Beyond the Target: U.S. Sanctions, Inventor Mobility, and the Unintended Knowledge Diffusion*

Di Lu
Shenzhen University

Frederick Yanzhe Wang
Peking University

Qi Wu
Peking University

Mingzhi Xu
Peking University

May 31, 2026

Abstract

This study examines the impact of U.S. Entity List sanctions on the mobility and productivity of individual inventors within Chinese firms. Using a staggered difference-in-differences (DiD) approach, we find that sanctions significantly increase the likelihood of job switching among inventors at targeted firms while no change on productivity measured by patents per innovator. However, the sanctioned firms suffer from significant brain drain alongside deteriorating innovation and financial performance. We also document a substantial knowledge diffusion effect. Departing inventors catalyze innovation at their new landing firms. This allows these firms to develop comparative advantages in technological domains where the inventors' former employers excelled. Such spillovers account for roughly 27 percent of the total decline in patenting caused by sanctions. Our findings highlight the unintended consequences of sanctions on talent redistribution, suggesting that human capital mobility plays a critical role in shaping the ultimate effectiveness of geopolitical trade restrictions.

Keywords: Inventor, Job Mobility, Economic Sanction, Entity List
JEL Classification: J61, F51, O31

*We are grateful to Sabrina Howell, Uta Schönberg, Yanru Han, Bo Hu, Ruochen Dai, Binkai Chen, Ting Chen, and seminar participants at the 13th AIEA-NBER Conference on Innovation and Entrepreneurship, Fudan University, the Central University of Finance and Economics, and Shenzhen University for their valuable comments and suggestions. We thank Shuyu Zhang for her excellent research assistance. Mingzhi Xu gratefully acknowledges financial support from the National Natural Science Foundation of China (Grant Nos. 72322007 and 72503230). Corresponding authors: Di Lu, Shenzhen Audencia Financial Technology Institute, Shenzhen University (ludi@szu.edu.cn); Frederick Yanzhe Wang, Institute of New Structural Economics, Peking University (yzwang2022@nsd.pku.edu.cn); Qi Wu, Guanghua School of Management, Peking University (qiwu@gsm.pku.edu.cn); and Mingzhi Xu, Institute of New Structural Economics, Peking University (mingzhixu@nsd.pku.edu.cn). All authors contributed equally. Any remaining errors are our own.

1 Introduction

In an era of intensifying geopolitical rivalry, economic restrictions have become a standard instrument of statecraft, used to slow technological upgrading in rival economies, shape industrial development, and restrict access to strategically important inputs and capabilities.[1] Yet whether such measures achieve their intended objectives is rarely clear. Their effects depend not only on policy design, but also on the adjustments they trigger across firms, workers, and markets (Morgan et al., 2023a; Hausmann et al., 2024).[2] Such adjustments can offset or redirect the initial disruption, so immediate effects are a poor guide to whether these measures actually weaken the capabilities they are meant to constrain. We provide new evidence that this form of containment may be incomplete.

Our setting is the post-2018 shift in U.S. policy from broad trade pressure to targeted technology sanctions against China (Felbermayr et al., 2025).[3] Tariffs imposed under Section 301 marked the initial escalation, but the more consequential development was the expanded use of the Entity List under the Export Administration Regulations amid growing technological decoupling. Once listed, firms could no longer obtain U.S.-origin technology without a specific license, typically reviewed under a presumption of denial. Recent policy research interprets this regime as part of a broader strategy of technological containment directed at sectors near the technological frontier (Branstetter, 2024; Allen, 2024). By curtailing access to frontier foreign technology in industries where innovation depends heavily on such inputs, this episode provides a natural setting in which to identify an adjustment margin that may blunt the intended effect of technology sanctions.

High-tech firms have been a primary target of these measures, given their central role in innovation. Our evidence indicates that sanctions materially weaken targeted Chinese firms by eroding profitability, slowing earnings growth, tightening liquidity, and increasing reliance on debt finance. This broader deterioration in firm performance is accompanied by

[1] Beyond the recent U.S.–China sanctions, other policy episodes underscore the range of objectives economic restrictions can serve (Lu et al., 2026). Following Russia's invasion of Ukraine, the European Union imposed export restrictions on dual-use and advanced-technology goods to weaken Russia's industrial and technological base (European Commission, 2025b,a). Indonesia's nickel ore export ban and domestic processing requirements were intended to shift production downstream and foster higher-value manufacturing (International Energy Agency, 2024; European Commission, 2022). The European Union has also restricted exports to Iran of components used in missile and unmanned aerial vehicle systems, thereby constraining access to inputs critical for military production (Council of the European Union, 2024).

[2] Morgan et al. (2023a) stress that sanctions are difficult to evaluate in terms of stated objectives, in part because success is often multidimensional and not readily observable. Hausmann et al. (2024) further show that their effects depend critically on how restrictions are designed, including which goods are targeted and how easily they can be substituted.

[3] Before 2018, additions to the U.S. Entity List were relatively infrequent and focused largely on traditional military-industrial and proliferation-related actors. Beginning in 2018, however, the Bureau of Industry and Security (BIS) sharply accelerated its use of the list, adding 668 Chinese entities between 2018 and 2022.

a decline in innovative activity, consistent with Cao et al. (2024). But part of that decline may reflect reallocation rather than pure loss. By narrowing research opportunities and lowering the returns to remaining at the firm, sanctions may induce inventors with portable expertise and strong outside options to leave. Innovation losses at the sanctioned firm may therefore capture not only weaker inventive activity, but also inventor exit. Since inventor mobility is a well-established channel of knowledge diffusion (Johnson et al., 2023; Reinmuth and Rockall, 2025), such reallocation can carry tacit knowledge beyond the sanctioned firm, and the resulting spillovers may mitigate the sanctions' intended technological effect.

To examine this adjustment systematically, we construct a novel inventor-level dataset from the universe of CNIPA patent records between 2014 and 2022, covering 1,642,454 inventors, 6,045,809 patents, and 4,939,429 inventor-year observations. The key advantage of these data is that they allow us to observe innovative-labor reallocation directly, rather than infer sanction effects solely from firm-level patenting outcomes. We implement an inventor-disambiguation algorithm that consolidates duplicate identities and links individuals across patent filings despite inconsistencies in names, affiliations, and reporting practices. This linkage recovers inventor employment histories and separates actual cross-firm mobility from changes in firms' patenting activity. Figure 1 provides a first look at this reallocation. Before the 2018 U.S. sanctions, moves from sanctioned Chinese firms to non-sanctioned firms were not the dominant flow; afterward, mobility shifts sharply in that direction. In the full data, sanctioned-to-non-sanctioned flows become roughly twice as large as flows in the opposite direction, and this asymmetry persists through 2022, the final year of our sample.

Using these inventor-level mobility histories, we treat the staggered listing of an inventor's current employer on the U.S. Entity List as a shock to the inventor's research environment. Since the sanction is imposed on the firm rather than the inventor, the timing of exposure is plausibly unrelated to the inventor's own productivity or mobility propensity. We estimate its effects in a staggered difference-in-differences framework with a rich set of fixed effects, so identification is based on within-inventor changes around the sanctioning event, net of broader differences across inventors, firms, and time. The results show a large and robust increase in mobility: exposure to sanctions raises the probability of leaving by 8.1 percentage points, roughly two-thirds of the pre-sanction mean of 12 percent. We find no long-term impact on inventors' productivity measured by patent applications per innovator.

The mobility response is stronger where the inventor-firm match appears fragile, and where sanctions most directly disrupt the firm's access to U.S. technology. Inventors with weaker collaborative ties and lower measured patent quality are especially likely to leave, so are inventors at firms with thinner knowledge bases and smaller co-inventor teams. The response is also substantially larger in firms more dependent on U.S.-sourced intermediate

inputs, underscoring that these are not generic firm shocks but sanctions transmitted through exposure to U.S. technology and supply chains. The reallocation is also broader than local job switching, as sanctions significantly increase moves across counties and cities, pointing to a national reshuffling of innovative talents.

Because departing inventors remain productive after moving, sanctions may shift not only research labor but also the technological knowledge embodied in it. We test this diffusion channel by asking whether recipient firms subsequently gain strength in the IPC classes tied to incoming inventors' prior expertise. Since such knowledge is often tacit and not directly observed, we adapt the revealed comparative advantage approach of Bahar and Rapoport (2018a), Bahar et al. (2024), and Lin et al. (2024) to patent portfolios. Specifically, we construct revealed comparative technological advantage (RTA) at the firm-IPC level and examine whether recipient firms' post-move patenting shifts toward technology classes in which the incoming inventor, and the sanctioned firm from which the inventor came, had previously specialized. If knowledge diffuses through inventor mobility, recipient firms should exhibit a disproportionate rise in RTA in those IPC classes relative to other fields and comparable firms. This test allows us to infer tacit knowledge transfer from changes in the direction of subsequent patenting, and to assess whether U.S. sanctions redirect technological capabilities across firms rather than simply reducing innovation at targeted firms.

We find recipient firms become more specialized in the fields where movers' previous employers held a comparative advantage, especially when those employers were sanctioned firms. In the full sample, inventor inflows raise recipient firms' IPC-level RTA by 2.6%; for inflows from sanctioned firms, the effect rises to 11.4%, more than four times as large. This diffusion is also visible in rising patent-portfolio similarity. As recipient firms absorb knowledge from sanctioned-origin movers, their innovation portfolios shift toward the similar technologies in which origin firms had been strong, increasing technological proximity to firms already competing in those fields. Thus, knowledge diffusion and tougher competition are two sides of the same process.

Finally, beyond diffusion across firms, inventor inflows also generate spillovers within recipient firms. We separate the movers' own output from the response of incumbent inventors and find that the arrival of new inventors raises patenting by non-moving colleagues. Each additional 100 incoming inventors increases incumbent patent filings by 6.1%, while the presence of any inflow raises them by 2.3%. The effect is much stronger when the movers come from sanctioned firms their arrival is associated with a 12.5% increase in incumbent patenting, more than five times the effect of any inventor inflow. Therefore, sanctions-induced mobility not only shifts recipient firms toward the technological fields of origin firms, but also raises the output of incumbent teams inside those firms. This amplifies the unintended

diffusion effects of sanctions.

The main findings are robust to a range of checks addressing the main threats to identification and measurement. To address nonrandom selection into sanctions, we re-estimate the analysis on a propensity-score-matched sample, pairing inventors at sanctioned firms with those at unsanctioned firms based on both individual and firm-level characteristics. To ensure that the control group is not contaminated by strategic responses of close rivals, we exclude technologically proximate peer firms. To address potential measurement error and sample selection bias arising from the intermittent nature of patent records, we use a Heckman selection model that corrects for selection in patent filing. To reduce timing error in patent-based mobility histories, we impose stricter windows for inferred job switches. The mobility effect remains positive and stable throughout. We further absorb city-year and county-year shocks, ruling out local innovation policies or regional confounders as the source of the estimates. Finally, alternative count-data models and balanced panels for inventor-level patenting leave the conclusion unchanged: sanctions increase inventor mobility, but do not reduce the post-move productivity of the inventors who leave.

We summarize the aggregate innovation effects with a simple accounting exercise. The origin-firm loss is based on our firm-level estimate of the sanction effect on patenting, capturing the total decline within targeted firms without decomposing direct inventor loss from broader organizational disruption. We then compare this loss with the gains generated when displaced inventors move to non-sanctioned firms. Quantitatively, sanctions reduce patenting at targeted firms by 875.88 patents, while inventor relocation generates 235.98 additional patents at recipient firms. These indirect gains save 26.94% of the patenting that would otherwise have been lost. The offset is strongest in thicker innovation markets, suggesting that agglomeration helps convert displaced human capital into local innovation resilience.

Our paper contributes to three strands of literature. First, it speaks to the literature on labor mobility as a channel of knowledge diffusion. A central insight of this work is that tacit knowledge is embodied in people and therefore moves with them across firms, regions, and countries.[4] Prior studies show that immigrants (Kerr, 2018; Bahar and Rapoport, 2018a), return migrants (Bahar et al., 2022), migrant inventors (Miguelez and Temgoua, 2020; Bahar et al., 2020a), farmers (Bazzi et al., 2016; Pellegrina and Sotelo, 2025), and political leaders (Lin et al., 2024) transmit knowledge that shapes specialization, productivity, and technology adoption across locations.[5] We extend this literature to sanctions-induced inventor

[4]Knowledge also diffuses through international trade (Cai et al., 2022; Akcigit and Melitz, 2022) and FDI (Lind and Ramondo, 2023; Abebe et al., 2022; Ahn et al., 2026).

[5]At a more disaggregated level, Hausmann and Neffke (2019) provide plant-level evidence from Germany that labor mobility facilitates technology diffusion across plants. Personal mobility can also affect other cross-border economic linkages, including FDI (Mayda et al., 2022) and demand for products from migrants'

reallocation across domestic firms. Patent evidence shows that displaced inventors transfer origin-firm capabilities to their new employers, making sanctions an unintended conduit of domestic technological diffusion.

Second, the paper contributes to the literature on sanctions, geoeconomic policy, and innovation. A growing body of work studies the economic effects of sanctions (for reviews, see Felbermayr and Sandkamp, 2020; Felbermayr et al., 2025; Hufbauer and Jung, 2020; Morgan et al., 2023b), while related research examines the U.S.-China trade war and its effects on firm outcomes, tariff exposure, technology decoupling, scientific research, and indirect trade linkages (Jiao et al., 2024; Kong et al., 2024; Huang et al., 2023; Benguria et al., 2022; Fajgelbaum and Khandelwal, 2022; Li et al., 2022; Han et al., 2024; Jia et al., 2024; Garred and Yuan, 2025). Existing evidence on domestic innovation responses remains mixed and is largely measured at the firm level.[6] We shift the focus from aggregate firm responses to the inventor-level adjustment margins behind them. This distinction is important because a decline in innovation at targeted firms need not imply a proportional loss of inventive capacity. We show that Entity List sanctions reduce innovation within targeted firms, but also induce inventors to move, remain productive, and transmit technological capabilities to other domestic firms.

Third, our paper relates to work on the institutional determinants of inventor mobility and firm innovation.[7] Existing studies emphasize tax policy as a driver of inventor location choices (Akcigit et al., 2016, 2019; Moretti and Wilson, 2017) and wartime displacement as a source of scientific reallocation and knowledge diffusion (Iaria et al., 2018; Moser et al., 2024). We study U.S. Entity List sanctions as a distinct institutional shock that redirects inventors from targeted firms to domestic recipients. Linking mobility histories to patent outcomes, we show that this reallocation has consequences beyond the movers themselves: recipient firms move toward the technologies of the origin firms, and incumbent inventors at those firms increase their own patenting, consistent with productivity gains from knowledge spillovers (Hsieh et al., 2019; Prato, 2024) and with evidence that labor mobility can enhance firm innovation (Braunerhjelm et al., 2020; Kaiser et al., 2015; Jin et al., 2025). Our contribution is to identify an unintended reallocation channel through which sanctions weaken targeted firms while redirecting inventive talent and embodied technological knowledge to domestic

home countries (Gould, 1994; Dunlevy and Hutchinson, 1999; Bailey et al., 2021).

[6]For example, Kong et al. (2024) find that tariff shocks reduce innovation among Chinese firms by limiting market expansion and innovation inputs, whereas Li et al. (2022) find that Chinese firms facing higher tariff exposure increase R&D expenditures. For U.S. firms, Autor et al. (2013) show that reduced trade-policy uncertainty lowers patenting, while Chen et al. (2021) find that lower tariff uncertainty with respect to Chinese imports improves innovation performance.

[7]See He and Tian (2020) for a review of the broader literature linking institutional environments to innovation incentives, firm-level innovative activity, and the allocation of inventive resources.

recipient firms.

The paper proceeds as follows. Section 2 describes the institutional background, introduces data and the construction of key variables, and presents motivating evidence. Section 3 estimates the effects of Entity List sanctions on inventor mobility and post-move innovation, heterogeneity analysis, robustness checks. We further explores the innovation consequences of mobility, including knowledge diffusion, technology convergence and within-firm spillovers. Section 4 calculates the aggregate offsets. Section 5 concludes.

2 Sanctions, Data, and Motivating Evidence

We begin with the institutional setting and data, tracing the escalation of U.S.-China technology sanctions and showing how patent, sanction, and firm-level records allow us to measure inventor mobility and subsequent shifts in innovation.

2.1 Recent U.S.-China Tech Tensions and Entity List Sanctions

The U.S.-China trade relationship has undergone a structural transformation since China's accession to the World Trade Organization (WTO), evolving from a period of deepening integration toward one defined by strategic decoupling. Initially, this friction was concentrated in the digital and intellectual property domains. Following its accession to the WTO, China gradually expanded market access in many sectors; however, it progressively restricted U.S.-based internet and cloud services – notably YouTube (2009), Facebook (2009), Twitter (2009), Google (2010), Instagram (2014), Telegram (2015), and WhatsApp (2017) – citing concerns over national security and domestic content control. Simultaneously, the U.S. raised persistent concerns about China's protection of intellectual property and treatment of foreign technology, specifically the practice of requiring U.S. firms to share source code or proprietary processes as a prerequisite for operating licenses (Qin, 2019). These concerns were formalized in the U.S. Trade Representative's 2018 Special 301 Report, which designated China as a "Priority Watch List" country for the 14th consecutive year and argued that policies associated with "Made in China 2025" could further intensify such practices.[8]

The bilateral relationship deteriorated sharply in 2018. In March of that year, the Trump administration launched tariff measures against China based on findings under Section 301 of the Trade Act, marking a broader escalation in trade and technology tensions. Beyond tariffs, the U.S. Department of Commerce began aggressively deploying the Entity List as

[8]The U.S. Trade Representative's 2018 Special 301 Report is an annual review that identified China as a major source of concern because of weak intellectual property protection and practices linked to forced technology transfer.

a tool of systemic technology containment. Once an entity is added to this list, the Export Administration Regulations (EAR) mandate that any transfer of U.S.-origin technology to that entity requires a specific license, which is typically reviewed under a "presumption of denial." The application of these sanctions shifted from pre-2018 traditional military non-proliferation to the post-2018 containment of dual-use technologies in high-tech industries. This post-2018 regime has targeted the "technological frontier" of the Chinese economy through the systematic inclusion of firms in sectors critical to the future of global supply chains, including semiconductors, artificial intelligence, quantum computing, and advanced manufacturing.

Landmark Enforcement Cases. High-profile enforcement actions during this period demonstrated the power of these tools to constrain targeted firms and disrupt global supply chains. In April 2018, the U.S. Department of Commerce imposed a denial order on ZTE, a major Chinese telecommunications equipment and smartphone producer, cutting off its access to U.S.-origin chips and software. The measure brought the firm close to suspension of operations and was lifted only after ZTE agreed to pay a $1.4 billion penalty and adopt extensive compliance measures. The U.S. soon broadened these restrictions: Fujian Jinhua was added in October 2018, followed by Huawei and over 70 affiliates in May 2019, along with several prominent Chinese universities and research institutes added throughout 2019 and 2020. These designations represent a severe idiosyncratic shock to the firms' production functions, as the cost of obtaining or substituting key technologies rises sharply if such substitution is even possible.

Since 2021, the scope and intensity of these controls have further tightened, focusing on systemic decoupling in advanced computing and semiconductor manufacturing. By leveraging the "Foreign Direct Product Rule," U.S. authorities effectively severed China's access to critical upstream inputs, such as electronic design automation (EDA) software and logic chips below the 16nm threshold. These rules impose stricter compliance requirements not only on the sanctioned entities but also on major international foundries like TSMC and Samsung that utilize U.S. equipment. By 2023, the extension of sanctions to supercomputers and semiconductor equipment institutionalized the Entity List as a persistent constraint on China's high-tech industries (Allen, 2022; Bureau of Industry and Security, 2022).

2.2 Data Description and Main Variables

To examine how sanctions shape innovation at the micro level, we link patent records to the U.S. Entity List and detailed firm-level data. The resulting comprehensive inventor–firm

dataset underlies our main measures and also offers motivating evidence on why inventor reallocation is a relevant margin of adjustment after sanctions.

Measuring Innovation and Inventor Mobility. Our primary innovation data come from inventor-level patent records from the China National Intellectual Property Administration (CNIPA) for 2014–2022. For each patent, we observe the patent number, title, assignee, inventor names, application date, International Patent Classification (IPC), and patent type. We supplement these records with patent-family and citation information from incoPat, a database widely used in innovation research (Kang et al., 2025).[9] Our analysis focuses on inventors affiliated with corporate assignees, which account for the patented technologies most relevant to the sanctions we study. Consistent with prior innovation research, such as He and Tian (2013) and Wang et al. (2021), we restrict attention to invention patents, which capture more substantive technological advances than utility-model or design patents; for simplicity, we refer to them as patents throughout the paper.[10] The resulting panel contains 1,642,454 inventors and 6,045,809 patents – accounting for 82.18% of total China's patents filed by domestic firm, corresponding to 4,939,429 inventor-year observations.

Patent records identify inventors and assignees, which typically correspond to employers. Following prior work (Melero et al., 2020; Liu et al., 2024), we infer inventor mobility from changes in an inventor's corporate assignee across applications. We retain inventors observed patenting in multiple years and restrict the sample to patents with a single corporate assignee, allowing each inventor-year to be assigned to an employer with minimal ambiguity. This restriction retains 92.74% of all firm-filed patents, suggesting limited sample loss. We code a job change when an inventor patents at different firms in different years, assigning the move to the first year in which the inventor appears at the new firm. For example, an inventor observed at firm A in 2017 and firm B in 2018 is recorded as moving from A to B in 2018. The resulting average annual mobility rate is about 12 percent, close to the 13–14 percent rate documented using U.S. patent and LinkedIn data (Ge et al., 2016).[11] Appendix A describes the cleaning procedure and mobility algorithm in detail.

