International Innovation Spillovers through Global Technology Ownership

[PRELIMINARY AND INCOMPLETE PLEASE DO NOT CIRCULATE] *

Christian Fons-Rosen† Şebnem Kalemli-Özcan‡ Bent E. Sørensen§
Carolina Villegas-Sanchez¶ Vadym Volosovych‖

Abstract: This paper examines how domestic firms benefit from cross-border knowledge diffusion via multinational ownership networks. Using firm-level data for ten European manufacturing economies (1999–2012), we construct shift-share measures of *firm-to-firm* foreign technology exposure based on the pre-sample patent-class overlap between domestic firms and the global ultimate owners (GUOs) of foreign affiliates operating in the same country-sector. A one-standard-deviation increase in foreign technology exposure raises citation-weighted domestic patenting by approximately 16 percent after four years and 39 percent after five years, with clean pre-trends. The response is concentrated among domestic firms exposed to technology-intensive foreign affiliates — those that themselves patent intensively in the host sector — and among firms with greater initial absorptive capacity. The innovation response propagates to measured total factor productivity, consistent with positive horizontal spillovers. Our findings imply that the domestic productivity dividend depends on the quality of inward FDI, not merely its volume, and on the absorptive capacity of local firms.

Keywords: Multinationals, Technology spillovers, Innovation, Productivity, Patent portfolios, Shift-share.
JEL Codes: F23, O30, O33, D24.

---

*This version: August 3, 2026. The first version circulated under the title: "Foreign Investment and Domestic Productivity: Identifying Knowledge Spillovers and Competition Effects." Affiliations: UC Merced; Brown University, CEPR, and NBER; University of Houston and CEPR; ESADE – Universitat Ramon Llull, ECB and CEPR; Erasmus University Rotterdam and Tinbergen Institute. Carolina Villegas-Sanchez acknowledges financial support from Banco Sabadell, MINECO grant FIG_MICIU_GC2021 and ICREA Academia.

†University of California at Merced, CEPR, and Max Planck Institute for Competition and Innovation. cfons-rosen@ucmerced.edu
‡Brown University.
§University of Houston.
¶ESADE, URL, ECB and CEPR. carolina.villegas@esade.edu
‖Erasmus University Rotterdam and Tinbergen Institute.

1 Introduction

Innovation is highly concentrated. The top 2,500 firms by R&D expenditure account for roughly 75 percent of global private R&D, and these frontier innovators are typically large multinational enterprises with extensive patent portfolios headquartered in a small number of countries.[1] Yet the knowledge embodied in their innovations, codified in patent disclosures that are publicly accessible worldwide, does not diffuse uniformly across the firms and countries positioned to benefit from it. Large and persistent differences in firm-level productivity coexist with a global patent system that makes frontier knowledge nominally accessible to all (Parente and Prescott, 1994). The question this paper addresses is whether multinational enterprises transmit frontier knowledge to domestic firms, and, if so, which domestic firms are able to use it.

This distinction between formal access and economic use is central to our argument. Patents disclose claims, technical descriptions, and technological classifications, but building on frontier inventions often requires tacit complements that cannot be read directly from a patent document: engineering judgment, knowledge of failed experiments, routines for integrating new techniques into production, and the absorptive capacity to recognise which elements of the frontier are relevant for one's own activities. These complements are more likely to be transmitted through observation, local engineering practice, and organisational proximity to firms that actively use the relevant technology, rather than solely through input-output relationships or supply chains.

We argue that foreign affiliates of large multinationals can act as local interfaces between the global knowledge frontier and domestic firms. The relevant stock of frontier knowledge resides at the level of the global ultimate owner, which aggregates technological capabilities developed across the multinational network. Domestic firms that are closer to this global knowledge stock in technology space are better positioned to understand, interpret, and build on it. But proximity to the multinational's global portfolio is not sufficient on its own. Spillovers should be strongest when the host-country affiliate also innovates, because affiliate-level innovation makes elements of the multinational's knowledge base locally visible, operationally relevant, and embedded in host-country engineering environments.

This mechanism implies a specific form of selective knowledge diffusion. Domestic firms do not simply imitate the patents of foreign multinationals. Rather, firms with related capabili-

[1] See the EU Industrial R&D Investment Scoreboard and Yeaple (2013) on the concentration of innovation among large multinationals.

ties use locally activated foreign knowledge as an input into their own innovative activity. The relevant response is recombination: domestic firms combine their pre-existing technological capabilities with opportunities revealed by the multinational's global portfolio and made actionable by the innovative activity of its local affiliate. The same foreign patent portfolio may therefore have very different effects across domestic firms, depending on their pre-existing technological proximity to the multinational owner and on whether the affiliate serves as an active local conduit for that knowledge.

To test this mechanism, we build a novel firm-level dataset combining financial records for manufacturing firms in ten European countries (1999–2012) with patent information from PATSTAT and detailed corporate ownership from Moody's Bureau van Dijk ORBIS. We construct a shift-share measure of *foreign technology exposure* for each domestic firm: for each patenting foreign affiliate in the same country-sector cell, we measure the domestic firm's pre-sample patent-class proximity to the affiliate's Global Ultimate Owner (GUO) and weight this proximity by the GUO's citation-weighted cumulative patent stock. This measure has the Bartik structure (Goldsmith-Pinkham et al., 2020): predetermined technological proximity shares interact with time-varying innovation shocks generated by the global R&D activity of multinational owners. Our identification compares domestic firms within the same country-sector-year cell that face the same local foreign-affiliate environment but differ in their pre-sample technological proximity to the GUOs of those affiliates. This design separates local exposure to foreign affiliates from a firm's ability to absorb and recombine the knowledge those affiliates make locally salient.

Our main findings are as follows. A one-standard-deviation increase in foreign technology exposure raises citation-weighted domestic patenting by approximately 16 percent after four years and 39 percent after five years. The response is absent at short horizons and builds gradually, consistent with a multi-year process of recognising technological opportunities and adapting internal R&D toward them. We find no evidence of differential pre-trends in the key foreign-technology exposure variable, and the result is stable across a battery of robustness checks. The patent response is concentrated among two sets of firms. First, it is driven by domestic firms exposed to technology-intensive affiliates — those that themselves patent intensively in the host sector — and is absent for foreign owners without technology-active local affiliates, confirming that the mechanism runs through affiliate innovation intensity rather than foreign ownership per se. Second, within the set of firms exposed to technology-intensive affiliates, the response is strongest among firms with above-median initial patent stocks, consistent with the view that R&D both generates innovation and builds absorptive

capacity (Cohen and Levinthal, 1989, 1990). The productivity results are consistent with this innovation response propagating to measured total factor productivity: TFP dynamics are positive among firms exposed to technology-intensive affiliates and flat for the rest.

Our design addresses three main identification concerns. First, country × four-digit-sector × year fixed effects absorb common shocks to local industries, including cyclical demand, sector-specific policy, and aggregate technology trends. Identification therefore comes from comparing domestic firms within the same local industry and year that differ in their pre-sample technological proximity to the global owners of foreign affiliates. Second, the exposure shares are constructed from pre-1999 patent portfolios and GUO assignments are held fixed at their initial values; since the affiliate set in each country-sector cell changes over time, the fixed GUO assignments ensure that variation in our exposure measure reflects changes in GUO global patent activity rather than changes in ownership structure or affiliate entry. Third, we control for the sectoral-proximity analogue of foreign technology exposure, so the coefficient on FOtech_{it-1} is identified net of product-market competition effects associated with the same GUO patent shocks. We further find no evidence of differential pre-trends in patenting and verify the robustness of the shift-share design using shift-share decomposition diagnostics, AKM inference, and BHJ overidentification tests (Goldsmith-Pinkham et al., 2020; Adão et al., 2019; Borusyak et al., 2022).

These results make three contributions to the literature.

First, cumulative innovation and the "shoulders of giants." A large body of work documents that innovation builds cumulatively on prior knowledge (Caballero and Jaffe, 1993; Furman and Stern, 2011; Acemoglu et al., 2016; Akcigit et al., 2021). Access to the frontier is, however, uneven: it depends on technological proximity, institutional linkages, and absorptive capacity. We contribute by identifying a specific organisational mechanism — the locally-innovating foreign affiliate — through which domestic firms gain practical access to the frontier patent portfolios of global innovators. Our results are consistent with affiliate-level innovation converting a publicly readable patent into one that domestic firms with complementary capabilities can build on.

Second, firm-to-firm technology spillovers and patent-space proximity. Jaffe (1986) uses technology-class overlap to measure technological proximity and R&D spillovers across firms, while Jaffe et al. (1993) uses patent citations to study the geographic localisation of knowledge spillovers. Bloom et al. (2013) operationalises this into a shift-share framework that separates technology spillovers from product-market competition effects, which we adopt and extend. Our contribution is to show that local affiliate innovation activates the cross-firm technology-

spillover channel: the same degree of patent-class proximity to a GUO is associated with no innovation response unless a technology-intensive affiliate is locally present. Proximity is necessary but not sufficient; the affiliate is what makes the channel operative.

Third, knowledge diffusion through multinational and firm networks. A growing body of work studies how knowledge diffuses within multinational firm networks (Keller and Yeaple, 2013; Bilir and Morales, 2020) and how productivity shocks propagate across geographically dispersed establishments connected through common ownership (Giroud et al., 2024, 2026). This literature emphasises *within-firm* knowledge flows. We study diffusion *beyond* firm boundaries: from a foreign affiliate to unaffiliated domestic firms in the same sector. This channel does not operate through vertical integration or corporate hierarchy but through the local knowledge environment created by a technology-active affiliate, and it implies that the benefits of multinational innovation are not fully internalised by the MNC but can spill over to independent domestic producers.

Our findings also speak to the long-running debate on horizontal FDI spillovers. The influential result of Aitken and Harrison (1999) — that within-sector foreign entry does not raise and may lower domestic TFP — emerges naturally from our framework as a special case: for domestic firms that are not technologically proximate to the local affiliate, or that lack the absorptive capacity to exploit proximity, the competitive channel can dominate. Our contribution is to show that conditioning on technological proximity and affiliate innovation intensity recovers a positive horizontal channel. The two results are not in tension: Aitken and Harrison estimate the average effect across all domestic firms; we identify the subset of firms for which the knowledge channel dominates.

The remainder of the paper is organised as follows. Section 2 describes the research setting, data, ownership structure, patent matching, and the construction of the foreign technology exposure measure. Section 3 presents the empirical strategy, identification assumptions, and shift-share validity checks. Section 4 reports the patent local projection results, pre-trend tests, robustness checks, and the LP-DiD discrete-treatment design. Section 5 analyses the two channels — technology-intensive affiliates and absorptive capacity — that determine which firms benefit. Section 6 examines whether the innovation response propagates to total factor productivity. Section 7 concludes.

2 Setting, Data, and Measurement

2.1 Research Setting

Figure 1 illustrates the structure of our research design and the ownership relationships that underpin it. Within a host-country sector, two types of firms can potentially transmit knowledge to independent domestic firms: foreign affiliates (red), owned by a Foreign Ultimate Owner (FUO) based abroad; and domestic affiliates (green), owned by a Domestic Ultimate Owner (DUO) in the same country. Our key object of interest is the foreign ownership channel. The FUO in Country B holds a frontier patent portfolio accumulated through its global R&D activities. Its local affiliate operates in the same four-digit NACE sector as two independent domestic firms, i and k. Firm i has high technological overlap with the FUO's patent portfolio ($tech_{i,\mathrm{FUO}}$ is large), while firm k has low overlap ($tech_{k,\mathrm{FUO}}$ is small). The central prediction is that, as the FUO's global patent stock expands, firm i — technologically proximate and embedded in the same local sector as the affiliate — is better positioned to recombine frontier knowledge into its own innovations. Firm k, despite facing the same affiliate, benefits less because its prior capabilities are further from the FUO's knowledge base. The domestic channel (DUO $\rightarrow$ Domestic Affiliate $\rightarrow$ firms i and k) operates in parallel and is captured by our domestic exposure controls DOtech and DOcomp, ensuring that the foreign technology channel is identified net of domestic knowledge spillovers.

The figure also clarifies what the spillover mechanism is *not*. The knowledge flow from the FUO to domestic firms does not run through direct ownership links, buyer-supplier relationships, or input-output connections. It runs through the local knowledge environment created by an affiliate that is itself actively innovating — making the FUO's frontier patent portfolio operationally relevant within the host-country sector. Domestic firms that are technologically close to the FUO can observe, interpret, and recombine elements of this frontier knowledge into their own R&D. Those that are technologically distant cannot, and may instead experience only the competitive effects of foreign entry. Technology-intensive affiliates are, in this sense, the organisational complement to technological proximity: codified frontier knowledge becomes economically useful only when both interact.

A concrete example illustrates the logic. Nokia Corporation, headquartered in Helsinki, held one of the world's largest wireless-technology patent portfolios during our sample period. Its German affiliate Nokia GmbH maintained an engineering centre in Ulm, actively developing and applying wireless modem technologies locally. A firm such as Rohde & Schwarz,

whose prior patents spanned adjacent radio-frequency, signal-processing, and measurement technologies, is the type of firm our measure would classify as technologically proximate to Nokia's global portfolio. In our framework, such a firm has high foreign technology exposure when Nokia's global knowledge stock is paired with a locally innovative affiliate. The relevant spillover is not imitation of Nokia's patents. It is recombination: Nokia's expanding frontier wireless technologies may create technical opportunities in testing, validation, and signal integration, while a domestic firm with complementary measurement capabilities may translate those opportunities into its own inventions. Consistent with technological proximity between the two firms, Rohde & Schwarz patent families during our sample period cite prior Nokia and Nokia Networks patent families in mobile communications and testing-related technologies.[2] More generally, patent-class proximity identifies firms close enough to understand and build on the frontier, while affiliate innovation intensity identifies cases in which that frontier is locally actionable.

2.2 The Firm Panel: ORBIS

Financial data come from Moody's Bureau van Dijk ORBIS (October 2015 vintage), which provides balance-sheet and income-statement information for European firms. Our sample covers the manufacturing sector in ten European economies — Austria, Belgium, Germany, Spain, Finland, France, Italy, Norway, Portugal, and Sweden — over the period 1999–2012. These countries are chosen for three reasons. First, they experienced a pronounced surge in inward FDI during this period: the stock of FDI into Europe rose from 34 to 42 percent of global FDI stock, the largest increase recorded during any episode of economic integration.[3] Second, ORBIS coverage of ownership links and financial statements is most complete for these economies. Third, EU/EEA market integration ensures institutional comparability across the sample.

The core financial variables are operating turnover, labour costs, intermediate materials, and tangible fixed assets. We deflate turnover, labour costs, and materials using two-digit NACE producer price indices, and capital using investment price deflators. Key cleaning decisions — winsorization at the 1st–99th percentile by country, treatment of firms changing sector codes, and construction of the TFP estimation sample — are described in Appendix A.1.

[2] We identified at least two Rohde & Schwarz patent families filed within our sample period (DE10320670A1, priority date 2003; DE102011005413A1, priority date 2011) whose citation lists include Nokia or Nokia Networks patent families in wireless communications. We report this as evidence consistent with the mechanism; we do not assert a causal link from Nokia's German affiliate specifically.

[3] Based on FDI inward stocks as reported in UN World Investment Report 2024, Annex Table 2.

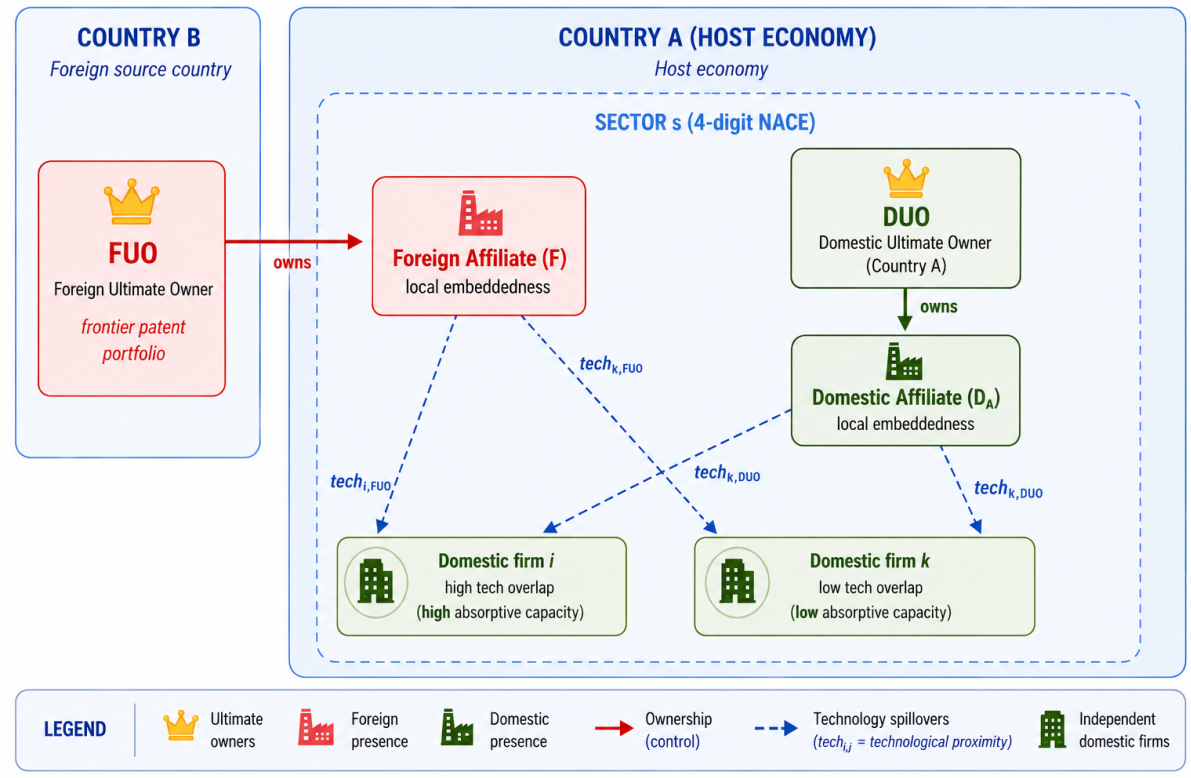


Figure 1: Research Setting: Global Ownership, Local Affiliates, and Domestic Knowledge Spillovers

Notes: The figure illustrates the ownership structure and knowledge-spillover channels in our research design. Within a host-country sector (four-digit NACE), a Foreign Ultimate Owner (FUO, Country B) owns a locally operating foreign affiliate. A Domestic Ultimate Owner (DUO, Country A) owns a domestic affiliate. Both types of affiliate co-locate in the same sector as two independent domestic firms, i and k. Dashed blue arrows represent potential technology spillovers, scaled by the pairwise technological proximity $tech_{i,j}$ between the domestic firm and the ultimate owner of the affiliate. Firm i has high technological overlap with the FUO's patent portfolio (high absorptive capacity); firm k has low overlap (low absorptive capacity). Red solid arrows denote ownership (control) links. Our main outcome is whether domestic firms in cells with a locally innovating foreign affiliate patent more and become more productive as the FUO's global patent stock expands, and whether this response is stronger for firm i than for firm k.

