumption. Since the HonestDiD framework is designed for sharp interventions, its application to our setting—where treatment unfolds gradually—is imperfect. We therefore interpret these results cautiously and primarily as a stress test rather than a formal validation (Appendix Figures A5, A7, and A10 Panels (b) and (e)).

Finally, we conduct placebo tests to assess whether our findings could plausibly arise by chance. In these tests, we randomly reassign treatment status across observations—preserving sample size and covariate structure—and re-estimate our main specifications 1000 times. This procedure generates a distribution of placebo treatment effects under the null hypothesis of no true effect. In every case, the treatment effect from our actual data lies in the far tail of this distribution, well outside the range of estimates produced under random assignment (Appendix Figures A5, A7, and A10 Panels (c) and (f)).

4.5 Disentangling the Mechanism

Until this point, we have treated the 2016 onset of geopolitical tension as a bundled "treatment" comprising a range of loosely correlated policy changes, political rhetoric, and social sentiment. In this section, we unpack that bundle to explore the underlying mechanism more directly. Given the especially pronounced effects for ethnically Chinese individuals—–both students making mobility decisions and researchers already working in the U.S.—–we ask: what explains these outcomes, and why were they concentrated among this group?

While no single root mechanism can be pinpointed, we find evidence consistent with a chilling effect shaped by fear, discomfort, and perceived hostility affecting ethnically Chinese scientists and students regardless of nationality or direct ties to China. This chilling effect appears to operate not through formal exclusion, but through anticipatory self-selection and disengagement, shaping decisions about where to study, where to work, and how (or whether) to continue participating in scientific research. Importantly, this chilling effect may not reflect *true* anti-Chinese sentiment but instead reflect the *perception* of anti-Chinese sentiment in the U.S. Below, we outline five categories of evidence consistent with this interpretation:

Timing of Effects Relative to Onset of Formal Policies. The effects we observe begin to emerge in 2016, before the rollout of most formal U.S. government actions targeting Chinese scholars or institutions. The most prominent formal policies aimed at Chinese nationals and institutions all begin in 2018: the China Initiative; new F-1 visa restrictions for students in sensitive STEM fields, introduced by the U.S. State Department; and the U.S.-China trade war. The fact that meaningful changes in enrollment, retention, and knowledge flows predate these formal policy shifts suggests that these formal policies alone are unlikely to be the primary

Table 8: Heterogeneous Effects by Direct Ties to China

	Enrolls in U.S.		Number of Pubs		
	(1) Treatment = Ethnically CN With CN Tie	(2) Treatment = Ethnically CN Without CN Tie	(3) Treatment = Ethnically CN With CN Tie	(4) Treatment = Ethnically CN Without CN Tie	(5) Treatment = Ethnically CN Non-CN First Name
Treatment	0.0386 (0.00684)	0.0261 (0.00517)			
Treatment × Post-2016	-0.0393 (0.00958)	-0.0226 (0.00904)	-0.0815 (0.00906)	-0.0865 (0.0183)	-0.0553 (0.0113)
Unit of Analysis	Doctoral Student	Doctoral Student	Researcher-Year	Researcher-Year	Researcher-Year
Fixed Effects	Y	Y	Y	Y	Y
Model	OLS	OLS	Poisson	Poisson	Poisson
Mean DV	0.233	0.236	1.083	1.079	1.084
Observations	123877	113336	1828988	1636859	1722322

Notes: Robust standard errors are in parentheses. For doctoral students in Columns (1) and (2), "CN Tie" is defined as having listed any education in China prior to the terminal degree. The control group is non-ethnically Chinese doctoral students, and standard errors are clustered at the field-year level. For researchers in Columns (3)-(5), "CN tie" is defined as meeting any of four criteria: having any Chinese coauthor, having any Chinese grant, citing China-produced science, or having a prior Chinese address. The control group is non-ethnically Chinese scholars matched using CEM. Standard errors are clustered at the person level.

driver of the observed changes among ethnically Chinese students and scholars. Instead, the data are consistent with a mechanism in which rhetoric, public sentiment, and perceived risk generated behavioral responses in advance of any concrete government action.

Effects Extend Beyond Those Directly Targeted. We find that the effects are not confined to individuals explicitly targeted by U.S. policy. First, as shown in Columns (1) and (2) of Table 8, the decline in U.S. doctoral enrollment is evident not only among ethnically Chinese students educated in China but also among ethnically Chinese students in other countries with no direct institutional ties to China. Second, mobility declines are steepest among students with strong outside options——specifically, those from top-ranked undergraduate institutions——as previously shown in Column (5) of Table 2. If formal restrictions were the main driver, we would expect a more uniform drop in the enrollment of nationally Chinese students. Instead, the disproportionate decline among top-tier students suggests that those with strong outside options may have self-selected away from the U.S. in response to perceived reputational or professional risk.

Third, as shown in Columns (3)–(4) of Table 8, we observe declines in productivity among U.S.-based ethnically Chinese scientists regardless of whether those individuals have had coauthors in mainland China, received Chinese government funding, cited China-produced research in their work, or previously held a mainland China address. This uniform decline across both affiliated and unaffiliated researchers points to a broader chilling effect in which ethnically Chinese scientists—irrespective of their direct connections to mainland China—may perceive themselves as vulnerable, leading to reduced productivity even in the absence of formal targeting. Notably, we also observe similar declines among researchers who are likely to be second-generation Chinese—those with ethnically Chinese last names but non-Chinese first names[31] in Column 5—suggesting that the chilling effect extends to individuals who may never have had direct ties to China.[32]

While our primary interpretation emphasizes perceived hostility, the decline in productivity among ethnically Chinese researchers in the U.S. may also reflect a compounding mechanism related to reduced access to Chinese graduate students. As shown earlier, Chinese trainees are increasingly opting for non-U.S. destinations. Given the central role that graduate students play in academic research—particularly in lab-based STEM fields—this shift may disproportionately affect faculty who have historically relied on Chinese students in their research teams. Importantly, prior work shows that ethnic ties and shared backgrounds play a meaningful role in advisor-student matching and academic science (e.g., Freeman and Huang, 2015; Gaule and Piacentini, 2018; Fry and Glennon, 2025), suggesting that ethnically Chinese faculty may be more affected by the declining pipeline of Chinese graduate students than their non-Chinese peers. As such, the chilling effect on productivity may be amplified not only by perceived hostility but also by the erosion of a key source of research support.

Heterogeneous Effects by Political Climate. We find that the chilling effect varies meaningfully with local political context. In particular, the effects are significantly

[31] For details on how we classified these researchers and for sensitivity tests to that classification, see Appendix E.

[32] These patterns hold when using impact-factor-weighted publication counts; results are available upon request.

larger in Republican-leaning (red) states, which are also perceived to have more anti-Chinese sentiment and restrictive policies targeting Chinese nationals.[33] As shown in Columns (1) and (2) of Table 9, the decline in doctoral enrollment among ethnically Chinese students is much more pronounced in these states than in Democrat-leaning (blue) states.

Importantly, we also observe evidence of a decline in interest before enrollment decisions are made. Using department-level Ph.D. application data from one red-state university and one blue-state university, Columns (3) and (4) of Table 9 show no relative decline in Chinese applications at the blue-state university but a sharp relative drop at the red-state university after 2016.[34] This provides more evidence of a chilling effect: the decline occurs at the application stage, before admission decisions or visa processes could plausibly intervene.

Together, these findings suggest that local political climate—shaped by rhetoric and policy—meaningfully influences student behavior, amplifying the chilling effect in certain states. They also underscore that the observed declines appear not to be driven by changes in acceptance or visa policy, but rather by a shift in where students choose to apply in the first place.

No Change in Visa Denials. The effects we observe are also unlikely to be driven by changes in visa policy, in the form of either formal changes in visa policy or changes in the rejection rate. Importantly, any changes would need to be both China-specific (or our control group would experience the same effect) and take effect around 2016. There were visa restrictions announced for Chinese students in 2018, but that was two years after we begin to see changes in U.S. doctoral enrollment. While formal policy changes came later, one might still wonder whether visa denials quietly increased during this period. However, Figure 4 of Bier (2024), based on data from Chen, Howell and Smith (2023), shows no evidence of a discon-

[33] For example, a 2023 Florida law halted recruitment of Chinese students and sparked widespread concern. See: Qin and Mazzei (2023). https://www.nytimes.com/2023/12/15/us/florida-law-chills-chinese-student-recruitment.html

[34] We requested doctoral application data from U.S. universities, using sunshine laws for public institutions and direct outreach for private ones. Only two universities provided department-level application data by nationality for the time period of interest. We restricted analysis to STEM doctoral programs in arts and science, engineering schools, and related divisions, excluding professional schools. This analysis was reviewed by the Boston University IRB (Protocol #7974X).

Table 9: Heterogeneous Effects by Political Climate

	Enrolls in U.S.		ln(# of Applications)	
	(1) Blue State Univ.	(2) Red State Univ.	(3) Blue State Univ.	(4) Red State Univ.
Treatment = Ethnically CN	0.0141 (0.00487)	0.0103 (0.00476)		
Treatment = Nationally CN			0.164 (0.177)	0.167 (0.112)
Treatment × Post-2016	0.00494 (0.00597)	-0.0348 (0.00566)	0.305 (0.249)	-0.381 (0.156)
Unit of Analysis	Doctoral Student	Doctoral Student	Degree Program-Year	Degree Program-Year
Fixed Effects	Y	Y	Y	Y
Model	OLS	OLS	OLS	OLS
Mean DV	0.124	0.138	3.181	2.659
Observations	92069	92069	377	709

Notes: Columns (1) and (2) present difference-in-differences estimates of the impact of Chinese ethnicity on the probability of enrolling in a U.S. university post-2016, split by its state's political leanings. Fixed effects are included for the doctoral cohort, field, prior country, and prior university ranking. In Columns (3) and (4), we regress the log-transformed number of applications to STEM doctoral programs at two universities on the treatment—the applicant being a Chinese national—the interaction term, as well as fixed effects for both year and degree program.

tinuity in F-1 visa rejection rates for Chinese students relative to other international students in 2016.

Perceptions: Survey Evidence and Behavioral Signatures of Fear. Finally, we draw on contemporaneous survey evidence, which reveal a high and disproportionate level of fear and uncertainty among Chinese and Chinese-American scientists working in the U.S. In one 2021 survey, 50.7% of Chinese scientists reported considerable fear of U.S. government surveillance, compared to just 11.7% of non-Chinese scientists (Lee and Li, 2021). Another survey found that 72% of U.S.-based scientists of Chinese descent "do not feel safe as academic researchers," 42% are "fearful of conducting research," and 61% have considered leaving the United States (Xie et al., 2023). These perceptions appear widespread even among those not directly affected by formal investigations, pointing to a generalized chilling effect that extends beyond policy enforcement.

This perception of being at risk also manifests in more subtle behavioral adaptations. For example, while we documented earlier that U.S.-based research teams on average show no significant change in their citations to Chinese research after

2016, we find in Table A15 that U.S.-based teams with ethnically Chinese coauthors are significantly less likely to cite Chinese work. This drop is difficult to explain in terms of scientific relevance or access and is instead consistent with a reputational or anticipatory avoidance mechanism, where individuals modify their citation practices to avoid signaling ties to China. Taken together, the survey evidence and this behavioral response offer both direct and indirect indications of a chilling effect driven by identity and perception rather than policy alone.

5 Discussion and Conclusion

Our results reveal that U.S.-China tensions, by the end of 2019, had already significantly disrupted international science, reducing cross-border talent and knowledge flows, and diminishing the productivity of ethnically Chinese scientists in the U.S. Specifically, ethnically Chinese graduate students became both less likely (15%) to attend a U.S.-based doctoral program and, if they did attend a U.S.-based program, were less likely (4%) to stay in the U.S. after graduation. China-based researchers cite U.S. science less frequently—particularly recent, frontier work—while U.S. scientists continue citing China-produced research at similar rates. Productivity impacts were asymmetric: U.S.-based ethnically Chinese researchers were 8-11% less productive and 10% more likely to stop publishing, while China-based researchers reliant on U.S. science did not see comparable declines.

These patterns suggest a "chilling effect" driven by perceived hostility and reputational risks, rather than formal policy alone. We find supporting evidence across five dimensions: (1) effects appear before most formal policies took effect, (2) they extend to ethnically Chinese researchers and students with no ties to China, (3) vary by local political climate, including early declines in applications to U.S. graduate programs in more conservative states, (4) cannot be explained by visa denials, and (5) align with survey data showing widespread fear among Chinese scientists in the U.S.

Our findings contribute to the literature on how conflict and nationalism affect science, offering new evidence that even relatively low-level geopolitical friction can meaningfully disrupt international science. These effects are likely a lower bound: the loss of top STEM talent may play out beyond our study window. In addition, tensions have only intensified since 2019, initially fueled by anti-Asian sentiment during the COVID-19 pandemic and now deepened by renewed nationalist policies under the second Trump administration and continued decoupling efforts by China. For instance, the Trump administration recently introduced targeted restrictions on Chinese nationals in STEM fields, including the revocation of F-1 visas for certain applicants and increased scrutiny of research visas connected to Chinese institutions. On the Chinese side, policies such as the "Delete America" directive continue to push for technological self-reliance and the removal of U.S. technology from Chinese systems.

These findings also carry clear policy implications. By documenting asymmetric declines in talent and knowledge flows, we provide early empirical evidence of scientific decoupling. Even in the absence of formal exclusion, we observe substantial behavioral shifts among ethnically Chinese researchers and trainees. More explicit barriers—–such as the recent revocation of F-1 visas for some international students—–are likely to intensify these effects. Policies aimed at safeguarding national research and innovation may inadvertently drive away talent, reduce access to frontier ideas, and weaken scientific productivity. Countries that maintain neutrality may benefit from redirected flows; indeed, Anglophone countries appear to be attracting ethnically Chinese trainees who might once have gone to the U.S., potentially shifting the geography of innovation.

