Disappearing Schumpeterian Growth: From Superstar Scaling to Arrested Dynamism

Ou Liu*

April 24, 2026

Abstract

This paper investigates the rising concentration among the largest U.S. firms and its shifting implications for productivity from 1985 to 2023. I document a fundamental regime shift in the drivers of concentration — from a period of "superstar scaling" to a modern era of "arrested dynamism" — marking the disappearance of Schumpeterian growth at the upper tail of the firm size distribution. Using a micro-audited Melitz-Polanec decomposition, I show that aggregate productivity drivers structurally inverted: alongside collapsing within-firm efficiency, reallocation surged and a severe negative selection drag emerged. Then, using CS-DiD event-study estimates, I found that while acquisitions historically boosted physical productivity, modern acquisitions led to persistent declines in both productivity and internal growth. To explain this reversal, I develop a model endogenizing the choice between synergistic integration and defensive shelving. I uncover two primary mechanisms: rising organizational costs that weaken internal scaling, and high markup stakes that incentivize acquirers to "shelf" targets to prevent cannibalization and markup erosion. My results imply that modern concentration reflects a strategic shift toward protecting market power at the expense of creative destruction, leading to a decoupling of concentration from productive efficiency.

1 Introduction

What determines the ultimate boundary of a firm, and does its massive scale drive or stifle aggregate economic growth? Coase (1937) famously posited that a firm expands only

*Rutgers University. ouliu@rutgers.edu

until the internal frictions of bureaucratic management outweigh the efficiency of transacting in the market. Yet, once a firm successfully overcomes these organizational costs to achieve very large scale, its macroeconomic impact is theoretically ambiguous. Schumpeter (1942) argues that large incumbents serve as core engines of growth, using their scale, resources, and organizational capacity to sustain innovation and displace less efficient producers through creative destruction. Conversely, Arrow (1962) emphasized the replacement effect, whereby market power weakens innovation because incumbents have less to gain from introducing new activity that cannibalizes existing rents. This tension between efficiency-driven expansion and rent-protective entrenchment raises a fundamental question: when economic activities become increasingly concentrated among a few giant firms, does this dominance reflect the triumph of Schumpeterian innovation, or the onset of Arrowian stagnation?

In recent years, this historical tension has re-emerged at the center of a modern macroeconomic debate over the sustained rise in corporate concentration. One prominent view interprets this trend through a Schumpeterian lens, arguing that concentration reflects efficient "superstar scaling," whereby highly productive firms organically expand and capture market share in increasingly winner-take-all environments (Autor, Dorn, Katz, Patterson, and Van Reenen, 2020). An alternative perspective emphasizes the role of entrenched market power, pointing to rising corporate markups and insulated incumbents (De Loecker, Eeckhout, and Unger, 2020). This dichotomy is sharpened by a central macroeconomic puzzle: despite the growing dominance of these large firms, aggregate productivity growth has persistently slowed and measures of business dynamism have steadily declined (Akcigit and Ates, 2021). If modern concentration simply reflects the expansion of highly productive superstars, why has their dominance failed to translate into stronger aggregate productivity growth?

To address this paradox, I show that the sources of top-firm concentration have undergone a structural shift over the past four decades. I establish this by constructing a novel accounting framework that partitions the aggregate sales share growth of the Top 0.01% of U.S. firms (which I refer to as the *top tier*) into two structurally distinct components: (i) *Mechanical M&A*, defined as the direct transfer of pre-existing sales of target firms to acquirers, and (ii) *Organic Growth*, defined as all remaining non-mechanical changes in sales. Crucially, this organic component captures not only pure internal expansion (e.g., via in-house R&D), but also any post-acquisition synergistic scaling.

To distinguish whether this concentration is driven by the expansion of incumbent leaders or by shifts in selection and turnover at the top of the firm size distribution, I further

decompose aggregate organic growth of the top tier into its underlying firm-level flows: incumbent firms that remain in the top tier (*Stayers*), upward mobility into the top tier (*Inflow*), and downward displacement (*Outflow*) or firm death (*Organic Exit*).

Implementing this decomposition requires accurately tracking firm transitions, which is challenging due to firm identifier changes, restructurings, and incomplete M&A linkages. I address this by constructing a micro-level audit of top-tier firm transitions, reconciling firm histories with M&A activity to distinguish true exit, acquisition-driven exit, and reporting artifacts. This ensures that the decomposition reflects underlying economic transitions rather than data artifacts.

Applying this micro-audited decomposition framework over rolling ten-year horizons from 1985 to 2023 reveals two central facts. First, the sources of concentration have shifted over time. Prior to the mid-2000s, top-share growth was driven primarily by organic expansion. Afterwards, organic growth has declined sharply, while mechanical M&A has remained persistently positive, implying that modern concentration increasingly reflects the continuation of mechanical consolidation operating in an environment of weakened organic expansion. Second, the decline in organic growth is driven primarily by a collapse in the expansion of incumbent top-tier firms. While entry and exit remain active, their contributions are relatively stable; the dominant change is the sharp contraction in the growth of incumbent leaders.

To assess whether these changes are efficiency-enhancing, I link concentration dynamics to physical productivity using a markup-adjusted Melitz–Polanec decomposition. This allows me to separate true productivity from revenue effects driven by market power. Two findings emerge. First, the primary source of aggregate productivity growth has shifted. Before 2005, within-firm productivity growth among incumbent leaders was the dominant contributor. After 2005, this component declines sharply, while reallocation becomes the main driver of productivity growth. Second, net entry into the top tier contributes negatively to aggregate productivity. This reflects not a lack of productive inflow, but a pattern of negative selection in the outflow: relatively high-productivity firms exit the top tier, while less productive incumbents remain.

This long-run transition unfolds across four distinct macroeconomic regimes:

- **Era 1 (1985–1994): Schumpeterian Growth.** Within-firm productivity growth was the largest positive component, while both inflow and outflow contributed positively. These dynamics align well with the Schumpeterian benchmark, where aggregate productivity is propelled by the robust scaling of incumbent leaders and ef-

ficient selection (relatively unproductive firms being displaced by more productive entrants).

- **Era 2 (1995–2004): Superstar Scaling.** Massive within-firm growth and positive reallocation drove aggregate productivity to its peak, but displacement dynamics began to change. Inflow turned negative – suggesting incumbent stayers were so exceptionally productive that entrants could not keep pace. Crucially, outflow also turned negative. This indicated that firms leaving the top tier were actually more productive than those remaining, marking the onset of negative selection.

- **Era 3 (2005–2014): Sclerotic Stagnation.** Within-firm productivity growth turned negative, signaling a breakdown in the scaling engine of incumbents. Furthermore, while this period featured high positive inflow, outflow became even more negative than in Era 2. This revealed a severely worsened negative selection mechanism, undermining aggregate efficiency despite the active influx of new firms.

- **Era 4 (2015–2023): Arrested Dynamism.** Aggregate productivity increases are now driven mainly by large positive reallocation, while within-firm growth is remarkably small. Inflow has shrunk, and outflow hovers near zero, indicating that the selection mechanism has stalled completely. Concurrently, the aggregate productivity drag is increasingly driven by a deeply negative *Exit: Acquired* term.

This modern regime presents a stark macroeconomic anomaly. Highly productive target firms are systematically exiting the top tier via acquisition. Yet, despite the sheer volume of ongoing consolidation executed by incumbent leaders, their within-firm productivity growth has entirely flatlined. This aggregate disconnect—the continuous disappearance of productive targets without a commensurate rise in incumbent efficiency—raises a critical micro-level question: when top incumbents execute major acquisitions, do they actually achieve efficiency-enhancing integration, or does the process structurally erode internal productivity?

Because the macroeconomic decomposition cannot observe post-merger integration trajectories, I estimate a doubly robust difference-in-differences (CS-DiD) event study on the physical productivity of incumbent stayers to isolate the causal mechanism. The results confirm a stark inversion in the economic function of M&A. Historically, acquisitions were efficiency-enhancing and supported internal growth. In Era 4, however, major acquisitions inflict a severe, persistent -9.1% structural penalty on the physical productivity of the acquirer. This presents an economic puzzle: if modern acquisitions systematically erode physical efficiency without conferring pricing power, why do rational top-tier in-

cumbents persistently execute them?

To solve this puzzle, I develop a continuous-time quality-ladder model that links firm-level jump-diffusion shocks to the evolution of the firm productivity distribution. Reduced-form data cannot observe the counterfactual threat of displacement or the internal organizational frictions that dictate a firm's strategy; a structural framework is required to infer these hidden motives. In the model, incumbents continuously invest in internal productivity growth while facing exogenous acquisition opportunities. Upon acquiring a target, the incumbent faces a fundamental trade-off. The firm can choose *synergistic integration*, which incorporates the target's technology and raises physical efficiency but cannibalizes the incumbent's legacy product, thereby diluting its markup. Alternatively, the firm can choose *defensive shelving*, neutralizing the target to perfectly preserve its legacy markup, but suffering a productivity rot penalty. The model proves that a firm endogenously chooses defensive shelving whenever the financial premium of preserving its legacy markup—relative to accepting a diluted post-integration markup—strictly outweighs the net physical productivity advantages of absorbing the target's technology.

Estimating this structural framework reveals that the endogenous optimal choice of top-tier firms flipped entirely: from synergistic integration in Era 1 to defensive shelving in Era 4. This transition was triggered by the secular evolution of key structural parameters. As the inferred cost of internal expansion escalated, internal growth weakened. Concurrently, as baseline markups rose over the decades, the sheer financial stakes of protecting legacy rents mechanically overwhelmed the technical benefits of integration.

Finally, the model highlights a policy tradeoff: conventional industrial policy often attempts to stimulate corporate dynamism by subsidizing R&D to lower the cost of internal expansion and boost internal growth. However, stimulating internal growth lowers the firm's effective discount rate, thereby increasing the value of incumbency. Consequently, this intervention paradoxically strengthens the incentive to protect legacy markups, unintentionally reinforcing the modern regime of Arrested Dynamism by driving firms to defensively shelve targets.

Related Literature. I contribute to three strands of the literature. First, this paper contributes to the longstanding debate on whether rising firm concentration reflects the entrenchment of market power or the efficient scaling of productive firms, a tension that traces back to classic industrial organization (e.g., Bain (1951); Demsetz (1973)) and has become central to recent macroeconomic debates over superstar firms, rising markups, and slowing productivity growth. One strand views concentration as the outcome of "su-

perstar" firms capturing a growing share of economic activity in response to technological change that generates winner-take-most markets, with more productive firms expanding and driving aggregate efficiency (Autor, Dorn, Katz, Patterson, and Van Reenen, 2020; Kwon, Ma, and Zimmermann, 2024; Choi, Levchenko, Ruzic, and Shim, 2024). Another emphasizes rising markups, declining investment, and weakening competition among dominant incumbents (De Loecker, Eeckhout, and Unger, 2020; Gutiérrez and Philippon, 2019; Grullon, Larkin, and Michaely, 2019). Despite this progress, existing work largely treats concentration as a single aggregate object, making it difficult to distinguish whether dominance reflects productivity-enhancing scaling or the preservation of market power.

I provide a unified solution by introducing a micro-founded accounting framework that jointly decomposes the growth of top sales share and the sources of aggregate productivity growth. Specifically, I combine a top-share accounting decomposition—separating mechanical consolidation through M&A from non-mechanical (organic) expansion—with a markup-adjusted Melitz–Polanec decomposition of physical productivity over the same horizons. While Covarrubias, Gutiérrez, and Philippon (2020) is close in spirit in arguing that concentration has shifted from efficiency-driven to entrenchment-driven regimes, my approach identifies the underlying firm-level margins driving this transition and directly maps them into productivity outcomes. This reveals a structural regime shift: earlier increases in concentration were driven by organic scaling and positive selection, whereas modern concentration is sustained by acquisition-driven consolidation, negative selection, and is largely decoupled from productivity growth. By probing the micro-level margins underlying concentration, the paper shows that concentration itself is not a sufficient statistic for efficiency: its aggregate consequences depend critically on the mechanisms through which dominance arises.

Second, this paper contributes to the literature on declining business dynamism and the slowdown of creative destruction. Recent work documents a broad-based decline in U.S. dynamism, including falling entry rates, reduced job reallocation, and weaker responsiveness of firms to productivity shocks (Decker, Haltiwanger, Jarmin, and Miranda, 2016, 2020; Haltiwanger, 2015). Macroeconomic frameworks interpret these trends through mechanisms such as a slowdown in knowledge diffusion from frontier firms to laggards (Akcigit and Ates, 2023), distortions in the selection between high- and low-type firms (Acemoglu, Akcigit, Alp, Bloom, and Kerr, 2018), or declining research productivity (Bloom, Jones, Van Reenen, and Webb, 2020). While these explanations rationalize aggregate patterns, they leave open where in the firm size distribution the breakdown of dynamism occurs and through which microeconomic channels it operates.

I provide a complementary perspective by identifying the upper tail of the firm productivity distribution as the locus where Schumpeterian dynamics weaken. Rather than attributing declining dynamism to a lack of entry or diffusion alone, I show that the deterioration in aggregate productivity growth is driven by a breakdown in incumbent expansion and the emergence of negative selection at the top. In particular, the net entry contribution turns negative not because of weak inflow, but because highly productive firms are systematically displaced from the top tier, while relatively less productive incumbents persist. By combining this empirical decomposition with a structural model that characterizes upper-tail dynamics, the paper identifies the microeconomic anatomy of declining dynamism at the top: the collapse of incumbent expansion, the displacement of highly productive firms, and the role of acquisition-driven consolidation in distorting selection.

Third, this paper contributes to the literature on mergers and acquisitions, innovation strategy, and the role of consolidation in shaping technological progress. A growing literature shows that acquisitions can be used not only to integrate productive assets but also to eliminate future competition. Cunningham, Ederer, and Ma (2021) document "killer acquisitions," in which incumbents acquire and discontinue overlapping innovation projects; Roldan-Blanco, Schmitz, and Fons-Rosen (2026) study how such acquisition behavior shapes innovation and welfare in equilibrium; and Argente, Lee, and Moreira (2024) highlight how cannibalization reduces incentives to introduce new products by eroding the value of existing lines.

I extend this literature along three dimensions. First, I provide economy-wide evidence that major acquisitions impose a persistent penalty on the physical productivity of acquiring firms in the modern period. Second, I uncover a distinct mechanism: acquisitions are used not only to eliminate future competitors, but also to protect the value of incumbents' existing market positions, even at the cost of internal efficiency. Third, I develop a structural model in which firms trade off efficiency-enhancing integration against defensive shelving. This framework generalizes the notion of "killer acquisitions" by showing that incumbents may optimally sacrifice productivity gains to preserve markup rents, thereby linking acquisition behavior to the breakdown of productive scaling and the emergence of arrested dynamism at the top of the firm distribution.

2 Data

2.1 Measuring Top Sales Shares

To identify the top 0.01% of firms, I combine firm-level financial data with aggregate economic denominators. I obtain annual sales for U.S.-headquartered firms from *Compustat North America* (1985–2023). The sample period intentionally concludes in 2023 due to the standard corporate reporting lag in Compustat, which otherwise results in missing values for the key variables required to execute the Melitz-Polanec decomposition analysis.

- **Total Sales (Denominator):** I proxy total aggregate sales using the Gross Output of Private Industries sourced from the Bureau of Economic Analysis (BEA).
- **Total Number of Firms (N):** The total number of firms is sourced directly from the Census *Business Dynamics Statistics* (BDS), which natively covers my entire core sample period.
- **Deflation:** All nominal figures are deflated to 2017 dollars using the U.S. GDP deflator (NIPA Table 1.1.9).

2.2 The Multi-Stage M&A Owner-Spine Pipeline

A central challenge in performing an anatomy on concentration is to accurately link transactional data to consolidated financial statements. I overcome this hurdle by constructing an "Owner-Spine" mapping between Refinitiv M&A deals and Compustat firm-level financial data using a tiered hierarchical matching protocol. To handle the complex legal entity structures of modern mega-firms, my pipeline independently matches both the direct asset and the ultimate parent entities for both acquirers and targets (detailed in Appendix A).

To properly account for mechanical growth through consolidation, I attribute the sales of acquired targets to the acquirer. If Refinitiv's reported target sales are missing, I proxy the value using the target's Compustat financials within a tight window around the deal year.

2.3 Micro-Audit and Structural Correction

Beyond the M&A linkage, I perform a systematic audit of every firm that "exits" the top 0.01% tier to construct a micro-audited dictionary. This allows me to distinguish between *Organic Outflows* (firms exiting the top-tier due to shrinkage of firm size) and *Data Artifacts* (e.g., internal entity collapses, CUSIP changes, or spinoffs). By isolating these artifacts into

a separate $D^{Artifact}$ term in the top sales share and Melitz-Polanec decomposition, I ensure that my results are driven by economic reality rather than database reorganizations.

2.4 Sample for Causal Analysis

For the event-study analysis, I merge the audited firm-level panel with firm-level markups and markup-adjusted total factor productivity (TFPQ) following the production function estimation approach of De Loecker, Eeckhout, and Unger (2020). The final dataset is a 10-year rolling panel centered on peak M&A years (from $t-4$ to $t+5$), allowing me to track the post-merger productivity and sales trajectories of "Stayers," "Inflows," and "Outflows" across different macroeconomic eras.

3 Decomposing the Growth of Top Sales Share

To understand whether growth in the top share accounted for by the Top 0.01% of firms (which I refer to as the "top tier") is driven by organic growth or mechanical M&A, I decompose the aggregate growth of this group into its fundamental firm-level flows. Crucially, I define organic (non-mechanical) expansion as growth not arising from the direct transfer of acquired sales. Accordingly, *organic growth* captures both pure internal expansion (e.g., in-house R&D) and any post-acquisition expansion not directly attributable to the mechanical transfer of target sales. In contrast, *mechanical M&A* isolates the zero-sum absorption of a target's pre-existing baseline sales.

At time $t \in \{0,1\}$, each firm i has total sales composed of organic and mechanical (M&A) stock: $y_{it}^{ALL} = y_{it}^{OG} + y_{it}^{MA}$. Let I_t denote the set of all firms at t, and let $P_t \subseteq I_t$ denote the top-p set (the Top 0.01%) at t ranked by y_{it}^{ALL}.

I decompose the dynamic evolution of the top tier by defining fundamental firm-level flow sets. Stayers ($S \equiv P_0 \cap P_1$) represent incumbent firms remaining in the top tier. Displacement ($X \equiv P_0 \setminus P_1$) captures legacy firms falling below the threshold, while Entry ($E \equiv P_1 \setminus P_0$) captures rising firms entering the top tier through organic expansion, and Births (B) capture new entities formed via M&A reorganizations. Crucially, I account for complete firm disappearances from the database through the death set (D). To partition this set, I conduct a baseline matching procedure—combining automated database cross-referencing between Compustat and Refinitiv M&A with manual entity-level verification—yielding an initial partition: $D = D^{MA} \cup D^{\neg MA}$. The mechanical exit set (D^{MA}) contains all target firms standardly linked to an M&A absorption. The unlinked exit set

($D^{\neg MA}$) captures all remaining disappearances.

In raw accounting data, some disappearances in $D^{\neg MA}$ reflect genuine firm death (e.g., bankruptcies or liquidations), while others arise from reporting-induced events such as identifier retirements, restructurings, or missed acquisition linkages. Because even a small number of such misclassifications can materially distort concentration dynamics in the top tier, I conduct a case-by-case audit of Top 0.01% firm exits. This procedure reallocates missed acquisitions to the mechanical M&A flow, retains genuine bankruptcies in the organic outflow, and isolates reporting-induced disappearances into a separate artifact term ($D^{Artifact}$).

To place all flows in common units, I define the baseline average firm size at $t = 0$ as $\bar{y}_0 \equiv Y_0/|I_0|$ and normalize $u_{it}^k \equiv y_{it}^k/\bar{y}_0$ for $k \in \{ALL, OG, MA\}$. Let $\Omega_t \equiv \sum_{i \in I_t} u_{it}^{ALL}$ denote aggregate sales in these units and define the aggregate scale factor $\Xi \equiv \Omega_0/\Omega_1$. The aggregate top share is $S_t = A_t^{ALL}/\Omega_t$, where $A_t^{ALL} = \sum_{i \in P_t} u_{it}^{ALL}$.

Top-share growth thus formally decomposes into four exact components (plotted in Figure 1). All components are constructed so that the decomposition holds exactly in the data.

$$\frac{S_1 - S_0}{S_0} = \underbrace{(\Xi - 1)}_{\text{Dilution from Total Output Growth}} + \underbrace{\Delta_{corr}^{OG}}_{\text{Organic Growth}} + \underbrace{\Delta_{corr}^{MA}}_{\text{Mechanical M\&A}} + \underbrace{\Delta^{Artifact}}_{\text{Data Artifacts (Excluded)}}. \tag{1}$$

3.1 The Corrected Threshold-Centered Decomposition

To extract the margins of organic growth at the top-tier boundary, I use a threshold-centered decomposition. By aligning the organic flows relative to the Pareto entry threshold $q_1^{(ALL)} \equiv \min_{i \in P_1} u_{i1}^{ALL}$, I isolate the organic scaling at the top-tier (plotted in Figure 2):

$$\Delta_{corr}^{OG} = \Xi \cdot \frac{1}{A_0^{ALL}} \Bigg[\underbrace{\sum_{i \in S} (u_{i1}^{OG} - u_{i0}^{OG})}_{\text{Stayers}} + \underbrace{\sum_{i \in E} \left(u_{i1}^{OG} - q_1^{(ALL)}\right)}_{\text{Inflow above Threshold}} + \underbrace{\sum_{i \in X} \left(q_1^{(ALL)} - u_{i0}^{ALL}\right)}_{\text{Outflow (Lost Top Mass)}} - \underbrace{\sum_{i \in D^{OG_true}} u_{i0}^{ALL}}_{\text{Organic Exit}} + \underbrace{\left(|E| - |X|\right) q_1^{(ALL)}}_{\text{Inflow - Outflow across Threshold}} \Bigg]. \tag{2}$$

Concurrently, the mechanical M&A component isolates the reallocation of pre-existing sales through acquisitions. Letting $D^{MA_total} = D^{MA} \cup D^{MA_missed}$ (plotted in Figure 3):

$$\Delta_{corr}^{MA} = \Xi \cdot \frac{1}{A_0^{ALL}} \Bigg[\underbrace{\sum_{i \in S} (u_{i1}^{MA} - u_{i0}^{MA})}_{\text{Stayer M\&A Roll-up}} + \underbrace{\sum_{i \in E} u_{i1}^{MA}}_{\text{Inflow via M\&A}} + \underbrace{\sum_{i \in B} u_{i1}^{ALL}}_{\text{Mechanical Births}} - \underbrace{\sum_{i \in D^{MA_total}} u_{i0}^{ALL}}_{\text{Absorbed Exits}} \Bigg]. \tag{3}$$

3.2 Sources of Top-Share Growth: Organic vs. Mechanical Components

Figure 1 presents the decomposition over rolling ten-year horizons from 1985 to 2023. The decomposition partitions top-share growth into three structural components: (i) *Mechanical M&A* (Δ_{corr}^{MA}), which captures share gains arising from the mechanical absorption of target revenue into the acquirer's consolidated sales; (ii) *Organic Growth* (Δ_{corr}^{OG}), which captures share gains attributable to firms expanding sales after removing the direct revenue transfer from acquired targets; and (iii) a *Dilution* term ($\Xi_t - 1$), which captures the mechanical decline in top-firm shares that would arise if top-tier firms merely kept pace with aggregate market growth. I also isolate a separate *Data Artifact* term ($\Delta^{Artifact}$), which captures discontinuities in Compustat coverage driven by corporate restructurings or reporting-format changes rather than genuine economic events, and exclude it from economic interpretation.

The evolution of actual top-share growth exhibits substantial variation over time. Growth rises steadily across horizons ending in the 1990s, peaks in horizons ending in the early 2000s, and then declines sharply. In horizons ending in the mid-2010s, top-share growth becomes negative, before partially recovering in later periods. These patterns indicate that the expansion of the top tier is highly uneven across decades, with a pronounced slowdown following the early 2000s.