This patent-based measure of mobility raises three measurement concerns. First, left truncation in inventors' patenting histories may lead to misclassification of moves near the

[9]The incoPat database covers more than 203 million patents from 173 countries, organizations, and regions. Its records are drawn from official and commercial intellectual property sources and include patent type, application date, applicants, applicant addresses, and citation information.

[10]China has three types of patents: invention, utility model, and design. Design patents protect aesthetic features and contain limited technological content, while utility-model patents generally reflect incremental or superficial improvements. Invention patents are therefore the most appropriate measure of technological innovation. See Wang et al. (2021) and Dai and Wang (2024).

[11]Liu et al. (2024) study inventor mobility among listed manufacturing firms in China over 2004–2012 and report an annual job-change rate of 9 percent.

beginning of the analysis period. We address this by using patent records dating back to 2000 to recover pre-sample patenting histories, while restricting the main analysis to 2014 onward. This allows us to measure prior experience and identify mobility more accurately. Second, common inventor names may confound the assignment of unique inventors. We therefore implement a two-step disambiguation procedure. We first trace inventors within the same county to identify intra-regional moves. For inventors who cannot be linked within a county, we match records with the same name and technology IPC sector across firms in different counties, thereby capturing inter-regional moves. We pool both sets of links in the analysis. Third, inventors who stop patenting after changing jobs may be incorrectly classified as stayers. This concern is mitigated by our requirement that inventors appear in the patent data at least twice.

Measuring Inventor Exposure to Sanctions. We construct sanction exposure using BIS Entity List additions from 2018 to 2022. After excluding entities registered in Hong Kong or Taiwan, financial firms, and firms under ST (i.e., special treatment) or PT (i.e., particular transfer) classifications, the list contains 304 mainland Chinese firms. We match these firms to patent assignees and identify 192 sanctioned entities in the patent data. These matched entities define the post-2018 sanctioned firms used in the analysis. A key challenge is that patent assignee names often differ from firm names in corporate databases and in the Entity List. We therefore harmonize names using a conservative fuzzy string-matching procedure, followed by manual verification.[12] We also obtain industry information for sanctioned firms from Bureau van Dijk's Orbis database.[13]

At the inventor level, we define sanction exposure as an indicator, Sanction_{it}, equal to one in the post-listing years after inventor i is employed by a firm added to the BIS Entity List, and zero otherwise.[14] This definition reflects the firm-level nature of Entity List restrictions: once a firm is listed, the restrictions alter the technological and organizational environment faced by all inventors employed there. Appendix Table A1 reports the annual distribution of Chinese firms added to the Entity List since 2018. The targeted firms are concentrated

[12]We implement approximate matching using Python's *RapidFuzz* package, based on Levenshtein edit distance. The algorithm measures similarity by the minimum number of single-character insertions, deletions, or substitutions required to transform one string into another (Levenshtein, 1966). This approach is well suited to Chinese firm-name matching and is widely used in related research (He et al., 2018). To limit false positives, we impose a similarity threshold of 0.85 and manually verify the final matches.

[13]Orbis provides standardized industry information for listed and unlisted firms worldwide. We use Orbis's name-matching tools, supplemented by manual validation, to link sanctioned firms to Orbis entries, achieving an 89.06% match rate. This high match rate supports the representativeness of Orbis for our sample and is consistent with recent work such as Bajgar et al. (2020).

[14]For example, if firm f is added to the Entity List in 2018, then $\text{Sanction}_{it} = 0$ before 2018 and $\text{Sanction}_{it} = 1$ from 2019 onward for inventors employed by that firm.

in business services, transportation equipment, and computer, communication, and other electronic equipment manufacturing, with a marked intensification of sanctions over time, especially in 2021.

Additional Firm-Level Controls and Import Exposure. We use additional firm-level data to construct comprehensive financial performance and to measure firms' pre-sanction exposure to sensitive imported inputs. We first draw on CSMAR, which provides detailed financial information for publicly listed Chinese firms, including R&D expenditure, firm size measured by log total sales, leverage, return on assets (ROA), sales growth, state ownership, firm age measured by years since listing, and cash flow measured as operating cash flow scaled by total assets. Following standard practice, we exclude financial firms and observations with missing financial data, and winsorize all continuous variables at the 1st and 99th percentiles. The resulting CSMAR panel contains 3,871 firms and 26,265 firm-year observations. [15]

We also use confidential Chinese customs transaction-level data, with 2016 as the latest available year, to measure firms' dependence on foreign input markets before the escalation of sanctions. These records report each import and export transaction at the firm level, with products classified by detailed Harmonized System (HS) codes, which we use as our product definition. For each import transaction, the data identify the importing firm, product, and source country, allowing us to measure sourcing exposure across narrowly defined products and origin markets. The customs data cover 7,572 imported products from 234 origin countries and regions. For the analysis, we classify products related to electric and electronic equipment, optics, media and broadcasting, and transport as sanction-sensitive, consistent with the industry concentration of U.S.-sanctioned firms in Table A1. We merge the customs records with the innovation data using a fuzzy name-matching procedure analogous to that used for sanctioned firms.

2.3 Motivating Evidence

We begin with several descriptive patterns that help motivate the empirical analysis. The goal here is not to establish causality, but to show why Entity List sanctions provide a useful setting for studying innovation and inventor reallocation.

[15] The CSMAR sample covers 73 industry classes out of the total 90 categories defined by the 2012 China Securities Regulatory Commission (CSRC) industry classification. Such wide industry coverage (81.11%) ensures representativeness of the sample firms. Appendix Table A2 presents the top 30 industry classes ranked by firm quantity. It reveals that CSMAR enterprises are predominantly concentrated in manufacturing sectors, especially electronic equipment, general and special machinery, chemical manufacturing and pharmaceutical production, alongside sizable proportions of information technology and business service entities.

The Innovation Importance of Sanctioned Firms The Entity List disproportionately covers firms close to the technological frontier of the Chinese economy. After consolidating subsidiaries, our sample contains 304 unique Chinese firms by 2022. Even before designation, these firms were substantially more innovative and larger than other firms in the same broad industries.[16] The importance of these firms is especially visible among publicly listed companies. Listed sanctioned firms account for about 60 percent of valid patents held by sanctioned entities as of June 2023.[17] These differences show that Entity List sanctions targeted firms with substantial innovative capacity, not marginal producers. Their designation placed a large stock of inventive human capital under constraint, making inventor mobility a natural channel through which technological capabilities could be reallocated domestically.

The restrictions also appear difficult to offset. Sanctioned firms may respond by strengthening compliance, diversifying suppliers, improving supply-chain transparency, or reducing exposure to high-risk jurisdictions. In practice, however, these responses are costly and incomplete. Legal and administrative relief is uncommon. Huawei challenged U.S. restrictions in 2019, and a small number of firms, including AMEC and IDG Capital, were later removed from restricted lists in 2024. ZTE was removed in 2018 only after paying substantial penalties and accepting extensive compliance obligations, a resolution that did not become a general precedent (Liangyu, 2018). For most firms, Entity List designation therefore represents a persistent constraint on access to key technologies, making it a meaningful shock to innovation and to the allocation of inventive labor.

Firm-Level Disruption after Entity List Designation. Although firm performance is not our central focus, we first document how Entity List designation is associated with changes in targeted firms' operating environment, in line with Cao et al. (2024). We estimate a simple staggered difference-in-differences specification,

$$Y_{ft} = \beta \cdot \text{Sanction}_{ft} + \boldsymbol{Z}'_{ft}\theta + \mu_f + \lambda_t + \varepsilon_{ft}, \tag{1}$$

where f indexes firms and t indexes years. The treatment variable, Sanction_{ft}, turns on when firm f is added to the BIS Entity List and remains equal to one thereafter. Since firms are not randomly selected for designation, we interpret β as the conditional association between Entity List placement and subsequent firm outcomes, rather than as the main causal estimate. The outcomes capture two dimensions of firm adjustment: innovation

[16] Appendix Figure A4 compares sanctioned and non-sanctioned firms across measures of innovation and scale. Sanctioned firms have higher patent output, broader technological scope, greater R&D spending, and larger assets, sales, and employment.

[17] Among sanctioned entities, 31 percent are publicly listed firms, and these firms hold 60 percent of patents associated with sanctioned entities, excluding Huawei.

and inventor reallocation in the full firm sample, and financial performance in the listed-firm sample.[18] The control vector $\boldsymbol{Z}_{ft}$ is chosen to match the outcome and sample.[19] All specifications include firm and year fixed effects, with standard errors clustered at the firm level. We estimate Poisson models for count outcomes with many zeros and OLS models for the remaining outcomes.

Appendix Table A3 reports the estimates. Panel A uses the full innovation sample and focuses on patenting and inventor flows. Entity List designation is followed by weaker innovation output and a clear loss of inventive labor. Sanctioned firms file about 1.5 fewer patents per year, experience roughly 23 additional inventor departures, and become 19.3 percentage points more likely to have a net inventor outflow. Inventor inflows do not increase enough to offset these departures. The decline in patenting therefore appears alongside a weakening of the human-capital base on which innovation depends. Panel B turns to listed firms, for which detailed financial accounts are available. The estimates point to broader operating pressure after designation. Profitability declines, with ROA falling by 2.2 percentage points and ROE by 4.0 percentage points. Employee earnings growth also falls, operating cash flow weakens, and leverage rises by 3.8 percentage points. The estimates point to more than a decline in firm performance. Entity List designation is followed by weaker fundamentals, lower patenting, and a net loss of inventors, suggesting that sanctions may reshape domestic innovation partly through the movement of inventive human capital across firms.

Inventor Outflows from Sanctioned Firms. The firm-level evidence above naturally lead to the question of whether sanctions also reallocate the inventors who embody firms' technological capabilities. Appendix Figure A5 shows that aggregate inventor mobility, measured by both the number of moves and the probability of moving, rises gradually in the early sample years before accelerating after 2020 and reaching its peak in 2022. The timing coincides with the escalation of Entity List sanctions and points to inventor mobility as a relevant margin of adjustment.

Aggregate mobility alone is not enough for our question. A rise in inventor moves could reflect general labor-market churn, sectoral expansion, or movement among firms unrelated

[18] The innovation and inventor-reallocation outcomes include annual patent counts, inventor inflows, inventor outflows, an indicator for net inventor outflow, and the number of net inventor outflows. The financial outcomes include return on assets (ROA), return on equity (ROE), employee earnings growth, operating cash flow scaled by total assets, and leverage scaled by total assets.

[19] For innovation outcomes, we control for technological diversity, measured by the number of IPC classes represented in firm f's patent applications in year t (Nayak et al., 2025), and patent stock, defined as an indicator for whether the firm filed any patents during the previous three years (Oettl and Agrawal, 2008). For listed-firm financial outcomes, we control for firm size, Tobin's Q, sales growth, firm age, and state ownership.

to sanctions. What matters here is whether inventors disproportionately leave sanctioned firms relative to the flow of inventors into them. We therefore disaggregate mobility by the sanction status of the origin and destination firms, distinguishing moves from sanctioned to non-sanctioned firms (S–N), from non-sanctioned to sanctioned firms (N–S), within non-sanctioned firms (N–N), and within sanctioned firms (S–S). Most moves occur among non-sanctioned firms, while moves between sanctioned firms are rare, making the comparison between S–N and N–S flows the central one. Appendix Figure A6 shows that this comparison changes sharply after sanctions begin. Before 2018, S–N and N–S flows were similar in magnitude. After 2018, S–N flows begin to dominate, and the gap widens substantially after 2020 as Entity List designations intensify. By 2021, inventor moves out of sanctioned firms were roughly twice as large as moves into them. The aggregate mobility series masks this directional imbalance, which motivates our micro-level analysis of how sanctions reshape inventor mobility and subsequent innovation.

3 Empirical Analysis

We now examine the micro-level adjustments behind these patterns. By following inventors before and after first exposure to Entity List sanctions, we test whether firm-level shocks lead to inventor departures, changes in individual innovation, and the reallocation of technological capabilities across firms.

3.1 Regression Specification

The empirical design exploits the staggered timing of firms' additions to the Entity List. We compare inventors before and after their first exposure to a sanctioned employer, while holding fixed persistent differences across inventors and firms. This within-inventor design is important because inventors differ systematically in ability, research orientation, and baseline propensity to move, all of which may also shape innovation. We estimate the following inventor-year specification:

$$Y_{ift} = \beta \cdot \text{Sanction}_{it} + \boldsymbol{X}'_{it}\gamma + \boldsymbol{Z}'_{ft}\theta + \delta_i + \mu_f + \lambda_t + \varepsilon_{ift} \tag{2}$$

where i indexes inventors, f indexes firms, and t indexes years. We study two outcomes. The first is inventor mobility. The variable FlowOut_{ift} is an indicator equal to one in the year in which inventor i leaves the firm with which the inventor was previously affiliated, and zero otherwise. This outcome captures whether exposure to sanctions increases departures from affected firms. The second outcome is inventor innovation, measured by the log number of

patent applications filed by inventor i in year t.[20] We use patent applications rather than grants because applications are observed closer to the timing of inventive activity and are less affected by approval lags, which would otherwise generate right-censoring near the end of the sample.[21] Our measure of sanction exposure, Sanction_{it}, equals one after inventor i is first exposed through employment at a firm added to the Entity List after 2018, and remains equal to one thereafter. This absorbing definition fixes exposure at the time of first contact with the sanction shock, rather than allowing exposure to change mechanically when inventors move across firms. The coefficient β therefore captures changes in inventor mobility and innovation following first exposure, including adjustments that occur after inventors leave the initially sanctioned employer.

The specification relies on rich fixed effects to compare changes within inventors and firms over time. Inventor fixed effects, δ_i, are especially important because they absorb permanent differences in inventor ability, research orientation, career ambition, and baseline propensity to switch employers. This helps ensure that the estimates are not driven by the fact that more mobile or more productive inventors are systematically selected into leaving firms. Firm fixed effects, μ_f, absorb persistent differences across employers, including technological specialization, organizational quality, and retention capacity, while year fixed effects, λ_t, account for aggregate changes in the innovation environment and labor market. Standard errors are clustered at the firm level. We further control for time-varying inventor and firm characteristics that may shape mobility and innovation. Following Liu et al. (2024), the inventor controls in $\boldsymbol{X}_{it}$ include experience, measured by the log number of years since the first patent application; average coauthor-team size in year t; and an indicator for patenting activity during the previous three years. The firm controls in $\boldsymbol{Z}_{ft}$ capture the employer's innovation environment, including technological diversity and recent patent stock.[22] Thus, the estimates compare changes among the same inventors, working at firms with similar observable innovation capacity, after accounting for persistent inventor, firm, and year-level differences.

[20] We use the log of patent counts rather than the $\log(1+y)$ transformation (Cohn et al., 2022) because all inventor-year observations in the estimation sample have patent records. We also report robustness checks using alternative patent-count models, including OLS and Poisson estimates. The results remain qualitatively unchanged.

[21] Patent applications also provide a cleaner measure around job transitions. In China, patent ownership typically belongs to the firm rather than the individual inventor. When an inventor changes employers, patents filed before the move remain with the previous firm. Applications therefore better capture changes in the inventor's current innovative activity.

[22] These firm-level controls are the same patent-based measures used in the motivating firm-level analysis reported in Panel A of Table A3. Technological diversity is measured by the number of IPC classes represented in the firm's patents, while patent stock is an indicator for whether the firm filed any patents during the previous three years.

Descriptive Statistics Appendix Table 1 reports the summary statistics for our sample. On average, 12 percent of the inventors in the sample switch jobs annually during the sample period. Each inventor files an average of 3.11 patent applications per year. Roughly 1 percent of the inventors in our sample undergone sanctioned employers. The average log number of years since the first patent application is 1.10. An inventor typically works in teams of about six people. 68 percent of inventors patented in the preceding three years relative to the focal year. In terms of employer characteristics, an average firm spans 10-11 technology areas as measured by IPC classes. Most of the firms have innovation experience over the three-year window, suggesting the continuity of innovation activities.

3.2 Main Results

Baseline Results. Table 2 reports the baseline estimates. Panel A uses the full inventor sample. Columns (1)–(3) examine inventor mobility, adding progressively richer controls. With firm and year fixed effects, Entity List exposure is associated with a 6.7 percentage-point increase in the probability that an inventor leaves the firm. Adding inventor fixed effects and inventor-level controls raises the estimate to 9.4 percentage points, indicating that the result is not driven by stable differences in inventors' underlying mobility propensity or research profiles. Once we further control for the firm's time-varying innovation environment through recent patenting and technological diversity, the estimate declines modestly to 8.1 percentage points and remains precisely estimated. This preferred specification implies a large economic effect: relative to the baseline annual mobility rate of 12 percent, sanction exposure raises the probability of moving by about 66.7 percent.

Columns (4)–(6) turn to inventor innovation, measured by the log number of annual patent applications. The estimates are small and statistically indistinguishable from zero across specifications, even as the mobility response remains large. Entity List exposure therefore appears to change where inventors work, rather than reducing their own patenting activity. This pattern is consistent with the idea that sanctions disrupt firms' access to technologies and organizational routines, while leaving much of the inventive capacity embodied in individual researchers intact.

Panel B restricts attention to inventors observed moving across firms. The mobility response is larger in this sample, indicating that sanction exposure shifts the movement of inventors already connected to the external labor market. Yet the innovation estimates remain small and statistically insignificant. Mobile inventors therefore appear to carry their productive capacity with them after exposure and relocation. Viewed alongside Panel A, the results suggest that sanctions reallocate inventive human capital across firms rather than

erode the productive capacity of individual inventors.

Dynamic Effect. The baseline estimates use the timing of first exposure to Entity List sanctions. We therefore examine whether inventor outcomes were already changing before formal designation and how they evolve afterward. This is useful for two reasons. If firms or inventors anticipated future sanctions, mobility or patenting could begin to adjust before the listing date. Similarly, if exposed and unexposed inventors were already following different trajectories, this would appear in the pre-period coefficients. We assess these patterns using the following event-study specification:

$$Y_{ift} = \sum_{s=-5, s\neq -1}^{3} \beta_s \cdot \text{Sanction}_i \times \mathbf{1}\{t - t_i = s\} + \boldsymbol{X}'_{it}\gamma + \boldsymbol{Z}'_{ft}\theta + \delta_i + \mu_f + \lambda_t + \nu_{ift} \tag{3}$$

where Sanction_i equals one if inventor i is ever exposed to Entity List sanctions after 2018, and zero otherwise. The variable t_i denotes the year in which inventor i's employer is added to the Entity List. The controls and fixed effects follow equation (2). We normalize the year before exposure, $s = -1$, to zero, so β_s traces inventor outcomes relative to the pre-sanction year. Pre-exposure coefficients show whether mobility or innovation was already changing before designation, while post-exposure coefficients describe the timing and persistence of adjustment after sanctions take effect.

Figure 2 reports the event-study estimates. We use the conventional two-way fixed-effects OLS specification as the baseline and compare it with alternative estimators designed for staggered adoption settings (Cengiz et al., 2019; Callaway and Sant'Anna, 2021; Sun and Abraham, 2021). The estimates are shown in the same figure using different markers. The pre-designation coefficients are small and statistically indistinguishable from zero for both outcomes. Inventors who are later exposed to sanctions do not exhibit differential trends in mobility or patenting before their employers are added to the Entity List. This pattern is reassuring for the timing-based design, since the adjustment we estimate does not appear to reflect earlier changes in inventor behavior or innovation trajectories.

After designation, the adjustment appears first and most clearly in inventor mobility. The probability that an inventor changes firms rises by about 10 percentage points immediately after sanction exposure. Although the effect declines somewhat over the next two years, it remains positive and returns to roughly 7 percentage points by the third post-sanction year. Individual patenting, shown in Panel (b), shows no comparable response. Across event years and estimators, the effects on patent applications remain small and statistically insignificant. The dynamic evidence mirrors the baseline estimates, showing that sanctions reallocate inventors across firms without measurably reducing their own patenting activity.

3.3 Heterogeneous Effects on Inventor Mobility

Sanctions do more than raise turnover when departing inventors carry capabilities that remain useful beyond the sanctioned firm. We therefore examine how mobility varies with inventor profiles, sanctioned-firm conditions, and the geographic scope of moves. These patterns help distinguish general labor-market churn from a systematic reallocation of technological capabilities across the domestic innovation system.

Inventor Profiles. Table 4 first examines whether the mobility response differs across inventors within sanctioned firms. Columns (1)–(3) show that departures are not evenly distributed across the inventor workforce. Column (1) uses an inventor's position in recent patent teams, measured by whether the inventor appeared as the first inventor on a patent during the previous three years. The negative interaction suggests that inventors more central to the firm's innovation teams are less likely to leave after sanction exposure. By contrast, inventors with more peripheral roles respond more strongly.

Columns (2) and (3) use recent patent quality to proxy for inventor capability, measured by citations per patent and patent-family size over the previous three years. The estimates again point to a stronger mobility response among inventors with weaker recent innovation records. This pattern is consistent with sanctioned firms retaining inventors who are more central to ongoing projects, while adjustment occurs more strongly among inventors whose roles are less deeply embedded in the firm's core innovation activities.

