The Global Incidence of Domestic Regulation

Parker Rogers and Fernando Stipanicic*

November 2, 2026

Preliminary

Please do not circulate

Abstract

Domestic regulation in a large market can affect innovation and product markets beyond the regulator's jurisdiction. We study this by linking FDA deregulation events that lower the cost of commercializing new medical technologies to global patents, inventor histories, and bilateral trade flows across 61 non-U.S. countries. Upon deregulation, foreign patenting in treated device types rises 55% relative to controls, inducing one foreign patent for every two it induces domestically. The response is concentrated among new inventors, in countries that import a higher share of treated devices from the U.S., and among teams without U.S.-based members. Patent text shows that induced foreign innovation moves toward and extends the U.S. technological path. U.S.-linked trade also increases by 39% for imports and 40% for exports over the long run. Together, these results are most consistent with a product-market channel, in which the expansion of U.S. products induces innovation abroad. The effects of domestic regulation on innovation and trade can therefore be much larger than domestic outcomes alone suggest. (JEL O31, O33, F14, I18, L51)

*Rogers: Indiana University. Email: paroger@iu.edu. Stipanicic: University of Oslo. Email: fernando.stipanicic@gmail.com. Rogers gratefully acknowledges support from the National Institute on Aging, grant number T32-AG000186. Stipanicic gratefuly acknowledges funding by the European Research Council (project no. 101141494). We thank Jorge Lemus, Kelli Marquardt, Lucy Wang and seminar participants at ASHEcon 2026 and MHEC 2026 for valuable feedback. We have used AI assistance in data analysis and manuscript editing.

1 Introduction

Innovators develop and sell technologies across borders, yet the regulations governing those technologies are set and evaluated nationally. That difference creates a potential wedge between regulatory jurisdiction and economic incidence. When the regulated market is large, changes in domestic regulation may alter the incentives and opportunities foreign inventors face, even when those consequences fall outside the typical scope of regulatory evaluation. To what extent do the effects of domestic regulation extend beyond the regulator's jurisdiction, and how do they propagate through innovation and product markets across countries?

We study these questions in the context of U.S. Food and Drug Administration (FDA) regulation of medical devices. Medical devices provide a useful setting because the U.S. is a large product market, accounting for roughly 40% of the global medical device market (Medicare Payment Advisory Commission, 2017), and a major source of innovation, filing nearly 70% of medical device patents in the categories we study.[1] Firms seeking access to this market, both domestic and foreign, design products to satisfy FDA requirements, while medical technologies diffuse internationally. We analyze 29 events between 1980 and 2015 in which the FDA lowered the regulatory burden of certain device types, a process known as down-regulation, substantially reducing premarket costs.[2] These events allow us to examine how regulation in a large market affects inventors and product markets abroad.

We estimate these effects using a stacked difference-in-differences design that compares affected device types to matched, never-reclassified device types within the same country. To implement this design, we construct a device-type-by-country panel spanning 61 non-U.S. countries. The panel links FDA down-regulation events and product definitions to global patent records, inventor locations and histories, patent text, FDA administrative data, HS-10 U.S. trade flows, and HS-6 bilateral global trade flows. Our main innovation measure counts

[1] Authors' analysis of patent data.

[2] We focus on down-regulation because moves in the opposite direction are rare. Only a handful of device types move from Class II to Class III during our sample period, and these changes are typically prompted by post-market safety concerns. The data therefore do not allow us to test whether increasing regulatory burden produces a symmetric decline in foreign innovation.

patent families (groups of patent filings covering the same invention across jurisdictions), weights them by family size—a standard proxy for patent value (Harhoff et al., 2003)—and assigns them across countries according to inventor residence. The estimates are stable across a variety of specifications and show little evidence of divergent pre-existing trends, consistent with an institutional process characterized by long internal delays and opaque FDA deliberations described as "unpredictable" (Makower et al., 2010; Powell, 2018).

The down-regulation of daily-wear soft contact lenses illustrates the general pattern. In 1994, after a decade-long internal delay, the FDA down-regulated daily-wear soft contact lenses while leaving extended-wear lenses heavily regulated. Patenting in daily-wear lenses subsequently quadrupled among both U.S.-based inventors (368%) and inventors outside the U.S. (349%), while patenting in extended-wear lenses remained flat (Figure 1, Panel a). In levels, the increase among U.S. inventors was roughly 2.6 times the foreign increase, close to the 2:1 ratio we estimate across all events. Product markets reacted as well, with U.S. imports of contact lenses increasing by 152% and exports by 288% (Panel b).

Across the full sample, we find that the foreign incidence of these regulatory reforms is large.[3] For roughly every two additional patents induced among U.S.-based inventors, down-regulation induces one additional patent abroad. Therefore, foreign innovation accounts for one-third of the total innovation response observed, despite falling outside the FDA's jurisdiction. After these events, patenting rises sharply for both U.S.- and foreign-based inventors; in proportional (Poisson) terms the U.S. response is roughly twice the foreign one (about +110% versus +55%), mirroring the two-to-one patent multiplier.[4] The foreign

[3]The U.S. is a lead market for medical devices: it is the largest single market (roughly 40% of the global total) and many national regulators reference or defer to FDA decisions, so U.S. requirements are frequently the binding constraint on a device's design and global launch. This regulatory deference does not, however, drive the foreign innovation response. Interacting treatment with our audited country-level FDA-reliance measure, an ordinal index of how far a country's device regulator formally relies on FDA decisions (Online Appendix B.2), yields a coefficient indistinguishable from zero (−0.066, s.e. 0.066), so the response does not concentrate in countries whose regulators rely on FDA decisions and is not simply foreign adoption of U.S.-cleared devices.

[4]The U.S. and non-U.S. Poisson figures are the average post-event effects from the event study on the region-aggregated matched panel behind Figure 2; the non-U.S. estimate (+59%) matches the country-level +55% reported below. This pattern is broadly consistent with the domestic response documented in Rogers (2023).

response emerges more gradually than the domestic response and continues to build over at least two decades. The magnitude of the foreign response raises a natural question: why does a reduction in U.S. regulatory costs generate innovation abroad?

Several mechanisms are consistent with such a response. Lower U.S. regulatory costs may simply increase the return to participating in the U.S. market, stimulating foreign innovation for U.S. consumers. Alternatively, down-regulation may increase the number of emerging U.S. technologies and products, which can then influence foreign inventors as they diffuse internationally. By contrast, inventors themselves may spread knowledge through international collaborations or relocation. These mechanisms have unique implications for where foreign innovation responds, what foreign inventors work on, and who innovates.

Cross-country heterogeneity helps uncover the mechanisms. Foreign patenting increases most in countries where the U.S. accounted for a larger share of pre-event imports in the treated device type. The estimated response is near zero in the bottom tercile of U.S. import exposure and rises to 0.24 log points in the top. By contrast, prior exports of treated devices to the U.S. are not associated with a differential response. This asymmetry strengthens the role of U.S. product exposure over an explanation based primarily on a higher return to serving a less regulated U.S. market. Import exposure may also capture broader technical sophistication or integration with the U.S. device sector, so this pattern alone does not isolate the central mechanism.

Patent text provides a second piece of evidence. In addition to increasing patenting within affected categories, down-regulation changes what scientific problems foreign inventors address. Foreign patent text becomes more similar to the U.S. technological frontier and moves in the same direction as subsequent U.S. technological change. Foreign innovation moves substantially toward the U.S. frontier and, more strongly, along the U.S. technological path. Yet the induced foreign patents remain distinct from contemporaneous U.S. prior art and move beyond the pre-event U.S. frontier. Therefore, foreign innovation follows the U.S. technological trajectory and extends it.

Down-regulation also changes who participates in innovation. These events add 2.8 new active inventors per country-device-year, a 112% increase relative to the counterfactual. By contrast, experienced inventors pivoting from other device types respond much less, while incumbents already working in treated devices change little. In addition, direct links to U.S. inventors are not necessary for the response, as teams without a U.S.-based inventor exhibit relatively large increases in innovation and pull toward the U.S. technological trajectory. Thus, the foreign innovation response reflects entry of new foreign inventors rather than reallocation of inventors already connected to the U.S. innovation ecosystem.

Product markets respond as well. U.S. imports and exports of treated devices increase by roughly 40% in the long run. Trade expands most along routes involving the U.S.; However, global bilateral data also show increased trade between foreign countries, suggesting that devices produced abroad in response to down-regulation reach markets beyond those connected to the United States. Hence, down-regulation increases, rather than redistributing, global trade flows. Domestic reclassification therefore affects trade well beyond the U.S. market. By lowering the cost of bringing a device to the lead market, it expands global trade in that device, including between foreign countries with no direct U.S. link.

Taken together, these results are most consistent with diffusion through U.S. product markets. Down-regulation increases foreign innovation most where exposure to U.S. products was greatest and reorients it toward the U.S. technological path, while the foreign response does not depend on direct collaboration with U.S.-based inventors. This evidence, by contrast, is less consistent with explanations based primarily on a higher return to serving the U.S. market (export exposure is not associated with the response), learning from U.S.-based inventors through co-invention, generic growth in affected categories (innovation moves toward the U.S. path), or strategic patenting alone (new inventors drive the response). We cannot, however, distinguish among the channels through which product-market diffusion operates. Imported products may convey technical knowledge, demonstrate commercially viable designs, or create compatibility incentives. Nor can we rule out that country-specific

import exposure may partly proxy for broader integration with the U.S. device market, although exposure is measured within device categories.[5] Nevertheless, the combined evidence points to product markets as an important link between domestic regulatory change and inventive activity abroad.

This paper contributes to three literatures. First, it extends work on induced medical innovation. Prior studies show that market size, insurance design, procurement, reimbursement, patent incentives, and regulatory approval shape pharmaceutical and medical innovation (Finkelstein, 2004; Acemoglu and Linn, 2004; Blume-Kohout and Sood, 2013; Clemens, 2013; Budish et al., 2015; Stern, 2017; Sampat and Williams, 2019; Agha et al., 2022; Clemens and Rogers, 2026; Ji and Rogers, 2024; Rogers, 2023). This literature has largely studied how domestic policy affects domestic innovation incentives. We show that the economic incidence of regulatory policy in a lead market extends internationally. A substantial share of induced innovation occurs abroad and follows the technological trajectory of the regulated market.

Second, the paper connects domestic regulation to the literature on trade, technology diffusion, and innovation (Keller, 2004). Existing work shows that access to export markets, import competition, imported inputs, and trade liberalization can generate technology upgrading and innovation (Lileeva and Trefler, 2010; Bustos, 2011; Bloom et al., 2016; Coelli et al., 2022; Ayerst et al., 2023). We study a different source of innovation linked to trade: a regulatory change in a large frontier market. In our setting, foreign innovation responds most in countries more exposed to U.S. treated-device imports. At the same time, product markets adjust as trade expands globally, and most strongly on U.S.-linked routes. The results highlight the far-reaching effects of domestic regulation on international product markets.

Third, the paper provides quasi-experimental evidence for theories of lead markets and global innovation. Such theories argue that large, technologically central markets can select

[5]Import exposure is measured within device categories as the share of a country's pre-period imports of the treated device originating in the United States, rather than as a country-level measure of U.S. trade integration. Export exposure is defined symmetrically, as the share of the country's pre-period exports of the treated device destined for the United States.

designs that later diffuse internationally (Vernon, 1966; Beise, 2004; Beise and Rennings, 2005). Evidence on patent citations shows that knowledge flows are geographically organized and often localized (Jaffe et al., 1993; Belenzon and Schankerman, 2013), while work on international spillovers emphasizes that knowledge flows across borders are an important component of the social return to innovation (Coe and Helpman, 1995; Keller, 2004; Jones and Williams, 1998; Bloom et al., 2013). We provide quasi-experimental evidence that a domestic regulatory shock in a lead market changes the quantity and direction of foreign innovation. Domestic regulation can therefore influence global technological trajectories even when the policy itself is designed and evaluated around domestic objectives.

The paper proceeds as follows. Section 2 provides institutional background. Section 3 develops the conceptual framework. Section 4 describes the data. Section 5 presents the empirical strategy. Section 6 reports results, and Section 7 concludes.

2 Background

The 1976 Medical Device Amendments organized FDA device regulation into three risk tiers. Class I (low risk) typically requires only facility registration and adherence to general manufacturing standards.[6] Class II (moderate risk) typically requires manufacturers to navigate the 510(k) "substantial equivalence" pathway by demonstrating that a new device is substantially equivalent to a legally marketed predicate. Although less intensive than clinical trial requirements, the 510(k) process is costly. Survey evidence estimates total FDA-related costs to be approximately $24 million per device, with average commercialization delays of ten months (Makower et al., 2010). Class III (high risk) typically requires premarket approval (PMA) backed by clinical evidence, costing roughly $75 million on average with years-long delays (Makower et al., 2010).

The FDA may reclassify device types to a lower class as postmarket evidence clarifies risk. This paper focuses on "down-regulation" from Class III to II. The $50 million difference

[6]Roughly 90% of Class I device types are exempt from premarket submissions.

between Class II and Class III review creates a large and discrete shock to the economic costs borne by innovators.

The timing of down-regulation is plausibly exogenous to innovation trends in treated device categories.[7] These events occur as the FDA reviews accumulated postmarket evidence and determines that special controls are sufficient to provide reasonable assurance of safety and effectiveness. The timing of these reviews reflects statutory mandates and accumulated safety and effectiveness evidence, rather than contemporaneous technological change. For example, vascular graft prostheses were down-regulated as part of the FDA Modernization Act (FDAMA), which established a deadline at the FDA for deciding whether to down-regulate pre-1976 Class III devices and relied on decades of safety data rather than recent technological advances. Contact lenses illustrate this pattern: daily-wear soft contact lenses were down-regulated in 1994 after a decade-long internal delay, while similar extended-wear lenses remained in Class III.

Lastly, both foreign and domestic manufacturers face the same FDA requirements to market products in the U.S. Regulatory changes therefore alter the commercialization costs for worldwide firms seeking access to the U.S. market and may also affect the incentives of firms designing products that are compatible with U.S.-origin devices.

3 Conceptual Framework

Down-regulation can affect foreign innovation through at least two broad channels. First, by lowering the cost of commercialization in the U.S., these events increase the expected return to developing products for the U.S. market. If this U.S.-market incentive is important, responses would concentrate in countries that were already more oriented toward exporting

[7] Reclassification is also unrelated to the extent to which treated devices are produced abroad. We measure this using each device's baseline (1989–1995) U.S. import share, defined as imports divided by the sum of imports and exports in the HS-10 trade data. Treated and matched-control devices have similar import shares (0.52 versus 0.57; $p = 0.27$), and the import share does not predict treatment in a linear-probability model (-0.074, s.e. 0.066). Thus, the FDA does not disproportionately down-regulate device categories in which the U.S. relies more heavily on foreign production.

treated device types to the U.S., reflecting more established supply chains or customer bases.[8]

Second, down-regulation may increase foreign exposure to new U.S. products and technical knowledge. By stimulating product entry in the U.S., these events provide foreign inventors more opportunities to observe, adapt, or build on emerging U.S. technologies. If this product-market channel is important, the response would concentrate in countries that imported more treated devices from the U.S. before deregulation, where exposure to new U.S. products was likely greater. This channel may also steer foreign innovation toward the technological path embodied by those emerging U.S. products.

We use this import–export exposure gradient to examine which mechanism is more consistent with the foreign response. Although we hold important observable differences between countries fixed in our analysis, baseline trade exposure is not randomly assigned. Thus, the comparison is not causal, but it helps distinguish between these competing interpretations of the observed heterogeneity. Online Appendix A presents an inventor-level model that formalizes these channels and derives the theoretical predictions examined below.

The framework yields three empirical predictions.

1. *Foreign innovation increases.* Down-regulation can increase foreign innovation by raising the expected returns to developing products in the treated device type and by making relevant technologies easier to identify and build upon. Tested in Section 6.1.

2. *The response is stronger in countries more exposed to U.S. products.* An increase in foreign innovation is consistent with several mechanisms, but cross-country heterogeneity helps distinguish among them. If diffusion through U.S. products is important, the response should be stronger in countries that imported more treated devices from the United States before down-regulation. If instead the response reflects a higher return to serving the U.S. market, it should be stronger in countries that exported more treated devices to

[8]We operationalize this orientation as a country's pre-period *export exposure*: the share of its exports of the treated device destined for the United States. Its counterpart, *import exposure*, is the share of a country's imports of the treated device originating in the United States. Both are measured over a fixed pre-period window and tested in Section 6.1.

the United States. We test these predictions in Section 6.1.

3. *Foreign innovation reorients toward, and extends, the U.S. technological path.* If down-regulation only raises the return to serving the U.S. market, foreign patenting may increase without becoming more aligned with U.S. technologies. Learning from emerging U.S. technologies, however, predicts both more patenting and greater alignment with the U.S. technological path. If foreign inventors build on rather than simply imitate new technologies, the new patents should move beyond the pre-deregulation U.S. frontier rather than converge on prior U.S. inventions. We test these predictions in Section 6.1.[9]

Online Appendix A develops additional mechanisms, including regulatory harmonization and strategic patenting, and provides supplemental tests of those mechanisms.

4 Data

We combine administrative, patent, inventor, and trade data to measure how U.S. down-regulation changes innovation and product markets abroad. The data include: (i) a history of FDA reclassification events, (ii) global patent records with bibliometric and textual information, (iii) linked inventor histories from PatentsView and PATSTAT, (iv) FDA submission, registration, and adverse event records, and (v) bilateral international trade flows. Our main analysis is organized as a stacked panel at the country-by-device-by-year level. Table 1 provides summary statistics.

Reclassification events. We compile a history of FDA reclassification events from 1980 to 2015. The main analysis includes all 29 treated device types. Eleven events followed petitions from device manufacturers. Because manufacturer petitions may be more common in commercially promising device categories, we also report robustness checks excluding these events.[10] Appendix Table A.2 lists all 29 treated device types, announcement years, and

[9] Online Appendix B.2.2 reports robustness and placebo checks for these measures.

[10] The FDA receives many reclassification petitions but only acts on relatively few. Manufacturer-

Federal Register citations. Appendix B.1 presents case studies showing that reclassification timing was driven by statutory directives or accumulated safety evidence, rather than by contemporaneous changes in domestic or foreign innovation. This institutional evidence supports the assumption that event timing is exogenous to foreign innovation trends.

Patent data and country assignment. We use the Dimensions database to build a global patent corpus linked to device types (Hook et al., 2018). For each device type, we extract keywords from its formal FDA description and pull granted patents whose titles or abstracts contain those terms. The resulting data include roughly 5.7 million patent records and 3.5 million patent families. Our main innovation outcome is a value-weighted count of patent families. To construct this outcome, we first collapse jurisdiction-level records to a single observation per family, keeping the earliest filing year.[11] Second, we weight each family by its size, measured as the number of distinct patent documents it contains across jurisdictions. Family size is a common proxy for invention value: more valuable inventions tend to be more protected in more countries, and family size predicts realized patent value about as well as forward citations (Harhoff et al., 2003; Lanjouw et al., 1998). This weighting therefore gives greater weight to inventions that receive broader international protection, while avoiding reliance on forward citations, which are sparse and uneven across the national patent offices in our foreign sample.[12] Third, we allocate each weighted family across countries according to the residence share of its inventors. For example, a family with two inventors in Germany and one in Japan contributes two-thirds to Germany and one-third to Japan. This inventor-based assignment measures where innovation occurs rather than where the resulting intellectual property is owned.[13]

petitioned events may nevertheless occur in more commercially active device types, where firms have signaled interest in regulatory change. Thus, we interpret estimates excluding these events as a conservative robustness check that removes the categories most likely to have been selected on private commercial information.

[11] Rare records without a family identifier are kept as single observations.

[12] The unweighted family count, which counts each family once regardless of size, yields a smaller but same-signed and statistically significant effect (0.057, s.e. 0.011) and is reported as a robustness check. Weighting by family size captures the stronger response among more widely protected inventions, whereas unweighted patent counts treat all patents equally regardless of value.

[13] Appendix Table A.3 repeats the main analysis using fractional assignment based on assignee geography and finds attenuated effects, consistent with response operating more strongly through the geography of

Inventor histories. We build inventor histories by linking PatentsView inventor identities to PATSTAT person and publication records. PatentsView provides longitudinal inventor identifiers and location histories, while PATSTAT expands coverage across patent offices and connects inventors to device patents outside the U.S. patent system. We use these histories for three analyses. First, we classify foreign patent teams according to whether any inventor had previously been observed at a U.S. address. Second, we decompose foreign patenting into new-to-device inventors, incumbents already active in the focal device type, and experienced device inventors entering from other device categories. Third, we construct measures of cross-border relocation using persistent changes in patent-reported inventor residence.[14] Appendix B.2 describes the inventor linkage, residence, and migration definitions.

Trade flows and exposure. We measure bilateral trade flows and country-specific baseline trade exposure using the BACI Database for International Trade Analysis (Gaulier and Zignago, 2010). BACI reports bilateral trade at the six-digit Harmonized System product-code level (HS-6) from 1996 to 2023. We link FDA device types to HS-6 codes in two stages. First, we use a classification algorithm to assign each device, based on its formal description, its most probable HS-6 heading.[15] Then, we hand-audit all treated device-type mappings by comparing each proposed heading with the FDA device description, HS heading text, U.S. Customs rulings, and the HS-2022/HS-1996 crosswalk.[16] Appendix Table A.4 reports the full derived device-to-HS-6 concordance. We use trade data for two purposes. First, we aggregate bilateral flows to the device-country-year level to measure each country's baseline connection to the U.S. market. Import exposure is the share of a country's imports in the relevant HS-6

R&D than through the ownership location of internationally connected firms.

[14]We keep the activity and migration measures separate because static country assignments are useful for counting inventor activity but do not provide valid evidence of physical movement.

[15]The initial classification is produced using `gpt-4o-mini`.

[16]Three treated device types receive specific mappings: rigid and soft daily-wear contact lenses and magnetic resonance imaging systems. The remaining 26 linkages are broader, including seven especially broad mappings for in vitro diagnostic devices. We classify mappings as *specific* when the HS-6 heading isolates the device and as *broad* when the heading also contains related goods. All control headings are medical-device related, and no control shares an HS-6 heading with a treated device. As a robustness check independent of both the classification model and the hand audit, we construct an alternative concordance by matching device and HS-heading descriptions using cosine similarity from `text-embedding-3-large`. The route-level trade estimates are similar under both concordances.

code from the United States. Export exposure is the share of a country's exports in that code that is shipped to the United States. Trade connectedness is the unweighted average of these two measures. Import exposure proxies for contact with U.S. products and technical knowledge, while export exposure proxies for exposure to U.S. demand. Second, we create a device-country-pair-year panel to study trade across global bilateral routes.[17] Appendix B.2 describes the concordance, the construction of the exposure measures, and a higher-resolution corroborative trade analysis.

Innovation direction outcomes. In addition to quantity, patent text allows us to examine whether foreign innovation changes direction. We embed patent titles and abstracts into a common technology space using OpenAI's `text-embedding-3-large` model (OpenAI, 2024), and then average the resulting vectors within each country-device-year cell. Intuitively, patents describing similar technologies lie closer together in this space, while patents describing different technologies lie farther apart. We use patent text rather than citations because citation coverage varies substantially across national patent offices, with roughly one-quarter of the foreign patents in our sample containing usable citation information.[18] We construct two families of direction outcomes. The first measures whether foreign patent text moves closer to the U.S. frontier within the same treated device type. The second measures whether foreign innovation moves in the same direction as U.S. patent text moves after down-regulation. Appendix B.2.2 provides the formal definitions and validation exercises.

5 Empirical Strategy

Our empirical strategy uses a stacked difference-in-differences design (Cengiz et al., 2019). For each down-regulation event, we construct a separate panel containing the treated device cohort, its admissible controls, and a common event-study window. We then stack these

[17] Because most FDA device types map to HS-6 headings that include related products, we interpret the estimates as changes in the broader product baskets containing treated devices rather than in the exact FDA-defined device types.

[18] Appendix B.2.4 discusses the embedding construction and novelty measures in greater detail.

event-specific panels and allow fixed effects to vary across stacks. The estimating equation is

$$Y_{c,p,t,r} = \gamma_{c,p,r} + \gamma_{c,t,r} + \beta_1 \, \mathrm{did}_{p,t,r} + \varepsilon_{c,p,t,r}, \tag{1}$$

where c indexes countries, p device types, t years, and r reclassification events. Each stack is defined at the country-event level, so treated and control device types are compared within the same country and event window. The outcome $Y_{c,p,t,r}$, such as $\log(1 + \text{patent count})$, is measured at the country-device-year level.

The country-device-event fixed effects $\gamma_{c,p,r}$ absorb time-invariant differences across device types within each country and event, while the country-year-event fixed effects $\gamma_{c,t,r}$ absorb shocks common to all device types in a country-year within that event window. The treatment indicator $\mathrm{did}_{p,t,r}$ equals one for the reclassified device type after the anticipation-adjusted event year and zero otherwise. Because each FDA down-regulation applies to the treated device type in every foreign country, $\hat{\beta}_1$ is identified from the differential post-event change for the treated device type relative to its matched controls within the same country and stack.[19] Standard errors are clustered at the country-by-treated-device-type level.

To study how the response varies with countries' exposure to U.S. product markets, we augment the baseline specification with interactions using baseline trade measures. We first estimate

$$Y_{c,p,t,r} = \gamma_{c,p,r} + \gamma_{c,t,r} + \beta_1 \, \mathrm{did}_{p,t,r} + \beta_2 \, (\mathrm{did} \times \mathrm{trade})_{c,p,t,r} + \beta_3 \, (\mathrm{did} \times \ln GDP)_{c,p,t,r} + \varepsilon_{c,p,t,r}. \tag{2}$$

Here $\mathrm{trade}_{c,r}$ is the trade connectedness of country c for the treated device type in cohort

[19] Because controls are never-reclassified device types, a control device type cannot itself become treated. Inventors who move from a control device type into the treated one after the event are the experienced-pivot entry margin we measure directly (Section 6.2); this margin is small relative to new-to-device entry, and the categories that supply pivots into treated devices expand rather than contract, so contamination of the control group is limited.

r, measured over a fixed 1996–2000 baseline window (described in Section 4). We also interact treatment with $\ln GDP_c$, the baseline log real GDP per capita for country c in 1994, from the Penn World Table, so that the trade-exposure coefficients capture differences in responsiveness beyond those associated with country income.

We next separate trade connectedness into import and export exposure:

$$\begin{aligned} Y_{c,p,t,r} = \gamma_{c,p,r} + \gamma_{c,t,r} + \beta_1 \,\mathrm{did}_{p,t,r} \\ + \beta_2\big(\mathrm{did}_{p,t,r} \times \mathrm{imp_from_US}_{c,r}\big) + \beta_3\big(\mathrm{did}_{p,t,r} \times \mathrm{exp_to_US}_{c,r}\big) \\ + \beta_4(\mathrm{did}_{p,t,r} \times \ln GDP_c) + \varepsilon_{c,p,t,r}. \end{aligned} \tag{3}$$

The import and export measures ($\mathrm{imp_from_US}_{c,r}$ and $\mathrm{exp_to_US}_{c,r}$) capture country c's baseline exposure to U.S. treated-device imports and exports, respectively, for the treated device type in cohort r. Both are measured over the fixed 1996–2000 window. Section 4 and Appendix B.2 describe their construction.

