MASTER OF SCIENCE IN

FINANCE

MASTER'S FINAL WORK

DISSERTATION

DO KILLER ACQUISITIONS EXTEND TO TECHNOLOGY? FIRM-LEVEL EVIDENCE FROM 25 YEARS OF M&A

LORIS IPPOLITI

JULY - 2026

MASTER OF SCIENCE IN
FINANCE

MASTER'S FINAL WORK
DISSERTATION

DO KILLER ACQUISITIONS EXTEND TO TECHNOLOGY? FIRM-LEVEL EVIDENCE FROM 25 YEARS OF M&A

LORIS IPPOLITI

SUPERVISION:
VICTOR BARROS

JULY - 2026

Glossary

AI – Artificial Intelligence

DiD – Difference-in-Differences

EBITDA – Earnings Before Interest, Taxes, Depreciation and Amortization

FDA – Food and Drug Administration

GAFAM – Google, Apple, Facebook (now Meta), Amazon, Microsoft

HHI – Herfindahl-Hirschman Index

ICT – Information and Communication Technology

ISIN – International Securities Identification Number

KA – Killer Acquisition

M&A – Mergers and Acquisitions

PE – Private Equity

R&D – Research & Development

ROA – Return on Assets

ROE – Return on Equity

ROIC – Return on Invested Capital

VC – Venture Capital

ABSTRACT

This thesis examines whether the "Killer Acquisition" (KA) phenomenon, documented by Cunningham, Ederer and Ma (2021) in pharmaceuticals, extends to technology, where incumbents may acquire nascent competitors to suppress competing innovation rather than realise synergies. Using a panel of 44,189 M&A transactions (2001–2025) enriched with Orbis data on 1,044,384 firm-year observations, the study applies a difference-in-differences design with deal and year fixed effects, decomposing the post-acquisition effect of sub-sector overlap into a contemporaneous (T=0) and two lagged (T+1, T+2) horizons.

Same-sub-sector targets show a significant revenue premium before and at announcement (T−2: +11.9%; T−1: +4.9%; T=0: +8.8%), suggesting that acquirers are selecting high-trajectory threats. Nonetheless, this premium reverses into a significant 6.0 per cent revenue underperformance one year later (T+1), with no corresponding improvement in EBITDA margin, ROA or ROIC. A joint test rejects the parallel-trends assumption, precluding a strict causal claim; results are therefore read as a structural pattern of post-acquisition performance decay congruent with the KA hypothesis, not definitive proof of anti-competitive intent. The effect intensifies where deal frequency is higher, while a revenue-based concentration index (HHI) shows no comparable power.

These results suggest that value-based merger-control thresholds may be ill-suited to detect sub-threshold technology acquisitions, warranting further antitrust scrutiny.

KEYWORDS: Killer Acquisition; Mergers and Acquisitions; Technology Sector; Innovation suppression.

JEL CODES: G34; L41; L86; O31; O33; C23.

Resumo

Esta dissertação examina se o fenómeno de “Killer Acquisition” (KA), documentado por Cunningham, Ederer and Ma (2021) no setor farmacêutico, se estende à tecnologia, onde incumbentes poderão adquirir concorrentes nascentes para suprimir inovação concorrente em vez de realizar sinergias. Com um painel de 44.189 transações de M&A (2001-2025) e dados da Orbis sobre 1.044.384 observações firma-ano, aplica-se um modelo de diferenças-em-diferenças com efeitos fixos de transação e de ano, decompondo o efeito pós-aquisição da sobreposição de subsetor num horizonte contemporâneo (T=0) e em dois horizontes desfasados (T+1, T+2).

Os alvos do mesmo subsetor apresentam um prémio de receita significativo antes e no anúncio (T−2: +11,9%; T−1: +4,9%; T=0: +8,8%), consistente com a seleção, por parte dos adquirentes, de ameaças com elevada trajetória de crescimento. Não obstante, este prémio inverte-se numa quebra de 6,0% da receita um ano depois (T+1), sem melhoria na margem EBITDA, no ROA ou no ROIC. Um teste conjunto rejeita o pressuposto das tendências paralelas, impedindo uma afirmação causal estrita; os resultados são lidos como um padrão estrutural de degradação pós-aquisição consistente com a hipótese KA, não como prova definitiva de intenção anticoncorrencial. O efeito intensifica-se onde a frequência de transações é maior; o HHI baseado na receita não revela poder comparável.

