

Technology M&As and Knowledge Diffusion¹

Zili Yang²

November 10, 2025

Abstract: This paper examines how technology mergers and acquisitions (tech M&As) affect the diffusion of target firms' pre-acquisition innovations in the United States. Using US patent and M&A data from 1980 to 2021. This study employs a difference-in-differences approach comparing successful acquisitions with exogenously failed deals; it finds that tech M&As significantly increase external diffusion of targets' technologies, as measured by patent citations, with effects concentrated within the acquirer's industry. Tech M&As do not diminish young firms' ability to cite and build upon acquired targets' patents, contradicting concerns about innovation foreclosure. To interpret these findings and quantify aggregate implications, I develop an idea flow model where firms improve productivity by choosing innovation intensity based on potential targets' technologies, with acquisitions affecting both the innovation step size in learning from targets and the cost of accessing them. The model, calibrated to the empirical estimates and US innovation data, reveals that doubling the 2015 tech M&A rate would increase annual productivity growth by five hundredths of a percentage point, with the diffusion channel contributing 40% of this increase. Surprisingly, relaxing restrictions on post-acquisition knowledge appropriation yields negligible growth effects: reduced spillovers from acquired targets are offset by increased innovation using independent technologies as acquisition values rise. These findings underscore the importance of incorporating diffusion effects and general equilibrium forces into antitrust policy for tech M&As.

JEL Codes: O31, O33, G34, L40, D24

Keywords: technology acquisitions, knowledge diffusion, patent citations, innovation spillovers, productivity growth, M&A

¹ **Acknowledgments:** I am especially grateful to my advisor Caroline Betts for invaluable guidance and support throughout this project. I also thank my committee members Pablo Kurlat and Thomas Chaney for their insightful comments and suggestions. Gerald Hoberg, Monica Morlacco, Daniel Sokol, and Florenta Teodoridis provided helpful guidance. I benefited from comments from the USC Econ Macro Reading Group participants (Marianne Andries, Robert Dekle, and Michael Droste) and conference participants at the Southern Economic Association 94th Annual Meeting in Washington, DC. All errors are my own.

² Department of Economics, University of Southern California, ziliyang@usc.edu

1. Introduction

When Google acquired Android for \$50 million in 2005, the small startup's mobile operating system technology transformed the smartphone industry, allowing innovation across thousands of companies worldwide. In contrast, when NVIDIA attempted to acquire Arm Ltd. in 2021, regulators blocked the \$40 billion deal, citing concerns that the acquisition would restrict competitors' access to Arm's essential processor designs (Federal Trade Commission, 2021). These cases demonstrate a central tension in technology mergers and acquisitions (tech M&As): such deals can accelerate or impede the diffusion of innovation through the economy. While tech M&A activity reached record levels exceeding \$1 trillion in deals in 2021 alone (M&ACommunity, 2024), the net effect of these transactions on innovation diffusion remains unclear; this matters for economic growth. Technological progress depends not only on creating innovations but also on knowledge spillovers that enable firms to learn from and advance upon existing technologies through cumulative innovation.

This paper addresses the following questions. First, do technology acquisitions increase or reduce the diffusion of targets' technologies to other firms, empirically and theoretically? Second, how much measured diffusion affects matter for aggregate productivity growth? This study attempts to comprehensively analyze tech M&As' role in shaping technology diffusion networks and their broader implications for economic development.

To identify the causal effect of acquisitions on technology diffusion, this study employs a difference-in-differences design using US patent and M&A data from 1980 to 2021. This paper compares changes in external citation patterns between patents attributable to firms acquired through M&A and similar patents from non-acquired firms, before and after acquisition events. Explicitly, it tracks whether third-party firms are more or less likely to cite or build upon patents of acquired firms relative to the control group. This patent citation analysis identifies whether acquisitions potentially create barriers to technology access, reflected in decreased external citations, or facilitate broader diffusion, indicated by increased citation activity..

Nevertheless, a key empirical challenge is endogenous selection—acquirers strategically choose targets based on unobservable characteristics that can affect future citation patterns. For example,

acquirers might systematically target firms whose technologies are already becoming more (or less) relevant to the broader industry, confounding the causal effect of acquisition. To strengthen identification, the study adopts patents from failed M&As as the control group, following Seru (2014), selecting deals that failed for reasons unrelated to targets' innovation characteristics, such as regulatory blocks on antitrust grounds or financing issues. Since these targets were chosen similarly for acquisition but remained independent, comparing successful versus failed acquisitions allows me to isolate the causal impact of acquisition completion on technology diffusion patterns, separate from selection effects.

Tech M&As are defined as acquisitions in which target companies received patents within the seven years before the acquisition announcement. To accurately track external citations, public firms' patent data from Arora et al. (2021) is utilized, which accounts for name and ownership changes in firm-patent matching. For private firms, this study applies fuzzy name matching combined with web searches, following Autor et al. (2020), to assign patents to firms over time.

Difference-in-differences regression analysis reveals that tech M&As significantly increase the diffusion of targets' technologies, evidenced by higher external patent citations. These findings prove robust after controlling for technology class and patent age fixed effects. Further analysis categorizes citing firms by industry and illustrates that firms within the acquirer's industry account for most of the increased patent citations. The diffusion effect of tech M&A measured by patent citations is also stronger in deals involving non-patenting acquirers than patenting acquirers.

Additionally, this research examines whether young external firms that cite the target's patents experience reduced technology access post-acquisition. Small and young firms are particularly vulnerable to exclusionary practices as they typically lack the resources and established patent portfolios to navigate complex licensing arrangements. If acquirers were strategically hoarding technologies or creating access barriers, we would expect these young citing firms to exhibit the largest decrease in citations post-acquisition. However, no evidence of decreased citations from young firms following acquisitions is found. This supports the main finding that tech M&As increase rather than reduce technology diffusion, even for the external firms most vulnerable to exclusion.

The positive diffusion effects extend beyond citation counts. Follow-on innovations citing acquired technologies show higher patent renewal rates, a measure of commercial value, indicating that acquisitions foster higher-quality derivative innovations.

To better understand the mechanisms underlying the increased diffusion from tech M&As and quantify their aggregate growth impact, a dynamic innovation model is developed where firms build upon other firms' technologies through knowledge diffusion. The model extends standard idea flow frameworks by differentiating innovation diffusion via independent firms versus acquired firms, corresponding to my empirical evidence comparing citation patterns between control and treated firms.

To investigate the innovation diffusion channel and how acquisitions generate different ideas through encounters with varied source technologies, the model features two groups of firms: potential acquirers and potential targets. Potential target firms represent technological opportunities that potential acquirers can build upon. In a steady state, potential targets are characterized by a stationary distribution over technology quality and acquisition status— independent or acquired. Potential acquirers can improve productivity through two channels, both modeled as exogenous Poisson processes: (1) knowledge diffusion, where acquirers randomly meet potential targets and can engage in cumulative innovation building on the potential target's technology, with innovation success contingent on the target's quality and acquisition status; and (2) direct acquisition, where acquisition opportunities arrive stochastically, permitting acquirers to purchase independent targets, generating immediate productivity gains and changing the target's status from independent to acquired.

The key innovation relative to extant theories is modeling how acquisition changes diffusion dynamics. Two opposing forces emerge post-acquisition: external firms benefit from positive spillovers from acquired technologies due to M&A-generated synergies. However, they could face increased learning costs if acquirers exert higher appropriation rates over target technologies. The model captures this mechanism through acquirers' choice of “impediment effort”—the degree to which acquirers invest in barriers that prevent other firms from building upon the acquired technology, with higher effort incurring increasing appropriation costs. This captures the enhanced degree of technological control that can follow ownership changes (see Bryan and Hovenkamp, 2020, for a differentiated oligopoly model rationalizing this effect in startup acquisitions). The

model calibrates impediment effort using differences in potential acquirers' appropriation rates between acquired target technology and internally developed technology.

Model parameters are calibrated to align with empirical estimates of tech M&A effects on target technology diffusion and other moments from firm and patent data. The calibrated model successfully replicates key non-targeted moments. For example, it generates an aggregate technology growth rate on a balanced growth path of 0.75 percent, which compares favorably to the long-run average 1.1 percent TFP growth rate observed in US data, with the difference reflecting that the model only captures the diffusion and M&A channels of technological progress. The model also produces annual external citation rates and cites patterns within their observed ranges for targets and acquirers. Finally, the model's implied appropriation costs do not exceed the upper bound of acquired intangible amortization based on observed US data from 2012-2015. Though imperfect, the model fits long-run averages of key variables characterizing US technology growth and diffusion patterns between firms' patenting networks. Using this benchmark model, the study conducts three counterfactual analyses to examine how tech M&As affect productivity growth through the diffusion channel.

First, the "synergy" benefits from tech M&As contribute approximately 0.01 percentage points to productivity growth on the balanced growth path, representing 1.38 percent of the total growth rate. Shutting down the "synergy" benefits annually reduces total citations by 0.01 per patent.

Second, doubling the baseline tech M&A rate—an exercise which an exogenous change in antitrust policy may rationalize—generates an additional 0.05 percentage points of annual growth, with 40 percent of this increase attributable to the diffusion channel. These exercises expose the non-negligible role of diffusion in tech M&As' impact on productivity growth, suggesting that merger policies should incorporate diffusion effects in their assessments.

Finally, relaxing regulatory constraints by exogenously lowering acquirers' appropriation costs produces minimal changes in growth rates or diffusion relative to the benchmark economy. While higher target technology appropriation restricts positive diffusion from acquired targets, potential acquirers compensate by growing their innovation efforts with independent targets, as higher acquisition values make independent targets more attractive. This equilibrium feedback highlights the need for careful policy design when attempting to enhance diffusion through appropriation regulation.

Related Literature

This paper contributes to several strands of literature examining innovation, mergers and acquisitions, and their implications for aggregate economic outcomes.

M&As and Innovation Outcomes. A substantial body of research examines how mergers affect innovation within merging firms, yielding mixed results. One strand emphasizes negative effects, often attributed to "killer acquisitions" in which incumbents acquire nascent competitors to suppress competing projects. Cunningham et al. (2021) provide evidence from the pharmaceutical industry that incumbents systematically acquire and discontinue rival innovations, while Seru (2014) shows that diversifying acquisitions reduce the novelty of corporate R&D as inventors become less productive after integration.

Conversely, a growing literature emphasizes positive innovation synergies. Liu (2022) finds that acquirers become more innovative after acquisitions with increased patent output. Bena and Li (2014) show that these innovation gains are driven by technological synergies between merging firms. Phillips and Zhdanov (2012) demonstrate both theoretically and empirically that an active acquisition market can strengthen small firms' innovation incentives: potential startups innovate more when acquisition opportunities exist, and large firms strategically acquire rather than replicate external R&D. These results suggest that M&As can serve as an outsourcing mechanism for innovation, fueling R&D investment across the broader ecosystem.

While most of this literature studies innovation within merged entities, this paper focuses on a distinct but critical dimension: the diffusion of acquired technologies to other firms. By measuring post-acquisition patent citations, I directly test whether tech M&As restrict or expand knowledge spillovers in the wider innovation network. This diffusion-based perspective supplements the M&A and innovation literatures and offers new empirical evidence on how acquisitions reshape the flow of ideas across firms. A recent study by Ederer, Seibel, and Simcoe (2025), of which I became aware after conducting my analysis, finds similar post-acquisition citation increases in digital markets, aligning with my empirical results on diffusion.

Market Power, Industry Dynamics, and Merger Waves. Recent research has revitalized debates over whether rising concentration and markups signal declining competition or "competition in

action." De Loecker et al. (2020) document increasing markups across U.S. industries since 1980, while Ganapati (2021) and Shapiro and Yurukoglu (2024) caution that such trends often reflect efficient expansion by innovative "superstar" firms rather than market power abuse. Industry-level evidence, such as Miller et al. (2023) in the cement sector, reinforces this view: technological advances and scale economies, rather than collusion, often drive consolidation and markup growth.

This context is closely related to the merger waves observed in technology sectors. Periods of intensive M&A activity, often linked to technological shifts, reshape both market structure and innovation dynamics (Mitchell & Mulherin, 1996; Harford, 2005). The surge in tech acquisitions during the 2010s exemplifies this trend, as large incumbents absorbed startups with complementary capabilities or disruptive technologies.

By focusing on external knowledge diffusion, this paper sheds new light on these debates. If post-acquisition diffusion expands, as the evidence here demonstrates, it implies that at least part of the observed consolidation in technology industries corresponds to productive reallocation rather than foreclosure. This study thus contributes to the macro-micro linkage between patent-level post-acquisition outcomes and aggregate patterns of markups, entry, and technological progress discussed in recent competition research.

Firm Boundaries, Complementary Assets, and Idea Flows. The modeling framework corresponds to Teece's (1986, 1997) view of acquisitions as coupling target innovations with the acquirer's complementary assets and dynamic capabilities—reflected in the model as a larger innovation step size when knowledge is integrated. It also complements Williamson (1975) by showing that shifts in firm boundaries not only reduce the acquirer's internal contracting frictions but also reprice outsiders' access to the technology via endogenous appropriation/impediment effort, thereby shaping external diffusion. Embedding these mechanisms in an idea-flow growth framework (e.g., Luttmer, 2012; Lucas and Moll, 2014; Perla and Tonetti, 2014) demonstrates that the diffusion channel of acquisitions operates through two levers—synergy-driven diffusion and appropriation—thereby clarifying how ownership-induced boundary changes shape long-run growth.

M&As, Diffusion, and Aggregate Growth. A growing macroeconomic literature analyzes how M&As affect aggregate economic outcomes through two primary channels: reshaping market structure (e.g., David, 2021; Cavenaile, Celik, and Tian, 2021; Cao and Zhu, 2025) and altering

innovation within merged entities (e.g., Rosen et al., 2021). This work approaches M&As from a separate angle—the diffusion perspective. This study examines how technology acquisitions alter knowledge diffusion patterns across firms and influence aggregate growth outcomes. This approach connects to the broader literature on knowledge diffusion as a driver of economic growth (e.g., Ayerst, 2022; Kim and Jo, 2024; Muratori, 2020) but uniquely emphasizes the role of how discrete corporate events (tech M&As) reshape these diffusion networks.

The model in the paper shows that diffusion effects account for a sizable fraction of tech M&As' contribution to productivity growth, highlighting a previously underappreciated aggregate consequence of acquisition activity. By quantifying the diffusion component of tech M&A-driven growth, the paper bridges micro-level innovation studies with macro-level productivity analysis.

