The Impact of Private Equity Ownership on Medical Technologies

Atul Gupta* Parker Rogers†

November 2026

Please Do Not Circulate.

Abstract

Private equity (PE) ownership could promote innovation by providing capital and expertise, but it may also redirect firms toward more certain, near-term improvements that increase resale value at the expense of riskier, long-run investments. Despite these competing predictions, empirical evidence on how PE ownership affects innovation remains limited. This paper examines these effects in medical technologies, a highly innovative sector with substantial PE activity. We build a novel dataset that links patents to a set of 470 medical device firms that received PE investment between 1995 and 2020 and match those firms to 879 never-treated comparison firms. Using a stacked difference-in-differences design, we find that PE ownership reduces patenting in the long-run: treated firms produce roughly 11% fewer patents and are 13 percentage points less likely to file any patent. The decline is concentrated in leveraged buyouts involving healthcare-specialized investors. Within acquired firms, patenting declines in technological fields that overlap with other companies in the same PE owner's portfolio, but not in other fields, suggesting that PE owners reduce competing innovation across their portfolio companies.

*The Wharton School, University of Pennsylvania and NBER. atulgup@wharton.upenn.edu.

†Indiana University, Kelley School of Business. paroger@iu.edu.

We thank Jay Sastry and Sita Kottilil for outstanding research assistance, Ashley Swanson, and seminar participants at UCLA Anderson and ASHEcon for helpful suggestions. The National Institute of Healthcare Management (NIHCM) and the Mack Institute for Innovation Management at Wharton generously funded this research. The authors and research assistants used AI assistance in data analysis and manuscript preparation.

1 Introduction

New product development requires firms to make uncertain, long-run investments. These investments span years of research and development (R&D), costly and uncertain regulatory review, and commercialization. Private equity (PE) ownership may weaken incentives to undertake long-run investments because PE funds typically operate on relatively short, predetermined horizons and, in the canonical leveraged buyout, place substantial acquisition debt on the target firm's balance sheet, constraining cash flow. PE owners may also have weaker incentives to develop technologies that compete with those of other companies in their portfolios. Alternatively, PE acquisition may expand access to capital and technical and managerial expertise, which could accelerate growth and innovation and improve the likelihood of success. Thus, the effects of PE ownership on innovation are theoretically ambiguous. The limited empirical evidence on this topic, however, remains mixed.

We study the impact of PE acquisition in the U.S. medical device industry. This setting is economically important and well-suited for measuring product innovation. Medical device firms collected over $200 billion in annual revenue in 2019 (Donahoe 2021), and medical devices have played an important role in health improvements over recent decades (Cutler, Rosen, and Vijan 2006). Innovation is crucial for remaining competitive in this industry, and new products often require regulatory clearance through the Food and Drug Administration (FDA) before commercialization, allowing us to measure innovation beyond patents using FDA submissions. The setting is also ideal for studying the impacts of PE acquisitions. Although much of the policy debate focuses on PE activity in care provision, we find that PE firms acquired roughly 10% of innovative device firms operating in the U.S. during 1995–2020, a rate higher than the corresponding PE ownership rates for hospitals, nursing homes, and physician groups (Singh, Reddy, Papanicolas, and Scheffler 2026). These ownership transitions could change the rate and direction of socially valuable product innovation, with consequences for health care spending and effectiveness. Yet, to our knowledge, no prior work has examined how PE ownership affects innovation in medical technologies.

We identify 770 firms targeted over 1995–2020 using proprietary deal records from PitchBook, Inc. We also use PitchBook to identify medical device firms that appear in the private-capital data but never experience a private-equity buyout, and use these firms as never-treated controls. We obtain information on patent grants from the US Patents and Trademarks Office (USPTO) and device submissions from the FDA, which we then link to unique firms. These linkages allow us to build a novel panel dataset for approximately 6,900 distinct medical device firms, which we consider the near-universe of innovative firms in this industry.

To estimate the effects of PE investment, we use a stacked difference-in-differences (DD) design in which each treated firm is compared against up to three matched never-treated firms (Cengiz, Dube, Lindner, and Zipperer 2019). Our primary analysis sample contains 470 treated firms with pre-deal patents matched to 879 unique comparison firms (with replacement). Both treated and matched comparison firms must be independent at the time of the deal. The stacked DD estimator is robust to treatment effect heterogeneity in the presence of staggered treatment (Baker, Larcker, and Wang 2022), the setting here. Identification rests on the standard assumption that the treated and comparison firms would evolve in parallel in the absence of PE investment. Although this assumption is untestable, we find reassuring evidence that the treated and comparison firms trend in parallel in a "testing" period prior to the window used to match firms.

Our central finding is that PE investment leads to a decline in long-run innovation. These declines typically emerge 5 or more years after the deal and progressively increase in magnitude. We report coefficients corresponding to the average treatment effects over years 0–6 (medium-run) and years 7–12 (long-run).[1] The long-run declines in patenting are large. For example, we find that targeted firms file 11% fewer total patents and are 13 percentage points less likely to file a patent over years 7–12. Adjusting patent counts for quality based on citations increases the magnitude of the decline. The effect is driven

1. Year 7 marks the start of the long-run horizon as it coincides with the maximum gestation lag for new medical device products, consistent with the prior literature (Ji and Rogers 2024).

primarily by the extensive margin (i.e., firms are much less likely to file any patent or FDA submission). These effects are both economically and statistically significant. In contrast, the medium-run effects are usually small and statistically insignificant. Taken together, PE acquisitions reduce measures of mid-to-late-stage innovation, materializing only after the typical PE holding period elapses.

Heterogeneity tests reveal that the decline in patenting is driven largely by approximately 57% of the deals that involve at least one healthcare-specialized PE sponsor. A PE fund is deemed as healthcare-specialized if a majority of their investment targets between 1995 and 2019 are in the health care sector. Intuitively, healthcare-specialized investors are much more likely to invest in multiple medical device companies (3.6 on average, versus 0.7 for other investors). As a result, 42% of their targets have patents in technological fields that overlap with those of an existing portfolio company, while only 13% of targets of non-healthcare investors have such overlap. We hypothesize that PE firms optimize research spending across firms in their portfolio and limit the duplication of investments into the same technology class, leading to a decline in innovation that is concentrated among targets of healthcare-specialized investors. We formally provide support for this conjecture by analyzing the effects for 100 deals where we observe the target patenting in both types of technological classes: those that overlap with other firms their PE sponsor has already invested in, and those that do not. For these targets, we find that the decline in patenting is concentrated in the overlapping technology classes. Over years 7–12, the probability of any patent in an overlapping class falls by 7.3 percentage points *more* than in a non-overlapping class *within* the same firm. Similarly, within-firm patent counts fall differentially by 8% in overlapping classes. This effect is relatively stable in magnitude regardless of the extent of overlap with existing portfolio firms. Further, PE investors can plausibly redirect the research efforts of targets only in deals where they enjoy decisive managerial control. Consistent with this, the decline in patenting and FDA submissions for targets of healthcare-specialized sponsors is much greater for the subset of leveraged buyout deals, in which the PE fund takes a controlling stake in the target

firm.

A welfare interpretation requires caution. It is possible that PE owners cut redundant or low-return R&D projects. However, when we account for the number of citations each patent receives, a measure of the scientific value of innovation, our patenting estimates increase in magnitude, suggesting that the marginal quality of innovation may worsen. Additionally, industry-level estimates from Bloom, Schankerman, and Van Reenen (2013) suggest that the social value of medical innovation is at least twice as large as the private value, as medical device firms do not capture the full value that their technologies generate for patients or other firms. Moreover, we find that acquired firms reduce medical patenting and patenting across all industries proportionally, suggesting that firms are not merely reallocating R&D efforts elsewhere. However, in this paper's current form, we do not measure other outcomes that PE acquisition could improve, such as firm survival and revenue. In ongoing analysis, we are examining these firm-level outcomes using the US Census' Longitudinal Business Database.

The paper contributes to several intersecting strands of the literature. First, we add to limited evidence on the effects of PE investment on innovation. The closest study is Lerner, Sorensen, and Strömberg (2011) (LSS), who examine the impacts of LBOs across manufacturing-heavy industries and find no evidence of reduced patenting. Although our results may appear to contradict their findings, important differences in the data and methodology may explain the divergence. First, we focus on medical device technologies, which is a small segment of their sample. Second, LSS does not use a DD design since they lack a never-treated comparison group. Instead, they focus on the change in innovation for treated firms between the periods before and after acquisition. Third, they focus on medium-run outcomes, following firms for five years after the deal. Indeed, we also find no decline in innovation output in the medium run; negative effects materialize only in the long run. More recently, Cumming, Peter, and Tarsalewska (2020) examine the effects of PE acquisition of publicly listed manufacturing firms, which are quite different from the mostly private medical device firms in our sample, using a DD design and find a decline in

both patent counts and citations. A related literature on PE's broader effects on operations, employment, productivity, and care quality (Davis et al. 2014; Bernstein and Sheen 2016; Bernstein, Lerner, Sorensen, and Strömberg 2017; Gupta, Howell, Yannelis, and Gupta 2024) provides a wider benchmark for interpreting our estimates.

Second, we contribute to the literature on medical innovation. Prior work has examined how regulatory pathway (Stern 2017; Grennan and Town 2020; Rogers 2024), intellectual-property scope (Budish, Roin, and Williams 2015; Williams 2013), demand and procurement (Finkelstein 2004; Clemens and Rogers 2026), and reimbursement design (Clemens and Gottlieb 2014; Dranove, Garthwaite, Heard, and Wu 2022; Ji and Rogers 2024) shape product development. These papers hold ownership fixed and vary policy; we hold policy fixed and vary ownership. The results show that firm ownership status affects the same invention and commercialization outcomes that the policy literature studies.

Third, we contribute to work on how ownership that spans related firms shapes innovation. Acquirers discontinue acquired drug projects that overlap their own products (Cunningham, Ederer, and Ma 2021); in theory, overlapping ownership can raise or lower R&D depending on whether owners internalize technology spillovers or product-market rivalry (López and Vives 2019); and common ownership by institutional investors has been linked to more innovation where inventions benefit other firms (Antón, Ederer, Giné, and Schmalz 2025). PE sponsors are concentrated owners with a voice in their companies' research budgets, and they often hold several companies in the same industry at once. We show that under their ownership targets retreat from the fields their sponsors' other companies occupy.

The remainder of the paper proceeds as follows. Section 2 provides institutional background. Section 3 develops the conceptual framework. Sections 4 and 5 describe the linked dataset and the empirical strategy. Section 6 reports the main results and Section 7 provides evidence on potential mechanisms. Section 8 concludes.

2 Background

This section describes the US medical-device industry, the structure of private-equity (PE) investment within it, and the deal-type variation in owner horizons that our empirical design exploits.

2.1 The US Medical-Device Industry

The US medical-device industry generates over $200 billion in annual sales and is among the more innovation-intensive segments of US health care. The Food and Drug Administration (FDA) regulates devices through a three-tier risk classification established by the 1976 Medical Device Amendments. Class I (low risk) devices generally require facility registration alone; Class II (moderate risk) devices typically proceed via the 510(k) "substantial equivalence" pathway; Class III (high risk) devices require premarket approval (PMA) supported by clinical evidence (Makower, Meer, and Denend 2010). The 510(k) pathway processes roughly three thousand submissions per year, while PMA originals number only thirty to fifty annually.[2] The two pathways differ in the type of innovation they reward: 510(k) activity is dominated by incremental product extensions, while PMA originals capture the breakthrough margin.[3]

A small set of large, publicly traded incumbents (Medtronic, Johnson & Johnson, Stryker, Boston Scientific) accounts for a sizable share of submissions, but the long tail of small and mid-size firms originates a large fraction of new device categories.[4] It is this long tail that supplies the population of PE-eligible targets.

2. Preparation of a PMA dossier routinely runs four to seven years from initiation, while a 510(k) submission can be prepared in months. Costs differ accordingly: Makower, Meer, and Denend (2010) estimate average costs of $24 million for 510(k) and $75 million for PMA.

3. Prior work on medical-device innovation has examined regulatory pathway design (Stern 2017; Grennan and Town 2020; Rogers 2024), intellectual-property scope (Budish, Roin, and Williams 2015; Williams 2013), demand and procurement shocks (Clemens and Rogers 2026), and reimbursement rules (Clemens and Gottlieb 2014; Dranove, Garthwaite, Heard, and Wu 2022; Ji and Rogers 2024). Prior coauthor papers on related questions include Rogers (2024), Clemens and Rogers (2026), Ji and Rogers (2024), and Gupta, Howell, Yannelis, and Gupta (2024). This paper complements that line by isolating the role of ownership rather than regulation, intellectual property, demand, or reimbursement.

4. In our analysis sample, the modal firm files one to three 510(k) submissions per year and maintains an active patent portfolio.

2.2 Private Equity in Medical Devices

PE investment in US medical-device firms expanded sharply between 1995 and 2020. PitchBook records over 1,400 PE-affiliated deals on US medical-device firms in our curated universe, with annual deal counts roughly tripling between the late 1990s and the late 2010s. Deals fall into four categories. *LBO* pairs a controlling stake with substantial leverage and an active operating playbook (Kaplan and Strömberg 2009); sponsors aim to exit within a fund's investment horizon via secondary sale, strategic acquisition, or IPO. *Growth* deals provide minority or non-controlling growth equity, carry little or no acquisition-related leverage, and typically come with board representation rather than operating control. *PIPE* (private investment in public equity) transactions inject capital into already-public firms.[5] *Buyout by Management* captures management-led acquisitions with PE financing. LBO and Growth together account for the large majority of medical-device deals. Sponsors often hold several device companies at once: among the treated firms in the within-firm design of Section 7.1 whose sponsor we can identify, 60% have a sponsor that held at least one other patenting device company on the date of the deal (a median of two).

The sectoral footprint tracks broader trends documented by Lerner, Sorensen, and Strömberg (2011): rising fund size and rising deal multiples.[6]

2.3 PE Ownership Structure

PE owners hold portfolio firms for a finite, deal-type-specific period. Kaplan and Strömberg (2009) report a median holding period across US buyouts of roughly five years, with substantial right skew and few holds extending beyond ten years. The two PE deal types, LBO and Growth, differ in how they own: an LBO pairs a controlling stake with substantial acquisition leverage, whereas PE Growth/Expansion supplies minority or non-controlling capital with

5. PIPE transactions behave more like financial than operating investments and constitute a small share of our sample.

6. See also Kaplan and Strömberg (2009) and Lerner, Sorensen, and Strömberg (2011) on the cross-industry evidence on PE and innovation. The medical-device sector is well suited to a within-industry test because deal flow is large and outcomes (patents, FDA submissions) are well measured at the firm level.

little leverage. LBO deals also provide the PE fund managerial control over the target firm that may be lacking when it has a minority stake.

3 Conceptual Framework

We develop a model of R&D choice under a portfolio owner (formal details in Appendix B). A PE sponsor typically holds several companies in the same industry at once, as discussed earlier. A target's research project can add value to the target while taking value from the sponsor's other companies—by duplicating research they already fund, or by producing a device that competes with theirs for the same hospitals and physicians. An independent owner has no reason to weigh these effects; a sponsor that owns both companies does. Prior research shows that acquisitions by incumbent firms can discourage the development of competing products (Cunningham, Ederer, and Ma 2021). We investigate whether similar incentives arise under private equity ownership, where multiple firms may remain operationally separate but share a common financial owner. Our findings suggest that PE ownership can redirect innovation away from technologies that compete with those developed by other portfolio companies. Cunningham, Ederer, and Ma (2021) examine the continuation of existing drug development projects and we look at subsequent patenting within the field, which could include new starts.

A firm has a set of potential R&D projects. Project j requires an up-front outlay c_j—predominantly the salaries of the scientists and engineers who develop it—and, if funded, delivers expected value v_j to the firm after a development lag τ_j. Projects are heterogeneous in cost, value, and payoff horizon: a line extension or repackaged 510(k) submission is inexpensive and pays off in one to two years, while a novel implantable backed by a PMA is expensive and pays off in seven to ten. Each project lies in a technology field f_j. If another company in the sponsor's portfolio is active in that field, the project reduces that company's value by $e_j > 0$; if no portfolio company is active in f_j, $e_j = 0$. What creates the loss is that

a portfolio company competes or does research in the field, not the scale of its activity there: a single product line can bear the full cost of a rival project. The owner internalizes a share $\omega_\theta \in [0,1]$ of that loss, where θ indexes the owner: an independent owner has $\omega = 0$, and a PE sponsor has $\omega > 0$, set by its influence over the target's research budget—full control in a buyout, board representation and investor rights in a growth-equity deal—so that $\omega_{\text{LBO}} \geq \omega_{\text{G}}$: a sponsor that controls the budget can act on the whole internalized loss, while a minority investor can act on it only as far as its board seats and investor rights reach. The firm funds a project if its value net of the internalized loss covers its cost:

$$v_j - \omega_\theta \, e_j \;\geq\; c_j, \qquad \text{equivalently} \qquad \frac{v_j}{c_j} \;\geq\; 1 + \omega_\theta \, \frac{e_j}{c_j}. \tag{1}$$

The funded set is $U_\theta = \{j : \; v_j/c_j \geq 1 + \omega_\theta e_j/c_j\}$. Equation (1) delivers five predictions.

(i) **Aggregate decline.** Because $e_j \geq 0$, a sponsor with $\omega > 0$ funds a subset of the projects an independent owner would fund, so innovation output weakly falls after the deal.

(ii) **Delayed decline.** The new owner's funding rule applies to projects started after closing, but projects already under way continue to generate patents for several years, so measured output falls only after the development lag: $|\hat{\beta}_{\text{Long}}| > |\hat{\beta}_{\text{Medium}}|$.

(iii) **Concentration in the sponsor's fields.** Projects in fields no portfolio company occupies ($e_j = 0$) face the same hurdle as before, so the decline concentrates in the fields the sponsor's other companies occupy.

(iv) **No selection on horizon, breadth, or impact.** The hurdle depends on a project's return and its overlap with the sponsor's other companies, not on its payoff horizon, technological breadth, or eventual impact; if those characteristics are unrelated to overlap, the decline carries no gradient in them.

(v) **Control, not leverage.** Acquisition debt does not enter Equation (1), so the decline should not rise with deal leverage. Deal type matters only through ω_θ: because $\omega_{\text{LBO}} \geq \omega_{\text{G}}$, the decline and its concentration in the sponsor's fields should be at least as large under LBOs, and should appear under growth only to the extent that board representation gives the investor influence over the research budget.

Two features of the mechanism are worth drawing out. First, it is selective along one dimension only. A financing constraint would cut research across the board, and a short horizon would cut the longest-dated projects first; portfolio internalization cuts the projects that overlap the sponsor's other companies and leaves the composition of the rest by horizon, breadth, and impact unchanged. Because funding is a discrete choice, the cut also shows up as fewer new projects rather than a uniform scaling-down of every project.