Sanctioned-Firm Conditions. A sanction shock is less likely to trigger inventor departures when the firm has the internal capacity to keep research projects viable. Firms with deeper knowledge bases and larger research teams may be better able to absorb disruptions, preserve ongoing projects, and retain inventors. Columns (4) and (5) examine this margin by interacting sanction exposure with two measures of internal innovation capacity: the firm's cumulative inventor base, *KnowledgeSize*, and the average number of co-inventors per patent, *TeamSize*. The negative interaction terms indicate that inventor outflows are weaker in firms with larger knowledge bases and more team-based research. This pattern suggests that internal technological depth and team embeddedness buffer the labor-market response to sanctions.

The opposite force should arise when the sanctioned firm depends on inputs directly exposed to U.S. restrictions. Columns (6)–(8) therefore interact sanction exposure with pre-sanction import dependence along three margins. The main measure, *Dependency_US*, captures the share of total imports accounted for by U.S. inputs in sanction-sensitive sectors. The other two measures provide placebo-style comparisons: *Dependency_USOther* captures

U.S. imports outside sanction-sensitive sectors, while *Dependency_Other* captures imports from non-U.S. origins. Column (6) shows that inventor outflows are significantly stronger among firms more dependent on U.S. inputs in sanction-sensitive sectors before designation. By contrast, Columns (7) and (8) show no comparable response for U.S. imports outside these sectors or for non-U.S. imports. The evidence therefore points to the specific technological and sourcing constraints imposed by Entity List restrictions, rather than to general import exposure or broader supply-chain dependence.

Regional Scope of Mobility. Finally, the Entity List sanctions may cause some inventors who previously could have moved but opted not to for various reasons to engage in job hopping. In reality, many inventors choose not to move due to specific geographic location preferences, especially home bias. The resulting job stickiness then stems from such attachment, inertia, switching costs, or risk aversion. However, the sanction shock may compel these inventors to finally make a move decision.

In this subsection, we examine geographic location preferences of inventors, specifically, whether they are likely to change jobs across counties or even provinces. If inventors exhibit some home bias or location preferences, they may accept a job at lower wages but in a preferred location. The results in Table 3 show the results. Columns 1 and 2 include all inventors. Column 1 shows that inventors are 1.6 percentage points more likely to switch jobs to a different county following a sanction, nearly 1.3 times the baseline inter-county job-changing probability of 0.0127. The large economic magnitude also applies to inter-city mobility, where exposure to sanctions doubles the inter-city job-changing probability, as indicated in column (2). Results remain qualitatively similar when limit to the subsample of mobile inventors in Columns 3 and 4. Altogether, our findings suggest that the breakdown of location preferences underscores the severity of the sanction shock and its role in triggering nationwide talent reallocation.

Viewed through these margins, the mobility response is not an indiscriminate rise in turnover. It is strongest where sanctions bind more tightly through dependence on critical U.S. inputs, and weaker where firms have deeper knowledge bases and larger research teams that can sustain ongoing projects. It is also concentrated among less embedded inventors and extends to cross-regional moves, suggesting that some displaced capabilities are portable beyond the original firm and local labor market. This creates a supply of displaced inventive capabilities whose effects on receiving firms we examine next.

3.4 Robustness Checks

The baseline estimates combine staggered sanction exposure with patent-based measures of inventor mobility and innovation. We therefore conduct a battery of robustness checks that vary the comparison group, the construction of inventor careers from patent records, and the empirical specification for mobility and patenting outcomes.

Balancing Treated and Comparison Inventors before Sanctions. A natural concern is that sanctioned firms, and the inventors they employ, differ from other firms and inventors even before designation. Entity List restrictions tend to target technologically important firms, whose inventors may have different pre-sanction productivity, experience, or mobility patterns. To make the comparison more balanced, we re-estimate equation (2) after propensity score matching. The matching uses pre-sanction averages from 2014–2017 of the firm- and inventor-level controls in the baseline specification, imposes a caliper of 0.01, and restricts the comparison to the region of common support. Appendix Table A5 reports the matched-sample estimates. Panel A balances treated and comparison observations on pre-sanction firm characteristics, while Panel B does so on pre-sanction inventor characteristics. In both cases, sanction exposure continues to predict higher inventor outflows, while the patenting estimates remain statistically indistinguishable from zero.[23] The baseline pattern therefore survives comparisons among firms and inventors with similar pre-sanction observables.

Addressing Contamination from Technological Peers. A related concern is that firms close to sanctioned firms in technology space may themselves adjust around the same time. Competitors in similar IPC fields could hire displaced inventors, change their own innovation strategies, or be indirectly affected by the same policy environment, making them a less clean comparison group. We therefore re-estimate the baseline specification after excluding control firms that share three or five core IPC technology domains with sanctioned firms. Appendix Table A6 shows that the estimates remain close to the baseline. The mobility response is therefore not driven by comparisons with technologically adjacent firms that may have been exposed to the same competitive or innovation-market adjustments.

Addressing Incomplete Inventor Histories. Patent records reveal inventor careers only when inventors file patent applications; they are not continuous employment histories. This feature matters for both mobility and innovation. If an inventor patents only intermittently, the exact timing of a job change may be measured with error because moves are

[23] The results are stable across 1:1, 1:2, and 1:3 nearest-neighbor matching specifications.

inferred from changes in patent assignees. If an inventor moves and then stops patenting, we may either fail to observe the move or lose the inventor from the post-move innovation sample. In both cases, the concern is that the patent-based panel may overrepresent inventors who remain active patenters after exposure.

We address this concern with a Heckman selection correction that models whether an inventor is observed patenting in a given year. In the first stage, we use prior-year patenting to predict current-year observation in the patent panel. This variable captures persistence in inventive activity, which is central to whether an inventor appears in patent data in the current year. We estimate the selection equation using both a linear probability model and a Probit model. Because a Probit model with inventor fixed effects is computationally demanding at this scale, we also implement a bootstrap procedure, repeatedly drawing subsamples of exposed and comparison inventors and estimating the two-step Heckman model with inventor fixed effects in the first stage.

Appendix Table A7 reports the selection-corrected estimates. Prior-year patenting strongly predicts current patenting in the first stage, confirming that the selection equation captures the activity-based structure of the patent panel. The second-stage estimates remain close to the baseline estimates. Appendix Figure A9 shows the corresponding bootstrap distributions: the mobility coefficients remain concentrated on positive values, while the patenting coefficients remain close to the baseline distribution. The results therefore do not appear to be driven by selective observation of inventors in patent records or by errors induced by gaps between patent filings.

Addressing Unobserved Job-Switch Dates. In the baseline analysis, a job change is dated to the first year in which an inventor is observed patenting with a new firm. This rule uses the earliest employer information available in the patent data, but it may misdate the actual move when there is a long gap between patent applications. Such timing error should be most relevant for inventors with infrequent patenting, whose employer transitions are observed only after a delay. We therefore tighten the mobility measure by restricting the sample to inventors whose observed interval between employer changes is short. Columns (1)–(3) of Appendix Table A8 impose increasingly strict windows, limiting the maximum observed job-change interval to less than five years, less than three years, and exactly one year. Across these restrictions, sanction exposure continues to raise the probability of job switching, with estimates between 6 and 8 percentage points. The mobility result is therefore not driven by long gaps between patent filings or by the imputed timing of job changes.

Addressing Local Time-Varying Shocks. Inventor mobility may also respond to changes in local economic conditions, industrial policies, or regional innovation environments that coincide with sanction timing. We therefore augment the baseline specification with increasingly demanding geographic-by-year fixed effects, including city-year and county-year fixed effects. These specifications compare inventors operating within the same local labor market in the same year, absorbing shocks that vary across regions over time. Appendix Table A9 reports the results. The coefficient on sanction exposure remains positive and precisely estimated for inventor outflows across all specifications. The estimates for patenting outcomes are also stable in magnitude and significance. Thus, the baseline findings are not explained by contemporaneous regional shocks or local policy changes that may affect inventor mobility and innovation.

Alternative Patent-Count Specifications. Finally, we examine whether the innovation results depend on how patent output is modeled. Our baseline measure uses the log number of patent applications in an unbalanced inventor-year panel. We therefore re-estimate equation (2) using patent counts directly, applying both OLS and fixed-effects Poisson models while keeping the remaining specification unchanged. Appendix Table A10 reports the results. Columns (1) and (2) use the full inventor-year sample; Columns (3) and (4) restrict the sample to inventors who patent in at least five years; and Columns (5) and (6) further restrict the sample to inventors observed patenting in every sample year. Across all samples and both count specifications, the sanction coefficient remains statistically insignificant. The absence of a patenting response is therefore not driven by the log transformation, the unbalanced panel structure, or the choice between linear and Poisson models.

3.5 Unintended Technology Transfer

The mobility response may do more than reallocate labor. Because displaced inventors carry tacit knowledge, sanctions can unintentionally transmit technological capabilities to recipient firms. We study this diffusion through two margins: shifts in recipient firms' technology fields and productivity gains among their incumbent inventors.

3.5.1 Inter-Firm Knowledge Diffusion

A central challenge in studying knowledge diffusion is that the knowledge acquired through prior employment is not directly observed. We therefore follow the logic of revealed comparative advantage, applied to technology fields, and infer knowledge transfer from changes in recipient firms' patenting across narrowly defined IPC classes (Bahar and Rapoport, 2018a;

Bahar et al., 2020a). The premise is that tacit capabilities should leave a technological footprint. If displaced inventors carry knowledge from their origin firms, recipient firms should increase innovation in the fields where those origin firms were previously strong. We therefore trace not only how much recipient firms patent, but where they patent. If incoming inventors transmit origin-firm capabilities, recipient firms should shift toward the technology fields in which sanctioned origin firms were strong. This field-level response provides a sharper test of knowledge diffusion than aggregate patenting alone.

Measuring Innovation Performance in Narrow IPC Classes. To measure technology-specific innovation performance, we use the Revealed Technological Advantage (RTA) index, following Soete (1987) and Bahar et al. (2020b).[24] For firm f, technology class p, and year t, we define

$$\text{RTA}_{fpt} \equiv \frac{m_{fpt} / \sum_p m_{fpt}}{\sum_f m_{fpt} / \sum_f \sum_p m_{fpt}}, \tag{4}$$

where m_{fpt} is the number of patent applications filed by firm f in technology class p in year t. The numerator is the share of firm f's patents in class p, while the denominator is the corresponding national share of patents in that class. The index therefore measures whether firm f is disproportionately active in technology class p relative to the national benchmark. Values above a benchmark level, conventionally one, indicate revealed specialization in that technology field. This measure is particularly well suited to this exercise because it captures the technological direction of innovation, not only its scale. While patent counts and citations measure the volume and quality of output, RTA shows whether recipient firms become more specialized in fields tied to incoming inventors' prior expertise. It therefore identifies the content of knowledge diffusion, not merely the amount of innovation produced.[25]

Specification for Knowledge Diffusion. We use the RTA measure to examine whether sanction-induced inventor mobility changes the technological direction of recipient firms. The key prediction is field-specific. If incoming inventors carry capabilities from their origin firms, recipient firms should become more specialized in the technology classes where those

[24] The RTA index is analogous to the Revealed Comparative Advantage (RCA) index introduced by Balassa (1965) and widely used in international trade and economic development (Bahar and Rapoport, 2018b; Bahar et al., 2024; Giorgi et al., 2025).

[25] OPPO provides illustrative evidence in Appendix Section B. After receiving inventors from sanctioned firms with optics-related strengths, OPPO's RTA in IPC G02 rose from below one in most pre-sanction years to 3.15 in 2019. In the narrower G02F1/1335 subgroup, patenting increased from one patent in 2014–2017 to ten patents after 2018, with 60% of the new patents linked to mobile inventors, including movers' own patents, joint patents with OPPO incumbents, and follow-on patents by incumbents who had collaborated with movers.

origin firms had revealed advantage at the time the inventors left. We estimate

$$\ln(\text{RTA}_{fpt}) = \beta \cdot \sum_{f'} \text{Flow}_{f' \to f,t} \times \mathbf{1}\{\text{RTA}_{f'p\tau(f',f,t)} \geq 1\} + \mu_{ft} + \lambda_{pt} + \delta_{fp} + \upsilon_{fpt}, \quad (5)$$

where f indexes recipient firms, p indexes IPC technology classes, and t indexes years. The variable $\text{Flow}_{f' \to f,t}$ denotes the number of inventors moving from origin firm f' to recipient firm f in year t. We evaluate the origin firm's technological specialization in the year these inventors leave, denoted by $\tau(f', f, t)$. The indicator $\mathbf{1}\{\text{RTA}_{f'p\tau(f',f,t)} \geq 1\}$ equals one if origin firm f' held revealed specialization in technology class p at that time. Summing across origin firms, the key regressor captures the inflow of inventors whose prior employers specialized in class p. Robust standard errors are clustered at the firm level.

The coefficient β captures whether recipient firms reorient innovation toward the technology fields in which the origin firms of incoming inventors were relatively strong. This comparison is disciplined by a rich set of fixed effects. Firm-year effects, μ_{ft}, absorb changes in the firm's overall innovative capacity, demand conditions, and research intensity in year t. Technology-year effects, λ_{pt}, remove common shocks to technology class p in that year. Firm-technology effects, δ_{fp}, absorb persistent differences in each firm's capabilities and research orientation across fields. The remaining variation comes from the allocation of innovation across technology classes within the same recipient firm and year, net of the firm's pre-existing specialization in each class. The estimate therefore does not rely on whether firms hiring more inventors innovate more overall. Instead, it asks whether innovation within a given firm-year shifts toward the fields embodied in the knowledge of incoming inventors. A positive β is thus consistent with the transfer of technological content from origin firms, rather than a mere increase in inventive labor.[26]

In practice, We restrict attention to firm–IPC cells with patenting activity, yielding a panel of 982,047 observations across 87,309 firms and 123 IPC classes. Appendix Figure A7 plots the distribution of log RTA. The measure varies widely, with a standard deviation of 1.67 and values ranging from -6.68 to 10.54. This dispersion shows that firms differ sharply in the technology fields in which they hold revealed advantage, providing substantial within-firm and across-field variation for the empirical analysis.

[26]This interpretation does not require inventor mobility to be random. Recipient firms may select inventors based on prior experience or firm-level conditions. The identifying assumption is more limited: conditional on the controls and fixed effects in equation (5), such selection is not driven by unobserved shocks that would have led the same firm, in the same year, to expand in precisely the technology classes associated with the origin-firm strengths of incoming inventors.

Main Findings. Panel A of Table 5 reports the estimates of equation (5). Column (1) includes firm-year and IPC-year fixed effects, so the comparison is made within recipient firms in a given year and within technology classes in a given year. The coefficient on inventor inflows from origin firms specialized in class p is positive and statistically significant. Column (2) adds firm-IPC fixed effects, which absorb each firm's persistent specialization in each technology class. The coefficient falls from 0.568 to 0.035, indicating that part of the raw association reflects stable firm-field advantages. Once all three sets of fixed effects are included in Column (3), the estimate remains positive and precisely estimated at 0.040. Thus, even after accounting for firm-year shocks, technology-year shocks, and persistent firm-field specialization, recipient firms become more specialized in the precise technology fields associated with incoming inventors' origin-firm strengths.

Columns (4)–(6) examine the extensive margin of inventor inflows. Instead of using the number of incoming inventors, we use an indicator for whether firm f receives at least one inventor from an origin firm with revealed specialization in class p. The estimates remain positive and statistically significant across the same sequence of fixed effects. In the fully saturated specification, the arrival of at least one such inventor is associated with a 2.6 percent increase in the recipient firm's RTA in the corresponding IPC class. The result is therefore not driven by the scaling of the count variable or by a small number of large flows; the presence of technologically relevant incoming inventors is itself associated with a shift in the recipient firm's innovation portfolio.

Panel B sharpens the interpretation by distinguishing inflows from sanctioned and non-sanctioned origin firms. This split is useful because inventor moves from sanctioned firms are more closely tied to the policy-induced displacement documented above, while other mobility flows may reflect ordinary hiring or strategic matching. Columns (1)–(3) show that inflows from sanctioned firms have a substantially larger association with recipient-firm specialization. In the fully saturated specification, receiving at least one inventor from a sanctioned origin firm is associated with about a 12 percent increase in RTA in the relevant technology class. By contrast, Columns (4)–(6) show that inflows only from non-sanctioned firms are associated with a much smaller effect, about 1.5 percent in the comparable specification.

The contrast between sanctioned-origin and non-sanctioned-origin inflows helps separate policy-induced knowledge reallocation from ordinary hiring. If the estimates mainly reflected recipient firms recruiting inventors in fields they already planned to expand, the effects should not differ so sharply by the sanction status of the origin firm. Instead, the much larger effect for sanctioned-origin inventors is consistent with the mechanism emphasized in this paper: sanctions displace inventors from technologically specialized firms, and their movement shifts recipient firms toward the technology fields embodied in those inventors' prior experience.

The evidence therefore points to inter-firm knowledge diffusion through inventor mobility, rather than a mere increase in recipient firms' overall inventive labor.

3.5.2 Technology Convergence and Innovation Competition

Inter-firm knowledge diffusion can also change the competitive structure of innovation. If recipient firms absorb technological capabilities from incoming inventors, their research portfolios may move closer to those of their industry peers. Such convergence matters because firms that innovate in more similar technology spaces are more likely to compete for the same technological opportunities, inputs, and product-market applications. We measure each firm's technological proximity to its industry peers using cosine similarity between patent-portfolio vectors:

$$\text{Similarity}_{ft} = \frac{\mathbf{u}_{ft} \cdot \mathbf{v}_{k(f),t}}{||\mathbf{u}_{ft}|| \cdot ||\mathbf{v}_{k(f),t}||}. \tag{6}$$

The vector $\mathbf{u}_{ft}$ describes firm f's patent portfolio across IPC subclasses in year t, with each element equal to the number of patents filed in a given subclass. The peer vector $\mathbf{v}_{k(f),t}$ is constructed analogously as the average patent portfolio of other firms in firm f's industry k. We define industries by the primary IPC class in which a firm filed the largest share of its patents during the pre-sanction period, 2014–2017. This IPC-based classification links firms to their technological rather than accounting industries. We exclude industry-year cells with only one firm, since peer similarity cannot be computed in those cases, leaving 121 IPC-based industries. We estimate the relationship between inventor inflows and technological proximity at the firm-year level. The dependent variable is Similarity_{ft}, and the key explanatory variable is the number of incoming inventors who held revealed technological advantage in at least one field before moving. We include the same time-varying firm controls as in Table 2, and cluster standard errors at the firm level.

Main Findings. Table 6 reports the results. Panel A includes firm, year, and industry fixed effects, absorbing persistent firm differences, aggregate shocks, and time-invariant differences across industries. A remaining concern is that U.S. technology restrictions may generate broad domestic substitution incentives within affected industries, pushing many firms toward similar research directions even without inventor mobility. Panel B therefore adds industry-year fixed effects, which absorb time-varying industry-level shocks, including sector-specific policy responses, technology trends, and changes in industry-wide patenting activity. Across both panels, inventor inflows are positively associated with technological proximity to industry peers. The count measure in Column (1) and the indicator for any inflow in Column (2) both imply that recipient firms become more technologically similar

to other firms in the same industry after absorbing mobile inventors. The effect is much stronger when the incoming inventors come from sanctioned firms. In Column (3), the coefficient for sanctioned-origin inflows is about 2.8–3.5 times larger than the baseline inflow effect in Column (2). By contrast, Column (4) shows that inflows only from non-sanctioned firms have a much smaller association, roughly 25–31 percent of the sanctioned-origin effect.

These results suggest that sanction-induced mobility extends beyond firm-level knowledge transfer. As inventors move from technologically advanced but constrained firms to other domestic firms, recipient firms become closer to their industry peers in technology space. This convergence is consistent with stronger innovation competition, especially when inflows come from sanctioned-origin inventors. Sanctions therefore appear to diffuse high-tech human capital in a way that reshapes firm specialization and narrows technological distance within domestic industries.

3.5.3 Spillover Effects on Incumbent Inventors

The preceding evidence shows that mobile inventors are associated with a reorientation of recipient firms' technology portfolios. That firm-level pattern, however, could partly reflect incoming inventors' own patenting. To examine whether knowledge also spreads within the receiving organization, we focus on incumbent inventors who were already employed at the destination firm and did not move. If incoming inventors transmit tacit expertise through collaboration, shared projects, or informal learning, these non-moving colleagues should become more active in the technology fields associated with the incoming inventors.

Spillovers are most likely under related specialization. Incumbents need enough technological overlap with incoming inventors to absorb their expertise, but enough difference for the inflow to introduce genuinely new capabilities. Table A4 shows this pattern. Sanctioned and non-sanctioned firms overlap in about 76% of IPC classes, indicating a broad common technological base. Yet their revealed strengths differ sharply.[27] The quality of incoming knowledge also matters. Figure A8 shows that inventors from sanctioned firms who later moved to recipient firms had stronger pre-move innovation records than incumbent inventors at their eventual destinations during 2015–2017. They produced more patents, received more forward citations, and held broader patent-family coverage. These differences point to an inflow of higher-quality technical experience that could plausibly spill over to incumbent colleagues in related technology fields.

[27] Sanctioned firms are concentrated in more technology-intensive and sanction-sensitive fields, such as H03 (electronic circuits and signal processing), G03 (photography and imaging technologies), and G02 (optical technologies). Non-sanctioned firms display more dispersed specialization across fields such as B81 (micro-mechanical technologies), B33 (additive manufacturing), and A44 (jewellery and personal adornments).