Antitrust and Innovation Policy. The findings in this paper contribute to policy debates regarding antitrust enforcement in technology markets. Bryan and Hovenkamp (2020) argue that incumbent acquisitions of startups could warrant regulatory intervention when they restrict knowledge diffusion, discussing remedies such as compulsory licensing of acquired technologies. This study quantitatively evaluates the potential harm arising from knowledge hoarding following acquisitions. The analysis reveals nuanced effects: while acquisitions can create appropriation frictions that limit diffusion, they simultaneously generate positive spillovers through technological synergies. These findings suggest that antitrust authorities should adopt a comprehensive approach to tech M&A oversight that considers both the restrictive and enhancing effects on knowledge diffusion.

Roadmap: The paper proceeds as follows. Section 2 discusses relevant mechanisms of the diffusion effect of tech M&As. Section 3 describes the data, sample construction, and empirical methodology. Section 4 presents empirical results on post-acquisition citation/diffusion patterns. Section 5 develops a theoretical model of idea flows and firm innovation, and section 6 outlines the structural estimation and conducts simulation exercises. Section 7 concludes.

2. Idea Development: Competing Mechanisms Linking Technology M&A to Diffusion

The effect of technology M&A on the diffusion of acquired technologies is theoretically ambiguous, as acquisitions can either facilitate or impede the spread of innovation. Building on prior work, two broad mechanisms emerge, emphasizing synergy-driven diffusion and highlighting appropriation-driven restriction. This section develops these competing mechanisms to motivate empirical analysis.

Synergy-driven diffusion (positive mechanism)

Technology-oriented M&A can develop the reach and impact of the acquired technology through several reinforcing channels. One is the technological synergy effect (Bena & Li, 2014; Liu, 2022; Li & Wang, 2023): combining related knowledge bases after acquisition can improve the technical performance, reliability, or applicability of the target's inventions, making them more attractive to potential adopters. A second is the acquisition of complementary assets—manufacturing capacity, distribution networks, regulatory expertise, and established customer relationships—that lower commercialization frictions and accelerate deployment (Teece, 1986). A third is the acquirer's network position, which can amplify visibility and compatibility through larger installed bases, brand reputation, and partner ecosystems. Collectively, these factors imply that promising development of the target's technology—through upgrades, broader reach, and new combinations—can stimulate greater external adoption and knowledge diffusion.

Appropriation-driven restriction (negative mechanism)

M&A can also lower diffusion if appropriate incentives dominate. A common channel is licensing reduction, where a dominant acquirer internalizes the technology and restricts rival access, blocking outsiders' adoption (Bryan & Hovenkamp, 2020). Another is strategic shelving or "killer acquisitions," where the acquirer discontinues competing product lines to protect its products (Cunningham, Ederer & Ma, 2021). Increased patent concentration and denser patent thickets can heighten bargaining frictions and litigation risk, further impeding spillovers (Akcigit & Ates, 2021, 2023). Integration can result in tacit-knowledge internalization, keeping key inventors and know-how within the merged entity through non-compete agreements; whereas inventor departures can sometimes spread knowledge to other firms (Seru, 2014; Kim, 2024), such mobility is uncertain and can be strategically limited. In this view, consolidation, control, and foreclosure mechanisms can outweigh technical improvements, leading to reduced citations, narrower licensing, and slower diffusion.

Since synergy-driven and appropriation-driven channels operate in opposite directions, the net impact of technology M&A on diffusion is theoretically ambiguous. This encourages an empirical test of whether, on average, technology M&A increases or decreases diffusion, followed by an examination of how the observed effect manifests across different citation patterns.

3. Data and Methods: Firm-to-firm Linkages and Technology M&As

This section describes the data sources, sample construction, and empirical methodology used to study the effect of technology M&As on the diffusion of targets' technology. A difference-in-differences approach is adopted using failed M&As as controls, building on the identification strategy pioneered by Seru (2014).

3.1 Primary Data Sources

The analysis combines two primary databases spanning 1980–2021, enabling comprehensive measurement of M&A activity and patent-based knowledge diffusion:

M&A Transaction Data. The SDC M&A database is employed, covering extensive M&A deals where targets are US firms from 1980 to 2021. Following established conventions in the M&A literature (Shenoy 2012; Zhang & Tong, 2021), this study also applies standard filters: deal value \geq \$1 million, acquirer seeking \geq 50% of target shares, excluding financial sector deals, repurchases, spinoffs, divestitures, recapitalizations, leveraged buyouts, and self-tender offers.

Patent and Citation Data: Patent information arises from the PatentsView USPTO database (1976–2022), containing all utility patents granted between 1976 and 2022 with 410,870 firm-patent assignees. Crucially, this includes citations made to US granted utility patents by US patents, facilitating measurement of knowledge diffusion through citation networks.

3.2 Firm-Patent Matching

Since there is no common identifier between the SDC and patent data, the study uses a multi-stage matching procedure combining validated datasets with systematic name-matching algorithms.

Stage 1: Public Firm Validated Matching

The procedure begins by identifying publicly traded firms using CUSIP codes from the SDC database, linked to PERMNO codes through the CRSP-Compustat merged database. For public firms with data covering 1980–2015, this study utilizes the comprehensive firm-patent matching dataset from Arora, Belenzon, and Sheer (2021). This dataset addresses fundamental challenges in patent-firm matching by systematically tracking corporate name changes, mergers, acquisitions, and subsidiary relationships that affect patent ownership attribution over time. Additional matches from Kogan, Papanikolaou, Seru, and Stoffman (2017) are incorporated to enhance coverage.

Stage 2: Systematic Name Matching for Remaining Firms

It has been decided that systematic name matching, following Autor et al. (2020), should be implemented for firms not covered by validated datasets- primarily private firms. This research first applies the standardization routine from Arora, Belenzon, and Sheer (2021) to SDC firm names and USPTO assignee names, removing corporate suffixes, standardizing abbreviations, and normalizing formatting. Then, the Stata *matchit* function is employed to perform fuzzy name matching on frequent stop words with a 0.7 similarity threshold.

For post-match checking, all exact matches are retained after manually verifying ambiguities and filtering fuzzy matches using NAICS97 industry codes before applying web-based verification. I keep the highest similarity scores above 0.9 for targets to reduce the search volume, given that targets typically correspond to single assignees. Following Autor et al. (2020), Bing web searches are conducted, combining firm and assignee names, analyzing each search's top five returned URLs. Matches are classified as validated when at least two 5 URLs reference the same entity, indicating that independent web sources confirm the firm-assignee relationship. This criterion helps identify subsidiary relationships, alternative business names, and corporate restructurings not captured in formal databases while providing an objective validation threshold. Target firms were also able to link to multiple assignees when justified.

3.3 Defining "Technology" M&As

Technology M&As (tech M&As) are defined as acquisitions of companies that obtained patents within seven years before the deal was announced. This definition captures firms with recent, active innovation portfolios while avoiding inclusion of inactive patent holders.

Given that the most dependable firm-patent matches derive from the Arora et al. (2021) dataset covering 1980–2015, the tech M&A sample concentrates heavily in the pre-2015 period (approximately 90% of deals), a pattern reflected in both the summary statistics and graphs in stylized facts. This concentration ensures sufficient post-acquisition observation periods through 2021 for my event study methodology.

Summary Statistics: Table 3.1 compares total M&A deals with the tech M&A subsample. While tech M&As constitute approximately 11% of total deals, they exhibit significantly higher deal values and patent counts than the SDC population. Public and high-tech sector targets represent a larger share of tech M&As. The industry composition of tech M&As is concentrated in high-tech and R&D-intensive sectors. As shown in Data Appendix Table DA.1, the largest target industries are Healthcare Equipment & Supplies (13.6%), Software (6.6%), and Semiconductors (6.6%), followed by Machinery and Pharmaceuticals. Acquirers exhibit a similar pattern, with notable shares in Healthcare Equipment, Pharmaceuticals, and Semiconductors.

Figure 3.1 illustrates an interesting target firm age distribution using founding year data from Ewens and Marx (2023). For easier direct matches, the sample firms are matched with single-organization patent records in their data to obtain founding years for the tech M&A targets (87% in the tech M&A sample have a founding year). The age distribution reveals concentration among "middle-aged" firms (7–30 years), with a shift toward younger firms in recent decades, suggesting increasing acquisition of earlier-stage technology companies. Conversely, startup acquisitions (0–6 years) remain a small fraction of tech M&A activity.

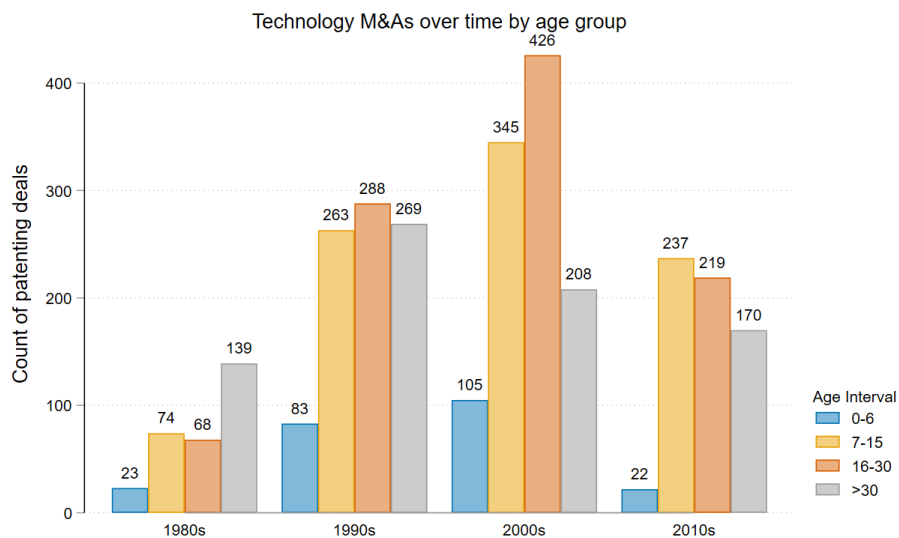


Figure 3.1: Age distribution for targets in Tech M&As

Table 3.1: Summary Statistics – M&A Sample and Tech M&As

	Full M&A Sample	Tech M&As	Tech Share
Deal characteristics:			
Number of deals	32917	3692	11.21%
Total deal value (\$ mi)	1.34×10^7	5.73×10^6	42.76%
Average deal value (\$ mi)	406.21	1553.08	382.3%
Same-industry deals (%)	48.07	44.66	-
Target characteristics:			
Private targets (%)	75.50	38.60	-
High-tech industry (%)	43.41	61.16	-
Patent Portfolio			
Number of targets with a patent (granted \geq 1976)	5422	3689	68.04%
Patents per target (mean)	2.6	4.1	157.69%
Patents per target (median)	1.0	1.5	150%

Citations per patent per year (average over targets)	0.656	0.664	101.22%
--	-------	-------	---------

Note: High-tech industry classifications and public/private status are from the SDC database. Same-industry deals defined by SDC Mid industry codes. Patent statistics conditional on target firms having patent portfolios, measured over the pre-acquisition period. Patents per target is the average patent counts per year for patents granted after 1976 (The mean is winsorized at the 95% level to address outliers.). Citation rates calculated using complete 5-year post-grant windows and the sample restricted to patents granted 1976–2016 to ensure full citation exposure by sample end (2021).

3.4 Identification Strategy

To identify the effects of tech M&As on the diffusion of the target’s technology, the core identification challenge stems from endogenous selection: acquirers systematically choose targets with high diffusion potential or technological alignment with their strategic goals. Simple comparisons of pre-/post-acquisition citation patterns would conflate treatment effects with selection effects.

Following Seru (2014), selection bias utilizing withdrawn M&A deals as controls is addressed. The significant insight is that deals failing for reasons unrelated to the targets' innovation potential provide valid counterfactuals for successful acquisitions.

Control Group Construction: Among 174 withdrawn tech M&A deals before 2015, news articles are manually reviewed for each withdrawal to identify reasons for failure. I exclude deals failing due to innovation-related concerns, disagreement over valuations, or reasons that cannot be determined.

Qualified Control Deals: This process yielded 63 qualified control deals where failure stemmed from exogenous factors:

Table 3.2: Summary of Failed Deals' Reasons for Qualified Controls

Failed Reason	Count	Share	Cumulative
Acquirer issue not related to tech (Unexpected funds issue, or neg shocks)	21	33.33%	33.33%

Regulation	18	28.57%	61.90%
Target reject is not related to the target tech	11	17.46%	79.37%
Other exogenous reasons (Macro shocks, adverse market conditions)	8	12.70%	92.06%
Competing offers	3	4.76%	96.83%
Management disagreement	2	3.17%	100%
Total	63	100%	

Each qualified control deal involves more than 20 patents for targets on average, ensuring sufficient patent portfolio size for meaningful analysis.

3.5 Sample Construction

The unit of analysis consists of each tech M&A target patent that the patent office granted at least two years before the M&A announcement year, but no longer than 7 years before. This gives me enough sample size and ensures the technology is fresh.

3.5.1 Matching refinement

The Initial Sample contains 128,764 treated patents and 1,693 control patents from qualified deals. The sample size of withdrawn deals is small, leading to greater sampling variability that could not be comparable with the treated ones. Furthermore, patent citations are not just a measure of individual patent character; they often reflect broader network effects. Withdrawn deals alone cannot fully balance differences in the network with potential citing firms. The above issues are addressed by following the matching refinement. Similar practice is seen in Bena and Li (2014).

Deal Grouping Strategy: Patents are grouped by target-acquirer industry pair and deal announcement year, yielding 2,455 treated versus 771 control patents. This approach controls industry composition and time trends—key dimensions affecting M&A patterns (Roberts and Whited, 2011).

Coarsened Exact Matching (CEM): Within each group, CEM is applied based on pre-event total citations and relative patent age in the event year. As matching on too many covariates creates a curse of dimensionality (Heckman et al. 1998) and causes the analysis to lose statistical power, the above two covariates are matched, which are strongly relevant to citation patterns and M&As.

The matching procedure yields 1,421 treated and 667 control patents with comparable pre-acquisition citation trajectories.

3.5.2 Matching Quality Assessment

Table 3.3 presents logistic regression results testing treatment prediction using patent characteristics in the matched sample. Five established patent-level measures are employed to reflect innovation quality and technological factors. Patent breadth and **RETech** from Bowen, Fresard, and Hoberg (2021) measure the scope of a patent's technological territory and whether it pertains to rapidly evolving technological areas, respectively. Originality follows Trajtenberg et al. (1997) as one minus a Herfindahl index based on cited patents' Cooperative Patent Classification (CPC) classes, capturing technological diversity. **New word** from Arts et al. (2021) represents a text-based novelty measure counting unique keywords in the patent document. None of the key patent quality measures (originality, breadth, novelty indicators, and claims count) significantly predict treatment status, confirming the balance of key characteristics of treatment and control patents.