Second, the welfare consequences are ambiguous. When e_j reflects research that duplicates a sibling's, dropping the project saves real resources. When e_j reflects business the project would take from a sibling, the sponsor's gain is a transfer, and the forgone project would have added competition and variety to the device market—the concern raised by killer acquisitions. Our estimates identify the cut, not its welfare sign.

The model treats each project's cost as avoidable. Research staff, however, are costly to hire and to let go, and much of their know-how is specific to the firm. If the firm keeps its inventors when it stops funding projects in the sponsor's fields, their salaries are no longer an avoidable cost, so some projects in the firm's other fields that did not clear the hurdle before now do (Appendix B, Proposition 5). The model then predicts a partial redirection: output falls in the sponsor's fields and weakly rises in the firm's other fields, by less than it falls in the sponsor's fields, so total output still falls; the inventors who worked in the sponsor's fields stay and move to other fields.

4 Data

We assemble a firm-year panel that links PitchBook deal records, USPTO patent grants and applications from PatentsView, and FDA regulatory submissions for US medical-device firms over 1995–2024. Building this panel requires resolving heterogeneous firm-name strings across three independent administrative systems — none of which share a common identifier — and then constructing matched pairs of PE-acquired ("treated") and never-PE-acquired ("control") firms with similar pre-treatment innovation trajectories.

4.1 Sources

PitchBook. PitchBook is the source for the universe of PE deals and for firm metadata — founding year, location, deal history, and deal type. We use two PitchBook extracts. The first is the full deal record of 174,726 deals spanning 1980–2025, from which we draw the control pool. The second is a curated medical-device PE deal list containing 1,998 deals across 1,464 unique firms, last refreshed in October 2022; this defines our treated population. Treated firms are PE acquisitions of US medical-device firms that appear in PitchBook's medical-device universe. Controls are drawn from the same universe but are not flagged as PE targets in the curated list.

PatentsView. PatentsView is the USPTO's bulk data repository of patent applications and granted patents. We extract filing date, grant date, CPC and IPC classifications, assignee identifiers, and forward citations. We use IPC subclasses for the technological-proximity matching caliper, CPC subclasses for the direction-of-innovation measures (concentration, new-subclass entry) and the core/non-core split, and the IPC class A61 to flag health-related patents. We align event time to filing year rather than grant year because filing dates are closer to the underlying R&D decision — grant lags average 2–3 years and are partly determined by examiner backlogs that have little to do with firm strategy. Forward citations are aggregated over a five-year window from grant and normalized within technology field

and grant year following Hall, Jaffe, and Trajtenberg (2001) (the HJT normalization).

FDA. The FDA data comprise 510(k) substantial-equivalence submissions, PMA premarket approvals, Manufacturer and User Facility Device Experience (MAUDE) adverse-event reports (device files, 1997–2019), recall records, and FDA establishment registrations. 510(k) is the dominant clearance pathway for Class II medical devices; PMA covers higher-risk Class III devices. We define a firm-year regulatory submission as the count of 510(k) and PMA submissions on which the firm appears as either applicant or contract manufacturer of record. Adverse-event reports come from the MAUDE database and are aggregated to the firm-year by manufacturer string. Establishment registrations are used for entity resolution, since they provide an additional set of firm-name strings to anchor the embedding cluster.

Census firm-level financial outcomes (ongoing). We have linked the matched sample to the Longitudinal Business Database housed at the US Census, to estimate effects on employment, payroll, revenue, and firm survival.

4.2 Entity Linkage

The three primary data sources use different firm-name conventions, and the same firm appears under many name variants in the FDA submissions, recalls, MAUDE, and registration files — legal entity names, doing-business-as names, subsidiary names, and registered trade names. Building a firm-year panel therefore requires us to resolve these strings to a common entity. This is a non-trivial step in the data construction that is described in full in Appendix D.

4.3 Matching Procedure

For each treated firm T with deal-anchor year a_T, we select $K = 3$ never-PE-treated controls by minimizing a standardized Euclidean distance over a common set of pre-period firm

characteristics, measured over the six years ending at the anchor:

$$d(T, C) \;=\; \sum_{f \in \mathcal{F}} \left(\tilde{x}_f^T - \tilde{x}_f^C \right)^2, \tag{2}$$

where $\mathcal{F}$ collects seven firm characteristics: the mean and the linear trend (slope) of the annual patent flow, the mean and the linear trend of the annual FDA-submission flow, an indicator for any venture-capital financing round, $\log(1 + \cdot)$ of cumulative pre-period patent stock, and firm age, measured as the number of years since the firm's first US patent (patenting tenure). Each characteristic x_f is standardized (z-scored across the candidate control pool for the treated firm's anchor cohort, denoted $\tilde{x}_f$), so all seven enter the distance on a common scale with equal weight. Summarizing the patent and FDA flows by their pre-window mean and slope, rather than the full year-by-year trajectory, matches on both the level and the pre-deal growth path of each outcome while keeping the distance low-dimensional.

We enforce a hard caliper on technological proximity: the cosine similarity of the two firms' patent portfolios across IPC technology subclasses (Jaffe 1989), evaluated at the deal-anchor year, must be at least 0.25. Controls below the caliper are removed from the candidate set before the distance minimization runs. We also exclude a small number of diversified conglomerates (firms with more than 200 patents in a single pre-period year) and require every candidate control to have filed at least one US patent in the ten years before the anchor.

Both treated and control firms in a pair must be independent in the treated firm's deal year D: neither owned by another company nor a business divested by another company at D. We classify each firm's status at D from its latest completed ownership event before D in the PitchBook deal record. A firm with no earlier event, or whose latest event left it independent (an initial public offering, a reverse merger, a spin-off, or a management buyout), is eligible. A firm whose latest event was an acquisition by another company, a take-private, a closure, or a private-equity deal is not; events in the same year resolve toward loss of independence, and a control is also ineligible after any completed private-equity deal before D.

We match $K = 3$ with replacement. The choice trades variance reduction against the risk of admitting poor-quality matches; with replacement, a strong control can serve multiple treated firms rather than forcing weaker matches into the sample. Controls reused across stacks are handled in the estimation step through firm-by-stack fixed effects. After the patent-active restriction the match yields 470 unique treated firms paired with 879 unique never-treated controls, producing 1,397 paired stacks. Appendix C reports full balance diagnostics, including the distribution of cosine values and age gaps.

4.4 Final Analysis Sample

Not every matched pair survives to the headline estimation sample. Two binding restrictions follow. First, we require both the treated firm and its matched control to have at least one US patent filed in the ten years preceding the deal anchor (event-time $[-10, -1]$); firms with no pre-period patenting have no patenting capacity to lose, so they do not enter a test of the patent-decline hypothesis. Second, FDA-side outcomes require both sides of the pair to have a non-empty FDA cluster mapping. We apply both restrictions uniformly across outcomes so that patent and FDA coefficients are directly comparable, and, like every filter in Table A1, identically to targets and controls. Table A1 reports every filter, with its count of targets and of never-PE firms.

The headline sample is the row in bold: 470 unique treated firms matched to 879 unique controls, producing 1,397 paired stacks. Of the 770 treated firms with a usable deal anchor, 2 are diversified conglomerates and 121 were not independent firms at the deal (Section 4.3), and of the remaining 647, 177 fail the patent-active filter or have no eligible control within the caliper. FDA outcomes are estimated on the FDA-linked subsample—both the treated firm and its matched control must be observable in the FDA data (i.e., appear in the FDA firm-cluster crosswalk)—which trims the sample to 230 treated firms and 357 controls (530 paired stacks).

4.5 Descriptive Statistics and Match Quality

Appendix Table A2 compares the estimation-sample targets with the never-PE-acquired firms that would be eligible controls, before matching, and Table 1 reports pre-period firm-year means for treated and matched controls across outcomes over event-time $[-5, -1]$, alongside the share of firm-years with any non-zero activity. Treated firms have higher levels of innovation activity than their controls in the pre-treatment period, which reflects selection into the PE-acquisition pool on observable innovation capacity.

Appendix Table A3 shows where treated firms and their matched controls file with the FDA in the decade before the deal. The two groups work in the same device specialties: cardiovascular, general and plastic surgery, and orthopedic devices account for 40% of the treated firms' activity and 38% of the controls', and the largest difference in any specialty is in orthopedic devices (12% against 9%). Appendix Table A4 lists the PE investors that back the most targets. No sponsor dominates the sample: the most frequent, RoundTable Healthcare Partners, backs 9 of the 347 baseline targets with a recorded PE investor (2.6%), and 7 of the 15 most frequent sponsors are healthcare-specialized.

5 Empirical Strategy

5.1 Estimating Equation

For each outcome Y we estimate,

$$Y_{ist} \;=\; \beta \cdot \mathbf{1}[\text{treated}_i \cdot t \geq 0] \;+\; \alpha_{is} \;+\; \gamma_{ts} \;+\; \varepsilon_{ist}, \tag{3}$$

where i indexes the firm, s indexes the matched stack, t is calendar year, and the indicator $\mathbf{1}[\text{treated}_i \cdot t \geq 0]$ equals one for the treated firm in stack s from the deal anchor onward and zero otherwise. The coefficient β identifies the average treatment effect on the treated under the conditional parallel-trends assumption.

Two features of the fixed-effect structure deserve emphasis. First, α_{is} is a *firm-by-stack* fixed effect, not a firm-only fixed effect. A control firm that is reused across multiple stacks enters each stack with its own fixed effect, so within-stack identification is not contaminated by cross-stack pooling. This is the standard stacked-DiD convention introduced by Cengiz, Dube, Lindner, and Zipperer (2019). Second, γ_{ts} is a *stack-by-year* fixed effect: each treated firm's counterfactual is the matched controls' value in the same calendar year within the same stack, so the comparison absorbs any within-stack common shock. Standard errors are clustered at the firm level (not the firm-by-stack composite) so that within-firm correlation propagates across the stacks in which a given control firm appears.

We estimate Eq. (3) on an event-time window of $[-10, +12]$; the ten pre-deal years are the years over which the patent-activity filter is applied ($[D-10, D-1]$), and Table 4 reports coefficients corresponding to the $[-12, 12]$ and $[-8, 12]$ windows as robustness checks.

For each outcome we report three ATT coefficients. *Overall* is the single-indicator specification with $t \in [0, 12]$, and reflects the average post-deal effect over the full estimation window. *Medium* ($t \in [0, 6]$) and *Long* ($t \in [7, 12]$) are estimated jointly within a single two-indicator regression that retains the same controls and fixed effects as the Overall specification. The medium/long split is a test of the delayed-decline prediction from the conceptual framework.

Identification rests on the conditional parallel-trends assumption: absent the PE acquisition, the matched controls' time path within the stack would be a valid counterfactual for the treated firm. Three design choices support this assumption.

First, matching on the pre-period level and growth of both innovation margins together with venture-capital financing history, patent stock, and firm age aligns the pre-period innovation paths of treated and control firms before any post-deal comparison is drawn. Second, the IPC-subclass cosine caliper restricts matches to firms operating in similar technology areas. This addresses the concern that medical-device sector shocks — say, a Center for Medicare and Medicaid Services coverage decision affecting cardiovascular devices

— could hit treated and control firms differentially if the matched control sits in a distant technology space. Third, the stack-by-year fixed effect γ_{ts} removes any within-stack common shock affecting treated and control firms in the same calendar year. The comparison in Eq. (3) is between deviations from this within-stack year mean, not between levels.

Furthermore, we present event study plots of dynamic treatment effects for each outcome of interest. This allows the reader to assess whether treated firms were trending differentially in the pre-treatment period. Reassuringly, we find that treatment and comparison firms were trending along similar lines in the pre-treatment period, including prior to the period used in matching (-10,-6).

6 Results

This section presents estimates of Equation (3), which captures the effect of private-equity acquisition on innovation and product safety for medical-device firms.

6.1 Patent Output: Quantity and Quality

Section 3 predicts that the effect of PE acquisition on patent output is concentrated at long horizons. Medical-device development cycles range from three to seven years for moderate-risk devices and longer for Class III breakthrough devices (Makower, Meer, and Denend 2010; Stern 2017), so a contraction of new-project initiation at the deal shows up in patent and FDA output only at long event-time. We test this prediction by estimating Equation (3) with two post-period indicators jointly: a Medium indicator for $t \in [0, 6]$ and a Long indicator for $t \in [7, 12]$. Both coefficients are reported in Table 2.

We report the Medium- ($t \in [0, 6]$) and Long-run ($t \in [7, 12]$) coefficients estimated jointly in the table. Each row reports three coefficients estimated on the uniform-filter sample of 470 unique treated firms matched to 879 unique control firms: the effect on log patents, log citation-weighted patents, and the probability of filing a patent. The estimates separate

the two horizons sharply. On log(1 + total patents), the Medium coefficient is −0.018 (SE 0.029), statistically indistinguishable from zero, while the Long coefficient is −0.111 (SE 0.052, significant at the 5% level), roughly an 11% decline in total patenting in years 7–12 relative to matched controls. The Long-run extensive ATT is −12.7pp on all-patents (against pre-period treated probability of 47%), corresponding to roughly a 27–28% relative decline in the probability of any patenting in years 7–12.

The quality margin declines at least as much as the count margin. Citation-weighted patents are normalized within CPC subclass × grant year following the HJT convention, stripping out cohort-level inflation in citation rates and field-specific citation practices. The Long-bin coefficient on citation-weighted patents is −0.135 (SE 0.045, significant at the 1% level), an approximately 13% decline in field-normalized citation-weighted output in years 7–12 post-deal. This exceeds the Long-bin coefficient on all-patents (−0.111, SE 0.052), so the decline is at least as large on the quality-weighted margin as on the raw count.[7] Because HJT normalization removes field-year compositional drift, the pattern is not a mechanical shift away from citation-rich fields; the projects treated firms forgo are at least as cited as the typical patent they continue to file.

We study dynamics by estimating Equation (3) as an event study. Figure 1 Panel (a) plots the event-time coefficients β_k for log(1 + total patents) on the uniform-filter sample. Pre-period coefficients are small and statistically indistinguishable from zero across $t \in [-5, -2]$. Joint pre-trend tests fail to reject zero on all three patent outcomes. A joint Wald test on the far-pre coefficients at $t \in [-10, -6]$ — entirely outside the matching window $[-5, -1]$ — also fails to reject the null hypothesis, indicating that treated and matched-control firms move in parallel deep into the pre-period, well before the periods used for matching. Post-deal, the coefficient is small and statistically insignificant through $t = 3$, turns negative at $t = 4$, and lies between −0.08 and −0.19 over years 7–12. The gradual post-period divergence

7. The decline is not a mechanical artifact of post-deal name changes. We attribute patents to firms through a union of assignee identifiers that follows renames and successor entities identically for treated and control firms.

is consistent with reduced project initiation after the deal rather than a discrete shift in patenting at the deal date. Panels (b) and (c) show similar patterns of long-run declines.

Medical-device patents follow the same pattern as all patents (shown in Appendix Figure A1 and Table A5): Medium −0.014 (SE 0.025, ns) versus Long −0.089 (SE 0.046, significant at the 10% level). Citation-weighted patents show the largest long-run decline: Medium −0.051 (SE 0.030, significant at the 10% level) against Long −0.135 (SE 0.045, significant at the 1% level), roughly a 13% decline in field-normalized citation-weighted output. Long-bin extensive-margin coefficients are −10.2pp (SE 3.0pp, significant at the 0.1% level), against pre-period treated probabilities of any filing of 36%.

The Long-bin coefficient on citation-weighted patents is at least as large in absolute value as the Long-bin coefficient for medical device patents. Because HJT normalization removes field-year compositional drift in citation rates, the citation-weighted decline is not driven by treated firms mechanically shifting away from citation-rich fields. The marginal project that treated firms forgo appears to be at least as cited as the typical patent the firm continues to file, but the evidence does not imply that remaining patents become lower quality.

6.2 FDA Submissions

We turn from patents to FDA regulatory submissions, which capture the firm's downstream commercialization pipeline.[8] Table 3 reports the corresponding ATTs.

Estimates indicate an FDA decline that is concentrated at long horizons. The Medium coefficients are small and not significant (−0.005 log1p; −2.8pp extensive, SE 2.4pp). At long horizons the probability of any submission falls by 7.6 percentage points (SE 3.6pp, significant at the 5% level); the log-count coefficient is −0.065 (SE 0.048), of the same sign

8. FDA outcomes are estimated on the *FDA-linked* subsample—firms for which both the treated unit and its matched control are observable in the FDA data (i.e., appear in the FDA firm-cluster crosswalk), the FDA analog of the patent-active restriction applied to the patent outcomes (Section 4.4); this trims the sample to 230 unique treated firms and 357 unique controls. Restricting to firms observable in the FDA records avoids attributing structural zeros to firms merely absent from those records rather than genuinely inactive, which would otherwise dilute the estimate toward zero. We apply this same firm-level FDA-linked restriction to every FDA outcome (submissions, recalls, and adverse events).

but not significant. Among firms observable in the FDA data, then, PE ownership lowers the likelihood of engaging the FDA in a given year over years 7–12—a decline somewhat smaller than on patents, of the same sign and long-run timing. Both FDA coefficients are, moreover, conservative: treated firms are on a mildly rising submission trajectory before the deal (a +0.006 to +0.008/year pre-trend), so differencing off that baseline understates the decline. Measured against the extrapolated pre-deal trend, the extensive-margin decline is −0.147 and the log decline −0.171 (both significant at the 5% level), and because acquisition *reverses* the pre-existing upward trend the pattern cannot be a continuation of pre-deal dynamics.

Figure 2 plots the event-time coefficients on log(1 + total FDA submissions). The series shows flat pre-period coefficients (the joint Wald test on the far-pre coefficients at $t \in [-10, -6]$ gives $p = 0.55$ on the log1p margin and $p = 0.73$ on the extensive margin) and the same delayed, gradual post-period divergence visible on patents, with the gap widening through $t = 12$.

We decompose the FDA pipeline using a firm-specific definition of regulatory specialty. For each firm, "core" regulation numbers are the pre-period top-80% by submission share; the remainder is "non-core." This decomposition is applied symmetrically to treated and matched controls, on the subset with non-zero pre-period FDA activity on both sides. If PE acquisition pushed firms back into established regulatory specialties, the core series should hold up better than the non-core series. Instead, the core series carries the larger decline: the Long-bin core ATT on log1p is −0.106 (SE 0.049, significant at the 5% level), while the non-core series falls less and not significantly (−0.075, SE 0.066). The contraction is if anything sharper in the firm's established regulatory areas—the opposite of a refocus-toward-core prediction, and consistent with a broad-based pullback rather than reallocation. The corresponding patent-side core/non-core estimates likewise show decline rather than reallocation, with no pivot toward core specialties.