We estimate firm-level total factor productivity using the Wooldridge (2009) GMM procedure, which addresses the simultaneity of input choices and productivity following Levinsohn and Petrin (2003) and Olley and Pakes (1996) and is robust to the identification critique of Ackerberg et al. (2015). The procedure uses intermediate inputs as the proxy variable and lagged inputs as instruments; we estimate separate production functions by two-digit sector and country. Estimated output elasticities are reported in Appendix B.3.

The full sample contains **471,581 firms** and 2.82 million firm-year observations, covering more than 75 percent of manufacturing output across the ten countries. The TFP estimation sample (requiring non-missing and positive values for all four inputs) comprises 463,914 firms. Firm-years range from 128,496 in 1999 to 252,147 at the 2007 peak.

2.3 Ownership Structure and the Affiliate Network

Ownership definitions. ORBIS records both *direct* and *ultimate* ownership links. Following Moody's classification, we define the **Global Ultimate Owner** (GUO) as the entity at the top of the ownership chain for which no other unique shareholder holds more than 25 percent of its equity — a threshold aligned with the EU Takeover Directive (2004/25/EC) and standard in the literature.[4] A GUO is *foreign* if it is incorporated in a country different from the firm it controls, and *domestic* otherwise. We refer to a foreign-owned subsidiary operating in our ten-country sample as a *foreign affiliate*; these are the organisational interfaces through which globally concentrated frontier knowledge becomes locally actionable.

Domestic firm definition. Our unit of analysis is *domestic firms*: firms that were never foreign-owned at any level — neither through direct equity stakes nor through the ultimate owner link — during the entire 1999–2012 sample period. Firms belonging to a domestic corporate group (where both the firm and its GUO are incorporated in the same country) are included, as these are typically small local groups without access to a foreign knowledge base. The formal three-condition definition is given in Appendix A.3.

Time-invariant initial GUO. Each target firm can in principle have multiple GUOs over the sample period as ownership restructuring occurs. For the purpose of constructing our exposure measure we assign a single, time-invariant *initial* GUO per target: the GUO first observed for that firm, with ties broken by frequency of appearance across the panel and, if still tied, alphabetically. As Panel B in Table 1 shows, 77.0 percent of firms in our sample have a single GUO throughout. Fixing the GUO at its initial value ensures that our exposure measure is predetermined with respect to sample-period ownership decisions, which is essential for identification (Section 3.2).

Ownership network statistics. Table 1 summarises the ownership network in our ten-country manufacturing sample. Of the 471,581 firms, 49,071 (10.4 percent) are linked to at least one GUO; of these, 71.2 percent have a domestic GUO and 28.8 percent a foreign GUO. There are 25,139 unique GUOs in total, of which 5,278 — 21.0 percent — own at least one foreign subsidiary in the sample. Figure 2 (discussed in Section 2.4) illustrates the stark difference

[4] We use the Moody's Independence Indicator to identify controlling owners. An entity is considered independent (and hence potentially the GUO) if no shareholder — other than unnamed individuals or public floats — holds more than 25 percent directly or in total. See Appendix A.3 for a full description of the independence indicator logic.

Table 1: Ownership Network: Ten-Country Manufacturing Sample

Entity	Count	Share
Panel A: Firms and their GUO links		
Firms with a GUO link	49,071	—
Domestic GUO	34,955	71.2%
Foreign GUO	14,116	28.8%
Unique GUOs	25,139	—
GUOs with ≥ 1 foreign subsidiary	5,278	21.0%
Panel B: GUO stability per target firm		
1 GUO throughout	37,795	77.0%
2 GUOs	8,790	17.9%
3 or more GUOs	2,486	5.1%

Notes: The sample covers manufacturing firms in Austria, Belgium, Germany, Spain, Finland, France, Italy, Norway, Portugal, and Sweden, 1999–2012. A GUO is defined as the entity at the top of the ownership chain for which no shareholder holds more than 25% of its equity (Moody's Independence Indicator A, B^+, or C). A GUO is classified as foreign if incorporated outside the host country. Panel B counts the number of distinct initial GUOs observed per target firm over the full sample period.

in innovation scale between domestic targets, foreign affiliates, and the GUOs that sit above them.

2.4 Patent Data: PATSTAT

Patent data come from the European Patent Office's PATSTAT database, linked to ORBIS firms via a unique BvD identifier bridge (August 2023 vintage). The matching pipeline, including BvD ID reconciliation across vintages, is described in Appendix A.2. Each patent is assigned to one or more of **35 technology classes** based on the International Patent Classification (IPC), using the OECD concordance between IPC codes and technical fields. Details on this classification are in Appendix B.1.

Outcome variable. Our main innovation outcome is the annual flow of new *patent families* for each firm, citation-weighted using forward citations received within a four-to-ten year window after the filing year. We use patent families rather than individual applications to avoid double-counting of the same underlying invention filed across multiple patent offices. Citation weighting captures patent quality, since more heavily cited patents represent more economically and technologically significant innovations (Hall et al., 2005).

Scale and concentration. Of the 471,581 manufacturing firms in our sample, **28,178** (6.0 percent) hold at least one patent directly (*patenting target firms*). Among GUOs, **3,917** patenting GUOs are linked to at least one target firm in the ten-country sample and collectively hold **3.75 million** unique patent families. The contrast in scale between targets and GUOs is striking: the median target holds 2 patent families with a maximum of 4,477, while the median GUO holds 12 families but the 90th percentile exceeds 1,960 and the 99th exceeds 108,000. Figure 2 plots the average cumulative stock of patents by group. This scale asymmetry — GUOs operate at a fundamentally different technological level than their affiliates and domestic neighbours — is central to the economic mechanism we study.

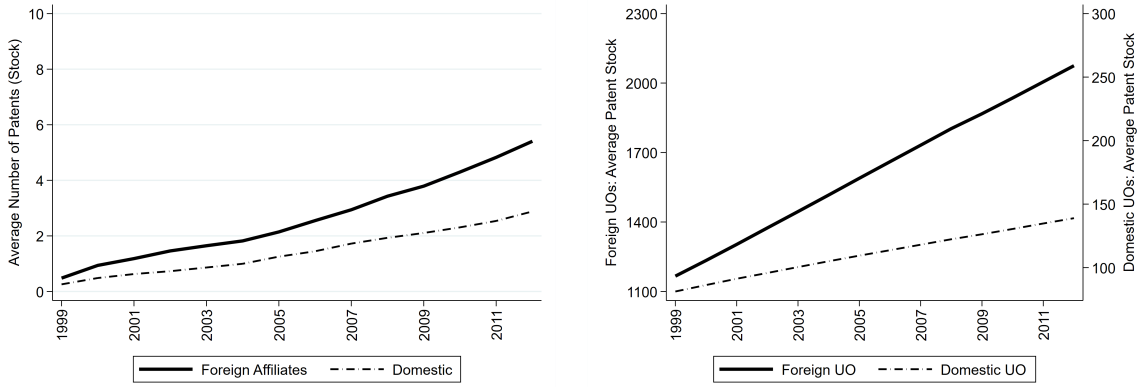


Figure 2: Cumulative Patent Stocks: Targets, Affiliates, and GUOs

Notes: Patent stock is the cumulative number of citation-weighted patent families filed up to year t since 1970. Left panel: domestic target firms vs. foreign affiliates (firms ever foreign-owned). Right panel: domestic GUOs (with country of incorporation inside the ten-country sample) vs. foreign GUOs (outside the sample). Vertical axis is in logs.

Geographic concentration. Figure 3 plots the geographic distribution of GUO patenting by country of incorporation. Bars report the number of patenting GUOs by country, while the dashed line reports each country's share of total GUO patent families. The figure shows that GUO innovation is highly concentrated across countries and firms. Japan alone accounts for approximately 57 percent of all GUO patent families despite having only 219 patenting GUOs in our sample, reflecting the concentration of patenting among a small number of large Japanese multinationals. The United States and Germany also account for large absolute patent stocks, but lower shares of total GUO families. This concentration implies that a small number of frontier innovators generate most of the potential foreign knowledge available to European domestic firms.

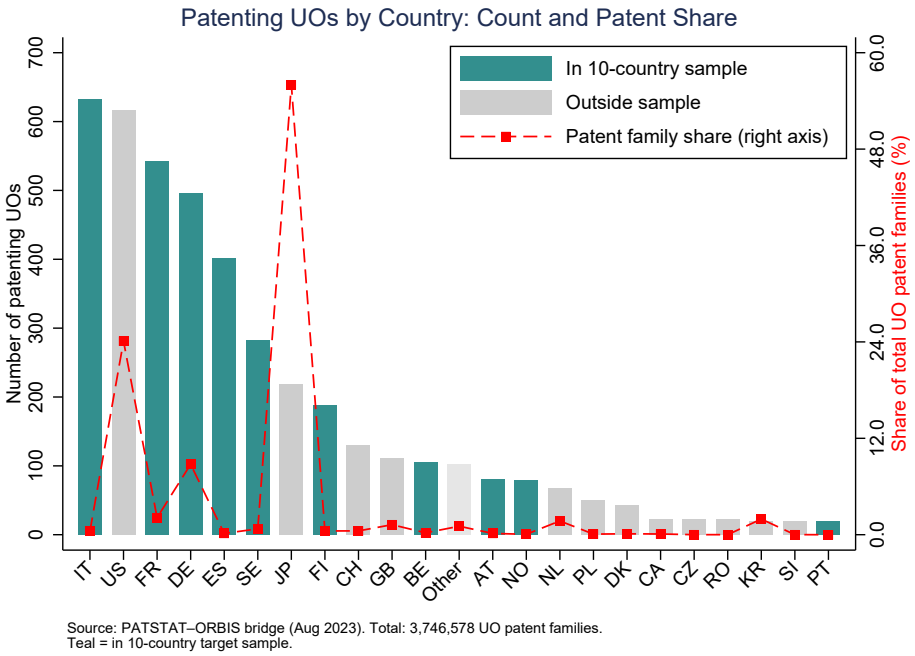


Figure 3: Geographic Concentration of GUO Patent Families

Notes: The figure reports the distribution of patenting Global Ultimate Owners (GUOs) and their patent families by GUO country of incorporation. Bars show the number of patenting GUOs. Teal bars indicate GUO countries within the ten-country sample; grey bars indicate GUO countries outside the sample. The dashed line reports each country's share of total GUO patent families, measured on the right axis. Total GUO patent families equal 3.75 million.

Patent stock growth: the shift. Figure 4 plots the evolution of the cumulative citation-weighted patent stock of foreign GUOs, indexed to 1999 = 100. This time-series variation captures the growth of the foreign knowledge stock that enters our exposure measure. In the baseline specification, we aggregate over all foreign GUOs: domestic firms are differentially exposed to this growth because they differ in their pre-sample technological proximity to the GUOs represented among patenting foreign affiliates in their country-sector cell. Thus, the relevant shift is the expansion of the foreign GUO knowledge stock, while the shares are domestic firms' predetermined technological proximity weights. In robustness exercises, we examine whether the effects differ by GUO country of origin, including for U.S. GUOs.

2.5 Technological Proximity

Construction. We follow Bloom et al. (2013) in measuring the technological distance between any pair of firms. For each firm (or GUO) i, we construct a *patent-class portfolio vector*

$$T_i = (T_{i1}, T_{i2}, \ldots, T_{i35}),$$

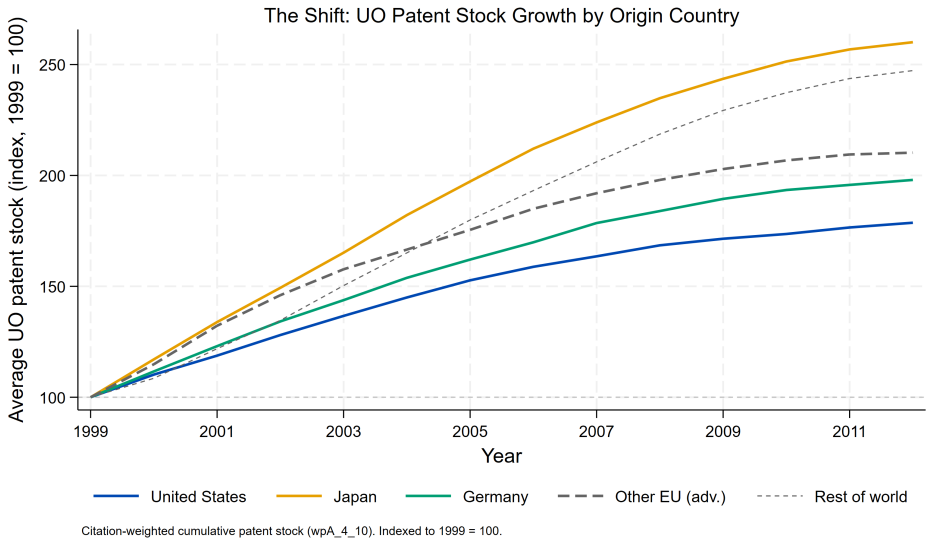


Figure 4: Growth in GUO Citation-Weighted Patent Stocks by Country

Notes: The figure plots the cumulative citation-weighted patent stock of foreign Global Ultimate Owners (GUOs), indexed to 1999 = 100. The baseline foreign technology exposure measure aggregates over all foreign GUOs associated with patenting affiliates in a domestic firm's country-sector cell. Cross-country series, when shown, are descriptive and illustrate heterogeneity in the growth of GUO knowledge stocks by country of incorporation; the baseline exposure measure does not use country-specific shocks separately.

where $T_{i\tau}$ is the share of firm i's patents in IPC technology class $\tau \in \{1, \ldots, 35\}$, computed using all patents filed *before* 1999 — that is, from the pre-sample period. Using pre-sample patents is essential: it ensures that the proximity weights are determined prior to any outcome variation and cannot be manipulated by domestic firms in response to sample-period productivity or innovation shocks.

Technological closeness for a firm pair (i, j) is the uncentered correlation of their portfolio vectors, following Jaffe (1986) and Bloom et al. (2013):

$$\text{TECH}_{ij} = \frac{T_i T_j'}{(T_i T_i')^{1/2}(T_j T_j')^{1/2}}. \tag{1}$$

This measure ranges from 0 to 1, is symmetric ($\text{TECH}_{ij} = \text{TECH}_{ji}$), and increases with greater overlap in the distribution of patents across technology classes. A value of 1 implies identical portfolio composition; a value of 0 implies no technology class in common.

In practice, we compute TECH_{ij} between each domestic target firm i and the GUO $u(j)$ of each foreign affiliate j operating in the same country and four-digit sector. We use the GUO's patent portfolio — not the affiliate's own patents — because the GUO represents the full technological base that an affiliate can draw on and potentially transmit to neighbouring domestic firms.

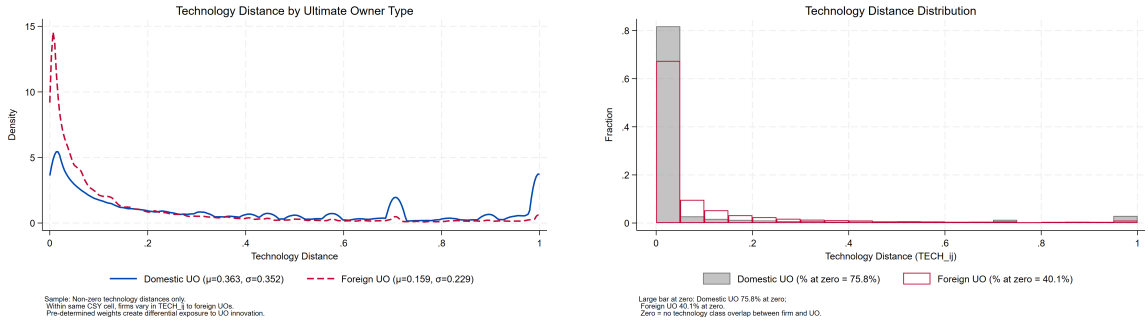


Figure 5: Distribution of Technological Proximity TECH_{ij}

Notes: Left panel: full distribution including zeros, separately for domestic-firm–domestic-GUO pairs (teal) and domestic-firm–foreign-GUO pairs (orange). The share of zero-proximity pairs is 75.8% for domestic GUOs and 40.1% for foreign GUOs. Right panel: kernel density estimate for the subsample of pairs with strictly positive proximity, with summary statistics.

Distribution and identifying variation. Figure 5 displays the distribution of TECH_{ij} for domestic-to-foreign-GUO pairs and domestic-to-domestic-GUO pairs. The most salient feature is the large mass at zero: 75.8 percent of domestic–domestic GUO pairs share no technology class overlap, compared with 40.1 percent of domestic–foreign GUO pairs. This reflects the fact that foreign GUOs tend to operate across a wider range of technology classes, creating more non-trivial proximity values with domestic firms. Among pairs with positive proximity, the distribution is right-skewed, with a median of approximately 0.07 and a long right tail reaching values above 0.8 for firms in closely related technological niches.

The identifying variation in our design comes from *within-sector* differences in TECH_{ij}: two firms in the same country-sector-year cell face the same set of foreign affiliates but differ in how close their pre-sample patent portfolios are to the GUOs of those affiliates. Figure 6 in Section **??** confirms that this within-sector variation is economically meaningful: the coefficient of variation of firm-level average TECH_{ij} exceeds zero in every two-digit sector and averages 0.40.