Estes resultados sugerem que os limiares de controlo de concentrações, calibrados pelo valor da transação, poderão estar mal adaptados para detetar aquisições tecnológicas abaixo desses limiares, justificando maior escrutínio antitrust.

PALAVRAS-CHAVE: Killer Acquisition; Fusões e Aquisições; Setor Tecnológico; Supressão de Inovação.

JEL CODES: G34; L41; L86; O31; O33; C23.

TABLE OF CONTENTS

LIST OF TABLES AND FIGURES

Acknowledgments

I wish to express my deepest gratitude to my supervisor, Professor Victor Barros, for his invaluable guidance, rigorous feedback, and unwavering patience throughout the development of this thesis. I am profoundly thankful for the time and care he dedicated to this project.

To my family: I am who I am because of you. You guided me, supported me, and shaped every part of who I have become, and a simple thank you will never be enough to repay that.

To Federica: you changed my life for the better. Thank you for showing me how colourful life can truly be.

To my friends: I am lucky to be surrounded by such beautiful souls, thank you.

AI DISCLAIMER

This master's thesis was developed with strict adherence to the academic integrity policies and guidelines set forth by ISEG – Lisbon School of Economics and Management.

The work presented is the result of my own research, analysis, and writing, unless otherwise cited. In the interest of transparency, I would like to disclose that I utilised AI tools throughout the development of this thesis. These tools supported my literature review by helping me organise and synthesise prior research, assisted in the development of the Stata code used for the empirical analysis, and contributed to the construction of the data visualisations presented in the Appendices. I also employed generative AI to improve the clarity, coherence, and grammatical correctness of the written text.

Nonetheless, I have ensured that the use of AI tools did not compromise the originality or integrity of my work. All analytical decisions, interpretations, and conclusions are my own, and all sources of information are appropriately cited in accordance with academic standards. The ethical use of AI in research and writing has been a guiding principle throughout the preparation of this thesis.

I understand the importance of maintaining academic integrity and take full responsibility for the content and originality of this work.

1. INTRODUCTION

The global technology sector has emerged as the most acquisitive industry of the past two and a half decades, recording 62,272 completed M&A transactions between January 2001 and December 2025, according to Bloomberg. While conventional corporate finance theory posits that acquisitions are primarily motivated by synergistic value creation (Renneboog and Vansteenkiste 2019), the "Killer Acquisition" (KA) framework, formally introduced by Cunningham, Ederer and Ma (2021), challenges this paradigm. It suggests that incumbents may acquire nascent competitors with the deliberate intention of discontinuing their innovation pipelines to pre-emptively protect existing monopoly rents. In this context, the transaction functions as a mechanism of exclusion rather than growth. Amazon's 2010 acquisition of Quidsi (Diapers.com), preceded by a documented price war intended to force the sale and followed by the platform's discontinuation in 2017 (Nadler), illustrates how incumbents may use acquisitions to pre-emptively eliminate market challengers (Motta and Peitz 2021).

This phenomenon has assumed critical policy relevance as regulatory bodies, including the European Commission and the U.S. Federal Trade Commission, question whether traditional merger control is adequate to detect innovation suppression (Crémer, De Montjoye and Schweitzer 2019). The challenge is amplified in digital markets characterised by network externalities and "winner-take-all" dynamics, where strategic acquisitions of potential entrants have become a defining competitive feature.

Despite the seminal contribution in the pharmaceutical sector (Cunningham, Ederer and Ma 2021), transposing the KA framework to technology presents difficulties due to fluid market boundaries and less observable development horizons. The theoretical conditions under which KAs emerge are nonetheless particularly plausible in post-2000 digital platform markets. These conditions are that the incumbent holds appropriable rents protected by entry barriers, the target's innovation is a demand-side substitute for those rents, and acquisition costs less than internal defensive R&D. In such markets, entry-for-acquisition has become the dominant startup exit route. This study addresses a central gap in the empirical literature, namely the absence of firm-level financial proxies for innovation suppression in technology. Section 3 details the research design developed to test this gap, spanning a 25-year horizon (2001–2025), a sample extending beyond GAFAM and above-threshold deals to the sub-threshold universe where KA incentives could be strongest, and an overlap-based identification strategy parallel to the pharmaceutical benchmark. Utilising a panel dataset of 44,189 transactions enriched with Orbis financial data, the study employs a difference-in-differences framework and an event study specification to identify structural patterns of post-acquisition revenue underperformance consistent with the Killer Acquisition hypothesis.