6.3 Robustness checks and Negative Controls

6.3.1 Robustness checks

We report four sensitivities for the long-bin estimates in Table 4. The first restricts to the subset of deals whose entire post-period $[0, 12]$ is observable—deals completed through 2008 for the patent outcomes and through 2012 for FDA submissions (Panel B). The citation-weighted and extensive-margin estimates are similar in magnitude and precision on this balanced subsample, while the patent counts coefficient is slightly smaller and loses precision in the smaller sample. Next, we examine sensitivity to dropping the patent filter restriction (Panel C). An advantage of this approach is that we use the full set of deals for both patent and FDA outcomes. The effects on patents remain similar in magnitude while the effects on FDA outcomes drop in size. For both sets of outcomes, the effects remain statistically significant. In Panels D and E we test sensitivity to shortening or extending the pre-treatment period by 2 years. These changes do not materially affect the estimates. The last check drops the data-horizon cap and keeps every calendar year the panel covers (Panel F): the long-run declines are then of similar size (-0.145 on citation-weighted, -0.105 on all-patents log1p, -8.4pp on the extensive margin) and more precisely estimated.

6.3.2 Negative-control outcomes

A causal reading of the decline requires that PE targets would have followed their matched controls absent the deal. The most plausible violations are that PE acquires firms whose technology is losing relevance, or firms in fields where innovation is shrinking. Both are testable with outcomes that such selection would move but PE ownership cannot. First, we count citations received in each year by the firm's *pre-deal* patents—those filed in the ten years before the deal—from patents that share no assignee with the firm, so that the firm's own post-deal patenting cannot enter. The value of a patent granted before the deal is fixed at the deal date, so if targets' technology were fading, outside citations to it would

decline relative to the controls' pre-deal patents. Second, we measure activity in each firm's technology fields: the number of US patents filed in each IPC subclass by assignees outside the sample, weighted by the firm's pre-deal subclass mix. Both outcomes are built on the same stacked firm-years as the headline, with the same censoring, and estimated with the headline specification.

Neither outcome declines (Table 5, Figure 3). Outside citations to the pre-deal stock are unchanged in the long run (+0.034, SE 0.070), as is the probability of receiving any (+0.013, SE 0.026). The probability has flat pre-periods (far-pre $p = 0.51$); the log count does not ($p = 0.01$). The confidence interval on the probability of any outside citation excludes a long-run decline as large as the headline effect on the probability of patenting (−0.127), and the interval on the log count excludes the headline citation-weighted decline (−0.135). Applicant-added and examiner-added citations both fail the far-pre test and we do not interpret them. Outside patenting in the targets' fields is also unchanged (+0.033, SE 0.057), with a flat pre-period ($p = 0.34$).

6.4 Product Safety

If PE ownership changes the pace of product development, a natural question is whether it also changes product safety. We measure safety with two FDA administrative records linked to the same firm clusters: device recalls and adverse-event reports from the MAUDE database, each aggregated to firm-year counts and estimated with Equation (3). Because both records are attributed through the FDA firm-cluster crosswalk, we estimate the safety outcomes on the *FDA-linked* subsample, pairs in which both the treated firm and its matched control appear in the FDA data (230 treated firms). Because recalls and adverse events are rare at the firm-year level, we read the extensive margin as our primary safety estimator.[9]

9. These outcomes are extremely sparse and heavy-tailed at the firm-year level: 97% of firm-years record no recall, and adverse-event counts are severely over-dispersed (a mean near four with a standard deviation above fifty and a maximum exceeding two thousand reports in a single firm-year). A rare binary event is well captured by the extensive-margin linear-probability specification, whereas count models are dominated by a handful of large firm-years and Poisson estimation is infeasible here—almost the entire panel separates on

Table 6 reports the ATTs and Figure 4 plots the event-study dynamics.

A firm's own safety reporting is not exogenous to ownership. After an acquisition the manufacturer may change how aggressively it initiates voluntary recalls or files adverse-event reports—to limit liability, warranty cost, or reputational exposure—so any measure that relies on manufacturer-initiated reports confounds a change in realized safety with a strategic change in disclosure prompted by the deal. In our data this endogeneity is visible: manufacturer-initiated recalls of the acquired firms' devices respond far more weakly than outside-initiated recalls of the same devices. We therefore take as our baseline the safety signals initiated outside the firm and not under its control—recalls initiated by a party other than the manufacturer (a distributor, importer, or the FDA) and adverse-event reports filed by users, patients, or distributors rather than by the firm (Panel A of Table 6). The manufacturer-initiated and all-reporter counterparts, which the firm can influence directly, are reported alongside as the endogenous channel (Panel B).

The outside-initiated evidence points to a modest deterioration in realized safety. Non-manufacturer recalls rise: the Medium-bin ATT on the probability of any outside-initiated recall is +2.5 percentage points (SE 1.0pp, significant at the 5% level) against a pre-deal base of 2%, with a matching increase on $\log(1 + \text{non-manufacturer recalls})$ (+0.032, SE 0.013, significant at the 5% level); in the long run the probability remains higher (+3.4pp, SE 1.8pp, significant at the 10% level) and the log coefficient is positive but not significant (+0.038). Non-manufacturer adverse-event reports, by contrast, show no significant change at any horizon. Because these signals do not depend on the acquired firm's own reporting, the recall increase is unlikely to reflect a strategic change in disclosure and instead points to a real, if modest, decline in realized product reliability that outside parties detect.

The manufacturer-initiated channel behaves as the endogeneity concern predicts. Recalls that the manufacturer initiates are essentially flat on the log margin at both horizons (Medium +0.001, Long −0.017, neither significant), a far weaker response than the outside-initiated

all-zero fixed-effect groups. We therefore lead with the extensive margin and report the $\log(1 + \cdot)$ intensive margin alongside as secondary.

recall increase, so the acquired firms do not initiate more of their own recalls even as outside parties do. Total adverse-event reports across all reporters rise more sharply in the medium run, but this increase is not robust to stripping out manufacturer-filed reports—the non-manufacturer adverse-event series above does not move. Taken together, the key safety finding is the increase in the probability of outside-initiated recalls.

6.5 Heterogeneity by Investor Type: Healthcare-Specialized vs. Other

We separately estimate Equation (3) on the subset of deals whose investor pool includes at least one healthcare-focused investor (HC-specialized) versus deals with no HC-focused investor. An investor is classified as healthcare-focused if a majority of its 1995–2019 investments are in healthcare companies; the deal-level indicator equals one if any such investor participates. It is computed from the investors that PitchBook records on each target's anchor deal and is defined for 446 of the 470 active treated firms; PitchBook records no investor on the anchor deal of the other 24, which are excluded. Healthcare-specialized targets patent somewhat less before the deal (1.64 against 2.51 patents per year) but are as likely to patent and file as many FDA submissions, and the two subsets differ little on the share of deals classified as LBOs (69% vs. 68%); Table 7 reports the full pre-deal comparison.

Figure 5 and Table 8 Panel A show the patent estimates by HC subset. Long-bin declines are larger in HC-specialized deals: -0.207 (SE 0.061, significant at the 0.1% level) on log1p patents and -16.8pp (SE 3.6pp, significant at the 0.1% level) on the extensive margin, against a log-count estimate near zero on the Non-HC subset (-0.011, SE 0.090) and a smaller extensive-margin decline that is not significant (-7.2pp, SE 5.5pp). The citation-weighted decline is likewise larger in the HC subset. The difference between the two subsets is significant at the 10% level on log patents (-0.197, SE 0.107) but not on the extensive margin; measured against each group's own pre-deal trend, the log-patent difference is -0.332 (SE 0.160), significant at the 5% level. The decline is thus concentrated in healthcare-specialized deals;

in the other deals neither margin falls significantly.

The Healthcare-Specialized and LBO indicators are essentially uncorrelated (Pearson $r = 0.01$ in the active sample), so the two cuts tag substantially different deals. Table 8 Panel B reports the effect within HC-specialized deals by LBO type. On both patent margins the decline is significant in the LBO subset (-0.261 on log patents and -21.6pp on the probability of any patent, both significant at the 0.1% level), but not in the non-LBO subset. The FDA estimates are imprecise: no cell is significant at the 5% level.

The next section explores channels that can help explain what drives the decline in long-run innovation after PE ownership, and why it appears to be concentrated in deals involving healthcare-specialized investors.

7 Potential Mechanisms

This section explores potential mechanisms that help explain the effects on long-run patenting and product innovation discussed in Section 6.

7.1 Technological Overlap With Other Portfolio Firms

A target patents in several technology classes (CPC main groups), and in some of them its PE sponsor's other portfolio companies were already active before the deal. We call these classes *overlapping* and compare the target's patenting in them with its patenting in its other classes, relative to the same comparison for its matched controls, which inherit the target's overlapping classes. We estimate the following triple difference model:

$$y_{isct} = \beta_{\mathrm{M}}\, T_{is}\, \mathrm{Ov}_c\, \mathrm{Med}_t + \beta_{\mathrm{L}}\, T_{is}\, \mathrm{Ov}_c\, \mathrm{Long}_t + \gamma_{\mathrm{M}}\, \mathrm{Ov}_c\, \mathrm{Med}_t + \gamma_{\mathrm{L}}\, \mathrm{Ov}_c\, \mathrm{Long}_t + \alpha_{ist} + \alpha_{isc} + \alpha_{ct} + \varepsilon_{isct}, \tag{4}$$

where y_{isct} is patenting by firm i in stack s, class c, and year t, α_{ist} are firm-by-stack-by-year effects, α_{isc} firm-by-stack-by-class effects, and α_{ct} class-by-year effects. The firm-by-stack-by-year effects absorb everything about the firm in a given year: its overall response to PE ownership, its lifecycle, and whatever led the sponsor to acquire it. Selection of targets on firm characteristics, observed or not, therefore cannot drive β_{L}, and the class-by-year effects absorb field-wide trends. What remains is whether, after the deal, the target shifts its patenting away from the fields its sponsor's other companies occupy, denoted by Ov_c. The portfolio is the set of other device companies the anchor-deal sponsors had invested in by the date of the deal, including companies they had since sold (investments are dated to the day, so one made later in the deal year counts as a later investment). Classes are current CPC main groups from PatentsView, which code every patent in the current version of the scheme, so a class means the same field in every year of the panel. Fields the target shares only with companies its sponsor first invested in after the deal could reflect the deal itself, so we set them aside for the target and its matched controls alike: they enter neither as overlapping nor as comparison classes. We drop the 27 stacks in which the target or a matched control has a patent record that incorporates a company acquired by the firm, because such a record places the firm in the acquired company's fields before the deal.

The design covers 443 targets. The sponsor is recorded for 325 of them; for the other 118 the anchor deal has no buyout or growth investor recorded in PitchBook, so these targets have no overlapping classes and do not identify the estimates. The estimates are identified from the 100 targets who have both overlapping and non-overlapping classes, each observed from up to ten years before the deal to twelve years after.

Table 9 and Figure 6 report the triple difference results. Before the deal, patenting in overlapping and other classes moves in parallel: no event-study coefficient is significant and joint tests do not reject on $t \in [-10, -6]$. After the deal the two diverge with the same delay as the firm-level decline. Over years 7–12, the probability of any patent in an overlapping class falls by 7.3 percentage points more than in the firm's other classes (SE 3.1pp, significant

at the 5% level), and patenting falls by 0.083 log points more (SE 0.038, significant at the 5% level). The second row of Panel A shows that counting the add-on platform as part of the portfolio leaves the estimates unchanged.

The gap between overlapping and non-overlapping classes is specific to the target's own sponsor. When each target is also assigned the portfolio of another sponsor that invested in the same year, that portfolio's fields add +0.001 on average to the long-run coefficient on the probability of patenting, while the real sponsor's fields keep their effect (−0.073 on average); the randomization-inference p-value—the share of 500 such draws in which the placebo coefficient is at least as negative as the actual estimate—is 0.016. The corresponding p-value for log patents is 0.044.

The gap is estimated precisely only for leveraged buyouts (Panel B): the long-run coefficient on the probability of patenting is −0.071 for LBO targets (significant at the 5% level) and −0.052 for the 27 non-LBO targets with overlapping classes (not significant). The difference (−0.020, SE 0.068) is not significant. The pattern is consistent with the role of control in the framework: a buyout sponsor controls the research budget, whereas a growth-equity investor typically holds a minority stake with board representation, so $\omega_{\mathrm{LBO}} \geq \omega_{\mathrm{G}}$. We read it as consistent with that ordering, not as evidence that growth-equity investors do not cut: the non-LBO estimate rests on 27 targets and as a result it is imprecisely estimated.

Panel C shows that the gap is concentrated in deals backed by healthcare-specialized investors and it differs significantly by investor type. Their portfolios contain more patenting device companies (3.6 per target, against 0.7 for other investors), so their targets far more often have overlapping classes: 42% of these targets have at least one against 13% for the targets of other investors. The probability of patenting in an overlapping class falls by 7.6 percentage points for the targets of healthcare-specialized investors with such classes (SE 3.0pp, significant at the 5% level); for the targets of other investors the estimate is +15.0pp (SE 10.5pp, not significant), and the difference (−22.6pp, SE 11.0pp) is significant at the 5% level; log patents show the same pattern (−0.092 and +0.140, a difference of −0.232, SE

0.111).

The within-firm evidence therefore identifies a channel through which PE ownership reduces innovation unrelated to acquisition debt: the target's own research retreats from the fields in which its sponsor already holds other companies. Restricting the overlap to companies the sponsor still held at the deal, and setting aside fields shared only with companies it had already sold, gives similar estimates (-0.098 and -0.086).

The design is also a check on the causal reading of the firm-level decline. Because the firm-by-year effects absorb everything about the firm—its characteristics, observed or not, and whatever led the sponsor to acquire it—a decline that follows the sponsor's portfolio within the firm cannot come from the selection of targets on firm characteristics, the standing concern with the matched comparison. The same effects absorb the firm's overall decline, so the design shows where within the firm the decline falls rather than how large it is; the firm-level estimates measure its size. The design cannot tell whether the sponsor curtails the research because it duplicates a sibling's or because it would compete with a sibling's products; each is a form of managing research across a portfolio, and each requires influence over the target's research budget, which a buyout sponsor has in full and a minority growth-equity investor only through its board seats and investor rights.

7.2 Reduced Hiring of Inventors

To probe whether the firm-level quality decline operates through the firm's inventor team, we define each firm's *incumbent cohort* as the inventors listed on its device (A61) patents in the ten years before the deal anchor, constructed symmetrically for treated and matched-control firms (controls inherit the treated firm's anchor year) and followed forward under the same own-exit censoring. Table 10 and Figure 7 report the results.

On this device-active subsample the firm's device output falls in the long run (-0.107, SE 0.057), and part of the decline registers in the incumbent team. The cohort's citation-weighted at-firm output falls (-0.058 in the medium run and -0.071 in the long run, SE 0.027 and

0.035, both significant at the 5% level), while its headcount still producing at the firm does not fall significantly (−0.034, SE 0.028, in the medium run; −0.037, SE 0.037, in the long run) and its at-firm output is lower but not significantly so in the long run (−0.028, SE 0.034; Figure 7). The larger margin is replacement: the probability that the firm hires an inventor new to the firm in a given year falls 6.0 percentage points in the long run (SE 2.9pp) against a pre-deal base of 32%. Incumbent inventors are no more likely to stop patenting or to move to another firm (Panel B of Table 10). We therefore read the inventor evidence as reduced hiring of new inventors, together with lower quality-weighted output among the incumbents who stay, rather than a loss of incumbents—consistent with a firm that stops staffing new projects. The fall in hiring is concentrated in the fields the sponsor's other companies occupy.

A stricter individual-level design re-matches each treated incumbent inventor to control incumbents on the inventor's own pre-deal trajectory, referenced at the same deal date (Panel C of Table 10). Under this design the per-inventor effects are small: per-inventor citation-weighted output falls 0.025 in the medium run (significant at the 10% level) and is unchanged in the long run (+0.001), against a pre-deal mean of 0.92, and the probability that an incumbent inventor patents at all falls 2.6 percentage points in the medium run (significant at the 5% level) and 2.1 in the long run (not significant) against a base of 40%. These effects are too small to account for the firm-level decline, which points again to fewer new inventors rather than lower productivity among those who remain.

7.3 Discussion

Overall, four features of the evidence are jointly informative. First, the patent and FDA declines are concentrated in the Long $[7, 12]$ bin; the Medium $[0, 6]$ effects are typically small and statistically insignificant. Second, the decline cannot be attributed to features specific to leveraged buyouts: growth-equity deals show declines of a similar size, although not statistically significant on their own. Third, across payoff horizon, technological breadth, and citation impact, and in a continuous semantic-embedding space, the decline is essentially

uniform. Fourth, within the firm, the decline concentrates in the fields the sponsor's other companies occupy, and not in the fields of other sponsors' companies.

The channel that fits these features best is the sponsor's management of research across its portfolio: the sponsor weighs the loss a target's project would impose on its other companies and stops funding projects in their fields. The within-firm concentration is precise only for buyouts, where the sponsor controls the research budget, so the channel would explain a decline in growth-equity deals only to the extent that board representation gives growth investors influence over research; other channels may contribute there. The change operates at the project-selection stage; projects already in the pipeline continue to generate patents for several years, so the output decline appears with a lag set by the device R&D cycle, even though the ownership change is delivered at closing. The platform-and-add-on strategies common in health care, software, and industrial technology create exactly such portfolios, and the killer-acquisition evidence from pharmaceuticals (Cunningham, Ederer, and Ma 2021) shows the same logic operating within a single acquirer.

8 Conclusion

This paper provides the first empirical evidence, to our knowledge, on the effects of private equity (PE) investment in medical device firms. We build a novel dataset to examine the effects on 470 firms that received PE investments between 1995 and 2020, approximately 10% of all device firms. Using a matched stacked difference-in-differences design, we find that PE investment results in large long-run declines in innovation output, measured using patents and FDA product submissions. The decline is primarily on the extensive margin, with targeted firms less likely to file a new patent or device for approval.

The temporal pattern—small effects in the first six years after the deal, a substantial decline in years 7–12—is consistent with a change in which projects the firm starts, delivered at closing but visible in output only after the multi-year device R&D cycle. The decline is

neutral in horizon, breadth, and impact but concentrated in the fields of the sponsor's other companies, the the combination one would expect from a sponsor that manages research across its portfolio. The mechanism is in principle present wherever a financial sponsor holds several companies whose research overlaps, as in the platform strategies common in health care and technology. The medical-device setting is useful because both invention and commercialization leave public, time-stamped records.

Although we have shed some light on the effects of PE investment in medical device firms, several avenues remain open for future work on this topic. We allude above to our limitation in precisely delineating the role of different mechanisms that operate simultaneously in this setting. With additional data on research inputs and variation in investment time horizons, one may have more success in separately quantifying the roles of leverage and finite-horizon truncation channels. We document the effect on product development, but another important aspect is product quality. Since these products are used to deliver health care services, changes in quality could meaningfully affect patient welfare. Examining the downstream effects on patient health will help quantify one important avenue through which PE investment in medical device technologies affects society.

References

Antón, Miguel, Florian Ederer, Mireia Giné, and Martin C. Schmalz. 2025. "Innovation: The Bright Side of Common Ownership?" *Management Science* 71 (5): 3713–3733.

