Cover note: NBER I3 Technical Working Group, December 2026
"The Return(s) of Superstar Scientists", Michela Giorcelli, Yujing Huang, Bo Li

The dataset. We reconstruct the population of U.S.-trained Chinese scholars from the first half of the twentieth century and follow each person across two research systems and two languages through 1976. The dataset brings together: (i) four biographical compilations of returnees, merged and deduplicated into 4,684 U.S.-trained returnees with education, dates of arrival and return, post-return careers, political events and honours; (ii) Yuan's (1961) guide to Chinese doctoral dissertations in America, 1905–1960, and the NSF's National Register of Scientific and Technical Personnel (eight waves, 1954–1970), which together identify 9,478 scholars who stayed in the United States; (iii) publication records in both OpenAlex (international) and CNKI (Chinese-language), matched through each person's name in Chinese characters, Hanyu Pinyin and Wade–Giles, with matching rules chosen on a hand-labeled audit of 190 matches; (iv) OCR transcriptions of declassified State Department, INS, FBI and CIA files on the 1951–1955 restrictions on Chinese students leaving the United States, including an FBI roster of 78 barred students; and (v) digitized annual directories of Chinese students in the United States. The linked individual panel covers 3,165 returnees and 2,914 stayers. The complete 565,780-record CNKI corpus for 1915–1976, a 3.5-million-paper OpenAlex reference sample and 38.7 million pre-1979 cross-field citation edges support field-level analysis. Section 3 and Appendix A of the paper document each construction step, including source compilation rules, matching rules, measured match precision and known limitations.

Value to the community. Chinese-language science from before the 1980s is poorly covered by English-language bibliometric databases. Because each scientist is linked to both corpora, the data show when output moved to another literature rather than stopped. In our application, returnees publish 6–9 percent less in OpenAlex after return but more in CNKI, with total output unchanged; a study using OpenAlex alone would mistake this move for a decline in productivity. The data support new work on return migration and brain drain, Cold War science and security policy, science under central planning, and the long-run formation of China's research system. We also provide three reusable tools: (i) a cross-lingual measure of how close each paper sits to the research frontier, after Kelly et al. (2021), scored against both the Chinese-language and the international literature in one multilingual embedding space; (ii) a crosswalk that assigns Chinese-language papers to OpenAlex's 244 subfields; and (iii) a table of historical Chinese names in characters, Pinyin and Wade–Giles, with the record-linkage code. The frontier measure and the crosswalk method transfer to other settings where researchers publish across languages, such as Soviet, East German, Japanese or Korean science.

Presentation focus. The talk would center on the data and tools rather than the findings. We would cover: (1) how we rebuilt the returnee and stayer populations from fragmentary directories, biographical compilations and declassified archives, and how the inclusion rules of each source shape coverage; (2) how we linked historical Chinese names across three romanization systems and two bibliometric databases, and where linkage failed (for example, single-character OCR errors that split one person into two records); (3) how we built and validated the cross-lingual frontier measure; and (4) our release plan, on which we would welcome feedback, especially on licensing and documentation. We would use the paper's results only to illustrate what the data make possible.

Dissemination. Alongside publication of the paper, we will deposit the person-level roster, the archival transcriptions, the name table, the scholar-to-publication links, the derived measures and the code in a public repository (openICPSR or Harvard Dataverse) with a DOI and a codebook. For CNKI papers we will release only identifiers and metadata, because CNKI's terms of use restrict redistribution.

The Return(s) of Superstar Scientists

Michela Giorcelli, Yujing Huang, Bo Li*

July 6, 2026

Abstract

This paper studies the long-run effects of high-skill return migration on scientific productivity and the direction of research, using U.S.-trained Chinese scientists during the 1950s Great Return Wave. Combining newly digitized historical archives with multilingual bibliometric data, we follow 4,684 returnees and 9,478 scholars who stayed in the United States across the international and Chinese-domestic literatures over four decades, and compare returnees with matched stayers from the same training pipeline. Return redirected output rather than reducing it: returnees published about 9 percent less internationally and 10 percent more domestically, with total output unchanged and no loss of research quality, and they shifted toward applied, mission-oriented work. In the fields they entered, output by other domestic scientists rose, including work close to the global frontier. Pricing the reallocation through a cross-field citation network, China gained only a quarter to a third of what the United States lost.

JEL Codes: O15, J61, O33, N35, I23

Keywords: return migration, human capital, economics of science, knowledge diffusion, scientific productivity, China

*Contact information: Giorcelli: University of California, Los Angeles, NBER, CEPR, CESifo, IZA, and J-PAL, 9262 Bunche Hall, 315 Portola Plaza, Los Angeles, CA 90095, USA; email: mgiorcelli@econ.ucla.edu. Huang: UCLA Anderson School of Management, 110 Westwood Plaza, Los Angeles CA 90095, USA; email: yujingh@g.ucla.edu. Li: Guanghua School of Management, Peking University; email: boli@gsm.pku.edu.cn. Fonda Hu, Leon Wang, and Zhenyu Luo provided excellent research assistance. We thank seminar participants at UCLA Anderson, the 11th Monash–Warwick–Zurich–CEPR Text-as-Data Workshop, and the 2026 CES Annual Conference in China for helpful comments and discussion. We gratefully acknowledge financial support from the UCLA Luskin Center for History and Policy.

1 Introduction

The allocation of scientific talent across countries shapes where frontier knowledge is produced and how it diffuses, and is a central input into long-run growth (Romer, 1990, 1994). A large body of evidence shows that this allocation matters: the emigration or dismissal of elite scientists imposes persistent losses on the places they leave, while inflows of foreign talent raise innovation in the places they join (Waldinger, 2010, 2012, 2016; Borjas and Doran, 2012; Moser et al., 2014). This evidence, however, is almost entirely about *outward* mobility from advanced economies. Far less is known about the consequences of *return* migration, when scientists trained at the frontier move back to their home countries, even though such flows are large and, for developing economies, potentially transformative. Returning scientists may reshape not only their own productivity but the direction of domestic research, the diffusion of frontier ideas, and the formation of scientific capacity.

Evidence on the long-run effects of return migration remains scarce, for two reasons. First, who returns is endogenous: scientists sort on ability and field, and political and institutional forces shape both the decision to leave and the environment returnees face at home. Second, measuring the consequences requires tracking individual scientists across countries, languages, and publication systems over decades, and distinguishing a scientist who stops producing science from one whose work has merely moved into a different scientific community.

This paper studies the long-run effects of high-skill return migration on scientific productivity and the direction of research, using the return of U.S.-trained Chinese scientists during the Great Return Wave of the 1950s. After the founding of the People's Republic of China in 1949, the new government mobilized overseas-educated Chinese to return and help build a modern nation, while the United States, in the grip of the early Cold War, moved to restrict the departure of scientists with sensitive training. The most prominent returnee was Qian Xuesen (Tsien Hsue-shen), a Caltech-trained aerodynamicist and co-founder of the Jet Propulsion Laboratory, whose U.S. security clearance was revoked in 1950 and who, after five years of detention, was released to China in 1955 as part of the Geneva negotiations that followed the Korean War; he went on to found China's ballistic-missile and space programs (Chang, 1995). Qian was not an isolated case. The wave also carried Deng Jiaxian, who led the design of China's first atomic bomb, the Caltech physicists Guo Yonghuai and Zhao Zhongyao, and the mathematician Hua Luogeng, among a cohort that would staff the founding faculties of China's universities and lead its strategic-weapons program (Wang, 2014). More broadly, this episode is one instance of a recurring twentieth-century phenomenon: geopolitical ruptures repeatedly reallocated frontier-trained scientists across

national research systems, reshaping the international organization of science. The early Cold War split a generation of such scientists between two emerging superpower systems, and studying how one of them absorbed this talent speaks to how the global scientific order is reorganized when talent is forced across political divides. Because frontier scientific capability is itself a strategic national asset, its reallocation is also a dimension of international power, alongside the economic interdependence studied by Liu and Yang (2026).

We assemble a comprehensive new dataset, digitized from multiple historical archives, that reconstructs the population of U.S.-trained Chinese scholars of this period and tracks their training, return decisions, institutional placements, and scientific output over four decades. We combine annual directories of Chinese students in the United States, doctoral dissertation records, declassified U.S. government files on exit-permit decisions, biographical dictionaries of returnees, and two multilingual bibliometric databases: OpenAlex, which indexes the international, largely English-language literature, and CNKI, which indexes the Chinese-domestic literature. Following the same scientists into *both* literatures is central to our design: it lets us distinguish the *redirection* of scientific output across research systems from its *disappearance*, a distinction that studies of a single corpus cannot draw. To characterize not only the volume but the content of research, we represent each paper in a shared, cross-lingual semantic space using multilingual language models, which lets us measure how close a scientist's work sits to the research frontier and how its direction shifts over time. We will release the linked dataset, the historical-name alias table, the crosswalk from CNKI subjects to OpenAlex subfields, and the code that computes the frontier-alignment measures as a public research resource; Appendix 8 documents each construction step and the release plan.

Return during the Great Wave was shaped by sharp, historically contingent political constraints. Under early-Cold-War security policy, most notably the 1951 order preventing the departure of aliens with scientific and technical training, the United States restricted the exit of Chinese scholars in fields deemed to have military or industrial value, and after 1957 domestic political turmoil in China stopped further returns. These constraints are context for our design rather than an instrument: we do not observe individual exit applications, so we do not instrument the return decision itself. Instead, we compare each returnee to an observably similar scientist who remained in the United States, matched within field on first-publication year and birth cohort from the same training pipeline, and trace the gap in their trajectories around the year of return with scholar, year, and field-by-year fixed effects. Matched groups are balanced on pre-return characteristics, including pre-return output that the matching does not target, and there is no differential pre-trend on any of six publication outcomes. At the field level, we compare subfields that received more or fewer returnees in

a stacked event study over the full Chinese-language corpus.

Our main findings can be summarized in four parts. First, return migration *redirected* scientific output rather than reducing it: relative to matched stayers, returnees publish about 9 percent less in the international, English-language literature after return and about 10 percent more in the domestic, Chinese-language literature, leaving their total publication count unchanged. Second, this reallocation came at no measurable cost to quality. Pooling each scholar's English- and Chinese-language output and scoring both against the same global reference, neither the textual importance of returnees' work nor its alignment with the research frontier moves after return. What changed was the *direction* of research rather than its caliber: a large-language-model reading of scholars' abstracts shows returnees shifting toward applied, mission-oriented work. Third, returnees reshaped the science produced around them. A subfield receiving one standard deviation more returnees (about 1.4 people) saw roughly 70 percent more output by *other*, non-returnee Chinese scientists over the following decade, and about 29 percent more of their papers close to the global frontier. Returnees became the faculty and scientific elite that plausibly carried this diffusion: 58 percent took university posts, and 17 of the 23 "Two Bombs, One Satellite" honourees are in our data. Fourth, pricing each scientist through a cross-field citation network, we find that the reallocation was net-negative for world science: China's gain offset only a quarter to a third of the social value the United States gave up, so it took roughly three scientists arriving in China to replace the value of one leaving the United States. Taken together, these results point to a principle that headcount obscures: the consequences of relocating scientific talent turn not on how many scientists move but on which strategic fields they seed and their position in the knowledge network. A small cadre concentrated in core strategic fields such as nuclear weapons and aerospace conferred capabilities out of all proportion to its size, and our welfare accounting shows that a scientist's value is governed by network centrality rather than raw output.

This paper contributes to four literatures. First, it adds to the study of how migration shapes knowledge production and spillovers, which has focused on brain drain, diaspora networks, and displacement shocks, largely in frontier economies and for outward mobility (Borjas and Doran, 2012; Docquier and Rapoport, 2012; Moser et al., 2014; Kerr et al., 2016; Waldinger, 2012, 2016; Iaria et al., 2018). A parallel theoretical literature models the global race for talent and the spillbacks that return migrants generate onto home collaborators (Prato, 2025). We provide the first large-scale evidence on *return* migration into an economy with weak preexisting scientific institutions and, by following scientists into a second literature, separate the redirection of output from its disappearance.

Second, the paper advances the literature on upper-tail human capital and the production

of useful knowledge. Growth models emphasize aggregate human capital and idea accumulation (Romer, 1990, 1994; Mokyr, 2005a,b), and recent work highlights the outsized role of scientific elites (Squicciarini and Voigtländer, 2015; Giorcelli, 2019; Bianchi and Giorcelli, 2020). We show that a single returning scientist can seed a field: the frontier-quality output of *other* domestic scientists rose in the fields returnees entered, a multiplier that individual-level productivity measures cannot capture. Our setting is a human-capital counterpart to the Soviet technology transfers that helped build China's industrial base in the same decade: Giorcelli and Li (2026) study the Soviet-aided "156 Projects" that brought industrial technology to 1950s China, while we trace the parallel inflow of frontier-trained scientists.

Third, the paper links return migration to the endogenous direction of research. Political incentives and institutional constraints are known to shape scientific agendas (Jia et al., 2019; Acemoglu et al., 2024), and foreign education can transmit organizational norms with long-run consequences (Spilimbergo, 2009; Kung and Wang, 2025). We document that returnees sharply redirected the *content* of their research toward applied, mission-oriented work relative to stayers, a shift a language-model reading of their bilingual publications detects but citation counts and broad field labels do not, clarifying how human-capital flows let states steer science toward strategic objectives while building domestic capacity. More broadly, the setting offers a rare quantitative window into science under central planning, in which a state that could direct where its scientists worked concentrated returnee talent in strategic-priority fields.

Fourth, the paper speaks to the economics of the allocation of talent. A long literature shows that how talent sorts across activities and fields is a first-order determinant of aggregate output and growth (Murphy et al., 1991; Hsieh et al., 2019; Bell et al., 2019). We add a network dimension: the social value of a scientist depends not only on the activity they choose but on where they sit in the global knowledge network (Liu and Ma, 2026). A frontier scientist embedded in a dense, self-sufficient research system generates more downstream spillovers than the same scientist transplanted into a frontier-dependent one, so reallocating talent across borders is not a zero-sum transfer, and forced relocation can destroy social value rather than merely redistribute it.

Beyond its academic contributions, the setting speaks to a debate the United States is having again. The 2018–2022 China Initiative, the 2020 restrictions on entry for Chinese graduate students and researchers (Presidential Proclamation 10043), and recurring proposals to screen researchers by nationality rest on the same implicit trade-off that governed the 1950s exit restrictions: whether the security gains from excluding suspect scientists outweigh the scientific costs (Jia et al., 2024; The White House, 2020). Our estimates put a magnitude on those costs, show that they fall on the excluding country as well as its rival, and indicate

that talent excluded from one research system is largely relocated rather than lost.

The remainder of the paper proceeds as follows. Section 2 describes the Great Return Wave; Section 3 the data and measurement; Section 4 the empirical strategy; Sections 5 and 6 the effects on individual scientists and on domestic scientific production; Section 7 the welfare accounting; and Section 8 concludes.

2 The Great Return Wave

2.1 Calling on Overseas-Educated Chinese for Nation-Building

Following the establishment of the People's Republic of China (PRC) in 1949, the new government confronted an economy that was overwhelmingly agrarian, technologically backward, and institutionally underdeveloped (Lardy, 1995). With limited domestic capacity in modern science and engineering, the state identified overseas-educated Chinese as a critical source of advanced human capital for industrialization and nation-building.

Beginning in late 1949, the PRC launched a coordinated campaign encouraging foreign-trained Chinese to return, commonly framed as a call to "serve the motherland and the people." The government established the Committee on Returned Students Affairs in December 1949 to facilitate return, providing travel subsidies, reception services, and centralized job placement in universities, research institutes, and state-owned enterprises (Zhang and Mo, 2018). These policies reflected a deliberate strategy to rapidly internalize frontier knowledge acquired abroad.

At the time, the stock of overseas Chinese students was sizable. In 1949, approximately 5,541 Chinese students were studying abroad, the product of earlier state-sponsored and private efforts to train elite students overseas (China Institute in America, 1954; Li, 2000). About 70 percent were enrolled in the United States, with the remainder primarily in the United Kingdom and Japan, and more than two-thirds pursued degrees in STEM fields.

2.2 Cold War Tensions and Barriers to Return

This initial momentum was disrupted by the rapid deterioration of relations between the PRC and the United States. After 1949, the United States refused to recognize the Communist government, severing diplomatic ties and leaving Chinese nationals without consular protection. The outbreak of the Korean War in 1950 further intensified security concerns and anti-communist sentiment.

Against this backdrop, U.S. authorities increasingly viewed Chinese students, especially those with advanced scientific and technical training, as potential national security risks. In

October 1951, the United States implemented administrative restrictions on the departure of aliens with scientific and technical training. In particular, U.S. immigration authorities formally prohibited Chinese students and scholars from departing the United States without authorization, with violations punishable by fines of up to $5,000, imprisonment of up to five years, or both, reflecting security concerns over the transfer of sensitive scientific and technical knowledge to Communist China (Li, 2000). Declassified intelligence reports emphasized the risk that returnees could transfer knowledge relevant to military, nuclear, and industrial programs (see Appendix for more details).

These policies generated substantial frictions to return. Archival records document interrogations by federal authorities, prolonged surveillance, repeated denials of exit permits, and pressure to adjust immigration status or naturalize (Han, 1993). While some students eventually returned after extended delays, others were prevented from leaving altogether despite documented intentions to do so. As a result, return outcomes during the early 1950s were shaped less by individual choice than by historically contingent political constraints.

2.3 The Geneva Conference

By the mid-1950s, the status of Chinese students in the United States became entangled with broader diplomatic negotiations. During the Geneva Talks of 1954–1955, Chinese students were discussed alongside American prisoners held in China, effectively becoming bargaining instruments in bilateral negotiations (Han, 1993). At this stage, many students were no longer viewed by U.S. authorities as strategically indispensable and had become administratively and politically costly.

The talks culminated in a joint announcement on September 10, 1955, affirming the right of nationals residing in the other country to return home if they wished. In principle, this agreement removed formal barriers to departure. In practice, return flows remained limited. U.S. records indicate that more than half of applicants withdrew their requests, reflecting persistent administrative obstacles, uncertainty, and the passage of time (Han, 1993). In parallel, the Refugee Relief Act of 1953 facilitated adjustment to permanent residence for many Chinese nationals, further reducing incentives to return.