Table 3.3: Logit Regression of Treatment Prediction

	Dept var: Treatment
	Coefficient
	(SE)
Originality	-0.655 (0.544)

Breadth	-1.154 (1.053)
New word	.066 (0.082)
RETech	-.049 (0.040)
Num claims	-.005 (0.007)
<hr/>	
Observations	1976

Note: I apply logistic regression with treatment status as the dependent variable. The regression also included matched strata fixed effects. Robust standard errors in parentheses. * significant at 10%; ** significant at 5%; *** significant at 1%.

4. Effects of Technology M&As on Diffusion of the Acquired Technology

This section presents the main empirical findings on how technology acquisitions affect the diffusion of targets' pre-acquisition patents. Event study methodology is applied with a difference-in-differences strategy to measure changes in external citation patterns.

4.1 Econometric Specification

An event study specification is implemented with a 14-year window (6 years pre-acquisition, 7 years post-acquisition):

$$y_{ijt} = c + \sum_{k=-4}^7 \delta_k \Delta z_{i,t-k} + \delta_{-5+}(1 - z_{i,t+5+}) + \alpha_i + \text{year}_{it} \times \text{pair}_j \text{ FE} + u_{ijt}$$

where y_{ijt} is the technology diffusion measure (number of non-self-citations received) for a patent i in a matched pair j at time t . $\Delta z_{i,t-k}$ indicates whether the unit i experienced an announced acquisition exactly k periods before period t . $year_{it} \times pair_j$ Are the time-varying pairs fixed effects that exploit variation within each pair? This study performs conditional Poisson regression to manage the count nature and zero-inflation in citation data. To stabilize estimates, I also bin lags beyond four periods together, as fewer patents reach ages of more than 4 years before the event.

Primary Diffusion Measure: Non-self-citations received by target patents starting from the grant year, excluding self-citations by the acquirer or target firm to focus on external knowledge buildup. This paper also scales citations by the mean citation counts within technology class and patent age to account for truncation bias (Lerner and Seru, 2022).

4.2 Baseline Diffusion Effect

4.2.1 Event Study Dynamics

Figure 4.1–4.2 presents event study plots showing citation dynamics after acquisition events. Both raw and normalized outcome measures show increased citations after tech M&As.

Citation patterns between treated and control patents exhibit parallel trends in pre-acquisition periods, supporting identification assumptions. The event study graphs confirm significant positive effects emerging within one year post-acquisition and persist for at least 5 years; the effects seem stable before year 6 and fade away in longer periods. Given that half of the withdrawn deals were withdrawn at date 0 and all sample deals were complete at or before date 1, it shows the timely response from other firms after this tech M&A shock. This dynamic is also robust to non-self-citations from all entities instead of just firm-to-firm citations.

The average number of non-self forward citations in the pre-event period is 0.34 for the raw counts and 0.41 for the normalized measure. A coefficient of 1 at event time 1 consequently signals that the acquisition induced more than a 1.5-fold increase in non-self forward citations relative to the pre-event baseline.




Figure 4.1: Event study with CEM matched sample on non-self-follow-on patents




Figure 4.2: Event study with CEM matched sample on mean citation normalized measures

4.2.2 Static Results

Table 4.1 exhibits static results showing that technology acquisitions significantly enhance the diffusion of the target's technology. The coefficient on $D(\text{acquired}) \times D(\text{Post})$ ranges from 0.884 to 1.323 across specifications, indicating 142–275% average increases in citation rates post-acquisition, consistent with the dynamics results. Results remain consistent across specifications with and without patent fixed effects, using raw and scaled citation measures, including various combinations of technology class, age, and year-pair fixed effects.

Table 4.1: Non-self-follow-on Patents as The Diffusion Measure

Dept var: # non-self-follow-on patents	Raw citations		Scaled citations	
	(1)	(2)	(3)	(4)
D(acquired)·D(Post)	1.323 ^{***} (.298)	1.284 ^{***} (.259)	.884 ^{***} (.327)	.891 ^{***} (.309)
Tech class F.E.	Yes	Yes	Yes	Yes
Patent F.E.	Yes	No	Yes	No
Age F.E.	Yes	Yes	Yes	Yes
Year×pair fixed effects	Yes	Yes	Yes	Yes
Observations	16492	23001	16492	23001

Note: I estimated a Poisson pseudo-maximum likelihood estimator. D(acquired) takes the value 1 for treated patents. D(Post) takes the value 1 after the acquisition year. The outcome variable is winsorized at 99%. Robust standard errors in parentheses (clustered at the target firm level). * significant at 10%; ** significant at 5%; *** significant at 1%.

The positive diffusion results indicate that the “synergy-driven” effects are dominant after the technology M&As, either through combining technology, complementary assets, or network distribution. Here are some examples from the sample data to further illustrate the positive diffusion results.

AMD horizontal integration with Monolithic Memories: AMD's 1987 acquisition of Monolithic Memories exemplifies the synergy through technology combination. Monolithic's programmable logic arrays—specialized circuits that could be programmed for different logical operations—were integrated directly into AMD's processor designs, transforming them from standalone components into core processor features. This technological integration created industry-wide knowledge spillovers as major semiconductor companies actively referenced Monolithic's foundational patents (4554640, 4476560) when developing their programmable logic capabilities. Within five years post-acquisition, leading competitors including Micron Technology

(4 citations, 1992), SIEMENS (2 citations, 1989), Altera Corporation (2 citations, 1988-1989), and National Semiconductor (1 citation, 1990) built upon Monolithic's innovations, producing more than 20 total citations that show how technology combination through M&A facilitates broader industry learning and innovation.

Johnson & Johnson (J&J) scope expansion acquisition of ALZA: J&J's 2001 acquisition of ALZA illustrates synergy through complementary assets and expanded distribution networks. ALZA's osmotic drug delivery systems (OROS)—controlled-release formulations using osmotic pressure—gained access to J&J's extensive pharmaceutical portfolio, regulatory expertise, and global distribution capabilities. This combination transformed ALZA's specialized contract-based technology into a platform serving diverse therapeutic applications across J&J's broad market reach. Within five years, leading pharmaceutical companies cited ALZA's core patents, including Human Genome Sciences (5 citations), Aventis Behring, and Novozymes Biopharma, as they developed controlled-release formulations. This citation pattern validates how combining complementary assets and leveraging expanded distribution networks transforms specialized technologies into industry-standard platforms, encouraging widespread adoption and knowledge building across numerous therapeutic areas.

4.3 Characterizing the Effects

The main regression results presented in the previous section demonstrate positive diffusion effects of target technology following tech M&A transactions. This section provides a deeper examination of the underlying mechanisms and heterogeneous effects across different dimensions.

4.3.1 Within-Industry Spillovers

First, it is examined whether positive spillovers from tech M&A benefit firms within the same industry as the acquirer or whether the effect extends across industry boundaries. To identify within-industry citing firms, this paper maps patents' CPC technology subclass to three-digit NAICS 97 industry codes (NAICS3) using the crosswalk provided by Goldschlag, Lybbert, and Zolas (2016). Based on each assignee's patenting history, the top three NAICS3 codes are obtained and compared with the acquirer's top three NAICS3 codes. If any codes overlap, the citing patents are classified as within-industry spillovers.

Table 4.2 exhibits the estimates of diffusion effects from tech M&As, distinguishing between within-industry and total citation counts. The concentration of the impact within acquirer industries suggests that technology M&As enhance diffusion primarily through industry-specific knowledge networks rather than broad cross-industry spillovers. This pattern supports theories of technology complementarity and industry-specific absorptive capacity, where organizations most effectively build upon technologies from related fields.

A supplementary finding emerges when restricting the sample to acquisitions by patenting firms. These transactions yield smaller diffusion effects (Poisson coefficients of 0.70–0.80) compared to the full-sample estimates in Table 3.1, which illustrate larger effects. This pattern reflects two potential mechanisms. First, patenting acquirers typically target firms in related technological fields, integrating acquired patents into their existing R&D portfolios. Such targets already show higher pre-acquisition citation rates (1.6 for patenting acquirer deals versus 1.4 for other deals), making the post-acquisition effect proportionally smaller. Second, non-patenting acquirers often lack in-house R&D capabilities and rely more heavily on complementary assets through licensing, partnerships, or spin-outs—channels that expose technology to broader user bases and boost non-self-citations (Teece, 1986). In contrast, patenting acquirers tend to integrate targets into internal R&D operations. They may acquire defensively to secure freedom to operate or bargaining leverage (Hall & Ziedonis, 2001), which can dampen external diffusion.

In the quantitative model in section 5, the acquirer’s knowledge hoarding behavior will partly offset the positive diffusion after tech M&A, which speaks to the lower diffusion effect for patenting acquirers observed in the data in an aggregate way.

Table 4.2: Diffusion Effect Conditional on Patenting Acquirer

Dept var: # non-self-follow- on patents	Within-industry citations		All citations	
	(1)	(2)	(3)	(4)

D(acquired)·D(Post)	.725*	.763**	.773*	.816**
	(.408)	(.380)	(.409)	(.371)
Tech class F.E.	Yes	Yes	Yes	Yes
Patent F.E.	Yes	No	Yes	No
Age F.E.	Yes	Yes	Yes	Yes
Year×pair fixed effects	Yes	Yes	Yes	Yes
Observations	8306	10703	8439	10718

Note: I estimated a Poisson pseudo-maximum likelihood estimator. D(acquired) takes the value 1 for treated patents. D(Post) takes the value 1 after the acquisition year. The outcome variable is winsorized at 99%. Robust standard errors in parentheses (clustered at the target firm level). * significant at 10%; ** significant at 5%; *** significant at 1%.

4.3.2 Impact on Young Firms

Young firms often occupy weak positions in technology markets relative to established competitors. One might expect that positive spillover effects may not benefit young firms and could result in decreased citations due to heightened protection following target integration. This section examines tech M&A's impact on citations from young firms (assignees).

Table 4.3 examines citations from young US assignees (firms ≤ 5 years old at citation time). Despite small sample sizes in Poisson regressions due to singleton drop, significant positive effects for young firm citations are discovered (coefficients = 1.161–1.449 in Poisson specifications). Although citations from young US assignees are infrequent (most targets receive zero citations from young US assignees), they also benefit on average from tech M&A events through increased citations to targets.

These results suggest that technology acquisitions do not reduce target technology access for young innovative firms. This finding addresses concerns that technology acquisitions might stifle entrepreneurial innovation by limiting technology access.

Table 4.3: Diffusion Effect from the US Young Assignees

Dept var: # non-self- follow-on patents	Poisson		OLS	
	(1)	(2)	(3)	(4)
D(acquired)·D(Post)	1.449 ^{**} (.621)	1.161 ^{***} (.386)	.028 [*] (.015)	.027 [*] (.014)
Tech class F.E.	Yes	Yes	Yes	Yes
Patent F.E.	Yes	No	Yes	No
Age F.E.	Yes	Yes	Yes	Yes
Year×pair fixed effects	Yes	Yes	Yes	Yes
Observations	2927	14757	23841	23841

Note: The outcome in the regression counts the non-self-citations from assignees whose maximum available age is less than 6 years (see section 2.3 for the data on ages). The outcome variable is winsorized at 99%. Robust standard errors in parentheses (clustered at the target firm level). * significant at 10%; ** significant at 5%; *** significant at 1%.

4.3.3 Quality of Follow-on Innovation

This section examines whether follow-on innovations based on target technology exhibit higher quality. Higher post-acquisition citations might reflect market attention without genuine technological improvement following integration. This paper examines citing patent quality using renewal decisions as a quality proxy rather than citation counts, since high diffusion can automatically increase the latter.

New patents must pay maintenance fees at 3.5, 7.5, and 11.5 years after grant, with a 6-month grace period. Non-payment indicates low commercial value. I Renewal data from the USPTO Maintenance Fee Event files is obtained. For each target patent i with R_{it} citations in the year t an aggregate quality measure for citing patents is defined Y_{it} :

$$Y_{it} = \frac{1}{R_{it}} \sum_{j \in C(i,t)} Renewed_{ij}$$

With $Renewed_{ij} = 1$ (*did not show Expire status within 4 yeras since grant or Reinstated after maintenace fee payment confirmed*)

This outcome is a fractional number. An OLS regression is run for the tech M&A sample comparing treated and control patents, confirming that predicted rates using OLS estimates remain between 0 and 1.

Figure 4.3 indicates around 20% increases in renewal rates for patents citing acquired technologies in the first three years of post-acquisition, suggesting higher-quality follow-on innovation. Results remain similar when examining twice-renewed citing patents.

Higher renewal rates indicate that acquisitions enable more commercially valuable innovations, building on target technologies. This suggests a more promising environment for acquired technology and supports theories of complementary assets and technology synergies following tech acquisitions.

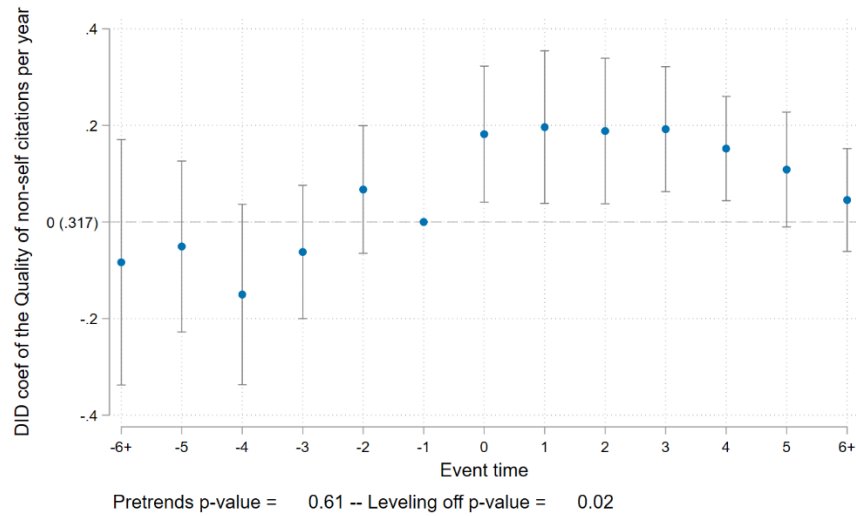


Figure 4.3: Event study graph for average renewal rate

4.3.4 Diffusion Environment

The positive spillover effects in previous analyses suggest that tech M&A does not impose stronger technology knowledge protection. To confirm this, the study examines whether acquisitions change the intellectual property enforcement environment using patent litigation data from the USPTO covering the post-2003 period.

After matching target patents in the regression sample to patent infringement data, limited enforcement effects are found: fewer than 5% (48/1,301) of target patents in the regression sample experience litigation. While most litigation cases involve treated rather than control patents, the low fraction indicates that enforcement environment changes do not significantly contribute to diffusion effects.

Due to the limited sample size, conducting a formal difference-in-differences regression is not feasible. However, matching results from patent infringement cases provide no strong support for concerns about knowledge hold-up following tech acquisitions based on the regression sample.