Baker, Andrew C, David F Larcker, and Charles CY Wang. 2022. "How much should we trust staggered difference-in-differences estimates?" *Journal of financial economics* 144 (2): 370–395.

Bernstein, Shai, Josh Lerner, Morten Sorensen, and Per Strömberg. 2017. "Private Equity and Industry Performance." *Management Science* 63 (4): 1198–1213.

Bernstein, Shai, and Albert Sheen. 2016. "The Operational Consequences of Private Equity Buyouts: Evidence from the Restaurant Industry." *Review of Financial Studies* 29 (9): 2387–2418.

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

Budish, Eric, Benjamin N. Roin, and Heidi Williams. 2015. "Do Firms Underinvest in Long-Term Research? Evidence from Cancer Clinical Trials." *American Economic Review* 105 (7): 2044–2085.

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

Clemens, Jeffrey, and Joshua D. Gottlieb. 2014. "Do Physicians' Financial Incentives Affect Medical Treatment and Patient Health?" *American Economic Review* 104 (4): 1320–1349.

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

Cumming, Douglas, Rejo Peter, and Monika Tarsalewska. 2020. "Public-to-private buyouts and innovation." *British Journal of Management* 31 (4): 811–829.

Cunningham, Colleen, Florian Ederer, and Song Ma. 2021. "Killer Acquisitions." *Journal of Political Economy* 129 (3): 649–702.

Cutler, David M, Allison B Rosen, and Sandeep Vijan. 2006. "The Value of Medical Spending in the United States, 1960–2000." *New England Journal of Medicine* 355 (9): 920–927.

Davis, Steven J., John Haltiwanger, Kyle Handley, Ron Jarmin, Josh Lerner, and Javier Miranda. 2014. "Private Equity, Jobs, and Productivity." *American Economic Review* 104 (12): 3956–3990.

Donahoe, Gerald. 2021. *Estimates of medical device spending in the United States.* Technical report. Advanced Medical Technology Association Report, June.

Dranove, David, Craig Garthwaite, Christopher Heard, and Bingxiao Wu. 2022. "The Economics of Medical Procedure Innovation." *Journal of Health Economics* 81:102549.

Finkelstein, Amy. 2004. "Static and dynamic effects of health policy: Evidence from the vaccine industry." *The Quarterly Journal of Economics* 119 (2): 527–564.

Grennan, Matthew, and Robert J. Town. 2020. "Regulating Innovation with Uncertain Quality: Information, Risk, and Access in Medical Devices." *American Economic Review* 110 (1): 120–161.

Gupta, Atul, Sabrina T. Howell, Constantine Yannelis, and Abhinav Gupta. 2024. "Owner Incentives and Performance in Healthcare: Private Equity Investment in Nursing Homes." *Review of Financial Studies* 37 (4): 1029–1077.

Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg. 2001. "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools." *NBER Working Paper No. 8498.*

Jaffe, Adam B. 1989. "Characterizing the "Technological Position" of Firms, with Application to Quantifying Technological Opportunity and Research Spillovers." *Research Policy* 18 (2): 87–97.

Ji, Yunan, and Parker Rogers. 2024. *The Long-Run Impacts of Regulated Price Cuts: Evidence from Medicare.* NBER Working Paper No. 33083. National Bureau of Economic Research, October.

Kaplan, Steven N., and Per Strömberg. 2009. "Leveraged Buyouts and Private Equity." *Journal of Economic Perspectives* 23 (1): 121–146.

Lerner, Josh, Morten Sorensen, and Per Strömberg. 2011. "Private Equity and Long-Run Investment: The Case of Innovation." *Journal of Finance* 66 (2): 445–477.

López, Ángel L., and Xavier Vives. 2019. "Overlapping Ownership, R&D Spillovers, and Antitrust Policy." *Journal of Political Economy* 127 (5): 2394–2437.

Makower, Josh, Aabed Meer, and Lyn Denend. 2010. *FDA Impact on U.S. Medical Technology Innovation: A Survey of over 200 Medical Technology Companies.* Technical report. Stanford University, Industry Report, November.

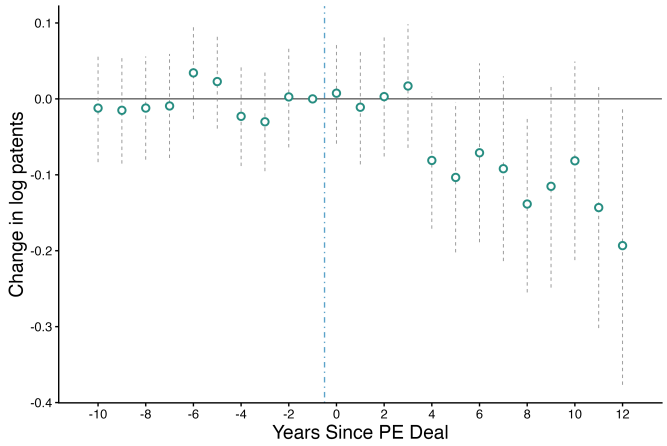
Rogers, Parker. 2024. "Regulating the Innovators: Approval Costs and Innovation in Medical Technologies." Working Paper, Indiana University.

Singh, Yashaswini, Megha Reddy, Irene Papanicolas, and Richard Scheffler. 2026. "Private equity investments in health care in OECD countries: an exploratory analysis." *Health Economics, Policy and Law,* 1–28.

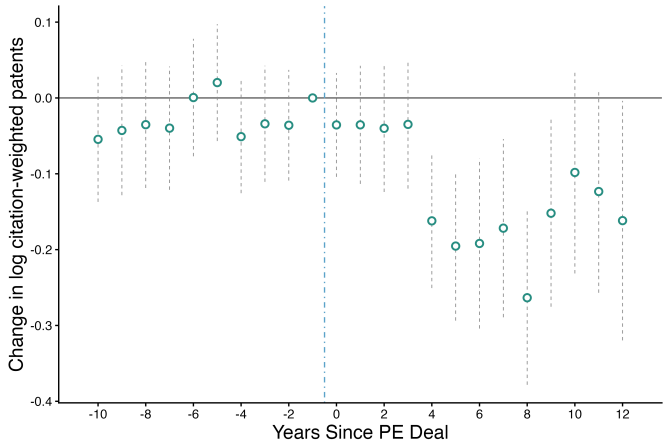
Stern, Ariel Dora. 2017. "Innovation Under Regulatory Uncertainty: Evidence from Medical Technology." *Journal of Public Economics* 145:181–200.

Williams, Heidi L. 2013. "Intellectual Property Rights and Innovation: Evidence from the Human Genome." *Journal of Political Economy* 121 (1): 1–27.

(a) (Log) Patent Count



(b) Citation-Weighted (log) Count



(c) Probability of Any Patent

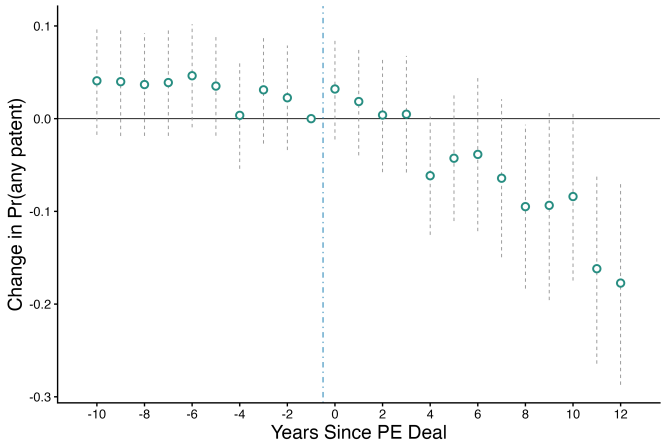
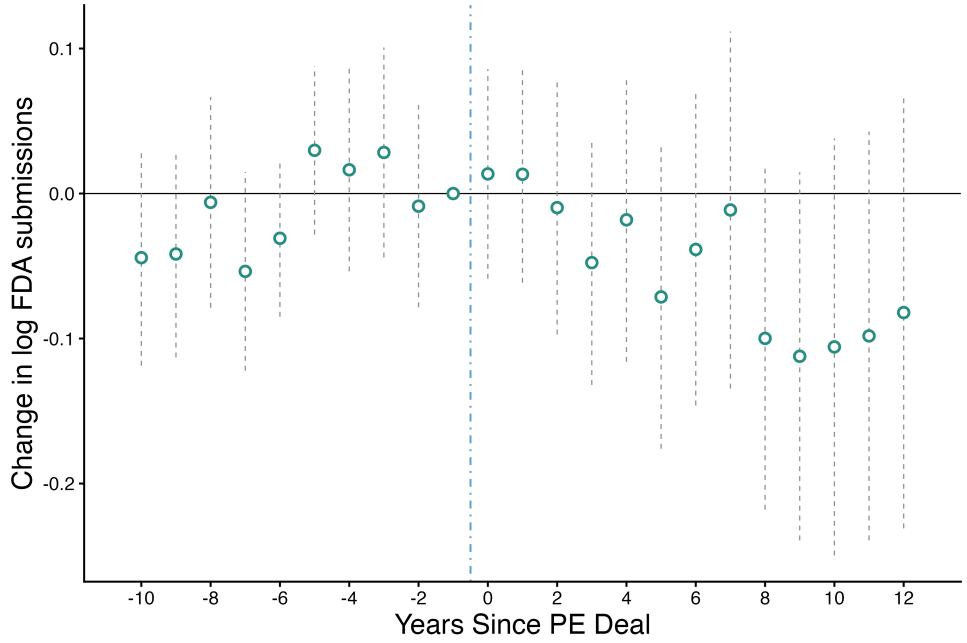


Figure 1: Effects on Total Patenting

Note: This figure plots dynamic effect coefficients from the event-study version of Equation (3). The outcome in Panel (b) is patent counts weighted by 5-year citation counts following the approach in Hall, Jaffe, and Trajtenberg (2001). The estimating sample includes 470 unique treated firms matched to 879 unique control firms (1,397 paired stacks). The coefficient β_{-1} is omitted and serves as the reference period. Standard errors are clustered at the firm level. Dashed grey lines indicate 95% confidence intervals.

(a) (Log) FDA Submissions



(b) Probability of Any FDA Submission

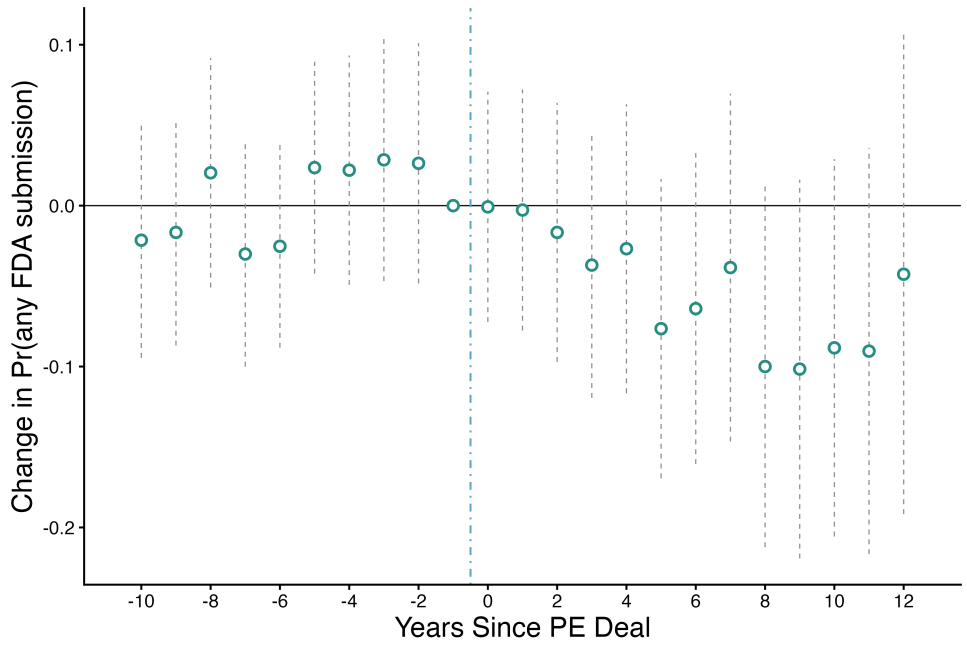
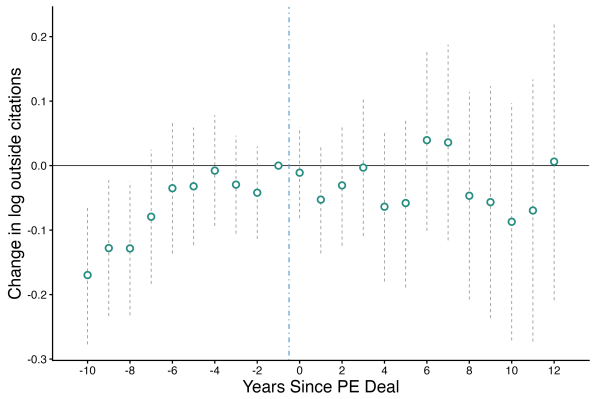


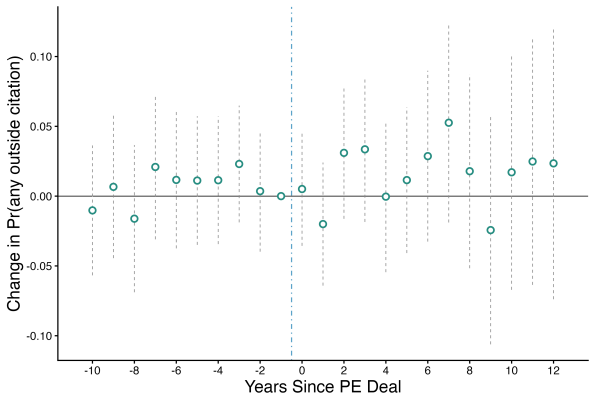
Figure 2: Effects on FDA Regulatory Submissions

Note: This figure plots the β_k coefficient estimates from the event-study version of Equation (3) for log(1 + total FDA submissions per firm-year), pooling 510(k) and PMA submissions. The estimating sample is the FDA-linked subsample: 230 unique treated firms matched to 357 unique control firms (530 paired stacks). The coefficient β_{-1} is omitted and serves as the reference period. Standard errors are clustered at the firm level. Dashed grey lines indicate 95% confidence intervals.

(a) Outside citations to pre-deal patents (log)



(b) Probability of any outside citation



(c) Outside patenting in the firm's fields (log)

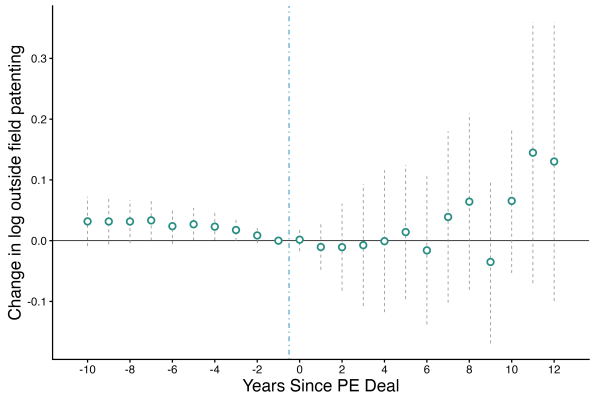


Figure 3: Negative-Control Event Studies

Note: Event-study coefficients β_k from Equation (3) for the negative-control outcomes in Table 5, on the patent-active sample. β_{-1} is the omitted reference period. Dashed grey lines are 95% confidence intervals.

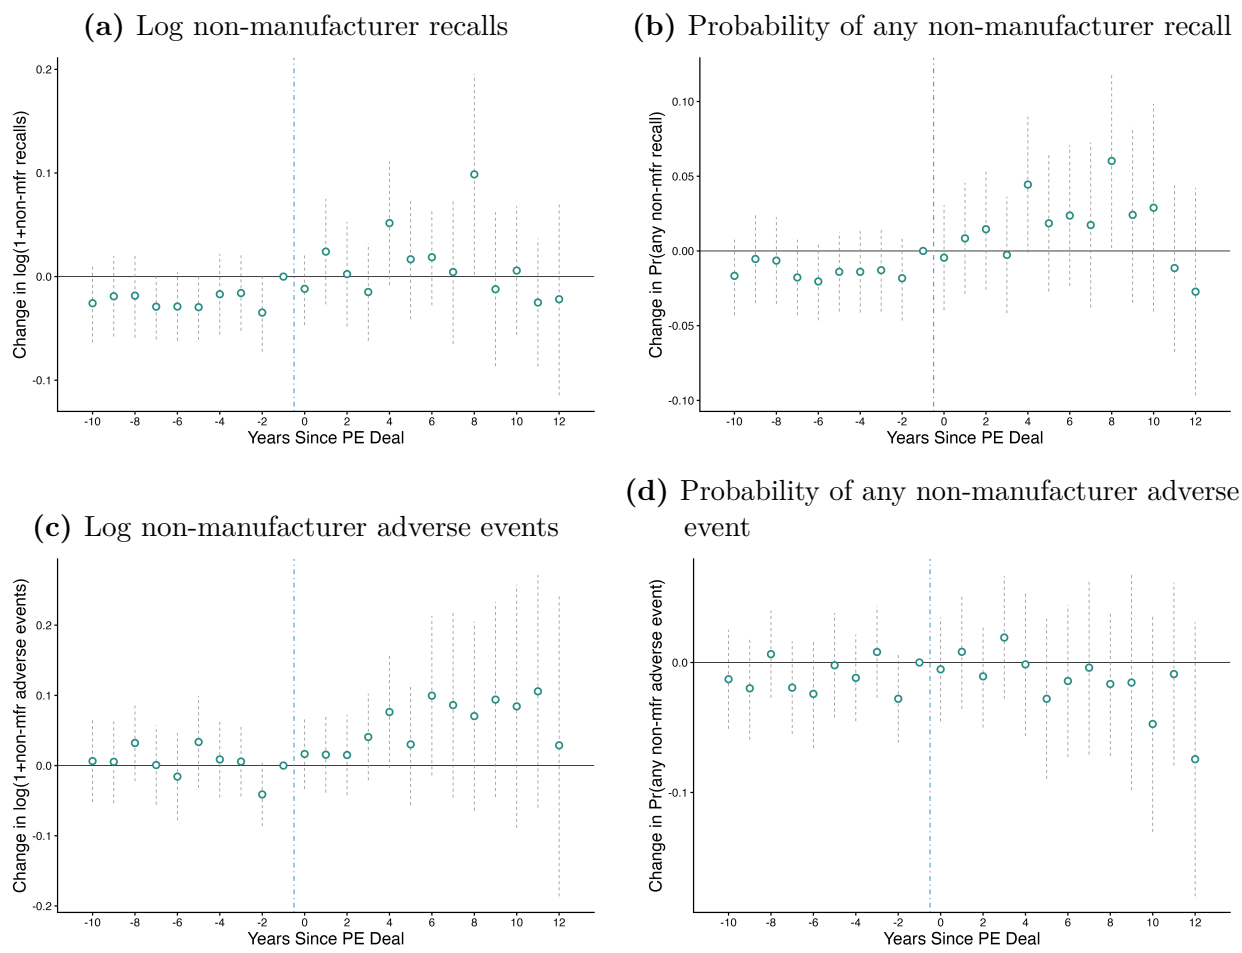


Figure 4: Effects on Product Safety (Outside-Initiated Baseline)

Note: Event-study coefficients β_k from Equation (3) for the baseline outside-initiated safety outcomes—non-manufacturer-initiated device recalls (top row) and non-manufacturer-filed MAUDE adverse-event reports (bottom row)—on the FDA-linked subsample (230 treated / 357 control), pairs in which both the treated firm and its matched control are observable in the FDA data, under the corrected FDA cluster attribution and own-exit censoring; adverse-event firm-years after the 2019 MAUDE source endpoint are dropped. These signals are initiated outside the acquired firm and so are not contaminated by any strategic change in its own reporting after the deal. β_{-1} is the omitted reference period. Dashed grey lines are 95% confidence intervals.