By the late 1950s, return migration had largely ceased. Chinese departures from the United States fell sharply, reaching only two in fiscal year 1958 (Kung, 1973). Archival evidence summarized by Wang (2014) indicates that, despite encouragement or, in some cases, force to stay from the U.S. side and subsequent political contestation from China side, the majority of Chinese students and scientists trained abroad ultimately remained in the United States. Of an estimated cohort of roughly 5,000 individuals, about 4,000 stayed,

including Chen-Ning Yang and Tsung-Dao Lee, who later received the 1957 Nobel Prize in Physics, while approximately 1,200 successfully contested U.S. restrictions and returned to China during the 1950s. Also, internal reports from Beijing Municipal Archives document that 1,954 Chinese students and scientists returned to mainland China between 1949 and 1958, of whom 1,244 (roughly 64 percent) returned from the United States. A majority of these returnees specialized in the natural sciences, broadly defined to include engineering, with the remainder concentrated in the social sciences or other fields. Combined with domestic political turmoil in China, these developments effectively ended the Great Return Wave until diplomatic normalization and the resumption of overseas study after 1978. The cohort shaped by this episode provides the institutional setting for our empirical analysis. Figure 1 summarizes the wave: annual returns spike after 1945, collapse under the 1950–1955 exit ban, and rebound after the 1955 Geneva talks.

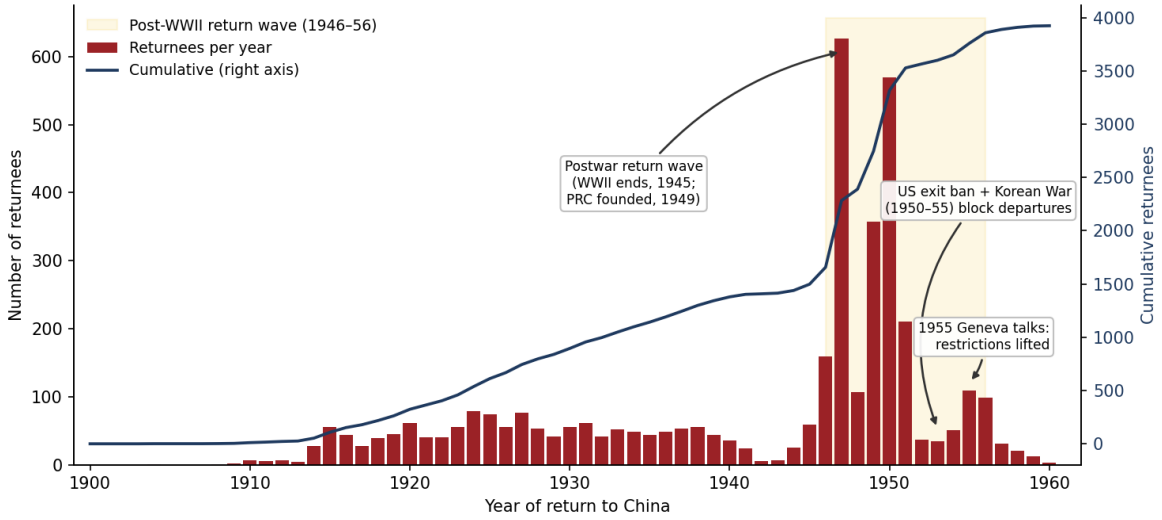


Figure 1: The Great Return Wave in numbers, 1900–1960

Notes: Annual number of overseas-educated Chinese scholars returning to China (bars, left axis) and the cumulative total (line, right axis), by year of return. The shaded band marks the post-war Great Return Wave (1946–1956). Returns surge after the end of World War II and the 1949 founding of the People's Republic, collapse during the 1950–1955 U.S. exit ban and the Korean War, and recover briefly after the 1955 Geneva talks lifted departure restrictions. Source: authors' biographical returnee database.

3 Data

In this section, we describe the data that we collected and digitized from several historical archives. Appendix 8 documents the full construction pipeline stage by stage, with record

counts, linkage rates, validation checks, known limitations, and the public-release plan.

3.1 Directory of Chinese Students in the United States

To construct a comprehensive dataset of Chinese students who pursued higher education in the United States during the first half of the 20th century, we digitize and extract individual-level information from a series of historical directories compiled by Chinese student organizations and educational institutions. These directories, published annually or biennially, were produced by associations such as the Chinese Students' Alliance, the Chinese Students Handbook Company, the Chinese Students' Christian Association, and later the China Institute in America. The directories span from 1911 to 1953, including the *Directory of the Chinese Students in the United States of America*, *The handbook of Chinese students in U.S.A.*, and the *Directory of Chinese Students in Colleges and Universities in the United States*. While the directories vary in format, we harmonize the data to create a unified structure that captures detailed information on students' names, universities attended, degrees obtained, and fields of study. The directories document the size and composition of the student population; they are not yet linked to the person-level returnee and stayer lists used in the analysis.

3.2 Biographical Database of Returnees

We manually collected and digitized individual-level data from four retrospective biographical compilations documenting Chinese scholars who studied abroad and returned to China between the 1840s and early 1960s. The sources were compiled for different audiences and under different inclusion rules, which determines how each one selects on prominence:

1. *Comprehensive Survey of Scientific and Technical Personnel Studying Abroad and Returning to China, 1840–1949* (*Linian chuguo/huiguo keji renyuan zonglan*), compiled by Ma Tsu-sheng and a team for government and research use. It profiles about 3,500 scientific and technical personnel (3,742 person-trip records) and aims to record *all* findable returnees, so completeness is its purpose.

2. *Guilai: A Directory of Students Who Returned from North America in the 1950s* (*Guilai: ershi shiji wushi niandai guiguo Beimei liuxuesheng minglu*), edited by Wang Delu (2023). It covers 1,806 returnees, of whom we parse 1,774. Inclusion follows criteria rather than merit, and the compilers reached obscure returnees through their families and work units. It records education, dates of arrival and return, post-return careers, political roles, and events such as detention and death during the Cultural Revolution.

3. Chu-yuan Cheng, *Scientific and Engineering Manpower in Communist China, 1949–1963* (U.S. National Science Foundation, NSF 65-14), prepared for U.S. policymakers. It profiles about 1,200 prominent scientists and engineers. Inclusion is based on prominence and visibility to U.S. observers, so it is the source most exposed to selection.

4. *Biographical Sketches of Chinese Science and Technology Experts* (*Zhongguo kexue jishu zhuanjia zhuanlue*), 56 volumes, compiled by the China Association for Science and Technology under the editorship of Qian Sanqiang for the scientific community. Inclusion is based on eminence, but with one volume per discipline, so the coverage is stratified by field.

The first two sources, which aim at completeness, form the backbone; the latter two add prominent scientists and variables such as Academy membership. After harmonizing names, deduplicating across sources, and keeping those trained in the United States who returned by about 1965, the merged database contains 4,684 returnees. Of these, 2,614 appear in the *Comprehensive Survey*, 1,537 in *Guilai*, 1,002 in Cheng, and 675 in the *Biographical Sketches* (a person can appear in several). Of the 4,030 with a recorded return year, 1,479 returned in 1949–1956.

We use this dataset to identify individuals who pursued graduate or postgraduate studies abroad and returned to China in the late 1940s and 1950s, either voluntarily or under geopolitical pressure. In addition to biographical details, the data include information on disciplinary specialization, institutional roles, and leadership positions within Chinese universities and research institutes. This historical sample allows us to characterize the early cohort of returnee scientists and compare them to individuals who remained abroad.

3.3 Comparison Group of Scientists Who Remained in the United States

To construct a comparison group of observationally similar scientists who did not return, we compile a person-level list of U.S.-trained Chinese scholars drawn from the same training pipeline as the returnees. We combine two sources. The first is Yuan (1961), *A Guide to Doctoral Dissertations by Chinese Students in America, 1905–1960*, compiled by Yuan Tung-li for the Sino-American Cultural Society. It is effectively a census of every U.S. doctorate earned by a Chinese student over the period. The second is the National Register of Scientific and Technical Personnel, the federal registry that the National Science Foundation was mandated to maintain for all U.S. scientists (National Science Foundation Act, §3(a)(8)). It was fielded in eight waves between 1954 and 1970, with 126,000 to 321,000 respondents

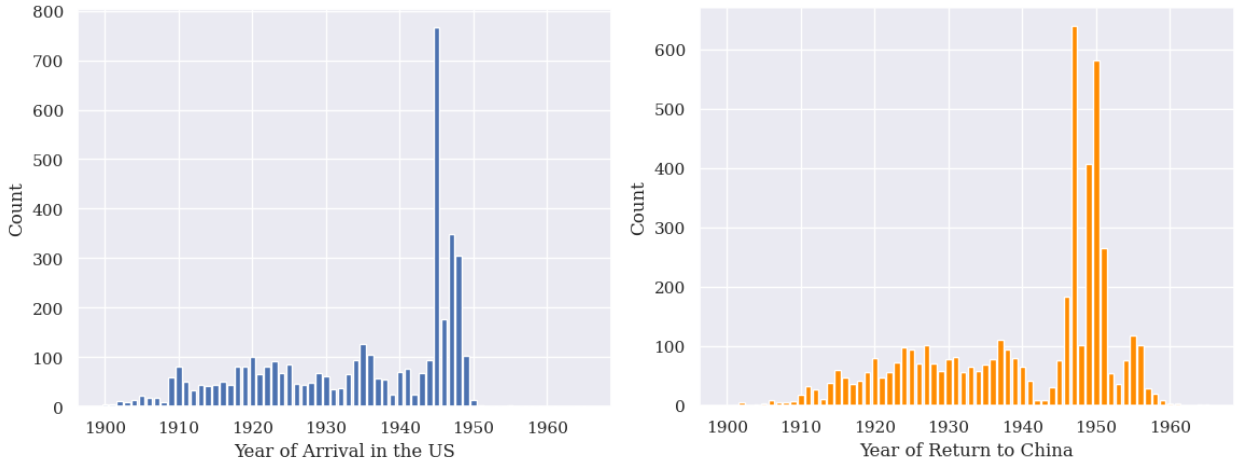


Figure 2: Year of arrival in the U.S. and year of return to China for U.S.-trained returnees in our sample.

Notes: Histograms of the year of first arrival in the United States (left) and the year of return to China (right) for the merged returnee database. About 80 percent of returnees trained in STEM fields (43 percent engineering, 19 percent biology, 14 percent physical sciences), 13 percent in the social sciences, and 3 percent in the humanities. Source: authors' biographical returnee database.

per wave. We harmonize the two pools into a common person-level schema, deduplicate across sources, and cross-filter against the returnee list to remove anyone recorded as having returned to China. This yields 9,478 stayers: 3,232 from the dissertation guide only, 6,212 from the National Register only, and 34 from both. Because the National Register has no person identifier, we merge records across waves only when they share an employer, degree institution, birth year, or degree year and field. Appendix Table 13 summarizes both sources. As with the returnees, each stayer is linked to their publication record in OpenAlex and CNKI. Comparing returnees to these stayers, who share the same educational pipeline but remained in the United States, is what allows us to isolate the effect of return migration from the characteristics of the scientists who returned.

3.4 Analysis Samples

The two populations together form the universe of U.S.-trained Chinese scholars whose return or degree year falls in or before 1960. For the individual-level analysis, we keep scholars matched to at least one research publication in OpenAlex or CNKI. After resolving over-merged identities, this *individual-productivity panel* contains 3,165 of the 4,684 returnees and 2,914 of the 9,478 stayers, observed at the scholar-year level from 1900 to 1976. Of these, 3,067 returnees and 2,754 stayers have at least one matched publication by 1976, and 2,514 and 1,549 have at least one corroborated publication. Most stayers never enter the panel because they have no matched publication at all. The matched design of Section 4 is

drawn from this panel. For the analyses of domestic scientific production, we instead use the full CNKI corpus of 565,780 Chinese-language records for 1915–1976 (564,632 papers with embeddable text), whose annual counts reproduce CNKI's own official record counts, aggregated to the subfield-year level. That corpus measures output by *other* domestic scientists, whom the scholar-level panel never observes.

3.5 Measures of Individual-Level Scientific Productivity and Research Agendas

We construct annual, individual-level measures of scholarly output by combining bibliometric, patent, and archival sources. Our measures capture not only the *volume* and *influence* of output, but also the *orientation* of research agendas, following methods developed in recent work on the economics of science and innovation.

Publications. We aggregate publications from two complementary databases. OpenAlex, the successor of the discontinued Microsoft Academic Graph, provides global coverage with standardized metadata and citation linkages (as in Moser et al., 2025), while the China National Knowledge Infrastructure (CNKI) captures nearly all Chinese-language academic literature. For CNKI, we follow the construction in Acemoglu et al. (2024), collecting titles, abstracts, authors, publication years, venues, and citation counts. For each scholar, we query under three naming conventions (Chinese characters, Hanyu Pinyin, Wade–Giles) to address historical variation in name usage. In OpenAlex we score romanized names (Pinyin, Wade–Giles, and Cantonese variants) and require the publication year to be plausible and the affiliation to fit the scholar's location at the time (U.S. before return, China after). In CNKI we match primarily on exact Chinese characters. Because romanized names collide often, every match must also be corroborated by the scholar's field or institution. We chose these rules by simulating alternatives on a hand-labeled audit of 190 matches; after them, measured precision is 92 percent in CNKI and 50 percent in OpenAlex. Our headline counts therefore use corroborated matches only, and uncorroborated matches are kept, flagged, for robustness. From these data we measure publications and citations per scholar-year in each corpus.

Field Classification and Research-Agenda Vectors. To characterize the *content* and *evolution* of research, we construct scholar-level research-agenda vectors using multilingual text embeddings. We adapt methods from Iaria et al. (2018), Bello et al. (2025), and Koffi (2025), combining English- and Chinese-language bibliometric sources to assemble a compre-

hensive corpus of publication titles and abstracts for each scholar–year. We first harmonize author identities, publication years, and venues across databases, and preprocess text by lowercasing, de-accenting, removing stopwords, and tokenizing. We then represent research content using multilingual sentence embeddings. Specifically, we embed each paper's title and abstract with LaBSE, a multilingual `SentenceTransformers` model (SciBERT and a Chinese-language SBERT serve as robustness checks), which map texts in different languages into a shared semantic space. For each scholar i in year t, we define an *agenda vector* as the mean embedding of all publications authored in that year. This vector provides a continuous, high-dimensional representation of the scientist's research orientation. To capture the *direction* of research along a dimension that embedding distances do not resolve—whether work becomes more basic or more applied—we complement these geometric measures with a direct-content measure in which a large language model reads each scholar's pre- and post-return titles and abstracts and scores both the magnitude and the basic-to-applied direction of the agenda shift (Section 5.3).

To assign research fields, we use two established subject taxonomies rather than constructing fields de novo. OpenAlex publications carry OpenAlex's native hierarchical classification, whose 244 subfields we adopt as the unit of analysis; Chinese-language CNKI publications carry FORD subject categories, which define the within-CNKI (domestic) frontier. To place both corpora in a common field space, we embed every title and abstract with multilingual LaBSE vectors and assign each CNKI paper to the OpenAlex subfield whose reference-paper centroid is nearest in cosine similarity. These same embeddings, following the measurement approach of Hager et al. (2024) and Moser et al. (2025), furnish the paper-level frontier-alignment outcomes used throughout. Each paper is assigned uniquely to a subfield, enabling the construction of scholar–year field shares and dominant research fields.

Outcome Variables. For each scholar-year we construct:

1. **Output:** $\log(1 + \text{publications})$ and $\log(1 + \text{citations})$, separately for OpenAlex and CNKI and summed across the two.

2. **Quality and direction (semantic):** the textual *importance* of a scholar's papers and their *frontier alignment*, the cosine similarity of the scholar-year agenda vector to the breakthrough centroid of its subfield-year (Kelly et al., 2021), constructed as described below.

3. **Agenda change (language model):** for each scholar, the magnitude of the shift between pre- and post-return research and its direction on a basic-to-applied axis (Sec-

tion 5.3), and the alignment of each paper with the priorities of the 1956 Twelve-Year Plan (Section 5.4).

Dual-Frontier Construction. A distinctive feature of our setting is that returnees publish in a domestic corpus (CNKI) while stayers publish in an international corpus (OpenAlex), and the two corpora point to different research frontiers. To leverage this, we construct *two* versions of every Kelly-style outcome. For each scholar's paper p in subfield s and year t, we compute backward similarity to papers in $s \times [t-5, t-1]$ and forward similarity to papers in $s \times [t+1, t+10]$ using semantic embeddings (LaBSE multilingual, with SciBERT and sbert-chinese as robustness checks), and then form $r \log q_p$ following Kelly et al. (2021), residualized on year fixed effects; a pair of papers counts as similar only when its cosine exceeds a threshold θ, set to the 95th percentile of cosines between random reference-sample paper pairs (model-specific). Each paper is scored against a citation-neutral reservoir of 200 papers per subfield-year, sampled independently of citations, so the resulting importance score is textual rather than citation-based. The *breakthrough* papers in a subfield-year are the top decile of the residualized score, and the frontier is their mean embedding; a scholar's frontier alignment is the cosine of their agenda vector to this centroid. We verify robustness across three breakthrough definitions (top-decile importance, top-decile citations, and top journal) and across embedding models: three on the OpenAlex side (LaBSE, SciBERT, and a Chinese-language SBERT) and two on the CNKI side (LaBSE and the Chinese SBERT, since SciBERT is English-only). We construct $r \log q_p$ within each of two reference universes:

1. *Domestic frontier* (CNKI reference): the 106,000-paper stratified CNKI sample, with subfields defined on FORD subcategories. This measures how each paper relates to the leading Chinese-language papers in its subfield.

2. *Global frontier* (OA reference): OpenAlex's 3.5-million-paper stratified reference, with each Chinese-language paper nearest-neighbor assigned to one of 244 OA subfields via centroid cosine similarity. This measures how each paper relates to the leading international papers in the same field.