As in any difference-in-differences design, identification requires that treated and control device types would have followed similar trends absent down-regulation. We assess this using a pooled stacked event study:

$$Y_{c,p,t,r} = \gamma_{c,p,r} + \gamma_{c,t,r} + \sum_{k(t,r)\neq -1} \beta_k \, 1\{\mathrm{Treated}\}_{p,r} \times I_{k(t,r)} + \varepsilon_{c,p,t,r}, \tag{4}$$

where $1\mathrm{Treated}_{p,r}$ indicates that device type p is treated in event r, and $I_{k(t,r)}$ indexes years relative to treatment, with $k = -1$ omitted.[20] We define $k(t, r) = 0$ using the anticipation-adjusted treatment year described below. We estimate analogous dynamic specifications for equations 2 and 3.

Constructing a credible counterfactual for foreign innovation is a central challenge. We

[20]Note that within a country a treated device type enters only its own event's cohort, so it carries a single cohort-by-device fixed effect $\gamma_{c,p,r}$; a control device type is reused across the cohorts of multiple events and therefore carries a separate cohort-by-device fixed effect in each. For example, daily-wear soft contact lenses (treated in 1994) enter only the 1994 contact-lens cohorts. In contrast, extended-wear soft contact lenses, used as an intuitive control, enter the cohorts of several treated events, each with its own fixed effect. This reuse is why we cluster on the treated device type.

use three distinct control groups to address different sources of potential bias. Our preferred specification uses five nearest-neighbor device types that are never reclassified and have similar pre-event patenting levels.[21] This approach most closely aligns outcome levels of treated and control device types before down-regulation.[22]

We use two alternative control groups, both covering the full sample of treated devices. The first consists of clinically similar device types: for each treated device, a never-reclassified device in the same FDA regulation panel but a different product family, of comparable pre-period patenting volume (Appendix Table A.6 lists the paired control for each treated device). These controls better capture common market conditions. The second uses device types that are down-regulated after the end of the relevant sample window. Because these devices are ultimately judged suitable for the same regulatory change, they provide a useful comparison against devices that remain at the higher classification for longer.

Inference follows the level at which treatment variation is observed. For patent outcomes, we cluster at the treated-device-by-country level. Because the identifying variation ultimately comes from a small number of policy changes—29 device reclassifications for the value-weighted count and 22 for the directional-alignment measure—we verify that the headline estimates survive inference at the level of the shock. Both survive clustering on the treated device alone, two-way clustering by device and country, and a wild cluster bootstrap: the count remains significant at the 5% level (wild bootstrap $p = 0.013$) and the directional-alignment effect at the 0.1% level (Appendix Table A.7). For bilateral HS-6 trade outcomes, where the treatment is a small number of device reclassifications, we report Conley–Taber confidence intervals with the reclassified device as the unit (Conley and Taber, 2011).[23]

[21] For each treated device type, we compute the pre-event mean of inventor-fractional patent counts over the five years preceding the announcement year and select the five nearest never-reclassified controls by squared Euclidean distance on these raw pre-event means. Controls must have at least one pre-period patent (like treated groups), and device types sharing the treated device's clinical indication are excluded from the candidate pool.

[22] Appendix Table A.5 reports pre-period covariate balance.

[23] The remaining outcomes, largely appendix material (e.g., HS-10 trade, contractor entry, and adverse events), are built from U.S.-side panels with no foreign-country dimension, so each treated device collapses to a single time-series and the cluster count effectively equals the number of treated device types. For those outcomes we report Conley–Taber 95% Monte Carlo confidence intervals (Conley and Taber, 2011), which

Down-regulations may be anticipated before formal enactment through petitions, proposed rules, advisory-panel activity, and Federal Register notices. For patent outcomes, we therefore define treatment as beginning three years before the formal reclassification year. This window captures the typical announcement–enactment delay over which firms could plausibly respond to regulatory signals before enactment.[24] Event time $k = 0$ corresponds to this announcement year, and the patent event-study window spans $k \in [-8, +20]$. The results are not sensitive to this choice. Appendix Table A.10 varies the anticipation window from zero to five years; the estimated effect remains positive and significant throughout, ranging from 0.109 to 0.154. Dating treatment to the formal reclassification year, with no anticipation adjustment, yields an estimate of 0.153 (s.e. 0.018). We use a three-year window because petitions precede formal reclassification by a median of three years, making earlier responses plausible. The stability of the estimates across alternative windows confirms that the results are not driven by this choice.

For outcomes that respond more directly to implementation, we use later treatment dates. Trade, outsourcing, and adverse event outcomes use a two-year anticipation window, while FDA submissions are dated to the enactment year because responsive submissions cannot precede the regulatory change that permits them.

Some outcomes, particularly the embedding-based direction measures, lack an intuitive scale in raw units. We therefore report their magnitude relative to the outcome's dispersion, in standard deviations, and, as a complementary translation, in years of pre-event technological drift. The years benchmark compares the estimated treated-control gap (i.e., the DiD estimate) with the pre-event rate of change among matched controls and asks how many years of that baseline drift would generate a change of the same size.[25]

construct the reference distribution from control-group residuals and do not require large-sample asymptotics over the number of treated units.

[24] Across the 29 treated reclassifications, the FDA proposed-rule lead has a median of 1 year and a maximum of 4 years; the petition-filing lead has a median of 3 years and a maximum of 5 years. Three years covers the modal petition-filing lead and approximates the maximum proposed-rule lead, and is conservative against earlier petition leakage. Pre-trend tests are computed relative to this anticipation-adjusted timing.

[25] Online Appendix B.2.2 provides the exact construction and alternative normalizations.

6 Results

We organize the results around three margins: innovation, inventors, and product markets. We first ask whether foreign patenting increases and moves toward the U.S. technological path. We then examine which inventors account for the response. Finally, we study how trade in treated devices changes after down-regulation. Throughout, we use cross-country differences in pre-event U.S. trade exposure to help distinguish among competing mechanisms.

6.1 Foreign Innovation and Technological Direction

We begin with the foreign patenting response. Down-regulation lowers the cost of bringing a treated device to the U.S. market, and U.S.-based inventors respond strongly. Figure 2 shows that U.S. patenting rises substantially after reclassification, consistent with the domestic response documented in Rogers (2023). We take this domestic response as the starting point and ask how far the innovation shock extends beyond the United States.

We first test Prediction 1, which asks whether down-regulation increases foreign innovation. Figure 3 plots the pooled stacked event study for foreign value-weighted patenting. Pre-event coefficients are close to zero, and a joint F-test fails to reject the null of no differential pre-trends. Foreign patenting rises after down-regulation and continues to build over the following decade.

The static DiD estimate is positive and significant across all three control groups, with coefficients between 0.122 and 0.225 (Table 2, Panel A).[26] In the matched-control specification, the Poisson event study implies a 55% increase in patenting (an 84% increase under a level translation of the log-point estimate). The long-run estimate for event time ≥ 7 is 0.189, indicating that the response continues to grow well after down-regulation.

[26]This headline estimate is reported in log points because the 61 non-U.S. countries differ in baseline patenting volume by orders of magnitude. Therefore, a level-based country-level regression would weight the regression toward the highest-volume countries and obscure the proportional response across countries. No single country drives the estimate: dropping any one country leaves it in the range 0.125–0.145, always significant at the 1% level (the most influential single country is Germany, at 0.125).

The foreign response is also large relative to the domestic effect. To compare magnitudes directly, we aggregate value-weighted patent counts across all non-U.S. countries and estimate the same design separately for the non-U.S. world and the United States.[27] In the long run, down-regulation induces 75.0 additional value-weighted foreign patents per device-type-year, compared with 162.7 additional U.S. patents. The implied multiplier is therefore 0.46,[28] indicating that every additional U.S. patent induced by down-regulation is accompanied by roughly half an additional patent abroad (Figure 2; Online Appendix **??**).

We next examine the factors that predict where the foreign response is largest. Because import exposure is measured over a fixed 1996–2000 window, our primary mechanism specification uses the 16 reclassifications that occur after 2000, for which the exposure window is entirely pre-treatment. On this predetermined-exposure sample the import interaction is 0.288 (s.e. 0.107); the corresponding export interaction is indistinguishable from zero. The full-sample interaction is nearly identical and more precisely estimated (0.294; Table 2, Panel C), and is positive and significant across all three control groups.[29] These results are consistent with Prediction 2.

Figure 4 shows how these differences evolve over time. After deregulation, the import-exposure gradient grows steadily, while the export-exposure gradient remains near zero. Thus, foreign innovation responds more strongly where prior exposure to U.S. products was greater, but not where access to the U.S. market was more important. Appendix Figure A.2 shows the same pattern without imposing a linear relationship: the response is small in the lowest tercile of pre-event U.S. import exposure and large and persistent in the middle and highest terciles.

This import–export asymmetry helps distinguish among mechanisms. If lower U.S. commercialization costs mainly raise the return to serving the U.S. market, we would expect a

[27] The pooled multiplier uses all non-U.S. countries rather than the 61-country sample used for the country-level interaction analysis. The latter requires baseline trade and GDP covariates, while the pooled comparison does not.

[28] Device-level bootstrap 95% confidence interval: $[0.23, 0.59]$

[29] The import gradient is robust to using total real GDP rather than per-capita real GDP as the country-size control, with the sign and significance unchanged.

stronger response among countries already oriented toward exporting treated devices to the United States. Instead, the response is concentrated in countries that imported more U.S. treated devices before down-regulation. This pattern is consistent with diffusion through U.S. products and the technical knowledge embodied therein.

We do not interpret the import gradient as identifying a single mechanism. Greater import exposure may also proxy for broader technological sophistication, supplier relationships, or integration with the U.S. device sector. Several additional results weaken these explanations. First, the response is strongest for imports of intermediate and final goods rather than capital goods (Appendix C.3.2). Second, the import gradient survives controlling for a country's human capital.[30] Third, the gradient is robust to controlling for a country's overall reliance on U.S. medical device imports.[31] These patterns suggest that exposure to U.S. products plays a stronger role than access to U.S. demand and broader country-level integration with the U.S. device sector.

Foreign innovation also changes direction, consistent with Prediction 3. We measure direction with four outcomes. All four are built from embeddings of patent titles and abstracts and are compared to U.S. patents in the same treated device type (Appendix B.2.2). The embedding space, which could be interpreted as a knowledge space, is useful for studying "location" of patents and "distance" (similarity) between patents. The first two measures study the average location (centroid) of a country's patents in the device type, which allow to know the average location of new knowledge produced. The last two measures study how close each foreign patent lies to (patent-specific) set of individual U.S. patents, which allows to study pairwise proximity. **Diagram XXX in Appendix YYY explains each of the**

[30] Adding the interaction between treatment and a country's human capital—the schooling-based Penn World Table index, measured in the baseline year—leaves the device-specific import gradient at 0.258 (s.e. 0.091), little changed from 0.296 and significant at the 1% level. Human capital is itself positively associated with the response (0.165, s.e. 0.033), but does not account for the import gradient.

[31] Adding the interaction between treatment and the country-level share of all medical-device imports originating in the United States leaves the device-specific import gradient at 0.23 (s.e. 0.10, $p < 0.05$), while the broader U.S. device-import share is not independently predictive of a patenting response (0.16, s.e. 0.16). The response therefore tracks which device types a country sources from the United States rather than its general orientation toward U.S. medical-device imports.

measures graphically.

Figure 5 plots event studies for the two centroid outcomes. The first, cosine similarity to the U.S., measures whether the foreign centroid lies closer to the U.S. centroid in the same year. In other words, it says whether foreign innovation converges in the knowledge space to the U.S. innovation. The second, directional alignment, measures whether the foreign centroid moves, relative to its own pre-event position, in a similar direction to how the U.S. centroid moved after reclassification. It equals the cosine of the angle between the two movements: +1 when the foreign centroid moves exactly along the same direction the U.S. moved, 0 when it moves in an orthogonal direction. Pre-event coefficients are near zero for both outcomes, and both rise after reclassification. Foreign patent text becomes more similar to U.S. patent text (DiD 0.020, s.e. 0.008) and moves is a similar direction as the U.S. moved (DiD 0.110, s.e. 0.012), the latter significant at the 0.1% level (Table 3, Panel A).[32] Because embedding outcomes have no natural units, we express magnitudes relative to the outcome's dispersion and to the pre-event drift of matched controls. The similarity effect is roughly 0.2 standard deviations, or about four years of pre-treatment convergence of matched controls. The directional-alignment effect is roughly 0.8 standard deviations, or about eleven years of pre-treatment drift of matched controls (Appendix B.2.2).

The directional response emerges earlier than the volume response. Within a few years of reclassification, directional alignment reaches much of its long-run level, while the patent count is still early in its decade-long build-up. We would expect this difference in timing: a normalized directional measure moves as soon as the first induced patents point toward the U.S. path, whereas a count rises only as new entry accumulates. We may be concerned that alignment reflects overall movement in embedding space rather than movement toward the treated U.S. path. A placebo that replaces the U.S. direction of the treated device type with U.S. directions of unrelated device types yields estimates centered near zero: across 50

[32]Re-embedding the patent corpus using PaECTER (Ghosh et al., 2024), a model trained specifically on patent text, produces a nearly identical directional-alignment DiD of 0.118 (Appendix B.2.2).

draws, the mean estimate is −0.001, and none exceeds 0.110 (Appendix B.2.2).[33]

Alignment with the U.S. path could reflect imitation of U.S. inventions, in other words, two patents being exactly in the same location of the knowledge space. Howewer, the previous analysis focuses on the average (centroid) across a country's relevant patents, rather than studying individual patents per se. The same same centroid could be obtained with very different patents. We therefore turn to study foreign individual patents distance to a set of nearest-neighbor U.S. patents (Appendix B.2.4). For each foreign patent family, we compute the distance (one minus cosine similarity) to its five most similar U.S. families in the same device type, and take the average across them. A foreign family that copies a U.S. invention scores close to zero. We do two different analysis by computing the distance against two sets of U.S. patents.

The first set contains all U.S. families filed at least two years before the foreign family.[34] Hence, the reference set of U.S. patents enlarges as more patents are filed in the U.S., which could be interpreted as the contemporaneous U.S. "knowledge stock". If foreign inventors copy U.S. designs, old or new, we would expect the distance to this set to fall. It does not: the DiD is 0.006 (s.e. 0.005, $p = 0.18$; Appendix Table A.8, Panel A).[35]

The second set freezes the U.S. reference at reclassification and contains only U.S. families filed at least two years before the event year. Hence, this is the U.S. knowledge stock at the time of reclassification. A rising distance to this set means that foreign patents locate

[33]As a complementary flow measure, Appendix Figure A.3 uses the annual change in the foreign centroid projected onto the U.S. direction. The estimate is positive but imprecise (0.007, s.e. 0.004; $p = 0.08$), consistent with down-regulation accelerating the movement of foreign innovation toward the U.S. path.

[34]Two years is the approximate lag between filing and public disclosure. The set of U.S. patents grows over time and includes U.S. patents filed after reclassification.

[35]U.S. patenting in the treated category rises after reclassification. A larger pool of U.S. patents mechanically shrinks nearest-neighbor distances even when the content of foreign patents does not change. Suppose we compute distance of patent A to the 5 closest U.S. patents. Then, as the set of U.S. patent grows -even randomly-, there would be a positive probability that one of the new patents would be closer to patent A than one of the existent 5 closest patents. Hence, we need to do the comparison by fixing the size (but not the patents themselves) of the patent pool to which we compute distance. We hold the size of the pool fixed by averaging the distance over ten random subsamples of 100 U.S. families. Without this correction, the distance falls (−0.009, $p = 0.09$). As imitation may be done to the "latest" patents, we repeat the analysis using as reference set the U.S. families filed within the preceding three, five, or eight years (Appendix B.2.4), also finding that distance does not fall.

in regions of technology space that no U.S. patent in the device type occupied at reclassification. The distance rises relative to controls (DiD 0.010, s.e. 0.005, $p = 0.04$; Appendix Table A.8, Panel B). The gap is about 0.13 standard deviations of the cross-sectional novelty distribution, or roughly seven years of normal advance among matched controls.[36] Taking together the last two results, foreign innovation distinguishes from previous U.S. innovation while at the same time it does not copy the new U.S. innovation.

In sum, average foreign innovation in treated devices moves closer to the U.S. in the knowledge space, and moves in the same direction as the U.S. moved, relative to control devices. However, the new foreign innovation, while being more similar on average, does not individually become more proximate to contemporaneous U.S. innovation, and differentiates from existing pre-deregulation innovation. Hence, the analysis suggest that while the new foreign knowledge shifts towards more U.S. similar knowledge, the new foreign knowledge differentiates from U.S. knowledge and does not imitate it. This combination of results is consistent with foreign inventors building on U.S. products (Appendix Table A.1).

The foreign patenting results are robust to dropping individual countries (footnote 26) or treated device types, alternative anticipation windows and panel-balance restrictions, assignee-based country assignment, alternative patent classifications, and Poisson estimation.[37] Appendix Tables A.9, A.10, and A.11 and Appendix Figure A.5 report these checks.

6.2 Inventor Entry, Teams, and Mobility

We next ask which inventors account for the foreign response. We first ask whether the response comes from new inventors or from inventors already working in the treated technology. We then examine whether the response depends on teams with U.S.-based inventors or on cross-border inventor movement.

We divide foreign inventors into three mutually exclusive groups: *new-to-device* inventors

[36] Results point in the same direction when we use the single nearest U.S. family (0.008, $p = 0.10$) or the ten nearest (0.011, $p = 0.03$) instead of five (Appendix Table A.8).

[37] Additionally, to show that these patents translated into commercialized products, Appendix Figure A.4 plots U.S. and non-U.S. FDA-submission event studies.

with no prior patenting in the device sector, *incumbents* already active in the treated device type, and *experienced pivots* with prior device patents but no prior patents in the treated category. Table 3 (Panels B and C) reports static log estimates of the entry decomposition by inventor type, and Appendix Figure A.7 plots the corresponding event studies.

New-to-device inventors account for most of the response. Table 3 reports a statistically significant DiD log estimate of 0.196 for new-to-device inventor counts (2.83 in levels), compared to 0.046 for experienced pivots (0.21 in levels).[38] Incumbents are relatively unaffected. The patent-level decomposition is similar: patents by new-to-device inventors rise by 0.115 log points, compared to 0.031 for experienced pivots and 0.011 for incumbents.

Therefore, the foreign innovation response is driven by an expansion of the set of foreign inventors working in treated technologies rather than an increase in activity among experienced innovators. This composition is consistent with a product-market channel. Greater exposure to expanding U.S. products may make entry into treated device categories more attractive or provide information that substitutes for prior experience. The response therefore comes primarily from new inventors rather than incumbents expanding existing work. This pattern mirrors the concentration of the domestic response among new firms in Rogers (2023) and is less consistent with mechanisms that depend on established relationships, such as knowledge transfer through U.S.-based co-inventors. New-to-device inventors account for 91% of treated-category patenting after down-regulation, while patents produced jointly by inventors and firms already active in the treated category account for less than one percent.[39]

The entry response does not dilute the quality of foreign innovation. New-to-device

[38]Experienced pivots could respond by shifting effort away from other device categories rather than increasing total inventive activity. Two descriptive patterns weigh against this interpretation. First, inventors who pivot into treated devices do not reduce patenting in the categories they previously worked in; their patenting in those source categories rises after the event. Second, the source categories that supply pivots into treated devices expand rather than contract relative to categories that supply none. We do not interpret either comparison causally, since the inventors and fields connected to treated devices are unusually dynamic. But the data show little evidence of the contraction in adjacent activity that a simple reallocation account would predict. This question concerns substitution across foreign device fields, distinct from the cross-border reorientation toward U.S. technologies discussed in Section 7.

[39]Cross-classifying patents by prior inventor and assignee experience shows that new inventors work predominantly at firms also new to the category (52% of treated-category patents) and at established device firms entering from adjacent categories (25%).

patent families are nearly as novel as those produced by experienced pivots or incumbents when measured by their distance from prior U.S. art, and the novelty gap between new entrants and established inventors is essentially the same in treated and control categories. Because entrants account for most of the induced patenting, the frontier-extension results above largely reflect innovation produced by these new entrants.

We then test whether the response depends on direct involvement by U.S.-based inventors. Appendix Figure A.8 compares patenting by foreign teams with and without at least one U.S.-based inventor. Teams without U.S.-based inventors respond just as strongly as teams with U.S.-based inventors. This result suggests that the diffusion of knowledge is not specific to U.S.-based inventors.

The same pattern appears in our analysis of the direction of foreign innovation. Teams with no U.S.-based inventor move significantly toward the U.S. trajectory, accounting for about 84% of the pooled directional-alignment effect (Appendix Figure A.9). Teams with U.S.-based inventors also move significantly toward the U.S. trajectory and have a somewhat larger point estimate, although the difference between the two groups is not statistically significant.

These results do not support direct co-invention with U.S.-based inventors as the main channel through which foreign inventors learn from and adapt to U.S. innovation. Instead, foreign teams reorient toward the U.S. technological path even without a U.S.-based inventor on the team.

Direct tests of cross-border inventor relocation lead to the same conclusion. We find that down-regulation does not increase inventor movement either into or out of the United States. Appendix D.1 reports the migration event studies.

Taken together, the inventor results show that down-regulation draws new foreign inventors into treated technologies and the foreign response does not require direct U.S. inventor participation or physical relocation. These results are most consistent with diffusion through products, disclosures, technical specifications, and other forms of knowledge that can travel

without inventor movement.

We also find that the induced patenting reflects genuine inventive activity rather than strategic filing around a newly profitable U.S. market. Four patterns support this claim: (i) foreign patents move along the U.S. technological trajectory; (ii) new-to-device inventors account for most of the response, and most work at firms that are themselves new to the device or pivoting into the treated category; (iii) the response concentrates among foreign inventors at foreign firms rather than U.S.-owned foreign R&D (Appendix C.3.1, Table A.12);[40] and (iv) the induced patents extend beyond the pre-event U.S. frontier rather than converge onto prior U.S. inventions. Together, these patterns are difficult to reconcile with a purely strategic-IP response.[41]

6.3 Trade and Product-Market Responses

For U.S. products to transmit the effects of down-regulation abroad, they must first reach foreign markets. Section 6.1 shows that foreign innovation responds most strongly in countries that previously sourced more of a treated device from the United States. We now ask whether down-regulation expands trade in treated devices, particularly U.S. exports to foreign markets. Beyond providing a channel through which the innovation response can spread abroad, the trade effects are important in their own right: domestic reclassification changes trade in treated devices well beyond the U.S. market. We use two complementary trade datasets. Global HS-6 data cover bilateral trade across the full international network, while U.S. HS-10 data provide a more granular mapping between individual FDA device

[40]The foreign-assignee, foreign-inventor cell yields $\hat{\beta} = 0.132^{**}$ versus $\hat{\beta} = 0.025^{+}$ for the U.S.-assignee, foreign-inventor cell ($\hat{\beta}_B - \hat{\beta}_A = +0.107^{**}$, $z = +6.00$), inconsistent with a story in which U.S.-headquartered firms file foreign-inventor IP to position around an expanded U.S. market.

[41]Most of the induced response (86%) comes from patent families with a granted U.S. patent, which firms could assert strategically in U.S. courts. We link these patents to the Stanford NPE Litigation Database (Miller et al., 2018), which classifies the plaintiffs in U.S. district-court patent suits filed between 2000 and 2024. We find little evidence that patent assertion entities, which earn revenue from enforcing patents rather than selling products, litigate these patents. Families ever litigated by these entities account for almost none of the induced response. After down-regulation, they make up 0.10% of treated-category patenting, compared to 0.16% in control categories. Families litigated by any plaintiff account for less than 0.5% of the induced response.

types and U.S. imports and exports.[42] Here, we focus on evidence using HS-6 routes and provide additional validation of our results using the finer HS-10 resolution in Appendix D.3.

Figure 6 plots route-level event studies and Table 4 reports the corresponding stacked DiD estimates. Trade increases along the major route types. U.S. exports to foreign destinations increase by 0.397 log points (49%), foreign exports to the U.S. by 0.463 log points (59%), and non-U.S.-to-non-U.S. exporter-side trade by 0.246 (28%), each statistically significant. The corresponding importer-side non-U.S.-to-non-U.S. estimate is smaller and statistically insignificant, at 0.130 (14%).

The effects amplify over time. In the long run, U.S. exports to foreign countries rise by 0.530 log points, foreign exports to the U.S. by 0.632, non-U.S.-to-non-U.S. trade by 0.350 on the exporter side and 0.176 on the importer side. While HS-6 categories are broader than individual FDA device types, capturing trade in product baskets containing treated devices, they allow us to follow trade impacts across the full global trade network.

More granular HS-10 trade data, which are available only for U.S.-linked trade, confirm the HS-6 results for these routes. At this higher resolution, imports of U.S. treated devices rise by 0.195 log points and foreign exports to the United States by 0.247. In the long run, the effects reach 0.329 and 0.336, corresponding to increases of roughly 39% and 40% (Table 4, Panel A; Appendix FigureA.10).[43] The agreement between the HS-6 and HS-10 estimates indicates that the global route results are not driven solely by neighboring products contained in the broader HS-6 categories.[44] The increase in trade value reflects more goods shipped rather than higher prices. Decomposing the response on each route into changes in physical quantities and unit values, we find little evidence of a price response. Quantities increase across all routes and significantly so along those with the largest trade effects (Appendix

[42]We report the trade response in customs value; Appendix Table A.14 decomposes it into physical-quantity and unit-value (price) components.

[43]The hand-audited HS-10 concordance is provided in Appendix Table A.13.

[44]Trade composition also changes. U.S.-linked routes gain relative importance over the post-event window, although the composition estimates are less precise than the volume effects. The share of treated-device trade on U.S.-linked routes rises by 0.034 (s.e. 0.035) after the event, so the data confirm the volume prediction but not the stronger share-tilt condition of the trade-volume prediction formalized in Online Appendix A.

Table A.14).

The bilateral data show that this expansion is globally distributed rather than concentrated in a few countries. Most trade routes expand, and the U.S. route sees the largest expansion on all routes, with the U.S. becoming a more central player in treated-device trade (Appendix Figure A.11). Country-specific bilateral routes both rise and fall. For example, German and Chinese exports to the U.S rise, while Japanese and Korean exports to the U.S. fall. Together, the findings suggest that down-regulation expands global trade, with the largest increases on routes involving the down-regulating country.[45]

Taken together, the trade results connect the exposure gradient to the product-market response. Down-regulation expands the flow of U.S. products abroad, and foreign innovation responds most strongly in countries with greater prior exposure to those products (Section 6.1).

Appendix D.2 reports two additional product-market outcomes. Down-regulation sharply increases contractor entry, where contractors are firms that manufacture on behalf of another firm. From a near-zero baseline, both foreign and domestic contractors show significant increases in U.S. entry. We then ask whether the accompanying expansion in foreign innovation and cross-border device flows is associated with worse safety outcomes. Adverse events reported for treated devices from foreign sources do not increase; Instead, the associated point estimates are negative, though statistically insignificant. Thus, the expansion in foreign innovation and trade is not accompanied by detectable deterioration in post-market safety. We interpret this as suggestive evidence that the global expansion in supply did not come at the expense of product quality.