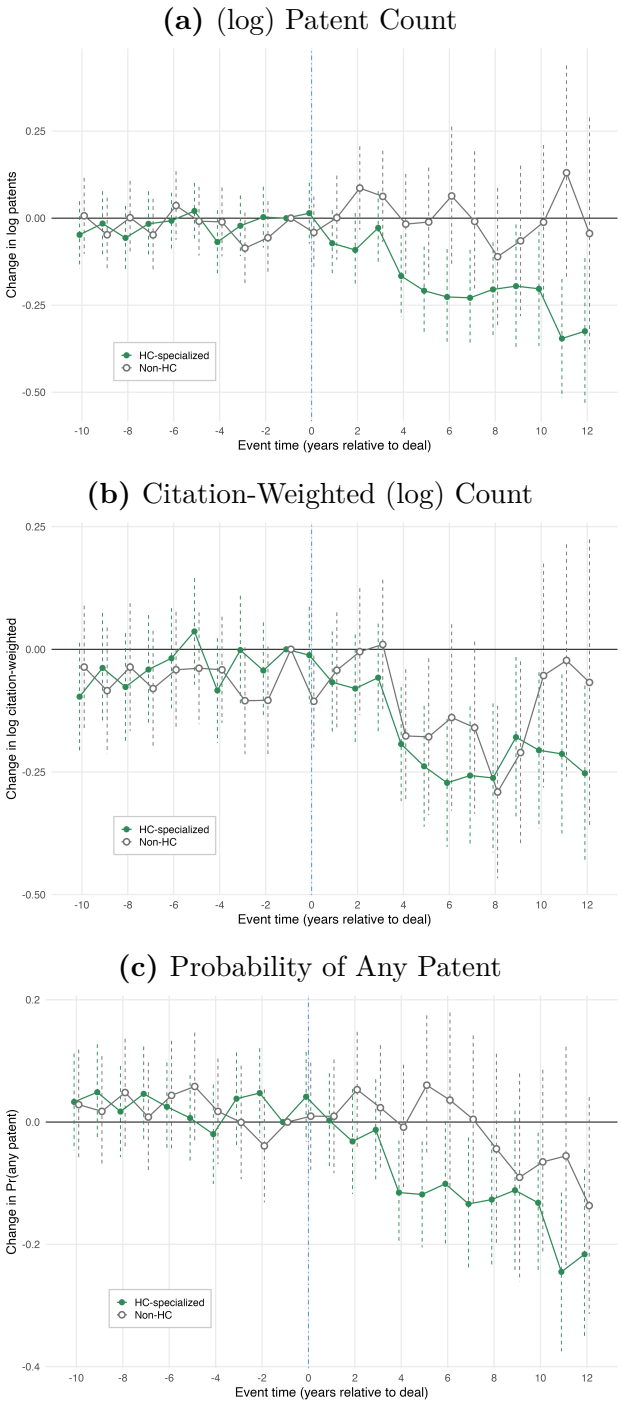
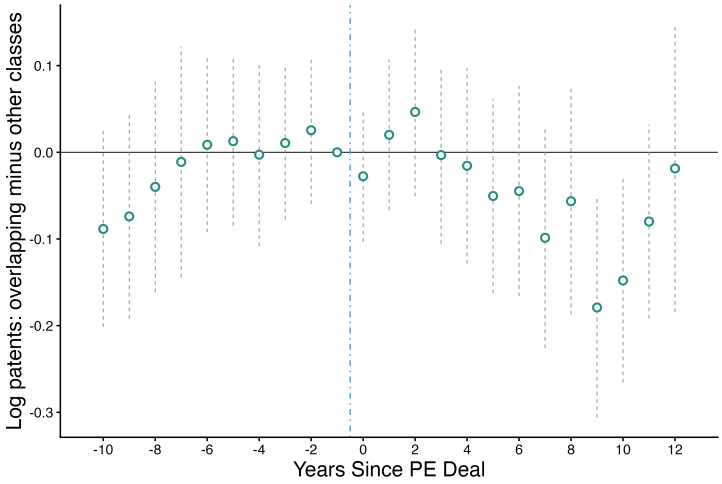


Figure 5: Heterogeneous Effects on Patenting by Healthcare-Specialized Investor

Note: This figure plots the β_k coefficient estimates from the event-study version of Equation (3) for total patenting (log1p, citation-weighted log1p, and extensive margin), estimated separately on the Healthcare-Specialized subset (deal has at least one HC-focused investor; 756 stacks) and the Non-HC subset (569 stacks). The coefficient β_{-1} is omitted and serves as the reference period. Standard errors are clustered at the firm level. Dashed grey lines indicate 95% confidence intervals.

(a) Triple difference: log patents in the class



(b) Triple difference: Pr(any patent in the class)

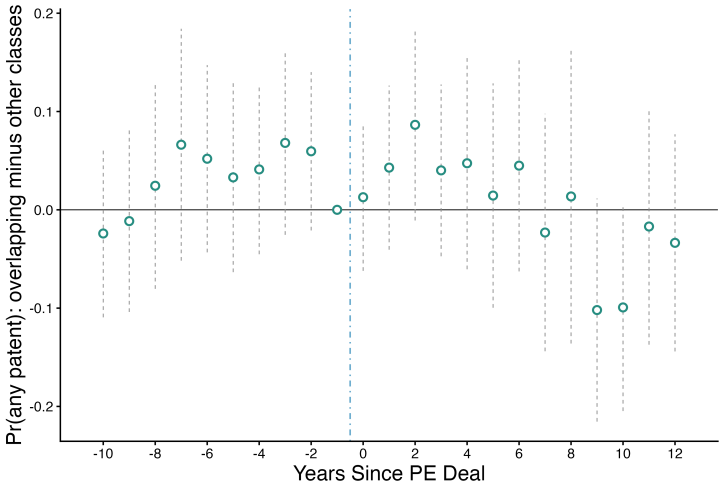
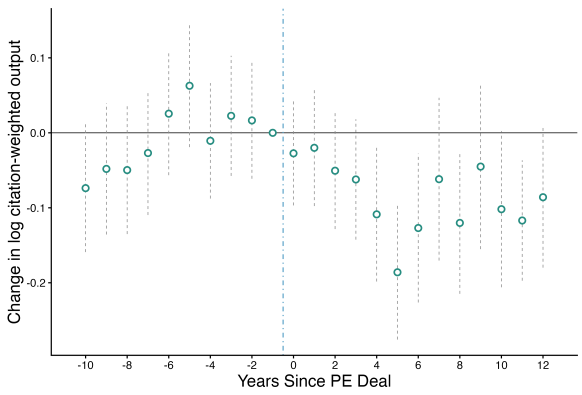
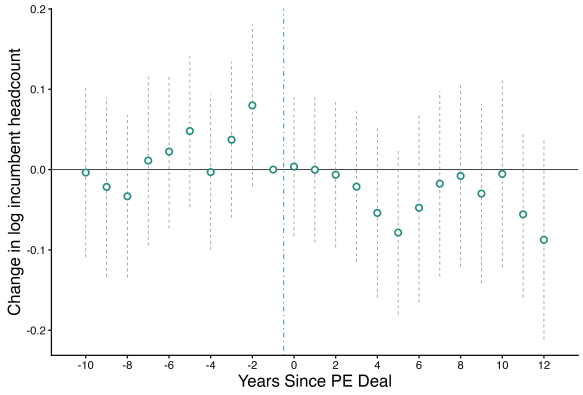


Figure 6: Within-Firm Evidence: Triple-Difference Event Studies

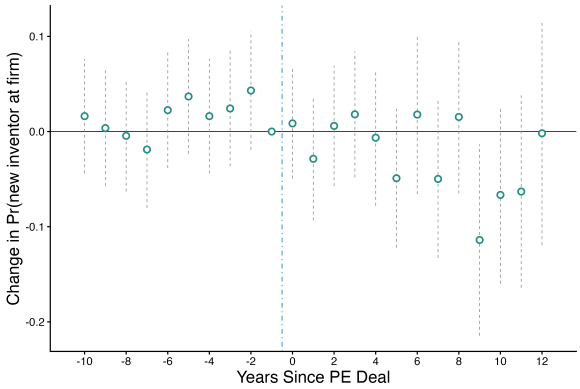
Note: Panels (a) and (b) plot the triple-difference event study: coefficients on treated × overlapping class × event time from the event-study version of Equation (4), with $t = -1$ as the reference period, firm-by-stack-by-year, firm-by-stack-by-class, and class-by-year fixed effects, and standard errors clustered at the firm level; patent years through 2020. A class is overlapping if a company the target's PE sponsor had invested in by the date of the deal patented in it during the ten years before the deal. Dashed lines are 95% confidence intervals.



(a) Incumbent-cohort citation-weighted output



(b) Incumbent headcount at the firm



(c) New-inventor hiring (Pr)

Figure 7: Inventor-Level Mechanism: Incumbent Cohort and Hiring

Note: Event-study coefficients on the matched sample (371 device-active episodes). The incumbent cohort is fixed at $[-10, -1]$ and defined symmetrically for treated and matched controls, with own-exit censoring. Panels (a)–(b) plot $\log(1 + y)$ for the cohort's citation-weighted at-firm device output and headcount; Panel (c) plots the linear-probability event study for hiring at least one new-to-the-firm inventor in the event year. Pre-period coefficients are flat. The incumbent team's quality-weighted output and headcount decline only slightly and imprecisely, while the probability of hiring a new inventor falls in the long run. β_{-1} is the omitted reference period. Dashed grey lines are 95% confidence intervals.

Table 1: Treated Firms vs. Matched Controls

	Treated ($N = 470$)	Matched controls ($N = 1{,}397$ stacks)	Difference
Mean patents per year	2.08	1.50	+0.58
Pr(any patent in window)	0.89	0.87	+0.01
Mean A61 patents per year	1.48	1.01	+0.46
Pr(any A61 patent in window)	0.70	0.69	+0.01
Mean FDA submissions per year	0.29	0.22	+0.07
Pr(any FDA submission in window)	0.33	0.27	+0.05
Pr(any VC round in window)	0.17	0.20	−0.03
Mean firm age (patenting tenure, years)	12.3	11.4	+0.9

Note: Pre-deal observable characteristics: PE-acquired (treated) firms versus their matched never-PE-acquired controls in the headline patent-active sample. Same row definitions as Table A2; pre-deal window is five years before the treated firm's deal anchor (matched controls inherit the same anchor year).

Table 2: Effects on Patenting

	Log patents	Citation-weighted (log)	Extensive (any patent)
Medium $[0, 6]$	-0.018	-0.051^{+}	-0.032^{+}
	(0.032)	(0.030)	(0.017)
Long $[7, 12]$	-0.111^{*}	-0.135^{**}	-0.127^{***}
	(0.052)	(0.045)	(0.031)
Pre-period treated mean	2.08 patents/yr	2.84 (HJT-norm cites)	Pr(any) = 0.47
Joint pre-trend p ($t \in [-10, -6]$)	0.59	0.63	0.67

Note: Medium $[0, 6]$ and Long $[7, 12]$ ATTs on patent output. Each cell reports the coefficient (standard error) from a single regression with both Medium and Long indicators, firm-by-stack and stack-by-year fixed effects, and standard errors clustered at the firm level. Sample: 470 unique treated firms × 879 unique control firms (1,397 paired stacks). Each outcome is used through its data horizon (Section 4.1): patents and outcomes built from them through filing year 2020, FDA submissions and recalls through 2024, and venture financing through 2025. Significance: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The bottom row reports the p-value from a joint Wald test that the three event-study pre-trend coefficients at $t \in [-10, -6]$ all equal zero (the matching window is $[-5, -1]$, so these far-pre coefficients are entirely out-of-sample for matching).

Table 3: Effects on FDA Product Applications

	Log FDA submissions	Extensive (any submission)
Medium $[0, 6]$	-0.005	-0.028
	(0.031)	(0.024)
Long $[7, 12]$	-0.065	-0.076^{*}
	(0.048)	(0.036)
Pre-period treated mean	0.63 submissions/yr	Pr(any) = 0.32
Joint pre-trend p ($t \in [-10, -6]$)	0.55	0.73

Note: Medium $[0, 6]$ and Long $[7, 12]$ ATTs on FDA regulatory submissions (510(k) plus PMA, fuzzy-deduped per the FDA-policies preprocessing convention). Each cell reports the coefficient (standard error) from a single regression with both Medium and Long indicators, firm-by-stack and stack-by-year fixed effects, and standard errors clustered at the firm level. FDA outcomes are estimated on the *FDA-linked* subsample—firms for which both the treated unit and its matched control are observable in the FDA data (i.e., appear in the FDA firm-cluster crosswalk), the FDA analog of the patent-active restriction applied to the patent outcomes (Section 4.4). Sample: 230 unique treated firms $\times$ 357 unique control firms (530 paired stacks). Each outcome is used through its data horizon (Section 4.1): patents and outcomes built from them through filing year 2020, FDA submissions and recalls through 2024, and venture financing through 2025. Significance: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The bottom row reports the p-value from a joint Wald test that the three event-study pre-trend coefficients at $t \in [-10, -6]$ all equal zero (the matching window is $[-5, -1]$, so these far-pre coefficients are entirely out-of-sample for matching).

Table 4: Robustness Checks

	Log patents	Cite-weighted (log)	Pr(Any patent)	Log FDA submissions	Pr(FDA submission)
A: Baseline					
Long $[7,12]$ ATT	-0.111^{*}	-0.135^{**}	-0.127^{***}	-0.065	-0.076^{*}
	(0.052)	(0.045)	(0.031)	(0.048)	(0.036)
Pre-period treated mean	2.08	2.84	0.47	0.59	0.32
# treated firms	470	470	470	230	230
# treated firms in Long bin	185	185	185	132	132
# paired stacks	1,397	1,397	1,397	530	530
B: Deals with the full post-period observed (patents: deals through 2008; FDA: through 2012)					
Long $[7,12]$ ATT	-0.091	-0.123^{*}	-0.116^{**}	-0.012	-0.037
	(0.064)	(0.055)	(0.038)	(0.059)	(0.043)
Pre-period treated mean	1.63	1.85	0.47	0.54	0.31
# treated firms	196	196	196	130	130
# treated firms in Long bin	119	119	119	71	71
# paired stacks	580	580	580	299	299
C: No patent filter sample					
Long $[7,12]$ ATT	-0.116^{**}	-0.105^{**}	-0.099^{***}	-0.035^{*}	-0.040^{**}
	(0.038)	(0.033)	(0.023)	(0.014)	(0.013)
Pre-period treated mean	1.52	2.07	0.35	0.28	0.13
# treated firms	647	647	647	647	647
# treated firms in Long bin	285	285	285	422	422
# paired stacks	1,941	1,941	1,941	1,941	1,941
D: Event-time window $[-12,12]$					
Long $[7,12]$ ATT	-0.109^{*}	-0.133^{**}	-0.125^{***}	-0.057	-0.074^{*}
	(0.053)	(0.046)	(0.032)	(0.049)	(0.036)
Pre-period treated mean	2.08	2.84	0.47	0.59	0.32
# treated firms	470	470	470	230	230
# treated firms in Long bin	185	185	185	132	132
# paired stacks	1,397	1,397	1,397	530	530
E: Event-time window $[-8,12]$					
Long $[7,12]$ ATT	-0.111^{*}	-0.137^{**}	-0.122^{***}	-0.072	-0.079^{*}
	(0.050)	(0.044)	(0.030)	(0.047)	(0.035)
Pre-period treated mean	2.08	2.84	0.47	0.59	0.32
# treated firms	470	470	470	230	230
# treated firms in Long bin	185	185	185	132	132
# paired stacks	1,397	1,397	1,397	530	530
F: No data-horizon cap					
Long $[7,12]$ ATT	-0.105^{***}	-0.145^{***}	-0.084^{***}	-0.041	-0.052^{*}
	(0.031)	(0.033)	(0.018)	(0.034)	(0.026)
Pre-period treated mean	2.08	2.84	0.47	0.59	0.32
# treated firms	470	470	470	230	230
# treated firms in Long bin	338	338	338	153	153
# paired stacks	1,397	1,397	1,397	530	530

Note: Each cell reports the Long $[7,12]$ ATT (standard error) on the indicated outcome from a single regression with both Medium $[0,6]$ and Long $[7,12]$ post-deal indicators, firm-by-stack and stack-by-year fixed effects, and standard errors clustered at the firm level (Equation 3). Panel A reports the headline patent-active sample (both treated and matched control filed at least one US patent in $[-10,-1]$). Every outcome is used through its data horizon (patents 2020, FDA submissions 2024; Section 4); "# treated firms in Long bin" counts the treated firms with at least one firm-year in $[7,12]$ inside the horizon. Panel B restricts to deals whose full post-period $[0,12]$ lies inside the horizon: deals through 2008 for the patent outcomes and through 2012 for the FDA outcomes. Panel C estimates on the no-filter sample, the full unified $K=3$ pool with no patent-active restriction. Panel D re-estimates Panel A on the event-time window $[-12,12]$. Panel E re-estimates Panel A without the data-horizon cap, keeping every calendar year the panel covers. Significance: $^{+}$ $p<0.10$, * $p<0.05$, ** $p<0.01$, *** $p<0.001$. The Panel C estimates should be read as a conservative lower bound in absolute value: roughly one-third of the no-filter sample consists of firms with zero patents and zero FDA submissions in the pre-period, which mechanically cannot generate negative post-deal coefficients (their outcome can only remain at zero or rise) and therefore attenuate the estimated average treatment effect toward zero. The Panel A patent-active filter excludes these firms by design, isolating the effect of PE acquisition on firms with prior innovation capacity.

