UNIVERSITY OF CAMBRIDGE
Faculty of Economics

JANEWAY INSTITUTE CAMBRIDGE

CAMBRIDGE WORKING PAPERS IN ECONOMICS
JANEWAY INSTITUTE WORKING PAPERS

The Granular Role of Mergers and Acquisitions in Aggregate Fluctuations

Seda Basihos
University of Cambridge

Abstract

This study investigates the role of firm-specific productivity shocks from mergers and acquisitions (M&A) in explaining aggregate fluctuations. I isolate these shocks by estimating a covariate-adjusted latent factor model on U.S. public firm data (1978–2023) under an event study approach. M&A activity generates significant short-run fluctuations in firm productivity, typically characterized by an initial decline, a rebound to a higher level, and eventual stabilization. A firm-size-weighted average of these shocks explains roughly one-fifth of variations in U.S. economic growth. This result is robust across multiple empirical validations and is not driven by merger waves or prevailing macroeconomic conditions.

Reference Details

2560	Cambridge Working Papers in Economics
2523	Janeway Institute Working Paper Series
Published	18 August 2025
Key Words	Aggregate fluctuations, granularity, mergers and acquisitions, productivity
JEL Codes	D22, E32, G34
Websites	https://www.econ.cam.ac.uk/publications/cwpe www.janeway.econ.cam.ac.uk/working-papers

The Granular Role of Mergers and Acquisitions in Aggregate Fluctuations

Seda Basihos*

18 August 2025

Abstract

This study investigates the role of firm-specific productivity shocks from mergers and acquisitions (M&A) in explaining aggregate fluctuations. I isolate these shocks by estimating a covariate-adjusted latent factor model on U.S. public firm data (1978–2023) under an event study approach. M&A activity generates significant short-run fluctuations in firm productivity, typically characterized by an initial decline, a rebound to a higher level, and eventual stabilization. A firm-size-weighted average of these shocks explains roughly one-fifth of variations in U.S. economic growth. This result is robust across multiple empirical validations and is not driven by merger waves or prevailing macroeconomic conditions.

Keywords: Aggregate fluctuations, granularity, mergers and acquisitions, productivity
JEL Classification: D22, E32, G34

I would like to thank M. Hashem Pesaran and Eric French for their invaluable feedback. All errors are mine.
* University of Cambridge, Faculty of Economics. Austin Robinson Building, Room 92, Cambridge UK, CB3 9DD. E-mail: sb2626@cam.ac.uk. Website: www.sedabasihos.com

1. Introduction

The seminal work of Gabaix (2011) establishes that firm-specific productivity shocks to large firms explain a considerable share of aggregate fluctuations. In deriving this result, Gabaix excludes extreme productivity movements often associated with major corporate events, such as mergers and acquisitions (M&A). If these M&A-driven shocks are as substantial as suggested, they may represent a meaningful yet unexplored source of aggregate fluctuations. This observation motivates the current study.

Empirical evidence from sector-specific studies documents significant productivity fluctuations for acquirers following M&A—typically an initial decline followed by a recovery (Schoar, 2002; Braguinsky et al., 2015). This J-curve pattern likely reflects adjustment dynamics after initial integration frictions (J. Xu, 2017). However, the evidence remains confined to individual industries, leaving the aggregate relevance of such movements an open question.

Anecdotal accounts similarly suggest that M&A is rarely smooth sailing for firms. For example, AT&T's recent $85 billion acquisition of Time Warner reportedly faced severe post-merger challenges, with some observers noting it "led them back to square one" (Financial Times, 2023). Likewise, IBM's $34 billion acquisition of Red Hat revealed cultural and operational tensions, described by executives as "two cultures working together, not coming together," even as the deal ultimately advanced IBM's hybrid cloud strategy (WSJ, 2019).

Consequently, given the granular structure of the U.S. economy, where large firms play a significant role and are predominantly acquirers, it becomes a policy-relevant question whether such productivity shocks are pervasive among large U.S. firms and, if so, to what extent aggregate fluctuations can be traced back to them.

This paper addresses this question by isolating M&A-driven productivity shocks across a broad sample of publicly traded U.S. firms, documenting that they explain a non-trivial share of variations in aggregate growth. Importantly, the current work extends the scope of the granular hypothesis beyond the purely random shock assumption by showing that firm-level strategic actions could be an important source of macroeconomic fluctuations.

The empirical analysis uses M&A transaction data from SDC Platinum (LSEG/Refinitiv), linked to U.S. public firms from Compustat. The sample includes completed U.S. domestic deals, to ensure consistency in regulatory and economic environments. Following Gabaix (2011), I measure firm productivity as the log of real sales per employee.

To estimate firm-specific productivity changes (1978–2023), I conduct a staggered event study in which firm-year observations are organized into rolling five-year cohorts centered on a focal merger year (Year 0), spanning two pre-merger years (Years -2, -1) and two post-merger years (Years +1, +2). Firms completing at least one M&A deal as acquirer or surviving entity in Year 0 with no other M&A within the cohort window are classified as treated. Firms with no M&A throughout the window are classified as controls. The outcome variable is the change in log productivity relative to baseline pre-merger year (Year -1).

Unlike the random shock profiles typically assumed in granular studies, however, M&A decisions are endogenous to firm characteristics and economy-wide conditions. I therefore do not assume random assignment. Following the potential outcomes framework (Rubin, 1974; Holland, 1986), the causal effect of M&A on productivity is defined as the difference between observed outcomes and counterfactual outcomes absent M&A.

Counterfactual outcomes are predicted using a three-step latent-factor model estimated on control firms' data. Treated-firm observations are excluded from estimation; for these firms, no-M&A outcomes at each event time are considered as unobserved and predicted out-of-sample using the control-based model. The approach models productivity changes (absent M&A) as a function of firm characteristics measured in the pre-event years and latent factors capturing common time-varying influences (e.g., macroeconomic conditions). Estimation is conducted separately for each event-time year (-2, 0, +1, +2), pooling data across cohorts to exploit cross-firm and calendar-time variations while maintaining event-time alignment.

First, productivity changes are regressed on firm covariates to capture systematic variation attributable to observables. Second, factor analysis is applied to residuals from the first step to extract common latent factors. Third, a final augmented regression is estimated that includes both the covariates and the extracted latent components, and fitted values from this model are then used to construct acquirers' counterfactual paths.

Identification relies on the assumption that M&A event is as-good-as random conditional on observed firm characteristics and latent factors that absorb aggregate conditions. Under this assumption, the difference between observed and predicted productivity in the merger year and in each subsequent post-merger year yields dynamic, firm-level estimates of productivity changes attributable to M&A.

The results reveal a stylized short-run dynamics consistent with earlier findings, where an initial productivity decline occurs during the M&A year, followed by a recovery with strong gains that gradually taper off. Within two years, productivity exceeds pre-merger baselines and converges to a higher plateau. The average decline is 4.9% points, with a subsequent around 8% points increase from the dip. While level effects persist, growth effects concentrate primarily in the immediate aftermath. Further results indicate that these fluctuations reflect integration-related capacity build-out: employment rises at merger closing and remains broadly stable, while sales ramp with a lag and ultimately outpace the increase in employment. A placebo test on the model identification confirms no spurious results.

Moreover, there is significant heterogeneity in productivity fluctuations across firm size categories. Large acquirers, for instance, experience milder initial declines in productivity relative to other firms in the sample, while showing strong subsequent recoveries. This suggests that pre-merger firm size moderates post-merger adjustment dynamics and that idiosyncratic M&A shocks to large firms may leave some granularity in aggregate outcomes.

Accordingly, I construct an annual granular measure as the size-weighted average of firm-level M&A shocks, using lagged firm market shares in GDP as weights. Shocks are derived from estimated cumulative productivity changes during the merger year and the following two years (peak adjustment period). Regressing GDP per capita growth on this measure shows that granular M&A shocks explain around one-fifth of aggregate fluctuations.

This explanatory power is largely robust against a set of statistical checks. These include: controlling for post-merger changes in firm size to rule out size-driven mechanical effects; injecting artificial noise into the estimated shocks to assess the stability of the granular measure's explanatory power; conducting placebo tests by randomly reassigning shocks across firms to evaluate the risk of spurious correlation with aggregate growth; and controlling for standard macroeconomic variables such as monetary policy changes, energy price shocks, and indicators of merger waves to address potential overlap with broader cyclical forces.

The results confirm that the granularity is not an artifact of specification, time-varying exposure, or omitted macro factors. Moreover, our granular measure is not predictable from its past values, yet it exhibits meaningful predictive power for near-term economic growth.

I organize the remainder of the paper as follows. Section 2 situates the study within the existing literature and outlines its contributions. Section 3 describes the data sources and structure. Section 4 details the research design, estimation strategy, and presents the results from the firm-level analysis. Section 5 reports the findings from the granularity analysis, along with an extensive set of robustness checks. Section 6 concludes.

2. Relation to Literature

The paper offers novel empirical support for the micro origins of macro movements in the context of M&A activity. Accordingly, it draws on and contributes to three main bodies of research: the granular hypothesis, the mergers and acquisitions literature, and the literature on estimation of individual treatment effects.

The main contribution of the paper is to the granularity literature (see, among others, Gabaix, 2011; Di Giovanni, Levchenko, and Mejean, 2014; Carvalho and Grassi, 2019; Gabaix and Koijen, 2024). While existing studies in this area mainly examine *random* idiosyncratic shocks, this paper analyzes shocks arising from deliberate, strategic firm actions—such as mergers and acquisitions. This extends the granular hypothesis by incorporating endogenous firm behavior with heterogeneous outcomes as a source of aggregate fluctuations. Perhaps the most related study is Chan and Qi (2025), which argues that M&A may amplify granular effects by increasing market concentration. Although that study addresses a different set of questions, I briefly discuss its insights on firm size changes in the analysis.

The M&A literature is long-standing and extensive, and this study builds on prior work examining the determinants of M&A activity (see, among others, Gort, 1969; Ravenscraft and Scherer, 1989; Lang, Stulz, and Walkling, 1991; Mitchell and Mulherin, 1996; Powell, 1997; Rhodes-Kropf and Viswanathan, 2004; Harford, 2005; Bena and Li, 2014; Becher, Jensen, and Liu, 2020). This line of research emphasizes that both firm-level characteristics and existing aggregate conditions play an important role in shaping acquisition decisions.

Efficiency outcomes of M&A have also been widely studied, though the literature varies in the metrics used for productivity. One strand uses financial indicators such as profitability or return on assets, yet findings in this area are often mixed (see, e.g., Bradley, Desai, and Kim, 1988; Healy, Palepu, and Ruback, 1992; Ghosh, 2001). Studies using operational or revenue-based productivity measures tend to yield more consistent results, mostly documenting post-merger productivity improvements relative to standalone firms or industry peers (see, e.g., Maksimovic and Phillips, 2001; Schoar, 2002; Braguinsky et al., 2015; Demirer and Karaduman, 2024).[1] While this study is not the first to explore productivity implications of M&A, it departs from earlier work in several important respects.

First, much of the prior research focuses on sector-specific firm data, whereas this study draws on data covering public firms across the economy. Second, these studies distinguish between acquirers and targets, analyzing their outcomes separately. This study does not adopt such a distinction; instead, it tracks firms with pre- and post-merger data. Given the focus on publicly listed firms, mostly act as acquirers, the analysis can be broadly interpreted from the acquirer's perspective, with post-merger financials consolidating those of the targets.

Among existing studies, the ones most closely related to the present work in terms of findings on productivity patterns around M&A are Schoar (2002) and Braguinsky et al. (2015). These studies document a temporary decline in productivity for acquirers around the merger event, alongside an improvement for targets that, on its own, is insufficient to offset the decline. However, in the post-merger period, the combined entity sees clear productivity

1. However, for instance, Blonigen and Pierce (2016) does not find significant post-merger effects.

gains, suggesting that integration eventually yields net efficiency improvements. The firm-level evidence in this study, based on U.S. public firms, is in line with these findings.

Third, an important stream of the literature studies long-run efficiency gains from M&A activity in general, highlighting value creation channels such as more efficient capital allocation or access to new technology (see, e.g., Jovanovic and Rousseau, 2002; Jovanovic and Rousseau, 2008; see also J. Xu, 2017). Although this study relies on long panel data, it focuses on short-run firm-level movements and how these map into aggregate variations.

Finally, most studies on M&A and productivity estimate an average post-merger effect using post-merger regressions or standard difference-in-differences designs. This paper instead identifies firm-level effects under potential outcomes using a latent factor model. This class of models has gained traction as flexible, data-generating tools (see, e.g., Bai, 2009; Hsiao, Steve Ching, and Ki Wan, 2012; Gobillon and Magnac, 2016; Y. Xu, 2017; Hsiao, 2025) and is especially convenient for estimating unit-level effects in settings with nonrandom treatment, staggered timing, and relatively few treated units (Samartsidis et al., 2019; Samartsidis et al., 2020). By applying these advances, the paper also makes a methodological contribution.

3. Data

The study uses M&A transaction data from LSEG/Refinitiv SDC Platinum (formerly Thomson Reuters). The sample is restricted to publicly traded U.S. acquirers involved in completed domestic deals announced between 1978 and 2023 with a reported transaction value of at least 10 million USD. I focus on domestic transactions to ensure a more controlled setting with common regulatory, tax, and labor market environments.

The M&A dataset (SDC Platinum) is linked to Compustat using CUSIP identifiers.[2] The final sample consists of publicly traded U.S. firms with available financial data from 1978 to 2023. Treated firms are acquirers (or post-merger surviving entities, regardless of name changes) that complete at least one M&A in a given year; control firms have no M&A activity in that year. Both groups satisfy the cohort-level restrictions described in Section 3.2.

3.1. Main Variables and Intermediate Dataset

The SDC data provide transaction completion dates, which I use to identify each firm's M&A event year. Since the original data are reported at a quarterly frequency, I aggregate them

2. Post-merger firms often have different current vs. historical CUSIPs. While Compustat backfills historical CUSIPs using current identifiers, SDC Platinum keeps deal-time identifiers. To recover such firms, I use historical CUSIP-to-gvkey mapping tables (via deal number) from Phillips and Zhdanov (2013).

to an annual level.[3] Accordingly, M&A activity is defined as a binary indicator equal to 1 if a firm acts as an acquirer at least once in a given year, and 0 otherwise.[4]

Before proceeding, a few clarifications are in order. First, throughout the paper, *acquirer* (or *acquiring firm*) and *treated firm*, as well as *control firm* and *untreated firm*, are used interchangeably. *Merger* and *M&A* are used synonymously. Second, although M&A deals vary in size (e.g., asset value) and type (e.g., horizontal vs. vertical), this study adopts a binary treatment indicator for tractability. Finally, all post-merger outcomes are evaluated using data from acquiring firms, since, as the surviving legal and operational entity, any observable post-merger productivity changes are recorded in their financial statements.

Firm-level productivity is measured as the log of sales revenue per employee, with sales revenue adjusted to 2017 prices. I use this measure as a proxy for productivity throughout the study because it relies solely on observed data and does not require the additional identifying assumptions required for alternative measures, such as the estimation of TFP. Note that at the aggregate level, I relate the size-weighted average of firm productivity shocks to GDP per capita growth—the dependent variable used in the granularity analysis.