One concern is that manufacturer-petitioned reclassifications may reflect private information about promising product areas. We therefore repeat the main analyses using the twelve events initiated independently by the FDA, which is the strictest subset of the non-

[45]Estimates are differences-in-differences against matched-control device baskets, with controls matched within exporter. The largest dollar increases are in U.S. exports (to Germany, $2.6 million; to China, $1.4 million) and in foreign exports to the United States (Germany and China, roughly $1.1 million each). The largest relative declines are in U.S.-bound exports from Japan and Korea (roughly $0.9 million each).

manufacturer events introduced above (Appendix Table A.15 reports both subsets). The results are broadly similar, although some estimates attenuate and become less precise in the smaller sample. Foreign patenting still increases, innovation continues to shift toward the U.S. technological path, new-inventor entry remains positive, and trade continues to expand. The import-exposure gradient also remains positive (Appendix Table A.15).

7 Conclusion

This paper shows that domestic regulation in a large market can shape innovation well beyond the regulator's borders. U.S. medical-device down-regulation induces roughly one foreign patent for every two it induces domestically, implying that a domestic evaluation overlooks about one-third of the innovation response. Foreign patenting rises, shifts toward the U.S. technological path, and extends beyond the pre-existing U.S. frontier. The response is driven largely by new foreign inventors rather than by U.S.-linked teams or inventor mobility, and it is accompanied by substantial growth in international trade. These effects also build over time rather than appearing as a one-time response to the regulatory change.

Taken together, these patterns point to a product-market channel. Foreign innovation responds most strongly in countries with greater prior exposure to U.S. products, but not in countries for which the U.S. is a more important export market (Figure 4). The response also remains strong among teams without U.S.-based inventors. These results favor a role for exposure to U.S. products over either a higher return to serving the U.S. market or knowledge transfer through U.S.-based collaborators. They do not distinguish among the mechanisms operating through those products. U.S. products may convey technical knowledge, demonstrate commercially viable designs, or create compatibility incentives. Import exposure may also partly capture broader integration with the U.S. device sector, though several checks in Section 6.1 weaken this reading.

These findings imply that the incidence of domestic regulation is broader than is typi-

cally measured. A regulator evaluating the consequences of administrative rules using only domestic innovation and safety outcomes misses important implications for foreign inventors, firms, and international product markets. In our setting, these responses are economically large. This incidence extends to product markets as well as innovation. Down-regulation expands trade in treated devices, including between foreign countries whose trade does not directly involve the United States. At the same time, the induced foreign innovation becomes more closely aligned with the technological path of the lead market, showing that domestic regulation can shape the direction of global innovation.

Several features of the setting limit our interpretation. First, while useful, patents are an imperfect measure of inventive value, and greater alignment with the U.S. technological path could reflect productive diffusion or reduced technological diversity. Nor can we fully distinguish new inventive activity from reallocation across all possible uses of scientific and engineering effort; however, the dominance of new-to-device inventors and the absence of contraction in the fields supplying experienced pivots are reassuring. Likewise, the safety estimates show no detectable deterioration among foreign-made devices but are too imprecise to rule out modest effects. Regarding trade, the most detailed HS-10 data are available only for U.S. imports and exports, limiting how precisely we measure the impacts of down-regulation on trade. Lastly, medical devices combine relatively short product cycles, incremental innovation, strong trade linkages, and close complementarities between technology and clinical practice. Whether domestic regulation produces similar global responses in industries with different innovation characteristics and product-market structures remains an open question.

More broadly, the results suggest that regulation in a technological lead market can function as a form of global innovation policy even when it is designed around domestic objectives.

References

Acemoglu, Daron and Joshua Linn, "Market Size in Innovation: Theory and Evidence from the Pharmaceutical Industry," *Quarterly Journal of Economics*, 2004, *119* (3), 1049–1090.

Agha, Leila, Soomi Kim, and Danielle Li, "Insurance Design and Pharmaceutical Innovation," *American Economic Review: Insights*, 2022, *4* (2), 191–208.

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt, "Competition and Innovation: An Inverted-U Relationship," *Quarterly Journal of Economics*, 2005, *120* (2), 701–728.

Ayerst, Stephen, Faisal Ibrahim, Gaelan MacKenzie, and Swapnika Rachapalli, "Trade and Diffusion of Embodied Technology: An Empirical Analysis," *Journal of Monetary Economics*, 2023, *137*, 128–145.

Beise, Marian, "Lead Markets: Country-Specific Drivers of the Global Diffusion of Innovations," *Research Policy*, 2004, *33* (6–7), 997–1018.

_ **and Klaus Rennings**, "Lead Markets and Regulation: A Framework for Analyzing the International Diffusion of Environmental Innovations," *Ecological Economics*, 2005, *52* (1), 5–17.

Belenzon, Sharon and Mark Schankerman, "Spreading the Word: Geography, Policy, and Knowledge Spillovers," *Review of Economics and Statistics*, 2013, *95* (3), 884–903.

Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 2013, *81* (4), 1347–1393.

_ **, Mirko Draca, and John Van Reenen**, "Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT, and Productivity," *Review of Economic Studies*, 2016, *83* (1), 87–117.

Blume-Kohout, Margaret E. and Neeraj Sood, "Market Size and Innovation: Effects of Medicare Part D on Pharmaceutical Research and Development," *Journal of Public Economics*, 2013, *97*, 327–336.

Budish, Eric, Benjamin N Roin, and Heidi Williams, "Do Firms Underinvest in Long-Term Research? Evidence from Cancer Clinical Trials," *American Economic Review*, 2015, *105* (7), 2044–85.

Bustos, Paula, "Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinean Firms," *American Economic Review*, 2011, *101* (1), 304–340.

CDRH, "About Manufacturer and User Facility Device Experience (MAUDE)," 2023.

Cengiz, Doruk, Arindrajit Dube, Attila Lindner, and Ben Zipperer, "The Effect of Minimum Wages on Low-Wage Jobs," *Quarterly Journal of Economics*, 2019, *134* (3), 1405–1454.

Clemens, Jeffrey, "The Effect of U.S. Health Insurance Expansions on Medical Innovation," Technical Report Working Paper 19761, National Bureau of Economic Research 2013.

_ **and Parker Rogers**, "Demand Shocks, Procurement Policies, and the Nature of Medical Innovation: Evidence from Wartime Prosthetic Device Patents," *Review of Economics and Statistics*, 2026, *108* (1), 75–89.

Coe, David T. and Elhanan Helpman, "International R&D Spillovers," *European Economic Review*, 1995, *39* (5), 859–887.

Coelli, Federica, Andreas Moxnes, and Karen-Helene Ulltveit-Moe, "Better, Faster, Stronger: Global Innovation and Trade Liberalization," *Review of Economics and Statistics*, 2022, *104* (2), 205–216.

Conley, Timothy G and Christopher R Taber, "Inference with "Difference in Differences" With a Small Number of Policy Changes," *The Review of Economics and Statistics*, 2011, *93* (1), 113–125.

Ensign, Lisa G. and K. Bretonnel Cohen, "A Primer to the Structure, Content and Linkage of the FDA's Manufacturer and User Facility Device Experience (MAUDE) Files," *eGEMs*, 2017, *5* (1), 12.

European Patent Office, "PATSTAT Global: Worldwide Patent Statistical Database," 2025. Spring 2025 edition.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer, "The Next Generation of the Penn World Table," *American Economic Review*, 2015, *105* (10), 3150–3182.

Finkelstein, Amy, "Static and Dynamic Effects of Health Policy: Evidence from the Vaccine Industry," *Quarterly Journal of Economics*, 2004, *119* (2), 527–564.

Gaulier, Guillaume and Soledad Zignago, "BACI: International Trade Database at the Product-Level. The 1994–2007 Version," Working Paper 2010-23, CEPII 2010.

Ghosh, Mainak, Michael E. Rose, Sebastian Erhardt, Erik Buunk, and Dietmar Harhoff, "PaECTER: Patent-level Representation Learning using Citation-informed Transformers," 2024.

Harhoff, Dietmar, Frederic M. Scherer, and Katrin Vopel, "Citations, Family Size, Opposition and the Value of Patent Rights," *Research Policy*, 2003, *32* (8), 1343–1363.

Hook, Daniel, Simon Porter, and Christian Herzog, "Dimensions: Building Context for Search and Evaluation," *Frontiers in Research Metrics and Analytics*, 2018, *3*, 23.

Imbens, Guido W. and Donald B. Rubin, *Causal Inference for Statistics, Social, and Biomedical Sciences: An Introduction*, Cambridge University Press, 2015.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson, "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quarterly Journal of Economics*, 1993, *108* (3), 577–598.

Ji, Yunan and Parker Rogers, "The Long-Run Impacts of Regulated Price Cuts: Evidence from Medicare," Technical Report Working Paper 33083, National Bureau of Economic Research 2024.

Jones, Charles I. and John C. Williams, "Measuring the Social Return to R&D," *The Quarterly Journal of Economics*, 1998, *113* (4), 1119–1135.

Keller, Wolfgang, "International Technology Diffusion," *Journal of Economic Literature*, 2004, *42* (3), 752–782.

Lanjouw, Jean O., Ariel Pakes, and Jonathan Putnam, "How to Count Patents and Value Intellectual Property: The Uses of Patent Renewal and Application Data," *Journal of Industrial Economics*, 1998, *46* (4), 405–432.

Lileeva, Alla and Daniel Trefler, "Improved Access to Foreign Markets Raises Plant-Level Productivity... For Some Plants," *Quarterly Journal of Economics*, 2010, *125* (3), 1051–1099.

Makower, Josh, Aabed Meer, and Lyn Denend, "FDA Impact on US Medical Technology Innovation; A Survey of Over 200 Medical Technology Companies," 2010.

Medicare Payment Advisory Commission, "An Overview of the Medical Device Industry," Technical Report, Medicare Payment Advisory Commission June 2017. Report to the Congress: Medicare and the Health Care Delivery System.

Miller, Shawn P., Ashwin Aravind, Bethany Bengfort, Clarisse De La Cerda, Matteo Dragoni, Kevin Gibson, Amit Itai, Charles Johnson, Deepa Kannappan, Emily Kehoe, Hyosang Kim, Katherine Mladinich, Roberto Pinho, John Polansky, and Brian Weissenberg, "Who's Suing Us? Decoding Patent Plaintiffs since 2000 with the Stanford NPE Litigation Dataset," *Stanford Technology Law Review*, 2018, *21* (2), 235.

OpenAI, "Embeddings," 2024.

Pierce, Justin R. and Peter K. Schott, "Concording U.S. Harmonized System Codes over Time," *Journal of Official Statistics*, 2012, *28* (1), 53–68.

Powell, Spenser F, "Changing Our Minds: Reforming the FDA Medical Device Reclassification Process," *Food & Drug LJ*, 2018, *73*, 177.

Rogers, Parker, "Regulating the Innovators: Approval Costs and Innovation in Medical Technologies," *Working Paper*, 2023.

Sampat, Bhaven and Heidi L. Williams, "How Do Patents Affect Follow-On Innovation? Evidence from the Human Genome," *American Economic Review*, 2019, *109* (1), 203–236.

Schott, Peter K., "The Relative Sophistication of Chinese Exports," *Economic Policy*, 2008, *23* (53), 5–49.

_ , "International Economics Resource Page," 2025.

Stern, Ariel Dora, "Innovation Under Regulatory Uncertainty: Evidence From Medical Technology," *Journal of Public Economics*, 2017, *145*, 181–200.

United States Patent and Trademark Office, "PatentsView Data Download Tables," 2025.

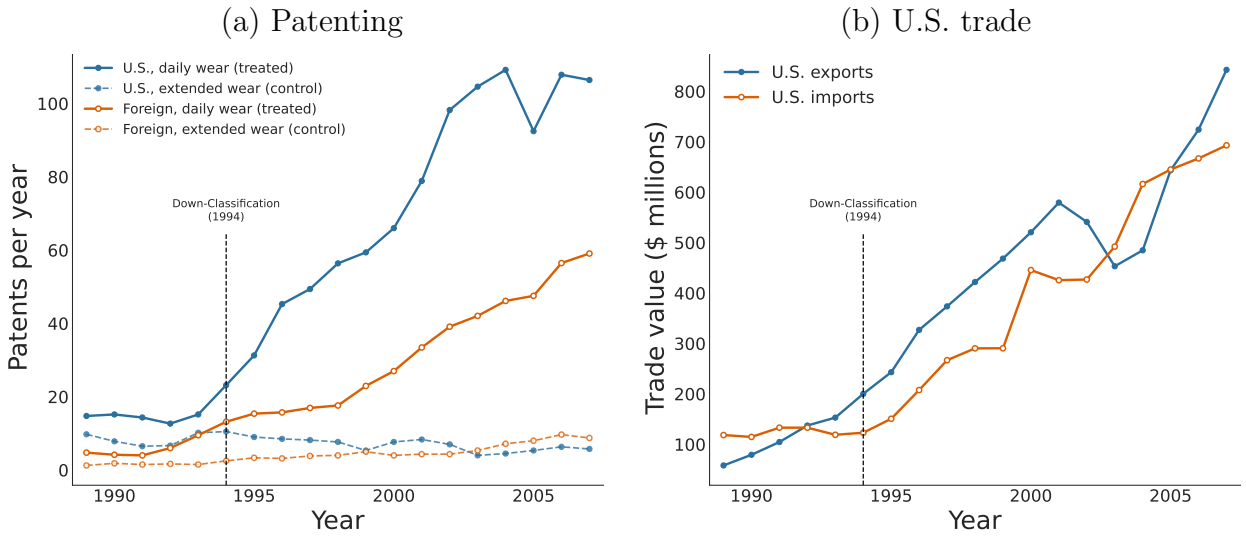
U.S. Food and Drug Administration, "510(k) Premarket Notification Database," 2025.

_ , "Premarket Approval (PMA) Database," 2025.

Vernon, Raymond, "International Investment and International Trade in the Product Cycle," *Quarterly Journal of Economics*, 1966, *80* (2), 190–207.

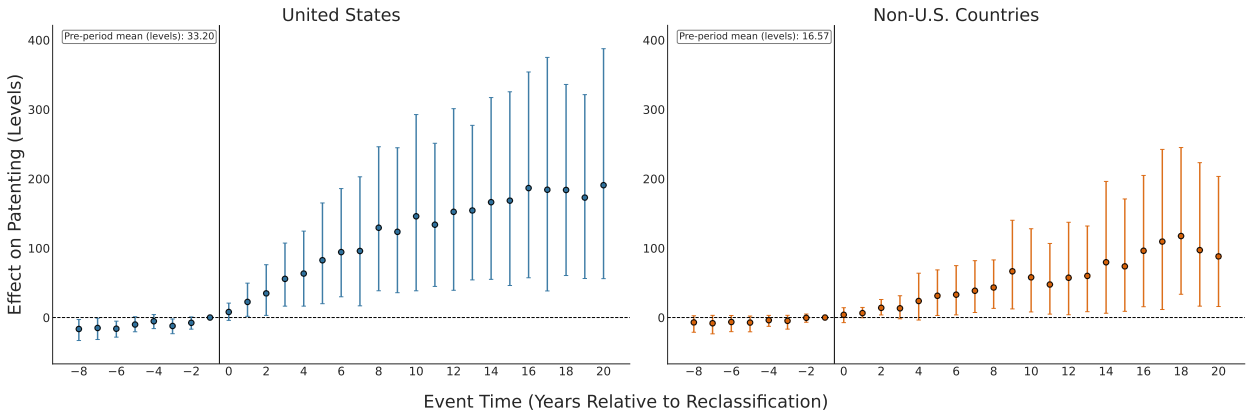
Figures and Tables

Figure 1: A Representative Down-Classification: Soft Contact Lenses



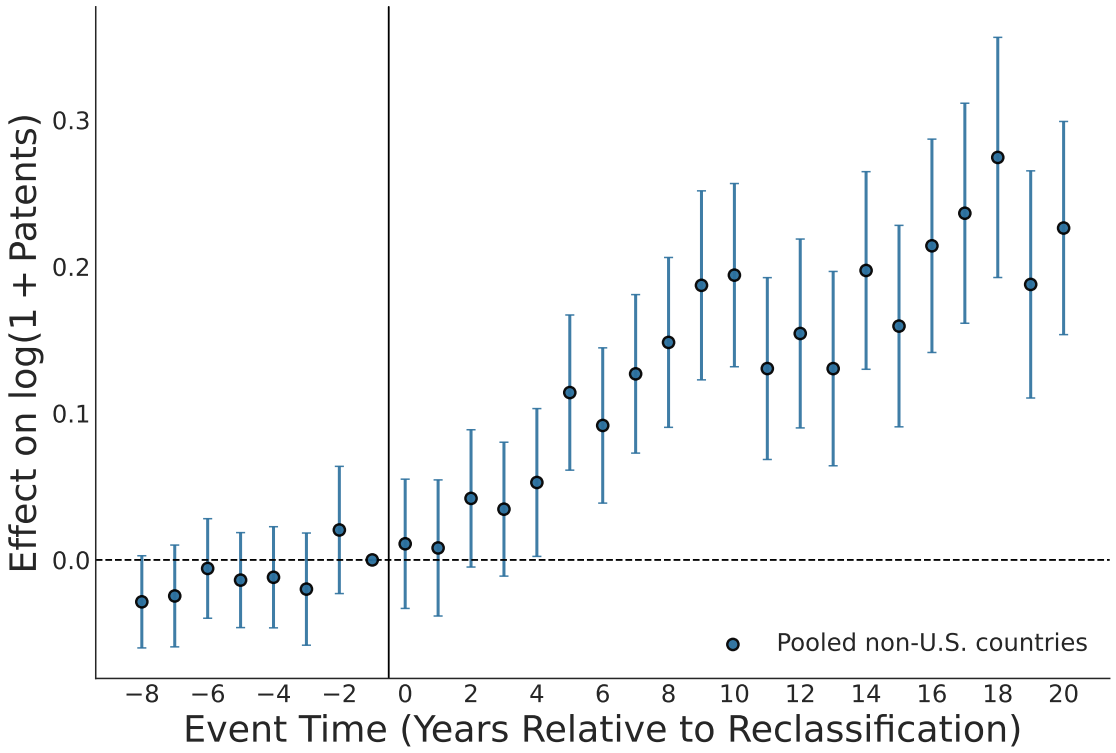
Note: In 1994 the FDA reclassified daily-wear soft contact lenses from Class III to II; the extended-wear variant stayed in Class III as control. *Panel (a)* plots inventor-fractional patents per year for U.S. (blue, filled) and foreign (orange, open) inventors, daily-wear (treated, solid) versus extended-wear (control, dashed). Treated patenting rises sharply after 1994 while the control barely moves; the U.S. rise exceeds the foreign, roughly two-to-one as in aggregate. *Panel (b)* plots U.S. contact-lens exports and imports (millions of dollars) at HS-10 heading 9001.30, which does not separate daily- from extended-wear, so no device-level trade control exists and the panel shows the whole category. Patents are inventor-fractional (non-U.S. pooled) from the global panel; trade from Schott HS-10 U.S. records.

Figure 2: U.S. and Non-U.S. Induced Patenting in Levels



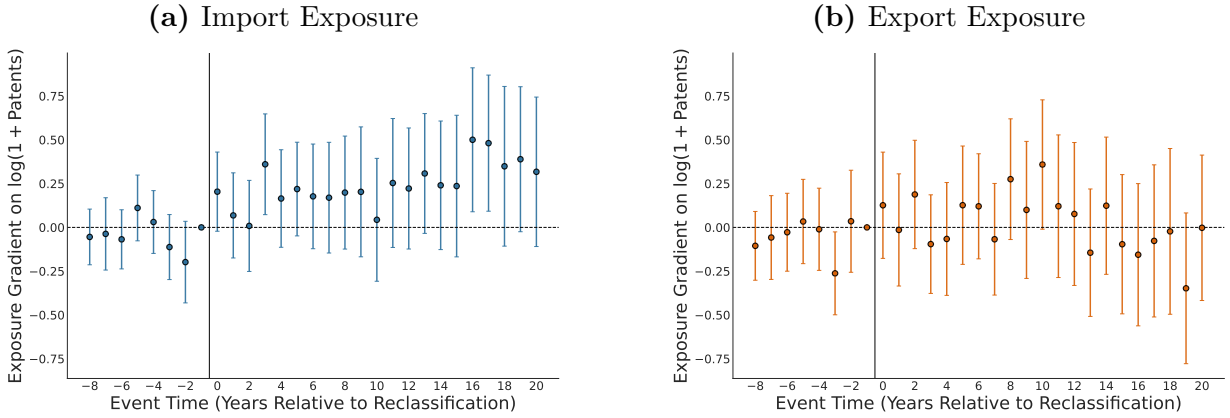
Note: Pooled region-comparison stacked event studies of induced patenting in *levels* around Class III-to-II reclassification, U.S. (left) and pooled non-U.S. world (right) on a common vertical scale. Each panel plots the average treated-minus-matched-control patent count against matched never-reclassified controls. Pre-event coefficients are flat in both regions; annotations give pre-period mean levels (33.2 U.S., 16.6 non-U.S.). Both responses build over two decades, with long-run DiD effects (event time ≥ 7) of 162.7 value-weighted patents per device-type-year (U.S.) and 75.0 (non-U.S.). The long-run levels multiplier is 0.46 (about half the domestic response), device-level bootstrap 95% CI $[0.23, 0.59]$, excluding zero and one (Section 6.1; Online Appendix **??**). Bars are device-level bootstrap 95% confidence intervals.

Figure 3: Ideas Margin: Pooled Stacked Event Study



Note: Pooled stacked DiD event study of value-weighted patenting in 61 non-U.S. countries (outcome: log(1+ patent count) at the country-by-device-by-year level). The design pools 29 Class III to II reclassification events with five matched never-reclassified control device types per treated event. Event time is anticipation-adjusted (the raw reclassification year is $k = 3$); $k = -1$ is the omitted reference period; the window is $[-8, +20]$. Standard errors are clustered at the treated-device-by-country level. Static DiD = 0.139 (s.e. 0.017), $N = 136{,}649$, pre-trend joint F-test $p = 0.31$.

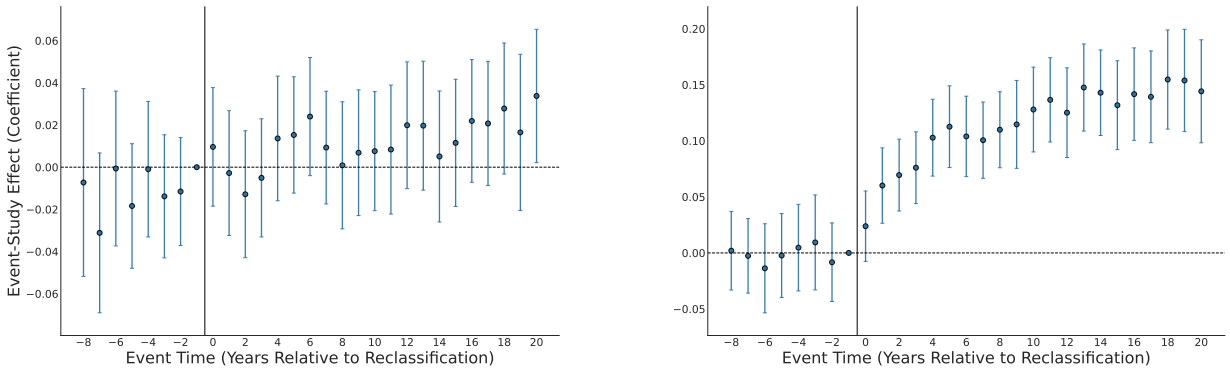
Figure 4: Innovation Responses by Import and Export Exposure



Note: Each panel plots the event-time path of an exposure-interaction coefficient from a single stacked DiD regression of value-weighted log(1 + patent count) on the treated-device indicator interacted with event time and, additionally, with pre-period U.S. import exposure, U.S. export exposure, and per-capita GDP (all measured at the country-by-treated-device level). *Panel (a)* plots the import-exposure gradient and *Panel (b)* the export-exposure gradient; because both interactions enter the same regression, each gradient is net of the other margin and of per-capita GDP. Points are coefficients relative to event time −1, bars are 95% confidence intervals, and the two panels share a common vertical axis. The specification uses matched controls with cohort × product-code and cohort × year fixed effects, with standard errors clustered at the treated-device-by-country level ($N = 136{,}649$). Collapsing the post-period to a single coefficient recovers the static interaction gradients (Table 2, Panel C): import +0.294 (s.e. 0.092) and export +0.072 (s.e. 0.093, not significant). Joint pre-trend F-tests do not reject ($p = 0.14$ import, $p = 0.27$ export). Pre-period exposure is the country-by-treated-device import (export) share from (to) the United States over a fixed 1996–2000 window under the HS-6 concordance (Online Appendix Table A.4). The response loads on import exposure and builds over the post-period, while the export gradient is flat at zero.

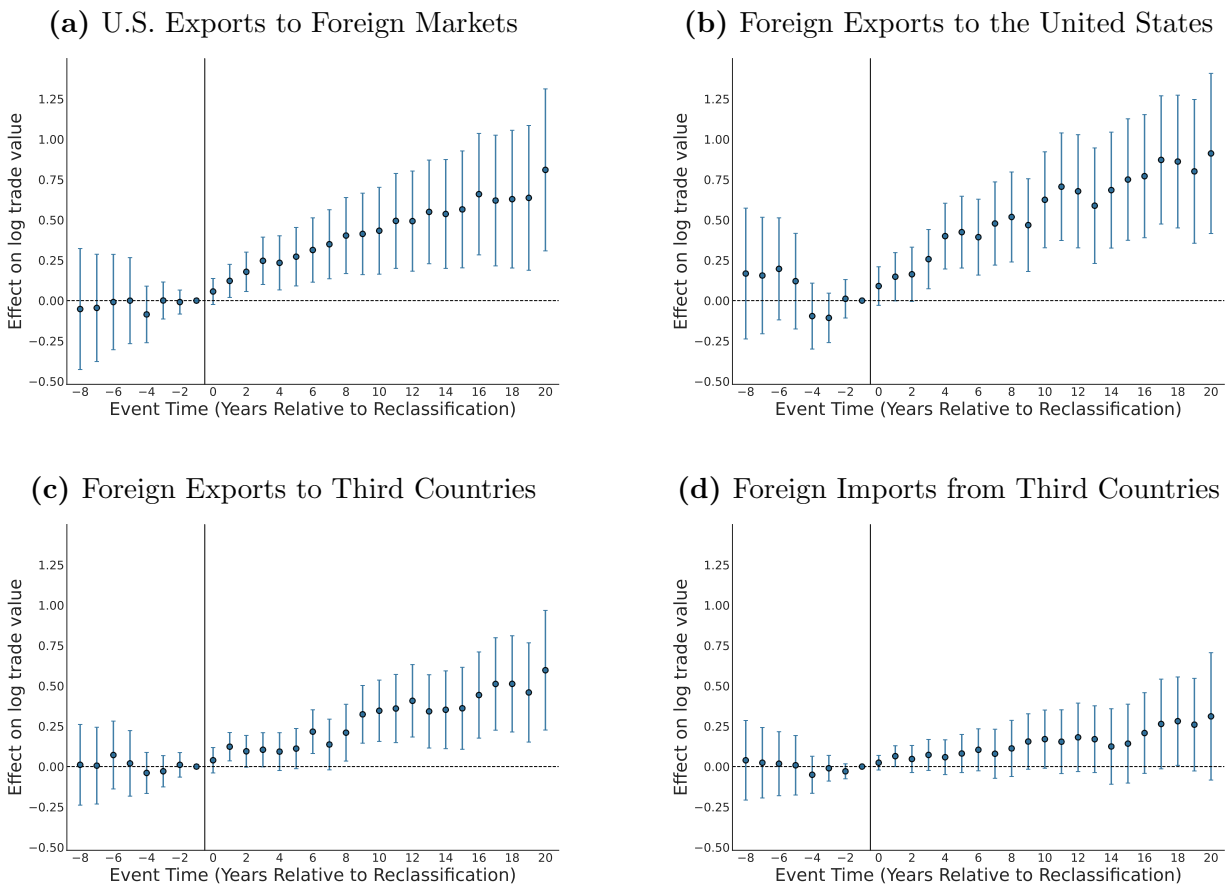
Figure 5: Ideas Margin: Innovation Direction

(a) Cosine Similarity to U.S. **(b)** Directional Alignment with U.S.



