THE IMPACT OF MERGERS ON CORPORATE INNOVATION

EVIDENCE FROM DECADES OF PATENT ACTIVITY

Mark Binnermark

Carl Björkman

Bachelor Thesis
Stockholm School of Economics
2025

The Impact of Mergers on Corporate Innovation: Evidence from Decades of Patent Activity

Abstract:

This paper revisits the long-standing question of whether mergers stimulate or stifle innovation activity. The dataset comprises 4,767 completed U.S. mergers and 1,654 withdrawn attempts announced from 1980 to 2018, resulting in 58,476 firm-year observations within balanced five-year windows. We construct three core innovation outcomes that capture innovation scale, innovation novelty, and private value of innovations. We find that merger completion reduces the scale of target-firm innovation, measured by the quantity of generated patents, and also reduces the novelty, measured by citation intensity of patents. However, private value, deduced from abnormal stock-market returns, market anticipation, and adjusted for inflation using the CPI, increases for target firms. Our findings pertaining to acquirers show that their innovation output improves across the board, over the full period studied. We show that the aggregate effect of mergers on innovation is nuanced and provide novel insights on different aspects, which are broadly consistent with the existing literature. Our findings are relevant for antitrust enforcement, the allocation of resources within internal capital markets, and the strategic design of post-merger R&D integration.

Keywords:

Citations, Mergers, M&A, Innovation, Patents

Authors:

Mark Binnermark (25560)
Carl Björkman (25672)

Tutor:

Mustafa Bulut, Postdoctoral Fellow, Swedish House of Finance

Examiner:

Ramin P. Baghai, Professor, Department of Finance

Acknowledgement:

The authors gratefully acknowledge the invaluable supervision and guidance provided by Mustafa Bulut, which has been essential to the completion of this research.

Bachelor Thesis
Bachelor Program in Business and Economics
Stockholm School of Economics

1 Introduction

The U.S. mergers-and-acquisitions (M&A) market is both extensive and persistent. In the year 2024, almost 14 thousand deals were announced, worth a total of $1.8 trillion, including 373 mega-deals exceeding a value of $1 billion each (FactSet, 2025). Although activity has ebbed since the 2021 peak, industry forecasters anticipate renewed growth in 2025 (Institute for Mergers, Acquisitions, and Alliances (IMAA), 2025). Against this background, findings of Seru (2014) showing that patent productivity and quality decrease dramatically for firms that were targeted in successful mergers announced between 1980 and 1998, appear troubling. This is especially true, as he finds no significant effect on acquirer outcomes. At the same time, innovation is described as a key durable source of competitive advantage and macroeconomic growth (He and Tian, 2018). However, whether the apparent 'innovation penalty' to mergers persists in the modern, intangible-intensive economy remains unknown. Subsequent evidence is mixed: Some studies find complementary effects after a merger based on technological overlap (Bena and Li, 2014), whereas others document innovation slowdowns concentrated on diversifying deals (Phillips and Zhdanov, 2013).

The troubling background presented above warrants the investigation of the following two research questions: (i) How does merger completion affect innovation outcomes for targets and acquirers? (ii) Has this effect changed between the 'Seru period' (1980-1998) and the modern period (1999–2018)?

We revisit the issue of mergers' effect on innovation outcomes using a comprehensive quasi-experimental approach that examines both targets and acquirers across nearly four decades, linking CRSP/Compustat fundamentals to the 2023 update of the KPSS patent file (Kogan, Papanikolaou, Seru, and Stoffman, 2017) and mergers from SDC Platinum announced between 1980 and 2018. Our methodology largely follows the quasi-experimental design of Seru (2014): Completed deals form the treatment group, withdrawn deals the counterfactual. Innovation is measured by (i) log patent counts (innovation scale), (ii) log truncation-adjusted citations per patent (innovation novelty)($CPatent_{it}$), and (iii) log CPI-deflated private patent value $\xi_{real,it}$. Difference-in-Differences (DiD) regressions include firm- and year-fixed effects, as well as controls for size, leverage, profitability, cash, industry concentration, and unrelated merger status. Our identification strategy is validated through parallel trends test, robustness to firm-specific linear trends, and heterogeneity analysis across innovation quality measures and technological overlap.

Main findings. We find that target firms experience a 44 percent decline in patent productivity and a 51 percent decline in citation-based quality after merger completion, which is practically identical in the modern era, confirming that the penalty is robust to changes in governance and technology. However, we also find that acquirers experience an 8.5 percent increase in patent output and an 11 percent increase in citation quality for the whole period, indicating that innovation is reallocated rather than destroyed. The improvement for acquirers is concentrated in the modern period and, therefore, does not contrast with the null findings by Seru (2014). Simultaneously, these findings are consistent with the 'buying innovation' motive (Zhao, 2009; Bena and Li, 2014) and previous studies that have found a positive effect of M&A on acquirer outcomes (Zhao, 2009; Sevilir and Tian, 2012). Interestingly, when measuring innovation output through the private patent value metric created by Kogan et al. (2017), we find a significant and positive treatment effect for both target firms and acquirers. This finding is indicative of a shifting focus towards strategic patenting that captures private rents, focusing on exploitative (or incremental) R&D rather than engaging in novel research. Such a shift might entail adverse social welfare effects. Our results remain consistent when exposed to several robustness checks.

Literature review

Our paper is part of a larger literature on the interconnection of finance and corporate innovation, which has seen an increasing interest among researchers since the beginning of the century, as presented by He and Tian (2018). Romer (1986) specifies endogenous technological change as a major driver of long-term economic growth. Moreover, Porter (1992) was among the first to establish the link between asset allocation systems, internal and external, and long-term competitive advantage on the global market. These articles highlight the importance of a system that fosters innovation. Building on the above is a more recent collection of papers examining the relationship between corporate finance, innovation within firms, and the performance of those firms. This literature seeks the optimal way to organize corporate structures for innovation conductivity, along with motivations to alter corporate structures and engage in mergers specifically, or M&A in general.

As a means of measuring innovation, this paper, along with much of the cited literature (Bena and Li, 2014; Seru, 2014; Zhao, 2009; Sevilir and Tian, 2012), uses R&D output. This takes the form of the number of generated patents and the number of future citations generated for each patent to measure the quality or novelty of those patents. A valuable contribution was made by Kogan et al. (2017) in creating a measure of the private value of patents. Our paper utilizes the number of generated patents and citations received, both of which are well-established in the literature, while making a more novel contribution to the literature with an empirical application of private values to investigate innovation outcomes.

Seru (2014), the closest article to ours, finds that mergers stifle innovation for the targeted firms and explains this by the inefficiencies associated with conglomerates and internal capital markets. Examples of the expected roots of such inefficiencies are information asymmetry and agency problems. Seru finds no effect on acquirers, while a limitation is that controls are directed at target firms. The contributions to the literature by Zhao (2009) together with Sevilir and Tian (2012) are very closely related to those of Seru (2014), in terms of coverage and methodology. The main difference these two have in common (relative to Seru), lies in their focus on the effect of M&A on acquirers, along with their motivation to engage in it. Sevilir and Tian (2012) observe that, oppositely from the perspective of Seru (2014) (mergers are detrimental to the targets), successful bidders in M&A, post deal, both successfully generate more patents and generate more citations per patent. Zhao (2009) goes on to describe that innovation is a motivation for acquisitions, where firms that have previously fallen behind are more likely to engage. The results also show that acquirers that were not as innovative before taking the acquisition benefit the most. Both articles agree that indirect buying innovation is a fundamental consideration that motivates acquirers to engage in M&A. We successfully contribute to reconciling previous empirical results on the effects of both targets and acquirers.

Bena and Li (2014) instead studies the effect of M&A (with full control as an outcome) through the perspective of the transaction pair. They find that complementarities, resulting from technological overlap, positively affect the likelihood of a pairing and also predict positive future innovation outcomes. Related to the findings of Zhao (2009), Bena and Li (2014) find that acquirers typically have larger patent portfolios, while incurring lower R&D expenses. In contrast, targets engage in more R&D but are slower to generate patents. The higher patent count, combined with lower R&D expenditure, could be explained by the 'buying innovation' motive. Our paper presents results on the explanatory value of technological overlap for target outcomes, investigating whether it alleviates the detrimental impact of mergers studied by Seru (2014).