Given that one objective of this study is to estimate the causal effects of M&A on firm-level productivity, the analysis also considers factors that may influence both a firm's decision to engage in M&A and its productivity trajectory post-merger. Thus, I use firm-level measures from Compustat as proxies for pre-merger firm characteristics. All variables are observed prior to the merger event, ensuring they are unaffected by the treatment outcome.

The first key measure is firm size, defined as a firm's sales-based market share in GDP. This variable is included based on strong empirical evidence that larger firms are more likely to engage in M&A activity (see, e.g., Powell, 1997; Cudd and Duggal, 2000; Maksimovic and Phillips, 2001). In addition, the analysis includes a set of financial ratios identified in previous research as important determinants of M&A activity among acquiring firms. These cover firm-level measures of capitalization, liquidity, valuation, and profitability.[5] To capture additional firm-specific characteristics related to innovation capacity, factor efficiency,

3. Working with annual data has practical advantages, as high-frequency firm-level data are often noisy or missing. For example, firms usually report employee counts only on an annual basis.

4. Firm-year observations involving serial M&A activity or consecutive acquisitions with at least a one-year gap are generally excluded to ensure clean event identification. An exception is made when one transaction clearly dominates the sequence—defined as the largest deal being more than three times the size of the next-largest deal (both measured relative to pre-merger assets). In such cases, the event year is assigned to the dominant deal, and other M&A years are treated as non-events. Assignments made under this rule account for around 15% of all treatment assignments in the sample. However, excluding these assignments does not significantly change the main results.

5. Capitalization is measured by Capital Structure Ratio; liquidity by Cash Ratio; valuation by P/E Ratio; and profitability by Gross Profit Ratio. See the Online Appendix A for relevant references.

and investment behavior, the analysis also includes R&D intensity, investment rate, capital intensity, and log labor productivity.[6]

Lastly, the dataset is restricted based on data availability for productivity levels and the aforementioned metrics. If any of these variables are missing for a given firm-year observation, the entire entry is considered as missing. A statistical summary of the financial ratios, including their construction and components, is provided in Online Appendix A.

3.2. Data Structure and Filtering

I organize firm-year data into rolling five-year cohorts to track outcomes before and after M&A events. In practice, M&A activity occurs at different times for different firms and in a sporadic manner—that is, M&A is not a routine activity for most firms. A cohort-based data structure provides a practical way to compare treated and untreated firms within a common event-time window, despite variation in the calendar years when M&A events occur.

Cohort Construction: Let $t \in \{1978, \dots, 2023\}$ denote calendar time. Each cohort c is defined by a center year t^*_c, such that the cohort index c maps directly to t^*_c (e.g., cohort 1: $t^*_c = 1979$). In total, there are 45 cohorts over the sample period. The event window for each cohort is defined as follows:

- *Standard cohorts* ($t^*_c \in \{1980, \dots, 2021\}$; 42 cohorts): symmetric five-year window $t \in \{t^*_c - 2, \dots, t^*_c + 2\}$.
- *Boundary cohorts* (3 cohorts): 1979 ($t \in \{1978, \dots, 1981\}$), 2022 ($t \in \{2020, \dots, 2023\}$), and 2023 ($t \in \{2021, \dots, 2023\}$).

For each cohort c, event time is defined as $w = t - t^*_c$, indicating a firm's relative position to the merger event year. The analysis uses the event-time window $w \in \{-2, -1, 0, +1, +2\}$, where available. For boundary cohorts, the event-time window is truncated due to data limitations at the start and end of the sample period: for 1979, $w \in \{-1, 0, +1, +2\}$; for 2022, $w \in \{-2, -1, 0, +1\}$; and for 2023, $w \in \{-2, -1, 0\}$.

Treatment Classifications and Inclusion Criteria: Within each cohort c:

- *Treated*: Firms undergoing M&A as acquirer *only* in t^*_c (retain status for $w \geq 0$).
- *Untreated*: Firms with no acquirer activity anywhere in the cohort's window.

6. R&D intensity is defined as R&D expenditures divided by total sales. The investment rate is calculated as capital expenditures over lagged capital stock, where capital stock is constructed using the perpetual inventory method. See Online Appendix A for further details.

Treatment assignment follows four principles:

1. *Multi-cohort eligibility*: Untreated firms may appear in multiple cohorts if they remain M&A-inactive across all corresponding event windows.
2. *Cohort independence*: Treatment status is determined independently for each cohort (exceptions noted below); a firm treated in cohort c may be untreated in cohort c'.
3. *Control status filtering*: Once a firm is treated in cohort c, it is never reclassified as untreated in any later cohort c' if any part of its original event window $[t_c^* - 2, t_c^* + 2]$ overlaps with the event window of cohort c', i.e., $[t_{c'}^* - 2, t_{c'}^* + 2]$.
4. *Flexible re-treatment*: A firm treated in cohort c may be eligible for re-treatment in a later cohort c', even if the two event windows are temporally adjacent (e.g., cohort c's endpoint coincides with cohort c''s start point), provided that the firm's merger year (t_c^*) does not fall within the event window of cohort c'.[7]

Lastly, within-cohort panels may be unbalanced, but firms must have non-missing data at $w = -1$ (baseline) and $w = 0$ (merger event year). Firms with missing data in either of these years, or those exiting via acquisition (i.e., targets), are excluded from the cohort.

Final Sample: The final sample contains 13,807 unique firms, of which 4,559 are assigned treated status at least once. It includes 6,752 M&A events, with each firm-year counted only once, even in cases involving multiple deals in that year. On average, there are approximately 150 treated firms per cohort, and the treated-to-untreated ratio is about 1:13.

4. Empirical Approach

I estimate the causal effects of M&A on firm-level productivity using an event-time study based on the *potential outcomes* framework (Rubin, 1974; Holland, 1986). The empirical strategy is designed to address several challenges in unit-level causal inference that are particularly relevant to this study: (i) the non-random nature of M&A events, (ii) their distributed timing over a long panel, and (iii) the relatively small number of acquiring firms compared to the available control pool at any given point in time.

The objective is to measure how firm productivity evolves around the M&A event relative to a plausible no-M&A counterfactual. Counterfactual benchmarks are constructed via out-of-

7. This relaxation helps preserve economically significant transactions and mitigates selection bias against more active acquirers. Any potential contamination from prior treatments will be addressed during estimation.

sample predictions from a latent factor model that captures both observed and unobserved systematic variations in firm productivity dynamics. This approach draws on the growing literature on causal inference in panel data models with interactive fixed effects (see, e.g., Bai, 2009; Y. Xu, 2017).

4.1. Data Structure in a Nutshell

The empirical analysis uses annual firm-level data (1978–2023) organized into rolling five-year cohorts indexed by $c \in \mathcal{C}$. Each cohort c is anchored to a center calendar time $t_c^* \in \mathcal{T}$, spanning the window $[t_c^* - 2, t_c^* + 2]$ (except for boundary cohorts, see Section 3.2). Within cohorts, event time $w \in \mathcal{W} \equiv \{-2, -1, 0, +1, +2\}$ is defined relative to t_c^*, where $w = 0$ denotes a merger event year. This setup maps each pair (c, w) to a unique calendar year t through the relation $t = t_c^* + w$.

Consequently, our setting features two complementary time dimensions: first, *calendar-time variation* across cohorts at a given event time w; second, *event-time variation* within cohorts through w. Identical event times therefore correspond to different calendar years across cohorts—for example, $w = 0$ falls in 1988 for the cohort with $t_c^* = 1988$, but in 2005 for the cohort with $t_c^* = 2005$. By anchoring treatment at t_c^* while maintaining event-time alignment, the design enables a dynamic event-time study and permits calendar-year controls (e.g., interactive fixed effects within each w-layer), if desired.

Treatment assignment follows Section 3.2: firms are assigned to cohort c as treated if they engage in M&A activity exclusively as acquirers at t_c^*, and as untreated if they exhibit no M&A activity during the cohort's window.

4.2. Framework

Let $i \in \mathcal{I}_c$, where $\mathcal{I}_c \subseteq \{1, \ldots, N\}$, index firms given calendar year $t \in \mathcal{T}$ and cohort $c \in \mathcal{C}$. Recall that each cohort c is centered at the merger year t_c^*. The binary treatment (i.e., M&A) indicator $D_{i,c,w}$ is formally defined as:

$$D_{i,c,w} = \begin{cases} 1 & \text{if firm } i \text{ undergoes an M\&A in year } t_c^* \text{ and } w \geq 0, \\ 0 & \text{otherwise} \end{cases}$$

Event time is defined as $w = t - t_c^*$, with $w = 0$ corresponding to a merger event year and $w = -1$ is the baseline period. Importantly, treatment status persists beyond the merger event year. Once a firm is treated at $w = 0$, it remains treated in periods $w = +1$ and $w = +2$. Thus, M&A effects are interpreted as cumulative impacts over the merger and post-merger window $w \in \{0, +1, +2\}$, relative to the pre-merger baseline.

Let $y_{i,t} \in \mathbb{R}$ denote log labor productivity for firm i in calendar year t. The outcome variable is defined as the change in log productivity relative to the cohort-specific pre-merger baseline:

$$\Delta y_{i,c,w} = y_{i,t_c^*+w} - y_{i,t_c^*-1} \tag{1}$$

where t_c^* is the center year of cohort c. This transformation ensures that productivity changes are measured relative to each firm's individual history, allowing for meaningful comparisons across firms irrespective of their starting productivity levels.

Let $\Delta y_{i,c,w}^{(1)}$ and $\Delta y_{i,c,w}^{(0)}$ denote the potential outcomes for firm i in cohort c at event time w under treatment (M&A) and control (no M&A), respectively. The firm-level causal effect of M&A is then defined as:

$$\mu_{i,c,w} = \Delta y_{i,c,w}^{(1)} - \Delta y_{i,c,w}^{(0)} \tag{2}$$

For treated firms ($D_{i,c,w} = 1$) in merger and post-merger periods ($w = 0, +1, +2$), we observe $\Delta y_{i,c,w}^{(1)}$ (the potential outcome under *M&A*) but not $\Delta y_{i,c,w}^{(0)}$ (the potential outcome under *no M&A*). The goal is to predict $\Delta y_{i,c,w}^{(0)}$ for these firms—that is, the productivity change that would have occurred without an M&A.

To do so, I draw on a set of firm-level covariates. Each firm i in cohort c is associated with a covariate vector $\mathbf{X}_{i,c} \in \mathbb{R}^p$, computed as the average over $w \in \{-2, -1\}$ (or just $w = -1$ if $w = -2$ is missing). This vector includes pre-merger firm characteristics that are predictive of M&A activity, such as market share, capitalization, liquidity, valuation, profitability, R&D intensity, capital intensity, investment rate, and baseline productivity, along with all pairwise interactions among these variables and a constant term (set to unity).

For firm i in cohort c, the change in productivity in the absence of M&A is specified by a dynamic factor model:

$$\Delta y_{i,c,w}^{(0)} = \mathbf{X}_{i,c}^\top \boldsymbol{\beta}_w + \boldsymbol{\lambda}_{i,w}^\top \mathbf{f}_{c,w} + \varepsilon_{i,c,w} \tag{3}$$

- $\boldsymbol{\beta}_w \in \mathbb{R}^p$ is a vector of coefficients capturing relationships between observed pre-merger firm characteristics and productivity changes at event time w,
- $\mathbf{f}_{c,w} \in \mathbb{R}^{K_w}$ is a vector of latent factors capturing aggregate conditions at event time w, shared by firms in cohort c,

- $\boldsymbol{\lambda}_{i,w} \in \mathbb{R}^{K_w}$ is a vector of firm-specific loadings (i.e., coefficients) reflecting heterogeneous sensitivities to $\mathbf{f}_{c,w}$, which may vary with w,

- $\varepsilon_{i,c,w}$ denotes a mean-zero error term.

Recall that $t = t_c^* + w$ is the calendar year associated with (c, w). Then, $\mathbf{f}_{c,w}$ is common to all firms in (c, w), while $\boldsymbol{\lambda}_{i,w}$ captures how firm i loads on those aggregate shocks at event time w. For a fixed w, the model is a standard interactive fixed effects (IFE) specification where cohort-specific common factors are actually calendar-time varying common factors through the mapping $t = t_c^* + w$; by allowing the loadings and factors to vary flexibly with w, the model generalizes to a dynamic IFE structure.

Assumption 1 (Non-collinearity) *Let $\mathbf{X}_{i,c}^{nc}$ denote $\mathbf{X}_{i,c}$ without the intercept. For every (c, w) with support, the conditional covariance of the non-constant covariates is positive definite,* $\mathrm{Var}\big(\mathbf{X}_{i,c}^{nc} \mid \mathbf{f}_{c,w}\big) \succ \mathbf{0}$.

Assumption 1 allows $\mathbf{X}_{i,c}^{\mathrm{nc}}$ to be correlated with the common factors $\mathbf{f}_{c,w}$; however, it rules out exact linear dependence among the non-constant covariates conditional on the factors.

Assumption 2 (Mean independence) *The error term is mean-independent of the firm covariates and common factors such that, for all $w \in \mathcal{W}$, $\mathbb{E}[\varepsilon_{i,c,w} \mid \mathbf{X}_{i,c}, \mathbf{f}_{c,w}] = 0$.*

Assumption 3 (Conditional ignorability) *Causal identification relies on the conditional ignorability assumption, such that, for all $w \in \mathcal{W}$, $\{\Delta y_{i,c,w}^{(0)}, \Delta y_{i,c,w}^{(1)}\} \perp\!\!\!\perp D_{i,c,w} \mid \mathbf{X}_{i,c}, \mathbf{f}_{c,w}$.*

Assumption 3 states that, conditional on pre-merger firm characteristics $\mathbf{X}_{i,c}$ and the common factors $\mathbf{f}_{c,w}$ (e.g., macroeconomic conditions), treatment assignment is as-good-as random.

Assumption 4 (Decaying carryover) *For any firm $i \in \mathcal{I}_c$ untreated in cohort c (i.e., $D_{i,c,0} = 0$) but with prior treatment in cohort $c' < c$, suppose the firm's productivity $y_{i,t}$ converges to a new post-merger level by year $t_{c'}^* + 2$. This implies growth effects from the prior M&A are largely realized within $[t_{c'}^*, t_{c'}^* + 2]$. Let $\Delta t = t_c^* - 2 - (t_{c'}^* + 2)$ with $\Delta t \geq 1$ (cohort non-overlap). Then the firm's expected productivity changes in cohort c satisfy*

$$\mathbb{E}\Big[\Delta y_{i,c,w}^{(0)} \mid \textit{prior ended at } t_{c'}^* + 2\Big] = \mathbb{E}\Big[\Delta y_{i,c,w}^{(0)} \mid \textit{no prior treatment}\Big] + \phi(\Delta t),$$

where $\phi(\Delta t)$ is a residual term bounded in sufficiently small magnitude by $\kappa \cdot h(\Delta t)$ for all Δt in the sample, with $\kappa > 0$ and $h(\cdot)$ a non-negative, decreasing function satisfying $h(\Delta t) \to 0$ as $\Delta t \to \infty$.