Note: Pooled stacked event studies of the direction of induced non-U.S. patenting (matched-control sample). *Panels (a) and (b)* report two direction outcomes, both unitless on $[-1, 1]$: cosine similarity to U.S. innovation (level convergence—foreign text becoming more similar to U.S. text in the category) and directional alignment (co-movement—foreign innovation tracking the U.S. post-deregulation trajectory). Static DiD: cosine 0.020 (s.e. 0.008), roughly 0.2 standard deviations of the outcome; alignment 0.110 (0.012), roughly 0.8. Relative to pre-period organic drift these gaps correspond to about four and eleven years of convergence (Online Appendix B.2.2). Outcomes from 3,072-dimensional patent-text embeddings (title+abstract); centroids are inventor-fractional weighted means. Points are DiD coefficients vs. $k = -1$; bars 95% CIs. Construction and robustness: Online Appendix B.2.2.

Figure 6: Products Margin: Route-Level HS-6 Trade Event Studies



Note: Conley–Taber event studies of log(HS-6 trade value) on the patent-country bilateral network (exporter-by-importer-by-device-by-year cells where both partners are in the patent panel). Values are BACI HS-6 (Gaulier and Zignago, 2010), mapped to FDA device types via the hand-audited concordance (Online Appendix Table A.4) for the treated devices and the broader candidate-control universe for controls. Points are DiD coefficients versus $k = -1$ with 95% Conley–Taber confidence intervals (Conley and Taber, 2011), with the reclassified device as the unit: each of the 17 treated devices is a distinct FDA regulatory action, and the reference distribution is formed by drawing placebo devices from the pool of never-reclassified controls and applying the treated timing. Treatment timing is the device-reclassification year (anticipation-adjusted). Long-run effects (event time ≥ 7) and joint pre-trend p-values are reported in Table 4.

Table 1: Summary Statistics

	Mean	SD
Panel A: Reclassification Events		
Class III to II treated device types (main matched sample)	29	
FDA-initiated Class III to II subset	12	
Announcement/event years	1983–2012	
Panel B: Stacked Panel		
Observations (matched controls)	136,649	
Countries	61	
Treated device types	29	
Control device types per treated	5	
Event-study window	$[-8, +20]$	
Panel C: Innovation Outcomes (Pre-Period)		
Inventor-fractional patent count	0.216	0.938
$\log(1 + \text{patent count})$	0.106	0.332
New-to-device inventor count	0.038	0.681
Experienced-pivot inventor count	0.011	0.246
Focal-code incumbent inventor count	0.732	3.635
Panel D: Trade Exposure (Pre-Period)		
Import share from U.S.	0.198	0.230
Export share to U.S.	0.142	0.215
Trade connectedness	0.170	0.186

Note: Panel A describes the reclassification sample. The main matched-control analysis uses the full estimable down-regulation sample; the FDA-initiated subset is used in supporting control-group specifications and robustness checks. Panel B describes the realized baseline stacked panel using matched controls. Panel C reports pre-treatment means and standard deviations of the primary innovation outcome. The pre-treatment mean of 0.106 is for $\log(1 + \text{patent count})$. Panel D reports pre-period trade exposure measures at the country-by-treated-device level. Import share from U.S. measures the fraction of a country's device-category imports originating in the United States. Export share to U.S. measures the fraction of device-category exports destined for the U.S. market. Trade connectedness combines both directional margins. Inventor-type counts are at the country-by-device-by-event-by-year cell level over the pre-event period; type assignment is described in Appendix B.2.

Table 2: Stacked DiD Estimates by Control Group

	(1) Matched	(2) Matched (LR)	(3) Intuitive	(4) Later-Treated
Panel A: Basic DiD				
DiD (value-weighted)	0.1392**	0.1890**	0.1222**	0.2247**
	(0.0168)	(0.0215)	(0.0177)	(0.0186)
DiD (unweighted)	0.0566**	0.0772**	0.0397**	0.1078**
	(0.0105)	(0.0134)	(0.0112)	(0.0116)
Panel B: Trade-Connectedness Heterogeneity				
DiD × Trade Connectedness	0.3798**	—	0.4657**	0.5731**
	(0.1028)		(0.1179)	(0.1273)
Panel C: Directional Trade Heterogeneity				
DiD × Import Share from U.S.	0.2943**	—	0.2891**	0.3632**
	(0.0921)		(0.0924)	(0.0964)
DiD × Export Share to U.S.	0.0721	—	0.1696	0.2003+
	(0.0927)		(0.1071)	(0.1169)
Panel D: Exogeneity Subsets				
Non-manufacturer (18 devices)	0.0577**	0.0917**	0.1283**	0.1308**
	(0.0180)	(0.0239)	(0.0189)	(0.0181)
FDA-initiated (12 devices)	0.0761**	0.1287**	0.2077**	0.1561**
	(0.0205)	(0.0275)	(0.0280)	(0.0212)
Implied effect in patent counts (matched controls)				
Treated patenting, pre-period	0.555			
Treated patenting, post-period	3.365			
Counterfactual patenting, post-period	1.834			
Implied increase	+55% Poisson (+84% level)			
Observations	136,649	136,649	371,581	274,681
Pre-treatment mean	0.1544	0.1544	0.1544	0.1544
Pre-trend F-test p-value	0.3105	0.3105	0.0755	0.3642

Note: Stacked difference-in-differences estimates of value-weighted (family-size-weighted) log(1 + patent count) for non-U.S. inventors, by control group, on the full sample of 29 treated devices. Column (1) uses five nearest-neighbor never-reclassified matched controls per treated device; Column (2) reports the long-run effect on the matched sample (average effect at event time ≥ 7 years, absorbing the short-run window); Column (3) uses clinical-neighbor controls (a never-reclassified device in the same FDA regulation panel as each treated device, patent-scale-matched); Column (4) uses later-treated controls (every device reclassified after the sample period). For Columns (3) and (4) each treated device is compared against the full control pool. Panel A reports the basic DiD, value-weighted (the baseline outcome throughout) and unweighted (each family counted once); Panels B and C the trade-connectedness and import/export-share interactions, each from a regression that also includes DiD and DiD× ln(per-capita GDP). Panel D re-estimates the count on the 18 non-manufacturer-petitioned and 12 FDA-initiated subsets, controls re-selected for the retained devices; every cell is significant at the 1% level with clean pre-trends. All specifications use inventor-fractional geography, a three-year anticipation window, and cohort-by-device and cohort-by-year fixed effects (a cohort is a country-by-treated-device pair); standard errors, in parentheses, cluster at the country-by-treated-device level. The lower block translates the Column (1) estimate into patent counts. $^{+}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$.

Table 3: Static Stacked DiD: Innovation Direction and Entry Decomposition

	Pre-period mean	$\hat{\beta}_{\text{DiD}}$, levels	$\hat{\beta}_{\text{DiD}}$, $\log(1+y)$
Panel A: Innovation direction†			
Cosine similarity to U.S. frontier	0.7271	0.0198*	—
		(0.0076)	
Directional alignment with U.S.	−0.0047	0.1102**	—
		(0.0119)	
Growth-matched controls (robustness)		0.1232**	—
		(0.0105)	
Alignment magnitude (auxiliary)	0.0000	0.0660**	—
		(0.0061)	
Negative Euclidean distance to U.S. (auxiliary)	−0.6420	0.0429**	—
		(0.0118)	
Sample size		8,263	
Panel B: Foreign inventor counts by inventor type			
New-to-device	0.044	2.8309**	0.1959**
		(0.9061)	(0.0172)
Experienced pivot	0.012	0.2141*	0.0462**
		(0.0988)	(0.0066)
Focal-code incumbent	0.750	0.0456+	0.0160**
		(0.0268)	(0.0041)
Sample size		147,030	
Panel C: Foreign patents by inventor type			
New-to-device	0.014	0.7507**	0.1153**
		(0.2077)	(0.0125)
Experienced pivot	0.005	0.1013*	0.0308**
		(0.0431)	(0.0052)
Focal-code incumbent	0.294	0.0221*	0.0107**
		(0.0089)	(0.0028)
Sample size		147,030	

Note: Static stacked DiD on the Class III-to-II matched-control panel (Section 5); standard errors clustered at the treated-device×country level in parentheses. *Panel A* reports four embedding-based direction outcomes: two headline (cosine similarity to the U.S. frontier; directional alignment with its post-deregulation trajectory) and two auxiliary (alignment magnitude; negative Euclidean distance). All are unitless on $[-1, 1]$ with no $\log(1+y)$ analog, and the pre-period mean is the treated-cell pre-event average; the smaller sample reflects the requirement of at least one focal-device patent per cell. The growth-matched robustness row re-matches controls on foreign patent-count pre-trajectories rather than the outcome's own, yielding 0.123 (s.e. 0.011) with clean pre-trends (F-test $p = 0.34$). *Panels B and C* decompose foreign inventor and patent counts into three pre-period inventor types; the levels coefficients, relative to the control post-period mean, imply treated-versus-control post-period gaps of 112%/89% (new-to-device), 93%/105% (experienced pivot), and 18%/20% (incumbent), preferred to $\exp(\hat{\beta})-1$ at near-zero pre-treatment counts. Pre-trend joint F-tests pass for every outcome, and the decomposition is robust to dropping the office-country inventor-geography fallback (about two-thirds of non-U.S. inventor-years). † Direction outcomes (Panel A) are unitless on $[-1, 1]$, with no $\log(1+y)$ analog. $^{+}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$.

Table 4: Products Margin: Treated-Device Trade Volumes

	Static DiD	Long-run effect	Long-run %	95% CI (long-run)
Panel A: U.S.-linked HS-10 trade values				
Foreign exports to the United States	0.247^{+}	0.336^{**}	40%	[0.07, 0.57]
Foreign imports of U.S. goods	0.195	0.329^{*}	39%	[0.02, 0.60]
Panel B: HS-6 U.S.-linked routes				
U.S. exports to foreign markets	0.397^{**}	0.530^{**}	70%	[0.21, 0.86]
Foreign exports to the United States	0.463^{**}	0.632^{**}	88%	[0.22, 0.89]
Panel C: HS-6 foreign-to-foreign routes				
Foreign exports to third countries	0.246^{*}	0.350^{**}	42%	[0.17, 0.61]
Foreign imports from third countries	0.130	0.176	19%	[−0.11, 0.32]

Note: Static and long-run (event time ≥ 7 years) stacked DiD estimates of log treated-device trade value, with the reclassified device dated three years before finalization; the long-run effect absorbs the first six post-event years and reports the mature response, translated to a level change (long-run %). Panel A uses U.S. trade value at the device-by-year level from Schott (2025), with HS-10 codes mapped to FDA device types via the hand-audited concordance in Appendix Table A.13; the confidence intervals are 95% Monte-Carlo Conley–Taber intervals and the stars derive from the symmetric-equivalent standard error. Panels B and C use HS-6 bilateral route values from BACI (Gaulier and Zignago, 2010) on the closed patent-country network, mapped via the concordance in Appendix Table A.4; inference is Conley–Taber with the reclassified device as the unit (each of the 17 treated devices is a distinct regulatory action; the reference distribution draws placebo devices from the never-reclassified control pool, $B = 999$). The response is significant on the U.S.-linked HS-10 flows and routes and on foreign exports to third countries, but not on foreign imports from third countries, consistent with a supply-side (production and export) response. Pre-event coefficients are jointly insignificant on all four routes (Conley–Taber $p = 0.38$–0.98; Figure 6) and, for the HS-10 flows, individually include zero across the pre-event window (Appendix Figure A.10). Foreign-contractor entry and adverse-event estimates appear in Appendix Table A.17. $^{+}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$.

Appendix Material: For Online Publication

A Mechanisms and Tests

In this appendix we describe how each mechanism discussed in Section 3 leaves a different pattern in the data, and we report which patterns we observe. We do not estimate a structural model. First, we describe how down-regulation changes the returns and costs of foreign inventors, and what this implies for the import- and export-exposure gradients. Second, we describe where in technology space induced foreign patents should locate under each mechanism.

A.1 Inventor Effort and Exposure to U.S. Trade

The model presented here focuses on a foreign inventor's effort on a technology and highlights two forces that affecting her decision: demand and cost shifters. We consider multiple demand shifters: improved access to U.S. markets, increased domestic demand due to complementarity (compatibility) with imported devices, and decreased domestic demand due to increased imports (substitution). Costs could shift due to increased knowledge spillovers. We will proxy exposure to such shifters with a country's pre-period import and export exposure to the U.S.

Consider an inventor located in country c (outside the U.S.) who chooses effort $z \geq 0$ on technology p and earns the return per unit of effort minus its cost:

$$\Pi_{cp}(z) \;=\; R_{cp}\, z \;-\; \tfrac{1}{2}\, K_{cp}\, z^2, \tag{A.1}$$

where R_{cp} is the marginal return and $K_{cp} > 0$ is the marginal cost. Optimal effort is then:

$$z^*_{cp} = R_{cp}/K_{cp},$$

with maximized surplus $\Pi^*_{cp} = R^2_{cp}/(2K_{cp})$.

We can express, in a reduced form, the returns and costs as function of the country's pre-period import exposure to the U.S. $m_c p$ and export exposure $x_c p$:

$$R_{cp} = R^0_{cp} + \rho_x x_{cp} + (\rho^c_m - \rho^s_m) m_{cp} \qquad K_{cp} = K^0_{cp} - \kappa_m m_c \tag{A.2}$$

With $rho_x, \rho^c_m, \rho^s_m, \kappa_m \geq 0$. ρ_m captures the increased competition of foreign imports which pushes local demand and hence returns down. ρ_x captures export demand due to easier entry into U.S. market. κ_m captures a cost reduction due to product-embodied knowledge spillovers.

Differentiating optimal effort $z^*_{cp} = R_{cp}/K_{cp}$ with respect to import and export exposure gives:

$$\frac{\partial z^*_{cp}}{\partial m_c} = \frac{\overbrace{\rho^c_m K_{cp}}^{\text{compatibility}} \overbrace{-\rho^s_m K_{cp}}^{\text{substitution}} + \overbrace{\kappa_m R_{cp}}^{\text{knowledge spillovers}}}{K^2_{cp}}, \qquad \frac{\partial z^*_{cp}}{\partial x_c} = \frac{\overbrace{\rho_x}^{\text{export demand}}}{K_{cp}}.$$

Import exposure has opposing effects on effort: (i) it decreases the returns due to local demand substituting away of local devices for foreign devices, (ii) it increases the returns due to the increased demand for complementary devices and compatibility and, (iii) decreases costs due to increased product-embodied knowledge spillovers. Export exposure raises effort due to an increased export demand.

The differential effect of import exposure on innovation might also contain other economic forces. Competition can also raise innovation: in the escape-competition mechanism of Aghion et al. (2005), firms that are technologically close to their rivals innovate more to escape competition. A positive import gradient is therefore also consistent with an escape-competition response. Two results are not supportive of this alternative. First, escape competition operates among incumbents close to the frontier, whereas the foreign response comes from new-to-device inventors and leaves incumbents unaffected (Section 6.2). Second, escape competition does not predict that foreign patents move along the U.S. technological path, which is what we find (Section 6.1).

A.2 Patents in the Technology Space

We aim to test whether foreign inventors respond to U.S. regulatory changes by innovating in new directions or by imitating existing U.S. designs. We do so by representing patents in the technology space creating embeddings of their title and abstract (Appendix B.2.2).[46]

We describe the location of foreign patents with three measures, each relative to U.S. patents in the same device type. The idea of the three measures is to capture different aspects of new patents in the technology space, first looking at the distance of the "average" patent and then looking at the distance of each individual foreign patent to the "neighboring" U.S. patents.

First, *directional alignment* asks whether the foreign centroid moves in the same direction as the U.S. centroid moved after the event (Appendix B.2.2). Second, the *distance to pre-event U.S. art* is the mean distance from each foreign family to its five nearest U.S. families filed at least two years before the reclassification, the approximate lag between filing and public disclosure. It rises when foreign patents locate in regions that pre-event U.S. art did not cover. Third, the *distance to all prior U.S. art* is the same measure computed against every U.S. family filed at least two years before the foreign family, including those filed after the reclassification. It falls when foreign patents locate next to existing U.S. designs, old or new (Appendix B.2.4).[47]

Different economic forces predict different combinations of these three measures. If down-regulation raises the return to serving the U.S. market, for foreign inventors, the share of returns coming from the U.S. rises, which could lead to a increase in directional alignment and a decrease in the distance to pre-event and all prior U.S. art. If foreign inventors imitate the new U.S. designs, locating nearby the new designs, then directional alignment rises and

[46]Embeddings are normalized to unit length. All patents therefore lie on the surface of a sphere, at the same distance from the origin. In the technology space two patents differ only in the angle between them.

[47]U.S. patenting in the treated category rises after deregulation. A larger pool of U.S. art mechanically shrinks nearest-neighbor distances even when the content of foreign patents does not change (when selecting the closest 5 patents, when the pool of U.S. patents increases, there is positive probability that one of those new patents is closer than an existing patent). We hold the size of the reference pool fixed by subsampling.

the distance to all prior U.S. art falls. If foreign inventors build on U.S. products, alignment rises and the distance to pre-event U.S. art rises, while the distance to all prior U.S. art does not fall, because the new patents extend the U.S. designs without imitating them. Finally, if foreign inventors respond to U.S. competition by moving into niches that U.S. producers do not serve, alignment does not rise and both distances rise.

Appendix Table A.1: Predicted and Observed Location of Induced Foreign Patents

	Directional alignment	Distance to pre-event U.S. art	Distance to all prior U.S. art
U.S. demand	0 or +	0 or +	0 or +
Imitation of U.S. designs	+	−	−
Building on U.S. products	+	+	0 or +
Differentiation from U.S. competitors	0 or −	+	+
Observed (DiD, s.e.)	+0.110 (0.012)	+0.010 (0.005)	+0.006 (0.005)

Note: Predicted sign of the stacked DiD estimate under each mechanism. Distances are density-fixed nearest-neighbor distances with $k = 5$ (Table A.8); directional alignment is defined in Appendix B.2.2 and estimated in Table 3, Panel A.

We observe the pattern predicted by building on U.S. products. Directional alignment rises by 0.110 (s.e. 0.012), the distance to pre-event U.S. art rises by 0.010 (s.e. 0.005, $p = 0.04$), and the distance to all prior U.S. art does not fall (+0.006, s.e. 0.005). In other words, foreign inventors move in the direction that U.S. innovation moved and into regions that U.S. art did not cover at reclassification, without locating next to existing U.S. designs. Imitation alone would lower the distance to all prior U.S. art, and differentiation alone would not raise alignment.

A limitation for the analysis is that imitation by some inventors could be offset by differentiation among others, leaving the distance to all prior U.S. art unchanged. A large differentiation response would, however, pull alignment down, whereas alignment rises by roughly 0.8 standard deviations. We therefore cannot exclude a small imitation response, but the data rule out imitation as the dominant one.

B Data, Measurement, and Institutional Background

B.1 Background on Class III to II Events

Following Rogers (2023, Appendix A.1), we describe four Class III to II reclassifications whose timing was driven mostly by process and was hard to predict. In each case the trigger was a U.S. statutory directive, a procedural step in agency review, or the accumulation of evidence already contained in PMA files, rather than contemporaneous technological change or foreign innovation trends. This pattern supports the identifying assumption of our spillover design: the timing of reclassification is exogenous to the foreign innovation responses we measure. For each case we report the Federal Register citation, the broad therapeutic area, and the year of announcement.

Daily-Wear Soft Contact Lenses

FR citation: 59 FR 10284 **Broad area:** Ophthalmic **Year Announced:** 1994

Rationale. The FDA based the down-regulation mainly on PMA evidence that had built up over the preceding decade, much of which was not public while the agency deliberated. The minutes of the October 1993 panel meeting record that "Key data that were needed to support reclassification were contained in PMAs and were not publicly available for use." Because of this constraint, the FDA ended up "waiting ten years to make this [reclassification] announcement" (Rogers, 2023). The reclassification covered only *daily-wear* lenses. Extended-wear lenses stayed in Class III, and they provide one of the intuitive control groups in our supporting specifications. In short, the timing reflected the point at which the evidence was mature enough, not a discrete technology shock.

Vascular Graft Prosthesis

FR citation: 64 FR 12774 **Broad area:** Cardiovascular **Year Announced:** 1999

Rationale. This reclassification was one of the steps taken after the Food and Drug Administration Modernization Act (FDAMA) of 1997 to settle the status of legacy device types, marketed before 1976, that had remained in Class III despite a long clinical record. FDAMA §515(i) required the agency to resolve the classification of preamendment Class III devices for which it had neither issued final regulations nor required PMAs. A statutory deadline and procedural requirements therefore set the timing, and the agency relied on decades of safety and effectiveness experience rather than on recent technological advances.

Tacrolimus Test System

FR citation: 67 FR 58329 **Broad area:** Clinical chemistry **Year Announced:** 2002

Rationale. The FDA reclassified the tacrolimus assay together with the analogous cyclosporine assay, whose reclassification a petition had initiated, and aligned the two actions. The agency judged that a single special-controls guidance, "Cyclosporine and Tacrolimus Assays," was enough to provide reasonable assurance of safety and effectiveness for both, given that the assays work in similar ways and play similar roles in clinical management. Hence, the timing came from harmonization with a related assay, not from a breakthrough specific to tacrolimus.

Vascular Embolization Device

FR citation: 69 FR 77899 **Broad area:** Cardiovascular **Year Announced:** 2002

Rationale. According to the panel records, the clinical information available decades earlier was not sufficient to establish performance standards for this device. Only after long experience with the device could the evidence support special controls and reliable performance standards. The reclassification was thus the end point of a slow accumulation of evidence, not a response to a sudden technological shock.

B.2 Data and Variable Construction

B.2.1 Data sources and panel construction

We assemble multiple datasets. Those on reclassification events, FDA submissions, adverse events and matching of devices to patents follow a similar procedure as in Rogers (2023), extending them to a global context when necessary (e.g. adverse events). Remaining datasets include matching devices to international trade (US to/from each other country using HS10 codes, and bilateral country-pair using HS6 codes), inventor histories and migration, countries' GDP and population.

Reclassification events. We compile a record of FDA reclassification events from the Code of Federal Regulations and Federal Register notices. The raw Class III-to-II event file contains 31 device-type rows. The main matched-control patent design uses 29 of these. We drop two rows: LTT, which the source classifies as a Class I-to-II correction rather than a Class III-to-II down-regulation; and HRC (contact-lens care products), which the source links to a related lens-care code (HRD) and which lacks sufficient pre-period activity and mixes submission pathways to support a distinct treated-device series. Of the 29 main-sample device types, twelve are FDA-initiated; the remaining seventeen are petition-driven or otherwise non-FDA-initiated. Appendix Table A.2 enumerates all 29 treated devices by initiation type, with device names, announcement years, and Federal Register citations.

FDA submissions. The FDA-submission appendix figures use the public 510(k) and PMA submission records (U.S. Food and Drug Administration, 2025a,b). We assign each submission to an FDA device type and submission year using the receipt date. For PMA records, we keep original applications, panel-track supplements, and design/component/specification/-material changes; this excludes routine supplements that are less directly tied to product innovation. We classify submissions as U.S. or non-U.S. using applicant address fields: 510(k) records report country codes directly, and PMA records are classified as U.S. when a state is reported, with non-U.S. PMA applicants geocoded from address fields where needed.

The submitted-device series deduplicates device names within device types using a fuzzy-match threshold of 75 and keeps the first observed submission date for each device-type-device-name pair. We then aggregate to annual device-type counts separately for U.S. and non-U.S. submissions.

Patent data. We use the Dimensions database (Hook et al., 2018) to identify medical-device-related patent families through a keyword-matching procedure. For each FDA-defined device type, the local recovery scripts extract terms from the regulation-number description and the device-name description, remove stopwords, punctuation, and duplicates, and assemble a single keyword string. For example, the device type "Implantable Pacemaker Pulse Generator" combined with the device name "Leadless Pacemaker" yields the keyword string "implantable pacemaker pulse generator leadless." We query Dimensions for granted patents whose titles or abstracts contain all keywords in the string, and repeat the procedure for

every device type. Patents that match the keyword sets of multiple FDA device types appear under each matched device type.

Inventor locations are geocoded from patent filing records to 61 countries using ISO-2 country codes. When a patent lists inventors in multiple countries, the patent is assigned fractionally to each country in proportion to its inventor share: a patent with two German inventors and one Japanese inventor assigns two-thirds credit to Germany and one-third to Japan. Assignee locations are coded separately, enabling comparisons between inventor-fractional (R&D geography) and assignee-fractional (ownership geography) definitions.

Adverse event data. The FDA's Manufacturer and User Facility Device Experience (MAUDE) database contains approximately 9.7 million device-related reports spanning 1988–2024 (CDRH, 2023; Ensign and Cohen, 2017). We construct device-by-year series by linking MAUDE's master event records, patient outcome records, and annual device identification records (the device files begin in 1997) using each report's unique MDR identifier. We use MAUDE's standard outcome codes for disability, death, hospitalization, and other serious injury. The primary outcome is the annual count of foreign serious-outcome reports for each device; a secondary outcome counts all foreign reports regardless of severity.

Global HS-6 trade data and exposure. Global product-market outcomes and baseline trade-exposure measures use BACI HS-6 trade flows (Gaulier and Zignago, 2010). We build a hand-audited FDA-device-type-to-HS-6 concordance (described below). For each treated and control device, the concordance assigns one or more HS-6 codes and aggregates BACI bilateral flows to the device-by-country-by-year level for the trade-exposure interactions and to the device-by-country-pair-by-year level for the bilateral route-level analyses.

The HS-6 concordance is built in two stages. An initial classification by a large language model (`gpt-4o-mini`) assigns each FDA device type—from its regulation-number, device-name, and part descriptions—to its most probable HS-6 heading. We then hand-audit the treated device types: for each treated device we review the candidate heading against the HS-6 heading text, U.S. Customs rulings, and the HS-2022 / HS-1996 cross-walk, and assign the device to one of two categories: *specific*, meaning the HS-6 heading uniquely identifies the device, or *broad*, meaning the heading covers the device together with closely related products at the same heading. Three devices receive a specific mapping (contact lenses HQD and LPL to HS 9001.30, magnetic resonance imaging LNH to HS 9018.13); the remaining 26 are broad. Control device types are not hand-audited: each retains its initial automated classification. We additionally require that a control share no HS-6 heading with any treated device. Appendix Table A.4 reports the full mapping. HS-1996 codes are used for the BACI-based bilateral and country-level trade outcomes; HS-2022 codes are used for HTSUS-aligned outcomes (see the U.S.-specific HS-10 panel below).