Sectoral proximity. Alongside TECH_{ij}, we construct an analogous *sectoral proximity* measure SEC_{ij} using firms' distributions across NACE sectors rather than patent classes. This follows the product-market closeness measure of Bloom et al. (2013), replacing patent-class portfolio vectors with sector-share vectors:

$$\mathrm{SEC}_{ij} = \frac{S_i S_j'}{(S_i S_i')^{1/2}(S_j S_j')^{1/2}},$$

where $S_i = (S_{i1}, \ldots, S_{iK})$ stacks firm i's revenue shares across NACE activities. SEC_{ij} captures product-market overlap: firms with similar sector portfolios compete in similar output markets. The correlation between TECH_{ij} and SEC_{ij} is low by construction, as patent-class overlap and sector overlap capture different dimensions of firm similarity. Details of the sectoral proximity construction are in Appendix B.2.

2.6 The Foreign Technology Exposure Measure

Shift-share construction. For domestic firm i in country c, four-digit sector s, and year t, we define *foreign technology exposure* as the following shift-share sum over all foreign affiliates k operating in the same country-sector cell:

$$\mathrm{FOtech}_{it} = \sum_{k \in \mathcal{F}(c,s)} \underbrace{\mathrm{TECH}_{i,\,u(k)}}_{\substack{\text{pre-sample} \\ \text{proximity share}}} \times \underbrace{\mathrm{PatStock}_{u(k),\,t}}_{\substack{\text{cumulative patent} \\ \text{stock of GUO}}} \tag{2}$$

where $\mathcal{F}(c, s)$ is the set of foreign-owned affiliates in country c and four-digit sector s, $u(k)$ denotes the initial GUO of affiliate k (assigned as described in Section 2.3), $\mathrm{TECH}_{i,u(k)}$ is the technological proximity between domestic firm i and GUO $u(k)$ computed from pre-sample patents, and $\mathrm{PatStock}_{u(k),t}$ is the cumulative citation-weighted patent stock of GUO $u(k)$ accumulated up to year t.

The structure of equation (2) has a clear Bartik interpretation (Goldsmith-Pinkham et al., 2020): the **shares** $\mathrm{TECH}_{i,u(k)}$ are pre-determined and firm-specific (measuring how much firm i could potentially absorb from GUO $u(k)$'s knowledge base), while the **shifts** $\mathrm{PatStock}_{u(k),t}$ are time-varying and GUO-specific (measuring how much new frontier knowledge has been accumulated by GUO $u(k)$ over time). A domestic firm sees its FOtech_{it} increase whenever GUOs that are technologically close to it expand their patent portfolios.

Why cumulative stocks rather than flows. We use the accumulated patent stock rather than annual patent flows for two reasons. First, a domestic firm can learn from the GUO's *entire* accumulated knowledge base, not only from innovations introduced in the most recent year; the relevant learning opportunity at time t is the full portfolio available for absorption. Second, stocks are considerably smoother than flows, reducing the influence of year-to-year volatility in filing activity that may be unrelated to the underlying pace of frontier R&D.

The four parallel measures. We construct three additional measures using the same shift-share structure:

$$\text{FOcomp}_{it} = \sum_{k \in \mathcal{F}(c,s)} \text{SEC}_{i,u(k)} \times \text{PatStock}_{u(k),t} \tag{3}$$

$$\text{DOtech}_{it} = \sum_{k \in \mathcal{D}(c,s)} \text{TECH}_{i,u(k)} \times \text{PatStock}_{u(k),t} \tag{4}$$

$$\text{DOcomp}_{it} = \sum_{k \in \mathcal{D}(c,s)} \text{SEC}_{i,u(k)} \times \text{PatStock}_{u(k),t} \tag{5}$$

where $\mathcal{D}(c,s)$ is the set of domestically-owned firms in the same country-sector cell. The *foreign sectoral exposure* FOcomp_{it} replaces patent-class proximity with sectoral proximity, capturing product-market competition from foreign-GUO innovation. The *domestic technology exposure* DOtech_{it} and *domestic sectoral exposure* DOcomp_{it} are the domestic-firm analogues. In all four measures, the patent-stock shift is identical; what varies is the proximity weight. Including FOcomp_{it}, DOtech_{it}, and DOcomp_{it} as controls in every regression ensures that our estimate of FOtech_{it} captures the technology-diffusion channel specifically, net of product-market competition from foreign firms, domestic knowledge spillovers, and domestic product-market exposure.

LP estimation sample. We restrict the estimation sample to domestic firms located in country-sector cells where at least one patenting foreign affiliate and its patenting GUO both operate. The resulting LP estimation sample contains **18,075** unique domestic firms (15,886 with non-missing lagged FOtech), yielding 12,999 firm-year observations at horizon $h = 0$ declining to 3,105 at $h = 5$ as the unbalanced panel thins at longer horizons.

Distribution of the exposure measure. The distribution of FOtech_{it} is highly right-skewed, consistent with the concentrated nature of frontier innovation documented in Section 2.4:

$$p_{25} = 0, \quad p_{50} = 0.0003, \quad p_{75} = 0.15, \quad p_{90} = 2.28, \quad p_{99} = 34.5$$

(mean $= 1.61$, standard deviation $= 8.85$). The large mass at or near zero reflects firms in cells with foreign affiliates whose GUOs are technologically distant, while the long right tail reflects a small number of firms operating in cells with affiliates of highly patent-intensive GUOs with whom they share substantial technology-class overlap. This skewness motivates

the discrete LP-DiD treatment design in Section 4.3, which defines a "large technology shock" as the first crossing of the 75th percentile.

3 Empirical Strategy

3.1 Specification

Our main interest is in the effect of foreign technology exposure on the innovation output of domestic firms (measured by citation-weighted patent flows).

Patent local projections. For domestic firm i in country c, four-digit sector s, and year t, we estimate a sequence of PPML regressions at horizons $h = 0, 1, \ldots, 5$:

$$y_{i,t+h} = \exp\Big(\beta_h \,\mathrm{FOtech}_{it-1} + \gamma_h \,\mathrm{FOcomp}_{it-1} + \delta_h \,\mathrm{DOtech}_{it-1} + \phi_h \,\mathrm{DOcomp}_{it-1} + \alpha_i + \lambda_{c,s,t+h}\Big) \cdot \varepsilon_{it}, \tag{6}$$

where $y_{i,t+h}$ is the citation-weighted patent flow of firm i at horizon $t+h$, α_i is a firm fixed effect, and $\lambda_{c,s,t+h}$ is a country $\times$ four-digit-sector $\times$ year fixed effect applied at horizon $t+h$. Standard errors are clustered at the firm level.

The PPML estimator is appropriate because the patent outcome is non-negative with a large mass of zeros. Unlike log-like transformations such as $\log(1+y)$ or $\mathrm{asinh}(y)$, whose treatment-effect interpretation depends on the units of measurement when zeros are present (Chen and Roth, 2024), PPML estimates an exponential conditional mean directly and is consistent under arbitrary misspecification of the conditional distribution, provided the conditional mean is correctly specified (Santos Silva and Tenreyro, 2006).

Each horizon h is estimated as a separate regression; we do not impose dynamics across horizons. This local projection approach (Jorda, 2005) directly traces the impulse-response function of domestic patenting to a shift in FOtech_{it-1}, without requiring correct specification of the full lag structure. Coefficients are reported in standardized form — the effect of a one-standard-deviation increase in each regressor, computed within the estimation sample at each horizon.

The fixed effect structure is as follows: firm fixed effects absorb all time-invariant differences in firm-level innovation propensity and productivity. Country $\times$ sector $\times$ year fixed effects absorb any shocks that are common to all firms in the same local industry-year cell: aggregate demand fluctuations, sector-specific technology trends, country-level policy changes, and the entry or exit of affiliates that affect all firms in a cell equally.

3.2 Identification

Source of variation. The key identifying variation comes from *within-cell* differences in FOtech_{it-1} across domestic firms in the same country-sector-year cell. Two domestic firms operating in the same local industry at the same time face identical sets of foreign affiliates, but differ in how technologically proximate they are to the GUOs of those affiliates. Because the proximity weights TECH_{ik} are fixed at their pre-sample values, any change in FOtech_{it} during the sample period is driven entirely by changes in the GUO's global patent stock. This design is analogous to a Bartik shift-share instrument (Goldsmith-Pinkham et al., 2020): the pre-determined exposure shares TECH_{ik} interact with the time-varying innovation shifts PatStock_{kt} driven by the GUO's global R&D. Critically, identification does not primarily come from affiliate entry or ownership restructuring: we hold foreign status fixed at the initial GUO classification throughout the sample. The variation we exploit is therefore periods in which the same foreign-owned affiliate becomes more innovation intensive because its GUO increases frontier patenting activity.

Threat 1: Correlated shocks within country-sector-year cells. Common productivity, demand, or policy shocks could affect both foreign-affiliate exposure and domestic patenting. The country $\times$ sector $\times$ year fixed effects $\lambda_{c,s,t+h}$ absorb all shocks common to firms in the same local industry-year cell. Identification is therefore immune to any shock — cyclical or structural — that affects the entire cell equally.

Threat 2: Endogenous affiliate location. Foreign GUOs may locate affiliates in sectors or countries where domestic innovation potential is already high, inducing a spurious correlation between FOtech_{it} and domestic patenting. The exposure shares TECH_{ik} are constructed from pre-sample (pre-1999) patent portfolios, before the start of the sample period over which outcomes are measured. They are predetermined with respect to any firm-level productivity or innovation shock during 1999–2012 and are outside the control of the domestic firm. Furthermore, because we hold ownership fixed at its initial classification, our design does not mechanically pick up periods in which previously unaffiliated domestic firms are acquired by foreign GUOs. While the set of foreign affiliates in each country-sector cell changes over time, fixing GUO assignments at their initial values ensures that variation in FOtech_{it} reflects changes in GUO global patenting activity rather than changes in ownership structure.

Threat 3: Confounding with product-market competition. A domestic firm that is technologically close to a foreign GUO may also compete in the same product market, conflating knowledge spillovers with competition effects. We include FOcomp_{it-1}, the exact parallel of FOtech_{it-1} but constructed using sectoral proximity SEC_{ij} in place of patent-class proximity TECH_{ij}. The coefficient on FOtech_{it-1} is therefore estimated net of any product-market exposure to the same set of foreign GUO patent shocks. The correlation between TECH_{ij} and SEC_{ij} is low by construction, since patent-class overlap and sector overlap capture distinct dimensions of firm similarity.

Threat 4: Domestic knowledge spillovers as a confound. Firms that are technologically close to foreign GUOs may also be close to domestic innovators, so any response in $y_{i,t+h}$ could reflect domestic spillovers rather than the foreign technology channel. We include DOtech_{it-1} and DOcomp_{it-1}, the domestic analogues of FOtech_{it-1} and FOcomp_{it-1}, which absorb the domestic spillover variation. The coefficient on FOtech_{it-1} therefore captures the incremental effect of foreign technology exposure over and above what domestic innovation activity can explain.

Threat 5: Pre-existing innovation trends. Domestic firms with higher foreign technology exposure may already be on steeper patenting trajectories before any change in FOtech_{it}, due to unobserved heterogeneity correlated with TECH_{ik}. We test this directly in Section 4.2 by extending the LP to negative horizons $h = -3, -2, -1$. We find that all three pre-horizon coefficients are statistically indistinguishable from zero, providing no evidence of differential pre-trends in patenting for firms with higher foreign technology exposure.

Threat 6: Limited within-cell variation in technology proximity. Our empirical design compares domestic firms within the same country-sector-year cell. Identification therefore requires meaningful variation in technological proximity across firms that face the same local foreign-affiliate environment. If TECH_{ij} varied only across sectors, or were nearly constant among firms within a sector, the fixed effects in equation (6) would absorb most of the relevant variation and the remaining estimates would be weakly identified or driven by a small number of cells.

We assess this concern directly in Figure 6, which plots the coefficient of variation of firm-level average TECH_{ij} within each two-digit NACE manufacturing sector. The coefficient of variation is positive in every sector, with a mean of 0.40, indicating substantial dispersion in

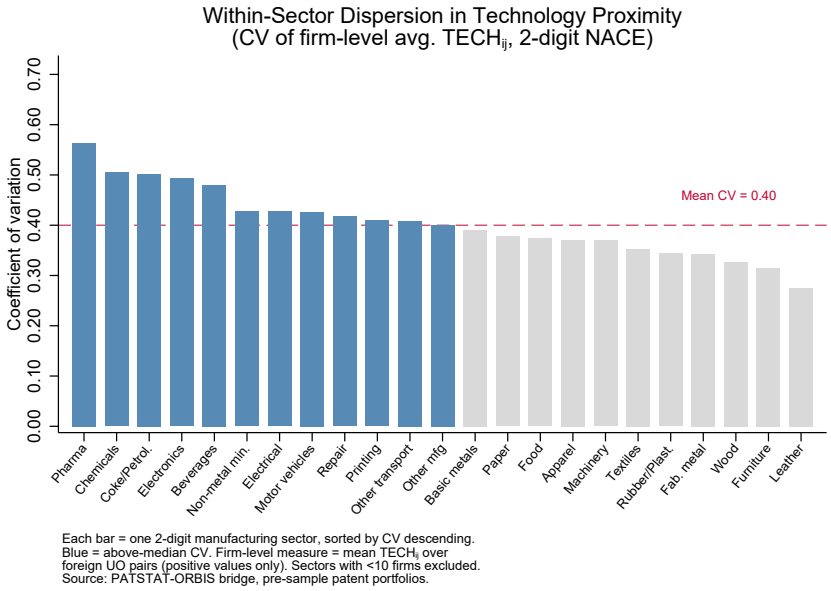


Figure 6: Within-Sector Variation in Technology Proximity

Notes: The figure plots the coefficient of variation (CV = SD/mean) of firm-level average TECH_{ij} within each two-digit NACE manufacturing sector. CV exceeds zero in every sector; mean CV = 0.40. Sectors are ordered by the number of domestic firms. Appendix C.4 presents the four-digit breakdown.

technological proximity among firms operating in the same broad sector. Dispersion is especially large in knowledge-intensive sectors such as Pharmaceuticals, Chemicals, and Electronics. Appendix C.4 repeats the exercise at the four-digit level and confirms that within-cell variation remains present at the level of the fixed effects used in the estimating equation. Thus, the estimates are not identified solely from cross-sector differences in exposure; they exploit within-sector differences in firms' pre-sample technological proximity to foreign GUOs.

4 Foreign Technology Exposure and Domestic Innovation

4.1 Baseline Local Projection Results

Figure 7 plots the PPML local projection coefficients on FOtech_{it-1} from equation (6) at horizons $h = -3$ through $h = 5$, with 95 percent confidence intervals. Table 2 reports the full set of coefficient estimates including the pre-horizon tests and all four spillover regressors.

The main result: delayed and persistent innovation response. The response of domestic patenting to foreign technology exposure is absent at short horizons and builds gradually over time. The coefficient on FOtech_{it-1} is statistically indistinguishable from zero at

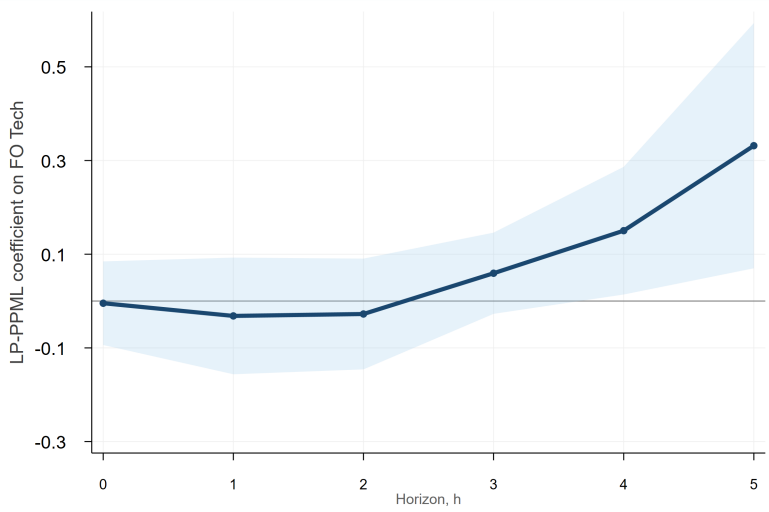


Figure 7: Local Projection Estimates: Foreign Technology Exposure and Domestic Patenting

Notes: The figure plots the PPML local projection coefficients $\hat{\beta}_h$ from equation (6) at horizons $h = 0, \ldots, 5$ for the foreign technology exposure variable FOtech_{it-1}, together with 95 percent confidence intervals (standard errors clustered by firm). Coefficients are standardized: effect of a one-standard-deviation increase in FOtech_{it-1} within the estimation sample at each horizon. The dependent variable is citation-weighted patent families. Fixed effects: firm; country $\times$ four-digit sector $\times$ year (forward). The full table including pre-horizon coefficients and all four regressors is reported in Table 2.

$h = 0, 1, 2$, and 3, but becomes positive and significant at $h = 4$ (coefficient 0.150, $p < 0.05$) and $h = 5$ (coefficient 0.332, $p < 0.05$). Since the specification uses PPML with an exponential conditional mean, these coefficients represent semi-elasticities. A one-standard-deviation increase in FOtech_{it-1} is associated with $e^{0.150} - 1 \approx 16$ percent higher citation-weighted patenting after four years and $e^{0.332} - 1 \approx 39$ percent higher patenting after five years. The raw (non-standardized) coefficients — reported in Appendix D.1 — imply that a one-unit increase in FOtech_{it-1} raises patenting by approximately 1 percent at $h = 4$ and 2.6 percent at $h = 5$; these translate to the standardized magnitudes above when scaled by the within-sample standard deviation of FOtech (approximately 15 units).[5]

The gradual build-up is economically meaningful and consistent with the technology-diffusion mechanism we have in mind. Learning from a frontier patent portfolio is not instantaneous: domestic firms must first identify which elements of the GUO's knowledge base are applicable to their own production processes, then invest in the complementary skills, equipment, and human capital needed to act on that knowledge. Only after this adaptation period — which our estimates place at roughly three to four years — do new patentable inno-

[5]The standardization is computed within the estimation sample at each horizon, so the scaling factor varies slightly across horizons due to sample composition changes.