4.4 Summary of Empirical Findings

4.4.1 Main Contributions

The empirical analysis yields key insights into the relationship between technology M&As and knowledge diffusion.

Positive diffusion effects: Technology acquisitions significantly increase external citation rates for target patents, countering concerns about “killer acquisitions” or knowledge hoarding by large acquirers.

Persistence: The positive diffusion effect persists for at least four years after acquisition, indicating a sustained improvement in technology accessibility rather than a short-lived integration outcome.

Industry concentration: The diffusion increase arises within the acquirer's industry, consistent with theories of absorptive capacity and technological complementarity.

Quality enhancement: Follow-on innovations citing acquired technologies exhibit higher commercial value, as measured by patent renewal rates, suggesting that acquisitions foster the development of higher-quality derivative innovations.

4.4.2 Policy Implications

From a policy perspective, these results indicate that, in the sample, technological M&As tend on average to increase rather than restrict knowledge diffusion. This highlights the need for antitrust and innovation policy assessments to account for potential anti-competitive risks and the pro-diffusion benefits of such acquisitions.

In the next few sections, a simple model is developed based on the main empirical findings, and it is employed to study the aggregate effects of tech M&As through the diffusion channel.

5. Theoretical Model: Idea Flows and Technology M&As

Objective of the Model: In this section, a parsimonious model of idea flows is developed to illustrate the mechanism linking technology acquisitions and knowledge diffusion in a dynamic setting and to quantify the aggregate growth impact of technology M&As at steady state. The structure is an idea-driven endogenous growth model augmented with an acquisition process, drawing on elements of quality ladder models and idea flow models (e.g., a blend of Lucas et al.

(2008)-style technology diffusion and the Klette and Kortum (2004) framework with a single product line). Rather than a general equilibrium growth model of the overall economy, the model focuses on a single-sector economy with particular attention to the technology sector. It is not intended to perfectly capture every aspect of M&A or knowledge diffusion activity (for instance, I do not model the detailed technology M&A search process), but rather to capture key channels (knowledge spillover versus hoarding) and provide a laboratory for counterfactual analysis.

5.1 Model Setup

Time is continuous, starting at $t = 0$, and the horizon is infinite. The economy consists of two groups of firms: potential acquirer firms (incumbents) and potential target firms. Potential acquirer firms can acquire and integrate new technologies to enhance their productivity. In contrast, potential target firms represent a pool of technological opportunities from which acquirer firms can draw insights.

There is a unit measure of potential acquirer firms, and each can interact with potential target firms through diffusion and tech M&A processes.

Diffusion Process

Potential acquirer firm i with productivity Z can learn from the potential target group through Poisson meetings. The firm experiences a Poisson arrival rate α to draw a random target j with quality h , and have the cumulative quality distribution of potential targets follow $G_I(t, h)$ at time t , where I is a binary variable indicating the source of technological opportunity:

- $I = 0$: Technology from potential target firms that are independent (independent target)
- $I = 1$: Technology from potential target firms following acquisition (acquired target)

The target did not exit after the acquisition since it represents a technology opportunity/blueprint that can be used partly as a “non-rivalrous” resource.

Upon meeting a potential target, the acquirer chooses the innovation rate i_{ind} (meet with an independent target) or i_{acq} (meet with an acquired target) with cost $k_{diff}(h)i_{ind}^\psi Z$ or $k'_{diff}(h)i_{acq}^\psi Z$, where $k_{diff}(h)$ is a scale factor that decreases in h (signifying that higher-quality ideas are easier to learn from or more fruitful given the fixed step size increase in productivity), and here I assume the specific form of $\frac{k_{diff}}{h}$. ψ is the elasticity of research output with respect to research spending. If innovating successfully, the acquirer's productivity Z rises by step size $\lambda_I - 1$. Different step sizes are allowed to reflect that the potential benefit from the combination with an acquired target could differ from that of independent targets.

Tech M&A Process

Technology M&As are an additional mechanism for potential acquirers to access external ideas. Acquisition opportunities arrive at a Poisson rate β , targeting only independent firms. Here, this study assumes the chance will always result in a successful acquisition. Upon the acquisition of a target with quality h , the incumbent's productivity Z_a jumps to:

$$Z_{new} = \lambda_1 \cdot Z_a h^\nu$$

Where $\nu \in (0,1)$ governs the target-specific contribution (This formulation follows the modeling approach for productivity evolution after M&A in David (2021)).

More importantly, λ_1 the scale parameter reflects the general improvement in the combined technology following acquisition. We can interpret this as the scale benefit from the continued development of the target's technology. Target status will transition from independent to acquired status after the tech M&A, and thus, the improvement could benefit other firms that subsequently encounter the acquired target.

The acquirer also incurs a linear cost kZ_a associated with the acquisition.

Impediment to Diffusion (Stronger Knowledge Hoarding)

Post-acquisition, the acquirer chooses the effort $\delta \in (0,1)$ to restrict knowledge flow to other firms. Higher δ represents stronger hoarding (through non-compete agreements, heavier patent thickets, etc.), increasing other firms' learning costs from acquired targets (here I model $k'_{diff}(h) = k_{diff}(h)(1 + \delta)$).

The acquirer solves:

$$\max_{\delta} \delta[\textit{benefit from } Z_a \textit{ to } Z_{new}] - I_{imp}(\delta)Z$$

with cost $I_{imp}(\delta) = \chi \left(\frac{1}{1-\delta} - 1 \right) Z$ reflecting escalating frictions as $\delta \rightarrow 1$. This creates a trade-off between monopoly advantages and costly knowledge protection.

Dynamics of Potential Target Firms

There is a measure M of potential targets, each indexed by j . I denote $\Phi_I(t, h)$ as the measure of potential targets with technology quality less than h in state I at time t . Then $G_I(t, h) = \frac{\Phi_I(t, h)}{M}$ by definition. I assume finite support for the potential targets distribution (from 1 to H_{max}) to correspond with the quality distribution observed in real data (will elaborate in the calibration part on the interpretation of the bounds and how I choose H_{max} in the calibration session).

The following stabilizing forces are incorporated to help us describe the stationary distribution of potential targets later:

- At each point in time, there is an exogenous birth of independent targets at a rate κ , drawn from the distribution $B_0(h)$.
- In the meantime, there is an exit rate μ such that potential targets exit the pool of possible targets.

The Firm (Potential Acquirer) Problem

For simplicity and clarity of exposition, each incumbent with a productivity level Z generates output Z , as well as profit Z (similar approaches are retained in Perla and Tonetti (2014), Lucas and Moll (2014), and Benhabib et al. (2021)). This assumption abstracts away cost structure and pricing decisions and allows us to focus on how productivity and innovation drive firm performance and long-term growth.

Firms discount at a rate $r > 0$. The firm's expected value function $V(Z)$ satisfies the Hamilton-Jacobi-Bellman equations:

$$\begin{aligned}
rV(Z, t) &= \left\{ Z + \alpha \left[\int_h g_0(h, t) \Delta V^{meet-ind} dh + \int_h g_1(h, t) \Delta V^{meet-acquired} dh \right] \right. \\
&\quad \left. + \beta \int_h \frac{g_0(h, t)}{prob(h \text{ from ind target})} [\Delta V^{acq}] dh + \frac{\partial V(Z, t)}{\partial t} \right\} \quad (*)
\end{aligned}$$

where:

- $\Delta V^{meet-ind} = \max_{i_{ind}} [i_{ind}(V(\lambda_0 Z, t) - V(Z, t)) - k_{diff}(h) i_{ind}^\psi Z]$
- $\Delta V^{meet-acq} = \max_{i_{acq}} \left[(i_{acq} V(\lambda_1 Z, t) - V(Z, t)) - (1 + \tilde{\delta}) k_{diff}(h) i_{acq}^\psi Z \right]$
- $\Delta V^{acq} = \max_{\delta} \{ \delta [V(\lambda_1 Z h^\nu, t) - V(Z, t)] - I_{imp}(\delta) Z \} + V(\lambda_1 Z h^\nu, t) - V(Z, t) - kZ$

$g_0(h, t)$ and $g_1(h, t)$ is the density of independent targets and the quality distribution of acquired targets. The three integral parts represent productivity improvement via meeting with an independent target, meeting with an acquired target, and making a tech acquisition of a target. In the first two cases, the potential acquirer chooses its innovation rate, while in the tech acquisition case, it chooses the impediment effort.

Law of Motion of the Productivity Distribution

Let $F(Z, t)$ represents the cumulative distribution for potential acquirers with productivity less than Z . The endogenous evolution of distribution is governed by:

$$\frac{\partial F(t,Z)}{\partial t} = -\beta \int_1^{h_{max}} \int_{\frac{z}{\lambda_1 h^\nu}}^z f(Z', t) dZ' b_0(h) dh -$$

$$\alpha \left\{ \int_{\frac{z}{\lambda_0}}^z \left[\int_h i_0(h, Z') g_0(h, t) dh \right] f(Z', t) dZ' + \int_{\frac{z}{\lambda_1}}^z \left[\int_h i_1(h, Z', \tilde{\delta}(h)) g_1(h, t) dh \right] f(Z', t) dZ' \right\} \quad (**)$$

The first line on the right-hand side is the outflow due to tech M&As, while the second and third line on the right-hand side captures outflows from meetings with potential targets' technology.

5.2 Normalization, Stationarity, and Balanced Growth Paths

Definition of Balanced Growth Path

I study economies in equilibrium on balanced growth paths (BGPs), where aggregate variables grow at constant rates, and the distribution of firms is stationary when properly normalized.

Definition 1—Balanced Growth Path: A balanced growth path is a number γ , and triple functions (ϕ, i, v) on R_+ such that:

$$f(Z, t) = e^{-\gamma t} \phi(Ze^{-\gamma t})$$

$$i(Z, t; h) = i(Ze^{-\gamma t}; h)$$

$$V(Z, t) = e^{\gamma t} v(Ze^{-\gamma t})$$

Define $z = Ze^{-\gamma t}$ as the detrended productivity. Thus $Y = EZ = e^{\gamma t} Ez = ye^{\gamma t}$, and z is the relative productivity along the BGP.

Stationary Target Distribution

I first obtained the stationary distribution of potential target groups to solve the system.

As $t \rightarrow \infty$:

$$M(t) \rightarrow \frac{\kappa}{\mu}, \quad \Phi_0(t, h) \rightarrow \frac{\kappa B_0(h)}{\mu + \beta} = \Phi_0(h), \quad \Phi_1(t, h) \rightarrow \frac{\beta}{\mu} \times \frac{\kappa B_0(h)}{\mu + \beta} = \Phi_1(h)$$

Therefore: $G_0(h) = \frac{\Phi_0(h)}{M} = \frac{\mu B_0(h)}{(\mu + \beta)}$, $G_1(h) = \frac{\Phi_1(h)}{M} = \frac{\beta B_0(h)}{\mu + \beta}$ and the associated density denoted as $g_0(h)$, $g_1(h)$ and $b_0(h)$.

Stationary Bellman Equations

To further simplify the stationary Bellman equations, I made the following two assumptions:

- A1, I assume the new idea birth distribution follows the mixing distribution of a point mass and a log-normal (based on patent assignees' quality distribution):

$$B_0(h) = (1 - p_{tail}) + p_{tail} \frac{\Phi_{stn}\left(\frac{\ln h}{\sigma}\right) - \Phi_{stn}\left(\frac{\ln h_{min}}{\sigma}\right)}{\Phi_{stn}\left(\frac{\ln H_{max}}{\sigma}\right) - \Phi_{stn}\left(\frac{\ln h_{min}}{\sigma}\right)}$$

With $h_{min} = 1$, $H_{max} \sim e^5$, and Φ_{stn} is the standard normal CDF.

- A2, Following Akcigit and Kerr (2018), I set the R&D cost curvature to 2 ($\psi = 2$).

Under the above assumptions, the stationary Bellman equations in normalized variables become:

$$\begin{aligned} (r - \gamma)v(z) + \gamma z v'(z) \\ = z + \alpha \left(\frac{\mu}{\mu + \beta} \frac{(v(\lambda_0 z) - v(z))^2}{4k_{diff} z} [(1 - p_{tail}) + p_{tail} E[h | 1 \leq h \leq H_{max}]] \right) \\ + \alpha \left(\frac{\beta}{\mu + \beta} \frac{(v(\lambda_1 z) - v(z))^2}{4k_{diff} z} [(1 - p_{tail}) \frac{1}{1 + \delta^*(1)} + p_{tail} E[\frac{h}{1 + \delta^*(h)} | 1 \leq h \leq H_{max}]] \right) \\ + \beta [(1 - p_{tail}) func(1) + p_{tail} E[func(h) | 1 \leq h \leq e^5]] \end{aligned} \quad (1)$$

Where $func(h) = (1 + \delta^*(h))[v(\lambda_1 z h^\nu) - v(z)] - kz - \chi \left(\frac{1}{1 - \delta^*(h)} - 1 \right) z$.

Stationary Kolmogorov Forward Equations

The stationary KFE for the stationary distribution $\phi(z)$ is:

$$\gamma z \phi(z) = \beta \int_1^{h_{max}} \int_{\frac{z}{\lambda_1 h^\nu}}^z \phi(z') dz' b_0(h) dh + \alpha \left[\int_{\frac{z}{\lambda_0}}^z I_0(z') \phi(z') dz' + \int_{\frac{z}{\lambda_1}}^z I_1(z') \phi(z') dz' \right] \quad (2)$$

Where $I_0(z') = \int_h i_0^*(h, z') g_0(h) dh$, and $I_1(z') = \int_h i_1^*(h, z', \widehat{\delta}^*(h)) g_1(h) dh$

With $i_{j=1}^*(h, z') = h \frac{v(\lambda_j z') - v(z')}{2(1 + \widehat{\delta}^*(h)) |_{j=1} k_{diff} z'}$, the solution according to the firms' optimization problem.

5.2.1 Value Function

Under the BGP assumption (with A1 and A2) and linear value function structure, I can solve for the value function in terms of parameter values.