Because the same scholar's papers receive both scores, we can read the gap directly: a paper that is highly aligned with the Chinese frontier (high CNKI-reference alignment) but distant from the global frontier (low OA-reference alignment) is locally important but internationally peripheral. The complementary case, alignment with the global frontier, captures *catching-up* to international science. Table 1 previews the unconditional gap; Section 6 traces how it evolves around return.

Table 1: Frontier alignment of scholar publications: unconditional means by reference universe

Comparison	Mean	SD	N
OA scholars vs OA frontier	0.392	0.153	13,181
CNKI scholars vs CNKI frontier	0.604	0.122	4,710
CNKI scholars vs OA frontier (catching-up)	0.356	0.133	4,837

Notes: Mean LaBSE cosine similarity between matched scholar publications and the breakthrough centroid in the corresponding subfield-year, reported separately for the three reference universes used in the paper. Row 1: OpenAlex publications by matched scholars compared to the OpenAlex breakthrough centroid. Row 2: CNKI publications by matched scholars compared to the CNKI breakthrough centroid (within-corpus alignment). Row 3: CNKI publications compared to the OpenAlex breakthrough centroid, the catching-up measure used in Section 6. SD denotes the cross-paper standard deviation. N is the number of paper-year observations with non-missing alignment in each row.

Validating the Embedding Space. Before using these vectors as outcome measures, we verify that they recover the substantive structure of science across both corpora. Figure 3 shows three side-by-side t-SNE projections at the field (macro-domain) level: the left panel uses only OpenAlex matched scholar papers (76,326 papers, sampled to 30,000 for tractability), the middle panel uses the full set of CNKI matched scholar papers (16,481 papers, all shown), and the right panel pools both corpora (92,807 papers, sampled to 30,000). Figure 4 repeats the exercise at the OpenAlex subfield level. In all three projections, papers from the same field or subfield cluster tightly, and substantively related categories (e.g. atomic and condensed-matter physics; electrical and mechanical engineering) occupy adjacent rather than disjoint regions of the space. The clustering structure is recovered separately within each corpus and is preserved when the two are pooled, supporting the use of LaBSE distances as the substrate for the frontier-alignment outcomes that follow.

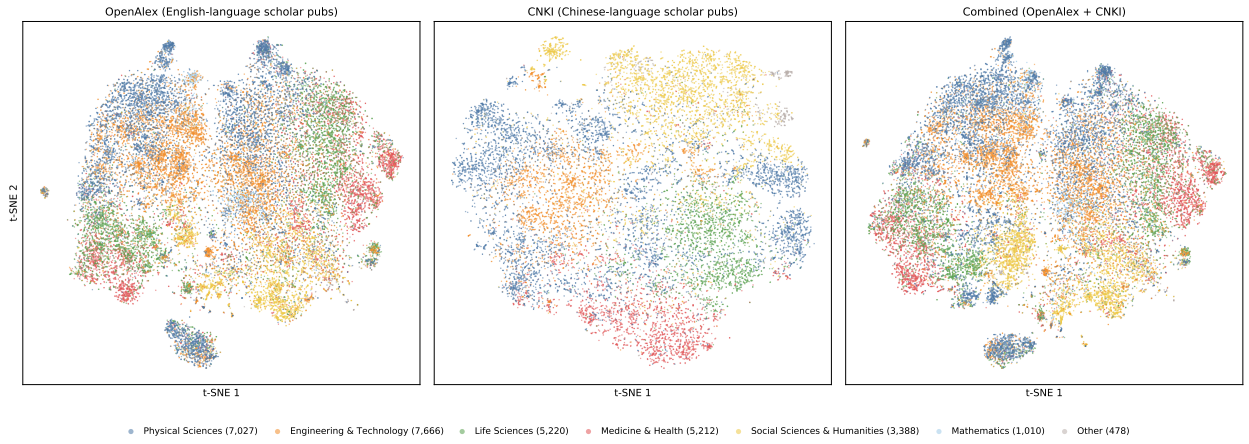


Figure 3: t-SNE map by macro-domain: OpenAlex, CNKI, combined

Notes: t-SNE projections of LaBSE paper embeddings, computed separately per corpus. Left: 30,000 OA matched scholar papers (stratified by field). Middle: all 16,481 CNKI matched papers. Right: pooled OA+CNKI, 30,000 stratified. Colors: macro-domain (CNKI FORD majors mapped to OA fields). Legend counts refer to the combined panel.

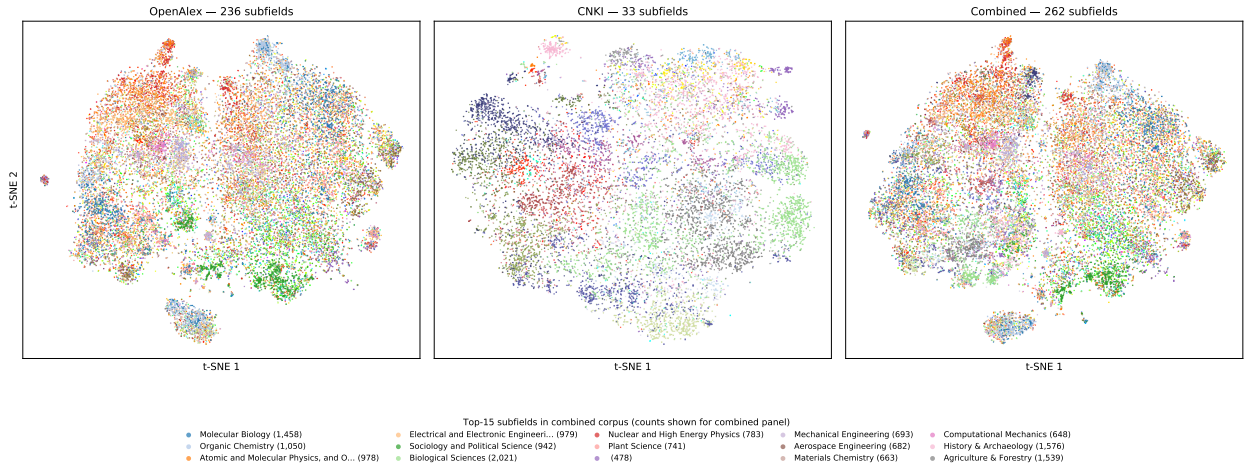


Figure 4: t-SNE map by OpenAlex subfield: OpenAlex, CNKI, combined

Notes: Same projections as Figure 3; points colored by OA subfield (every subfield assigned a distinct color; panel header reports the count). Legend lists the top-15 subfields in the combined corpus; remaining subfields are colored but unlabeled. CNKI subfields are mapped to OA subfields via centroid cosine similarity.

4 Empirical Strategy

What the design identifies. We estimate how a scientist's output and research agenda change *after* returning to China, relative to observationally similar scientists from the same training pipeline who stayed in the United States. The U.S. exit restrictions of 1951–1955

and the closure of China after 1957 are the historical context for this comparison rather than an instrument. Table 2 summarizes the declassified State Department accounting of the restrictions: of the roughly 434 students who applied to leave between 1951 and April 1954, 310 were permitted and 124 were temporarily denied. We have digitized these records (Appendix A.3), but we do not observe exit applications for the full population, so we cannot instrument the return decision itself. The main threat to our design is therefore that returnees were on a different career trajectory regardless of return. We address it with matching, rich fixed effects, and tests for pre-trends on every outcome. The welfare accounting in Section 7 does not rest on this design: it prices the observed allocation of returnees across fields and requires no counterfactual productivity.

Table 2: Chinese Students in the United States (1946–1954)

Item	Number (approx.)
1. Total number of Chinese who arrived in U.S. on student visas or as students on other visas	
1946–50	4,800
1951–54	1,500
Total	6,300
2. Departed from U.S. before June 1951 (Korean War)	800
3. Applied to depart 1951–April 1954 (restrictions in effect)	434
4. Permitted to depart 1951–54	310
5. Temporarily denied permission to depart because of U.S. security regulations 1951–April 1954	124
6. Of those in 5 above, number who wished to depart when canvassed June 1954	62
7. Of those in 6 above:	
a. Departed U.S.	15
b. Notified, but not yet departed	12
c. Not notified they could leave	35
8. Additional students applying for permission to depart since June 30, 1954 (cases in process)	36
9. Total departed U.S. 1946–54	1,100
10. Remaining in U.S. as of December 1954	5,200

Notes: Counts as reported in the source, which rounds most figures. *Source:* U.S. Department of State, "Memorandum for the President: Status of Chinese Students Who Have Been Refused Permission to Leave the United States," March 3, 1955, declassified memorandum.

Matched returnee–stayer pairs. From the individual-productivity panel (Section 3.4), we match each returnee whose return year falls in 1940–1957 to one stayer. The matching is an optimal (Hungarian) assignment without replacement that first maximizes the number of pairs and then minimizes the total distance. A pair is feasible only if the two scholars

share a coarse research field and their first publications are no more than seven years apart. Within that set, the distance penalizes differences in first-publication year and birth year, and adds a penalty when the fine discipline differs. This yields 1,324 pairs, the maximum attainable under these constraints. Coarse field matches exactly in every pair, and 74.6 percent of pairs also match on fine discipline. Each matched stayer inherits its returnee's return year as a placebo event date $\tau = 0$. Outcomes are tracked through 1976, when the Cultural Revolution ended.[1]

Table 3 reports balance. All normalized differences lie within the conventional $|ND| < 0.25$ band. Pre-return output, which the matching does not target, is balanced ($ND = 0.069$), so this balance is not mechanical.

Table 3: Covariate balance: returnees versus matched stayers

	Returnees	Matched stayers	Difference	Normalized diff.	N
Birth year	1918.4	1919.9	−1.5	−0.242	517
Year of first publication	1956.3	1958.3	−2.1	−0.201	1,324
Pre-period output, mean log(1 + pubs)	0.034	0.026	+0.008	+0.069	1,324
Any pre-period publication	0.116	0.100	+0.016	+0.051	1,324

Notes: Headline matched sample of 1,324 returnee–stayer pairs (1,324 distinct stayers). Coarse field is matched exactly in all pairs and fine discipline in 74.6 percent. Birth year is observed for both members in 39.0 percent of pairs; N is the number of pairs with the variable observed for both members. The normalized difference is the difference in means divided by the square root of the average of the two group variances. Balance is tighter still on the sample restricted to corroborated publication matches (maximum $|ND|$ of 0.181, 933 pairs).

Individual-level event study. On the matched sample we estimate

$$y_{it} = \alpha_i + \lambda_t + \theta_{f(i),t} + \sum_{k \neq -1} \delta_k \, \mathbf{1}\{\tau_{it} = k\} \cdot \text{Returnee}_i + u_{it}, \tag{1}$$

where y_{it} is an outcome of scholar i in calendar year t, τ_{it} is event time relative to the (placebo) return year, α_i are scholar fixed effects, and $\theta_{f(i),t}$ are fixed effects for the scholar's own broad field by calendar year, which absorb the calendar-year effects λ_t. We use the scholar's own field rather than the pair's because 25.4 percent of pairs match only on coarse field; field-by-year shocks do not difference out within those pairs. The event window is $\tau \in [-5, 15]$ and the reference period is $\tau = -1$. Standard errors are clustered by matched pair. The pooled analogue replaces the event-time indicators with a single $\text{Returnee}_i \times \text{Post}_{it}$ term, where $\text{Post}_{it} = \mathbf{1}\{\tau_{it} \geq 0\}$.

[1] After 1976 the treated and control environments are no longer separable: "reform and opening" began in 1978, and by 1978–79 China was again sending students abroad, mainly to the United States and Japan. The domestic corpus itself nearly disappears in 1967–69 and recovers only from 1970.

Field-level spillovers. To study the domestic system beyond the returnees themselves, we assign each returnee to their modal CNKI subfield and date each subfield's treatment by the arrival of its first returnee, t_s^0. We then estimate a stacked event study in which each arrival cohort s is compared with subfields that never received a returnee:

$$y_{jts} = \sum_{k=-5,\, k\neq -1}^{20} \beta_k \left(\text{Dose}_j \times \mathbf{1}\left[t - t_s^0 = k\right] \right) + \alpha_{js} + \vartheta_{ts} + \varepsilon_{jts}. \tag{2}$$

Here y_{jts} is the log of one plus a count of papers in subfield j and year t, measured on the full 564,632-paper CNKI corpus. Dose_j is the standardized number of returnees whose modal subfield is j (zero for never-treated subfields). α_{js} and ϑ_{ts} are subfield and year fixed effects estimated within each stack. Standard errors are clustered by subfield. Because 97 percent of the 17,920 stacked subfield-year cells belong to never-treated subfields, one standard deviation of dose corresponds to about 1.4 returnees. We report magnitudes in those terms throughout.

5 Effects on Individual Scholars

5.1 Returnees Redirect Output

Table 4 reports the pooled version of equation (1) for six outcomes: $\log(1 + \text{publications})$ and $\log(1 + \text{citations})$, each measured in OpenAlex only, in CNKI only, and summed across the two corpora. OpenAlex captures internationally circulating, largely English-language scholarship; CNKI captures Chinese-language journal articles. The decomposition therefore separates effects on international publishing from effects on domestic publishing.

Returnees publish less in OpenAlex after return (column 1) and more in CNKI (column 2), and their total publication count does not change (column 3). Citations move in the same direction but with a larger international decline (columns 4–6). A study that observed only OpenAlex, as nearly all bibliometric work on this period does, would conclude that return sharply reduced both productivity and impact. Once both corpora are observed, the effect on total output disappears: return *reallocated* scholarly output across research systems rather than reducing it. The lower citation count largely reflects venue rather than research quality, because Chinese-language publication mechanically reaches a smaller citing audience. We test research quality directly in Section 5.2.

The pooled coefficients are -0.057 for OpenAlex and $+0.071$ for CNKI, about 6 percent less and 7 percent more; the total is small and insignificant. Because the effects build over time after return, the pooled coefficient understates the typical post-return gap. Figure 5

Table 4: Effect of return on publications and citations

	log(1 + publications)			log(1 + citations)		
	(1) OpenAlex	(2) CNKI	(3) Total	(4) OpenAlex	(5) CNKI	(6) Total
Returnee × Post	−0.057***	0.071***	0.010	−0.144***	0.081***	−0.068***
	(0.007)	(0.006)	(0.009)	(0.019)	(0.010)	(0.020)
Mean of dep. var.	0.073	0.043	0.114	0.152	0.046	0.196
Scholar FE	Yes	Yes	Yes	Yes	Yes	Yes
Field × year FE	Yes	Yes	Yes	Yes	Yes	Yes
Matched pairs	1,324	1,324	1,324	1,324	1,324	1,324
Observations	55,608	55,608	55,608	55,608	55,608	55,608

Notes: Pooled difference-in-differences version of equation (1), estimated on the 1,324 matched returnee–stayer pairs (return years 1940–1957) over event time $\tau \in [-5, 15]$. Post $= \mathbf{1}\{\tau \geq 0\}$; matched stayers carry their returnee's return year as a placebo event date. The regressions include scholar fixed effects and fixed effects for the scholar's broad field by calendar year. Standard errors clustered by matched pair are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

traces these dynamics. Averaged over the post-return event times, the OpenAlex coefficients are about −0.09 log points (about 9 percent fewer publications), the CNKI coefficients about +0.09 (about 10 percent more), and the total essentially zero (+0.0005). The pre-return coefficients are flat. Across the six outcomes, and across both the full and the corroborated publication samples, only 1 of 24 pre-period coefficients is significant at the 5 percent level, fewer than the 1.2 expected by chance. Citations follow the same pattern with larger magnitudes (Figure 6).

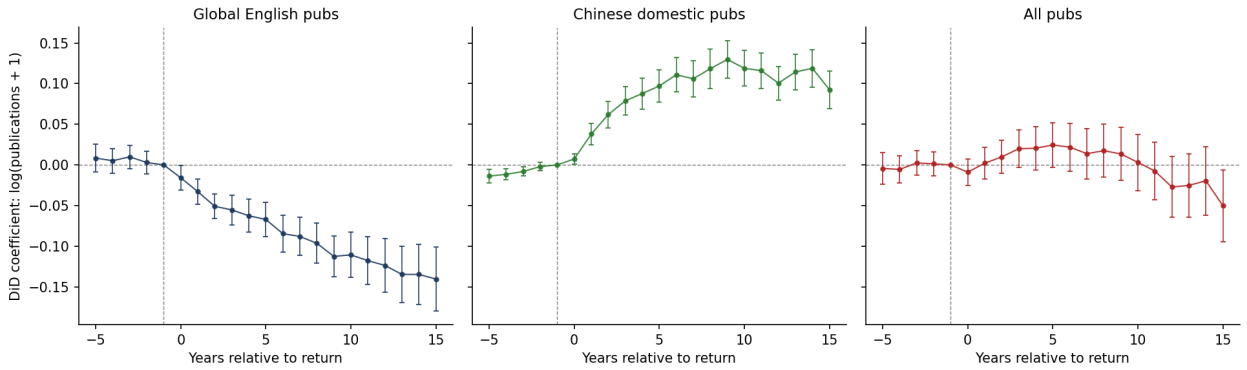


Figure 5: Event study, log(1 + publications)

Notes: Coefficients $\hat{\delta}_k$ from equation (1), estimated on 1,324 matched returnee–stayer pairs, with scholar fixed effects and broad-field-by-year fixed effects. Reference period $\tau = -1$. Left: OpenAlex (international) publications. Middle: CNKI (Chinese-language) publications. Right: both corpora. Bars are 95 percent confidence intervals from pair-clustered standard errors.