Decomposing this evolution reveals a clear shift in the underlying sources of concentration. In horizons ending in the late 1990s and early 2000s, organic growth is the dominant component, reaching its highest levels in the sample and accounting for the bulk of the increase in top-share. Beginning in the mid-2000s, however, the organic component declines sharply, falling to low or near-zero levels in horizons ending in the mid-2010s, before partially recovering thereafter. By contrast, the mechanical M&A component remains positive throughout the sample and varies more gradually over time. Although it is also elevated in horizons ending in the late 1990s and early 2000s, it does not experience the same collapse as organic growth. As a result, the relative importance of mechanical

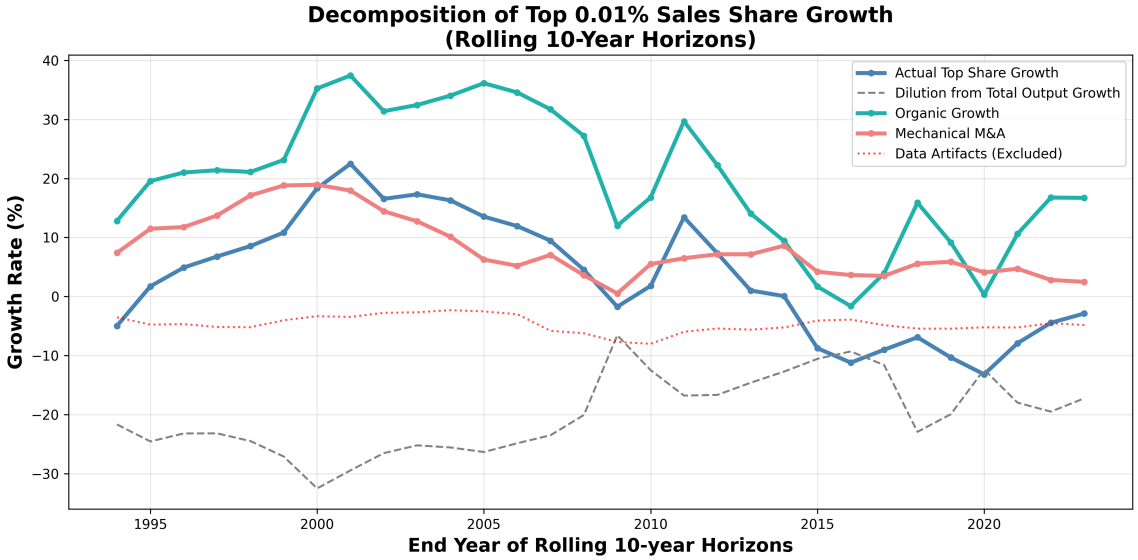


Figure 1: Decomposition of Top 0.01% Sales Share Growth (Rolling 10-Year Horizons). Each series shows the annualized contribution to the growth rate of the Top 0.01% sales share over the ten-year window ending in the year indicated on the horizontal axis. *Actual Top Share Growth* (steelblue) is the left-hand side of the identity. *Organic Growth* (teal) and *Mechanical M&A* (coral) are the two structural forces. *Dilution* (gray, dashed) is the counterfactual growth that would occur if top-tier firms merely kept pace with aggregate output. *Data Artifacts* (red, dotted) are isolated discontinuities excluded from inference. All quantities are expressed as percentage growth rates.

consolidation increases precisely when organic expansion weakens. The dilution term $(\Xi - 1)$ is negative throughout the sample, reflecting aggregate market expansion, and offsets a substantial portion of the positive contributions from organic growth and M&A. The artifact term is small by construction and excluded from economic interpretation.

Taken together, these results indicate a shift in the sources of top-share growth: early increases in concentration are associated with strong organic scaling among incumbent firms, whereas in later periods, concentration dynamics increasingly reflect the persistence of mechanical consolidation in an environment of weakened organic expansion. This distinction is central to the theoretical framework developed below, where firms face a tradeoff between innovation-driven growth and acquisition-based expansion.

Organic Components. Figure 2 decomposes the organic contribution to top sales share growth into its underlying margins. *Stayers* captures the revenue expansion of incumbent firms that remain in the top tier, encompassing both pure internal growth (e.g., via in-house R&D) and the synergistic scaling realized after acquisitions. The remaining four margins capture the dynamic churning of firms at the top-tier boundary. *Inflow above Threshold* and *Inflow - Outflow across Threshold* jointly capture the contribution of firms

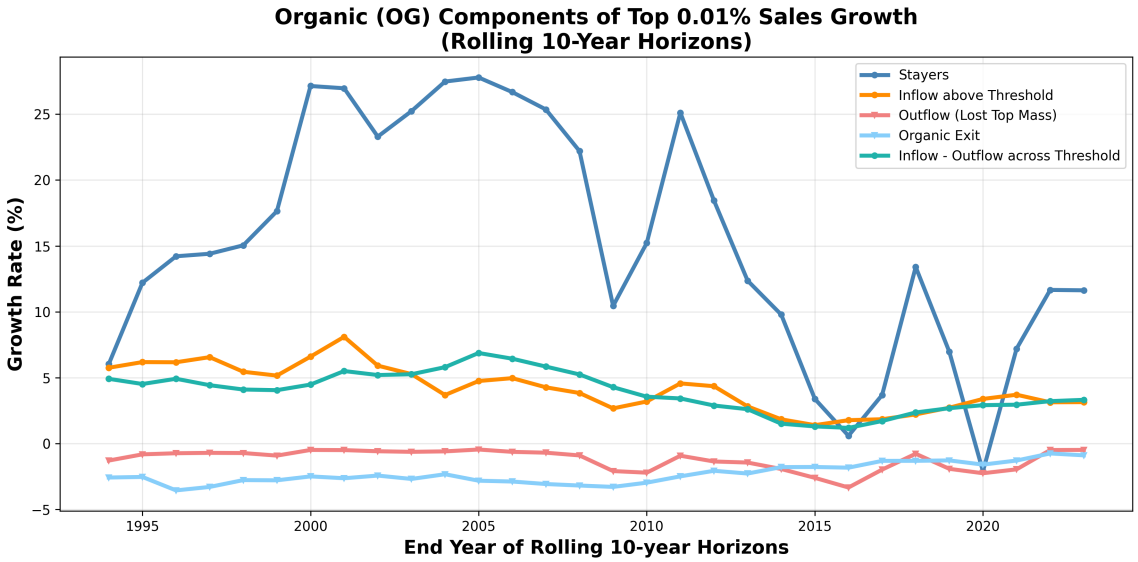


Figure 2: Sub-Components of Organic Growth (Rolling 10-Year Horizons). *Stayers* (steelblue) capture revenue growth of firms continuously in the top tier. *Inflow above Threshold* (orange) is the contribution of firms organically entering the top tier. *Outflow (Lost Top Mass)* (coral) captures firms that organically fall out of the top tier. *Organic Exit* (skyblue) captures firms that exit Compustat entirely through bankruptcy or dissolution. *Inflow minus Outflow across Threshold* (teal) is the net organic churning term. All quantities are annualized percentage contributions over rolling ten-year horizons.

that enter the top tier through organic expansion. Conversely, *Outflow (Lost Top Mass)* and *Organic Exit* capture the contribution of firms that exit the top tier, either through contraction or firm disappearance.

The *Stayers* component accounts for a large share of the variation in aggregate organic growth over time. Its contribution is elevated in horizons ending in the late 1990s and early 2000s, and declines substantially in later periods. By contrast, the upward churning margins (*Inflow above Threshold* and *Inflow - Outflow across Threshold*) remain positive and comparatively stable throughout the sample. On the downside, the relative importance of the two exit-related margins shifts over time: in earlier horizons, *Organic Exit* contributes more to the overall drag, while in later periods *Outflow (Lost Top Mass)* becomes more prominent. Taken together, these patterns indicate that variation in aggregate organic growth is driven primarily by incumbent expansion within the top tier rather than by changes in threshold churning.

Mechanical M&A Components. Figure 3 decomposes the mechanical M&A contribution into four margins: *Stayer Roll-up* (mechanical transfer of target revenue through acquisitions undertaken by firms already in the top tier), *Inflow via M&A* (firms entering the

top tier through acquisition-driven expansion), *Mechanical Births* (new firm identifiers that appear following acquisition-related restructuring), and *Absorbed Exits* (firms that disappear from the top tier after being acquired).

Several patterns emerge. First, both *Stayer Roll-up* and *Inflow via M&A* are positive throughout the sample. *Stayer Roll-up* is quantitatively larger, peaking in horizons ending in the early 2000s and gradually declining thereafter, while remaining positive in later periods. *Inflow via M&A* is smaller in magnitude and relatively stable over time.

Second, *Mechanical Births* is positive and sizable throughout the sample, with particularly large contributions in horizons ending in the late 1990s and early 2000s, followed by a decline to lower but still positive levels. In contrast, *Absorbed Exits* is negative in all periods, with its largest absolute magnitudes also occurring in horizons ending in the early-to-mid 2000s. Because these two margins reflect opposite sides of acquisition-related restructuring, I summarize their net effect by the *Net Mechanical Reorganization* series. This net component is strongly positive in horizons ending in the late 1990s and early 2000s, turns negative in horizons ending in the mid-to-late 2000s, and becomes modestly positive again in more recent periods.

Taken together, the figure shows that mechanical top-share growth reflects incumbent roll-up, acquisition-driven entry, and restructuring-related reorganization, with the latter exhibiting substantial time variation across horizons.

3.3 Aggregate Productivity Dynamics

The accounting decomposition in Section 3.2 establishes *what* drove top-share growth, showing that modern concentration increasingly reflects mechanical consolidation rather than organic scaling. However, accounting flows alone cannot determine whether this concentration is associated with genuine productivity gains or with the suppression of productive challengers. To answer this, I link the accounting identity to the physical productivity dynamics of firms.

Standard revenue-based total factor productivity (TFPR) is insufficient for this task because it conflates physical efficiency with markups (De Loecker, Eeckhout, and Unger, 2020). When a firm expands through M&A and gains market power, TFPR can rise mechanically even if underlying technology does not improve. To isolate physical efficiency, I use a markup-adjusted measure of productivity (TFPQ).

I define aggregate top-tier productivity as $\Phi_t = \sum_{i \in P_t} s_{it} \phi_{it}$, where s_{it} is the firm's sales

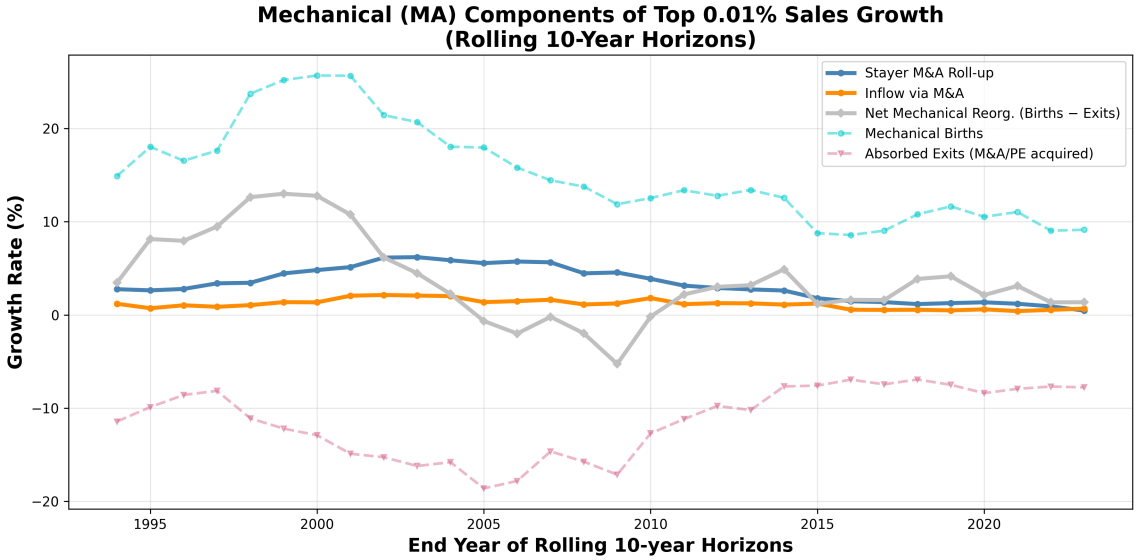


Figure 3: Sub-Components of Mechanical M&A Growth (Rolling 10-Year Horizons). *Stayer M&A Roll-up* (steelblue) is the contribution of acquisitions made by firms already in the top tier. *Inflow via M&A* (orange) captures firms that enter the top tier by acquiring their way to scale. *Mechanical Births* (turquoise) captures newly formed consolidated entities that immediately appear in the top tier. *Absorbed Exits* (palevioletred) captures top-tier firms that are themselves acquired by another party, removing them as independent entities. All quantities are annualized percentage contributions over rolling ten-year horizons.

share within the top tier and ϕ_{it} is its log TFPQ. Following Melitz and Polanec (2015), I decompose changes in aggregate productivity across rolling ten-year horizons into:

$$\Delta\Phi = \underbrace{\sum_{i\in S} \bar{s}_i \Delta\phi_i}_{\text{Within-Firm (Stayers)}} + \underbrace{\sum_{i\in S} \Delta s_i (\bar{\phi}_i - \bar{\Phi}_S)}_{\text{Reallocation (Market Share Shift)}} + \underbrace{s_{E_{all},1}(\Phi_{E_{all},1} - \Phi_{S1}) - s_{X_{all},0}(\Phi_{X_{all},0} - \Phi_{S0})}_{\text{Net Entry (Turnover)}}, \tag{4}$$

where S denotes stayers, E_{all} entrants into the top tier, and X_{all} exits.[1]

Shifting Sources of Top-Tier Productivity Growth. Figure 4 presents the cumulative Melitz–Polanec decomposition, where each point corresponds to a rolling ten-year horizon indexed by its ending year. The figure shows a reorientation in the sources of productivity growth at the top of the distribution.

The *Within-Firm* component rises steadily over horizons ending in the late 1990s and early 2000s, peaking around the 2004 endpoint. This pattern indicates that, in earlier decades, productivity growth among top-tier firms was primarily driven by sustained improve-

[1] Appendix C provides the full derivation and my finer decomposition into organic and M&A-driven margins.

ments within incumbent firms. After this peak, however, the Within-firm contribution declines markedly and eventually stabilizes at a substantially lower level, implying that incumbent firms contribute much less to aggregate TFPQ growth in more recent decades.

In contrast, the *Reallocation* component exhibits a persistent upward trend throughout the sample and becomes the dominant contributor in later periods. This pattern indicates that an increasing share of aggregate productivity growth reflects shifts in market share toward firms that already have relatively high productivity levels, rather than broad-based improvements in firm-level efficiency. Taken together, these trends point to a transition from productivity growth driven by internal technological progress to growth driven by the redistribution of economic activity toward existing high-productivity firms.

Finally, the *Net Entry* component turns negative early in the sample (for horizons ending in the early 1990s), reaches its most negative levels for horizons ending in the early 2000s, and remains persistently negative thereafter. This sustained negative contribution suggests that entry and exit dynamics increasingly act as a drag on aggregate productivity at the top of the distribution.

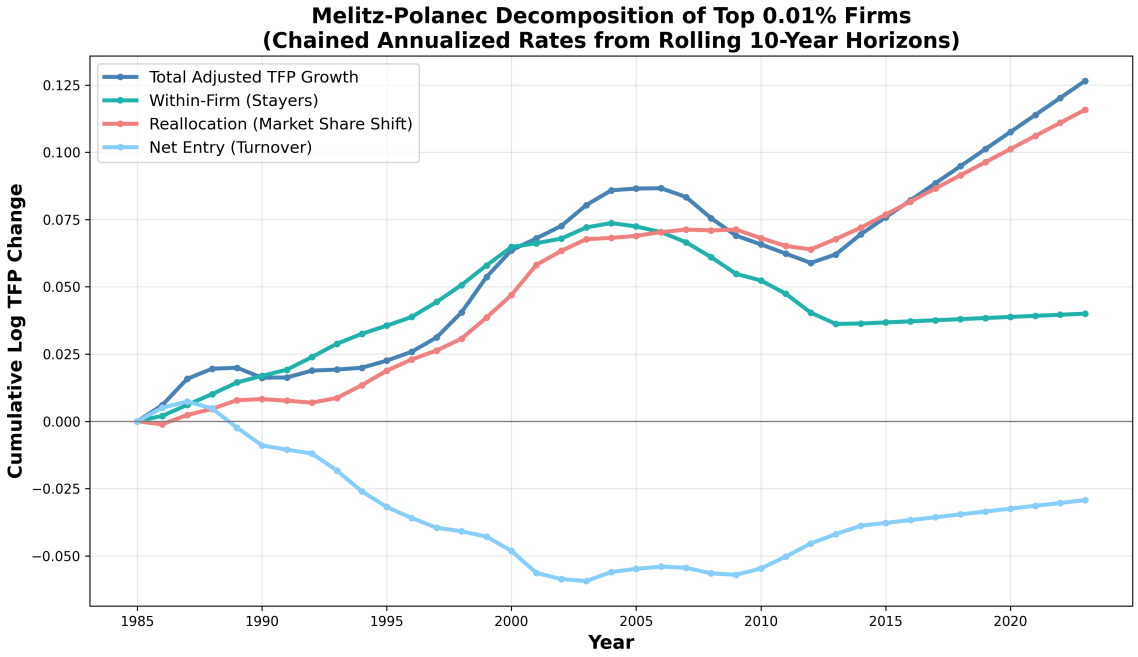


Figure 4: Melitz–Polanec TFP Decomposition of the Top 0.01% (Cumulative, Chained). Series are constructed by chaining annualized rates from rolling ten-year horizon Melitz–Polanec decompositions estimated on the Top 0.01% firm panel with markup-adjusted TFP. *Total Adjusted TFP Growth* (steelblue) is the aggregate change. *Within-Firm* (teal) captures average TFP improvement among stayer firms. *Reallocation* (coral) captures market share shifts toward initially more productive firms. *Net Entry* (skyblue) captures the TFP contribution of entry and exit turnover.

Exit Dynamics and the Decline of Schumpeterian Selection. To understand the source of the negative Net Entry contribution, Figure 5 decomposes this term into its underlying margins. The decomposition shows that the persistent negative Net Entry contribution is not driven by a disappearance of productive inflow. *Organic Inflow* remains positive and rises steadily across horizons, while M&A-driven inflow is modest and non-negative.

Instead, the negative Net Entry term reflects two distinct exit-related patterns. First, the *Organic Outflow* component becomes increasingly negative and then stabilizes at a persistently negative level in recent decades. Because this term is positive when exiting firms are less productive than stayers, its negative sign implies the opposite: firms being pushed out of the top tier are, on average, more productive than those that remain. This stands in contrast to the standard Schumpeterian benchmark, in which exit reflects the cleansing of low-productivity firms.

Second, the *Exit: M&A Targets* component also becomes increasingly negative over the sample, indicating that firms removed from the top tier through acquisition tend to have relatively high productivity. On its own, however, this pattern does not establish whether such acquisitions are efficiency-enhancing or efficiency-reducing, since the Melitz–Polanec decomposition records the disappearance of the target from the independent-firm distribution but does not reveal whether its productive capacity is subsequently integrated into the acquirer.

Taken together, these patterns show that the persistent negative Net Entry contribution reflects the growing importance of exit margins, with the most striking departure from the Schumpeterian benchmark coming from the increasingly negative *Organic Outflow* term. This breakdown is consistent with a framework in which concentration increasingly operates through mechanisms that reallocate and remove firms, rather than through productivity-enhancing selection among competitors.

Acquired Exits and Shrinking Targets. The increasingly negative contribution of acquired exits is consistent with a shift in the nature of M&A targets. As documented in Appendix D, acquisitions have increasingly involved much smaller targets over time.

In the Melitz–Polanec framework, acquisitions mechanically remove the target's standalone contribution from the productivity distribution. When these targets have relatively high productivity, their removal contributes negatively to the acquired-exit component. This pattern is consistent with mechanisms in which acquisitions disproportionately affect productive firms at early stages of growth (Cunningham, Ederer, and Ma, 2021), although the decomposition itself does not reveal whether these acquisitions ultimately enhance or

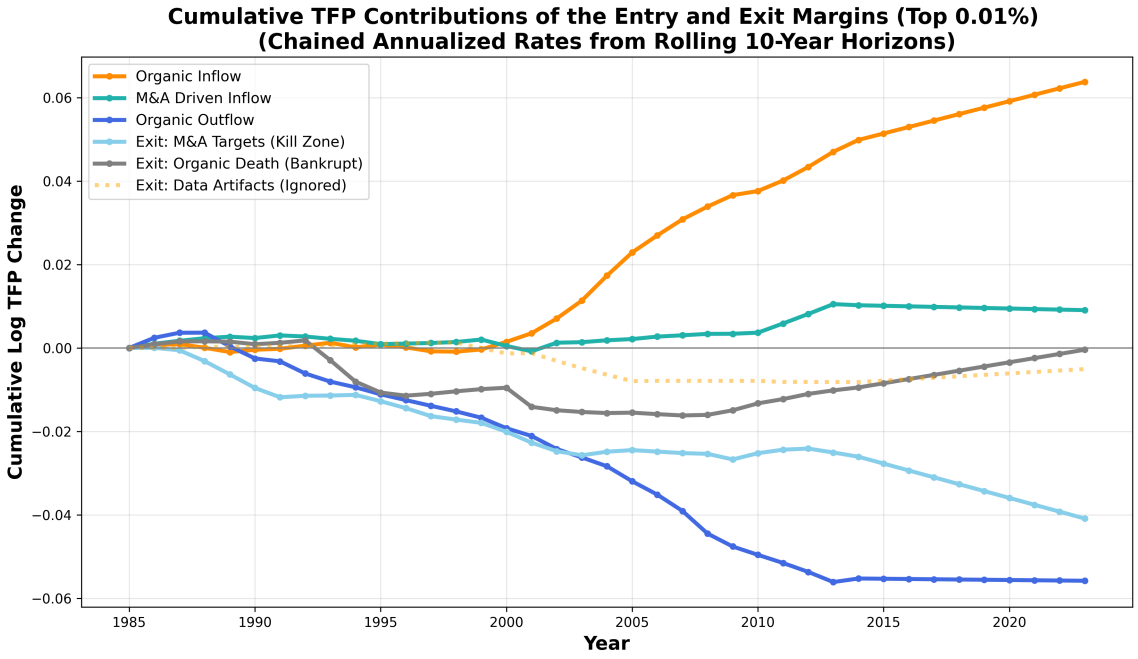


Figure 5: Cumulative TFP Contributions of Entry and Exit Margins (Top 0.01%). Each series is a chained cumulative sum of annualized contributions from rolling ten-year Melitz–Polanec decompositions. The breakdown isolates the structural drag driven by *Organic Outflows* and the absorption of productive *M&A Targets*.

reduce aggregate efficiency.

3.4 The Four Eras of US Concentration

Figure 6 summarizes these dynamics across four eras and reveals a sequence of distinct productivity regimes governing the evolution of top-tier firms. The common scale across panels highlights a broad shift in the composition of productivity growth: earlier eras feature substantial *Within-Firm* contributions, whereas later eras are increasingly characterized by positive *Reallocation* despite weak or negative *Within-Firm* growth.

In the 1985–1994 period, the decomposition exhibits a canonical Schumpeterian pattern. *Within-Firm* productivity growth is strongly positive and is the largest positive component, while both *Organic Inflow* and *Organic Outflow* contribute positively, indicating that entry brings in relatively productive firms and exit removes relatively unproductive ones. *Exit: M&A Targets* is negligible. *Reallocation* is modestly negative, implying that market share is not yet systematically shifting toward the most productive firms. In summary, this period reflects a regime in which aggregate productivity growth is driven by internal improvements and effective selection, rather than by the concentration of market share.