Table 2: Foreign Technology Exposure and Domestic Patenting: Local Projection Estimates

	Pre-trend test			Post-horizons					
	$h{=}-3$	$h{=}-2$	$h{=}-1$	$h{=}0$	$h{=}1$	$h{=}2$	$h{=}3$	$h{=}4$	$h{=}5$
Panel A: Foreign technology and competition exposure									
FOtech_{it-1}	0.071	−0.039	−0.007	−0.005	−0.032	−0.028	0.059	0.150**	0.332**
	(0.112)	(0.085)	(0.053)	(0.046)	(0.064)	(0.060)	(0.044)	(0.070)	(0.134)
FOcomp_{it-1}	−0.219	−0.156**	−0.028	−0.029	0.052	0.113*	0.104	0.172***	−0.034
	(0.144)	(0.073)	(0.076)	(0.040)	(0.077)	(0.061)	(0.105)	(0.056)	(0.067)
Panel B: Domestic technology and competition exposure									
DOtech_{it-1}	0.037	0.007	−0.011	0.038	0.022	−0.082*	0.143	0.649**	−0.078
	(0.023)	(0.053)	(0.078)	(0.055)	(0.084)	(0.048)	(0.158)	(0.324)	(0.145)
DOcomp_{it-1}	0.045**	0.003	−0.029	0.011	0.047	0.076**	0.006	−0.030	−0.319
	(0.021)	(0.031)	(0.036)	(0.041)	(0.060)	(0.038)	(0.085)	(0.100)	(0.372)
Observations	7,035	8,549	10,384	12,999	9,479	7,309	5,636	4,184	3,105
Firms	1,148	1,311	1,523	1,837	1,408	1,134	929	758	607
Pseudo R^2	0.372	0.385	0.351	0.365	0.344	0.370	0.324	0.362	0.351

Notes: PPML local projections. Dependent variable: citation-weighted patent families (four-to-ten year citation window, family-adjusted). All coefficients are standardized: effect of a one-standard-deviation increase in each regressor within the estimation sample at each horizon. Fixed effects: firm; country × four-digit sector × year (at horizon $t+h$ for post-horizons; at $t-|h|$ for pre-horizon tests). Standard errors clustered by firm in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$. The estimation sample covers domestic firms in country-sector-year cells where at least one patenting foreign affiliate and its patenting GUO both operate. The raw (non-standardized) version of this table is in Appendix D.1.

vations emerge.

The other regressors. The foreign sectoral exposure FOcomp_{it-1} enters with a positive and significant coefficient at $h = 4$ (0.172, $p < 0.01$) but is near zero at other horizons and turns negative at $h = 5$. Results are imprecise and not robust across specifications. The domestic technology exposure DOtech_{it-1} has a large positive coefficient at $h = 4$ (0.649, $p < 0.05$) but is insignificant at most other horizons, suggesting heterogeneous and less precisely estimated domestic knowledge spillovers. The domestic sectoral measure DOcomp_{it-1} is small and only marginally significant at $h = 2$. Importantly, the inclusion of these three controls ensures that the coefficient on FOtech_{it-1} captures the technology-diffusion channel specifically, net of competing explanations.

4.2 Pre-Trend Tests and Robustness

Pre-trend tests. A key threat to identification is that domestic firms with higher foreign technology exposure were already on steeper patenting trajectories before any change in

FOtech_{it}. Table 2 directly addresses this by extending the local projection to negative horizons $h = -3, -2, -1$. For FOtech_{it-1}, all three pre-horizon coefficients are statistically indistinguishable from zero (0.071, -0.039, and -0.007 respectively), with standard errors larger than or comparable to the point estimates. There is no evidence of a pre-existing upward or downward trend in patenting for firms with higher foreign technology exposure. The effect of FOtech_{it-1} therefore emerges only at longer post-horizons, consistent with a genuine causal effect operating through a slow knowledge-diffusion channel.

Robustness across specifications. Table 3 reports the FOtech_{it-1} coefficient across a range of alternative specifications at each horizon from $h = 0$ to $h = 5$. The baseline result is highly stable.

The first two rows confirm that the standardization is not driving the result. The raw (non-standardized) coefficient of 0.0096 at $h = 4$ and 0.0257 at $h = 5$ correspond to a 1 and 2.6 percent increase in patenting per unit of FOtech_{it-1}; multiplied by the within-sample standard deviation of approximately 15, these recover the baseline standardized magnitudes. Dropping the FOcomp, DOtech, and DOcomp controls barely affects the coefficients (row 3), confirming that the technology spillover estimate is not dependent on the particular control structure.

Restricting the sample to firms with at least ten employees (row 4: EMPL$\geq$10) yields *larger* coefficients at the relevant horizons (0.173^{**} at $h = 4$, 0.402^{**} at $h = 5$), suggesting that the technology-diffusion mechanism is if anything stronger among firms large enough to have the internal capacity to act on foreign knowledge. Limiting to affiliates classified as industrial firms (type *FI*, row 5) leaves coefficients virtually unchanged, ruling out financial holding structures as a driver. The specification using only US-owned affiliates (row 6: F_{us}) produces a somewhat smaller but still significant response at $h = 5$ (0.214^{**}), consistent with US multinationals — one of the most innovation-intensive groups in our sample — being an important but not exclusive source of spillovers.

The split by affiliate ownership type is informative about the mechanism. Large foreign *business groups* (row 7: F_B) — affiliates embedded in multi-plant conglomerate networks — produce larger and more precisely estimated coefficients (0.204^{***} at $h = 4$, 0.387^{***} at $h = 5$). In contrast, non-business-group foreign owners (row 8: F_{NB}, which includes individual investors, financial intermediaries, and sovereign wealth funds) produce negative or insignificant coefficients at $h = 4$ and $h = 5$. This contrast is consistent with the recombination mechanism: codified frontier knowledge becomes actionable only when the local affiliate actively

Table 3: Robustness: FOtech_{it-1} Coefficient Across Specifications

Check	Description	h=0	h=1	h=2	h=3	h=4	h=5
Baseline	Standardized, all controls	−0.005	−0.032	−0.028	0.059	0.150**	0.332**
	Raw (non-standardized)	−0.0003	−0.0021	−0.0018	0.0039	0.0096**	0.0257**
	Raw, no FOcomp/DOtech/DOcomp controls	−0.0005	−0.0026	−0.0024	0.0038	0.0089*	0.0247***
Sample	EMPL$\geq$ 10	−0.004	0.017	−0.014	0.069	0.173**	0.402**
F_I	Industrial affiliates only	−0.006	−0.032	−0.027	0.059	0.151**	0.321**
F_{us}	US-owned affiliates only	−0.063	−0.111*	−0.136**	0.018	0.066	0.214**
F_B	Large foreign business groups	0.006	−0.036	−0.049	0.062	0.204***	0.387***
F_{NB}	Non-business-group owners	0.096	0.133	0.299*	−0.111	−0.313**	−0.230

Notes: Each row reports the FOtech_{it-1} coefficient from a separate PPML local projection. Baseline specification uses all four controls (FOtech, FOcomp, DOtech, DOcomp), standardized coefficients, and the full sample. Rows 1–3 vary the scaling or control set; remaining rows vary the sample or affiliate classification. Financial groups (F_F) are not shown due to insufficient observations. Dynamic controls variants (adding lagged outcome to the specification) are in Appendix D.2. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

deploys technology, a condition satisfied for operating subsidiaries within business groups but not for passive financial investors holding equity stakes without local technological activity.

4.3 LP-DiD: Discrete Treatment Design

Motivation. The distribution of FOtech_{it} is highly right-skewed — with $p_{25} = 0$, $p_{50} = 0.0003$, $p_{75} = 0.15$, $p_{90} = 2.28$, and $p_{99} = 34.5$ (mean 1.61, standard deviation 8.85) — reflecting the episodic nature of frontier innovation: most domestic firms face little foreign technology exposure in most years, while a small number experience large, discrete surges when a frontier GUO in their technology space substantially expands its patent portfolio. This distributional feature motivates a complementary design in which we study what happens to domestic patenting around discrete threshold-crossing events (Dube et al., 2025). The LP-DiD design makes parallel-trends assumptions transparent, allows visual inspection of pre-treatment dynamics, and is robust to heterogeneous treatment effects across cohorts (Callaway and Sant'Anna, 2021).

Treatment definition. We discretize foreign technology exposure using the 75th percentile of the lagged exposure measure within the estimation sample:

$$D_{it}^{75} = \mathbf{1}\{L.\text{FOtech}_{it} > p_{75}(L.\text{FOtech})\}. \tag{7}$$

A firm enters treatment in the first year it crosses this threshold from below:

$$T_i = \min\{t : D_{it}^{75} = 1,\ D_{i,t-1}^{75} = 0\}.$$

This "first crossing" design focuses attention on the transition from low to high foreign technology exposure and avoids contamination from repeated treatment entries (Dube et al., 2025). Treatment is staggered across firms and years; the histogram of first-crossing years in Appendix Figure D.1 shows that treated firms are distributed broadly across 2000–2012 with no dominant cohort, ruling out identification from a single macroeconomic event.

Estimation and control group. We estimate the following event-study regression for event times $h = -3, \ldots, 5$:

$$\begin{aligned} y_{it} = & \sum_{h=-3}^{5} \tau_h \, \mathbf{1}\{t - T_i = h\} + \rho \, y_{i,t-1} \\ & + \theta_1 \, \mathrm{DOtech}_{i,t-1} + \theta_2 \, \mathrm{FOcomp}_{i,t-1} + \theta_3 \, \mathrm{DOcomp}_{i,t-1} \\ & + \lambda_{c,s,t} + \varepsilon_{it}. \end{aligned} \tag{8}$$

where $y_{it} = \mathrm{asinh}(\text{citation-weighted patent flow}_{it})$ and $\lambda_{c,s,t}$ is a country $\times$ four-digit sector $\times$ year fixed effect. The lagged outcome $y_{i,t-1}$ is included as a control for pre-existing patenting trends, and the three spillover controls DOtech, FOcomp, and DOcomp absorb confounding variation from domestic and sectoral technology shocks.

The control group excludes: (i) firms that are treated within the next five years (to prevent contamination from anticipatory responses); (ii) observations in first-treatment years (to avoid mechanical confounds at the crossing point itself); and (iii) firms that cross the threshold only once and then revert ("one-off crossers"), since these may reflect measurement noise rather than a genuine permanent increase in exposure. The control group therefore consists of **never-treated firms** and not-yet-treated firms in years at least five periods before their eventual treatment. Pre-treatment balance on key observables is shown in Appendix Figure D.2: the kernel density estimates of pre-treatment patent flows, $L.\mathrm{FOtech}$, log employment, and TFP levels overlap closely between treated and control firms, supporting the parallel-trends assumption.

Results. Figure 8 plots the estimated event-study coefficients $\hat{\tau}_h$ from equation (8). Three features stand out.

First, the pre-treatment coefficients at $h = -3, -2, -1$ are all close to zero and statistically insignificant, confirming no differential pre-trend in patenting between firms that will expe-

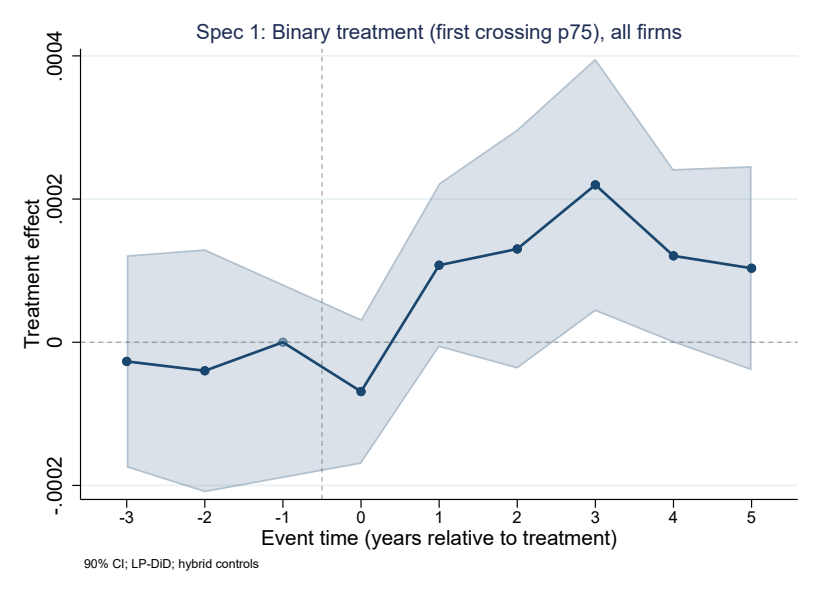


Figure 8: LP-DiD Event Study: Innovation Response to a Large Foreign Technology Shock

Notes: LP-DiD event-study estimates from equation (8). Treatment: first year domestic firm's lagged FOtech_{it} crosses the 75th percentile ($p_{75} = 0.15$). Outcome: asinh(citation-weighted patent flow). Event window: $[-3, 5]$. Control group: never-treated firms and firms not yet treated in the next five years, excluding first-treatment-year observations and one-off crossers. Controls: lagged outcome, $\mathrm{DOtech}_{i,t-1}$, $\mathrm{FOcomp}_{i,t-1}$, $\mathrm{DOcomp}_{i,t-1}$. Fixed effects: country $\times$ four-digit sector $\times$ year. Standard errors clustered by firm; 90 percent confidence intervals shown. Treatment timing histogram and pre-treatment balance checks are in Appendix D.4.

rience a large foreign technology shock and those that will not. This corroborates the PPML pre-trend tests in Section 4.2 and provides additional design-based support for the parallel-trends assumption.

Second, following the treatment event, innovation responses emerge gradually and accumulate over time, with the peak response at approximately $h = 3$. The coefficients rise from near zero at $h = 0$ and $h = 1$ to become positive and economically meaningful at $h = 2$ and $h = 3$, with the confidence interval excluding zero by the third post-treatment year. Effects remain positive at longer horizons, suggesting persistent technology adoption rather than a short-lived patent-timing response.

Third, the magnitude of the peak effect — approximately 1.5 percent of a standard deviation in annual citation-weighted patenting — is consistent with the PPML local projection estimates from Section 4.1. A discrete jump into the upper quartile of the FOtech distribution is a large shock, and the resulting innovation response is economically significant given that the outcome is the level of patenting, not its growth rate.

The delayed and persistent pattern is the signature of a genuine knowledge-diffusion mechanism. The delayed response is itself informative: it reflects the time required for do-

mestic firms to recognise which elements of the frontier are relevant to their own capabilities, invest in the complementary R&D, and translate recombined knowledge into patentable innovations — the same process that explains why publicly disclosed knowledge does not diffuse uniformly across firms. We explore which features of the affiliate environment determine whether this process succeeds in Section 5.

5 Channels

5.1 Technology-Intensive Affiliates as the Conduit

Hypothesis. The baseline results establish that domestic patenting rises following an increase in the patent stock of technologically proximate foreign GUOs. But the mechanism we proposed in Section 2 requires more than a large frontier patent portfolio. A GUO's patents are publicly readable everywhere; what activates their potential for local spillovers is the presence of a foreign affiliate that embeds the GUO's knowledge in the host-country engineering environment through its own applied R&D activity. The prediction is therefore that the patent response should be concentrated among domestic firms whose country-sector environment contains a technology-intensive affiliate — one that itself patents intensively — and absent for firms whose local affiliates are passive or financially oriented, however innovative their GUO may be.

Affiliate technology-intensity split. We test this prediction using the LP-DiD framework from Section 4.3. Among all treated domestic firms (those that experience a first crossing above p_{75} of $L.\mathrm{FOtech}_{it}$), we split at the crossing year T_i by whether the firm is exposed to a technology-intensive affiliate: a foreign affiliate in the same country-sector cell whose own citation-weighted patent stock at T_i exceeds the 75th percentile of the within-sector affiliate patent distribution. Formally:

$$\text{CT-exposed}_i = \mathbf{1}\big\{\text{max citation-weighted patents of affiliates in cell}(c, s, T_i) > p_{75}\big\},$$

$$\text{Non-CT-exposed}_i = 1 - \text{CT-exposed}_i.$$

We estimate the LP-DiD event study of equation (8) separately for CT-exposed treated firms (against never-treated controls) and Non-CT-exposed treated firms (against the same control group). Both groups experience the same treatment event — a large crossing of the FOtech

threshold — so the split isolates whether the innovation response depends on the innovation intensity of the locally embedded affiliate rather than on the size of the GUO patent shock alone.

Results. Figure 9 presents the two event-study plots side by side.

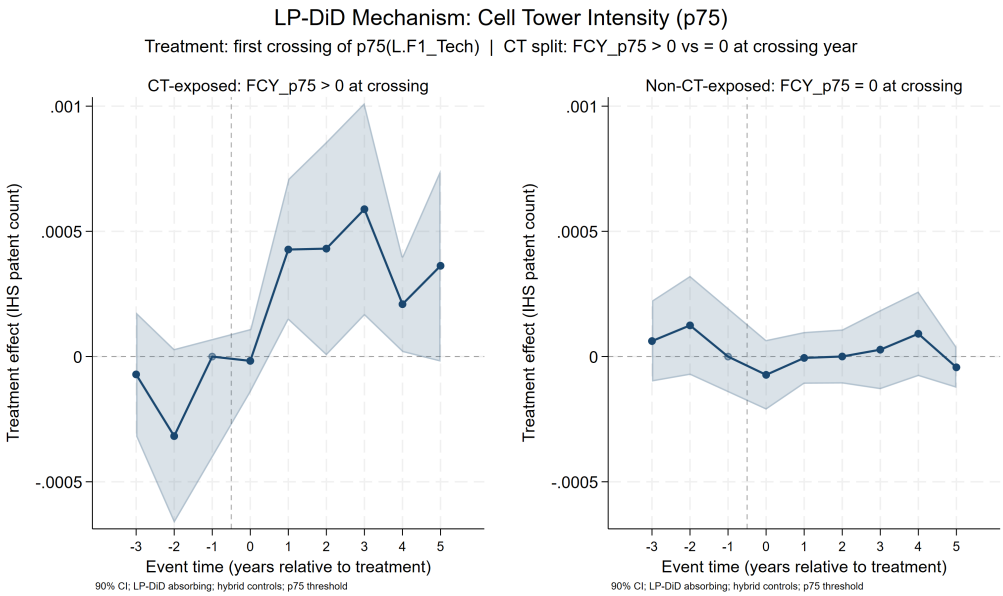


Figure 9: Technology-Intensive Affiliates Drive the Innovation Response

Notes: LP-DiD event studies estimated separately for CT-exposed (left panel, $N = 2{,}605$ treated firms) and Non-CT-exposed (right panel, $N = 2{,}487$ treated firms) domestic firms, against a common control group of $N = 12{,}809$ never-treated firms. Treatment: first crossing of $p_{75}(L.\text{FOtech}_{it})$. Technology-intensity split: at the crossing year T_i, a firm is classified as exposed to a technology-intensive affiliate if foreign affiliates in its country-sector cell have citation-weighted patent stock above p_{75} within the sector. Outcome: asinh(citation-weighted patent flow). Controls: lagged outcome, $\text{DOtech}_{i,t-1}$, $\text{FOcomp}_{i,t-1}$, $\text{DOcomp}_{i,t-1}$; country $\times$ four-digit sector $\times$ year FEs. Standard errors clustered by firm; 90 percent confidence intervals.