Proposition 1: Under appropriate parameter values and BGP assumptions, the (detrended) value function of a potential acquirer firm with productivity z can be expressed by $v(z) = Dz$, where D satisfies:

$$\left[\frac{\text{diff_term}}{A(D)} \right] D^2 + \left(\frac{\text{TechMA_term} - r}{B(D)} \right) D + \left[\frac{1 - \beta k - \beta B_0}{C(D)} \right] = 0 \quad (1')$$

With $D > 0$ and:

$$\text{diff_term} = \alpha \frac{\mu}{\mu + \beta} \frac{(\lambda_0 - 1)^2}{4k_{diff}} E_1 + \alpha \frac{\beta}{\mu + \beta} \frac{(\lambda_1 - 1)^2}{4k_{diff}} E_2$$

$$\text{TechMA_term} = \beta B_1$$

where:

- $E_1 = (1 - p_{tail}) + p_{tail}E[h|1 \leq h \leq H_{max}]$,
- $E_2 = (1 - p_{tail})\frac{1}{1+\delta^*(1)} + p_{tail}E[\frac{h}{1+\delta^*(h)}|1 \leq h \leq H_{max}]$
- $B_1 = (1 - p_{tail})(1 + \delta^*(1))[\lambda_1 - 1] + p_{tail}E[(1 + \delta^*(h))(\lambda_1 h^\nu - 1)|1 \leq h \leq H_{max}]$
- $B_0 = (1 - p_{tail})\chi\left(\frac{1}{1-\delta^*(1)} - 1\right) + p_{tail}E\left[\chi\left(\frac{1}{1-\delta^*(h)} - 1\right)|1 \leq h \leq H_{max}\right]$

Equation (1') is not an exact quadratic equation in D since the optimal impediment effort ($\delta^*(h)$) also depends on the value function. I keep the smaller value of D if two solutions arise, it corresponds to the stable solution.

Proof in the appendix A.1.

5.2.2 Growth Rate and Stable Distribution

I can directly solve the endogenous steady state distribution using equation (2) with the condition of unity integration for CDF. However, I can obtain an analytical approximation for the growth rate by imposing Pareto tail assumptions for $\phi(z)$.

Proposition 2: Under assumptions (A3), the BGP growth rate γ can be approximated by:

$$\gamma = \beta\theta[\lambda_1^{\frac{1}{\theta}} \cdot E_{mix}[h^{\frac{\nu}{\theta}}] - 1] + \alpha\theta[I_0\left(\lambda_0^{\frac{1}{\theta}} - 1\right) + I_1\left(\lambda_1^{\frac{1}{\theta}} - 1\right)] \quad (2')$$

Assumptions (A3):

1. **Pareto tail assumption for initial distribution:** $\lim_{z \rightarrow \infty} \frac{1-F(z,0)}{z^{-\frac{1}{\theta}}} = c$
2. **Linear value function:** $v(z) = Dz$

where $I_0 = \frac{D\mu}{\mu+\beta} \frac{\lambda_0-1}{2k_{diff}} E_1$, $I_1 = \frac{D\beta}{\mu+\beta} \frac{\lambda_1-1}{2k_{diff}} E_2$ and $E_{mix}\left[h^{\frac{\nu}{\theta}}\right] = (1 - p_{tail}) + p_{tail}E\left[h^{\frac{\nu}{\theta}}|1 \leq h \leq e^5\right]$

Assumption 1 in A3 is a boundary condition for (2) to be uniquely determined.

Proof in the appendix A.2.

5.2.3 Mechanism Discussion

The model generates endogenous growth through the interaction of three key mechanisms:

Technology Acquisition Channel: Tech M&As directly boost the acquirer's productivity through $\lambda_1 Zh^v$, creating immediate productivity gains that depend on both acquirer and target characteristics.

Knowledge Diffusion Channel: The diffusion process allows firms to learn from both independent (λ_0) and acquired (λ_1) targets through Poisson meetings, but acquired targets might provide additional learning opportunities due to the positive spillover effect from the tech M&As.

Impediment Effects: Post-acquisition knowledge hoarding (δ) builds a trade-off where acquirers can extract more value from acquired knowledge but simultaneously reduce aggregate diffusion, affecting economy-wide growth.

The balanced growth rate depends on the relative strength of these channels. Higher M&A rates (β) boost growth through direct acquisition effects and indirectly through diffusion from targets acquired. The magnitude of the indirect impact is determined by positive spillovers from the technology combination and negative impediment effects from the acquirer. The model thus captures both the private benefits and social costs of technology acquisitions from a technology diffusion perspective in determining aggregate innovation outcomes.

6. Quantitative Analysis

This segment explains how the model is quantified. I discipline the calibration of key parameters to match the model's balanced growth path values, drawing on data from the 1980–2015 period. I

will first discuss the calibration strategy. Next, I will demonstrate how well the model fits the data. Finally, I will apply the calibrated parameters to do counterfactual analysis and explore the effect of tech M&A on growth through the diffusion channel.

6.1 Calibration

Data sources. I use multiple datasets to calibrate the model and construct targeted moments. Patent assignees and their patents and citation data from the USPTO identify potential targets. The sample of public firms with non-zero R&D expenditures during the sample period from Compustat maps to acquirers in the model. Tech M&A transaction data comes from SDC Platinum and is consistent with empirical analysis.

I calibrate in three steps to ensure proper identification and computational tractability.

Step 1: In the first step, eight parameter values are from three sources: (1) externally sourced values from prior studies, (2) parameters statistically estimated from the sample data (e.g., distributional parameters), and (3) parameters directly observed in the data without estimation (e.g., relative firm size).

For externally sourced parameters, I fix the interest rate at $r = 0.4$, a standard long-run macro value. For the target contribution ν David (2021) treats post-M&A productivity as a convex combination of acquirer and target components and assigns the target a weight of 0.5. In the specification, the combined technology is linear in the acquirer component; to keep the total post-M&A improvement comparable after shifting more weight to the acquirer, I set ν to 0.4.

Distributional parameters $(\theta, c, \mu_h, \sigma_h)$ are estimated directly from 2015 firm-level data by fitting a statistical model. The estimated distributional parameters for the potential target and potential acquirers are not stable during the sample period; thus, I use the most recent year in the sample for the calibration exercise. For potential acquirers of the Pareto tail (see appendix B for the tail fit), I estimate the maximum likelihood of sales per patent stock. I use the patent assignees' average five-year citation since granted count to represent the target's quality. I model the exogenous quality kernel $B_0(h)$ as a two-part mixture: a point mass at the benchmark (median) quality, normalized

to 1, and a truncated log-normal tail for $h > 1$. Specifically, I map all assignee qualities below the benchmark to $h = 1$; I fit the remaining portion with a log-normal distribution truncated to $[1, e^5]$. After this normalization, the support of target quality is $[1, e^5]$.

For the same reason as above, I also use single-year cross-sectional data to directly calculate M and κ . I set M using the relative size of potential acquirers over potential targets in 2015 and κ The number of patent assignees (normalized by the size of potential acquirers) whose first innovation occurred in 2015.

Step 2: Separately identify exogenous process parameters. Three key process parameters are calibrated using their direct analytical relationships to observable moments from sample data. The target exit rate is determined from the relative mass of potential targets and acquirers ($\frac{\kappa}{\mu} = M$), the acquisition rate from the ratio of independent target assignees over all assignees ($\frac{\mu}{\beta + \mu}$, relatively stable during the sample period), and the diffusion meeting rate from the share of zero R&D firms (relatively stable during the sample period).

Step 3: Jointly calibrated internal parameters. Due to their interdependence, the remaining five parameters governing core innovation and knowledge flow mechanisms require simultaneous identification. I use indirect inference to jointly estimate these parameters by targeting five key moments calculated from patent and financial data from 1980 to 2015.

This paper uses zero non-self-citation share for independent assignees (1985–2010 average to avoid truncation issue, see appendix B) to pin down step size $\lambda_0 - 1$. A higher step size will result in a higher innovation rate and thus a lower share for zero citation assignees. I use the DID estimates in the empirical part (Counterfactual citation ratio between acquired and independent patents) to identify step size $\lambda_1 - 1$. The DID estimates result in 1.41 times higher non-self-citations in the first 5 years post-acquisition, thus for 10-year annual average incidence ratio, the counterfactual non-self-citation ratio is 1.705 ($IRR_{10} = \frac{5 \cdot IRR_1 + 5 \cdot IRR_2}{10} = \frac{5 \times 2.41 + 5 \times 1}{10}$), revealing higher acquisition learning premium λ_1 than baseline learning factor λ_0 . To identify the impediment cost χ , I exploit the difference in appropriation rates between acquired and independent patents using a 14-year self-citation rate (6 years pre and 7 years after). The financial moments pin down cost parameters through the R&D-to-profits ratio for diffusion and acquirer

abnormal returns around M&A announcements. I target 4% for acquirer surplus based on recent structure estimates, removing the revelation effects (Wang, 2018).

For the joint calibration, I minimize the distance between model-implied moments and data counterparts using:

$$\min \sum_{i=1}^5 \frac{|model(i) - data(i)|}{\frac{1}{2}|model(i)| + \frac{1}{2}|data(i)|}$$

Table 6.1 presents an overview of the calibrated parameters. The cost for diffusion innovation is lower than for imitation or acquisition, and the step size for combining with acquired targets is twice as large as for independent targets. The tech acquisition rate is lower than the diffusion rate. These patterns reflect that most patents remain independent, with only a small fraction changing ownership through tech acquisition.

Table 6.1: Calibration Results

Parameter	Description	Method	Value
r	Real interest rate	External	0.04
v	Target's contribution to productivity increases after the acquisition	External	0.40
θ	Tail parameter of acquirers' prod. Distribution	Direct estimation	1.39
μ _h , σ _h	Target quality distribution parameters	Direct estimation	(0, 1.36)
M	Relative target mass	Direct calculation	58.8
μ	Target exit rate	Separate calibration	0.12
β	Acquisition rate	Separate calibration	0.0013
α	Diffusion meeting rate	Separate calibration	3.22

Parameter	Description	Method	Value
$\lambda_0 - 1$	Step size for successful innovation by combining independent targets	Joint calibration	0.0008
$\lambda_1 - 1$	Step size for successful innovation by combining acquired targets	Joint calibration	0.0017
k_{diff}	Scale of the cost for diffusion innovation	Joint calibration	0.01
k	Acquisition cost parameter	Joint calibration	5.87
χ	Impediment cost scale	Joint calibration	1.014

6.2 Model Fit and Validation

I use a method of moments approach to infer key model parameters. The model fit is presented for both targeted and non-targeted moments, delivering confidence that the model is consistent with observed data patterns, particularly regarding patent diffusion and technology growth.

Table 6.2 reveals the calibration results, demonstrating an effective fit across all targeted moments. The model strictly matches key empirical regularities, including the zero citation share for independent patents (0.148 vs. 0.15 in data), the citation ratio between acquired and independent targets (1.717 vs. 1.705), and the appropriation rate difference (0.25 vs. 0.25).

Table 6.2: Targeted Moments in Model and Data

Targeted Moments	Key Parameter	Model Value	Data Value
Share of zero R&D firms	α	0.04	0.04
Independent target share	β	0.985	0.985
Relative mass targets/acquirers	μ	58.8	58.8

Targeted Moments	Key Parameter	Model Value	Data Value
Zero non-self-citation share for independent assignees	λ_0	0.148	0.15
Citation ratio (acquired/independent)	λ_1	1.717	1.705
R&D-to-profits ratio (diffusion)	k_{diff}	0.10	0.10
Abnormal returns (acquirer)	k	0.04	0.04
Appropriation rate difference	χ	0.25	0.25

Non-targeted moments

The model fit for non-targeted moments is summarized in Table 6.3.

Table 6.3: Non-targeted Moments in Model and Data

Non-targeted Moments	Model Value	Data Value
TFP growth rate	0.75%	1.1%
Annual external citation rates for targets (range)	1.2	0.2–1.4
Annual external citing patents for potential acquirers (range)	8	4–12
Acquired intangibles amortization intensity (range)	0.17%	0.13%–0.57%

Growth rate in the economy. In the model economy, output is driven by productivity growth from firm innovation. The calibrated model generates an aggregate technology growth rate of 0.75% under the balanced growth path (Proposition 2), somewhat lower than the average TFP growth for

the non-farm business sector in the US economy since 1948 (1.1%, data calculated based on BLS Annual total factor productivity and related measures for major industries, linked SIC-NAICS). This difference is reasonable given that the model captures only technological growth's diffusion and tech M&A channels. I verify this result by directly solving equation (2) using calibrated parameters, yielding a similar growth rate, confirming that Proposition 2 provides a good analytical approximation.

Expected non-self-citation level for targets and acquirers. The model generates citation rates of approximately 1.2 for independent target patents, matching the observed average number of non-self-citations received by independent targets in one year, which ranges from 0.2 to 1.4. While citation rates for independent targets exhibit temporal variation, the model produces a value within the historical range of stable periods around 2000. Dynamics in the target pool contribute to trends before and after 2000, though our model does not capture these dynamics. The model also indicates that potential acquirers receive an average of 8 external citing patents per year, consistent with the observed data shown in Figure 6.2.

Acquired intangible amortization intensity. Amortization value reflects the cost of developing and using intangible assets over their lifespan, resulting in expenses being recognized in the financial statements. I use the direct acquired intangible amortization measure to capture the upper bound expense on impediment efforts for potential acquirers. Based on Compustat intangible asset data, I construct the amortization intensity of acquired intangibles using annual depreciation rates ($\delta = \text{amortization} / (\text{amortization} + \text{intangibles})$), tracking acquired intangible stock dynamics, and computing acquired-only amortization as a share of total profits. Figure 6.3 depicts the values for 2012-2015, with data before 2011 unavailable. The empirical amortization intensity ranges from 0.0013 to 0.0057, and my model-implied cost of 0.0017 falls in the lower part of this range, suggesting reasonable estimates of the cost for impediment effort.

Tail distribution from public firms (potential acquirers). Figure 6.4 compares the steady state distribution for potential acquirers with the Pareto tail estimated from 2015 data. The model-implied tail approximates the data at the far end of the productivity distribution but is thinner. In typical idea flow models, productivity distributions exhibit fat Pareto tails due to the infinite support of source distributions. The model features finite support for the target quality distribution, resulting in thinner tails than the typical Pareto distribution. Although the tail estimate deviates

from observed behavior, this has a minimal impact on the growth rate since the probability mass in the tail is sufficiently small to alter the balanced growth rate. Since the focus here is on aggregate averages rather than distributional outcomes, the model would require modifications when capturing the inequality in the acquirer's productivity.