Important questions remain. We do not fully quantify how these changes have affected the quality composition of talent in the U.S., or precisely trace how mobility shifts impact knowledge flows and productivity. Future work should examine these linkages and assess the broader welfare implications. While this paper highlights several adverse consequences for science, those costs must ultimately be balanced against national security concerns, particularly in domains where collaboration may pose real risks. More granular field classifications could help policymakers balance openness and security.

References

Abramitzky, Ran, and Isabelle Sin. 2014. "Book Translations as Idea Flows: The

Effects of the Collapse of Communism on the Diffusion of Knowledge." *Journal of the European Economic Association*, 12(6): 1453–1520.

Aghion, Philippe, Celine Antonin, Luc Paluskiewicz, David Stromberg, Xueping Sun, Raphael Wargon, and Karolina Westin. 2023. "Does Chinese Research Hinge on US Coauthors? Evidence from the China Initiative." Working Paper, Collège de France.

Agrawal, Ajay, John McHale, and Alexander Oettl. 2017. "How stars matter: Recruiting and peer effects in evolutionary biology." *Research Policy*, 46(4): 853–867.

Becker, Sascha O., Volker Lindenthal, Sharun Mukand, and Fabian Waldinger. 2021. "Scholars at Risk: Academic Networks and High-Skilled Emigration from Nazi Germany." *The Warwick Economics Research Paper Series (TWERPS)*.

Bernstein, Shai, Rebecca Diamond, Abhisit Jiranaphawiboon, Timothy McQuade, and Beatriz Pousada. 2022. "The Contribution of High-Skilled Immigrants to Innovation in the United States." NBER Working Paper #30797.

Bier, David J. 2024. "A Record Quarter of a Million International Students Denied Visas, 36 Percent of Applicants." Cato at Liberty, Cato Institute. https://www.cato.org/blog/record-quarter-million-international-students-denied-visas-36-applicants.

Black, Grant C., and Paula E. Stephan. 2010. "The Economics of University Science and the Role of Foreign Graduate Students and Postdoctoral Scholars." In *American Universities in a Global Market*. 129–161. University of Chicago Press.

Borjas, George J., and Kirk B. Doran. 2012. "The Collapse of the Soviet Union and the Productivity of American Mathematicians*." *The Quarterly Journal of Economics*, 127(3): 1143–1203.

Chen, Mingyu, Jessica Howell, and Jonathan Smith. 2023. "Best and brightest? The impact of student visa restrictiveness on who attends college in the US." *Labour Economics*, 84: 102385.

Freeman, Richard R., and Wei Huang. 2015. "Collaborating with People Like Me: Ethnic Coauthorship within the United States." *Journal of Labor Economics*, 33(S1).

Fry, Caroline, and Britta Glennon. 2025. "In Good Company: Coethnic Advisors and Career Paths of Immigrant Ph.D. Students." NBER Working Paper #33782.

Ganguli, Ina. 2017. "Saving Soviet Science: The Impact of Grants When Government R&D Funding Disappears." *American Economic Journal: Applied Economics*, 9(2): 165–201.

Ganguli, Ina, and Fabian Waldinger. 2023. "War and Science in Ukraine." NBER Working Paper #31449.

Gaule, Patrick, and Mario Piacentini. 2018. "An advisor like me? Advisor gender and post-graduate careers in science." *Research Policy*, 47(4): 805–813.

Herzog, Christian, Daniel Hook, and Stacy Konkiel. 2020. "Dimensions: Bringing down barriers between scientometricians and data." *Quantitative Science Studies*, 1(1): 387–395.

Hook, Daniel W., Simon J. Porter, and Christian Herzog. 2018. "Dimensions: Building Context for Search and Evaluation." *Frontiers in Research Metrics and Analytics*, 3. Publisher: Frontiers.

Hsiehchen, David, Magdalena Espinoza, and Antony Hsieh. 2015. "Multinational teams and diseconomies of scale in collaborative research." *Science Advances*, 1(8).

Hunt, Jennifer, and Marjolaine Gauthier-Loiselle. 2010. "How Much Does Immigration Boost Innovation?" *American Economic Journal: Macroeconomics*, 2(2): 31–56.

Iacus, Stefano M., Gary King, and Giuseppe Porro. 2012. "Causal Inference without Balance Checking: Coarsened Exact Matching." *Political Analysis*, 20(1): 1–24.

Iaria, Alessandro, Carlo Schwarz, and Fabian Waldinger. 2018. "Frontier Knowledge and Scientific Production: Evidence from the Collapse of International Science*." *The Quarterly Journal of Economics*, 133(2): 927–991.

Jia, Ruixue, Margaret E. Roberts, Ye Wang, and Eddie Yang. 2024. "The impact of US–China tensions on US science: Evidence from the NIH investigations." *Proceedings of the National Academy of Sciences*, 121(19): e2301436121.

Kerr, William R. 2008. "Ethnic Scientific Communities and International Technology Diffusion." *The Review of Economics and Statistics*, 90(3): 518–537.

Lee, Jenny, and Xiaojie Li. 2021. "Racial Profiling Among Scientists of Chinese Descent and Consequences for the U.S. Scientific Community." Committee of 100 and the University of Arizona.

Miguelez, Ernest. 2018. "Inventor Diasporas and the Internationalization of Technology." *The World Bank Economic Review*, 32(1): 41–63.

Moser, Petra, Alessandra Voena, and Fabian Waldinger. 2014. "German Jewish Émigrés and US Invention." *American Economic Review*, 104(10): 3222–55.

Moser, Petra, and Shmuel San. 2020. "Immigration, Science, and Invention. Lessons from the Quota Acts." Available at SSRN: https://papers.ssrn.com/abstract=3558718.

Oettl, Alexander, and Ajay Agrawal. 2008. "International labor mobility and knowledge flow externalities." *Journal of International Business Studies*, 39(8): 1242–1260.

Rambachan, Ashesh, and Jonathan Roth. 2023. "A More Credible Approach to Parallel Trends." *Review of Economic Studies*, 90(5): 2555–2591.

Roth, Jonathan. 2022. "Pretest with Caution: Event-Study Estimates after Testing for Parallel Trends." *American Economic Review: Insights*, 4(3): 305–322.

Sim, Suzan, and A-Sung Hong. 2025. "National Conflict and High-Skilled Immigrants' Workplace Efforts: Evidence from the U.S.–China Conflict." *Organization Science*. Publisher: INFORMS.

Singh, Vivek Kumar, Prashasti Singh, Mousumi Karmakar, Jacqueline Leta, and Philipp Mayr. 2021. "The journal coverage of Web of Science, Scopus and Dimensions: A comparative analysis." *Scientometrics*, 126(6): 5113–5142.

Thelwall, Mike. 2018. "Dimensions: A competitor to Scopus and the Web of Science?" *Journal of Informetrics*, 12(2): 430–435.

Torvik, Vetle Ingvald, and Sneha Agarwal. 2016. "Ethnea – an instance-based ethnicity classifier based on geo-coded author names in a large-scale bibliographic database." Conference Paper, International Symposium on Science of Science.

USCET. 2023. "Three Decades of Chinese Students in America, 1991-2021." US-China Education Trust. https://chinadatalab.ucsd.edu/uscet/three-decades-of-chinese-students-in-america-1991-2021/.

Waldinger, Fabian. 2010. "Quality Matters: The Expulsion of Professors and the Consequences for PhD Student Outcomes in Nazi Germany." *Journal of Political Economy*, 118(4).

Waldinger, Fabian. 2012. "Peer Effects in Science: Evidence from the Dismissal of Scientists in Nazi Germany." *The Review of Economic Studies*, 79(2): 838–861.

Waldinger, Fabian. 2016. "Bombs, Brains, and Science: The Role of Human and Physical Capital for the Creation of Scientific Knowledge." *The Review of Economics and Statistics*, 98(5): 811–831.

Xie, Yu, Xihong Lin, Ju Li, Qian He, and Junming Huang. 2023. "Caught in the crossfire: Fears of Chinese-American scientists." *Proceedings of the National Academy of Sciences of the United States of America*, 120(27): e2216248120.

Supplemental Appendix

A Data Construction

In this section, we provide additional details on the construction of the datasets used for analysis. Figure A1 provides an overview of these constructed datasets, and Table A2 describes how each is leveraged in our empirical analyses.

Table A1: Summary of Data Sources & Sample Construction

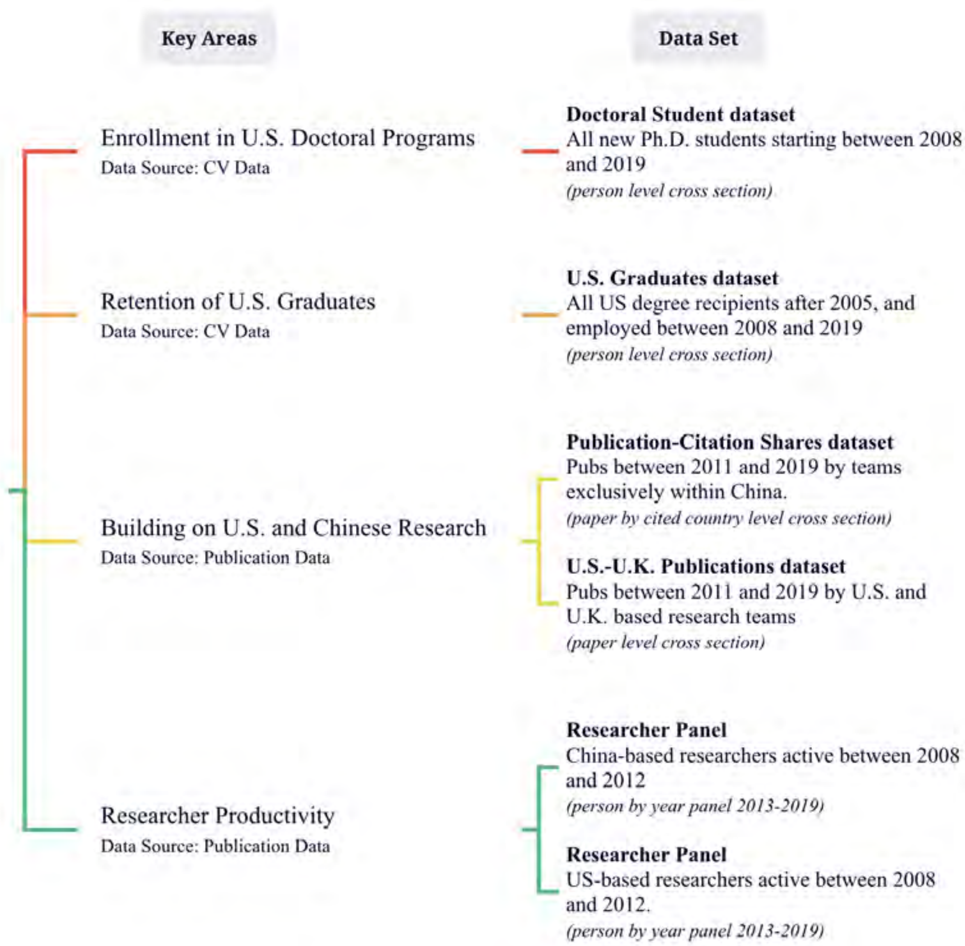

Notes: This figure summarizes the data sources and constructed data sets used across all analyses. Data sets for analyses of trainee mobility (enrollment and retention) derive from ORCID. Data sets for analyses of knowledge flows and researcher productivity derive from Dimensions from Digital Science.

Table A2: Applied Differences-in-Differences Across Analyses

		Sample	Main DV	Treated	Control
Student Mobility	*China*	Doctoral Students	Enrolls in U.S.	Ethnically Chinese	Non-Ethnically Chinese
	U.S.	U.S. Graduates	Job in U.S.	Ethnically Chinese	Non-Ethnically Chinese
Knowledge Flows	*China*	Publication-Citation Shares	Size	Citing U.S.	Citing U.K.
	U.S.	U.S.-U.K. Publications	Citation Share to Chinese Research	U.S. Publication	U.K Publication
Researcher Productivity	*China*	China-Based Researchers	# of publications	Predominantly Citing U.S.	Predominantly Citing U.K.
	U.S.	U.S.-Based Researchers	# of publications	Ethnically Chinese	Non-Ethnically Chinese

Notes: This table describes the basic components of each differences-in-differences analysis conducted in Section 4. in each analysis. The treatment period begins in 2016. More detailed considerations (e.g., fixed effects) are described alongside each analysis.

A.1 Enrollment and Retention Dataset

We use ORCID data to analyze enrollment in U.S. doctoral programs and post-graduation job placement, as described in Section 3.1 of the main text. We summarize key steps and note additional construction details not included in the main body. Detailed summary statistics on both samples are provided in Table A3(a) and Table A3(b).

From the 2022 ORCID release, we parse educational and employment histories from publicly available CVs. We restrict to individuals with complete educational records (n = 1.8M)—i.e., if each of their education spells includes both start and end dates—and who are in STEM fields (n = 836,495). Degrees are classified based on keywords in the degree title (e.g., a doctoral degree would contain "Ph.D.," "PhD," "Doctoral", etc.).

Doctoral Student Dataset. We identify individuals whose unique terminal degree[35] is a doctorate and who have at least one prior degree listed. We infer nation-

[35]Based on which began last for a given individual.