Among the **2,605 CT-exposed treated firms** (51.2 percent of all treated firms), the pre-treatment coefficients are close to zero and statistically insignificant, and a clear post-treatment patent buildup emerges from $h = 1$ onward, peaking around $h = 3$–4 and remaining positive at $h = 5$. The response mirrors the aggregate LP-DiD result from Figure 8 in both timing and magnitude.

Among the **2,487 Non-CT-exposed treated firms** (48.8 percent), the coefficients are near zero throughout the entire event window. The confidence intervals exclude any economically meaningful positive response at all horizons from $h = 0$ to $h = 5$. There is no evidence of a patent response for firms that experience the same large increase in GUO-level frontier knowledge but do not face a technology-intensive affiliate in their local environment.

The contrast between the two panels is sharp. Both groups start from comparable pre-treatment patenting trajectories, so the difference in post-treatment dynamics is not attributable to differential pre-trends or selection on initial innovation activity. Robustness using a PPML specification for this split is reported in Appendix D.5. The result implies that GUO portfolio size creates the *potential* for spillovers, but that potential is realised only when a locally innovating affiliate makes the frontier knowledge actionable. This is consistent with the recombination mechanism: domestic firms need both the technological proximity to understand the frontier and a local conduit through which to observe and build on it.

5.2 Absorptive Capacity

Hypothesis. Even among firms exposed to a technology-intensive affiliate, the ability to absorb frontier knowledge is not uniform. Cohen and Levinthal (1990) emphasise that a firm's capacity to recognise, assimilate, and exploit external knowledge depends critically on its own prior knowledge base. Firms with stronger initial innovation capabilities are better equipped to decode the technical content of the GUO's patent portfolio, identify applicable elements, and invest in the complementary R&D needed to build on them. We test this absorptive capacity hypothesis by asking whether the patent response to FOtech_{it-1} is stronger for domestic firms with a higher initial stock of innovative activity.

Proxy and specification. We measure absorptive capacity using each domestic firm's citation-weighted patent flow in its first observed year in the sample, $\mathrm{PatFlow}_i^{base}$, and define a binary indicator:

$$\mathrm{HighAC}_i = \mathbf{1}\{\mathrm{PatFlow}_i^{base} > \mathrm{median}(\mathrm{PatFlow}^{base})\}, \quad \mathrm{LowAC}_i = 1 - \mathrm{HighAC}_i.$$

Using the first observed year ensures the proxy is predetermined with respect to sample-period dynamics and does not change over time, eliminating concerns about reverse causality from productivity shocks to patenting behaviour. We then estimate the PPML local projection

with the four spillover measures fully interacted with the absorptive capacity indicator:

$$y_{i,t+h} = \exp\Big(\beta_h^H \left[\widetilde{\text{FOtech}}_{it-1} \times \text{HighAC}_i\right] + \beta_h^L \left[\widetilde{\text{FOtech}}_{it-1} \times \text{LowAC}_i\right] + \delta_h^H \left[\widetilde{\text{DOtech}}_{it-1} \times \text{HighAC}_i\right] + \delta_h^L \left[\widetilde{\text{DOtech}}_{it-1} \times \text{LowAC}_i\right] + \gamma_h \widetilde{\text{FOcomp}}_{it-1} + \phi_h \widetilde{\text{DOcomp}}_{it-1} + \alpha_i + \lambda_{c,s,t+h}\Big) \cdot \varepsilon_{it}. \tag{9}$$

where tildes denote standardization. The test of absorptive capacity is $H_0 : \beta_h^H = \beta_h^L$, with the prediction that $\beta_h^H > \beta_h^L$.

Results. Table 4 reports the estimated coefficients from equation (9). The pattern is striking and consistent across specifications.

Table 4: Absorptive Capacity Amplifies the Innovation Spillover

	h=0	h=1	h=2	h=3	h=4	h=5
FOtech × High AC	0.002	−0.007	−0.022	0.029	0.120**	0.405**
	(0.027)	(0.038)	(0.036)	(0.025)	(0.050)	(0.176)
FOtech × Low AC	−0.092	−0.195**	0.115	0.189	−0.456*	−0.669
	(0.072)	(0.087)	(0.186)	(0.191)	(0.255)	(0.484)
DOtech × High AC	0.029	−0.010	−0.064	0.166**	0.174	0.154
	(0.035)	(0.079)	(0.047)	(0.082)	(0.339)	(0.183)
DOtech × Low AC	0.014	0.041	−0.053	0.068	0.505**	−0.099
	(0.048)	(0.060)	(0.034)	(0.102)	(0.239)	(0.092)
Wald p: High = Low (FOtech)	0.160	**0.024**	0.459	0.403	**0.039**	**0.040**
Observations	12,999	9,479	7,309	5,636	4,184	3,105
Firms	1,837	1,408	1,134	929	758	607

Notes: PPML local projections from equation (9). High AC: baseline patent flow above sample median; Low AC: at or below median. FOcomp and DOcomp are included as uninteracted controls (not shown). All coefficients are standardized: effect of a one-standard-deviation increase in the corresponding interaction variable. Fixed effects: firm; country × four-digit sector × year. Standard errors clustered by firm. Basic specification and EMPL≥ 10 robustness in Appendix D.3. $^*p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Among **High-AC firms**, the coefficient on FOtech_{it-1} rises from near zero at short horizons to 0.120^{**} at $h = 4$ and 0.405^{**} at $h = 5$. These translate to $e^{0.120} - 1 \approx 13$ percent higher citation-weighted patenting after four years and $e^{0.405} - 1 \approx 50$ percent higher patenting after five years for a one-standard-deviation increase in foreign technology exposure. Among **Low-AC firms**, the coefficient is negative or statistically insignificant at all horizons, with a significantly negative estimate at $h = 1$ (-0.195^{**}) and $h = 4$ (-0.456^{*}) suggesting that firms at or below the median initial patent stock may actually experience adverse competitive effects from the same foreign technology shock.

The Wald test of equality between High-AC and Low-AC coefficients rejects at $h = 1$ ($p = 0.024$), $h = 4$ ($p = 0.039$), and $h = 5$ ($p = 0.040$). The pattern of rejection is consistent with the timing story: at very short horizons both groups are unaffected, but as the learning-and-adoption window opens, the differential capacity to exploit frontier knowledge produces diverging innovation trajectories.

Interpretation. Taken together, these results sharpen the picture from the affiliate technology-intensity split. Foreign technology diffusion operates through two sequential conditions: (i) a technology-intensive affiliate must be locally present to transmit tacit knowledge from the GUO's portfolio, and (ii) the domestic firm must have sufficient prior innovative capacity to absorb and build on that knowledge. Technological proximity creates the exposure channel; absorptive capacity determines whether exposure translates into new patentable ideas. Firms below the median patent stock — which in our sample is a low bar, as the median first-year patent flow is a handful of families — do not benefit and may in fact be crowded out by the superior innovation intensity of the foreign affiliate. Full robustness tables including the Basic specification and the EMPL≥ 10 sample are in Appendix D.3. The two sequential conditions together explain why the same publicly disclosed patent portfolio generates very different outcomes across the firms and countries positioned to benefit from it.

6 From Innovation to Productivity

The patent evidence in Sections 4 and 5 shows that domestic firms innovate more following exposure to frontier GUO knowledge when a technology-intensive affiliate is locally present and the domestic firm has sufficient absorptive capacity. A natural question is whether this innovation response propagates to measured total factor productivity. The prior literature on horizontal FDI spillovers is often pessimistic: Aitken and Harrison (1999) find negative within-sector TFP effects of foreign entry, a finding echoed in many subsequent studies. Our framework provides a way to reconcile these results. For domestic firms that are not technologically proximate to foreign affiliates, or that lack the absorptive capacity to benefit from them, competitive forces may dominate — not only through product-market competition but also through input-market competition, especially for skilled labour. For such firms, the net TFP effect of foreign presence may be zero or negative. By contrast, for firms that are technologically proximate to a locally innovating affiliate, the knowledge channel should be stronger and may dominate the competitive channel.

We test this prediction in two steps. We first estimate an average static specification in the full sample, asking whether foreign technology exposure is associated with higher TFP once we condition on the same controls and fixed effects used in the patent analysis. We then use the LP-DiD affiliate technology-intensity split to compare dynamic TFP responses for firms exposed to technology-intensive affiliates versus those that are not, mirroring the patent analysis and asking whether productivity gains are concentrated among the same firms for which the innovation response is strongest.

The identification concerns discussed in Section 3.2 also apply here. We address domestic-spillover confounds by controlling for DOtech_{it-1} and DOcomp_{it-1}, and pre-existing trends by testing for differential pre-trends in the LP-DiD specification. We also verify that the shift-share structure is not driven mechanically by a small number of influential GUOs. Following Goldsmith-Pinkham et al. (2020), we decompose the identifying variation into GUO-specific contributions and show that dominant GUOs carry signs consistent with the technology-diffusion mechanism. Following Adão et al. (2019), we re-cluster standard errors at the dominant-GUO level; the AKM standard errors are approximately $1.17\times$ the firm-clustered baseline, and the main patent and TFP results remain significant under GUO-level clustering. Following Borusyak et al. (2022), a BHJ overidentification test using the top-five GUO-specific instruments yields a Hansen J p-value of 0.12, providing no evidence against the common-shocks assumption. Full diagnostics are in Appendix C.3, C.1, and C.2.

6.1 Average TFP Effects

Table 5 reports two-step FGLS estimates of the static specification:

$$\mathrm{TFP}_{it} = \beta_F \,\mathrm{FOtech}_{it-1} + \gamma_F \,\mathrm{FOcomp}_{it-1} + \beta_D \,\mathrm{DOtech}_{it-1} + \gamma_D \,\mathrm{DOcomp}_{it-1} + \alpha_i + \lambda_{c,s,t} + u_{it}, \tag{10}$$

where TFP_{it} is log TFP estimated by the Wooldridge (2009) procedure. Given that our sample pools manufacturing firms across ten countries with potentially different firm-level volatility, we estimate equation (10) by feasible GLS: the first step recovers firm-level residual variances from OLS; the second step re-estimates using inverse-variance weights $\omega_i = 1/\hat{\sigma}_i^2$, giving less weight to firms with noisier productivity residuals.[6] Standard errors are clustered at the firm level.

Columns (1)–(2) use the full sample; columns (3)–(4) restrict to firms with at least ten em-

[6]Concretely, $\hat{\sigma}_i^2 = T_i^{-1} \sum_t \hat{\varepsilon}_{it}^2$. To prevent extreme weights for a small number of long-lived firms, we winsorize the lower tail at the 5th percentile.

Table 5: Foreign Technology Exposure and TFP: Average FGLS Estimates

	Full sample		EMPL≥10	
	(1)	(2)	(3)	(4)
FOtech_{it-1} (std)	0.0021^{*}	0.0018^{*}	0.0029^{**}	0.0025^{**}
	(0.0011)	(0.0011)	(0.0012)	(0.0012)
DOtech_{it-1} (std)	0.0003	0.0003	0.0002	0.0001
	(0.0005)	(0.0005)	(0.0006)	(0.0006)
FOcomp_{it-1} (std)	—	0.0001	—	0.0002
		(0.0002)		(0.0002)
DOcomp_{it-1} (std)	—	-0.0001	—	-0.0001
		(0.0001)		(0.0002)
Observations	94,546	94,546	74,597	74,597
Firms	13,281	13,281	10,260	10,260
R^2	0.995	0.995	0.996	0.996

Notes: Two-step FGLS estimates of equation (10). Dependent variable: log TFP (Wooldridge-LP estimator). All coefficients are standardized: effect of a one-standard-deviation increase in each regressor. Fixed effects: firm; country × four-digit sector × year. Standard errors clustered by firm. Columns (1) and (3) include only FOtech and DOtech; columns (2) and (4) additionally include FOcomp and DOcomp. First-step OLS residual variances used to construct inverse-variance weights; lower tail winsorized at the 5th percentile. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

ployees. Within each sample, columns (1) and (3) include only FOtech and DOtech; columns (2) and (4) add FOcomp and DOcomp to control for product-market competition. All coefficients are standardized.

The coefficient on FOtech_{it-1} is positive and significant across all four columns: 0.0021^{*} and 0.0018^{*} in the full sample, 0.0029^{**} and 0.0025^{**} among larger firms. A one-standard-deviation increase in foreign technology exposure is associated with a contemporaneous TFP level approximately 0.2 percent higher in the full sample and 0.25–0.3 percent higher among firms with at least ten employees. These per-year magnitudes are modest, and the static specification captures only the average contemporaneous relationship: it should be read as consistent with the innovation channel, not as a precise estimate of the long-run productivity effect. The DOtech_{it-1} coefficient is positive but statistically insignificant throughout. The competition controls FOcomp_{it-1} and DOcomp_{it-1} are indistinguishable from zero, indicating that product-market exposure to the same patent shocks does not generate measurable contemporaneous TFP effects.

6.2 Dynamic TFP Response: Affiliate Technology-Intensity Split

The static FGLS estimates pool all domestic firms and capture only the average contemporaneous effect. The more informative test of the mechanism is dynamic: does the TFP response follow the same pattern as the patent response — gradual, delayed, and concentrated among firms exposed to technology-intensive affiliates?

Specification. We apply the LP-DiD framework of equation (8) with log TFP as the outcome and the same affiliate technology-intensity split as Section 5.1. One modification relative to the patent LP-DiD is the addition of a pre-treatment TFP growth control: we include $L_1.\text{TFP}_{it} - L_4.\text{TFP}_{it}$ as a regressor to absorb differential pre-existing trends in productivity growth across treated and control firms. All other design features — event window $[-3, 5]$, never-treated control group, country $\times$ sector $\times$ year fixed effects, firm-clustered standard errors — are identical to Section 4.3. Full specification details are in Appendix D.5.

Results. Figure 10 displays the event-study estimates for firms exposed to technology-intensive affiliates (left) and non-CT-exposed firms (right). For CT-exposed firms, the pre-treatment coefficients at $h = -3, -2, -1$ are flat and statistically indistinguishable from zero, confirming parallel productivity trajectories before the treatment event. After treatment, TFP rises gradually, becoming positive and economically meaningful by $h = 2$–3 and remaining elevated through $h = 5$. The timing is broadly consistent with the patent response; TFP gains appear slightly earlier, consistent with the interpretation that firms begin applying recombined knowledge to production processes before formally filing patents on the resulting innovations. For non-CT-exposed firms, TFP dynamics are flat throughout the entire event window at all pre- and post-treatment horizons.

Three results from the two panels together are worth emphasising. First, the innovation response propagates to productivity: positive TFP dynamics among CT-exposed firms confirm that the patent buildup reflects genuine improvements in production capability, not pure patent-filing activity. Second, both outcomes — patents and TFP — are driven by the same firms, ruling out explanations in which productivity improves through channels unrelated to the knowledge mechanism, such as input sourcing or competition-induced rationalisation. Third, the contrast between the two panels provides a direct empirical counterpart to the prior horizontal FDI spillover literature: the non-CT-exposed group — firms facing foreign presence without a locally innovating affiliate — shows exactly the flat TFP dynamics that Aitken and Harrison (1999) and related studies document as the average effect across all domestic firms.

Technological proximity, mediated by affiliate innovation intensity, is what converts foreign presence from a competitive threat into a productivity dividend. More broadly, these results suggest that the organisational structure of multinational networks — not just their size or sectoral composition — determines how much of the global innovation frontier becomes usable by firms in host countries.

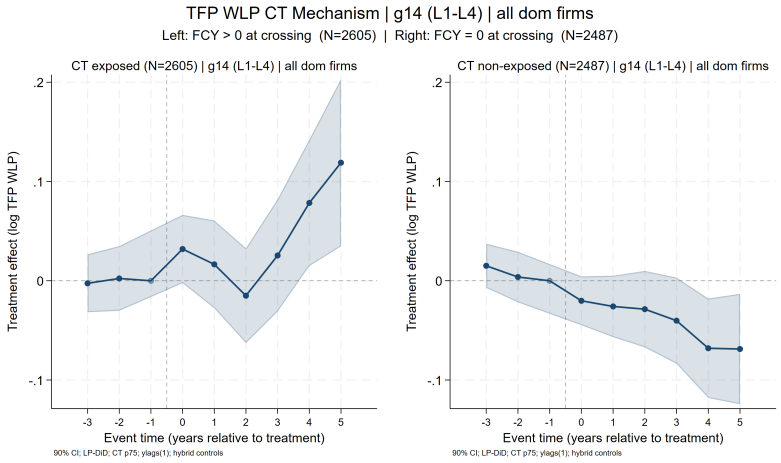


Figure 10: Technology-Intensive Affiliates Drive the TFP Response

Notes: LP-DiD event studies of the TFP response to a large foreign technology shock, estimated separately for CT-exposed (left) and non-CT-exposed (right) domestic firms. Outcome: log TFP (Wooldridge-LP estimator). Treatment and technology-intensity split definitions as in Figure 9. Additional control: pre-treatment TFP growth $L_1.\text{TFP} - L_4.\text{TFP}$. Never-treated firms as controls; country × four-digit sector × year fixed effects; standard errors clustered by firm; 90 percent confidence intervals. Full specification details and robustness to De Loecker–Warzynski TFP in Appendix D.5.