Assumption 4 requires that firms used as controls are credible benchmarks, even if they were acquirers in the past. Accordingly, any control firm with a prior M&A treatment must have completed the bulk of its post-merger adjustment and converged to a new productivity level by year $t^*_{c'} + 2$. After this convergence point, the firm resumes a new growth trajectory from its new productivity level, with any lingering deviations in growth rates decaying sufficiently fast over time. This requires that growth effects (temporary rate changes during adjustment) do not persist beyond the adjustment window, while allowing permanent level shifts.

4.3. Estimation Procedure

I predict counterfactual outcomes using a three-step procedure.[8]

Step 1: Initial Regression

For each event time w, I estimate a separate regression of the outcome $\Delta y_{i,c,w}$ on firm covariates $\mathbf{X}_{i,c}$. Each regression is estimated using data only from control firms, pooled across all cohorts at that specific event time w.[9] This yields a coefficient vector $\tilde{\boldsymbol{\beta}}_w$:

$$\tilde{\boldsymbol{\beta}}_w = \arg\min_{\boldsymbol{\beta}_w} \sum_{i,c:D_{i,c,0}=0} \left(\Delta y_{i,c,w} - \mathbf{X}_{i,c}^\top \boldsymbol{\beta}_w\right)^2 \tag{4}$$

Using the estimated coefficients, I compute first-stage residuals as:

$$u_{i,c,w} = \Delta y_{i,c,w} - \mathbf{X}_{i,c}^\top \tilde{\boldsymbol{\beta}}_w = \boldsymbol{\lambda}_{i,w}^\top \mathbf{f}_{c,w} + \varepsilon_{i,c,w} \tag{5}$$

These residuals capture variation in outcomes not explained by observables. Thus, the remaining variation is assumed to be a combination of latent common shocks (cohort-specific, $\mathbf{f}_{c,w}$) and firm-specific sensitivities (loadings, $\boldsymbol{\lambda}_{i,w}$), along with noise $\varepsilon_{i,c,w}$.

Step 2: Factor Extraction

To estimate the latent common factors $\mathbf{f}_{c,w}$ and firm-specific loadings $\boldsymbol{\lambda}_{i,w}$ at each event time w, I stack the first-stage residuals into a matrix $\mathbf{U}_w \in \mathbb{R}^{N\times|\mathcal{C}|}$ with firms i as rows and cohorts c as columns. By construction, $\mathbf{U}_w$ has missing entries either because a firm is unobserved in a cohort or because it is treated there; concretely, $U_w[i,c] = u_{i,c,w}$ if $D_{i,c,0} = 0$ (untreated in cohort c) and $U_w[i,c] = \text{NA}$ if $D_{i,c,0} = 1$ (treated in cohort c) or since it is simply missing.

8. See Y. Xu (2017) and Gabaix and Koijen (2024) for similar implementations

9. To keep estimation strictly control-only, I exclude all observations from firm–cohort pairs that are treated in that cohort (including their pre-merger rows with $w < 0$).

Hence, each column contains only control residuals if available, and treated cells are missing and never enter factor extraction.

However, in our data structure, nearly all treated firms appear as untreated in at least one non-overlapping cohort, ensuring each firm contributes some residuals. For each w, the loadings $\boldsymbol{\lambda}_{i,w}$ are firm-specific and constant across cohorts, while the factors $\mathbf{f}_{c,w}$ are common across firms within cohort c. This structure allows the estimation of a firm's loadings even when it is observed only in a subset of cohorts, by utilizing available control data. Eventually, for the extraction of factor structure, I proceed as follows:

- Matrix Completion: for each w, I apply a regularized PCA to impute missing entries in $\mathbf{U}_w$, based on the assumption that the residual matrix has a low-rank structure.[10]
- After completion, again for each w, I apply standard PCA to the imputed matrix. This yields a firm-specific loading matrix $\tilde{\boldsymbol{\Lambda}}_w \in \mathbb{R}^{N \times K_w}$ and a cohort-specific factor matrix $\tilde{\mathbf{F}}_w \in \mathbb{R}^{|\mathcal{C}| \times K_w}$, where K_w denotes the number of latent factors selected for that w.
- I compute a composite latent component for each firm–cohort pair (i, c) at event time w, defined as $\tilde{L}_{i,c,w} = \sum_{k=1}^{K_w} \tilde{\lambda}_{i,w}^{(k)} \tilde{f}_{c,w}^{(k)}$, where $\tilde{\lambda}_{i,w}^{(k)}$ is the loading of firm i for factor k and $\tilde{f}_{c,w}^{(k)}$ is cohort c's realization of factor k.

The number of factors is selected using the eigenvalue-ratio (ER) criterion (Ahn and Horenstein, 2013). If the initially selected factors explain less than 50% of the total residual variance at a given w, additional factors are added in descending order of explanatory power until this threshold is met.[11]

Note that while the latent component captures systematic covariation patterns, its scale is arbitrary due to PCA's mathematical properties. For empirical relevance, the component will be rescaled in Step 3 to align with the outcome variable's metric, without changing the underlying factor structure.

Step 3: Counterfactual Prediction

For each event time w, I then estimate the model again only using control observations pooled across all cohorts at that specific event time w:

$$\Delta y_{i,c,w} = \mathbf{X}_{i,c}^{\top} \boldsymbol{\beta}_w + \gamma_w \cdot \tilde{L}_{i,c,w} + \text{error}, \tag{6}$$

10. A *low-rank* matrix is one whose rank is much smaller than its dimensions. This implies that the data can be well-approximated by a few underlying systematic factors. I use the `missMDA` package in `R`

11. Mild over-selection is preferable to under-selection in interactive fixed-effects settings.

yielding estimates $\hat{\boldsymbol{\beta}}_w$ and $\hat{\gamma}_w$. The latent component is then rescaled as $\widehat{\boldsymbol{\lambda}}_{i,w}^{\top}\widehat{\mathbf{f}}_{c,w} := \hat{\gamma}_w \tilde{L}_{i,c,w}$. For each treated firm–cohort pair ($D_{i,c,0} = 1$) and event time $w = \{-2, 0, +1, +2\}$, I use these estimates to construct the counterfactual change in productivity in the absence of M&A:

$$\widehat{\Delta y}_{i,c,w}^{(0)} = \mathbf{X}_{i,c}^{\top}\hat{\boldsymbol{\beta}}_w + \hat{\boldsymbol{\lambda}}_{i,w}^{\top}\hat{\mathbf{f}}_{c,w}, \tag{7}$$

and compute the individual productivity changes driven by M&A as:

$$\hat{\mu}_{i,c,w} = \Delta y_{i,c,w}^{(1)} - \widehat{\Delta y}_{i,c,w}^{(0)} \tag{8}$$

4.4. Aggregation of M&A Effects

The pooled average productivity change at event time w is computed across all acquiring firms and cohorts:

$$\hat{\vartheta}_w = \frac{1}{n_w} \sum_{c \in \mathcal{C}} \sum_{\substack{i \in \mathcal{I}_c \\ D_{i,c,0}=1}} \hat{\mu}_{i,c,w}, \tag{9}$$

where $n_w = \sum_{c \in \mathcal{C}} n_{c,w}$ is the total number of acquiring firm-event time observations at w. Alternatively, the pooled effects can be obtained via regression:

$$\hat{\mu}_{i,c,w} = \sum_{\tau \in \{-2,0,+1,+2\}} \theta_\tau \cdot \mathbf{1}\{w = \tau\} + \epsilon_{i,c,w} \quad \text{(relative to } w = -1\text{)} \tag{10}$$

with robust standard errors clustered at the cohort level.[12]

4.5. Main Results from Micro Data

Figure 1 presents the estimated average productivity changes due to M&A. The upper panel plots the baseline results using a five-year cohort window (window = 5), and the lower panel presents the results for different categories of firm size.

I classify firms based on their pre-merger market share in GDP within each cohort. Firms in the bottom 50% are classified as small, those in the 50th–75th percentiles as medium, and the top 25% as large. Note that investigating firm size heterogeneity is not just an

12. All specifications use robust standard errors clustered at the cohort level. Inferences are nearly identical when clustering by cohort-event time (c, w). Note that clustering by firm is not preferred in this setting, as it yields too few repeated firm observations in each cluster because treated firms appear sparingly and typically belong to only 1–2 cohorts. This results in anti-conservative (spuriously small) standard errors.

econometric exercise here but also aims to inform subsequent analysis on granularity.

The results reveal a clear pattern. In the merger year ($w = 0$), acquirers experience a statistically significant average decline in productivity of about 4.9% points. Productivity then recovers: by $w = +1$ it begins to rebound, and by $w = +2$ it stands roughly 3.2% points above the pre-merger baseline.

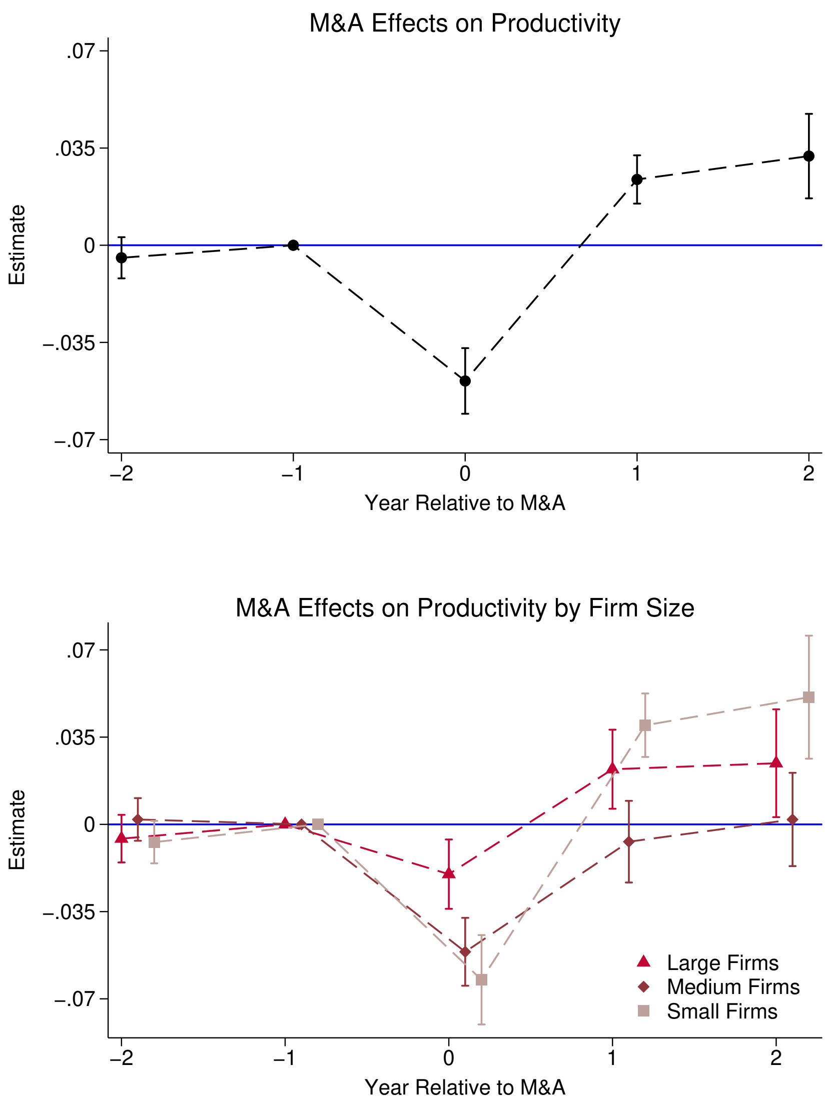

Figure 1: Event-Time Effects of M&A on Firm Productivity

Notes: The upper panel presents average M&A effects on firm productivity (log). The lower panel shows heterogeneous effects by firm size groups. Robust standard errors are clustered at the cohort level and computed using pooled data across cohorts for each event time. The same estimation procedure is applied separately to each firm size subgroup. The vertical lines refer to 90% CI.

These dynamics suggest that M&A activity may temporarily disrupt firm performance, likely due to integration frictions or transitional inefficiencies, but eventually leads to economically significant productivity gains within a two-year horizon. These results are consistent with empirical evidence from industry studies (e.g., Schoar, 2002; Braguinsky et al., 2015) documenting post-merger recoveries after initial productivity declines.

The lower panel of the figure confirms a heterogeneity in productivity changes by firm size. While pre-merger trajectories appear flat across all size groups, substantial divergence emerges after the M&A. In the event year, small firms experience the sharpest drop in productivity, around 6.2% points, followed by medium-sized firms at 5.1% points and large firms at 2.0% points. The milder decline observed among larger firms likely reflects their greater ability to absorb integration inefficiencies, due to stronger management, previous M&A experience, or scale advantages. In the post-merger period, however, both large and small firms show clear productivity gains. By year $w = +2$, small firms reach an estimated gain of about 5.1% points above baseline, while large firms show quick improvement and end the window roughly 2.5% points higher than pre-merger levels. In contrast, medium-sized firms display a weaker recovery, with productivity gain close to zero by year $w = +2$.

Ultimately, two main takeaways emerge from this analysis. First, acquiring firms experience significant fluctuations in productivity due to M&A. Second, firm size plays a critical role in moderating these post-merger effects. The magnitude of these productivity changes, and their variation in intensity and recovery patterns across the firm-size distribution, naturally raises an important question: do these short-run, firm-level movements have implications at the aggregate level? I take up this question in Section 5.

4.6. Decomposition of Productivity Dynamics

To shed light on the mechanisms behind this J-curve response of firm labor productivity to M&A, I decompose productivity into its two components, changes in log revenue and changes in log employment, and estimate event-time effects for each outcome separately. Both series are expressed as differences relative to the pre-merger base period ($w = -1$). They are obtained using the same counterfactual procedure applied in the main analysis, run independently for each outcome. Note that because the counterfactuals are estimated separately, the difference of the two estimated components approximates—but does not necessarily equal—estimated changes in the log of the revenue-to-employment ratio. The purpose is interpretive rather than quantitative equivalence. The results are reported in Figure Bx.1 in Online Appendix B.

The decomposition shows that employment rises immediately at the merger event year ($w = 0$) and then remains relatively stable over the event window, indicating no evidence of systematic layoffs. Revenue increases more slowly at first, which accounts for the short-run dip in labor productivity, but subsequently accelerates and ultimately surpasses the employment increase toward the end of the window. This pattern is consistent with integration and capacity expansion occurring first (e.g., absorbing staff, overlapping roles, and harmonizing systems), followed by the realization of demand-side synergies (e.g., cross-selling, a wider product line, and pricing improvements). While large and small firms follow a similar qualitative path as the average, medium-sized firms, in line with the main analysis, display convergence of revenue and employment by the end of the window; neither margin clearly dominates, and thus there is no evident productivity gain.