Table 7 examines whether this knowledge remains confined to the movers or diffuses to colleagues inside the recipient firm. The dependent variable is patenting by incumbent, non-moving inventors in the same technology class p associated with the incoming inventors' origin-firm strengths. Column (1) shows that 100 incoming inventors from origin firms specialized in class p are associated with a 6.1 percent increase in incumbent patenting in that field. Column (2) reaches the same conclusion using an indicator for the arrival of any such inventor. The source of the inflow matters sharply. When incoming inventors come from sanctioned firms, Column (3) shows a 12.5 percent increase in incumbent patenting in the relevant field, far larger than the 2.3 percent effect for general inflows. By contrast, Column (4) shows that inflows only from non-sanctioned firms have a small and imprecise effect. This pattern is difficult to reconcile with a purely mechanical increase in firm patenting from hiring additional inventors. The response occurs among incumbent colleagues and in the specific fields linked to incoming inventors' prior expertise. The evidence points to within-firm spillovers, as incumbent colleagues absorb and redeploy capabilities brought by sanctioned-origin inventors.

4 Aggregate Innovation Effects of Sanctions

The analysis so far leaves a natural accounting question. Sanctions clearly damage targeted firms, both by restricting access to key technologies and by pushing inventors out. Yet the same outflows also move knowledge into other domestic firms, where they generate spillovers. We therefore use a back-of-the-envelope exercise to compare these two forces and gauge how much of the innovation loss at sanctioned firms is absorbed through unintended knowledge diffusion elsewhere in the economy.

Simple Accounting Method. To discipline the magnitude of the offsetting innovation effects, we use a simple accounting decomposition. Let M denote mobile inventors, N non-sanctioned receiving firms, and S sanctioned origin firms. We write the aggregate change in annual patent applications as

$$\Delta P = \sum_{i \in M} (y_{i,T} - y_{i,T_0}) + \sum_{n \in N} \Phi_n - \sum_{s \in S} \mathcal{L}_s, \tag{7}$$

where y_{i,T_0} and $y_{i,T}$ denote inventor i's patent output before and after sanctions, respectively. The first term captures whether mobile inventors become less productive after moving. Table 2 shows little change in their own patenting, so this component is near zero. Because movers' own patenting remains stable, the accounting reduces to a comparison between receiving-firm

spillover gains and innovation losses at sanctioned origin firms.

The second term in (7) captures spillover gains at non-sanctioned receiving firms. As shown in Section 3.5, these gains arise when incoming inventors raise incumbent patenting in the technology fields where the movers' origin firms held revealed technological advantage. Table 7 provides the empirical basis for the calculation. To keep the exercise conservative, we count only these field-specific inflows rather than assigning a general productivity gain to all inventor mobility. For receiving firm n, we define

$$\Phi_n = \sum_{p \in \mathcal{P}_n} \kappa_{np} \times \text{Influx}_{n,p}. \tag{8}$$

where p indexes IPC technology classes, and $\mathcal{P}_n$ is the set of fields in which receiving firm n is active. The variable $\text{Influx}_{n,p}$ measures the inflow of inventors from sanctioned firms that held revealed technological advantage in class p.[28] The parameter κ_{np} converts the estimated spillover into a patent-count gain. We construct it by multiplying the coefficient from Column (1) of Table 7, 0.061, by the average pre-existing patent output of incumbent inventors at firm n in class p. This maps the estimated increase in incumbent patenting into additional patents generated at the receiving firm. Appendix C illustrates the calculation for one receiving firm.

The third term captures innovation losses at sanctioned origin firms that go beyond the direct output of departing inventors. When inventors leave, the origin firm may lose not only their own patenting, but also the collaborative complementarities they provided to remaining colleagues. These losses may arise through disrupted research teams, weakened co-authorship networks, and reduced support for high-impact inventors whose productivity depends on specialized collaborators. Because such organizational losses cannot be separately observed in the reduced-form estimates, we approximate them with a back-of-the-envelope residual. For sanctioned firm s, we define

$$\mathcal{L}_s = \Delta y_s - \sum_{i \in M_s} y_{i,T_0} \tag{9}$$

where Δy_s is the post-sanction decline in annual patent applications at sanctioned firm s, and M_s is the set of inventors who leave that firm. Thus, $\sum_{i \in M_s} y_{i,T_0}$ captures the direct patenting

[28] As shown in Table A1, U.S. sanctions began in 2018 and intensified after 2021. Inventor mobility also rises sharply after 2021, as shown in Figure A5, while outward flows from sanctioned to non-sanctioned firms become increasingly asymmetric, as shown in Figure A6. We therefore measure $\text{Influx}_{n,p}$ as the mean annual inflow during 2021–2022. The calculation is conservative because it counts only inventors from sanctioned firms with revealed technological advantage in field p, excludes the post-move output of the movers themselves, and attributes gains only to incumbent inventors at receiving firms.

capacity lost through inventor departures, while the residual $\mathcal{L}_s$ captures the remaining decline in firm innovation after netting out this direct brain-drain component.[29] Appendix C illustrates the calculation for one sanctioned firm.

Quantification Results. Aggregating the firm-level components puts the competing forces in perspective. Spillover gains at non-sanctioned receiving firms amount to 235.98 patents, while indirect losses at sanctioned origin firms amount to 875.88 patents. The net effect is therefore a decline of 639.90 patents. Sanctions reduce aggregate patenting overall, but the loss is partly absorbed by the reallocation of inventive human capital to other domestic firms.

Table 8 also clarifies the source of the offset. Mobile inventors do not experience a measurable decline in their own patenting, so the main countervailing force comes from spillovers at receiving firms. These spillovers recover 26.94 percent of the estimated indirect loss at sanctioned origin firms. Put differently, sanctions still impose a substantial innovation cost on targeted firms, but roughly one quarter of that loss reappears as additional patenting elsewhere in the domestic innovation system. The aggregate effect is therefore not simply the destruction of inventive capacity. It is a net loss partly cushioned by unintended learning gains generated through inventor mobility.

Agglomeration as Innovation Resilience. The offsetting gains from inventor mobility are unlikely to arise evenly across space. Tacit knowledge is easier to redeploy where displaced inventors can find related firms, specialized labor markets, and dense collaboration networks. Agglomeration therefore provides a natural source of innovation resilience: it does not prevent the shock to sanctioned firms, but it helps the surrounding economy absorb displaced human capital and convert it into new patenting.

The county-level pattern in Figure 3 is consistent with this mechanism. Counties with fewer than 50 innovative firms generate almost no measurable spillover gains. As the local firm base thickens, the gains rise steadily, reaching more than eight additional patents in counties with over 400 innovative firms. This pattern echoes the standard agglomeration logic (Lucas Jr, 1988; Kerr and Kominers, 2015; Moretti, 2021): dense innovation clusters provide more potential employers, more related research teams, and more opportunities for knowledge exchange, making displaced inventors' expertise easier to absorb and redeploy.

[29] We measure y_{i,T_0} using the patent application count of inventor i at the focal sanctioned firm in 2017, the year before the first major U.S. sanctions in 2018. Consistent with the spillover-gain calculation, we measure post-sanction changes using 2021–2022. Because $\mathcal{L}_s$ may also include other sanction-related disruptions, such as input restrictions and supply-chain shocks, we interpret it as an upper-bound measure of organizational loss.

The spatial distribution in Figure 4 reinforces the same interpretation. Spillover gains are concentrated in China's major innovation clusters, especially the Beijing–Tianjin–Hebei region, the Yangtze River Delta, and the Pearl River Delta, where high-tech firms, research institutions, and specialized labor markets are densely co-located. Innovation-sparse regions in the northwest and northeast, by contrast, show limited gains, consistent with thinner local networks and weaker absorptive capacity. The geography of spillovers clarifies why the aggregate loss is only partly absorbed. Inventor mobility generates the largest learning effects in places where the local innovation system is dense enough to receive and redeploy specialized knowledge.

5 Conclusion

In recent years, deglobalization has increasingly become a widespread phenomenon. As China and the United States decouple economically and technologically, there is growing concern about the impact of U.S. sanctions on China's high-tech sector. This paper addresses a critical gap in the literature by assessing the impact of the U.S. Entity List sanctions through the lens of inventor mobility. We find that the Entity List sanctions increase the likelihood of job changes, while having no long-term impact on their innovation output. As talents redistribute across firms, we also find unintended knowledge diffusion in critical technological fields, which improves the productivity of nonmoving inventors in the receiving firms. In total, such spillovers account for about 27 percent of total patenting loss due to sanctions. This paper highlights the important role of local inventor mobility in shaping the consequences of economic sanctions.

References

Abebe, G., McMillan, M., and Serafinelli, M. (2022). Foreign direct investment and knowledge diffusion in poor locations. *Journal of Development Economics*, 158:102926.

Ahn, J., Kim, C., Li, N., and Manera, A. (2026). Knowledge diffusion through fdi: Worldwide firm-level evidence. *Journal of International Economics*, page 104205.

Akcigit, U., Baslandze, S., and Stantcheva, S. (2016). Taxation and the international mobility of inventors. *American Economic Review*, 106(10):2930–2981.

Akcigit, U., Grigsby, J., Nicholas, T., and Stantcheva, S. (2019). Taxation and innovation in the 20th century. *Economics*, 16(17).

Akcigit, U. and Melitz, M. J. (2022). International trade and innovation. In *Handbook of International Economics*, volume 5, pages 377–404. Elsevier.

Allen, G. C. (2022). Choking off China's access to the future of AI. Technical report, Center for Strategic and International Studies (CSIS).

Allen, G. C. (2024). The true impact of allied export controls on the u.s. and chinese semiconductor manufacturing equipment industries. Technical report, Center for Strategic and International Studies.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–2168.

Bahar, D., Choudhury, P., and Rapoport, H. (2020a). Migrant inventors and the technological advantage of nations. *Research Policy*, 49(9):103947.

Bahar, D., Choudhury, P., and Rapoport, H. (2020b). Migrant inventors and the technological advantage of nations. *Research Policy*, 49(9):103947.

Bahar, D., Hauptmann, A., ÃzgÃ¼zel, C., and Rapoport, H. (2022). Migration and Knowledge Diffusion: The Effect of Returning Refugees on Export Performance in the Former Yugoslavia. *Review of Economics and Statistics*, pages 1–50.

Bahar, D., Hauptmann, A., ÃzgÃ¼zel, C., and Rapoport, H. (2024). Migration and knowledge diffusion: The effect of returning refugees on export performance in the former yugoslavia. *Review of Economics and Statistics*, 106(2):287–304.

Bahar, D. and Rapoport, H. (2018a). Migration, knowledge diffusion and the comparative advantage of nations. *Economic Journal*, 128(612):F273–F305.

Bahar, D. and Rapoport, H. (2018b). Migration, knowledge diffusion and the comparative advantage of nations. *The Economic Journal*, 128(612):F273–F305.

Bailey, M., Gupta, A., Hillenbrand, S., Kuchler, T., Richmond, R., and Stroebel, J. (2021). International trade and social connectedness. *Journal of International Economics*, 129:103418.

Bajgar, M., Berlingieri, G., Calligaris, S., Criscuolo, C., and Timmis, J. (2020). Coverage and representativeness of orbis data. *OECD Working Paper*.

Balassa, B. (1965). Trade liberalisation and "revealed" comparative advantage 1. *The manchester school*, 33(2):99–123.

Bazzi, S., Gaduh, A., Rothenberg, A. D., and Wong, M. (2016). Skill transferability, migration, and development: Evidence from population resettlement in indonesia. *American Economic Review*, 106(9):2658–98.

Benguria, F., Choi, J., Swenson, D. L., and Xu, M. J. (2022). Anxiety or pain? the impact of tariffs and uncertainty on Chinese firms in the trade war. *Journal of International Economics*, 137:103608.

Branstetter, L. (2024). Export controls and u.s.-china technology competition in an interdependent world. Technical report, Brookings Institution. Brookings policy report; author affiliation: Carnegie Mellon University and NBER.

Braunerhjelm, P., Ding, D., and Thulin, P. (2020). Labour market mobility, knowledge diffusion and innovation. *European Economic Review*, 123:103386.

Bureau of Industry and Security (2022). Implementation of additional export controls: Certain advanced computing and semiconductor manufacturing items; supercomputer and semiconductor end use. *Federal Register*, 87(197):62186–62215.

Cai, J., Li, N., and Santacreu, A. M. (2022). Knowledge diffusion, trade, and innovation across countries and sectors. *American Economic Journal: Macroeconomics*, 14(1):104–45.

Callaway, B. and Sant'Anna, P. H. C. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.

Cantwell, J. and Mudambi, R. (2005). Mne competenceâcreating subsidiary mandates. *Strategic management journal*, 26(12):1109–1128.

Cao, Y., Farronato, C., Tartarini, F., and Ye, J. (2024). Us entity list restrictions slow the innovation of chinese firms and their us collaborators.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454.

Chen, T., Gao, H., and Wang, Y. (2021). Tariff uncertainty and firm innovation: Evidence from the US-China permanent normal trade relation. *Journal of Empirical Finance*, 62:12–27.

Cohn, J. B., Liu, Z. C., and Wardlaw, M. I. (2022). Count (and count-like) data in finance. *Journal of Financial Economics*, 146(2):529–551.

Council of the European Union (2024). Iran: Eu widens restrictive measures in view of iran support of the russian war of aggression against ukraine. Accessed April 22, 2026.

Dai, X. Y. and Wang, M. Q. (2024). Unintended effects of tax incentives on firms' strategic patenting. *Journal of Economic Behavior Organization*, 219:1–24.

Dunlevy, J. A. and Hutchinson, W. K. (1999). The impact of immigration on american import trade in the late nineteenth and early twentieth centuries. *Journal of Economic History*, 59(4):1043–1062.

European Commission (2022). European commission welcomes panel report in wto dispute against indonesia on raw materials. https://ec.europa.eu/commission/presscorner/api/files/document/print/en/ip_22_7314/IP_22_7314_EN.pdf. Accessed April 22, 2026.

European Commission (2025a). Eu adopts 17th sanctions package against russia. https://ec.europa.eu/commission/presscorner/api/files/document/print/en/ip_25_1249/IP_25_1249_EN.pdf. Accessed April 22, 2026.

European Commission (2025b). Export-related restrictions: Frequently asked questions concerning sanctions adopted following russia's military aggression against ukraine. https://finance.ec.europa.eu/document/download/586f9b62-50c3-46e3-9940-de73ae96a24b_en?filename=faqs-sanctions-russia-export-related-restrictions-russia_en.pdf. Accessed April 22, 2026.

Fajgelbaum, P. D. and Khandelwal, A. K. (2022). The economic impacts of the US-China trade war. *Annual Review of Economics*, 14(1):205–228.

Felbermayr, G., Morgan, T. C., Syropoulos, C., and Yotov, Y. V. (2025). Economic sanctions: Stylized facts and quantitative evidence. *Annual Review of Economics*, 17.

Felbermayr, G. and Sandkamp, A. (2020). The trade effects of anti-dumping duties: Firm-level evidence from China. *European Economic Review*, 122:103367.

Garred, J. and Yuan, S. (2025). Relocation from China (with Chinese characteristics). *Journal of Development Economics*, 176:103510.

Ge, C., Huang, K., and Png, I. P. (2016). Engineer/scientist careers: Patents, online profiles, and misclassification bias. *Strategic Management Journal*, 37(1):232–253.

Giorgi, J., Plunket, A., and de Waldemar, F. S. (2025). Inter-regional highly skilled worker mobility and technological novelty. *Research Policy*, 54(1):105113.

Gould, D. M. (1994). Immigrant links to the home country: empirical implications for us bilateral trade flows. *Review of Economics and Statistics*, pages 302–316.

Han, P., Jiang, W., and Mei, D. (2024). Mapping US-China technology decoupling: Policies, innovation, and firm performance. *Management Science*.

Hausmann, R. and Neffke, F. M. (2019). The workforce of pioneer plants: The role of worker mobility in the diffusion of industries. *Research Policy*, 48(3):628–648.

Hausmann, R., Schetter, U., and Yildirim, M. A. (2024). On the design of effective sanctions: The case of bans on exports to russia. *Economic Policy*, 39(117):109–158.

He, J. and Tian, X. (2013). The dark side of analyst coverage: The case of innovation. *Journal of Financial Economics*, 109(3):856–878.

He, J. J. and Tian, X. (2020). Institutions and innovation. *Annual Review of Financial Economics*, 12(1):377–398.

He, Z., Tong, T. W., Zhang, Y., and He, W. (2018). Constructing a chinese patent database of listed firms in china: Descriptions, lessons, and insights. *Journal of Economics Management Strategy*, 27(3):579–606.

Hsieh, C.-T., Hurst, E., Jones, C. I., and Klenow, P. J. (2019). The allocation of talent and US economic growth. *Econometrica*, 87(5):1439–1474.

Huang, Y., Lin, C., Liu, S., and Tang, H. (2023). Trade networks and firm value: Evidence from the US-China trade war. *Journal of International Economics*, 145:103811.

Hufbauer, G. C. and Jung, E. (2020). What's new in economic sanctions? *European Economic Review*, 130:103572.

Iaria, A., Schwarz, C., and Waldinger, F. (2018). Frontier knowledge and scientific production: Evidence from the collapse of international science. *The Quarterly Journal of Economics*, 133(2):927–991.

International Energy Agency (2024). Prohibition of the export of nickel ore. https://www.iea.org/policies/16084-prohibition-of-the-export-of-nickel-ore. Accessed April 22, 2026.

Jia, R., Roberts, M. E., Wang, Y., and Yang, E. (2024). The impact of usâchina tensions on us science: Evidence from the nih investigations. *Proceedings of the National Academy of Sciences*, 121(19):e2301436121.

Jiao, Y., Liu, Z., Tian, Z., and Wang, X. (2024). The impacts of the u.s. trade war on chinese exporters. *The Review of Economics and Statistics*, 106(6):1576–1587.

Jin, Z., Kermani, A., and McQuade, T. (2025). Native-immigrant entrepreneurial synergies. Technical report, National Bureau of Economic Research.

Johnson, M. S., Lipsitz, M., and Pei, A. (2023). Innovation, inventor mobility, and the enforceability of noncompete agreements. NBER Working Paper 31487, National Bureau of Economic Research.

Kaiser, U., Kongsted, H. C., and Rønde, T. (2015). Does the mobility of R&D labor increase innovation? *Journal of Economic Behavior & Organization*, 110:91–105.

Kang, Y., Liu, R., and Yang, B. (2025). Bridging the ivory tower and industry: How university science parks promote university-industry collaboration? *Research Policy*, 54(6):105241.

Kerr, W. R. (2018). Heterogeneous technology diffusion and ricardian trade patterns. *World Bank Economic Review*, 32(1):163–182.

Kerr, W. R. and Kominers, S. D. (2015). Agglomerative forces and cluster shapes. *Review of Economics and Statistics*, 97(4):877–899.

Kong, D., Liu, C., Narayan, P. K., and Sharma, S. S. (2024). The US-China trade war and corporate innovation: Evidence from China. *Financial Management*.

Levenshtein, V. (1966). Binary codes capable of correcting deletions, insertions, and reversal. *Soviet Physics Doklady*, 10(8):707–710.

Li, L. S. Z., Liu, Y., and Yuan, J. (2022). The effect of the US-China trade war on chinese corporate innovation: A curse or a blessing? *The World Economy*.

Liangyu (2018). U.S. suspends export control deal with China's ZTE. Retrieved June 20, 2025.

Lin, J. Y., Liu, S., Xu, M., and Zhang, H. (2024). Learn in order to practice: The effect of political rotation on local comparative advantage in china. *Available at SSRN*.

Lind, N. and Ramondo, N. (2023). Global innovation and knowledge diffusion. *American Economic Review: Insights*, 5(4):494–510.

Liu, G., Liu, Y., Yang, J., and Zhang, Y. (2024). Labor contract law and inventor mobility: evidence from china: G. liu et al. *Journal of Population Economics*, 37(1):7.

Lu, Y., Wu, J. S., and Yang, Y. (2026). Does foreign investors' information access vary with geopolitical tensions? evidence from corporate conference calls. *Journal of Accounting Research*, 64(1):417–475.

Lucas Jr, R. E. (1988). On the mechanics of economic development. *Journal of monetary economics*, 22(1):3–42.

Mayda, A.-M., Parsons, C., Pham, H., and Vézina, P.-L. (2022). Refugees and foreign direct investment: Quasi-experimental evidence from us resettlements. *Journal of Development Economics*, 156:102818.

Melero, E., Palomeras, N., and Wehrheim, D. (2020). The effect of patent protection on inventor mobility. *Management Science*, 66(12):5485–5504.

Miguelez, E. and Temgoua, C. N. (2020). Inventor migration and knowledge flows: A two-way communication channel? *Research Policy*, 49(9):103914.

Moretti, E. (2021). The effect of high-tech clusters on the productivity of top inventors. *American Economic Review*, 111(10):3328–3375.

Moretti, E. and Wilson, D. J. (2017). The effect of state taxes on the geographical location of top earners: Evidence from star scientists. *American Economic Review*, 107(7):1858–1903.

Morgan, T. C., Syropoulos, C., and Yotov, Y. V. (2023a). Economic sanctions: Evolution, consequences, and challenges. *Journal of Economic Perspectives*, 37(1):3–30.

Morgan, T. C., Syropoulos, C., and Yotov, Y. V. (2023b). Economic sanctions: Evolution, consequences, and challenges. *Journal of Economic Perspectives*, 37(1):3–29.

Moser, P., Voena, A., and Waldinger, F. (2024). German jewish Émigrés and US invention. *American Economic Review*, 104(10):3222–3255.

Nayak, D., Moreira, S., and Mudambi, R. (2025). Restrictive immigration policies and mne innovation. *Journal of International Business Studies*, 56(1):84–104.