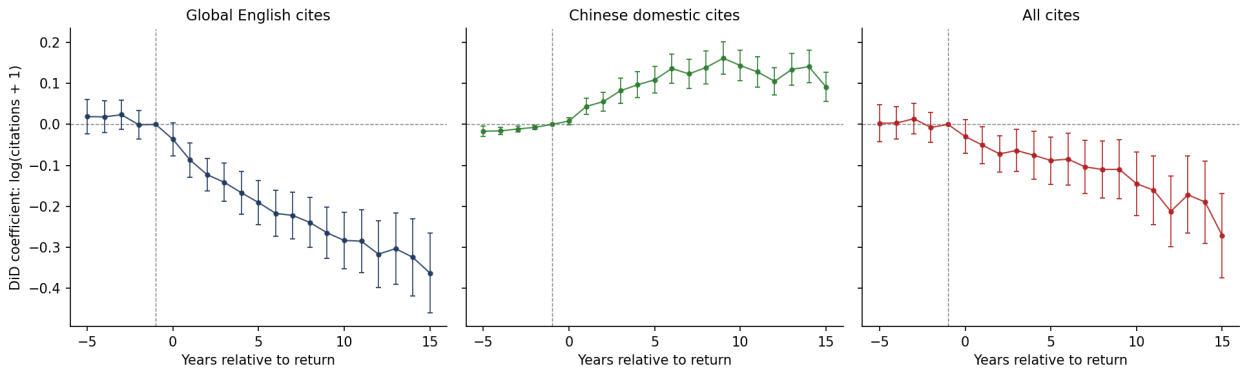


Figure 6: Event study, $\log(1 + \text{citations})$

Notes: Same specification as Figure 5, with $\log(1 + \text{citations})$ as the outcome.

One concern is that the domestic increase merely reflects the growth of CNKI's coverage. Calendar-year fixed effects difference out any corpus-wide coverage trend. Moreover, returnee-authored work is a stable 2–4 percent of the CNKI corpus from 1950 onward, even as the corpus grows from about 8,000 to 48,000 records a year.

5.2 Research Quality Does Not Fall

The reallocation could hide a decline in quality if returnees published weaker work after return. Measuring quality corpus by corpus cannot answer this question. Returning *is* a move from one literature to the other, so a returnee's pre-return work is almost entirely in OpenAlex and their post-return work almost entirely in CNKI: no window exists in which the same corpus-specific measure is observed on both sides of return. We therefore *pool* each scholar's entire output. For scholar i in year t, the pooled measure is the publication-weighted average

$$x_{it} = \frac{n_{it}^{\mathrm{OA}}\, x_{it}^{\mathrm{OA}} + n_{it}^{\mathrm{CNKI}}\, x_{it}^{\mathrm{CNKI}}}{n_{it}^{\mathrm{OA}} + n_{it}^{\mathrm{CNKI}}},$$

where both components are scored against the same OpenAlex reference, so they lie on one scale. Pooling raises the number of returnees observed before return from a handful to 95 and the number of stayers observed from 5 to 523, which turns the comparison into a real difference-in-differences. We study two outcomes: the paper-level *importance* score of Section 3, which measures quality, and *frontier alignment*, which measures closeness in content to the breakthrough frontier. The two are only weakly correlated across scholar-years (0.16 in OpenAlex and 0.12 in CNKI), so alignment is better read as the direction of research than as its quality.

Neither measure moves after return (Table 5 and Figure 7). Pre-trends are flat for both:

Table 5: Effect of return on research quality, pooled across corpora

	(1) Importance, $\log(FS/BS)$	(2) Frontier alignment
Returnee × Post	0.171	−0.007
	(0.112)	(0.026)
Mean of dep. var.	0.047	0.356
Scholar FE	Yes	Yes
Field × year FE	Yes	Yes
Scholars	1,409	1,409
Observations	4,252	4,248

Notes: Pooled difference-in-differences version of equation (1) on the matched sample. The outcomes are publication-weighted averages of each scholar's OpenAlex and CNKI measures, both scored against the OpenAlex reference corpus. Column 1: the textual importance score of Kelly et al. (2021), log of forward over backward similarity, residualized on year. Column 2: cosine similarity of the scholar-year agenda vector to the subfield-year breakthrough centroid. Scholar-years with no publication have no measure and are dropped. Standard errors clustered by matched pair are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

the largest pre-period $|t|$ is 0.82 for importance and 1.76 for alignment. No individual post-return coefficient is significant. This is a clean, well-identified null, and it is a demanding test: both halves of each scholar's output are benchmarked against *world* science, so a paper that was important inside China but invisible abroad scores low. Returnees changed *where* they published, not the caliber of what they produced.

5.3 Redirecting the Research Agenda

The count and frontier-alignment results show that return reallocates *where* scholars publish without lowering how much or how close to the frontier they publish. They do not, however, speak to *what* returnees work on. The planning regime that directed returnee placements toward defense, heavy industry, and applied engineering (Section 6) plausibly redirected not just careers but the substantive content of research, pulling scholars away from the basic questions they pursued in the United States and toward mission-oriented, applied work. Semantic frontier-alignment, which measures distance to a subfield-year breakthrough centroid, is silent on this basic-versus-applied dimension: a paper can be equally frontier-proximate whether it develops fundamental theory or solves an applied engineering problem.

To measure the direction of research directly, we use a large language model to read each scholar's actual publications. For every scholar we assemble the titles and abstracts of their pre-return and post-return output, pooling English-language OpenAlex and Chinese-

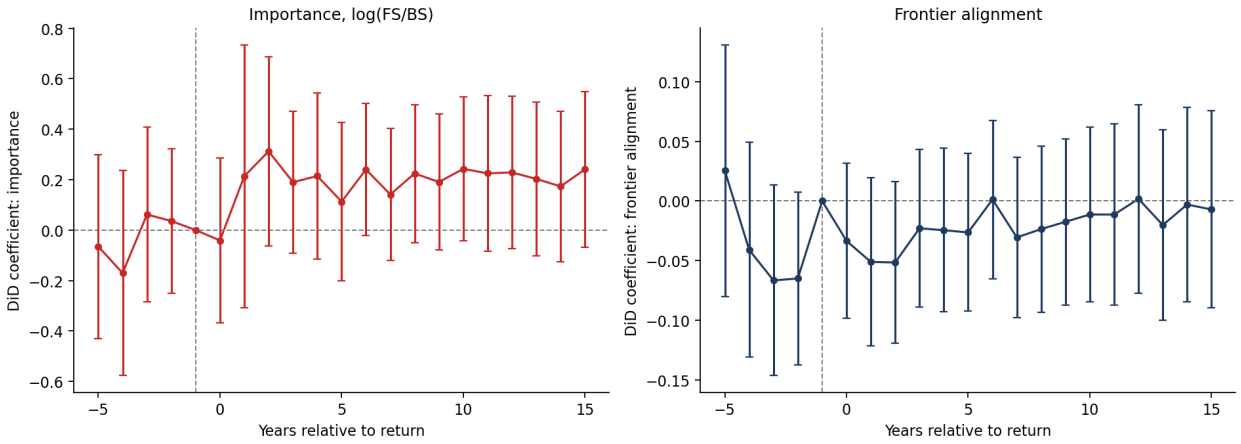


Figure 7: Research quality around return, pooled across corpora

Notes: Coefficients $\hat{\delta}_k$ from equation (1) with the pooled importance score (left) and pooled frontier alignment (right) as outcomes; 4,252 scholar-years from 1,409 scholars. Reference period $\tau = -1$. Bars are 95 percent confidence intervals from pair-clustered standard errors.

language CNKI text, and prompt `claude-sonnet` to compare the two bodies of work and return two scores: the *magnitude* of the agenda shift (0 = identical research program, 100 = unrelated) and its *direction* on a basic-to-applied axis (−100 = decisively more basic, +100 = decisively more applied). This direct-content measure is complementary to the embedding-based measures: rather than locating a scholar in a fixed semantic space, it asks a multilingual reader to judge how, and in which direction, the research program moved. We restrict the sample to the 244 scholars (171 returnees and 73 stayers) with enough bilingual abstract text on both sides of return to support the comparison. Unlike the count and frontier-alignment results, this comparison is drawn from the full universe of returnees and stayers rather than from the matched pairs: returnees are assigned their own return year and stayers a common placebo year of 1946, the median returnee return year. Of the 244 scholars, 133 also appear in the matched sample.

Table 6 reports the unconditional difference between returnees and stayers in the pre-to-post change. Returnees redirect their research decisively toward applied work: the returnee-stayer gap in the applied-direction score is +21.7 ($p < 0.01$), against a stayer mean change of only +13.6. Returnees also change their agenda somewhat more in overall magnitude (+8.1, $p < 0.05$), and are 16 percentage points more likely to move toward the applied end of the spectrum. The redirection is thus not simply that returnees revise their research more than comparable stayers as careers mature; it is that the revision is systematically *toward* applied, mission-oriented science, exactly the margin that citation counts and frontier-alignment cannot see.

Table 6: Return redirects research toward applied work

	(1) Agenda shift (0–100)	(2) Toward applied (−100/+100)
Returnee × Post	8.10**	21.73***
	(3.53)	(4.30)
Stayer mean change	59.1	13.6
Observations	244	244

Notes: Each scholar's pre-return and post-return publications (pooled English-language OpenAlex and Chinese-language CNKI titles and abstracts) are read by `claude-sonnet`, which returns the magnitude of the agenda shift (column 1; 0 = unchanged, 100 = unrelated) and its direction on a basic-to-applied axis (column 2; −100 = more basic, +100 = more applied). The sample is the 244 scholars (171 returnees, 73 stayers) with sufficient bilingual abstract coverage on both sides of return, drawn from the full universe rather than the matched pairs; stayers are assigned a common placebo year of 1946, the median returnee return year. Estimates are unconditional returnee-stayer differences with no matching, pair fixed effects, or covariates. Each entry is the estimated returnee-stayer difference in the pre-to-post change, with heteroskedasticity-robust standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

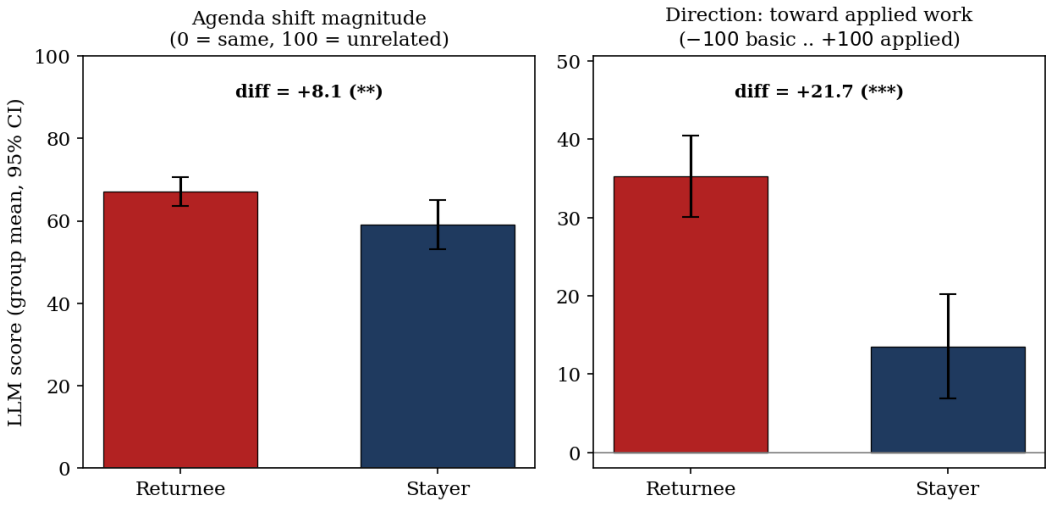


Figure 8: Returnees redirect their research toward applied work

Notes: Group means of the language-model agenda scores, which compare each scientist's own papers before and after the return year (a placebo year for stayers). Left: magnitude of the agenda shift (0 = same research program, 100 = unrelated). Right: direction on the basic-to-applied axis (−100 = more basic, +100 = more applied). The sample is the 244 scholars with enough bilingual abstract text on both sides of return. Bars are 95 percent confidence intervals; the reported differences correspond to Table 6.

In standardized units, the direction effect is about twice the magnitude effect: 0.63 versus 0.33 standard deviations. An embedding-based measure that uses no language model gives the same answer. We define drift as one minus the cosine similarity between a scholar's pre- and post-return topic shares, computed from the same LaBSE embeddings. Returnees drift more than stayers at every topic granularity, from 52 to 864 topics. The gap, 0.20–0.40 standard deviations, matches the language-model magnitude effect on the same scientists, and the two measures correlate at 0.21–0.34. Embeddings cannot, however, express *direction*: basic versus applied is not an axis of the semantic space, and direction is where the result lies. This is why the language-model measure leads.

At the topic level, this redirection is a movement *within* a scholar's research, from its basic toward its applied end, rather than a change in *which* topics they pursue.

The single clearest case is Qian Xuesen, whose career motivates the exercise. At Caltech his research was *basic* aerodynamics, propulsion, and structural mechanics; after his 1955 return it pivots to *applied*, mission-oriented work: cybernetics, systems science, and the methodology of "technical science", the intellectual scaffolding for China's missile and space programs (Table 7). His broad subfield, aerospace engineering, is unchanged across return, so a subfield-level measure registers no movement; the redirection is visible only in the content of the work itself.

Table 7: Qian Xuesen's research before and after return

Before return, Caltech, 1939–52 *Basic aerodynamics & structures*	**After return**, China, 1956–65 *Applied control & mission science*
1939. Two-Dimensional Subsonic Flow of Compressible Fluids	1957. Engineering Cybernetics
1942. A Theory for the Buckling of Thin Shells	1957. On Technical Science
1946. Similarity Laws of Hypersonic Flows	1957. Methodological Problems in Technical Science
1948. Wind-Tunnel Testing Problems in Superaerodynamics	1957. Introduction to Physical Mechanics
1952. Applications of Jet Propulsion Power Plants	1957. Long-Range Spaceflight
	1964. Large-Scale Scientific Experimentation

Notes: Representative publications of Qian Xuesen (H. S. Tsien) before and after his 1955 return to China, illustrating the basic-to-applied redirection measured in Table 6. Chinese-language titles are given in English translation. His OpenAlex subfield (aerospace engineering) is identical on both sides; the shift is in research content, not field.

5.4 Applied Toward What? Alignment with the Twelve-Year Plan

The 1956 Twelve-Year Plan for Science and Technology named twelve priority areas: atomic energy; radio electronics, semiconductors, and computers; jet and rocket technology; automation and precision instruments; petroleum and scarce minerals; alloys and new metallurgy; fuel utilization and heavy organic synthesis; new power and large-scale machinery; development of the Yellow and Yangtze rivers; mechanization and chemicals in agriculture; control of the major endemic diseases; and basic theory in the natural sciences. We ask a language model to score each paper from 0 to 100 against the plan's own twelve Chinese-language priority statements. The prompt is blind: it does not mention return, the U.S. restrictions, or whether the author stayed. About 13,900 works are scored. The measure validates well: it correctly classifies all six face-validity cases, has a test–retest correlation of 0.964 when the list of priorities is shuffled, agrees with a binary recode 98.3 percent of the time, and correlates at 0.50 with our hand-coding of strategic fields.[2]

Figure 9 shows that returnees' plan alignment rises after return and stays elevated, at 6 to 10 points above that of matched stayers. We read this pattern as *descriptive* rather than causal, for two reasons. First, the plan is a Chinese-language document, and the CNKI share of the pooled measure rises from 0.06 before return to 0.92 after it, so the post-return rise cannot be separated from the switch to Chinese-language publishing. Second, the dip at $\tau = -2$ is an artefact of normalizing to $\tau = -1$ in cells of only 17–25 scholars. In levels, returnees are more plan-aligned than stayers at 20 of 21 event times, including 6 to 11 points *before* return. The defensible fact is therefore a level difference that points to selection: the scientists who returned were already working closer to what China would prioritize. This is unsurprising, because the plan was drafted with input from returnees, among them Qian Xuesen.

6 Effects on Domestic Scientific Production

Returnees did not publish less in total or produce lower-quality work after return. We now ask whether their arrival reshaped the wider Chinese scientific system, and in particular the output of the domestic scientists who worked alongside them. In the early 1950s many returnees entered fields that were nascent or effectively nonexistent in China, including modern physics, electronics, aeronautics, and applied engineering. Placements were set by the central government in line with national priorities. After 1957, when China was cut off

[2] An alternative measure based on LaBSE cosine similarity to the plan text fails the same validation. It correlates at −0.10 with the hand-coding, passes only 2 of the 6 face-validity cases, and suffers from an anisotropic embedding space in which unrelated subfields have a mean pairwise cosine of 0.53.

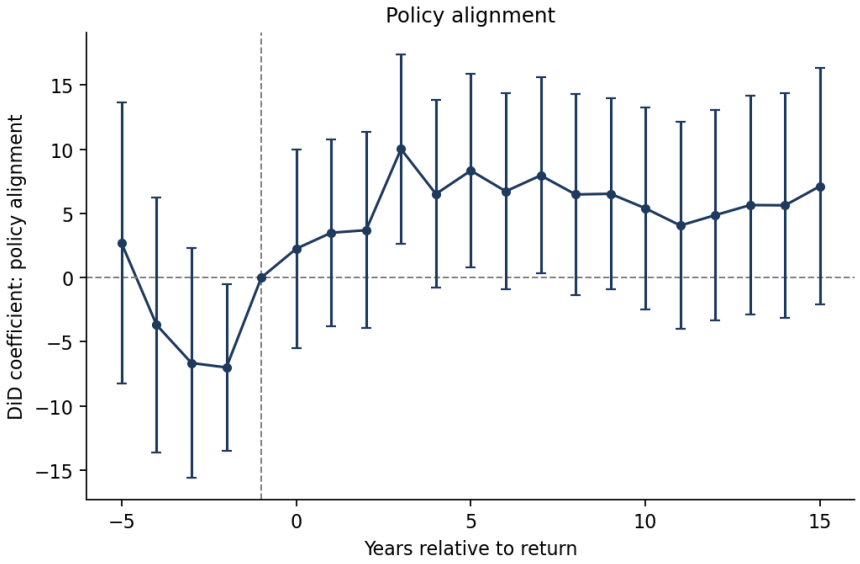


Figure 9: Alignment with the 1956 Twelve-Year Plan around return

Notes: Coefficients $\hat{\delta}_k$ from equation (1), where the outcome is the publication-weighted average of each scholar's paper-level scores (0–100) against the twelve priority areas of the 1956 plan. Scores are assigned by a language model under a blind prompt. Reference period $\tau = -1$; 95 percent confidence intervals from pair-clustered standard errors. The estimates are descriptive; see the text for the confound with the language of publication.

from Western science until the reforms of 1978, returnees became a primary conduit through which frontier methods, topics, and tacit knowledge reached the domestic system.