7 Conclusion

Publicly disclosed patents are nominally accessible everywhere, yet frontier knowledge does not diffuse uniformly across the firms and countries positioned to benefit from it. This paper shows that access to the global innovation frontier is mediated by two organisational conditions: the local presence of a technology-intensive foreign affiliate that makes the frontier patent portfolio actionable, and the prior innovative capacity of the domestic firm that allows it to recognise and recombine the relevant elements. When both conditions are met, domestic firms patent substantially more and become more productive. When either is absent, they do not.

The first condition is affiliate technology intensity. We show that the innovation response to foreign technology exposure is concentrated among domestic firms whose country-sector

environment contains a locally innovating affiliate — one that itself patents intensively. Firms facing the same large increase in GUO frontier knowledge but without such an affiliate show no measurable response. Foreign ownership alone is not the channel; technology-active local presence is. The second condition is absorptive capacity. Even among firms exposed to technology-intensive affiliates, only those with above-median initial patent activity benefit: they exhibit around 13 percent higher citation-weighted patenting four years after a one-standard-deviation increase in foreign technology exposure, and approximately 50 percent higher after five years. Low-capacity firms show no positive response and may experience adverse competitive effects. Technological proximity creates the exposure opportunity; absorptive capacity determines whether that opportunity translates into new ideas.

Taking the results together: a frontier GUO expands its patent portfolio; a technology-intensive affiliate embeds that knowledge locally, creating opportunities for recombination; technologically proximate domestic firms with sufficient innovative capacity build on it and eventually patent around it; and the resulting innovation propagates to measured total factor productivity. The LP-DiD TFP results — which show positive dynamics for firms exposed to technology-intensive affiliates and flat dynamics for the rest — confirm that each link in this chain operates. These findings speak directly to the long-running debate on horizontal FDI spillovers. The canonical result of Aitken and Harrison (1999) — that within-sector foreign presence does not raise, and may lower, domestic TFP — emerges naturally as a special case: for firms not technologically close to the local affiliate, or lacking the absorptive capacity to exploit proximity, the competitive channel dominates. Conditioning on technological proximity and affiliate innovation intensity recovers the positive channel that aggregate sector-level analyses miss.

The policy implications follow directly. FDI promotion and screening policies have traditionally focused on the volume of inward investment. Our results suggest that the composition of FDI matters more than the quantity. Attracting a technology-active multinational affiliate in a sector where domestic firms have overlapping patent portfolios generates measurable innovation and productivity gains; attracting a financial holding company with an innovative parent but no local R&D activity does not. Targeted attraction of frontier technology carriers, and complementary policies that strengthen the absorptive capacity of domestic firms through investment in R&D and skilled engineering talent, are what move the needle for host-country productivity growth.

Our results carry broader implications for theories of knowledge diffusion. The effective boundaries of an innovation system are determined not by geography alone but by the structure of corporate ownership networks: a domestic firm in Europe can access the patent fron-

tier of a Japanese multinational because of the presence of a technology-active subsidiary in its own local sector, not because of physical proximity to Tokyo. This decoupling of knowledge access from geographic distance implies that countries and firms can integrate into global innovation networks without being located in the hubs where frontier R&D is conducted — but only when the organisational and technological conditions for recombination are in place. Several questions remain open. Our sample ends in 2012, before the most recent wave of digital-technology diffusion and the rise of new innovation leaders in Asia. Whether the mechanism operates with similar force in post-GFC global value chains, or in service sectors where patents are less central to innovation, remains to be examined. Labour mobility — engineers moving between foreign affiliates and domestic firms — is a natural complementary channel that our patent-based design cannot directly identify. We leave these extensions to future research.

References

Acemoglu, D., U. Akcigit, and W. Kerr (2016): "Innovation Network," *Proceedings of the National Academy of Sciences*, 113.

Ackerberg, D. A., K. Caves, and G. Frazer (2015): "Identification Properties of Recent Production Function Estimators," *Econometrica*, 83, 2411–2451.

Adão, R., M. Kolesár, and E. Morales (2019): "Shift-Share Designs: Theory and Inference," *Quarterly Journal of Economics*, 134, 1949–2010.

Aitken, B. J. and A. E. Harrison (1999): "Do Domestic Firms Benefit from Direct Foreign Investment?," *American Economic Review*, 89, 605–618.

Akcigit, U., D. Hanley, and N. Serrano-Velarde (2021): "Back to Basics: Basic Research Spillovers, Innovation Policy, and Growth," *The Review of Economic Studies*, 88, 1–43.

Bilir, L. K. and E. Morales (2020): "Innovation in the Global Firm," *Journal of Political Economy*, 128, 1566–1625.

Bloom, N., M. Schankerman, and J. Van Reenen (2013): "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 81, 1347–1393.

Borusyak, K., P. Hull, and X. Jaravel (2022): "Quasi-Experimental Shift-Share Research Designs," *The Review of Economic Studies*, 89, 181–213.

Caballero, R. J. and A. B. Jaffe (1993): "How High Are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth," in *NBER Macroeconomics Annual 1993*, MIT Press, vol. 8, 15–86.

Callaway, B. and P. H. C. Sant'Anna (2021): "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 225, 200–230.

Chen, J. and J. Roth (2024): "Logs with Zeros? Some Problems and Solutions," *Quarterly Journal of Economics*, 139, 891–936.

Cohen, W. M. and D. A. Levinthal (1989): "Innovation and Learning: The Two Faces of R&D," *The Economic Journal*, 99, 569–596.

——— (1990): "Absorptive Capacity: A New Perspective on Learning and Innovation," *Administrative Science Quarterly*, 35, 128–152.

Dube, A., D. Girardi, Ò. Jordà, and A. M. Taylor (2025): "A Local Projections Approach to Difference-in-Differences," *Journal of Applied Econometrics*, 40, 741–758.

Furman, J. L. and S. Stern (2011): "Climbing Atop the Shoulders of Giants: The Impact of Institutions on Cumulative Research," *American Economic Review*, 101, 1933–1963.

Giroud, X., S. Lenzu, Q. Maingi, and H. Mueller (2024): "Propagation and Amplification of Local Productivity Spillovers," *Econometrica*, 92, 1589–1619.

Giroud, X., E. Liu, and H. M. Mueller (2026): "Innovation Spillovers across U.S. Tech Clusters," *Journal of Financial Economics*, 179, 104264.

Goldsmith-Pinkham, P., I. Sorkin, and H. Swift (2020): "Bartik Instruments: What, When, Why, and How," *American Economic Review*, 110, 2586–2624.

Hall, B. H., A. Jaffe, and M. Trajtenberg (2005): "Market Value and Patent Citations," *RAND Journal of Economics*, 36, 16–38.

Jaffe, A. B. (1986): "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits, and Market Value," *American Economic Review*, 76, 984–1001.

Jaffe, A. B., M. Trajtenberg, and R. Henderson (1993): "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quarterly Journal of Economics*, 108, 577–598.

Jorda, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95, 161–182.

Keller, W. and S. R. Yeaple (2013): "The Gravity of Knowledge," *American Economic Review*, 103, 1414–1444.

Levinsohn, J. and A. Petrin (2003): "Estimating Production Functions Using Inputs to Control for Unobservables," *The Review of Economic Studies*, 70, 317–341.

Olley, G. S. and A. Pakes (1996): "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64, 1263–1297.

Parente, S. and E. C. Prescott (1994): "Barriers to Technology Adoption and Development," *Journal of Political Economy*, 102, 298–321.

Petrin, A., T. K. White, and J. P. Reiter (2011): "The Impact of Plant-level Resource Reallocations and Technical Progress on U.S. Macroeconomic Growth," *Review of Economic Dynamics*, 14, 3–26.

Santos Silva, J. M. C. and S. Tenreyro (2006): "The Log of Gravity," *The Review of Economics and Statistics*, 88, 641–658.

Wooldridge, J. M. (2009): "On Estimating Firm-Level Production Functions Using Proxy Variables to Control for Unobservables," *Economics Letters*, 104, 112–114.

Yeaple, S. R. (2013): "The Multinational Firm," *Annual Review of Economics*, 5, 193–217.

International Innovation Spillovers through Global Technology Ownership

Fons-Rosen, Kalemli-Özcan, Sørensen, Villegas-Sanchez, Volosovych

Online Appendix

Contents

1 Introduction **1**

2 Setting, Data, and Measurement **5**

2.1 Research Setting 5

2.2 The Firm Panel: ORBIS 6

2.3 Ownership Structure and the Affiliate Network 8

2.4 Patent Data: PATSTAT 9

2.5 Technological Proximity 11

2.6 The Foreign Technology Exposure Measure 14

3 Empirical Strategy **16**

3.1 Specification 16

3.2 Identification 17

4 Foreign Technology Exposure and Domestic Innovation **19**

4.1 Baseline Local Projection Results 19

4.2 Pre-Trend Tests and Robustness 21

4.3 LP-DiD: Discrete Treatment Design 23

5 Channels **26**

5.1 Technology-Intensive Affiliates as the Conduit 26

5.2 Absorptive Capacity 28

6 From Innovation to Productivity **30**

6.1 Average TFP Effects 31

6.2 Dynamic TFP Response: Affiliate Technology-Intensity Split 33

7 Conclusion **34**

A Data Construction **41**

A.1 ORBIS: Raw Panel, Identifier Cleaning, and Sample Selection 41

A.2 ID Reconciliation: Patent–ORBIS Bridge and GUO Panel 42

A.3 Ownership Network: GUO Assignment, Domestic Firm Definition, and Summary Statistics 43

A.4 PATSTAT: Technology Classification, Granted Patents, and Citation Network . 44

A.5 Dataset Map: Files Produced by the Data Pipeline 45

B Measurement Details **46**

B.1 IPC Technology Class Structure 46

B.2 Sectoral Proximity: Construction and Formula 46

B.3 TFP Estimation: Wooldridge (2009) Procedure and Output Elasticities 47

C Shift-Share Validity: Full Tables **48**

C.1 AKM Inference: GUO-Clustered Standard Errors 48

C.2 BHJ Overidentification Test 49

C.3 Rotemberg (2019) Decomposition 50

C.4 Within-Sector Variation in Technology Proximity: Four-Digit Breakdown . . . 50

D Additional Results **51**

D.1 Patent LP: Full Table (Non-Standardized Coefficients) 51

D.2 Dynamic Controls Variant 51

D.3 Absorptive Capacity: Full Tables 52

D.4 LP-DiD: Treatment Timing and Pre-Treatment Balance 52

D.5 TFP LP-DiD: Specification Details and Robustness 53

D.6 Firm-to-Firm Citations: Domestic Firms Citing the GUO 54

A Data Construction

A.1 ORBIS: Raw Panel, Identifier Cleaning, and Sample Selection

Source and raw data. The raw data file is `all_data_final_tfp_2015oct16.dta` (Bureau van Dijk ORBIS, October 2015 vintage). The file covers European manufacturing firms with at least basic financial data available. Before any sample restrictions the ten-country sample contains 2,821,948 firm-year observations for 471,581 unique firms.

BvD identifier reconciliation. BvD identifiers change over time when firms are re-coded, creating spurious discontinuities in firm histories. We reconcile identifiers using BvD's own change-tracking tool (August 2023 vintage). The reconciliation proceeds in four steps. (i) We remove auxiliary identifiers: observations ending in the letter 'C' (277 observations dropped) and strip trailing 'U' characters to obtain a raw identifier `oldid`. (ii) We merge the ID-change crosswalk; of approximately 1.26 million matched observations, all changed IDs are replaced with their new values (`newid`), which becomes `ID_CLEAN`. (iii) We resolve duplicate mappings: 10,282 observations (862 firms) are dropped where the mapping is not one-to-one. (iv) The resulting crosswalk `Step0_0_Link_IDNUMBER_IDCLEAN.dta` covers 721,023 firms across all countries (548,275 in the ten-country sample) and satisfies `isid` in both directions. This crosswalk is used in all downstream merges to ensure identifier consistency across the ORBIS financial panel, the ownership network, and the patent datasets.

Sector codes. We drop observations with missing or generic NACE Rev. 2 codes (category 'C' only, 181 observations). We flag firms that change four-digit sector during the sample (32,699 firms in the ten-country sample) and drop those that change three or more times; firms changing two-digit sector (19,666 firms) are flagged but not dropped. Each firm's primary sector is defined as the four-digit NACE code appearing most frequently across its panel observations.

Financial variables and winsorization. The key financial variables are operating turnover (OPER_TURN), total labour costs (COSTEMPL), intermediate materials (MATERIAL), and intangible fixed assets (TGINTGFIXEDASSETS). Each is deflated to real terms: turnover, labour costs, and materials using two-digit NACE producer price indices (suffix `_PPI_2dgt`); capital using a country-level investment deflator. All four PPI-deflated variables are winsorized at the 1st–99th percentile *within each country* to reduce the influence of extreme out-

liers arising from data entry errors or firm-level reporting inconsistencies.

TFP sample. The TFP estimation sample (indicator `tfp_sample` = 1) requires all four deflated inputs — COSTEMPL_PPI, OPER_TURN_PPI, MATERIAL_PPI, and TGINTGFIXEDASSET — to be non-missing and strictly positive. This yields 463,652 firms (2,748,137 observations) in the ten-country sample.

Foreign ownership indicators. Using the GUO and direct-owner links, we construct an indicator `dumfirm_fo_imuo` that equals one if the firm is ever foreign-owned through either an immediate direct owner or a GUO during the sample period. This identifies 8,041 firms ever foreign-owned through GUO links and 16,871 through combined immediate-plus-GUO links. Domestic firms are those for which this indicator equals zero throughout the sample, as defined in Appendix A.3.

A.2 ID Reconciliation: Patent–ORBIS Bridge and GUO Panel

Patent–ORBIS bridge. The raw PATSTAT–ORBIS bridge (ALL_GRANTED_ORBIS_LINK_unique. May 2019 vintage) links approximately 500,000 BvD identifiers to patent application IDs globally. Because BvD identifiers evolved between the 2019 patent-bridge vintage and the 2023 ownership-update vintage, we update all identifiers in the bridge to the August 2023 BvD vintage before any merging. The update proceeds as follows: (i) we extract the unique BvD identifiers from the bridge and submit them to BvD's change-tracking tool; (ii) of approximately 30,000 returned records, roughly 10,000 have the status 'changed' and receive new identifiers (`newid`); (iii) old identifiers are replaced with `newid` where a match exists; (iv) the updated file (ALL_GRANTED_ORBIS_LINK_unique_ids0823.dta) satisfies `isid ID_CLEAN appln_id`, confirming that the bridge is one-to-one in identifiers and patent applications.

GUO panel identifier update. The GUO panel is sourced from `all_data_final_tfp_2015oct1` We apply the same August 2023 BvD change-tracking update to GUO identifiers (UOID_CLEAN), generating a cleaned GUO panel (`Step0_1_Link_BVDID_UO_panel.dta`). Nine observations are dropped where the firm's own identifier equals its GUO identifier — these are cases where the firm is incorrectly recorded as its own controlling owner. The cleaned GUO panel is then merged with the main ORBIS panel via ID_CLEAN, and the initial GUO assignment (UOID_CLEAN_init) is generated as described in Appendix A.3.

A.3 Ownership Network: GUO Assignment, Domestic Firm Definition, and Summary Statistics

Moody's Independence Indicator. ORBIS assigns each firm an Independence Indicator that summarises its degree of independence from external shareholders. The relevant categories are: **A** (independent: no shareholder holds more than 25 percent directly or in total); **B** (one or more shareholders hold between 25 and 50 percent); **C** (one shareholder holds more than 50 percent in total); **D** (one shareholder holds more than 50 percent directly). Confidence levels (A+, A, A$^-$; B+, B, B$^-$; C+) reflect the number of known shareholders. We define a GUO as the highest entity in the ownership chain that is itself classified as A, B, or independent by nature (state or individual). Entities with indicators C or D have a controlling owner above them and are therefore not GUOs.

GUO assignment algorithm. Each target firm may be linked to more than one GUO across the sample period due to ownership restructuring (acquisitions, mergers, demergers). We assign a single time-invariant initial GUO (UOID_CLEAN_`init`) per target using the following algorithm: (i) identify all GUO links observed for the firm across all years; (ii) keep the GUO first observed in the panel (lowest `minyear`); (iii) if multiple GUOs are first observed in the same year, keep the one appearing most frequently across all panel years; (iv) if frequency is tied, keep the alphabetically first GUO identifier. The resulting mapping satisfies `isid ID_CLEAN`, confirming a unique GUO per target.

Domestic firm definition. A firm is classified as a *domestic firm* if and only if it satisfies all three of the following conditions throughout the entire 1999–2012 sample period:

1. It is not directly foreign-owned (DUMFO_`firm` = 0);
2. It has no foreign GUO (UOFO_`firm` = 0);
3. Either it has no GUO at all (UODO_`firm` = 0), or its GUO is incorporated in the same country (UODO`Local_firm` = 1).

Condition (iii) ensures that firms belonging to domestic corporate groups — where both the firm and its GUO are in the same country — are included, since these groups do not expose their members to foreign technology. Firms that switch from domestic to foreign ownership at any point are excluded from the domestic firm sample entirely.

GUO type and country statistics. Table A.1 reports the breakdown of GUO entity types and top countries of origin among the 25,139 unique GUOs in the ten-country sample. Trade

Table A.1: GUO Entity Types and Top Countries of Origin

GUO type	Targets	Share
C – Trade & industry company	42,002	85.6%
F – Financial company	3,708	7.6%
B – Bank	1,464	3.0%
E – Pension / mutual fund	687	1.4%
S – Public authority / state	595	1.2%
Other	615	1.3%
Total	49,071	100%

Top GUO countries	Share
France	27.6%
Sweden	17.2%
Spain	11.9%
Italy	11.3%
Belgium	5.7%
Germany	5.1%
Norway	4.6%
United States	3.4%

Notes: Left table: GUO entity types among 49,071 target firms with at least one GUO link in the ten-country sample, using the initial GUO assignment. Right table: share of unique GUOs by country of incorporation (top eight countries). Results are robust to excluding financial GUOs (types B, E, F, P) from the analysis.

and industry companies (type C) account for 85.6 percent of GUOs by count; financial companies (F) and banks (B) together account for a further 10.6 percent. Among the top GUO countries by unique GUO count, France (27.6 percent), Sweden (17.2 percent), and Spain (11.9 percent) dominate because our sample includes large manufacturing sectors in these countries with high rates of domestic group formation.