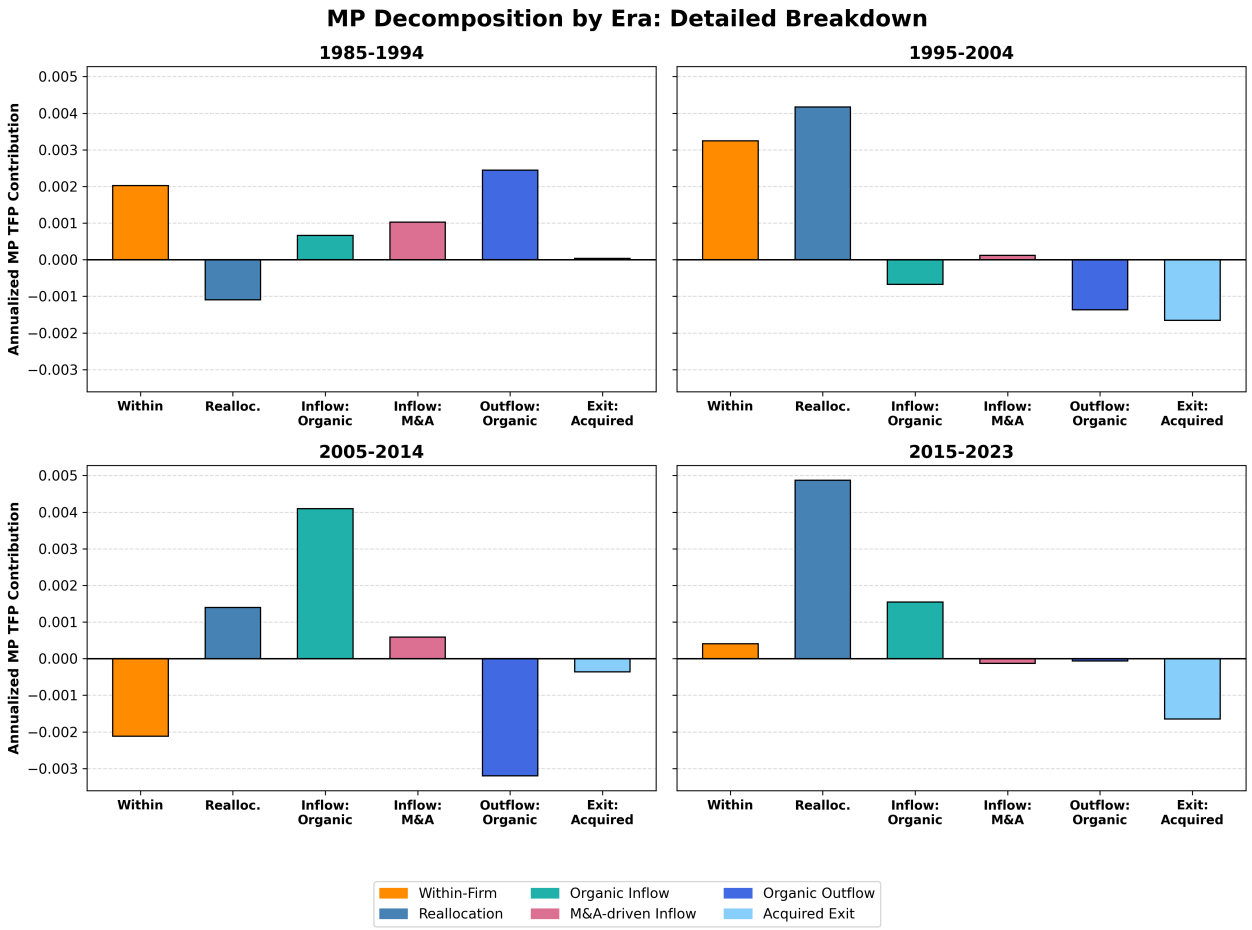


Figure 6: Melitz–Polanec TFP Decomposition of the Top 0.01% by Era (Annualized). This figure illustrates the annualized productivity components across four distinct macroeconomic periods, highlighting the transition from Schumpeterian internal growth in Era 1 (1985–1994) to reallocation and adverse exit dynamics in Era 4 (2015–2023).

In the 1995–2004 period, this pattern transitions into a superstar scaling regime. Both *Within-Firm* and *Reallocation* are large and positive, with *Reallocation* slightly exceeding *Within-Firm*, indicating that market share increasingly flows toward more productive firms as incumbents continue to generate strong internal productivity growth. Although *Exit: M&A Targets* is negative, its magnitude remains limited relative to the large positive contributions of *Within-Firm* and *Reallocation*. *Organic Inflow* is negative, while *Inflow: M&A* is only slightly positive, implying that aggregate productivity growth in this era is driven much more by incumbent performance than by new entry into the top tier. *Organic Outflow* is also negative, suggesting an early weakening of selection dynamics even within this otherwise high-productivity period.

In the 2005–2014 period, the decomposition reveals a clear turning point. The *Within-Firm* component becomes negative, indicating a deterioration in productivity among incumbent firms. At the same time, *Organic Outflow* becomes strongly negative, implying that firms exiting the top tier are, on average, more productive than those that remain. *Organic Inflow* is the largest positive component in this era, but its contribution is insufficient to offset the negative effects from *Within-Firm* and *Organic Outflow*. *Reallocation* remains positive, but now operates in an environment where incumbent productivity is declining. This period marks a breakdown in both internal productivity growth and selection, even as productive entry remains active.

In the 2015–2023 period, the economy enters a regime of arrested dynamism. *Within-Firm* productivity growth is close to zero, indicating a near-complete stagnation of incumbent efficiency gains. *Reallocation* remains strongly positive and is the dominant contributor to aggregate productivity growth, far exceeding the near-zero *Within-Firm* component. *Organic Inflow* remains positive but is much smaller than in the previous period, while *Inflow: M&A* is slightly negative. At the same time, exit margins continue to exert negative effects: *Organic Outflow* is close to zero but remains weakly negative, and *Exit: M&A Targets* is substantially negative, indicating that relatively productive firms are being removed from the top-tier distribution through acquisition. In this regime, aggregate productivity dynamics are driven primarily by reallocation, with neither internal productivity growth nor favorable selection contributing meaningfully.

Taken together, Figure 6 reveals a sequence of four distinct regimes: Schumpeterian growth in the 1980s, superstar scaling in the 1990s and early 2000s, a turning-point period after 2005 in which internal productivity growth becomes negative even as productive inflow remains strong. Finally, in the most recent decade (2015–2023), the data reveal a regime of arrested dynamism, in which reallocation remains strong despite near-zero incumbent

productivity growth and persistently adverse exit margins. These cross-era moments are central to the paper's logic: they show that the changing nature of concentration is not simply a shift in the level of market share, but a shift in the underlying productivity regime that governs how firms rise, remain, and disappear at the top of the distribution.

This distinction is also critical for the sections that follow. The Melitz–Polanec decomposition identifies which margins account for the aggregate productivity transition, but it does not by itself isolate the firm-level mechanism behind the collapse of *Within-Firm* growth or the negative contribution of exit margins. The next section therefore turns to event-study evidence to determine whether modern acquisitions improve firm performance, as in earlier eras, or instead contribute to the anti-Schumpeterian dynamics documented here. These same moments also provide the key empirical targets for the structural model developed below.

3.5 Linking Concentration and Productivity Across Horizons

The accounting decomposition isolates the sources of top-share growth, while the Melitz–Polanec decomposition characterizes the margins of productivity dynamics. I now link these two perspectives directly by examining how accounting flows map into productivity outcomes across identical rolling ten-year horizons.

Figures 7 and 8 plot the relationship between top-share growth components and the corresponding Melitz–Polanec productivity terms. Each point represents a ten-year horizon, indexed by its ending year, thereby aligning the accounting and productivity decompositions within the same temporal units.

A clear set of empirical regularities emerges. Horizons with stronger organic top-share growth exhibit stronger *Within-Firm* productivity growth, while horizons with stronger mechanical M&A-driven growth exhibit weaker, and often negative, *Net Entry* contributions. These relationships are descriptive rather than causal, but they show that the same periods associated with concentration growth are also those associated with shifts in productivity dynamics.

From Correlations to Mechanisms. The combined evidence from the Melitz–Polanec decomposition and the cross-horizon mapping points to a central tension. On the one hand, the rise of *Reallocation* and the persistence of positive inflow indicate that high-productivity firms continue to attract market share and that entry into the top tier remains active. On the other hand, the increasingly negative *Organic Outflow* component shows

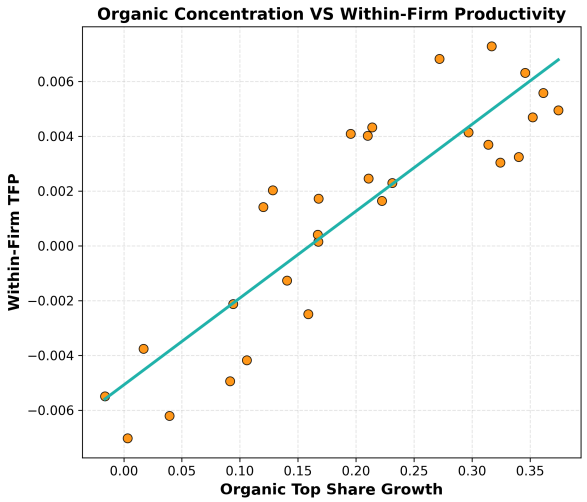


Figure 7: Organic Scaling vs. Within-Firm. Scatter plot linking the accounting and productivity decompositions across horizons. It demonstrates that periods with higher organic top-share growth systematically exhibit stronger internal (Within-Firm) physical efficiency gains.

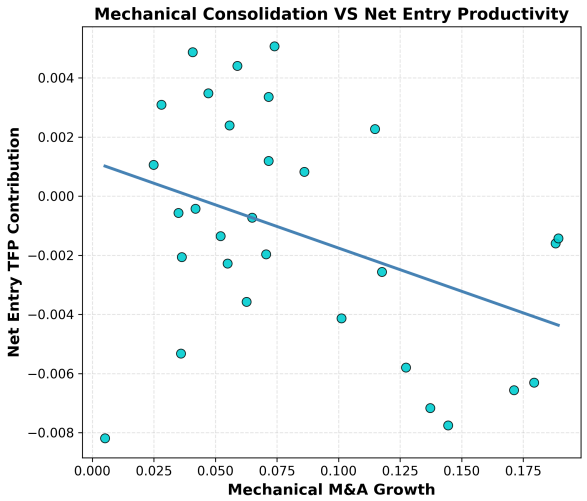


Figure 8: Mechanical M&A vs. Net Entry. Scatter plot linking the accounting and productivity decompositions. Horizons dominated by mechanical M&A consolidation correspond to periods where entry and exit dynamics (Net Entry) heavily drag down aggregate productivity.

that firms displaced from the top tier are not low-productivity laggards, but instead relatively productive firms. This constitutes a clear departure from the Schumpeterian benchmark in which exit reflects the cleansing of inefficient producers.

The scatter plots reinforce this interpretation by showing that the deterioration in Net Entry dynamics is systematically associated with periods of M&A-driven concentration. However, the aggregate decomposition cannot distinguish between competing mechanisms. In particular, it cannot determine whether acquisitions enhance efficiency through integration and synergies, or instead reduce dynamism by removing productive competitors from the independent-firm distribution.

This distinction is central to interpreting the anti-Schumpeterian patterns documented above. If acquisitions improve efficiency, the negative Net Entry contribution may simply reflect measurement, as targets are absorbed into more productive firms. If instead acquisitions suppress competition, then the observed patterns reflect a genuine decline in the process of creative destruction.

To disentangle between these mechanisms, I move beyond aggregate accounting and exploit firm-level variation. In the next section, I estimate a doubly robust event-study design that traces the dynamic effects of M&A on firm-level productivity and growth.

This approach allows me to directly test whether acquisitions generate efficiency gains for incumbents or instead lead to the erosion of productivity consistent with the anti-Schumpeterian dynamics identified in the aggregate data.

3.6 Event-Study Evidence Across Eras

The aggregate Melitz–Polanec decomposition documents a shift from within-firm productivity growth toward reallocation and adverse selection dynamics. I now investigate the firm-level mechanism underlying this transition using a doubly robust difference-in-differences estimator developed by Callaway and Sant'Anna (2021).

To isolate the causal effect of large-scale consolidation, I must carefully define the treatment. I define an "M&A event" as the first year in which a firm's total number of completed acquisitions reaches or exceeds the 75th percentile of annual deal frequency (calculated across the pooled distribution of all active top-tier firm-years in my sample). Crucially, to ensure a clean pre-treatment window and avoid capturing the midpoint of rolling acquisition sprees, I strictly require that the firm engaged in zero M&A transactions in the two preceding years ($t-1$ and $t-2$). Furthermore, to isolate the behavior of entrenched incumbents, I focus on firms that remain continuously in the top tier ("Stayers"). As shown in Appendix G, including entrants introduces substantial mechanical compositional noise, obscuring the true treatment effects on incumbent top-tier firms. A detailed summary of the panel characteristics is provided in Appendix G.1; crucially, the data confirms strong pre-treatment baseline balance. Prior to their acquisitions, treated Stayers exhibited nearly identical organic growth trajectories and higher profitability relative to the untreated control pool, fundamentally ruling out pre-existing adverse selection.

Formally, I estimate event-time average treatment effects (ATT_k) using the doubly robust CS-DiD procedure, which conceptually maps to the following dynamic event-study specification:

$$Y_{i,t} = \alpha_i + \tau_t + \sum_{k=-4, k\neq -1}^{5} \beta_k \cdot D_{i,t}^k + \Gamma' X_{i,t-1} + \varepsilon_{i,t} \tag{5}$$

where $Y_{i,t}$ represents the outcome of interest (log markup-adjusted TFPQ, log organic sales, log total sales, or log markups) for firm i in year t. The term $D_{i,t}^k$ is an indicator equal to one if firm i is k years relative to its first peak M&A treatment year. The terms α_i and τ_t represent firm and year fixed effects, respectively. To ensure the parallel trends assumption holds conditionally, the doubly robust estimation explicitly balances a vector of pre-treatment covariates, $X_{i,t-1}$, which includes lagged log total assets, lagged cash-to-assets ratio, lagged return on assets (ROA), and lagged R&D intensity. Standard errors are

clustered at the firm level.

Table 1 reports average treatment effects across horizons, estimating the impact of an M&A event on physical productivity (TFPQ), organic sales growth, total sales growth, and markups for incumbent Stayers. Figure 9 and Figure 10 present the dynamic event studies for these key outcomes across representative eras.

Table 1: Stayers Average Post-Treatment Effects ($t \in [2, 5]$): Era-by-Era Shift

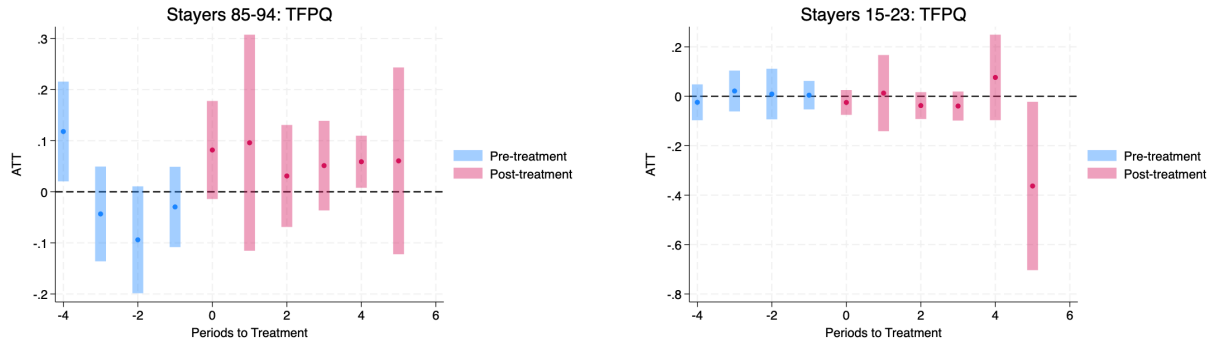
Outcome	1985–1994	1995–2004	2005–2014	2015–2023
TFPQ	0.050	-0.021	-0.033	-0.091**
	(0.038)	(0.087)	(0.060)	(0.042)
Log Markup	-0.046**	0.042	-0.015	-0.001
	(0.020)	(0.044)	(0.073)	(0.078)
Organic Sales Growth	0.060	0.024	0.072	-0.061
	(0.050)	(0.099)	(0.064)	(0.042)
Total Sales Growth	0.033	0.026	0.082	-0.041**
	(0.043)	(0.088)	(0.069)	(0.020)

Notes: This table presents the average of the CS-DiD event study coefficients for periods $t \in [2, 5]$ post-treatment for Stayers, estimated independently within each macroeconomic era. The transition from positive long-run scaling in the 1990s to significantly negative physical productivity (TFPQ) and sales growth in the modern era isolates the structural regime shift in top-tier M&A. Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

The Reversal of M&A Efficacy. The event studies reveal a stark reversal in the physical efficacy of M&A across eras. In the 1985–1994 period, acquisitions by top-tier incumbents were highly synergistic. As shown in Figure 9a, post-merger TFPQ rises over the medium run, generating a $+5.0\%$ average increase in physical productivity, consistent with gradual post-merger integration and operational restructuring. In this period, external M&A served to push the firm's internal technological frontier outward.

By the modern 2015–2023 era, this dynamic has completely inverted (Figure 9b). Large-scale acquisitions now inflict a massive -9.1% structural penalty on physical productivity, with TFPQ declining sharply at longer horizons.

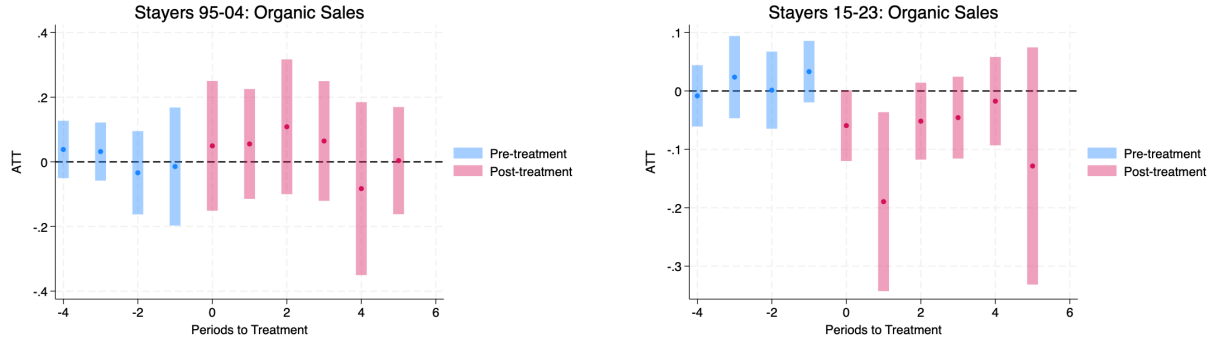
Importantly, pre-treatment dynamics are remarkably flat in both periods, supporting the identifying assumption. The absence of systematic pre-trends suggests that the post-treatment divergence reflects causal effects rather than selection on pre-existing trends. To further ensure this modern collapse is not driven by unobservable adverse selection



(a) Era 1 (1985–1994): Productivity Gains (b) Era 4 (2015–2023): Productivity Decline

Figure 9: Dynamic Effects of M&A on Physical Productivity (TFPQ). Event-study estimates for Top 0.01% Stayers. Early-period acquisitions are consistent with delayed productivity improvements, while modern acquisitions are associated with severe declines in physical efficiency at longer horizons. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D^k_{i,t} + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).

(i.e., the possibility that only declining firms execute M&A today), Appendix G.3 replicates this design on a strictly Propensity Score Matched sample. The physical rot penalty remains highly significant.



(a) Era 2 (1995–2004): Organic Expansion (b) Era 4 (2015–2023): Organic Contraction

Figure 10: Dynamic Effects of M&A on Organic Growth. Event-study estimates for Top 0.01% Stayers. The internal scaling engine breaks down in the modern era following major acquisitions. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D^k_{i,t} + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).

Markups and the Motive for Defensive M&A. Figure 10 corroborates the productivity penalty: modern acquisitions heavily depress the firm's internal organic scaling engine (-6.1%), whereas they historically boosted it.

This raises a central economic question: If modern M&A destroys physical productivity

and chokes organic scaling, why do apex firms continue to aggressively pursue it? The treatment effects on Log Markup provide a crucial clue. As shown in Table 1, across the modern eras, the post-treatment change in markups is statistically indistinguishable from zero (-0.001 in Era 4). This indicates that modern acquisitions are not primarily vehicles for acquiring *new* pricing power. Rather, as I formalize in my structural model below, they are deployed defensively to *protect* the titan's pre-existing, highly elevated baseline market power from disruptive challengers.

From Reduced-Form Evidence to Economic Mechanisms. Taken together, the evidence points to a fundamental change in the role of M&A. In earlier periods, acquisitions were consistent with efficiency-enhancing integration and sustained internal growth. In the modern era, acquisitions by top incumbents are associated with economically meaningful declines in physical efficiency and weaker internal growth, executed primarily to shield legacy moats.

The aggregate and firm-level evidence jointly suggest that the nature of concentration has shifted from a process driven by technological scaling to one increasingly characterized by defensive behavior and weakened dynamism. In the next section, I formalize this mechanism in a continuous-time structural framework, demonstrating how increases in internal organizational costs and the rising stakes associated with high markups mathematically force firms to transition from synergistic integration toward defensive acquisition behavior.

4 A Structural Model of Concentration and Productivity

To provide a structural framework for interpreting the empirical transition from productivity-driven concentration to the regime of arrested dynamism documented above, I develop a continuous-time model of firm dynamics. The model characterizes the incumbent firm's joint choice of internal expansion and acquisition behavior. By linking firm-level jump-diffusion processes to the evolution of the macroeconomic firm size distribution, I construct a sequential mapping from the model's structural parameters to the empirical moments from the Melitz–Polanec decomposition and the event-study estimates.

4.1 Economic Environment and Firm Dynamics

Following standard quality-ladder models (Grossman and Helpman, 1991; Aghion and Howitt, 1990), consumers have Constant Elasticity of Substitution (CES) preferences over

a continuum of intermediate goods. Each incumbent firm in the top tier holds a monopoly in its product line indexed by j, pricing at a markup $\mu_{j,t}$ over marginal cost. As derived in Appendix E, this yields instantaneous operating profits that are proportional to the firm's physical productivity ($\phi_{j,t}$):

$$\pi_{j,t} = \pi(\mu_{j,t})\phi_{j,t},$$

where $\pi(\mu_{j,t}) = 1 - 1/\mu_{j,t}$ is increasing in the markup. Incumbents operate two margins of expansion: an internal margin (in-house productivity growth) and an external margin (M&A).

Internal Margin: The firm invests to improve its own fundamental productivity $\phi_{j,t}$. Achieving a gross internal productivity growth rate g_c requires paying a convex investment cost $C(g_c) = \frac{1}{2} \cdot c_k \cdot g_c^2 \phi_{j,t}$, where c_k represents the inferred cost of internal expansion, which can be interpreted as reflecting organizational complexity or adjustment frictions. Simultaneously, the firm is subject to a continuous obsolescence rate δ_t, capturing the gradual erosion of its relative technological position due to competitive pressures.

External Margin: The firm faces an observed M&A execution rate $\lambda_{total,t}$, capturing the frequency of acquisition activity. Upon acquisition, the firm experiences a discrete productivity jump, denoted by J_t, representing the log change in productivity associated with the acquisition, whose magnitude depends on the firm's integration decision. In addition, acquisitions affect the firm's pricing environment through changes in product scope and competitive exposure. I capture this channel parsimoniously through adjustments in the markup. The firm therefore makes a discrete choice $I \in \{0, 1\}$ upon acquisition:

- **Integration ($I = 1$):** The firm incorporates the target's technology, yielding an expected productivity gain of γ. In doing so, the firm expands its effective product space and internalizes previously external varieties. This expansion of the firm's product scope introduces substitution across the firm's own product lines. Because the firm internalizes this cannibalization when setting prices, its effective residual demand becomes more elastic, leading to a reduction in the optimal markup to a post-merger level $\mu_{int} < \mu_t$. Thus, integration trades off higher physical efficiency against a dilution of market power.

- **Defensive Shelving ($I = 0$):** The firm acquires the target without integrating its technology, effectively removing it as an independent competitor while maintaining separation in production. This preserves the incumbent's existing markup μ_t, as the competitive structure faced by its core product line is unchanged. However, foregoing integration implies that the organizational capital of the target is not ef-

ficiently deployed within the firm. I capture this misallocation through a negative productivity effect ("rot") of magnitude $-\tau$, reflecting internal disruption, duplicated structures, and foregone synergies.

Law of Motion of Firm Productivity. Firm productivity therefore evolves through the combination of continuous internal investment and discrete acquisition shocks. In log terms, the law of motion of a firm j's productivity is

$$d \ln \phi_{j,t} = (g_c - \delta_t)\, dt + J_t\, dN_t, \tag{6}$$

where g_c is the firm's continuous gross internal productivity growth rate, δ_t is the obsolescence rate, and N_t is a Poisson process with intensity $\lambda_{total,t}$ governing acquisition events. At the arrival of an acquisition event (i.e., when $dN_t = 1$), productivity undergoes a discrete multiplicative jump, $\phi_{j,t} \to \phi_{j,t} e^{J_t}$, where the realized jump J_t depends on the firm's integration choice:

$$J_t = \begin{cases} \gamma & \text{if } I = 1 \text{ (synergistic integration)}, \\ -\tau & \text{if } I = 0 \text{ (defensive shelving)}. \end{cases}$$

Here γ and $-\tau$ denote log productivity changes.

