Quietly Ahead: The Diverging R&D Productivity of Public and Private Firms*†

Daniel Bias[1] and Alexander Ljungqvist[2]

[1]Vanderbilt University

[2]Stockholm School of Economics, Swedish House of Finance, and CEPR

August 28, 2026

Abstract

Using 80 years of patent data, we show that unlisted firms are "quietly ahead" of stock market–listed firms: less valuable innovators for most of the twentieth century, they lead in innovation value today. The reversal is not explained by patent characteristics or by the multiples markets attach to them, but by where firms innovate: from the 1980s, unlisted firms reallocated R&D toward fields with rising valuations and away from declining ones. We date a structural break in 1985 in the dispersion of innovation values across technology fields—the year the reversal begins—while persistence remains high throughout. A simple model explains both halves: listed firms hold a scale advantage in turning research effort into value, unlisted firms redirect that effort more cheaply, and rising dispersion tips the balance. Young, first-time and VC-backed patenters drive the shift, and unlisted firms' share of aggregate R&D value-add has nearly tripled since 1985 despite listed firms' faster R&D growth.

JEL classification: O31, O32, O33, L25, O41.
Keywords: Innovation, Technological change, Public and private firms, R&D allocation, Organizational form, Economic growth.

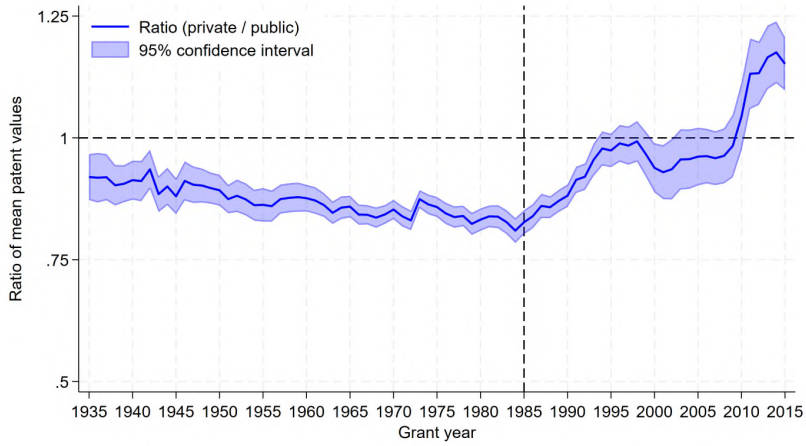
*We are grateful to Joan Farre-Mensa, Gerard Hoberg, Michelle Lowry, Gordon Phillips, and Leandro Sanz for helpful comments. Ljungqvist gratefully acknowledges generous funding from the Knut & Alice Wallenberg Foundation (grant no. KAW 2023.0270), the European Research Council (grant no. 101142217), and the Thule Foundation at Skandia.

†Funded by the European Union. Views and opinions expressed are, however, those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council Executive Agency. Neither the European Union nor the granting authority can be held responsible for them.

1. Introduction

Innovation—much of it corporate—is the engine powering economic growth.[1] We document a shift in the relative roles of public (stock market listed) and private (unlisted) firms that standard narratives do not predict. Figure I tells the basic story. Abstracting from estimation details, the figure plots the ratio of the mean patent values of private and public firms over the eight decades through 2015, with 95% confidence intervals. In 1935, the mean value of patents granted to private firms was 92% of the mean value of patents granted to public firms. This ratio drifted downward for half a century, reaching a low of 81% in 1985. Over the next three decades, however, the trend reverses: private and public firms reach parity in the decade through 1994, remain at parity through 2010, and by 2015 private-firm patents are on average 15% more valuable.

Figure I: The Reversal in the Relative Value of Private- and Public-Firm Innovation, 1935-2015



Notes: This figure plots the ratio of mean patent values for private firms relative to public firms from 1935 to 2015. The vertical dashed line marks 1985, the date at which we independently estimate a structural break in the dispersion of innovation values across technology fields (Section 4). Details of our estimation approach can be found in Section 2. The 95% confidence intervals are based on standard errors clustered at the technology-field level.

The basic pattern in Figure I is not an artifact of how we measure innovation value. It is robust to alternative measures: for example, we find similar reversals in the numbers of breakthroughs and disruptive innovations per R&D dollar. In short, private firms have gone from clear laggards in the mid-twentieth century to quiet leaders in the innovation race today.

[1]See Aghion and Howitt (1998) for a canonical survey of theory linking firm-level innovation to long-run growth; see Bloom, Van Reenen, and Williams (2019) for empirical evidence.

The reversal in Figure I raises a natural question: why did private firms start overtaking public firms? Using a set of counterfactual exercises, we show that the reversal is not driven by two tempting explanations. Specifically, neither shifts in the characteristics of patented innovations (such as citation intensity or the incidence of breakthroughs) nor changes in the "valuation multiples" investors attach to those characteristics can explain the reversal. Nor can the decline in IPO activity since the mid-1990s. Instead, the key lies in where firms choose to innovate within the technological opportunity landscape. Starting around the mid-1980s, private firms increasingly shifted innovative effort toward technology fields whose market valuations were rising and away from fields where valuations were stagnating or falling. In this sense, their innovation portfolios adjust more nimbly to signals about how the technological opportunity landscape is evolving. As a result, private firms move from being overrepresented in low-value fields to being overrepresented in high-value fields.

To interpret both the reallocation patterns and their timing, we develop a simple model of technological reallocation in which firms allocate innovative effort across fields within a technological opportunity landscape that evolves over time. Organizational form confers two advantages that pull in opposite directions: public firms are more productive per unit of inventive effort—scale, complementary assets, established research infrastructure—while private firms redirect that effort at lower cost. Persistence determines whether reallocation is worth undertaking at all, since chasing transitory opportunities earns nothing; given persistence, dispersion determines what it is worth. The model therefore delivers the reversal as a threshold: relative innovation value flips sign once cross-field dispersion crosses a level set by the size of the productivity gap and the difference in adjustment frictions. On this account the reversal is not a change in organizational form, but a change in the environment in which pre-existing organizational differences are priced.

Why should organizational nimbleness begin to pay in the mid-1980s and not before? Because that is when the technological opportunity landscape changed shape. That such a shift occurred somewhere in the late twentieth century is widely argued; its date, as far as we are

aware, has not previously been estimated.[2] Applying the multiple-break methodology of Bai and Perron (1998) to the cross-field dispersion of innovation values, we date a structural break to 1985, with a 95% confidence interval of 1984–1986. The date is not imposed, and it is identified off the comparison group: field-level values are constructed from public-firm patents alone, so no private-firm data enter the dispersion series. Cross-field dispersion is flat over the three decades preceding the break and rises by more than an order of magnitude over the fifteen years that follow it. Persistence, by contrast, is high throughout the eighty-year sample and exhibits no comparable break: what changes in 1985 is not how long a field's advantage lasts, but how much that advantage is worth. That distinction matters for the interpretation, because opportunities becoming more durable would not by itself reverse the ranking—it is the rise in what reallocation is worth that does. Seen through the lens of the model, this shift moves the economy across the threshold at which rapid reallocation toward high-value technologies becomes decisive—and it does so in the same year in which the reversal in Figure I begins.

Which private firms drive the shift? A striking share of the reversal is driven by first-time patenters and young private firms. Geography also matters. Private firms in California outperform public firms during the crucial decades of the 1970s and 1980s, while private firms in Massachusetts maintain an advantage from the 1970s through 2015. These patterns suggest an interaction between organizational form, firm life cycle, and local innovative environments.

Finally, we explore aggregate macro patterns. We construct a time series of private firms' share of total R&D value-add, defined as the aggregate value of corporate patents net of R&D

[2]The macro and innovation literatures emphasize three broad mechanisms. General-purpose technologies such as information and communication technologies increased the dispersion of returns across downstream applications by creating highly scalable opportunities in some domains while leaving others largely unaffected (Bresnahan and Trajtenberg 1995; Helpman 1998). At the same time, innovation became more tightly linked to frontier science, increasing the persistence and skewness of technological payoffs across fields (Rosenberg 1982; Jones 2009). More recent work further suggests that the pace of new technology creation accelerated beginning in the 1970s, while innovation outcomes became increasingly concentrated across technologies and firms, contributing to a more uneven technological opportunity landscape (Kelly et al. 2021; Hassan, Kalyani, and Restrepo 2026). We view this work as complementary to ours. Its concern is with establishing *that* the character of technological opportunity changed and with the mechanisms through which it did so; where it is quantitative, it measures related but distinct objects—the rate at which new technologies arrive, or the concentration of innovation outcomes across firms and technologies—rather than the dispersion of market-assessed value across fields. Perhaps as a consequence, the timing implied by these accounts ranges from the early 1970s to the mid-1990s.

spending. Since 1985, that share has nearly tripled. This macro shift is not due to a surge in private-firm patenting or faster growth in private-firm R&D spending relative to public-firm spending—indeed, the opposite is true. Instead, the rising contribution of private firms reflects the steadily increasing value of the patents private firms are granted.

This paper documents a fundamental shift in the organization of innovative activity in the United States. Over much of the twentieth century, public firms dominated high-value corporate innovation. Since the mid-1980s, however, private firms have emerged as leading producers of high-value patented innovation. We show that this reversal is not explained by shifts in the characteristics of patented inventions or by changes in the market "multiples" attached to those characteristics. Instead, it reflects an increasingly large advantage of private firms in reallocating inventive effort toward rising-value regions of the technological opportunity landscape and away from stagnating or declining areas—an advantage that becomes decisive only after 1985, the year in which we date a structural break in the dispersion of technological opportunity across fields. By linking organizational form to the speed and direction of technological reallocation, our findings illuminate how technological opportunity and market pricing interact to shape long-run patterns of innovation and economic growth.

Our paper contributes to the Schumpeterian growth literature surveyed by Aghion and Howitt (1998) and developed in subsequent work (Klette and Kortum 2004; Jones 2005; Akcigit and Kerr 2018). While this tradition emphasizes that long-run growth is driven by quality-improving innovation and by the reallocation of inventive effort toward high-opportunity technologies, canonical models typically abstract from heterogeneity in organizational form, treating firms as identical producers of innovation. We provide long-horizon empirical evidence that the structure and identity of innovating firms—private versus public, young versus established—systematically shape the value of innovative output.[3] Consistent with the reallocation mecha-

[3]A related literature studies how going public affects a firm's innovative output. It finds that the transition from private to public markets decreases patent quality among U.S. biotech firms (Aggarwal and Hsu 2014), U.S. firms more broadly (Bernstein 2015), Chinese firms (Cong and Howell 2021), and European firms (Larrain et al. 2025). Gao, Hsu, and Li (2018) further show that patents by public firms rely more on existing knowledge and are less likely to be in new technology classes than those of large private firms during 1997–2008. Our focus is distinct in that we compare the value of patents produced by all private and public U.S. firms over a long

nism emphasized in Schumpeterian models, we show that the return to shifting R&D toward rising-value fields rose sharply after 1985 and that private firms were positioned to capture it, yielding a long-run reversal in the relative value of private- and public-firm innovation. By demonstrating that organizational form governs the speed and direction of technological reallocation, our findings extend a core channel in Schumpeterian growth theory.

Our results also speak to broader literatures in corporate finance and venture capital. In corporate finance, a large body of work examines how ownership structure (Aghion, Van Reenen, and Zingales 2013), financial constraints (Brown, Fazzari, and Petersen 2009; Acharya and Xu 2017), and market pressures (Baber, Fairfield, and Haggard 1991; Dechow and Sloan 1991) shape firms' innovative activity. Studies document that public firms face pressures to meet short-term earnings targets (Stein 1989; Asker, Farre-Mensa, and Ljungqvist 2015) and may underinvest in long-term innovation relative to private firms, while governance, incentive design, and managerial horizons can influence the risk profile and quality of innovation (Holmström 1989; Lerner, Sorensen, and Strömberg 2011). We complement this literature by showing that organizational form predicts not only the amount or type of innovative activity, but the value of innovation, and that this gap has widened markedly over time.

Finally, our paper connects to the literature on venture capital and entrepreneurial innovation. Venture capital has been shown to accelerate commercialization, scale frontier technologies, and support young firms in high-opportunity sectors.[4] We document that the reversal in private-firm innovation value is disproportionately driven by young and VC-backed firms, consistent with an amplifying role for venture capital in the post-1980s innovation regime. In this way, our evidence bridges Schumpeterian growth models—where entrepreneurial entrants drive

time horizon and study the evolution of their relative innovation performance.

[4]A large literature links venture capital to greater innovative activity and technological progress. Early work shows that patenting increases with VC investment (Kortum and Lerner 2000), and subsequent studies document that VC financing is associated with higher-impact innovation and entrepreneurial activity (Mollica and Zingales 2007; Bernstein, Giroud, and Townsend 2016). Venture capital also shapes the *type* of innovation pursued: firms backed by more failure-tolerant VCs produce patents with greater impact (Tian and Wang 2011), and periods of abundant VC funding are associated with riskier, more innovative projects and higher-quality patents (Nanda and Rhodes-Kropf 2013; Howell et al. 2020). At the aggregate level, VC-backed firms account for a disproportionate share of R&D spending among U.S. public companies (Gornall and Strebulaev 2021).

creative destruction—with modern corporate finance models emphasizing the role of specialized financiers in fostering high-growth innovation.

Together, these literatures underscore the importance of institutional context, investor structure, and organizational form in shaping innovative performance. Our paper contributes a new fact of first-order importance: over the past century, the identity of the marginal innovator has shifted decisively, with private firms emerging as the leading producers of high-value patented innovation.

The remainder of the paper proceeds as follows. Section 2 describes the data and the construction of patent-level and field-level measures of innovation value. Section 3 introduces a simple model of technological reallocation in which a productivity advantage for public firms and a speed advantage for private firms trade off, and derives the level of cross-field dispersion at which the ranking between them reverses. Section 4 dates a structural break in 1985 in the dispersion of innovation values across the technological opportunity landscape, and shows that persistence is high and stable throughout the sample. Section 5 shows that private firms respond more quickly to these changing opportunities, reallocating inventive effort toward rising-value fields more rapidly than public firms. Section 6 examines the aggregate implications, documenting a substantial rise in private firms' share of economy-wide R&D value-add. Section 7 concludes.

2. Data and Measurement

We construct a comprehensive dataset of U.S. corporate patents and their market values in three steps. First, we identify all patents granted to U.S. firms between 1926 (the beginning of CRSP coverage) and 2020, and classify each patent by the organizational form of the assignee, distinguishing between public (stock market listed) and private (unlisted) firms. Second, following Kogan et al. (2017), we estimate the market values of patents granted to public firms using stock price responses to patent grant announcements. Third, following Kline et al. (2019), we impute the values of patents granted to private firms using comparable information from

public firms.

2.1. Identifying Public- and Private-firm Patents

Since its founding in 1836, the United States Patent and Trademark Office has issued more than eleven million patents. Our objective is to identify patents granted to U.S. corporations and to distinguish between those held by public firms and those held by private firms. We exclude patents assigned to individuals, universities, hospitals, non-profit organizations, government entities, and foreign corporations, as our focus is on domestic corporate innovation.

We construct the sample using three conceptual filters. First, we identify whether a patent assignee is a firm rather than a non-corporate entity. Second, among corporate assignees, we determine whether the firm is domestic or foreign, excluding patents assigned to foreign firms even when those firms file patents in the United States or are cross-listed on U.S. exchanges. Third, for domestic firms, we classify each patent according to whether the assignee is publicly listed or privately held at the time of the patent grant.

To implement the first filter, we use assignee information from Google Patents, the USPTO PatentsView database, the USPTO Patent Assignment database, and textual parsing of patent documents. While PatentsView provides standardized assignee classifications from 1976 onward, for earlier patents we rely on our own classification procedures and on information from Fleming et al. (2019).

To determine whether a corporate assignee is domestic or foreign, we combine location information from PatentsView, the USPTO Assignment database, Fleming et al. (2019), and the PatentCity database of Bergeaud and Verluise (2024). A particular challenge arises when foreign firms patent through U.S. subsidiaries, which appear as domestic assignees in patent records. To address this, we use listing information for foreign firms from LSEG Datastream and perform record linkage between patent assignees and internationally listed firms, allowing us to exclude patents filed by U.S. subsidiaries of foreign corporations.

Finally, we classify domestic corporate patents as public- or private-firm patents. Our start-

ing point is the latest patent–firm linking table from Kogan et al. (2017), which links patents granted since 1926 to firms in the CRSP database. We refine and extend this table in several steps. First, we match patent assignees to firms listed on the American Stock Exchange prior to 1962, before CRSP coverage of AMEX firms begins. Second, using corporate ownership trees from the National Establishment Time Series (NETS) database, we identify patents granted to subsidiaries of public firms over the period 1989–2019.[5] Third, we manually review approximately 3.2 million patent–firm links, identifying an additional 450,000 public-firm patents that are coded as private in Kogan et al.'s latest linking table, reclassifying roughly 28,000 of Kogan et al.'s public-firm patents as private or something else, and correcting 20,000 erroneous CRSP permco links.

Our final analysis sample consists of 3,462,008 patents granted to U.S. public and private firms between 1935 and 2015. Public firms account for 2,260,195 patents (65.3%), while private firms account for 1,201,813 patents (34.7%). We classify a firm as publicly traded if it has a common share listed on the NYSE, AMEX, NASDAQ, a regional exchange, or the over-the-counter market.