A.4 PATSTAT: Technology Classification, Granted Patents, and Citation Network

Technology classification. Each patent application in PATSTAT is assigned one or more IPC-based technical field codes from two source files (`TECHNICAL_FIELD_PART1.txt` and `TECHNICAL_FIELD_PART2.txt`), together covering approximately 80 million observations. We collapse these files to one record per application (`isid appln_id`), retaining the minimum and maximum technical field codes where multiple codes are assigned. We then aggregate to the patent family level (`docdb_family_id`) by assigning the modal technical field across all applications in the family, yielding a single technology class per family.

Granted patent flag. PATSTAT records whether a patent was granted via the publication kind code (`publn_kind`), which varies by patent authority. We construct an indicator `d_publn_kind_B_FamID` equal to one if at least one application in the family received a granted publication code, using the following authority-specific mappings: CN (B, C, U), DE (C), EP/FR/GB/KR (B), JP (B, Y), US (B, C, E, H, R), WO (A), and standard B codes for most

other authorities with additional codes for FI (B, C, L), HU (B, T, V), PL (B, T), and SE (B, C, D, E, T). Approximately 80 percent of patent families in our matched sample receive the granted flag.

Patenting targets and GUOs. Of the 471,581 manufacturing firms in our sample, 28,178 (6.0 percent) are matched to at least one PATSTAT patent family. Their median portfolio size is 2 families (mean 9.3; p99 = 120; max = 4,477), with filing years spanning 1900 to 2012 and a median first filing year of 1998. Among GUOs, 4,272 are matched to at least one patent; of these, 3,917 are linked to at least one target firm in the ten-country sample. GUO portfolio sizes are substantially larger: the median GUO holds 12 families (mean 2,798; p90 = 1,960; p99 $>$ 108,000; max $>$ 10 million). Among patenting GUOs, 2,715 have their country of incorporation inside the ten-country sample (domestic GUOs) and 1,202 outside (foreign GUOs).

Citation network. We construct the citation network from the PATSTAT TLS212 table of citing–cited patent pairs. Starting from approximately 273 million raw citing–cited pairs, we impose the following restrictions: (i) both the citing and the cited patent must be matched to a BvD identifier (reducing to $\sim$107 million pairs); (ii) the citing application year must fall within 1999–2017 (reducing to $\sim$92 million); (iii) we deduplicate at the patent-family-pair level (reducing to $\sim$52 million). The resulting file (`Step0_3_Link_citing_cited.dta`) is used to construct citation counts for patent quality weighting and for the firm-to-firm citation analysis in Appendix D.6. We also construct an indicator for applicant-filed citations (`d_citn_origin_appl_FamDiad` = 1 if the citation origin code is 'APP'), which identifies citations added by the patent applicant rather than by the examiner.

A.5 Dataset Map: Files Produced by the Data Pipeline

Table A.2 lists the main output files produced by pipeline steps 0.0–0.3 and the content of each. The final analytic datasets for the LP regressions are assembled by linking the ORBIS financial panel (`Step0_0_main_dataset`) with the initial GUO assignment (`Step0_1_Link_BVDID_UO_u`) and the patent flow panel (`Step0_3_Flow_NumPat`). The target–GUO–patent file (`Step0_3_main_`) is the input for computing TECH_{ij}: a `joinby` on `UOID_CLEAN_init` expands each of the 13,326 eligible targets to all patent families of its GUO. The resulting file contains 37.3 million observations, with a median of 27 GUO patent families per target (mean 2,798; max 170,590).

B Measurement Details

B.1 IPC Technology Class Structure

The International Patent Classification (IPC) organises the entire space of patented inventions into a hierarchical taxonomy. At the highest level, the IPC divides technology into eight main **sections**: A (Human necessities), B (Performing operations; transporting), C (Chemistry; metallurgy), D (Textiles; paper), E (Fixed constructions), F (Mechanical engineering; lighting; heating; weapons; blasting), G (Physics), and H (Electricity). Each section is subdivided into classes, subclasses, groups, and subgroups, yielding tens of thousands of terminal codes. For computing TECH_{ij}, we aggregate to the 35 technical fields defined by the OECD concordance between IPC codes and economic technology categories (**?**). This concordance is specifically designed for economic analysis and groups IPC codes into fields of broadly comparable technological scope. The 35 fields cover all eight IPC sections and include, among others, Food chemistry, Pharmaceuticals, Biotechnology, Semiconductors, IT methods for management, Telecommunications, Audio-visual technology, Engines and turbines, and Metallic materials.

Figure B.1 illustrates the hierarchical structure of the IPC using the Espacenet taxonomy. The use of 35 aggregated fields rather than the full IPC hierarchy is standard in the innovation-spillovers literature (Bloom et al., 2013; **?**) and ensures that the technology proximity measure TECH_{ij} captures meaningful technological relatedness rather than superficial code similarity.

B.2 Sectoral Proximity: Construction and Formula

The sectoral proximity measure SEC_{ij} parallels TECH_{ij} but uses firms' distributions across NACE activities rather than patent classes.

Sector-share vectors. For each firm or GUO i, we construct a sector-share vector $S_i = (S_{i1}, \ldots, S_{iK})$, where S_{ik} is the share of firm i's activity assigned to NACE sector k. Because ORBIS does not report firm-specific revenue by sector, we follow **?** and use the reported NACE codes directly. A firm's primary (four-digit) sector receives a weight of two-thirds; if secondary NACE codes are also reported, these share the remaining one-third equally. For example, if a firm reports one primary and two secondary sectors, the primary receives weight 0.667 and each secondary receives weight 0.167. Firms reporting only a primary sector have S_i concentrated entirely on that sector. Duplicate firm-sector pairs are dropped to prevent any

sector from contributing more than a weight of one. For GUOs, sector information is drawn from Orbis Online GUO sector-code files, since GUO records in the offline ORBIS vintage may be less complete.

Formula. Sectoral proximity is the uncentered correlation of the sector-share vectors, analogous to equation (1):

$$\mathrm{SEC}_{ij} = \frac{S_i S_j'}{(S_i S_i')^{1/2}(S_j S_j')^{1/2}}. \tag{11}$$

SEC_{ij} ranges from 0 to 1 and equals 1 for firms with identical sector portfolios. It is symmetric ($\mathrm{SEC}_{ij} = \mathrm{SEC}_{ji}$) and captures the degree of product-market overlap between the two firms.

Relationship to TECH_{ij}. While TECH_{ij} and SEC_{ij} both use uncentered correlations of share vectors, they capture substantively different dimensions of firm similarity. Technological proximity measures overlap in the knowledge produced (which IPC technology classes a firm's patents span), while sectoral proximity measures overlap in what is produced (which output markets a firm participates in). The two measures are intentionally kept separate: in our main specifications, FOtech_{it} and FOcomp_{it} are included simultaneously, so that the coefficient on FOtech_{it} captures pure knowledge-diffusion effects net of product-market competition from the same set of foreign GUO patent stocks.

B.3 TFP Estimation: Wooldridge (2009) Procedure and Output Elasticities

Methodology. We estimate firm-level TFP using the GMM procedure of Wooldridge (2009), which builds on the proxy-variable approach of Olley and Pakes (1996) (OP) and Levinsohn and Petrin (2003) (LP) while addressing the identification critique of Ackerberg et al. (2015). The Ackerberg et al. critique points out that if the flexible input (labour) is chosen after the proxy variable (investment or intermediate inputs) is determined, the labour coefficient is not identified in the OP/LP first stage. The Wooldridge (2009) GMM approach resolves this by combining two moment conditions — one identifying the proxy function and one identifying the production function — in a single estimating system, without imposing constant returns to scale.

The log production function is:

$$y_{it} = \beta_0 + \beta_\ell \, \ell_{it} + \beta_k \, k_{it} + \omega_{it} + \varepsilon_{it}, \tag{12}$$

where y_{it} is log real output (operating turnover deflated by the two-digit NACE PPI), ℓ_{it} is log employment (number of employees), k_{it} is log real capital stock (intangible fixed assets deflated by investment price index), ω_{it} is log physical TFP, and ε_{it} is a classical measurement or shock term.

The key assumptions are: (i) productivity follows a first-order Markov process, $\omega_{it} = f(\omega_{it-1}) + a_{it}$, where a_{it} is the productivity innovation; (ii) the proxy variable — here, intermediate materials (m_{it}, deflated) — satisfies $\omega_{it} = g(k_{it}, m_{it})$ for some invertible function g; (iii) a_{it} is uncorrelated with current capital and lagged labour, capital, and materials. We approximate $g(\cdot)$ and $f(\cdot)$ by third-degree polynomials and use lagged levels of inputs as GMM instruments, following the implementation of Petrin et al. (2011).

We estimate the production function separately by two-digit NACE sector and country, dropping sector-country cells with fewer than 50 firms or cells where the estimated labour elasticity is negative. TFP is then recovered as $\hat{\omega}_{it} = y_{it} - \hat{\beta}_\ell \, \ell_{it} - \hat{\beta}_k \, k_{it}$.

Estimated output elasticities. Table B.1 reports summary statistics for the estimated output elasticities across all sector-country cells. The mean labour elasticity is 0.719 and the mean capital elasticity is 0.083, consistent with estimates in the European manufacturing literature. The labour share dominates, reflecting the labour-intensive nature of European manufacturing, particularly among small and medium-sized firms.

C Shift-Share Validity: Full Tables

C.1 AKM Inference: GUO-Clustered Standard Errors

In a shift-share design, the identifying variation is generated by the interaction of pre-determined exposure shares with time-varying shifts. When the shifts are correlated across firms that share the same dominant GUO, standard errors clustered at the firm level may understate inference uncertainty. Adão et al. (2019) derive an alternative variance estimator — the AKM standard error — that accounts for this correlation by clustering at the level of the dominant shift (here, the GUO with the highest $|\alpha_k|$ weight for each firm).

We implement AKM standard errors by re-estimating the baseline PPML local projection at each horizon, clustering by the dominant GUO rather than by firm. Table C.1 reports the ratio of AKM to firm-clustered standard errors for the FOtech_{it-1} coefficient at each post-horizon. The AKM standard errors are approximately $1.17\times$ the firm-clustered standard er-

rors on average across horizons. Firm-clustered standard errors are therefore modestly anti-conservative relative to the AKM benchmark, but the inflation is small and does not overturn the baseline inference: the $h = 4$ and $h = 5$ results remain significant under GUO-level clustering. Two-way standard errors (clustering simultaneously at the firm and GUO levels) yield results nearly identical to the GUO-only cluster, since the two dimensions are nearly collinear — a firm's variation in FOtech_{it} is dominated by a single GUO exposure for most firms in the sample.

C.2 BHJ Overidentification Test

Borusyak et al. (2022) propose a framework for causal inference in shift-share designs that treats the shifts — here, the GUO patent stock changes — as the source of quasi-experimental variation. Under the BHJ framework, the identifying assumption is that the shifts are as-good-as-randomly assigned across GUOs conditional on a set of controls (in our case, country $\times$ sector $\times$ year fixed effects). A key diagnostic is an overidentification test: if each individual GUO-specific shift generates a consistent IV estimate, then instruments constructed from different subsets of GUOs should all point to the same parameter.

We implement the BHJ overidentification test using the top five GUOs by absolute Rotemberg weight as five separate instruments. Each instrument is the GUO-specific contribution to FOtech_{it-1}: $z_{it}^{k} = \mathrm{TECH}_{i,k} \times \mathrm{PatStock}_{k,t}$ for k in the top-5 Rotemberg set. We then estimate the patent LP via IV (using the five instruments in place of the full FOtech_{it-1} measure) and test whether the five instruments are jointly consistent.

Table C.2 reports the key diagnostics at the $h = 4$ horizon (where the main effect is most precisely estimated). The Hansen J overidentification p-value of $\mathbf{0.12}$ provides no evidence against the null hypothesis that all five GUO-specific instruments are valid, supporting the identifying assumption that GUO patent stock growth is exogenous to domestic firm patenting conditional on the fixed effects. The Cragg-Donald F-statistic of 78.7 confirms instrument relevance under classical (i.i.d.) standard errors; the Kleibergen-Paap F-statistic of 7.5 is weaker under clustered standard errors and falls below the standard Stock-Yogo 10% critical value of 26.9. However, as Borusyak et al. note, the BHJ framework is primarily a *consistency* test rather than a strength test, and the relevant diagnostic is the overidentification test rather than the first-stage F-statistic.

C.3 Rotemberg (2019) Decomposition

? provides a decomposition of the shift-share estimator into GUO-specific contributions. For each GUO k, the Rotemberg weight α_k measures the sensitivity of the estimator to that GUO's shift; GUOs with large $|\alpha_k|$ account for most of the identifying variation. The Rotemberg decomposition serves two purposes: it identifies which GUOs are driving the result (a transparency check), and it allows us to verify that the dominant sources of variation carry the sign expected under the technology-diffusion hypothesis.

Table C.3 lists the top five GUOs by $|\alpha_k|$. The signs of α_k are predominantly positive, consistent with a technology-diffusion interpretation: GUOs whose patent stocks grow faster (relative to their cross-sectional average) generate larger positive impulses to domestic firm patenting. Only one of the top five GUOs carries a negative weight.

Beyond the top five, the correlation between $|\alpha_k|$ and the GUO's patent stock level is 0.40, indicating that high-innovation GUOs drive a disproportionate share of the identifying variation. This is a desirable feature of the design: the variation we rely on comes precisely from the frontier innovators most likely to generate genuine knowledge spillovers. The result would be harder to interpret if low-innovation GUOs dominated the variation.

C.4 Within-Sector Variation in Technology Proximity: Four-Digit Breakdown

Figure 6 in Section **??** shows the coefficient of variation (CV = SD/mean) of firm-level average TECH_{ij} within two-digit NACE sectors, confirming meaningful within-sector variation in technological proximity at a broad level of aggregation. Here we verify that this variation is preserved at the finer four-digit level at which the fixed effects in our regression are defined.

Figure C.1 plots the CV across the top 40 four-digit NACE sectors by domestic firm count. The mean CV at the four-digit level is 0.38, nearly identical to the 0.40 found at the two-digit level. The CV exceeds zero in every sector plotted, confirming that within every four-digit industry cell there is meaningful dispersion in how technologically proximate different domestic firms are to the foreign GUOs in their environment. Since the country $\times$ four-digit sector $\times$ year fixed effects in equation (6) absorb all variation common to firms in the same four-digit cell, the identifying variation is precisely this within-cell dispersion in TECH_{ij} — and Figure C.1 confirms it is alive and economically substantial at the finest level of the fixed-effect structure.

The dispersion tends to be highest in knowledge-intensive sectors such as Pharmaceu-

ticals (NACE 21), Computer and electronic products (NACE 26), and Chemicals (NACE 20), where firms can differ substantially in which narrow technological niches they occupy even within the same four-digit category.

D Additional Results

D.1 Patent LP: Full Table (Non-Standardized Coefficients)

Table D.1 reproduces the baseline local projection results of Table 2 using raw (non-standardized) coefficients. Because PPML with an exponential conditional mean estimates semi-elasticities, each coefficient $\hat{\beta}_h$ implies that a one-unit increase in the corresponding regressor multiplies expected patenting by $e^{\hat{\beta}_h}$. The key magnitudes for FOtech_{it-1} are: at $h = 4$, a coefficient of 0.0096 implies $e^{0.0096} - 1 \approx 1.0$ percent higher patenting per unit of FOtech_{it-1}; at $h = 5$, a coefficient of 0.0257 implies $e^{0.0257} - 1 \approx 2.6$ percent. To convert to the standardized units reported in the main text, multiply by the within-estimation-sample standard deviation of FOtech_{it-1}, which is approximately 15.6 at $h = 0$ and varies modestly across horizons as sample composition changes. Note that this exceeds the full-population standard deviation of 8.85 reported in Section 2.6 because the LP estimation sample is restricted to country-sector cells where patenting affiliates and their GUOs both operate, which selects for cells with higher and more variable foreign technology exposure.

D.2 Dynamic Controls Variant

A potential concern with the baseline local projection specification is that firm fixed effects alone may be insufficient to absorb persistent heterogeneity in firm-level patenting trends. In this appendix we augment the baseline specification with a lagged outcome control, which more directly absorbs any firm-specific trends in patenting that are not captured by the fixed effects.

Table D.2 reports the FOtech_{it-1} coefficient from two variants. *Panel B* includes the lagged inverse hyperbolic sine of the patent flow ($\mathrm{asinh}(y_{i,t-1})$) as an additional regressor. *Panel C* instead uses the lagged $\log(\text{citation-weighted patents}+1)$ transformation. Both controls serve the same purpose — absorbing persistence in patenting — but differ in how they handle the mass of zero-patent firm-years.

The key finding is that including a lagged outcome control does not weaken the main result; if anything, it sharpens the estimates at medium horizons. In Panel B, the FOtech_{it-1}

coefficient becomes significant at $h = 3$ (0.165, $p < 0.05$) as well as at $h = 4$ and $h = 5$, in addition to the $h = 4$ and $h = 5$ significance in the baseline. Panel C is virtually identical to Panel B. The magnitude at $h = 5$ (0.382^{**} in both panels) is close to the baseline (0.332^{**}). Since the dynamic control specification requires $t \geq 2$ observations per firm, the sample starts at $h = 1$ (9,479 observations) rather than $h = 0$.

These results confirm that the baseline finding is not an artefact of omitted firm-level patenting trends: controlling directly for lagged patenting activity leaves the technology spillover estimate intact and in some horizons improves its precision.