The empirical analysis reveals a structural inflection in target performance, whereby same-sub-sector acquisitions exhibit a revenue premium of approximately 12 per cent two years before announcement. This is compatible with acquirers selecting high-trajectory targets, followed by a reversal into a significant 6.0 per cent revenue underperformance one year post-acquisition, a pattern accompanied by flat EBITDA margins and ROA that discounts conventional efficiency-driven synergy explanations. This inflection constitutes a trajectory of performance decay consistent with the Killer Acquisition hypothesis, though the rejection of parallel trends precludes a strictly causal interpretation. The thesis's principal contribution is to extend this evidence to the technology sector using firm-level financial proxies over a 25-year, sub-threshold-inclusive sample, supplying the empirical infrastructure absent from the literature to date.

The study is structured as follows: Section 2 reviews the relevant literature; Section 3 details the data and econometric strategy; Section 4 discusses the empirical results; and Section 5 concludes with policy implications.

2. LITERATURE REVIEW

2.1 Theoretical foundations

The Killer Acquisition hypothesis of Cunningham, Ederer and Ma (2021) rests on two complementary theoretical primitives. Arrow (1972) demonstrated that a monopolist faces systematically weaker incentives to invest in process innovation than a competitive entrant, a phenomenon he named the replacement effect, because its gain is merely the increment above pre-existing monopoly rent. Schumpeter (2013) argued the inverse, namely that monopoly rents are the necessary reward for innovation, making incumbents the natural engines of technological progress. The Arrow-Schumpeter contrast is the logical seed of the Killer Acquisition literature, since incumbents that prefer the status quo can rationally adopt acquisition followed by discontinuation of a target's development as a rent-protecting strategy.

Gilbert and Newbery (1982) formalised the complementary efficiency effect: an incumbent monopolist will willingly pay more for an innovative entrant than any rival acquirer, because it internalises the full loss of monopoly rents that successful entry would impose. This pre-emption result proves robust to asymmetric information environments (Gans and Stern 2000) and innovation settings (Reinganum 1983), establishing its generality beyond the model.

Three conditions emerge from this literature: the incumbent must hold appropriable rents protected by entry barriers; the target's innovation must be a demand-side substitute for those rents; and acquisition must cost less than internal defensive R&D.

2.2 The pharmaceutical benchmark

The empirical contribution of Cunningham, Ederer and Ma (2021) is fundamental to this study. The authors construct a model in which an incumbent decides whether to acquire a target and whether to continue its development, generating three testable predictions: KAs are more probable when the target overlaps directly with the acquirer's portfolio; their likelihood increases in the bidder's market power; and they cluster below antitrust notification thresholds to evade regulatory scrutiny. Empirically, the study exploits pharmaceutical data to construct a dependent variable, post-acquisition trial discontinuation. They estimate that 5.3 to 7.4 per cent of deals are compatible with a Killer Acquisition rationale.

The technology sector satisfies all three enabling conditions identified above. What it does not replicate is the measurement environment that enables the identification strategy of Cunningham, Ederer and Ma (2021). Digital products iterate continuously rather than following observable trial stages; innovation units are diffuse and embedded in multi-product platforms rather than delineated by patent filings. Moreover, there exist no disclosure requirements equivalent to FDA trial registrations that would render post-acquisition development decisions publicly observable. Methodological transplant is therefore impossible; the hypothesis must instead be operationalised through a dependent variable whose post-acquisition trajectory is measurable in digital markets.

2.3 Extension to digital markets

Katz (2021) provides the first argument that Killer Acquisition incentives are stronger in technology than in pharmaceuticals, demonstrating that big tech firms compete for the market rather than in the market: the acquisition of an emerging competitor eliminates not merely a current product substitute, but an entire future market trajectory. Motta and Peitz (2021) formalise this distinction. They demonstrate that mergers between an incumbent platform and a potential competitor are categorically more harmful to competition than mergers involving complementary targets, even when the target is pre-revenue and its threat is probabilistic rather than realised. Separately, Kamepalli, Rajan and Zingales (2020) document that the mere expectation of such acquisition suppresses startup investment *ex ante*, as switching costs and network externalities reduce the stand-alone commercial value of entrant products in the eyes of consumers who anticipate eventual platform integration, discouraging venture capital-backed entry even absent any explicit anticompetitive conduct by the incumbent.