6.1 More Returnees, More Output by Others

Figure 10 reports the stacked event study of equation (2), with the log of domestic publications in a subfield-year as the outcome. The left panel counts all domestic output. The right panel counts only papers by *others*, excluding every paper written by any returnee, including their coauthored papers. Coefficients are flat before the first returnee arrives and rise steadily afterward. Over the first decade after arrival (event times 1–10), a subfield receiving one standard deviation more returnees, about 1.4 people, produced about 56 percent more domestic output in total ($\hat{\beta} = 0.444$) and about 69 percent more output by others ($\hat{\beta} = 0.526$). Averaging across such subfields, this is about 25 additional papers a year by non-returnee scientists.

The effect is *larger* once returnees' own papers are removed. The gain therefore does not come from returnees publishing more themselves: frontier knowledge diffused into the surrounding domestic community. The effect also builds over time, from +39 percent in event years 0–4 to +107 percent in years 11–20. With the logarithm of the returnee count as

the dose instead of its standardized level, the first returnee is associated with an 86 percent increase in others' output and the tenth with only 9 percent. Returns diminish steeply, so the field-creating effect of the first arrivals carries the result. The two specifications agree closely at the extensive margin.

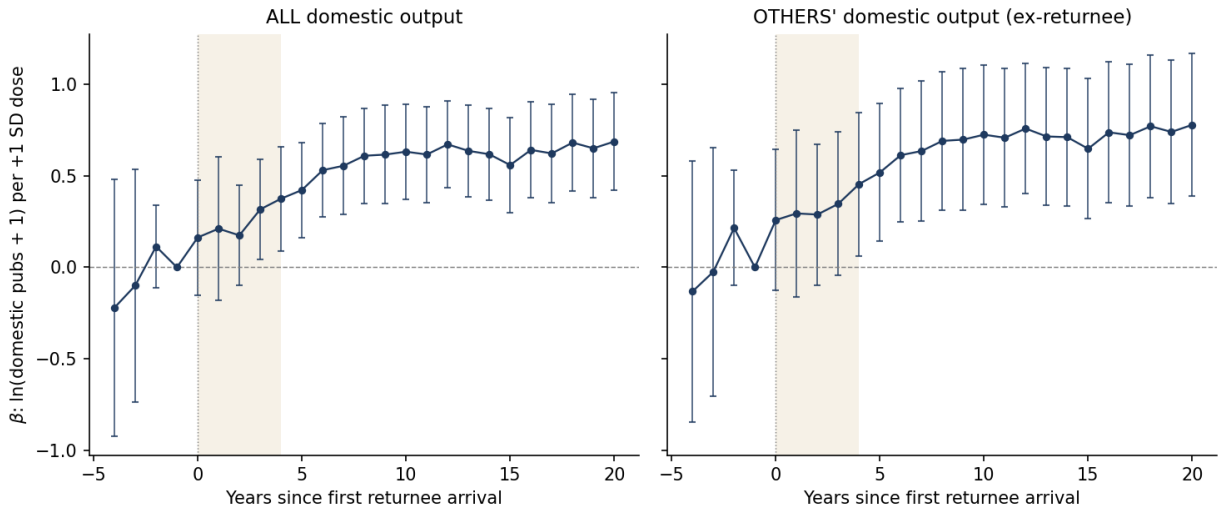


Figure 10: Domestic spillover: more returnees, more output by others

Notes: Coefficients $\hat{\beta}_k$ from equation (2) on the full CNKI corpus. The outcome is ln(1 + domestic publications) in a subfield-year: all output (left) and output excluding every paper with a returnee author (right). Dose is the standardized number of returnees whose modal subfield is j; one standard deviation is about 1.4 returnees. Each arrival cohort is compared with never-treated subfields, with subfield and year fixed effects within each stack. Event time is measured from the arrival of the subfield's first returnee, and the reference period is $\tau = -1$. The shaded band marks event years 0–4. Bars are 95 percent confidence intervals from standard errors clustered by subfield.

6.2 More Frontier-Proximate Output

The right test of capacity building is not whether the *average* domestic paper moved closer to the frontier. A field-creating returnee does not change papers that would have been written anyway; the question is whether new frontier-proximate work came to exist. We therefore count others' papers that lie close to a frontier. Figure 11 shows two such counts: papers near the *global* frontier, and papers near the research agenda that returnees themselves carried back from the United States. Over the first decade, one standard deviation of dose is associated with about 29 percent more papers near the global frontier ($\hat{\beta} = 0.258$), roughly 16 additional papers over the decade in the average treated subfield. The count near the returnee agenda rises by about 8 percent ($\hat{\beta} = 0.081$), roughly 6 additional papers. The global-frontier count has flat leads. The returnee-agenda count has negative leads at $\tau = -4$

and −3, so we treat its post-arrival rise as suggestive only. Returnees seeded frontier activity rather than raising averages: the mean alignment of others' papers is essentially unchanged (Appendix A.4).

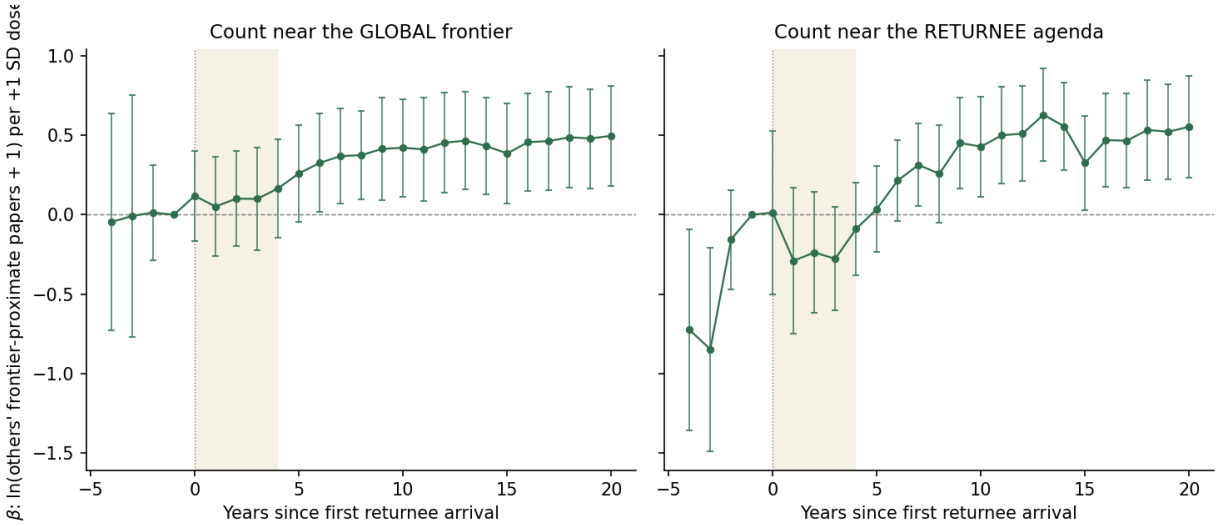


Figure 11: Domestic spillover: more frontier-proximate output by others

Notes: Same design as Figure 10. The outcome is $\ln(1 + \text{count})$ of non-returnee papers in a subfield-year whose embedding lies close to the global (OpenAlex) breakthrough frontier (left) or to the returnees' own research-agenda centroid (right). Bars are 95 percent confidence intervals from standard errors clustered by subfield.

6.3 Through What Channel? Returnees Became the Faculty and the Elite

Our biographical data record where returnees worked and what positions they held, which points to the channel for these spillovers (Figure 12). Among the 3,822 returnees with post-return employment recorded, 58 percent held university posts, 30 percent worked in other research institutes, 22 percent at the Chinese Academy of Sciences, and 21 percent in government ministries, which ran the defence research of the period. These categories overlap because careers span several posts. Leadership is coded from free-text career histories, so we report two definitions: 26 percent held some leadership post, and 7 percent headed an institute, university, or bureau outright, excluding deputies. Returnees also formed China's scientific elite. Of the 23 "Two Bombs, One Satellite" honourees who led the nuclear and space programs, 17 are in our data and 12 of those were elected to the Academy, against a base rate of 6.3 percent among all returnees. The Academy's elections themselves show the period's disruption: no elections were held between 1958 and 1979, and the 130 returnees elected in

1980 were recognized for work done across that gap. Returnees in strategic fields are also less documented on both margins. Their employment is recorded for 74 percent, against 88 percent elsewhere, and more of them have no matched publication. Publication-based measures therefore understate precisely the classified mission work these fields produced.

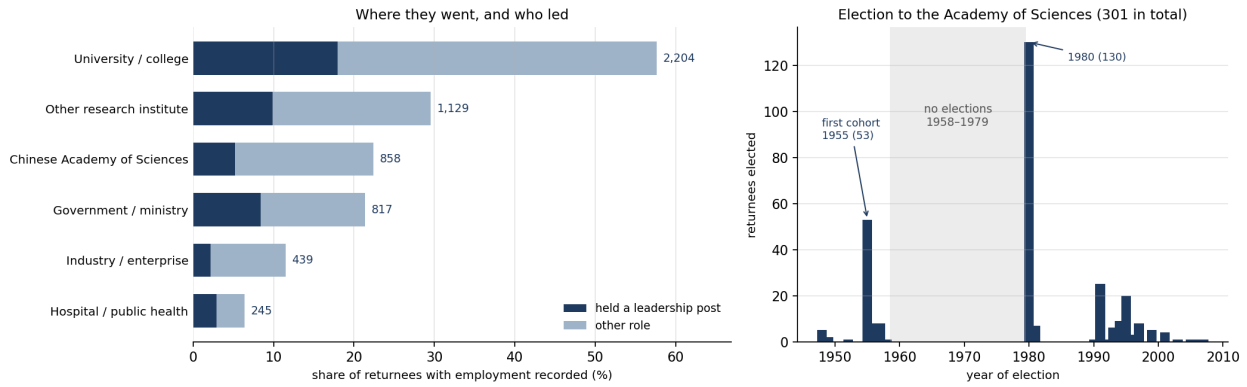


Figure 12: Where returnees worked and who led

Notes: Left: share of returnees with post-return employment recorded who ever held a post in each sector (categories overlap), split by whether they held a leadership post (broad keyword definition). Counts of returnees are shown at the end of each bar. Right: returnees elected to the Chinese Academy of Sciences by year of election (301 in total); no elections were held in 1958–1979. Source: authors' biographical returnee database.

6.4 Strategic Fields: Convergence Toward the Returnee Frontier

The catching-up is concentrated in the fields the state designated as strategic. Figure 13 tracks, for the strategic subfields prioritized in the 1956 Twelve-Year Plan, the share of domestic papers close to the returnee research frontier. The comparison runs from the early postwar years (1949–1954) to the early-1960s peak (1958–1964). The share rises in every strategic subfield, most strongly in electronic materials, process chemistry, metals and alloys, and nuclear engineering. The returnee agenda was a reachable target for the domestic system in a way that the global frontier was not. Figure 14 shows the same convergence in volume: over time, Chinese-language output grows into the regions of the semantic map where the global frontier is densest in the strategic fields.

7 The Welfare Cost of Redirection

The reduced-form results establish that return redirected scientific output from the international to the domestic literature without lowering its quality, and that returnees seeded

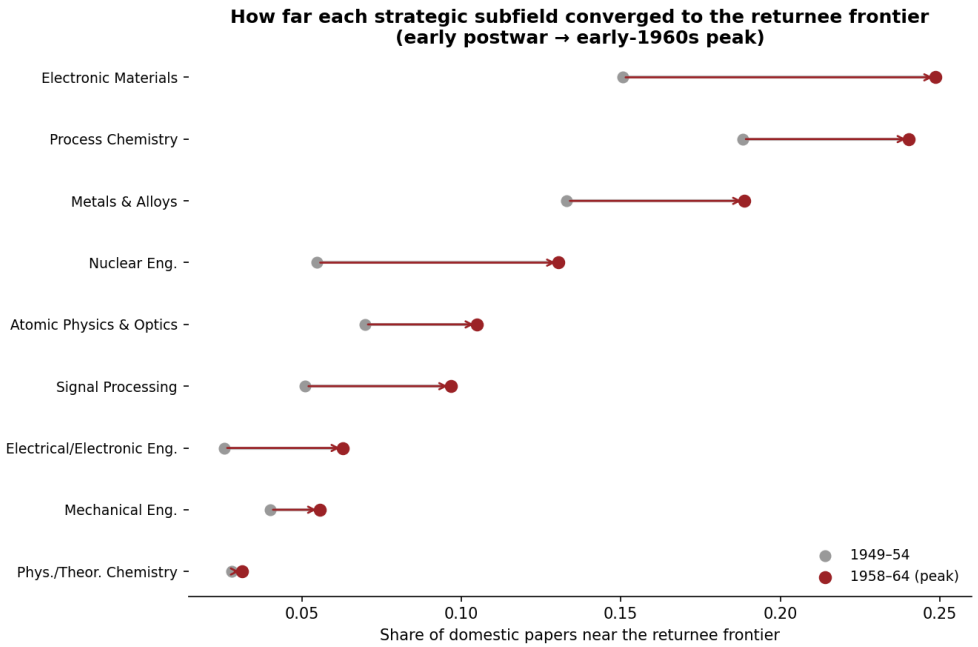


Figure 13: Convergence of strategic subfields to the returnee research frontier

Notes: For each strategic subfield, the share of domestic (CNKI) papers whose embedding lies close to the research-agenda centroid of returnees in that subfield, in 1949–1954 (grey) and at the 1958–1964 peak (red). Source: authors' analysis of the full CNKI corpus.

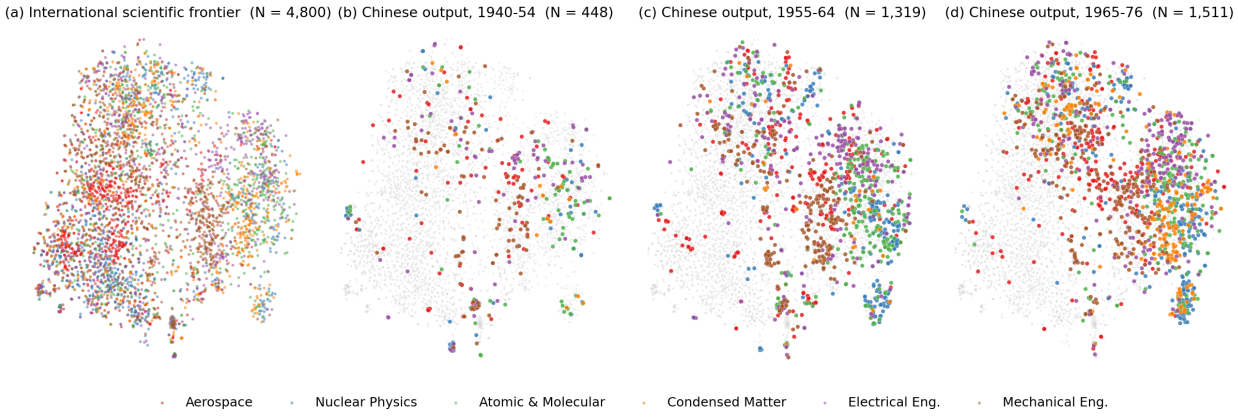


Figure 14: Volume of scientific output in the semantic space of strategic subfields

Notes: Two-dimensional t-SNE projection of LaBSE paper embeddings restricted to the strategic subfields. The first panel shows the global English-language frontier (top-cited OpenAlex papers); the remaining panels show the hexagonal density of Chinese-language CNKI publications in successive periods. Color intensity reflects the local density of papers.

output among the domestic scientists around them. We now ask what these movements were worth in aggregate. When a scientist is reallocated from the U.S. research system to the Chinese one, is world science better off, worse off, or merely rearranged? This accounting does not depend on the individual-level design of Section 4: it prices the *observed* allocation of returnees across fields and requires no counterfactual productivity.

What is a field worth? We build on the innovation-network model of Liu and Ma (2026). Fields are nodes in an innovation network Ω whose (i, j) entry ω_{ij} is the share of field i's citations that go to field j. A field's social value equals its own output plus the discounted chain of spillovers it sets off across the network:

$$\gamma' = \frac{\rho}{\rho + \lambda}\, \beta' \Big(I - \frac{\Omega}{1 + \rho/\lambda} \Big)^{-1}, \tag{3}$$

where β is each field's share of global research output. The inverse matrix is a spillover multiplier: it credits a field for the downstream work it enables, so central fields are worth more than their raw output. We set the discount rate $\rho = 0.05$ and the knowledge-obsolescence rate $\lambda = 0.17$. We construct Ω from a scan of the pre-1979 OpenAlex snapshot, which yields 38.7 million cross-field citation edges among 244 subfields.

Two countries. Following the open-economy extension of Liu and Ma (2026), each country keeps a share x^c of a field's spillovers at home, its *self-sufficiency*, and leaks the rest to the world frontier. We stack Ω by country into a $2K \times 2K$ network $\hat{\Omega}$ and solve for country-specific values $\hat{\gamma}_{\mathrm{CN},f}$ and $\hat{\gamma}_{\mathrm{US},f}$. A returnee is a unit of research reallocated from a (U.S., field) node to a (China, field) node:

$$\text{China gain} = \tfrac{\psi\lambda}{\rho^2}\, s_f\, \hat{\gamma}_{\mathrm{CN},f}, \qquad \text{U.S. loss} = \tfrac{\psi\lambda}{\rho^2}\, s_f\, \hat{\gamma}_{\mathrm{US},f}, \qquad \text{net}_{\text{world}} = \text{gain} - \text{loss},$$

where s_f is the returnees' output share in field f and $\psi = 0.06$ is a scale parameter that cancels in the ratio of gain to loss. We set $x_{\mathrm{US}} = 0.70$ following Liu and Ma (2026). We set $x_{\mathrm{CN}} \in [0.10, 0.45]$ field by field according to how close Chinese-language output in that field sits to the global frontier. The accounting thus needs four ingredients: the network Ω and its centrality structure, each field's output share β, where the returnees went (s_f), and how much each country keeps (x^c). This static accounting complements the dynamic two-country migration-growth model of Prato (2025), in which return migrants transmit knowledge across borders; we leave the full dynamic transition to future work.