Table A3: Summary Statistics for Mobility Analysis Datasets

	Mean	SD	p25	p50	p75	Count
Ph.D. first year	2,012.79	2.88	2,010	2,013	2,015	129,223
Ph.D. last year	2,017.22	2.89	2,015	2,017	2,019	129,223
Treatment = ethnically Chinese	0.16	0.37	0	0	0	129,223
Origins						
Prior degree in U.S.	0.18	0.38	0	0	0	129,223
Prior degree in China	0.11	0.32	0	0	0	129,223
Prior degree in India	0.10	0.30	0	0	0	129,223
Prior degree in U.K.	0.07	0.26	0	0	0	129,223
Destinations						
Enrolls in U.S. university	0.24	0.43	0	0	0	129,223
Enrolls in U.K. university	0.09	0.29	0	0	0	129,223
Enrolls in Chinese university	0.07	0.25	0	0	0	129,223
Enrolls in Indian university	0.08	0.27	0	0	0	129,223
Enrolls in non-U.S. anglo. university	0.17	0.37	0	0	0	129,223
Research fields						
Science & Engineering program	0.79	0.41	1	1	1	129,223
Medicine/Health program	0.21	0.41	0	0	0	129,223

(a) Doctoral Students Dataset (Enrollment)

	Mean	SD	p25	p50	p75	Count
Job first year	2,014.92	3.16	2,013	2,016	2,018	50,890
Graduation year	2,014.46	3.30	2,012	2,015	2,017	50,890
Ethnically CN	0.18	0.38	0	0	0	50,890
Terminal degree is Ph.D.	0.71	0.45	0	1	1	50,890
Lag from degree to job	0.46	0.80	0	0	1	50,890
Destinations						
Job in U.S.	0.85	0.35	1	1	1	50,890
Job in U.K.	0.01	0.10	0	0	0	50,890
Job in China	0.02	0.13	0	0	0	50,890
Job in non-U.S. anglo. country	0.03	0.16	0	0	0	50,890
Research fields						
Science & Engineering program	0.77	0.42	1	1	1	50,890
Medicine/Health program	0.23	0.42	0	0	0	50,890

(b) U.S. Graduates Dataset (Retention)

Notes: Panels (a) and (b) describe basic summary statistics for observations in the Doctoral Students and U.S. Graduates data sets, respectively. In Panel (a), the unit of observation is a doctoral student originating from anywhere in the world. In Panel (b), the unit of observation is a U.S. graduate seen to subsequently take a job anywhere in the world.

ality from the location of this prior degree. Enrollment location (U.S., U.K., China, etc.) is also derived from the terminal degree institution. To proxy for university quality, we match the prior (i.e., the pre-doctoral degree) institutions to 2023 research rankings from SCIMago.

U.S. Graduates Dataset. We identify individuals whose unique terminal degree is from a U.S. institution and who lists a job within three years of graduation. Location of post-graduate employment is based on the job entry immediately following the U.S. degree.

A.2 Knowledge Flows Dataset

We use bibliometric data provided by Dimensions from Digital Science to construct citation-based measures of cross-border knowledge flows, as outlined in Section 3.2 of the main text. Dimensions is a comprehensive bibliometric database that that includes metadata and citation links for over 140 million publications across fields and countries (Thelwall, 2018; Singh et al., 2021).[36] Its extensive coverage and algorithmic author disambiguation make it well suited to longitudinal analyses of researcher behavior.[37]

Publication references serve as a large-scale and observable proxy for scientific knowledge flows in the form of trace data (Iaria, Schwarz and Waldinger, 2018). As such, we derive datasets from Dimensions describing the citation behavior of publications written by Chinese, U.S., and U.K. research teams between 2011 and 2019, proceeding in two steps: (1) identifying publications by research teams in each nation and (2) quantifying the degree to which each cites papers from foreign countries.

First, for each publication in Dimensions, we create flags for whether all of the authors with address data list their first affiliation address as being in China, the U.S., or the U.K. We refer to these publications as being written by China-based, U.S., or U.K. research teams respectively. We discard publications where authors list addresses from more than one of these three countries (e.g., publications where all authors claim a Chinese address but one author also claims a U.S. address). Further, we discard publications that are not categorized as belonging to a STEM field.

We then employ four measures of how publications build on research produced in the U.S., the U.K., or China. Given a focal (citing) publication, we identify the (cited) papers on its reference list. For each cited paper, we assign an origin country based on the first affiliation address of its corresponding author. If the cited paper has no corresponding author, but involves researchers exclusively

[36] More information about the Dimensions database can be found in Hook, Porter and Herzog (2018) and Herzog, Hook and Konkiel (2020).

[37] Most publications in Dimensions are in English, which should be considered when interpreting patterns of international knowledge flows. See https://www.digital-science.com/tldr/article/in-the-spotlight-english-as-the-lingua-franca-in-science/ for more details.

Table A4: Summary Statistics for Knowledge Flows Datasets

	Mean	SD	p25	p50	p75	Count
Treated = citing U.S.	0.50	0.50	0	.5	1	4,237,614
Share of raw references	0.13	0.15	0	.0667	.2	4,237,614
Share of recent references	0.09	0.15	0	0	.125	4,042,500
Share of frontier references (1%)	0.19	0.28	0	0	.333	3,332,908
Share of recent frontier references (1%)	0.16	0.29	0	0	.25	2,303,014
Citing paper attributes						
Publication year	2,015.69	2.56	2,014	2,016	2,018	4,237,614
Number of fields	1.24	0.46	1	1	1	4,237,614
Field: biology	0.11	0.31	0	0	0	4,237,614
Field: biomedical	0.21	0.41	0	0	0	4,237,614
Field: chemistry	0.20	0.40	0	0	0	4,237,614
Field: engineering	0.38	0.48	0	0	1	4,237,614
Field: health	0.02	0.15	0	0	0	4,237,614
Field: physics	0.07	0.26	0	0	0	4,237,614
Science & Engineering program	0.81	0.40	1	1	1	4,237,614
Medicine/Health program	0.22	0.42	0	0	0	4,237,614

(a) Summary Statistics for Publication-Citation Shares Dataset

	Mean	SD	p25	p50	p75	Count
Publication year	2,015.05	2.59	2,013	2,015	2,017	2,833,730
Number of fields	1.25	0.48	1	1	1	2,833,730
U.S. publication	0.83	0.38	1	1	1	2,833,730
Cite share to CN	0.02	0.04	0	0	.0137	2,833,730
Recent cite share to CN	0.03	0.07	0	0	0	2,756,324
Frontier cite share to CN (1%)	0.01	0.06	0	0	0	2,349,159
Recent frontier cite share to CN (1%)	0.02	0.09	0	0	0	1,826,557
Share of references citing U.S.	0.47	0.22	.325	.48	.622	2,833,730
Share of references citing U.K.	0.08	0.14	0	.0408	.103	2,833,730
Research fields						
Field: biology	0.14	0.35	0	0	0	2,833,730
Field: biomedical	0.48	0.50	0	0	1	2,833,730
Field: chemistry	0.08	0.27	0	0	0	2,833,730
Field: engineering	0.13	0.34	0	0	0	2,833,730
Field: health	0.15	0.35	0	0	0	2,833,730
Field: physics	0.06	0.24	0	0	0	2,833,730
Science & Engineering program	0.48	0.50	0	0	1	2,833,730
Medicine/Health program	0.57	0.50	0	1	1	2,833,730

(b) Summary Statistics for U.S.-U.K. Publications Dataset

Notes: Panels (a) and (b) describe basic summary statistics for observations in the Publication-Citation Shares and U.S.-U.K. Publications data sets, respectively. In Panel (a), the unit of observation is a publication-citation share to either U.S.- or U.K.-produced papers. In Panel (b), the unit of observation is a publication authored by a U.S. or U.K. research team.

from one country, we assign that country. Otherwise, we leave the cited paper's origin country blank. Then, for each citing paper, we use the assigned locations of its references to construct four measures of usage per country. First, we calculate the "raw" usage of a given country's research by dividing the number of cited papers assigned to that country by the total number of cited papers. Second, we calculate the "recent" usage of a country's research by dividing the number of cited papers assigned to that country and published within five years of the citing publication by the total number of similarly recent cited papers. Third, we calculate the "frontier" usage of a country's research by dividing the number of cited papers assigned to that country and landing in the top 1%, 3%, or 5% of their field's citation distribution (using Dimensions' field citation ratio measure) by the total number of cited papers at similar frontiers. Finally, we calculate the "recent frontier" usage of a country's research by dividing the number of cited papers assigned to that country and satisfying both of these restrictions by the total number of cited papers satisfying these restrictions.

In Table A4(a), we provide summary statistics on the dataset used for analyzing the changing usage of U.S.-produced research by China-based research teams (the *Publication-Citation Shares Dataset*). An observation in this dataset is a publication-citation share, where the citation share represents either references to U.S.-produced papers or references to U.K.-produced papers. These citation shares come from the publications of China-based research teams.

In Table A4(b), we provide summary statistics on the dataset used for analyzing the changing usage of China-produced research by U.S. research teams (the *U.S.-U.K. Publications Dataset*). An observation in this dataset is a publication written by a U.S. or U.K. research team, and the outcome of interest is the share of their references citing China-produced papers.

A.3 Productivity Dataset

We use the Dimensions database to construct a strongly balanced panel of STEM researchers between 2008 and 2019[38] (*the Researcher Panel*) and measure changes in publication productivity before and after the rise in U.S.-China tensions. As mentioned above, Dimensions provides comprehensive publication metadata and author disambiguation, which supports author-level tracking over time and across institutions. We require that researchers have at least one publication during our sample period: between 2013 and 2019 to be included in the panel.

We identify China-based and U.S.-based researchers based on their modal affiliation country between 2008 and 2012 and assign them to academic fields us-

[38]We choose 2019 as the end year to avoid shocks to scientific productivity associated with COVID-19.

ing the modal field of their publications. We restrict to researchers operating in STEM fields. For each researcher-year observation, we include the number of publications, the number of U.S.-based publications, and the impact-factor-weighted versions of both.

China-based Researcher Panel. We further filter the Researcher Panel to a subsample of China-based researchers who published five or more publications between 2008 and 2012 (our pre-analysis period), to ensure a sufficient number of observations to identify treatment status. For each China-based researcher, we calculated the share of references on their 2008–2012 publications that cite the U.S., U.K., and China, locating cited papers using the same approach as in Section A.2. We define the treatment group as China-based researchers whose pre-2013 references heavily relied U.S. papers (above the 75th percentile within their field for the U.S. and below the 25th for the U.K.) and the control group as those whose references predominantly cited U.K. papers (above the 75th percentile for the U.K. and below the 25th for the U.S.). To increase comparability, we match treated and control researchers using the coarsened exact matching (CEM) described in Iacus, King and Porro (2012).

The matching covariates include: the number of publications between 2008–2012 (binned into deciles); proxied career age as of 2012 (in four bins); the number of active publishing years from 2008–2012; whether the researcher has a university affiliation; whether the researcher is located in a Tier 1 City, or a New Tier 1 City; and both the level and growth rates of total and impact-factor-weighted publications from 2008–2012. The CEM algorithm matched 11,975 individuals, out of which 5,982 who heavily relied on scientific papers from the U.S. and 5,993 who heavily relied on papers from the U.K. We use matching weights generated by the CEM algorithm in the associated regression analyses. Balance statistics are reported in Tables A5 and A6.

U.S.-based Researcher Panel. We filter the Researcher Panel to U.S.-based researchers. We define ethnically Chinese researchers[39] as the treatment group and match these researchers to non-ethnically Chinese researchers the same field using

[39]See Section A.4 for ethnicity imputation details.

Table A5: China-based Researcher Panel Descriptive Statistics: CEM Matching Variables

	(1) All	(2) Matched	(3) Matched-Untreated	(4) Matched-Treated
Num Pubs in 2008-2012	13.20	8.081	8.135	8.027
	(10.31)	(4.124)	(4.147)	(4.101)
Career Age	7.282	5.895	5.953	5.837
	(4.401)	(3.449)	(3.533)	(3.363)
Num Active Years in 2008-2012	3.807	3.429	3.429	3.430
	(1.077)	(1.012)	(1.013)	(1.011)
1[University]	0.655	0.657	0.657	0.657
	(0.475)	(0.475)	(0.475)	(0.475)
1[Tier 1 Cities]	0.351	0.348	0.326	0.370
	(0.477)	(0.476)	(0.469)	(0.483)
1[New Tier 1 Cities]	0.344	0.351	0.351	0.352
	(0.475)	(0.477)	(0.477)	(0.478)
Growth Rate of Num Pubs 2008-2012	-0.00771	-0.00910	-0.00838	-0.00982
	(0.113)	(0.107)	(0.108)	(0.105)
Num Pubs	4.001	1.959	1.933	1.986
	(5.222)	(1.689)	(1.692)	(1.686)
Growth Rate of IF-wt Pubs 2008-2012	-0.00520	-0.00789	-0.00675	-0.00904
	(0.149)	(0.143)	(0.142)	(0.144)
Num Impact Factor Weighted Pubs	8.714	3.801	3.603	3.999
	(13.27)	(3.637)	(3.499)	(3.761)
Observations	136,736	11,977	5,994	5,983

Notes: Standard deviation in parentheses.

Table A6: China-based Researcher Panel Descriptive Statistics: Outcome Variables

	(1) All	(2) Matched	(3) Matched-Untreated	(4) Matched-Treated
Num Pubs	4.001	1.959	1.933	1.986
	(5.222)	(1.689)	(1.692)	(1.686)
Num Pubs in US-based journals	1.180	0.599	0.513	0.686
	(1.879)	(0.793)	(0.718)	(0.854)
Num Impact Factor Weighted Pubs	8.714	3.801	3.603	3.999
	(13.27)	(3.637)	(3.499)	(3.761)
Impact Factor Weighted US-based Pubs	3.001	1.328	1.057	1.599
	(5.951)	(1.969)	(1.622)	(2.230)
Observations	136,736	11,977	5,994	5,983

Notes: Standard deviation in parentheses.