Oettl, A. and Agrawal, A. (2008). International labor mobility and knowledge flow externalities. *Journal of international business studies*, 39(8):1242–1260.

Pellegrina, H. S. and Sotelo, S. (2025). Migration, specialization, and trade: Evidence from brazil's march to the west. *Journal of Political Economy*, 133(12):3993–4049.

Prato, M. (2024). The global race for talent: Brain drain, knowledge transfer, and growth. *The Quarterly Journal of Economics*, 140(1):165–238.

Qin, J. Y. (2019). Forced technology transfer and the US-China trade war: Implications for international economic law. *Journal of International Economic Law*, 22(4):743–762.

Reinmuth, K. and Rockall, E. (2025). Innovation through inventor mobility: Evidence from noncompete agreements. *American Economic Journal: Applied Economics*, 17(4):1–35.

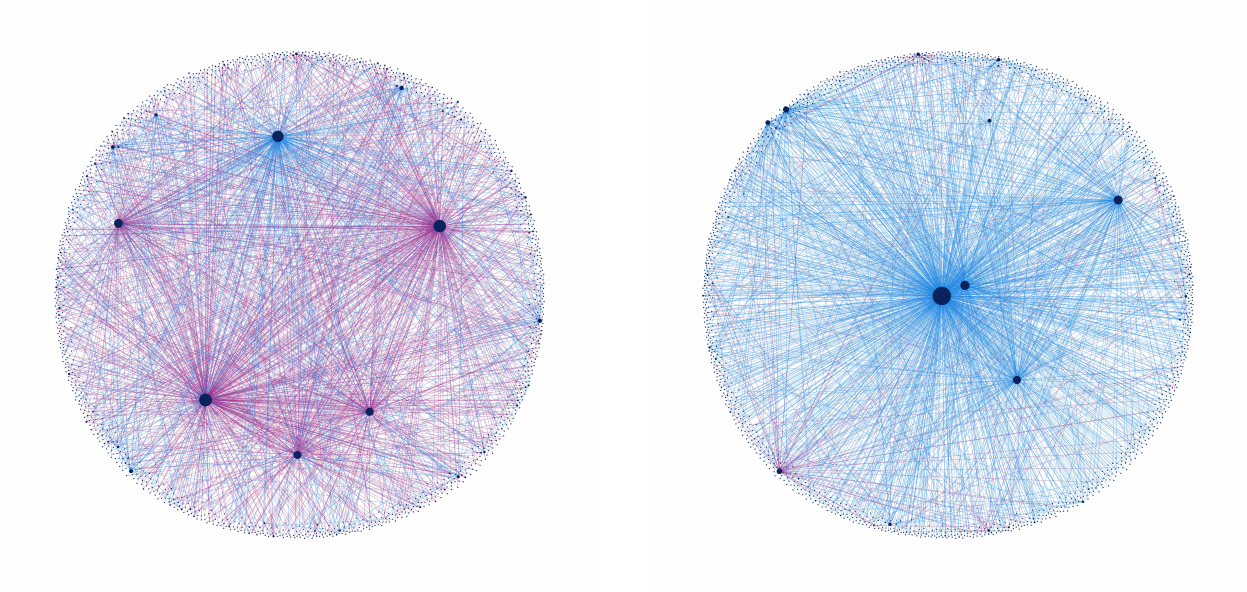
Soete, L. (1987). The impact of technological innovation on international trade patterns: the evidence reconsidered. *Research policy*, 16(2-4):101–130.

Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.

Wang, Y. B., Stuart, T., and Li, J. Z. (2021). Fraud and innovation. *Administrative Science Quarterly*, 66(2):267–297.

Figures and Tables

Figure 1: Inventor Mobility Dynamics Before vs. After Sanction

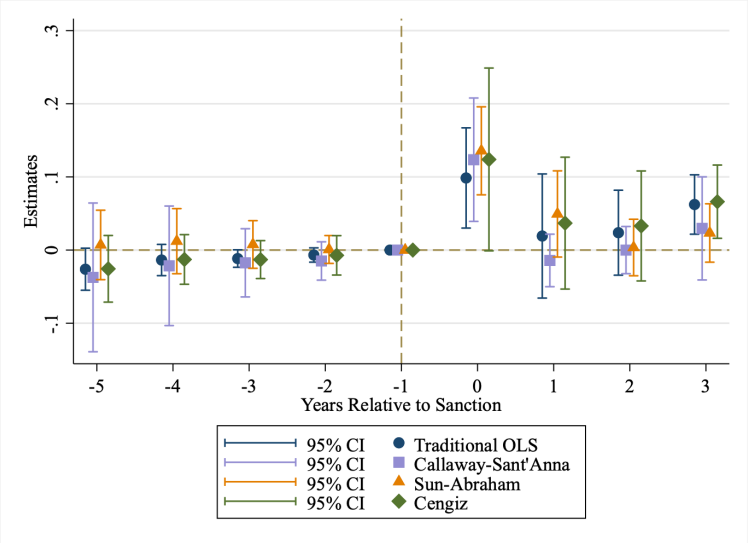


(a) Before Sanction

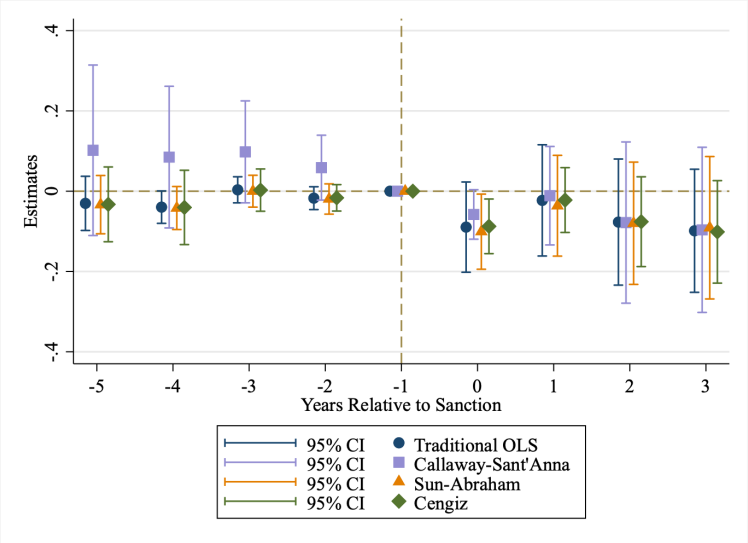
(b) After Sanction

Notes. This figure presents the inventor flow between a subsample of firms before and after the Entity List sanction. Each node denotes a firm, with the node size proportional to the number of mobile inventors. The edges reflect the inventor flow across firms. Royal blue edges capture inventor mobility from sanctioned to non-sanctioned firms, whereas purple edges denote all other inter-firm mobility flows. To alleviate computational costs and improve network readability, we use a subsample of 40 randomly selected firms, evenly split between 20 non-sanctioned and 20 sanctioned firms, where the former are those on the Entity List since 2018, and the latter are firms never on the Entity List up to 2022. Panel (a) depicts patterns of inventor job reallocation prior to the imposition of U.S. sanctions (2018), while panel (b) shows the corresponding patterns in the postâsanction period. We employ the force-directed layout via Fruchterman-Reingold algorithm to visualize the network of inventor mobility.

Figure 2: Event Study



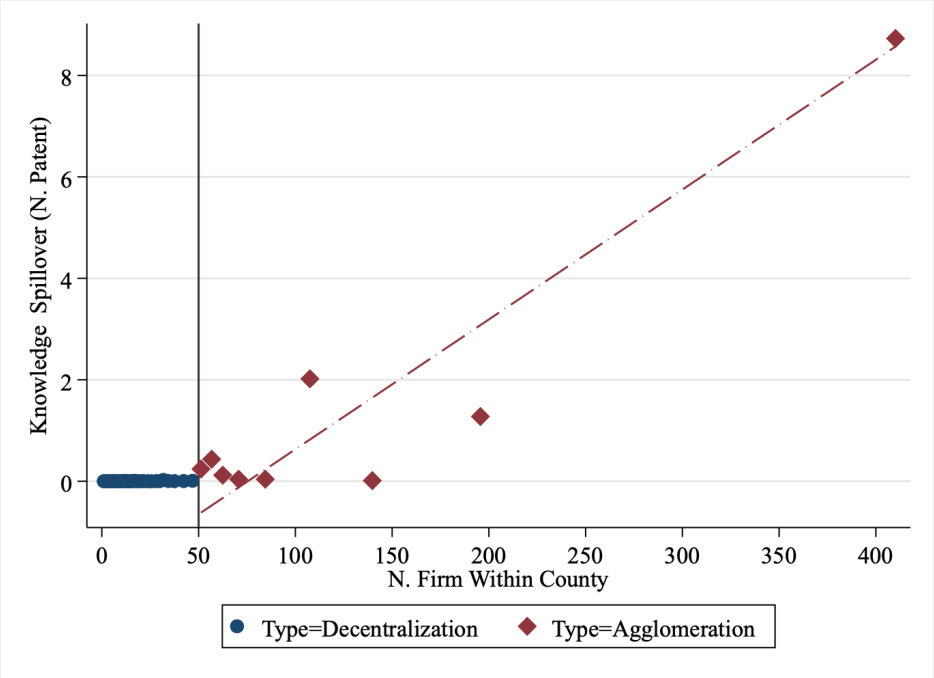
(a) Inventor job switching



(b) Patent Application

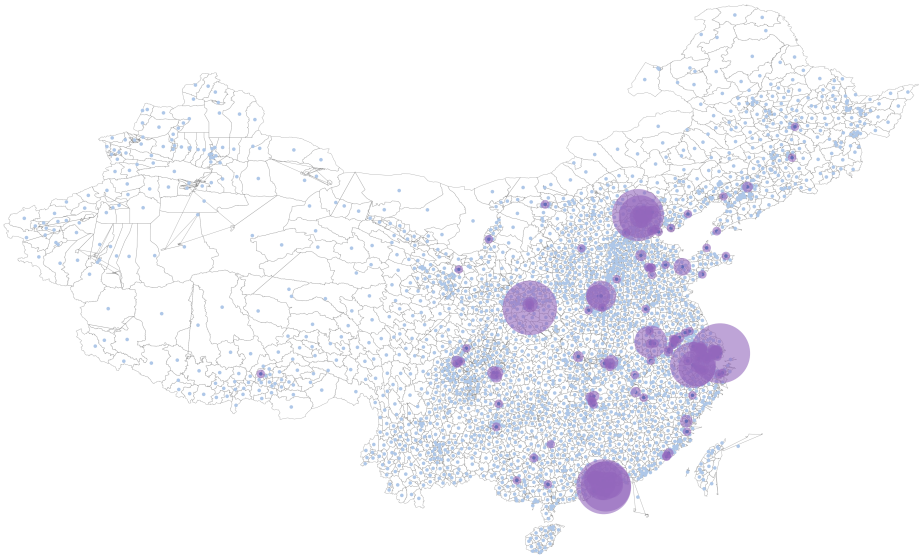
Notes. This figure reports coefficients from the event study that examines the dynamic impact of the Entity List trade sanction on inventor job changes and innovation performance. The unit of observation is an inventor-year pair. The treated group is inventors who have ever experienced sanctions, and the control group is inventors who do not experience sanctions during the sample period. We report coefficient estimates using OLS specification (navy dot), following Callaway and Sant'Anna (2021) (square), Sun and Abraham (2021) (triangle), and Cengiz et al. (2019) (diamond), surrounded by the 95% confidence interval. In Panel (a), the dependent variable is the probability of an inventor changing jobs in a given year. In Panel (b), the dependent variable is the log number of patent applications by an inventor in a given year. The regressions include controls for firm characteristics, inventor characteristics, inventor fixed effects, and year and firm fixed effects. Individual controls include the log years since first patent application, the average size of the co-authors in a given year, and whether inventor has actively filed patents in the past three years. Firm controls include the number of IPC technology classes on patents applied by a firm in a given year and whether firm j has patented in the past three years. Standard errors are clustered at the firm level.

Figure 3: County-Level Knowledge Diffusion Gains and Firm Count



Notes. This binscatter plot shows the relationship between the number of innovative firms in a county (x-axis) and county-level knowledge spillover gains (y-axis). Number of bins = 100. Counties with fewer than 50 innovative firms are classified as "decentralized" (low-agglomeration) regions, while those with 50 or more firms are classified as "agglomeration" regions. Knowledge spillover gains in patents are calculated using Equation 8 in Section 4 for each firm and aggregated to counties. The dashed line represents the linear fit for the agglomeration subsample.

Figure 4: Spatial Distribution of County-Level Knowledge Diffusion Gains



Notes. The map plots the geographic distribution of spillover gains across over 2,800 counties in China. Knowledge spillover gains in patents are calculated using Equation 8 in Section 4 for each firm. Light blue points denote counties with no measurable spillover gains, while purple points represent counties with positive gains; the size of purple bubbles around points is proportional to the magnitude of the county-level spillover gain.

Table 1: Summary Statistics

	Mean	SD	Min	Max
Probability of i changing jobs in year t	0.12	0.32	0.00	1.00
N. patents applied in year t	3.11	6.22	1.00	937.00
Ever employed in a sanctioned firm	0.01	0.11	0.00	1.00
Inventor-year controls				
Log years since first patent application	1.10	0.84	0.00	3.14
Average coauthor size in year t	5.68	3.85	1.00	75.00
Whether filed a patent in the past three years	0.68	0.47	0.00	1.00
Firm-year controls				
N. IPCs a firm applied in year t	10.53	15.14	1.00	89.00
N. patents firm patented in the past three years	0.90	0.30	0.00	1.00

Notes. The table reports summary statistics of main variables in our inventor-firm-year analysis. N = 4,939,429.

Table 2: Mobility and Innovation Effect of U.S. Sanctions

Dependent variable:	FlowOut (probability of switching jobs)			ln(number of patents applied)		
	(1)	(2)	(3)	(4)	(5)	(6)
A. All Inventors						
Sanction	0.067*	0.094***	0.081***	0.009	0.063	-0.066
	(0.04)	(0.03)	(0.03)	(0.02)	(0.05)	(0.07)
Inventor Controls	No	Yes	Yes	No	Yes	Yes
Firm Controls	No	No	Yes	No	No	Yes
Inventor FE	No	Yes	Yes	No	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,961,605	4,939,429	4,939,429	4,961,605	4,939,429	4,939,429
R^2	0.256	0.597	0.609	0.218	0.600	0.604
Dependent variable:	FlowOut (probability of switching jobs)			ln(number of patents applied)		
	(1)	(2)	(3)	(4)	(5)	(6)
B. Focusing on Mobile Inventors						
Sanction	0.116***	0.260***	0.208***	0.006	0.028	-0.067
	(0.03)	(0.04)	(0.04)	(0.03)	(0.05)	(0.04)
Inventor Controls	No	Yes	Yes	No	Yes	Yes
Firm Controls	No	No	Yes	No	No	Yes
Inventor FE	No	Yes	Yes	No	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,671,181	1,645,124	1,645,124	1,671,181	1,645,124	1,645,124
R^2	0.259	0.491	0.530	0.241	0.579	0.583

Notes. This table examines the impact of the Entity List trade sanctions on inventor job changes and innovation performance. The unit of observation is an inventor-year pair. The dependent variable is the probability of an inventor changing jobs in a given year in Columns 1 to 3; the log number of patent applications by an inventor in a given year in Columns 4 to 6. "Sanction" is a dummy variable taking one for all years after inventor was first employed by a firm included on the Entity List post-2018 and zero otherwise. Individual controls include the log years since first patent application, the average size of the co-authors in a given year, and whether inventor has actively filed patents in the past three years. Firm controls include the number of IPC technology classes on patents applied by a firm in a given year and whether firm j has patented in the past three years. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.

Table 3: U.S. Sanctions and Inventors' Inter-Regional Mobility

Dependent variable:	All Inventors		Mobile Inventors	
	(1)	(2)	(3)	(4)
FlowOut (probability of switching jobs)	Inter-county	Inter-city	Inter-county	Inter-city
Sanction	0.016***	0.014***	0.044***	0.039***
	(0.01)	(0.01)	(0.01)	(0.01)
Inventor Controls	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes
Inventor FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	4,939,429	4,939,429	1,645,124	1,645,124
R^2	0.453	0.455	0.438	0.439

Notes: This table reports the effects of sanctions on inventor mobility across regions using specification in Column 3 Table 2 with different dependent variables. The dependent variable is a dummy variable indicating whether the inventor changes job to a different county in Columns 1 and 3, and to a different city in Columns 2 and 4. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.

Table 4: Heterogeneous Effect of U.S. Sanctions

Dependent variable: FlowOut (probability of switching jobs)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Sanction	0.114***	0.090***	0.118***	0.740**	0.205***	0.009	0.030**	0.108
	(0.04)	(0.03)	(0.03)	(0.33)	(0.06)	(0.01)	(0.01)	(0.17)
Sanction*FirstAuthor	-0.056***							
	(0.02)							
Sanction*Citation		-0.003***						
		(0.00)						
Sanction*PatentFamilySize			-0.017***					
			(0.00)					
Sanction*KnowledgeSize				-0.059**				
				(0.03)				
Sanction*TeamSize					-0.038***			
					(0.01)			
Sanction*Dependency_US						0.885***		
						(0.33)		
Sanction*Dependency_USOther							0.011	
							(0.12)	
Sanction*Dependency_Other								-0.083
								(0.19)
Inventor Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Inventor FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,939,429	4,939,429	4,939,429	4,939,429	4,939,429	1,463,816	1,463,816	1,463,816
R^2	0.609	0.609	0.609	0.610	0.609	0.812	0.812	0.812

Notes: This table presents heterogeneous effects from the specification corresponding to Column 3 in Table 2 with the "Sanction" dummy interacted with inventor or firm characteristics. Three measures of inventor ability are considered: "FirstAuthor" flags first authors on patents granted in the last three years, "Citation" is the average number of citations per patent in the last three years, and "PatentFamilySize" is the average number of family members per inventor's patent in the last three years. Firm characteristics include "KnowledgeSize", the cumulative number of inventors on patents assigned to the focal firm by year t, which accounts for the size of knowledge production function (Cantwell and Mudambi, 2005); and "TeamSize", the average number of co-inventors across all patents of the focal firm in year t. "Dependency_US" is the share of imports from the U.S. in sanction-related industries (communication, EEE, transport, and media & broadcasting) among total imports. "Dependency_USOther" is the share of imports from the U.S. in non-sanction-related industries. "Dependency_Other" is the share of imports aside from the U.S. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.

Table 5: Knowledge Diffusion of Inventor Mobility on Destination Firm Technology Advantage

Dependent variable: ln(RTA)	(1)	(2)	(3)	(4)	(5)	(6)
A. Overall Effect of Inventor Mobility						
Number of inventor influx (x100)	0.568***	0.035***	0.040***			
	(0.17)	(0.01)	(0.01)			
1{Influx > 0}				0.367***	0.024***	0.026***
				(0.01)	(0.01)	(0.01)
Firm-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-IPC FE	No	Yes	Yes	No	Yes	Yes
IPC-Year FE	Yes	No	Yes	Yes	No	Yes
Observations	982,047	982,047	982,047	982,047	982,047	982,047
R^2	0.727	0.912	0.914	0.728	0.912	0.914
Dependent variable: ln(RTA)	(1)	(2)	(3)	(4)	(5)	(6)
B. Heterogeneous Effect by Firm of Origin						
1{Influx from sanctioned firms > 0}	0.856***	0.050**	0.114***			
	(0.08)	(0.02)	(0.02)			
1{Influx from only non-sanctioned firms > 0}				0.281***	0.018***	0.015**
				(0.01)	(0.01)	(0.01)
Firm-Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm-IPC FE	No	Yes	Yes	No	Yes	Yes
IPC-Year FE	Yes	No	Yes	Yes	No	Yes
Observations	982,047	982,047	982,047	982,047	982,047	982,047
R^2	0.727	0.912	0.914	0.727	0.912	0.914

Notes. This table reports knowledge diffusion of inventor mobility to the destination firms. The dependent variable in both Panel A and B is the log value of a firm's specific revealed technological advantage (RTA) in class p in year t. In Panel A, "Number of inventor influx" is the number of inventors originating from firms with comparative advantage in technology class p into the destination firm, scaled down by a factor of 100. "**1**{Influx > 0}" is a dummy variable indicating at least one inflow of such inventors to a destination firm in year t. In Panel B, "**1**{Influx from sanctioned firms > 0}" indicates whether there is at least one inventors originating from a sanctioned firm which has comparative advantage in area p moving into the destination firm in year t. "**1**{Influx from only non-sanctioned firms > 0}" flags the case when the inventor only comes from non-sanctioned firms. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.