Similar to the closely related literature (Seru, 2014; Bena and Li, 2014; Zhao, 2009; Sevilir and Tian, 2012), our paper focuses on the study of publicly listed firms on U.S. exchanges. Hence, it is essential to consider our paper in the context of the existing literature on the differences between public and private firms in their innovation-related activities. This is particularly important in estimating the external validity of both the related literature and our findings. Gao, Hsu, and Li (2017) find that public companies, relative to private ones, do more exploitative research that builds incrementally on existing research. Their findings support the notion that the shorter investment horizon of investors in the stock market is a primary contributing factor. The results of (Acharya and Xu, 2017) align with this message, as the effect of public listing is beneficial only when there are external capital needs, which the authors suggest is related to short-term stock market behavior.

Our contributions to the literature above lie in part in testing the effect of mergers on innovation as studied by Seru (2014). Although the truncation bias of the patent metrics (mainly citations) for that period is less of a concern at this point (Lerner and Seru, 2022), our results support the validity of the findings presented by Seru (2014). Furthermore, we demonstrate that this effect persists in more recent years and over the full 39-year period of mergers studied. Furthermore, we link the findings of Seru with those of Sevilir and Tian (2012) and Zhao (2009), since for the modern half and the full period studied, we find that successful acquirers experience significantly better post-merger innovation outcomes than failed acquirers in the control group. Our results contribute to a combined finding, as mergers appear to transfer innovation from target firms to acquiring firms. Moreover, by empirically implementing private patent values (Kogan et al., 2017) as a measure of innovation outcomes, we contribute to the literature with additional nuance on the effects of mergers.

The remainder of the paper is organized as follows. Section 2 develops our main research questions into three testable hypotheses. Section 3 describes the data, sample construction, and key variables, including the KPSS patent value measure and our field-normalized citation and patent metrics. Section 4 presents the methodology employed in the paper, while Section 5 reports the results, including empirical estimates, as well as robustness and mechanisms. Finally, Section 6 concludes.

2 Hypothesis development

This study adopts an empirical, quasi-experimental approach to uncover the causal impact of mergers on innovation. Building on previous theory and evidence (Seru, 2014) and insights into the valuation of patents of Kogan et al. (2017), we present three testable hypotheses. The first two test and extend the negative findings by Seru (2014) to a longer sample period; the third introduces patent *private value* as an additional outcome metric.

2.1 Mergers affect R&D outcomes

Hypothesis $\mathbf{H}_1$. *Merger completion alters a target firm's innovation output—as measured by patent counts (innovation scale) and citation-weighted counts (innovation novelty)—relative to withdrawn deals.*

The null hypothesis of no effect was rejected by Seru (2014), interpreting the post-merger decline in citations as evidence that the conglomerate headquarters redistribute capital away from high-risk and high-novelty projects. Replicating Seru's design with an extended window allows us to test whether modern mergers face the same problems, which could support the attribution made by Seru to the inner workings of conglomerates. Instead, a *positive* effect would be consistent with the intended strengths of internal capital markets, where the HQs use their information advantage to allocate resources more efficiently (Porter, 1992). Using separate productivity metrics (patent) and novelty metrics (citation) allows us to distinguish between quantity-driven and quality-driven channels.

2.2 Effects are stable across periods

Hypothesis $\mathbf{H}_2$. *The merger effect on innovation is unchanged between the Seru era (1980–1998) and the modern era (1999–2018).*

If conglomerate frictions and inefficiencies are the root cause, the DiD coefficient should be equally negative in both sub-samples, unless conglomerates have become less interventionist. By reestimating the panel model separately for each era (Section 4.3.2 we test whether the innovation penalty has *weakened* amid shifts in corporate governance and technology.

2.3 Patent *private value* follows citations

Hypothesis $\mathbf{H}_3$. *The sign of the merger effect on patent <u>private</u> value mirrors that on citation-weighted value.*

Kogan et al. (2017) shows that the market-based private value reacts more strongly to firm growth opportunities than only citations. Because the measure embeds the abnormal stock price response in patent grant, it is not exposed to the same truncation bias as citations (Hall, Jaffe, and Trajtenberg, 2001). Finding opposite signs for citations versus private value would imply that conglomerates *shift* are innovative toward firm-specific rents at the expense of public scientific impact. If this were the case, it would be consistent with the short-lived findings.

3 Data, sample construction, and key variables

In Section 3.1, we describe all sources of data used for the sample construction, along with the identifiers used to merge firm observations across databases. Section 3.2. delineates the process of sample construction, outlining the filtering approaches used to create the test sample. Once the sample is created, we provide an overview in Section 3.3, which includes simple descriptive statistics and counts of observations. Lastly, in Section 3.4, we outline the creation of the key variables used in the regressions, as well as for the validation of the sample construction.

3.1 Data description

We source M&A transaction data from SDC Platinum, covering the announcement years 1980 to 2018. This database provides six-digit CUSIP identifiers for all targets and acquirers engaged in M&A deals. We link these deal-specific CUSIPs to permanent security-level PERMNO identifiers using the CRSP/Compustat Merged database linking table (accessed via WRDS). To preserve data integrity, we address potential overlap by retaining only one deal per target firm within a short time window, thereby avoiding bias from clustered acquisitions.

All financial data come from the CRSP/Compustat Merged database (covering 1960–2024), except for business segment information, which we extract from Compustat Segments (covering 1976–2018). For patent data, we rely on Kogan et al. (2017) and their 2023 update, specifically the KPSS_2023 and MatchedPatentCPC_2023 files, which are accessed through GitHub. The KPSS_2023 dataset contains patent-wise observations including patent numbers, forward citation counts, issue and filing dates, PERMNO identifiers, and CPI-adjusted patent values in 1982 $ millions (covering patents from 1926 through 2023). We merge CPC technology class data with patents using patent numbers, then link the complete patent dataset to our financial data using PERMNO identifiers.

3.2 Sample construction

From the full dataset, we draw on transactions with announcement dates between 1980 and 2018, with a 2018 cut-off to ensure data availability and sufficient time for patent outcomes of the last deals to be observed. We apply a pre-filter to the SDC data to include only deals with publicly listed U.S. target firms, in line with Seru (2014). Our first step in creating the sample is to include only transactions where the deal form was a merger. Next, we drop all deals without a friendly deal attitude (Seru, 2014; Bena and Li, 2014), as it is argued that firms targeted in hostile takeovers may be incentivized to change their R&D policy in irreversible ways, which may confound the results.

Following the approach of Seru (2014), we look to Savor and Lu (2009) for additional sample criteria. The first criterion pertains to the three-year cut-off for the deal data, although we already meet this requirement with our 2018 cut-off. Second, drawing on acquirer information in the SDC Platinum data, we exclude any deals where the acquirer is not a US-incorporated public firm. Next, we only retain merger pairs where financial data is available on the acquirer in CRSP/Compustat, by matching the acquirer to a PERMNO. Once this is done, we only keep deals where the market value of the target exceeds five percent of the acquirer's in the year preceding the announcement year, to focus the sample on deals of material impact. We obtain firm market value data from the CRSP/Compustat Merged database. Unlike much of the related M&A literature, we do not restrict our sample to cash-only or stock-only transactions. Since the choice of consideration form is orthogonal to our focus on post-merger innovation outcomes, excluding mixed-consideration deals would offer no analytical benefit.

Finally, we only include the first bid by a given acquirer for the specific target within a bidding cycle to avoid the same contested deal being repeated multiple times. This would otherwise affect the control group, particularly, and bias our estimates.

After assembling our full test sample (see Section 4.2), we apply two final filters. Completed deals from the treatment group and withdrawn deals from the control group, each subject to the eligibility criteria outlined above. Unlike Seru (2014), Bena and Li (2014), and Savor and Lu (2009), we do not manually verify that withdrawn transactions reflect R&D-orthogonal failures—both because we lack exhaustive access to Factiva / Lexis-Nexis and because our 1980-2018 sample is much larger. Although this choice may introduce additional selection risk, our control group (1,654 deals) significantly exceeds the sample sizes in prior studies (170 in Seru; 60 in Bena and Li), providing enhanced statistical power. We mitigate any residual imbalance through balance diagnostics and pre-trend analyses, as outlined in Section 4.2.

Additionally, in our main quasi-experiment, we estimate three regressions using identical specifications, differing only in the sample periods. We construct three distinct subsamples based on merger announcement dates. The first, which we refer to as the Seru period (1980-1998), replicates the time frame examined in Seru (2014). The Extended Period encompasses all announcement years in our sample, spanning from 1980 to 2018. Lastly, we have the Post Seru period, which includes the more recent half of our sample period, from announcement years 1999 to 2018.