Taking expectations over the Poisson process implies

$$\mathbb{E}[d \ln \phi_{j,t}] = (g_c - \delta_t)\, dt + \lambda_{total,t} \mathbb{E}[J_t]\, dt. \tag{7}$$

Given these primitives, the firm's dynamic optimization problem can be written as a continuous-time control problem. The firm chooses its continuous gross internal productivity growth rate g_c and its M&A policy $I \in \{0, 1\}$ to maximize its value function $V(\mu_t, \phi_t)$, subject to the Hamilton–Jacobi–Bellman (HJB) equation:

$$rV(\mu_t, \phi_t) = \max_{g_c \geq 0, I \in \{0,1\}} \left\{ \pi(\mu_t)\phi_t - \frac{1}{2} c_k g_c^2 \phi_t + (g_c - \delta_t)\phi_t \frac{\partial V}{\partial \phi_t} + \lambda_{total,t} \mathbb{E}\Big[V(\mu_{post}, \phi_t e^{J_t}) - V(\mu_t, \phi_t) \Big] \right\}, \tag{8}$$

where the post-acquisition markup μ_{post} depends on the firm's discrete choice:

$$\mu_{post} = \begin{cases} \mu_{int} & \text{if } I = 1 \text{ (integration)}, \\ \mu_t & \text{if } I = 0 \text{ (defensive shelving)}. \end{cases}$$

Given the proportional structure of profits and the law of motion for productivity, we conjecture a value function linear in productivity of the form $V(\mu_t, \phi_t) = v(\mu_t)\phi_t$. Substituting this form into the HJB verifies that it is consistent with equilibrium.

We summarize the acquisition-induced change in continuation value by the expected proportional change in firm value, Ω_t, evaluated at the post-acquisition state $(\mu_{post}, \phi_t e^{J_t})$:

$$\Omega_t \equiv \mathbb{E}\left[\frac{V(\mu_{post}, \phi_t e^{J_t}) - V(\mu_t, \phi_t)}{V(\mu_t, \phi_t)}\right] = \mathbb{E}\left[\frac{v(\mu_{post})}{v(\mu_t)} e^{J_t} - 1\right].$$

This formulation separates the effect of acquisitions on productivity from their effect on the markup state. Substituting into the HJB yields the normalized problem:

$$rv(\mu_t) = \max_{g_c \geq 0, I \in \{0,1\}} \left\{\pi(\mu_t) - \frac{1}{2}c_k g_c^2 + v(\mu_t)(g_c - \delta_t) + \lambda_{total,t}\Omega_t v(\mu_t)\right\}. \tag{9}$$

To take this object to the data, we operationalize the markup component of Ω_t through a tractable approximation to the continuation-value ratio $v(\mu_{post})/v(\mu_t)$, which depends on the firm's integration decision.

Under **integration** ($I = 1$), acquisitions increase physical productivity but reduce the incumbent's markup due to cannibalization across product lines. We capture this markup dilution through a reduced-form proportional haircut on continuation value, denoted by $\chi \in (0, 1)$, so that

$$\frac{v(\mu_{post})}{v(\mu_t)} \approx 1 - \chi.$$

The exact implied change in firm value is therefore

$$\Omega_t(I = 1) = (1 - \chi)e^{\mathbb{E}[J_t | I=1]} - 1.$$

Under the model, $\mathbb{E}[J_t \mid I = 1] = \gamma$, so this can be written as

$$\Omega_t(I = 1) = (1 - \chi)e^{\gamma} - 1.$$

In implementation, we use the first-order approximation

$$\Omega_t(I = 1) \approx (1 - \chi)(1 + \gamma) - 1,$$

which is accurate for the empirically relevant range of acquisition effects.

Under **defensive shelving** ($I = 0$), the incumbent preserves its existing markup, so that $v(\mu_{post})/v(\mu_t) \approx 1$. In this case,

$$\Omega_t(I = 0) = e^{\mathbb{E}[J_t|I=0]} - 1.$$

Under the model, $\mathbb{E}[J_t \mid I = 0] = -\tau$, so

$$\Omega_t(I = 0) = e^{-\tau} - 1 \approx -\tau.$$

Thus, Ω_t provides a unified representation of acquisition effects that combines productivity changes with the endogenous response of markups, conditional on the firm's integration decision.

Optimal Policies. The firm's problem yields two optimality conditions corresponding to the continuous control g_c and the discrete acquisition choice I.

The first-order condition for internal investment g_c is obtained from equation (9):

$$-c_k g_c + v(\mu_t) = 0 \quad \Rightarrow \quad g_c^* = \frac{v(\mu_t)}{c_k}. \tag{10}$$

This condition equates the marginal cost of internal expansion to its marginal value in terms of increased firm valuation.

The discrete policy choice compares continuation values under integration and defensive shelving. The firm chooses defensive shelving ($I^* = 0$) whenever

$$v(\mu_t)e^{-\tau} > v(\mu_{int})e^{\gamma} \quad \Longleftrightarrow \quad \frac{v(\mu_t)}{v(\mu_{int})} > e^{\gamma+\tau}. \tag{11}$$

This condition highlights the trade-off between preserving incumbent market power and realizing productivity gains from integration.

4.2 The Dynamics of Firm Productivity Distribution

To connect the firm-level dynamics to aggregate outcomes, we consider the evolution of the cross-sectional distribution of firm productivity. Let $p(x, t)$ denote the density of firms over (log) productivity $x = \ln \phi$. Given the jump-diffusion process defined above, the distribution evolves according to the Kolmogorov Forward Equation (KFE),

or Fokker–Planck equation:

$$\frac{\partial p(x,t)}{\partial t} = -\frac{\partial}{\partial x}\Big[(g_c - \delta_t)p(x,t)\Big] - \lambda_{total,t}p(x,t) + \lambda_{total,t}\int p(x-J,t)\,dF_J(J), \tag{12}$$

where F_J denotes the cumulative distribution function of the log productivity jump J_t. The first term captures the continuous drift in productivity, while the latter two terms capture the exit and re-entry of mass induced by acquisition-driven jumps.

We focus on stationary environments corresponding to each empirical era. In this case, the upper tail of the firm size distribution is well approximated by a Pareto distribution with tail exponent ζ_t. A key property of the Pareto distribution is that its shape is invariant to proportional scaling, which implies a tight link between the drift of incumbent firms and the rate at which firms exit the top tail.

In particular, the stationarity condition implies a mapping between the effective downward drift induced by obsolescence and the mass of firms exiting the top tier. As shown in Appendix E, this yields:

$$\delta_t = \frac{m_{org,t}}{\zeta_t}, \tag{13}$$

where $m_{org,t}$ denotes the observed mass of firms that organically exit the Top 0.01% in a given period.

4.3 Model Estimation via Sequential Structural Mapping

A key advantage of the framework is that the model's structural parameters can be disciplined using a sequence of moment conditions that map directly to observable objects in the data. Leveraging the decomposition results from Sections 3.4 and 3.6, we construct a set of empirical counterparts to the model's sufficient statistics, allowing each parameter to be identified using a small subset of moments rather than through joint numerical estimation.

Our identification strategy proceeds sequentially. We first recover the obsolescence rate δ_t from the stationary-tail condition implied by the Kolmogorov Forward Equation. Given δ_t, we use the observed net productivity growth of incumbent firms to infer gross internal productivity growth rate g_c, accounting for both obsolescence and acquisition-driven jumps. We then evaluate the normalized HJB equation to recover firm value $v(\mu_t)$ and the implied cost of internal expansion c_k. Finally, we discipline the acquisition component by combining event-study estimates of $\mathbb{E}[J_t]$ with the selection patterns implied by the Melitz–Polanec decomposition, which together inform both the magnitude and interpre-

tation of acquisition-induced value changes.

We begin by pinning down the obsolescence rate δ_t using the stationary-tail condition derived from the Kolmogorov Forward Equation. Given δ_t, we recover the firm's gross internal productivity growth rate (g_c) using the decomposition of expected productivity growth. The expected net organic growth of incumbent firms (g_t^W), measured in the data as the weighted average log productivity growth of Melitz–Polanec stayers, can be written as:

$$g_t^W = \int_{Z_t}^{\infty} \Big(g_c - \delta_t + \lambda_{total,t} \mathbb{E}[J_t] \Big) \frac{f(\phi)}{1 - F(Z_t)} d\phi = g_c - \delta_t + \lambda_{total,t} \mathbb{E}[J_t], \tag{14}$$

where the equality follows because the drift and jump terms are common across firms in the top tail. Rearranging yields a direct mapping from observed moments to g_c, conditional on δ_t, the M&A execution rate $\lambda_{total,t}$, and the acquisition effect $\mathbb{E}[J_t]$.

Given g_c and δ_t as disciplined by the data, we recover the firm's steady-state normalized valuation by evaluating the normalized HJB equation along the equilibrium path. The optimal internal investment condition implies that the marginal cost equals the marginal benefit ($c_k \cdot g_c = v(\mu_t)$), allowing us to substitute out the cost parameter. Evaluating the HJB along the equilibrium path yields:

$$v(\mu_t) = \frac{\pi(\mu_t)}{r + \delta_t - \frac{1}{2} \cdot g_c - \lambda_{total,t} \Omega_t}, \tag{15}$$

where Ω_t denotes the expected proportional change in enterprise value induced by acquisition, incorporating both productivity and markup effects. This expression links firm valuation to observable moments and disciplined assumptions on acquisition outcomes. In the empirical implementation, the markup-dilution parameter χ is calibrated from the CS-DiD markup regression, using the estimated post-acquisition decline in markups under integration as the empirical counterpart to the continuation-value haircut associated with markup dilution.

Disciplining M&A Effects Across Eras.

To discipline the M&A component, we combine the event-study estimates with evidence from the Melitz–Polanec decomposition presented in Section 3.4, in particular the behavior of selection as captured by the organic outflow term.

In Era 1, the organic outflow contribution is positive, indicating that firms exiting the top tier are systematically less productive than those that remain. This pattern is consis-

tent with efficient selection at the boundary and a Schumpeterian environment in which high-productivity firms expand and replace less efficient incumbents. Although the event-study estimate of the M&A effect is imprecise, it is economically large and aligns with this pattern of efficient restructuring. We therefore retain the positive estimate of $\mathbb{E}[J_t]$ for this period as a disciplined calibration consistent with both micro-level evidence and aggregate selection dynamics.

In Eras 2 and 3, the organic outflow contribution turns negative, implying that firms exiting the top tier are, on average, more productive than those that remain. This reversal indicates a breakdown in efficient selection and suggests increasing distortions in the allocation of activity within the top tail. At the same time, the event-study estimates of M&A effects are statistically imprecise, and the contribution of acquisitions to aggregate productivity is quantitatively small. Taken together, these patterns provide limited evidence for systematic productivity-enhancing or productivity-destroying acquisition effects at the aggregate level. We therefore adopt a conservative calibration of $\mathbb{E}[J_t] = 0$ for these periods, treating the M&A margin as approximately neutral in expectation.

By Era 4, the organic outflow contribution is approximately zero, consistent with a regime in which selection at the top of the distribution is largely muted. In this environment, the event-study estimates indicate a statistically significant and economically large negative effect of acquisitions on productivity. We therefore directly use the estimated value of $\mathbb{E}[J_t] < 0$ for this period, capturing the role of acquisition-driven productivity losses in sustaining incumbent dominance.

This joint use of micro-level treatment effects and macro-level selection patterns provides a disciplined and internally consistent approach to calibrating the M&A component of the model across regimes.

Empirical Moments and Structural Estimates.

The empirical moments used in the calibration are reported in Table 2. Table 3 reports the macro transition parameters—specifically the Pareto exponent and the implied obsolescence rate—used in the stationary-tail mapping. The resulting structural parameters obtained from the sequential mapping are presented in Table 4.

Table 2: Moments for Empirical Calibration by Era

Moment	1985–1994	1995–2004	2005–2014	2015–2023
Micro-Level Firm Dynamics				
Baseline Markup (μ_t)	1.588	1.693	1.650	1.830
Net Organic Growth (g_t^W)	0.0258	0.0277	0.0087	0.0178
Empirical M&A Execution Rate ($\lambda_{total,t}$)	0.0037	0.0102	0.0023	0.0091
Causal M&A Shocks				
CS-DiD TFPQ Shock ($\mathbb{E}[J_t]$)	0.050	0.000	0.000	-0.091

Notes: This table reports the empirical moments used to discipline the firm dynamics model. Markups are calculated following De Loecker et al. (2020), restricted to Melitz–Polanec "Stayers". The net organic growth measure (g_t^W) is constructed as the weighted average log productivity growth of Melitz–Polanec "Stayers" and serves as the empirical counterpart to incumbent productivity growth in the model. The empirical M&A execution rate ($\lambda_{total,t}$) is computed from Stayer firm-years. The CS-DiD TFPQ shock corresponds to the doubly robust average treatment effect on the treated. For Eras 2 and 3, the estimated M&A effects are imprecise and their implied aggregate contributions are quantitatively small. In addition, the Melitz–Polanec decomposition indicates negative organic outflow contributions in these periods, suggesting a breakdown in efficient selection at the top of the distribution. We therefore adopt a conservative calibration setting $\mathbb{E}[J_t] = 0$ for these eras, treating the M&A margin as approximately neutral in expectation.

Table 3: Macro Structural Estimates: Shift in Upper Tail Dynamics

Parameter	1985–1994	1995–2004	2005–2014	2015–2023
Pareto Exponent (ζ_t)	1.329	1.196	1.120	1.058
Natural Obsolescence Rate (δ_t)	0.0387	0.0293	0.0373	0.0138

Notes: This table reports the macroeconomic transition parameters used in the structural mapping across the four eras. The Pareto exponent (ζ_t) characterizes the stationary shape of the Top 0.01% tail distribution. The obsolescence rate (δ_t) captures the continuous downward drift that moves incumbents toward the boundary of the top tail. Under the stationary-tail approximation, the boundary-flow condition implies the mapping $\delta_t = m_{org,t}/\zeta_t$, where $m_{org,t}$ is the empirical mass of firms that organically exit the top tier.

Table 4: Sequential Structural Identification of Firm Dynamics

Parameter	1985–1994	1995–2004	2005–2014	2015–2023
Gross Internal Productivity Growth Rate (g_c)	0.064	0.057	0.046	0.032
Normalized Value Function ($v(\mu_t)$)	6.55	8.06	6.13	9.37
Relative Internal Bureaucratic Friction ($c_k / c_{k,Era1}$)	1.00	1.39	1.31	2.84
Defensive Shelving Premium Threshold		**15.5%**		

Notes: This table reports the analytically recovered structural parameters of the model, obtained through the sequential mapping from empirical moments. Baseline Markup (μ_t) and Net Organic Growth (g_t^W) are estimated directly from the micro data. Gross internal productivity growth rate (g_c) is recovered using the accounting identity $g_c = g_t^W + \delta_t - \lambda_{total,t}\mathbb{E}[J_t]$, combining the observed net growth rate with the estimated obsolescence rate (δ_t) and the calibrated acquisition effect. The normalized value function ($v(\mu_t)$) is computed from the HJB equation using the implied enterprise value jump (Ω_t), which incorporates both the physical productivity effect of acquisitions and the associated change in markup. In Era 1, Ω_t reflects both the estimated productivity gains ($\gamma = 0.050$) and an inferred markup dilution ($\chi = 0.046$). For Eras 2 and 3, the M&A component is set to zero as a conservative benchmark, consistent with imprecise micro estimates and evidence of distorted selection in the Melitz–Polanec decomposition. In Era 4, the value jump reflects the estimated negative productivity effect under defensive acquisitions. The discount rate is set to $r = 0.05$. The inferred cost of internal expansion (c_k) is recovered from the optimal investment condition and normalized relative to its Era 1 value.

Identifying Regimes of Firm Growth

The structural estimates in Table 4, interpreted jointly with the Melitz–Polanec decomposition, identify distinct regimes of top-tier firm growth driven by the evolving balance between gross internal productivity growth rate (g_c), selection dynamics, and acquisition activity.

Era 1 (1985–1994): Schumpeterian Growth.
Era 1 exhibits the hallmark features of a Schumpeterian regime. The Melitz–Polanec decomposition shows strong within-firm productivity growth alongside a positive contribution from organic outflows, indicating that firms exiting the top tier are systematically less productive than those that remain. This pattern reflects efficient selection and active reallocation at the boundary of the distribution. Consistent with this environment, the model recovers a high level of gross internal productivity growth ($g_c = 0.064$). While the event-study estimate of M&A effects is imprecise, it is economically large and aligns with a regime of efficiency-enhancing restructuring. Taken together, the evidence suggests that concentration during this period was driven by productivity growth and creative destruction.

Era 2 (1995–2004): Superstar Scaling.
Era 2 is characterized by exceptionally strong within-firm productivity growth alongside high levels of reallocation, but a sharp deterioration in selection dynamics. The organic outflow term turns negative, implying that firms exiting the top tier are, on average, more

productive than those that remain. This pattern indicates that while resources are actively reallocated across firms, the direction of reallocation is distorted. The model reflects this configuration: gross internal productivity growth remains high ($g_c = 0.057$), sustaining rapid expansion among incumbents, while the contribution of acquisitions is muted. Given the lack of statistically precise M&A effects and their limited aggregate contribution, we conservatively treat the acquisition margin as neutral ($\mathbb{E}[J_t] = 0$). The resulting pattern is one of concentration driven by internal scaling combined with active but inefficient reallocation, rather than by Schumpeterian selection.

Era 3 (2005–2014): Sclerotic Stagnation.
Era 3 marks a broad-based weakening of both internal expansion and selection. Net organic growth declines sharply ($g_t^W = 0.0087$), and the implied gross internal productivity growth rate falls to $g_c = 0.046$, while the inferred cost of scaling rises ($c_k = 1.31\times$ its Era 1 level). At the same time, the organic outflow term becomes strongly negative, indicating further deterioration in selection, and the M&A execution rate drops to its lowest level in the sample. Although point estimates of acquisition effects are negative, their quantitative importance is limited due to low acquisition activity. Taken together, this period is best understood as one of stagnation driven by rising organizational frictions and weakening reallocation, rather than by active consolidation.

Era 4 (2015–2023): Arrested Dynamism.
By Era 4, the economy transitions into a regime of arrested dynamism. Gross internal productivity growth weakens further ($g_c = 0.032$) as the inferred cost of scaling rises sharply ($c_k = 2.84\times$ its Era 1 level), while selection at the top of the distribution effectively shuts down, as reflected in the near-zero organic outflow contribution. In contrast to earlier periods, acquisitions become both frequent and systematically productivity-reducing, with a large and statistically significant negative effect ($\mathbb{E}[J_t] < 0$). The model interprets this pattern as a shift toward defensive acquisition behavior, in which incumbents preserve their market position at the cost of internal efficiency. Concentration in this regime is no longer associated with productivity growth, but instead reflects the suppression of both internal expansion and competitive reallocation.

4.4 Counterfactuals and the Cost of Arrested Dynamism

The sequential structural mapping of firm behavior across regimes allows us to quantify the macroeconomic consequences of arrested dynamism and evaluate the effects of targeted policy interventions.

Counterfactual 1: Restricting Defensive Shelving.
Consider a policy environment in Era 4 in which antitrust authorities can effectively identify and restrict acquisitions that result in the shelving of acquired technologies. In this case, the top incumbent is constrained to abandon defensive shelving ($I^* = 0$) and instead implements integration ($I^* = 1$). This shift raises aggregate productivity by restoring efficiency-enhancing reallocation through integration and eliminating the negative productivity effects associated with shelving. In terms of the model, the economy moves away from a regime characterized by negative acquisition-induced productivity effects toward one in which acquisitions no longer destroy productive capacity.

Counterfactual 2: R&D Subsidies and the Innovation Paradox.
Alternatively, industrial policy often aims to stimulate corporate dynamism through subsidies to internal R&D. We model this intervention as a reduction in the cost of internal expansion, lowering c_k from its elevated Era 4 level (2.84× its Era 1 baseline) toward earlier-period values.

While this policy successfully increases gross internal productivity growth rate (g_c), it has unintended consequences for acquisition behavior. In equilibrium, a higher g_c reduces the effective discount rate in the firm's valuation expression $(r + \delta_t - 0.5 \cdot g_c - \lambda_{total,t}\Omega_t)$, thereby increasing the normalized value of incumbency $v(\mu_t)$.

Because the firm's acquisition decision depends on the relative valuation of preserving the legacy markup versus integrating new technologies, an increase in $v(\mu_t)$ raises the value of maintaining existing market power. As a result, the firm becomes more likely to satisfy the defensive shelving condition, strengthening the incentive to neutralize rather than integrate acquired firms.

Thus, policies that stimulate internal expansion without addressing the incentives governing acquisition behavior can unintentionally reinforce defensive consolidation. In this sense, subsidizing R&D in highly concentrated environments may exacerbate the very forces underlying arrested dynamism.

5 Conclusion

The rise in corporate concentration over the past four decades has raised central questions about the nature of business dynamism in the US economy. This paper combines a decomposition of firm-level productivity dynamics with a continuous-time structural framework to study how the mechanisms underlying top-share growth have evolved

from 1985 to 2023.

My results point to a shift in the sources of concentration over time. In earlier decades, concentration was associated with strong internal productivity growth and patterns of selection consistent with efficient reallocation. In contrast, more recent periods exhibit weaker internal expansion, a deterioration in selection dynamics, and a growing role for acquisition activity with negative or muted productivity effects. Interpreted through the model, these patterns are consistent with a transition from a regime in which concentration reflects productive scaling and reallocation toward one in which it is increasingly associated with the entrenchment of incumbent positions.

The empirical evidence from event studies and aggregate productivity decompositions indicates that the effects of acquisitions have changed over time. While earlier periods are consistent with efficiency-enhancing restructuring, the most recent decade shows acquisitions with statistically significant negative effects on productivity. Within the model, this shift reflects a change in the tradeoff between integrating new technologies and preserving market power: integration can raise physical productivity, but it also lowers the incumbent's markup. As the value of preserving markup rises, firms have stronger incentives to avoid integration and instead adopt acquisition strategies that protect market power at the cost of productive efficiency.

These findings have implications for how concentration is interpreted in both empirical and policy contexts. First, they suggest that rising concentration need not have uniform effects on productivity, as its underlying drivers can differ across periods. Second, they highlight the importance of considering both internal firm dynamics and selection mechanisms when evaluating aggregate productivity outcomes. Finally, the counterfactual analysis illustrates that policies aimed at stimulating innovation may interact with firms' acquisition incentives in non-trivial ways, particularly in environments where incumbents have substantial market power.

Taken together, the results underscore that understanding the consequences of concentration requires distinguishing between regimes in which it reflects productive scaling and those in which it is associated with weakened reallocation and slower productivity growth.

References

ACEMOGLU, D., U. AKCIGIT, H. ALP, N. BLOOM, AND W. KERR (2018): "Innovation, reallocation, and growth," *American economic review*, 108(11), 3450–3491.

AGHION, P., AND P. HOWITT (1990): "A model of growth through creative destruction," .

AKCIGIT, U., AND S. T. ATES (2021): "What Happened to US Business Dynamism?," *Journal of Political Economy*, 129(8), 2123–2167.

AKCIGIT, U., AND S. T. ATES (2023): "What happened to US business dynamism?," *Journal of Political Economy*, 131(8), 2059–2124.

ARGENTE, D., M. LEE, AND S. MOREIRA (2024): "The life cycle of products: Evidence and implications," *Journal of Political Economy*, 132(2), 337–390.

ARROW, K. J. (1962): "Economic Welfare and the Allocation of Resources for Invention," in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, pp. 609–626. Princeton University Press.

AUTOR, D., D. DORN, L. F. KATZ, C. PATTERSON, AND J. VAN REENEN (2020): "The fall of the labor share and the rise of superstar firms," *The Quarterly journal of economics*, 135(2), 645–709.

BAIN, J. S. (1951): "Relation of profit rate to industry concentration: American manufacturing, 1936–1940," *The Quarterly Journal of Economics*, 65(3), 293–324.

BLOOM, N., C. I. JONES, J. VAN REENEN, AND M. WEBB (2020): "Are Ideas Getting Harder to Find?," *American Economic Review*, 110(4), 1104–1144.