Taken together, the results on average point to a sales-led increase in labor productivity relative to the pre-merger level, rather than to improvements driven by headcount reductions.

4.7. Placebo Test

To evaluate the performance of the factor model, I conduct a placebo test using a subset of control firms that are *ex ante* similar to the treated firms based on observable pre-merger characteristics. For each cohort c, I estimate a logit model to compute firm-level propensity scores (i.e., predicted probabilities of M&A activity) based on pre-merger covariates and their pairwise interactions, including a constant. Specifically, this is given by the equation $\Pr(\text{Merger}_{i,c} = 1 \mid \mathbf{X}_{i,c}) = \exp(\boldsymbol{b}_c^\top \mathbf{X}_{i,c})/\{1 + \exp(\boldsymbol{b}_c^\top \mathbf{X}_{i,c})\}$, where $\mathbf{X}_{i,c}$ denotes the vector of respective firm covariates.

Each treated firm is then matched to a control firm within the same cohort based on the estimated scores, using nearest-neighbor matching without replacement. These matched controls, while not actually treated, are used as placebo-treated units. I withhold outcome data of these placebo firms during the entire event window and estimate the factor model using the remaining control units. The actual treated firms are excluded from this exercise.

Since the model accounts for aggregate shocks (via estimated latent factors) and adjusts for productivity variation associated with firm characteristics, and because the placebo firms were not involved in an actual M&A deal, any remaining deviation between predicted and observed outcomes should reflect only random idiosyncratic noise. Consistent with this identification assumption, I find that the estimated dynamic effects for placebo firms do not exhibit a systematic pattern and are statistically insignificant throughout the event time window (see Table Bx.3 in the Online Appendix B).

5. M&A-Driven Shocks and Granularity

5.1. Main Findings

The findings in the previous section show that acquiring firms typically experience substantial productivity fluctuations within a short period of time. This raises the question whether these firm-level movements could contribute to aggregate fluctuations—an especially relevant consideration in highly granular economies like the United States.

Naturally, firm-level shocks may cancel out if productivity gains at some firms offset losses at others. Whether this occurs, however, depends on the distribution of firm size among acquirers and the share of the economy they represent at a given point in time.

Firm size, s, is often distributed according to Zipf's Law, where the tail probability behaves as $\mathbb{P}(s > \omega) \sim \omega^{-1/\zeta}$ for large ω, with $\zeta > 0$ as the tail index and the canonical case being $\zeta \approx 1$ (see Axtell, 2001; Gabaix, 2009).[13] The treated firms in this study exhibit a similar distribution.[14]

To quantify tail thickness, I estimate the tail index using the Hill estimator (Hill, 1975), given by $\widehat{\xi}_m = \frac{1}{m}\sum_{i=1}^{m}\left(\ln s_{(i)} - \ln s_{(m+1)}\right)$, and the extreme value index is $\widehat{\zeta}_m = 1/\widehat{\xi}_m$, where m is the number of top-order statistics (largest observations), and $s_{(1)} \geq s_{(2)} \geq \cdots \geq s_{(M)}$ are the order statistics with $s_{(i)}$ denoting the i-th largest value.[15]

Consequently, higher values of $\widehat{\xi}_m$ (or lower $\widehat{\zeta}_m$) are indicative of heavier tails, meaning a higher likelihood of firms with very large market shares. As shown in the left panel of Figure 2, estimated tail indices $\widehat{\zeta}_m$ range between 0.5 and 1.5 across selected years and values of m, converging toward 0.5 when more top firms are added to the estimation (i.e., larger m). This implies a persistently heavy upper tail in firm size distribution. The decline in tail indices over time suggests increasing concentration among top acquirers.

To check whether these firms' combined size is large enough to matter, I compute their market share in each year. A firm is included if its merger year ($w = 0$) or either of the two post-merger years ($w = +1, +2$) falls within the calendar year, in line with the productivity analysis window. As shown in the right panel of Figure 2, treated firms account for around 10% of U.S. GDP on average. While modest by the standards of typical granular studies, this figure covers only firms actively adjusting post-merger. In any given year, then, roughly

13. This implies that firm size is inversely proportional to rank: $s_{rank} \propto \text{rank}^{-\zeta}$.

14. Firm size $s_{i,t-1}$ is defined as firm i's lagged market share: $s_{i,t-1} = \text{Sales}_{i,t-1}/\text{GDP}_{t-1}$.

15. This assumes $s > 0$, which holds by construction. Under standard conditions, the Hill estimator is consistent and asymptotically normal, with variance $\text{Var}(\widehat{\xi}_m) \approx \xi^2/m$ (see Hill, 1975; Embrechts, Klüppelberg, and Mikosch, 1997). The choice of m is guided by visual inspection of the Hill plot, selecting a region where estimates stabilize and variance declines (see Danielsson et al., 2016).

one-tenth of the economy is exposed to substantial, firm-level productivity shocks.

All these observations suggest that M&A shocks might leave some granularity in the macro data. To further investigate this, I now turn to a size-weighted analysis of these shocks, examining whether they inject fluctuations into aggregate productivity growth.

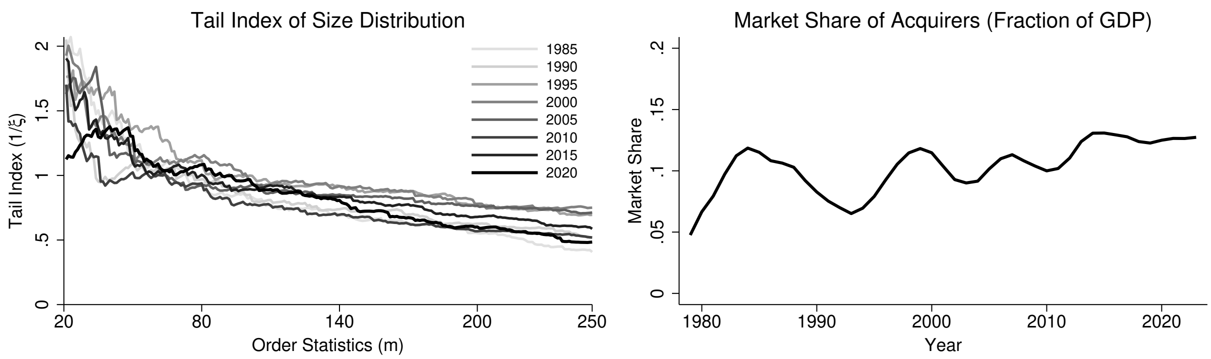

Figure 2: Fat-Tailedness of Size Distribution and Market Share of Acquiring Firms

Notes: Each calendar year in the figure includes acquiring firms observed during their merger or post-merger periods ($w = 0, +1, +2$), provided any of these phases fall within that year. The left panel shows the estimated tail index for selected years. A downward slope (e.g., from the top 50 to top 250 firms) indicates heavier tails in the firm size distribution, while flattening suggests convergence of the Hill estimator once a sufficient number of top firms (m) are included. The right panel shows the total market share held by acquiring firms during their merger year and two subsequent post-merger years.

From Section 4, I have firm-level cumulative changes relative to the base period—under the baseline structure, $\hat{\mu}_{i,c,w}$ for $w = 0, +1, +2$. Using these estimates, I compute year-over-year changes in log productivity to obtain M&A-driven productivity shocks on a yearly basis: $\Delta\hat{\mu}_{i,c,0}$ is the change from $w = -1$ to $w = 0$; $\Delta\hat{\mu}_{i,c,1}$ from $w = 0$ to $w = +1$; and $\Delta\hat{\mu}_{i,c,2}$ from $w = +1$ to $w = +2$. Accordingly, I use annual data on firm size and productivity shocks for each acquiring firm in a given calendar year, provided that year corresponds to the M&A event year or one of the two post-merger years.

The sample period of interest is 1979–2023. For notational convenience, let $\widehat{\Delta\mu}_{it}$ denote the estimated firm-specific productivity shock for firm i in year t, and let $\mathcal{M}_t \subseteq \{1, \ldots, N^{\text{merger}}\}$ present the set of firms that are in their merger or post-merger period at time t. Following the granular aggregation logic of Gabaix (2011), the aggregate M&A-driven productivity shock, i.e., a granular measure, is defined as follows:

$$\Gamma_t = \sum_{i \in \mathcal{M}_t} \frac{\text{Sales}_{i,t-1}}{\text{GDP}_{t-1}} \cdot \widehat{\Delta\mu}_{it} \tag{11}$$

where $s_{i,t-1} \equiv \text{Sales}_{i,t-1}/\text{GDP}_{t-1}$ denotes firm i's lagged market share.

Recall that M&A-driven productivity shocks are identified under an orthogonality assumption—that they are firm-specific and not driven by underlying aggregate conditions. This implies that the granular measure is not expected to show strong systematic comovement over time. Consistent with this, I do not detect statistically significant serial correlation (see Ljung–Box and Breusch–Godfrey LM tests in Table Bx.4 in Online Appendix B). Accordingly, within our sample there is no evidence that future values of the granular measure are predictable from past values, which is reassuring for our identification assumptions.

I next estimate a distributed lag model for U.S. GDP per capita growth on this granular measure of M&A shocks. The model is specified as follows:

$$g_t = a_0 + \sum_{l=0}^{L} \alpha_l \cdot \Gamma_{t-l} + e_t \tag{12}$$

where g_t denotes real GDP per capita growth (source: BEA, Chained 2017 Dollars). Results in Table 1 indicate that the granular measure of M&A-driven shocks explains around 20% of variations in aggregate productivity growth.

The coefficient estimates show three patterns: a statistically significant positive contemporaneous coefficient ($\hat{\alpha}_0 > 0$ at 5% level), a statistically significant negative first-lag coefficient ($\hat{\alpha}_1 < 0$ at 5% level), and an insignificant long-run multiplier ($\sum_{k=0}^{K} \hat{\alpha}_k \approx 0$). This implies an immediate pass-through of the productivity changes to aggregate growth followed by mean reversion within a year. This is consistent with M&A generating firm-specific productivity *level* effects rather than permanent *growth* effects. Relatedly, the statistically insignificant long-run multiplier implies no persistent effect on trend growth.

The granular measure is constructed as an annual cross-sectional aggregation of firm-specific shocks in a given year t. This construction implies that, for example, when firms in disruption phases dominate the cross-section, the measure declines, thereby dampening aggregate growth. Conversely, when these same firms are in recovery phases in subsequent years, the measure increases, boosting growth. Consequently, the positive contemporaneous and negative lagged coefficients capture, respectively, the impulse and mean-reverting components of this micro-to-macro relationship. A null long-run effect confirms that these firm-level shocks are transitory deviations from the growth trend, rather than permanent structural changes.

In short, these results are intriguing. While M&A shocks exhibit stylized post-event dynamics in the microdata, the size-weighted aggregate of these shocks behaves as an exogenous innovation in our macro data—not forecastable from its own lags—yet it has statistically significant explanatory power for economic growth. The subsequent sections explore the robustness through a series of complementary statistical checks.

Table 1: Regressions of GDP per capita Growth on the Granular Measure

	(1)	(2)	(3)
Γ_t	1.051**	0.983**	1.066**
	(0.436)	(0.430)	(0.430)
Γ_{t-1}	-0.957**	-0.985**	-1.083**
	(0.436)	(0.427)	(0.427)
Γ_{t-2}		-0.403	-0.449
		(0.430)	(0.427)
Γ_{t-3}			-0.697
			(0.434)
Intercept	0.0167***	0.0173***	0.0173***
	(0.0026)	(0.0026)	(0.0026)
Long-run Multiplier	0.095 (0.639)	-0.405 (0.820)	-1.164 (0.946)
N	44	43	42
R^2	0.219	0.247	0.298
Adj. R^2	0.181	0.189	0.222

Notes: The study period is 1979-2023. The dependent variable is GDP per capita Growth. Γ_t denotes the current value of the granular measure of M&A-driven shocks, while Γ_{t-1},Γ_{t-2},Γ_{t-3} denote its lagged values. The long-run multiplier refers to the cumulative effect of the granular measure on growth. Standard errors are in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

5.2. Robustness Checks

M&A-Driven Revenue Transfers: The existing literature highlights that M&A can increase market concentration, often by increasing profits or revenue transfers from targets to acquirers (see, e.g., Kim and Singal, 1993; Grullon, Larkin, and Michaely, 2019; Philippon, 2019; Levonyan and Mengano, 2022). Moreover, a recent theoretical work by Chan and Qi (2025) argues that mergers increase the fat-tailedness of the overall firm-size distribution in the U.S. economy, thereby amplifying granular fluctuations.

As discussed in Section 4.6, firm revenues rise following M&A. Because firm market share is used as an aggregation weight, these changes could affect the granular analysis. While this does not invalidate the core findings, since both productivity and market share respond to M&A, attributing granularity solely to productivity shocks would be incomplete.

I first reconstruct the granular measure using pre-merger market shares from each firm-cohort rather than the lagged shares in specification (11). This approach fixes market shares throughout the merger and post-merger periods ($w = 0, +1, +2$). As Table 2 Panel A shows, this modification does not drastically affect the explanatory power.

As a further check, I examine whether the findings hold when using counterfactual market shares—that is, as if firms had not undertaken any M&A activity. To implement this, I repeat the factor analysis model, this time regressing base-period differenced absolute revenue changes on firm-level covariates and time-varying common factors in order to estimate counterfactual sizes.[16] Aggregating productivity shocks using the counterfactual market shares (one-year lagged) yields very similar results (R^2 values) to those obtained from the baseline analysis (compare Panel B with Table 1). To sum up, the results suggest that granularity is driven mainly by productivity shocks rather than by changes in market share.

A natural question is why changes in firm revenues would *not* meaningfully affect granularity. One explanation is that such increases do not substantially alter market structure in the short run. Figure Bx.3 in Online Appendix B compares the realized total market share of acquiring firms with two benchmarks: (i) the total based on pre-merger firm sizes and (ii) the counterfactual values estimated from the factor model. Both benchmarks lie slightly below the realized total, indicating that sales revenues rise following M&A; however, these gains are not large enough to materially change the acquirers' combined market share. This is especially plausible for large acquirers, for whom there are few, if any, larger target entities from which to absorb substantial share. As a result, M&A effects on relative size remain limited in the short run.