CEM on a rich set of pre-2016 characteristics.

The matching variables include: total publications 2008–2012; proxied career age as of 2012; the number of active publishing years 2008–2012; whether the researcher has a university affiliation; fraction of the coauthors that are foreign; number of distinct foreign coauthors; whether the researcher has a Chinese coau-

Table A7: U.S.-based Researcher Panel Descriptive Statistics : CEM Matching Variables

	(1) All	(2) Matched	(3) Matched-Untreated	(4) Matched-Treated
Num Pubs in 2008-2012	8.968	3.669	3.526	4.394
	(15.90)	(4.801)	(4.639)	(5.490)
Career Age	10.29	4.481	4.366	5.062
	(9.795)	(4.359)	(4.228)	(4.930)
Num Active Years in 2008-2012	2.823	1.972	1.925	2.209
	(1.503)	(1.191)	(1.162)	(1.300)
1[University]	0.440	0.422	0.414	0.463
	(0.496)	(0.494)	(0.493)	(0.499)
Fraction of Foreign Coauthors	0.127	0.0258	0.0212	0.0486
	(0.210)	(0.101)	(0.0919)	(0.135)
Num of Distinct Foreign Coauthors	4.103	0.409	0.323	0.844
	(11.96)	(1.827)	(1.574)	(2.732)
[1em] 1[Have Foreign Address]	0.261	0.0385	0.0306	0.0756
	(0.439)	(0.193)	(0.172)	(0.264)
Fraction of Foreign Coauthors	0.127	0.0258	0.0212	0.0486
	(0.210)	(0.101)	(0.0919)	(0.135)
1[Have Foreign Funding]	0.286	0.0601	0.0475	0.123
	(0.452)	(0.238)	(0.213)	(0.329)
1[Have Coauthors with Foreign Funding]	0.810	0.738	0.720	0.830
	(0.393)	(0.440)	(0.449)	(0.376)
Growth Rate of Num Pubs 2008-2012	-0.0347	-0.0568	-0.0536	-0.0728
	(0.349)	(0.288)	(0.284)	(0.307)
Num Pubs	2.108	0.929	0.898	1.087
	(3.721)	(1.358)	(1.337)	(1.451)
Growth Rate of IF-wt Pubs 2008-2012	-0.0585	-0.0946	-0.0889	-0.123
	(0.561)	(0.473)	(0.464)	(0.519)
Num Impact Factor Weighted Pubs	6.563	2.579	2.444	3.259
	(15.39)	(4.530)	(4.381)	(5.166)
Growth Rate of Num CN Collab 2008-2012	0.00710	0.000178	0.000101	0.000569
	(0.162)	(0.0254)	(0.0198)	(0.0439)
Num Chinese Collaborator	0.228	0.0183	0.00675	0.0764
	(1.808)	(0.256)	(0.126)	(0.558)
Growth Rate of Num Collab 2008-2012	-0.0581	-0.113	-0.107	-0.142
	(0.683)	(0.622)	(0.614)	(0.661)
Num Collaborator	9.360	4.115	3.924	5.075
	(17.56)	(5.997)	(5.783)	(6.901)
Observations	749,589	273,316	228,139	45,177

Notes: Standard deviation in parentheses.

thor; whether the researcher listed any foreign address; whether the researcher listed any foreign funding; and whether the researcher has coauthors with foreign funding. We also match on the level and growth rates (2008–2012) of publications; impact-factor-weighted publications; number of collaborators; and number of Chinese collaborators. The algorithm matched 231,296 individuals, out of which

Table A8: U.S.-based Researcher Panel Descriptive Statistics: Outcome Variables

	(1) All	(2) Matched	(3) Matched-Untreated	(4) Matched-Treated
Num Pubs	2.108	0.929	0.898	1.087
	(3.721)	(1.358)	(1.337)	(1.451)
Num Pubs in US-based journals	1.197	0.564	0.546	0.655
	(2.394)	(0.940)	(0.928)	(0.990)
Num Impact Factor Weighted Pubs	6.563	2.579	2.444	3.259
	(15.39)	(4.530)	(4.381)	(5.166)
Impact Factor Weighted US-based Pubs	4.133	1.702	1.619	2.123
	(11.43)	(3.445)	(3.348)	(3.873)
Observations	749,589	273,316	228,139	45,177

Notes: Standard deviation in parentheses.

129,032 individuals are non-ethnically Chinese and 29,587 individuals are ethnically Chinese. The generated CEM weights are applied in the regression analysis. Balance statistics are reported in Tables A7 and A8.

A.4 Ethnicity & Field Imputation

This paper uses measures of ethnicity and scientific field distilled from self-reported names and departments, respectively. To impute ethnicity, we employ the Python package *ethnicseer* using individuals' full names.[40]

To infer scientific fields for ORCID researchers, we leverage academic department data self-reported on education spells. We use the FoR (Field of Research) framework from ANZSRC (Australian and New Zealand Standard Research Classification) to map unstructured department data onto 22 distinct research fields. When a researcher's listed department contains multiple relevant substrings, we rely on the last one to infer field.

B Empirical Context and Research Design Checks

This section provides additional context and validation for our research design.

Geopolitical Context and Policy Timing The China Initiative was formally launched in 2018, but investigations targeting Chinese scientists had already begun years

[40]We use *ethnicseer* because it can classify ethnicity with the granularity our analysis requires (e.g., Chinese ethnicity instead of Asian). Reassuringly, Torvik and Agarwal (2016) find that *ethnicseer* agrees with *Ethnea* 94% of the time for ethnically Chinese names.

Figure A1: China Initiative Cases Collected by *MIT Technology Review*'s China Initiative Database

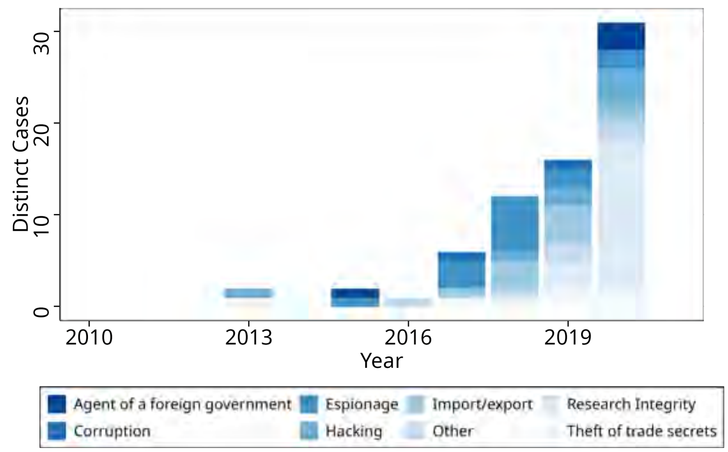

Notes: This figure graphs the number of distinct cases listed in the MIT Technology Review's China Initiative database for the years 2010 to 2021. The count of cases are displayed according to the category of legal issue in the case. We bin cases involving any type of espionage, research integrity issue, or hacking into respective categories. Distinct cases are defined as legal cases with distinct court docket numbers. This figure utilizes data from the following article: Eileen Guo, Jess Aloe, and Karen Hao, "The US crackdown on Chinese economic espionage is a mess. We have the data to show it," *MIT Technology Review*, December 2, 2021, https://www.technologyreview.com/2021/12/02/1040656/china-initative-us-justice-department/.

earlier. Figure A1 tallies these cases, as documented by the MIT Technology Review, and shows that enforcement actions linked to the China Initiative were underway well before the initiative's formal launch in 2018. This contributed to a growing sense of fear and uncertainty among Chinese and Chinese-American scientists, consistent with the early behavioral changes we document.

ORCID Sample Representativeness Our mobility analysis relies on ORCID CVs. There are good reasons to expect that researchers who create ORCID profiles differ systematically from those who do not. To assess this selection, we compare researchers in the Dimensions database with and without ORCID profiles. As shown in Table A10(a), ORCID users tend to be more research-active—they publish more, receive more grants, and have more affiliations. This selection is not a major concern for our study, as we are most interested in research-active individuals—those

Table A9: Inspecting Selection Bias in ORCID Data

	No ORCID 10,125,703 (84.6%)	Has ORCID 1,841,893 (15.4%)	Total 11,967,596 (100.0%)
Years active after 2008	2.735	5.570	3.171
Publications (lifetime)	4.989	15.308	6.577
Pre-2016 publications	1.801	3.072	1.997
Post-2016 publications	3.187	12.230	4.579
SJR-weighted publications (lifetime)	10.334	34.811	14.102
Pre-2016 SJR-weighted publications	3.675	6.822	4.160
Post-2016 SJR-weighted publications	6.659	27.989	9.942
# of grants (lifetime)	0.071	0.252	0.099
# of research organizations (lifetime)	1.074	2.132	1.237
STEM field	0.755	0.816	0.765
HASS field	0.101	0.161	0.110
Missing field	0.143	0.023	0.125
Missing publications	0.077	0.008	0.066
Ethnically CN	0.198	0.133	0.188

(a) Split by ORCID Status

	Non-ORCID Non-Eth. CN 8,121,699 (67.9%)	Non-ORCID Eth. CN 2,004,004 (16.7%)	ORCID Non-Eth. CN 1,597,553 (13.3%)	ORCID Eth. CN 244,340 (2.0%)
Years active after 2008	2.703	2.865	5.497	6.045
Publications (lifetime)	4.829	5.641	14.752	18.941
Pre-2016 publications	1.798	1.812	3.004	3.513
Post-2016 publications	3.028	3.828	11.742	15.424
SJR-weighted publications (lifetime)	9.784	12.566	32.683	48.726
Pre-2016 SJR-weighted publications	3.667	3.708	6.600	8.279
Post-2016 SJR-weighted publications	6.117	8.858	26.084	40.447
# of grants (lifetime)	0.073	0.060	0.246	0.296
# of research organizations (lifetime)	1.019	1.296	2.082	2.454
STEM field	0.721	0.894	0.797	0.938
HASS field	0.119	0.029	0.177	0.058
Missing field	0.160	0.077	0.026	0.004
Missing publications	0.078	0.073	0.009	0.004
Has ORCID	0.000	0.000	1.000	1.000
Ethnically CN	0.000	1.000	0.000	1.000

(b) Split by ORCID and Chinese Ethnicity

Notes: These tables provide summary statistics for researchers whose first publication in Dimensions was after 2008. Panel (a) divides the sample by whether or not the researcher in Dimensions lists an ORCID profile. Panel (b) additionally divides the sample by whether or not they are ethnically Chinese. In this table, an ethnically Chinese researcher is one whose last name in Dimensions is among China's top 150 most frequent surnames.

likely to contribute to the global scientific enterprise and to be responsive to geopolitical shifts. However, this also means that our findings may not generalize to the broader population of all STEM graduates or early-career researchers.

A potential threat to validity would arise if selection into ORCID varies by ethnicity in a way that interacts with our outcomes. Table A10(b) addresses this concern, showing that while ethnically Chinese researchers are modestly underrepresented, their research profiles among ORCID users are not systematically different from non-Chinese users in ways that would bias our estimates.

Figure A2: National Research Similarities Between the U.S. and U.K.

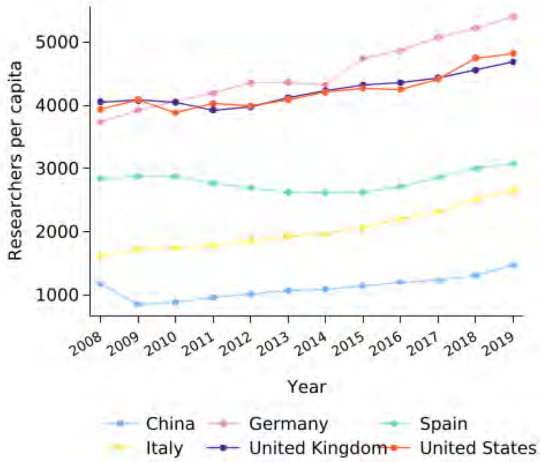

(a) Researchers in R&D (per million people)

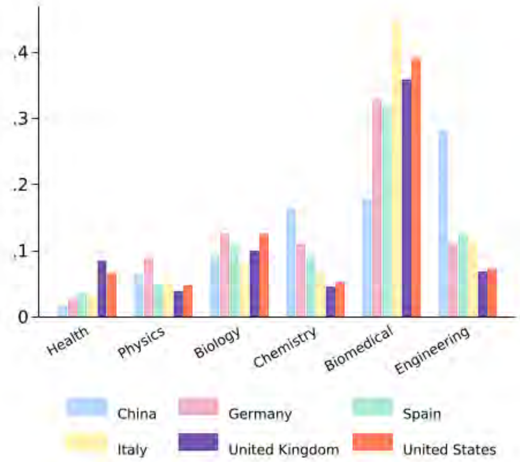

(b) National Research Attention to Various Fields

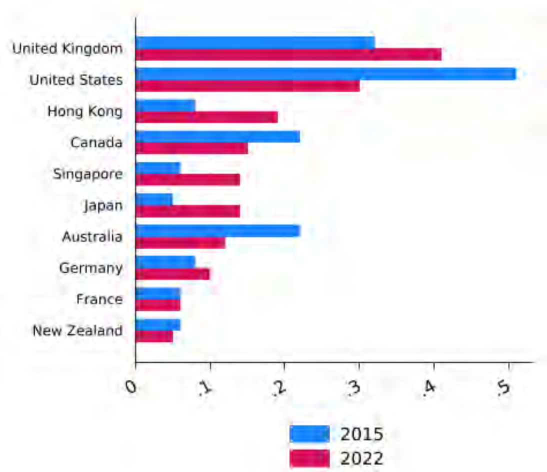

(c) Chinese Students' Preferred Overseas Study Destinations

Notes: Panel (a) presents the number of researchers (per million people) for various countries between 2008 and 2019. The data is sourced from the World Bank via the UNESCO Institute for Statistics (UIS). Panel (b) graphs the fraction of countries' impact-factor-weighted publications belonging to each of the selected STEM fields. Publications are weighted by the SJR score of their publishing journal. The data is constructed using Dimensions from Digital Science. Panel (c) graphs the results of a survey gathering destination country preferences for (nationally) Chinese students studying abroad. The data is sourced from Statista.