Table 5: Negative-Control Outcomes

Outcome	Pre-period treated mean	Medium $[0, 6]$	Long $[7, 12]$	Far-pre p
Outside citations to pre-deal patents (log)	21.73	+0.037	+0.034	0.01
		(0.042)	(0.070)	
Pr(any outside citation to pre-deal patents)	0.60	+0.004	+0.013	0.51
		(0.014)	(0.026)	
Examiner-added outside citations (log)	2.66	$+0.105^{***}$	$+0.103^{**}$	0.00
		(0.024)	(0.033)	
Applicant-added outside citations (log)	8.85	+0.043	+0.062	0.02
		(0.041)	(0.083)	
Outside patenting in the firm's pre-deal fields (log)	3,187.01	−0.027	+0.033	0.34
		(0.031)	(0.057)	

Note: Headline specification (Equation 3) on the patent-active sample with firm-by-stack and stack-by-year fixed effects and firm-clustered standard errors. Outside citations are citations received in the calendar year (citing patent's filing year) by the firm's patents filed in $[D - 10, D - 1]$, excluding citing patents that share any assignee with the firm; count outcomes enter as $\log(1 + \cdot)$. Field activity is $\sum_s w_s N_{st}$, where w_s is the firm's pre-deal share of patents in IPC subclass s and N_{st} is the number of US patents filed in s in year t by assignees outside the sample. The pre-period mean is the treated firms' raw mean over $[-5, -1]$ (a share for the extensive row). "Far-pre p" is the joint Wald test that the event-study coefficients at $t \in [-10, -6]$ are zero. Each outcome is used through its data horizon (Section 4.1): patents and outcomes built from them through filing year 2020, FDA submissions and recalls through 2024, and venture financing through 2025. Significance: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 6: Effects on Product Safety

	Pre-mean	Medium $[0, 6]$	Long $[7, 12]$
A. Baseline: outside-initiated safety signals			
Log non-mfr recalls	0.09	$+0.032^{*}$	+0.038
		(0.013)	(0.023)
Pr(any non-mfr recall)	0.02	$+0.025^{*}$	$+0.034^{+}$
		(0.010)	(0.018)
Log non-mfr adverse events	0.19	+0.028	+0.073
		(0.021)	(0.065)
Pr(any non-mfr adverse event)	0.06	+0.008	−0.011
		(0.013)	(0.028)
B. Manufacturer-initiated / all-reporter channels (endogenous to the deal)			
Log manufacturer recalls	0.08	+0.001	−0.017
		(0.016)	(0.023)
Pr(any manufacturer recall)	0.02	+0.011	−0.014
		(0.011)	(0.013)
Log adverse events, all reporters	5.01	$+0.141^{*}$	+0.131
		(0.060)	(0.124)
Pr(any serious adverse event)	0.04	+0.017	+0.031
		(0.011)	(0.024)

Note: Medium $[0, 6]$ and Long $[7, 12]$ ATTs on product-safety outcomes, same estimator and own-exit censoring as Table 2. Recalls and adverse events are attributed to firms through the FDA firm-cluster crosswalk, so the estimates are restricted to the *FDA-linked* subsample—pairs in which both the treated firm and its matched control are observable in the FDA data (230 treated firms)—the same firm-level FDA-linked restriction applied to the submission outcomes (Section 6.2); firms not linked to the FDA data would otherwise contribute structural zeros that dilute the estimates toward zero. Panel A is our baseline: safety signals initiated *outside* the acquired firm—device recalls initiated by a party other than the manufacturer (a distributor, importer, or the FDA) and MAUDE adverse-event reports filed by users, patients, or distributors rather than the firm. These do not depend on the firm's own reporting decisions and so are not contaminated by any strategic change in disclosure or recall behavior following the deal (Section 6.4). Panel B reports the manufacturer-initiated and all-reporter counterparts, which the firm can influence directly. The Pre-mean column reports the treated-firm mean of each outcome over the pre-period $[-5, -1]$ (a count per firm-year for the log rows, a probability for the extensive-margin rows); non-manufacturer recalls are the any-role count net of manufacturer-initiated recalls. Recalls are used through 2024, the last complete year; adverse-event firm-years after the 2019 MAUDE source endpoint are dropped rather than coded as zeros. Seriousness is defined by patient outcome codes (death/life-threatening/hospitalization/disability). Standard errors clustered at the firm level. Significance: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 7: Treated Firms by Healthcare-Specialization of Investor Pool

	HC-specialized ($N = 254$ firms)	Non-HC ($N = 192$ firms)	HC − Non-HC
Mean patents per year	1.64	2.51	−0.87
Pr(any patent in window)	0.88	0.89	−0.01
Mean A61 patents per year	1.12	1.80	−0.68
Pr(any A61 patent in window)	0.72	0.67	+0.04
Mean FDA submissions per year	0.29	0.31	−0.03
Pr(any FDA submission in window)	0.35	0.30	+0.05
Mean firm age at deal (patenting tenure, years)	11.8	13.3	−1.5
% LBO deal type	68.9	68.2	+0.7

Note: Pre-deal observable characteristics of PE-acquired (treated) firms, split by whether the deal's investor pool included at least one healthcare-focused investor. Pre-deal window: five years before the deal anchor. Computed on the active patent-active sample; healthcare-specialized targets are somewhat smaller pre-deal patenters (1.64 vs. 2.51 patents/yr) but file a similar number of FDA submissions, and healthcare-specialized deals are hardly tilted toward LBO (69% vs. 68%); the simple correlation between the LBO and Healthcare-Specialized indicators is 0.01. The HC subset spans 254 unique treated firms and the Non-HC subset 192; the remaining 24 active treated firms have no investor recorded in PitchBook on their anchor deal and are excluded from both subsets (no treated firm has deals classified to both subsets).

Table 8: Heterogeneity by Investor Type: Healthcare-Specialized vs. Other (Patent Output)

A: By investor type

	Healthcare-specialized		Non-HC		Difference
	Medium $[0,6]$	Long $[7,12]$	Medium $[0,6]$	Long $[7,12]$	Long $[7,12]$
A: (Log) patents					
ATT	-0.065^{+}	-0.207^{***}	$+0.038$	-0.011	-0.197^{+}
	(0.037)	(0.061)	(0.043)	(0.090)	(0.107)
Pre-period mean	1.64		2.51		
B: Pr(any patent)					
ATT	-0.054^{*}	-0.168^{***}	$+0.005$	-0.072	-0.096
	(0.024)	(0.036)	(0.026)	(0.055)	(0.066)
Pre-period mean	0.47		0.47		

B: Healthcare-specialized deals, by deal type

	LBO		Non-LBO		Difference
	Medium $[0,6]$	Long $[7,12]$	Medium $[0,6]$	Long $[7,12]$	Long $[7,12]$
Log patents	-0.120^{**}	-0.261^{***}	$+0.070$	-0.046	-0.216
	(0.041)	(0.061)	(0.073)	(0.152)	(0.159)
Pre-period mean	1.41		2.15		
Pr(any patent)	-0.109^{***}	-0.216^{***}	$+0.082^{+}$	-0.023	-0.193^{*}
	(0.027)	(0.036)	(0.043)	(0.091)	(0.095)
Pre-period mean	0.42		0.57		

Note: Heterogeneity by healthcare-specialization of the investor pool, using an indicator equal to one if at least one linked investor in the deal is healthcare-focused. An investor is classified as healthcare-focused if a majority of its 1995–2019 investments are in healthcare companies. Each cell is the coefficient (standard error) from a single regression with both Medium and Long indicators, firm-by-stack and stack-by-year fixed effects, and standard errors clustered at the firm level. HC subset: 254 treated firms (756 paired stacks). Non-HC subset: 192 treated firms (569 paired stacks). Pairs whose anchor deal has no investor recorded in PitchBook are excluded (72 of 1,397 paired stacks, from 24 unique treated firms). The final column reports the difference in Long $[7,12]$ effects (HC-specialized minus non-HC) from a pooled regression in which every treatment term is interacted with the healthcare-specialization indicator, with standard errors clustered at the firm level. Each outcome is used through its data horizon (Section 4.1): patents and outcomes built from them through filing year 2020, FDA submissions and recalls through 2024, and venture financing through 2025. Significance: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Table 9: Within-Firm Evidence: Patenting in the Sponsor's Fields

	Log patents in class		Pr(any patent in class)	
	Medium $[0,6]$	Long $[7,12]$	Medium $[0,6]$	Long $[7,12]$
A: Baseline				
Treated × overlapping class	$+0.010$	-0.083^{*}	$+0.009$	-0.073^{*}
	(0.028)	(0.038)	(0.021)	(0.031)
Add-on platform counted in the portfolio	$+0.008$	-0.085^{*}	$+0.008$	-0.073^{*}
	(0.028)	(0.038)	(0.021)	(0.031)
Joint pre-trend p ($t \in [-10,-6]$)	0.11		0.15	
Randomization-inference p, Long $[7,12]$		0.044		0.016
B: By deal type				
LBO targets	$+0.023$	-0.087^{+}	$+0.021$	-0.071^{*}
	(0.031)	(0.050)	(0.023)	(0.035)
Non-LBO targets	-0.020	-0.081	-0.019	-0.052
	(0.063)	(0.070)	(0.048)	(0.058)
Difference		-0.005		-0.020
		(0.095)		(0.068)
C: By investor type				
Healthcare-specialized investors	$+0.022$	-0.092^{*}	$+0.019$	-0.076^{*}
	(0.032)	(0.043)	(0.023)	(0.030)
Other investors	-0.044	$+0.140$	-0.038	$+0.150$
	(0.060)	(0.099)	(0.047)	(0.105)
Difference, healthcare-specialized minus other		-0.232^{*}		-0.226^{*}
		(0.111)		(0.110)

Note: Each cell is the treated × overlapping-class × horizon coefficient (standard error) from Equation (4), on firm × class × year observations for treated firms and their matched controls, with firm-by-stack-by-year, firm-by-stack-by-class, and class-by-year fixed effects and standard errors clustered at the firm level. A class is a CPC main group (PatentsView current CPC) in which the firm patented during $[D-10, D-1]$; it is overlapping if a company the sponsor had invested in by the date of the deal patented in it over the same years; classes shared only with companies the sponsor first invested in after the deal are dropped for the target and its controls, and stacks in which the target or a control has a patent record that merges a company it acquired are dropped; controls inherit their stack's overlapping classes. Randomization-inference p: the share of 500 draws, each assigning every target the portfolio of another sponsor that invested in the same year and estimated jointly with the real overlap, in which the placebo coefficient is at least as negative as the first row. Panels B and C interact every overlap term with an indicator for LBO targets and with indicators for deals whose investor pool includes a healthcare-focused investor and deals without one (Section 6.5); the targets whose investor type is not recorded have no overlapping classes. Differences are Long $[7,12]$ contrasts from the same regression. Patent years through 2020; 327,619 observations; 100 treated firms have both overlapping and other classes. Significance: $^{+}$ $p<0.10$, * $p<0.05$, ** $p<0.01$, *** $p<0.001$.

Table 10: Inventor-Level Mechanism (Incumbent Device-Inventor Cohort)

	Pre-mean	Medium $[0, 6]$	Long $[7, 12]$
A: Firm-matched incumbent cohort			
Firm device output (all inventors)	1.87	-0.018	-0.107^{+}
		(0.031)	(0.057)
Incumbent at-firm output	1.80	-0.016	-0.028
		(0.025)	(0.034)
Incumbent citation-weighted	2.78	-0.058^{*}	-0.071^{*}
		(0.027)	(0.035)
Incumbent headcount	3.27	-0.034	-0.037
		(0.028)	(0.037)
New-inventor hiring (Pr)	0.32	-0.018	-0.060^{*}
		(0.017)	(0.029)
B: Attrition destination			
Active anywhere (device)	0.38	-0.023^{+}	$+0.003$
		(0.013)	(0.024)
Moved to another firm	0.15	-0.001	$+0.023$
		(0.010)	(0.021)
C: Strict inventor-matched (per-individual robustness)			
Per-inventor at-firm quality	0.92	-0.025^{+}	$+0.001$
		(0.013)	(0.019)
Per-inventor active (extensive)	0.40	-0.026^{*}	-0.021
		(0.011)	(0.020)

Note: Incumbent cohort = inventors on the firm's A61 device patents in $[-10, -1]$, defined symmetrically for treated and matched controls (controls inherit the treated deal year), own-exit censored. Panel A is firm-matched (same pairs as Table 2); Panel B decomposes cohort activity by destination; Panel C re-matches each incumbent inventor to control incumbents on the inventor's own pre-deal trajectory. "New-inventor hiring (Pr)" is the probability that at least one inventor makes a first-time appearance on a firm device patent in the year. The "Pre-mean" column reports the treated-firm mean of each outcome over the pre-period $[-5, -1]$ (a count for output/headcount rows, a probability for the extensive-margin rows). All estimates are $\log(1 + \cdot)$ for count outcomes and linear-probability for extensive-margin outcomes. Standard errors clustered at the firm level. Each outcome is used through its data horizon (Section 4.1): patents and outcomes built from them through filing year 2020, FDA submissions and recalls through 2024, and venture financing through 2025. Significance: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

Additional Figures

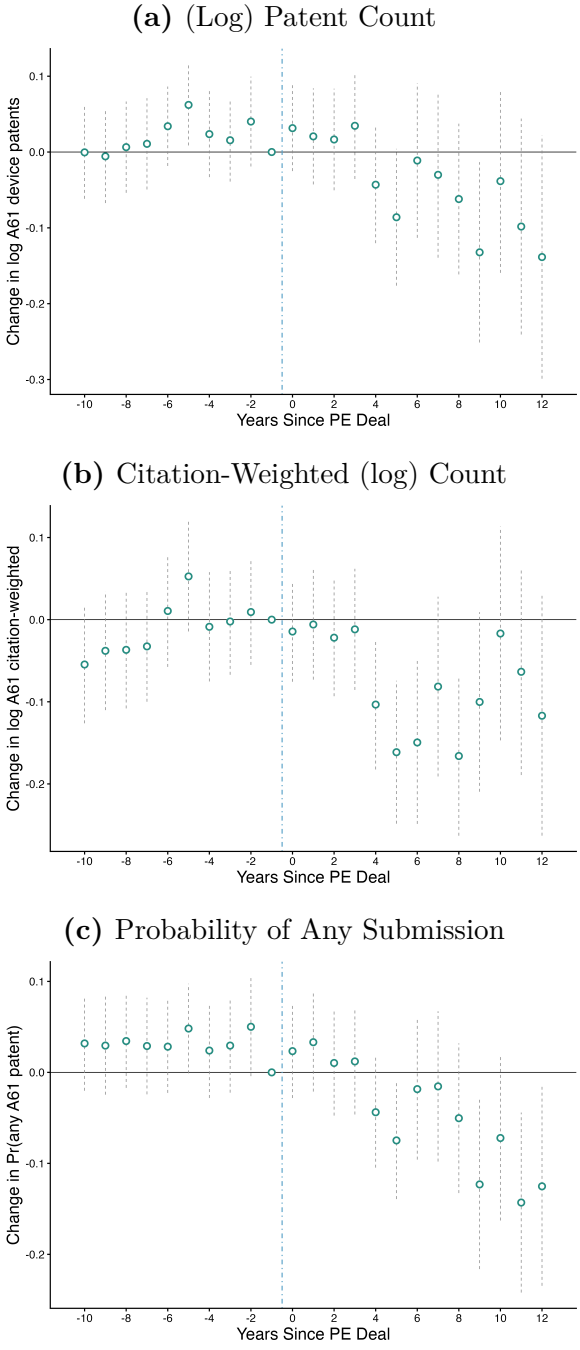


Figure A1: Effects on Medical Device Patenting

Note: This figure plots dynamic effect coefficients from the event-study version of Equation (3). It is similar in format to Figure 1 except that the outcomes correspond specifically to medical device patents, i.e., those filed in category A61. The outcome in Panel (b) is patent counts weighted by 5-year citation counts following the approach in Hall, Jaffe, and Trajtenberg (2001). The estimating sample includes 470 unique treated firms matched to 879 unique control firms (1,397 paired stacks). The coefficient β_{-1} is omitted and serves as the reference period. Standard errors are clustered by firm to account for cross-stack correlations within firm. Dashed grey lines indicate 95% confidence intervals.

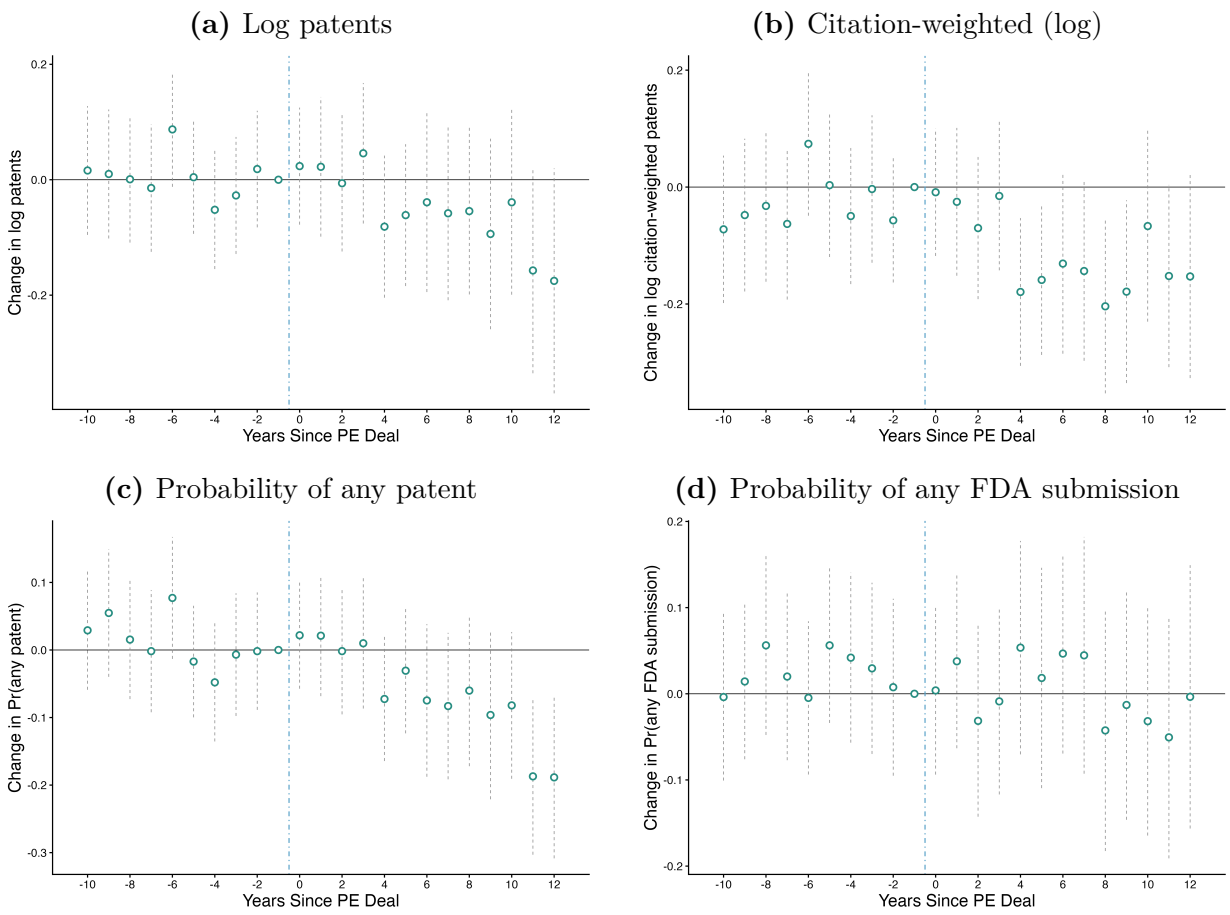


Figure A2: Event Studies on the Subsample of Deals Through 2013

Note: Event-study coefficients β_k from Equation (3) on the subsample of deals whose full post-period $[0, 12]$ lies within the data horizon (Table 4, Panel B): deals through 2008 for the patent outcomes, 196 treated firms (580 paired stacks), and deals through 2012 for the FDA outcome, 130 treated firms (299 paired stacks), which is estimated on the FDA-linked sample. β_{-1} is the omitted reference period. Firm-by-stack and stack-by-year fixed effects; standard errors clustered at the firm level. Dashed grey lines are 95% confidence intervals.