3.3 Sample overview

Observing our treatment group of 4,767 merger transactions, Table 2 shows a breakdown of deals for each announcement year, along with information on whether the parties to the deal get granted at least one patent during the sample period. Bena and Li (2014) produces a similar table, and we observe that almost half of the deals include at least one party with at least one patent, and in about a fifth of deals, both targets and acquirers get granted a patent at some point.

Furthermore, from the annual breakdown in Table 2 we can deduce procyclical tendencies for the sample, where a substantial proportion of the deals is centered around the middle of the sample period, leading up to the dot-com bubble, as well as a clear drop after the Great Recession. From an analytical perspective, it may be beneficial that the sample is not skewed towards more recent observations, as even with controls and adjustments, truncation bias can still have a significant impact on the results (Lerner and Seru, 2022). The original sample of Seru (2014) (1980–1998) is disproportionately weighted toward the later years of that window, which can introduce truncation bias when post-event outcomes are censored by the sample end. By contrast, we extend our horizon through 2018 and employ a symmetric $\pm$5-year event panel, ensuring balanced pre- and post-merger coverage and mitigating this bias. Moreover, our filtered set of innovation-relevant transactions represents a declining share of overall U.S. M&A activity over time, (Institute for Mergers, Acquisitions, and Alliances (IMAA), 2025).

Table 1 shows that a substantial share of our sample firms belong to SIC 6 (financial services), industries traditionally characterized by minimal R&D intensity and patenting activity. Consistent with this sectoral mix, many firms record zero patent grants. This pattern reflects the underlying industry composition rather than any anomaly in our data. Seru (2014) does not exclude any industry from the quasi-experiment sample, and neither do we. In total, we conducted our main tests on a sample of 46,118 treatment and 8,225 control group firm-year observations.

Table 1: One-Digit SIC Composition of Filtered Panel (1980–2018)

SIC	Industry	Firm–Years	% of FYs	Unique Firms	% of Firms	Avg FY/Firm
0	Agriculture, Forestry & Fishing	176	0.3	23	0.3	7.7
1	Mining & Construction	2 788	4.8	432	5.1	6.5
2	Manufacturing (nondurables)	7 016	12.0	944	11.1	7.4
3	Manufacturing (durables)	12 933	22.1	1 868	21.9	6.9
4	Transportation & Public Utilities	4 377	7.5	681	8.0	6.4
5	Wholesale & Retail Trade	4 947	8.5	735	8.6	6.7
6	Finance, Insurance & Real Estate	16 428	28.1	2 251	26.4	7.3
7	Services	7 487	12.8	1 204	14.1	6.2
8	Public Administration	2 041	3.5	347	4.1	5.9
9	Non-classifiable	283	0.5	40	0.5	7.1

Table 2: Annual Breakdown of Included Mergers (Full Test Sample)

Year	All Deals	Acquirers with Patent	Targets with Patent	Acquirers *or* Targets Patent	Acquirers *and* Targets Patent
1980	4	1	2	2	1
1981	39	13	23	25	11
1982	50	9	26	27	8
1983	78	13	38	39	12
1984	156	42	85	90	37
1985	110	25	55	60	20
1986	85	23	39	44	18
1987	116	22	53	57	18
1988	112	31	55	60	26
1989	128	31	71	72	30
1990	147	37	73	79	31
1991	115	26	52	57	21
1992	114	27	39	46	20
1993	116	17	42	45	14
1994	175	26	64	67	23
1995	208	38	63	68	33
1996	250	37	88	96	29
1997	249	41	79	86	34
1998	274	44	99	107	36
1999	242	58	85	102	41
2000	223	46	80	92	34
2001	202	43	71	82	32
2002	121	33	49	53	29
2003	120	26	53	56	23
2004	97	23	43	46	20
2005	107	36	57	63	30
2006	101	35	52	58	29
2007	126	40	63	69	34
2008	109	33	47	55	25
2009	73	25	33	37	21
2010	59	19	35	37	17
2011	69	17	30	30	17
2012	70	25	40	43	22
2013	68	21	32	35	18
2014	81	20	34	39	15
2015	112	37	55	58	34
2016	104	27	47	49	25
2017	80	21	31	33	19
2018	77	24	32	36	20
Total	4 767	1 112	2 015	2 200	927

Notes: "Acquirers or Targets Patent" counts deals where at least one side patents during the sample window; "Acquirers and Targets Patent" requires both parties to do so.

3.4 Patent data and innovation metrics

Our patent data come from the 2023 release of the Kogan et al. (2017) database (hereafter *KPSS 23*).[1] The file links patents granted between 1926 and 2023 to the CRSP/Compustat PERMNOs using a hybrid of administrative identifiers and probabilistic matching, covering more than 90.00 % of patents assigned to publicly listed firms. Relative to legacy NBER files, KPSS 23 (i) extends the time span well beyond our sample window, (ii) provides market-based private-value estimates for each patent, and (iii) supplies harmonized Cooperative Patent Classification (CPC) codes and an updated citation network Kogan et al. (2017).

We construct three core innovation outcomes that capture citation-weighted knowledge flow, field-adjusted patenting intensity, and market valuation. A fourth, unadjusted citation measure is retained solely for robustness checks. All variables are winsorized at the 1 percent level to mitigate extreme outliers,[2] monetary figures are deflated to 1982 dollars using the Consumer Price Index (CPI).

As a means of measuring innovation, it is well established in the adjacent literature (Bena and Li, 2014; Seru, 2014; Zhao, 2009; Sevilir and Tian, 2012), to use R&D output. This takes the form of the number of generated patents and the number of future citations generated for each patent to measure the quality (also referred to as novelty or scientific value) of those patents. The motivation to use this approach is derived from Hall, Jaffe, and Trajtenberg (2005), which presents a strong link between patent citations and profitable firm growth. In terms of the treatment of patent metrics, we receive our main insights from Hall et al. (2001), Seru (2014), and Lerner and Seru (2022). This primarily pertains to truncation bias stemming from lags between patent applications and eventual grants, as well as even greater lags associated with citations received by patents reported more recently.

As noted by Hall et al. (2005), down-the-line citations of patents are highly correlated with market reaction informed patent value and more informative compared to what economists can deduce from historical data on citations. Kogan et al. (2017) builds on this idea by constructing their private patent value. They also find that their measure is strongly positively correlated with future received citations and even more closely related to firm performance than citations. That said, despite the correlation, they do not measure the same thing.

[1] We drop design, plant, and re-issue patents (<1 percent of total) and assign each remaining patent to the first-listed assignee when multiple PERMNOs are recorded.

[2] Two-sided at the 1 percent tails *after* deflation and normalization.

Field-Normalised Patent Output ($\boldsymbol{Patent_{it}}$). To capture sheer inventive activity net of field inflation, we scale patent counts by contemporaneous field averages, similar to Seru (2014)

The assignment of firms to technology fields follows a systematic approach where we first identify each firm's patents through their PERMNO identifiers. For each patent p, we observe its associated CPC codes, which define its technological domain, denoted as the set $\mathcal{C}_p$. Each patent can be classified into multiple CPC classes, and firms often have patents across multiple technology classes, with $\mathcal{P}_{i,t}$ representing the set of patents for firm i in year t. Patents are weighted based on both their application year and technology class before being attributed to their respective firms in the panel.

To construct comparable measures of innovation output across firms, technologies, and time, we implement a comprehensive normalization procedure that accounts for both temporal and technological heterogeneity. First, for each CPC class c and year t, we calculate $\bar{N}_{c,t}$ as the mean patent count across all firms active in that technological field during application year t. This field-year specific normalization addresses two critical empirical challenges: (1) the secular growth in patenting activity, particularly pronounced in high-technology sectors, and (2) the variable lag between patent applications and grants, which typically ranges from two to three years and could otherwise introduce temporal biases in our measures.