CALLAWAY, B., AND P. H. SANT'ANNA (2021): "Difference-in-differences with multiple time periods," *Journal of econometrics*, 225(2), 200–230.

CHOI, J., A. A. LEVCHENKO, D. RUZIC, AND Y. SHIM (2024): "Superstars Or Supervillains?: Large Firms in the South Korean Growth Miracle," Discussion paper, National Bureau of Economic Research.

COASE, R. H. (1937): "The Nature of the Firm," *Economica*, 4(16), 386–405.

COVARRUBIAS, M., G. GUTIÉRREZ, AND T. PHILIPPON (2020): "From good to bad concentration? US industries over the past 30 years," *NBER Macroeconomics Annual*, 34(1), 1–46.

CUNNINGHAM, C., F. EDERER, AND S. MA (2021): "Killer acquisitions," *Journal of political economy*, 129(3), 649–702.

DE LOECKER, J., J. EECKHOUT, AND G. UNGER (2020): "The Rise of Market Power and the Macroeconomic Implications," *The Quarterly Journal of Economics*, 135(2), 561–644.

DECKER, R. A., J. HALTIWANGER, R. S. JARMIN, AND J. MIRANDA (2016): "Declining business dynamism: What we know and the way forward," *American Economic Review*, 106(5), 203–207.

——— (2020): "Changing business dynamism and productivity: Shocks versus responsiveness," *American Economic Review*, 110(12), 3952–3990.

DEMSETZ, H. (1973): "Industry structure, market rivalry, and public policy," *The Journal of Law and Economics*, 16(1), 1–9.

GROSSMAN, G. M., AND E. HELPMAN (1991): "Trade, knowledge spillovers, and growth," *European economic review*, 35(2-3), 517–526.

GRULLON, G., Y. LARKIN, AND R. MICHAELY (2019): "Are US industries becoming more concentrated?," *Review of Finance*, 23(4), 697–743.

GUTIÉRREZ, G., AND T. PHILIPPON (2019): "Fading stars," in *AEA Papers and Proceedings*, vol. 109, pp. 312–316. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

HALTIWANGER, J. (2015): "Job creation, job destruction, and productivity growth: The role of young businesses," *Annual Review of Economics*, 7(1), 341–358.

KWON, S. Y., Y. MA, AND K. ZIMMERMANN (2024): "100 years of rising corporate concentration," *American Economic Review*, 114(7), 2111–2140.

MELITZ, M. J., AND S. POLANEC (2015): "Dynamic Olley-Pakes productivity decomposition with entry and exit," *The RAND Journal of Economics*, 46(2), 362–375.

ROLDAN-BLANCO, P., T. SCHMITZ, AND C. FONS-ROSEN (2026): "The Effects of Startup Acquisitions on Innovation and Economic Growth," Discussion paper.

SCHUMPETER, J. A. (1942): *Capitalism, Socialism and Democracy*. Harper & Brothers, New York.

A Empirical Appendix: Detailed Data Construction

A.1 Aggregate Macroeconomic Denominators

Aggregate Gross Output: To construct a continuous nominal and real gross output series spanning 1947 to 2024, I splice historical BEA data releases (1947–1997) with the modern BEA publications (1997–2024). Specifically, I isolate the gross output of Private Industries (BEA Line TGO105-A) to serve as the aggregate sales denominator for my concentration metrics.

Firm Universe (N): To calculate the top 0.01% threshold accurately across the 1985–2023 core sample period, I merge the panel with the continuous count of total U.S. firms provided by the Census *Business Dynamics Statistics* (BDS).

A.2 M&A Pipeline and Target Sales Fallback Rules

My "Owner-Spine" pipeline processes Refinitiv M&A transactions through four distinct hierarchical lanes to maximize Compustat linkage coverage:

1. `acq_assetfirst`: Prioritizes matching the acquirer's direct asset-level identifiers (e.g., 6-digit CUSIP).

2. `acq_parentfirst`: Prioritizes matching the acquirer's ultimate parent-level identifiers.

3. `tgt_assetfirst`: Prioritizes matching the target's direct asset-level identifiers.

4. `tgt_parentfirst`: Prioritizes matching the target's ultimate parent-level identifiers.

Target Sales Proxy Logic: To construct the mechanical M&A sales add-on, I implement a strict fallback hierarchy for missing Refinitiv target sales (`salesltm`). When `salesltm` is unavailable, I search for the target asset's `GVKEY` sales (`sale` or `revt`) in Compustat within a $\pm$2-year window surrounding the deal's effective year. If the matched target reports in a non-USD currency, I apply the contemporaneous exchange rate conversion factor (`curuscn`) before integrating the proxy sales into the mechanical growth accounting.

B Empirical Appendix: Proof of the Corrected Top-Share Decomposition

Here I provide the mathematical derivation of the threshold-centered top-share decomposition.

B.1 Proof of the Top-Share Decomposition Identity

Recall the definition of the top share at time t: $\mathsf{S}_t = A_t^{ALL}/\Omega_t$. The percentage growth in the top share from $t = 0$ to $t = 1$ is:

$$\frac{\mathsf{S}_1 - \mathsf{S}_0}{\mathsf{S}_0} = \frac{\frac{A_1^{ALL}}{\Omega_1} - \frac{A_0^{ALL}}{\Omega_0}}{\frac{A_0^{ALL}}{\Omega_0}} = \frac{A_1^{ALL}}{A_0^{ALL}} \left(\frac{\Omega_0}{\Omega_1} \right) - 1 \tag{16}$$

Substituting the aggregate scale factor $\Xi \equiv \Omega_0/\Omega_1$:

$$\frac{\mathsf{S}_1 - \mathsf{S}_0}{\mathsf{S}_0} = \Xi \left(1 + \frac{A_1^{ALL} - A_0^{ALL}}{A_0^{ALL}} \right) - 1 = (\Xi - 1) + \Xi \cdot \frac{A_1^{ALL} - A_0^{ALL}}{A_0^{ALL}} \tag{17}$$

The term $(\Xi - 1)$ captures aggregate dilution. The true firm dynamics are contained within the absolute numerator change $A_1^{ALL} - A_0^{ALL}$.

B.2 Algebraic Expansion of the Corrected Dynamics

To map the empirical flows to the theoretical jump-diffusion processes, I expand the total change using the disjoint set partitions: $P_1 = S \cup E \cup B$ and $P_0 = S \cup X \cup D$. As detailed in the main text, my micro-auditing corrects the raw death set (D) by re-partitioning it into structurally verified true outcomes: $D = D^{OG_true} \cup D^{MA_total} \cup D^{Artifact}$.

The total absolute change in the top tier is therefore:

$$\begin{aligned} A_1^{ALL} - A_0^{ALL} = & \left(\sum_S u_{i1}^{ALL} + \sum_E u_{i1}^{ALL} + \sum_B u_{i1}^{ALL} \right) \\ & - \left(\sum_S u_{i0}^{ALL} + \sum_X u_{i0}^{ALL} + \sum_{D^{OG_true}} u_{i0}^{ALL} + \sum_{D^{MA_total}} u_{i0}^{ALL} + \sum_{D^{Artifact}} u_{i0}^{ALL} \right) \end{aligned} \tag{18}$$

To isolate organic technological scaling from mechanical consolidation, I decompose $u^{ALL} = u^{OG} + u^{MA}$ specifically for surviving sets (S and E). For exiting and displaced sets (X and D), I deduct their *entire* baseline size (u^{ALL}), as their mass is entirely lost from the respective distributions. I group these terms into three components:

1. The Data Artifact Term ($\Delta A_{Artifact}$)
I isolate and drop the data artifact term:

$$\Delta A_{Artifact} = -\sum_{i \in D^{Artifact}} u_{i0}^{ALL} \tag{19}$$

2. The Mechanical M&A Term (ΔA_{MA})
Grouping the remaining M&A-specific components yields the pure mechanical consolidation flows:

$$\Delta A_{MA} = \sum_{i \in S}(u_{i1}^{MA} - u_{i0}^{MA}) + \sum_{i \in E} u_{i1}^{MA} + \sum_{i \in B} u_{i1}^{ALL} - \sum_{i \in D^{MA_total}} u_{i0}^{ALL} \tag{20}$$

3. The Organic Growth Term (ΔA_{OG})
Grouping the remaining organic changes yields:

$$\Delta A_{OG} = \sum_{i \in S}(u_{i1}^{OG} - u_{i0}^{OG}) + \sum_{i \in E} u_{i1}^{OG} - \sum_{i \in X} u_{i0}^{ALL} - \sum_{i \in D^{OG_true}} u_{i0}^{ALL} \tag{21}$$

To center the dynamic entry and displacement around the Pareto boundary (the marginal threshold $q_1^{(ALL)}$), I add and subtract $q_1^{(ALL)}$ to the respective sums:

$$\begin{aligned} \sum_{i \in E} u_{i1}^{OG} - \sum_{i \in X} u_{i0}^{ALL} &= \sum_{i \in E}(u_{i1}^{OG} - q_1^{(ALL)} + q_1^{(ALL)}) - \sum_{i \in X}(u_{i0}^{ALL} - q_1^{(ALL)} + q_1^{(ALL)}) \\ &= \sum_{i \in E}(u_{i1}^{OG} - q_1^{(ALL)}) + \sum_{i \in X}(q_1^{(ALL)} - u_{i0}^{ALL}) + (|E| - |X|)q_1^{(ALL)} \end{aligned} \tag{22}$$

Substituting this threshold-centered expression back into ΔA_{OG}:

$$\Delta A_{OG} = \sum_{i \in S}(u_{i1}^{OG} - u_{i0}^{OG}) + \sum_{i \in E}(u_{i1}^{OG} - q_1^{(ALL)}) + \sum_{i \in X}(q_1^{(ALL)} - u_{i0}^{ALL}) - \sum_{i \in D^{OG_true}} u_{i0}^{ALL} + (|E| - |X|)q_1^{(ALL)} \tag{23}$$

B.3 Conclusion of the Proof

Summing the three parsed components recovers the total exact change identically:

$$\Delta A_{OG} + \Delta A_{MA} + \Delta A_{Artifact} = A_1^{ALL} - A_0^{ALL} \tag{24}$$

Multiplying these absolute changes by the scalar (Ξ / A_0^{ALL}) completes the exact identities for Δ_{corr}^{OG}, Δ_{corr}^{MA}, and $\Delta^{Artifact}$ provided in the main text (Equations 1, 2, and 3).

C Details of the Melitz–Polanec Productivity Decomposition

This section details the exact derivation of the Melitz–Polanec (Melitz and Polanec, 2015) decomposition deployed in Section 3.3, including our novel sub-decomposition of the Net Entry term to align perfectly with the micro-audited flow sets defined in Section 3.

Let ϕ_{it} denote the log markup-adjusted physical productivity (TFPQ) of firm i at time t. Aggregate productivity for the top tier P_t is the share-weighted average: $\Phi_t = \sum_{i \in P_t} s_{it} \phi_{it}$, where $s_{it} = Y_{it} / \sum_{j \in P_t} Y_{jt}$.

Following Olley and Pakes (1996), aggregate productivity can be rewritten as the sum of unweighted average productivity ($\bar{\phi}_t$) and the covariance between market share and productivity:

$$\Phi_t = \bar{\phi}_t + \sum_{i \in P_t} (s_{it} - \bar{s}_t)(\phi_{it} - \bar{\phi}_t). \tag{25}$$

To measure changes over time $t \in \{0, 1\}$, we partition the firms into three mutually exclusive sets: Stayers ($S = P_0 \cap P_1$), all Entrants ($E_{all} = P_1 \setminus P_0$), and all Exits ($X_{all} = P_0 \setminus P_1$). The Melitz and Polanec (2015) dynamic decomposition of the change in aggregate productivity ($\Delta\Phi = \Phi_1 - \Phi_0$) is exactly:

$$\Delta\Phi = \underbrace{\sum_{i \in S} \bar{s}_i \Delta\phi_i}_{\text{Within-Firm}} + \underbrace{\sum_{i \in S} \Delta s_i (\bar{\phi}_i - \bar{\Phi}_S)}_{\text{Reallocation}} + \underbrace{s_{E_{all},1}(\Phi_{E_{all},1} - \Phi_{S1})}_{\text{Entry Contribution}} - \underbrace{s_{X_{all},0}(\Phi_{X_{all},0} - \Phi_{S0})}_{\text{Exit Contribution}}. \tag{26}$$

Sub-Decomposition of Net Entry into Organic and Mechanical Flows. To strictly map this productivity framework to our exact M&A accounting audit, we further disaggregate the comprehensive Entry and Exit sets into their underlying structural components.

Unlike the continuous sales accounting in Section 3, productivity is a firm-level attribute. Therefore, we partition the discrete set of entering firms ($E_{all} = P_1 \setminus P_0$) into two mutually exclusive sets based on their growth strategy: E_{org} captures *Organic Inflow* (entrants scaling entirely through internal growth), and E_{MA} captures *M&A Driven Inflow* (entrants that relied on acquisitions to cross the top-tier threshold).

Similarly, the departing set ($X_{all} = P_0 \setminus P_1$) is partitioned into three economic sets: $X_{demoted}$ captures *Organic Outflow* (legacy firms that survive in the economy but are demoted below the threshold), X_{MA} captures *Exit: M&A Targets* (highly productive firms eliminated from the top tier because they were absorbed by another entity), and X_{dead} isolates true *Organic Death* (bankruptcies).[2]

Expanding the Entry and Exit contributions over these mutually exclusive subsets yields the exact empirical margins plotted in Figure 5:

$$\begin{aligned}\text{Net Entry Drag} = &\underbrace{s_{E_{org},1}(\Phi_{E_{org},1} - \Phi_{S1})}_{\text{Organic Inflow}} + \underbrace{s_{E_{MA},1}(\Phi_{E_{MA},1} - \Phi_{S1})}_{\text{M\&A Driven Inflow}} \\ &- \underbrace{s_{X_{demoted},0}(\Phi_{X_{demoted},0} - \Phi_{S0})}_{\text{Organic Outflow (Demoted)}} - \underbrace{s_{X_{MA},0}(\Phi_{X_{MA},0} - \Phi_{S0})}_{\text{Exit: M\&A Targets}} - \underbrace{s_{X_{dead},0}(\Phi_{X_{dead},0} - \Phi_{S0})}_{\text{Exit: Organic Death}}.\end{aligned} \tag{27}$$

This granular formulation isolates exactly how defensive M&A suppresses the aggregate distribution by stripping high-TFPQ targets (X_{MA}) out of the ecosystem without transferring their underlying physical efficiency to the acquirer's Within-firm term.

D Empirical Evidence for the Kill Zone: Shrinking Target Sizes

In Section 3.3, we document that the negative drag on the aggregate Net Entry productivity term is driven by a surge in "Acquired Exits"—highly productive targets being mechanically absorbed by apex incumbents. This appendix provides the direct institutional evidence that this phenomenon reflects a structural shift from synergistic integration toward defensive "kill zone" acquisitions (Cunningham, Ederer, and Ma, 2021).

[2] As in our accounting decomposition, firms that disappear due to pure reporting illusions ($X_{artifact}$) are strictly isolated and excluded from our aggregate productivity inference.

The Plunge in Target Size. During the Superstar Scaling era of the 1990s, M&A by top-tier firms was characterized by the integration of large, synergistic peers. Figure 11 demonstrates that over the last two decades, this paradigm has completely collapsed. Across all cuts of the data, the median ratio of target sales to acquirer sales has plunged by one to two orders of magnitude. In the modern era, the Top 0.01% overwhelmingly execute acquisitions of tiny startups—frequently valued at less than 0.1% of the acquirer's own size. Because these targets are too small to move the needle on the acquirer's physical technology frontier, their primary economic function is the preemptive elimination of future competitive threats.

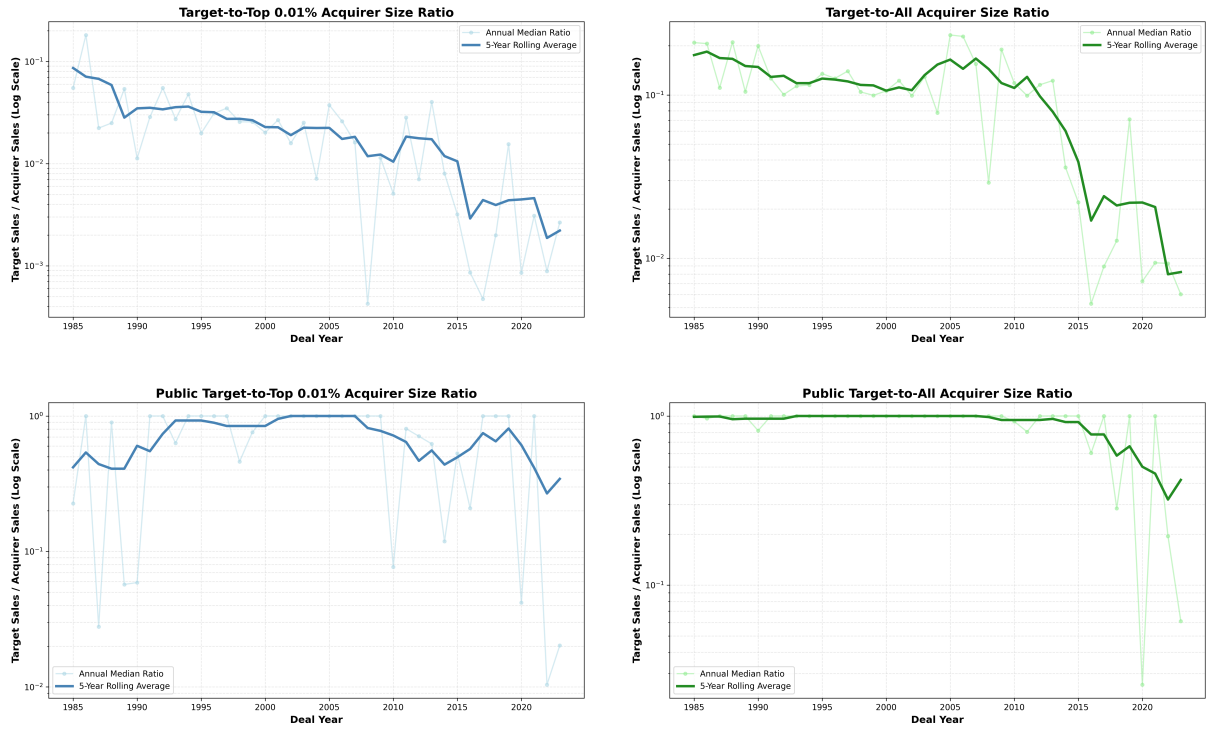


Figure 11: The Plunge in Target-to-Acquirer Size Ratios (Log Scale). All panels plot the median ratio of target sales to acquirer sales for completed deals, smoothed with a five-year centered rolling average (bold line) over annual medians (light dots). *Top row*: all targets (public and private combined). *Bottom row*: publicly listed targets only (active Compustat firm in the deal year or prior year). *Left column*: Top 0.01% acquirers only. *Right column*: all acquirers (economy-wide baseline). Across all four cuts, the ratio has fallen by one to two orders of magnitude since the mid-1980s, with the steepest decline occurring after 2005.

The Vanishing Public Target. Concurrently, the composition of these targets has shifted aggressively out of the public domain. Figure 12 illustrates the "vanishing public target" paradox. While aggregate M&A dollar volumes have reached record highs in the modern era, the share of these deals involving publicly listed (Compustat) targets has plummeted.

Today, acquisitions by the Top 0.01% are overwhelmingly concentrated in the private markets. By sheltering these acquisitions in the opaque private startup ecosystem, apex firms can more easily execute defensive consolidation away from the scrutiny of traditional public-market antitrust enforcement.

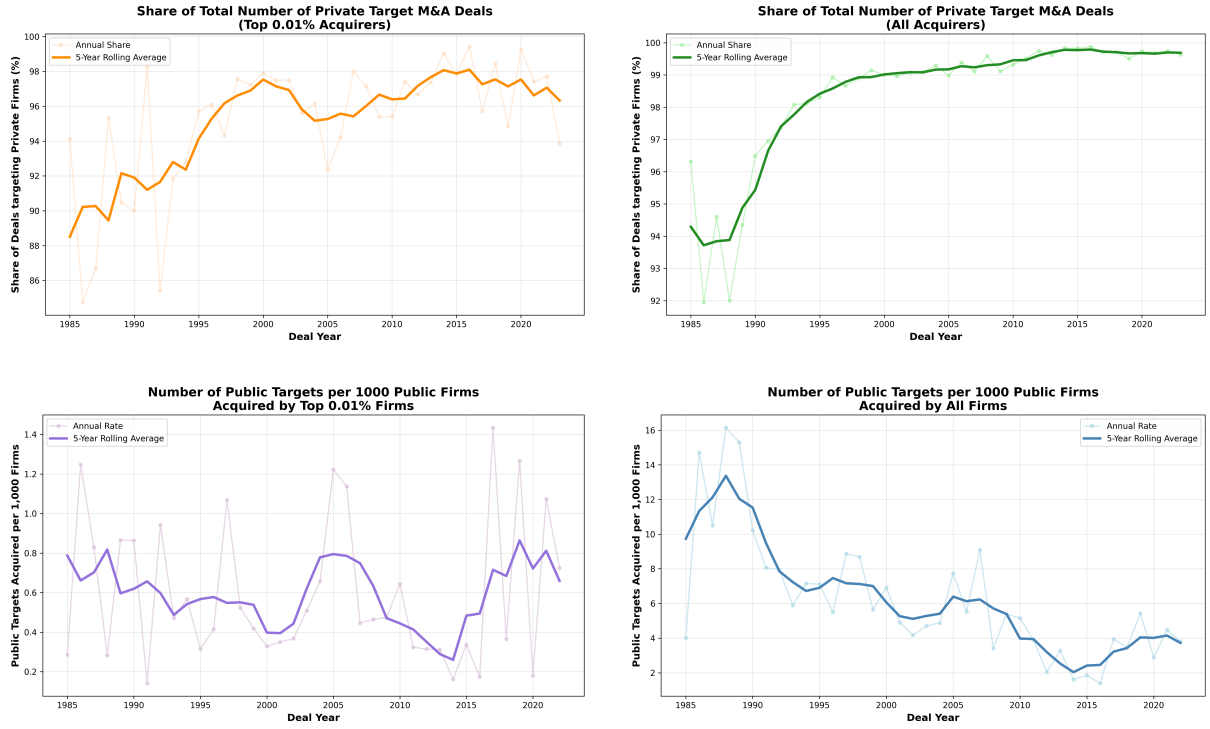


Figure 12: The Vanishing Public Target: Private-Deal Share and Public Acquisition Rates. *Top row*: share of completed M&A deals in which the target does not appear in Compustat (i.e., is privately held), computed annually and smoothed with a five-year centered rolling average. *Bottom row*: number of publicly listed targets acquired per 1,000 active Compustat firms, smoothed with a five-year centered rolling average; truncated at 2022. *Left column*: Top 0.01% acquirers only. *Right column*: all acquirers (economy-wide baseline). Together, the panels show that the M&A market has shifted overwhelmingly toward private targets while the rate at which public firms are absorbed has declined since the mid-2000s merger waves, despite rising aggregate M&A dollar volumes.

E Theory Appendix: Detailed Solution to the Firm Dynamics Model

This appendix provides the complete mathematical micro-foundations and step-by-step solutions to the continuous-time firm dynamics model presented in Section 4. I derive the firm's profit function from primitive Constant Elasticity of Substitution (CES) preferences, set up the Hamilton-Jacobi-Bellman (HJB) equation, perform the Fokker-Planck integra-

tion over the Pareto tail, and map the sequential calibration strategy that unites the micro and macro frameworks.