Table 6: Effect of Inventor Mobility on Technological Similarity within Industry

Dependent Variable: Similarity between firm technology portfolio and industry counterparts	(1)	(2)	(3)	(4)
A. Firm, Year, and Industry Fixed Effects				
Number of inventor influx (×100)	0.015***			
	(0.00)			
1{Influx > 0}		0.016***		
		(0.00)		
1{Influx from sanctioned firms > 0}			0.056***	
			(0.01)	
1{Influx from only non-sanctioned firms > 0}				0.014***
				(0.00)
Firm Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Industry FE	Yes	Yes	Yes	Yes
Observations	521,578	521,578	521,578	521,578
R^2	0.606	0.606	0.606	0.606
	(1)	(2)	(3)	(4)
B. Firm and Industry-Year Fixed Effects				
Number of inventor influx (×100)	0.012***			
	(0.00)			
1{Influx > 0}		0.015***		
		(0.00)		
1{Influx from sanctioned firms > 0}			0.042***	
			(0.01)	
1{Influx from only non-sanctioned firms > 0}				0.013***
				(0.00)
Firm Controls	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry-Year FE	Yes	Yes	Yes	Yes
Observations	521,575	521,575	521,575	521,575
R^2	0.619	0.619	0.619	0.619

Notes. This table reports competition effect of inventor mobility on the technological similarity within industry. The dependent variable is cosine similarity between focal firm technology portfolio and its industry counterparts. "Number of inventor influx" is the number of inventors originating from firms with comparative advantage in at least one technology class into the destination firm, scaled down by a factor of 100. "**1**{Influx > 0}" is a dummy variable indicating at least one inflow of such inventors to a destination firm in year t. "**1**{Influx from sanctioned firms > 0}" indicates whether there is at least one inventors originating from a sanctioned firm which has comparative advantage in at least one area moving into the destination firm in year t. "**1**{Influx from only non-sanctioned firms > 0}" flags the case when the inventor only comes from non-sanctioned firms. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.

Table 7: Spillover Effects of Inventor Mobility on New Colleagues Innovation Performance

Dependent variable: ln(number of patents applied by colleague)	(1)	(2)	(3)	(4)
Number of inventor influx (x100)	0.061*			
	(0.03)			
1{Influx > 0}		0.023***		
		(0.01)		
1{Influx from sanctioned firms > 0}			0.125***	
			(0.02)	
1{Influx from only non-sanctioned firms > 0}				0.011
				(0.01)
Firm-Year FE	Yes	Yes	Yes	Yes
Firm-IPC FE	Yes	Yes	Yes	Yes
IPC-Year FE	Yes	Yes	Yes	Yes
Observations	586,086	586,086	586,086	586,086
R^2	0.927	0.927	0.927	0.927

Notes. This table reports knowledge diffusion of inventor mobility on the innovation performance of non-moving colleagues in the destination firm. The dependent variable is the log number of patent applications filed by colleague in the receiving firm in technology class p in year t. "Number of inventor influx" is the number of inventors originating from firms with comparative advantage in technology class p into the destination firm, scaled down by a factor of 100. "**1**{Influx > 0}" is a dummy variable indicating at least one inflow of such inventors to a destination firm in year t. "**1**{Influx from sanctioned firms > 0}" indicates whether there is at least one inventors originating from a sanctioned firm which has comparative advantage in area p moving into the destination firm in year t. "**1**{Influx from only non-sanctioned firms > 0}" flags the case when the inventor only comes from non-sanctioned firms. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.

Table 8: Back-of-the-Envelope Decomposition of Aggregate Innovation Effects

Adjustment Margin	Contribution to ΔP	Share of indirect loss
Movers' own patenting response	≈ 0	–
Receiving-firm spillover gains	+235.98	26.94%
Indirect losses at origin firms	−875.88	100.00%
Net effect	−639.90	−73.06%

Notes. The table reports annual patent-equivalent changes. Receiving-firm spillover gains are computed as $\sum_{n \in N} \Phi_n = 235.98$, and indirect losses at sanctioned origin firms as $\sum_{s \in S} \mathcal{L}_s = 875.88$. The last column normalizes each component by the estimated indirect loss at origin firms, 875.88. Thus, 26.94% = 235.98/875.88 and −73.06% = −639.90/875.88. A negative value indicates a net loss relative to the estimated indirect loss borne by sanctioned origin firms.

Appendix

A Details for Patent Cleaning Procedure and Mobility Algorithm

We begin with the patent cleaning procedure, describing how we restrict and standardize the patent dataset for our analysis. We then move on to the inventor mobility algorithm, detailing the construction of both intra-county and cross-county mobility trajectories using standardized inventor records, detailed geographic information, and technical classifications. To confirm the reliability of our identification strategy, we supplement these algorithmic steps with illustrative case studies and examine the overall distribution of mobility events in our sample.

A.1 Patent Cleaning Procedure

We begin by restricting the patent to applications filed by domestic assignees only, excluding all patents with foreign applicants. Given our focus on the domestic reallocation and diffusion of human capital induced by U.S. sanctions, this restriction ensures that we trace inventor mobility within the Chinese innovation system and avoids conflating cross-border moves with domestic job changes. Next, we restrict the sample to patents with a single corporate assignee, so that each inventor-year can be assigned to an employer with minimal ambiguity. This restriction retains 92.74% of all firm-filed patents, suggesting limited sample loss.

A major challenge in constructing the assignee-firm link is firm name disambiguation, which arises from frequent firm name changes and minor spelling variations across patent applications. For example, *Piesat Information Technology Co., Ltd.*, a high-tech enterprise founded in 2008 that specializes in satellite remote sensing and geographic information systems, has appeared under multiple names in patent records. These include its original name, *Beijing Huadimap Information Technology Co., Ltd.*, its intermediate name since year 2012 *Beijing Piesat Information Technology Co., Ltd.*, and its current name from year 2020 onward *Piesat Information Technology Co., Ltd.*, alongside a minor variant spelling of the intermediate name. Despite these differences, all refer to the same legal entity, as verified by its Taxpayer Identification Number: 91110108671708384H.

To address this firm-type disambiguation challenge, we rely on the Taxpayer Identification Number (TIN) provided in the incoPat database, supplemented with cross-verification using *Aiqicha* and *Qichacha*, the widely used Chinese commercial databases that provide official business registration records – including TIN codes, historical name changes, and corporate

profiles for legal entities registered in China. We match all variant assignee names to their corresponding TIN, creating a single, consistent identifier for each legal entity. This step is critical for our subsequent inventor mobility algorithm, as it ensures that we do not incorrectly identify moves between different aliases of the same firm as true job changes.

Note that an alternative identifier in the Chinese context is Business Registration Number (e.g., 110108010770237 for the example "*Piesat Information Technology Co., Ltd.*" above), but it has lower coverage in our patent dataset. The TIN, by contrast, provides a unique and stable identifier across different names and forms, allowing us to link most of patent applications belonging to the same firm regardless of name changes or minor spelling differences. Specifically, after cleaning and disambiguation, our patent database contains 596,866 unique firm assignees (enterprises). Between 2014 and 2022, these assignees filed a total of 6,121,145 patents as the sole applicant under the China National Intellectual Property Administration (CNIPA), covering approximately 94% of all single-firm-assignee patents in the same period. Put differently, only roughly 6% of single-firm-assignee patents were excluded because the corresponding firm could not be matched to a valid TIN code, preserving the representativeness of our patent sample for subsequent analysis.

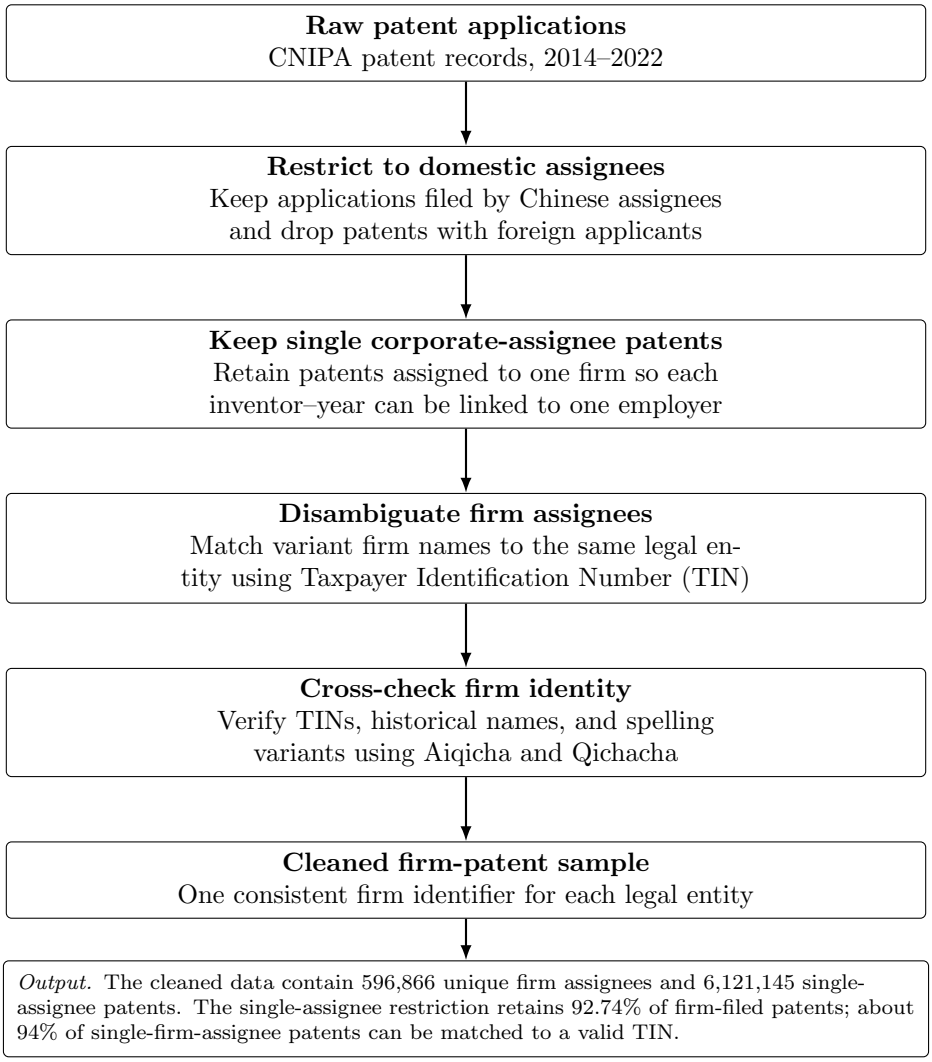
The patent-cleaning algorithm is summarized in Appendix Figure A1. We restrict the sample to domestic, single-assignee patents and harmonize firm identities using Taxpayer Identification Numbers, supplemented by business registry checks. This produces a stable employer identifier for each inventor–year observation and reduces false mobility events caused by firm-name changes or noisy assignee records.

A.2 Inventor Mobility Algorithm.

Before formally introducing the mobility algorithm, we first detail the inventor-level cleaning procedure. While the firm-level cleaning described above ensures accurate assignee-firm links, inventor-level records may introduce additional sources of ambiguity that require careful standardization – such as inventor names in inconsistent formats. Standardizing these records is therefore essential to construct reliable career trajectories and identify valid mobility events.

Inventor Cleaning Procedure. Under the patent application practices of CNIPA, a small share of inventors choose not to disclose personal names, marked as blank entries or labels indicating anonymous application. We eliminate all such unnamed inventors from the sample. Besides, some inventor records contain redundant symbols, dots and numerical prefixes attached to full names, which essentially refer to identical individuals. We standardize all name formats uniformly, such as revising ".Guo Jianxin" and "2 Chen Guangquan"

Figure A1: Patent Cleaning Procedure



Notes. This flowchart summarizes the firm-level patent cleaning procedure. The key step is to construct a stable legal-entity identifier using TINs, so that changes in firm names or spelling variants are not misclassified as inventor mobility.

into standard full name forms. This standardization effectively prevents misjudgment of personnel turnover and false identification of distinct inventors.

We further extract city-level and county-level residential information of inventors based on corporate geographic records from incoPat and cross-check with business inquiry platforms such as *Aiqicha* and *Qichacha*. Inventors and patent applications originating from Hong Kong, Macao and Taiwan are excluded, narrowing our research scope to mainland China. Meanwhile, we sort out core technological domains each inventor engages in according to patent IPC classification codes. Taken together, standardized inventor names, detailed geographic locations and technical fields jointly serve as core identification criteria to track cross-firm inventor mobility.

Intra-county Mobility Algorithm. Following prior work (Melero et al., 2020; Liu et al., 2024), we infer inventor mobility from changes in an inventor's corporate assignee across patent applications. We code a job change when an inventor patents at different firms in different years, assigning the move to the first year in which the inventor appears at the new firm. For example, an inventor observed at firm A in 2017 and firm B in 2018 is recorded as moving from A to B in 2018. We first trace inventors within the same county to identify intra-regional moves. Specifically, we adopt a name-city-county triple matching strategy for inventor identification in order to eliminate recognition errors. We illustrate this procedure with a typical inventor case: inventor *Pan Yishan* conducted all her innovative activities within *Baiyun District, Guangzhou*. Her patents were exclusively affiliated with *Guangzhou Discovery Cosmetics Trade Co., Ltd.* (TIN: 914401015566588876) up to 2020, and have been assigned to *Guangzhou Fajue Biotechnology Co., Ltd.* (TIN: 91440101MA5APWC139) since 2021. Accordingly, she experienced one intra-regional job mobility event occurring in 2021.

The joint adoption of city- and county-level geographic information is critical for identification accuracy, given the frequent duplication of county-level district names across different Chinese cities. For instance, *Gulou* District appears in four distinct prefecture-level cities: *Nanjing* (Jiangsu Province), *Xuzhou* (Jiangsu Province), *Kaifeng* (Henan Province), and *Fuzhou* (Fujian Province). Similarly, several prefecture-level cities host districts named *Xinhua* District, *Chengguan* District, and *Xincheng* District. As such, solely relying on county-level labels would mistakenly match inventors with identical names across geographically distant, identically named districts, resulting in false inventor duplication and spurious mobility observations. By anchoring identification at the city-county level, we precisely lock each inventor's fixed geographic scope and resolve cross-city homonymous district confounding issues.

Beyond name uniqueness, fine-grained county-level geographic matching further improves

the precision of the mobility identification. Empirically, it is generally rare for a single inventor to work for multiple distinct firms within the same county in the same year. This feature ensures that any cross-year assignee switch observed within a fixed city-county unit reliably reflects genuine job mobility, rather than concurrent multi-firm employment or data noise. Altogether, combined with standardized inventor names and refined city-county geographic information, this matching rule effectively avoids misidentification of homonymous inventors and achieves precise identification of intra-regional inventor mobility.

Inter-county Mobility Algorithm. In terms of inter-county mobility, we first track inventors who appear only once each year in patent archives, and integrate scattered patent records to reconstruct their full career paths. Take inventor *Cai Yongbang* as an example. Before 2019, his patents were exclusively assigned to *Guangdong Sinid Technology Co., Ltd.* (TIN: 914406000615333374) in *Sanshui District, Foshan*. From 2019 onward, he filed patents for *Dongguan Glorystar Laser Technology Co., Ltd.* (TIN: 914419005536016661) located in *Dongcheng District, Dongguan*. This clear shift in affiliated company and geographic location captures a typical cross-regional job mobility event.

Next, we additionally identify cross-regional inventor mobility by integrating standardized inventor names and their IPC technological classification. We adopt the broadest one-digit IPC sector code for technical positioning. Individual inventors commonly engage in multiple refined technological subclasses throughout their careers, even after workplace relocation. Taking inventor *Pan Yishan* as an example, her research focus centers on cosmetics and skincare formulations. Before the job transition, her patented work covered cosmetic compositions (A61K), hair care preparations (A61Q) and dermatological therapeutic agents (A61P). After moving to the new firm in 2021, she expanded her scope while still staying within the same broad technological boundary, namely IPC Sector A. Put differently, all aforementioned subclasses fall into Sector A, which represents human necessities. The same applies to inventor *Cai Yongbang*, whose research activities belong to IPC Sector B covering operations and transportation. To sum up, these cases demonstrate the merits of adopting sector-level codes for inventor trajectory recognition: the consistent broad technological positioning effectively connects working records across different firms, enabling reliable tracking of cross-regional inventor mobility.

It is worth noting that if an inventor has multiple affiliations which specialize in a specific technological area in a given year, we assign the inventor to the firm with which they have the longest tenure over their patenting career. To avoid misclassification, we drop the small number of cases in which multiple firms are tied for this primary affiliation. We illustrate this with inventor(s) *Du Jianjun*, whose work falls within IPC Sector H (electricity, including electron-

ics and communications). In the early years of his career (2009-2012), all of his patents were filed with *Shenzhen Chao San Technology Co., Ltd.* (TIN: 91440300774117464B), located in *Longgang District, Shenzhen*. In year 2012, another inventor also named *Du Jianjun*, working in the same technological sector H, appeared in patents assigned to *Shenzhen Xin'an Communication Technology Co., Ltd.* (TIN: 9144030057767320XH) in *Pingshan District, Shenzhen*. Under our rule of assigning each inventor to the firm with the longest tenure, we treat the First *Du Jianjun* as remaining with his original employer and continuing his patenting activity at *Shenzhen Chao San Technology Co., Ltd.* in 2012, rather than erroneously coding a mobility event. In this way, the Second *Du Jianjun*, appearing at *Shenzhen Xin'an Communication Technology Co., Ltd.* in 2012, is treated as a distinct inventor.

Such assignment rule is validated by subsequent patent records: from year 2013 onward, the First *Du Jianjun* continued to file patents annually with his original firm, confirming no actual job change occurred. The Second *Du Jianjun* also filed patents consistently with *Shenzhen Xin'an Communication Technology Co., Ltd.* in years 2016 and 2017, confirming he was indeed a separate individual. Overall, this tenure-based assignment strategy helps eliminate false positive mobility events by prioritizing long-term affiliation records, and thus provides a conservative estimate of inventor mobility that avoids overstating cross-firm transitions in cases of homonymous inventors working in the same technological sector.

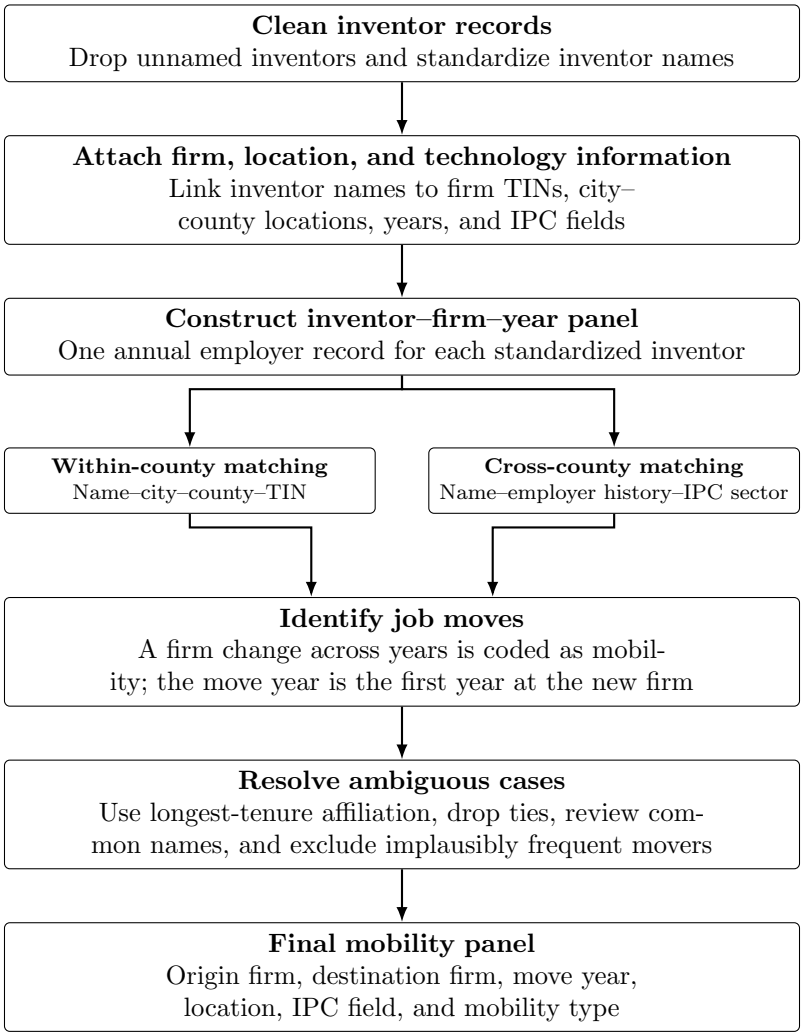
While the single-sector rule resolves majority of the ambiguity in cross-regional mobility, several multi-sector cases may require further checks. A typical example involves two inventors both named *Luo Wenjuan*. The first inventor was affiliated with *Beijing Didi Unlimited Technology Development Co., Ltd.* (TIN: 911101080673385595) in *Haidian District, Beijing*, with patents primarily falling into IPC Sector G (physics, including computing and telecommunications). The second inventor, meanwhile, was associated with *Shanghai NIO Investment Co., Ltd.* (TIN: 91310000329555773R) in *Jiading District, Shanghai,* whose patents were concentrated in IPC Sector B (operations and transportation). Neither inventor filed patents with these two original firms after 2021. In 2022, however, a new inventor also named *Luo Wenjuan* appeared at *Zhejiang Tongxin Engineering Project Management Co., Ltd.* (TIN: 91330304MA2HCKKC1A), whose patent portfolio spans both IPC Sector G and Sector B. Given that both original inventors ceased filing patents at their prior firms in the same year and the new inventor's combined technological scope matches both previous profiles, we conservatively treat both as moving to the third firm in 2022.

For extremely common Chinese names such as *Zhang Wei, Wang Wei*, and *Li Wei*, additional checks are needed to reduce false positives in mobility identification. Take *Wang Wei* as an example: given the extreme prevalence of this name, multiple distinct inventors with the same name may appear even within the same IPC sector in the same year, creating

potential for misclassification. After constructing career trajectories for these inventors using the matching procedures outlined above, we manually review cases with highly frequent mobility events. To reduce false positives and provide a relatively conservative estimate, we leverage the empirical fact that Chinese inventors rarely switch jobs frequently and impose a threshold, excluding cross-regional inventor trajectories with mobility rates exceeding 50% during the sample period. These exclusions mainly affect highly common-name cases and do not materially alter the overall distribution. Overall, despite these edge cases, the distribution of mobility events in our final sample aligns well with real-world patterns in China. Between 2014 and 2022, the vast majority of mobile inventors – approximately 80% – switched employers only once. Most of the remaining inventors moved just twice, and those with more than three job changes account for less than 1% of all mobile inventors. This distribution confirms the validity of our mobility identification and reflects the low-frequency turnover behavior observed in Chinese high-tech labor markets.