For each patent p, we compute its normalized value as:

$$\frac{1}{|\mathcal{C}_p|} \sum_{c \in \mathcal{C}_p} \frac{1}{\bar{N}_{c,t}} \tag{1}$$

where $|\mathcal{C}_p|$ represents the number of CPC classes assigned to the patent. This formulation ensures that a patent's contribution is weighted equally across its technological classifications, preventing the over-counting of patents that span multiple fields. The firm-level normalized patent count aggregates these patent-specific measures across all patents in the firm's portfolio:

$$\sum_{p \in \mathcal{P}_{i,t}} \frac{1}{|\mathcal{C}_p|} \sum_{c \in \mathcal{C}_p} \frac{1}{\bar{N}_{c,t}} \tag{2}$$

This normalized measure offers several advantages for our empirical analysis. By accounting for both temporal variations in patenting intensity and the multiplicity of technological classifications, we obtain innovation metrics that are comparable across different technology domains and time periods. The normalization mitigates potential biases from technological fields with systematically higher patenting rates or processing times, while properly weighting the contribution of patents that span multiple technological domains. This approach provides a more accurate representation of firms' relative innovative activity within their respective technological spaces, which is particularly important in our merger and acquisition context, where firms often operate across multiple technological domains.

Citation-Intensity ($\boldsymbol{CPatent_{it}}$). Following Seru (2014), we also normalize forward citations by technology–year as a means to mitigate truncation and field heterogeneity:

$$CPatent_{it} \;=\; \frac{1}{N_{it}} \sum_{p \in \mathcal{P}_{it}} \frac{\text{Cites}_p}{\overline{\text{Cites}}_{c(p),\,y(p)}}, \tag{3}$$

where $\mathcal{P}_{it}$ is firm i's set of patents *applied for* in year t; $N_{it} = |\mathcal{P}_{it}|$; $c(p)$ and $y(p)$ denote the CPC three-digit field and application year of patent p; and the bar indicates the cross-sectional mean. This implicit field–year fixed effect addresses (i) long-lag citation truncation (Hall et al., 2001) and (ii) systematic citation rate differences across technologies.

Market-Based Patent Value ($\boldsymbol{\xi_{\text{real},it}}$). Our most economically salient measure aggregates the grant-date private values of individual patents:

$$\xi_{\text{real},it} \;=\; \sum_{p \in \mathcal{P}_{it}} \xi_p^{\text{KPSS}}, \tag{4}$$

Where ξ_p^{KPSS} is the CPI-deflated abnormal return-based value (in 1982 USD millions) from KPSS23 (Kogan et al., 2017). Unlike citation metrics that accumulate over decades, ξ_{real} captures the contemporaneous assessment of future cash flow contributions by investors and therefore embeds information on obsolescence risk and competitive positioning.

Raw Citation Intensity ($\boldsymbol{RawCite_{it}}$). For comparison with the early mergers literature, we also compute

$$RawCite_{it} \;=\; \frac{\sum_{p \in \mathcal{P}_{it}} \text{Cites}_p}{N_{it}}. \tag{5}$$

Because it ignores both truncation and technology effects, we use this variable only in robustness exercises (Section 5.2). Our headline results are based on the normalized measure in (3).

Unadjusted Patent Count ($\boldsymbol{RawPatent_{it}}$). We retain $RawPatent_{it} = N_{it}$ as a control variable. The KPSS 23 dataset's comprehensive coverage through 2023 provides a complete window for our $\pm$5-year panel, ensuring that patent applications from our sample period (through 2018) have had sufficient time to reach grant status. This extended coverage window substantially mitigates right-truncation concerns that have affected earlier studies in this literature.

Construction of standard controls (Size, Leverage, Herfindahl industry concentration, and the *unrelated-merger* dummy) is detailed in Appendix A.

4 Methodology

In Section 4.1, we begin by presenting our identification strategy and describing the empirical procedures that operationalize it. Section 4.2 then details our distributional evidence, while Section 4.3 details the econometric specification that underpins our main regression analysis and guides subsequent inference.

4.1 Identification Strategy

This section presents our empirical framework for identifying the causal impact of merger completion on firm-level innovation. We develop a quasi-experimental design that exploits variation in merger outcomes among firms that announce consolidation attempts, allowing us to isolate the treatment effect of actual merger completion from the strategic decision to engage in it.

4.1.1 Quasi-experimental Design

Our identification strategy builds on the insight that while the decision to announce a merger is endogenous, the ultimate completion versus withdrawal of announced deals provides exogenous variation in treatment assignment. Although complete exogeneity is not guaranteed, it is assumed to hold approximately (Seru, 2014; Bena and Li, 2014). This approach addresses the fundamental selection problem in merger research: firms that pursue acquisitions may differ systematically from those that do not, making it difficult to identify causal effects of mergers on innovation. This design maintains the initial strategic decision to pursue consolidation while exploiting variation in deal outcomes that is largely orthogonal to the target firms' ex-ante innovation potential.

The validity of our identification strategy rests on two key assumptions:

Parallel Trends Assumption. In the absence of merger completion, target firms in treatment and control groups would exhibit similar innovation trajectories. While we cannot directly test this assumption post-treatment, we provide extensive validation through pre-treatment trend analysis and event study plots that show no significant differences in innovation trends between groups prior to merger announcements.

Conditional Exogeneity. Conditional on announcing a merger and observable firm characteristics, the factors determining completion versus withdrawal are uncorrelated with unobserved determinants of innovation outcomes. Our analysis of completion determinants (Table 3b) reveals that while some firm characteristics predict completion probability, innovation measures show limited or no predictive power, supporting this assumption.

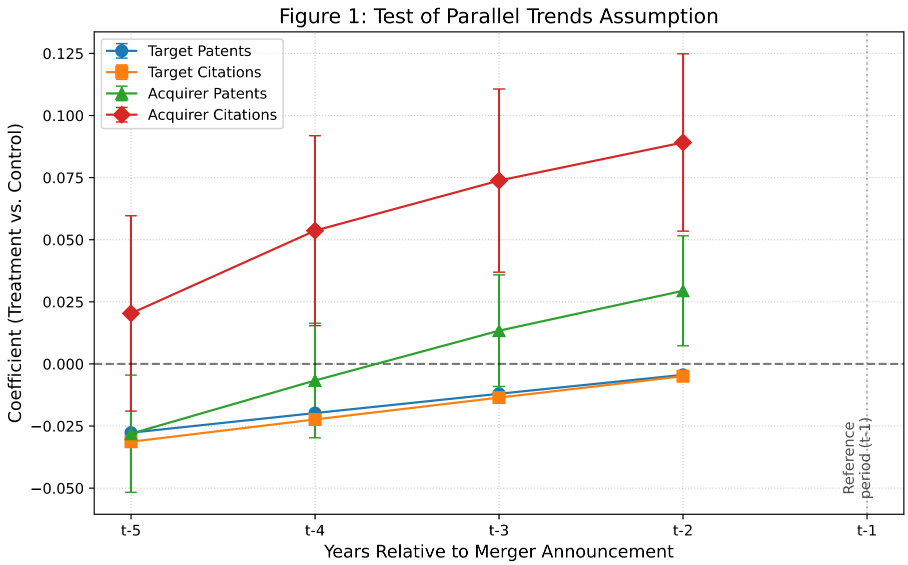

Figure 1: plots coefficients from event-time regressions examining innovation outcomes in the pre-merger period (t-5 to t-1). The dependent variables are patent counts and forward citations for both targets and acquirers. The coefficients represent differences between treatment (completed mergers) and control (withdrawn deals) firms, estimated from regressions with firm fixed effects and t-1 as the reference period. Vertical bars denote 95 percent confidence intervals with standard errors clustered at the firm level.

4.1.2 Treatment and Control Definition

We operationalize treatment assignment based on deal completion status within a 12-month window following the announcement:

Treatment Group ($T_i = 1$): Mergers completed within 12 months of announcement, where completion is defined by the effective date recorded in SDC Platinum. This timeframe encompasses the vast majority of successful deals, excluding those that face prolonged regulatory review or financing difficulties.

Control Group ($T_i = 0$): Mergers officially withdrawn or terminated prior to completion, as indicated by withdrawal dates in SDC Platinum. Control group firms experience the same announcement effects and initial integration planning as treatment firms but do not undergo actual ownership change or operational integration.

This design is appealing for several reasons. Both treatment and control firms experience identical announcement effects and initial integration planning, ensuring that any differential outcomes reflect actual consummation rather than merger intent or market expectations. The restriction to announced deals mitigates selection concerns that plague comparisons between acquiring and nonacquiring firms, as both groups have revealed similar strategic preferences and growth ambitions.