The ingredients. Figure 15 shows the network: almost every field cites the same handful of hub fields, and this concentration is what makes the spillover multiplier bite. Innovation centrality is highly skewed (Figure 16). The hubs are molecular biology, atomic and molecular physics, and organic and materials chemistry. The most central field warrants 3.8, 7.1, and 17.5 times the research effort of the 5th-, 20th-, and 50th-ranked subfields, so *which* field a scientist works in dominates their social value. Returnees were concentrated in the central, high-value fields rather than the periphery: across subfields, the returnee output share correlates at +0.56 with centrality (Figure 17). A large amount of social value therefore moved. Finally, China's self-sufficiency lies far below the U.S. benchmark in every field (Figure 18): China could absorb only a fraction of each returnee's spillovers, and the rest leaked to a frontier it could not yet use.

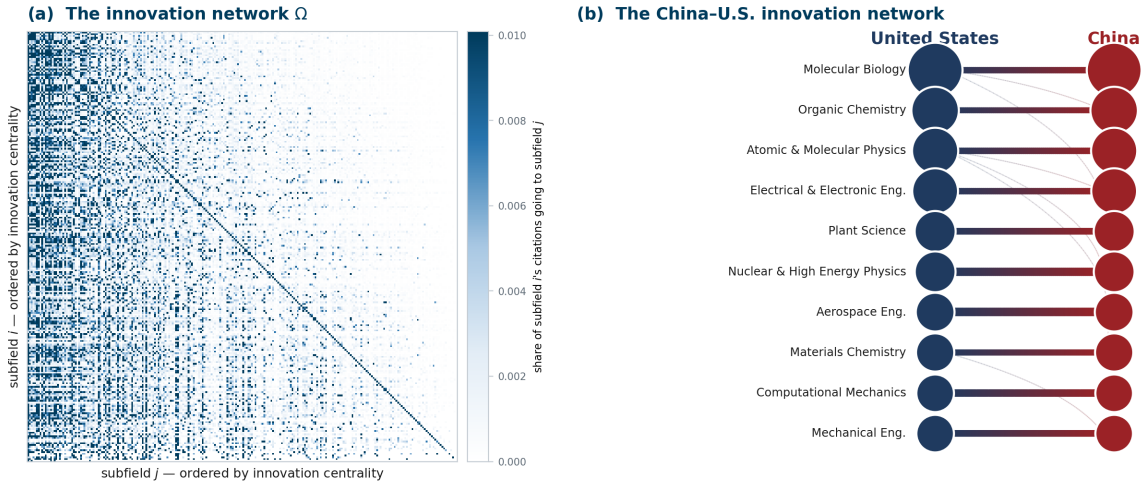
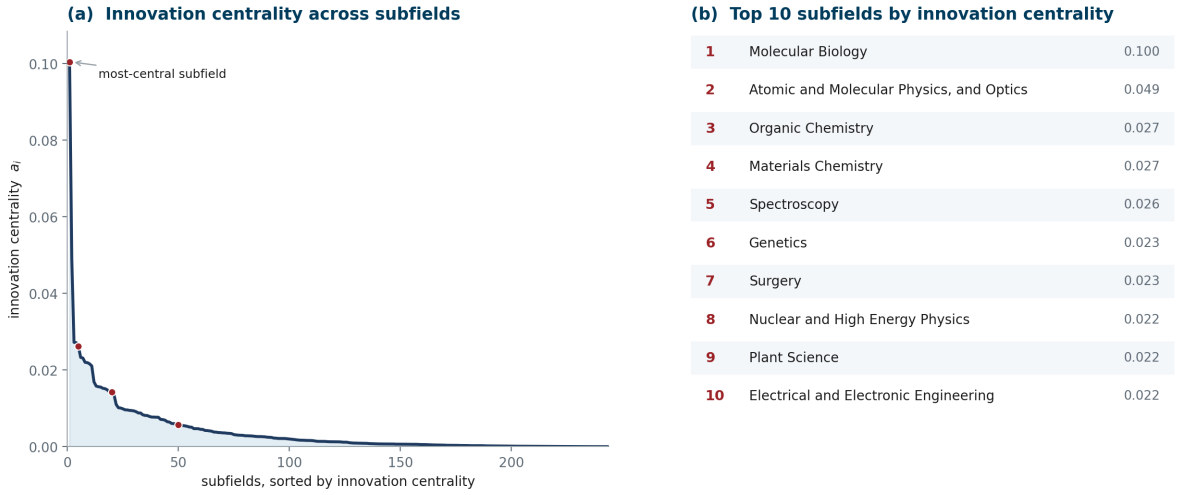


Figure 15: The innovation network and cross-country knowledge flow

Notes: (a) The cross-field citation network Ω built from 38.7 million pre-1979 citation edges among 244 subfields, with fields ordered by innovation centrality; darker cells denote larger flows. (b) The ten largest science and engineering fields in the two-country network: circle size is the field's share of research output, and bar width is China's reliance on the U.S. frontier, $1 - x_{\mathrm{CN}}$ (a thicker bar means China kept less). Source: authors' analysis of OpenAlex and CNKI.

Result. China's gain offsets only a quarter to a third of the U.S. loss (Figure 19). Per returnee, China gains +0.015 and the United States loses −0.045 in the relative units of the model, a net world loss of about −0.030. The U.S. loss exceeds China's gain in every field in which returnees were concentrated (Table 8). Scientific value was destroyed in transit, not merely transferred. Because the levels depend on the scale parameter, we base our conclusions on the ratio, which does not: it took about three scientists arriving in China to replace the social value of one leaving the United States.



(b) Top 10 subfields by innovation centrality

1	Molecular Biology	0.100
2	Atomic and Molecular Physics, and Optics	0.049
3	Organic Chemistry	0.027
4	Materials Chemistry	0.027
5	Spectroscopy	0.026
6	Genetics	0.023
7	Surgery	0.023
8	Nuclear and High Energy Physics	0.022
9	Plant Science	0.022
10	Electrical and Electronic Engineering	0.022

Figure 16: Innovation centrality across fields

Notes: Innovation centrality is the dominant left eigenvector of Ω. (a) The distribution across subfields. (b) The most central subfields. Source: authors' analysis of the pre-1979 OpenAlex citation network.

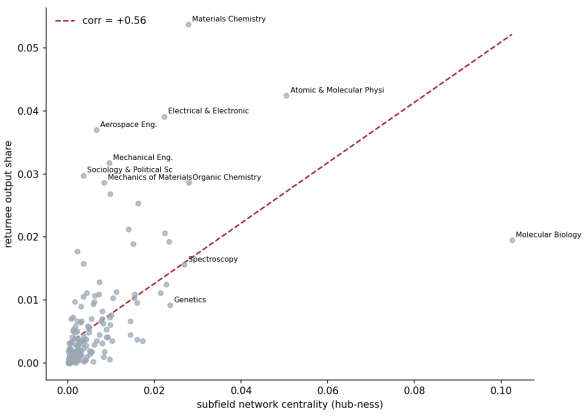


Figure 17: Returnees and centrality

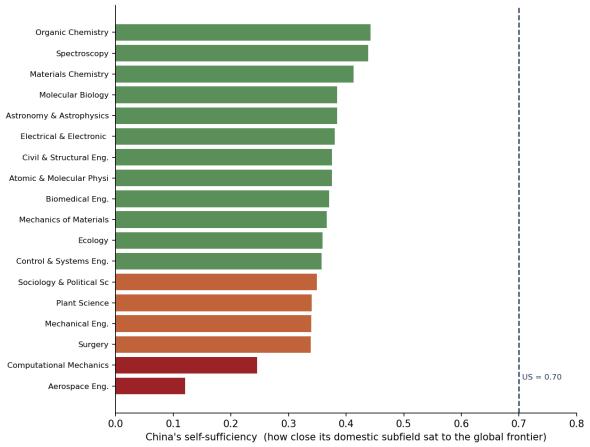


Figure 18: China's self-sufficiency

Notes: Figure 17: each point is a subfield; the horizontal axis is innovation centrality and the vertical axis is the returnee output share s_f (correlation +0.56). Figure 18: China's self-sufficiency x_{CN} by field, the share of a field's spillovers China keeps at home. It lies below the U.S. benchmark of 0.70 in every field. Source: authors' analysis.

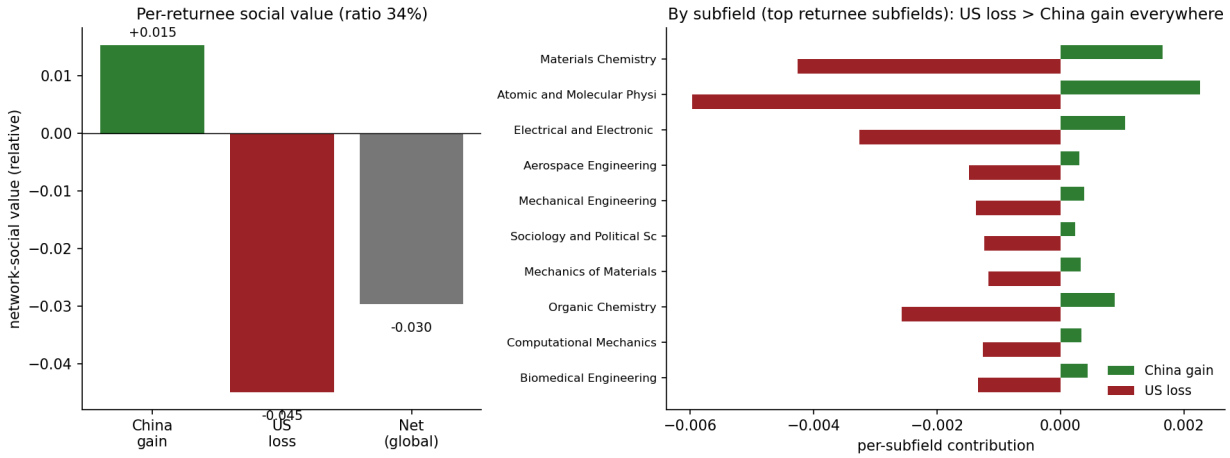


Figure 19: China's gain versus the U.S. loss per returnee

Notes: Left: per-returnee social value in the relative units of the model: China's gain, the U.S. loss, and the net world effect. The ratio of gain to loss is 0.34. Right: per-subfield contributions for the subfields in which returnees were most concentrated. Computed on the stacked (country × field) citation network. Source: authors' analysis.

Table 8: What a returnee was worth, by field: China's gain versus the U.S. loss

Field	Returnee share s_f	China self-sufficiency x_{CN}	China gain	U.S.
Materials Chemistry	0.054	0.41	0.0017	0.0
Atomic and Molecular Physics, and Optics	0.042	0.38	0.0023	0.0
Electrical and Electronic Engineering	0.039	0.38	0.0010	0.0
Aerospace Engineering	0.037	0.12	0.0003	0.0
Mechanical Engineering	0.032	0.34	0.0004	0.0
Sociology and Political Science	0.030	0.35	0.0002	0.0
Mechanics of Materials	0.029	0.37	0.0003	0.0
Organic Chemistry	0.029	0.44	0.0009	0.0

Notes: Per-returnee social value, in the relative units of the model, for the eight subfields with the largest returnee c shares. China self-sufficiency x_{CN} is the share of a field's spillovers China keeps at home; the U.S. benchmark is 0.70. V are computed on the stacked (country × field) network of Liu and Ma (2026). Source: authors' analysis.

Robustness. The ratio of gain to loss is 0.34 on the citation network and between 0.26 and 0.28 on three alternative innovation networks built from text rather than citations (Table 9). Letting China use its *own* network, estimated from the CNKI text, lowers the ratio in every case to 0.22–0.28. The shared-network baseline is therefore the conservative end. This is not a density artefact: thinning the U.S. network to the same density moves the ratio in the opposite direction, and randomly rewiring China's network recovers only part of the gap, so the placement of links matters.

The natural objection is that frontier knowledge is worth more where it is scarce, which is exactly the case of building a field from nothing. We scale China's value in field f by $1 + \pi\,\mathrm{gap}_f$, where gap_f is one minus the closeness of China's field f to the global frontier and π is the intensity of the scarcity premium. A threefold premium still leaves China's gain at half the U.S. loss (the ratio rises from 0.33 to 0.50). Breaking even requires a premium of roughly nine times on the scarcest field ($\pi^\star \approx 8.1$), well outside any plausible range (Figure 20).

Table 9: Ratio of China's gain to the U.S. loss across innovation networks

China block	Citation	Embedding	Text-VAR (abstracts)	Text-VAR (all)
Ω shared (baseline)	0.341	0.268	0.276	0.263
Ω^{CN} from CNKI text	0.282	0.221	0.228	0.217
U.S. network thinned to same density	0.352	0.279	0.293	0.281
Ω^{CN} rewired at random	0.307	0.260	0.268	0.257

Notes: Each cell is the ratio of China's per-returnee gain to the U.S. per-returnee loss. Columns vary how the innovation network is built: from citations, from embedding similarity, or from a vector autoregression on text (using abstracts only, or all text). Rows vary the China block of the stacked network: the same network as the U.S. (baseline); a network estimated from CNKI text; the baseline with the U.S. network thinned to the density of the CNKI network; and the CNKI network with links rewired at random. Source: authors' analysis.

What China gained. The gain is real even though it is smaller than the loss. Measuring dependence from text, because mid-century Chinese journals carry no reference lists, a subfield with one standard deviation more returnees relied about 0.02 less on U.S. work, against a base of 0.28. U.S. reliance on Chinese work does not move, and the sign holds under every weighting we try.

Interpretation and caveats. We interpret these magnitudes as illustrative rather than as welfare estimates in a national-accounting sense. They come from calibrating the model of Liu and Ma (2026) to the networks and embeddings we construct, and we base our conclusions on the sign and the ratio, both of which are robust to the modeling choices we vary. The exercise counts only knowledge spillovers visible in publications. The institutions returnees

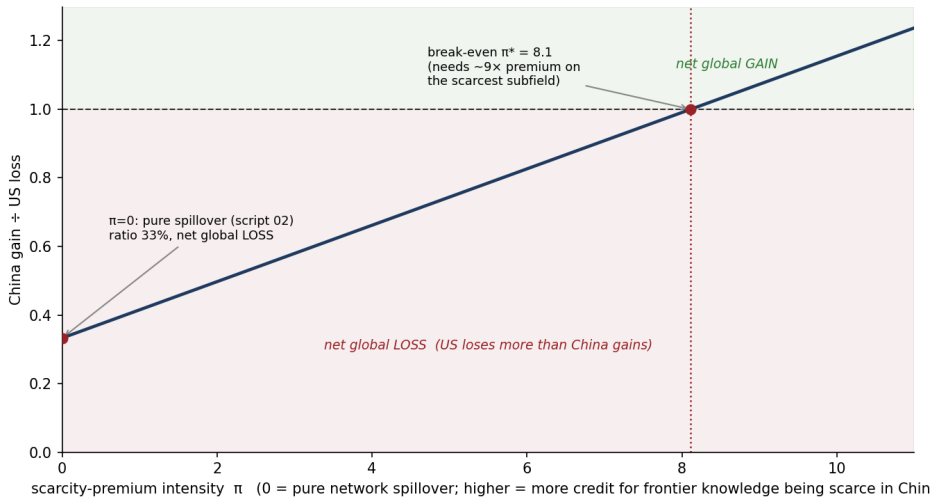


Figure 20: The net world loss survives a scarcity premium

Notes: Ratio of China's gain to the U.S. loss when China's value in field f is scaled by $1+\pi\,\text{gap}_f$, with $\text{gap}_f \in [0,1]$ equal to one minus the closeness of China's field f to the global frontier. The ratio stays below one, a net world loss, until $\pi^\star \approx 8.1$. Source: authors' analysis.

built, the students they trained, and classified mission work appear in no bibliographic record, so the China-side gain is, if anything, understated.

8 Conclusion

The Great Return Wave of 1946–1956 brought well over a thousand U.S.-trained Chinese scientists back to a country in the middle of building a planned research system. We reconstruct the population of U.S.-trained Chinese scholars, 4,684 returnees and 9,478 stayers, and follow their publications across the international (OpenAlex) and domestic (CNKI) literatures through 1976. We measure the quality and direction of each paper in a single cross-lingual semantic space. Observing the same person in both literatures lets us separate output that moved to another literature from output that was lost, which single-corpus analyses cannot do.

Three findings stand out. First, return reallocated publishing channels rather than reducing total output. Relative to matched stayers, returnees published about 9 percent less in OpenAlex and 10 percent more in CNKI, and their total publications were unchanged. Second, this switch came at no measurable cost to quality. When each scholar's English- and Chinese-language output is pooled and scored against the same global reference, neither importance nor frontier alignment moves after return. What changed was the direction of research, which moved toward applied, mission-oriented work. Third, returnees reshaped

the system around them. A subfield receiving one standard deviation more returnees, about 1.4 people, saw roughly 70 percent more output by *other* domestic scientists over the next decade, including more work near the global frontier. The strategic subfields named in the 1956 Twelve-Year Plan converged toward the returnees' own research frontier. Returnees staffed the universities and led the institutes, and they account for 17 of the 23 scientists honoured for China's nuclear and space programs.

Finally, pricing the reallocation through the cross-field citation network turns these reduced-form patterns into a welfare statement. China's gain offset only a quarter to a third of what the United States lost, because a frontier scientist generates more knowledge spillovers embedded in a dense, self-sufficient research system than transplanted into a frontier-dependent one. Forced relocation therefore did not simply transfer scientific value across the Cold War divide; it destroyed some of it in transit. This divergence cost, robust to a generous scarcity premium, is the aggregate counterpart to the field-level catching-up we document: China gained, the United States lost more, and world science was left poorer.