Use of the U.K. as a Comparison Country We use U.K.-authored publications and researchers as a comparison group to benchmark trends in U.S. citation flows and publication activity. We believe that the U.K. is a suitable control due to similar researchers per capita (Figure A2(a)), field composition (Figure A2(b)), and prominence as a destination for Chinese STEM students (Figure A2(c)). These parallels help isolate U.S.-specific effects from broader global shifts.

C Additional Mobility Results and Robustness Checks

To probe the robustness of our mobility findings, we conduct a range of additional tests, each described in the main text. These include restricting to students and graduates with no lag between prior degree and enrollment or graduation and job (Table A11, Columns (1)–(2)), limiting the enrollment sample to individuals from typical U.S. sender countries (Table A11, Column (3)), and testing whether increased domestic university quality in China explains declining U.S. enrollment (Figure A3). We also evaluate potential crowd-out by Indian applicants (Figure A4(a)) and rule out broader international trends by comparing Chinese and Indian retention patterns post-graduation (Figure A4(b)). Table A12 shows robustness to restricting our sample to those we are certain are doctoral students. Finally, we show formal diagnostic checks to assess the identifying assumptions underlying our difference-in-differences design (Figure A5).

Table A11: Lag Sensitivity Tests and Focusing On Sender Countries

	(1) Enrolls in U.S.	(2) Job in U.S.	(3) Enrolls in U.S.
Treatment = Ethnically CN	0.0331 (0.00497)	0.00592 (0.00748)	0.0512 (0.00773)
Treatment × Post-2016	-0.0379 (0.00810)	-0.0385 (0.00960)	-0.0396 (0.00879)
Unit of Analysis	Doctoral Student	U.S. Graduate	Doctoral Student
Field FE	Y	Y	Y
Cohort FE	Y	Y	Y
Prior Country FE	Y	-	Y
Model	OLS	OLS	OLS
Sample	No Lag to PhD	No Lag to Job	Africa, Asia, LatAm Only
Mean DV	0.259	0.863	0.117
Observations	91940	45537	63902

Notes: Standard errors clustered at the field-year level are in parentheses. The dependent variable is in the column heading. Column (1) restricts to students beginning their doctoral degree less than one year following the completion of their prior degree. Column (2) restricts to U.S. graduates beginning their first post-graduation job less than one year following the completion of their degree. Column (3) restricts to only students with prior degrees from typical U.S. sender countries. The analysis period is 2008–2019, and the post-treatment period is 2016–2019.

Figure A3: Event Study for Propensity to Enroll in a Chinese University

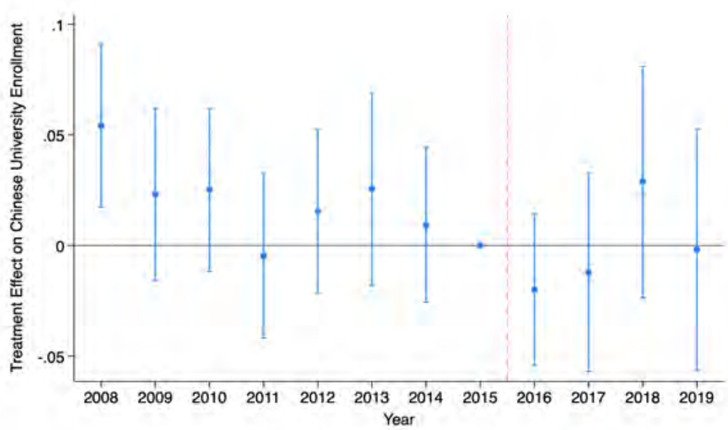

Notes: This plot reports event-study coefficients from a regression predicting enrollment in a Chinese university. The treated group is ethnically Chinese doctoral students, and the control group is non-ethnically Chinese doctoral students. The regression includes cohort, prior degree country, and field fixed effects. Standard errors are clustered at the field-year level.

Figure A4: Mean U.S. Enrollment and Retention by Nationality

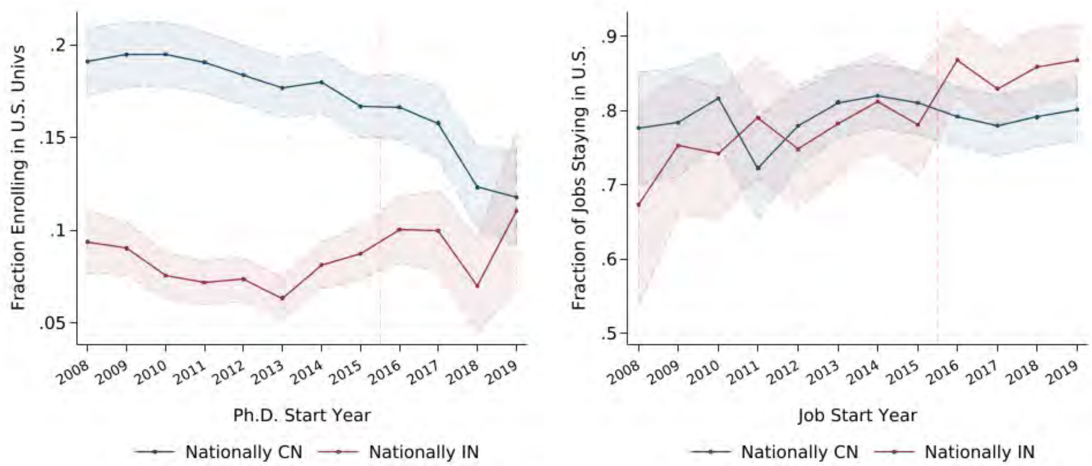

(a) Doctoral Students Enrolling in U.S. **(b)** Post-Graduation Jobs in U.S.

Notes: Panel (a) plots the raw fraction of new doctoral students enrolling in U.S. universities by nationality. Panel (b) plots the raw fraction of post-graduation jobs taken with U.S. employers by the graduate's nationality.

Table A12: Main Treatment Effects on Mobility among Ethnically Chinese U.S. Graduates (Ph.D. Only)

	(1) Job in U.S.	(2) Job in U.K.	(3) Job in Anglo.
Ethnically CN=1	0.0206	-0.00825	-0.0158
	(0.00689)	(0.00223)	(0.00351)
Ethnically CN=1 × Post-2016=1	-0.0304	0.00396	0.00820
	(0.00955)	(0.00264)	(0.00424)
Field FE	Y	Y	Y
Job Year FE	Y	Y	Y
Model	OLS	OLS	OLS
Mean DV	0.845	0.0121	0.0297
Observations	36015	36015	36015

Notes: Robust standard errors in parentheses without clustering. The dependent variable is in the column heading. The analysis sample is all U.S. graduates subsequently taking jobs after 2008. The analysis period is 2008–2019, where the post treatment period is 2016–2019.

Figure A5: Statistical Robustness Checks for Mobility Analyses

(a) Enrollment *pretrends* Test

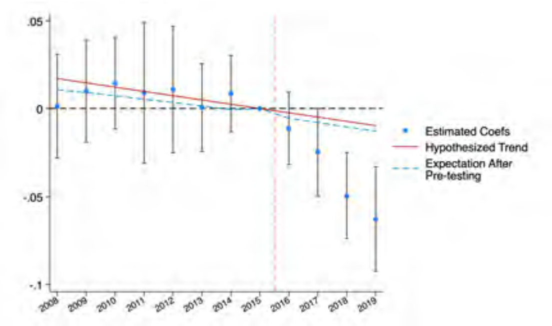

(b) Enrollment *HonestDiD* Test

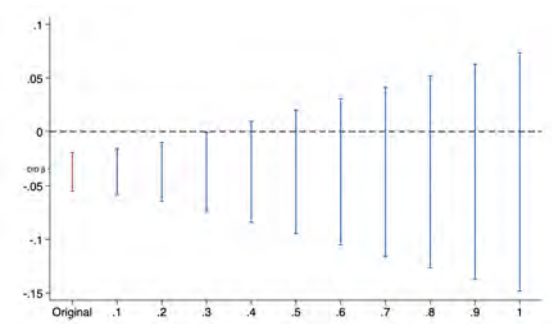

(c) Enrollment *ritest* Test

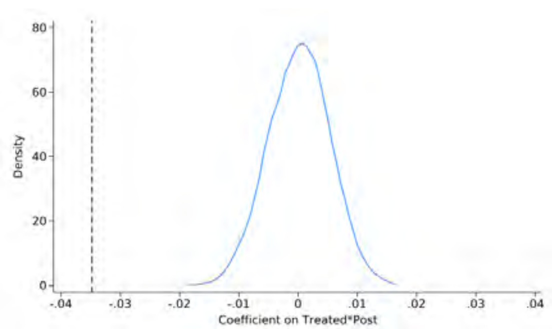

(d) Retention *pretrends* Test

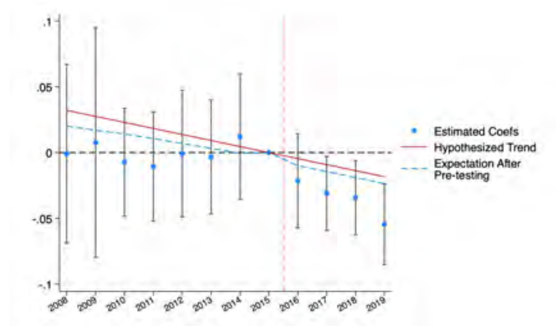

(e) Retention *HonestDiD* Test

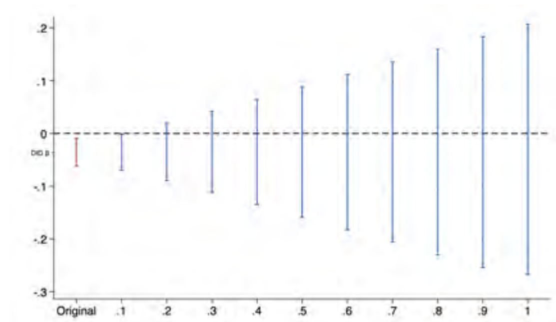

(f) Retention *ritest* Test

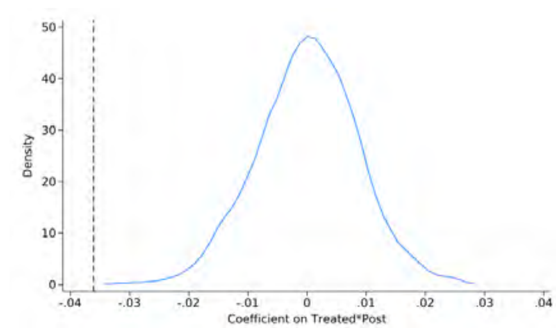

Notes: These figures report the outcomes of three statistical tests assessing the robustness of the mobility analyses' main results. We visualize the possibility of a parallel trends violation in Panels (a) and (d) (pretrends test), conduct a sensitivty analysis of the difference-in-differences estimate in Panels (b) and (e) (HonestDiD test), and assess the validity of our treatment construct in Panels (c) and (f) (permutation test).

D Additional Knowledge Flows Results and Robustness Checks

We conduct several additional robustness checks to confirm that our knowledge flow results are not sensitive to modeling choices or alternate definitions. As described in the main text, we show that our findings are robust to varying the threshold for defining "frontier" research (Table A13), adjusting fixed effects specifications (Table A14), and using alternate dependent variables for citation shares (Figure A6). Formal pretrend diagnostics, HonestDiD bounds, and placebo permutation tests further support our identification strategy (Figure A7).

Beyond robustness, Table A15 and Figure A8 (which presents the event-study plot corresponding to Column (1) of Table A15) provide evidence of a distinct pattern: U.S. publications with at least one ethnically Chinese author show a significant post-2016 decline in citations to China-produced research. This effect is absent in the broader U.S. sample and is consistent with the mechanism discussed in Section 4.5—specifically, a chilling effect manifesting as anticipatory disengagement or reputational risk avoidance among ethnically Chinese scientists.

Table A13: Main Treatment Effects on Knowledge Flows among Publications by China-based Teams (Other Frontier Thresholds)

	DV: Share Size			
	(1) Frontier (3%)	(2) Recent Frontier (3%)	(3) Frontier (5%)	(4) Recent Frontier (5%)
Treated = citing U.S.=1	0.247	0.211	0.231	0.194
	(0.0188)	(0.0223)	(0.0182)	(0.0211)
Treated = citing U.S.=1 × Post-2016=1	-0.0130	-0.0273	-0.0124	-0.0255
	(0.00463)	(0.00456)	(0.00448)	(0.00446)
Citing Paper FE	Y	Y	Y	Y
Model	OLS	OLS	OLS	OLS
Mean DV	0.172	0.140	0.161	0.130
Observations	3718356	2939138	3854408	3224622

Notes: Robust standard errors in parentheses with standard errors clustered at the field level. The dependent variable is in the column heading. The analysis sample is citations shares to U.S. or U.K. papers among Chinese publications. The analysis period is 2011–2019, where the post treatment period is 2016–2019. "Treated" refers to citation shares to U.S. papers, with those to U.K. papers serving as the control group.