4.1.3 Selection Analysis and Completion Determinants

To examine potential selection into treatment and validate our identification strategy, we estimate the determinants of deal completion using a logistic regression:

$$\Pr\big(\text{Success}_{it} = 1\big) = \Lambda\Big(\alpha + \boldsymbol{\gamma}'\mathbf{X}_{i,t-1} + \beta_1 \text{ R\&D Productivity}_{i,t-1} + \mu_t\Big) \quad (6)$$

where $\Lambda(\cdot)$ is the logistic cumulative distribution function, $\mathbf{X}_{i,t-1}$ contains pre-announcement firm characteristics (Size, Leverage, EBITDA/Assets, industry concentration), and R&D Productivity bundles our innovation measures (Patent, CPatent, RD/Sales).

This analysis serves two purposes: (1) identifying systematic factors that predict completion probability, which we subsequently control for in our main specifications, and (2) testing whether unobserved innovation capacity drives both completion likelihood and post-merger innovation outcomes, which would violate our identification assumptions.

4.1.4 Time-period analysis

We construct a ±5-year window around the announcement for every target. Our baseline specification uses the full 1980–2018 sample, combining all observations into a single estimation of the equation (7).

To gauge whether merger effects changed over time, we additionally re-estimate the model separately for two sub-periods: the *Seru period* (1980–1998) and the *modern period* (1999–2018). This chronological split serves two purposes: (i) it allows direct comparison to the original findings of Seru (2014), and (ii) it tests whether more recent mergers, conducted under different corporate governance regimes and technology management practices, exhibit different innovation outcomes.

The exact specification, as with acquirer metrics, tests whether completion changes acquirer innovation capacity across all three samples.

4.1.5 Validation of sample construction and controls

Table 3a presents comprehensive balance tests comparing treatment and control groups during the pre-announcement period ($t \in [-5, -1]$). While we observe some statistically significant differences in firm characteristics—treatment firms are larger (Size: 6.17 vs. 2.14) but less leveraged (22 percent vs. 24 percent)—innovation measures show no statistically significant differences at conventional levels.

Table 3: Sample validation: summary statistics and balance tests

(a) Panel A: Summary statistics

	Treatment (T)			Control (C)			T − C		
Variable	Mean	Std	Median	Mean	Std	Median	Diff	t	p
Size	5.74	2.14	5.66	6.13	2.47	6.06	−0.38	−6.11	0.00 ***
Leverage	0.22	0.22	0.17	0.24	0.21	0.21	−0.02	−3.37	0.00 ***
EBITDA/Assets (%)	5.94	20.16	8.96	7.54	17.70	10.03	−1.60	−3.38	0.00 ***
$(\text{RD/Sales})^{a}$ (%)	33.87	128.83	5.82	11.96	46.46	4.70	21.91	7.02	0.00 ***
CPatent^{a}	1.31	1.47	0.88	1.23	1.26	0.87	0.08	1.66	0.10 *
Patent^{a}	0.24	0.53	0.08	0.27	0.57	0.10	−0.03	−1.42	0.16

Observations: 10 277 (T), 1 787 (C)

Note: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

(b) Panel B: Probability of deal succeeding (logit)

	Prob(Success = 1)			
	(1)	(2)	(3)	(4)
Size	-0.044^{***}	-0.044^{***}	-0.044^{***}	-0.044^{***}
	(0.015)	(0.015)	(0.015)	(0.015)
Leverage	-0.299^{**}	-0.299^{**}	-0.299^{**}	-0.299^{**}
	(0.119)	(0.119)	(0.119)	(0.119)
RD	-0.000^{**}	-0.000^{**}	-0.000^{**}	-0.000^{**}
	(0.000)	(0.000)	(0.000)	(0.000)
Patent	0.186^{*}	0.186^{*}	0.186^{*}	0.186^{*}
	(0.100)	(0.100)	(0.100)	(0.100)
CPatent	0.006	0.006	0.006	−0.005
	(0.028)	(0.028)	(0.028)	(0.031)
EBITDA/Assets	-0.003^{*}	-0.003^{*}	-0.003^{*}	-0.003^{*}
	(0.002)	(0.002)	(0.002)	(0.002)
HI		−0.022	−0.022	−0.022
		(0.204)	(0.204)	(0.204)
$\text{Dummy}^{\text{mseg}}$			−0.004	−0.004
			(0.053)	(0.053)
$\text{Dummy}^{\text{unrelated}}$				0.032
				(0.053)
HI^{acq}				−0.207
				(0.226)
$\text{CPatent}^{\text{acq}}$				0.033
				(0.039)
Observations	12,069			
Time fixed effects	Yes			

Note: Robust standard errors in parentheses. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

Patent-based innovation metrics exhibit similar distributions across groups: normalized patent counts (0.24 vs. 0.27, p=0.16) and citation-weighted measures (3.1 vs. 1.23, p=0.10) show no statistically significant differences at the 5 percent level. This pattern supports the validity of our control group for identifying causal effects on innovation outcomes, as firms do not systematically differ in their pre-announcement innovation productivity.

The observed differences in financial characteristics likely reflect the types of firms that attract acquisition interest and complete transactions. Larger firms may face fewer financing constraints and regulatory hurdles, leading to higher completion rates. In contrast, the lower leverage in treatment firms may reflect stronger balance sheets that facilitate deal completion. Our inclusion of these characteristics as time-varying controls in equation (7) ensures that systematic differences in observable firm attributes do not bias our treatment effect estimates. Looking back at Figure 1, with the controls in place, the target innovation output the target innovation output of treatment and control firms traces virtually identical trajectories throughout the five-year pre-announcement window—none of the event-time coefficients differ significantly from zero—offering a visual confirmation that the parallel-trends assumption holds. Moreover, additional justification for the inclusion of observable firm characteristics as controls, in line with Seru (2014), is presented by Lerner and Seru (2022), showing that many of them are correlated with truncation bias in citation measures.

4.2 Distributional evidence

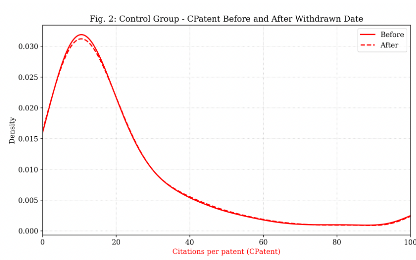

Figure 2: Target Citations (Control)

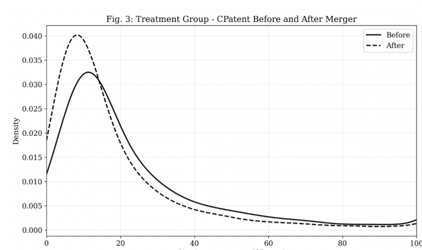

Figure 3: Target Citations (Treated)

Notes: Epanechnikov kernel, bandwidth $h = 0.35$; sample sizes $N_{\text{control}} = 8{,}225$ firm–year observations, $N_{\text{treated}} = 46{,}118$. Solid lines are five-year **pre-event** densities; dashed lines are five-year **post-event** densities. Control curves overlap, whereas treated targets shift left.

The above figures show the kernel densities of citations per patent ($CPatent_{it}$), for patents generated during five years **before** the event (solid) and **after** the event (dashed). $CPatent_{it}$ is the truncation-adjusted measure of citations per patent by a given firm in a given year (see equation (3) in Section 3.4). The only controls employed are time-fixed effects.

Figure 2 illustrates the kernel density of $CPatent_{it}$ for the control group both before and after the merger event. Control distributions overlap almost perfectly, exhibiting nearly the same characteristics.

Figure 3 illustrates the kernel density of $CPatent_{it}$ for the treatment group both before and after the merger event. The treatment group sees a clear shift to the left, toward fewer citations per patent. Mass concentrates at 0–15 citations, and the right tail thins. A Kolmogorov–Smirnov test rejects equality of the two groups at the 1 percent level, corroborating the regression evidence.

4.3 Econometric Specification

4.3.1 Baseline Difference-in-Differences Model

Our baseline specification employs a difference-in-differences (DiD) framework that compares innovation outcomes between treatment and control groups before and after merger completion:

$$Y_{it} = \alpha + \beta_1 \text{After}_{it} + \beta_3(\text{After}_{it} \times T_i) + \boldsymbol{\delta}'\mathbf{Z}_{it} + \mu_i + \gamma_t + \varepsilon_{it} \tag{7}$$

where Y_{it} represents firm i's innovation outcome in year t, taking values $\log(1 + \text{Patent}_{it})$, $\log(1 + \text{CPatent}_{it})$, or $\log(1 + \xi_{\text{real},it})$. The log transformation accommodates the highly skewed distribution of innovation measures while preserving zero values through the $\log(1 + x)$ specification.