Nevertheless, it bears repeating that this study focuses on the immediate, direct effects of M&A. Longer-run implications for market structure—whether through competition-distorting channels or a general rise in merger activity over time—lie beyond our scope. By design,

16. Note that the common latent components are re-estimated based on revenue changes.

our analysis isolates abrupt, firm-specific responses to M&A. Within this scope, productivity shocks appear to be the primary source of the observed granularity

Table 2: Regressions on the Granular Measure Constructed with Counterfactual Firm Sizes

	Panel A: Pre-Merger Sizes		Panel B: Counterfactual Sizes	
	(1)	(2)	(3)	(4)
$\Gamma_{\mathrm{t}}^{\mathrm{S}}$	0.936*	0.854*	1.004*	0.921*
	(0.465)	(0.457)	(0.479)	(0.469)
$\Gamma_{\mathrm{t-1}}^{\mathrm{S}}$	-1.062**	-1.134**	-1.153**	-1.200**
	(0.466)	(0.458)	(0.479)	(0.467)
$\Gamma_{\mathrm{t-2}}^{\mathrm{S}}$		-0.511		-0.546
		(0.457)		(0.469)
Intercept	0.0167***	0.0173***	0.0167***	0.0173***
	(0.0026)	(0.0026)	(0.0026)	(0.0026)
N	44	43	44	43
R^2	0.206	0.241	0.214	0.250
Adj. R^2	0.168	0.183	0.175	0.192

Notes: The study period is 1979-2023. The dependent variable is GDP per capita Growth. Γ_t^S denotes the current value of the granular measure (of M&A-driven shocks) constructed using pre-merger firm size from data or counterfactual firm size from the factor model, while Γ_{t-1}^S,Γ_{t-2}^S denote its lagged values. Standard errors are reported in parentheses. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Time-Varying Exposure to M&A: One validity concern is whether cyclical patterns in merger activity, rather than genuine productivity shocks, could drive the granularity result. Although firm-level shocks are constructed to be orthogonal to aggregate conditions, time-varying exposure through two channels may mechanically correlate the aggregate shock measure with macroeconomic outcomes.

First, on the extensive margin, the relative number of acquiring firms may increase during expansionary periods (e.g., merger waves). Second, on the intensive margin, the relative weight of acquiring firms, measured using their lagged market shares when constructing the granular measure, may increase during booms if larger firms are more likely to engage in M&A. Even though market shares are lagged, this weighting can still be correlated with macroeconomic conditions, so the aggregate measure may capture cyclical exposure patterns rather than purely underlying productivity shocks.

However, two empirical features help mitigate this concern. First, merger waves typically unfold over several years, meaning that firm-specific shocks tend to be staggered rather than synchronized. Relatedly, because treated firms enter and move through the post-merger period at different times, they contribute shocks of varying signs at any given moment—some in the early phase (negative shocks), others in recovery (positive shocks), and some already stabilized (minimal shocks). This temporal dispersion implies that any confounding exposure-outcome relationship is a priori ambiguous in the context of our granular design.

Nevertheless, I formally test for exposure-driven effects by controlling for both the lagged total GDP share of acquiring firms, $S_{t-1}^{\text{merger}} = \sum_{i \in \mathcal{M}_t} s_{i,t-1}$, and the relative number of acquiring firms in the sample in a given year, i.e., the M&A intensity N_t^{merger}/N_t– along with their further lagged values.

Table 3 reports the respective results. Merger intensity is correlated with aggregate productivity growth (Column 2). However, when it is included alongside the granular measure, the latter's incremental explanatory power is essentially similar to that in specifications where it enters alone (Column 3 vs. Column 1). The same holds across alternative specifications. The intercept becomes statistically insignificant once the total market share of the acquiring firms is included, indicating that this control captures average growth variation previously absorbed by the constant term.

These results suggest that cyclical exposure, whether through the relative scale of acquiring firms or the prevalence (extensive margin) of M&A events, is unlikely to drive the granularity.

Macroeconomic Controls: I would like to test whether the granular measure of M&A shocks retains explanatory power after controlling for standard macroeconomic variables. This section can also be viewed as a separate check to see whether the estimated granularity is due to actual firm-level M&A shocks or to underlying macroeconomic conditions.

To maintain parsimony given the limited time-series length, and so as not to stifle degrees of freedom, I control for the most commonly used indicators of economic growth. These controls are also used in Gabaix (2011), so the specification is broadly in line with that study. Eventually, I control for the short-term real interest rate (measured by the three-month T-bill rate adjusted for year-over-year inflation), the term spread (defined as the five-year Treasury constant maturity rate minus the federal funds rate), and lagged GDP growth to capture persistence.[17] I also include energy price shocks from Hamilton (2003), defined as net increases relative to the maximum price in the preceding three years.[18]

Table 3: Explanatory Power of the Granular Measure with Merger Exposure Controls

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Γ_t	1.051**		0.989**		1.155**		1.155***
	(0.436)		(0.410)		(0.436)		(0.412)
Γ_{t-1}	-0.957**		-1.051**		-0.944**		-0.916**
	(0.436)		(0.411)		(0.436)		(0.411)
N_t^{merger}/N_t		0.0037**	0.0038**			0.0047**	0.0046***
		(0.0016)	(0.0014)			(0.0019)	(0.0017)
N_{t-1}^{merger}/N_{t-1}		-0.0030*	-0.0031**			-0.0043**	-0.0043**
		(0.0016)	(0.0014)			(0.0019)	(0.0017)
S_{t-1}^{merger}				0.196	0.252	-0.015	0.045
				(0.185)	(0.165)	(0.197)	(0.172)
S_{t-2}^{merger}				-0.063	-0.093	0.177	0.147
				(0.163)	(0.146)	(0.187)	(0.166)
Intercept	0.0167***	0.0117**	0.0115**	0.0021	-0.0007	-0.0032	-0.0060
	(0.0026)	(0.0048)	(0.0043)	(0.0123)	(0.0111)	(0.0120)	(0.0106)
N	44	44	44	44	44	44	44
R^2	0.219	0.120	0.346	0.041	0.279	0.169	0.402
Adj. R^2	0.181	0.077	0.279	-0.006	0.205	0.084	0.305

Notes: The study period is 1979-2023. The dependent variable is GDP per capita Growth. S_{t-1}^{merger}, S_{t-2}^{merger} denote lagged total market shares of acquiring firms in a given year, N_t^{merger}/N_t and N_{t-1}^{merger}/N_{t-1} denote current and lagged merger intensities in the sample in a given year. Γ_t denotes the current value of the granular measure of M&A-driven shocks, while Γ_{t-1} denote its lagged values. Standard errors are reported in parentheses. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

Table 4 reports results from regressions that horse-race the granular measures against these variables. The adjusted R^2 increases from 0.181 in the baseline with macro controls (Column 2) to 0.360 when the granular measures are added (Column 3). Conditional on oil-price shocks and the full set of macro controls, the granular measures continue to improve model

17. Gabaix (2011) also includes monetary policy shocks from Romer and Romer (2004), but these are only available through 2007.

18. As the original series for energy price shocks is reported quarterly, I sum the quarterly shocks at the annual level.

fit. This finding suggests that M&A-driven shocks provide additional explanatory power for growth even conditional on conventional macroeconomic variables.

Table 4: Explanatory Power of the Granular Measure with Macroeconomic Controls

	(1)	(2)	(3)	(4)	(5)
Γ_t	1.051**		0.728*		0.553
	(0.436)		(0.395)		(0.392)
Γ_{t-1}	-0.957**		-1.229***		-1.266***
	(0.436)		(0.431)		(0.436)
g_{t-1}		-0.011	0.109	-0.090	0.018
		(0.147)	(0.137)	(0.142)	(0.128)
r_t		-0.058	0.061	0.091	0.184
		(0.149)	(0.137)	(0.145)	(0.130)
r_{t-1}		0.216	0.114	0.082	0.035
		(0.147)	(0.133)	(0.152)	(0.136)
Term Spread$_t$		-0.582	-0.415	-0.594*	-0.444
		(0.360)	(0.322)	(0.329)	(0.291)
Term Spread$_{t-1}$		1.118***	0.954***	0.893**	0.877***
		(0.317)	(0.286)	(0.338)	(0.318)
Oil$_t$				0.00016	0.00027
				(0.00021)	(0.00021)
Oil$_{t-1}$				-0.00057***	-0.00046**
				(0.00020)	(0.00018)
Intercept	0.0167***	0.0115**	0.0091**	0.0184***	0.0131**
	(0.0026)	(0.0044)	(0.0040)	(0.0057)	(0.0057)
N	44	44	44	44	44
R^2	0.219	0.276	0.464	0.426	0.588
Adj. R^2	0.181	0.181	0.360	0.314	0.479

Notes: The study period is 1979–2023. The dependent variable is GDP per capita growth and g_{t-1} denotes its lagged value; r_t, r_{t-1} refer to the current and lagged real interest rate. Oil$_t$, Oil$_{t-1}$ indicate current and lagged energy price shocks. Γ_t, Γ_{t-1} denote the current and lagged values of the granular measure of M&A-driven shocks. Standard errors are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Random Idiosyncratic Shocks: A limitation of the analysis is the implicit assumption that acquiring firms face only M&A shocks during the intervention window. In practice, observed productivity may also contain idiosyncratic fluctuations unrelated to mergers, which can contaminate the estimated firm–level shocks. While such noise tends to wash out in average treatment effects, it is more problematic for a granular analysis because the granular measure is firm-size–weighted and the size distribution of acquiring firms is fat–tailed.

Large random shocks to big firms can distort the explanatory power of the granular series in either direction: granularity may be *attenuated* if noise obscures true M&A shocks (the more likely case), or *amplified* if it happens to move in the same direction. There is no direct way to identify the extent to which true granularity is masked by such random shocks. However, I will undertake several robustness checks to assess the credibility of the granular measure here. Before turning to formal analysis, I would like to offer some intuitive reasoning that M&A shocks could be the genuine driver of our results.

First, acquiring firms account for only about 10% of aggregate activity in a given year, yet the resulting R^2 values are relatively high, which points beyond sheer noise. Consistent with this, I employ a placebo granular measure based on placebo-treated firms matched to actual acquirers in the earlier firm-level analysis. By design, these shocks consist solely of random idiosyncratic noise, and the measure explains far less of aggregate fluctuations than the M&A-based measure.[19] These results are reported in Online Appendix B (Table Bx.5).

Herein, I conduct two robustness tests on the granular measure. The first is a permutation-based placebo test, in which productivity shocks are randomly reassigned across firms to assess whether the observed granularity could arise spuriously. The second is a stress test, where random variation is added to the shock series to evaluate the sensitivity of aggregate results to signal distortion. Together, these tests ask: does the real firm–shock pairing matter (i.e., systematic firm-level relationships), and does it still matter when shocks are deliberately made noisier? If the actual granular measure consistently outperforms the permuted versions and continues to do so under plausible noise stress, the observed granularity is unlikely to be an artifact of random idiosyncratic shocks.

(i) Synthetic Granular Shocks. Recall that the estimated firm-level productivity shocks are quite large on average — that alone might generate some explanatory power (high R-squared), even if the assignment of shocks to firms is arbitrary. This is because in this

19. Placebo-treated firms account for roughly 9% of GDP in a given year, and their size distribution is right-fat-tailed (see Figure Bx.2 and Table Bx.2 in the Online Appendix). Under the granular hypothesis, a size-weighted aggregation of such shocks should explain part of aggregate fluctuations. In our analysis, they explain about 7% of aggregate fluctuations, compared with around 19% for the M&A-driven shocks.

economy, large firms exist and receive large weights, so even random alignment between firm size and shock magnitude can produce non-trivial macro-level variations, even without needing any structural causality behind.

This mechanism lies at the heart of the granular hypothesis's assumption of shock randomness, and it is precisely this point where I put the granular measure to the test, as my setup departs from this assumption by arguing that firm-level dynamics are causally driven rather than purely random. Eventually, I keep the estimated (yearly) shocks intact, and thus their distribution, but break their alignment with firm identities by randomly reassigning shocks across treated firms in the same cohort–event year (c, w). This process simulates a world in which shocks are arbitrary events and thus statistically indistinguishable noise.

I implement this placebo exercise using a Monte Carlo-style permutation. Formally, for each draw $r = 1, \ldots, R$ (with $R = 1000$), I reshuffle firm shocks within each (c, w) group. Let $\pi_{c,w}^{(r)}$ be a random derangement (a bijection with no fixed points) on the firms in group (c, w), so that $\pi_{c,w}^{(r)}(i) \neq i$ for all i (i.e., no firm keeps its original shock), and define the assigned shock for firm i as $\Delta\tilde{\mu}_{i,c,w}^{(r)} = \Delta\hat{\mu}_{\pi_{c,w}^{(r)}(i),c,w}$ for $w \in \{0, +1, +2\}$. Using these reshuffled firm shocks and the same lagged market share weights as in the baseline, I construct a synthetic granular measure for each year and re-estimate the GDP per capita growth regression with the identical specification. This procedure yields a distribution of placebo R^2 values, capturing the macro-level explanatory power that could arise purely from randomly reassigned (non-informative) shocks.

Figure 3 shows the resulting histogram. The dashed line marks the adjusted R^2 from the real (non-shuffled) granular measure. For the specification with current and one lag terms (left panel), only 2.3% of the placebo regressions produced a higher R^2 than the actual one, and this is 2.6% for the specification with two lags. Put differently, around 97.5% of the synthetic granular measures performed worse than the baseline in explaining aggregate variations.

First, this implies that the measured granularity is not a mechanical outcome of the shock distribution. Second, the correct alignment between firm and shock matters, indicating an underlying structure rather than randomness. Taken together, these results strengthen the interpretation that the granularity is driven by real shocks arising from systematic firm-level relationships rather than random noise.

(ii) Stochastic perturbation of shock values. Suppose that the estimated M&A shocks contain some degree of random noise, such that $\Delta\hat{\mu}_{i,c,w} = \Delta\mu_{i,c,w}^{\text{true}} + v_{i,c,w}, \quad$ for $w = 0, +1, +2$, where $v_{i,c,w} \sim$ i.i.d. $\mathcal{N}(0, \sigma_\nu^2)$. To test whether the observed granularity persists under noisier conditions, I simulate shock contamination by injecting Gaussian noise. For each simulation

$r = 1, \ldots, R$ (with $R = 1000$), I generate a new noise-added shock, $\Delta\tilde{\mu}^{(r)}_{i,c,w} = \Delta\hat{\mu}_{i,c,w} + \epsilon^{(r)}_{i,c,w}$, where $\epsilon^{(r)}_{i,c,w} \sim \mathcal{N}\left(0, (0.25 \cdot \hat{\sigma}_{c,w})^2\right)$, and $\hat{\sigma}_{c,w}$ is the cross-sectional standard deviation of $\Delta\hat{\mu}_{i,c,w}$ for cohort c at event time w.[20]

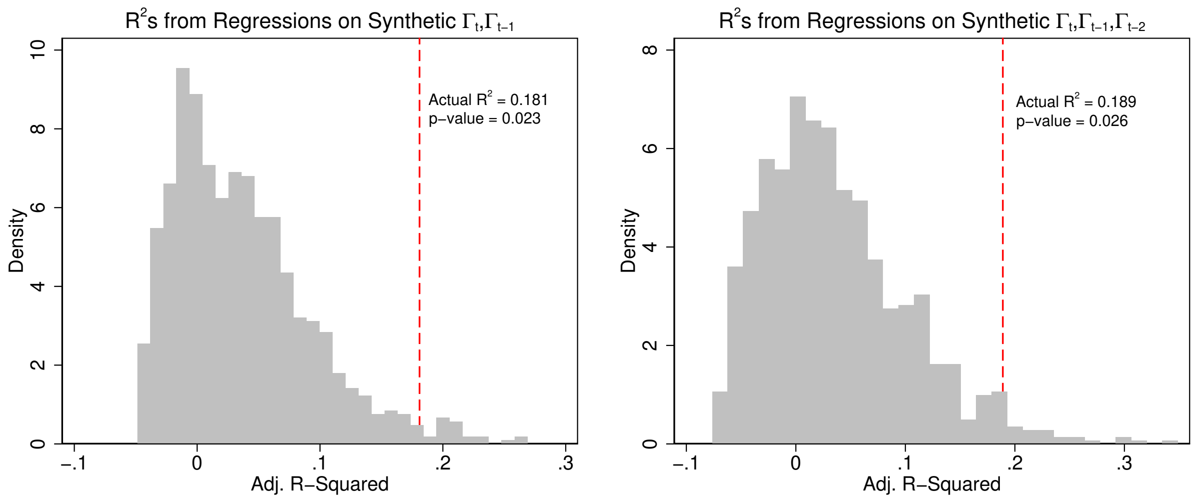

Figure 3: R^2 Distribution from Regressions Using Synthetic Granular Measures of Shocks

Notes: The figure shows the distribution of adjusted R^2 from regressions using synthetic granular measures constructed from randomly reassigning firm-level M&A shocks within each cohort-year. The dashed lines indicate the adjusted R^2 from the baseline regression using the true shock alignment. p-values denote the share of simulations with higher R^2 than the baseline.