2.2. Estimating the Value of Public Patents

To estimate the economic value of patents granted to public firms, we build on the methodology of Kogan et al. (2017), who infer patent values from stock market reactions to patent grant announcements. The key insight is that equity prices are forward-looking and therefore reflect the private value that investors assign to newly granted intellectual property. Specifically, Kogan et al. identify patent grant dates using USPTO records and measure abnormal stock returns in a narrow event window following the grant. Under the assumption that patent grants convey new information about monopoly rents associated with the innovation, the stock price reaction isolates the component of firm value attributable to the patent. Filtering these returns from background volatility and adjusting for the market's expectation of patent grant yields a

[5]This step captures patents assigned to subsidiaries whose names bear little resemblance to their publicly listed parent, and corrects errors arising from fuzzy name matching in extant linking tables.

patent-level estimate of private economic value that is comparable across firms, industries, and time.

Rather than relying on the original linking table in Kogan et al. (2017), we reestimate the entire valuation procedure using our revised and corrected patent–firm matches described in Section 2.1. As a result, our estimates of public-firm patent values are internally consistent with our classification of organizational form and are directly comparable to the values we later impute for private firms.

Apart from the updated firm–patent linkages, we follow Kogan et al. (2017) closely in terms of event-window construction, return filtering, and distributional assumptions. We therefore inherit the desirable properties of their measure—most notably, its grounding in asset prices and its ability to capture economically meaningful heterogeneity in innovation value—while ensuring that the public-firm benchmark is measured as accurately as possible.

2.3. Imputing the Value of Private-Firm Patents

Because private firms are not publicly traded, their patent values cannot be inferred directly from stock market reactions. We therefore impute the value of patents granted to private firms using observable patent characteristics, drawing on the relationship between those characteristics and market-based patent values estimated for public firms.

Our approach follows Kline et al. (2019), who extrapolate the patent valuation framework of Kogan et al. (2017) to unlisted firms. In our implementation, we estimate the mapping from patent characteristics to economic value using the reestimated public-firm patent values described in Section 2.2, ensuring full consistency between the valuation of public- and private-firm patents.

Specifically, we regress the market value of patents granted to public firms on a set of patent characteristics that have been shown to predict innovation value.[6] To allow the valuation

[6]The characteristics include forward citations over a five-year window, the number of independent claims, measures of generality and originality (Trajtenberg, Henderson, and Jaffe 1997; Hall, Jaffe, and Trajtenberg 2001), an indicator for breakthrough inventions (Kelly et al. 2021), and an indicator for strategic patents (Veihl

relationship to evolve over time and across technologies, we estimate this model recursively using rolling ten-year windows and include grant-year fixed effects and technology-field fixed effects at the Cooperative Patent Classification (CPC) main group level.[7]

Formally, for patent i granted in year t, we estimate the following Poisson specification:

$$v_{it} = \exp(\beta_t X_i + \pi_j + \phi_t + \varepsilon_{it}), \tag{1}$$

where v_{it} denotes the market value of the patent as inferred from stock prices, X_i is the vector of patent characteristics, π_j are technology-field fixed effects, and ϕ_t are grant-year fixed effects. Estimation over rolling windows yields a time series of coefficient estimates which we use to compute the imputed value $\tilde{v}_{it}$ of each private-firm patent granted in year t.

The Online Appendix provides additional detail and diagnostics for the imputation procedure. Figure IA.1 reports the rolling-window coefficient estimates from equation (1). While magnitudes vary over time, patent values are consistently increasing in citations, claims, generality, and originality, and are higher for breakthrough and strategic patents. Figure IA.2 reports diagnostics on model fit and coverage, including pseudo-R^2 measures, the number of public patents and technology fields used in estimation, and the incidence of excluded singletons. Finally, Figure IA.3 validates the imputation by comparing out-of-sample imputed values to market-based patent values for public firms after absorbing grant-year fixed effects, showing a close within-year alignment between the two measures.

Our data allow us to observe how the value of innovation varies over time, both across organizational forms and across technology fields, and how firms reallocate inventive effort in response to changing technological opportunities. The next section introduces a simple model

2022). In early estimation windows beginning with 1926–1935, citation-based measures are unavailable; in these windows, the specification excludes citations, generality, and originality, which enter once the data coverage permits.

[7]CPC main groups are granular technology fields defined by an eight-character code. The firms in our sample patent in 7,211 distinct CPC main groups. One example is C12N15/00, "Mutation or genetic engineering; DNA or RNA concerning genetic engineering, vectors, e.g. plasmids, or their isolation, preparation or purification; Use of hosts therefor."

to organize these patterns and clarify the role of organizational nimbleness.

3. A Simple Model of Technological Reallocation

Schumpeterian growth theory (Aghion and Howitt 1992, 1998) emphasizes that long-run growth depends on the direction of technological change and on how rapidly inventive effort is reallocated toward high-opportunity technologies. Canonical models assume homogeneous firms and treat innovation productivity as independent of organizational form. This section develops a stylized model in which organizational form matters along two dimensions that pull in opposite directions: public firms are more productive per unit of inventive effort, while private firms redirect that effort at lower cost. Neither advantage dominates for all parameter values. Which one prevails depends on the technological environment, and a change in that environment reverses the ranking.

The model is deliberately minimal, and its purpose is twofold. First, it makes the reversal we document a derived implication of a shift in the opportunity landscape rather than an assertion about organizational form. Second, it tells us what to measure: the model implies that the relevant feature of the environment is the dispersion of technological opportunities across fields, conditional on those opportunities being persistent enough to be worth pursuing. Sections 4 and 5 take these two objects to the data.[8]

3.1. Environment

There are two technology fields, $j \in \{A, B\}$. In each period t, field j has a "technology opportunity value" $v_{j,t}$ summarizing the expected market return to an innovation in that field, for example through higher patent citations, commercialization prospects, or stock-market multiples. Field values evolve according to

$$v_{A,t} = \mu + \epsilon_t, \qquad v_{B,t} = \mu - \epsilon_t, \qquad \epsilon_t = \rho\,\epsilon_{t-1} + u_t, \quad u_t \sim \text{i.i.d.}\,(0, \sigma_u^2), \quad \rho \in [0, 1), \tag{2}$$

[8]Formal propositions and proofs are in Appendix A.

so that ϵ_t is the opportunity gap between the two fields. We refer to its unconditional variance $\sigma_\epsilon^2 = \sigma_u^2/(1-\rho^2)$ as the *dispersion* of technological opportunities and to ρ as their *persistence*. Section 4 constructs empirical counterparts to both.

Two types of firms operate: private firms (indexed P) and public firms (indexed U). Each allocates its inventive effort across the two fields, with $s_{i,j,t} \in [0,1]$ and $\sum_j s_{i,j,t} = 1$ denoting the share that type i devotes to field j. Effort is committed one period in advance: the firm chooses $s_{i,t}$ using information available at $t-1$, reflecting the lag between redirecting a research program and observing what it produces. This timing is what gives persistence its role. If opportunity gaps were transitory, a firm reallocating today would be chasing noise, and speed would earn nothing.

3.2. Firm Problem

Firm type i chooses $s_{i,t} = (s_{i,A,t}, s_{i,B,t})$ to maximize the expected value of its R&D portfolio net of concentration and adjustment costs:

$$\mathbb{E}_{t-1}\big[\Pi_{i,t}\big] \;=\; \theta_i \;+\; \sum_j s_{i,j,t}\,\mathbb{E}_{t-1}[v_{j,t}] \;-\; \frac{\chi}{2}\sum_j s_{i,j,t}^2 \;-\; \frac{\kappa_i}{2}\sum_j \big(s_{i,j,t} - s_{i,j,t-1}\big)^2. \tag{3}$$

The parameter θ_i is a field-invariant productivity shifter: the value a firm of type i extracts from a unit of inventive effort regardless of where that effort is directed. We assume $\theta_U > \theta_P$ and write $A \equiv \theta_U - \theta_P > 0$. This captures the advantages of public-firm scale—larger and more specialized research staffs, deeper complementary assets in manufacturing and distribution, established patent portfolios that raise the value of marginal inventions, and lower costs of capital for long-horizon projects—and reflects the fact that through much of the twentieth century, the large corporate laboratory was the dominant locus of commercially valuable invention. Section 5 provides direct evidence on this level difference: measured net of R&D spending, the average private-firm patent is worth roughly 40% of the average public-firm patent in 1985.

The parameter $\chi > 0$ captures diminishing returns to concentrating inventive effort within a single field. As a firm piles researchers, budget, and managerial attention into one technology,

the marginal project it funds there is progressively less valuable and within-field spillovers are progressively exhausted. Because $\sum_j s_{i,j,t}^2$ is minimized at the balanced portfolio, the term penalizes concentration; it also ensures that the firm's problem has an interior solution, so that effort shares respond smoothly to opportunities rather than jumping to corners.

The parameter $\kappa_i > 0$ captures organizational inertia, that is, frictions associated with shifting scientific direction, reorganizing research teams, or revising project portfolios. We assume $\kappa_P < \kappa_U$: public firms face higher effective costs of changing R&D direction due to (i) the need to disclose private information to dispersed shareholders (an Arrow 1962–style appropriation problem), (ii) reliance on verifiable hard information before executing a pivot, or (iii) stronger stock-price and career penalties for unsuccessful pivots.[9]

The two forms of heterogeneity work in opposite directions: public firms are more productive per unit of effort but slower to redirect it. We treat the two types as a simplification; nothing in what follows depends on there being only two, and the results hold for any ordering of κ across firms. This matters for interpreting the cross-sectional evidence in Section 5.5, where the reversal is concentrated among the private firms plausibly facing the lowest adjustment costs.

3.3. Reallocation and Innovation Value

Because effort shares sum to one, it is convenient to work with the portfolio tilt $x_{i,t} \equiv s_{i,A,t} - \frac{1}{2} \in [-\frac{1}{2}, \frac{1}{2}]$. Substituting (2) into (3) and using $\mathbb{E}_{t-1}[\epsilon_t] = \rho\,\epsilon_{t-1}$, the problem becomes strictly concave in $x_{i,t}$, and the first-order condition delivers a partial-adjustment rule:

$$x_{i,t} \;=\; \phi_i\, x_{i,t-1} \;+\; (1-\phi_i)\, x_t^*, \qquad x_t^* \;\equiv\; \frac{\rho\,\epsilon_{t-1}}{\chi}, \qquad \phi_i \;\equiv\; \frac{\kappa_i}{\chi+\kappa_i} \in (0,1). \tag{4}$$

[9]We treat the higher adjustment cost of public firms as a reduced-form representation of several microfoundations. In an Arrow (1962) information environment, justifying a pivot to outside investors requires public firms to disclose more of their informational advantage, raising the effective cost of changing direction relative to closely held private firms. Governance and financing frictions further amplify this gap: dispersed shareholders and independent boards rely more heavily on hard, verifiable signals (Holmström 1979; Stein 2002), and public managers face stronger stock-price and career penalties for failed pivots (Stein 1989; Holmström 1999; Gibbons and Murphy 1992). Appendix B provides formal microfoundations showing how each of these mechanisms yields a higher effective adjustment friction κ_U for public firms than κ_P for private firms.

Firms move toward a target x_t^* that is proportional to the *forecast* opportunity gap, closing a fraction $1 - \phi_i$ of the distance each period. Since ϕ_i is increasing in κ_i, private firms adjust faster. Two features of (4) are worth noting. Persistence enters the decision rule directly: when $\rho = 0$, the target is zero and no firm reallocates, so reallocation is worth undertaking only to the extent that opportunity gaps endure. And because $x_{i,t}$ is a weighted average of past targets, effort shares mean-revert rather than drifting, so the constraint $s_{i,j,t} \in [0, 1]$ binds only when opportunity gaps are large relative to χ.[10]

Expected innovation value for type i is $V_{i,t} = \theta_i + \sum_j s_{i,j,t} v_{j,t} = \theta_i + \mu + 2\,\epsilon_t x_{i,t}$, which exceeds the field average exactly to the extent that the portfolio tilt is aligned with the realized opportunity gap. Iterating (4) and taking expectations gives

$$\mathbb{E}[V_{i,t}] \;=\; \theta_i + \mu \;+\; \frac{2\rho^2\sigma_\epsilon^2}{\chi} \cdot \frac{1-\phi_i}{1-\phi_i\rho}, \tag{5}$$

which is decreasing in κ_i: nimbler firms align their portfolios more closely with high-value fields and produce more valuable innovations. The reward to doing so is increasing in dispersion, and it is zero when $\rho = 0$.

3.4. The Reversal

Let $\Delta \equiv \mathbb{E}[V_{P,t}] - \mathbb{E}[V_{U,t}]$ denote the private-firm advantage in innovation value. Differencing (5) across types gives

$$\Delta \;=\; -A \;+\; \Lambda(\rho)\,\sigma_\epsilon^2, \qquad \Lambda(\rho) \;\equiv\; \frac{2}{\chi} \cdot \frac{\rho^2(1-\rho)\,(\phi_U - \phi_P)}{(1-\phi_P\rho)(1-\phi_U\rho)} \;>\; 0. \tag{6}$$

The private-firm advantage is the sum of a fixed productivity deficit and a nimbleness premium that scales linearly with dispersion. Private firms therefore lead in innovation value if and only

[10]The problem is concave, so the constrained optimum is the projection of (4) onto $[-\frac{1}{2}, \frac{1}{2}]$. Imposing the constraint leaves the results below essentially unchanged: in simulations calibrated to a reversal, the dispersion threshold in (7) shifts by a few percent, while the post-reversal gap is attenuated, so the interior solution overstates how fast the gap widens but not when it opens.

if the technological landscape is sufficiently uneven:

$$\Delta > 0 \quad \Longleftrightarrow \quad \sigma_\epsilon^2 \; > \; \sigma^{2*} \; \equiv \; \frac{A}{\Lambda(\rho)}. \tag{7}$$

An economy that crosses this threshold exhibits precisely the pattern documented in Figure I: private-firm innovation value below public-firm value while dispersion is low, above it once dispersion is high. The reversal requires no change in organizational form, no change in κ_P or κ_U, and no change in the productivity gap A. It requires only that the environment move.

The persistence term $\Lambda(\rho)$ governs how much dispersion is required, and it is non-monotonic. It vanishes at $\rho = 0$, where reallocation is pure noise-chasing, and again as $\rho \to 1$, where opportunity gaps are permanent and the slower firm eventually catches up regardless of its adjustment cost. Nimbleness is therefore most valuable at intermediate persistence, with Λ maximized around $\rho \approx 0.8$ for plausible friction parameters. This has a direct implication for our empirical strategy. Persistence enters the model as a condition that must hold for the mechanism to operate at all, not as a margin along which the reversal is generated: raising ρ while holding dispersion fixed does not produce a sign flip, and at high ρ it works against one.[11] Accordingly, Section 4.2 asks whether persistence is high enough for reallocation to be economically meaningful, rather than whether it shifts, while Section 4.1 tests for a break in dispersion.

3.5. Interpretation

The model attributes the reversal to a race between two organizational advantages whose relative value is set by the technological environment. The parameter θ_i captures what public listing buys—scale, complementary assets, patient capital, and the accumulated research

[11] Conditional on σ_ϵ^2, persistence is a substitute for speed rather than a complement, because a sluggish portfolio remains well aligned with an opportunity that has not moved. Appendix A shows that the two readings of a change in persistence differ: holding the arrival variance σ_u^2 fixed, so that greater persistence also raises realized dispersion, Δ is strictly increasing in ρ; holding σ_ϵ^2 fixed, Λ is single-peaked and vanishes at both $\rho = 0$ and $\rho \to 1$.

infrastructure of the mid-century corporate laboratory—while κ_i captures what it costs in organizational inertia, disclosure frictions, and the internal coordination required to redirect a large research program. What decides between them is how much dispersion in technological opportunity there is to exploit.

Read this way, the mid-century dominance of public firms is not evidence that the public form was inherently superior for innovation, nor is the post-1985 reversal evidence that it became inferior. Both are consistent with a stable ranking of θ and κ and a moving environment. Rising dispersion over the post-war period is plausibly linked to the ICT revolution and the arrival of general-purpose technologies, the growth of specialized scientific domains, and the intensified market pricing of technological fields. The emergence of venture capital following the 1979 ERISA Simplification Act and especially the clarification of the Prudent Man Rule (Gompers and Lerner 2001) operates on the other side of (7), lowering κ_P for young private firms and hence lowering the dispersion required for the reversal to occur.

The model delivers a sharp implication that organizes the empirical work that follows: the private-firm advantage in innovation value should be increasing in the dispersion of technological opportunities, and the point at which private firms overtake public firms should coincide with a discrete increase in that dispersion. Testing this implication is the main focus of the paper.

4. The Evolution of the Technological Opportunity Landscape

The model in Section 3 implies that the private-firm advantage in innovation value takes the form $\Delta = -A + \Lambda(\rho)\,\sigma_\epsilon^2$, where σ_ϵ^2 is the dispersion of technological opportunities across fields and $\Lambda(\rho)$ is a positive function of their persistence. The two features of the environment therefore enter in different ways. Persistence is a condition for the mechanism to operate at all: if opportunity gaps were transitory, firms reallocating today would be chasing noise, $\Lambda(\rho)$ would be zero, and organizational differences in adjustment speed would be irrelevant. Given

that reallocation is worth undertaking, dispersion determines the stakes, and a discrete increase in dispersion moves the economy across the threshold in (7).

This section characterizes both features of the landscape. We show that field-level differences in innovation value are persistent throughout the sample, and that their cross-field dispersion increases discretely in the mid-1980s. Section 5 then examines the behavioral implication, namely that private firms respond to this change by reallocating inventive effort toward higher-value fields.