Letina, Schmutzler and Seibel (2024) add a critical policy qualification: blanket prohibition on incumbent-startup acquisitions deters startup entry almost as severely as it removes the killing incentive, implying that welfare-superior regulation requires micro-level evidence on acquisition motives. This institutional demand is articulated by Crémer, De Montjoye and Schweitzer (2019), whose report identified the absence of longitudinal outcome data on acquired firms as the primary gap preventing effective digital merger enforcement, and reinforced by Federico, Morton and Shapiro (2020), who argue that standard antitrust frameworks are inadequately equipped to evaluate acquisitions in innovation-intensive markets without firm-level longitudinal evidence.

Cabral (2024) quantifies the welfare stakes, finding that a balance-of-harms merger standard would generate approximately 15 per cent welfare gains, but the underlying models are parameterised on assumptions about target-outcome depression that are evidenced for pharmaceuticals and weakly supported for technology.

Koski, Kässi and Braesemann (2026), applying a difference-in-differences design to 742 product markets over 2003-2021, statistically confirm post-acquisition declines in both entry rates and Venture Capital (VC) financing. Kim (2022), using data on US tech acquisitions, identifies that M&A increases spin-out formation by displaced target employees, representing a market-level counterforce to competitive elimination.

Collectively, ecosystem-level studies establish the macro-signature of Killer Acquisition behaviours, but leave the micro-mechanism under-identified, since the unit of observation is the market, not the acquired firm.

2.4 Empirical evidence in technology

At a firm level, the broader M&A and innovation literature establishes the empirical prior that motivates the dependent variable construction of the present thesis. Seru (2014) documents that absorbed targets innovate less than comparable independent firms, as measured by patent counts and forward patent citations, and attributes this decline to the loss of project-level autonomy within the acquirer's hierarchical structure rather than to the pre-existing selection of less innovative targets. His finding is directly applicable to the KA context, as absorption into an incumbent conglomerate generates suppression of innovation output even absent a deliberate discontinuation motive. Evidence specific to the largest technology acquirers reinforces this picture while also exposing the limits of a GAFAM-only research design. Gautier and Lamesch (2021) found that a large majority of GAFAM-acquired products are discontinued post-deal, yet applying a competitive-substitution criterion reduces confirmed KAs to only one. Affeldt and Kesler (2021), exploiting GAFAM mobile app acquisitions, show that competing apps reduce update frequency post-deal. Jin, Leccese and Wagman (2023) demonstrate that restricting samples to GAFAM misrepresents the full cross-section of horizontal technology acquisitions, as private equity (PE) now accounts for a comparable acquisition share.

This thesis extends this logic from patent-based proxies to comprehensive firm-level financial performance trajectories, and from a GAFAM-only lens to a sample spanning the full population of technology acquirers. Specifically, given the unobservable nature of R&D reporting in sub-threshold private tech firms, this thesis employs operating revenue trajectories, return on assets (ROA), EBITDA margins, and Return on Invested Capital (ROIC) as financial proxies to evaluate whether post-acquisition environments reflect growth, stagnation, or structural operational suppression.

2.5 Counter evidence and the limits of existing research

Ivaldi, Petit and Unekbas (2023) examine *ex post* outcomes for all European Commission-reviewed ICT merger cases and find no systematic evidence of product discontinuation, competitor weakening or entry suppression. The critical caveat is structural: their sample is filtered by notification thresholds. Since Cunningham, Ederer and Ma (2021) argue that KAs cluster below such thresholds precisely to evade scrutiny, the null result is biased against detecting the phenomenon by design. Their study establishes only that among the largest, most visible deals, anticompetitive patterns are not observable *ex post*; the inference cannot extend to the sub-threshold universe where the hypothesis predicts effects to be concentrated. Barnett (2023) offers the most comprehensive critical appraisal, arguing that policy enthusiasm for stricter big tech enforcement rests on a narrow evidentiary base, while under-weighting the role of acquisition as a monetisation mechanism that sustains VC investments. His most pointed challenge is an external-validity concern: the Cunningham, Ederer and Ma (2021) result may be a product of pharmaceutical data environments and may not generalise to technology markets, rendering its application to digital merger policy an empirical overreach without independent corroboration. This critique is warranted as a characterisation of the current record, and it is the gap the present thesis addresses.