From the distorted shocks, I construct granular measures and use them in the regression of per capita growth. If the explanatory power were to collapse too easily, it would suggest that the original M&A shocks are not informative to begin with. However, as shown in Figure 4, the adjusted R^2 values cluster somewhere close to the baseline value, suggesting that the granular measure is not overly sensitive to modest disturbances in the shock estimates.

I extend this exercise using more aggressive levels of noise (reported in Online Appendix B, see Figure Bx.4), and find that heavier contamination consistently attenuates the explanatory power of M&A shocks, as expected. This suggests that the adjusted R^2 values in the main specifications may represent a conservative estimate of the true macroeconomic variation driven by M&A activity, if the granular measure contains some degree of random noise.

20. This introduces mild disturbance, sufficient to test robustness without overwhelming the signal.

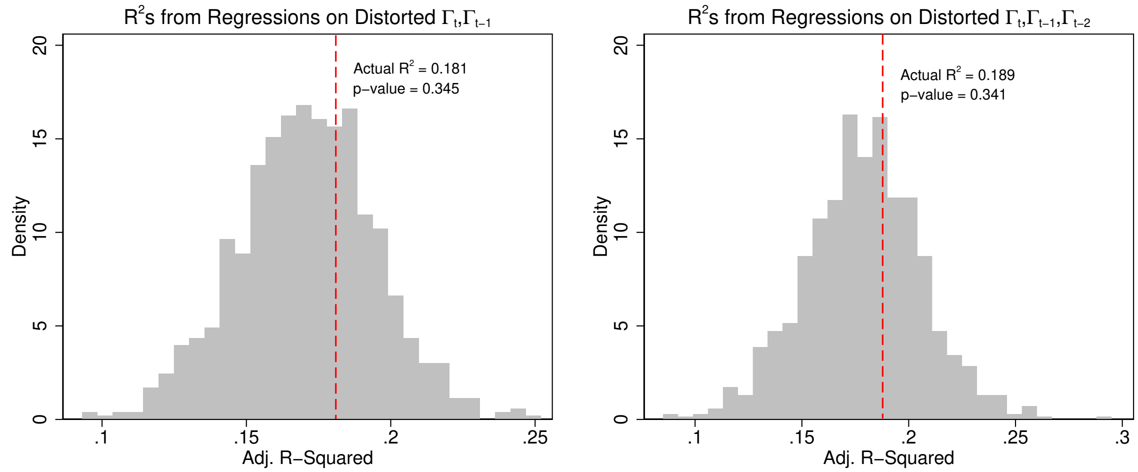

Figure 4: R^2 Distribution from Regressions Using Granular Measures of Distorted Shocks

Notes: The figure shows the distribution of adjusted R^2 from regressions using granular measures based on noise-added M&A shocks. The dashed lines indicate the adjusted R^2 from the baseline regression using the true shock profiles. p-values denote the share of simulations with higher R^2 than the baseline.

5.3. Predictive Power of M&A Shocks

Thus far, I have examined the explanatory power of the granular measure. I now turn to its predictive content—that is, whether lagged values of the granular M&A shock measure can help forecast future growth. To evaluate this, I estimate a series of regressions in which current output growth per capita is regressed on lags of the granular measure, conditional on its own lags and standard macroeconomic predictors (e.g., interest rates, the term spread, and energy price shocks).

The results (Table Bx.6 in Online appendix B) show that the first lag of the granular measure consistently enters with a statistically significant coefficient across all specifications, whereas the second lag remains insignificant. Eventually, the first lag alone explains approximately one-tenth of variations in the next year's growth. Furthermore, when included alongside standard predictors, the granular measure largely preserves its predictive power. These findings suggest that firm-level M&A dynamics contain valuable information for anticipating next-year growth and may warrant greater attention in policy monitoring.

6. Conclusion

This paper investigates how mergers and acquisitions (M&A) contribute to aggregate growth fluctuations. M&A is an extraordinary event for most firms and, when undertaken, is seldom smooth. Yet systematic evidence on firm-level productivity dynamics around M&A is scarce, and their aggregate implications are even less understood. This paper addresses that gap through the lens of the granular hypothesis.

My empirical analysis draws on M&A transaction data from SDC Platinum, combined with financial data from Compustat. I focus on completed domestic deals involving U.S. public firms between 1978 and 2023. I identify M&A-driven productivity fluctuations at the firm level by estimating a latent factor (interactive fixed-effects) model on a cohort-based panel under the potential outcomes framework.

On average, firms exhibit an adjustment pattern following M&A: a sharp productivity decline in the merger year, followed by a recovery within two years that lifts productivity above pre-merger levels. Despite this persistence in levels, the shocks are growth-rate neutral over the longer horizon and resemble level-shifting temporary impulses. The estimates also reveal heterogeneity in both the depth of the initial drop and the pace of recovery, depending on pre-merger firm size.

A granular measure, constructed as the firm-size-weighted average of these shocks, explains roughly one-fifth of variations in output (per capita) growth. Its explanatory power is largely robust across a range of checks, including controls for post-merger changes in firm size, exposure to merger activity (i.e., merger waves), as well as various stress and placebo tests. Moreover, the contribution of M&A shocks to aggregate fluctuations remains significant when conditioning on standard macroeconomic variables.

The findings indicate that the granularity associated with M&A shocks is not spuriously driven; these firm-level events represent a distinct and measurable source of aggregate fluctuations. Additionally, these shocks, when aggregated, demonstrate not only explanatory power but also predictive power for near-term future growth.

Finally, by extending the analysis of firm-level shocks beyond purely random movements to include tractable or observable firm decisions, this study broadens the scope of the granular origins of macroeconomic fluctuations. Thus, the findings of this study offer practical insights for macroeconomic modeling that incorporate heterogeneous firm structures, emphasizing the role of micro-level strategic behavior in shaping aggregate outcomes. Future research could further expand this perspective by examining other forms of strategic firm behavior. Examples that come to mind include large-scale capital replacement events (i.e., investment spikes) or employee layoffs at the firm level, both of which often involve immediate and irreversible changes and may leave measurable traces in aggregate data.

References

Ahn, Seung C, and Alex R Horenstein. 2013. "Eigenvalue Ratio Test for the Number of Factors." *Econometrica* 81 (3): 1203–1227.

Axtell, Robert L. 2001. "Zipf Distribution of US Firm Sizes." *Science* 293 (5536): 1818–1820.

Bai, Jushan. 2009. "Panel Data Models with Interactive Fixed Effects." *Econometrica* 77 (4): 1229–1279.

Becher, David, Tyler K Jensen, and Tingting Liu. 2020. "Acquisitions and Funding Conditions." *Journal of Corporate Finance* 65:101760.

Bena, Jan, and Kai Li. 2014. "Corporate Innovations and Mergers and Acquisitions." *The Journal of Finance* 69 (5): 1923–1960.

Blonigen, Bruce A, and Justin R Pierce. 2016. *Evidence For the Effects of Mergers on Market Power and Efficiency,* Working Paper 22750. National Bureau of Economic Research.

Bradley, Michael, Anand Desai, and E Han Kim. 1988. "Synergistic gains from corporate acquisitions and their division between the stockholders of target and acquiring firms." *Journal of Financial Economics* 21 (1): 3–40.

Braguinsky, Serguey, Atsushi Ohyama, Tetsuji Okazaki, and Chad Syverson. 2015. "Acquisitions, Productivity, and Profitability: Evidence from the Japanese Cotton Spinning Industry." *American Economic Review* 105 (7): 2086–2119.

Carvalho, Vasco M, and Basile Grassi. 2019. "Large Firm Dynamics and the Business Cycle." *American Economic Review* 109 (4): 1375–1425.

Chan, Jackie ML, and Han Steffan Qi. 2025. "Mergers, Firm Size, and Volatility in a Granular Economy." *Review of Economic Dynamics* 55:101254.

Compustat. 2024. (Compustat North America: Fundamentals Quarterly & Fundamentals Annual. Data creator: S&P Global Market Intelligence).

Cudd, Mike, and Rakesh Duggal. 2000. "Industry Distributional Characteristics of Financial Ratios: An Acquisition Theory Application." *Financial Review* 35 (1): 105–120.

Danielsson, Jon, Lerby Murat Ergun, Laurens de Haan, and Casper G de Vries. 2016. "Tail Index Estimation: Quantile Driven Threshold Selection." *Available at SSRN 2717478.*

Demirer, Mert, and Ömer Karaduman. 2024. *Do Mergers and Acquisitions Improve Efficiency? Evidence from Power Plants,* Working Paper 32727. National Bureau of Economic Research.

Di Giovanni, Julian, Andrei A Levchenko, and Isabelle Mejean. 2014. "Firms, Destinations, and Aggregate Fluctuations." *Econometrica* 82 (4): 1303–1340.

Embrechts, Paul, Claudia Klüppelberg, and Thomas Mikosch. 1997. "Modelling Extremal Events for Insurance and Finance." Springer.

Gabaix, Xavier. 2009. "Power Laws in Economics and Finance." *Annual Review of Economics* 1 (1): 255–294.

———. 2011. "The Granular Origins of Aggregate Fluctuations." *Econometrica* 79 (3): 733–772.

Gabaix, Xavier, and Ralph SJ Koijen. 2024. "Granular Instrumental Variables." *Journal of Political Economy* 132 (7): 2274–2303.

Ghosh, Aloke. 2001. "Does Operating Performance Really Improve Following Corporate Acquisitions?" *Journal of Corporate Finance* 7 (2): 151–178.

Gobillon, Laurent, and Thierry Magnac. 2016. "Regional Policy Evaluation: Interactive Fixed Effects and Synthetic Controls." *Review of Economics and Statistics* 98 (3): 535–551.

Gort, Michael. 1969. "An Economic Disturbance Theory of Mergers." *The Quarterly Journal of Economics* 83 (4): 624–642.

Grullon, Gustavo, Yelena Larkin, and Roni Michaely. 2019. "Are US Industries Becoming More Concentrated?" *Review of Finance* 23 (4): 697–743.

Hamilton, James D. 2003. "What is an oil shock?" *Journal of Econometrics* 113 (2): 363–398.

Harford, Jarrad. 2005. "What drives merger waves?" *Journal of Financial Economics* 77 (3): 529–560.

Healy, Paul M, Krishna G Palepu, and Richard S Ruback. 1992. "Does corporate performance improve after mergers?" *Journal of Financial Economics* 31 (2): 135–175.

Hill, Bruce M. 1975. "A Simple General Approach to Inference About the Tail of a Distribution." *The Annals of Statistics,* 1163–1174.

Holland, Paul W. 1986. "Statistics and Causal Inference." *Journal of the American Statistical Association* 81 (396): 945–960.

Hsiao, Cheng. 2025. "A selective review of panel approaches to construct counterfactuals." *Empirical Economics,* 1–20.

Hsiao, Cheng, H Steve Ching, and Shui Ki Wan. 2012. "A Panel Data Approach for Program Evaluation: Measuring the Benefits of Political and Economic Integration of Hong Kong with Mainland China." *Journal of Applied Econometrics* 27 (5): 705–740.

Jovanovic, Boyan, and Peter L Rousseau. 2002. "The Q-theory of Mergers." *American Economic Review* 92 (2): 198–204.

———. 2008. "Mergers As Reallocation." *The Review of Economics and Statistics* 90 (4): 765–776.

Kim, E. Han, and Vijay Singal. 1993. "Mergers and Market Power: Evidence from the Airline Industry." *American Economic Review* 83 (3): 549–569. Accessed March 24, 2025.

Lang, Larry HP, René M Stulz, and Ralph A Walkling. 1991. "A test of the free cash flow hypothesis: The case of bidder returns." *Journal of Financial Economics* 29 (2): 315–335.

Levonyan, Vardges, and Paolo Mengano. 2022. *Mergers and Markups.* Available at SSRN: https://ssrn.com/abstract=4466826 or http://dx.doi.org/10.2139/ssrn.4466826.

LSEG/Refinitiv. 2024. (SDC Platinum: Mergers & Acquisitions (U.S.) Accessed via Refinitiv Eikon with permission of the Marshall Library, University of Cambridge).

Maksimovic, Vojislav, and Gordon Phillips. 2001. "The Market for Corporate Assets: Who Engages in Mergers and Asset Sales and Are There Efficiency Gains?" *The Journal of Finance* 56 (6): 2019–2065.

Mitchell, Mark L, and J Harold Mulherin. 1996. "The impact of industry shocks on takeover and restructuring activity." *Journal of Financial Economics* 41 (2): 193–229.

Philippon, Thomas. 2019. "The Great Reversal: How America Gave Up On Free Markets." Harvard University Press.

Phillips, Gordon M, and Alexei Zhdanov. 2013. "R&D and the Incentives from Merger and Acquisition Activity." *The Review of Financial Studies* 26 (1): 34–78.

Powell, Ronan G. 1997. "Modelling Takeover Likelihood." *Journal of Business Finance & Accounting* 24 (7-8): 1009–1030.

Ravenscraft, David J, and Frederic M Scherer. 1989. "The Profitability of Mergers." *International Journal of Industrial Organization* 7 (1): 101–116.

Rhodes-Kropf, Matthew, and Steven Viswanathan. 2004. "Market Valuation and Merger Waves." *The Journal of Finance* 59 (6): 2685–2718.

Romer, Christina D, and David H Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review* 94 (4): 1055–1084.

Rubin, Donald B. 1974. "Estimating causal effects of treatments in randomized and nonrandomized studies." *Journal of Educational Psychology* 66 (5): 688.

Samartsidis, Pantelis, Shaun R Seaman, Silvia Montagna, André Charlett, Matthew Hickman, and Daniela De Angelis. 2020. "A Bayesian multivariate factor analysis model for evaluating an intervention by using observational time series data on multiple outcomes." *Journal of the Royal Statistical Society Series A: Statistics in Society* 183 (4): 1437–1459.