Table A14: Main Treatment Effects on Knowledge Flows among Publications by China-based Teams (Fixed Effect Configurations)

	Share Size		
	(1) Raw	(2) Raw	(3) Raw
Treatment = Citing U.S.	0.182 (0.0137)	0.182 (0.0137)	0.182 (0.0137)
Post-2016	-0.00853 (0.00129)		
Treatment = Citing U.S. × Post-2016	-0.0138 (0.00438)	-0.0138 (0.00438)	-0.0138 (0.00438)
Unit of Analysis	Pub-Cite Share	Pub-Cite Share	Pub-Cite Share
Fixed Effects	Field-Year Trends	Field-Year FE	Field & Year FE
Model	OLS	OLS	OLS
Mean DV	0.126	0.126	0.126
Observations	4237614	4237614	4237614

Notes: Robust standard errors in parentheses with standard errors clustered at the field level. The dependent variable is in the column heading. The analysis sample is citation shares to U.S. or U.K. papers among Chinese publications. The analysis period is 2011–2019, where the post treatment period is 2016–2019. "Treated" refers to citation shares to U.S. papers, with those to U.K. papers serving as the control group.

Table A15: Main Treatment Effects on Knowledge Flows among U.S. Research Teams with Ethnically CN Researchers

	DV: Share Size			
	(1) Raw	(2) Recent	(3) Frontier	(4) Recent Frontier
Treated = U.S. publication=1	-0.00513 (0.00158)	-0.00752 (0.00232)	-0.00358 (0.00123)	-0.00467 (0.00216)
Treated = U.S. publication=1 × Post-2016=1	-0.00650 (0.00236)	-0.00872 (0.00359)	-0.00441 (0.00195)	-0.00861 (0.00355)
Field & Year FE	Y	Y	Y	Y
Model	OLS	OLS	OLS	OLS
Sample	Has Eth. CN Author	Has Eth. CN Author	Has Eth. CN Author	Has Eth. CN Author
Mean DV	0.0263	0.0408	0.0174	0.0266
Observations	652433	644473	586249	471328

Notes: Robust standard errors in parentheses with standard errors clustered at the field level. The dependent variable is in the column heading. The sample is all publications with U.S. or U.K. research teams involving at least one ethnically Chinese coauthor during the analysis period. The analysis period is 2011–2019, and the post treatment period is 2016–2019. "Treated" refers to publications in the U.S., with U.K. publications serving as the control group.

Figure A6: Event-Study Coefficients for Knowledge Flows (Recent, Frontier, and Recent Frontier)

(a) CN Pubs - Recent Cites

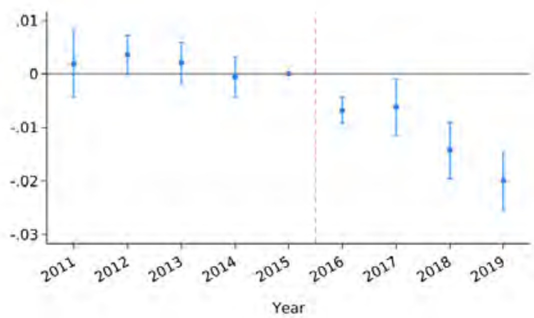

(d) U.S. Pubs - Recent Cites

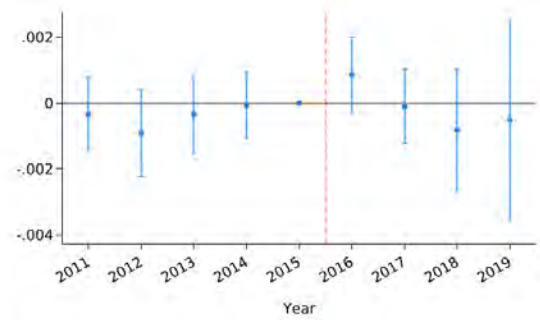

(b) CN Pubs - Frontier Cites

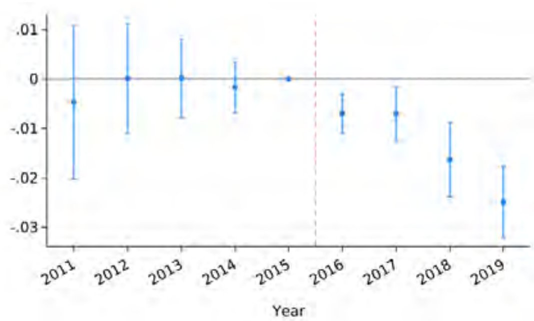

(e) U.S. Pubs - Frontier Cites

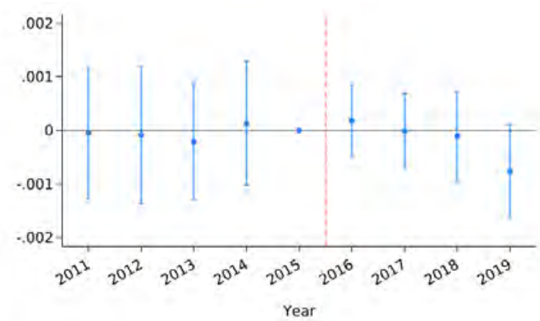

(c) CN Pubs - Recent Frontier Cites

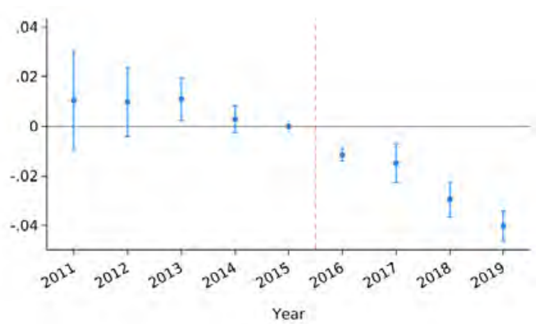

(f) U.S. Pubs - Recent Frontier Cites

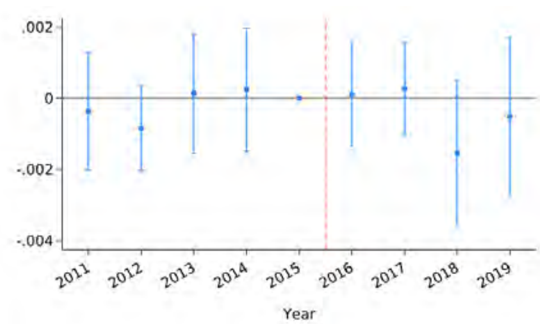

Notes: These plots report event-study coefficients from regressions predicting adjusted sizes of citation shares on Chinese (U.S.) publications to U.S. (Chinese) papers. In Panels (a)-(c), the treated group is citation shares to U.S. papers, and the control group is citation shares to U.K. papers. In Panels (d)-(f), the treated group is U.S. publications, and the control group is U.K. publications.

Figure A7: Statistical Robustness Checks for Chinese Reliance on U.S. Science Analysis

(a) *pretrends* Test

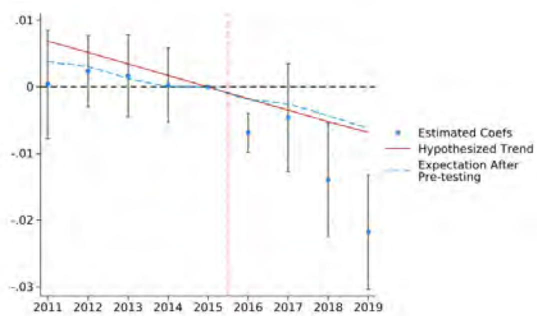

(b) *HonestDiD* Test

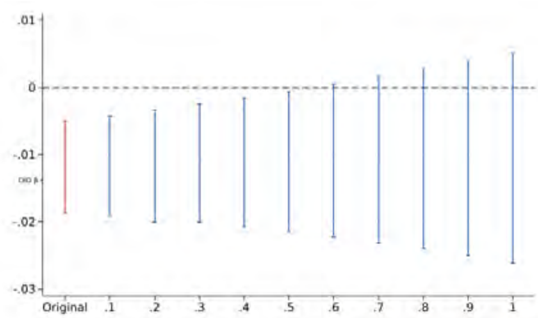

(c) *ritest* Test

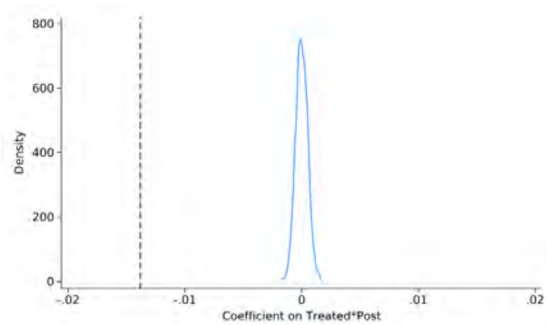

Notes: These figures report the outcomes of three statistical tests assessing the robustness of the China-to-U.S. knowledge flows section's main results. We visualize the possibility of a parallel trends violation in Panel (a) (pretrends test), conduct a sensitivty analysis of the difference-in-differences estimate in Panel (b) (HonestDiD test), and assess the validity of our treatment construct in Panel (c) (permutation test).

Figure A8: Main Treatment Effects on Knowledge Flows among U.S. Research Teams with Ethnically CN Researchers

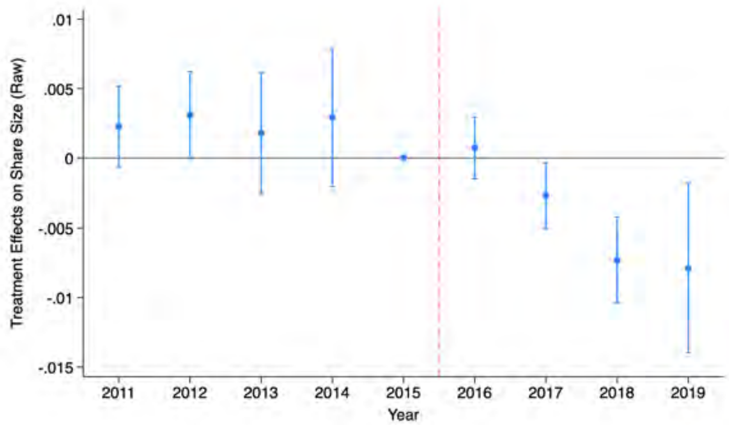

Notes: This plot reports event-study coefficients from a regression predicting the unadjusted share of references to Chinese papers on U.S. and U.K. publications. The treated group is U.S. publications with at least one ethnically Chinese coauthor, and the control group is U.K. publications with at least one ethnically Chinese coauthor. The regressions include fixed effects for publication year and research field. Standard errors are clustered at the field level.

E Additional Productivity Results and Robustness Checks

We conduct a range of additional checks to validate our main productivity results. For China-based researchers, we show that the null result is robust to alternative definitions of U.S. reliance—both in terms of threshold (Table A16 Column (1)–(2)) and whether we use recent or raw citation shares (Table A16 Column(3)-(6)). Kaplan-Meier plots (Figure A9) visually confirm the patterns in extensive margin estimates from the Cox models in Tables 6 and 7. Formal pretrend diagnostics, HonestDiD bounds, and placebo permutation tests further support our identification strategy (Figure A10).

In additional analyses performed, available upon request, we provide evidence that supports the finding that productivity declines are not limited to ethnically Chinese researchers with clear links to China. We focus on U.S.-based researchers with Chinese surnames but non-Chinese first names — individuals likely to be second-generation or otherwise socially assimilated. Identifying this group presents challenges; machine learning-based ethnicity classifiers like *ethnicseer* (which we use in the rest of the paper), which perform well when predicting

Table A16: China-based Researcher Main Treatment Effects by Varying Reliance Definition

	U.S.-Reliant = Cite Share to US above 50pct	U.S.-Reliant = Cite Share to US above 90pct	U.S.-Reliant = Recent Cite Share to US above 75pct			
	(1)	(2)	(3)	(4)	(5)	(6)
	Pubs	Pubs	Pubs	U.S. Pubs	IF wted Pubs	IF wted U.S. Pubs
U.S.-Reliant =1 × Post-2016=1	-0.006	0.017	-0.021	0.004	-0.011	-0.009
	(0.007)	(0.057)	(0.022)	(0.032)	(0.024)	(0.036)
Indiv FE	Y	Y	Y	Y	Y	Y
Year FE	Y	Y	Y	Y	Y	Y
Model	Poisson	Poisson	Poisson	Poisson	Poisson	Poisson
CEM	Y	Y	Y	Y	Y	Y
Mean DV	3.675	1.531	2.029	0.781	4.105	1.788
Observations	411,845	14,336	71,218	54,943	71,218	54,943

Notes: The analysis sample is the China-based researcher panel. Treatment (U.S.-Reliant) is defined using alternative measures, as indicated in the column headers. Regressions are weighted by the CEM matching weights. Standard errors are in parentheses and clustered at the person level.

overall ethnicity from *full* names, are not designed to separately assess surname and given name. As a result, they struggle to accurately estimate ethnicity separately for first and last name.

To address this, we adopt a dictionary-based approach that explicitly evaluates the first and last name independently. This allows for more precise detection of second-generation immigrants who often retain their ethnic surname but adopt an English given name. However, the method's effectiveness depends on the breadth of the underlying name dictionaries. Reassuringly, our results—first presented in Table 8 using the most aggressive dictionary—remain robust to more conservative versions that include only the most common Chinese names. As we move from a more conservative to a more aggressive name-matching approach, we reduce false negatives—correctly identifying more individuals with Chinese names (especially those with rare, diaspora, or non-standard romanizations)—but this comes with a modest increase in false positives. In all cases, the control group remains fixed, ensuring that any changes in estimated effects are attributable solely to adjustments in treatment group classification. Across specifications, we continue to find a significant post-2016 decline in both publication counts and impact-factor-weighted output, consistent with a chilling effect that operates through perceived identity, not formal affiliation. Tables showing these results are available upon request.