Throughout, we measure the technological opportunity landscape using patents granted to *public* firms only. Because private-firm patent values are imputed while public-firm values are not, this restriction grounds our characterization of the environment in observed market valuations. It also ensures that the timing we document is identified entirely from the comparison group: no private-firm data enter the construction of the series in this section, so any coincidence between changes in the landscape and the reversal in Figure I cannot be an artifact of a common construction.

4.1. Dispersion Across Technology Fields

To quantify how the dispersion of technological opportunities evolves over time, we construct a time series of the cross-sectional variance of innovation values across technology fields. In each year, we compute the value of a technology field as the mean Kogan et al. (2017) value of patents granted to public firms in that field, and we then estimate the cross-field variance of these field-level values for each year in the sample. To quantify sampling uncertainty, we construct 95% confidence intervals using bootstrap standard errors clustered at the technology-field level.

To assess whether the evolution of dispersion exhibits discrete shifts rather than gradual drift, we test for structural breaks in the variance series using the multiple-break methodology of Bai and Perron (1998). Specifically, we model the variance series as an autoregressive process of order 1 and estimate the model in first differences to ensure stationarity. We follow the sequential testing procedure proposed by Bai and Perron, which tests for the presence of multiple

breaks at unknown dates. Break dates are chosen to minimize the sum of squared residuals subject to a trimming factor that enforces a minimum separation between breaks. Consistent with Bai and Perron (2006), we set the trimming factor to 20% of the sample length to allow for serial correlation. Because the estimation is conducted in first differences, the procedure identifies changes in the process governing dispersion—that is, shifts in its trend—rather than one-time jumps in its level.

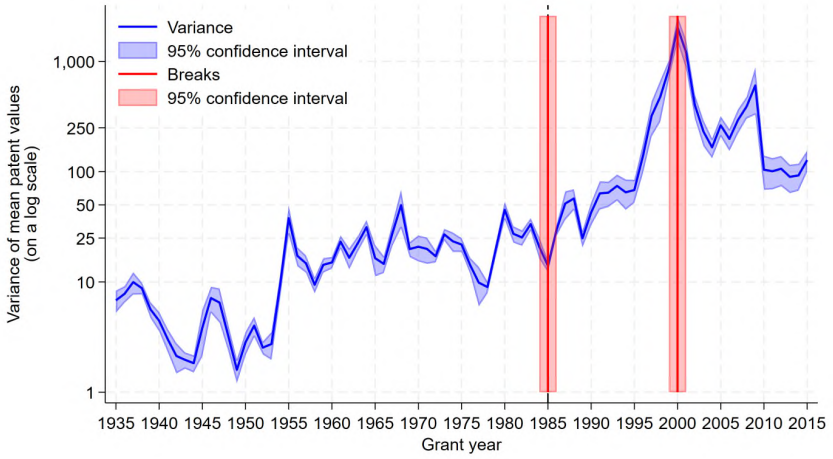
Figure II plots the resulting variance series together with 95% confidence intervals. Vertical red lines indicate the structural breaks identified by the Bai and Perron (1998) procedure, with shaded regions denoting 95% confidence intervals for the estimated break dates. The primary break of interest occurs in 1985, with a 95% confidence interval of 1984–1986. Cross-field dispersion is flat, and if anything mildly declining, over the three decades preceding the break; from 1985 onward it rises by more than an order of magnitude over the following fifteen years. The break therefore marks the point at which the technological opportunity landscape begins a sustained divergence, and its timing coincides closely with the onset of the reversal in relative innovation value between private and public firms documented in Figure I.

A second break occurs around 2000, coinciding with the bursting of the dot-com bubble, after which dispersion declines but remains elevated relative to much of the twentieth century. To verify that the mid-1980s break is not an artifact of the dot-com period, we re-estimate the break timing excluding post-1999 observations. As Figure IA.4 in the Online Appendix) shows, this yields a structural break in 1986, with a 95% confidence interval spanning 1985–1987. Thus, the timing of the first break is unaffected by the behavior of the series at the end of the sample.

4.2. Persistence of Technological Opportunities

While dispersion determines the stakes associated with reallocating innovative effort across technology fields, persistence determines whether such reallocation is worth undertaking. If differences in field-level innovation value are short-lived, rapid movement across fields amounts

Figure II: Dispersion of Innovation Value Across Technology Fields



Notes: This figure plots the cross-sectional variance of technology-field innovation values over time. Field-level values are computed as the mean non-imputed market value of public-firm patents granted within a technology field in a given year. The blue shaded band denotes 95% confidence intervals for the estimated variance, constructed using bootstrap standard errors clustered at the technology-field level with 100 iterations. Vertical lines indicate estimated structural breaks in the variance series, identified using the multiple-break methodology of Bai and Perron (1998) applied to an autoregressive specification estimated in first differences; the red shaded bands denote 95% confidence intervals for the estimated *break dates*. Break dates are chosen to minimize the sum of squared residuals subject to a trimming factor of 20%, following Bai and Perron (2006). The vertical axis is plotted on a logarithmic scale.

to chasing noise, and in the model the return to nimbleness is exactly zero. If they are persistent, firms that reallocate more quickly can systematically align their innovation portfolios with higher-value opportunities.

Establishing that persistence is high, and that it is stable, does two things for the argument. First, it verifies the condition under which the mechanism operates: $\Lambda(\rho)$ is bounded away from zero, so differences in adjustment speed translate into differences in innovation value. Second, and less obviously, the stability of persistence is what isolates dispersion as the source of the reversal. The model implies that an increase in persistence, holding dispersion fixed, does not generate a sign flip in relative innovation value; if anything, greater persistence allows slower-adjusting firms to catch up. Evidence that persistence does not shift therefore rules out the competing explanation that technological opportunities simply became easier to forecast, and leaves the increase in dispersion documented in Section 4.1 as the operative change in the environment.

We assess persistence using two complementary approaches. First, we present a nonpara-

metric visualization of persistence across the distribution of technology fields. Panel (a) of Figure III displays a heatmap of relative field-level innovation values over time, with technology fields ordered by their mean value in 1935. This visualization allows us to trace whether fields that are relatively high- or low-value at the beginning of the sample remain so subsequently, without imposing parametric structure or focusing on average behavior alone. The heatmap reveals substantial stability in relative rankings over time, with low- and especially high-value fields remaining visually distinct over long horizons. Because the panel necessarily conditions on fields that exist in 1935 and are observed repeatedly thereafter, it excludes technology fields that emerged later in the sample and therefore describes the persistence of the established landscape rather than of the landscape as a whole.[12]

Second, we construct summary measures of persistence based on rank stability. Panels (b) and (c) of Figure III report the fraction of technology fields that remain in the same value decile from one year to the next, plotted as five-year centered moving averages. Panel (b) reports average persistence across all fields, while Panel (c) focuses on persistence among fields initially ranked in the top decile. Both measures indicate that persistence is quantitatively meaningful and stable over time.

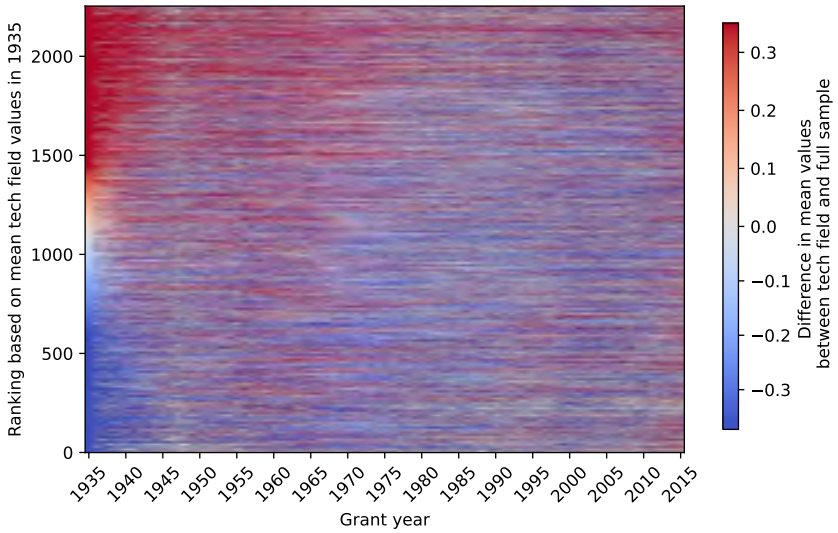
4.3. Persistent Heterogeneity and the Return to Nimbleness

The evidence in Sections 4.1 and 4.2 points to a shift in the technological environment with first-order implications for the value of organizational nimbleness. Throughout the sample period, differences in innovation value across technology fields are sufficiently persistent for reallocation to be economically meaningful. What changes around the mid-1980s is not the durability of these differences, but their magnitude.

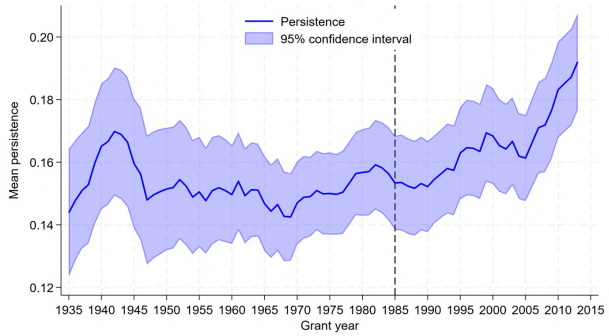
In terms of the model, the two subsections identify the two objects that determine the

[12]This restriction is conservative for our purposes in one respect and not in another. It selects on fields that survived, which tends to overstate persistence; but it also excludes precisely the newly emerging fields associated with the post-1980 increase in dispersion, and so understates the dynamism of the opportunity landscape. Panels (b) and (c) impose no such restriction and yield the same conclusion.

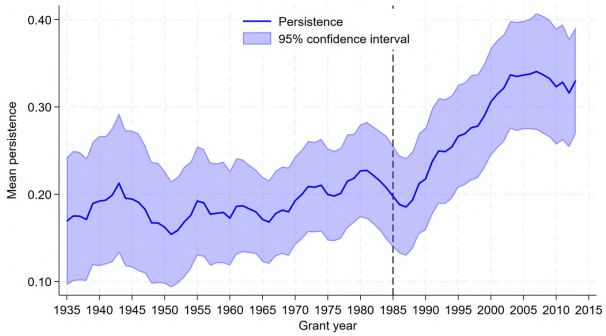
Figure III: Technology Fields: Persistence of Values



Panel (a) Ranking based on mean tech field values in 1935



Panel (b) Average persistence



Panel (c) Persistence in the top decile

Notes: This figure illustrates the persistence of technology-field innovation values over time. Panel (a) presents a heatmap of relative field-level values, with grant years on the x-axis and technology fields on the y-axis, sorted by their mean innovation value in 1935. Colors indicate the deviation of a field's mean value from the cross-field average in a given year. The sample includes 2,254 technology fields with at least one patent in 1935 and in at least 25 additional years between 1935 and 2015; the color scale is truncated at the 25th and 75th percentiles. Panel (b) reports the average persistence of field-level values over quinquennial periods, where persistence is defined as remaining in the same value decile at the beginning and end of the period. Panel (c) reports persistence for fields initially ranked in the top decile. In both panels, we plot 5-year centered moving averages. Shaded bands denote 95% confidence intervals, constructed using bootstrap standard errors clustered at the technology-field level.

private-firm advantage $\Delta = -A + \Lambda(\rho)\,\sigma_\epsilon^2$. Persistence pins down $\Lambda(\rho)$, which is stable; dispersion pins down σ_ϵ^2, which rises by more than an order of magnitude after 1985. Holding Λ fixed, the increase in σ_ϵ^2 raises the nimbleness premium and, once dispersion crosses the threshold $\sigma^{2*} = A/\Lambda(\rho)$ in (7), reverses the ranking of private- and public-firm innovation value. The

timing of the increase in dispersion therefore moves the economy into a regime in which pre-existing organizational differences in adjustment speed translate into systematic differences in innovation value.

This mechanism provides a natural interpretation of the reversal we document. Private firms need not become more nimble over time for their relative innovation value to rise. Rather, a change in the structure of the technological opportunity landscape—toward more heterogeneous fields, with persistence unchanged—makes pre-existing differences in organizational flexibility increasingly consequential. As we show in Section 5 through counterfactual exercises, this change in the technological opportunity landscape plays a central role in generating the subsequent reversal in the relative value of private- and public-firm innovation.

5. Reallocation, Innovation Value, and the Private-Firm Reversal

This section documents the empirical reversal in innovation value between private and public firms and examines its sources. We begin by formalizing the baseline comparison underlying Figure I, describing the construction of our relative-innovation-value measure deferred in the introduction. We then show that the reversal is robust to alternative measures of innovative performance that do not rely on imputed patent values. Next, we use counterfactual decompositions to assess competing explanations for the reversal. These exercises demonstrate that the shift is not driven by changes in the characteristics of patented innovations or by changes in how markets price those characteristics, but instead by differences in the technology fields in which firms choose to innovate. Finally, we explore the cross-sectional sources of the reversal, showing that it is disproportionately driven by young private firms and first-time patenters, consistent with the role of organizational nimbleness emphasized in the model.

5.1. Baseline Measure of Relative Innovation Value

Our baseline measure compares the average value of patents granted to private and public firms in each year. Because patent values are highly right-skewed, we summarize innovation value using the geometric mean. Specifically, for a set of patents granted in year t, the geometric mean value is

$$\mathrm{GM}(\tilde{v}_t) = \exp\left(\frac{1}{N_t}\sum_{i=1}^{N_t}\log(\tilde{v}_{it})\right),$$

where $\tilde{v}_{it}$ denotes the imputed value of patent i and N_t is the number of patents in the group. Our baseline statistic is the ratio of the geometric mean value of patents granted to private firms to that of patents granted to public firms in year t,

$$\frac{\mathrm{GM}(\hat{v}_t^P)}{\mathrm{GM}(\hat{v}_t^U)},$$

where superscripts P and U continue to index private- and public-firm patents, respectively. This statistic is the object plotted in Figure I and serves as our primary measure of relative innovation value over time.

As noted in the introduction, this measure reveals a pronounced reversal in relative innovation value over time. Private firms trail public firms for much of the mid-twentieth century, start catching up in the mid-1980s, reach parity in the mid 1990s, and subsequently emerge as the leading producers of high-value patented innovation.[13]

5.2. Alternative Measures

We next examine whether the private-firm reversal depends on our baseline measure of patent values. To do so, we consider a set of alternative measures of innovative performance that vary in both their economic construction and their reliance on imputed patent values. We begin with an accounting-based measure of R&D value-add, which allows us to assess whether

[13]Figure IA.5 in the Online Appendix shows that using the ratio of arithmetic instead of geometric means results in nearly identical levels and dynamics.

private firms simply increase R&D spending to generate higher innovation value. We then turn to patent-based outcome measures—breakthrough and disruptive innovations per dollar of R&D—that do not rely on market-based valuation or imputation. Finally, we present a nonparametric field-level test that uses only non-imputed public-firm patent values to examine whether private firms disproportionately enter high-value technological fields. Across all four exercises, we obtain the same qualitative pattern as in the baseline, confirming that the reversal is not driven by any single measure or by the imputation procedure.

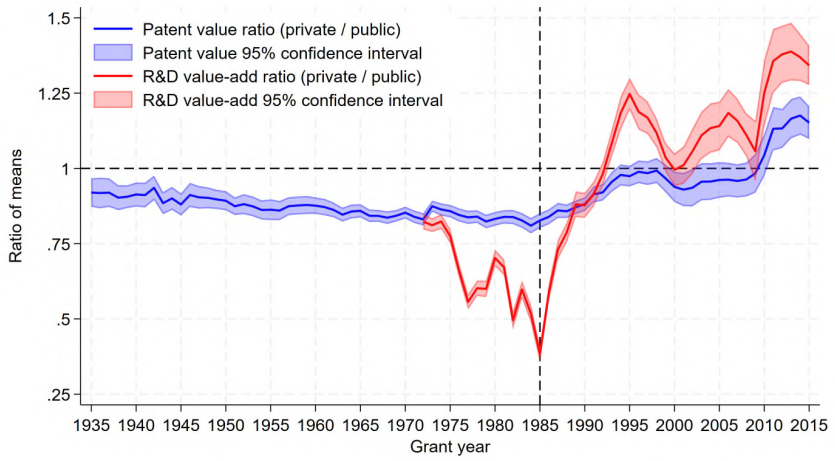
To construct the alternative measures in this section, we require estimates of aggregate R&D spending by public and private firms. The measurement of public-firm R&D expenditures is facilitated by changes in U.S. accounting standards in the 1970s. In 1972, the Accounting Principles Board and the Securities and Exchange Commission made the disclosure of R&D expenses mandatory for public firms, and in 1975 the Financial Accounting Standards Board required firms to expense all R&D expenditures in the year incurred rather than capitalizing them (Nix and Nix 1992). These reforms allow us to observe firm-level R&D spending for publicly listed U.S. firms from 1972 onward using Compustat data. We aggregate these firm-level expenditures to obtain annual public-firm R&D spending and infer private-firm R&D spending by subtracting public-firm R&D from the National Science Foundation's estimates of total U.S. corporate R&D expenditures.

Our first robustness exercise uses these data to construct a measure of *R&D value-add*. For each firm type and year, we compute the aggregate market value of patented innovation using the Kogan-style patent values and subtract aggregate R&D expenditures, normalizing by the number of patents granted. This measure captures the net-of-R&D-cost value of the average patent and directly addresses the concern that private firms may generate higher innovation value simply by spending more on R&D.