Samartsidis, Pantelis, Shaun R. Seaman, Anne M. Presanis, Matthew Hickman, and Daniela De Angelis. 2019. "Assessing the Causal Effect of Binary Interventions from Observational Panel Data with Few Treated Units." *Statistical Science* 34 (3).

Schoar, Antoinette. 2002. "Effects of Corporate Diversification on Productivity." *The Journal of Finance* 57 (6): 2379–2403.

Xu, Jianhuan. 2017. "Growing Through the Merger and Acquisition." *Journal of Economic Dynamics and Control* 80:54–74.

Xu, Yiqing. 2017. "Generalized synthetic control method: Causal inference with interactive fixed effects models." *Political Analysis* 25 (1): 57–76.

ONLINE APPENDIX

The Granular Role of Mergers and Acquisitions in Aggregate Fluctuations

Seda Basihos

University of Cambridge, E-mail: sb2626@cam.ac.uk

Appendix A Descriptions and Summary of Data

Table Ax.1: Means, Standard Deviations, and Observations of Base-Period Differenced log of Productivity

	Treat		Overall
	0	1	
$\tau = -2$	-0.03277	-0.04214	-0.03338
	[0.41342]	[0.32131]	[0.40808]
	92,036	6,388	98,424
$\tau = -1$	0.00000	0.00000	0.00000
	[0.00000]	[0.00000]	[0.00000]
	97,631	6,752	104,383
$\tau = 0$	0.02136	-0.03740	0.01756
	[0.43818]	[0.36749]	[0.43420]
	97,631	6,752	104,383
$\tau = +1$	0.03749	0.03353	0.03724
	[0.51996]	[0.40009]	[0.51299]
	87,444	6,106	93,550
$\tau = +2$	0.05105	0.05200	0.05111
	[0.56633]	[0.44174]	[0.55910]
	79,353	5,497	84,850
Overall	0.01409	-0.00099	0.01311
	[0.42997]	[0.34105]	[0.42478]
	454,095	31,495	485,590

Notes: The table presents basic statistics for the base-period differenced log of productivity by pre-merger years, $w = -2, -1$, the merger event year, $w = 0$, and post-merger years, $w = +1, +2$, and treatment status. The base period is $w = -1$. Sample means are reported directly, with standard deviations shown in brackets. The sample covers the period 1978-2023.

Main Variables and Financial Ratios

** Some preliminary notes for readers: All firm-level data on financial ratios and productivity are sourced from Compustat. If any of the variables listed below have missing values, or if an observation is dropped for any of the reasons outlined, the entire firm-year observation is considered as missing and excluded from the final dataset.

Labor Productivity. This is measured as the logarithm of the ratio of sales revenue to the number of employees. Observations with non-positive values for either variable are excluded. In the data, there may be reporting errors, especially in employee counts, as some firms report rough estimates based on approximate turnover rates. To mitigate the influence of such misreporting, I winsorize the sales-to-employee ratio at the 1% tails of its yearly distribution in the raw data.

Market Share. This is calculated as a firm's sales revenue (inflation-adjusted using the implicit GDP deflator) divided by real GDP (2017 constant prices). Although public firms follow standardized reporting, extreme values or occasional inconsistencies in revenue reporting may still occur. To mitigate the influence of such outliers while preserving the distribution's fat tails, I winsorize market share at the 0.5% tails of its yearly distribution in the raw data. This is applied regardless of whether market share is constructed using current or lagged sales.

Capital Structure Ratio. This measures a firm's leverage or capital mix between debt and equity. A lower capital ratio (i.e., more equity relative to debt) suggests that the firm has room to take on additional debt. For example, Becher, Jensen, and Liu (2020) finds that firm leverage is a significant determinant of M&A decisions for acquirers during periods of favorable aggregate funding conditions (see also Krishnan and Yakimenko, 2022).

Cash Ratio. This ratio measures how easily a firm can meet its short-term obligations using only the most liquid assets (cash and near-cash). Several empirical studies find that liquidity is an important determinant of both M&A decision and post-merger financial performance (see, e.g., Lang, Stulz, and Walkling, 1991; Oler, 2008; Duchin, 2010; Chen et al., 2020).

Gross Profitability. A higher gross margin means the firm retains more revenue after covering the direct cost of production, which indicates strong pricing power, cost efficiency, or both. This measure follows an accounting-based approach to assessing profitability. While the literature evaluates profitability using both accounting measures and stock return-based metrics, prior research consistently finds that acquirers tend to be highly profitable in the pre-merger period (see, e.g., Ravenscraft and Scherer, 1989; Healy, Palepu, and Ruback, 1992; Brar, Giamouridis, and Liodakis, 2009). Note (1): This metric shows extremely large negative values in the far left tail of its distribution. I trim the bottom 1.5% of firm-cohort observations to remove these outliers. Since this is a trimming procedure (not winsorization), it is applied after cohort construction to preserve the timing information of M&A events at the firm level. Following this step, any firm lacking data at either the merger event year or the baseline period is fully excluded from the respective cohort.

P/E Ratio. This measure is the valuation of firm stocks. Shleifer and Vishny (2003) argues that firms with overvalued stock are more likely to acquire. Moreover, previous research documents that firm-specific and market-wide over- or under-valuation is an important determinant of M&A waves (see, e.g., Rhodes-Kropf and Viswanathan, 2004; Dong et al., 2006; Bouwman, Fuller, and Nain, 2009). Note (1): the denominator of this ratio uses diluted earnings per share (EPS), which includes all potential shares from options, convertible debt, etc. Note (2): This metric shows extremely large negative values in the far left tail of its distribution. I trim the bottom 1.5% of firm-cohort observations to remove these outliers. Since this is a trimming procedure (not winsorization), it is applied after cohort construction to preserve the timing information of M&A events at the firm level. Following this step, any firm lacking data at either the merger event year or the baseline period is fully excluded from the respective cohort.

R&D Intensity. This is measured as the ratio of R&D expenditures to sales revenue. Negative values are dropped for this ratio, but zero values are kept. If R&D expenditure is not reported, the value is considered as missing. Prior research suggests that larger firms may engage in M&A activity as a means of outsourcing R&D investment, often by acquiring more innovative firms (see, e.g., Phillips and Zhdanov, 2013; Bena and Li, 2014).

Capital Intensity. To construct firm-level capital intensity, I first compute the capital stock for each firm using the perpetual inventory method (PIM), where capital stock is given by $K_{it} = \left((1-\delta_{KJ})\frac{K_{it-1}}{P^I_{it-1}} + \frac{I_{it}}{P^I_{it}}\right) P^I_{it}$. The initial capital stock is set using the earliest available observation of the firm's net book value of capital (PPENT). δ_{KJ} are 2-digit industry-level economic depreciation rates from the Bureau of Economic Analysis (BEA), and I use the period average of these rates. Investment I_{it} is capital expenditures (CAPX) (see, e.g., Bloom, 2009 and Rajan and Zingales, 1998). The investment deflator P^I_{it} is the Quality-Adjusted Relative Price of Investment Goods (DiCecio, 2009), with 2017 as the base year. Finally, capital intensity is measured as the logarithm of the ratio of constructed capital stock to the number of employees in the firm.

Investment Rate. The investment rate is measured as capital expenditure scaled by lagged PIM-constructed capital, $\frac{I_{it}}{K_{it-1}}$. This ratio is winsorized at 1% tails of yearly observations to exclude any negative construction or extreme values.

Table Ax.2: Descriptions for Variables used in Firm Matching and Factor Analysis Model

Variable	Category	Acronym	Definition
Labor Productivity	Productivity		Sales Revenue (2017 prices) per Employee (log)
Market Share	Firm Size		Sales-Based Market Share in GDP
Capital Structure Ratio	Capitalization	capital_ratio	Total Long-term Debt as a fraction of the sum of Total Long-term Debt, Common/Ordinary Equity and Preferred Stock
Cash Ratio	Liquidity	cash_ratio	Cash and Short-term Investments as a fraction of Current Liabilities
Gross Profitability	Profitability		Gross Margin (Sales – COGS) as a fraction of Total Sales
P/E Ratio	Valuation	pe_exi	Price-to-Earnings, excl. Extraordinary Items (diluted)
R&D Intensity	Other	RD_SALE	R&D expenses as a fraction of Sales
Capital Intensity	Other		Net Property, Plant & Equipment (PP&E) per Employee (log)
Investment Rate	Other		Gross Investment adjusted by previous year PP&E

Notes: All financial ratios with Compustat acronyms are accessed quarterly and converted to the annual level by taking averages.

Table Ax.3: Descriptive Statistics for Financial Ratios and Productivity

	Mean	Std. Deviation	Observation
Labor Productivity (log)	6.065	0.887	104,383
Market Share (%)	0.0109	0.0371	104,383
Capital Structure Ratio	0.408	77.898	104,383
Cash Ratio	1.535	9.838	104,383
Gross Profitability	0.0635	2.717	104,383
P/E Ratio	3.827	100.210	104,383
R&D Intensity	0.0068	0.460	104,383
Capital Intensity (log)	3.857	1.504	104,383
Investment Rate	0.374	0.760	104,383

Notes: The table reports summary statistics for firm-level ratios used as proxies for pre-merger characteristics. These ratios are averaged over event times $w = -2, -1$ (using only $w = -1$ if data for $w = -2$ is missing) and are time-invariant within each cohort, meaning the same values are used across all firm–w pairs within a given cohort. Reported means and standard deviations are calculated across all firm–cohort observations.

Appendix B Additional Tables and Figures

Table Bx.1: Estimates for M&A Effects on Productivity – Full Sample and by Firm Size

	Full Sample	Large Firms	Medium Firms	Small Firms
	(1)	(2)	(3)	(4)
-2	-0.00450	-0.00573	0.00197	-0.00712
	(0.00450)	(0.00580)	(0.00520)	(0.00518)
0	-0.04887***	-0.01997**	-0.0511***	-0.06236***
	(0.00718)	(0.00846)	(0.00828)	(0.01090)
+1	0.02369***	0.02212**	-0.00694	0.03978***
	(0.00529)	(0.00964)	(0.00996)	(0.00774)
+2	0.03210***	0.02453*	0.00196	0.05102***
	(0.00925)	(0.01315)	(0.01137)	(0.01500)
N (all)	485,590			
N (treated)	31,495	7,956	7,841	15,698

Notes: One year before M&A ($w = -1$) is taken as the reference category. Robust standard errors (clustered at the cohort level) are reported in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

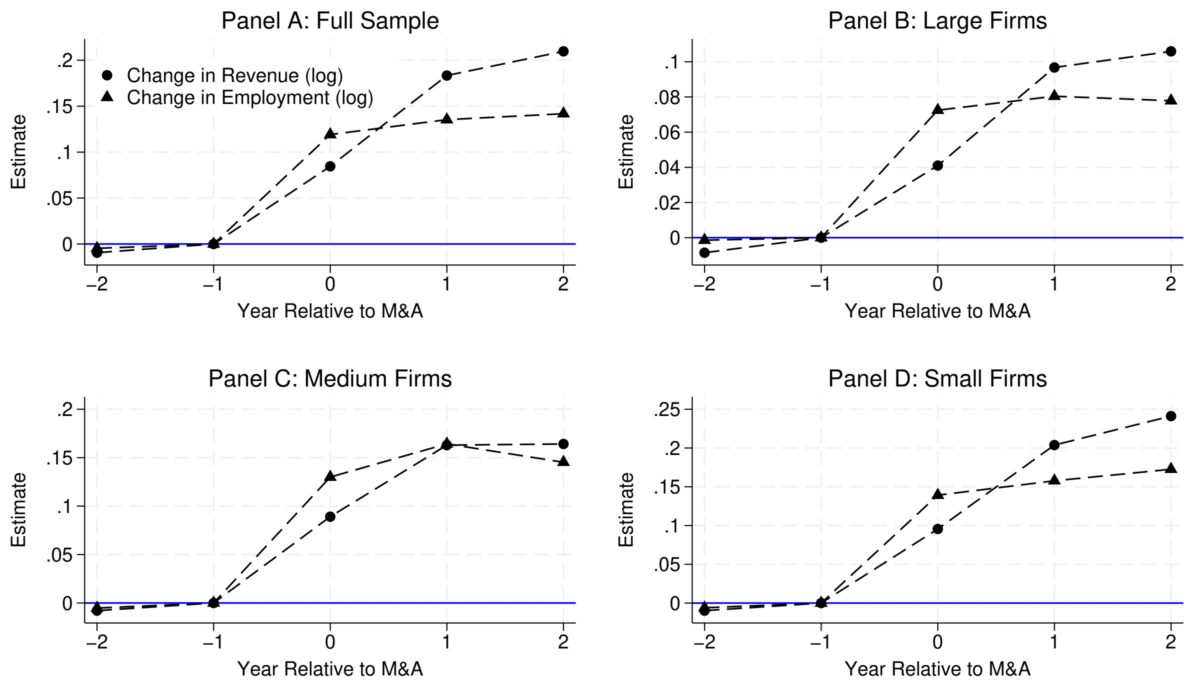

Figure Bx.1: Event-Time Effects of M&A on Log Revenue and Log Employment

Notes: Panels show average M&A effects on log revenue and log employment for the full sample and by firm-size category.

Table Bx.2: Firm Size Mean Difference – Treated (1) vs Placebo Treated (2)

Group	Obs	Mean	Std. Dev.	95% CI	
				Lower	Upper
1	31,495	0.02451	0.06100	0.02383	0.02518
2	31,482	0.02251	0.06361	0.02181	0.02321
Combined	62,977	0.02351	0.06232	0.02302	0.02400
Difference		0.00200		0.00103	0.00297

Note: The table reports a two-sample t-test with equal variances comparing the mean firm size between treated firms and placebo-treated firms in the sample.

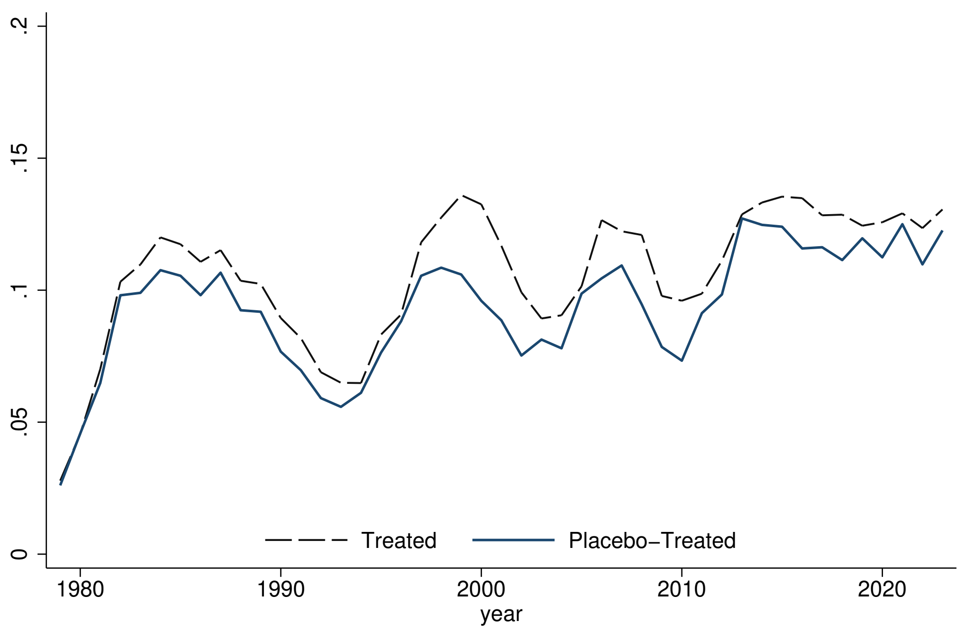

Figure Bx.2: Market Shares of Treated and Placebo-Treated Firms

Notes: The figures reflect data for acquiring firms (treated), defined as those in their merger event year or in either of the two post-merger years, provided any of these periods fall within the given calendar year. The same approach is applied to placebo-treated firms, considering them as if they were in their merger or post-merger period.