E.1 Preferences, Demand, and Limit Pricing

Consider a continuous-time economy where a representative consumer has Constant Elasticity of Substitution (CES) preferences over a continuum of intermediate goods $j \in [0, 1]$. The aggregate consumption index Y_t is given by:

$$Y_t = \left(\int_0^1 \left(q_{j,t} x_{j,t} \right)^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}, \quad \sigma > 1 \tag{28}$$

where $x_{j,t}$ is the quantity consumed of variety j, and $q_{j,t}$ represents the objective quality (or physical efficiency) of the variety. The elasticity of substitution between varieties is denoted by σ. Let E_t be total aggregate expenditure and P_t be the ideal price index. The standard demand function for variety j is:

$$x_{j,t} = q_{j,t}^{\sigma-1} \left(\frac{p_{j,t}}{P_t} \right)^{-\sigma} \frac{E_t}{P_t} \tag{29}$$

A Stayer incumbent possesses a fundamental productivity $\phi_{j,t}$, yielding a marginal cost of $w/\phi_{j,t}$, where w is the nominal wage. Assuming the incumbent engages in limit pricing against a competitive fringe, it charges a markup $\mu_{j,t}$. The firm's price is:

$$p_{j,t} = \mu_{j,t} \frac{w}{\phi_{j,t}} \tag{30}$$

The firm's operating profit is defined as revenue minus variable costs. Normalizing aggregate demand terms, the profit flow can be written as proportional to the firm's productivity:

$$\pi(\phi_{j,t}, \mu_{j,t}) = \phi_{j,t} \left(1 - \frac{1}{\mu_{j,t}} \right) = \phi_{j,t} \cdot \pi(\mu_{j,t}), \tag{31}$$

where $\pi(\mu_{j,t})$ is increasing in the markup. This linearity arises from the CES demand system combined with limit pricing.

E.2 Law of Motion for Productivity

Firm productivity follows a jump-diffusion process combining continuous internal investment and discrete acquisition shocks. In log terms, the law of motion of a firm j's productivity is given by:

$$d \ln \phi_{j,t} = (g_c - \delta_t)\, dt + J_t\, dN_t, \tag{32}$$

where N_t is a Poisson process with intensity $\lambda_{total,t}$.

At each acquisition event (i.e., when $dN_t = 1$), productivity evolves according to:

$$\phi_{j,t} \to \phi_{j,t} e^{J_t}, \tag{33}$$

where J_t denotes the realized log productivity jump induced by the acquisition. Equivalently, $e^{J_t} - 1$ corresponds to the percentage change in productivity associated with the event. The distribution of J_t depends on the firm's acquisition decision. Under integration ($I = 1$), J_t is centered around γ, while under defensive shelving ($I = 0$), it is centered around $-\tau$. Here γ and $-\tau$ denote log productivity changes.

Taking expectations over the jump process yields:

$$\mathbb{E}[d \ln \phi_{j,t}] = (g_c - \delta_t)\, dt + \lambda_{total,t} \mathbb{E}[J_t]\, dt. \tag{34}$$

E.3 Derivation of the Normalized HJB Equation

Let $V(\mu_t, \phi_t)$ denote the enterprise value of the firm. The firm optimally chooses gross internal investment g_c to maximize this value. Achieving a gross internal growth rate g_c requires the firm to pay a convex investment cost $C(g_c) = \frac{1}{2} \cdot c_k \cdot g_c^2 \phi_t$, where c_k parameterizes the inferred cost of internal expansion.

The unnormalized HJB equation is:

$$rV(\mu_t, \phi_t) = \max_{g_c \geq 0, I \in \{0,1\}} \left\{ \pi(\mu_t)\phi_t - \frac{1}{2} c_k g_c^2 \phi_t + (g_c - \delta_t)\phi_t \frac{\partial V}{\partial \phi_t} + \lambda_{total,t} \mathbb{E}\Big[V(\mu_{post}, \phi_t e^{J_t}) - V(\mu_t, \phi_t) \Big] \right\}, \tag{35}$$

where δ_t reflects the continuous natural obsolescence rate, $\lambda_{total,t}$ is the observed M&A execution rate, and the post-acquisition markup μ_{post} depends on the firm's discrete choice:

$$\mu_{post} = \begin{cases} \mu_{int} & \text{if } I = 1 \text{ (integration)}, \\ \mu_t & \text{if } I = 0 \text{ (defensive shelving)}. \end{cases}$$

Because the limit-pricing markup μ_t summarizes the firm's distance to the competitive frontier, I normalize the value function by productivity and conjecture:

$$V(\mu_t, \phi_t) = v(\mu_t)\phi_t. \tag{36}$$

Substituting into the HJB yields:

$$rv(\mu_t)\phi_t = \max_{g_c \geq 0, I \in \{0,1\}} \left\{ \pi(\mu_t)\phi_t - \frac{1}{2}c_k g_c^2 \phi_t + v(\mu_t)(g_c - \delta_t)\phi_t + \lambda_{total,t}\mathbb{E}\left[v(\mu_{post})\phi_t e^{J_t} - v(\mu_t)\phi_t\right] \right\}. \tag{37}$$

Dividing through by ϕ_t, and defining the expected proportional change in enterprise value from an acquisition as

$$\Omega_t \equiv \mathbb{E}\left[\frac{v(\mu_{post})}{v(\mu_t)} e^{J_t} - 1\right],$$

I obtain the normalized HJB equation:

$$rv(\mu_t) = \max_{g_c \geq 0, I \in \{0,1\}} \left\{ \pi(\mu_t) - \frac{1}{2}c_k g_c^2 + v(\mu_t)(g_c - \delta_t) + \lambda_{total,t}\Omega_t v(\mu_t) \right\}. \tag{38}$$

The object Ω_t captures the expected proportional change in enterprise value from an acquisition, incorporating both the physical productivity jump and the change in markup induced by the firm's integration decision. To bring this object to the data, I adopt a tractable approximation for the continuation-value ratio $v(\mu_{post})/v(\mu_t)$ that depends on the firm's acquisition strategy.

Under **integration** ($I = 1$), acquisitions may increase physical productivity while lowering the incumbent markup. I capture this effect through a reduced-form proportional haircut on continuation value, denoted by $\chi \in (0, 1)$, so that

$$\frac{v(\mu_{post})}{v(\mu_t)} \approx 1 - \chi.$$

The exact effective firm-value jump is then

$$\Omega_t(I = 1) = (1 - \chi)e^{\mathbb{E}[J_t | I=1]} - 1.$$

Under the model primitives, $\mathbb{E}[J_t \mid I = 1] = \gamma$, so this can be written as

$$\Omega_t(I = 1) = (1 - \chi)e^{\gamma} - 1.$$

In implementation, I use the first-order Taylor approximation

$$e^{\mathbb{E}[J_t|I=1]} \approx 1 + \mathbb{E}[J_t \mid I = 1],$$

so that

$$\Omega_t(I = 1) \approx (1 - \chi)(1 + \gamma) - 1.$$

Under **defensive shelving** ($I = 0$), the incumbent preserves its legacy markup, implying $v(\mu_{post})/v(\mu_t) \approx 1$. The exact effective firm-value jump is therefore

$$\Omega_t(I = 0) = e^{\mathbb{E}[J_t|I=0]} - 1.$$

Under the model primitives, $\mathbb{E}[J_t \mid I = 0] = -\tau$, so that

$$\Omega_t(I = 0) = e^{-\tau} - 1,$$

which I analogously approximate by

$$\Omega_t(I = 0) \approx -\tau.$$

In intermediate cases where acquisition effects are statistically imprecise and their aggregate contribution is small, I impose the neutral benchmark $\Omega_t = 0$.

Optimal Policies. The firm's problem yields two optimality conditions corresponding to the continuous control g_c and the discrete acquisition choice I.

Optimal Internal Investment. The first-order condition with respect to continuous internal investment g_c equates marginal cost and marginal benefit:

$$-c_k g_c + v(\mu_t) = 0 \quad \Rightarrow \quad g_c^* = \frac{v(\mu_t)}{c_k}. \tag{39}$$

This condition reflects the linear scaling of firm value in productivity: the marginal value of increasing productivity is $v(\mu_t)$, while the marginal cost is $c_k g_c$.

Optimal M&A Policy. The discrete acquisition decision compares continuation values under integration and defensive shelving. The firm prefers defensive shelving ($I^* = 0$) if

and only if:

$$v(\mu_t)e^{-\tau} > v(\mu_{int})e^{\gamma} \quad \Longleftrightarrow \quad \frac{v(\mu_t)}{v(\mu_{int})} > e^{\gamma+\tau}. \tag{40}$$

This condition highlights the fundamental tradeoff between preserving incumbent markups and realizing productivity gains from integration. Defensive shelving is optimal when the value of maintaining market power outweighs the efficiency gains from incorporating acquired technologies.

E.4 Fokker-Planck Integration and the Macro Distribution

Let $p(x,t)$ denote the probability density function of firms over log productivity $x = \ln \phi$. Given the firm-level jump-diffusion process, the evolution of the distribution is governed by the Kolmogorov Forward Equation (KFE), or Fokker–Planck equation:

$$\frac{\partial p(x,t)}{\partial t} = -\frac{\partial}{\partial x}\Big[(g_c - \delta_t)p(x,t)\Big] - \lambda_{total,t}p(x,t) + \lambda_{total,t}\int p(x-J,t)\,dF_J(J), \tag{41}$$

where F_J denotes the cumulative distribution function of the log productivity jump J_t, so that $dF_J(J)$ represents the probability measure over jump realizations.

Stationary Pareto Tail. I focus on stationary environments corresponding to each empirical era, where $\partial p(x,t)/\partial t = 0$. Empirically, the upper tail of the firm size distribution is well approximated by a Pareto distribution. In log terms, this implies an exponential density:

$$p(x) = \zeta_t e^{-\zeta_t(x - x_{\min,t})}, \quad x \geq x_{\min,t}, \tag{42}$$

where ζ_t is the Pareto tail exponent and $x_{\min,t}$ is the threshold defining the Top 0.01%.

Boundary Flow Condition. To characterize the flow of firms exiting the top tail, consider the lower boundary at $x_{\min,t}$. Firms cross this boundary due to continuous downward drift induced by obsolescence. The instantaneous drift in log productivity is $(g_c - \delta_t)$, so the downward component of drift is governed by δ_t, which captures the rate at which firms lose relative position in the productivity distribution.

The outflow of firms across the boundary is given by the probability flux at $x_{\min,t}$. In a diffusion setting, the boundary outflow induced by deterministic drift is:

$$m_{org,t} = \delta_t \cdot p(x_{\min,t}), \tag{43}$$

which follows from evaluating the drift-induced probability current at the boundary.

Evaluating the Density at the Threshold. I work with the conditional distribution of firms in the upper tail (Top 0.01%), so that the density $p(x)$ is normalized over the truncated support $x \geq x_{\min,t}$:

$$\int_{x_{\min,t}}^{\infty} p(x)\, dx = 1.$$

Under this normalization, the Pareto density implies:

$$p(x) = \zeta_t e^{-\zeta_t (x - x_{\min,t})}, \tag{44}$$

so that the density evaluated at the threshold is:

$$p(x_{\min,t}) = \zeta_t. \tag{45}$$

Because both the boundary mass $m_{org,t}$ and the density $p(x_{\min,t})$ are defined with respect to the same truncated distribution, this normalization does not affect the identification of δ_t.

Mapping to Empirical Outflow. Combining the boundary flow condition with the Pareto density yields:

$$m_{org,t} = \delta_t \cdot \zeta_t, \tag{46}$$

which implies the structural mapping:

$$\delta_t = \frac{m_{org,t}}{\zeta_t}. \tag{47}$$

This mapping links the model-implied obsolescence rate to the empirically observed mass of firms that organically exit the Top 0.01% percentile.

E.5 Sequential Structural Mapping

To recover the structural parameters governing firm dynamics across eras, I construct a sequential identification strategy between aggregate moments from the Melitz–Polanec decomposition and firm-level acquisition effects from the event-study estimates. Each step uses a small set of empirically observed moments to discipline the corresponding model object.

Step 1: Macro Obsolescence Rate (δ_t)
Data Moment Used: MP Organic Demographic Outflow Mass

From the law of motion for productivity, expected log productivity growth for incumbents satisfies:

$$\mathbb{E}[d\ln\phi_t] = (g_c - \delta_t)dt + \lambda_{total,t}\mathbb{E}[J_t]dt. \tag{48}$$

Aggregating over the stayer distribution yields the net organic growth term:

$$g_t^W = \int_{Z_t}^{\infty} \left(g_c - \delta_t + \lambda_{total,t}\mathbb{E}[J_t]\right) \frac{f(\phi)}{1 - F(Z_t)} d\phi = g_c - \delta_t + \lambda_{total,t}\mathbb{E}[J_t], \tag{49}$$

where the equality follows because the drift and jump terms are common across firms in the top tail.

The expected growth identity alone does not separately identify δ_t. To pin down δ_t, I use the stationary-tail condition derived from the Kolmogorov Forward Equation. As shown in equation 47, the Pareto upper-tail approximation implies $\delta_t = m_{org,t}/\zeta_t$, where $m_{org,t}$ denotes the empirical mass of firms exiting the top tail and ζ_t is the Pareto exponent.

Step 2: Gross Internal Productivity Growth Rate (g_c)
Data Moments Used: MP Stayer Growth Rate (g_t^W), Acquisition Rate ($\lambda_{total,t}$), Acquisition Effect ($\mathbb{E}[J_t]$)

I recover gross internal productivity growth using the model-implied decomposition of expected productivity growth for incumbent firms. From the law of motion,

$$\mathbb{E}[d\ln\phi_{j,t}] = (g_c - \delta_t)\,dt + \lambda_{total,t}\mathbb{E}[J_t]\,dt, \tag{50}$$

so that expected net productivity growth reflects the combined contribution of internal expansion, obsolescence, and acquisition-driven jumps.

Aggregating over the distribution of incumbent firms in the top tail yields:

$$g_t^W = g_c - \delta_t + \lambda_{total,t}\mathbb{E}[J_t], \tag{51}$$

where g_t^W is the empirical counterpart measured as the weighted average log productivity growth of Melitz–Polanec stayers:

$$g_t^W = \frac{1}{T} \sum_{i \in \text{Stayers}} \left[\frac{1}{2}\left(s_{i,start} + s_{i,end}\right)\right] \left(\ln\phi_{i,end} - \ln\phi_{i,start}\right). \tag{52}$$

Rearranging yields a closed-form mapping for gross internal productivity growth:

$$g_c = g_t^W + \delta_t - \lambda_{total,t}\mathbb{E}[J_t]. \tag{53}$$

Calibration Note: For Eras 2 and 3, the micro-level estimates of $\mathbb{E}[J_t]$ are statistically imprecise and their implied aggregate contribution is small. I therefore impose a conservative structural null ($\mathbb{E}[J_t] = 0$), so that $g_c = g_t^W + \delta_t$ in these periods.

Step 3: Normalized Value Function ($v(\mu_t)$)
Data Moments Used: Profit Function ($\pi(\mu_t)$), Discount Rate (r), Firm Value Jump (Ω_t)

Given g_c and δ_t as disciplined by the data, I recover the firm's steady-state normalized valuation by evaluating the normalized HJB equation along the equilibrium path. The optimal internal investment condition implies:

$$c_k \cdot g_c = v(\mu_t), \tag{54}$$

which allows me to substitute out the cost parameter and express firm value directly in terms of observable moments.

A key object in the valuation equation is the expected proportional change in firm value induced by acquisition, defined as

$$\Omega_t \equiv \mathbb{E}\left[\frac{v(\mu_{post})}{v(\mu_t)}e^{J_t} - 1\right].$$

This object captures both the productivity effects of acquisitions and the endogenous response of markups through the continuation-value ratio.

To take this object to the data, I operationalize the continuation-value ratio $v(\mu_{post})/v(\mu_t)$ using a tractable approximation that depends on the firm's integration decision.

Under **integration** ($I = 1$), acquisitions increase physical productivity but reduce the incumbent's markup due to cannibalization across product lines. I capture this effect through a reduced-form proportional haircut on continuation value, denoted by $\chi \in (0, 1)$, so that

$$\frac{v(\mu_{post})}{v(\mu_t)} \approx 1 - \chi.$$

The exact implied change in firm value is therefore

$$\Omega_t(I=1) = (1-\chi)e^{\mathbb{E}[J_t|I=1]} - 1.$$

Under the model primitives, $\mathbb{E}[J_t \mid I=1] = \gamma$, so this can be written as

$$\Omega_t(I=1) = (1-\chi)e^{\gamma} - 1.$$

In implementation, I use the first-order approximation

$$\Omega_t(I=1) \approx (1-\chi)(1+\gamma) - 1,$$

which is accurate for the empirically relevant range of acquisition effects.

Under **defensive shelving** ($I=0$), the incumbent preserves its existing markup, implying $v(\mu_{post})/v(\mu_t) \approx 1$. In this case,

$$\Omega_t(I=0) = e^{\mathbb{E}[J_t|I=0]} - 1.$$

Under the model primitives, $\mathbb{E}[J_t \mid I=0] = -\tau$, so

$$\Omega_t(I=0) = e^{-\tau} - 1 \approx -\tau.$$

In intermediate regimes where acquisition effects are statistically insignificant in the event-study estimates and their aggregate contribution is small, I impose a conservative neutral benchmark $\Omega_t = 0$.

Substituting the definition of Ω_t into the normalized HJB yields:

$$rv(\mu_t) = \pi(\mu_t) - \frac{1}{2}v(\mu_t)g_c + v(\mu_t)g_c - v(\mu_t)\delta_t + \lambda_{total,t}v(\mu_t)\Omega_t. \tag{55}$$

Rearranging gives the closed-form solution for firm value:

$$v(\mu_t) = \frac{\pi(\mu_t)}{r + \delta_t - 0.5 \cdot g_c - \lambda_{total,t}\Omega_t}. \tag{56}$$

This expression provides a direct mapping from observable moments to firm valuation, with acquisition behavior entering through Ω_t.

Calibration Note: In the empirical implementation, the markup-dilution parameter χ is cal-

ibrated using the CS-DiD markup regression, where the estimated post-acquisition reduction in markups under integration provides the empirical counterpart to the continuation-value haircut associated with markup dilution.

Step 4: Inferred Cost of Internal Expansion (c_k)

Finally, given $v(\mu_t)$ and g_c, I recover the implied cost of internal expansion:

$$c_k = \frac{v(\mu_t)}{g_c}. \tag{57}$$

While I interpret c_k as capturing organizational frictions or scaling costs, it is structurally identified here as the parameter required to rationalize observed internal growth rates given the value of incumbency. I normalize c_k relative to Era 1 to highlight changes in the cost of scaling over time.

F Appendix Figures

F.1 Complete Event-Study Evidence Across Eras

This appendix presents the complete set of dynamic treatment effects estimated using the Callaway and Sant'Anna (2021) Doubly Robust Difference-in-Differences estimator. Each panel reports event-time coefficients (β_k) relative to the year immediately preceding the M&A event ($t-1$), corresponding to the dynamic specification:

$$Y_{i,t} = \alpha_i + \tau_t + \sum_{k=-4, k\neq -1}^{5} \beta_k \cdot D_{i,t}^k + \Gamma' X_{i,t-1} + \varepsilon_{i,t} \tag{58}$$

where pre-treatment controls $X_{i,t-1}$ include lagged log total assets, lagged cash ratio, lagged ROA, and lagged R&D intensity. Shaded bars denote 95% confidence intervals, with standard errors clustered at the firm level.

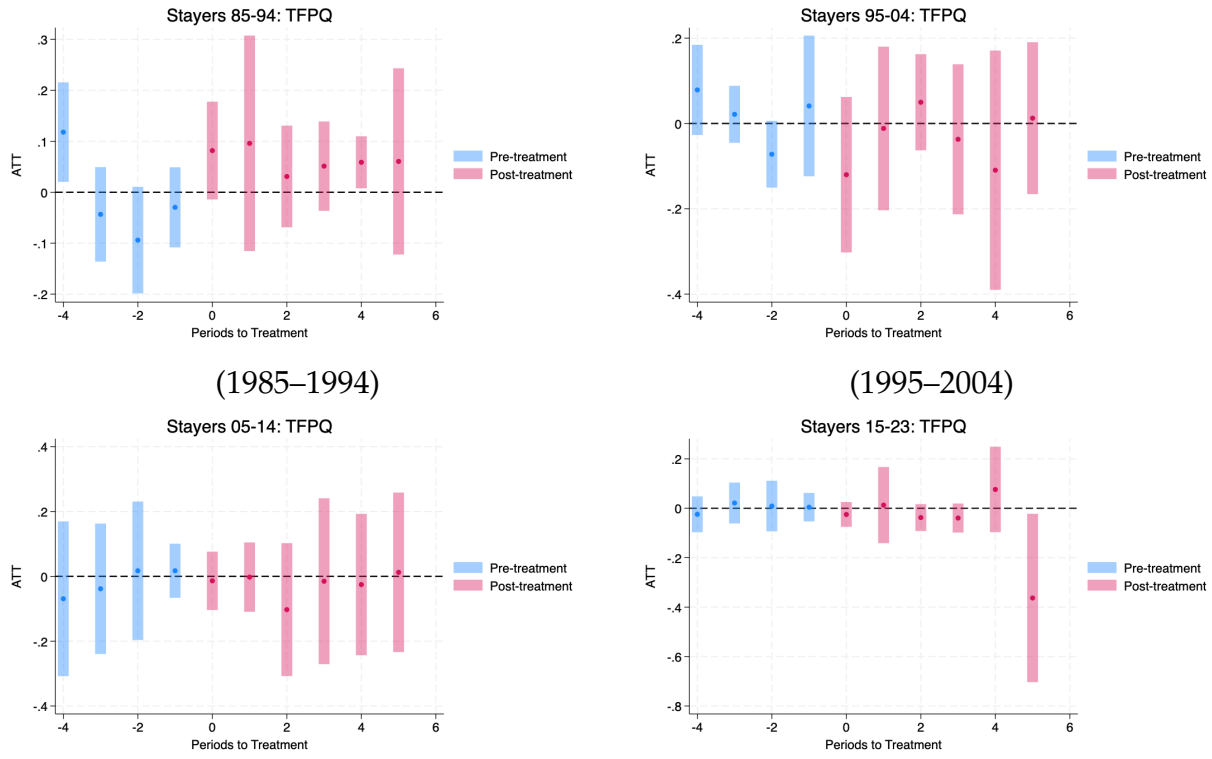
To rigorously isolate the structural productivity and scaling behavior of entrenched incumbents, Section F.2 restricts the sample exclusively to continuous top-tier incumbents ("Stayers"). For full econometric transparency, Section F.3 reports the unadjusted baseline encompassing all firms. Finally, Section F.4 reports the rigorously matched robustness checks to rule out selection bias.

Throughout all sections, the figures are presented in a consistent theoretical order: **Physical Productivity (TFPQ)** and **Organic Sales Growth** (the core drivers of physical scaling), followed by **Total Sales Growth** and **Log Markup** (the core drivers of accounting scale and market power).

Pre-Trend Evaluation: Pre-treatment trends firmly support my identifying assumption across the specifications. The estimates for $t-3$ through $t-1$ are remarkably tight and statistically indistinguishable from zero. This confirms that the parallel trends assumption strictly holds, allowing me to interpret the post-treatment divergence as the causal effect of the M&A events.

F.2 Baseline Dynamics: Top-Tier Incumbents (Stayers)

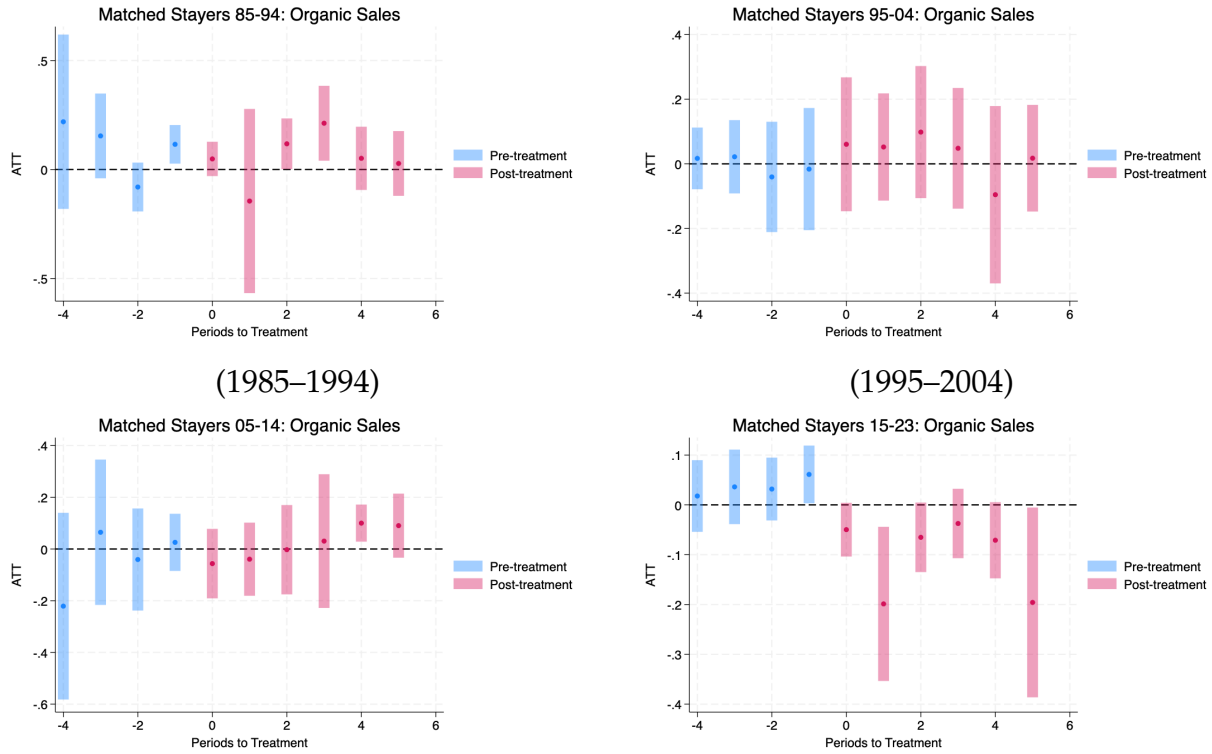
This section provides the visual event-study mapping for the core results discussed in the main text. Note the stark divergence in TFPQ and Organic Growth between Era 1 (synergistic scaling) and Era 4 (defensive rot), while Log Markups remain largely flat post-treatment, confirming M&A is used to protect—rather than generate—market power.