Finally, Appendix Figure A2 summarizes the inventor mobility algorithm. We standardize inventor names, link each inventor–year to a firm TIN, location, and IPC field, and infer mobility from changes in corporate assignees across years. Separate rules for within-county and cross-county moves, together with conservative checks for common names and multiple affiliations, help ensure that the final panel captures genuine cross-firm inventor mobility.

Figure A2: Inventor Mobility Identification Algorithm



Notes. The algorithm infers inventor mobility from changes in corporate assignees across patent applications. Within-county moves are identified using standardized inventor names, city–county locations, and firm TINs. Cross-county moves additionally use broad IPC sectors and employer histories to reduce false matches among inventors with common names. Ambiguous cases are resolved conservatively by assigning inventors to their longest-tenure firm, dropping tied affiliations, and excluding cross-regional trajectories with unusually high mobility rates.

B Illustrative Evidence from OPPO

The OPPO case illustrates how the RTA measure captures the direction of knowledge diffusion. OPPO is a major global smartphone producer, with 114.3 million smartphone shipments and an 8.3 percent global market share in 2019. In this market, optical technologies are central to product differentiation, especially through camera modules, zoom systems, image stabilization, sensors, and display-related imaging functions. OPPO's own technology releases also emphasize this margin: its 2021 Future Imaging Technology Launch highlighted RGBW sensors, continuous optical zoom, five-axis optical image stabilization, and under-screen camera technologies.

This context makes the change in OPPO's optics portfolio especially informative. IPC class G02 covers optical technologies, a field closely related to smartphone imaging and display performance. Before the escalation of U.S. sanctions, OPPO had little revealed specialization in this field. Its G02 RTA was only 0.29 in 2014, with an average value of 0.47 across the pre-sanction period of 2014-2017. After 2019, however, OPPO received a sizable inflow of optics-focused, external inventors from sanctioned firms. For example, it recruited an average of about eight G02 specialists each year from Huawei between 2019 and 2022, and brought in roughly one optical expert annually from DJI over the 2021-2022 window. All these inventors boasted proven comparative advantages (RTA ≥ 1) in optical technologies before their job transitions. Consistent with the arrival of this expertise, OPPO's G02 RTA notably rose to an average value of 2.91 between 2019 and 2022, moving the firm from weak revealed presence to clear specialization in optics.

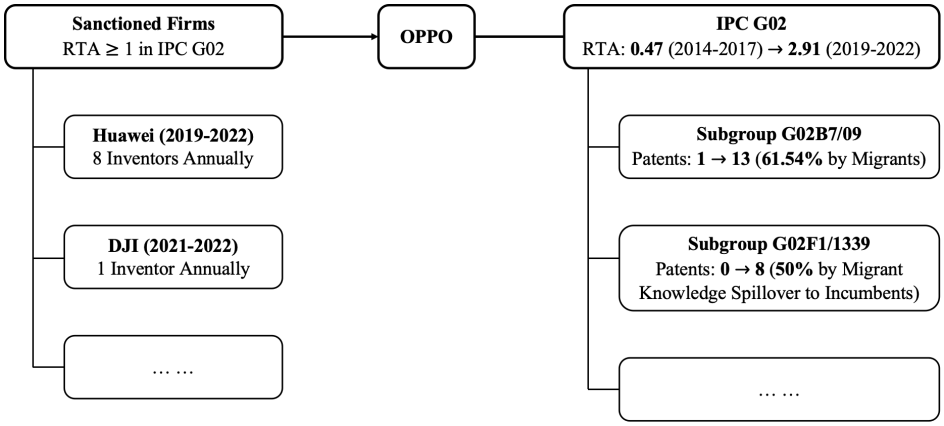
We decompose and concretely illustrate the RTA pattern via the IPC subgroup G02B7/09, covering optical lens and anti-shake modules critical to OPPO's smartphone camera business. During the pre-sanction period, OPPO only filed one sporadic patent here in year 2016 with no continuous research. After 2019's inflow of inventors from sanctioned firms, OPPO filed patents in this detailed subgroup every year from 2019 to 2022, totaling 13 applications. Eight of these patents (61.54%) were contributed by the mobile inventors. This granular evidence echoes the jump in OPPO's overall G02 RTA index, demonstrating that newly recruited optical specialists lifted the firm's technological advantage by independently producing original optical inventions.

Similarly, we turn to a second narrow optical subgroup, G02F1/1339, to document an additional spillover channel through which migrant inventors lift the firm's technological specialization: migrant knowledge diffusion to OPPO's incumbent researchers. The IPC subgroup, G02F1/1339, covers liquid crystal display light modulation components, a core technology underpinning OPPO's high-end smartphone screen products. Specifically, OPPO

submitted zero patents within this field across the entire pre-sanction window of 2014-2017, signaling a complete lack of internal research capacity here before talent inflows. After 2019, however, the firm sustained annual filings in IPC G02F1/1339, accumulating eight total patents by the end of 2022. Among these eight applications, two were jointly produced teams mixing migrant optical inventors and long-standing OPPO incumbents. A following two patents were filed independently by those same incumbent researchers, who evidently absorbed relevant technical know-how from their mobile team partners. In other words, four out of eight total patents (50%) in this disaggregated area stemmed from cross-employee knowledge spillovers initiated by externally recruited inventors. Taken together, the G02F1/1339 case complements the evidence from G02B7/09, revealing that migrant inventors raise a firm's RTA not merely via their own individual output, but also by upgrading the innovative capacity of their in-house colleagues.

Note that the OPPO case summarized in Appendix Figure A3 is illustrative rather than causal evidence, whose value is to make the RTA mechanism concrete. The inflow of inventors from sanctioned firms was followed not only by more patenting at OPPO, but by a sharp movement into the specific optical field in which the origin firms had technological strength. This field-level reorientation is precisely the pattern the RTA analysis is designed to capture, and it is consistent with inventor mobility transmitting tacit technological capabilities across firms.

Figure A3: The RTA Mechanism: OPPO Case



Notes. This figure illustrates a stylized case of the RTA mechanism: OPPO received knowledge from migrating inventors, gaining comparative advantage in optical technologies (G02) and its subfields G02B7/09 and G02F1/1339.

C Aggregate Effect

C.1 Example of Spillover Effect at Receiving Firms

We take Shenzhen Urban Transport Planning Center Co., Ltd. (SUTPC), a listed company in Shenzhen, as a stylized example to illustrate the calculation of firm-specific spillover gains. SUTPC specializes in intelligent traffic control, vehicle-road communication, and smart transportation system leveraging big data and digital twin technologies. Its core inventions align closely with IPC H04 (wireless communication and network technologies) and G06 (computing, data processing, and artificial intelligence). The SUTPC's non-moving stayers annually file 47.831 and 5.539 patents on average within G06 and H04 area, respectively. In the meantime, SUTPC receives an average number of 0.260 (0.285) hundred inventors who departure from sanctioned firms and hold technological comparative advantage within field G06 (H04). We thus combine there facts and measure the spillover gain by SUTPC as follows:

$$\begin{aligned}
\Phi_{SUTPC} &= \sum_{p \in \mathcal{P}_n} \bar{\beta}_{spillover,SUTPC,p} \times \text{Influx}_{SUTPC,p} \\
&= \underbrace{6.1\% \times 47.831 \times 0.260}_{p=\text{G06}} + \underbrace{6.1\% \times 5.539 \times 0.285}_{p=\text{H04}} \\
&= 0.759 + 0.096 \\
&= 0.855
\end{aligned}$$

Note that the computing procedure of other non-sanctioned firms' spillover gains is highly similar as that of SUTPC. For instance, OPPO holds 85.698 patents on average within IPC G02 (refer to the stylized case analysis in Section B for more details). The average inflow inventor count that OPPO receives annually from sanctioned firms between 2021 and 2022 is 0.145 hundred persons. Therefore, the spillover gain of OPPO within G02 is calculated as follows: $\bar{\beta}_{spillover,OPPO,G02} \times \text{Influx}_{OPPO,G02} = 6.1\% \times 85.698 \times 0.145 = 0.758$. The same algorithm applies to all other 57 IPCs that OPPO engages in, of which 18 enjoy positive spillover effect. For simplicity, we do not unfold the detailed but boring calculation process for OPPO case and illustrate another firm SUTPC in detail who obtains spillover only within two IPCs.

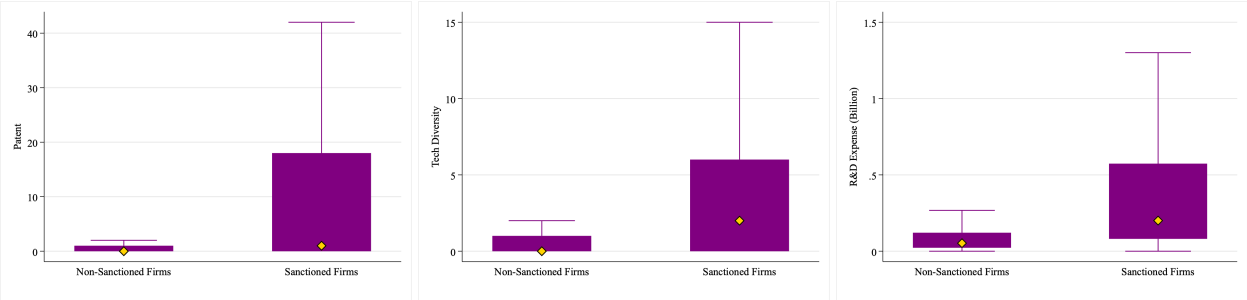
C.2 Example of Synergy Loss at Sanctioned Firms

Consider Guangzhou Haige Communications Group as an example to elaborate the calculation procedure for organizational/synergy loss. Haige is a publicly listed Chinese defense electronics firm specializing in wireless communications, BeiDou navigation and aerospace technologies. It was added to the U.S. Entity List in 2020, and the resulting access restrictions and market disruptions severely undermined its innovation capacity. Under the shock of such sanction, the firm experienced a decline in annual patent applications of 16.5 and a net outflow of 20 inventors per year between 2021 and 2022. Combined with 0.1232 average patent applications per inventor affiliated with Haige in year 2017, we can infer the synergy loss of Haige as follows:

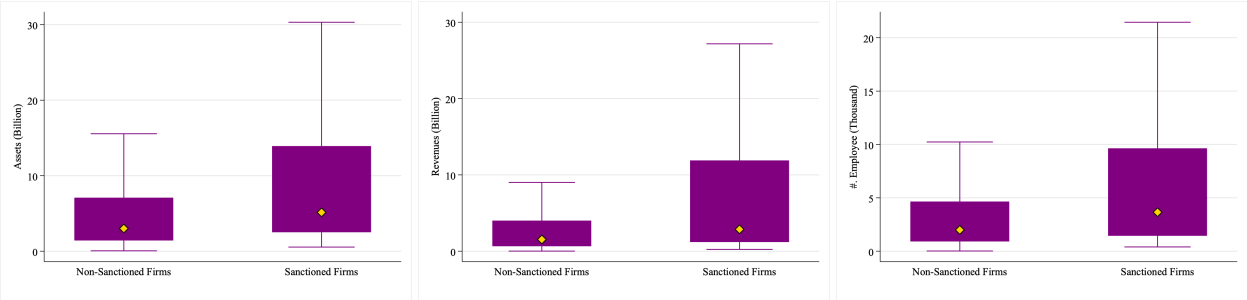
$$\begin{aligned}\mathcal{L}_{Haige} &= \Delta P_{Haige} - \sum_{m \in M_{Haige}} y_{m,pre} \\ &= 16.5 - 20 \times 0.1232 \\ &= 14.036\end{aligned}$$

D Appendix Figures

Figure A4: Box Plot: Non- vs. Sanctioned Firms



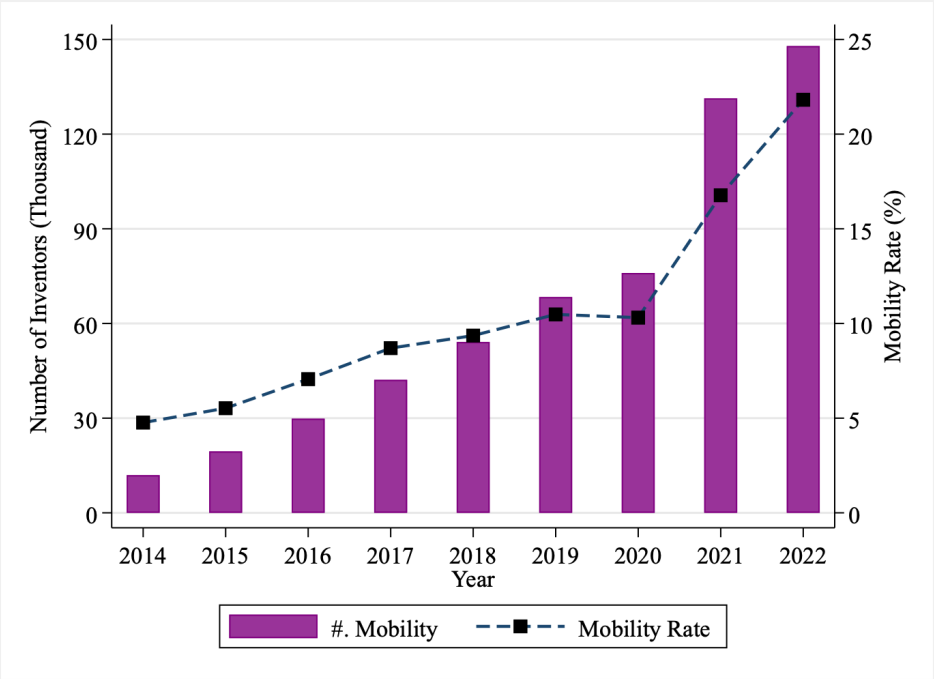
(a) Patent (b) Tech Diversity (c) R&D Expense



(d) Asset (e) Revenue (f) Employee

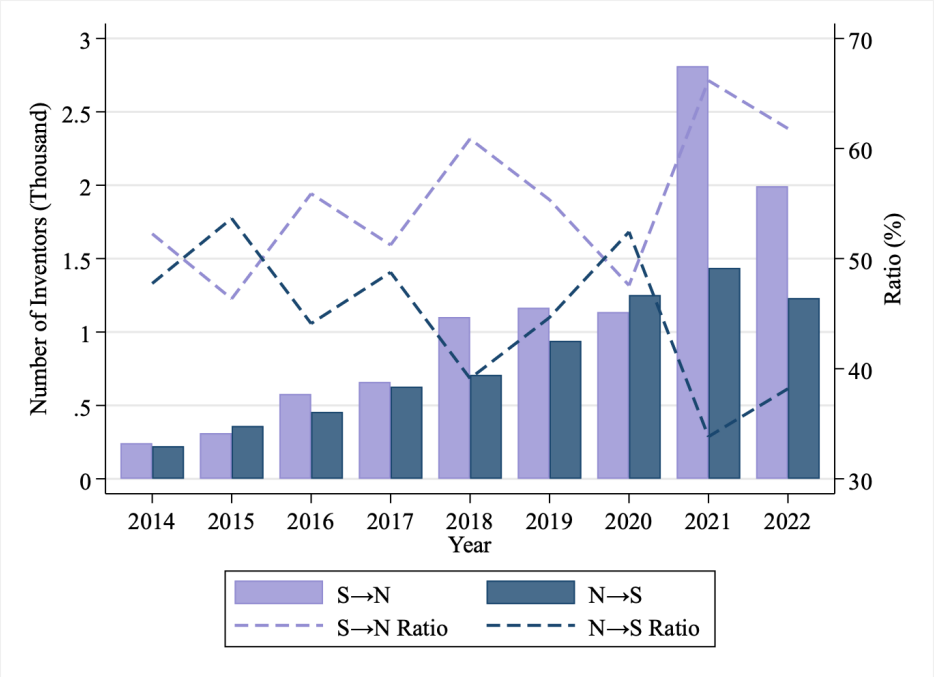
Notes. To draw a whole picture of innovation/financial characteristics of both non- and sanctioned sample firms, we select several critical firm-level variables and draw box plots. The purple box spans the interquartile range (IQR), bounded by the 25th percentile (lower hinge) and 75th percentile (upper hinge), with a diamond inside marking the median. The whiskers extend from the box to the upper and lower adjacent values, defined as the maximum and minimum observations within $Q3 + 1.5 \times IQR$ and $Q1 - 1.5 \times IQR$, respectively, which serve as thresholds to separate typical data from potential outliers. Panels (a) and (b) show innovation performance comparison between non- and sanctioned firms before U.S. sanctions (2014-2017) using the patent dataset. Tech Diversity denotes the number of IPCs where the focal firm conducts. Panels (c)-(f) compare other financial variables using the CSMAR database (2014-2017). We use Asset, Revenue, and Employee as proxies for firm size. Overall, this graph mainly tells that sanctioned firms are (1) more innovation-intensive with higher patents, patent breadth, and R&D expenditure, and (2) larger in terms of assets, revenue, and employees.

Figure A5: Inventor Mobility across Firms, 2014–2022



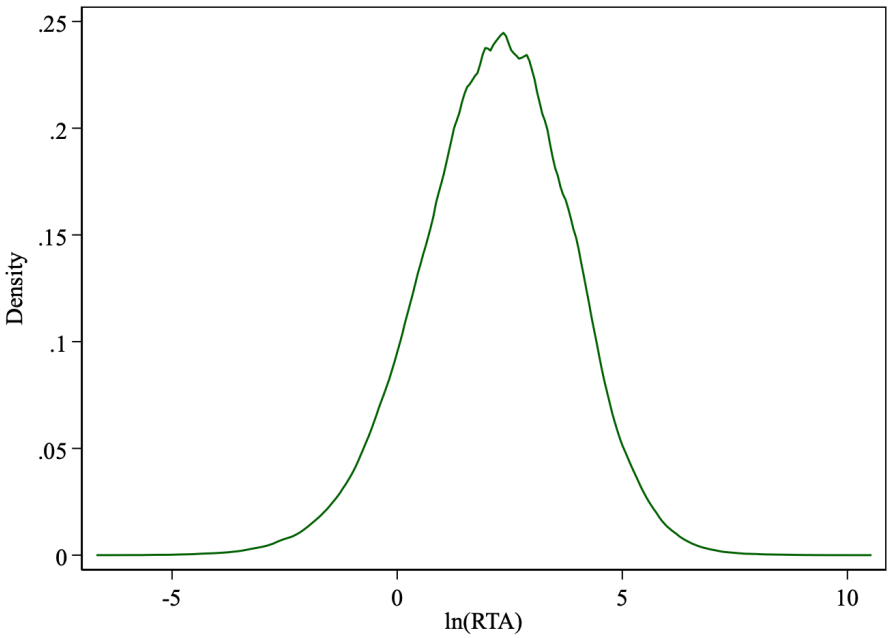
Notes. The figure plots aggregate inventor mobility from 2014 to 2022. The bars report the number of inventors who move across firms in each year, measured in thousands and shown on the left axis. The dashed line reports the annual mobility rate, defined as the share of active inventors who change firms in that year, shown on the right axis. Inventor moves are identified from changes in the corporate assignee associated with an inventor's patent applications across years.

Figure A6: Inventor Flows Out of and Into Sanctioned Firms, 2014–2022



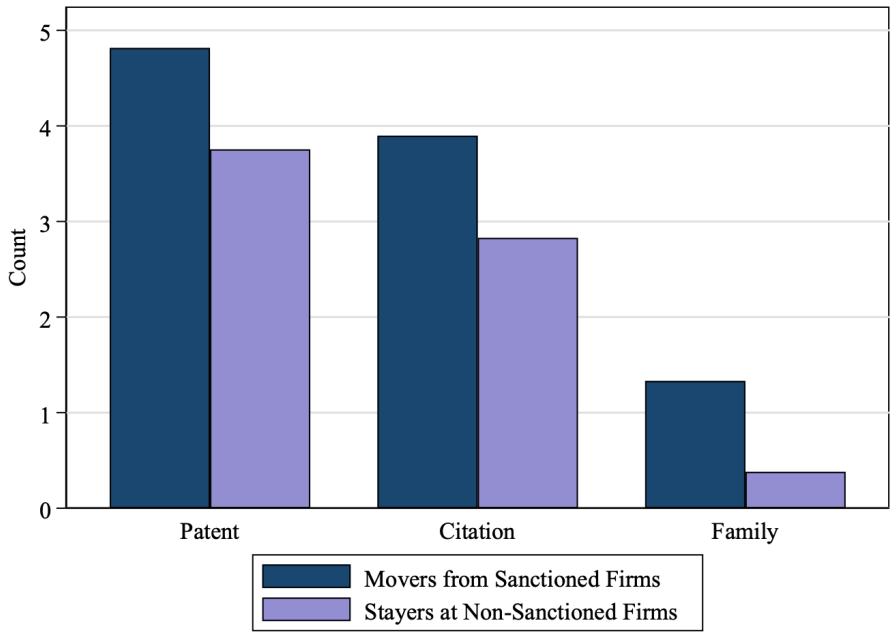
Notes. The figure reports annual inventor flows between sanctioned and non-sanctioned firms. Sanctioned firms are firms added to the BIS Entity List from 2018 onward; non-sanctioned firms are those never listed through 2022. $S \to N$ denotes inventors moving from sanctioned to non-sanctioned firms, while $N \to S$ denotes inventors moving in the opposite direction. The $S \to N$ ratio is the number of $S \to N$ movers divided by the number of $N \to S$ movers; the $N \to S$ ratio is the reciprocal. Inventor moves are identified from changes in corporate assignees across patent applications.