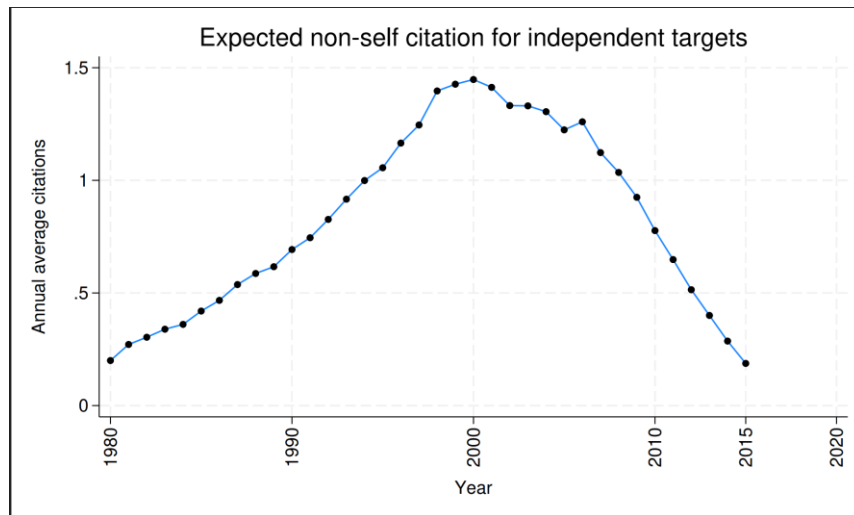


Figure 6.1: Citation received by independent targets over time

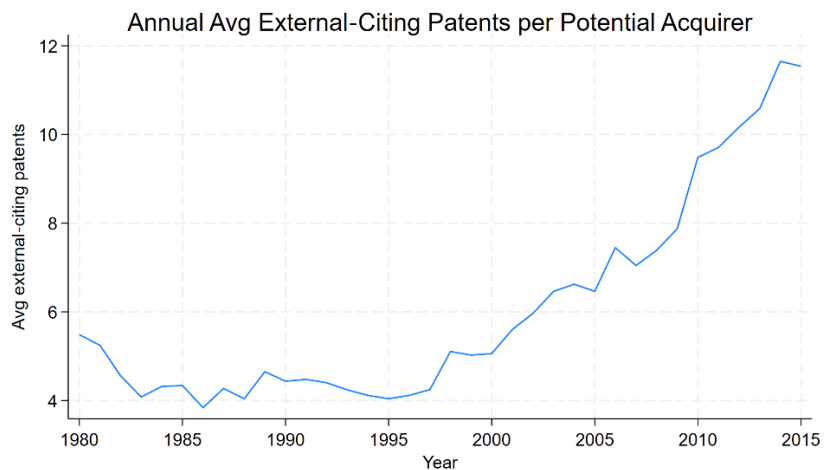


Figure 6.2: Overall average external citing patents across potential acquirers

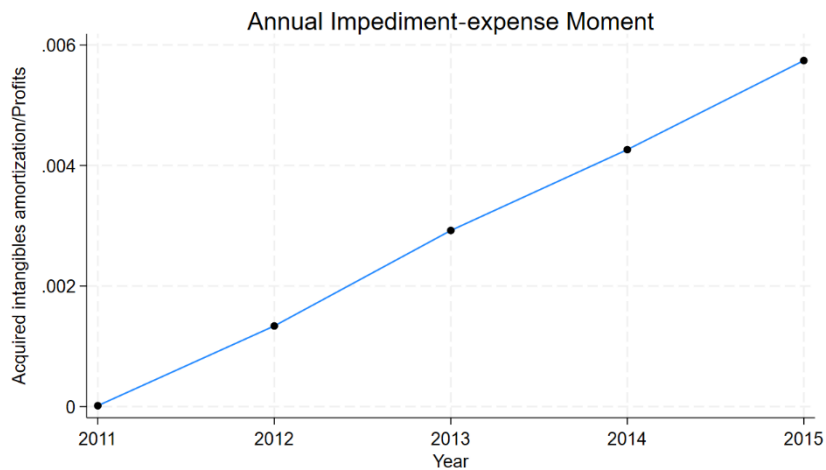


Figure 6.3: Acquired intangible amortization intensity

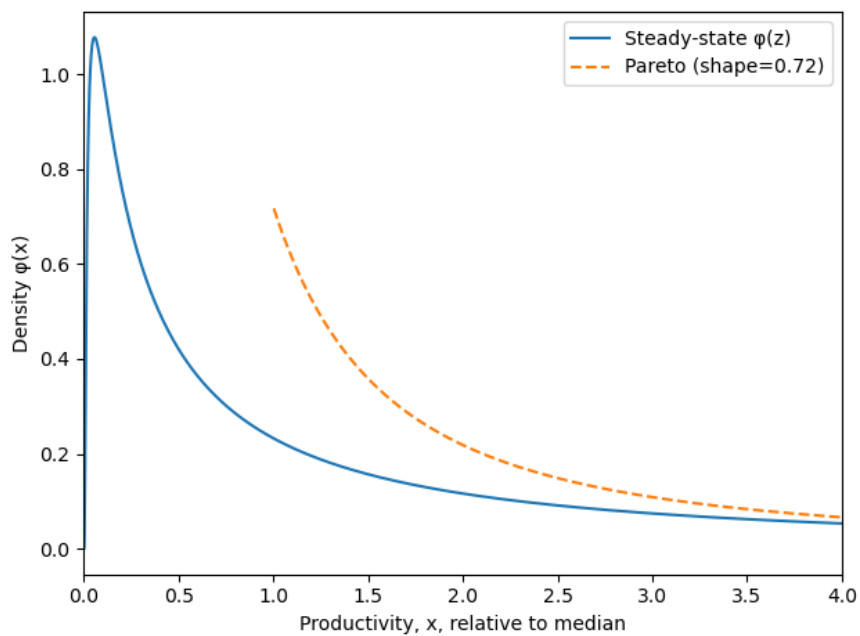


Figure 6.4: Productivity density in Model and Data

Based on calibrated parameters, the model has reasonable predictions about the steady state technology growth rate and average diffusion patterns measured by patent citation data. In the

next section, a counterfactual analysis will be conducted to learn the aggregate effects of tech M&A on diffusion and growth.

6.3 Aggregate Effects of Technology M&A on Diffusion and Growth

This section presents counterfactual experiments using the calibrated model to decompose the aggregate effects of technology M&As on economic growth and knowledge diffusion. By systematically varying key parameters while holding others constant, I isolate three distinct channels through which technology acquisitions affect aggregate outcomes: (1) the innovation synergy effect, (2) the knowledge appropriation mechanism, and (3) the acquisition frequency channel.

6.3.1 Counterfactual 1: Eliminating the M&A Learning Premium

In the baseline calibration, acquired targets exhibit a non-self-citation ratio of 1.7 relative to independent targets, doubling the effective productivity step size when building on acquired targets. This first counterfactual isolates the contribution of acquisition-specific positive spillovers by eliminating additional innovation boosts from tech M&As. For example, acquisitions where targets' technologies are discontinued, put in maintenance mode, or left undeveloped by acquirers—like Google's acquihires (Sparrow), EA's studio shutdowns (Westwood), or Oracle's legacy software acquisitions—provide no enhanced development that would stimulate greater external knowledge diffusion.

I reduced the acquisition jump parameter λ_1 such that the innovation success probability from acquired targets equals that from independent targets ($i_{acq} = i_{ind}$). This adjustment guarantees that acquisitions no longer confer additional technology spillovers for outsiders beyond ordinary Poisson-meeting learning, effectively setting the relative spillover gain ($\lambda_1 - \lambda_0$) to zero.

Figure 6.5 explains how the spillover parameter (λ_1) affects the relative citation ratio between acquired and independent targets. To eliminate M&A's diffusion advantage, λ_1 must fall from its baseline of 1.0016 to 1.0009—nearly matching the independent learning parameter (λ_0). Table 6.4 states the resulting counterfactual effects on growth and diffusion rates.

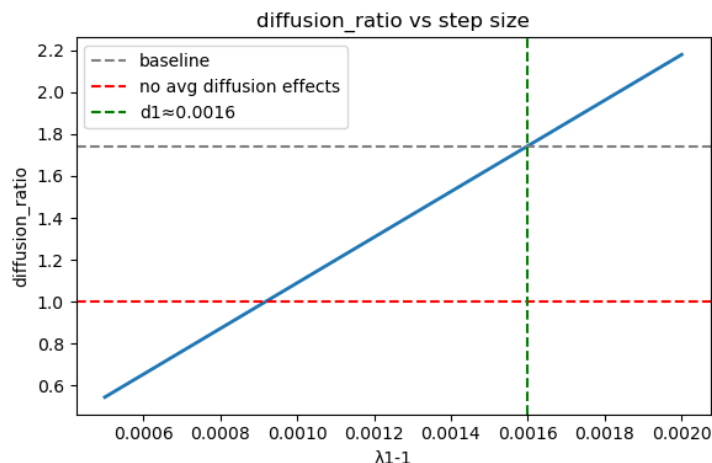


Figure 6.5: Relationship between non-self-citation ratio and spillover parameter

Table 6.4: No Tech M&A Learning Premium Effect

Parameter	Baseline	Counterfactual	Change
Spillover parameter (λ_1)	1.0016	1.0009	-44%
Value function (D)	27.54	27.48	-0.22%
Diffusion rate (Λ)	1.21	1.20	-0.83%
Growth rate (γ)	0.72%	0.71%	-1.38%

The results demonstrate that synergy-driven positive spillovers contribute approximately 1.4% to the total growth rate in the economy. Diffusion rate is calculated as the average annual citations per target patent. A diffusion change of 0.01 means each patent generates 0.01 fewer citations per year across the economy's target patents. While small per patent, when multiplied by thousands of patents over the years, represents a meaningful change in economy-wide knowledge flows.

6.3.2 Counterfactual 2: Strengthening Post-Merger Knowledge Appropriation

This experiment investigates how enhanced ability to protect acquired knowledge affects aggregate outcomes. In practice, acquirers could invest in various mechanisms—legal protections, or technical barriers—to limit competitors' access to acquired technologies. Acquirers can more clearly keep acquired knowledge exclusively when new laws and regulations make it easier and cheaper for acquirers to block competitors/outside from learning acquired technologies. Examples include stronger trade secret protection (like the 2016 Defend Trade Secrets Act), relaxed patent disclosure requirements, or court decisions allowing companies to refuse technology licensing.

This paper reduces the diffusion-impediment cost parameter χ such that the acquirer's optimal appropriation rate δ^* doubles from its baseline level (21% to 42%). Figure 6.6 displays the effort response to impediment cost reduction. The kink in the response function reflects the mass point where agents transition from corner to interior solutions for optimal extra appropriation rate (impediment effort).

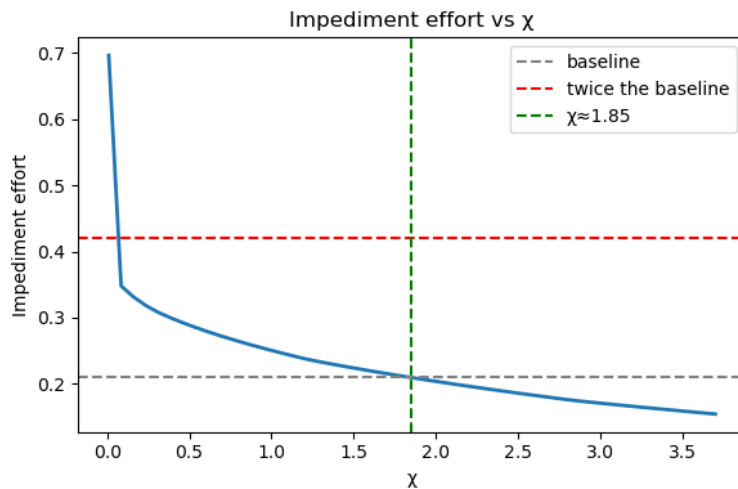


Figure 6.6: Relationship between impediment effort and impediment cost

Table 6.5: Strengthening Post-merger Knowledge Appropriation

Parameter	Baseline	Counterfactual	Change
Cost parameter (χ)	1.84	0.036	-98%
Appropriation rate (δ)	0.21	0.42	+100%
Value function (D)	27.54	27.64	+0.36%
Diffusion rate (A)	1.21	1.21	0%
Growth rate (γ)	0.7217%	0.7218%	0%
Growth rate (from acquired target diffusion)	0.02%	0.018%	-10%

Surprisingly, enhanced appropriation ability has little impact on growth and diffusion rates. This counterintuitive result emerges from two offsetting forces:

1. Direct effect: Higher appropriation reduces spillovers from acquired targets, limiting technology diffusion
2. Indirect effect: Increased acquisition values (due to better appropriation) stimulate more innovation building upon independent targets (due to higher firm value), offsetting the direct negative effect on diffusion and innovation.

By decomposing the source of growth using (2'), it can be seen that growth from building upon independent targets rises slightly, while growth from building upon acquired targets falls (last row in Table 6.5). The near-zero net effect suggests that post-merger knowledge hoarding, despite its distributional consequences, does not harm aggregate growth.

6.3.3 Counterfactual 3: Varying Acquisition Frequency

The last counterfactual exercise explores how the frequency of M&A opportunities affects aggregate outcomes. Since acquired targets generate higher spillovers than independent targets in the baseline calibration, increasing the acquisition rate β shifts the steady state distribution toward

more acquired targets, potentially accelerating productivity growth through enhanced diffusion. Although in our model the acquisition frequency is exogenously determined, policies such as reducing M&A friction through tax incentives for IP acquisitions, streamlined regulatory approval (higher Hart-Scott-Rodino (HSR) Act thresholds), or safe harbors for emerging technology deals could be used to effectively increase the tech M&As.

Two scenarios are considered: halving and doubling the baseline acquisition rate. These variations capture potential policy interventions that restrict or facilitate tech M&A activity. Table 6.6 shows the results.

Table 6.6: Effects of Acquisition Frequency on Growth and Diffusion

Scenario	Tech M&A Rate (β)	Growth Rate (γ)	Diffusion (Δ)
Low rate (50% baseline)	0.00065	0.70%	1.200
Baseline	0.0013	0.72%	1.209
High rate (200% baseline)	0.0026	0.77%	1.225

Doubling the M&A rate improves growth by 0.05 percentage points, with knowledge spillovers from acquired targets contributing 40% (0.02% to 0.04%) of this gain and direct productivity improvements within merged entities contributing 60%. While the direct channel dominates, the diffusion effect remains economically significant. Conversely, halving the M&A rate decreases growth comparably to eliminating all learning premiums (Counterfactual 1). Since policymakers can more readily influence acquisition frequency through regulatory changes than engineer technological synergies between firms, these results propose that facilitating technology M&As—provided they maintain positive synergies—offers an effective lever for enhancing long-run growth.

The counterfactuals in section 6.3 explain that the growth effect of tech M&A through the diffusion channel matters. Policy should therefore facilitate synergy-creating acquisitions while cautiously managing appropriation incentives, recognizing that M&A benefits extend beyond the merging firms through diffusion effects.

6.4 Sensitivity Analysis

This study evaluates the robustness of the main results by changing key calibration targets and examining their impact on policy outcomes. This sensitivity analysis serves two purposes: first, to establish bounds on the quantitative predictions given uncertainty in the empirical moments; second, to verify that the qualitative policy implications are not artifacts of specific parameter choices. I focus the sensitivity analysis on the acquisition frequency counterfactual (doubling/halving β) as it generates the largest policy effects and is most relevant for regulatory decisions. The robustness of the other counterfactuals and more parameters are discussed in the appendix C.