The key variables are:

- $\text{After}_{it} = \mathbf{1}[t \geq 0]$: indicator for post-announcement periods
- $T_i = \mathbf{1}[\text{Deal Completed}]$: treatment indicator for completed mergers
- $\text{After}_{it} \times T_i$: interaction term capturing the treatment effect
- $\mathbf{Z}_{it}$: time-varying controls including Size, Leverage, EBITDA/Assets, Cash/Assets, industry concentration (Herfindahl index), and diversification indicators
- μ_i: firm fixed effects absorbing time-invariant heterogeneity and the level effect of treatment T_i
- γ_t: fiscal year fixed effects capturing aggregate trends and market conditions
- ε_{it}: idiosyncratic error term

The coefficient of interest, β_3, captures the average treatment effect of merger completion on innovation outcomes. Under our identification assumptions, β_3 provides an unbiased estimate of the causal impact of merger completion. Note that β_2 (the main effect of treatment T_I), commonly included in standard DiD specifications, is intentionally excluded from our model. This omission is methodologically correct because our firm fixed-effects μ_I absorbs all time-invariant firm characteristics, including the treatment indicator T_i. Since each firm's treatment status remains constant across all periods in our panel, T_i would be perfectly collinear with the firm fixed effects, making separate identification impossible and unnecessary.

All controls part of $\mathbf{Z}_{it}$ are based on levels reported for the fiscal year preceding the merger announcement year. Furthermore, we cluster standard errors at the firm level to account for serial correlation within firms and potential correlation across multiple merger events involving the same firm (Bertrand and Mullainathan, 2004). This clustering approach is conservative and accounts for any residual dependence in innovation outcomes within firms over time.

5 Empirical results

This section quantifies the causal impact of merger completion on innovation, examining how corporate consolidation affects both the quantity and quality of R&D output. We begin by presenting baseline difference-in-differences estimates that exploit the quasi-experimental setting of completed versus withdrawn mergers.

The analysis unfolds in several parts. First, we examine the main treatment effects for both targets and acquirers across different time periods (Section 5.1), presenting estimates of the effect on $Patent_{it}$ and $CPatent_{it}$. Next Section 5.2 presents estimates of the treatment effect on private values ($\xi_{\text{real},it}$) of patents generated by both targets and acquirers. Section 5.3 presents key robustness checks, including specifications with firm-specific linear trends, and quality-weighted patent outcomes. We then explore heterogeneity across technological overlap and time periods (Section 5.4), which yields insights into how merger effects have evolved over the past four decades. We translate our coefficient estimates into economically interpretable magnitudes (Section 5.5) before concluding with a discussion of implications for competition policy and innovation incentives (Section 5.6).

5.1 Main treatment effect

Table 4 presents our main DiD results across three observation windows, revealing asymmetries between target and acquirer innovation outcomes following merger completion.

Targets. Across 1980-2018, the post-merger decline in target innovation is both large and highly significant. Even for withdrawn deals, patenting falls $\approx 18\%$ $\left[e^{-0.196}-1\right]$ after the announcement year. Completion triggers an additional 44% reduction, and citations exhibit an overall 51% collapse—evidence of a broad dampening of R&D productivity at the target level. Estimates for the targets are similar across all three observation windows, in terms of significance, sign, and magnitude for both citations and patent counts.

Acquirers. Control-group acquirers also reduce patenting ($\beta_1 = -0.227^{***}$), but the DiD treatment effect is positive, as shown by the interaction term estimate ($\beta_3 = 0.082^{**}$). Citation results are directionally similar, with the interaction (0.106^{*}) significant at the 10

5.2 Private value treatment effect

The positive treatment effect for acquirers presented in Table 4 for the modern and full periods is similarly reflected in Table 5 ($\beta_3 = 0.220^{***}$) in terms of the valuation of patents generated post-treatment. Interestingly, the treatment effect for target firms shifts sign ($\beta_3 = 0.155^{***}$), being positive in contrast to the negative treatment effect estimated for $CPatent_{it}$ and $Patent_{it}$. Both treatment estimates reflect the full sample period 1980-2018 and are significant at the 1 percent level.

Table 4: Difference-in-Differences Results Across Time Periods

Dependent variable	**Seru (1980–1998)**	**Modern (1999–2018)**	**Full (1980–2018)**
Target patents [log(1+Patents)]			
After (β_1)	-0.185^{***}	-0.202^{***}	-0.196^{***}
	(0.003)	(0.002)	(0.002)
After$\times T$ (β_3)	-0.583^{***}	-0.577^{***}	-0.578^{***}
	(0.006)	(0.006)	(0.004)
Target citations [log(1+Citations)]			
After (β_1)	-0.216^{***}	-0.223^{***}	-0.220^{***}
	(0.001)	(0.001)	(0.001)
After$\times T$ (β_3)	-0.722^{***}	-0.716^{***}	-0.718^{***}
	(0.003)	(0.002)	(0.002)
Acquirer patents [log(1+Patents)]			
After (β_1)	-0.073^{*}	-0.242^{***}	-0.227^{***}
	(0.040)	(0.045)	(0.032)
After$\times T$ (β_3)	0.037	0.128^{***}	0.082^{**}
	(0.042)	(0.047)	(0.034)
Acquirer citations [log(1+Citations)]			
After (β_1)	-0.225^{***}	-0.874^{***}	-0.668^{***}
	(0.082)	(0.070)	(0.055)
After$\times T$ (β_3)	0.022	0.167^{**}	0.106^{*}
	(0.087)	(0.076)	(0.059)
Observations	29 975	28 501	58 476
R^2 (Target patents)	0.671	0.678	0.673
R^2 (Target citations)	0.713	0.712	0.711
R^2 (Acquirer patents)	0.005	0.019	0.012
R^2 (Acquirer citations)	0.009	0.071	0.052
Firm fixed effects	Yes	Yes	Yes
Year fixed effects	Yes	Yes	Yes

Notes: Robust standard errors in parentheses. $^{*}\,p < 0.10$, $^{**}\,p < 0.05$, $^{***}\,p < 0.01$. All patent and citation variables are transformed using $\log(1+x)$ to reduce skewness.

Table 5: Difference-in-Differences estimates for market-based patent value ξ_{real}

	Target	**Acquirer**
After (β_1)	-0.409^{***}	-0.229^{***}
	(0.031)	(0.026)
After$\times T$ (β_3)	0.155^{***}	0.220^{***}
	(0.033)	(0.029)
Observations	58 476	58 476

Note: Dependent variable is $\log(1 + \xi_{\text{real}})$. Standard errors (clustered by firm level) in parentheses. $^{***}p < 0.01$.

5.3 Robustness and Mechanism Tests

To rule out spurious effects driven by heterogeneous trajectories, we augment the baseline model with a firm-specific linear trends.

DiD with Linear Time Trends. Augmented specification with firm-specific linear trends:

$$Y_{it} = \alpha + \beta_1 \,\text{After}_{it} + \beta_3 \big(\text{After}_{it} \times T_i\big) + \varphi_i \,\text{Trend}_{it} + \boldsymbol{\delta}' \mathbf{Z}_{it} + \mu_i + \gamma_t + \varepsilon_{it} \quad (8)$$

where $\text{Trend}_{it} = t - t_0$ and φ_i varies freely across firms.

Table 6 shows that our main results remain robust to this more demanding specification. For target firms, we continue to find significant effects on both patent counts ($\beta_3 =$ -0.259, p $<$ 0.01) and citations ($\beta_3 =$ -0.313, p $<$ 0.01). The acquirer results are particularly noteworthy, showing positive and significant treatment effects for both patents ($\beta_3 = 0.173$, p $<$ 0.01) and citations ($\beta_3 = 0.324$, p $<$ 0.01), suggesting our findings are not driven by pre-existing firm-specific trajectories.

Table 6: DiD with Linear Time Trends

Variable	After	After×T	Observations
Patent	0.071***	-0.259***	58,476
	(0.004)	(0.006)	
CPatent	0.095***	-0.313***	58,476
	(0.005)	(0.007)	
Patent_acq	-0.177***	0.173***	58,476
	(0.028)	(0.032)	
CPatent_acq	-0.367***	0.324***	58,476
	(0.045)	(0.048)	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

5.3.1 Quality-Weighted Innovation Outcomes

To examine whether merger completion affects innovation quality beyond aggregate measures, we analyze two complementary dimensions that capture distinct aspects of innovative output value. This investigation addresses the possibility that mergers may alter the *distribution* of innovation quality rather than simply reducing overall activity, a distinction critical for understanding the efficiency implications of consolidation.