These results speak to a longstanding question in the economic history of socialist science: whether centrally planned research systems can absorb foreign frontier knowledge, and where in the field-space the absorption occurs. Soviet-era emulation strategies described by Lardy (1995) and the science-policy template the early PRC inherited from them (Wang, 2014) predicted convergence in disciplines tied to industrial and military priorities, and convergence stalls at the periphery. Our subfield estimates match that prediction: the catching-up signal loads on engineering and applied physics, not on basic physics or the humanities. The pattern is consistent with a broader literature documenting how political institutions shape the direction, rather than only the level, of innovation (Acemoglu et al., 2024; Moscona, 2024; Moser et al., 2025), and complements work tracing the long-run effects of mass cross-border knowledge transfers in market economies (Moser et al., 2014; Giorcelli, 2019). Returnees in our setting were not a marginal input that the planner could deploy uniformly; they were the binding constraint on absorptive capacity in the priority subfields, and the data trace exactly that pattern. The lesson extends beyond this setting. When a rising power draws frontier scientists into a planned research system, what it gains, and what the losing system forgoes, is governed by the strategic fields those scientists seed and their position in the global knowledge network, not by their number: a strategically concentrated cadre built the atomic bomb and the missile program, capabilities that are themselves a source of international power, complementing the economic interdependence emphasized by Liu and Yang (2026). Forced relocations of this kind reorganize the international scientific order, and, as our accounting shows, typically at a net loss to the world.

The empirical strategy used here is portable, and three extensions are natural. The

first is to follow the 1978 reopening cohort, where the same individuals are observed in a market environment and the channel-substitution mechanism should weaken if the planner's preferences were what concentrated catch-up in engineering. The second is to apply the dual-frontier measure to other socialist science systems with bilingual bibliographies (the post-1949 USSR, the GDR, post-1949 Cuba), where the predicted heterogeneity across subfields differs by planning regime. The third is to map the coauthor network that returnees built around themselves. The cross-corpus panel records every domestic collaborator on every CNKI paper and every international collaborator on every OpenAlex paper, so each returnee appears simultaneously as the seed of a Chinese coauthor cluster and as a persisting node in the U.S. network they trained in. Tracing how those two networks evolve, whether the domestic cluster widens as it absorbs the returnee, whether the international tie persists or atrophies under closure, and whether bridging scientists drive the subfield-level catching-up we document, is the natural next step, and the embedding-based outcomes used here generalize directly to coauthor-pair similarity. Whether returnees pull a domestic system toward the frontier, and where in the field-space they do so, appears to depend less on the scientists themselves than on the institutional environment that decides where to put them and on the networks that form around them once they arrive.

References

Acemoglu, Daron, David Y. Yang, and Jie Zhou, "Power and the Direction of Research: Evidence from China's Academia," 2024. Draft manuscript.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen, "Who Becomes an Inventor in America? The Importance of Exposure to Innovation," *The Quarterly Journal of Economics*, 2019, *134* (2), 647–713.

Bello, Piera, Alessandra Casarico, and Debora Nozza, "Research Similarity and Women in Academia," *The Economic Journal*, 2025, p. ueaf113.

Bianchi, Nicola and Michela Giorcelli, "Scientific education and innovation: from technical diplomas to university STEM degrees," *Journal of the European Economic Association*, 2020, *18* (5), 2608–2646.

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

Chang, Iris, *Thread of the Silkworm*, New York: Basic Books, 1995.

China Institute in America, *A Survey of Chinese Students in American Universities and Colleges in the Past One Hundred Years: In Commemoration of the One Hundredth Anniversary of the Graduation of the First Chinese from an American University, Yung Wing, B.A., Yale 1854*, New York: China Institute in America, 1954. Preliminary report. Jointly sponsored by the National Tsing Hua University Research Fellowship Fund and the China Institute in America.

Docquier, Frédéric and Hillel Rapoport, "Globalization, brain drain, and development," *Journal of economic literature*, 2012, *50* (3), 681–730.

Giorcelli, Michela, "The long-term effects of management and technology transfers," *American Economic Review*, 2019, *109* (1), 121–152.

_ **and Bo Li**, "Technology transfer and early industrial development: evidence from the Sino-Soviet alliance," *Review of Economic Studies*, 2026.

Hager, Sebastian, Carlo Schwarz, and Fabian Waldinger, "Measuring science: Performance metrics and the allocation of talent," *American Economic Review*, 2024, *114* (12), 4052–4090.

Han, Yelong, "An untold story: American policy toward Chinese students in the United States, 1949—1955," *Journal of American-East Asian Relations*, 1993, pp. 77–99.

Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow, "The Allocation of Talent and U.S. Economic Growth," *Econometrica*, 2019, *87* (5), 1439–1474.

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

Jia, Ruixue, Huihua Nie, and Wei Xiao, "Power and publications in Chinese academia," *Journal of Comparative Economics*, 2019, *47* (4), 792–805.

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

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy, "Measuring Technological Innovation over the Long Run," *American Economic Review: Insights*, 2021, *3* (3), 303–320.

Kerr, Sari Pekkala, William Kerr, Çağlar Özden, and Christopher Parsons, "Global talent flows," *Journal of Economic Perspectives*, 2016, *30* (4), 83–106.

Koffi, Marlene, "Innovative Ideas and Gender (In) Equality," *American Economic Review*, 2025, *115* (7), 2207–2236.

Kung, James Kai-sing and Alina Wang, "Overseas students, ideology, and the fall of dynastic China," *Journal of Economic Growth*, 2025, pp. 1–40.

Kung, S. W., *Chinese in American Life: Some Aspects of Their History, Status, Problems, and Contributions*, Westport, CT: Greenwood Press, 1973. Reprint of the edition originally published by University of Washington Press, Seattle.

Lardy, Nicholas, *Emulating the Soviet Model, 1949–1957*, Cambridge University Press, 1995.

Li, Tao, *A Historical Record of Overseas Education in China since 1949*, Beijing: Higher Education Press, 2000.

Liu, Ernest and David Y. Yang, "International Power," 2026. Working paper.

_ **and Song Ma**, "Innovation Networks and R&D Allocation," *American Economic Review*, 2026. Conditionally accepted.

Mokyr, Joel, "The intellectual origins of modern economic growth," *The Journal of Economic History*, 2005, *65* (2), 285–351.

_ , "Long-term economic growth and the history of technology," in "Handbook of economic growth," Vol. 1, Elsevier, 2005, pp. 1113–1180.

Moscona, Jacob, "Flowers of Invention: Patent Protection and Productivity Growth in U.S. Agriculture," *Working paper*, 2024.

Moser, Petra, Alessandra Voena, and Fabian Waldinger, "German Jewish Émigrés and U.S. Invention," *American Economic Review*, 2014, *104* (10), 3222–3255.

_ **, Sahar Parsa, and Shmuel San**, "Immigration and Innovation: Lessons from the Quota Acts," 2025. Working paper.

Murphy, Kevin M., Andrei Shleifer, and Robert W. Vishny, "The Allocation of Talent: Implications for Growth," *The Quarterly Journal of Economics*, 1991, *106* (2), 503–530.

Prato, Marta, "The Global Race for Talent: Brain Drain, Knowledge Transfer, and Growth," *The Quarterly Journal of Economics*, 2025.

Romer, Paul M, "Endogenous technological change," *Journal of political Economy*, 1990, *98* (5, Part 2), S71–S102.

_ , "The origins of endogenous growth," *Journal of Economic perspectives*, 1994, *8* (1), 3–22.

Spilimbergo, Antonio, "Democracy and foreign education," *American economic review*, 2009, *99* (1), 528–543.

Squicciarini, Mara P and Nico Voigtländer, "Human capital and industrialization: Evidence from the age of enlightenment," *The Quarterly Journal of Economics*, 2015, *130* (4), 1825–1883.

The White House, "Proclamation on the Suspension of Entry as Nonimmigrants of Certain Students and Researchers from the People's Republic of China (Proclamation 10043)," 2020. May 29.

Waldinger, Fabian, "Quality matters: The expulsion of professors and the consequences for PhD student outcomes in Nazi Germany," *Journal of political economy*, 2010, *118* (4), 787–831.

_ , "Peer effects in science: Evidence from the dismissal of scientists in Nazi Germany," *The review of economic studies*, 2012, *79* (2), 838–861.

_ , "Bombs, brains, and science: The role of human and physical capital for the creation of scientific knowledge," *Review of Economics and Statistics*, 2016, *98* (5), 811–831.

Wang, Zuoyue, "The Cold War and the Reshaping of Transnational Science in China," in Naomi Oreskes and John Krige, eds., *Science and Technology in the Global Cold War*, Cambridge, MA: MIT Press, 2014, pp. 343–369.

Yuan, Tung-li, *A Guide to Doctoral Dissertations by Chinese Students in America, 1905–1960*, Washington, DC: Sino-American Cultural Society, 1961.

Zhang, Jun and Yueyun Mo, "The Historical Evolution of China's Overseas Students' Return Policies since the Founding of the People's Republic of China," *Party History Research and Teaching*, 2018, (5), 60–69.

Appendix A: Data Construction

A.1 Data Sources

Directories of Chinese Scholars in the United States

Across the early twentieth century, different organizations voluntarily took on the responsibility of creating a record of Chinese students at American universities. These organizations oftentimes used these directories for the two-pronged purpose of fostering community ties among Chinese students and amplify marketing efforts. As the responsibility for creating directories for Chinese students was informally passed from organization to organization, the methods of collection began to change as well. Due to the lack of a centralized database, statistics on Chinese students were often a collation of information gathered from different methods, including outreach to universities, personal interviews, and records from local student clubs. The prevalence of name changes and ambiguity around classifications of origin made it difficult to maintain consistent records over the years. Table 10 summarizes the directories we draw on; Figure 21 and Table 11 report the annual counts of recorded students.

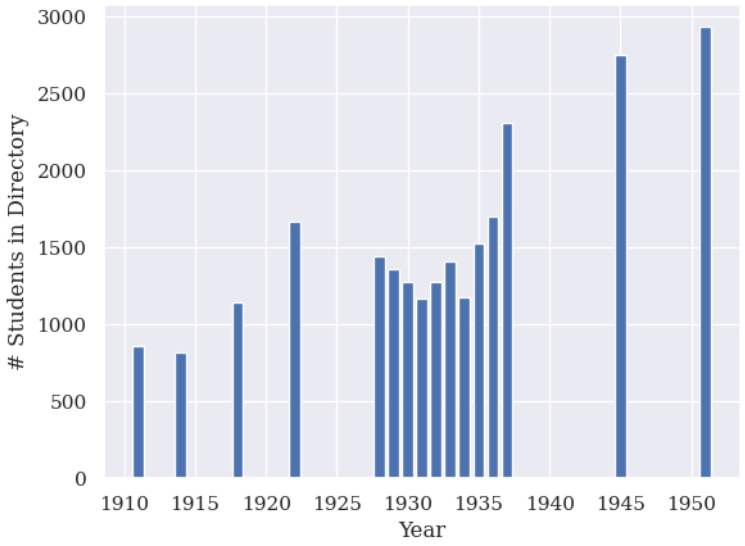


Figure 21: Annual Count of Students in Directory

Table 10: **Chinese Student Directories in the United States (1911–1952)**

Organization	**Years Covered**	**Variables Available**	**Data Collection Methods**
Chinese Students' Alliance in the U.S.A.	1911–12, 1914–15, 1921, 1929–33	Name, School, Major/Field of Study, Degree Pursued, Mailing Address	Compiled and edited by a handbook committee and local club officers, with support from alumni and professional societies.
Chinese Students' Christian Association	1935–40	Name, School, Major/Field of Study, Mailing Address	Aggregated registrar records, correspondence with individual students, collaboration with student groups, institutional catalogs, and bulletin mailings.
The China Institute in America	1941–45, 1949–52	Name, School, Degree Program, Graduation Year, Student Status (undergrad/grad), Major Field	Data obtained from schools (winter semester returns), Chinese student groups, and independent graduate surveys. Limitations noted: incomplete due to mid-year enrollment changes, name inconsistencies, and exclusions based on origin.

year	No. of Entries	Truncated Names
1911	861	387
1914	821	593
1929	1357	586
1930	1281	375
1931	1168	433
1932	1281	374
1933	1413	420
1935	1702	275
1937	2312	682
1945	2750	1078
1951	2934	292

Table 11: Count of Students by Directory

The first step of our data collection focused on the retrieval of all directories of Chinese students in the United States, which served as a time stamp of their activity abroad. We then constructed a dataset of all scholars by concatenating the directories.

Below are a few challenges that we experienced in trying to merge data across disparate sources:

- truncated or anglicized names (more common for first names), which created significant misalignment between the dataset of returnees and the dataset of Chinese scholars in the US
- heterogeneity in academic schema, making it difficult to group scholars under field of study
- unclear timelines across directories; although we could not find overlapping directories across organizations, we also did not find any formal documents corroborating our claim that each organization operated dis-jointly
- multi-year gaps in directories left huge holes in our data, making it difficult to corroborate timelines of certain scholars whose time abroad fell within one of these periods

Returnees

We compile a dataset of Chinese scholars who studied abroad and returned to China from four biographical compilations covering 1840–1963. The Comprehensive Survey of Scientific and Technical Personnel Studying Abroad and Returning to China, 1840–1949 profiles 3,500 scientific and technical personnel, documenting pre-departure domestic education, foreign

study details, and post-return employment; the data, originally structured in tabular form by field, were digitized and cleaned for misspellings and syntactical inconsistencies. The Scientific and Engineering Manpower in Communist China, 1949–1963 contains information on 1,200 prominent scientists and engineers, including birth year, educational background, specialization, political status, and Chinese Academy of Sciences (CAS) membership, compiled from a wide range of official and published materials. *Guilai*, the directory of students who returned from North America in the 1950s (ed. Wang Delu, 2023), covers 1,806 returnees, with detailed variables on education, migration dates, post-return careers, political roles, and notable life events, based on interviews, organizational records, personal communications, and archival materials. Finally, the 56-volume *Biographical Sketches of Chinese Science and Technology Experts* (China Association for Science and Technology, ed. Qian Sanqiang) devotes one volume to each discipline and adds field-stratified coverage of eminent scientists. For *Guilai*, we parsed the entries using GPT-4 API calls to classify variables related to time abroad and post-return activities. Table 12 summarizes the scope, variable coverage, and data collection methods for each source.

Table 12: **Primary Sources on Returnee Scholars to China (1840–1963)**

Source	**Time Period Covered**	**Variables Available**	**Data Collection Methods**
A Comprehensive Survey of Scientific and Technical Personnel Studying Abroad and Returning to China, 1840–1949	1840–1949	Name, Domestic School (before going abroad), Year of Departure, Country of Study Abroad, Foreign Institution, Academic Discipline, Major, Year of Return, Domestic Work Location (after return)	Profiles 3,500 scientific and technical personnel in mainland China that had studied abroad. Compiled from biographical books, publications, and archival materials on students who studied abroad, particularly those in science and engineering.
Cheng, Scientific and Engineering Manpower in Communist China, 1949–1963 (NSF 65-14)	1949–1963	Name, Year of Birth, Educational Background (highest degree held), Academic Position (as of 1964), Special Field (basic training), Social and Political Status (as of 1964), CAS Membership Indicator	Profiles 1,200 prominent scientists and engineers in mainland China. Data compiled from Chinese official documents and books, periodicals and newspapers, statistical yearbooks, scientific journals, leadership speeches, and translated materials.
Guilai: Directory of Students Who Returned from North America in the 1950s (ed. Wang Delu, 2023)	1950s	Name, Year of Birth, Year of Death, Domestic Education (Affiliation, Degree, Year, Location), Foreign Education (U.S. Affiliation, Degree, Year, Location), Date of Arrival in the U.S., Date of Return to China, Post-return Employment and Destination, Employment in China before Studying Abroad, Academic Discipline, Academic Specialization, Political Roles, CAS Academician Status, Whether Detained, Whether Died During	Profiles 1,806 returnees. Based on five types of sources: (1) Interviews, (2) information from acquaintances, (3) records from organizations, (4) public materials (e.g., memoirs), (5) archival letters/photos.

Stayers

We build the comparison group from two sources that cover the same U.S. training pipeline as the returnees. Table 13 summarizes them. We harmonize the two into a common person-level schema, deduplicate across them, and remove anyone who appears in the returnee database, which leaves 9,478 stayers.

Table 13: **Primary Sources on Stayers in the United States**

Source	**Time Period Covered**	**Variables Available**	**Data Collection Methods**
Yuan Tung-li, *A Guide to Doctoral Dissertations by Chinese Students in America, 1905–1960* (Sino-American Cultural Society, 1961)	1905–1960	Name, Dissertation Title, University, Degree, Year of Degree, Field	Compiled for libraries as a census of every doctorate earned by a Chinese student at a U.S. university; inclusion does not depend on later prominence.
National Register of Scientific and Technical Personnel (National Science Foundation)	1954–1970	Name, Year of Birth, Education, Field and Specialty, Employer and Affiliation	Federal registry that the National Science Foundation Act (§3(a)(8)) mandated for all U.S. scientists; eight survey waves with 126,000 to 321,000 respondents each.

Publication Data

To measure scholar productivity and characterize research specialization, we assemble a comprehensive publication panel by merging English- and Chinese-language bibliometric sources. Our construction proceeds in two steps: (i) for each scholar, we recover a broad candidate set of publications using multiple name variants; (ii) we remove false matches using a text-based classifier and structured metadata checks, yielding a clean scholar–year corpus of titles and abstracts.

Publications. For each scholar, we collect publications from OpenAlex and CNKI. From OpenAlex we retain `work_id`, title, abstract, publication year, venue (journal/proceedings),

document type, authorship roster (including `author_id` and author order), author affiliations (institution identifiers and countries), concepts/keywords (with hierarchical levels), `referenced_works` (outgoing citations), and `cited_by_count`. From CNKI we add Chinese-language works with titles, abstracts, venues, publication years, and author affiliations. Table 14 summarizes the core fields retained from OpenAlex.