6.4.1 Sensitivity to the Zero Citation Share

The zero citation share ($M1$) captures the fraction of patents receiving no forward citations within the observation window. This moment is particularly important as it identifies the baseline rate of innovation success and knowledge diffusion in the absence of M&A activity. Variations are observed in $M1$ from its baseline value of 0.15 to alternative values spanning the historical range observed in the data:

- $M1 = 0.10$: Lower value from 1985–2010 data (periods of high citation intensity)
- $M1 = 0.15$: Baseline calibration (median historical value)
- $M1 = 0.20$: Upper value from 1985–2010 data (periods of low citation intensity)

These values reflect substantial temporal variation in citation patterns, potentially driven by changes in patenting practices, technological opportunities, or USPTO examination procedures.

Table 6.5 presents the sensitivity of counterfactual experiments to alternative calibrations of M1. The Table reports both the recalibrated benchmark values and the effects of varying acquisition frequency (β) on growth and diffusion outcomes.

Table 6.5: Sensitivity Analysis for Zero Citation Moment (M1)

Calibration Target	Benchmark Values (Growth Rate, Diffusion)	Δ Growth Rate (Half β)	Δ Growth Rate (Double β)	Δ Diffusion (Half β)	Δ Diffusion (Double β)
0.10	0.70%, 1.88	-0.02 pp	+0.05 pp	-0.01	+0.03
0.15 (Baseline)	0.72%, 1.21	-0.02 pp	+0.05 pp	-0.01	+0.01
0.20	0.78%, 0.87	-0.03 pp	+0.04 pp	-0.01	+0.01

Note: Changes in growth rate are in percentage points (pp). Diffusion changes are in levels.

It may be perceived as counterintuitive that higher zero citation rates generate lower diffusion but higher growth. This reflects that the growth rate is more responsive to innovation step sizes than diffusion costs. When M1 expands from 0.10 to 0.20, maintaining calibration targets requires both costlier diffusion (k_{diff} increases 5-fold) and larger innovation jumps (λ values double). The growth impact of bigger innovation steps dominates the negative effect of reduced spillovers, yielding higher growth (0.78% vs 0.70%) despite lower diffusion (1.88 to 0.87). The main results for this sensitivity analysis are as follows:

Policy effects are bounded in absolute terms. Despite varying M1 by 100% (from 0.10 to 0.20), growth effects remain tightly bound: halving β reduces growth by 0.02–0.03 percentage points, while doubling β increases it by 0.04–0.05 percentage points. This variation of at most ± 0.01 percentage points provides reliable bounds for policy evaluation. Similarly, diffusion effects are stable: halving β consistently reduces the citation rate by 0.01 (a 0.5–1.2% decrease depending on baseline), while doubling β increases it by 0.01–0.03 (a 0.8–1.6% increase).

Relative sensitivity is moderate despite some variation. While certain effects show non-trivial relative changes, for instance, the diffusion impact of doubling β varies from 0.01 to 0.03 (a threefold difference), and the growth effect of halving β ranges from -0.02 to -0.03 pp (a 50% relative change), most policy effects remain stable. The growth impact of doubling β (0.04–0.05 pp) and the diffusion effect of halving β (consistently -0.01) show minimal variation. All specifications preserve the sign and economic significance of impact.

6.4.2 Robustness to other moments and summary

In the Appendix C, sensitivity results to the R&D-to-profits ratio (M3), distribution parameters in different years, and the external parameter measure (ν) are presented. All qualitative and quantitative results remain unchanged.

In summary, the sensitivity analysis confirms that the paper's economic conclusions are robust to calibration uncertainty. Although certain parameter values (such as diffusion cost k_{diff} or innovation step size) may have relatively large changes with different target values, the economic implications from the counterfactual analysis are stable.

7. Conclusion

This paper examines the effect of technology M&As on the diffusion of target firms' technologies using comprehensive patent and M&A data from 1980 to 2021. It was found that tech M&As significantly enhance the external diffusion of targets' technologies, as measured by increased patent citations from other firms. These positive spillovers are particularly strong for acquisitions by non-patenting acquirers and among firms within the acquirer's industry. Through structural estimation, I quantify that doubling the tech M&A rate would increase aggregate productivity growth by up to 0.05 percentage points annually, with the diffusion channel accounting for approximately 40% of this effect.

While this study documents the importance of the diffusion channel in tech M&As, several limitations point to productive avenues for future research. First, although this study identifies significant diffusion effects, the precise mechanisms driving these "synergies" remain

underexplored. Future work could develop more granular measures of acquirers' network capabilities, scale advantages, and technological complementarities to decompose how much enhanced diffusion stems from technological improvements versus market-based synergies.

Second, the limited coverage in the LSEG SDC M&A database constrains the empirical results. As Jin et al. (2024) demonstrate, specialized databases such as 451 Research (S&P Global) contain more comprehensive coverage of technology acquisitions, particularly for smaller deals that could exhibit different diffusion patterns. Replicating this analysis with more complete data would strengthen external validity and potentially reveal richer heterogeneous effects.

Third, the theoretical model simplifies assumptions by treating potential targets and acquirers as distinct groups with an exogenous target distribution to isolate the diffusion channel. While this approach provides clarity, it abstracts from important firm dynamics and strategic interactions. Future theoretical work could endogenize the selection into target status, allow firms to transition between groups, and incorporate richer strategic considerations in acquisition and innovation decisions. Such extensions would better capture the dynamic interplay between M&A activity, technological progress, and market structure, providing deeper insights into how tech M&As shape aggregate growth and innovation distribution across firms.

Reference:

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., & Kerr, W. (2018). Innovation, reallocation, and growth. *American Economic Review*, *108*(11), 3450-3491.

Akcigit, U., & Ates, S. T. (2021). Ten facts on declining business dynamism and lessons from endogenous growth theory. *American Economic Journal: Macroeconomics*, *13*(1), 257-298.

Akcigit, U., & Ates, S. T. (2023). What happened to US business dynamism? *Journal of Political Economy*, *131*(8), 2059-2124.

Alvarez, F. E., Buera, F. J., & Lucas, R. E. (2008). *Models of idea flows* (No. w14135). National Bureau of Economic Research.

Arora, A., Belenzon, S., & Sheer, L. (2021). Knowledge spillovers and corporate investment in scientific research. *American Economic Review*, *111*(3), 871-898.

Arts, S., Hou, J., & Gomez, J. C. (2021). Natural language processing to identify the creation and impact of new technologies in patent text: Code, data, and new measures. *Research Policy*, *50*(2), 104144.

Autor, D., Dorn, D., Hanson, G. H., Pisano, G., & Shu, P. (2020). Foreign competition and domestic innovation: Evidence from US patents. *American Economic Review: Insights*, *2*(3), 357-374.

Ayerst, S. (2022). The diffusion of new general-purpose technologies. In the *Working Paper*.

Bena, J., & Li, K. (2014). Corporate innovations, mergers, and acquisitions. *The Journal of Finance*, *69*(5), 1923-1960.

Benhabib, J., Perla, J., & Tonetti, C. (2021). Reconciling models of diffusion and innovation: A theory of the productivity distribution and technology frontier. *Econometrica*, *89*(5), 2261-2301.

Bowen III, D. E., Frésard, L., & Hoberg, G. (2023). Rapidly evolving technologies and startup exits. *Management Science*, *69*(2), 940-967.

Bryan, K. A., & Hovenkamp, E. (2020). Antitrust limits on startup acquisitions. *Review of Industrial Organization*, *56*(4), 615-636.

- Cao, L., & Zhu, L. (2025). Mergers and acquisitions and the aggregate markup. *Macroeconomic Dynamics*, 29, e115.
- Cavenaile, L., Celik, M. A., & Tian, X. (2019). Are markups too high? competition, strategic innovation, and industry dynamics. *Competition, Strategic Innovation, and Industry Dynamics* (August 1, 2019).
- Cunningham, C., Ederer, F., & Ma, S. (2021). Killer acquisitions. *Journal of Political Economy*, 129(3), 649-702.
- David, J. M. (2021). The aggregate implications of mergers and acquisitions. *The Review of Economic Studies*, 88(4), 1796-1830.
- De Loecker, J., Eeckhout, J., & Unger, G. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2), 561-644.
- Ederer, F., Seibel, R., & Simcoe, T. (2025). *Digital (Killer?) acquisitions*. Working Paper.
- Ewens, M., & Marx, M. (2023). *Entrepreneurial patents*. Working Paper.
- Fons-Rosen, C., Roldan-Blanco, P., & Schmitz, T. (2021). The effects of startup acquisitions on innovation and economic growth. *Max Planck Institute for Innovation & Competition Research Paper*, (23-02).
- Ganapati, S. (2021). Growing oligopolies, prices, output, and productivity. *American Economic Journal: Microeconomics*, 13(3), 309-327.
- Goldschlag, N., Lybbert, T. J., & Zolas, N. J. (2016). An ‘algorithmic links with probabilities’ crosswalk for uspc and cpc patent classifications with an application towards industrial technology composition. *US Census Bureau Center for Economic Studies paper no. CES-WP-16-15*.
- Hall, B. H., & Ziedonis, R. H. (2001). The patent paradox revisited: an empirical study of patenting in the US semiconductor industry, 1979-1995. *RAND Journal of Economics*, 101-128.
- Harford, J. (2005). What drives merger waves?. *Journal of Financial Economics*, 77(3), 529-560.
- Heckman, J. J., Ichimura, H., & Todd, P. (1998). Matching as an econometric evaluation estimator. *The Review of Economic Studies*, 65(2), 261-294.

- Jin, G. Z., Leccese, M., & Wagman, L. (2024). M&A and technological expansion. *Journal of Economics & Management Strategy*, 33(2), 338-359.
- Jo, K., & Kim, S. (2024). Heterogeneous Innovations and Growth Under Imperfect Technology Spillovers.
- Kim, J. D. (2024). Startup acquisitions as a hiring strategy: Turnover differences between acquired and regular hires. *Strategy Science*, 9(2), 118-134.
- Klette, T. J., & Kortum, S. (2004). Innovating firms and aggregate innovation. *Journal of Political Economy*, 112(5), 986-1018.
- Kogan, L., Papanikolaou, D., Seru, A., & Stoffman, N. (2017). Technological innovation, resource allocation, and growth. *The Quarterly Journal of Economics*, 132(2), 665-712.
- Lerner, J., & Seru, A. (2022). The use and misuse of patent data: Issues for finance and beyond. *The Review of Financial Studies*, 35(6), 2667-2704.
- Li, K., & Wang, J. (2023). Inter-firm inventor collaboration and path-breaking innovation: Evidence from inventor teams post-merger. *Journal of Financial and Quantitative Analysis*, 58(3), 1144-1171.
- Liu, O. (2022). The Dynamics of Firm Size Inequality: The Role of Acquisition and Innovation. Available at SSRN <https://doi.org/10.2139/ssrn.4279062>
- Lucas Jr, R. E., & Moll, B. (2014). Knowledge growth and the allocation of time. *Journal of Political Economy*, 122(1), 1-51.
- Luttmer, E. G. (2012). Technology diffusion and growth. *Journal of Economic Theory*, 147(2), 602-622.
- Miller, N., Osborne, M., Sheu, G., & Sileo, G. (2023). Technology and market power: The united states cement industry, 1974-2019. *Georgetown McDonough School of Business Research Paper*, (4041168).
- Mitchell, M. L., & Mulherin, J. H. (1996). The impact of industry shocks on takeover and restructuring activity. *Journal of Financial Economics*, 41(2), 193-229.

- Muratori, U. (2020). Knowledge diffusion, markups, and cohorts of firms. *Unpublished Manuscript*.
- Perla, J., & Tonetti, C. (2014). Equilibrium imitation and growth. *Journal of Political Economy*, 122(1), 52-76.
- Phillips, G. M., & Zhdanov, A. (2013). R&D and the incentives from merger and acquisition activity. *The Review of Financial Studies*, 26(1), 34-78.
- Roberts, M. R., & Whited, T. M. (2013). Endogeneity in empirical corporate finance. In *Handbook of the Economics of Finance* (Vol. 2, pp. 493-572). Elsevier.
- Seru, A. (2014). Firm boundaries matter: Evidence from conglomerates and R&D activity. *Journal of Financial Economics*, 111(2), 381-405.
- Shapiro, C., & Yurukoglu, A. (2024). Trends in Competition in the United States: What Does the Evidence Show? NBER Working Paper No. w32762, Available at SSRN: <https://ssrn.com/abstract=4908573>
- Shenoy, J. (2012). An examination of the efficiency, foreclosure, and collusion rationales for vertical takeovers. *Management Science*, 58(8), 1482-1501.
- Teece, D. J. (1986). Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy. *Research Policy*, 15(6), 285-305.
- Teece, D. J., Pisano, G., & Shuen, A. (1997). Dynamic capabilities and strategic management. *Strategic Management Journal*, 18(7), 509-533.
- Trajtenberg, M., Henderson, R., & Jaffe, A. (1997). University versus corporate patents: A window on the basicness of invention. *Economics of Innovation and New Technology*, 5(1), 19-50.
- Wang, W. (2018). Bid anticipation, information revelation, and merger gains. *Journal of Financial Economics*, 128(2), 320-343.
- Williamson, O. E. (1975). Markets and hierarchies: analysis and antitrust implications: a study in the economics of internal organization. *University of Illinois at Urbana-Champaign's Academy for Entrepreneurial Leadership Historical Research Reference in Entrepreneurship*.

Zhang, Y., & Tong, T. W. (2021). How vertical integration affects firm innovation: Quasi-experimental evidence. *Organization Science*, 32(2), 455-479.

Federal Trade Commission. (2021, December 2). FTC sues to block \$40 billion semiconductor chip merger [Press release]. <https://www.ftc.gov/news-events/news/press-releases/2021/12/ftc-suesblock-40-billion-semiconductor-chip-merger>

M&ACommunity. (2024, August 15). The largest tech acquisitions in history. <https://mnacommunity.com/insights/largest-tech-mergers-and-acquisitions/>

Appendix:

Data Appendix:

Industry Composition of tech M&A Sample

Table DA.1 presents the industry composition of target and acquirer firms in the full sample of tech M&A transactions between 1980 and 2021.