For the stacked patent analysis, baseline trade exposure enters as country×cohort-level measures. *Import share from U.S.* is the fraction of country c's device-category imports originating in the United States, computed over a fixed 1996–2000 window, the earliest five years the bilateral trade data cover.[48] *Export share to U.S.* is the fraction of country c's

[48]This fixed window predates the reclassification for the 16 events after 2000 and postdates it for those before 1996, since BACI coverage begins only in 1996. For the 16 after-2000 events, which we take as

device-category exports destined for the U.S. *Trade connectedness* is the unweighted average of the import-share and export-share measures.

U.S. HS-10 import and export data. The U.S.-specific import and export event studies use Peter Schott's processing of official U.S. trade statistics at the HS-10 level (Schott, 2008; Pierce and Schott, 2012), which reports annual trade values by partner country and HS-10 product code from 1989 through 2024. We construct the FDA-device-type-to-HS-10 concordance in two steps. First, for each treated device we hand-audit a candidate HS-10 by reviewing the HTSUS heading text against the FDA device description, reassigning the candidate to a more device-specific code where appropriate; eleven of the 29 mappings are reassigned in this step (for example, the antimicrobial-susceptibility analyzer LON is moved from a generic analytical-instrument heading to HS 9027.80). Appendix Table A.13 reports the final mapping, mapping-confidence rating, and audit rationale for each device. Second, we expand each anchor HS-10 to a *family* that aggregates codes whose digit-level definitions change over the sample period and that share the device's six-digit heading. This family construction preserves a balanced 1989–2024 trade series even when individual HS-10 sublines were created, retired, or renumbered.

We then audit the family for each of the 29 treated devices to verify that the constituent HS-10 codes fall within the same broad HS-2 chapter as the initially matched HS-10 to each device (the device's HS-10 anchor). Twenty-eight devices pass cleanly, with families of 1 to 24 HS-10 codes consistent with the device's clinical function. One device fails: the salivary-system stimulator (LTF) maps to a family of 1,212 HS-10 codes that spans fifteen unrelated HS-2 chapters covering chemicals, base metals, and electrical equipment. We drop LTF on this basis. After dropping LTF and the six in-vitro-diagnostic devices whose HS-10 anchors are post-2017 reclassification sublines with no continuous historical predecessor, 22 of the 29 treated devices enter the HS-10 event studies. The HS-10 layer provides the finer, device-resolution robustness reported in Appendix D.3; the broader HS-6 layer above is the primary product-market evidence, providing bilateral-route coverage that HS-10 cannot, at the cost of capturing closely related goods at the same heading.

FDA establishment registration (contractor data). The FDA maintains a public database of establishments registered to manufacture, import, or export medical devices. We use the processed registration/listing file produced in the related DME Innovation replication workflow from archived FDA registration snapshots. That workflow combines registration, owner-operator, listing, device-type, proprietary-name, and establishment-type files for July 2009, June 2010, May 2011, February 2013, May 2020, and 2022. The processed file contains 1,226,205 establishment-product observations, 5,762 device types, 119 countries, and first-observed years spanning 1973–2022. Establishment type codes distinguish contractors (code 1), importers (codes 3 and 12), exporters (code 10), and foreign entrants (code 5, restricted to non-U.S. firms). We aggregate to the device-type×year level by taking the first appearance of each unique owner×product combination.

the primary mechanism sample in Section 6.1, the window is therefore entirely pre-treatment; restricting instead to the 20 events from 1996 onward strengthens the import interaction (0.414), and the export-share interaction stays insignificant across these cuts.

Inventor histories, prior U.S. exposure, and migration. We combine PatentsView inventor identities with PATSTAT person and publication records (United States Patent and Trademark Office, 2025; European Patent Office, 2025). PatentsView supplies longitudinal inventor identifiers, patent filing dates, inventor address locations, inventor sequence numbers, and patent-level application information. PATSTAT Spring 2025 supplies application metadata, person records, inventor-application links, U.S. grant-publication links, and CPC technology profiles. The linkage maps PatentsView U.S. patent grants to PATSTAT U.S. B publications, then matches inventors within the same patent using inventor sequence numbers. The direct PatentsView inventor-to-PATSTAT person links are expanded through PATSTAT's standardized person clusters to recover additional person identifiers for the same inventor.

The time-varying residence panel first assigns each inventor-year a dominant country from PatentsView address evidence. The dominant country is the country with the largest number of unique patents for that inventor-year, provided its patent share is at least 0.7; otherwise the inventor-year country is set to missing. PATSTAT supplements add conservative pre-U.S. and post-U.S. residence evidence for linked inventors. These supplements require nonmissing PATSTAT person-country information, CPC overlap with the inventor's PatentsView technology profile, and a stricter 0.8 dominant-country threshold. The final supplemented residence panel contains 13.4 million inventor-year observations for 4.2 million inventors across 207 countries. It is used to classify whether inventors on non-U.S. teams had prior U.S. address histories.

Observed moves are defined from this time-varying residence panel. A move from country A to country B occurs in year t when the inventor's dominant country changes from A in year $t-1$ to B in year t, the countries differ, and the destination-country assignment either persists into year $t+1$ or is supported by at least three destination-country patents over years t and $t+1$. The migration falsification keeps U.S.-involved moves, separately tracking U.S.-to-non-U.S. and non-U.S.-to-U.S. movements. The final supplemented move file contains 122,613 moves by 100,440 inventors, with 60,637 U.S.-to-non-U.S. moves and 61,976 non-U.S.-to-U.S. moves.

Inventor entry decomposition. The entry decomposition asks which types of non-U.S. inventors produce the patenting response. For this purpose, we use expanded PATSTAT-based inventor-product histories to improve coverage of device patenting outside the United States.

Inventors are classified relative to each treated device-type cohort. *New-to-device* inventors have no prior device patenting in the relevant pre-period. *Incumbents* have prior patenting in the focal FDA device type. *Experienced pivots* have prior patents in other device types but no prior patents in the focal treated code.

GDP data. Country-level GDP enters the interaction specifications as a control. We use per-capita real GDP from the Penn World Table version 11.0 (Feenstra et al., 2015). This is real GDP at constant 2017 national prices divided by population, taken at a fixed 1994 baseline year and entered in logs.

FDA-reliance measure. For the reliance interaction, R_c is a country-by-epoch measure of formal reliance on FDA device decisions, ordinal on $\{0, 1, 2\}$: 0 is no formal reliance, 1 partial or recognized reliance (FDA approval expedites but does not replace local review), and 2 strong reliance (FDA approval sufficient or nearly sufficient for local marketing). We code it from primary sources—WHO regulatory-capacity surveys, IMDRF documentation, and national regulators—with an AI-assisted extraction pass, and human-audit every country-epoch cell against the source documents.[49]

Patent embeddings (innovation direction). To measure the direction of innovation, we embed patent titles and abstracts using `text-embedding-3-large` (OpenAI, 2024), which returns 3,072-dimensional vectors. In robustness we also use a patent-specific embedding model PaECTER Ghosh et al. (2024). Deatils on computation of variables with these data are provided in Appendix Section B.2.2.

B.2.2 Direction of Innovation: Embedding Measures

Each patent is a point in a high-dimensional technology space, in which nearby points describe similar inventions. We summarize a country's patenting in a device and year by the center of mass of these points, its centroid. Two measures capture *level*: whether the foreign centroid is closer to the contemporaneous U.S. centroid, in angle (cosine) or in straight-line distance (Euclidean). Two measures capture *direction*: after deregulation the U.S. innovation moves along a direction u_r in this space, and we ask whether the foreign centroid moves along the same direction (directional alignment) and how far (magnitude alignment).

Denote reclassification event r, p_r the treated device type, T_r its event year, country c, device type p, and year t within the (c, r) cohort. Let $\mu_{cpt} \in \mathbb{R}^K$ denote the country-device-type centroid and μ_{US,p_rt} the U.S. centroid for the treated device type in the same year, where $K = 3{,}072$ is the dimension of the OpenAI `text-embedding-3-large` vectors (OpenAI, 2024) that embed each patent's title and abstract. A centroid is the inventor-fractional weighted mean of these patent-family embeddings:

$$\mu_{cpt} = \frac{\sum_i w_{icpt} e_{icpt}}{\sum_i w_{icpt}},$$

where e_{icpt} is the embedding vector of patent family i in cell (c, p, t) and w_{icpt} is the inventor-country fractional weight assigned to country c. The centroid summarizes the mean location of innovation content in the high-dimensional technology space.

First, *cosine-to-U.S.* is

$$\text{cosine_to_us}_{rcpt} = \frac{\mu_{cpt} \cdot \mu_{US,p_rt}}{\|\mu_{cpt}\| \, \|\mu_{US,p_rt}\|}.$$

Higher values indicate that the composition of a country's innovation is more similar to the contemporaneous U.S. innovation.

[49] Across the 187 non-missing rows, 154 are coded 0, 29 are coded 1, and 4 are coded 2; the ranking is sensible on inspection, with EU member states that maintain independent regulators at the bottom and several Middle Eastern, Latin American, and Southeast Asian regulators at the top.

Second, *negative Euclidean distance to U.S.* is

$$\text{neg_euclidean_to_us}_{rcpt} = -\left\|\mu_{cpt} - \mu_{US,p_r t}\right\|.$$

An increase indicates convergence in absolute embedding distance (note the negative sign). The measure is the straight-line gap between the two centers of mass, so, unlike cosine, it depends on absolute closeness and not only on angle.

Third, we construct a cohort-specific U.S. post-deregulation direction:

$$\Delta_r^{US} = \bar{\mu}_{US,p_r}^{post}(T_r) - \bar{\mu}_{US,p_r}^{pre}(T_r), \qquad u_r = \frac{\Delta_r^{US}}{\|\Delta_r^{US}\|},$$

where the pre-period average uses the 15 years before treatment and the post-period average uses the 15 years beginning in the treatment year, with years weighted by total inventor weight. The raw vector Δ_r^{US} records how far and in which direction the U.S. innovation moved after the event. For country-product (c, p), we define movement relative to its own pre-period centroid:

$$d_{rcpt} = \mu_{cpt} - \bar{\mu}_{cp}^{pre}(T_r).$$

The *directional alignment* outcome is

$$\text{align_dir}_{rcpt} = \frac{d_{rcpt}}{\|d_{rcpt}\|} \cdot u_r \in [-1, 1].$$

The measure captures direction only: $+1$ means movement in the same direction as the U.S. shift, 0 orthogonal movement, and -1 opposite movement. Figure A.1 illustrates the construction.

Fourth, *magnitude alignment* keeps the same U.S. direction but does not normalize country movement:

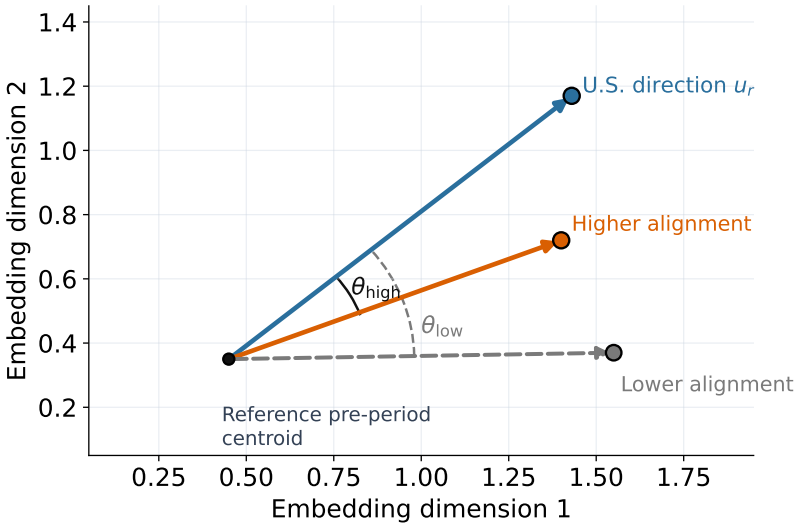
$$\text{align_mag}_{rcpt} = d_{rcpt} \cdot u_r.$$

B.2.3 Effect-size benchmarks for direction outcomes

we benchmark dynamic magnitudes in years of pre-period drift. Let $\Delta = \overline{g}_{post} - g_{-1}$ denote the change in the treated-minus-control gap from the reference pre-period ($k = -1$) to the average post period. Within each matched-control unit (country $\times$ device type $\times$ event), we compute adjacent-year first differences of the outcome over the pre-window $k \in [-8, -1]$ and average them across all such observations to obtain $\overline{\delta}_{pre}$, the typical annual change in the outcome among untreated units absent treatment. The benchmark is

$$\text{years ahead} = \frac{\Delta}{\overline{\delta}_{pre}}.$$

The interpretation is a level-gap statement, not an elapsed-time statement: by the post-window, treated cells sit where untreated cells would arrive only after that many additional years of pre-period organic drift in the same outcome. The denominator is the matched-control-only pre-period drift rate—the pace at which untreated cells converge on the outcome



Appendix Figure A.1: Directional Alignment: An Illustration

Note: Illustrative two-dimensional depiction of the 3,072-dimensional embedding space. Each foreign country-device centroid moves from its pre-period position (the origin of the arrows) in some direction. Directional alignment is the cosine of the angle θ between that movement, d_{rcpt}, and the U.S. post-deregulation shift direction, u_r: $\text{align_dir}_{rcpt} = \frac{d_{rcpt}}{\|d_{rcpt}\|} \cdot u_r = \cos\theta \in [-1, 1]$. The orange vector, at a small angle θ_{high} to u_r, has higher alignment than the gray vector, at a large angle θ_{low}; $+1$ is exact alignment, 0 orthogonal, and -1 opposite. The measure depends only on direction, not on how far the centroid moves.

absent treatment—which the pre-trend tests support as the treated units' counterfactual pace under parallel trends.

For the pooled Class III to II direction estimates in Figure 5, the cosine-to-U.S. effect equals roughly 4.1 years ahead of pre-period organic convergence; the directional-alignment effect equals roughly 11.3 years ahead of pre-period drift along the U.S. trajectory. The same translation gives roughly 5.1 years for negative Euclidean distance to the U.S. and roughly 14 years for magnitude alignment along the U.S. direction; the magnitude-alignment number is large because the pre-period drift rate that anchors its denominator is small. In standard-deviation terms these effects are roughly 0.2 (cosine), 0.8 (directional alignment), 0.2 (negative Euclidean distance), and 0.8 (magnitude alignment) standard deviations of the respective outcomes; unlike the years figure, the standardized magnitude does not depend on the drift denominator, so magnitude alignment is no longer an outlier.

B.2.4 Novelty of induced foreign patents

The centroid outcomes in Appendix B.2.2 summarize *where* the mean of foreign innovation sits relative to the U.S. frontier, but a centroid is a mean: two clouds of points can share a center of mass without any of one cloud's points landing on the other's. They therefore cannot tell a foreign response that *extends* existing U.S. work from one that *replicates* it. We complement them with a nearest-neighbor novelty measure—a minimum-distance rather than a mean statistic—that asks, family by family, whether induced foreign patents duplicate prior U.S. inventions.

Measure. Let $e_i \in \mathbb{R}^K$ (K = 3,072) denote the unit-normalized embedding of foreign patent family i in treated device type p, filed in year t_i, so that $e_i \cdot e_j$ is the cosine similarity between families i and j. Let $\mathcal{P}_p^{US}(t_i) = \{\, j : \text{U.S.-inventor family in device } p,\ t_j \leq t_i - \tau \,\}$ be the set of prior U.S. art, where the publication lag $\tau = 2$ years approximates the gap between filing and public disclosure, so that i can only be compared against U.S. work plausibly observable at its filing. Excluding i's own family, let $\mathcal{N}_k(i)$ be the k members of $\mathcal{P}_p^{US}(t_i)$ most similar to i. Nearest-neighbor novelty is

$$n_{i,k}^{\exp} = 1 - \frac{1}{k} \sum_{j \in \mathcal{N}_k(i)} e_i \cdot e_j,$$

the distance from i to its k closest prior U.S. patents. Higher values mean the family sits further from existing U.S. work; a family that copies a specific prior U.S. invention scores near zero. We report $k \in \{1, 5, 10\}$ and require at least ten prior U.S. families to score a cell.

Density correction. The reclassification itself thickens $\mathcal{P}_p^{US}$: as U.S. patenting in the treated category surges, the prior-art cloud densifies, and a denser cloud shrinks nearest-neighbor distances by pure sampling, even with no change in underlying content. That densification would bias the novelty measure toward finding false convergence in exactly the treated cells. We remove it by standardizing the reference size: for each family we draw $B = 10$ random subsets of fixed size $N = 100$ from $\mathcal{P}_p^{US}(t_i)$, compute the nearest-neighbor distance within each, and average. The subsampled measure holds the prior-art pool size constant across cohorts and years, so differential densification cannot drive it.

Two reference windows. The measure above uses an *expanding* reference: prior U.S. art up to $t_i - \tau$, which advances with each foreign filing. It tests whether induced patents converge on the *contemporaneous* U.S. frontier. We pair it with a *fixed* reference that freezes the U.S. frontier at the cohort's event year: $\mathcal{P}_p^{US}(T_r - \tau)$, U.S. art in device p filed before reclassification. Because this reference is fixed before the event, no treatment-induced patent can enter it, so the fixed-frontier measure is immune not only to pool *size* (as the subsample already is) but to pool *composition*: it asks whether induced patents advance *beyond* where the field stood at reclassification. Novelty is aggregated to the country-device-type-year panel as the inventor-fractional weighted mean of n_i over non-U.S. families, paralleling the centroid construction, and enters the identical stacked DiD (Equation 1).

Accounting for expanding U.S. patent pool for computing nearest neighbor distances. Table A.8 reports the full set of estimates across neighbor counts and both reference windows; Section 6.1 discusses the headline density-fixed measures and Figure A.12 plots them. The raw-versus-subsampled contrast is the diagnostic for the densification confound. On the expanding window the raw measure is negative (-0.009, $p = 0.09$) while the density-fixed measure is not ($+0.006$, 95% CI $[-0.003, +0.015]$): the reclassification thickens the prior-art pool and a denser pool mechanically shrinks nearest-neighbor distances, which the fixed-size subsample removes. On the fixed frontier—which has no post-event densification to remove because its reference is sealed before the event—the raw and subsampled

measures agree almost exactly (+0.0096 vs. +0.0100). The correction thus does work precisely on the window where a confound is present and nothing on the window where it is absent.

Restricting reference patents to more recent periods. Because the diffusion mechanism concerns the current U.S. frontier rather than the full stock of past designs, we also restrict the expanding reference to the most recent U.S. art—families filed within the three, five, or eight years preceding $t_i - \tau$. The density-fixed estimate stays between +0.005 and +0.006 (s.e. ≈ 0.005) across the three windows, none distinguishable from zero and with clean pre-trends, and the raw measure holds near -0.009; the copying null does not depend on how far back the reference reaches.

Patent citations: a data limitation. A natural alternative novelty measure is forward or backward citations. We use text embeddings instead because citation records are the binding data constraint in this international device sample: only about a quarter of foreign treated-device families carry PATSTAT citations, and coverage is uneven across patent offices and filing years, so a citation-based novelty measure would be defined on a selected and office-dependent subsample and confound genuine novelty with reporting completeness. Patent titles and abstracts, by contrast, exist for essentially every family, so the embedding measure is defined everywhere the outcome is.

C Robustness and Mechanisms

C.1 Sample Composition

This appendix section reports three diagnostics on the matched-control design: covariate balance between treated devices and their matched controls in the year before the anticipation window, robustness of the headline to alternative IPC- and CPC-based definitions of the patent corpus, and leave-one-out robustness of the headline DiD to dropping each treated device type.

C.1.1 Pre-period covariate balance

Appendix Table A.5 reports covariate balance at the cohort level. Each cohort is a (country, treated device type) pair; for each cohort we restrict to the balance year four years before reclassification, one year before the manuscript's three-year anticipation window opens. The Treated column reports the mean across cohorts of the (country,treated device type, balance year) cell; the Matched-controls column reports the mean across cohorts of the within-country control-device type average at the same year. Across the seven device type-level covariates that vary within a cohort, every normalized difference is well below the 0.25 Imbens-Rubin imbalance threshold (range $|\text{ND}| \in [0.002, 0.049]$).

C.1.2 Robustness to the patent-corpus definition

We discipline the patent-to-FDA-device linkage with the IPC- and CPC-restriction robustness reported in Appendix Table A.11. We re-estimate the headline matched-control DiD on two restricted corpora. Column (2) restricts to patent families that carry at least one IPC code in any A61 medical subclass except A61K (pharmaceutical preparations), or G02B/G02C (optical subclasses that classify contact-lens technologies); drug-delivery devices remain in the sample via A61M. Column (3) takes a stricter, device-specific cut: we hand-curate a natural CPC subclass for each treated device using the FDA device description and clinical function, and restrict to the 22 of 29 devices whose matched-patent corpus has that subclass present in at least 40% of patents with CPC metadata. The 22 included devices and their natural CPC subclasses are sutures and bone cement on A61L (GAM, GAO, GAP, GAS, LOD, LPK); orthopedic and vascular implants on A61F (KWZ, MAL, MAX); diagnostic and imaging devices on A61B (KRD, LNH, LNS, MUA, MUE); antibody-based tumor-marker and immunoassay tests on G01N33 (LOJ, LOL, LTJ, LTK, MLM); microbial nucleic-acid amplification on C12Q (MWA); electrotherapy stimulation on A61N (LTF); and wheelchairs on A61G (IMK). The seven excluded devices are HQD and LPL (contact lenses, whose corpus is dominated by lens-care formulations rather than lens optics), LKD (cutaneous CO_2 monitor), LON (antimicrobial-susceptibility analyzer), LQK (home uterine activity monitor), MKW (cyclosporine immunoassay), and MPN (tissue adhesive); each of these returns a heterogeneous patent corpus that spans multiple unrelated CPC subclasses.

C.1.3 Leave-one-out robustness across treated devices

Appendix Table A.9 reports the headline DiD coefficient re-estimated 29 times, each time dropping the cohorts whose treated device equals the row's device. All 29 single-device leave-one-out estimates remain positive and significant at $p < 0.01$, with a range of $[+0.099, +0.151]$. The two highest-influence devices remain LTJ (prostate-specific-antigen test) and MKW (cyclosporine immunoassay), both in-vitro-diagnostic immunoassays, which shift the headline by about two full-sample standard errors when dropped individually. The lower panel tests whether the headline survives dropping these influential devices simultaneously: dropping LTJ and MKW jointly yields $+0.067$ ($p < 0.001$), and additionally dropping the extracorporeal shock-wave lithotripter (LNS) yields $+0.079$ ($p < 0.001$). The headline is positive and significant in every case, and the magnitude falls from $+0.139$ to about $+0.07$–0.08 once the two most influential immunoassay devices are excluded.

C.1.4 FDA-initiated robustness

The headline pools all 29 Class III-to-II reclassifications drawn from the universe of events in the Federal Register. Eleven were petitioned by the device's own manufacturer, who controls both the filing date and the underlying R&D, product-launch, and marketing plans; this generates an anticipation channel the matching design cannot fully absorb when treated and matched-control devices have different pre-event growth profiles. The FDA initiated twelve of the 29 events on its own accord, with no manufacturer petition. We re-estimate every main-text outcome on those twelve FDA-initiated events.

Appendix Table A.15 re-estimates all five margins on the two exogenous subsets, and Appendix Figure A.6 plots the corresponding pooled patenting event studies. On the twelve FDA-initiated events the basic count DiD is 0.076—about 55% of the 0.139 headline, significant at the 1% level—and the joint pre-trend test clears more cleanly than the full sample ($p = 0.629$ versus 0.311), consistent with removing the petition-driven growth mismatch. The count response also remains significant on the broader 18-event non-manufacturer-initiated subset (0.058; Column 2), which drops only the eleven manufacturer-petitioned events. The mechanism margins survive on both exogenous subsets: directional alignment with the U.S. trajectory stays significant at the 1% level (0.085 on the FDA-initiated events, 0.084 on the non-manufacturer subset), new-to-device inventor entry stays significant at the 1% level (0.142 and 0.132 log points), and U.S.-to-non-U.S. export trade remains positive and significant (0.404 and 0.462) with clean pre-event coefficients. The import-share interaction keeps its sign and order of magnitude (0.161 on the FDA-initiated events, 0.175 on the non-manufacturer subset, against 0.294 in the full sample) but loses significance on so few events.

C.2 Direction of innovation

This appendix section presents different robustness checks on the directional alignment outcomes, constructed on the technology space.

Cumulative direction outcomes may absorb gradual lead-market movement that predates the policy into the event-specific baseline. We report two diagnostics that address this concern. First, the annual direction-movement outcome replaces cumulative movement from the pre-period centroid with the one-year change in the country-product centroid:

$$\text{annual_dir}_{rcpt} = (\mu_{cpt} - \mu_{cp,t-1}) \cdot u_r.$$

This outcome asks whether foreign innovation accelerates toward the U.S. post-deregulation direction after reclassification. The corresponding static DiD estimate is 0.0072 (s.e. 0.0041). A cosine version that normalizes the one-year country movement vector produces a similar positive estimate of 0.0101 (s.e. 0.0057).

Second, the wrong-vector placebo keeps the real event timing and matched cohort sample but replaces the focal direction u_r with randomly chosen U.S. direction vectors from non-focal device types observed in the same event year. This placebo tests whether the direction result reflects generic movement in embedding space rather than movement toward the treated U.S. technological path. Across 50 placebo draws, the mean static DiD estimate is -0.0009, the 95th percentile is 0.0100, and none of the placebo estimates exceeds the true focal-direction estimate of 0.110.

Embedding-model robustness. The main analysis represents patent text with general-purpose 3,072-dimensional embeddings. To confirm that the directional reorientation is not an artifact of a general-purpose model, we re-embed all patent titles and abstracts with PaECTER (Ghosh et al., 2024), a sentence-embedding model pretrained specifically on patent text, and re-run the identical pipeline—the same matched cohorts, 64-country

sample,[50] anticipation window, and specification. The directional-alignment DiD is 0.118 (s.e. 0.014), against 0.110 (s.e. 0.012) under the general-purpose embeddings.

Fixed pre-period anchor. The directional-alignment outcome measures each foreign country-device centroid's movement away from its own pre-period position, which the baseline forms as the mean over the full pre-event window. To confirm that the result does not depend on including the years just before the event in this reference—which mechanically dampens the estimated pre-event leads—we re-form the anchor over an earlier fixed window (relative years -10 to -5) and estimate on relative years ≥ -4, so the four years before the event enter as freely estimated leads. On the same matched cohorts and specification, the directional-alignment DiD is 0.097 (s.e. 0.012), against 0.110 (s.e. 0.012) in the baseline, and the joint test of the pre-event leads is far from significant ($F = 0.39$, $p = 0.76$). Fixing the reference discards roughly one-fifth of the estimation sample but leaves the reorientation result essentially unchanged.

C.3 Mechanism Discrimination Tests

This appendix addresses two alternative mechanisms: financial-market integration and patent-disclosure learning. The financial-integration test, reported below, partitions the foreign-inventor family pool into assignee × inventor cells; it uses the same matched stacked panel, inventor-fractional geography, log(1+count) outcomes, cohort×device-type and cohort×year fixed effects, and country×treated-device-type clustering as Table 2.

On disclosure learning, a same-device-type foreign-to-U.S. citation test finds foreign citations to U.S. patents in treated device types rising after the event ($\hat{\beta} = 0.008$, s.e. 0.003, $p = 0.002$). We note, however, that PATSTAT records citations for only about a quarter of foreign treated-device families, and that coverage is itself strongly correlated with both import exposure. We are hence cautious when interpreting this result.