Table Bx.3: Results for Placebo-Treated Firms

	RCM Results (1)	Observed Data (2)
-2	-0.0007184	-0.027443***
	(0.0064363)	(0.0055931)
0	0.0089206	0.0101272*
	(0.0065909)	(0.0060151)
+1	0.0028209	0.0164045**
	(0.0069582)	(0.0081133)
+2	0.0018819	0.0239341**
	(0.0094231)	(0.0098381)
N (all)	454,095	
N (treated)	31,482	

Notes: RCM refers to the Rubin Causal Model estimates, where both observed and latent common constructs are extracted from the productivity changes of placebo-treated firms. The observed data provide the direct mean and standard deviation of productivity changes for those firms. One year before M&A ($w = -1$) is taken as the reference category. Robust standard errors (clustered at the cohort level) are reported in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Table Bx.4: Tests for Serial Correlation in Granular Measure Γ_t

Lag	Breusch-Godfrey χ^2	p-value	Ljung-Box Q	p-value	Conclusion
1	0.004	0.9514	0.2747	0.6002	No AC
2	0.207	0.9016	1.3129	0.5187	No AC
3	0.481	0.9231	2.0337	0.5655	No AC
4	0.560	0.9674	3.3827	0.4959	No AC
5	6.475	0.2627	3.3968	0.6391	No AC

Notes: Breusch–Godfrey test: H_0: no serial correlation in residuals at the specified lag order. Ljung–Box test: H_0: no autocorrelation up to the specified lag order. AC = Autocorrelation.

Table Bx.5: Baseline and Placebo: Adjusted R^2 and Long-run Multipliers

	Baseline			Placebo		
$\sum_{l=0}^{1} \alpha_l \Gamma_{t-l}$	$l \in \{0,1\}$	$l \in \{0,1,2\}$	$l \in \{0,1,2,3\}$	$l \in \{0,1\}$	$l \in \{0,1,2\}$	$l \in \{0,1,2,3\}$
	(1)	(2)	(3)	(4)	(5)	(6)
Adj. R^2	0.181	0.189	0.222	0.073	0.073	0.095
Long-run	0.095	-0.405	-1.164	0.972	0.457	-0.192
Multiplier	(0.639)	(0.820)	(0.946)	(0.766)	(0.945)	(1.050)
Observation	44	43	42	44	43	42

Notes: The sample period is 1979–2023. The dependent variable is GDP per capita growth. The index l denotes the lag specification. "Baseline" columns report regressions on the granular measure constructed from M&A-driven shocks, while "Placebo" columns use a granular measure based on random idiosyncratic shocks to placebo-treated firms. Standard errors are reported in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

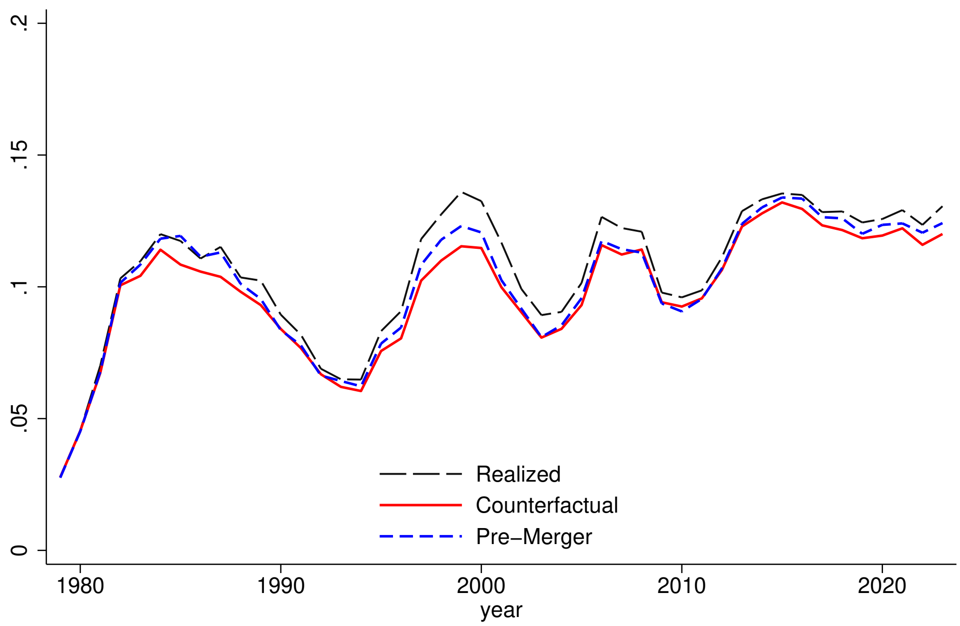

Figure Bx.3: Realized and Counterfactual Market Shares of Acquirers

Notes: The figures reflect data for acquiring firms, defined as those in their merger event year or in either of the two post-merger years, provided any of these periods fall within the given calendar year. *Realized* refers to the actual total sales of acquirers as a share of GDP in that year. *Pre-Merger* refers to what their share of GDP would have been based on their sales before the merger. *Counterfactual* denotes the factor-model estimate of the market share the firms would have had, absent M&A.

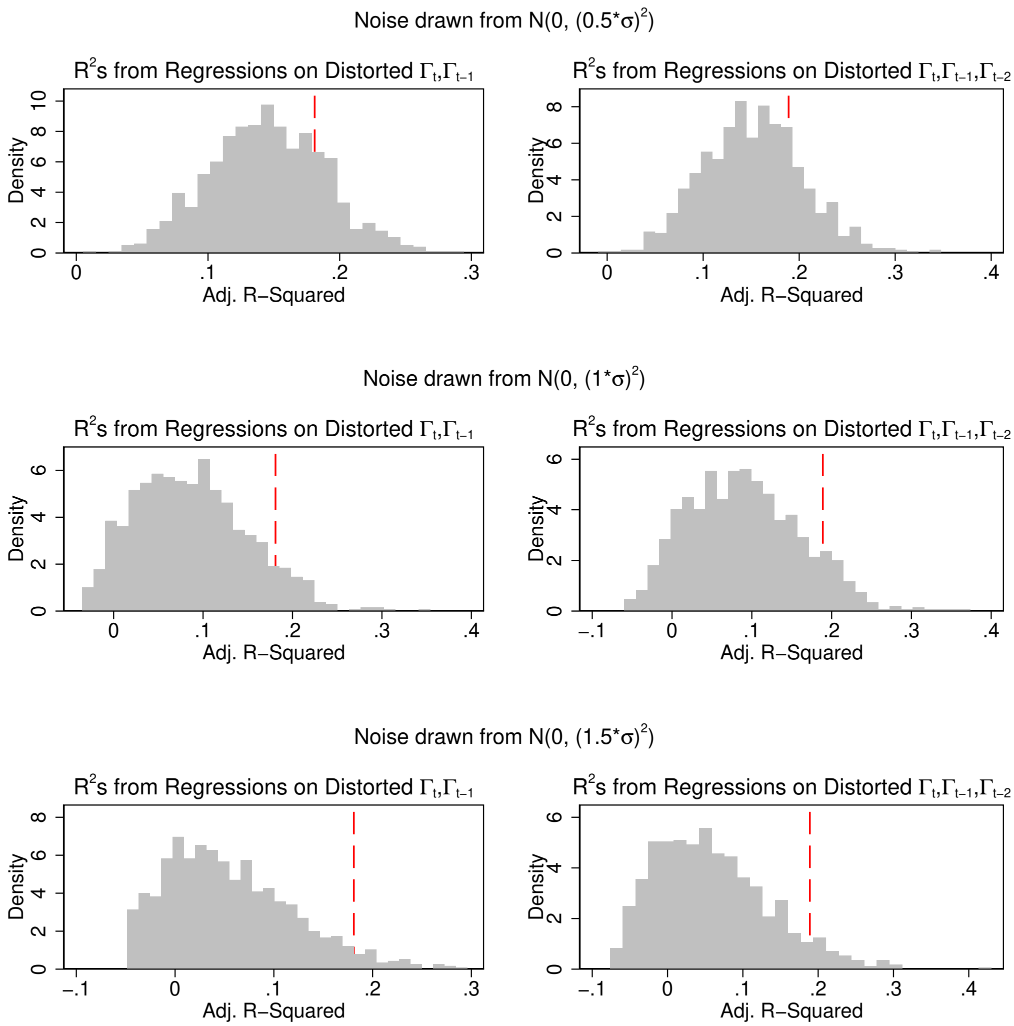

Figure Bx.4: R^2 Distribution from Regressions Using Granular Measures of More Heavily Distorted Shocks

Notes: This figure shows the distribution of adjusted R^2 values from regressions using granular measures constructed from noise-added M&A shocks. Noise is drawn from $\mathcal{N}(0, \sigma^2)$ with standard deviations set to 0.5, 1.0, and 1.5 times the cross-sectional standard deviation of the estimated shocks. The dashed lines indicate the adjusted R^2 from the baseline regression using the true shock profiles.

Table Bx.6: Predictive Power of the Granular Measure with Macroeconomic Controls

	(1)	(2)	(3)	(4)	(5)	(6)
Γ_{t-1}	-1.036**		-1.017**	-1.071**		-0.942*
	(0.459)		(0.433)	(0.447)		(0.478)
Γ_{t-2}				-0.552		-0.220
				(0.447)		(0.576)
g_{t-1}		0.045	0.164		-0.035	0.051
		(0.132)	(0.135)		(0.181)	(0.207)
g_{t-2}					0.158	0.151
					(0.152)	(0.151)
r_{t-1}		0.064	0.077		0.048	0.036
		(0.093)	(0.088)		(0.156)	(0.158)
r_{t-2}					-0.004	0.055
					(0.166)	(0.165)
Term Spread$_{t-1}$		0.476*	0.512**		0.552	0.493
		(0.262)	(0.249)		(0.368)	(0.360)
Term Spread$_{t-2}$					-0.336	-0.130
					(0.396)	(0.422)
Oil$_{t-1}$		-0.00048**	-0.00036*		-0.00054**	-0.00035
		(0.00019)	(0.00019)		(0.00022)	(0.00026)
Oil$_{t-2}$					-0.00033	-0.00022
					(0.00026)	(0.00026)
Intercept	0.0165***	0.0159***	0.0124**	0.0172***	0.0202***	0.0147*
	(0.0028)	(0.0049)	(0.0048)	(0.0027)	(0.0069)	(0.0073)
N	44	44	44	43	43	43
R^2	0.108	0.332	0.416	0.146	0.352	0.422
Adj. R^2	0.087	0.263	0.340	0.104	0.199	0.241

Notes: The study period is 1979-2023. The dependent variable is GDP per capita Growth. g_{t-1}, g_{t-2} denote lagged values of per capita growth, r_{t-1}, r_{t-2} denote lagged values of real interest rate, and Oil$_{t-1}$, Oil$_{t-2}$ refer to lagged energy price shocks. Γ_{t-1}, Γ_{t-2} denote the lagged values of the granular measure of M&A-driven shocks. Standard errors are reported in parentheses. * p<0.10, ** p<0.05, *** p<0.01.

References

Becher, David, Tyler K Jensen, and Tingting Liu. 2020. "Acquisitions and Funding Conditions." *Journal of Corporate Finance* 65:101760.

Bena, Jan, and Kai Li. 2014. "Corporate Innovations and Mergers and Acquisitions." *The Journal of Finance* 69 (5): 1923–1960.

Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 (3): 623–685.

Bouwman, Christa HS, Kathleen Fuller, and Amrita S Nain. 2009. "Market Valuation and Acquisition Quality: Empirical Evidence." *The Review of Financial Studies* 22 (2): 633–679.

Brar, Gurvinder, Daniel Giamouridis, and Manolis Liodakis. 2009. "Predicting European Takeover Targets." *European Financial Management* 15 (2): 430–450.

Chen, Jia, Ya-Chun Gao, Qiang Li, and Yong Zeng. 2020. "Cash holdings, M&A Decision and Risk Premium." *Physica A: Statistical Mechanics and Its Applications* 537:122571.

Compustat. 2024. (Compustat North America: Fundamentals Quarterly & Fundamentals Annual. Data creator: S&P Global Market Intelligence).

DiCecio, Riccardo. 2009. "Sticky wages and sectoral labor comovement." *Journal of Economic Dynamics and Control* 33 (3): 538–553.

Dong, Ming, David Hirshleifer, Scott Richardson, and Siew Hong Teoh. 2006. "Does Investor Misvaluation Drive the Takeover Market?" *The Journal of Finance* 61 (2): 725–762.

Duchin, Ran. 2010. "Cash Holdings and Corporate Diversification." *The Journal of Finance* 65 (3): 955–992.

Healy, Paul M, Krishna G Palepu, and Richard S Ruback. 1992. "Does corporate performance improve after mergers?" *Journal of Financial Economics* 31 (2): 135–175.

Krishnan, CNV, and Vasiliy Yakimenko. 2022. "Market Misreaction? Leverage and Mergers and Acquisitions." *Journal of Risk and Financial Management* 15 (3): 144.

Lang, Larry HP, René M Stulz, and Ralph A Walkling. 1991. "A test of the free cash flow hypothesis: The case of bidder returns." *Journal of Financial Economics* 29 (2): 315–335.

Oler, Derek K. 2008. "Does acquirer cash level predict post-acquisition returns?" *Review of Accounting Studies* 13:479–511.

Phillips, Gordon M, and Alexei Zhdanov. 2013. "R&D and the Incentives from Merger and Acquisition Activity." *The Review of Financial Studies* 26 (1): 34–78.

Rajan, Raghuram G., and Luigi Zingales. 1998. "Financial Dependence and Growth." *American Economic Review* 88 (3): 559–586.

Ravenscraft, David J, and Frederic M Scherer. 1989. "The Profitability of Mergers." *International Journal of Industrial Organization* 7 (1): 101–116.

Rhodes-Kropf, Matthew, and Steven Viswanathan. 2004. "Market Valuation and Merger Waves." *The Journal of Finance* 59 (6): 2685–2718.

Shleifer, Andrei, and Robert W Vishny. 2003. "Stock Market Driven Acquisitions." *Journal of Financial Economics* 70 (3): 295–311.