Table 14: Core Publication Fields Retained

Entity	**Fields**
Work	id, title, abstract, year, venue, type, concepts (id, level), referenced_works, cited_by_count
Authorship	work_id, author_id, author_position, affiliation_id, affiliation_country
Author	author_id, canonical name, aliases, ORCID (if any)
Institution	inst_id, name, country, historical name variants

Person Lists and Name Variants. Both bibliographic pipelines read the same two person lists, so every publication is attached to one canonical person identifier. Each person carries a name in Chinese characters where the source records one (returnees; most stayers are recorded only in romanized form), in Hanyu Pinyin, and in Wade–Giles; we generate Wade–Giles for every person because U.S. journals of the 1940s–1960s used it. Person identifiers are issued from an append-only registry so that they are stable across rebuilds. Figure 22 shows the distribution of publication counts across matched scholars.

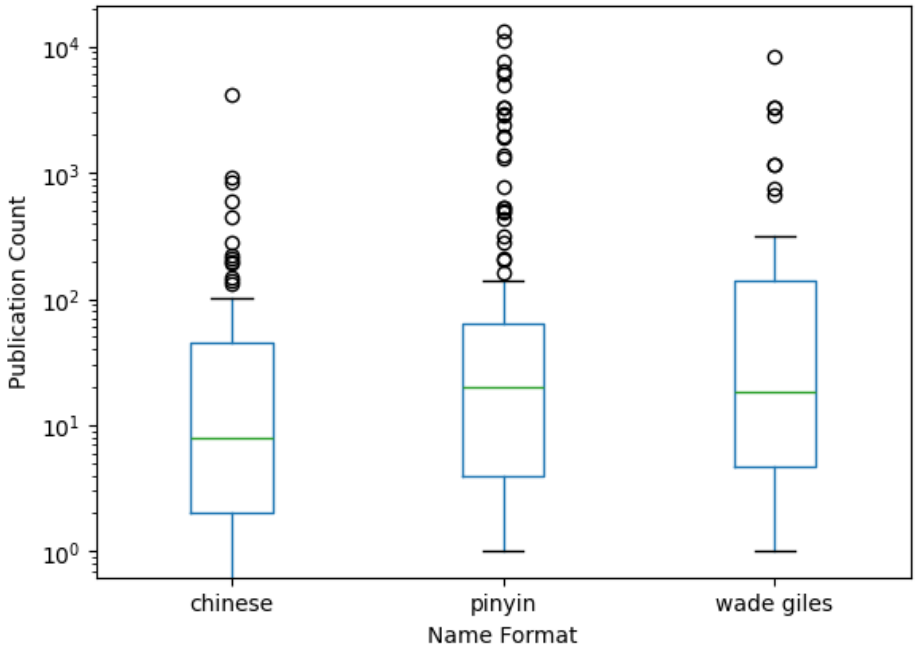


Figure 22: Publication Counts for Matched Scholars

OpenAlex Matching. We scan the full OpenAlex snapshot for works published up to 1979 with at least one Chinese-named author, identified from a list of about 470 romanized Chinese surnames in Pinyin, Wade–Giles, Cantonese, Hokkien, and anglicized forms. This step is deliberately high-recall (about 4.5 million authorships). For each authorship we look up candidate persons by canonical surname, score the romanized full name (threshold 85), require the publication year to be plausible, and check the affiliation against the period: a returnee's pre-return work should be affiliated in the United States and post-return work in China, and a stayer's in the United States throughout. We do not use OpenAlex author identifiers, whose pre-1979 disambiguation is unreliable, and we block initials-only matches.

CNKI Matching. Most CNKI authors are recorded in Chinese characters, so the primary path is an exact character match; a fuzzy Pinyin path handles the minority of romanized records. We keep journal articles published up to 1976.

Corroboration and Measured Precision. Common names collide, especially in romanized form. Every OpenAlex match, and every CNKI match on a name of two or fewer characters, must therefore be corroborated by the scholar's field or institution. We did not choose these rules by judgment: we simulated alternative rules on a hand-labeled audit of 190 matches and adopted the rules with the best precision–recall trade-off. Under them,

measured precision is 92 percent in CNKI (recall 71 percent) and 50 percent in OpenAlex (recall 76 percent). Uncorroborated matches are retained with a flag rather than deleted, so every analysis can be run on corroborated matches (the default) or on all matches. A final identity-resolution pass flags person identifiers whose matched name variants share no common surname (over-merges), and applies hand-verified corrections, such as reassigning Qian Xuesen's Caltech papers to his returnee record.

Dissertations. We harmonize five dissertation source files, including Yuan (1961), into a universe of 4,626 dissertations and match both returnees and stayers to it with the same name machinery. A contradicted year rejects a match, and a global one-to-one assignment ensures that each dissertation is linked to at most one scholar. Because 86 percent of returnees have no pre-return publication, the dissertation title is often the only record of their pre-return research.

Citations. OpenAlex citations are the `cited_by_count` of each work in the snapshot; CNKI citations are CNKI's own citation counts. We sum them by scholar and publication year. The two are on different scales, which is why we analyze them in logs and separately by corpus.

Frontier-Measurement Sensitivity. We compute the embedding-based outcomes under three sentence-embedding models on the OpenAlex side (multilingual LaBSE, SciBERT, and a Chinese-language SBERT) and two on the CNKI side (LaBSE and the Chinese SBERT, since SciBERT is English-only), and under three breakthrough definitions used to build each subfield-year frontier (top-decile importance, top-decile citations, and top-journal publication). LaBSE, which places Chinese and English text on one scale, is the headline model.

Reference Corpora and the Cross-Lingual Field Crosswalk. Three reference corpora support the field-level and frontier measures. (i) The *full CNKI corpus* of 564,632 Chinese-language papers is the population in which we measure spillovers to non-returnee domestic scientists (Section 6.1). (ii) A *stratified CNKI reference sample* of 106,000 papers, stratified by FORD subject category and year, defines the domestic frontier. (iii) A *stratified OpenAlex reference sample* of 3.5 million papers across 244 subfields defines the global frontier. Separately, a scan of `referenced_works` in the pre-1979 OpenAlex snapshot yields 38.7 million cross-field citation edges, from which we build the innovation network used in Section 7. We embed every title and abstract in all corpora with multilingual LaBSE. Each CNKI paper is then assigned to the OpenAlex subfield whose reference-paper centroid is

nearest in cosine similarity. This assignment is the CNKI-to-OpenAlex crosswalk, and it lets us score Chinese- and English-language papers against a common global frontier. Figures 3 and 4 show that the assignment recovers the structure of fields separately within each corpus and when the two corpora are pooled.

A.2 Data Availability and Public Release

We will make the dataset available as an open research resource. It will be deposited with a persistent DOI in a public repository (openICPSR or Harvard Dataverse) under a CC-BY license, and the construction code will be posted on GitHub. The release will contain:

1. **Person-level roster.** All returnees and stayers with stable person identifiers; source provenance for each record; training (institution, degree, field, years); arrival and return dates; the harmonized field; post-return employment (both the original text and the coded sector and leadership flags); and political-event and honours indicators. It comes with a codebook.

2. **Archival exit-permit records.** OCR transcriptions of the declassified State Department, Immigration and Naturalization Service, FBI, and CIA files we collected, including an FBI roster of 78 students barred from leaving.

3. **Historical-name table.** For each person, the name in Chinese characters (where recorded), Hanyu Pinyin and Wade–Giles, together with the surname canonicalization rules used for matching.

4. **Scholar-to-publication linkage.** OpenAlex work identifiers for each matched scholar, with the match-confidence bucket and the corroboration flag. For CNKI, we release bibliographic identifiers and metadata (title, venue, year) but not full text or abstracts, because CNKI's terms of use restrict redistribution.

5. **Derived measures.** Paper-level Kelly importance and dual-frontier alignment scores; the CNKI-to-OpenAlex subfield crosswalk; subfield-by-year panels of output and returnee exposure; and the LLM agenda scores, with the prompts and model versions used.

6. **Code.** Scripts covering the full pipeline, from digitized sources to the derived measures, including the embedding and frontier computations. Users can therefore apply the dual-frontier measure to other bilingual corpora.

We will release the package alongside publication of the paper.

A.3 U.S. Exit Restrictions on Chinese Students during the Early Cold War

This subsection documents the policy rationale and implementation of U.S. exit restrictions on Chinese students during the early Cold War. Drawing on declassified CIA, State Department, and FBI materials, we show that U.S. authorities treated the overseas training of Chinese students, particularly in technical and scientific fields, as a potential national security risk and acted deliberately to restrict their return to the People's Republic of China (PRC).

Fields Deemed Strategically Sensitive. Across intelligence and diplomatic documents, a consistent set of disciplines was identified as posing heightened risk if U.S.-trained Chinese students returned to China. These fields were closely linked to weapons development, military-industrial capacity, and advanced infrastructure.

Table 15: Fields Targeted by U.S. Security Concerns

Category	**Fields Mentioned or Inferred**
Nuclear Science	Nuclear physics, nuclear engineering
Weapons and Missile Technology	Aeronautics, aerospace engineering, guidance systems
Advanced Engineering	Mechanical, electrical, civil engineering
Electronics and Computing	Electronics, communications, radar, control systems
Chemical and Biological Sciences	Chemistry, chemical engineering, biomedicine
Metallurgy and Materials	Metallurgy, materials science
Mathematics and Physics	Mathematics, applied physics
Geophysics	Geophysics (e.g., military mapping and exploration)

Sources: CIA, *Science and Technology in Communist China* (1966); National Intelligence Estimate 13-2-62, *Chinese Communist Advanced Weapons Capabilities* (1962); State Department and FBI memoranda.

Government Actions. Declassified State Department and FBI documents from the early 1950s show that U.S. authorities deliberately detained or blocked the departure of technically trained Chinese students. These actions were motivated by three main considerations:

- preventing the transfer of U.S.-acquired technical knowledge to China's weapons and industrial programs;

- using detained students as leverage in negotiations over American prisoners held in China;
- responding to intelligence assessments warning that returnees trained in sensitive disciplines posed a national security risk.

Surveillance and Enforcement Criteria. Students were subject to heightened monitoring and scrutiny if they:

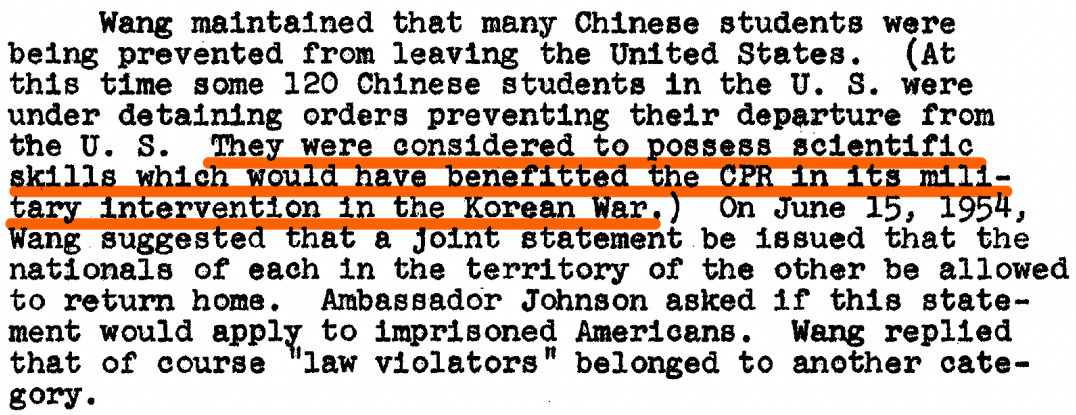
- pursued advanced training in strategically sensitive fields;
- were suspected of affiliation with the Chinese Communist Party or participation in pro-PRC student organizations;
- expressed intentions or made observable preparations to return to mainland China;
- accessed or conducted research linked to U.S. military or classified projects.

Strategic Motivation. The overarching objective of these restrictions was to limit the repatriation of U.S.-acquired technical expertise that could accelerate China's military and industrial modernization. U.S. authorities explicitly linked student return to concerns over nuclear weapons development, missile technology, and long-run strategic competition during the early Cold War.

Chinese Activity

The September 15, 1954, Summary mentioned a letter being sent to President Eisenhower by a group of Chinese students in this country urging permission to return to China. The students are being restrained temporarily from leaving the United States pursuant to a Presidential directive dated January 17, 1953, aimed at preventing the departure to Communist-dominated countries of alien students possessing scientific and technical training deemed of possible military value to a potential enemy. This topic also has figured prominently in news releases bearing on relations between the United States and Communist China, particularly in dispatches of the past few days pertaining to the recent visit of Dag Hammarskjold, Secretary General of the United Nations, with CHOU En-lai, Premier of Communist China.

Figure 23: FBI memorandum expressing concern over Chinese students "possessing scientific and technical training deemed of possible military value to a potential enemy." J. Edgar Hoover to Harold E. Stassen, January 19, 1955.



Wang maintained that many Chinese students were being prevented from leaving the United States. (At this time some 120 Chinese students in the U. S. were under detaining orders preventing their departure from the U. S. They were considered to possess scientific skills which would have benefitted the CPR in its military intervention in the Korean War.) On June 15, 1954, Wang suggested that a joint statement be issued that the nationals of each in the territory of the other be allowed to return home. Ambassador Johnson asked if this statement would apply to imprisoned Americans. Wang replied that of course "law violators" belonged to another category.

Figure 24: State Department analysis of negotiations with Chinese representatives, June 1954–September 1955, concerning the detention and release of Chinese nationals in the United States.

Implications for Identification. This historical record underscores that U.S. exit restrictions were (i) sharply targeted toward specific technical fields, (ii) implemented through administrative discretion rather than individual behavior, and (iii) motivated by geopolitical considerations external to scholars' productivity or preferences. These features support our use of permit and denial decisions as plausibly exogenous variation in return migration and motivate field-based heterogeneity analyses in the main text.

A.4 Subfield Mean-Alignment Spillover

As an additional test of domestic spillover, we ask whether returnee arrival raised the *mean* frontier alignment of a subfield's publications, complementing the extensive-margin and strategic-convergence results in Section 6. We implement equation (4) at the subfield-year level: Returnee_j is an indicator for whether subfield j ever received a 1940–1957 returnee, and Post_{jt} equals one in years on or after the first 1940–1957 returnee's arrival in j. Subfields without a 1940–1957 arrival enter as never-treated controls. The pre-1940 returnee stock $\text{cumret}_{j,t}^{\text{pre-1940}}$ enters as a covariate X_{jt}:

$$y_{jt} = \alpha_j + \vartheta_t + \delta\,\text{Post}_{jt} \times \text{Returnee}_j + \gamma\,\text{cumret}_{j,t}^{\text{pre-1940}} + \varepsilon_{jt}, \tag{4}$$

where y_{jt} is the mean frontier alignment of subfield-j publications in year t, computed separately for OpenAlex and CNKI publications against the OpenAlex global frontier. Standard errors are clustered at the subfield level. As a continuous specification, we replace the indicator with the cumulative count of 1940–1957 returnees who had published in j by year t,

denoted $\text{cumret}_{j,t}^{1940\text{-}57}$:

$$y_{jt} = \alpha_j + \lambda_t + \gamma\, \text{cumret}_{j,t}^{\text{pre-1940}} + \delta\, \text{cumret}_{j,t}^{1940\text{-}57} + \varepsilon_{jt}, \tag{5}$$

with both regressors and outcome standardized so that δ is interpretable per standard deviation of cumulative-returnee exposure.

Table 16 reports the estimates. The binary specification of (4) produces a positive coefficient of 0.019 on the alignment of OpenAlex publications to the global frontier (column 1, $t = 1.40$) and a smaller, indistinguishable-from-zero coefficient of 0.011 on the alignment of CNKI publications (column 2). The continuous specification of (5) delivers similar magnitudes: a one-standard-deviation increase in cumulative 1940–1957 returnees raises OpenAlex-side alignment by 0.020 standard deviations (column 3) but does not move the CNKI-side alignment (column 4). The point estimates are economically modest at the aggregate panel level; restricting to subfields whose first 1940–1957 returnee actually arrived during the cohort window and to the event-time window $[-10, 15]$ around that arrival, the OpenAlex coefficient on $\text{cumret}_{j,t}^{1940\text{-}57}$ rises to 0.069 ($t = 1.68$, $p = 0.09$). The mean-alignment spillover therefore operates primarily through the international-publishing arm of returnees' own work and the work of their domestic collaborators who appear in OpenAlex; the field-creating spillover onto other domestic scientists shows up instead on the extensive margin, as documented in the main text.

Table 16: Subfield-level spillover: 1940–1957 returnees and frontier alignment

	Eq. (4): Binary Post×Returnee		Eq. (5): Continuous (per SD)	
	(1)	(2)	(3)	(4)
	OA vs OA fr.	CNKI vs OA fr.	OA vs OA fr.	CNKI vs OA fr.
Treatment	0.019	0.011	0.020	−0.005
	(0.014)	(0.021)	(0.014)	(0.020)
Pre-1940 cumret control	Yes	Yes	Yes	Yes
Subfield fixed effects	Yes	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes	Yes
Subfields	211	176	211	176
Observations	5,427	2,427	5,427	2,427

Notes: Subfield-year regressions. Columns 1–2 estimate equation (4): $\text{Returnee}_j = 1$ if subfield j ever received a 1940–1957 returnee, $\text{Post}_{jt} = 1$ in years on or after the first such arrival. Subfields with no 1940–1957 arrival serve as never-treated controls. Columns 3–4 estimate equation (5) with the cumulative count of 1940–1957 returnees in subfield j by year t, standardized; the outcome is also standardized so that the coefficient is per standard deviation. Columns 1 and 3 take the mean alignment of OpenAlex publications to the OpenAlex global frontier as the outcome; columns 2 and 4 take the mean alignment of CNKI publications to the same frontier. All specifications include subfield and year fixed effects and partial out the pre-1940 returnee stock. Standard errors clustered at the subfield level (columns 1–2) or analytic OLS (columns 3–4) in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.