Our quality-weighted analysis employs two theoretically motivated measures:

1. **Citations per Patent** ($\text{CitesPerPat}_{it} = \text{CPatent}_{it}/\text{Patent}_{it}$): This measure captures the average impact of a firm's innovation portfolio, providing insight into whether completed mergers affect the intensive margin of innovation quality. A decline in this measure would indicate that mergers reduce the per-unit impact of remaining patents, suggesting efficiency losses in R&D resource allocation.

2. **High-Impact Patent Indicator** (HighImpact_{it}): A binary measure equal to one when a firm's annual citation count exceeds the sample median, capturing whether firms maintain breakthrough innovation capability. This measure isolates the extensive margin of quality, testing whether mergers specifically erode firms' ability to produce highly influential innovations that drive technological progress.

Table 7 presents our difference-in-differences estimates for these quality measures. The results reveal economically meaningful heterogeneity in how merger completion affects different dimensions of innovation quality. While we find no statistically significant effect on average citations per patent ($\beta_3 = -0.000$, $p > 0.10$), completed mergers generate a substantial 3.7 percentage point reduction in the probability of producing high-impact innovations ($\beta_3 = -0.037^{***}$, $p < 0.01$).

This pattern suggests that merger completion, particularly in the right tail of the innovation quality distribution—the breakthrough innovations that drive long-term competitive advantage and technological progress —has a significant impact. The preservation of average citation quality alongside the erosion of high-impact innovation capability indicates that mergers may particularly disrupt the organizational conditions and research networks necessary for generating transformative innovations, while leaving routine innovative activities relatively unaffected.

These findings have important implications for antitrust policy and corporate strategy. The selective impact on breakthrough innovations suggests that the welfare costs of merger-induced innovation decline may be concentrated in the most technologically significant research areas, potentially slowing industry-wide technological progress even when average innovation metrics show modest effects.

Table 7: DiD with Quality-Weighted Patent Outcomes

Dependent Variable	After	After×T	Observations
Citations per Patent	0.000	-0.000	58,476
	(0.000)	(0.000)	
High-Impact Patents	-0.009***	-0.037***	58,476
	(0.002)	(0.003)	

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 8: Heterogeneity by Technological Overlap

Dependent Variable	After	After×T	After×T×TechOverlap
Patents	-0.075***	-0.253***	-0.003
	(0.002)	(0.006)	(0.007)
Citations	-0.087***	-0.307***	-0.005
	(0.002)	(0.007)	(0.007)

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Technological Overlap. We measure the similarity of target and acquirer patent portfolios using:

$$\text{TechOverlap}_i = \frac{\sum_c s_{ic}^{\text{tar}}\, s_{ic}^{\text{acq}}}{\sqrt{\sum_c (s_{ic}^{\text{tar}})^2}\,\sqrt{\sum_c (s_{ic}^{\text{acq}})^2}},$$

where s_{ic} is firm i's share of patents in CPC field c over the five pre-announcement years. We then estimate triple-difference models interacting this measure with the treatment effect.

To explore potential mechanisms, we investigate whether technological overlap between acquirers and targets moderates the effect on target outcomes. Table 8 employs a triple-difference specification incorporating our TechOverlap measure, which captures similarity in pre-merger patent portfolios. Surprisingly, we find no significant moderating effect of technological overlap (After $\times$ T $\times$ TechOverlap coefficients are statistically insignificant), suggesting that the innovation impacts of mergers operate through channels beyond pure technological synergies. While Bena and Li (2014) show that technological overlap improves innovation outcomes for merger pairs, we do not find that it mitigates the detrimental effect on target outcomes.

5.4 Time-series heterogeneity

Stability for targets. Returning to Table 4, the incremental post-merger drop in citations are virtually unchanged (–0.722 versus –0.716) between the Seru and modern eras, and patent counts tell the same story.

Changing acquirer response. Acquirer patents show no incremental effect pre-1999 but a significant 0.128^{***} boost thereafter, hinting at better technology integration in contemporary M&A.

The model R^2 values remain high (0.67-0.71), indicating that the fixed effects of firm and year plus controls capture most variation within the firm.

5.5 Economic magnitude and interpretation

To assess the economic significance of our findings, we translate the log-coefficients into percentage changes relative to the control group. The treatment effect $\beta_3 = -0.578$ for target patents implies that completion reduces patent output by approximately 44 percent relative to withdrawn deals $[e^{-0.578} - 1 = -0.439]$. For targets citations, the treatment effect $\beta_3 = -0.716$ corresponds to a 51 percent reduction relative to control firms.

For acquirers, positive treatment effects present the $\beta_3 = 0.082$ coefficient for patents in the full sample, indicating that completion increases acquirer patent output by 8.5 percent relative to withdrawn deals $[e^{0.082} - 1 = 0.085]$. This effect strengthens considerably in the modern period ($\beta_3 = 0.128$), where completed mergers boost acquirer patents by 13.6 percent relative to withdrawals.

These differential effects highlight the redistributive nature of merger-induced innovation changes: the innovation losses experienced by targets are partially offset by gains at acquiring firms. However, assessing the net effect across the combined entity requires careful consideration of firm size differences. Since acquirers are systematically larger than targets (Table 3a shows acquirers average 6.1 versus 5.7 in log assets), the 8.5 percent acquirer gain may represent a larger absolute innovation increase than the 44 percent target decline in absolute terms. A precise assessment of the aggregate innovation effect would require weighting these percentage changes by firms' respective pre-merger innovation levels. The asymmetric treatment effects also suggest that organizational integration capabilities have improved over time, allowing modern acquirers to preserve better and leverage acquired innovation assets more effectively.

5.6 Discussion and implications

Our findings in Table 4 provide strong support for our first two hypotheses. The DiD estimates demonstrate that successful mergers have a significant effect on the R&D outcomes of target firms, with treatment effects that are both economically and statistically significant. The negative coefficients (–0.578 for $Patent_{it}$, –0.718 for $CPatent_{it}$) reaffirm the conclusion drawn by Seru (2014), that completed mergers systematically depress both the quantity and quality of target-firm innovation. The persistence of these effects into the 21^{st} century confirms our second hypothesis regarding temporal stability. This challenges the view that modern corporate governance has mitigated the inefficiencies associated with conglomerates and internal capital markets. While our primary focus is on the temporal persistence rather than the mechanisms underlying conglomerate inefficiency, our empirical evidence aligns with the view that organizational complexity impedes innovation productivity.

For acquirers, we document a distinct pattern. The *positive* interaction term for modern-era acquirers (0.128 for patents, 0.167 for citations) indicates improved innovation outcomes relative to control firms (withdrawn acquirers). This suggests enhanced post-merger integration capabilities, consistent with the proposition that contemporary acquirers more effectively absorb target technologies. These findings align with prior work on acquirer R&D outcomes (Zhao, 2009; Sevilir and Tian, 2012). In particular, we obtain these results despite controlling only for target-firm characteristics. The absence of a significant treatment effect for acquirers during the Seru (2014) period (1980–1998) ensures our findings do not contradict his. Collectively, the evidence points toward a nuanced interpretation: conglomerate headquarters continue to *reallocate* resources away from standalone target R&D programs, but acquiring entities now partially productively redeploy these assets.

Table 5 reveals a notable divergence: merger completion positively impacts KPSS-based patent values for both targets ($+0.155$, $p < 0.01$) and acquirers ($+0.220$, $p < 0.01$), even as it depresses citation-based metrics for targets. This contrast suggests that private, market-based valuations of patent portfolios can diverge from citation counts. Consistent with Kogan et al. (2017), we interpret this as evidence that firms often patent for strategic reasons that generate economic rents without necessarily stimulating broader knowledge production. Patents that establish entry barriers or protect market position may command substantial private value while contributing minimally to scientific advancement.