Table DA.1. Industry Composition of Tech M&A Deals (1980–2021)

Target Industry	Count	Percent	Acquiror Industry	Count	Percent
Healthcare Equipment & Supplies	502	13.61%	Healthcare Equipment & Supplies	416	11.28%
Software	245	6.64%	Pharmaceuticals	283	7.67%
Semiconductors	243	6.59%	Semiconductors	254	6.89%
Machinery	226	6.13%	Machinery	223	6.04%
Pharmaceuticals	216	5.86%	Software	209	5.67%
Other Industrials	185	5.01%	Other Industrials	201	5.45%
Computers & Peripherals	166	4.50%	Computers & Peripherals	160	4.34%
Electronics	157	4.26%	Electronics	145	3.93%
Telecommunications	138	3.74%	Aerospace & Defense	120	3.23%
Equipment					
Biotechnology	132	3.58%	Biotechnology	119	3.23%
Top 10 industries total	2210	59.91%		2130	57.74%

Note: This table reports the industry distribution of target and acquirer firms in the full sample of tech M&A transactions between 1980 and 2021. Industry classifications are based on the SDC Platinum database's mid-industry codes. While this appendix presents the full tech M&A sample, the regression sample is more concentrated in Pharmaceuticals, Oil & Gas, and Healthcare Equipment due to the availability of comparable failed deals used to define the control group.

Model Appendix:

A Derivation and proofs

A.1 proof of proposition 1

I use a guess-and-verify approach. Insert $v(z) = Dz$ into equation (1), I obtain:

$$LHS = (r - \gamma)Dz + \gamma Dz = rDz$$

The right-hand side components are:

- Meeting with independent target: $\alpha \frac{\mu}{\mu+\beta} \frac{D^2(\lambda_0-1)^2}{4k_{diff}} E_1 z$
- Meeting with acquired target: $\alpha \frac{\beta}{\mu+\beta} \frac{D^2(\lambda_1-1)^2}{4k_{diff}} E_2 z$
- Acquisition term: $[\beta B_1 \cdot D - \beta k - \beta B_0]z$

where E_1, E_2, B_1, B_0 are the relevant terms defined in proposition 1.

Thus:

$$RHS = z + z \cdot \left[\alpha \frac{\mu}{\mu+\beta} \frac{D^2(\lambda_0-1)^2}{4k_{diff}} E_1 + \alpha \frac{\beta}{\mu+\beta} \frac{D^2(\lambda_1-1)^2}{4k_{diff}} E_2 + \beta B_1 \cdot D - \beta k - \beta B_0 \right]$$

Equating the left- and right-hand sides yields:

$$\left[\alpha \frac{\mu}{\mu+\beta} \frac{(\lambda_0-1)^2}{4k_{diff}} E_1 + \alpha \frac{\beta}{\mu+\beta} \frac{(\lambda_1-1)^2}{4k_{diff}} E_2 \right] D^2 + (\beta B_1 - r)D + [1 - \beta k - \beta B_0] = 0$$

This expression for D is identical to that in Proposition 1.

Q.E.D.

A.2 proof of proposition 2

Using the linear value function $v(z) = Dz$, I simplify equation (2) to obtain:

$$\gamma z \phi(z) = \beta \int_1^{h_{max}} I_a b_0(h) dh + \alpha \left[I_0 \int_{\frac{z}{\lambda_0}}^z \phi(z') dz' + I_1 \int_{\frac{z}{\lambda_1}}^z \phi(z') dz' \right] \quad (A.1)$$

where $I_0 = \frac{D\mu}{\mu+\beta} \frac{\lambda_0-1}{2k_{diff}} E_1$, $I_1 = \frac{D\beta}{\mu+\beta} \frac{\lambda_1-1}{2k_{diff}} E_2$, and $I_a = \int_{\frac{z}{\lambda_1 h^\nu}}^z \phi(z') dz'$.

For large z , I assume $\Phi(z) = 1 - cz^{-\frac{1}{\theta}}$ and $\phi(z) = \frac{1}{\theta} cz^{-\frac{1}{\theta}-1}$. Therefore:

$$\int_b^a \phi(z') dz' = \Phi(a) - \Phi(b) = cb^{-\frac{1}{\theta}} - ca^{-\frac{1}{\theta}} \quad (A.2)$$

Express both side of equation (A.1) using (A.2):

$$LHS = \gamma \frac{1}{\theta} cz^{-\frac{1}{\theta}}$$

$$RHS = \beta \int_1^{h_{max}} cz^{-\frac{1}{\theta}} ((\lambda_1 h^\nu)^{\frac{1}{\theta}} - 1) b_0(h) dh + \alpha cz^{-\frac{1}{\theta}} \sum_0^1 I_j [\lambda_j^{\frac{1}{\theta}} - 1]$$

Eliminating $cz^{-\frac{1}{\theta}}$ from both sides and defining $E_{mix}[h^{\frac{\nu}{\theta}}] = (1 - p_{tail}) + p_{tail}E[h^{\frac{\nu}{\theta}} | 1 \leq h \leq e^5]$:

$$\gamma = \beta \theta [\lambda_1^{1/\theta} \cdot E_{mix}[h^{\nu/\theta}] - 1] + \alpha \theta [I_0(\lambda_0^{1/\theta} - 1) + I_1(\lambda_1^{1/\theta} - 1)]$$

Q.E.D.

B Distribution fit and Quantification moments

B.1 Pareto tail fitting for potential acquirers

To fit the tail distribution of potential acquirers' productivity, I construct a productivity proxy using sales per patent stock for each Compustat firm in my sample. I assume that firms with zero patents are in the lower tail of the distribution (25% have no patent stock) and therefore do not affect tail estimation. I then estimate the tail parameters (c , θ) using log-log linear regression.

Figure B1 demonstrates that Pareto parameters provide superior fitting for the distribution above the 90th percentile compared to lognormal fitting. While I use 2015 data in the main analysis, the shape parameter is slightly smaller in earlier years ($1/\theta = 0.89, 0.96, 0.98$ for 2000, 1990, and 1985, respectively), but this choice does not affect my main conclusions.

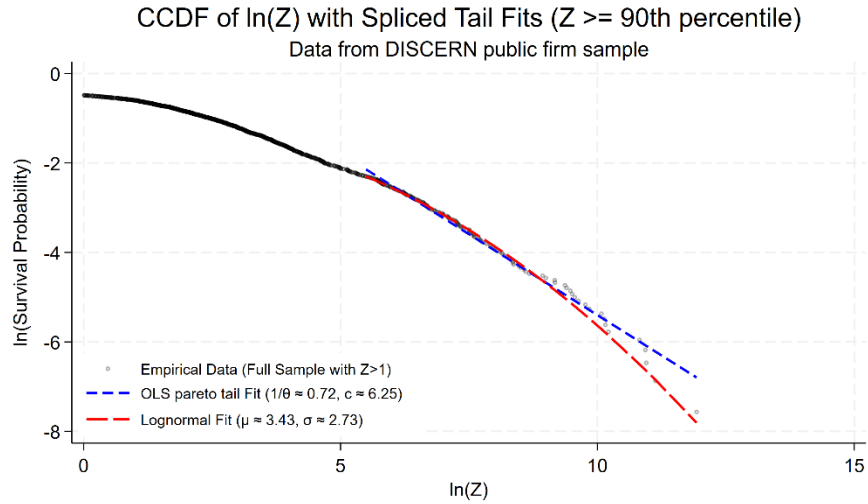


Figure B1: Pareto tail fitting for potential acquirers using 2015 data

B.2 Tail fitting for potential targets

As discussed in the quantification section, I use a log-normal distribution to fit the quality of target technology beyond the median quality level.

Figures B2-B3 show the fitting of the above-median distribution for citation-based quality measures using both Pareto and log-normal distributions. The log-normal distribution aligns better with the original data points, with even better fits for earlier years.

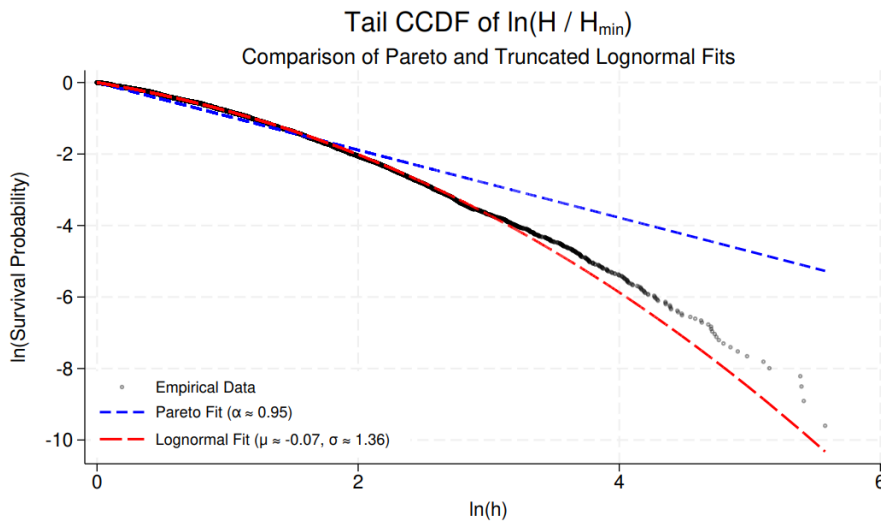


Figure B2: Tail fitting for potential targets using 2015 data

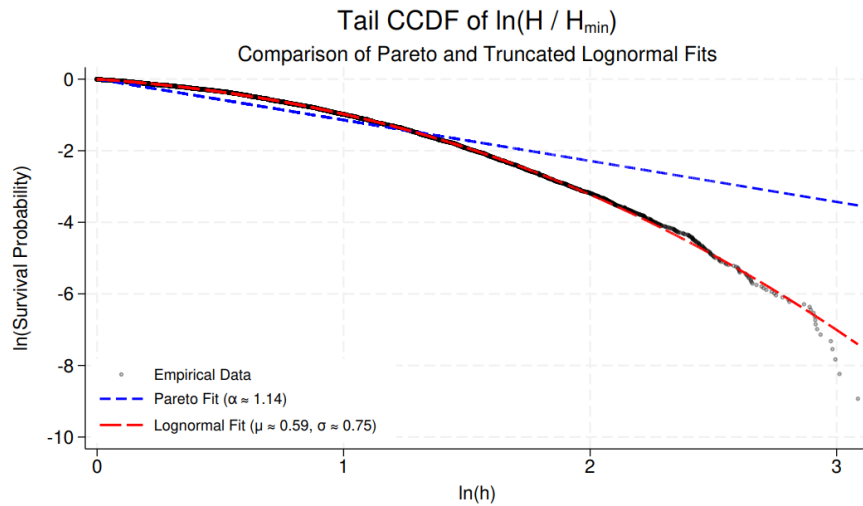


Figure B3: Tail fitting for potential targets using 1985 data

B.3 Data moments for zero non-self-citation share for independent assignees

To help identify λ_0 , I calculate the share of patent assignees that receive zero non-self-citations for each year. For each assignee-year, I first count their newly granted patents' 10-year citations, averaging to the assignee-year level. I then collapse over patent grant years to obtain the cross-sectional share of zero non-self-citation assignees across different years.

Figure B4 shows the values for each grant year. Since I examine 10-year citations, data after 2010 are less reliable as they only capture 5-year citations in my sample. I use the average value from 1985-2010 as the final moment in my calibration exercise.




Figure B4: Fraction of assignees with zero 10-year citations at each grant-year

C Other sensitivity results

In this appendix, I present additional sensitivity analyses to assess the robustness of the main calibration results. Specifically, I examine the effects of varying (i) the R&D-to-profits ratio (moment M3), (ii) the external parameter ν , and (iii) the source distribution parameters using alternative years.

Sensitivity to the R&D-to-Profits Ratio (Moment M3)

I first vary the calibration target for the R&D-to-profits ratio around its baseline value of 0.1, considering values of 0.12 (high, corresponding to the upper bound observed between 1990 and 2015) and 0.08 (low, corresponding to the lower bound). As reported in Table C1, benchmark growth rates and diffusion levels change with the calibration target. However, the comparative statics with respect to changes in the acquisition arrival rate (β) remain essentially unchanged. Doubling β continues to raise the growth rate by approximately 0.05 percentage points and the diffusion rate by around 0.01, while halving β produces the opposite effect.

Sensitivity to the External Parameter ν

Next, I examine sensitivity to the external parameter ν , which I set to 0.4 in the baseline. I consider values of 0.5 (a 25% increase) and 0.3 (a 25% decrease). The results, summarized in Table C2, show that small changes in ν barely affect the benchmark growth rate or diffusion levels. More

importantly, the quantitative effects of varying β remain stable: doubling β raises growth by around 0.05 percentage points and diffusion by around 0.01.

Sensitivity to the Source Distribution Parameters

Finally, I re-estimate the model using alternative source distribution parameters corresponding to earlier years. The baseline calibration uses the 2015 distribution (0,1.36). I also consider the 2005 distribution (0.37,1.03) and the 1985 distribution (0.59,0.75). As shown in Table C3, benchmark growth and diffusion levels differ slightly across these cases, but the direction and magnitude of the comparative statics with respect to β remain virtually identical.

Across all three exercises, I find that the quantitative implications of the model are robust to alternative parameter choices. Although benchmark values shift modestly depending on the calibration target, the counterfactual results remain unaffected.

Table C1: Sensitivity Analysis for R&D-to-Profits Ratio (M3)

Calibration Target	Benchmark Values (Growth Rate, Diffusion)	Δ Growth Rate (Half β)	Δ Growth Rate (Double β)	Δ Diffusion (Half β)	Δ Diffusion (Double β)
0.10 (Baseline)	0.72%, 1.21	-0.02 pp	+0.05 pp	-0.01	+0.01
0.12	0.96%, 1.30	-0.02 pp	+0.05 pp	-0.01	+0.004
0.08	0.69%, 1.31	-0.03 pp	+0.04 pp	-0.01	+0.01

Table C2: Sensitivity Analysis for External Parameter ν

Calibration Target	Benchmark Values (Growth Rate, Diffusion)	Δ Growth Rate (Half β)	Δ Growth Rate (Double β)	Δ Diffusion (Half β)	Δ Diffusion (Double β)
0.3	0.71%, 1.21	-0.02 pp	+0.04 pp	-0.01	+0.01

Calibration Target	Benchmark Values (Growth Rate, Diffusion)	Δ Growth Rate (Half β)	Δ Growth Rate (Double β)	Δ Diffusion (Half β)	Δ Diffusion (Double β)
0.4 (Baseline)	0.72%, 1.21	-0.02 pp	+0.05 pp	-0.01	+0.02
0.5	0.73%, 1.21	-0.03 pp	+0.06 pp	-0.01	+0.01

Table C3: Sensitivity Analysis for Source Distribution Parameters

Calibration Target	Benchmark Values (Growth Rate, Diffusion)	Δ Growth Rate (Half β)	Δ Growth Rate (Double β)	Δ Diffusion (Half β)	Δ Diffusion (Double β)
(0.371,0.03)	0.78%, 1.05	-0.02 pp	+0.04 pp	-0.01	+0.01
(0,1.36) (Baseline)	0.72%, 1.21	-0.02 pp	+0.05 pp	-0.01	+0.01
(0.59,0.75)	0.75%, 0.93	-0.022 pp	+0.05 pp	-0.01	+0.01