(1985–1994) (1995–2004)

(2005–2014) (2015–2023)

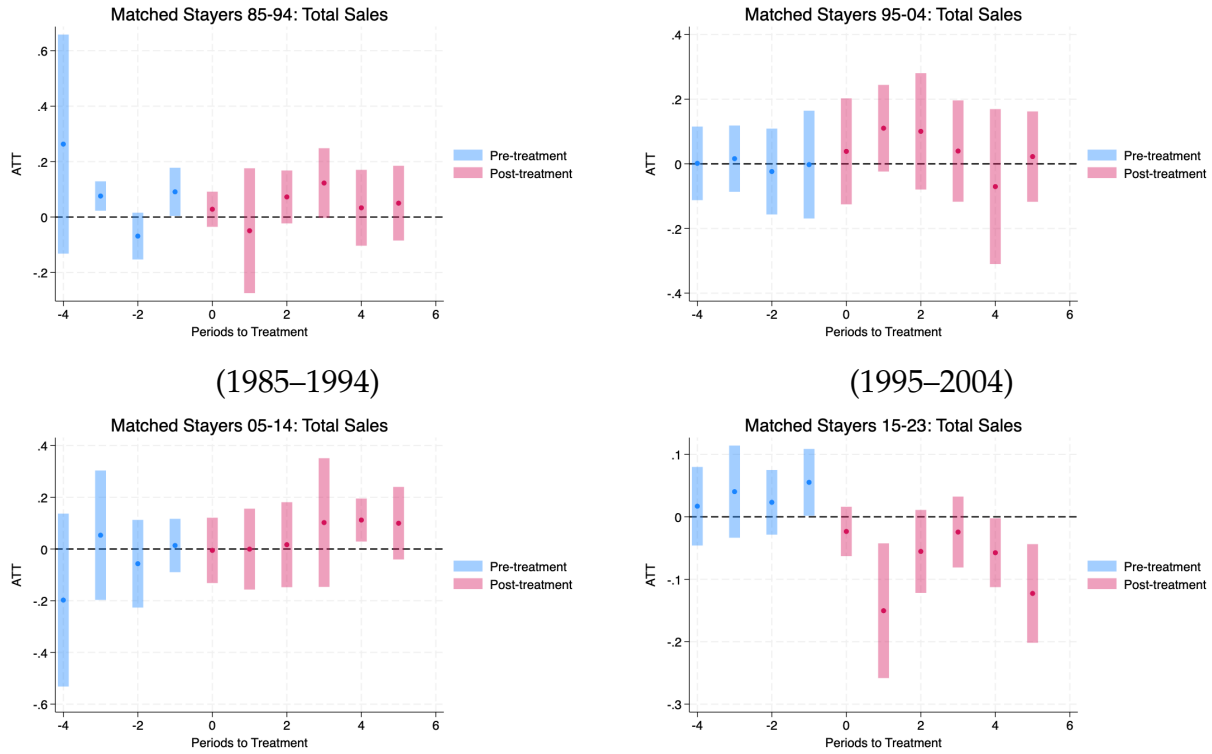
Figure 13: Stayers: Physical Productivity (TFPQ). Highlights the structural shift from synergistic leaps in Era 1 to massive defensive rot in Era 4. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D_{i,t}^k + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).



(1985–1994) (1995–2004)

(2005–2014) (2015–2023)

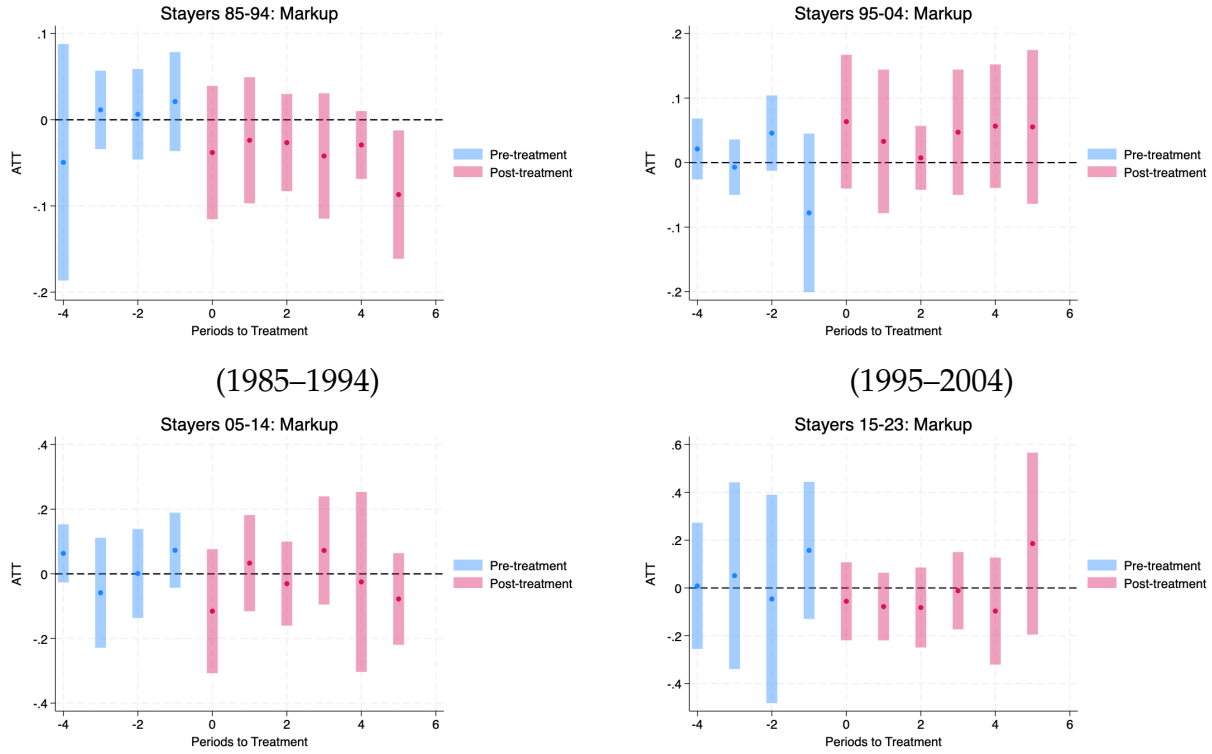
Figure 14: Matched Stayers: Organic Sales Growth. Event study estimates on organic scaling for the propensity score matched sample. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D^k_{i,t} + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).



(1985–1994) (1995–2004)

(2005–2014) (2015–2023)

Figure 15: Matched Stayers: Total Sales Growth. Event study estimates on total scaling for the propensity score matched sample. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D^k_{i,t} + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).



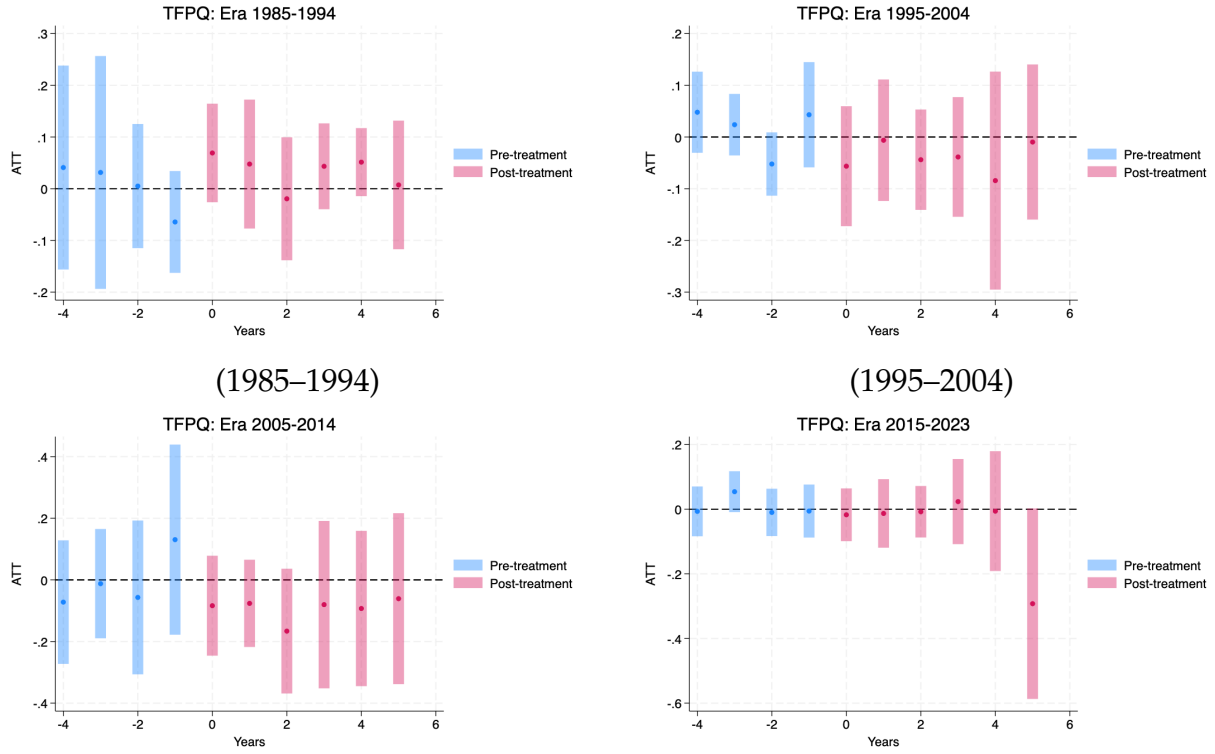
(1985–1994) (1995–2004)

(2005–2014) (2015–2023)

Figure 16: Stayers: Log Markup. Flat post-treatment dynamics confirm that M&A serves to protect existing baseline markups rather than directly generating new pricing power. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D_{i,t}^k + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).

F.3 Unadjusted Baseline Dynamics (All Firms)

This section reports the unadjusted baseline including all firms (Stayers, Inflows, and Outflows). As described in the text, incorporating Inflows introduces mechanical compositional volatility. Notice how Organic Sales Growth mathematically crashes post-treatment in all eras, an artificial base-effect caused when entrants absorb massive M&A targets. This confirms the necessity of restricting my main analysis strictly to Stayers to recover true structural parameters.



(1985–1994) (1995–2004)

(2005–2014) (2015–2023)

Figure 17: All Firms Baseline: Physical Productivity (TFPQ). The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D_{i,t}^k + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).

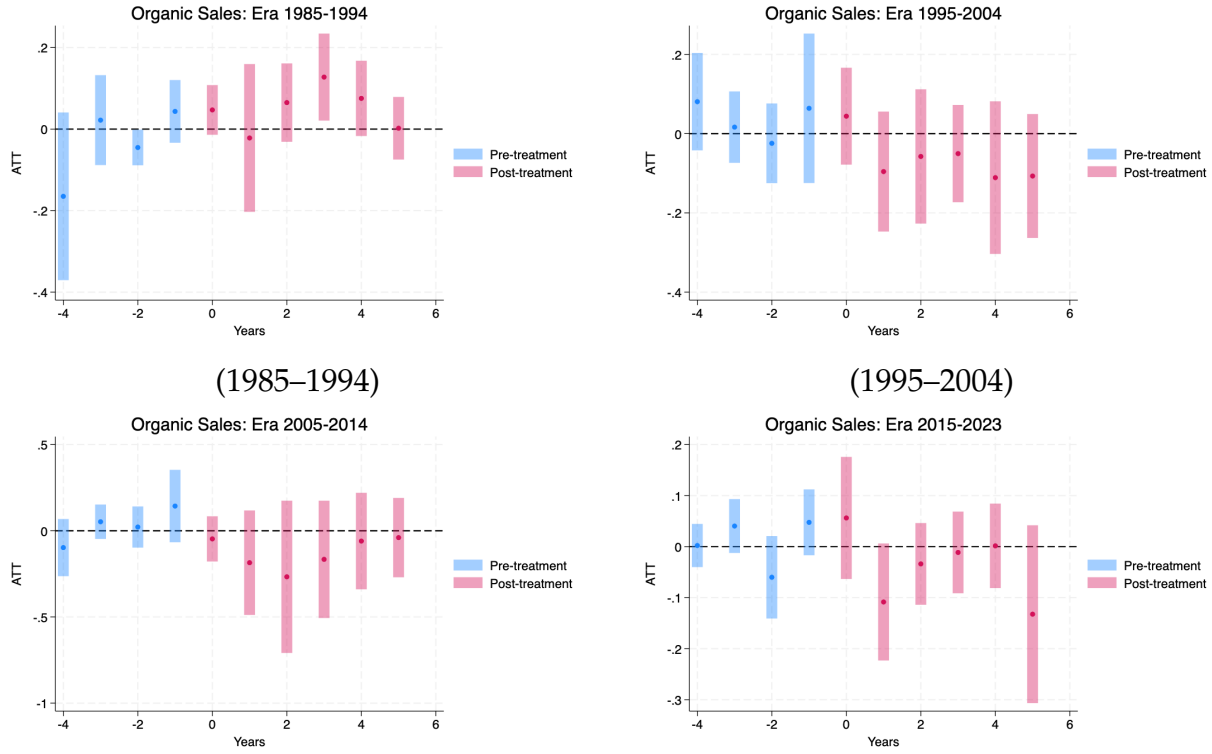
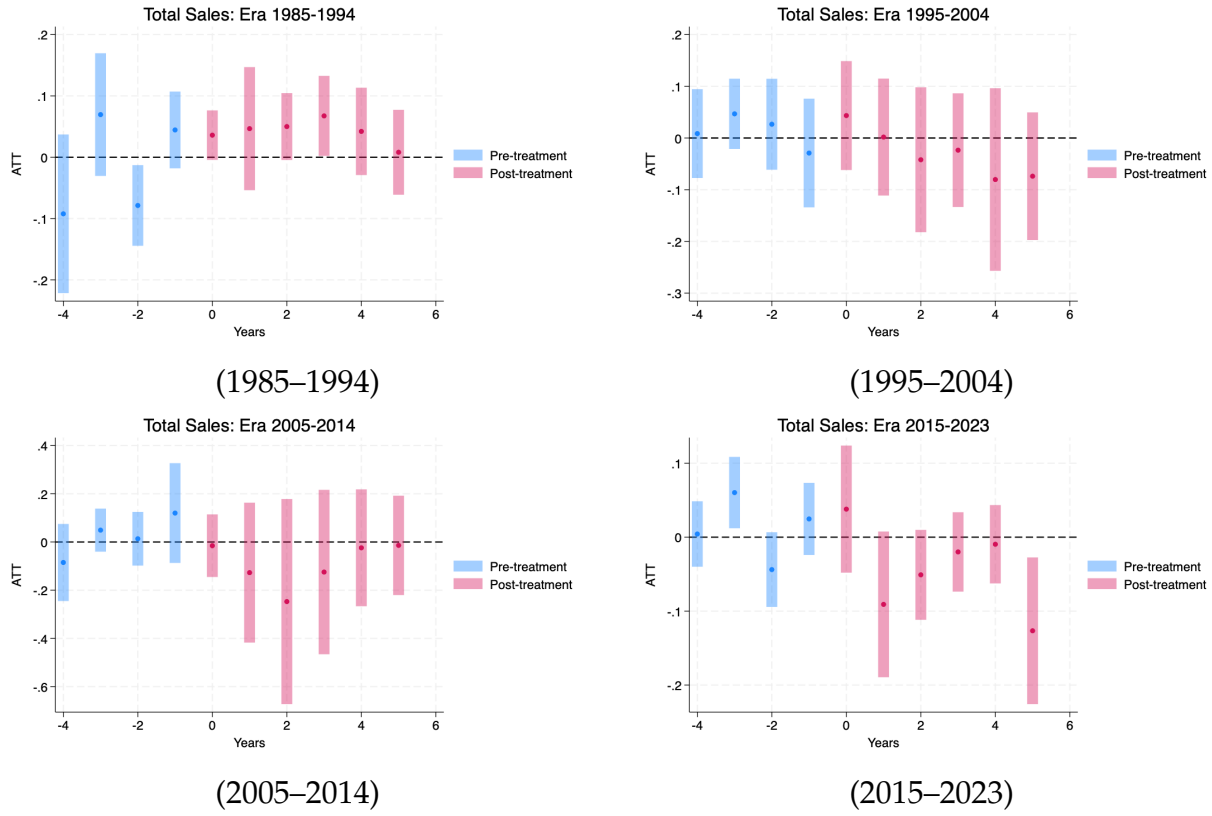


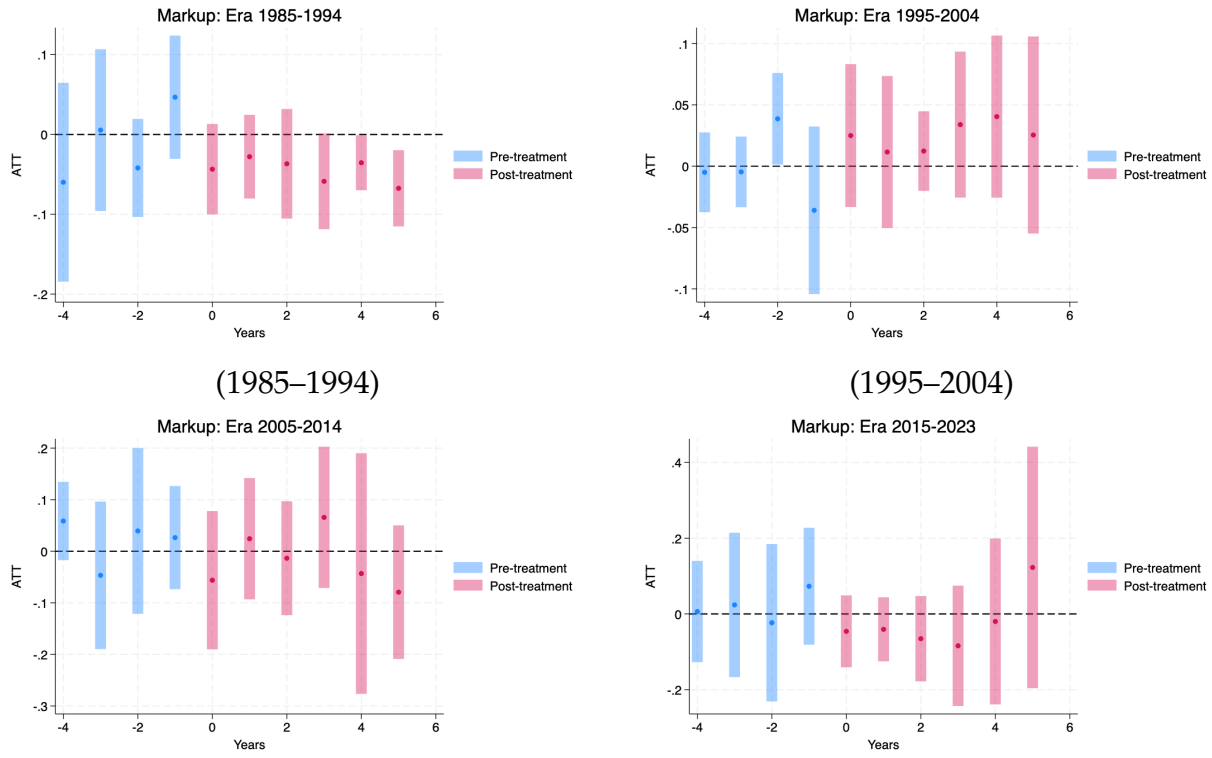
Figure 18: All Firms Baseline: Organic Sales Growth. Highlights the artificial post-treatment crash induced by rapid M&A entrants (Inflows). The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D_{i,t}^k + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).



(1985–1994) (1995–2004)

(2005–2014) (2015–2023)

Figure 19: All Firms Baseline: Total Sales Growth. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D^k_{i,t} + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).



(1985–1994) (1995–2004)

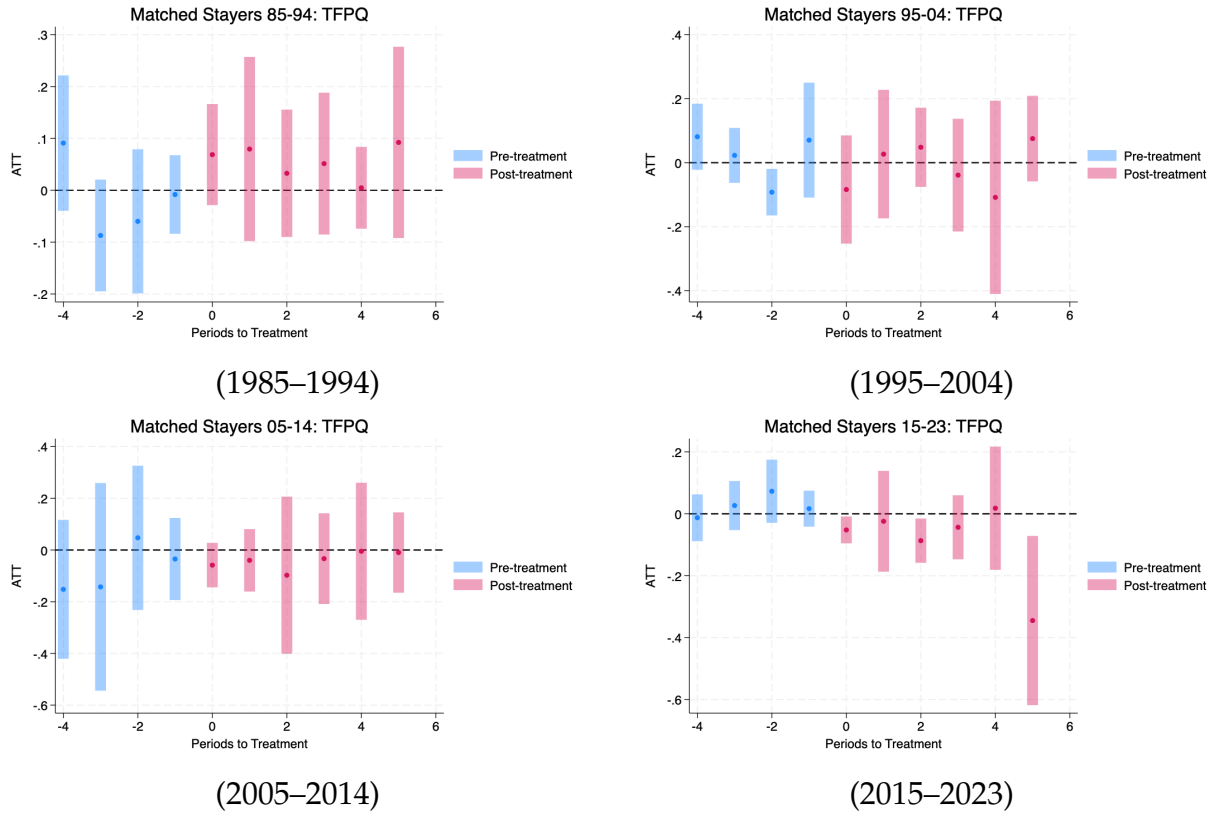
(2005–2014) (2015–2023)

Figure 20: All Firms Baseline: Log Markup. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D_{i,t}^k + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).