This divergence may reflect the short-term focus of stock market investors driving incremental research that generates more certain near-term value (Gao et al., 2017). Combined with insights from Seru (2014) on information asymmetry and agency problems within conglomerates, this pattern becomes more intelligible. Agency conflicts within merged entities may amplify the existing short-term emphasis of public firms, incentivizing managers to pursue safer, incremental innovations with higher market valuations but lower citation impact. This interpretation aligns with the observed post-merger patterns across our different innovation metrics. Whether the net social welfare effect is positive or negative hinges on the externality between the reduction in target-level innovation quantity and system-wide knowledge production efficiency. Particularly concerning is the significant reduction in high-impact innovation capability, indicating that mergers specifically erode firms' ability to produce breakthrough innovations that exceed median citation performance.

6 Conclusion

With the prevalence of M&A near all-time highs, this paper revisits findings on mergers being detrimental to innovation outcomes. Innovation output measured by proxies for R&D output in the form of number of patents, measuring R&D productivity, and number of citations per patent to measure R&D quality. These measures have been well-established in the literature over the past 20 years. Prior research has covered transactions announced between 1980 and 1998. With a larger and more recent dataset, combining patent, financial, and merger data, we investigate the innovation outcomes related to almost four decades of mergers (1980-2018). This paper employs a DiD methodology involving a quasi-experiment where the treatment group comprises acquired (target) firms and their respective acquirers, while the control group consists of parties to failed mergers.

Our findings align with prior research, showing a significant adverse treatment effect on the innovation outcomes of acquired target firms. These findings remain robust throughout the entire observation period, as well as when considering only more recent years. Moreover, we demonstrate that our results exhibit remarkable temporal stability across different institutional eras, with treatment effects remaining consistent between the earlier Seru period (1980-1998) and the modern period (1999-2018), indicating that our findings are not artifacts of specific time periods or changing market conditions. Furthermore, our paper goes beyond studying only the effect on the targeted firms. We simultaneously present significant results on a positive treatment effect on both patent counts and citations for acquirers, driven by observations in the more recent half of our full observation period.

Another additional contribution is our employment of the private value of patents as a measure of innovation output. Our results using this measure are significant for both targets and acquirers; however, the estimated treatment effect for targets switches sign, turning positive. We reconcile this finding by examining parts of the literature that argue for the prevalence of short-term focus in stock markets, along with agency problems within conglomerates. It can be argued that these would drive the public to favor merger pairs that focus on short-term, exploitative innovation and seek strategic rents, which stock markets may reward. The key to understanding this argument is that the measure of private value is primarily constructed based on stock market reactions.

One benefit of the approach, which utilizes market-informed patent values, is that a market anticipation component is "built in". In turn, this makes the measure less susceptible to truncation bias, which arises from grant lags and lags in citations received by patents in the future. Patent count and citation metrics are also adjusted to better account for this bias.

Our findings contribute to the research on the less established application of the private value of patents as a measure to compare R&D output. With this, researchers will be able to leverage the benefits of such a measure while ensuring it is treated correctly for comparability. Based on our findings using more established patent metrics, it appears that there are still issues related to mergers and innovation conductivity that warrant further investigation, particularly as markets and firms adapt. Our findings, however, are more nuanced than those of prior papers, which often focus almost singularly on the effect on either target firms or acquirers.

All findings taken together, we believe that understanding the underlying mechanisms in action can have a material impact in the end, ensuring future endogenous economic growth.

References

Viral V. Acharya and Zhaoxia Xu. Financial dependence and innovation: The case of public versus private firms. *Journal of Financial Economics*, 124: 223–243, 2017. doi: 10.1016/j.jfineco.2017.02.006.

Jan Bena and Kai Li. Corporate innovations and mergers and acquisitions. *Journal of Finance*, 69(5):1923–1960, 2014. doi: 10.1111/jofi.12051.

Marianne Bertrand and Sendhil Mullainathan. Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American Economic Review*, 94(4):991–1013, 2004.

FactSet. Number of merger and acquisition transactions in the united states from 2022 to 2024, by deal value. Statista dataset, 2025. URL https://www.statista.com/statistics/245977/number-of-munda-deals-in-the-united-states/. Accessed 5 May 2025.

Huasheng Gao, Po-Han Hsu, and Kai Li. Innovation strategy of private firms. *Journal of Financial and Quantitative Analysis*, 53(1):1–32, 2017. doi: 10.1017/S0022109017000898.

Bronwyn H. Hall, Adam B. Jaffe, and Manuel Trajtenberg. The nber patent citations data file: Lessons, insights and methodological tools. Working Paper 8498, National Bureau of Economic Research, 2001.

Bronwyn H. Hall, Adam B. Jaffe, and Manuel Trajtenberg. Market value and patent citations. *RAND Journal of Economics*, 36(1):16–38, 2005. doi: 10.2307/1594028.

Jie He and Xuan Tian. Finance and corporate innovation: A survey. *Asia-Pacific Journal of Financial Studies*, 47(2):165–212, 2018. doi: 10.1111/ajfs.12225.

Institute for Mergers, Acquisitions, and Alliances (IMAA). United states – m&a statistics. IMAA online statistics portal, 2025. URL https://imaa-institute.org/mergers-and-acquisitions-statistics/united-states-ma-statistics/. Accessed 6 May 2025.

Leonid Kogan, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman. Technological innovation, resource allocation, and growth. *Quarterly Journal of Economics*, 132(2):665–712, 2017. doi: 10.1093/qje/qjx002.

Josh Lerner and Amit Seru. The use and misuse of patent data: Issues for finance and beyond. *Review of Financial Studies*, 35(5):2667–2704, 2022. doi: 10.1093/rfs/hhab089.

Gordon M. Phillips and Alexei Zhdanov. R&d and the incentives from merger and acquisition activity. *Review of Financial Studies*, 26(1):34–78, 2013. doi: 10.1093/rfs/hhs118.

Michael E. Porter. Capital disadvantage: America's failing capital investment system. *Harvard Business Review*, 70(5):65–82, 1992.

Paul M. Romer. Increasing returns and long-run growth. *Journal of Political Economy*, 94(5):1002–1037, 1986. doi: 10.1086/261420.

Pavel Savor and Qi Lu. Do stock mergers create value for acquirers? *Journal of Finance*, 64(3):1061–1097, 2009. doi: 10.1111/j.1540-6261.2009.01458.x.

Amit Seru. Firm boundaries matter: Evidence from conglomerates and r&d activity. *Journal of Financial Economics*, 111(2):381–405, 2014. doi: 10.1016/j.jfineco.2013.01.014.

Merih Sevilir and Xuan Tian. Acquiring innovation. SSRN Working Paper, 2012. URL http://dx.doi.org/10.2139/ssrn.1731722.

Xiang Zhao. Technological innovation and acquisitions. *Management Science*, 55(7):1170–1183, 2009. doi: 10.1287/mnsc.1080.1046.

Appendix

Appendix A: Variable definitions

For all variables described in this section, in later sections, any variable VarX presented in this section may appear as VarX_{it} or VarX_{i}. Lower case it denotes a firm-year observation, such as sales for firm i in year t, whereas only i denotes a firm-wise average within the observation window. The following variables are used in subsequent parts of this paper, and are all constructed in line with Seru (2014):

Assets: Total asset (in $ million) for firm i for year t or average of observation window. (Data: CRSP/Compustat Merged (column: `ta`))

Size: Log of Assets (in $ million) of firm i for year t or average of observation window. (Data: CRSP/Compustat Merged (column: `ta`))

Leverage: Total debt of firm i divided by its total assets, in year t or average of observation window. (Data: CRSP/Compustat Merged (column: `ta`, `dlc+dltt`))

EBITDA: Earnings before interest, taxes, depreciation, and amortization (in $ million) of firm i in year t or the observation window average (Data: CRSP/Compustat Merged (column: `ebitda`))

RD: Firm i R&D expenditure in year t or average over observation window (in $ million). (Data: CRSP/Compustat Merged (column: `xrd`))

Sales: Sales in year t or average across observation window (in $ million) by firm i. (Data: CRSP/Compustat Merged (column: `sale`))

HI: Herfindahl Index of firm i in year t constructed based on sales at a four-digit SIC (Data: CRSP/Compustat Merged (columns: `sale`, `sic`))

$\textbf{Dummy}^{\textbf{mseg}}$: This dummy variable takes the value one when the potential acquirer is a diversified firm.

$\textbf{Dummy}^{\textbf{unrelated}}$: A dummy variable that takes a value of one if the potential acquirer is in a different two-digit SIC as the target (Data: CRSP/Compustat Merged (column: `sic`))