Figure A9: Kaplan-Meier Survival Curve

China-based Researchers

(a) DV: Pubs

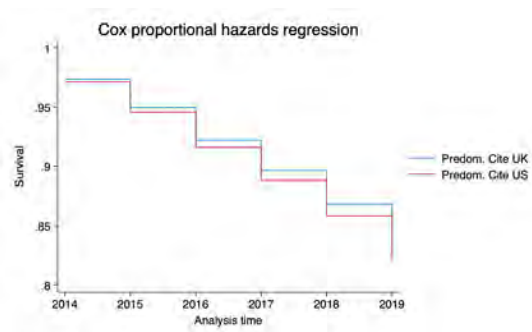

(b) DV: U.S. Pubs

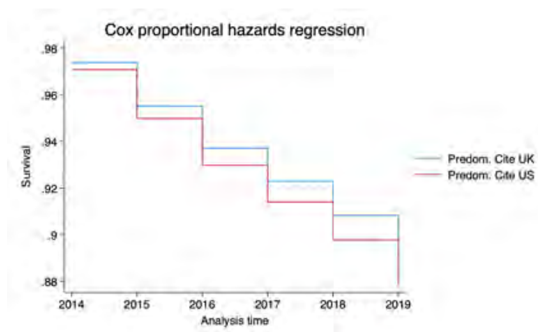

U.S.-based Researchers

(c) DV: Pubs

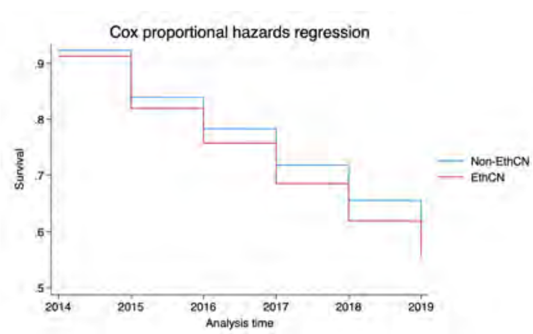

(d) DV: U.S. Pubs

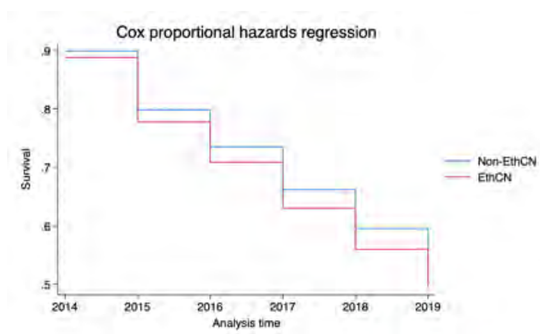

Notes: The Kaplan-Meier curves display the estimated survival probability over time. The failure event is to stop publishing anywhere in the world in Panel (a) and (c), and stop publishing in U.S.-based journals in Panel (b) and (d).

Figure A10: Statistical Robustness Checks for U.S.-based Researcher Productivity Analyses

(a) *pretrends* Test
DV: Num Pubs

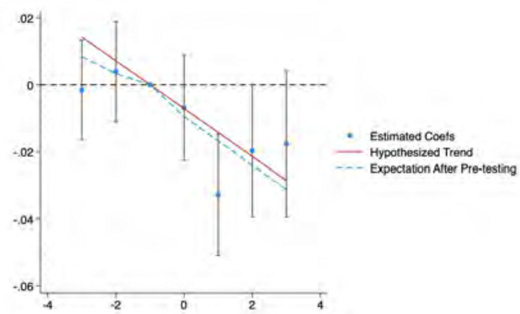

(d) *pretrends* Test
DV: IF wt Pubs

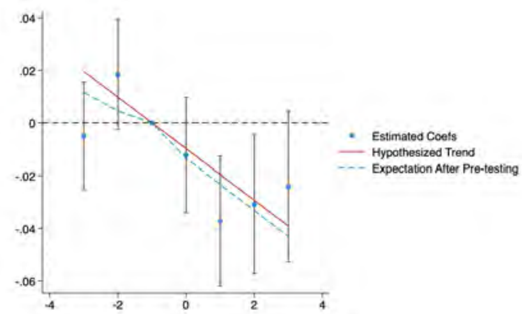

(b) *HonestDiD* Test
DV: Num Pubs

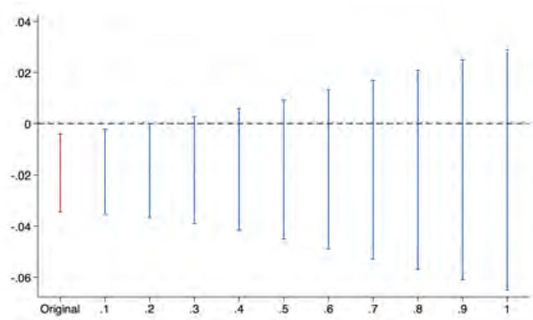

(e) *HonestDiD* Test
DV: IF wt Pubs

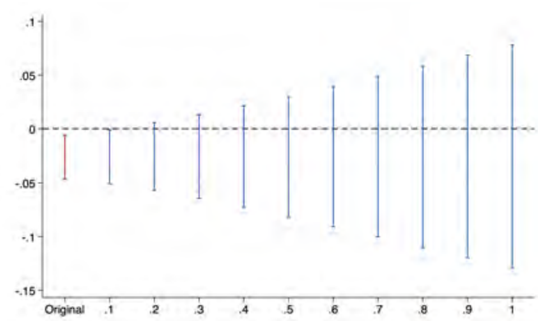

(c) *ritest* Test
DV: Num Pubs

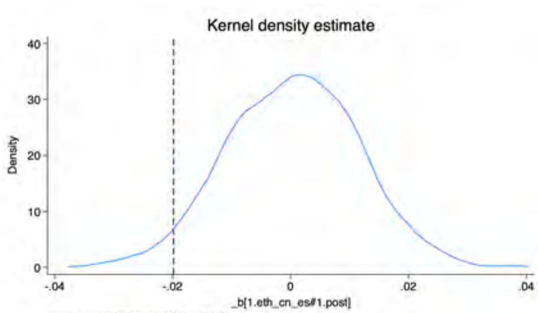

(f) *ritest* Test
DV: IF wt Pubs

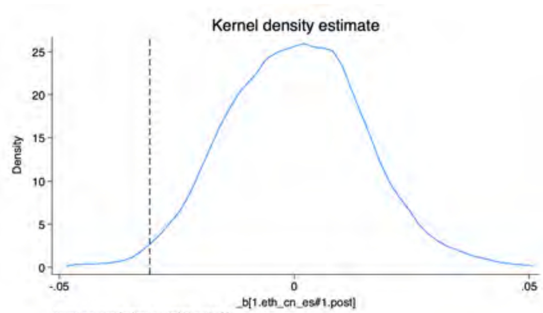

Notes: These figures report the outcomes of three statistical tests assessing the robustness of the productivity analyses' main results. The dependent variables are number of publications, and impact-weighted publications. We visualize the possibility of a parallel trends violation in Panels (a) and (d) (pretrends test), conduct a sensitivty analysis of the difference-in-differences estimate in Panels (b) and (e) (HonestDiD test), and assess the validity of our treatment construct in Panels (c) and (f) (permutation test).

Building a Wall Around Science

Supplement

Robert Flynn   Britta Glennon   Raviv Murciano-Goroff   Jiusi Xiao

Figure S1: Fraction of Nationally Chinese Doctoral Students Coming From Top Chinese Universities

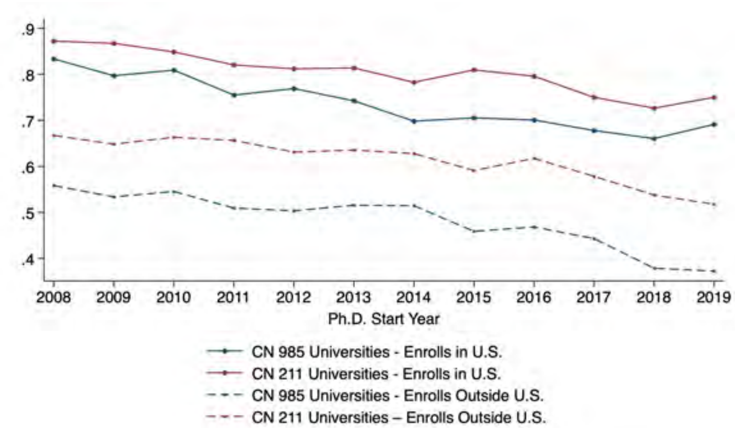

Notes: This figure plots the raw fraction of new doctoral students from China whose prior degree was at a top Chinese university, split by whether or not they enrolled in a U.S. doctoral program. A doctoral student comes from a top Chinese university if their prior degree university is associated with China's "985" or "211" education projects. We observe the fraction of Chinese doctoral students with these credentials to be falling over time.

Figure S2: Additional Event-Study Plots for Productivity Impacts

China-based Researchers

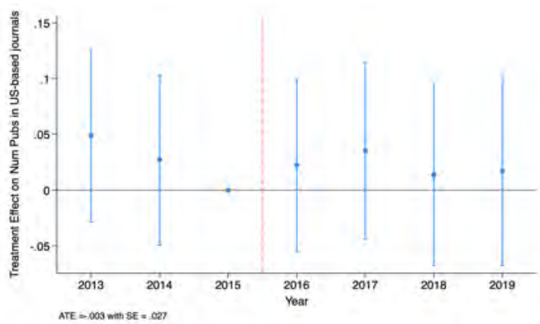

(a) DV: U.S. Pubs

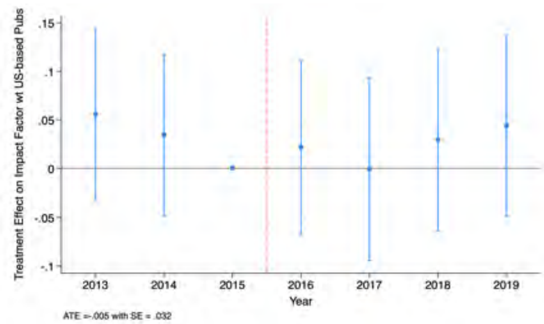

(b) DV: IF wt U.S. Pubs

U.S.-based Researchers

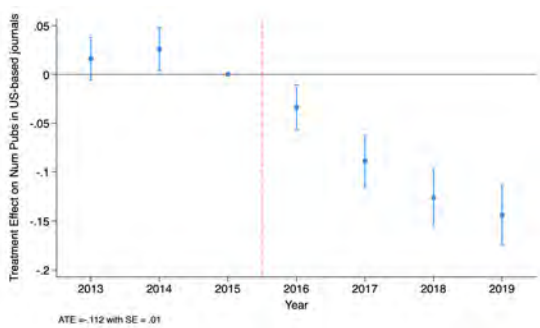

(c) DV: U.S. Pubs

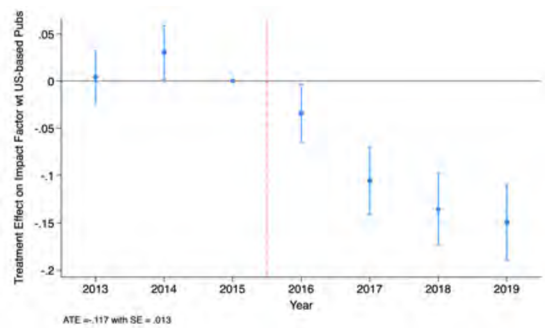

(d) DV: IF wt U.S. Pubs

Notes: This plot reports event-study coefficients from Poisson regressions. The dependent variable is indicated in the subfigure titles. Panels (a) and (b) use the China-based researcher panel; the treated group is U.S.-reliant, and the control group is U.K.-reliant researchers. Panels (c) and(d) use the U.S.-based researcher panel; the treated group is the U.S.-based ethnically Chinese researchers, and the control group is the matched non-ethnically Chinese researchers. Regressions include individual and year fixed effects. Standard errors are clustered at the individual level.

Table S1: Main Treatment Effects on Knowledge Flows among Publications by China-based Teams (Fixed Effect Configurations) - Recent

	Share Size		
	(1) Raw	(2) Raw	(3) Raw
Treatment = Citing U.S.	0.125 (0.0150)	0.125 (0.0150)	0.125 (0.0150)
Post-2016	-0.00682 (0.00141)		
Treatment = Citing U.S. × Post-2016	-0.0140 (0.00256)	-0.0140 (0.00256)	-0.0140 (0.00256)
Unit of Analysis	Pub-Cite Share	Pub-Cite Share	Pub-Cite Share
Fixed Effects	Field-Year Trends	Field-Year FE	Field & Year FE
Model	OLS	OLS	OLS
Mean DV	0.0853	0.0853	0.0853
Observations	4042500	4042500	4042500

Notes: Robust standard errors in parentheses with standard errors clustered at the field level. The dependent variable is in the column heading. The analysis sample is citation shares to U.S. or U.K. papers among Chinese publications. The analysis period is 2011–2019, where the post treatment period is 2016–2019. 'Treated' refers to citation shares to U.S. papers, with those to U.K. papers serving as the control group.

Table S2: Main Treatment Effects on Knowledge Flows among Publications by China-based Teams (Fixed Effect Configurations) - Frontier

	Share Size		
	(1) Raw	(2) Raw	(3) Raw
Treatment = Citing U.S.	0.281 (0.0186)	0.281 (0.0186)	0.281 (0.0186)
Post-2016	-0.00611 (0.000875)		
Treatment = Citing U.S. × Post-2016	-0.0142 (0.00514)	-0.0142 (0.00514)	-0.0142 (0.00514)
Unit of Analysis	Pub-Cite Share	Pub-Cite Share	Pub-Cite Share
Fixed Effects	Field-Year Trends	Field-Year FE	Field & Year FE
Model	OLS	OLS	OLS
Mean DV	0.193	0.193	0.193
Observations	3332908	3332908	3332908

Notes: Robust standard errors in parentheses with standard errors clustered at the field level. The dependent variable is in the column heading. The analysis sample is citation shares to U.S. or U.K. papers among Chinese publications. The analysis period is 2011–2019, where the post treatment period is 2016–2019. 'Treated' refers to citation shares to U.S. papers, with those to U.K. papers serving as the control group.