F.4 Propensity Score Matched Stayers (Robustness)

To address concerns regarding selection bias—specifically, the hypothesis that firms executing M&A in the modern era might simply be worse-performing firms—this section replicates the event-study design on a strictly restricted, Propensity Score Matched (PSM) sample. Treated Stayers are matched to untreated Stayers based on pre-treatment trajectories.

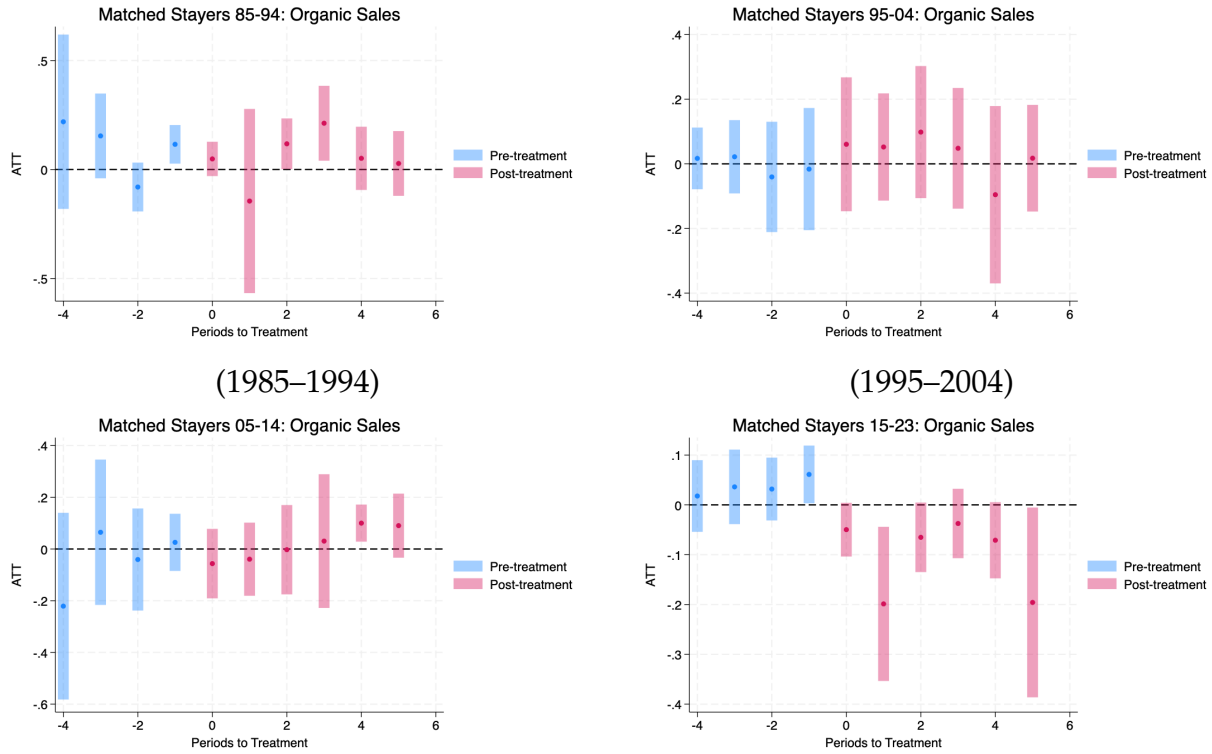
As shown in the figures below, the persistence of the negative TFPQ treatment effect in Era 4 (Figure 21) confirms that the physical rot is a causal penalty of defensive shelving, not an artifact of adverse selection.



(1985–1994) (1995–2004)

(2005–2014) (2015–2023)

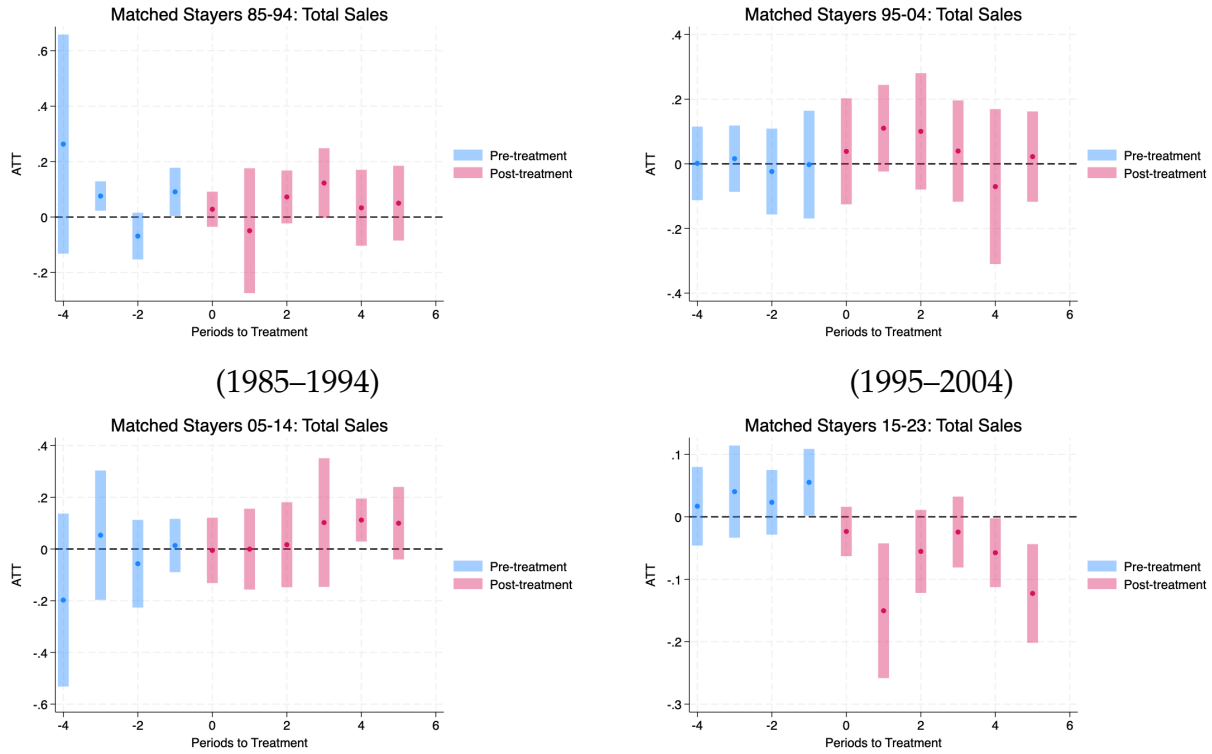
Figure 21: Matched Stayers: Physical Productivity (TFPQ). Pre-trend matched sample confirming the causal shift from synergistic leaps in Era 1 to defensive rot in Era 4. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D^k_{i,t} + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).



(1985–1994) (1995–2004)

(2005–2014) (2015–2023)

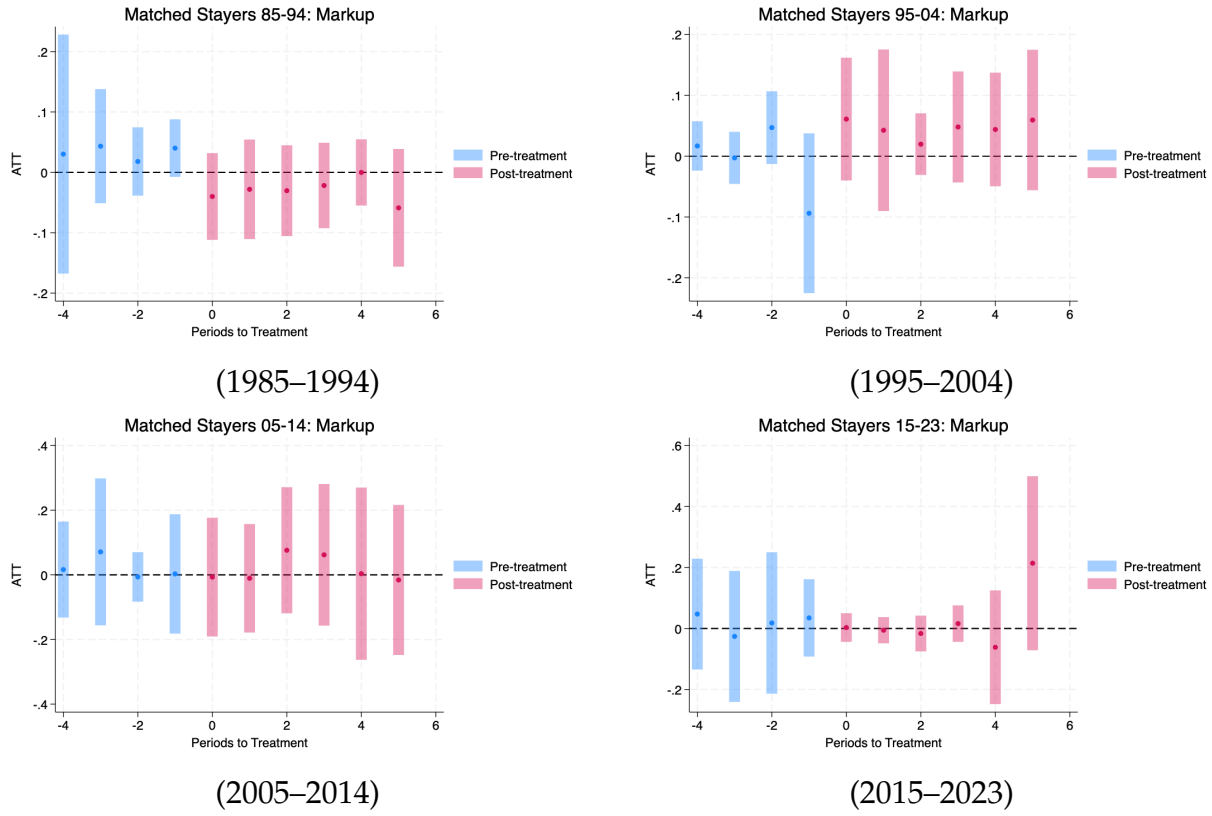
Figure 22: Matched Stayers: Organic Sales Growth. Event study estimates on organic scaling for the propensity score matched sample. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D_{i,t}^k + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).



(1985–1994) (1995–2004)

(2005–2014) (2015–2023)

Figure 23: Matched Stayers: Total Sales Growth. Event study estimates on total scaling for the propensity score matched sample. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D^k_{i,t} + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).



(1985–1994) (1995–2004)

(2005–2014) (2015–2023)

Figure 24: Matched Stayers: Log Markup. Event study estimates on pricing power for the propensity score matched sample. The plotted coefficients represent event-time average treatment effects mapping to the dynamic specification $Y_{i,t} = \alpha_i + \tau_t + \sum_{k \neq -1} \beta_k D^k_{i,t} + \Gamma' X_{i,t-1} + \varepsilon_{i,t}$, balancing pre-treatment covariates $X_{i,t-1}$ (lagged log assets, cash ratio, ROA, and R&D intensity).

G Appendix Tables

G.1 Panel Summary Statistics

Table 5 presents the summary statistics for the core variables used in the doubly robust CS-DiD event-study analysis (Section 3.6). Because volatile entrants (Inflows) and exiting firms (Outflows) systematically violate parallel trends due to mechanical compositional noise, my core causal analysis strictly isolates continuous top-tier incumbents ("Stayers").

Accordingly, this table partitions the sample to reflect the core Difference-in-Differences identification strategy. Panel A reports the characteristics of the Treated Stayers *strictly* in their pre-treatment window ($t-4$ to $t-1$) prior to their peak M&A event. Panel B reports the baseline characteristics of the Untreated Control pool (top-tier firms that never executed a peak M&A event during the sample period).

Crucially, the table demonstrates strong pre-treatment baseline balance. In the years leading up to their massive M&A sprees, Treated Stayers exhibited nearly identical organic sales growth trajectories (4.6% vs 4.7%) compared to the control pool, while simultaneously boasting higher profitability (ROA) and stronger baseline markups. This fundamentally rules out adverse selection—the treated firms were not underperforming or failing prior to their acquisitions.

G.2 Event-Study Robustness and Compositional Decompositions

A known econometric challenge with estimating the dynamic effects of entering the absolute apex of the firm size distribution is extreme compositional volatility. Firms that briefly spike into the Top 0.01% via a massive one-off acquisition ("Inflows") experience mechanically artificial drops in subsequent organic growth due to mathematical base-effects.

Table 6 proves this empirically. When evaluating the pooled 40-year baseline, the "All Firms" average suggests M&A destroys -5.0% of organic growth. However, this is entirely driven by the severe -17.2% mechanical drop in the volatile "Inflow" group. To ensure our causal claims regarding structural regime shifts are not driven by this compositional noise, the main text rigorously restricts analysis to continuous incumbents ("Stayers").

Furthermore, the markup results in this table show that M&A has a near-zero average treatment effect on pricing power across all subgroups. This reinforces our theoretical assumption: modern M&A is a defensive tool used to safeguard existing baseline markups,

Table 5: Summary Statistics for the Event-Study Panel

	N	Mean	p10	Median	p90
Panel A: Treated Stayers (Pre-Treatment Window: -4*to* − 1)					
Deflated Total Sales	244	19013.716	5833.267	15910.295	32909.114
Total Assets	244	27592.072	3036.852	18228.795	56053.472
Cash Ratio	244	0.094	0.008	0.054	0.213
Return on Assets (ROA)	244	0.122	0.023	0.122	0.211
R&D Intensity	244	0.016	0.000	0.000	0.051
Organic Sales Growth (Log)	239	0.046	-0.092	0.042	0.188
Total Sales Growth (Log)	239	0.046	-0.083	0.035	0.170
Log Physical Productivity (TFPQ)	240	3.794	2.742	3.584	5.210
Log Markups	244	0.508	0.114	0.330	1.128
Panel B: Untreated Control Pool (Top 0.01% Firms)					
Deflated Total Sales	16,445	14302.853	5661.141	10573.596	32909.114
Total Assets	16,445	16702.015	1.000	9294.168	56053.472
Cash Ratio	16,445	0.071	0.000	0.033	0.191
Return on Assets (ROA)	16,445	0.110	0.000	0.105	0.220
R&D Intensity	16,445	0.012	0.000	0.000	0.031
Organic Sales Growth (Log)	15,712	0.047	-0.128	0.034	0.240
Total Sales Growth (Log)	15,712	0.051	-0.114	0.034	0.231
Log Physical Productivity (TFPQ)	13,482	3.885	2.895	3.756	5.156
Log Markups	14,516	0.445	0.096	0.334	0.917

Notes: This table reports the observation count (N), mean, 10th percentile (p10), median, and 90th percentile (p90) for the primary covariates and outcome variables evaluated in the CS-DiD specification. To demonstrate pre-treatment comparability and baseline balance, Panel A evaluates Treated Stayers in the pre-treatment window ($t-4$ to $t-1$) prior to their peak M&A event. Panel B captures the baseline characteristics of the Untreated Control pool (continuous top-tier firms that never crossed the peak M&A threshold). Deflated Total Sales and Total Assets are expressed in millions of 2017 U.S. dollars. Continuous ratios and growth variables are winsorized at the 1st and 99th percentiles to minimize the influence of extreme accounting anomalies.

rather than a mechanism to synthetically generate new markups.

Table 6: Average Post-Treatment Effects ($t \in [2, 5]$): Baseline & Trajectories

Outcome	All Firms (Baseline)	Stayers	Inflows	Outflows
TFPQ	-0.048	-0.040	-0.009	0.009
	(0.030)	(0.038)	(0.054)	(0.021)
Log Markup	-0.004	0.007	-0.052	0.060
	(0.022)	(0.033)	(0.044)	(0.039)
Organic Sales Growth	-0.050	0.010	-0.172**	-0.057
	(0.036)	(0.043)	(0.087)	(0.084)
Total Sales Growth	-0.011	0.015	-0.156*	-0.071
	(0.030)	(0.037)	(0.086)	(0.044)

Notes: This table decomposes the pooled 40-year CS-DiD event-study baseline. It isolates the average treatment effect on the treated (ATT) for periods $t \in [2, 5]$ post-treatment. By decomposing the baseline into entrenched incumbents (Stayers) versus volatile entrants (Inflows), we eliminate compositional noise. Standard errors are in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

G.3 Propensity Score Matched Robustness

A primary empirical challenge when evaluating the treatment effect of M&A on incumbent productivity is selection bias. If the firms executing defensive M&A in the modern era (Era 4) are fundamentally different—or on systematically worse trajectories—than those that do not, the observed -9.1% TFPQ drop might reflect adverse selection rather than the causal rot of defensive shelving.

To strictly isolate the causal treatment effect and eliminate observable selection bias, we execute a Propensity Score Matching (PSM) robustness check. We restrict the control group exclusively to the top quintile of the distribution and match treated Stayers to untreated Stayers based on their pre-treatment size, baseline markup, and pre-treatment TFPQ trajectory.

Table 7 reports the doubly robust CS-DiD estimates estimated strictly on this matched sample. The results confirm our baseline findings. In Era 4, the matched treatment effect on TFPQ remains a statistically significant structural penalty. This confirms that the severe physical rot documented in the modern era is the direct causal consequence of integrating mismatched targets to protect legacy markups, not a selection artifact.

Table 7: Matched Stayers Average Post-Treatment Effects ($t \in [2, 5]$): Era-by-Era Shift

Outcome	1985–1994	1995–2004	2005–2014	2015–2023
TFPQ	0.045	-0.006	-0.036	-0.114***
	(0.035)	(0.085)	(0.066)	(0.034)
Log Markup	-0.028	0.043	0.031	0.038
	(0.026)	(0.042)	(0.108)	(0.028)
Organic Sales Growth	0.102*	0.017	0.055	-0.092**
	(0.056)	(0.101)	(0.041)	(0.042)
Total Sales Growth	0.070	0.023	0.083*	-0.065***
	(0.053)	(0.087)	(0.046)	(0.022)

Notes: This table reports the average post-treatment effects from the Callaway and Sant'Anna (2021) doubly robust estimator, restricted to a Propensity Score Matched sample. Treated Stayers are matched to untreated Stayers based on pre-treatment observables and trajectories to strictly eliminate selection bias. Standard errors are clustered at the firm level. *, **, and *** denote significance at the 10%, 5%, and 1% levels, respectively.

G.4 Compustat Micro-Audit Dictionary

Because the extreme right tail of the firm size distribution dictates the aggregate macroeconomic indices, standard algorithmic cleaning of Compustat data is insufficient. A single false entry, unrecorded M&A consolidation, or entity identifier shift among the top 20 firms can generate massive, artificial swings in aggregate productivity and concentration metrics.

Table 8 details the master dictionary of our manual, firm-by-firm micro-audit. We manually investigated the SEC filings (10-Ks), historical press releases, and M&A databases for the largest firm-year anomalies in the dataset. This table documents the specific structural artifacts—such as post-merger accounting delays, internal tracking stock collapses, and unrecorded subsidiary spin-offs—that were manually corrected to ensure our Melitz–Polanec decomposition reflects genuine economic activity rather than data architecture flaws.

Table 8: Expert Mapping for Top $D^{\neg MA}$ Contributors

Entity ID	Company Name	Corrected Set	Historical Fact / Reason
$\boldsymbol{D^{OG_true}}$			
GVKEY_143972	WorldCom	D^{OG_true}	Filed for Chapter 11 bankruptcy in 2002.
GVKEY_030128	Lehman Brothers	D^{OG_true}	Filed for Chapter 11 bankruptcy in 2008.
GVKEY_016243	Washington Mutual	D^{OG_true}	Failed and seized by the FDIC in 2008.
GVKEY_001045	AMR Corp (American Airlines)	D^{OG_true}	Filed for Chapter 11 bankruptcy in 2011.
GVKEY_006307	Sears Holdings	D^{OG_true}	Filed for Chapter 11 bankruptcy in 2018.
GVKEY_009155	Rite Aid	D^{OG_true}	Filed for Chapter 11 bankruptcy in 2023.
GVKEY_004769	Fleming Companies	D^{OG_true}	Filed for bankruptcy in 2003 and liquidated.
GVKEY_004584	Farmland Industries	D^{OG_true}	Filed for bankruptcy and liquidated in 2002.

Continued on next page

Table 8 – continued from previous page

Entity ID	Company Name	Corrected Set	Historical Fact / Reason
GVKEY_005301	A&P (Great Atlantic & Pacific Tea)	D^{OG_true}	Filed for Chapter 11 bankruptcy in 2010.
GVKEY_007021	Marcor (Montgomery Ward)	D^{OG_true}	Parent of Montgomery Ward; failed in 1994.
GVKEY_004183	Eastern Air Lines	D^{OG_true}	Liquidated in 1991.
GVKEY_003452	Continental Airlines (Old)	D^{OG_true}	Filed for Chapter 11 in 1990.
GVKEY_001301	Allied Stores	D^{OG_true}	Filed for bankruptcy in 1990.
GVKEY_008308	Pan Am Corp	D^{OG_true}	Liquidated in 1991.
GVKEY_012391	Pan American World Airways	D^{OG_true}	Liquidated in 1991.
$\boldsymbol{D^{MA_missed}}$			
GVKEY_014489	Dell Inc Michael Dell in 2013.	D^{MA_missed}	LBO by Silver Lake
GVKEY_022260	HCA Healthcare KKR in 2006.	D^{MA_missed}	LBO by Bain Capital
GVKEY_139001	Georgia-Pacific	D^{MA_missed}	Taken private by Koch Industries in 2005.
CUSIP_UP_53656T	Tech Data Corp	D^{MA_missed}	LBO by Apollo Global in 2020.
CUSIP_UP_42907P	AT T Corp	D^{MA_missed}	Acquired by SBC Communications in 2005.
CUSIP_UP_07557L	Amoco	D^{MA_missed}	Merged with BP in 1998.
CUSIP_UP_360271	Bank One	D^{MA_missed}	Acquired by JPMorgan Chase in 2004.
CUSIP_UP_55275Z	Wachovia	D^{MA_missed}	Acquired by Wells Fargo in 2008.
CUSIP_UP_68153T	21st Century Fox	D^{MA_missed}	Acquired by Disney in 2019.
CUSIP_UP_31998Z	Raytheon	D^{MA_missed}	Merged with United Technologies in 2020.

Continued on next page

Table 8 – continued from previous page

Entity ID	Company Name	Corrected Set	Historical Fact / Reason
GVKEY_003066	Citicorp	D^{MA_missed}	Merged with Travelers in 1998.
GVKEY_007267	Merrill Lynch	D^{MA_missed}	Acquired by Bank of America in 2009.
GVKEY_004060	DuPont	D^{MA_missed}	Merged with Dow Chemical in 2017.
GVKEY_014456	BP America	D^{MA_missed}	Merged with Amoco in 1998.
$\boldsymbol{D^{Artifact}}$			
GVKEY_005048	GE Capital Corp	$D^{Artifact}$	Massive internal de-designation; folded into parent.
CUSIP_UP_55977Q	IBM	$D^{Artifact}$	CUSIP/Reporting Entity change; firm remains active.
CUSIP_UP_86652W	Bank of America	$D^{Artifact}$	Internal entity collapse; firm remains active.
CUSIP_UP_50270Q	JPMorgan Chase	$D^{Artifact}$	Internal entity collapse; firm remains active.
GVKEY_121633	Northwestern Mutual	$D^{Artifact}$	Mutual life insurance coverage/format drop.
GVKEY_127774	New York Life	$D^{Artifact}$	Mutual life insurance coverage/format drop.
GVKEY_020942	NBCUniversal Media	$D^{Artifact}$	Subsidiary of Comcast; identifier drop.
GVKEY_004840	Ford Motor Credit	$D^{Artifact}$	Subsidiary shift; parent remains active.
GVKEY_025550	Kraft Foods Group	$D^{Artifact}$	Spinoff reorganization artifact (Mondelez).
GVKEY_007241	CVS Health	$D^{Artifact}$	Format change drop; firm remains highly active.
GVKEY_001722	Archer-Daniels-Midland	$D^{Artifact}$	Format change drop; firm remains highly active.

Continued on next page

Table 8 – continued from previous page

Entity ID	Company Name	Corrected Set	Historical Fact / Reason
GVKEY_002176	Berkshire Hathaway	$D^{Artifact}$	Format drop to Financials; highly active.
GVKEY_125134	CITGO	$D^{Artifact}$	Foreign ownership (PDVSA) SEC coverage drop.
GVKEY_011456	Weyerhaeuser	$D^{Artifact}$	Reclassified as REIT; format drop.
GVKEY_001356	Alcoa Inc	$D^{Artifact}$	Split into Alcoa Corp and Arconic; identifier retired.