C.3.1 Financial-Integration Test: Assignee × Inventor Decomposition

If U.S. capital flowing abroad to fund foreign R&D were the operative channel, the response should concentrate in patents with U.S. assignees and foreign inventors. We split the foreign-inventor patent family pool into three mutually exclusive cells based on the assignee-country composition of each family: (A) U.S. assignee only with at least one foreign inventor; (B) foreign assignee only with at least one foreign inventor; (C) mixed (both U.S. and foreign assignees) with at least one foreign inventor. We stack the three cells into a single panel with cell-specific DiD interactions and a globally unique cohort indexing, then estimate cell-specific DiD coefficients in one regression.

Table A.12 reports the result. Most of the result comes from B foreign assignee-foreign inventor ($\hat{\beta} = 0.132$, s.e. 0.019, $p < 10^{-11}$), while the effect in cases A and C is much smaller ($\hat{\beta} \approx 0.03$, marginally significant). The Wald test for $\beta_B - \beta_A$ is +0.107 ($z = 6.00$, $p < 10^{-8}$)

[50] The embedding-based direction analyses use the broader set of 64 non-U.S. countries with patent-text data, rather than the 61-country panel of the trade- and GDP-interacted count regressions, which is restricted to the countries carrying the predetermined covariates those interactions require.

and grows to +0.136 ($z = 6.49$, $p < 10^{-10}$) once the trade-connectedness and R_c regulatory-channel controls are added. Based on this result, we argue that financial integration is not the main driver of the foreign response.

C.3.2 Capital, Intermediate, and Final-Goods Decomposition

The trade-heterogeneity result shows that the response is stronger with greater U.S. import exposure but does not separate within types of imports. We open up imports based on BEC end-use classification: capital goods, intermediate goods and final goods. Each of them would, respectively, be more likely to be associated with different mechanisms: capital-equipment compatibility (capital), product-embodied knowledge transmission (intermediates), and demonstration-and-imitation (final). We construct country-level shares of each end-use category from BACI HS-96 trade flows over the 1996–2000 baseline window via the UN BEC Rev. 4 concordance: capital goods (BEC 41, 521); intermediate inputs (BEC 111, 121, 21, 22, 31, 32, 322, 42, 53); final consumption goods (BEC 112, 122, 51, 522, 61, 62, 63, 12).

Table A.16 reports six specifications. Column (1) reproduces the product-level import-share interaction ($\hat{\beta}_{\text{ImpShare}} = +0.324$, $p < 0.01$). Columns (2)–(4) enter each country-level share alone: the capital-share interaction is negative (-0.671, $p < 0.01$), while the intermediate-share ($+0.303$, $p < 0.05$) and final-consumption-share ($+0.558$, $p < 0.05$) interactions are both positive. Column (5) enters all three shares jointly; with the budget shares competing for the same variation, none is individually significant—the capital-share interaction stays negative (-0.744), the final-consumption interaction stays positive ($+0.411$), and the intermediate-share interaction turns slightly negative (-0.099) (standard errors ≈ 0.47). Column (6) replaces shares with log(1 + level) of imports in each end-use category: the capital-good interaction is not significant (-0.053, $p = 0.23$), the final-consumption interaction is positive and significant ($+0.076$, $p = 0.02$), and the intermediate-input interaction is positive but small and imprecise ($+0.027$, $p = 0.36$).

D Additional Outcomes and Trade

D.1 Migration of inventors

In extensive-margin linear probability models, we find that the probability of any move to the U.S. does not increase after deregulation (first moves: 0.004, s.e. 0.026; all moves: -0.010, s.e. 0.037). U.S.-to-non-U.S. first moves are also not statistically significant and exhibit problematic pre-trends. Hence, while this test is limited, we do not find clear evidence of cross-border migration causing the innovation response.

D.2 Outsourcing and Safety Outcomes

The product-market response may carry offsetting consequences. A lower U.S. compliance threshold can shift manufacturing toward outside contractors, and a less stringent classification can allow more defective products to reach the market. We test both directly.

Outsourcing. Contractor entry rises sharply after reclassification. The static stacked DiD estimate for the probability that any new manufacturing contractor registers with the FDA in a treated device category is 0.170 (95% CI: [0.070, 0.213]) under matched controls (Panel A of Table A.17). The pre-period base rate is essentially zero (0.002), so an event that almost never occurred before reclassification becomes a rather common post-period occurrence in treated categories.

Safety. Lower premarket evidence requirements raise the concern that more lightly reviewed products generate more adverse events. We do not find evidence for this case in the data. Static stacked DiD estimates for serious adverse events (-0.973, 95% CI $[-7.190, 0.752]$) and total adverse-event counts (-11.054, $[-43.545, 9.043]$) attributed to foreign-source manufacturers are pointwise negative and imprecise (Panel B of Table A.17).

D.3 Device-Resolution (HS-10) Trade

This appendix reports the device-resolution (HS-10) version of the U.S.-linked trade result summarized in Section 6.3.

Figure A.10 plots HS-10 event studies for U.S.-linked trade in treated device types: foreign exports to the U.S. (left panel) and foreign imports of U.S. goods (right panel). The HS-10 panel uses the hand-audited concordance from Appendix Table A.13, which maps each device to an anchor HS-10 and aggregates HS-10 codes whose digit-level definitions changed across the sample period so that the device-level trade series is balanced from 1989 to 2024. We keep 22 of the 29 treated device types in the HS-10 event studies. We drop the salivary-system stimulator (LTF, whose underlying HS-10 family aggregates products from fifteen unrelated HS-2 chapters and contaminates the device-specific signal) and six in-vitro-diagnostic devices whose HS-10 anchors lack a continuous historical predecessor. Table 4 (Panel A) reports static Conley–Taber stacked DiD estimates from device-type×year regressions with cohort×device-type and cohort×year fixed effects, using matched controls. The static import effect is 0.195 (95% CI: $[-0.158, +0.486]$); the static export effect is 0.247 ($[-0.046, +0.499]$).[51]

The event-study trajectories rise steadily over two decades: imports climb from near zero at event time 0 to +0.5–0.6 by event time 16–20, and exports from near zero to +0.4–0.5 over the same horizon. Estimating a long-run coefficient that absorbs the event-time-$\in [0, 6]$ window and identifies the average effect at event time ≥ 7 years confirms this pattern: the long-run import coefficient is +0.329 ($[+0.022, +0.604]$) and the long-run export coefficient is +0.336 ($[+0.073, +0.575]$). Both exceed their static counterparts by 35–70%, and the long-run export interval excludes zero (Table 4, Panel A). The product-market response continues to build for at least two decades after the regulatory shock.

We note that HTSUS revisions over the sample period created, retired, and renumbered HS-10 sublines; the hand-audited concordance described above and in Appendix Table A.13 addresses this by aggregating each device's anchor HS-10 with codes whose digit-level def-

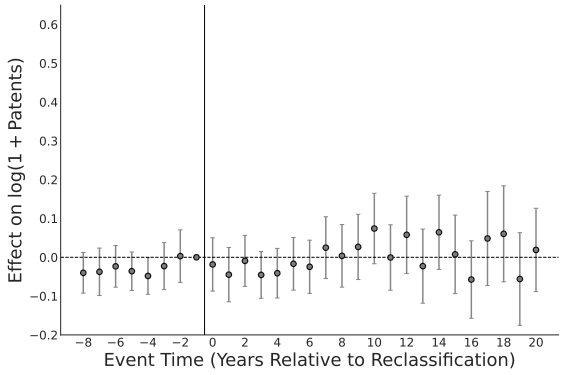
[51] We decompose this value response into physical-quantity and unit-value (price) components in Appendix Table A.14, using the bilateral HS-6 routes, where a physical quantity is reported for 97–99% of trade value. On the intensive margin the price component is a tight null and the response comes from the quantity side.

initions changed but whose six-digit heading is shared (this follows closely what is done in Pierce and Schott (2012)).

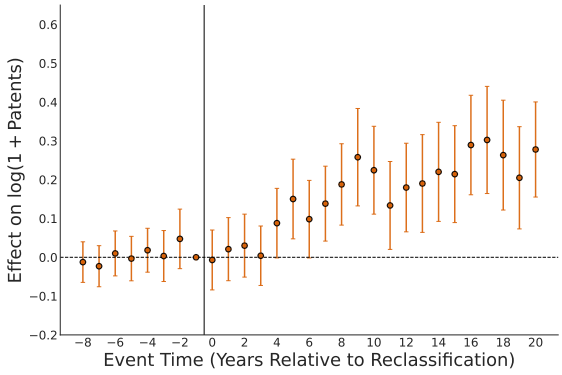
E Appendix Figures and Tables

Appendix Figure A.2: Ideas Margin: Binned Event Studies by Import Exposure

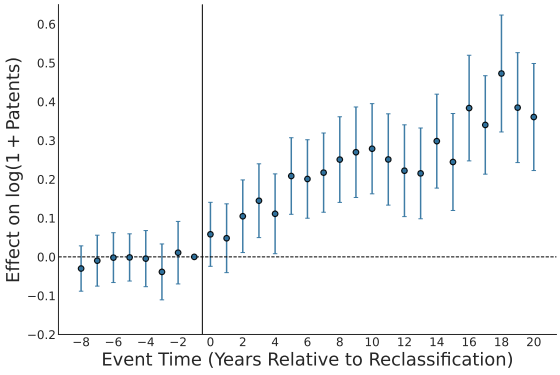
(a) Import Low



(b) Import Mid



(c) Import High

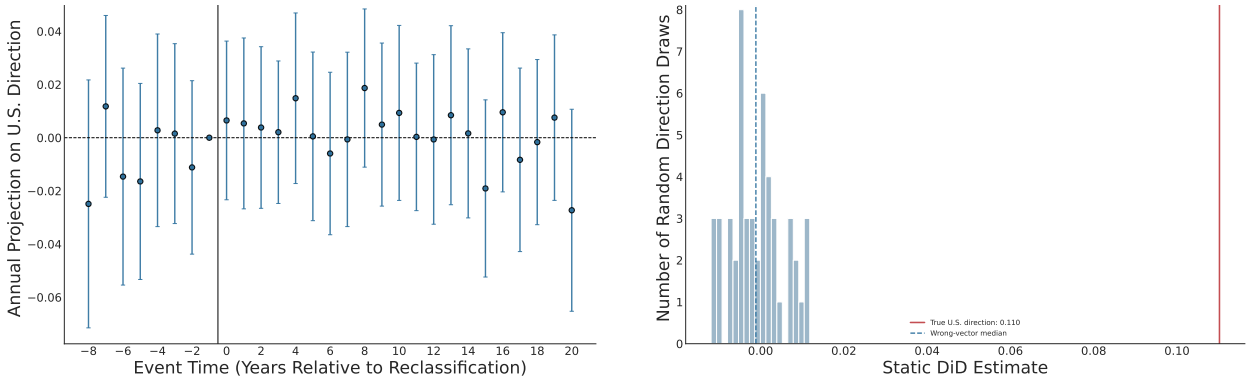


Note: Non-parametric, import-tercile companion to Figure 4. Each panel reports the pooled stacked DiD event study of Figure 3 restricted to a tercile of pre-period import share from the U.S., on the value-weighted $\log(1 + \text{patent count})$ outcome. Pre-period import share is the country-by-treated-device fraction of treated device-type imports originating in the United States, computed over a fixed 1996–2000 window under the HS-6 concordance (Online Appendix Table A.4). Terciles are constructed at the country-by-treated-device level. The static DiD coefficients are +0.025 (low, s.e. 0.021), +0.148 (mid, s.e. 0.031), and +0.240 (high, s.e. 0.033); the long-run DiD coefficients at event time ≥ 7 years are +0.044, +0.210, and +0.301 respectively. The response concentrates in the high-import-exposure tercile and amplifies over the long run.

Appendix Figure A.3: Direction Robustness Diagnostics

(a) Annual movement toward the U.S. direction

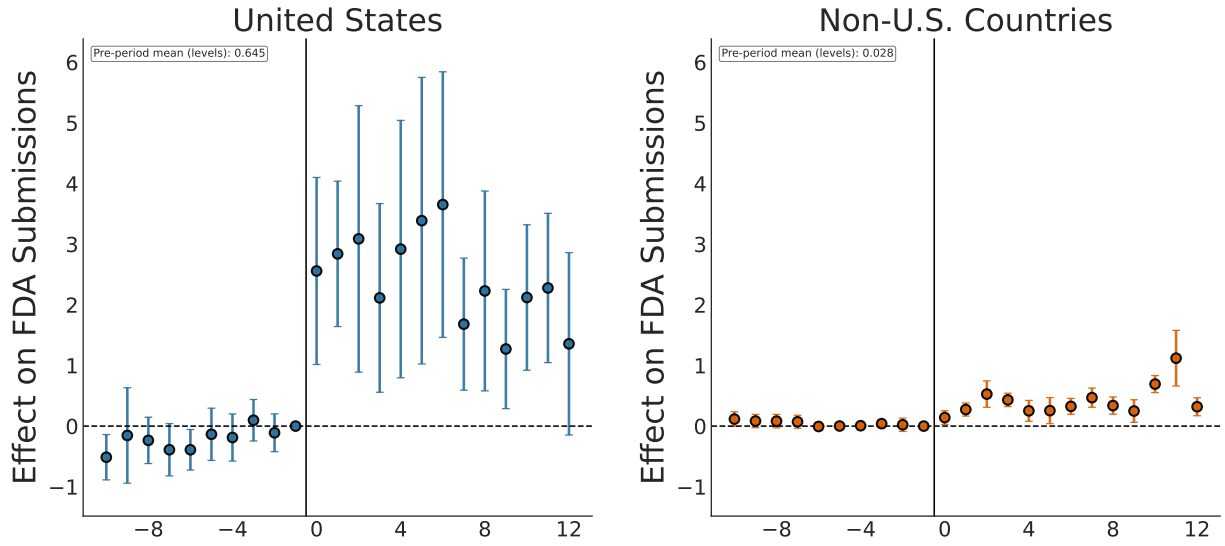
(b) Wrong-vector placebo for directional alignment



Note: Two robustness diagnostics for the directional-alignment result in Section 6.1. Panel (a) plots the one-year change in the foreign country-product patent-text centroid projected onto the treated U.S. pre/post direction, asking whether foreign innovation accelerates toward the U.S. direction after reclassification rather than only whether cumulative movement rises. Panel (b) plots a placebo distribution that keeps the real event timing and matched sample but replaces each treated U.S. direction with a randomly chosen non-focal U.S. device-type direction; the vertical red line is the true focal estimate, and the placebo distribution centered near zero never reaches the true estimate across 50 draws.

Appendix Figure A.4: FDA Submissions Event Studies

(a) U.S.-based applicants **(b)** Non-U.S.-based applicants

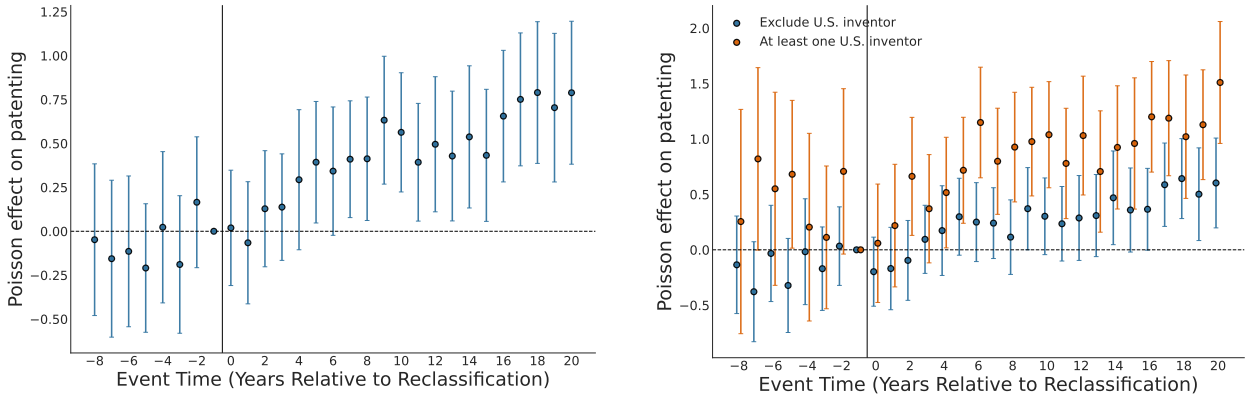


Note: Pooled stacked event studies of FDA medical-device submissions (510(k) and PMA) around Class III-to-II reclassifications, measured relative to matched never-reclassified control device types. Panel (a) covers U.S.-based applicants and Panel (b) non-U.S.-based applicants; the outcome is the annual count of submissions for each treated device type, and the pre-period mean (in levels) is annotated in each panel. Points are Sun–Abraham event-study coefficients with 95% confidence intervals. Both series rise sharply after the event, indicating that reclassification raises submission activity on both sides of the U.S. border.

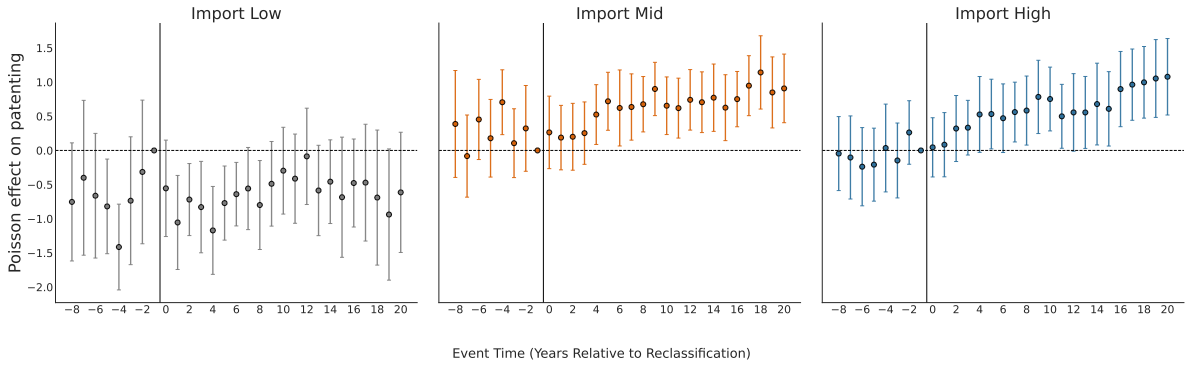
Appendix Figure A.5: Poisson Robustness for the Headline Patenting Effects

(a) Pooled non-U.S. event study

(b) Team split: no-U.S. vs. at-least-one-U.S. inventor



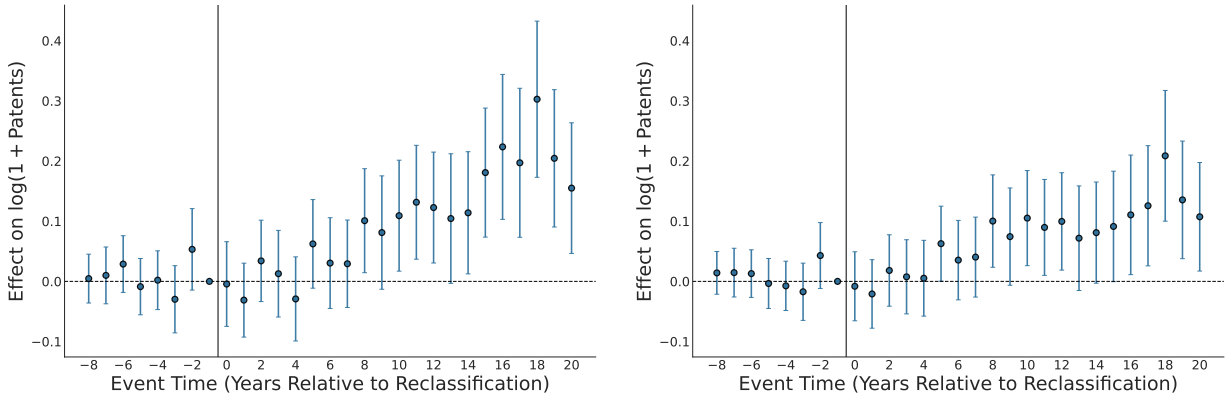
(c) Binned by pre-period U.S. import exposure (terciles)



Note: Poisson counterparts to the OLS event-study figures referenced in the main text. Panel (a) re-estimates the value-weighted pooled non-U.S. baseline. Panel (b) re-estimates the team-composition split, separating patents with no U.S.-based inventor from those with at least one U.S.-based inventor. Panel (c) re-estimates the import-exposure binned event study, splitting countries into terciles of pre-period U.S. import exposure. The Poisson estimates align qualitatively with the OLS baselines reported in the main text, confirming that the headline results are not artifacts of the $\log(1 + \text{count})$ transformation.

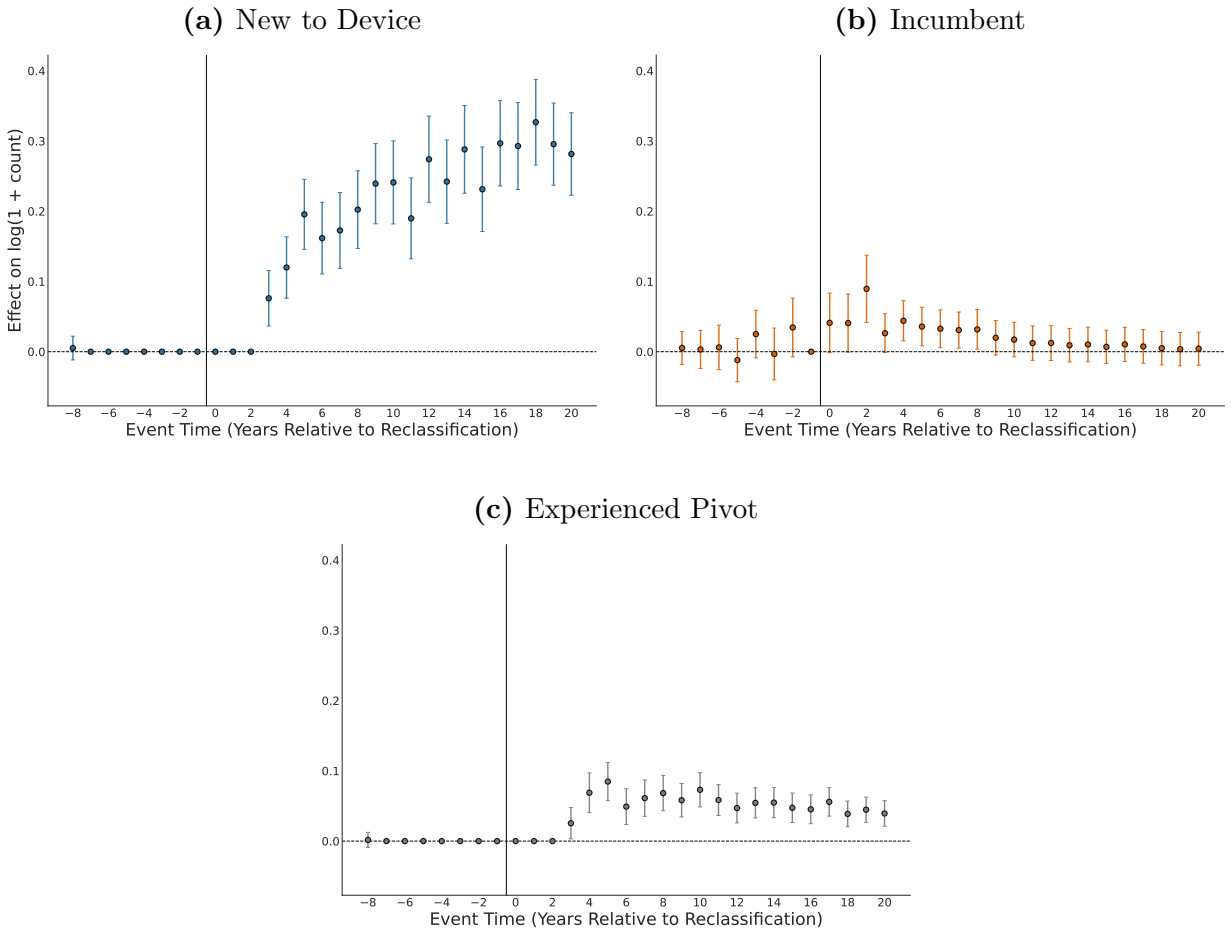
Appendix Figure A.6: Exogeneity Subsets: Pooled Patenting Event Studies

(a) FDA-initiated events (12) **(b)** Non-manufacturer-initiated events (18)



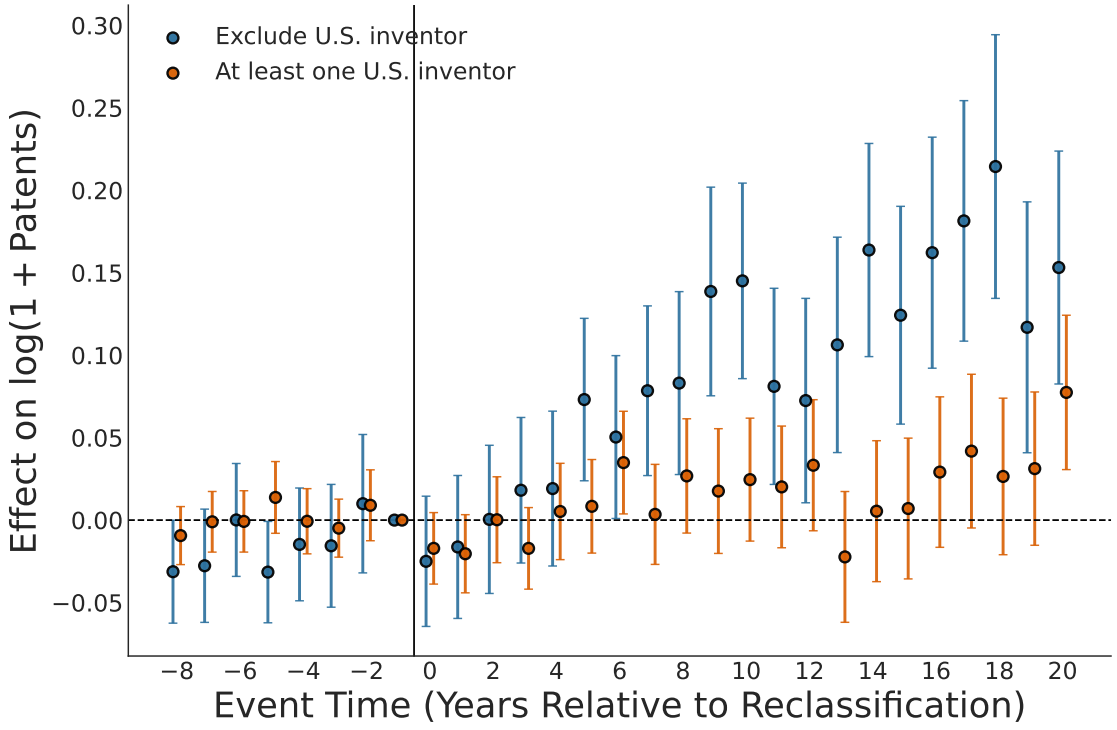
Note: Pooled stacked event studies of foreign value-weighted log(1 + patent count) on the two exogenous subsets that progressively drop reclassifications whose timing a device manufacturer could have influenced, the dynamics behind Panel A of Appendix Table A.15. *Panel (a)* keeps only the twelve reclassifications the FDA initiated on its own accord; *Panel (b)* drops the eleven manufacturer-petitioned events, leaving eighteen. Both use matched controls, a three-year anticipation window, cohort × product-code and cohort × year fixed effects, and standard errors clustered at the treated-device-by-country level; the two panels share a common vertical axis. Pre-event coefficients stay near zero and the response emerges after reclassification on both subsets. The static DiD is 0.076 (FDA-initiated) and 0.058 (non-manufacturer-initiated), each significant at the 1% level, and the joint pre-trend test clears more cleanly than the full sample.