Additional Tables

Table A1: From the PitchBook Universe to the Analysis Sample

Stage	PE targets	Never-PE firms	Matched pairs
PitchBook innovative medical-device universe	823	6,074	—
Targets with a usable deal anchor	770	—	—
North American headquarters	770	3,545	—
Not a diversified conglomerate (no year with more than 200 patents)	768	3,537	—
Independent at the deal: no earlier acquisition, closure or PE deal (controls: at 2011)	728	2,990	—
Not a divested business of another company at the deal (controls: at 2011)	647	2,971	—
Patent-active in $[D-10, D-1]$, with a control within the technology caliper (analysis sample)	**470**	**879**	**1,397**
FDA outcomes: both sides in the FDA data	230	357	530
FDA core/non-core outcomes: pre-deal FDA activity on both sides	143	—	—
No patent-activity filter (robustness)	647	1,154	1,941

Note: Each row applies one more filter. Every filter applies identically to targets and to never-PE firms; only the deal anchor is target-specific. A target is screened in its own deal year D. A never-PE firm's eligibility depends on the deal year of the stack it would join, so the two independence rows count never-PE firms eligible in 2011, the median deal year of the analysis-sample targets; every later row counts firms in the matched sample. Independent at the deal: no acquisition, closure, or private-equity deal completed before D. Divested business: the firm was a business that another company sold at D (hand review of PitchBook records). Conglomerate: any year with more than 200 patents. The analysis sample keeps stacks in which the target and at least one matched control filed a US patent in $[D-10, D-1]$, with controls within the matcher's technology caliper (Section 4.3). Subsample rows: firms in the FDA crosswalk on both sides (FDA outcomes); stacks with pre-deal FDA activity on both sides (FDA core/non-core outcomes, targets only); and the matched sample without the patent-activity filter (Table 4, Panel C).

Table A2: Treated Firms vs. Eligible Never-PE-Acquired Firms, Before Matching

	Treated ($N = 470$)	Eligible never-PE firms ($N = 1{,}441$)	Difference
Mean patents per year	2.08	1.35	+0.73
Pr(any patent in window)	0.89	0.88	+0.01
Mean A61 patents per year	1.48	0.83	+0.64
Pr(any A61 patent in window)	0.70	0.66	+0.04
Mean FDA submissions per year	0.29	0.21	+0.08
Pr(any FDA submission in window)	0.33	0.28	+0.05
Pr(any VC round in window)	0.17	0.45	−0.28
Mean firm age (patenting tenure, years)	12.3	9.7	+2.5

Note: Pre-deal observable characteristics of the 470 estimation-sample targets and of the never-PE-acquired firms that would be eligible controls, before matching. Targets are independent at the deal and patent-active (at least one US patent in $[D-10, D-1]$); each is measured over the five years before its own deal year, $[D-5, D-1]$, the window of Table 1. The comparison firms are the never-PE-acquired firms in the innovative medical-device universe that meet the matcher's criteria at the reference year 2011, the median deal year of the targets: North American headquarters, no year with more than 200 patents, independent at 2011 under the rule of Section 4.3, and at least one US patent in 2001–2010; they are measured over 2006–2010. FDA submissions are 510(k) and PMA submissions through the firm-to-FDA crosswalk (zero for firms with no FDA record), as in the matching. Firm age is years since the firm's first US patent at the reference year. Differences are computed from unrounded means; means are across firms.

Table A3: FDA Regulatory Specialty Composition

Part	% treated	% control
Cardiovascular Devices	15.48%	13.99%
General and Plastic Surgery Devices	12.55%	15.58%
Orthopedic Devices	12.13%	8.74%
Radiology Devices	8.37%	7.15%
General Hospital and Personal Use Devices	7.11%	7.31%
Neurological Devices	6.69%	7.47%
Gastroenterology-Urology Devices	6.28%	8.27%
Anesthesiology Devices	6.28%	4.77%
Physical Medicine Devices	6.28%	3.18%
Dental Devices	5.02%	4.77%
Clinical Chemistry and Clinical Toxicology Devices	3.77%	3.50%
Obstetrical and Gynecological Devices	2.93%	4.45%
Ophthalmic Devices	2.09%	2.86%
Hematology and Pathology Devices	2.09%	1.91%
Ear, Nose, and Throat Devices	1.67%	3.18%
Immunology and Microbiology Devices	1.26%	2.86%
Total	100.00%	100.00%

Note: Distribution of treated and matched-control FDA activity across the 16 device Parts of 21 CFR (862–892), in the ten years before the deal, $[D-10, D-1]$ (controls take their stack's deal year). Sample: the firm-level FDA-linked matched sample used for the FDA outcomes (230 treated firms, 530 matched control units); the 170 treated and 371 control units with at least one classified submission in the window enter the table. A unit is active in a Part if it filed at least one 510(k) or PMA submission there; each column is that group's unit-by-Part cells over its total cells, so it sums to 100%. A submission's Part is the first three digits of its regulation number; a PMA submission without one (75% of PMA records) takes the Part of its FDA review panel, which agrees with the regulation-number Part on 99.4% of the PMA records that carry both. 510(k) records without a regulation number (2.8%) are not classified.

Table A4: Top PE-Firm Investors in the Treated Sample

		Baseline sample		Independent at the deal	
Investor	HC-specialized	Targets	%	Targets	%
RoundTable Healthcare Partners	Yes	9	2.6%	21	3.7%
Warburg Pincus	No	8	2.3%	13	2.3%
Audax Private Equity	No	8	2.3%	9	1.6%
HealthpointCapital	Yes	8	2.3%	9	1.6%
The Riverside Company	No	7	2.0%	11	1.9%
DW Healthcare Partners	Yes	7	2.0%	8	1.4%
Cortec Group Management Services	No	6	1.7%	8	1.4%
KRG Capital Partners	No	6	1.7%	8	1.4%
TPG	No	5	1.4%	8	1.4%
Altaris	Yes	5	1.4%	7	1.2%
Ampersand Capital Partners	Yes	5	1.4%	6	1.1%
Galen Partners	Yes	5	1.4%	6	1.1%
The Carlyle Group	No	5	1.4%	6	1.1%
Apax Partners	No	5	1.4%	5	0.9%
OrbiMed	Yes	5	1.4%	5	0.9%

Note: PE investors (PitchBook primary type PE/Buyout or Growth/Expansion) on the targets' anchor deals, ranked by the number of baseline targets they back. Investors are those on the target's anchor deal in the full-coverage investor file, the one the within-firm design of Section 7.1 uses: the anchor deal is the target's first effective PE deal, and 347 of the 470 baseline targets (569 of the 729 targets that no other company owned at the deal) have at least one recorded PE investor on it. Shares are over those targets; a deal can have several investors, so shares sum to more than 100%. An investor is healthcare-specialized if more than half of its 1995–2019 deals with an industry-classified company involve a company PitchBook classifies in the Healthcare sector, the rule behind the deal-level healthcare indicator of Section 6.5.

Table A5: ATT on A61 Medical-Device Patent Output

	Log A61 patents	A61 cite-weighted (log)	A61 extensive (any)
Medium $[0, 6]$	-0.014	-0.038	-0.029^{+}
	(0.025)	(0.026)	(0.017)
Long $[7, 12]$	-0.089^{+}	-0.078^{+}	-0.102^{***}
	(0.046)	(0.042)	(0.030)
Pre-period treated mean	1.48 patents/yr	2.27 (HJT-norm cites)	Pr(any) = 0.36
Joint pre-trend p ($t \in [-10, -6]$)	0.70	0.39	0.83

Note: Medium $[0, 6]$ and Long $[7, 12]$ ATTs on A61 medical-device patent output. Each cell reports the coefficient (standard error) from a single regression with both Medium and Long indicators, firm-by-stack and stack-by-year fixed effects, and standard errors clustered at the firm level. Sample as in Table 2. Each outcome is used through its data horizon (Section 4.1): patents and outcomes built from them through filing year 2020, FDA submissions and recalls through 2024, and venture financing through 2025. Significance: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$. The bottom row reports the p-value from a joint Wald test that the three event-study pre-trend coefficients at $t \in [-10, -6]$ all equal zero (out-of-sample for the $[-5, -1]$ matching window).

Table A6: Treated Firms by Deal Type: LBO vs. Non-LBO

	LBO ($N = 272$)	Non-LBO ($N = 144$)	Difference
Mean patents per year	4.44	3.26	+1.18
Pr(any patent in window)	0.94	0.96	−0.02
Mean A61 patents per year	3.54	2.61	+0.93
Pr(any A61 patent in window)	0.76	0.86	−0.10
Mean FDA submissions per year	2.67	3.10	−0.43
Pr(any FDA submission in window)	0.64	0.68	−0.03
Mean firm age at deal (years)	21.9	13.6	+8.2

Note: Pre-deal observable characteristics of PE-acquired (treated) firms, split by deal type. The LBO subset comprises Buyout/LBO and Investor Buyout by Management deals; the Non-LBO subset comprises PE Growth/Expansion and PIPE deals. Pre-deal window: five years before the deal anchor. The two subsets partition the sample exactly (LBO: 272 unique treated firms, 789 stacks; Non-LBO: 144 unique treated firms, 416 stacks).

A Institutional Detail

This appendix collects extended institutional background that supports the main text but would crowd its readability. It documents the FDA regulatory pathways whose filings we use as innovation and safety outcomes, and the structure of private-equity involvement in the medical-device sector.

FDA regulatory pathways

FDA device oversight is organized by risk class (Stern 2017). *Class I* devices (low risk—bandages, examination gloves) are subject to general controls and are largely exempt from premarket review. *Class II* devices (moderate risk—infusion pumps, diagnostic catheters) clear the market through the **510(k)** premarket-notification pathway. *Class III* devices (high risk, life-sustaining or implantable—heart valves, implantable defibrillators) require **premarket approval (PMA)**. The two pathways differ sharply in evidentiary burden and gestation, which is why they anchor the long- versus short-horizon distinction in Section 3.

510(k) substantial equivalence. A 510(k) submission demonstrates that a new device is "substantially equivalent" to a legally marketed *predicate* device in intended use and technological characteristics; it does not, in general, require new clinical trials. The sponsor identifies one or more predicates, and FDA review—typically on the order of three to six months—turns on comparative rather than de novo evidence. 510(k) clearance is the dominant route for Class II devices and accounts for the large majority of device submissions, so it captures incremental and line-extension innovation well. Devices with no suitable predicate but low-to-moderate risk may instead use the *De Novo* classification request, which creates a new Class I/II type and can then serve as a predicate for later 510(k)s.

Premarket approval (PMA). PMA is the most stringent pathway and applies to Class III devices. It requires affirmative evidence of safety and effectiveness, ordinarily including one or more clinical investigations, and review is substantially longer and more resource-intensive than 510(k) review, frequently involving an advisory-panel evaluation. Approvals are indication-specific, and post-approval commitments (for example, condition-of-approval studies) are common. Because a PMA device's development cycle routinely spans several years from initiation to approval, PMA activity is the margin most exposed to a truncated ownership horizon. We count both 510(k) and PMA submissions as regulatory-pipeline output (Section 4.1).

Adverse events (MAUDE). The Manufacturer and User Facility Device Experience (MAUDE) database records post-market adverse-event reports involving medical devices. Reporting is *mandatory* for manufacturers and importers (deaths, serious injuries, and certain malfunctions) and *voluntary* for health-care professionals, patients, and consumers. Reports therefore arrive from heterogeneous sources with different incentives; this is precisely why the safety analysis (Section 6.4) separates manufacturer-filed reports—which the acquired firm can influence—from reports filed by users, patients, and distributors, which it cannot.

Recalls. A device recall is a correction or removal action for a product that violates FDA law or is defective. FDA classifies recalls by the severity of the health hazard: *Class I* (reasonable probability of serious injury or death), *Class II* (temporary or medically reversible harm), and *Class III* (unlikely to cause harm). A recall may be initiated by the manufacturer, by a distributor or importer, or at FDA's request; the initiating party is recorded, which underlies the non-manufacturer-recall measure used as the baseline safety outcome.

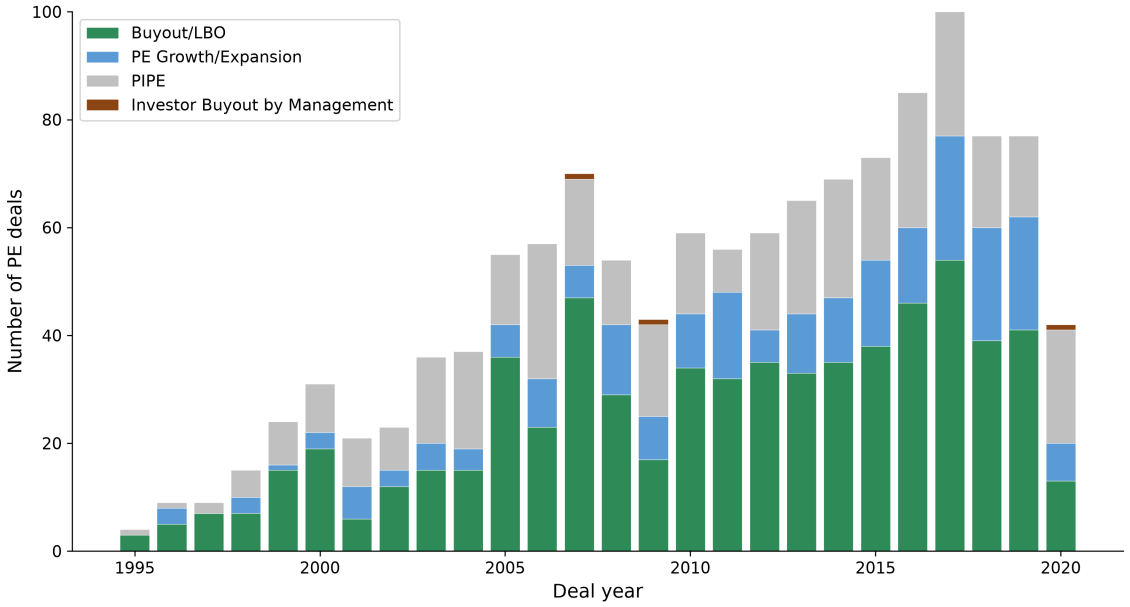
Establishment registration. Firms that manufacture or distribute devices in the United States must register their establishments with FDA and list their marketed devices annually. We use the FDA establishment-registration files only for entity resolution: they supply an additional set of firm-name strings—legal names, registered trade names, and establishment addresses—that anchor the embedding-based firm clusters described in Appendix D.

Private equity in medical devices

Private-equity activity in the sector is drawn from PitchBook (Section 4.1); the curated medical-device deal list contains 1,998 transactions across 1,464 unique firms. Deals fall into four PitchBook types—Buyout/LBO, PE Growth/Expansion, PIPE, and Investor Buyout by Management—with Buyout/LBO and PE Growth/Expansion accounting for the large majority (Section 2.2). Buyout/LBO deals pair a controlling stake with substantial acquisition leverage; PE Growth/Expansion deals provide minority or non-controlling growth capital with little acquisition leverage (Kaplan and Strömberg 2009). Deal flow into the sector rose markedly over the sample, with annual counts roughly tripling between the late 1990s and the late 2010s, consistent with the sector-level trends documented by Lerner, Sorensen, and Strömberg (2011).

Figure A3 plots the annual flow of the four PE deal types into the medical-device target firms in our universe over 1995–2020 (1,250 deals across 762 firms). The count rises from a handful of deals per year in the late 1990s to roughly seventy to one hundred in the late 2010s—the tripling noted above—with Buyout/LBO the largest single type throughout. Among completed exits of these firms, strategic sales (mergers and acquisitions) account for 51%, initial public offerings 27%, and secondary buyouts 23%.[10]

Figure A3: Private-Equity Deal Flow into Medical-Device Target Firms, 1995–2020



Note: Annual counts of the four PitchBook PE deal types—Buyout/LBO, PE Growth/Expansion, PIPE, and Investor Buyout by Management—for the medical-device target firms in our universe (823 firms, 762 of them with a qualifying PE deal in the window; 1,250 deals over 1995–2020), tabulated from the PitchBook deal record. The 2020 bar reflects a partial deal-year at the data cutoff. Deals of other types (venture financing, debt, public secondary offerings) and non-target firms are excluded.

B Conceptual Framework: Full Derivations

This appendix gives the full derivation of the portfolio-internalization model summarized in Section 3.