Figure IV plots the ratio of average R&D value-add for private and public firms from 1972 through 2015, alongside our baseline relative innovation value series. The reversal is even more pronounced when R&D costs are netted out. Private firms' R&D value-add reaches a low point

in 1985—when the average net value of a private-firm patent is roughly 40% that of a public-firm patent—before rising sharply thereafter. By the early 1990s, private firms reach parity, and by 2015 their average R&D value-add exceeds that of public firms by approximately 30%. These patterns indicate that the baseline measure is, if anything, conservative: the growing private-firm advantage reflects higher net innovation value rather than increased R&D spending.

Figure IV: Private vs. Public R&D value-add



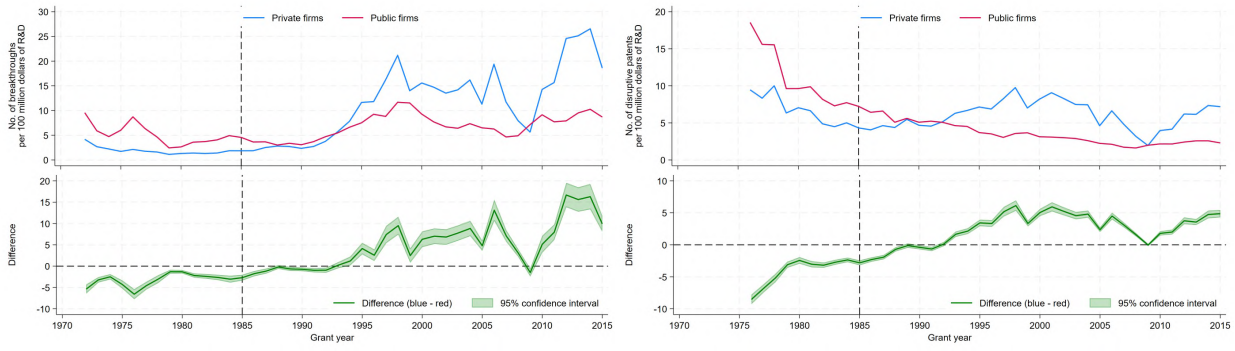
Notes: This figure compares the baseline measure of relative innovation value—the ratio of the geometric mean value of patents granted to private firms relative to public firms—with an alternative measure based on R&D value-add. R&D value-add is defined as the aggregate market value of patented innovation net of aggregate R&D expenditures, normalized by the number of patents granted to each firm type in a given year. Public-firm R&D expenditures are obtained from Compustat and are available from 1972 onward following mandatory disclosure and expensing requirements. Private-firm R&D expenditures are inferred by subtracting public-firm R&D from the National Science Foundation's estimates of total U.S. corporate R&D spending. The figure shows that the reversal in relative innovation value is even more pronounced when innovation value is measured net of R&D costs. The 95% confidence intervals are based on standard errors clustered at the technology-field level.

Our second and third robustness exercises use patent-based outcome measures that do not rely on market-based valuation or imputation. Specifically, we compute the number of breakthrough patents—defined following Kelly et al. (2021)—and the number of disruptive patents—defined following Park, Leahey, and Funk (2023) and Macher, Rutzer, and Weder (2024)—per $100 million of R&D spending for private and public firms. These measures capture high-impact and paradigm-shifting innovation outcomes using observable patent characteristics alone.

Figure V reports the resulting time series. Panels (a) and (b) plot, respectively, the number

of breakthrough and disruptive patents per $100 million of R&D for private and public firms, along with the difference between the two series and corresponding 95% confidence intervals. Consistent with the baseline results, private firms produce significantly fewer breakthrough and disruptive patents per dollar of R&D in the early years. This pattern reverses in the mid-1980s: private firms reach parity by the late 1980s and, from the early 1990s onward, consistently outperform public firms on both measures.

Figure V: Public vs. Private Breakthroughs and Disruptive Patents



Panel (a) Breakthrough patents per R&D $ Panel (b) Disruptive patents per R&D $

Notes: This figure presents alternative, outcome-based measures of innovative performance for public and private firms, scaled by real R&D spending (in constant 1982 dollars). Panel (a) reports the number of breakthrough patents—defined following Kelly et al. (2021)—per $100 million of R&D for private and public firms, as well as the difference between the two series. Panel (b) reports the corresponding measure for disruptive patents, based on the data shared by Macher, Rutzer, and Weder (2024) and measured as patents in the top decile of their disruptiveness measure adjusted for grant-year fixed effects. R&D expenditures for public firms are obtained from Compustat and are available from 1972 onward; private-firm R&D expenditures are inferred by subtracting public-firm R&D from the National Science Foundation's estimates of total U.S. corporate R&D spending. Shaded bands denote 95% confidence intervals for the private–public differences, clustered at the technology-field level.

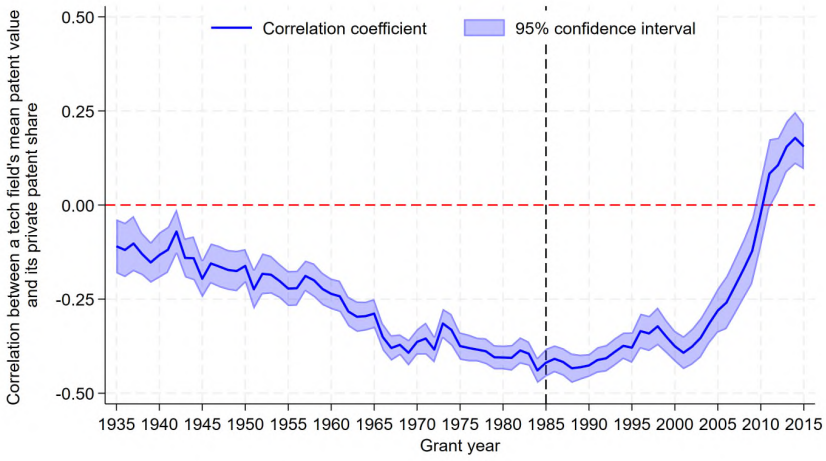
As a final and particularly stringent robustness exercise, we examine whether private firms disproportionately enter high-value technology fields using a measure that does not rely on any imputation of private-firm patent values. For each year, we compute the cross-sectional correlation between a technology field's innovation value and its private-firm share of patents. Field-level innovation value is measured as a ten-year moving average of non-imputed patent values based exclusively on public-firm patents, ensuring that the measure reflects only observed market valuations. We use moving averages to smooth short-run fluctuations in patent values

and focus on persistent differences across fields.

The analysis is conducted on annual cross sections of technology fields. Correlation estimates are weighted by the number of public and private patents granted in each field-year. This weighting reflects both the precision of field-level value estimates and the fact that patents are very unevenly distributed across fields, ensuring that the analysis captures the behavior of the average patent rather than the average technology field. The resulting correlation therefore provides a nonparametric test of whether private firms increasingly concentrate their innovative activity in higher-value technology fields.

Figure VI reports the resulting correlation coefficients over time. The pattern exhibits a pronounced U-shape. Over much of the twentieth century, the correlation between technology-field values and private-firm patent shares becomes increasingly negative, reaching a low point around the mid-1980s. Thereafter, the correlation rises steadily, indicating that private firms increasingly concentrate their patenting activity in higher-value technology fields.

Figure VI: Private-Firm Entry and Non-Imputed Technology-Field Values



Notes: This figure plots annual estimates of the correlation between a technology field's innovation value and its private-firm share of patents. Field-level innovation value is measured as the ten-year moving average of non-imputed patent values based solely on public-firm patents. Correlations are estimated using year-by-year cross sections of technology fields and are weighted by the number of public and private patents granted in each field-year. Shaded bands denote 95% confidence intervals based on standard errors clustered at the technology-field level.

Unlike the baseline and the other alternative measures, the correlation turns positive only in the late 2000s. This delayed crossover is not surprising given the stringency of the test: because field values are constructed exclusively from public-firm patent valuations and smoothed

over a ten-year window, the measure responds slowly to changes in private-firm innovation behavior. Nonetheless, the upward shift in the correlation after the mid-1980s reinforces the central conclusion that private firms increasingly reallocate inventive effort toward higher-value technological opportunities.

5.3. IPO Activity and Mechanical Explanations

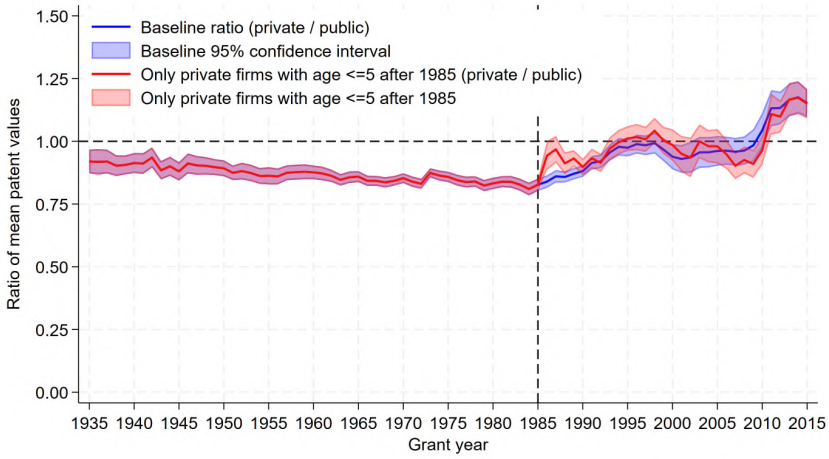
A potential concern is that the reversal in relative innovation value we document reflects mechanical changes in firms' listing status rather than changes in innovative behavior. In particular, the well-documented decline in U.S. IPO activity since the late 1990s raises the possibility that high-quality innovative firms increasingly stay private, so that innovation that would previously have been attributed to public firms is mechanically reclassified as private. Under this view, the observed reversal would arise from shifts in the composition of public and private firms rather than from a change in how firms allocate innovative effort.

To address this concern, we construct a counterfactual that isolates innovation by firms for which IPO timing is unlikely to play a role. Specifically, we recompute the relative innovation value series restricting the private-firm sample to firms that are at most five years old at the time a patent is granted. These young firms are, by construction, too early in their life cycle for delayed IPO decisions to account for their classification as private, and they cannot represent mature public-firm substitutes that simply remain private for longer. The public-firm benchmark is unchanged.

Figure VII reports the results. The reversal in relative innovation value remains clearly visible, with timing and magnitude similar to the baseline series. If anything, the catch-up of private firms appears slightly sharper among young firms. This evidence rules out IPO-related mechanical explanations: delayed IPOs and changes in listing behavior operate through mature firms remaining private, whereas the reversal we document is already present at the point of entry. The findings therefore indicate that the rise of private-firm innovation value reflects changes in innovative behavior and opportunity allocation, rather than shifts in how innovation

is classified across public and private firms.

Figure VII: The Relative Value of *Young*-Private- and Public-Firm Innovation



Notes: This figure compares the baseline measure of relative innovation value—the ratio of the geometric mean value of patents granted to private firms relative to public firms—with an alternative measure that restricts the post-1985 private-firm sample to firms that are at most five years old at the time of patent grant. Firm age is defined as the patent grant year minus the firm's self-reported founding year from the National Establishment Time Series (NETS) database. The 95% confidence intervals are based on standard errors clustered at the technology-field level.

With the reversal robust across a wide range of measures, and not mechanically driven by compositional changes, the next section examines its sources by asking whether it reflects changes in the characteristics of patented innovations, the valuation of those characteristics, or shifts in the technology fields in which firms choose to innovate.

5.4. Counterfactual Decomposition of the Reversal

The analysis so far establishes two key facts. First, the relative value of innovation produced by private and public firms reverses sharply and robustly beginning in the mid-1980s. Second, this reversal coincides with a change in the technological environment, marked by a sharp increase in the dispersion of innovation values across technology fields. We now turn to the question of *why* this reversal occurs. In particular, we ask whether it reflects changes in the characteristics of patented innovations, changes in how markets value those characteristics, or changes in where firms direct their innovative effort across technologies. To address this question, we exploit the structure of our patent valuation framework to decompose innovation

value into distinct components and to assess their respective roles in generating the observed reversal. This decomposition allows us to evaluate whether the reversal reflects changes in what firms invent, how innovation is priced, or how firms reallocate inventive effort across technology fields as the opportunity landscape evolves.

Our approach is based on the observation that patent value can be decomposed into four elements: the characteristics of the patented invention (what firms patent), the valuation multiples attached to those characteristics (how markets price them), firms' technology-field choices (where firms innovate), and the underlying value of those technology fields (what those fields are worth). In our imputation model, patent value is expressed as a function of observable patent characteristics with time-varying coefficients, along with technology-field effects that capture differences in innovation value across fields. This structure allows us to construct counterfactual innovation value paths by holding one component fixed at its 1985 level while allowing the remaining components to evolve as observed in the data. The logic of the exercise is straightforward. If a counterfactual that fixes a given component closely tracks the observed reversal, then changes in that component cannot explain the reversal. Conversely, if fixing a component eliminates or substantially attenuates the reversal, then changes in that component are central to its emergence.

We implement a sequence of counterfactual exercises that use the structure of our valuation model to rule out candidate explanations for the reversal. We begin by examining whether changes in patent characteristics or in the pricing of those characteristics can account for the pattern. Specifically, we construct counterfactuals that hold fixed, at their 1985 levels, either the valuation multiples attached to patent characteristics or the distribution of patent characteristics themselves, while allowing all other components of patent value to evolve as observed. As shown in Panel (a) of Figure VIII, both counterfactual paths closely track the data. This implies that neither changes in what firms patent nor changes in how markets price patent characteristics can explain the rising relative value of private-firm innovation. In both cases, the near coincidence of the counterfactual and observed series rules out patent characteristics and their pricing as primary drivers of the reversal.

We next turn to counterfactuals that isolate the role of technology fields, shown in Panel (b) of Figure VIII. This panel reports three counterfactual paths. The yellow path holds technology-field values (multiples) fixed at their 1985 levels while allowing firms to reallocate across fields. The red path holds firms' technology-field choices fixed at their 1985 levels while allowing field values to evolve. The purple path holds both field values and field choices fixed at their 1985 levels. Comparing these paths allows us to separately assess the roles of changing technology-field opportunities and firms' reallocation across those opportunities.

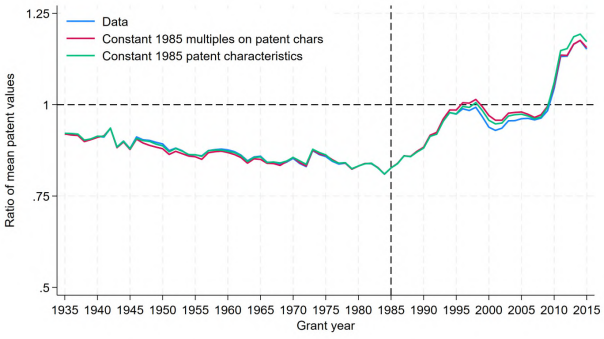
The results reveal that both components are important. Holding technology-field values fixed largely eliminates the post-1985 increase in the relative value of private-firm innovation. This exercise effectively freezes the technological opportunity landscape at its pre-reversal configuration, corresponding to a world in which the structural break in cross-field dispersion documented in Section 4.1 does not occur. In the model, greater dispersion raises the payoff to reallocating innovative effort toward high-value opportunities, thereby increasing the value of organizational nimbleness. The flattening of the private–public innovation gap under this counterfactual is consistent with this prediction: only when opportunities become sufficiently uneven does nimbleness pay off.

Holding firms' field choices fixed attenuates—but does not eliminate—the reversal, implying that changes in where firms innovate also play a central role. When both field values and field choices are fixed at their 1985 levels, the relative value of private- and public-firm patents remains essentially flat, closely resembling the pre-1985 period.

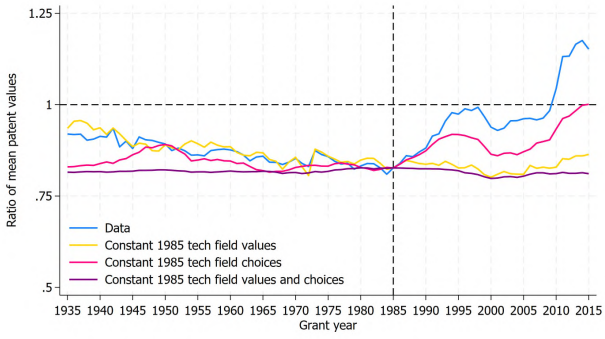
Taken together, these counterfactuals show that the reversal emerges from the interaction between an increasingly uneven technological opportunity landscape and organizational differences in the ability to reallocate innovative effort toward high-value fields.