Table S3: Main Treatment Effects on Knowledge Flows among Publications by China-based Teams (Fixed Effect Configurations) - Recent Frontier

	Share Size		
	(1) Raw	(2) Raw	(3) Raw
Treatment = Citing U.S.	0.247 (0.0231)	0.247 (0.0231)	0.247 (0.0231)
Post-2016	-0.00718 (0.00141)		
Treatment = Citing U.S. × Post-2016	-0.0321 (0.00553)	-0.0321 (0.00553)	-0.0321 (0.00553)
Unit of Analysis	Pub-Cite Share	Pub-Cite Share	Pub-Cite Share
Fixed Effects	Field-Year Trends	Field-Year FE	Field & Year FE
Model	OLS	OLS	OLS
Mean DV	0.162	0.162	0.162
Observations	2303014	2303014	2303014

Notes: Robust standard errors in parentheses with standard errors clustered at the field level. The dependent variable is in the column heading. The analysis sample is citation shares to U.S. or U.K. papers among Chinese publications. The analysis period is 2011–2019, where the post treatment period is 2016–2019. 'Treated' refers to citation shares to U.S. papers, with those to U.K. papers serving as the control group.

Figure S3: Statistical Robustness Checks for U.S. Reliance on Chinese Science Analysis

(a) *pretrends* Test

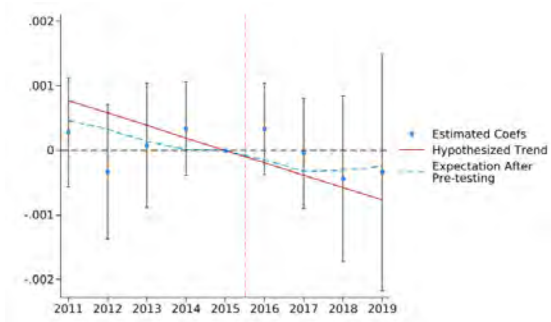

(b) *HonestDiD* Test

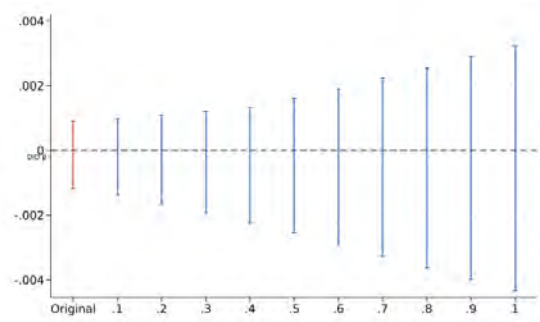

(c) *ritest* Test

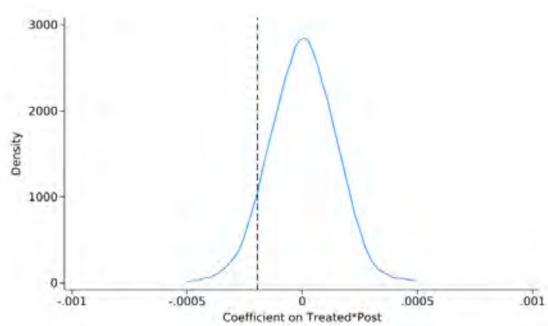

Notes: These figures report the outcomes of three statistical tests assessing the robustness of the knowledge flows section's main results. We visualize the possibility of a parallel trends violation in Panel (a) (pretrends test), conduct a sensitivty analysis of the difference-in-differences estimate in Panel (b) (HonestDiD test), and assess the validity of our treatment construct in Panel (c) (permutation test).

Figure S4: Statistical Robustness Checks for China-based Researcher Productivity Analyses

(a) *pretrends* Test
DV = Num Pubs

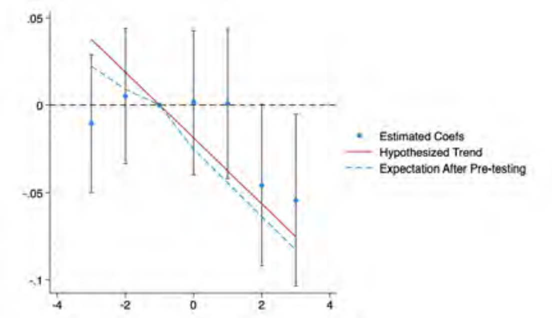

(d) *pretrends* Test
DV = IF wt Pubs

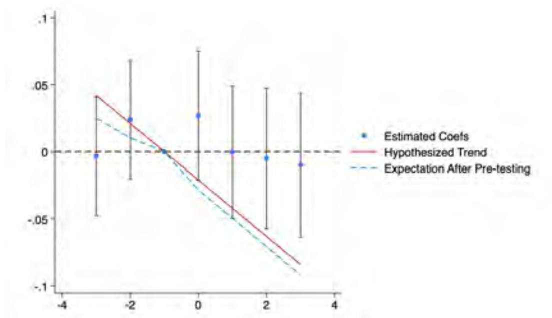

(b) *HonestDiD* Test
DV = Num Pubs

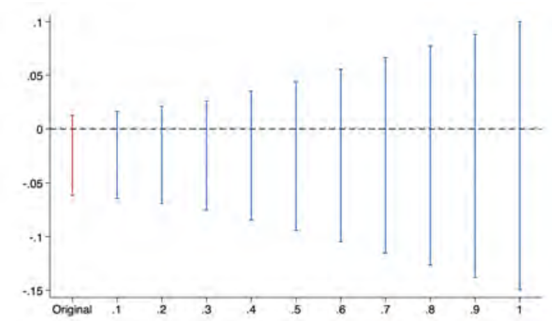

(e) *HonestDiD* Test
DV = IF wt Pubs

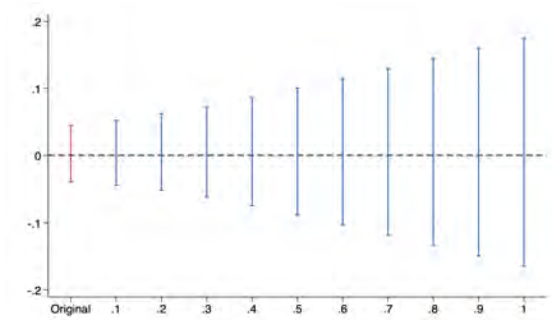

(c) *ritest* Test
DV = Num Pubs

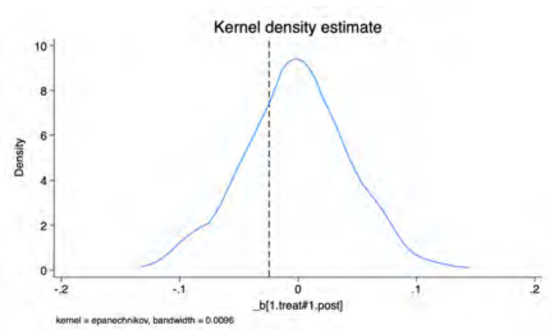

(f) *ritest* Test
DV = IF wt Pubs

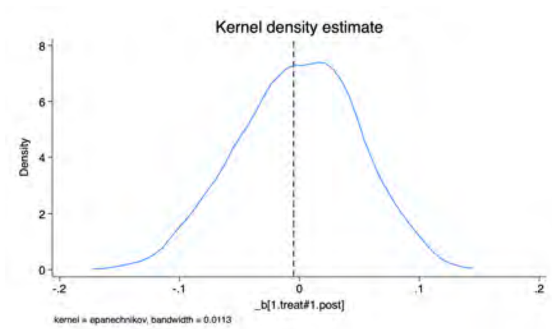

Notes: These figures report the outcomes of three statistical tests assessing the robustness of the China-based productivity analyses' main results. The dependent variable is number of publications and number of impact-weighted publications. We visualize the possibility of a parallel trends violation in Panels (a) and (d) (pretrends test), conduct a sensitivty analysis of the difference-in-differences estimate in Panels (b) and (e) (HonestDiD test), and assess the validity of our treatment construct in Panels (c) and (f) (permutation test).

Figure S5: Mechanism Event Study for U.S.-based Researchers Productivity

Treatment = Ethnically CN with China Tie

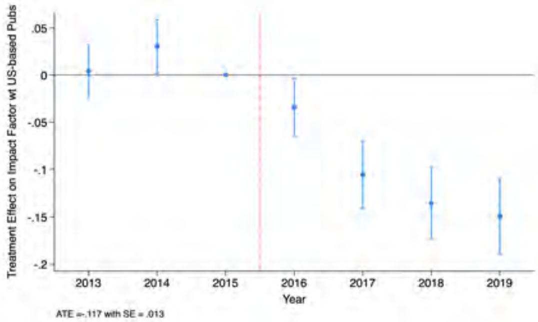

(a) DV: Pubs

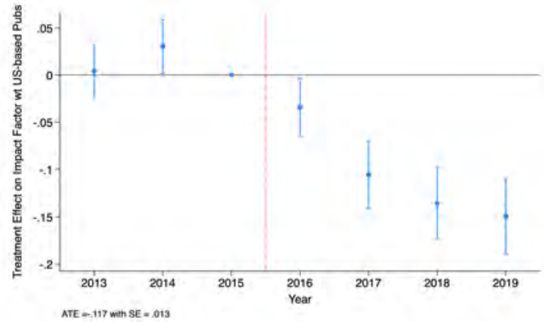

(b) DV: IF wt Pubs

Treatment = Ethnically CN without China Tie

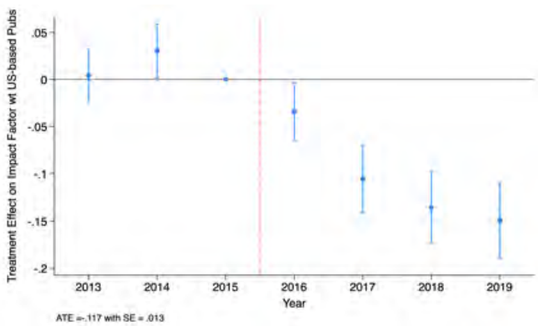

(c) DV: Pubs

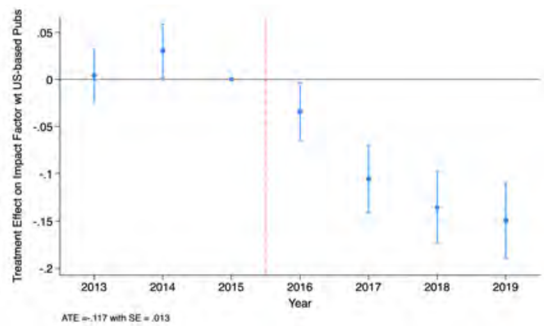

(d) DV: IF wt Pubs

Treatment = Ethnically CN with Non-Chinese First Name

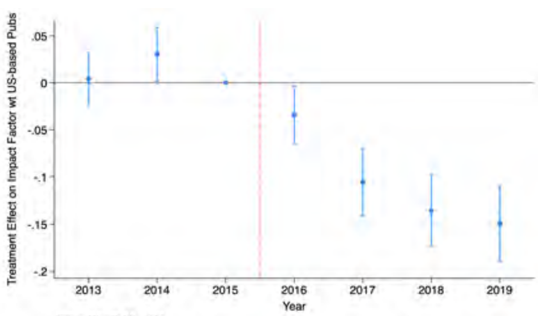

(e) DV: Pubs

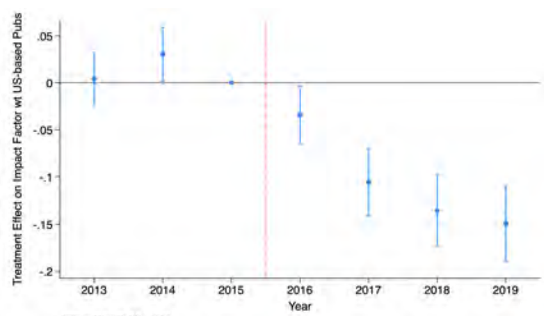

(f) DV: IF wt Pubs

Table S4: U.S.-based Researcher Mechanism Robustness

	Impact Factor wt Pubs		
	(1) Treatment = Ethnic CN With China Tie	(2) Treatment = Ethnic CN Without China Tie	(3) Treatment = Ethnic CN Non-CN First Name
Treatment=1 × Post-2016=1	-0.081	-0.123	-0.062
	(0.011)	(0.023)	(0.014)
Indiv FE	Y	Y	Y
Year FE	Y	Y	Y
Model	Poisson	Poisson	Poisson
CEM	Y	Y	Y
Mean DV	3.201	3.160	3.187
Observations	1,829,121	1,636,726	1,722,322

(a) DV: Impact wt Pubs

	Pubs		IF wt Pubs	
	(1) Treatmet = Ethnic CN Non-CN First Name (Conservative)	(2) Treatmet = Ethnic CN Non-CN First Name (Middle)	(3) Treatmet = Ethnic CN Non-CN First Name (Conservative)	(4) Treatmet = Ethnic CN Non-CN First Name (Middle)
Treatment=1 × Post-2016=1	-0.078	-0.069	-0.081	-0.076
	(0.009)	(0.010)	(0.012)	(0.012)
Indiv FE	Y	Y	Y	Y
Year FE	Y	Y	Y	Y
Model	Poisson	Poisson	Poisson	Poisson
CEM	Y	Y	Y	Y
Mean DV	1.078	1.081	3.181	3.184
Observations	1,826,468	1,776,677	1,826,468	1,776,677

(b) Non-Chinese First Name Results by Different Threshold Definitions

Notes: The analysis sample is the U.S.-based researcher panel. Regressions are weighted by the CEM matching weights. Standard errors are in parentheses and clustered at the person level.