Figure A7: Distribution of RTA index



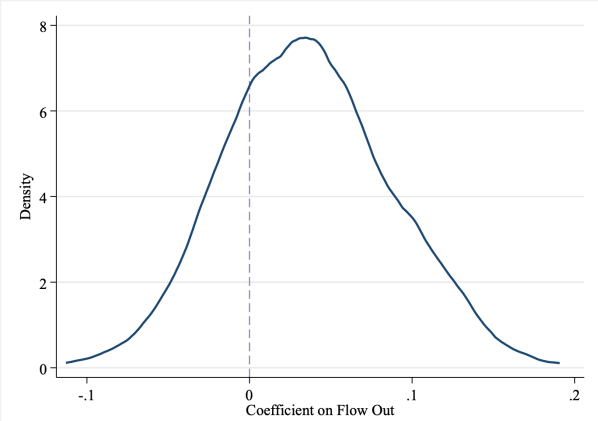
Notes. This figure plots the kernel density of the log Revealed Technological Advantage (RTA) index across firm-IPC-year observations in our sample. The horizontal axis shows log values of RTA index, while the vertical axis represents kernel density. The RTA index captures a firm's relative technological specialization in a given IPC class, constructed as in 4. The distribution shows substantial variation, with a wide range of values indicating pronounced heterogeneity in firms' revealed technological comparative advantages across different fields. This dispersion provides rich identifying variation for our subsequent empirical analyses.

Figure A8: Movers vs. Stayers during Pre-Move Period (2015-2017)

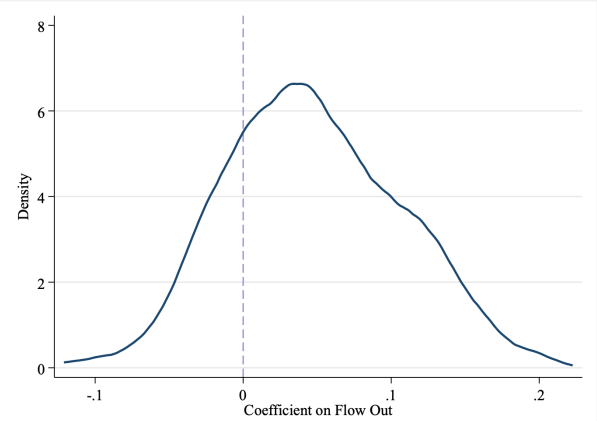


Notes. This bar chart compares the average innovative productivity of two inventor groups during the pre-sanction period (2015-2017). The dark blue bars represent inventors who moved out of sanctioned firms, while the light purple bars represent incumbent stayers at non-sanctioned recipient firms. The horizontal axis lists three innovation metrics: patent applications, forward citations, and international patent family counts; the vertical axis reports the average count for each measure. Movers from sanctioned firms exhibit consistently higher productivity across all metrics, generating more patents, receiving more citations, and holding broader international patent families than stayers. This pattern confirms the pre-existing technological superiority of mobile inventors, supporting the premise that their arrival may generate significant knowledge spillovers within recipient firms.

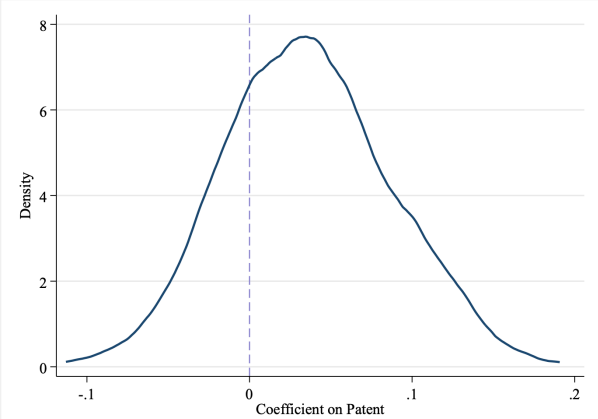
Figure A9: OLS vs. Heckman Estimates by Random Sampling



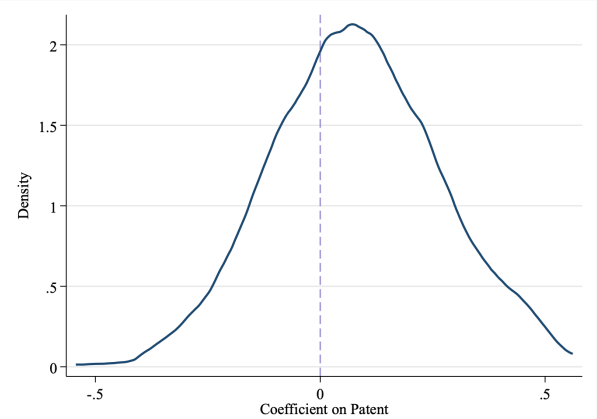
(a) Flow Out by OLS



(b) Flow Out by Heckman



(c) Patent by OLS



(d) Patent by Heckman

Notes. This figure presents the distributions of key coefficients from our OLS (see Table 2 for more details) and Heckman two-stage selection (see Appendix Table A7 for more details) procedures, based on 500 repeated estimations. We randomly sample 100 inventors from both the control and treatment groups and re-run each specification. The horizontal axis shows the estimated coefficient values, while the vertical axis displays the kernel density of these estimates. Specifically, Panels (a) and (c) plot the coefficient distributions for the probability of an inventor changing jobs in a given year (FlowOut) and the log number of patent applications by an inventor in a given year (Patent) from OLS regressions, while panels (b) and (d) show the corresponding distributions from the Heckman two-stage model. The distributions of coefficients are very similar across OLS and Heckman specifications. This close alignment confirms that our main findings are robust to sample selection concerns and alleviates the worry that patent-related results may be severely upward-biased due to selection into patenting.

E Appendix Tables

Table A1: Industry Distribution of Matched Sanctioned Firms by Entity List Year

Industry Name	2018	2019	2020	2021	2022	Total
Biotechnology and Life Sciences	0	0	2	4	0	6
Business Services	0	3	5	16	2	26
Chemicals, Petroleum, Rubber & Plastic	0	0	2	0	0	2
Communications	0	3	3	6	0	12
Computer Hardware	0	1	1	5	0	7
Computer Software	2	4	3	14	3	26
Industrial, Electric & Electronic Machinery	4	5	6	16	6	37
Media & Broadcasting	0	3	1	3	8	15
Metals & Metal Products	0	0	0	1	0	1
Mining & Extraction	0	0	0	1	0	1
Miscellaneous Manufacturing	0	0	0	6	0	6
Retail	0	1	1	0	3	5
Textiles & Clothing Manufacturing	0	0	1	0	0	1
Transport Manufacturing	0	0	2	2	0	4
Wholesale	0	2	2	15	3	22
Total	6	22	29	89	25	171

Notes. The table reports the industry distribution of matched Chinese firms added to the BIS Entity List between 2018 and 2022. Industry classifications come from Bureau van Dijk's Orbis database and correspond to BvD sector categories, which are based on SIC, NACE, and NAICS classifications. We identify 304 mainland Chinese sanctioned firms after excluding entities registered in Hong Kong or Taiwan, financial firms, and firms under ST or PT classifications. Of these, 192 are matched to patent assignees, and 171 have non-missing Orbis industry information and are included in the table.

Table A2: Top 30 Industry Distribution of Matched CSMAR Frims

Industry Code	Industry Description	Firm Count
C39	Computer, Comm. & Electronic Equipment	536
C35	Special Purpose Machinery	333
I65	Software & Information Technology Services	295
C26	Chemical Raw Materials & Products	287
C38	Electrical Machinery & Equipment	282
C27	Pharmaceutical Manufacturing	258
C34	General Purpose Machinery	187
C36	Automobile Manufacturing	138
C29	Rubber & Plastic Products	110
C30	Non-metallic Mineral Products	90
C33	Metal Products	88
C40	Instrumentation Manufacturing	82
M74	Professional & Technical Services	77
C32	Non-ferrous Metal Smelting	70
N77	Environmental & Public Facilities Management	67
C37	Railway, Ship & Aerospace Equipment	64
C14	Food Manufacturing	62
E48	Construction Installation	58
C13	Food Processing	50
D44	Electricity & Heat Production & Supply	46
C17	Textile Industry	41
C22	Paper & Paper Products	38
C15	Beverage Manufacturing	36
I64	Internet & Related Services	31
C18	Textile, Apparel & Accessories	31
C21	Wood Processing & Wood Products	30
C31	Ferrous Metal Smelting	27
C28	Chemical Fiber Manufacturing	27
F52	Retail Trade	26
E50	Decoration & Other Construction	24

Notes. Industry classification and descriptions follow the 2012 industry standard issued by the China Securities Regulatory Commission. All industry codes represent industry classes. The sample firms retrieved from the CSMAR database cover 73 industries in total, and this table presents in detail the top 30 industries ranked by firm quantity.

Table A3: Effect of U.S. Sanctions on Firm-Level Outcomes

	(1) N. patents	(2) N. inflow	(3) N. outflow	(4) 1{Net outflow}	(5) N. net outflow
A. Innovation Outcomes					
Sanction	-1.498**	12.030	23.079**	0.193***	11.049
	(0.76)	(7.70)	(10.78)	(0.04)	(7.20)
Firm Controls	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	1,441,211	1,441,211	1,441,211	1,441,211	1,441,211
R^2	0.770	0.519	0.554	0.246	0.249
	(1) ROA	(2) ROE	(3) Growth of Earnings	(4) Cash/Asset	(5) Leverage/Asset
B. Financial/Accounting Outcomes					
Sanction	-0.022***	-0.040**	-0.007**	-0.013**	0.038**
	(0.01)	(0.02)	(0.00)	(0.01)	(0.02)
Firm Controls	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes
Observations	26,265	24,643	23,303	26,257	26,265
R^2	0.622	0.515	0.199	0.618	0.812

Notes. This table examines the impact of the Entity List trade sanctions on firm innovation and financial performance. The unit of observation is a firm-year pair. Panel A uses all firms with at least one patent application during the sample period, while Panel B focuses on listed firms with public financial performance information. The dependent variable is the probability of an inventor changing jobs in a given year in Columns 1 to 3; the log number of patent applications by an inventor in a given year in Columns 4 to 6. "Sanction" is a dummy variable taking one for all years after inventor was first employed by a firm included on the Entity List post-2018 and zero otherwise. Individual controls include the log years since first patent application, the average size of the co-authors in a given year, and whether inventor has actively filed patents in the past three years. Firm controls include the number of IPC technology classes on patents applied by a firm in a given year and whether firm j has patented in the past three years. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.001, respectively.

Table A4: Top 10 IPC Classes by Revealed Technological Advantage

IPC	IPC Description	Rank
Panel A: Sanctioned Firms		
G10	Musical acoustics and audio technologies	1
H03	Electronic circuits and signal-processing technologies	2
G03	Photography, imaging, and holography technologies	3
G11	Information storage and memory technologies	4
C03	Glass, mineral wool, and related material technologies	5
C23	Metallic coating and surface-treatment technologies	6
G02	Optics and optical instrument technologies	7
G08	Signalling, traffic control, and alarm systems	8
B29	Plastics working and shaping technologies	9
H05	Electric techniques not otherwise provided for	10
Panel B: Non-Sanctioned Firms		
B81	Micro-structural and micro-mechanical technologies	1
F41	Weapons and related launching technologies	2
G16	Specialized ICT applications, including health informatics	3
B33	Additive manufacturing and 3D printing technologies	4
C30	Crystal growth and crystalline material technologies	5
B82	Nanotechnology and nano-structure technologies	6
A44	Jewellery, haberdashery, and personal adornments	7
G04	Horology and time-measuring technologies	8
F17	Gas or liquid storage and distribution systems	9
B63	Ships, marine vessels, and waterborne equipment	10

Notes: Panel A reports the top ten IPC classes for sanctioned firms, while Panel B for non-sanctioned firms. Rankings are based on mean RTA index during the pre-sanction period, namely 2015-2017. Rare IPC classes with fewer than 10 firm observations over 2015-2017 are excluded before ranking. IPC descriptions are abbreviated for readability.

Table A5: Robustness: PSM

	1:1 Matching		1:2 Matching		1:3 Matching	
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	FlowOut	log n. patents	FlowOut	log n. patents	FlowOut	log n. patents
A. Matching with Firm-Level Characteristics						
Sanction	0.027**	-0.049	0.026**	-0.069	0.020*	-0.080
	(0.01)	(0.04)	(0.01)	(0.04)	(0.01)	(0.05)
Inventor Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
Inventor FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	107,166	107,166	112,117	112,117	117,256	117,256
R^2	0.373	0.621	0.373	0.623	0.375	0.622
	1:1 Matching		1:2 Matching		1:3 Matching	
	(1)	(2)	(3)	(4)	(5)	(6)
Dependent variable:	FlowOut	log n. patents	FlowOut	log n. patents	FlowOut	log n. patents
B. Matching with Inventor-Level Characteristics						
Sanction	0.036**	-0.022	0.032***	0.135	0.031***	0.134
	(0.01)	(0.04)	(0.01)	(0.16)	(0.01)	(0.16)
Inventor Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
Inventor FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	105,819	105,819	115,278	115,278	116,737	116,737
R^2	0.378	0.621	0.382	0.631	0.383	0.631

Notes. This table performs robustness checks of Panel A in Table 2 using propensity score matching based on firm and inventor pre-sanction characteristics during 2014-2017. The unit of observation is an inventor-year pair. The dependent variable is the probability of an inventor changing jobs in a given year in Columns 1, 3, and 5; the log number of patent applications by an inventor in a given year in Columns 2, 4, and 6. "Sanction" is a dummy variable taking one for all years after inventor was first employed by a firm included on the Entity List post-2018 and zero otherwise. Individual controls include the log years since first patent application, the average size of the co-authors in a given year, and whether inventor has actively filed patents in the past three years. Firm controls include the number of IPC technology classes on patents applied by a firm in a given year and whether firm j has patented in the past three years. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.

Table A6: Robustness: Excluding Competing Firms

	5 Main IPCs		3 Main IPCs	
	(1)	(2)	(3)	(4)
Dependent variable:	FlowOut	log n. patents	FlowOut	log n. patents
Sanction	0.080***	-0.065	0.079***	-0.072
	(0.03)	(0.07)	(0.03)	(0.07)
Inventor Controls	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes
Inventor FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	4,901,019	4,901,019	4,812,813	4,812,813
R^2	0.609	0.604	0.611	0.604

Notes. This table performs robustness checks of Panel A in Table 2 by excluding firms in the control groups that share the same top three or five core IPC domains as treated firms. The unit of observation is an inventor-year pair. The dependent variable is the probability of an inventor changing jobs in a given year in Columns 1 and 3; the log number of patent applications by an inventor in a given year in Columns 2 and 4. "Sanction" is a dummy variable taking one for all years after inventor was first employed by a firm included on the Entity List post-2018 and zero otherwise. Individual controls include the log years since first patent application, the average size of the co-authors in a given year, and whether inventor has actively filed patents in the past three years. Firm controls include the number of IPC technology classes on patents applied by a firm in a given year and whether firm j has patented in the past three years. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.

Table A7: Robustness: Heckman Two-Step Estimation

	Linear Probability Model		Probit Model	
	(1)	(2)	(3)	(4)
Dependent variable:	FlowOut	log n. patents	FlowOut	log n. patents
Sanction	0.081***	-0.073	0.081***	-0.067
	(0.03)	(0.07)	(0.03)	(0.06)
Inverse Mills Ratio	0.321**	-4.429***	0.395***	-0.360***
	(0.15)	(0.32)	(0.01)	(0.02)
First Stage:				
Whether file patents in $t-1$	0.012***		0.327***	
	(0.00)		(0.00)	
Inventor Controls	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes
Inventor FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	4,939,429	4,939,429	4,939,429	4,939,429
R^2	0.609	0.604	0.615	0.604

Notes. This table performs robustness checks of Panel A in Table 2 using Heckman two-stage selection model. In the first stage, the outcome variable is an indicator which takes the value of one if inventor has patent records in the focal year, and zero otherwise. We use a dummy for whether an inventor files any patent application in the prior year as an instrument. In line with our baseline specification in Table 2, we control individual characteristics, including the log years since first patent application. We exclude time-varying firm controls and firm fixed effects because assigning employers to inventor-year observations without patent records would require subjective imputation, which introduces potentially severe measurement errors. Standard errors are clustered at the firm level. Both linear probability model and probit model are employed in the first stage for robustness. In the linear probability specification, we absorb year and inventor fixed effects. In the Probit model, we do not include inventor fixed effects because of computational burden. We address this issue in Appendix Figure A9 by conducting a bootstrap-style Heckman procedure. N = $9,859,456$ in the first stage since we fill in the observations without patent records. In the second stage, the unit of observation is an inventor-year pair. The dependent variable is the probability of an inventor changing jobs in a given year in Columns 1 and 3; the log number of patent applications by an inventor in a given year in Columns 2 and 4. "Sanction" is a dummy variable taking one for all years after inventor was first employed by a firm included on the Entity List post-2018 and zero otherwise. Individual controls include the log years since first patent application, the average size of the co-authors in a given year, and whether inventor has actively filed patents in the past three years. Firm controls include the number of IPC technology classes on patents applied by a firm in a given year and whether firm j has patented in the past three years. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.

Table A8: Robustness: Sensitivity Test of Job Change Time

Dependent variable: FlowOut (probability of switching jobs)	(1)	(2)	(3)
Sanction	0.078***	0.071***	0.057***
	(0.03)	(0.02)	(0.01)
Inventor Controls	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes
Inventor FE	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
Observations	4,814,980	4,472,113	4,088,218
R^2	0.600	0.592	0.591

Notes. This table performs robustness checks of Panel A in Table 2 for the sensitivity of job change timing. Specifically, Columns (1) to (3) restrict the sample to inventors whose maximum job change interval is less than 5 years, 3 years, and exactly 1 year, respectively. The unit of observation is an inventor-year pair. The dependent variable is the probability of an inventor changing jobs in a given year in all of the three columns. "Sanction" is a dummy variable taking one for all years after inventor was first employed by a firm included on the Entity List post-2018 and zero otherwise. Individual controls include the log years since first patent application, the average size of the co-authors in a given year, and whether inventor has actively filed patents in the past three years. Firm controls include the number of IPC technology classes on patents applied by a firm in a given year and whether firm j has patented in the past three years. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.

Table A9: Robustness: High-Dimensional FE

Dependent variable:	FlowOut (probability of switching jobs)		log n. patents	
	(1)	(2)	(3)	(4)
Sanction	0.064***	0.082***	-0.097	-0.093*
	(0.02)	(0.02)	(0.07)	(0.05)
Inventor Controls	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes
Inventor FE	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
City-year FE	Yes	No	Yes	No
County-year FE	No	Yes	No	Yes
Observations	3,935,500	3,901,236	3,935,500	3,901,236
R^2	0.672	0.681	0.601	0.612

Notes. This table uses the same specification as in Table Table 2 adding city-year fixed effects in Columns 1 and 3, and more granular county-year fixed effects in columns 2 and 4. The unit of observation is an inventor-year pair. The dependent variable is the probability of an inventor changing jobs in a given year in Columns 1 and 2; the log number of patent applications by an inventor in a given year in Columns 3 and 4. "Sanction" is a dummy variable taking one for all years after inventor was first employed by a firm included on the Entity List post-2018 and zero otherwise. Individual controls include the log years since first patent application, the average size of the co-authors in a given year, and whether inventor has actively filed patents in the past three years. Firm controls include the number of IPC technology classes on patents applied by a firm in a given year and whether firm j has patented in the past three years. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.

Table A10: Robustness: Alternative Econometric Specification

	All sample		Inventors patent at least 5 years		Inventors patent every year	
Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)
Number of patents applied	Poisson	OLS	Poisson	OLS	Poisson	OLS
Sanction	-0.134	-0.368	-0.133	-0.366	-0.122	-0.189
	(0.10)	(0.50)	(0.11)	(0.58)	(0.11)	(0.93)
Inventor Controls	Yes	Yes	Yes	Yes	Yes	Yes
Firm Controls	Yes	Yes	Yes	Yes	Yes	Yes
Inventor FE	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes	Yes	Yes
Observations	4,939,429	4,939,429	1,351,560	1,351,560	95,324	95,324
R^2	0.438	0.613	0.443	0.534	0.515	0.541

Notes. This table reports results on patent application using alternative models. The dependent variable is the number of patents applied in a year. Columns (1) and (2) report the estimate using Poisson model and OLS, respectively. Columns (3) to (4) restrict the sample to active inventors who patent for at least 5 years during the sample period (2014-2022). Columns (5) to (6) further restrict to inventors who patent every sample year. The unit of observation is an inventor-year pair. "Sanction" is a dummy variable taking one for all years after inventor was first employed by a firm included on the Entity List post-2018 and zero otherwise. Individual controls include the log years since first patent application, the average size of the co-authors in a given year, and whether inventor has actively filed patents in the past three years. Firm controls include the number of IPC technology classes on patents applied by a firm in a given year and whether firm j has patented in the past three years. Standard errors are clustered at the firm level and reported in parentheses. *, **, and *** represent significance levels of 0.10, 0.05, and 0.01, respectively.