Panel (c) of Figure VIII attributes the reallocation underlying the reversal by disaggregating the technology-field counterfactuals by firm type. Specifically, we separately hold fixed the 1985 technology-field distribution of public firms and that of private firms, while allowing field values to evolve as observed. Holding constant the field distribution of public firms has little effect on

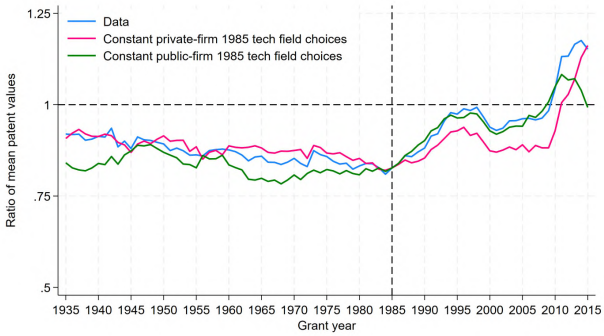
Figure VIII: Counterfactual Decompositions



Panel (a) Patent Characteristics Counterfactuals



Panel (b) Technology-Field Values and Reallocation



Panel (c) Attribution of Reallocation: Private vs. Public Firms

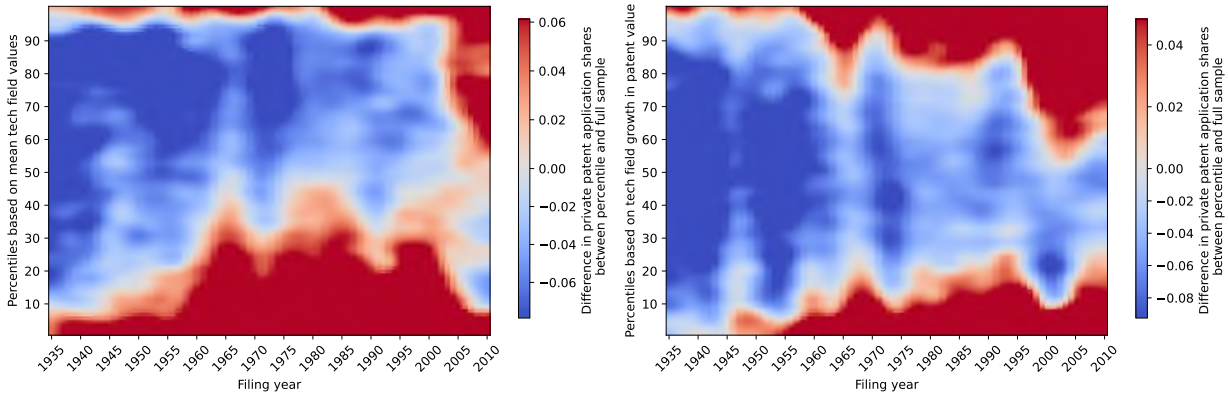
Notes: This figure presents counterfactual paths for the ratio of mean patent values of private firms relative to public firms. Counterfactuals are constructed using the patent valuation model described in Section 2.3, which decomposes patent value into patent characteristics, time-varying valuation multiples, and technology-field effects capturing field-level innovation value. Panel (a) presents two counterfactuals that hold fixed, at their 1985 levels, either the valuation multiples attached to patent characteristics or the distribution of patent characteristics themselves, while allowing all other components of patent value to evolve as observed. Panel (b) reports counterfactuals that hold fixed technology-field values, technology-field allocations, or both at their 1985 levels. Panel (c) separately holds fixed the 1985 technology-field allocations of public firms and of private firms, allowing field values to evolve in both cases.

the relative innovation value series, which closely tracks the data through the late 2000s. By contrast, holding constant the 1985 field distribution of private firms largely eliminates the post-1985 increase in relative innovation value. This asymmetry indicates that the reversal is driven primarily by changes in where private firms choose to innovate rather than by public firms exiting high-value fields. Consistent with the model's emphasis on organizational nimbleness, private firms actively reallocate toward increasingly valuable technologies, while public firms'

field composition plays a comparatively minor role.[14]

The counterfactual exercises summarize technology-field reallocation through changes in *average* field exposure. We now complement this analysis with a nonparametric visualization that illustrates how private and public firms' patenting shifts across the *distribution* of technology-field opportunities. Figure IX therefore provides a distributional view of the reallocation mechanism emphasized in the model, allowing us to trace how innovative effort moves across the technological opportunity landscape over time without imposing parametric structure or aggregating across fields.

Figure IX: Private-Firm Innovation Across Technology-Field Opportunities



Panel (a) Private Innovation Share by Technology-Field Value

Panel (b) Private Innovation Share by Technology-Field Value Growth

Notes: This figure presents nonparametric evidence on the relationship between technology-field opportunities and private firms' innovation activity. Each panel displays a heatmap showing how the private share of patenting varies across technology fields over time. The x-axis reports patent filing (rather than grant) years, to more closely capture the point in time when a firm might pivot. Each cell reports the difference between the private share of patents in a given technology-field percentile and the private share in the full sample in that year. In Panel (a), technology fields are ranked by their mean patent value; in Panel (b), fields are ranked by the subsequent five-year growth rate of mean patent value. Field rankings are based on non-imputed patent values estimated for public firms following Kogan et al. (2017). Warmer colors indicate fields in which private firms are overrepresented relative to the aggregate, while cooler colors indicate underrepresentation. The color scale is truncated at the 25th and 75th percentiles, and the heatmaps are smoothed using a two-dimensional Gaussian filter with bandwidth $\sigma = 2$. The sample covers patent applications filed between 1935 and 2010, reflecting incomplete coverage of patent filings in later years.

[14]To further examine the timing of private firms' reallocation, Figure IA.6 in the Online Appendix reports counterfactuals that hold fixed private firms' technology-field distributions at different benchmark years. These results show that the most consequential reallocations occurred between the late 1980s and late 1990s.

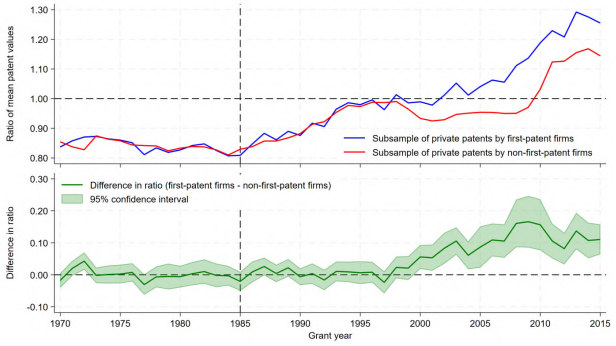
The heatmap in Panel (a) ranks technology fields by their contemporaneous innovation value and shows the relative over- and under-representation of private firms across the distribution in shades of red and blue, respectively. Prior to the mid-1980s, private firms are disproportionately concentrated in lower-value fields and underrepresented in the upper tail of the field-value distribution. Beginning in the mid-1980s (and especially after 2000), this pattern reverses sharply: private firms become increasingly overrepresented in the highest-value fields and underrepresented in the lowest-value fields, consistent with a broad re-sorting of innovative activity across the technological landscape. This shift coincides with the structural break in cross-field dispersion documented in Section 4.1 and mirrors the timing of the reversal in relative innovation value. The heatmap therefore provides a distributional counterpart to the average reallocation identified in the counterfactual analysis. Viewed through the lens of the model, it visually illustrates the mechanism emphasized earlier: when technological opportunities become more uneven, nimbler private firms increasingly concentrate innovative effort in the upper tail of the opportunity distribution.

Panel (b) focuses on a forward-looking measure of innovation opportunities by ranking technology fields based on subsequent growth in innovation value. Here again, private firms' patenting tilts increasingly toward fields that experience higher future value growth, particularly after the mid-1980s. This pattern indicates that private firms' reallocation is not limited to avoiding declining technologies but reflects a systematic shift toward emerging, high-growth areas of the frontier. Consistent with the maintained hypothesis that private firms are nimbler than public firms, these visualizations reinforce the conclusion that the reversal in relative innovation value is driven by active reallocation across technology fields, with private firms increasingly positioning their inventive activity in the upper tail of the opportunity distribution as the dispersion of technological opportunities rises.

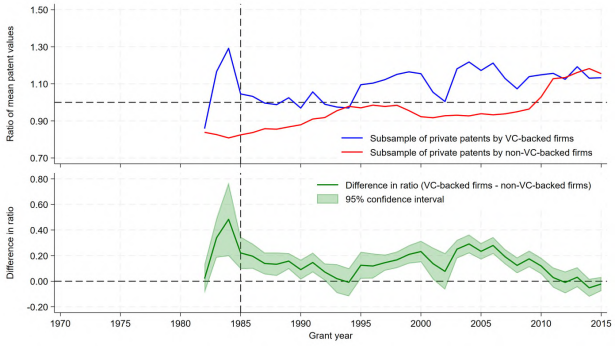
5.5. Cross-Sectional Evidence

We next examine which private firms drive the reallocation toward high-value technology fields documented above. Figure X reports the relative innovation value of private firms across several dimensions that proxy for organizational flexibility and exposure to frontier innovation environments.

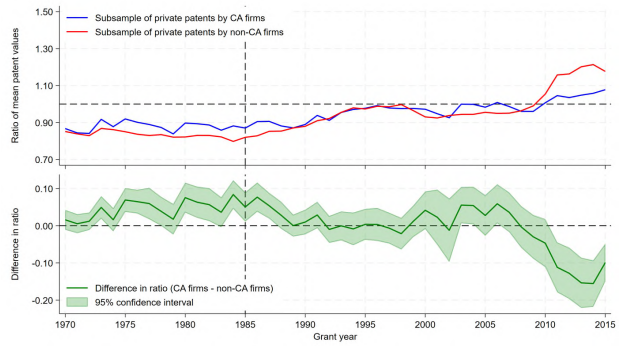
Figure X: Cross-section of Private Firms: Who Drives the Reversal?



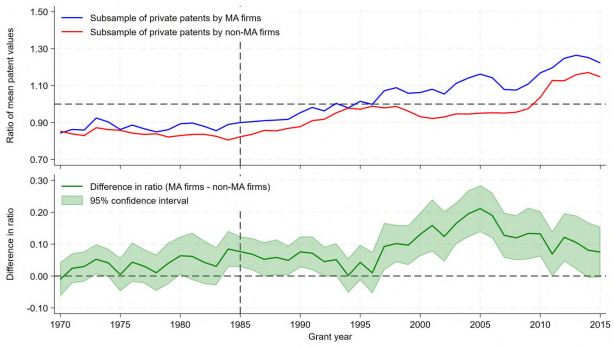
Panel (a) First-time



Panel (b) VC-backed



Panel (c) California



Panel (d) Massachusetts

Notes: This figure presents the evolution of the relative innovation value of private firms across subsamples defined by firm characteristics and location. Each panel plots the ratio of the mean patent value of private firms to that of public firms for two mutually exclusive subsamples (top), along with the difference between the two ratios (bottom). Panel (a) compares first-time private patenters to repeat private patenters. Panel (b) compares VC-backed and non-VC-backed private firms. A firm is classified as VC-backed if it received venture capital financing on or before the patent grant date, based on VentureSource or VentureXpert. Panels (c) and (d) compare private firms located in California and Massachusetts, respectively, to private firms located elsewhere in the United States. Shaded areas in the lower sub-panels indicate 95% confidence intervals. Patent values are constructed as described in Section 2.

The reversal is concentrated among first-time patenters (many of which are presumably startups). Panel (a) shows that private firms patenting for the first time move from clear

laggards to clear leaders relative to public firms beginning in the mid-1980s. This pattern indicates that entry and early-stage experimentation play a disproportionate role in the rise of private-firm innovation value, consistent with lower adjustment frictions among firms without legacy technology portfolios.

Panels (b) through (d) highlight the importance of financing and geography. VC-backed private firms outperform public firms earlier and by a wider margin than the average private firm, suggesting that access to risk-tolerant capital amplifies the ability to pursue high-value technological opportunities. Similarly, private firms located in California and Massachusetts—regions characterized by dense scientific infrastructure and active VC markets—exhibit especially strong gains in relative innovation value, with California-based private firms losing their edge after 2010.

Taken together, the cross-sectional patterns reinforce the interpretation that the reversal is driven by a subset of private firms that are particularly nimble: young entrants operating in frontier innovation ecosystems and supported by specialized finance. These firms are best positioned to respond when technological opportunities become more dispersed, completing the link between the aggregate reversal and the underlying reallocation mechanism.

Summary. Taken together, the evidence in this section supports the model's central mechanism: once dispersion in technology-field values rises, private firms—especially young and VC-backed entrants—reallocate more nimbly toward high-value opportunities, generating the observed reversal.

6. Macroeconomic Patterns

The rise in the relative value of private-firm innovation would be of limited economic importance if it did not translate into a meaningful reallocation of aggregate innovation-value creation. We therefore examine whether the micro-level reversal we document aggregates to the macro level.

Figure XI, Panel (a) shows that the shift in innovation leadership extends to the aggregate level. It plots the private-firm share of aggregate R&D value-add, defined following Ebrahimian and Ljungqvist (2025) as the value of patented innovation net of R&D expenditures. Since the mid-1980s, this share has nearly tripled. The timing of this increase closely mirrors the reversal in relative patent values documented earlier.

Panels (b) and (c) clarify the source of this macro reallocation. Panel (b) shows that the private-firm share of patenting remains relatively stable over time. Public firms continue to account for the majority of patents granted throughout the sample. The growing contribution of private firms is therefore not driven by an increase in the quantity of innovation.

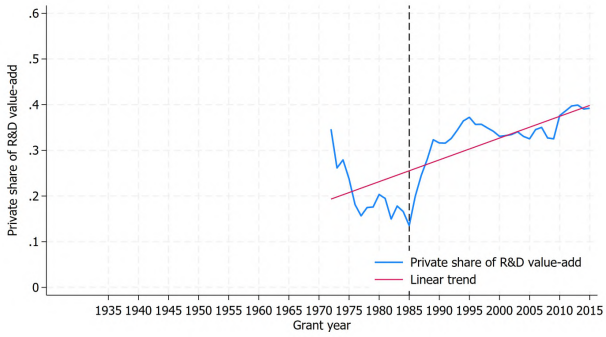
Panel (c) shows that the private-firm share of aggregate R&D spending declines sharply beginning in the mid-1980s. Public firms account for an increasing fraction of total corporate R&D expenditures over time. The rise in the private-firm share of R&D value-add therefore reflects a pronounced divergence between innovative inputs and innovative outputs: private firms generate a growing share of innovation value while accounting for a shrinking share of R&D spending.

Taken together, these patterns point to a macro reallocation of innovative value rather than innovative effort. As technological opportunities become more dispersed, private firms—particularly those able to pursue opportunities nimbly—generate substantially more value per dollar of R&D. A rising share of innovation-driven value creation thus originates outside public equity markets. This shift has implications for how innovation-driven growth is measured (as the behavior of private firms is relatively opaque and thus understudied) and how it is understood (as private firms are not widely investible).

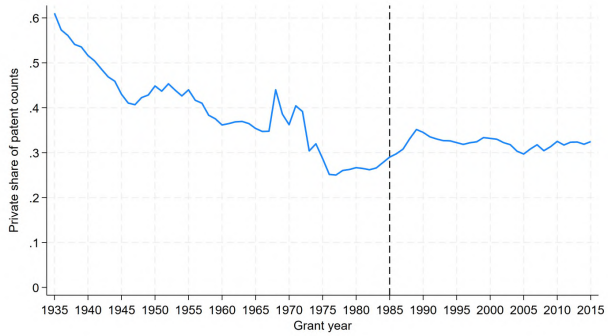
7. Conclusion

This paper documents a long-run reversal in the relative value of innovation produced by private and public firms in the United States. After lagging public firms for much of the twentieth century, private firms have emerged since the mid-1980s as the leading producers of high-value

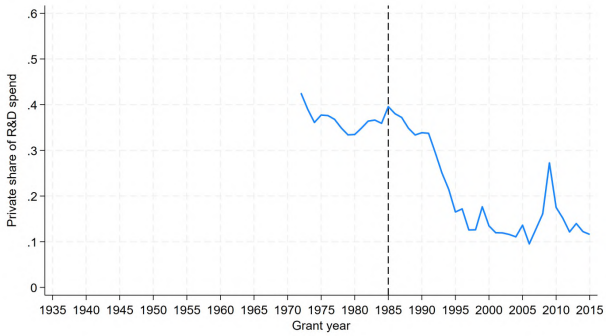
Figure XI: Private Firms and the Aggregate Value of Innovation



Panel (a) Private share of R&D value-add



Panel (b) Private patent share



Panel (c) Private share of R&D

Notes: This figure reports the aggregate implications of the rise in private-firm innovation value. Panel (a) shows the private-firm share of aggregate R&D value-add, defined as the total value of patented innovation net of R&D expenditures. Panel (b) shows the private-firm share of total patents granted. Panel (c) shows the private-firm share of aggregate corporate R&D spending. Patent values are constructed as described in Section 2. Aggregate R&D spending for public firms is obtained from Compustat beginning in 1972, following the introduction of mandatory R&D disclosure and expensing requirements. Aggregate R&D spending for private firms is estimated as total corporate R&D spending reported by the National Science Foundation minus aggregate R&D spending by public firms.

patented innovation. We show that this shift is not driven by changes in patent characteristics, by how markets price innovation, or by greater innovative scale. Instead, it reflects differences in how firms reallocate innovative effort across technology fields as the opportunity landscape evolves.

To interpret these patterns, we develop a simple model in which firms differ in their ability to adjust innovative activity across technologies. The model highlights a key interaction: when technological opportunities are relatively uniform or short-lived, organizational nimbleness confers little advantage; when opportunities become more dispersed and persistent, by contrast,

differences in technology-field values widen and rapid reallocation generates higher innovation value. Consistent with this mechanism, we document a sharp increase in the dispersion of technology-field values beginning in the mid-1980s, alongside high and rising persistence. Private firms respond more quickly to these shifts, reallocating toward higher-value fields and generating greater value per dollar of R&D.

The consequences of this reallocation extend beyond individual patents. While accounting for a stable share of corporate patenting, private firms now generate a growing share of aggregate R&D value-add through a combination of producing more valuable inventions and doing so at lower R&D cost. The rise of private-firm innovation is therefore not merely a micro-level reordering of innovators, but a macro-level shift in where economically valuable technological progress originates.