Appendix Figure A.7: People Margin: Country-Level Entry Decomposition by Inventor Type



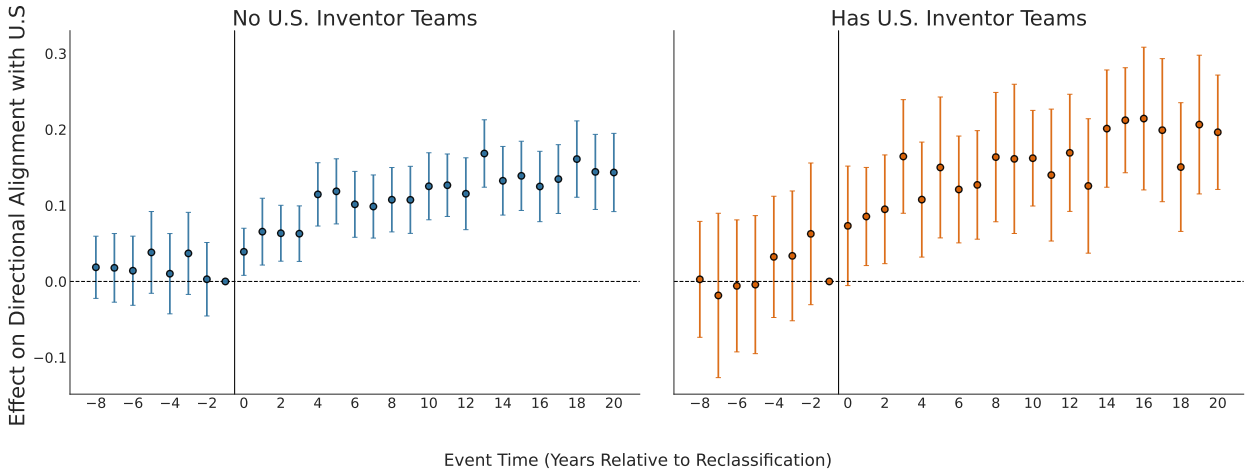
Note: Inventor-count event studies (outcome: $\log(1 + \text{inventor count})$ at the country-by-device-by-event-by-year level) decomposed into three mutually exclusive inventor types defined by pre-period patent history. *New-to-device* inventors have no prior device patenting; *focal-code incumbents* have prior patenting in the treated device type; *experienced pivots* have prior patents in other device types but not in the focal treated code. Type assignments are event-specific and persistent: an inventor classified in event r retains that classification across all years of event r. Event time is anticipation-adjusted, so the raw reclassification year appears at relative year $k = 3$. Panel (b) captures any extensive-margin activity by inventors who were already patenting in the focal device code before deregulation: a focal-code incumbent enters country c's count in year t only when she actively files in device p that year. The post-reclassification rise therefore reflects previously dormant incumbents resuming activity in the focal code. First-ever device experience is captured separately in panel (a). See Table 3 (Panel B) for static DiD estimates by inventor type.

Appendix Figure A.8: Ideas Margin: Pooled Event Study by Team Composition



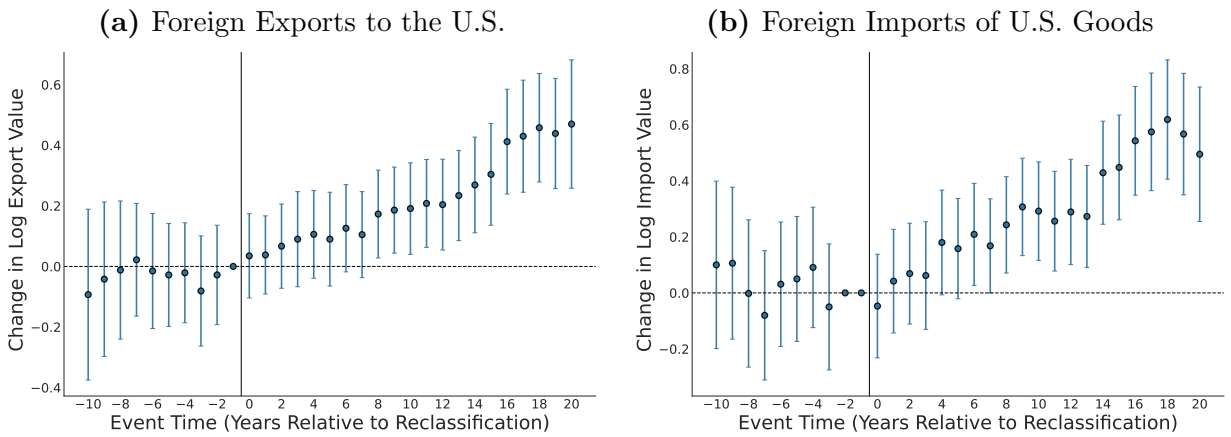
Note: Pooled stacked event studies of value-weighted patenting by team composition. "Exclude U.S." restricts to patent teams with no U.S.-based inventors; "At Least One U.S." requires at least one U.S.-based inventor on the team. Markers are offset horizontally for visibility.

Appendix Figure A.9: Ideas Margin: Direction by Team Composition



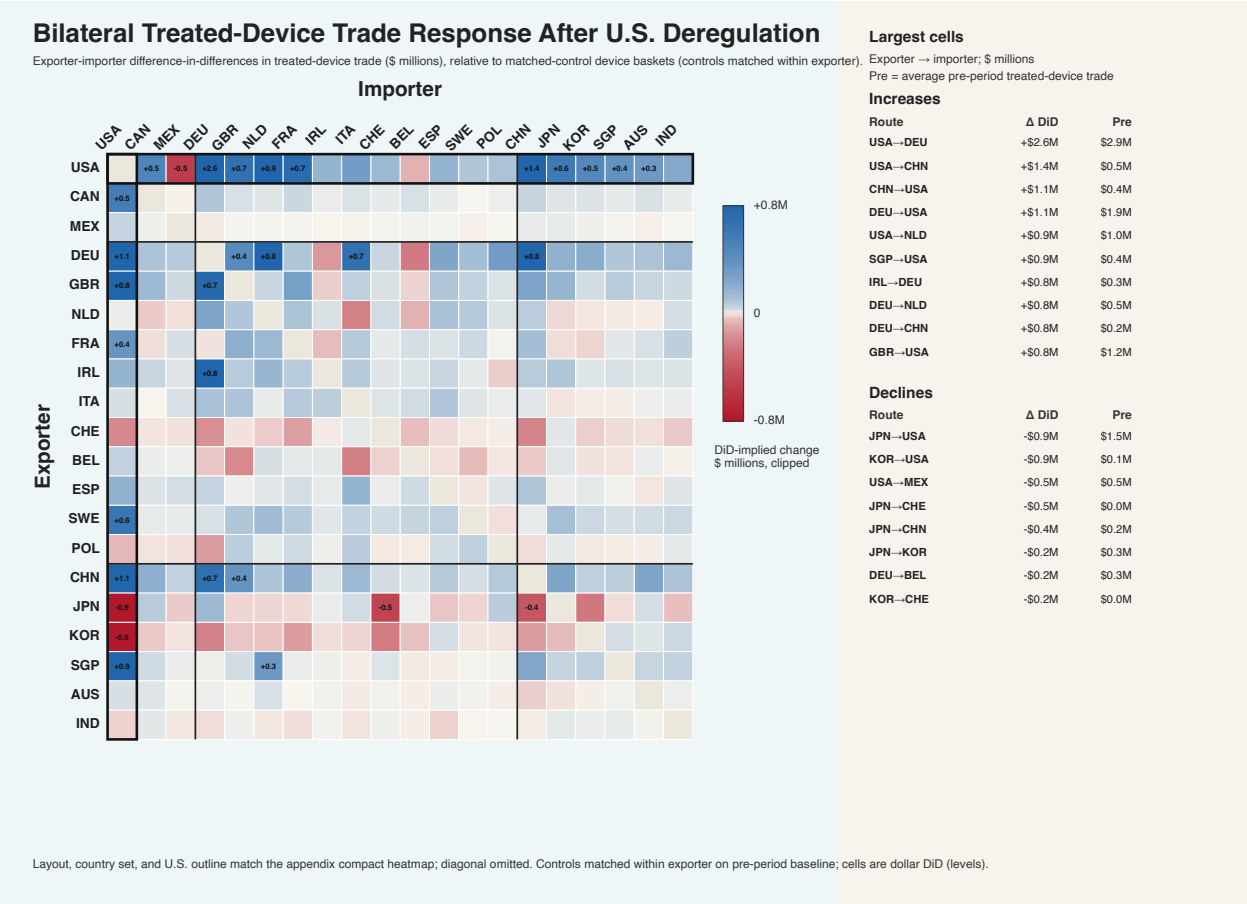
Note: Stacked event studies of directional alignment with the U.S. post-deregulation trajectory, split by team composition: teams with no U.S.-based inventor (left) versus at least one U.S.-based inventor (right). Timing is anticipation-adjusted ($k = 0$ is the formal reclassification year minus three), with $k = -1$ the omitted reference. Whiskers are 95% confidence intervals, clustered at the treated-device-by-country level. Directional alignment is positive and significant for both team types—static DiD 0.092 (s.e. 0.013) for teams with no U.S.-based inventor and 0.124 (s.e. 0.020) for U.S.-linked teams, gradual for the former and more front-loaded for the latter—so the reorientation toward the U.S. frontier is not an artifact of U.S. co-invention.

Appendix Figure A.10: HS-10 U.S.-Linked Trade Event Studies



Note: Conley–Taber stacked event studies of log(HS-10 U.S. trade value) around Class III to II reclassification. HS-10 trade values are from Schott (2008); Pierce and Schott (2012), with HS-10 codes mapped to FDA device types via the HS-10 concordance described in Online Appendix B.2 (full mapping in Online Appendix Table A.13). Each device's HS-10 family is the union of the Schott reclassification connected component of the audited anchor HS-10 and all HS-10 codes whose first six digits match; this preserves a balanced 1989–2024 panel under HTSUS revisions. The left panel shows foreign exports to the U.S. (static DiD = 0.247, 95% Monte Carlo CI $[-0.046, 0.499]$; long-run DiD at event time $\geq 7 = 0.336$, $[+0.073, +0.575]$); the right panel shows foreign imports of U.S. goods (static DiD = 0.195, $[-0.158, +0.486]$; long-run DiD = 0.329, $[+0.022, +0.604]$). The post-period coefficients rise approximately monotonically through event time 20. Conley–Taber 95% Monte Carlo confidence intervals (Conley and Taber, 2011) are constructed from control-group residuals; we use them here because the HS-10 panel collapses to a single time-series per device and leaves roughly 29 effective clusters, making cluster-asymptotic inference unreliable.

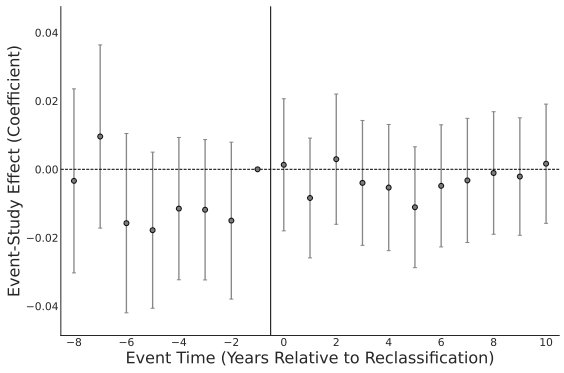
Appendix Figure A.11: Products Margin: Bilateral HS-6 Trade Response



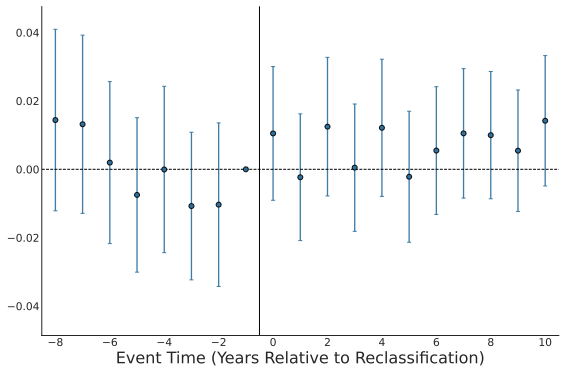
Note: Cells report exporter-importer differences-in-differences in treated-device trade (BACI, HS-6), summed over the post-event window, relative to matched-control device baskets; controls are matched within exporter on pre-period trade. HS-6 codes are mapped to FDA device types via the hand-audited concordance (Online Appendix Table A.4) for the 29 treated devices and the broader candidate-control universe for control device types. The U.S. row and column are outlined; the diagonal is omitted; the side panel reports the largest cells and pre-period treated-device baselines. Treated-device trade expands on U.S.-linked routes and among U.S. exports; among foreign suppliers, individual routes both rise and fall—the largest increases include German and Chinese exports to the United States, and the largest decreases are in Japanese and Korean exports to the United States.

Appendix Figure A.12: Ideas Margin: Novelty

(a) Distance to Contemporaneous U.S. Art



(b) Distance from Frozen Pre-Event Frontier



Note: Pooled stacked event studies of the novelty of induced non-U.S. patenting (matched-control sample). *Panels (a) and (b)* report nearest-neighbor novelty—embedding distance (1−cosine) from each induced foreign family to prior U.S. art in the device type: (a) contemporaneous reference (flat = does not converge onto existing U.S. designs); (b) frontier frozen at reclassification (rising = extends beyond it). Foreign innovation does not copy the contemporaneous frontier and extends past the pre-event frontier—frontier extension, not imitation. Static DiD: contemporaneous novelty +0.006 (s.e. 0.005, ns), frozen +0.010 (0.005). Outcomes from 3,072-dimensional patent-text embeddings (title+abstract); centroids are inventor-fractional weighted means. Points are DiD coefficients vs. $k = -1$; bars 95% CIs; (a) and (b) share a vertical scale. Construction and robustness: Online Appendix B.2.4.

Appendix Table A.2: Class III to II Down-Regulated Device Types: FDA-Initiated and Petitioned

A. FDA-Initiated Down-Regulations (Main Sample, $N = 12$)

Device Type	Device Name	Year Announced	FR Citation
HQD	Lens, Contact (Other Material) – Daily	1994	59 FR 10284
KRD	Device, Vascular, for Promoting Embolization	2002	69 FR 77899
LKD	Monitor, Carbon-Dioxide, Cutaneous	1988	54 FR 27160
LNS	Lithotriptor, Extracorporeal Shock-Wave, Urological	2000	65 FR 48612
LPL	Lenses, Soft Contact, Daily Wear	1994	59 FR 10284
LTF	Stimulator, Salivary System	2015	80 FR 72586
MAL	Graft, Vascular, Synthetic/Biologic Composite	1994	66 FR 18542
MAX	Intervertebral Fusion Device with Bone Graft, Lumbar	2007	72 FR 32172
MLM	Enzyme Immunoassay, Tacrolimus	2002	67 FR 58329
MUA	Bone Sonometer	2008	73 FR 40969
MUE	Full Field Digital, X-Ray Mammography System	2010	75 FR 68203
MWA	System, Nucleic Acid Amplification, Mycobacterium Tuberculosis Complex	2014	79 FR 31027

B. Petition-Driven and Other Down-Regulations ($N = 17$)

Device Type	Device Name	Year Announced	FR Citation
GAM	Suture, Absorbable, Synthetic, Polyglycolic Acid	1991	56 FR 47151
GAO	Suture, Nonabsorbable	1986	65 FR 19836
GAP	Suture, Nonabsorbable, Silk	1993	58 FR 57558
GAS	Suture, Nonabsorbable, Synthetic, Polyester	1991	56 FR 24685
IMK	Wheelchair, Stair Climbing	2014	79 FR 20782
KWZ	Prosthesis, Hip, Constrained, Cemented or Uncemented, Metal/Polymer	2002	67 FR 21173
LNH	System, Nuclear Magnetic Resonance Imaging	1988	53 FR 5078
LOD	Bone Cement	2001	67 FR 46855
LOJ	Kit, Test, Alpha-Fetoprotein for Testicular Cancer	1997	62 FR 66005
LOL	Hepatitis A Test (Antibody and IgM Antibody)	2006	71 FR 6679
LON	System, Test, Automated, Antimicrobial Susceptibility, Short Incubation	2002	68 FR 5827
LPK	Bone Grafting Material, for Dental Bone Repair (TCP granules)	2005	70 FR 21949
LQK	Home Uterine Activity Monitor	2001	66 FR 14076
LTJ	Prostate-Specific Antigen (PSA) for Prostate Cancer Management	1997	62 FR 66005
LTK	Test, Epithelial Ovarian Tumor-Associated Antigen (CA125)	1997	62 FR 66005
MKW	Cyclosporine Test System	2002	67 FR 58329
MPN	Tissue Adhesive for the Topical Approximation of Skin	2008	73 FR 31033

Note: The table enumerates the 29 treated FDA device types used in the main matched-control Class III-to-II patent design. Panel A lists the 12 events initiated by the FDA on its own accord; Panel B lists the 17 petition-driven and other reclassifications. The Year Announced column reports the reclassification year used to anchor event time in the analysis (the manuscript's `year_reclassified`), drawn from the FDA's Code of Federal Regulations and Federal Register record. The FR Citation column gives the corresponding Federal Register volume and page reference. For events that proceeded through both a proposed and a final rule, the year and citation may refer to different milestones in the rulemaking process; we use the announcement / initiation year throughout the empirical work because that timing is the earliest point at which the public could anticipate the reclassification. The FDA-initiated subset is used in the FDA-initiated robustness in Appendix C.1.4.

Appendix Table A.3: Definition Comparison: Inventor, Assignee, and Hybrid Geography

	Estimate	p-value	Pre-mean	N
Panel A. Basic DiD				
Inventor fractional (R&D geography)	0.1392**	<0.001	0.5555	136,649
	(0.0168)			
Assignee fractional (ownership geography)	−0.0004	0.9759	0.1725	113,197
	(0.0125)			
Hybrid fractional	0.1161**	<0.001	0.3510	139,769
	(0.0139)			
Panel B. DiD × Trade Connectedness				
Inventor fractional (R&D geography)	0.3798**	<0.001	0.5555	136,649
	(0.1028)			
Assignee fractional (ownership geography)	0.2020*	0.0250	0.1725	113,197
	(0.0900)			
Hybrid fractional	0.3465**	<0.001	0.3510	139,769
	(0.0922)			
Panel C. DiD × Import Share from U.S.				
Inventor fractional (R&D geography)	0.2943**	0.0015	0.5555	136,649
	(0.0921)			
Assignee fractional (ownership geography)	0.0484	0.4755	0.1725	113,197
	(0.0678)			
Hybrid fractional	0.2505**	0.0012	0.3510	139,769
	(0.0771)			
Panel D. DiD × Export Share to U.S.				
Inventor fractional (R&D geography)	0.0721	0.4368	0.5555	136,649
	(0.0927)			
Assignee fractional (ownership geography)	0.1602^{+}	0.0562	0.1725	113,197
	(0.0838)			
Hybrid fractional	0.0857	0.3500	0.3510	139,769
	(0.0916)			

Note: The headline matched-control stacked DiD re-estimated under three geographic-assignment definitions on the value-weighted log(1+patent count) outcome, each with cohort-by-device and cohort-by-year fixed effects and standard errors (in parentheses) clustered at the treated-device-by-country level. Inventor-fractional assigns each family across countries by inventor residence shares (R&D geography); assignee-fractional by the residence of the patent assignee (ownership geography); hybrid splits the difference. Panel A reports the basic DiD; Panels B–D the trade-connectedness, import-share, and export-share interactions, on the refined HS-6 concordance with a ln(per-capita GDP) control, as in Table 2. The response is large under inventor and hybrid geography and attenuates toward zero under assignee geography, consistent with the spillover operating through where research is performed rather than where the resulting intellectual property is owned. $^{+}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$.

Appendix Table A.4: FDA Device Type to HS-6 Concordance

PC	Device name	HS-6 (2022)	HS-6 (1996)	HS-6 bin description
Specific: HS-6 heading isolates the device				
HQD	Lens, Contact (Other Material) - Daily	900130	900130	Contact lenses
LNH	System, Nuclear Magnetic Resonance Imaging	901813	901813	MRI apparatus
LPL	Lenses, Soft Contact, Daily Wear	900130	900130	Contact lenses
Broad: HS-6 heading covers the device together with closely related goods				
GAM	Suture, Absorbable, Synthetic, Polyglycolic Acid	300610	300610	Sterile surgical suture material
GAO	Suture, Nonabsorbable	300610	300610	Sterile surgical suture material
GAP	Suture, Nonabsorbable, Silk	300610	300610	Sterile surgical suture material
GAS	Suture, Nonabsorbable, Synthetic, Polyester	300610	300610	Sterile surgical suture material
IMK	Wheelchair, Stair Climbing	871390	871390	Carriages for disabled persons
KRD	Device, Vascular, For Promoting Embolization	901890	901890	Medical/surgical instruments, residual
KWZ	Prosthesis, Hip, Constrained, Cemented Or Uncemented, Met.	902131	902111	Artificial joints
LKD	Monitor, Carbon-Dioxide, Cutaneous	901819	901819	Electro-diagnostic apparatus, residual
LNS	Lithotriptor, Extracorporeal Shock-Wave, Urological	901890	901890	Medical/surgical instruments, residual
LOD	Bone Cement	300640	300640	Dental cements / bone reconstruction cements
LON	System, Test, Automated, Antimicrobial Susceptibility, Sh.	902780	902780	Laboratory analyzer
LPK	Bone Grafting Material, For Dental Bone Repair	300640	300640	Dental cements / bone reconstruction cements
LQK	Home Uterine Activity Monitor	901819	901819	Electro-diagnostic apparatus, residual
LTF	Stimulator, Salivary System	901849	901849	Other dental instruments
MAL	Graft, Vascular, Synthetic/Biologic Composite	902139	902130	Artificial parts of the body
MAX	Intervertebral Fusion Device With Bone Graft, Lumbar	902110	902119	Orthopaedic appliances
MPN	Tissue Adhesive For The Topical Approximation Of Skin	300610	300610	Sterile surgical suture material (adhesives)
MUA	Bone Sonometer	901812	901812	Ultrasonic scanning apparatus
MUE	Full Field Digital, System, X-Ray, Mammographic	902214	902214	Medical X-ray apparatus
Broad: HS-6 reagent heading (in vitro diagnostic devices)				
LOJ	Kit, Test, Alpha-Fetoprotein For Testicular Cancer	382219	382200	Diagnostic reagents
LOL	Hepatitis A Test (Antibody And Igm Antibody)	382219	382200	Diagnostic reagents
LTJ	Prostate-Specific Antigen (Psa) For Management Of Prostat.	382219	382200	Diagnostic reagents
LTK	Test, Epithelial Ovarian Tumor-Associated Antigen (Ca125)	382219	382200	Diagnostic reagents
MKW	Cyclosporine	382219	382200	Diagnostic reagents
MLM	Enzyme Immunoassay, Tracrolimus	382219	382200	Diagnostic reagents
MWA	System, Nucleic Acid Amplification, Mycobacterium Tubercu.	382219	382200	Diagnostic reagents

Note: Concordance from FDA product codes to six-digit Harmonized System (HS-6) headings for the trade outcomes in Section 6.3 and Appendix Table 4 Panel B. Each mapping is hand-validated against the HS-6 text, U.S. Customs rulings, and the HS-2022-to-HS-1996 crosswalk (HS-1996 is the BACI vintage). Three headings isolate the device cleanly: contact lenses (HQD, LPL) to HS 9001.30 and MRI (LNH) to HS 9018.13. The rest capture the device with closely related goods—surgical sutures (HS 3006.10), artificial joints (HS 9021.31), medical X-ray apparatus (HS 9022.14), etc.—so those route estimates reflect broader product baskets, not the exact FDA subtype. The seven in-vitro-diagnostic devices map to HS 3822.19 (diagnostic reagents), broader than the FDA test alone; we retain them and report robustness without them in Appendix Table A.11.

Appendix Table A.5: Pre-Period Covariate Balance by Cohort

Covariate	Treated	Matched controls	Diff.	|Norm. diff.|	t-stat
Panel A. PC-level covariates that vary across treated and matched controls within a cohort					
Patent count, 4 years before reclassification	0.646	0.519	0.128	0.049	1.62
log(1 + patent count), 4 years before reclassification	0.166	0.158	0.008	0.017	0.67
Average patent count, 4–8 years before reclassification	0.409	0.357	0.053	0.037	2.40
log(1 + average patent count, 4–8 years before reclassification)	0.173	0.163	0.010	0.024	1.70
Any patents 4 years before reclassification (0/1 indicator)	0.112	0.112	0.001	0.002	0.08
Any patents in the 4–8 year pre-window (0/1 indicator)	0.226	0.231	-0.005	0.013	-0.58
Pre-window years with positive patenting (out of 5)	0.458	0.446	0.011	0.012	0.78
Panel B. Country-level covariates (constant within a cohort by design; reported for sample description)					
Pre-period trade connectedness with U.S.	0.216	0.216	–	–	–
Pre-period imports from U.S., share of total	0.297	0.297	–	–	–
Pre-period exports to U.S., share of total	0.135	0.135	–	–	–
ln(per-capita real GDP) in baseline year	10.142	10.142	–	–	–
U.S. FDA-reliance score	0.447	0.447	–	–	–

Note: Pre-period covariate balance computed at the cohort level. Each cohort is a (country, treated device-type) pair; 891 of the 900 cohorts have an observation in the balance year. For each cohort, the balance year is year of reclassification minus 4, one year before the 3-year anticipation window used in the manuscript event studies. Panel A patent-count covariates use the value-weighted (family-size-weighted), deduped count that is the paper's main outcome. Treated column reports the mean across cohorts of the (country, treated-PC, balance-year) cell; matched-controls column reports the mean across cohorts of the cohort's within-country control-PC average at the same year. Diff. is the paired treated-minus-control difference averaged across cohorts. Normalized difference follows Imbens and Rubin (2015): $(\bar{x}_t - \bar{x}_c)/\sqrt{(s_t^2 + s_c^2)/2}$. Cells flagged with † exceed the 0.25 Imbens-Rubin imbalance threshold. Panel A covariates vary across treated and control within each cohort and constitute the genuine balance test. Panel B covariates are country-level measures (trade exposure, GDP, FDA reliance) that are constant within a cohort by construction; they balance trivially in the cohort design and are reported for sample description.