Setup

A target has a set of potential R&D projects $\mathcal{J}$. Project j requires an up-front outlay $c_j > 0$ and, if funded, yields expected value $v_j > 0$ to the target, in present value, realized after a

10. Deal counts and exit modes are tabulated from the PitchBook deal record for the target firms (Section 4.1); disclosed transaction sizes and fund-level sponsor concentration are available only for a subset of deals and are not shown. The figure covers the innovative-universe target firms, a subset of the broader curated deal list.

development lag τ_j. Each project lies in a technology field f_j and has other characteristics x_j—payoff horizon, technological breadth, eventual impact. The sponsor holds other companies $k \in \mathcal{K}$, and company k is active in field f with intensity $a_{kf} \geq 0$. Funding project j lowers the value of the sponsor's other companies by

$$e_j = \sum_{k \in \mathcal{K}} \phi\Big(a_{kf_j}\Big), \qquad \phi(0) = 0, \quad \phi(a) > 0 \text{ for } a > 0,$$

through research that duplicates theirs, business taken from their products, or both. The model restricts only whether a portfolio company is active in the field; it places no restriction on how the loss varies with the intensity of that activity. An owner of type θ values the target's projects net of a share $\omega_\theta \in [0, 1]$ of that loss, with $\omega_{\text{INDEP}} = 0$ and $\omega_{\text{PE}} > 0$ reflecting the sponsor's influence over the target's research budget. Projects are evaluated one at a time, so the owner solves

$$\max_{S \subseteq \mathcal{J}} \sum_{j \in S} \Big(v_j - \omega_\theta e_j - c_j\Big). \tag{A1}$$

The funding rule

A project enters the optimal set iff $v_j - \omega_\theta e_j \geq c_j$, which is Eq. (1); the funded set is

$$U(\omega) = \Big\{ j : \; v_j/c_j \;\geq\; 1 + \omega\, e_j/c_j \Big\}. \tag{A2}$$

Proposition 1 (Contraction). *$U(\omega') \subseteq U(\omega)$ for $\omega' > \omega$; in particular $U_{PE} \subseteq U_{INDEP}$. The forgone set $U_{INDEP} \setminus U_{PE} = \{j : \; 1 \leq v_j/c_j < 1 + \omega_{PE} e_j/c_j\}$ is nonempty only if some project has $e_j > 0$.*

Proof. Since $e_j \geq 0$, the hurdle $1 + \omega e_j/c_j$ is nondecreasing in ω, so any project that clears it at ω' clears it at ω. A project with $e_j = 0$ faces the hurdle 1 for every ω. □

Proposition 2 (Concentration). *The forgone set contains no project in a field that no portfolio company occupies, so the decline is confined to the fields in which the sponsor's other companies are active.*

Proof. Proposition 1 with $e_j = 0$: such projects face the hurdle 1 under every owner. □

Proposition 3 (No selection on horizon, breadth, or impact). *If x_j is independent of $(v_j/c_j,\, e_j/c_j)$, the probability that a project is forgone does not vary with x_j, so the decline carries no gradient in payoff horizon, breadth, or impact.*

Proof. The forgone set in Proposition 1 is defined by $(v_j/c_j, e_j/c_j)$ alone; under independence, $\Pr(\text{forgone} \mid x)$ equals its unconditional value. □

Proposition 4 (Deal type and leverage). *Acquisition debt does not enter Eq.* (A1)*; the owner's type matters only through ω_θ. If a buyout sponsor's control of the research budget gives it at least the influence that board representation gives a growth-equity investor, $\omega_{LBO} \geq \omega_G$, then the projects forgone under a growth-equity deal are also forgone under a buyout, so the decline and its concentration in the sponsor's fields are weakly larger under buyouts; the two coincide only if board representation confers the same influence as control. Within buyouts the decline does not vary with leverage.*

Proof. Leverage does not enter (A1). The forgone set grows with ω (Proposition 1), so $\omega_{\text{LBO}} \geq \omega_{\text{G}}$ implies $U_{\text{INDEP}} \setminus U_{\text{G}} \subseteq U_{\text{INDEP}} \setminus U_{\text{LBO}}$, and the forgone projects lie only in fields with $e_j > 0$ (Proposition 2). □

Proposition 5 (Partial redirection with retained staff). *Let each project's cost be $c_j = ws_j + k_j$, where s_j is its research staff, w the wage, and k_j its other costs, and suppose the independent owner's funded set employs the firm's staff $\bar{S}$ exactly. Suppose the PE owner retains $\bar{S}$, so that the wage bill of retained staff is sunk, and that the projects it forgoes remain unprofitable to it even with free staff, $v_j - \omega e_j < k_j$. Then (a) the PE owner funds U_{PE} plus a set $\mathcal{R}$ of projects with $k_j \leq v_j - \omega e_j < c_j$ that uses at most the staff released by the forgone projects; (b) projects in $\mathcal{R}$ satisfy $\omega e_j \leq v_j - k_j$ and so lie disproportionately in fields where e_j is small, outside the sponsor's fields; and (c) if all projects use the same staff s, then $|\mathcal{R}| \leq |U_{INDEP} \setminus U_{PE}|$, so the number of funded projects weakly falls.*

Proof. Projects in U_{PE} employ $\bar{S} - \Delta S$, where $\Delta S = \sum_{j \in U_{\text{INDEP}} \setminus U_{\text{PE}}} s_j$ is the staff released. With the wage bill sunk, a project staffed from the released pool is worth funding iff $v_j - \omega e_j \geq k_j$; the forgone projects fail this by assumption, so the pool is filled by other projects with the highest $v_j - \omega e_j - k_j$ among those with $k_j \leq v_j - \omega e_j < c_j$, up to ΔS, which is (a). Part (b) follows because the condition $v_j - \omega e_j \geq k_j$ tightens as e_j rises. For (c), with common s the pool holds $\Delta S / s = |U_{\text{INDEP}} \setminus U_{\text{PE}}|$ projects. □

Proposition 5 is the within-firm counterpart of the inventor evidence in Section 7.2: inventors who worked in the sponsor's fields stay and move to other fields, the firm brings fewer new inventors into the sponsor's fields, output in its other fields holds up, and output in the sponsor's fields and in total falls.

Timing

A project funded at date 0 realizes its output at $0 + \tau_j$. The new owner's rule applies to projects started from closing onward, but the output of projects already under way is unaffected, so measured patent and FDA output is flat at short event-time and falls only once the forgone projects would have matured—the Long-bin loading of prediction (ii). Because the forgone set is neutral in τ (Proposition 3), the timing of the aggregate decline reflects the average development lag rather than a selective loss of the longest-horizon output.

Welfare

Write $e_j = d_j + b_j$, where d_j is the value of research at a sibling that the target's project would duplicate and b_j the business it would take from a sibling's products. If the loss is duplication ($b_j = 0$), the target's project adds little beyond what the sibling already produces, and dropping it saves resources. If the loss is business stealing ($d_j = 0$), the sibling's loss is a transfer to the target and to buyers; the project's social value exceeds v_j by the consumer surplus from added competition and variety, so the sponsor's rule forgoes projects with positive social value—the concern raised by killer acquisitions (Cunningham, Ederer, and Ma 2021). We treat the welfare sign as open: the estimates identify the forgone projects, not the split of e_j.

C Matching Diagnostics

This appendix reports the balance diagnostics referenced in Sections 4.3 and 4.5 for the headline matched sample (470 treated firms, 879 controls, 1,397 paired stacks). Table A7 summarizes the three matching dimensions with a hard or soft target: technological proximity (the IPC-subclass cosine, with a hard caliper at 0.25), firm age, and pre-period patent stock.

Table A7: Matched-Pair Balance Diagnostics

Matching dimension	Median	Mean	p10	Share within band	
IPC-subclass cosine	0.50	0.55	0.28	all ≥ 0.25 (caliper)	
|Age gap| (years)	2	3.1	—	$\leq$ 5yr: 84.0%	$\leq$ 10yr: 94.2%
|Patent-stock gap|	2.0	—	—	p25: 0	p90: 24

Note: Diagnostics for the 1,397 matched treated–control pairs. The IPC-subclass cosine is the technological-proximity measure on which the 0.25 caliper binds during matching; its median, mean, and 10th percentile reproduce the main-text figures (Section 4.5). The age gap is the absolute difference in *years since first patent* between the treated firm and its matched control (see note below on the age variable). The patent-stock gap is the absolute difference in cumulative pre-anchor patent stock. A minority of pairs bind near the technology caliper: 15.6% of matched pairs have a cosine within 0.05 of the 0.25 floor (cosine in $[0.25, 0.30)$) and 26.9% within 0.10; the 10th percentile (0.28) sits just above the floor.

The age variable. The "firm age" entering the matching distance (Eq. (2)) is measured as *years since the firm's first US patent* (anchor year minus first-patent year) rather than calendar founding age. Because it tracks the firm's position in its innovation lifecycle, years-since-first-patent aligns the pre-period patenting paths of treated and matched controls—precisely the balance the parallel-trends design requires. A firm's calendar founding age, recoverable for a subset of firms, is largely orthogonal to its patenting history (among the 346 targets with a recorded founding year, the correlation of log pre-deal patent stock with years since founding is 0.25, against 0.58 with years since first patent); matching on it balances demographic age but degrades pre-period patenting balance, so it is not used.

Sensitivity to matching choices. The main text makes two design choices: the number of matched controls per treated firm ($K = 3$) and matching with replacement. Table A8 re-runs the matcher at $K = 1$ and $K = 5$, and without replacement (each control used at most once), and re-estimates the headline Long$[7, 12]$ ATTs. Raising the number of controls to $K = 5$ gives somewhat larger patent coefficients (-0.139 for the count, -0.159 for citation-weighted output, -0.137 on the extensive margin) and a similar FDA decline (-0.072, significant at the 5% level). Matching *without* replacement also yields somewhat larger patent declines (-0.166 for the count and for citation-weighted output, and -0.156 on the extensive margin, all significant at the 1% level) and a similar FDA decline (-0.078, significant at the 5% level), but this match fails the far-pre test on the count ($p = 0.08$) and the extensive margin ($p < 0.001$): when each control can serve only once, later cohorts receive worse matches. The one-nearest-neighbor match ($K = 1$) gives estimates close to the baseline (-0.102 on the count, not significant; -0.087 on the extensive margin, -0.164 on citation-weighted output,

and −0.124 on FDA submissions, each significant at the 5% level), with standard errors 30–50% larger than the baseline's; its pre-period coefficients are not jointly significant. The patent coefficients therefore do not depend on the number of matched controls; the baseline's $K = 3$ trades the variance reduction of additional controls against match quality.

Table A8: Sensitivity to Matching Choices (Long $[7, 12]$ ATT)

Matching specification	Log patents	Cite-weighted	Pr(any patent)	Pr(FDA sub.)
Baseline ($K = 3$, with replacement)	-0.111^{*}	-0.135^{**}	-0.127^{***}	-0.076^{*}
$K = 1$	-0.102	-0.164^{*}	-0.087^{*}	-0.124^{*}
$K = 5$	-0.139^{**}	-0.159^{***}	-0.137^{***}	-0.072^{*}
Without replacement ($K = 3$)	-0.166^{**}	-0.166^{***}	-0.156^{***}	-0.078^{*}

Note: Long$[7, 12]$ ATT from the headline two-indicator specification (Eq. (3)), re-estimated after re-matching under each alternative. Patent outcomes on the patent-active sample; Pr(FDA sub.) on the FDA-linked sample. "Without replacement" assigns each never-treated control to at most one treated firm (greedy, cohorts in anchor-year order). Far-pre joint tests ($t \in [-10, -6]$) do not reject at the 10% level under $K = 1$ or $K = 5$ on any margin ($p \geq 0.11$); without replacement they reject on the log count ($p = 0.08$) and the extensive margin ($p < 0.001$). Each outcome is used through its data horizon (Section 4.1): patents and outcomes built from them through filing year 2020, FDA submissions and recalls through 2024, and venture financing through 2025. Significance: $^{+}$ $p < 0.10$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

D Data Construction and Entity Linkage

This appendix documents the entity-resolution pipeline at the level of detail needed for replication; parameter values below are those committed in the replication package. The three primary sources (PitchBook, PatentsView, FDA) share no common firm identifier, so the panel is built by resolving heterogeneous firm-name strings to a common entity, symmetrically for treated and control firms.

Data horizons. The outcome sources end, or stop being complete, at different dates, and the event-time grid runs past them: a firm acquired in 2019 is followed to 2031. Patents are the binding case. A patent enters PatentsView only once granted, two to three years after filing on average, so the most recent filing years are incomplete: US grants by filing year number 367,500 for 2020 but 330,800 for 2021 and 74,500 for 2024. This is the truncation problem familiar from patent data (Hall, Jaffe, and Trajtenberg 2001). Coding the missing years as zeros would attenuate the long-run estimates, because years in which nearly every firm records nothing carry no information about the treatment effect yet enter the within-firm comparison. We therefore use each outcome only through its data horizon, set by one rule applied to each source's full annual series: the horizon is the latest full calendar year whose count is at least 90% of the mean of the three preceding years. The rule gives 2020 for patents and every outcome built from them (counts, citations, technology classes, inventors), 2024 for FDA submissions (the source ends in May 2025) and for recalls, 2025 for venture financing, and 2019 for MAUDE adverse events, whose device files end in that year. Calendar years after an outcome's horizon are dropped rather than filled with zeros, and every deal is kept

(Appendix D, Table A9).

Embedding-based name clustering

Firm-name strings from the FDA 510(k), PMA, MAUDE, recall, and establishment registration files are normalized (case-folding, punctuation and legal-suffix stripping) and embedded with OpenAI's *text-embedding-3-large* model, which returns a 3,072-dimensional vector per name. Embeddings are L_2-normalized and indexed in an exact FAISS inner-product index (`IndexFlatIP`), so inner products equal cosine similarities; for each name we retain its 50 nearest neighbors. Clusters are formed by single-link Union-Find under a fixed cosine-threshold schedule: at the initial threshold 0.83 all above-threshold neighbor pairs are merged, and any resulting cluster exceeding 30 names is recursively re-clustered at the next threshold in $\{0.88, 0.92, 0.95, 0.97, 0.985\}$. The size-triggered tightening breaks chains formed by collisions on common tokens—for example, two distinct firms sharing a parent brand name. The procedure resolves 124,848 normalized strings to 76,572 clusters (largest cluster 87 names). The threshold schedule and the 30-name size cap constitute the calibration; the package does not include a separate hand-labeled validation set, so no validation precision statistic is reported.

Firm-to-cluster crosswalk and symmetry

Each PitchBook firm—treated and control alike—is mapped to its union of FDA clusters through an alias graph built from the firm's canonical-applicant, legal-name, former-name, and also-known-as fields. A candidate alias is scored against cluster names by a `rapidfuzz` token-set ratio: a score of at least 92 is auto-accepted, a score below 78 is rejected, and the intermediate band (78–92) is adjudicated by a language model (`gpt-5.4-mini`). The model is instructed to select a candidate only if it is "clearly the same real-world firm," to return `no_match` otherwise, and to be conservative with generic words, returning a strict-JSON decision; a `match` with a chosen candidate is accepted, a `no_match` is rejected, and ambiguous or malformed responses are routed to manual review. The frozen manual-decision files in the package record 19 control-side and 166 FDA-cluster adjudications. The scoring and acceptance rule is identical on both sides, so no alias is admitted for a treated firm that would not also be admitted for a control; an asymmetric alias graph would inflate treated firms' measured FDA activity relative to controls and mechanically bias the post-deal comparison. The crosswalk is rebuilt from the firm's FDA *names* rather than stored cluster identifiers, because the cluster-id vocabulary was renumbered between covariate-file construction and the current submissions source; name-based linkage is vintage-independent and restores treated firms' post-rename submissions that a stale id-join would drop.

PatentsView assignee linkage

Patents are attributed to firms through the published PatentsView assignee-disambiguation identifiers, which are on a stable vintage. We link PitchBook firms to assignee identifiers by a `rapidfuzz` token-sort-ratio fuzzy match (shortlist threshold 75) against PatentsView assignee names, combining the PitchBook company-name, legal-name, former-name, and

also-known-as fields; mid-band and ambiguous scores are reviewed case by case, and the resulting decisions are frozen in the package. Because the assignee identifiers are stable across vintages, 99.9% of union identifiers resolve to a name in the source and no re-vintaging is required. Patents are aggregated to the firm-year over the union of a firm's assignee identifiers, following renames and successor entities identically for treated and control firms. Before aggregation we remove links that credit a firm with another organization's patents. The project's audit of every firm–assignee link proposes a class for each: the firm itself, a small or ambiguous assignee left for review, a carve-out, a different company, or a university or government body. We review the 186 links proposed as a different company or a university or government body under three rules fixed before any estimate was computed. (a) The university/government class is a keyword match that also fires on firms' own names (for example, College Park Industries and Clinical Innovations), so a link in that class is kept when the assignee's name matches the firm's own. (b) Every remaining link whose assignee matches one of the firm's PitchBook names (former, legal, or also-known-as) is verified by hand and kept when the assignee is the firm under that name (11 links, for example Covidien and Tyco Healthcare). (c) All others are dropped, such as a small orthopedic firm linked to Stryker's Howmedica Osteonics. The rules apply identically to targets and controls; 150 links are removed and firms keep their own patents. Firms with no linked patents are removed from the sample through the patent-activity filter, on whichever side they sit.

Control-pool contamination

Two mechanisms keep firms that are themselves PE-owned out of the never-treated comparison. First, during universe construction the target flag is reconciled against the curated PE-deals list, so no matched control carries the target flag; the package's reconciliation log documents the firms moved between the target and control sets under the qualifying-deal timing rule. Second, the symmetric own-exit censoring (Section 5) drops any control at its own subsequent PE buyout: the 38 matched controls that carry a completed leveraged-buyout or growth-equity deal in the PitchBook record are all present in the exit-event file and censored at that deal, so their PE-owned firm-years are excluded from estimation; the remaining controls with a deal record hold only PIPE stakes. These two steps supersede an earlier "misflagged-target" reconciliation described in prior drafts, which addressed the same concern on an intermediate data vintage.

Data horizons

Table A9 reports the horizon rule of Section 4.1 applied to each outcome source. The rule is applied to the source's full annual series rather than to our sample, so it measures the completeness of the source and not the composition of the sample: the horizon is the latest full calendar year whose count is at least 90% of the mean of the three preceding years. Real declines in the middle of a series (venture financing after its 2021 peak, year-to-year swings in recalls) do not move the horizon; a shortfall at the end of the series does. Each estimator drops an outcome's firm-years after its horizon at the point where it applies the event-time window, so the stacked panels themselves are unchanged and the uncapped estimates can be reproduced exactly (Table 4, Panel F).

Table A9: Data Horizons by Outcome Source

Outcome family	Source series	Horizon	Evidence
Patents and every outcome built from them	US patent grants by filing year	2020	2020: 367,541 (98%); 2021: 330,843 (88%); 2024: 74,525
FDA submissions (510(k) and PMA)	Submissions by date received	2024	2024: 5,003 (94%); 2025: 1,018 (source ends May)
Recalls	Recall events by initiation date	2024	2024: 3,262; 2025: 826 (partial year)
Venture-capital financing	PitchBook VC rounds by deal date	2025	2025: 15,946 (95%); 2026: 1,410 (partial year)
Adverse events	MAUDE device files	2019	Source ends in 2019

Note: The horizon is the latest full calendar year whose count in the source's full annual series is at least 90% of the mean of the three preceding years; percentages in the last column are that ratio. Patents are dated by filing year and enter the data only once granted, so recent filing years are incomplete. Outcomes built from patents include citation-weighted and A61 counts, core and non-core counts, CPC concentration and entry, the breadth, horizon, impact and semantic-scope measures, the inventor-cohort outcomes, the negative-control outcomes, and the sponsor-portfolio design. Firm-years after an outcome's horizon are dropped; no deal is dropped.

Data availability

The FDA sources are public. The openFDA 510(k), PMA, and recall/enforcement records, together with the MAUDE device files (1997–2019), were extracted on 21 July 2026; the frozen registration firm-by-product-code panel used only for entity resolution carries the same vintage. PatentsView bulk patent data—the grant, assignee, and IPC tables—are public and were extracted on 23 July 2026. PitchBook deal, company, and investor tables were extracted on 21 July 2026 (foreign-deal supplements on 17 August 2026), and the curated treated-firm seed is dated 13 December 2025. PitchBook data are proprietary and require a subscription license: the curated medical-device deal list and firm universe cannot be redistributed, but the derived firm-year panels used in estimation can be posted. Every source is recorded in the replication package with its file size and a SHA-256 hash.