Our findings speak to a growing literature documenting rising dispersion in economic outcomes and its implications for innovation and growth. A growing body of work emphasizes that technological change since the late twentieth century has become increasingly uneven, generating large gains in some domains while leaving others relatively unaffected (e.g., Autor et al. 2020; Gordon 2016). While much of this literature focuses on firm-level outcomes such as concentration, productivity, or market power, our evidence highlights a complementary dimension of unevenness operating at the level of technological opportunity. We show that the dispersion of innovation value across technology fields increases sharply beginning in the mid-1980s, raising the stakes associated with where firms choose to innovate. In this environment, firms that can reallocate innovative effort more rapidly toward high-value fields gain a systematic advantage. Viewed through this lens, the rise of private-firm innovation reflects not only changes in firm behavior, but a shift in the technological landscape itself—one in which organizational flexibility translates into persistent differences in innovation value.

More broadly, our findings underscore that innovation outcomes depend on the interaction between organizational form and the technological environment. Who innovates matters—but when it matters depends on the nature of the technology frontier. As scientific and technolog-

ical opportunities continue to evolve, understanding how institutions shape the allocation of innovative effort will remain central to understanding long-run economic growth.

References

Acharya, Viral, and Zhaoxia Xu. 2017. "Financial Dependence and Innovation: The Case of Public Versus Private Firms." *Journal of Financial Economics*, 124(2): 223–243.

Aggarwal, Vikas A., and David H. Hsu. 2014. "Entrepreneurial Exits and Innovation." *Management Science*, 60(4): 867–887.

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth Through Creative Destruction." *Econometrica*, 60(2): 323–351.

Aghion, Philippe, and Peter Howitt. 1998. *Endogenous Growth Theory.* Cambridge, MA:MIT Press.

Aghion, Philippe, John Van Reenen, and Luigi Zingales. 2013. "Innovation and Institutional Ownership." *American Economic Review*, 103(1): 277–304.

Akcigit, Ufuk, and William R. Kerr. 2018. "Growth through Heterogeneous Innovations." *Journal of Political Economy*, 126(4): 1374–1443.

Arrow, Kenneth J. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, ed. Richard R. Nelson, 609–626. Princeton, NJ:Princeton University Press.

Asker, John, Joan Farre-Mensa, and Alexander Ljungqvist. 2015. "Corporate Investment and Stock Market Listing: A Puzzle?" *Review of Financial Studies*, 28(2): 342–390.

Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen. 2020. "The Fall of the Labor Share and the Rise of Superstar Firms." *Quarterly Journal of Economics*, 135(2): 645–709.

Baber, William R., Patricia M. Fairfield, and James A. Haggard. 1991. "The Effect of Concern about Reported Income on Discretionary Spending Decisions: The Case of Research and Development." *The Accounting Review*, 66(4): 818–829.

Bai, Jushan, and Pierre Perron. 1998. "Estimating and Testing Linear Models with Multiple Structural Changes." *Econometrica*, 66(1): 47–78.

Bai, Jushan, and Pierre Perron. 2006. "Multiple Structural Change Models: A Simulation Analysis." In *Econometric Theory and Practice: Frontiers of Analysis and Applied Research*, ed. Dean Corbae, Steven N. Durlauf, and Bruce E. Hansen, 212–238. Cambridge University Press.

Bergeaud, Antonin, and Cyril Verluise. 2024. "A New Dataset to Study a Century of Innovation in Europe and in the U.S." *Research Policy*, 53(1): 104903.

Bernstein, Shai. 2015. "Does Going Public Affect Innovation?" *Journal of Finance*, 70(4): 1365–1403.

Bernstein, Shai, Xavier Giroud, and Richard R. Townsend. 2016. "The Impact of Venture Capital Monitoring." *Journal of Finance*, 71(4): 1591–1622.

Bloom, Nicholas, John Van Reenen, and Heidi Williams. 2019. "A Toolkit of Policies to Promote Innovation." *Journal of Economic Perspectives*, 33(3): 163–184.

Bresnahan, Timothy F., and Manuel Trajtenberg. 1995. "General Purpose Technologies: 'Engines of Growth'?" *Journal of Econometrics*, 65(1): 83–108.

Brown, James R., Steven M. Fazzari, and Bruce C. Petersen. 2009. "Financing Innovation and Growth: Cash Flow, External Equity, and the 1990s R&D Boom." *Journal of Finance*, 64(1): 151–185.

Cong, Lin William, and Sabrina T. Howell. 2021. "Policy Uncertainty and Innovation: Evidence from Initial Public Offering Interventions in China." *Management Science*, 67(11): 7238–7261.

Dechow, Patricia M., and Richard G. Sloan. 1991. "Executive Incentives and the Horizon Problem: An Empirical Investigation." *Journal of Accounting and Economics*, 14(1): 51–89.

Ebrahimian, Mehran, and Alexander Ljungqvist. 2025. "Private Innovation and Wealth Inequality." Stockholm School of Economics Working Paper.

Fleming, L., H. Greene, G. Li, M. Marx, and D. Yao. 2019. "Government-funded Research Increasingly Fuels Innovation." *Science*, 364(6446): 1139–1141.

Gao, Huasheng, Po-Hsuan Hsu, and Kai Li. 2018. "Innovation Strategy of Private Firms." *Journal of Financial and Quantitative Analysis*, 53(1): 1–32.

Gibbons, Robert, and Kevin J. Murphy. 1992. "Optimal Incentive Contracts in the Presence of Career Concerns." *Journal of Political Economy*, 100(3): 468–505.

Gompers, Paul, and Josh Lerner. 2001. "The venture capital revolution." *Journal of economic perspectives*, 15(2): 145–168.

Gordon, Robert J. 2016. *The Rise and Fall of American Growth: The U.S. Standard of Living Since the Civil War.* Princeton, NJ:Princeton University Press.

Gornall, Will, and Ilya A Strebulaev. 2021. "The Economic Impact of Venture Capital: Evidence from Public Companies." *Unpublished Working Paper*.

Hall, Bronwyn H, Adam B Jaffe, and Manuel Trajtenberg. 2001. "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools." National Bureau of Economic Research Working Paper 8498.

Hassan, Tarek Alexander, Aakash Kalyani, and Pascual Restrepo. 2026. "The Skill Premium in Times of Rapid Technological Change." National Bureau of Economic Research NBER Working Paper 34939.

Helpman, Elhanan, ed. 1998. *General Purpose Technologies and Economic Growth.* Cambridge, MA:MIT Press.

Holmström, Bengt. 1979. "Moral Hazard and Observability." *Bell Journal of Economics*, 10(1): 74–91.

Holmström, Bengt. 1982. "Moral Hazard in Teams." In *Bell Journal of Economics*324–340. RAND Corporation.

Holmström, Bengt. 1989. "Agency Costs and Innovation." In *The Incentive Problem in Innovation*, ed. D. Bos, and P. Mössner, 51–66. Springer.

Holmström, Bengt. 1999. "Managerial Incentive Problems: A Dynamic Perspective." *The Review of Economic Studies*, 66(1): 169–182.

Howell, Sabrina T, Josh Lerner, Ramana Nanda, and Richard R Townsend. 2020. "How Resilient is Venture-Backed Innovation? Evidence from Four Decades of U.S. Patenting." National Bureau of Economic Research Working Paper 27150.

Jones, Benjamin F. 2009. "The Burden of Knowledge and the "Death of the Renaissance Man": Is Innovation Getting Harder?" *Review of Economic Studies*, 76(1): 283–317.

Jones, Charles I. 2005. "Growth and Ideas." *Handbook of Economic Growth*, 1B: 1063–1111.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. "Measuring Technological Innovation over the Long Run." *American Economic Review: Insights*, 3(3): 303–20.

Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *Journal of Political Economy*, 112(5): 986–1018.

Kline, Patrick, Neviana Petkova, Heidi Williams, and Owen Zidar. 2019. "Who Profits from Patents? Rent-Sharing at Innovative Firms." *Quarterly Journal of Economics*, 134(3): 1343–1404.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. 2017. "Technological Innovation, Resource Allocation, and Growth." *Quarterly Journal of Economics*, 132(2): 665–712.

Kortum, Samuel, and Josh Lerner. 2000. "Assessing the Contribution of Venture Capital to Innovation." *RAND Journal of Economics*, 31(4): 674–692.

Larrain, Borja, Gordon M Phillips, Giorgo Sertsios, and Francisco Urzúa I. 2025. "The Effects of Going Public on Firm Profitability and Strategy." *The Review of Financial Studies*, 38(8): 2467–2514.

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.

Macher, Jeffrey T., Christian Rutzer, and Rolf Weder. 2024. "Is There a Secular Decline in Disruptive Patents? Correcting for Measurement Bias." *Research Policy*, 53(5): 104992.

Mollica, Marcos, and Luigi Zingales. 2007. "The Impact of Venture Capital on Innovation and the Creation of New Businesses." *Unpublished Working Paper*.

Nanda, Ramana, and Matthew Rhodes-Kropf. 2013. "Investment Cycles and Startup Innovation." *Journal of Financial Economics*, 110(2): 403–418.

Nix, Paul E., and David E. Nix. 1992. "A historical review of the accounting treatment of research and development costs." *Accounting Historians Journal*, 19(1): 51–78.

Park, Michael, Erin Leahey, and Russell J Funk. 2023. "Papers and Patents Are Becoming Less Disruptive Over Time." *Nature*, 613(7942): 138–144.

Rosenberg, Nathan. 1982. *Inside the Black Box: Technology and Economics.* Cambridge:Cambridge University Press.

Stein, Jeremy C. 1989. "Efficient Capital Markets, Inefficient Firms: A Model of Myopic Corporate Behavior." *Quarterly Journal of Economics*, 104(4): 655–669.

Stein, Jeremy C. 2002. "Information Production and Capital Allocation: Decentralized vs. Hierarchical Firms." *Quarterly Journal of Economics*, 117(5): 189–228.

Tian, Xuan, and Tracy Yue Wang. 2011. "Tolerance for Failure and Corporate Innovation." *The Review of Financial Studies*, 27(1): 211–255.

Trajtenberg, Manuel, Rebecca Henderson, and Adam Jaffe. 1997. "University Versus Corporate Patents: A Window on the Basicness of Invention." *Economics of Innovation and New Technology*, 5(1): 19–50.

Veihl, Maria. 2022. "The Dark Side of Patents: Strategic Patenting, Product Market Entry and Competitor Innovation." *Academy of Management Proceedings*, 2022(1): 17819.

A. Proofs

A.1. Setup

This appendix states and proves the results reported in Section 3. We restate the environment for completeness.

There are two technology fields $j \in \{A, B\}$ with period-t values

$$v_{A,t} = \mu + \epsilon_t, \qquad v_{B,t} = \mu - \epsilon_t,$$

where the opportunity gap follows a stationary $AR(1)$ process,

$$\epsilon_t = \rho\, \epsilon_{t-1} + u_t, \qquad u_t \sim \text{i.i.d.}\,(0, \sigma_u^2), \qquad \rho \in [0, 1),$$

with unconditional variance $\sigma_\epsilon^2 = \sigma_u^2/(1-\rho^2)$ and autocovariance $\mathbb{E}[\epsilon_t \epsilon_{t-k}] = \rho^k \sigma_\epsilon^2$. Firm types $i \in \{P, U\}$ choose effort shares $s_{i,j,t} \in [0,1]$ with $\sum_j s_{i,j,t} = 1$ to maximize

$$\mathbb{E}_{t-1}\big[\Pi_{i,t}\big] = \theta_i + \sum_j s_{i,j,t}\, \mathbb{E}_{t-1}[v_{j,t}] - \frac{\chi}{2} \sum_j s_{i,j,t}^2 - \frac{\kappa_i}{2} \sum_j \big(s_{i,j,t} - s_{i,j,t-1}\big)^2,$$

with $\chi > 0$, $\kappa_i > 0$, $\kappa_P < \kappa_U$, and $A \equiv \theta_U - \theta_P > 0$. Effort is committed one period in advance, so $s_{i,t}$ is measurable with respect to information dated $t-1$.[15] Throughout we write

$$x_{i,t} \equiv s_{i,A,t} - \tfrac{1}{2} \in [-\tfrac{1}{2}, \tfrac{1}{2}], \qquad \phi_i \equiv \frac{\kappa_i}{\chi + \kappa_i} \in (0, 1),$$

and note that ϕ_i is strictly increasing in κ_i, so that $\phi_P < \phi_U$.

Reduction. Since $s_{i,B,t} = 1 - s_{i,A,t}$, we have $\sum_j s_{i,j,t} \mathbb{E}_{t-1}[v_{j,t}] = \mu + 2\rho\epsilon_{t-1} x_{i,t}$, $\sum_j s_{i,j,t}^2 = \tfrac{1}{2} + 2x_{i,t}^2$, and $\sum_j (s_{i,j,t} - s_{i,j,t-1})^2 = 2(x_{i,t} - x_{i,t-1})^2$. The firm's problem is therefore equivalent

[15] While econometricians proxy ϵ using stock-market reactions observed at the time of patent grant, researchers and engineers working at the technological frontier often receive earlier signals about these technological opportunities, for example through scientific breakthroughs. The one-period commitment lag should be read as the interval between acting on such a signal and observing the market value of the resulting invention.

to

$$\max_{x_{i,t}\in[-\frac{1}{2},\frac{1}{2}]} \left\{ \theta_i + \mu - \tfrac{\chi}{4} \right\} + 2\rho\,\epsilon_{t-1}x_{i,t} - \chi x_{i,t}^2 - \kappa_i\big(x_{i,t} - x_{i,t-1}\big)^2. \tag{A.8}$$

The objective is strictly concave, with second derivative $-2(\chi + \kappa_i) < 0$.

A.2. Optimal Reallocation

Proposition 1 (Partial adjustment)**.** *The optimal effort allocation satisfies*

$$x_{i,t} = \phi_i\, x_{i,t-1} + (1-\phi_i)\, x_t^*, \qquad x_t^* \equiv \frac{\rho\,\epsilon_{t-1}}{\chi},$$

whenever the constraint $x_{i,t} \in [-\frac{1}{2}, \frac{1}{2}]$ is slack, and equals the projection of this expression onto $[-\frac{1}{2}, \frac{1}{2}]$ otherwise. The adjustment speed $1 - \phi_i = \chi/(\chi + \kappa_i)$ is strictly decreasing in κ_i, so private firms close a strictly larger fraction of the gap to the target each period.

Proof. Differentiating (A.8) with respect to $x_{i,t}$ and setting the derivative to zero gives

$$2\rho\epsilon_{t-1} - 2\chi x_{i,t} - 2\kappa_i\big(x_{i,t} - x_{i,t-1}\big) = 0 \quad \Longrightarrow \quad (\chi + \kappa_i)\, x_{i,t} = \rho\epsilon_{t-1} + \kappa_i x_{i,t-1},$$

and dividing by $\chi+\kappa_i$ yields the stated rule, using $(1-\phi_i)/\chi = 1/(\chi+\kappa_i)$. Because the objective is strictly concave and the feasible set is a closed interval, the constrained maximizer is the projection of the unconstrained maximizer onto that interval. Finally, $\partial\phi_i/\partial\kappa_i = \chi/(\chi+\kappa_i)^2 > 0$. □

Lemma A.1 (Moving-average representation)**.** *In the stationary solution with the constraint slack,*

$$x_{i,t} = \frac{(1-\phi_i)\rho}{\chi} \sum_{k\geq 0} \phi_i^k\, \epsilon_{t-1-k}.$$

Proof. Iterate Proposition 1 backwards and use $|\phi_i| < 1$ to discard the initial condition. □

Two remarks follow. First, because the weights $(1-\phi_i)\phi_i^k$ are non-negative and sum to one, $x_{i,t}$ is a convex combination of the targets $\{\rho\epsilon_{t-1-k}/\chi\}$; effort shares therefore mean-revert

rather than accumulating shocks, and $|x_{i,t}| \leq \rho \sup_t |\epsilon_t|/\chi$. Second, the constraint binds only when opportunity gaps are large relative to χ, and imposing it changes the results below only quantitatively. In simulations calibrated so that the economy crosses the reversal threshold of Proposition 3, the threshold itself moves by between one and four percent across the calibrations we examine, while the post-reversal gap is attenuated; the interior solution therefore overstates the rate at which the gap widens with dispersion but not the point at which it opens.

A.3. Innovation Value

Proposition 2 (Expected innovation value)**.** *Let $V_{i,t} = \theta_i + \sum_j s_{i,j,t} v_{j,t}$. Then*

$$\mathbb{E}\big[V_{i,t}\big] = \theta_i + \mu + \frac{2\rho^2\sigma_\epsilon^2}{\chi} \cdot \frac{1-\phi_i}{1-\phi_i\rho},$$

which is strictly decreasing in κ_i for any $\rho \in (0,1)$, and equal to $\theta_i + \mu$ when $\rho = 0$.