D.3 Absorptive Capacity: Full Tables

Table D.3 reports the full absorptive capacity results side by side for the Basic and Hybrid-full specifications. The Basic specification includes only FOtech and DOtech (each interacted with High/Low AC); the Hybrid-full specification adds FOcomp and DOcomp as uninteracted controls, which is the specification reported in the main text. The main result — that High-AC firms drive the positive response at $h = 4$ and $h = 5$ while Low-AC firms show no positive response — is stable across both specifications.

Table D.4 verifies robustness to restricting the sample to firms with at least ten employees. The High-AC coefficient remains significant at $h = 4$ and $h = 5$, with magnitudes (0.134^{***} and 0.472^{**}) somewhat larger than in the full-sample baseline, while the Low-AC coefficient remains insignificant or negative throughout.

D.4 LP-DiD: Treatment Timing and Pre-Treatment Balance

Treatment timing. Figure D.1 plots the histogram of first-crossing years — the year in which each treated domestic firm first records $L.\mathrm{FOtech}_{it} > p_{75}$. Treatment events are distributed broadly across the entire 2000–2012 sample window, with no single year accounting for more than 15 percent of treated firms. The distribution has modest peaks around 2002–2003 and 2007–2008, which coincide with periods of accelerated patenting by Japanese and US multinationals (visible in Figure **??**), but there is no dominant cohort. This staggered structure rules out identification from a single macroeconomic event or policy change and supports the generalisability of the LP-DiD estimates across years.

Pre-treatment balance. Figure D.2 plots kernel density estimates of key pre-treatment characteristics separately for treated firms (those that eventually cross p_{75}) and the control

group (never-treated firms and not-yet-treated firms at least five years before their eventual treatment), evaluated over all pre-treatment observations. Five variables are shown: (a) $\operatorname{asinh}(\text{patent flow})$; (b) $\operatorname{asinh}(L.\mathrm{FOtech}_{it})$; (c) log employment; (d) TFP (Wooldridge-LP estimate); and (e) TFP (De Loecker–Warzynski estimate).

The distributions overlap closely in all five panels. Treated and control firms are similar in their pre-treatment patent activity (panel a), their baseline foreign technology exposure level (panel b), their employment size (panel c), and their productivity levels (panels d and e). There is no systematic pattern of treated firms being more innovative, larger, or more productive than controls prior to treatment. This similarity supports the parallel-trends assumption underlying the LP-DiD design and is consistent with the clean pre-horizon coefficients in Figure 8.

D.5 TFP LP-DiD: Specification Details and Robustness

Specification. The TFP LP-DiD specification follows equation (8) with the following modifications. The outcome is log TFP (Wooldridge-LP estimator, `TFP_WLP`). A pre-treatment TFP growth control, $L_1.\mathrm{TFP}_{it} - L_4.\mathrm{TFP}_{it}$, is added as a regressor to absorb differential pre-existing trends in productivity. This is motivated by the observation that TFP is a levels variable that may exhibit persistent trends driven by firm-specific factors unrelated to foreign technology exposure; including the growth control ensures the post-treatment dynamics we estimate are not confounded by pre-existing divergences. All other features of the design are identical to Section 4.3: event window $[-3, 5]$, never-treated control group (firms that never cross p_{75} of $L.\mathrm{FOtech}$), exclusion of first-treatment-year observations and one-off crossers, country $\times$ four-digit sector $\times$ year fixed effects, and standard errors clustered at the firm level. The cell tower split is applied at the crossing year T_i as described in Section 5.1.

Robustness to TFP estimator. Figure 10 in the main text uses the Wooldridge-LP TFP estimator. To check robustness, we repeat the TFP LP-DiD using the De Loecker and Warzynski (2012) TFP estimator (`TFP_DLW`), which uses a different control function approach to address input endogeneity. The qualitative pattern is identical: CT-exposed firms show a gradual post-treatment TFP increase with flat pre-trends, while Non-CT-exposed firms show no response. The timing and magnitude of the CT-exposed response are similar to the Wooldridge-LP baseline, confirming that the TFP result is not sensitive to the production function estimator.

Magnitude. The peak TFP response among CT-exposed firms (at $h \approx 3$–4) corresponds to an increase of approximately 0.5–0.8 percentage points in log TFP relative to the pre-treatment level. Over a five-year event window, this accumulation is consistent with the gradual learning-and-adoption mechanism suggested by the patent LP results: firms that begin innovating more three to five years after the treatment event (Section 4.1) would be expected to translate those innovations into measured productivity improvements with a similar or slightly shorter lag, as process improvements often precede formal patent filings.

D.6 Firm-to-Firm Citations: Domestic Firms Citing the GUO

[Appendix D.6 to be written.]

IPC divisions

The IPC has a systematic and hierarchical structure. Classification becomes more detailed at every further level or subdivision, as you can see in this example:

Level	Symbol	Description
Section	A	Human necessities
Class	A21	Baking; equipment for making or processing doughs; doughs for baking
Subclass	A21C	Machines or equipment for making or processing doughs; handling baked articles made from dough
Group	A21C1	Mixing or kneading machines for the preparation of dough
Subgroup	A21C1/06	with horizontally-mounted mixing or kneading tools; Worm mixers

Figure B.1: IPC Hierarchical Structure

Notes: The figure illustrates the four-level hierarchy of the International Patent Classification: sections (A–H), classes (e.g., A01), subclasses (e.g., A01B), groups, and subgroups. We aggregate to 35 technical fields using the OECD concordance for the construction of TECH_{ij}.

Table B.1: Summary Statistics of Production Function Output Elasticities

	Labour elasticity ($\hat{\beta}_\ell$)	Capital elasticity ($\hat{\beta}_k$)
Mean	0.719	0.083
Median	0.725	0.079
Standard deviation	0.099	0.057
Maximum	0.943	0.573
Minimum	0.133	0.004

Notes: Summary statistics computed across sector-country estimation cells using the Wooldridge (2009) GMM procedure. Cells with negative or missing elasticities are excluded (primarily NACE 12 "Tobacco products" and NACE 19 "Coke and refined petroleum products", which have very few observations). Labour input is measured as the number of employees; capital as intangible fixed assets deflated by the investment price index; output as operating turnover deflated by the two-digit NACE PPI.

Table C.1: AKM vs. Firm-Clustered Standard Errors for FOtech_{it-1}

	h=0	h=1	h=2	h=3	h=4	h=5	Average
Firm-clustered SE	0.0455	0.0635	0.0603	0.0442	0.0695	0.1335	—
AKM (GUO-clustered) SE	0.0518	0.0724	0.0715	0.0521	0.0804	0.1567	—
Ratio (AKM / firm)	1.14	1.14	1.19	1.18	1.16	1.17	1.17

Notes: Standard errors for the FOtech_{it-1} coefficient from the baseline PPML local projection (Table 2), clustered at the firm level (row 1) and at the dominant GUO level following Adão et al. (2019) (row 2). The dominant GUO is the GUO with the largest absolute Rotemberg weight $|\alpha_k|$ for each firm. The ratio (row 3) is the AKM inflation factor; a ratio above 1 implies that firm-clustered SEs are modestly anti-conservative relative to the AKM benchmark. Statistical significance of the $h = 4$ and $h = 5$ results is maintained under GUO-level clustering.

Table C.2: BHJ Overidentification Test Diagnostics (at $h = 4$)

Diagnostic	**Value**
Number of instruments (top-5 GUOs by $\lvert\alpha_k\rvert$)	5
Hansen J overidentification p-value	0.12
Cragg-Donald F-statistic (i.i.d.)	78.7
Kleibergen-Paap F-statistic (clustered)	7.5
Stock-Yogo 10% critical value (5 instruments)	26.9

Notes: BHJ overidentification test following Borusyak et al. (2022). Five instruments are the top-five GUO-specific contributions to FOtech_{it-1} ranked by absolute Rotemberg weight $|\alpha_k|$. Estimated at $h = 4$ (baseline PPML LP). The Hansen J p-value of 0.12 passes the overidentification test, supporting the exogeneity of GUO patent stock growth. The low KP F-statistic under clustering reflects the high correlation across the five instruments (all loading on the same right-tail GUOs) rather than weak identification per se; the Cragg-Donald F-statistic of 78.7 confirms strong first-stage relevance under classical assumptions.

Table D.1: Patent LP: Non-Standardized Coefficients (Full Table)

	Pre-trend test			Post-horizons					
	$h{=}-3$	$h{=}-2$	$h{=}-1$	$h{=}0$	$h{=}1$	$h{=}2$	$h{=}3$	$h{=}4$	$h{=}5$
				(1)	(2)	(3)	(4)	(5)	(6)
FOtech_{it-1}				-0.0003	-0.0021	-0.0018	0.0039	0.0096^{**}	0.0257^{**}
				(0.0030)	(0.0042)	(0.0039)	(0.0029)	(0.0044)	(0.0104)
FOcomp_{it-1}				-0.0053	0.0107	0.0250^{*}	0.0236	0.0390^{***}	-0.0083
				(0.0075)	(0.0156)	(0.0135)	(0.0239)	(0.0128)	(0.0163)
DOtech_{it-1}				0.0070	0.0043	-0.0169^{*}	0.0309	0.1295^{**}	-0.0199
				(0.0100)	(0.0166)	(0.0098)	(0.0342)	(0.0647)	(0.0371)
DOcomp_{it-1}				0.0030	0.0165	0.0440^{**}	0.0036	-0.0157	-0.1594
				(0.0111)	(0.0213)	(0.0221)	(0.0500)	(0.0525)	(0.1857)
Observations				12,999	9,479	7,309	5,636	4,184	3,105
Firms				1,837	1,408	1,134	929	758	607
Pseudo R^2				0.365	0.344	0.370	0.324	0.362	0.351

Notes: PPML local projections. Dependent variable: citation-weighted patent families. Coefficients are non-standardized: a unit increase in the regressor multiplies expected patenting by $e^{\hat{\beta}}$. Pre-horizon columns are omitted here as the magnitudes without standardization are not directly comparable across horizons; the standardized pre-horizon results are in Table 2. Fixed effects: firm; country × four-digit sector × year (forward). Standard errors clustered by firm in parentheses. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$. Multiply FOtech coefficient by within-sample SD ≈ 15.6 to recover the standardized magnitudes in Table 2.

Table D.2: Dynamic Controls Variant: FOtech_{it-1} Coefficient (Standardized)

	Panel B: Lagged IHS patent control						Panel C: Lagged log$(+1)$ patent control					
	$h{=}1$	$h{=}2$	$h{=}3$	$h{=}4$	$h{=}5$		$h{=}1$	$h{=}2$	$h{=}3$	$h{=}4$	$h{=}5$	
FOtech_{it-1}	-0.013	-0.059	-0.021	0.165^{**}	0.166^{*}	0.382^{**}	-0.012	-0.058	-0.021	0.165^{**}	0.167^{*}	0.382^{**}
	(0.073)	(0.075)	(0.068)	(0.080)	(0.095)	(0.174)	(0.073)	(0.074)	(0.068)	(0.080)	(0.095)	(0.175)
Observations	9,479	6,948	5,289	3,902	2,909	2,046	9,479	6,948	5,289	3,902	2,909	2,046
Firms	1,408	1,077	875	715	575	451	1,408	1,077	875	715	575	451

Notes: PPML local projections with lagged outcome control. Panel B adds $\text{asinh}(y_{i,t-1})$ as a regressor; Panel C adds $\log(y_{i,t-1}+1)$. All other specification details identical to Table 2. Coefficients are standardized: effect of a one-standard-deviation increase in FOtech_{it-1}. Sample begins at $h = 1$ as the lagged control requires one additional period. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table D.3: Absorptive Capacity: Basic and Hybrid-Full Specifications

	Basic						Hybrid-full					
	h=0	h=1	h=2	h=3	h=4	h=5	h=0	h=1	h=2	h=3	h=4	h=5
FOtech × High AC	0.001	−0.015	−0.029	0.026	0.087*	0.346*	0.002	−0.007	−0.022	0.029	0.120**	0.405**
	(0.027)	(0.037)	(0.035)	(0.027)	(0.046)	(0.178)	(0.027)	(0.038)	(0.036)	(0.025)	(0.050)	(0.176)
FOtech × Low AC	−0.091	−0.174**	0.201	0.263	−0.045	−0.332	−0.092	−0.195**	0.115	0.189	−0.456*	−0.669
	(0.073)	(0.085)	(0.176)	(0.209)	(0.225)	(0.296)	(0.072)	(0.087)	(0.186)	(0.191)	(0.255)	(0.484)
DOtech × High AC	0.028	−0.015	−0.066	0.163**	0.160	0.130	0.029	−0.010	−0.064	0.166**	0.174	0.154
	(0.034)	(0.079)	(0.046)	(0.082)	(0.337)	(0.183)	(0.035)	(0.079)	(0.047)	(0.082)	(0.339)	(0.183)
DOtech × Low AC	0.016	0.043	−0.049	0.066	0.498**	−0.105	0.014	0.041	−0.053	0.068	0.505**	−0.099
	(0.047)	(0.059)	(0.033)	(0.101)	(0.239)	(0.091)	(0.048)	(0.060)	(0.034)	(0.102)	(0.239)	(0.092)
Obs	12,999 / 9,479 / 7,309 / 5,636 / 4,184 / 3,105						(same)					

Notes: PPML local projections from equation (9). Standardized coefficients throughout. High AC: baseline patent flow above sample median; Low AC: at or below median. Basic specification: only FOtech and DOtech interacted with AC; Hybrid-full: additionally includes FOcomp and DOcomp as uninteracted controls. Fixed effects: firm; country × four-digit sector × year. Standard errors clustered by firm. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

Table D.4: Absorptive Capacity: EMPL$\geq$ 10 Robustness (Hybrid-Full)

	h=0	h=1	h=2	h=3	h=4	h=5
FOtech × High AC	0.003	0.018	−0.015	0.036	0.134***	0.472**
	(0.029)	(0.041)	(0.037)	(0.026)	(0.052)	(0.215)
FOtech × Low AC	−0.112	−0.134	0.169	0.217	−0.400	−0.504
	(0.074)	(0.090)	(0.207)	(0.205)	(0.271)	(0.476)
Observations	11,414	8,471	6,535	5,071	3,847	2,864
Firms	1,621	1,254	1,008	833	691	549

Notes: As Table D.3 (Hybrid-full), restricted to firms with at least ten employees. Only the FOtech × AC rows shown; DOtech × AC, FOcomp, and DOcomp are included but omitted for brevity. Standardized coefficients. $^{*}p < 0.1$; $^{**}p < 0.05$; $^{***}p < 0.01$.

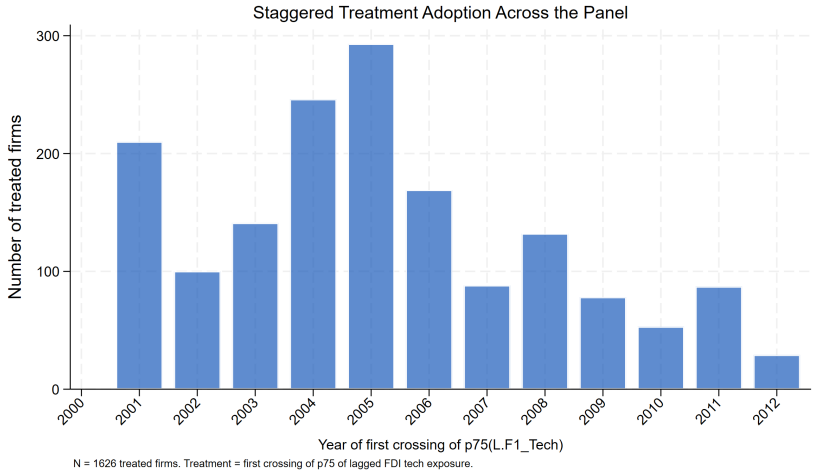


Figure D.1: Distribution of First-Crossing Years (LP-DiD Treatment Timing)

Notes: Histogram of the year in which each treated domestic firm first records $L.\text{FOtech}_{it} > p_{75}(L.\text{FOtech})$. Treatment is staggered across 2000–2012 with no dominant cohort. "One-off crossers" (firms that exceed the threshold in only one year) and firms treated in the first sample year are excluded from the analysis.

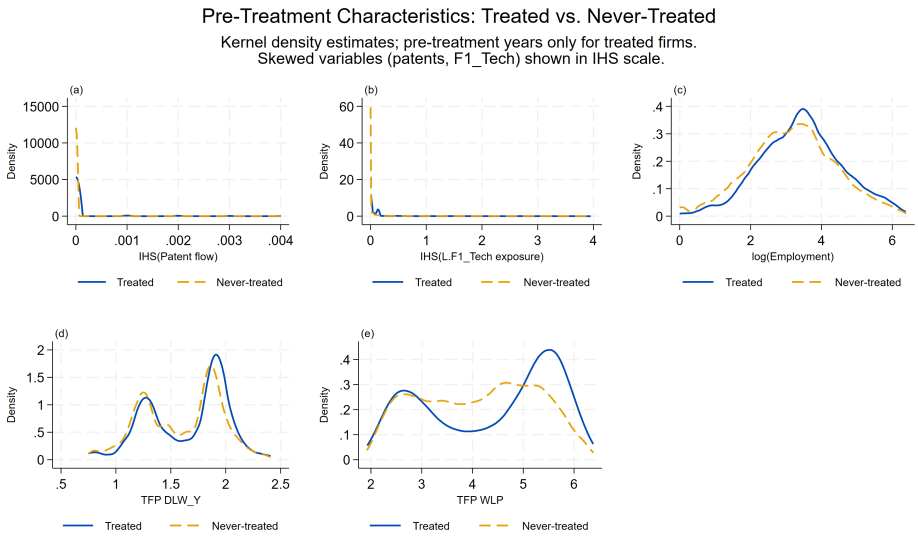


Figure D.2: Pre-Treatment Characteristics: Treated vs. Control Firms

Notes: Kernel density estimates of pre-treatment characteristics for treated firms (solid) and the control group (dashed), evaluated over all pre-treatment firm-year observations. Panels (a)–(b): innovation and exposure in IHS scale; Panel (c): log employment; Panels (d)–(e): TFP under Wooldridge-LP and De Loecker–Warzynski estimators respectively. The close overlap across all panels supports the parallel-trends assumption.