Appendix Table A.6: Intuitive (Clinically-Matched) Control for Each Treated Device Type

Treated device type	Intuitive control device type	Medical specialty
Cyclosporine immunoassay	Atomic-absorption calcium assay	Clinical Chemistry
Tacrolimus test system	Anticoagulant blood-collection tubes	
Alpha-fetoprotein (AFP) tumor-marker test	Hepatitis A antigens	Immunology & Microbiology
Automated antimicrobial-susceptibility system	Anaerobic transport culture media	
Hepatitis A antibody test	Escherichia coli antigens	
Mycobacterium tuberculosis nucleic-acid amplification test	Fluorescent tuberculosis antiserum	
Ovarian tumor-marker (CA-125) test	Cytomegalovirus DNA reagents	
Prostate-specific antigen (PSA) test	Automated zone reader	
Cutaneous carbon dioxide monitor	Thorpe-tube flowmeter	Anesthesiology
Synthetic/biologic vascular graft	Small-diameter vascular graft prosthesis	Cardiovascular
Vascular embolization device	Pacemaker repair/replacement material	
Dental bone-grafting material	Dental bur	Dental
Electrical salivary stimulation system	Intranasal expiratory-resistance valve	
Extracorporeal shock-wave lithotripter	Dialyzer reprocessing system	Gastro.–Urology
Nonabsorbable suture	Surgical light connector	General & Plastic Surgery
Silk nonabsorbable suture	Surgeon's gloves	
Synthetic absorbable (polyglycolic-acid) suture	Polymeric surgical mesh	
Synthetic polyester nonabsorbable suture	Surgical gown	
Topical tissue adhesive	Removable skin staple	
Home uterine-activity monitor	Assisted-reproduction accessory	Obstetrics–Gynecology
Daily-wear rigid gas-permeable contact lens	Extended-wear rigid gas-permeable contact lens	Ophthalmic
Daily-wear soft contact lens	Extended-wear soft contact lens	
Bone cement	Orthopedic caliper	Orthopedic
Constrained hip prosthesis	Orthopedic template	
Intervertebral body fusion device	Carpal-lunate wrist prosthesis	
Stair-climbing wheelchair	Interactive rehabilitation exercise device	Physical Medicine
Bone sonometer	Combined PET/MRI tomographic imager	Radiology
Full-field digital mammography system	Post-biopsy breast compression device	
Magnetic resonance imaging system	Stationary X-ray system	

Note: For each of the 29 treated Class III-to-II device types, the table reports the paired *intuitive* control used in the clinically-matched robustness specification (Table 2, column 3): a never-reclassified device in the same FDA medical-specialty panel (3-digit regulation part), from a different regulation-number family, selected as the nearest such device by pre-event patenting scale. The headline nearest-neighbor controls instead match on pre-event patenting level and slope and exclude same-indication devices. Device descriptions correspond to FDA product classifications; the headline result is robust to these clinically-matched controls (Table 2).

Appendix Table A.7: Robustness to Shock-Level Inference

	Value-weighted count	Directional alignment
Coefficient	0.1392	0.1102
Standard error, clustering by:		
Treated device × country (baseline)	(0.0168)	(0.0119)
Treated device	(0.0543)	(0.0205)
Treated device and country (two-way)	(0.0580)	(0.0210)
Wild cluster bootstrap p-value	0.013	< 0.001
Treated devices (clusters)	29	22

Note: Shock-level inference for the two main-text patent headlines—the value-weighted (family-size-weighted) count of Table 2 and the directional-alignment measure of Table 3. The point estimate is invariant to the clustering scheme; only the standard error changes. The baseline clusters at the treated-device-by-country level; the next two rows cluster at the level of the policy shock (the treated device) and two-way by treated device and country. The wild cluster bootstrap-t uses Webb six-point weights with the null imposed ($B = 1999$). We do not apply Conley–Taber inference here because both outcomes have well above the number of treated units (29 and 22) at which its few-treated asymptotics are needed; it is reserved for the bilateral trade routes (Table 4) and the 12-event FDA-initiated subset. Clustering at the treated device inflates the count standard error roughly threefold (the t-statistic falls from 8.3 to 2.6) but leaves it significant, while the directional-alignment estimate is essentially unchanged.

Appendix Table A.8: Novelty of induced foreign patents: nearest-neighbor distance to prior U.S. art

	DiD	s.e.	p	Pre-trend p
Panel A. Expanding window (copying test): distance to contemporaneous prior U.S. art				
Density-fixed, $k = 5$ (headline)	+0.006	0.005	0.18	0.55
Raw pool, $k = 5$	−0.009	0.005	0.09	0.35
Density-fixed, $k = 1$	+0.002	0.004	0.69	0.50
Panel B. Fixed pre-event frontier (extension test): distance to frozen U.S. frontier				
Density-fixed, $k = 5$ (headline)	+0.010	0.005	0.04	0.46
Raw frozen frontier, $k = 5$	+0.010	0.005	0.04	0.56
Density-fixed, $k = 1$	+0.008	0.005	0.10	0.61
Density-fixed, $k = 10$	+0.011	0.005	0.03	0.34

Note: Stacked difference-in-differences estimates (Equation 1) of the nearest-neighbor novelty of induced non-U.S. patent families, matched-control sample. The outcome is the embedding distance (1−cosine) from each foreign family to its k nearest prior U.S.-inventor families in the same device type; higher is more novel. Panel A uses an expanding reference (U.S. art filed up to two years before each foreign filing); Panel B a reference frozen two years before the cohort's event year. "Density-fixed" averages the distance over ten random size-100 subsamples of the prior-art pool; "raw" uses the full pool. Family-level novelty is aggregated to the country-device-type-year panel by inventor-country fractional weights. Standard errors clustered by matched-control unit; pre-trend p is the joint test of leads $k \in [-8, -1]$. Panel A N = 20,697; Panel B N = 16,386. The Panel B frozen-frontier effect (+0.010) is about 0.13 standard deviations of the cross-sectional novelty distribution.

Appendix Table A.9: Leave-One-Out Robustness: Patenting Headline by Dropped Treated Device

Dropped device	Cohorts dropped	$\hat{\beta}_{\text{DiD}}$	(SE)	Δ vs. full	Influence
LTJ	51	0.0990**	(0.0149)	−0.0402	2.39
MKW	56	0.1121**	(0.0167)	−0.0271	1.61
LNS	25	0.1513**	(0.0170)	+0.0121	0.72
MAL	49	0.1294**	(0.0173)	−0.0098	0.58
LKD	33	0.1481**	(0.0173)	+0.0089	0.53
GAM	26	0.1481**	(0.0171)	+0.0089	0.53
HQD	23	0.1471**	(0.0171)	+0.0079	0.47
LOL	39	0.1468**	(0.0172)	+0.0076	0.45
GAS	11	0.1467**	(0.0168)	+0.0075	0.45
GAO	19	0.1465**	(0.0170)	+0.0073	0.43
MPN	32	0.1335**	(0.0172)	−0.0057	0.34
MWA	28	0.1444**	(0.0171)	+0.0052	0.31
LOJ	32	0.1436**	(0.0173)	+0.0044	0.26
KWZ	23	0.1435**	(0.0171)	+0.0043	0.26
MLM	28	0.1350**	(0.0172)	−0.0042	0.25
LQK	31	0.1430**	(0.0173)	+0.0038	0.23
GAP	8	0.1427**	(0.0169)	+0.0035	0.21
LNH	47	0.1358**	(0.0174)	−0.0034	0.20
LPK	38	0.1359**	(0.0173)	−0.0033	0.20
MAX	25	0.1418**	(0.0172)	+0.0026	0.15
LPL	43	0.1413**	(0.0174)	+0.0021	0.13
LOD	35	0.1412**	(0.0171)	+0.0020	0.12
MUA	7	0.1412**	(0.0169)	+0.0020	0.12
KRD	33	0.1374**	(0.0173)	−0.0018	0.11
IMK	29	0.1410**	(0.0171)	+0.0018	0.11
LON	33	0.1403**	(0.0172)	+0.0011	0.07
LTF	24	0.1402**	(0.0171)	+0.0010	0.06

Appendix Table A.10: Headline Sensitivity: Sample Restrictions and Anticipation Window

	Basic DiD	DiD × Import	DiD × Export	N
Headline	0.1392**	0.2943**	0.0721	136,649
	(0.0168)	(0.0921)	(0.0927)	
Panel A. Sample restriction (anticipation = 3, headline)				
Balanced ≥10y post (26 events)	0.1476**	0.2952**	0.0823	128,748
	(0.0178)	(0.0973)	(0.0979)	
Balanced ≥15y post (23 events)	0.1433**	0.3192**	0.0662	119,123
	(0.0190)	(0.1047)	(0.1020)	
Panel B. Anticipation window (all 29 events)				
Anticipation = 0	0.1525**	0.2953**	0.0138	129,710
	(0.0182)	(0.1002)	(0.0980)	
Anticipation = 1	0.1544**	0.2937**	0.0390	132,447
	(0.0177)	(0.0953)	(0.0936)	
Anticipation = 2	0.1369**	0.2786**	−0.0037	134,860
	(0.0170)	(0.0926)	(0.0925)	
Anticipation = 4	0.1217**	0.2906**	0.1105	137,468
	(0.0168)	(0.0871)	(0.0989)	
Anticipation = 5	0.1090**	0.2338**	0.1660	137,635
	(0.0169)	(0.0881)	(0.1030)	

Note: Family-size-weighted counterpart to the headline-sensitivity table. Each row is a separate stacked DiD of the value-weighted $\log(1 + \text{patent count})$ on the indicated DiD interaction, with cohort-by-device and cohort-by-year fixed effects and standard errors (parentheses) clustered at the treated-device-by-country level. Panel A fixes anticipation at three years and requires at least ten (or fifteen) post-treatment years; Panel B fixes the sample at all 29 events and varies anticipation from zero to five. The basic DiD stays positive and significant throughout (0.109–0.154). The import-share interaction, on the refined HS-6 concordance of Table 2, stays positive and significant at the 1% level in every specification; the export-share interaction is small and insignificant throughout. $^{+}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$.

Appendix Table A.11: IPC/CPC-Based Patent-Sample Robustness

	(1) Full corpus (baseline)	(2) A61 medical IPC (ex-A61K + G02B/G02C)	(3) CPC-natural subclass (22 of 29 treated PCs)
Panel A: Basic DiD			
DiD	0.1392**	0.1201**	0.1632**
	(0.0168)	(0.0190)	(0.0211)
Panel B: Directional Trade Heterogeneity			
DiD × Import Share from U.S.	0.2943**	0.2547*	0.4431**
	(0.0921)	(0.1057)	(0.1282)
DiD × Export Share to U.S.	0.0721	0.0884	0.1408
	(0.0927)	(0.1176)	(0.1358)
Observations	136,649	126,107	94,885
Cohorts	900	831	649

Note: All three columns re-estimate the headline matched-control stacked DiD (value-weighted log(1 + patent count) outcome, cohort-by-device and cohort-by-year fixed effects, standard errors clustered at the treated-device-by-country level), varying only the patent corpus: Column (1) the full keyword corpus; Column (2) families with an A61 IPC code other than A61K (or G02B/G02C); Column (3) the 22 devices whose corpus carries the device's natural CPC subclass in at least 40% of CPC-tagged patents. Panel A reports the basic DiD, Panel B the jointly estimated import- and export-share interactions on the refined HS-6 trade concordance of Table 2. The DiD stays positive and significant across all three corpora, with the import interaction positive and significant and the export interaction small and insignificant. $^{+}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$.

Appendix Table A.12: Assignee × Inventor Decomposition: Financial-Integration Test

	$\hat{\beta}_{\text{DiD}}$	(s.e.)
Panel A: Basic DiD (one regression with cell-specific coefficients)		
Cell A: U.S. assignee, foreign inventor (financial integration)	0.0245^{+}	(0.0134)
Cell B: Foreign assignee, foreign inventor (independent foreign firm)	0.1319^{**}	(0.0192)
Cell C: Mixed (both U.S. and foreign assignees)	0.0318^{+}	(0.0165)
Panel B: With trade-connectedness and R_c regulatory-channel controls		
Cell A	-0.2226	(0.2099)
Cell B	-0.0869	(0.2111)
Cell C	-0.2206	(0.2132)
Panel C: Cross-cell Wald tests		
Basic: $\beta_B - \beta_A$	$+0.1074^{**}$	($z = +6.00$)
Basic: $\beta_C - \beta_A$	+0.0073	($z = +0.50$)
With R_c: $\beta_B - \beta_A$	$+0.1358^{**}$	($z = +6.49$)
With R_c: $\beta_C - \beta_A$	+0.0021	($z = +0.12$)

Note: The patent family pool is split into three cells based on the (assignee, inventor) geography. All cells require at least one foreign inventor. Each row is a coefficient from a single pooled regression that stacks all three cells, with cell-specific DiD interactions and cohort×PC and cohort×year fixed effects; standard errors clustered on country×treated PC. Outcome is the value-weighted (family-size-weighted) log(1 + patent count) with inventor-fractional geography. Panel B adds DiD × Trade, DiD × ln(per-capita GDP), DiD × R_c, and DiD × Trade × R_c, the regulatory-channel specification, with R_c the audited FDA-reliance measure (Appendix B.2). Panel C reports Wald tests for equality of cell-specific DiDs derived from the joint covariance. Under the financial-integration channel, $\beta_A \geq \beta_B$: U.S.-owned patents filed by foreign inventors should drive the response. The data show the opposite. $^{+}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$.

Appendix Table A.13: FDA Device Type to HS-10 Concordance for U.S. Trade

Product code	**Device name**	**Anchor HS-10**	**Mapping confidence** *V. High / High / Med.-High / Med.*	**Trade coverage** *Full / Sparse / Dropped*	**Audit rationale**
GAM	Suture, Absorbable, Synthetic, Polyglycolic Acid	3006.10.01.00	High	Sparse	Sterile absorbable surgical suture (HS 3006.10)
GAO	Suture, Nonabsorbable	9018.32.00.00	Medium	Full	Stainless-steel suture; no clean HS-10 for steel suture wire; needles heading (HS 9018.32) is the closest
GAP	Suture, Nonabsorbable, Silk	3006.10.01.00	High	Sparse	Silk nonabsorbable suture (HS 3006.10)
GAS	Suture, Nonabsorbable, Synthetic, Polyester	3006.10.01.00	High	Sparse	Sterile polyester surgical suture (HS 3006.10)
HQD	Lens, Contact (Other Material) - Daily	9001.30.00.00	Very High	Full	Contact lenses (exact HS heading)
IMK	Wheelchair, Stair Climbing	8713.90.00.60	High	Sparse	Stair-climbing wheelchair; carriages for disabled persons (HS 87.13)
KRD	Device, Vascular, For Promoting Embolization	9018.90.80.00	Medium	Full	Intravascular embolization device; residual medical/-surgical instruments heading
KWZ	Prosthesis, Hip, Constrained, Cemented or Uncemented, Metal/Polymer	9021.31.00.00	Very High	Sparse	Hip joint prosthesis; artificial joints (HS 9021.31)
LKD	Monitor, Carbon-Dioxide, Cutaneous	9018.19.55.00	High	Full	Cutaneous CO2 monitor; specific to patient monitoring (HS 9018.19.55)
LNH	System, Nuclear Magnetic Resonance Imaging	9018.13.00.00	Very High	Sparse	Magnetic resonance imaging apparatus (exact HS heading)
LNS	Lithotriptor, Extracorporeal Shock-Wave, Urological	9018.90.75.40	Medium-High	Full	Extracorporeal shock-wave lithotripter (HS 9018.90.75.40)
LOD	Bone Cement	3006.40.00.00	Very High	Full	Bone cement (exact HS heading)
LOJ	Kit, Test, Alpha-Fetoprotein for Testicular Cancer	3002.15.00.11	Medium-High	Sparse	AFP tumor-marker antibody-based diagnostic; HS 3822.19 as robustness alternative
LOL	Hepatitis A Test (Antibody and IgM Antibody)	3002.15.00.11	High	Sparse	Hepatitis A antibody/serology diagnostic
LON	System, Test, Automated, Antimicrobial Susceptibility, Short Incubation	9027.80.00.00	Medium	Sparse	Automated antimicrobial-susceptibility analyzer (HS 9027.80 instrument family)
LPK	Bone Grafting Material, for Dental Bone Repair	9021.90.81.00	Medium	Sparse	Dental bone grafting material; broader than cement (hydroxyapatite / tricalcium phosphate / collagen)
LPL	Lenses, Soft Contact, Daily Wear	9001.30.00.00	Very High	Full	Contact lenses (exact HS heading)
LQK	Home Uterine Activity Monitor	9018.19.55.00	High	Full	Home uterine activity monitor; patient-monitoring electrodiagnostic apparatus
LTF	Stimulator, Salivary System	9018.49.80.80	Medium	*Dropped*	Intraoral salivary stimulator; raw Schott family spans unrelated HS-2 chapters (see appendix text)
LTJ	Prostate-Specific Antigen (PSA) for Management of Prostate Cancers	3002.15.00.11	Medium-High	Sparse	PSA immunoassay/tumor-marker diagnostic

continued on next page

(continued from previous page)

Product code	Device name	Anchor HS-10	Mapping confidence	Trade coverage	Audit rationale
LTK	Test, Epithelial Ovarian Tumor-Associated Antigen (CA-125)	3002.15.00.11	Medium-High	Sparse	CA-125 epithelial-ovarian-tumor antibody test
MAL	Graft, Vascular, Synthetic/Biologic Composite	9021.39.00.00	High	Sparse	Vascular graft prosthesis; artificial body-part implant
MAX	Intervertebral Fusion Device with Bone Graft, Lumbar	9021.10.00.50	High	Sparse	Intervertebral fusion device; implanted spinal device (HS 9021.10)
MKW	Cyclosporine	3822.19.00.80	Medium	Sparse	Cyclosporine immunoassay; clinical chemistry / toxicology reagent
MLM	Enzyme Immunoassay, Tacrolimus	3002.15.00.11	Medium-High	Sparse	Tacrolimus enzyme immunoassay
MPN	Tissue Adhesive for the Topical Approximation of Skin	3006.10.01.00	Very High	Sparse	Tissue adhesive for surgical wound closure (HS 3006.10)
MUA	Bone Sonometer	9018.12.00.00	High	Sparse	Bone densitometry by ultrasound (HS 9018.12)
MUE	Full-Field Digital, System, X-Ray, Mammographic	9022.14.00.00	High	Sparse	Full-field digital mammography (HS 9022.14)
MWA	System, Nucleic Acid Amplification, Mycobacterium tuberculosis Complex	3822.19.00.80	High-Medium	Sparse	Mycobacterium tuberculosis nucleic-acid amplification test (HS 3822.19 diagnostic reagents)

Note: Concordance from FDA medical-device product codes to ten-digit Harmonized Tariff Schedule of the United States (HTSUS) codes, used to construct the U.S.-specific import and export event studies in Appendix D.3 and Appendix Table 4 Panel A. Each device is mapped to an anchor HS-10 by reviewing the HTSUS textual description against the FDA device description; the trade panel aggregates all HS-10 codes whose definitions changed across the sample period and that share the device's six-digit heading, so that the device-level trade series is balanced from 1989 to 2024 even when individual HS-10 sublines were created, retired, or renumbered. *Mapping confidence* grades how cleanly the anchor HS-10 identifies the FDA device: *Very High* indicates the HS-10 heading uniquely identifies the device; *High* indicates the heading captures the device alongside closely related products; *Medium-High* and *Medium* indicate the best available HS-10 is plausible but broader than the FDA subtype. *Trade coverage* describes the resulting time-series support: *Full* means the device records positive imports and exports for at least 30 of the 36 sample years; *Sparse* means coverage is partial in at least one direction, typically because the anchor HS-10 is a post-2017 subline whose historical predecessor is grouped under a broader heading; *Dropped* indicates the LTF salivary stimulator, whose underlying Harmonized System family aggregates products from fifteen unrelated HS-2 chapters and contaminates the device-specific trade signal beyond repair. After dropping LTF and the six in-vitro-diagnostic devices whose HS-10 anchors lack a continuous historical series, 22 of the 29 treated devices enter the HS-10 event studies.

Appendix Table A.14: Quantity and Price Decomposition of the Route Trade Response

	Static DiD		Long-Run DiD	
	Quantity	Price	Quantity	Price
U.S. exports to foreign markets	0.203	0.049	0.290^{+}	0.059
	[-0.060, 0.472]	[-0.037, 0.172]	[-0.051, 0.615]	[-0.056, 0.196]
Foreign exports to the United States	0.473^{**}	-0.039	0.640^{**}	-0.014
	[0.052, 0.588]	[-0.126, 0.157]	[0.131, 0.797]	[-0.124, 0.189]
Foreign exports to third countries	0.280^{**}	-0.070	0.391^{**}	-0.084
	[0.122, 0.471]	[-0.140, 0.032]	[0.184, 0.624]	[-0.167, 0.048]
Foreign imports from third countries	0.023	0.050	0.042	0.067
	[-0.231, 0.174]	[-0.067, 0.152]	[-0.280, 0.250]	[-0.090, 0.181]

Note: Decomposition of the bilateral HS-6 route trade-value response (Table 4) into a physical-quantity component and a unit-value (price) component. *Quantity* is the stacked DiD on log(quantity) (BACI physical quantity in metric tons) and *Price* is the DiD on log(value/quantity), the log unit value, each estimated with the same specification and Conley–Taber inference as the route value estimates (reclassified device as the unit, placebo devices drawn from the control pool, 999 draws). The decomposition is on the intensive margin—cells with positive value *and* positive quantity, where price is defined and $\log(\text{value}) = \log(\text{quantity}) + \log(\text{price})$ holds exactly; this margin carries 97–99% of route trade value. Ninety-five percent Conley–Taber confidence intervals in brackets; significance from Conley–Taber p-values ($^{+}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$). Long-run effects are for event time > 6 years. On every route the price (unit-value) component is a tight null—its interval excludes any sizable repricing—so the significant route value responses (Table 4) are volume responses, not changes in unit values. The quantity component is positive on every route and significant on the two routes with the strongest value response (foreign exports to the United States and to third countries).

Appendix Table A.15: Exogeneity Robustness: The Response Survives on the Non-Manufacturer-Initiated and FDA-Initiated Subsets, Across All Margins

	(1) Full (29 events)	(2) Non-mfr-initiated (18 events)	(3) FDA-initiated (12 events)
Panel A: Patenting count (value-weighted $\log(1+y)$*)*			
DiD	0.139**	0.058**	0.076**
	(0.017)	(0.018)	(0.021)
Panel B: Import-share interaction			
DiD × Import Share from U.S.	0.294**	0.175	0.161
	(0.092)	(0.109)	(0.130)
Panel C: Innovation direction — alignment with U.S. (unitless)			
Directional alignment	0.110**	0.084**	0.085**
	(0.012)	(0.017)	(0.021)
Panel D: New-to-device inventor entry ($\log(1+y)$*)*			
New-to-device inventors	0.196**	0.132**	0.142**
	(0.017)	(0.020)	(0.023)
Panel E: Trade — U.S. exports to non-U.S. (HS-6, log *value)*			
Static DiD	0.397**	0.462**	0.404**
	(0.038)	(0.046)	(0.049)
Anticipation window (years)	3	3	3
Pre-trend F-test p (count)	0.311	0.716	0.629

Note: Robustness of the response across all five margins to two nested subsets that progressively drop reclassifications whose timing a device manufacturer could have influenced. A petition from a device's own manufacturer can carry private information about commercial pipelines and product-launch timing. Of the 29 reclassifications, 11 were petitioned by a named manufacturer: Column (2) drops those 11, leaving 18 events (12 FDA-initiated, 2 trade-group, 4 consultant), and Column (3) keeps only the 12 the FDA initiated on its own accord. Column (1) reports the full-sample estimates carried in the main text. Every column uses a three-year anticipation window, inventor-fractional geography, cohort-by-device and cohort-by-year fixed effects, and standard errors clustered at the treated-device-by-country level (treated-device-by-country-pair for the HS-6 trade route in Panel E). Panel A is the value-weighted $\log(1+\text{patent count})$ DiD (full-sample estimate from Table 2); Panel B the import-share interaction ($\hat{\beta}_2$ on DiD × Import Share); Panel C the directional-alignment outcome, unitless on $[-1,1]$ with no log analog; Panel D new-to-device foreign inventor counts, $\log(1+y)$; Panel E the static DiD on log HS-6 U.S.-to-non-U.S. export value. Sample sizes vary by outcome (largest for the patent-count panel). The patenting count attenuates from the full sample to the exogenous subsets—from 0.139 to 0.058 (18 events) and 0.076 (12 events)—but stays significant at the 1% level, and the pre-trend test clears more cleanly on both exogenous subsets than on the full sample. Innovation direction, new-to-device entry, and export trade remain positive and significant on both exogenous subsets, with clean pre-trends (direction and entry joint F-test $p \geq 0.63$; export-route pre-event coefficients jointly insignificant, $p \geq 0.51$). The import-share interaction keeps its sign and order of magnitude but loses significance on the thinner subsets, where the event count is too small to resolve the finer cross-country heterogeneity. $^{+}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$.

Appendix Table A.16: Where Within the Import Bundle Does the Response Load? Capital, Intermediate, and Final-Goods Decomposition

	(1) ImpShare	(2) CapShare	(3) IntShare	(4) ConShare	(5) Triple shares	(6) Triple log-levels
DiD	+0.140**	+0.140**	+0.139**	+0.140**	+0.140**	+0.141**
	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)	(0.017)
DiD × ImpShare	+0.324**	—	—	—	—	—
	(0.085)					
DiD × CapShare	—	-0.671**	—	—	-0.744	—
		(0.144)			(0.467)	
DiD × IntShare	—	—	+0.303*	—	-0.099	—
			(0.119)		(0.479)	
DiD × ConShare	—	—	—	+0.558*	+0.411	—
				(0.231)	(0.470)	
DiD × log Cap. imports	—	—	—	—	—	-0.053
						(0.044)
DiD × log Int. imports	—	—	—	—	—	+0.027
						(0.029)
DiD × log Con. imports	—	—	—	—	—	+0.076*
						(0.033)

Note: Family-size-weighted counterpart to the BEC end-use decomposition. Each column regresses the value-weighted $\log(1 + \text{patent count})$ on the DiD term and country-level share interactions (capital, intermediate, and final-consumption shares of U.S. imports, from BACI HS-96 mapped to BEC Rev. 4), with cohort-by-device and cohort-by-year fixed effects and standard errors clustered at the treated-device-by-country level. Column (1) is the product-level import-share interaction; (2)–(4) enter each share alone; (5) enters all three jointly; (6) uses $\log(1+\text{level})$ of the three import categories. The capital-goods interaction carries no positive signal (negative in every specification), while the response loads on intermediate and final-consumption exposure, consistent with product-embodied transmission rather than capital-equipment compatibility. $^{+}p < 0.10$, $^{*}p < 0.05$, $^{**}p < 0.01$.

Appendix Table A.17: Outsourcing and Safety Outcomes: Conley–Taber Stacked DiD

	$\hat{\beta}_{\text{DiD}}$	95% CI	Pre-period mean
Panel A: Contractor Entry (Extensive Margin)			
Any new contractor	0.170*	[0.069, 0.213]	0.002
Foreign contractor	0.091*	[0.024, 0.128]	0.002
Domestic (U.S.) contractor	0.156*	[0.064, 0.193]	0.002
Panel B: Adverse Events			
Event count	−11.054	[−43.545, 9.043]	1.87
Serious events	−0.973	[−7.190, 0.752]	0.35

Note: Sample is the Class III-to-II matched-control panel of Section 5. Each entry is a static Conley–Taber stacked DiD (static form of Equation 1) on a device-by-year panel. Panel A is an extensive-margin indicator for any new contractor in the treated device's FDA registry, split into foreign and domestic (U.S.); the pre-period mean is the cell-level probability of appearance under Class III. Panel B reports MAUDE counts of adverse and serious adverse events for foreign-source devices, pre-period means as annual cell counts. Both Panel B coefficients are negative with 95% confidence intervals covering zero—no evidence of safety deterioration after down-regulation. Bracketed intervals are 95% Monte Carlo; * marks intervals excluding zero. Intuitive- and later-treated-control columns are omitted: only two such control device types survive the outcome-specific restrictions.