Proof. Using $s_{i,B,t} = 1 - s_{i,A,t}$,

$$\sum_j s_{i,j,t} v_{j,t} = \mu + 2\epsilon_t x_{i,t},$$

so $\mathbb{E}[V_{i,t}] = \theta_i + \mu + 2\mathbb{E}[\epsilon_t x_{i,t}]$. Substituting Lemma A.1 and using $\mathbb{E}[\epsilon_t \epsilon_{t-1-k}] = \rho^{k+1}\sigma_\epsilon^2$,

$$\mathbb{E}\big[\epsilon_t x_{i,t}\big] = \frac{(1-\phi_i)\rho}{\chi} \sum_{k\geq 0} \phi_i^k \rho^{k+1} \sigma_\epsilon^2 = \frac{(1-\phi_i)\rho^2\sigma_\epsilon^2}{\chi} \sum_{k\geq 0} (\phi_i\rho)^k = \frac{(1-\phi_i)\rho^2\sigma_\epsilon^2}{\chi\,(1-\phi_i\rho)},$$

where the geometric sum converges because $\phi_i\rho < 1$. For monotonicity, let $g(\phi) \equiv (1-\phi)/(1-\phi\rho)$. Then

$$g'(\phi) = \frac{-(1-\phi\rho) + \rho(1-\phi)}{(1-\phi\rho)^2} = \frac{-(1-\rho)}{(1-\phi\rho)^2} < 0 \quad \text{for } \rho < 1,$$

and ϕ_i is increasing in κ_i by Proposition 1. Setting $\rho = 0$ gives $x_{i,t} = 0$ for all t and the final claim. □

The interpretation is that innovation value exceeds the field average exactly to the extent that

a firm's portfolio tilt is aligned with the realized opportunity gap. Speed improves alignment, and alignment is what nimbleness buys. When $\rho = 0$, the forecastable component of the gap is zero, no firm reallocates, and organizational differences are irrelevant.

A.4. The Reversal

Proposition 3 (Reversal threshold)**.** *Let* $\Delta \equiv \mathbb{E}[V_{P,t}] - \mathbb{E}[V_{U,t}]$*. Then*

$$\Delta = -A + \Lambda(\rho)\,\sigma_\epsilon^2, \qquad \Lambda(\rho) \equiv \frac{2}{\chi} \cdot \frac{\rho^2(1-\rho)(\phi_U - \phi_P)}{(1-\phi_P\rho)(1-\phi_U\rho)} > 0 \quad \textit{for } \rho \in (0,1),$$

and consequently

$$\Delta > 0 \iff \sigma_\epsilon^2 > \sigma^{2*} \equiv \frac{A}{\Lambda(\rho)}.$$

Both signs of Δ *are attainable for fixed* $(\theta_P, \theta_U, \kappa_P, \kappa_U, \chi, \rho)$*, so a rise in dispersion across the threshold reverses the ranking of private- and public-firm innovation value without any change in organizational primitives.*

Proof. Differencing Proposition 2 across types,

$$\Delta = (\theta_P - \theta_U) + \frac{2\rho^2\sigma_\epsilon^2}{\chi}\big[g(\phi_P) - g(\phi_U)\big], \qquad g(\phi) = \frac{1-\phi}{1-\phi\rho}.$$

Placing the bracket over a common denominator,

$$g(\phi_P) - g(\phi_U) = \frac{(1-\phi_P)(1-\phi_U\rho) - (1-\phi_U)(1-\phi_P\rho)}{(1-\phi_P\rho)(1-\phi_U\rho)} = \frac{(\phi_U - \phi_P)(1-\rho)}{(1-\phi_P\rho)(1-\phi_U\rho)},$$

since the terms in $\phi_P\phi_U\rho$ cancel in the numerator. Substituting and using $A = \theta_U - \theta_P$ gives the stated expression. Positivity of Λ on $(0,1)$ follows from $\phi_P < \phi_U$ and $\phi_i\rho < 1$. The threshold characterization is immediate because Δ is affine and strictly increasing in σ_ϵ^2. Finally, $\Delta < 0$ for σ_ϵ^2 small and $\Delta > 0$ for σ_ϵ^2 large, establishing attainability of both signs. □

The threshold σ^{2*} is increasing in the public-firm productivity advantage A and in the curvature χ, and decreasing in the friction gap $\phi_U - \phi_P$. A reduction in κ_P—for instance through the

emergence of specialized private financing—therefore lowers the dispersion required for the reversal to occur, and the model admits two non-exclusive routes to a reversal at a given date: a rise in σ_ϵ^2 or a fall in σ^{2*}.

A.5. Dispersion and Persistence

Proposition 4 (Comparative statics)**.**

(i) Dispersion. $\partial\Delta/\partial\sigma_\epsilon^2 = \Lambda(\rho) > 0$ *for all* $\rho \in (0,1)$.

(ii) Persistence, holding realized dispersion fixed. $\Lambda(0) = 0$, $\lim_{\rho\to 1}\Lambda(\rho) = 0$, *and* $\Lambda(\rho) > 0$ *on* $(0,1)$. *Hence* Λ *attains an interior maximum, and* Δ *is not monotone in* ρ.

(iii) Persistence, holding the arrival variance fixed. *Writing* $\sigma_\epsilon^2 = \sigma_u^2/(1-\rho^2)$,

$$\Delta = -A + \frac{2\sigma_u^2}{\chi}(\phi_U - \phi_P)\cdot\frac{\rho^2}{(1+\rho)(1-\phi_P\rho)(1-\phi_U\rho)},$$

which is strictly increasing in ρ *on* $(0,1)$ *for all admissible parameters.*

Proof. (i) is immediate from Proposition 3, since Δ is affine in σ_ϵ^2 with slope $\Lambda(\rho)$.

(ii) The factor ρ^2 vanishes at $\rho = 0$ and the factor $(1-\rho)$ vanishes at $\rho = 1$, while the denominator $(1-\phi_P\rho)(1-\phi_U\rho)$ is bounded away from zero on $[0,1]$ because $\phi_i < 1$. Continuity and strict positivity on the interior then imply an interior maximum.

(iii) Substituting $\sigma_\epsilon^2 = \sigma_u^2/[(1-\rho)(1+\rho)]$ into $\Lambda(\rho)\sigma_\epsilon^2$ cancels the factor $(1-\rho)$ and yields the stated expression. Write $D(\rho) \equiv (1+\rho)(1-\phi_P\rho)(1-\phi_U\rho)$ and let $S \equiv \phi_P + \phi_U$ and $P \equiv \phi_P\phi_U$, so that

$$D(\rho) = 1 + (1-S)\rho + (P-S)\rho^2 + P\rho^3, \qquad D'(\rho) = (1-S) + 2(P-S)\rho + 3P\rho^2.$$

The expression is increasing in ρ if and only if $\frac{d}{d\rho}\big[\rho^2/D(\rho)\big] > 0$, that is, if and only if $2D(\rho) - \rho D'(\rho) > 0$. Direct computation gives

$$2D(\rho) - \rho D'(\rho) = 2 + (1-S)\rho - P\rho^3.$$

Since $\phi_i \in (0,1)$ we have $S < 2$ and $P < 1$, so for $\rho \in (0,1)$,

$$2 + (1 - S)\rho - P\rho^3 > 2 + \rho - 2\rho - \rho^3 = 2 - \rho - \rho^3 > 0.$$

Hence the derivative is strictly positive throughout. □

Parts (ii) and (iii) are not in conflict; they answer different questions. A rise in persistence holding the arrival variance fixed also raises realized dispersion, and by part (i) that favors nimble firms. Conditional on realized dispersion, however, persistence is a substitute for speed rather than a complement: a sluggish portfolio remains well aligned with an opportunity gap that has not moved, and in the limit $\rho \to 1$ the slower firm converges to the same allocation as the faster one, so the entire nimbleness premium dissipates. Persistence therefore enters the model as a condition for the mechanism to operate—it must be bounded away from zero for reallocation to be worth undertaking at all—rather than as a margin along which a reversal is generated. This is why Section 4.2 establishes that persistence is high and stable rather than testing for a break in it, while Section 4.1 tests for a break in dispersion.

For plausible friction parameters the maximum in part (ii) occurs at high persistence. Setting $\chi = 1$, Λ peaks at $\rho \approx 0.74$ for $(\kappa_P, \kappa_U) = (0.3, 1)$, at $\rho \approx 0.81$ for $(1, 3)$, and at $\rho \approx 0.86$ for $(2, 6)$. The persistence levels documented in Section 4.2 therefore place the economy in the region where the return to nimbleness is largest.

A.6. Forward-Looking Firms

The baseline treats the firm as solving a sequence of one-period problems. Because effort is committed in advance, persistence already enters the decision rule through the forecast $\rho\epsilon_{t-1}$, so the qualitative results do not depend on firms internalizing the future path of opportunities. We record here that they survive when firms do.

Let type i maximize $\mathbb{E}_0 \sum_{t \geq 0} \beta^t \Pi_{i,t}$ with $\beta \in (0,1)$. The Euler equation associated with

(A.8) is

$$\rho\epsilon_{t-1} - \chi x_{i,t} - \kappa_i\big(x_{i,t} - x_{i,t-1}\big) + \beta\kappa_i\,\mathbb{E}_{t-1}\big[x_{i,t+1} - x_{i,t}\big] = 0,$$

and the stationary linear solution takes the form $x_{i,t} = a_i x_{i,t-1} + b_i(\rho)\,\epsilon_{t-1}$, where $a_i \in (0,1)$ is the stable root of

$$\beta\kappa_i a^2 - \big(\chi + \kappa_i + \beta\kappa_i\big)a + \kappa_i = 0, \qquad b_i(\rho) = \frac{a_i\rho}{\kappa_i\big(1 - a_i\beta\rho\big)}.$$

Three properties hold. First, at $\beta = 0$ we have $a_i = \phi_i$ and $b_i(\rho) = (1-\phi_i)\rho/\chi$, so the baseline is recovered exactly. Second, repeating the argument of Proposition 2 with loading b_i and inertia a_i gives

$$\mathbb{E}\big[V_{i,t}\big] = \theta_i + \mu + \frac{2\rho\,\sigma_\epsilon^2\,b_i(\rho)}{1 - a_i\rho},$$

so that Δ remains affine in σ_ϵ^2. The stable root a_i is strictly increasing in κ_i and independent of ρ, and numerical solution confirms that expected innovation value is strictly decreasing in κ_i throughout the admissible parameter range; the slope of Δ in σ_ϵ^2 therefore remains strictly positive and the threshold characterization of Proposition 3 is unchanged in form. Third, $b_i(\rho)$ is strictly increasing in ρ, with $\partial b_i/\partial\rho = a_i/[\kappa_i(1 - a_i\beta\rho)^2] > 0$: a firm that internalizes the future responds more aggressively to a given opportunity gap when gaps are persistent, because the adjustment cost is paid once while the benefit accrues repeatedly. This channel reinforces rather than offsets the comparative static in Proposition 4(iii).

We retain the baseline in the main text because it delivers the transparent threshold in Proposition 3, in which the private-firm advantage is affine in dispersion and the reversal is a single crossing.

B. Microfoundations for the Adjustment-Cost Asymmetry

This appendix provides three complementary microfoundations for the assumption that public firms face higher effective adjustment costs than private firms, $\kappa_U > \kappa_P$. Each microfoundation produces a reduced-form adjustment-cost term of the form

$$\frac{\kappa_i^{\text{eff}}}{2} \sum_j (s_{i,j,t} - s_{i,j,t-1})^2, \quad \kappa_U^{\text{eff}} > \kappa_P^{\text{eff}},$$

consistent with the main model in Section 3. Throughout this appendix we suppress the productivity term θ_i and the concentration term $\frac{\chi}{2} \sum_j s_{i,j,t}^2$ from the objective, since neither depends on the change in effort shares and neither affects the derivation of κ_i^{eff}.

B.1. Arrow's Information Paradox and Appropriation Losses

To convince outside investors and analysts that a pivot to a new technology field is worthwhile, a public firm has to disclose more about what it has learned about emerging technological opportunities. By Arrow's (1962) classic paradox, that disclosure (a) makes the information more valuable to others and (b) less appropriable by the disclosing firm. For a private firm with concentrated owners and no need to persuade anonymous investors, disclosure can be kept internal, so this "leakage" cost is much smaller. We capture this by adding to the firm's payoff an Arrow-type appropriation loss that is increasing in how much the firm changes direction and that is steeper for public firms.

Formally, let a firm reveal an information set $\mathcal{I}_{i,t}$ to external investors when adjusting its R&D portfolio. Let the fraction of information lost through disclosure be $\alpha_i \in (0, 1)$, with

$$\alpha_U > \alpha_P,$$

capturing that public firms lose more informational advantage.

Suppose that losing information reduces the expected surplus from the pivot by

$$L_i = \frac{\alpha_i}{2} \sum_j (s_{i,j,t} - s_{i,j,t-1})^2.$$

Then the firm's objective becomes

$$\Pi_{i,t} = \sum_j s_{i,j,t} v_{j,t} - \frac{\kappa}{2} \sum_j (s_{i,j,t} - s_{i,j,t-1})^2 - L_i.$$

The effective adjustment cost is therefore

$$\kappa_i^{\text{eff}} = \kappa + \alpha_i, \qquad \kappa_U^{\text{eff}} > \kappa_P^{\text{eff}}.$$

where κ is a "technological" adjustment cost (reassign people, labs, etc.), common to both firm types. A public firm's pivot is doubly costly: the firm has to reorganize, and it also bleeds some of its informational advantage when it explains the pivot to the market. A private firm only pays the first part. Arrow's (1962) information paradox therefore provides a natural foundation for higher adjustment frictions in public firms.

B.2. Governance, Verifiability, and Soft vs. Hard Information

Public firms have dispersed shareholders, independent boards, committees, analyst coverage, etc. A big change in R&D direction typically requires more formal approval, more meetings, more documentation, and more career risk for managers. Private firms with concentrated owners can act on soft information (an internal signal that "field A is hot") quickly. We implement this microfoundation by drawing on the literature on verifiability in contracting (Holmström 1979, 1982) and on the economics of hard versus soft information (Stein 2002).

Specifically, we assume that public firms must obtain approval for major strategic pivots from independent boards, which rely on verifiable, "hard" information. Private firms can instead act on "soft" internal information, such as engineer intuition or preliminary scientific feedback.

This creates differences in the precision of the effective signals used for reallocation decisions. Let firms observe imperfect signals of the opportunity gap,

$$\tilde{\epsilon}_{i,t} = \epsilon_t + \eta_{i,t},$$

where $\eta_{i,t}$ is noise. Suppose public firms must rely on hard, verifiable signals and therefore have

$$\mathrm{Var}(\eta_{U,t}) > \mathrm{Var}(\eta_{P,t}).$$

Writing $x_{i,t} = s_{i,A,t} - \frac{1}{2}$ as in Section 3, optimal partial adjustment with respect to the signal yields

$$x_{i,t} = \phi_i^{\text{eff}}\, x_{i,t-1} + \big(1 - \phi_i^{\text{eff}}\big)\frac{\rho\, \tilde{\epsilon}_{i,t-1}}{\chi}, \qquad \phi_i^{\text{eff}} = \frac{\kappa_i^{\text{eff}}}{\chi + \kappa_i^{\text{eff}}},$$

where the effective adjustment cost is

$$\kappa_i^{\text{eff}} \equiv \frac{\kappa}{\lambda_i},$$

and λ_i is a Kalman-gain-type responsiveness parameter increasing in signal precision. Public firms have lower λ_U because the board only acts when signals are very strong (and verifiable), effectively dampening responsiveness. Because $\lambda_P > \lambda_U$, we have

$$\kappa_U^{\text{eff}} > \kappa_P^{\text{eff}}, \qquad \text{and hence} \qquad 1 - \phi_U^{\text{eff}} < 1 - \phi_P^{\text{eff}} :$$

public firms close a smaller fraction of the gap to their target each period. Thus governance constraints and reliance on hard information (Holmström 1979, 1982; Stein 2002) generate slower responsiveness for public firms.

B.3. Stock-Market Penalties, Downside Risk, and Cautious Reallocation

A pivot that doesn't work out is more painful for a public firm: the stock price drops, managers face career risk, covenants may bite, activists may show up. Private owners can tolerate more variance in outcomes and are less exposed to noisy short-term market reactions.

So public firms have stronger ex ante aversion to big reallocations. The idea that public firms face strong stock-price, financing, and managerial-career penalties for failure is central in Stein (1989), Holmström (1999), and Gibbons and Murphy (1992).

We model these mechanisms as a mean–variance decision problem, where pivot failures increase payoff variance. Let innovative payoff be

$$\Pi_{i,t} = \sum_j s_{i,j,t} v_{j,t} + \xi_{i,t},$$

where $\xi_{i,t}$ captures pivot-specific risk, with variance

$$\mathrm{Var}(\xi_{i,t}) = \sigma^2_{\text{pivot}} \sum_j (s_{i,j,t} - s_{i,j,t-1})^2.$$

Suppose public managers solve a mean–variance problem:

$$U_U = \mathbb{E}[\Pi_{U,t}] - \frac{\gamma_U}{2}\mathrm{Var}(\Pi_{U,t}), \qquad \gamma_U > \gamma_P,$$

while private owners have a lower risk penalty γ_P.

A quadratic approximation yields the objective

$$U_i \approx \sum_j s_{i,j,t} v_{j,t} - \frac{1}{2}\left(\kappa + \gamma_i \sigma^2_{\text{pivot}}\right) \sum_j (s_{i,j,t} - s_{i,j,t-1})^2.$$

Thus, the effective adjustment cost is

$$\kappa_i^{\text{eff}} = \kappa + \gamma_i \sigma^2_{\text{pivot}}, \qquad \kappa_U^{\text{eff}} > \kappa_P^{\text{eff}}.$$

Intuitively, because public firms internalize stronger stock-market and career penalties for failed pivots, they act as if they face higher adjustment costs, reoptimizing their R&D portfolios more slowly than private firms. Stock-market discipline, career concerns, and downside-risk exposure (Stein 1989; Holmström 1999; Gibbons and Murphy 1992) therefore imply greater sluggishness

for public firms.

B.4. Summary

Each microfoundation—appropriation losses from disclosure, governance-induced signal attenuation, and differential downside-risk penalties—generates a reduced-form adjustment friction that is higher for public firms than for private firms. The core implication of the main text, $\kappa_U > \kappa_P$, is therefore consistent with multiple economic primitives.

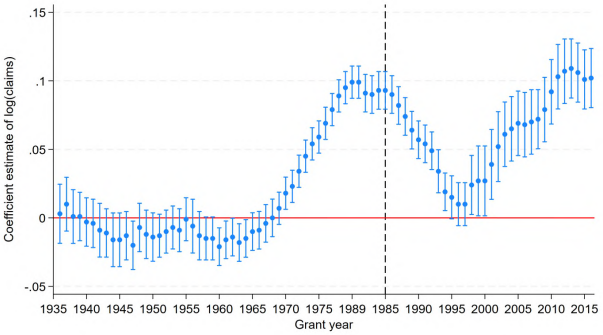
Online Appendix

for

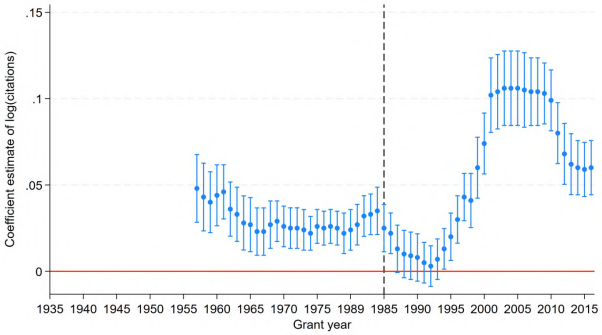
Quietly Ahead: The Diverging R&D Productivity of Public and Private Firms

(NOT INTENDED FOR PUBLICATION)

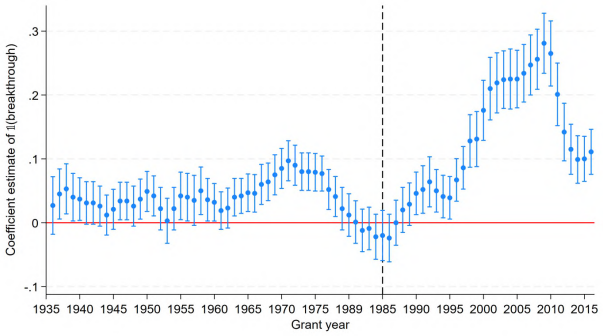
Figure IA.1: Patent-Value Imputation Model Coefficient Estimates



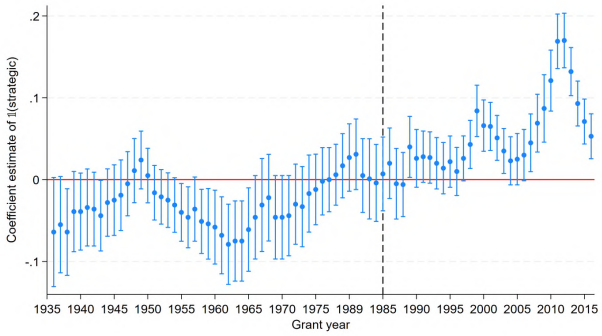
Panel (a) Claims



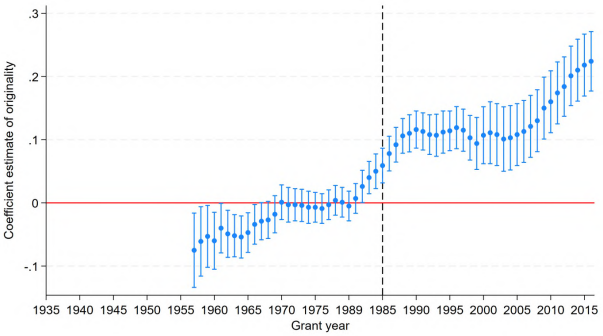
Panel (b) Citations



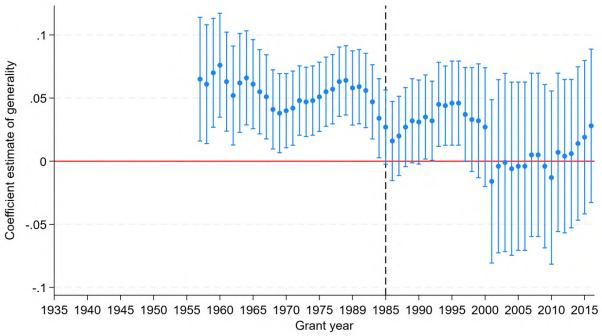
Panel (c) Breakthroughs



Panel (d) Strategic patents



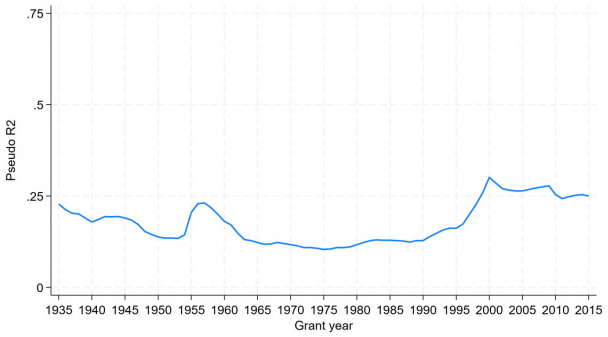
Panel (e) Originality



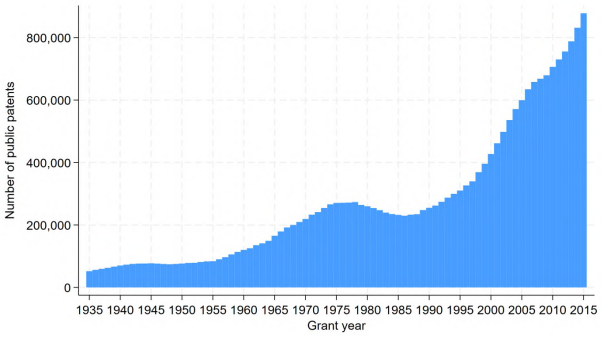
Panel (f) Generality

Notes: This figure plots annual coefficient estimates from the patent-value imputation model in equation 1, estimated using rolling ten-year windows. Data on citations, originality, and generality are unavailable prior to 1947, when the USPTO began reporting citations to prior art on the front page of patents. The blue bars indicate 95% confidence intervals based on standard errors clustered at the technology-field level.

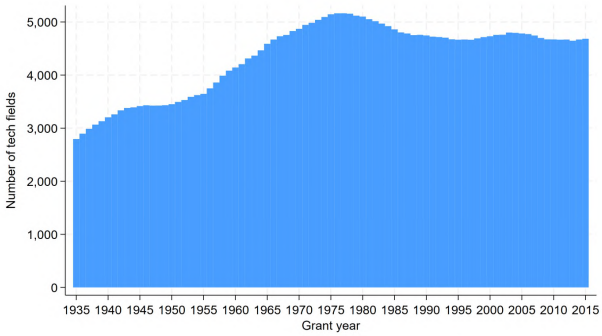
Figure IA.2: Patent-Value Imputation Model: Diagnostics



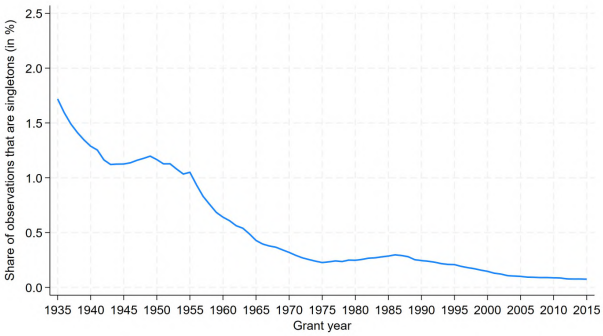
Panel (a) Pseudo $R2$



Panel (b) No. of public patents



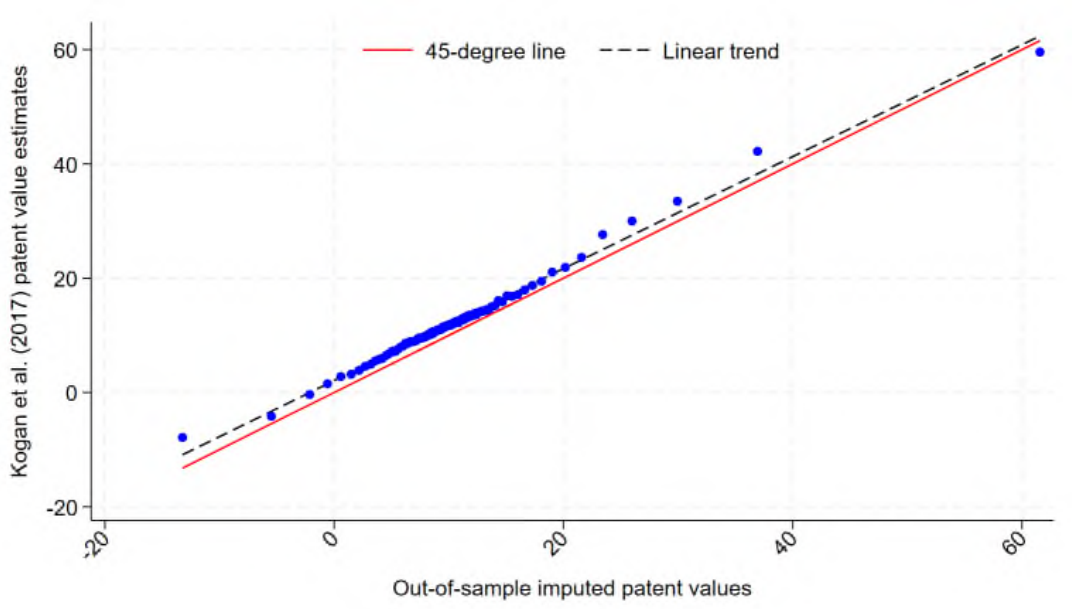
Panel (c) No. of tech fields



Panel (d) Share of singletons

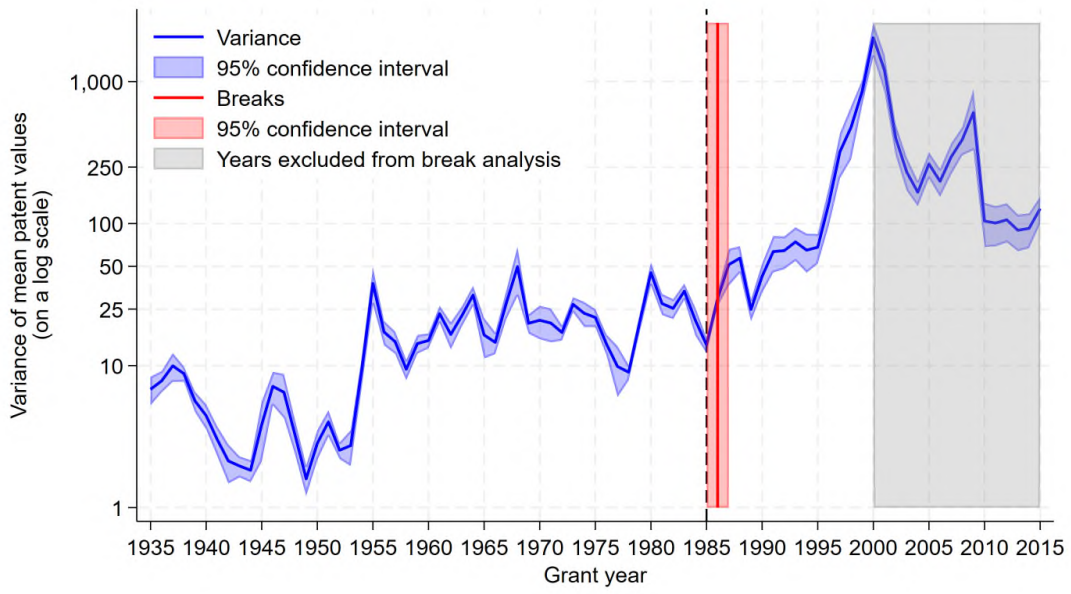
Notes: This figure presents diagnostics of the imputation models (equation (1)) estimated over rolling windows of ten years. Panel (a) presents the Pseudo $R2$ of the imputation models, Panel (b) the number of public patents used to estimate the imputation models, Panel (c) the number of distinct tech fields, and Panel (d) the number of patents that are singletons and therefore excluded.

Figure IA.3: Patent-Value Imputation Model:
Imputed vs. Kogan et al. (2017) Patent Values



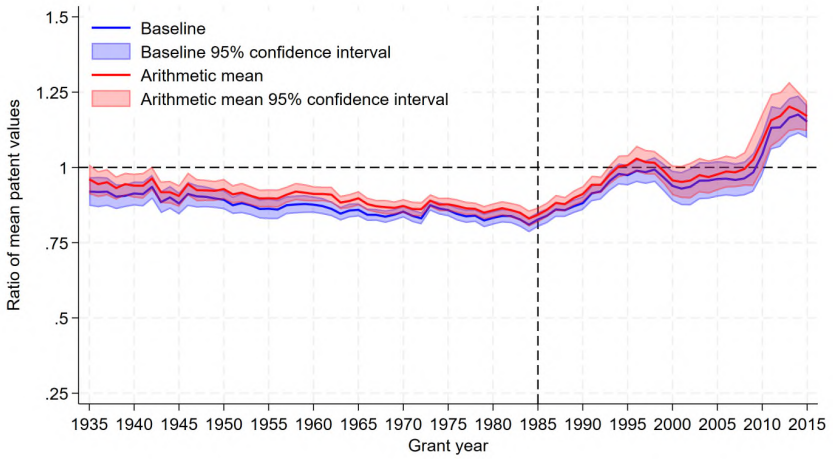
Notes: This figure presents a binned scatter plot comparing out-of-sample imputed patent values to the Kogan et al. (2017) patent-value estimates for public-firm patents. Imputed values are constructed using coefficients estimated from preceding rolling 10-year windows, ensuring that each observation is imputed out of sample with respect to its grant year. Both measures are residualized with respect to grant-year fixed effects, so the comparison reflects cross-sectional variation within year. Each point represents the mean of a bin of observations. The solid line denotes the 45-degree line, and the dashed line shows the fitted linear relationship. The slope of the fitted line is 1.01 using the binned data and 0.94 in the patent-level microdata.

Figure IA.4: Dispersion of Innovation Value Across Technology Fields: Excluding the Period After 1999



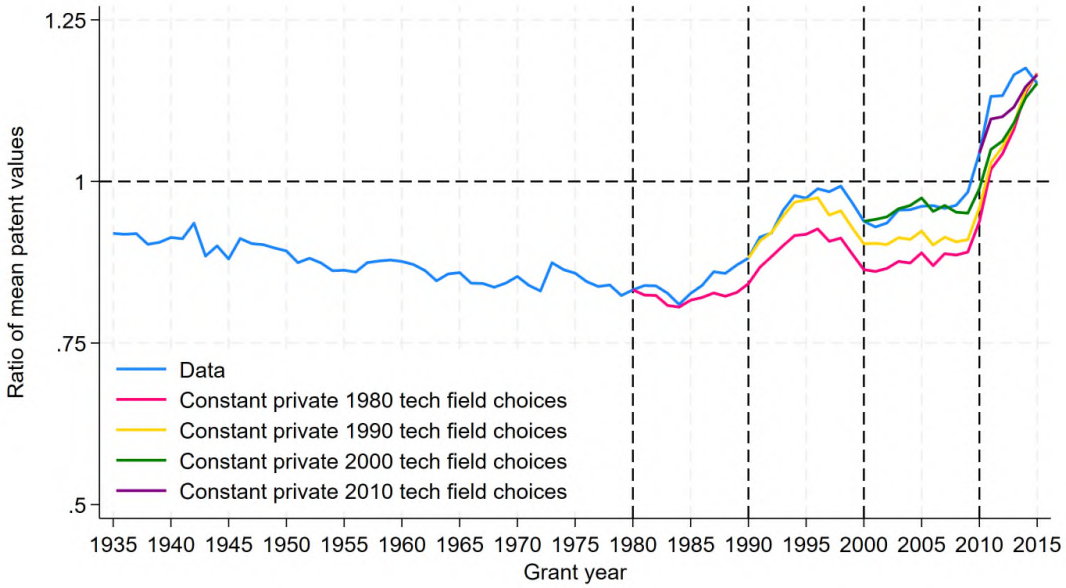
Notes: This figure plots the cross-sectional variance of technology-field innovation values over time. Field-level values are computed as the mean market value of non-imputed public-firm patents granted within a technology field in a given year. Shaded bands denote 95% confidence intervals, constructed using bootstrap standard errors clustered at the technology-field level. Vertical lines indicate estimated structural breaks in the variance series, identified by testing for one structural break using the methodology of Bai and Perron (1998) applied to an autoregressive specification estimated in first differences. Break dates are chosen to minimize the sum of squared residuals subject to a trimming factor of 20%, following Bai and Perron (2006). Unlike in Figure II in the text, the Bai and Perron test is instructed to ignore post-1999 observations to verify that the mid-1980s break is not mechanically driven by the dot-com period.

Figure IA.5: The Relative Value of Private- and Public-Firm Innovation: Geometric versus Arithmetic Means



Notes: This figure compares our baseline measure of relative innovation value—the ratio of the *geometric* mean value of patents granted to private firms relative to public firms—with an alternative version based on the *arithmetic* mean. Shaded areas indicate 95% confidence intervals based on standard errors clustered at the technology-field level.

Figure IA.6: Counterfactual Decompositions



Notes: This figure examines the timing of private firms' technology-field reallocation. Each counterfactual holds fixed the distribution of private firms' technology-field allocations at the beginning of a given decade (1980, 1990, 2000, or 2010), while allowing technology-field values to evolve. Public firms' technology-field allocations are allowed to evolve in all cases.