2.6 Research gap and contribution

The methodological landscape of the literature reviewed above reveals a body of knowledge that is theoretically mature but empirically fragmented at the firm-level in the technology sector. This fragmentation defines a single central gap motivating the present thesis, namely the absence of firm-level financial proxies capable of capturing the post-acquisition trajectory of acquired technology firms. Existing studies rely on binary product survival, competitor innovation response, categorical deal classification, or market-level proxies such as entry rates and VC flows. None directly captures the financial trajectory of the acquired firm itself. Responding to the institutional mandate of Crémer, De Montjoye and Schweitzer (2019) and the methodological agenda of Federico, Morton and Shapiro (2020), this thesis introduces a battery of firm-level financial metrics – operating revenue, EBITDA margins, ROA, and ROIC – which, unlike sporadic and frequently undisclosed R&D expenditures, provide a robust, continuous reflection of the target's operational vitality and commercial exploitation post-acquisition.

Addressing this gap with the rigour it requires entails three complementary design choices. First, a temporal extension: existing studies are constrained to at most one decade, with none spanning the full arc from the post-dot-com era through the current AI transition. Covering 2001–2025 permits an assessment of whether KA incentives intensify as market concentration deepens across successive technology cycles. Second, a broadened sample: existing firm-level evidence skews toward GAFAM and deals above regulatory notification thresholds, simultaneously overstating the distinctiveness of big tech behaviour and excluding the sub-threshold universe where the Killer Acquisition motive is theoretically most concentrated. This thesis expands scope to a cross-acquirer sample irrespective of acquirer identity or review status, compatible with the cross-sectional evidence of Jin, Leccese and Wagman (2023). Third, a comparable identification strategy: an overlap-based measure conceptually parallel to the therapeuticoverlap measure of the pharmaceutical benchmark enables a direct test of the external validity of the Killer Acquisition hypothesis in technology.

3. SAMPLE AND METHODOLOGY

3.1 Data sources

The primary source of transaction-level data is Bloomberg, accessed via its dedicated Mergers and Acquisitions search function, which aggregates deal information from regulatory filings, press releases, and proprietary intelligence feeds. The search universe was restricted to the technology sector according to Bloomberg's proprietary industry classification taxonomy, which assigns each firm to a broad industry sector and a finer industry sub-sector. This two-tier hierarchy is central to the identification strategy, as it enables the construction of competitive overlap measures at two levels of granularity. The temporal scope spans all transactions with an announced date between 1 January 2001 and 31 December 2025, a 25-year horizon chosen to capture multiple macroeconomic cycles and successive platform generations, across which KA incentives are theoretically predicted to intensify (Kamepalli, Rajan and Zingales 2020; Motta and Peitz 2021). The sample was restricted to completed transactions only, to isolate genuine transfers of corporate control. After automated deduplication and removal of records lacking critical classification data, the final transaction-level sample comprises 62,272 completed technology sector acquisitions, for each of which the Bloomberg industry sector and sub-sector are recorded for both bidder and target firms.

Longitudinal target-level financial data are sourced from the Moody's Orbis Bureau van Dijk database. Target data were retrieved via the Orbis batch search function, which accepts company name, country and ISIN as matching keys. Given Orbis's technical constraint of a maximum of 1,000 records per batch session, a dedicated VBA macro was used within Excel to automate the batching and submission process. The batch procedure returned confirmed matches for 44,189 target companies, corresponding to a match rate of 70.96 per cent. This match rate is not missing at random. Orbis's underlying national filing sources provide systematically deeper coverage of large, listed, and EU-domiciled firms than of small, private firms outside Europe's mandatory disclosure regimes (Kalemli-Özcan *et al.* 2024). This pattern plausibly concentrates non-matches among the small, private, non-EU-domiciled targets. To the extent that this holds, the matched sample is biased toward the population in which KA incentives are theoretically weaker. This is because the rents being protected are most threatened by exactly the small, nascent challengers Orbis is least likely to capture. Therefore, the post-acquisition effects estimated in Section 4 should be read as a conservative, lower-bound approximation of the underlying phenomenon rather than an inflated one. For each matched entity, an initial panel dataset covering 1997 to 2025 with 29 financial variables per firm-year was constructed and subsequently exported to Stata for further processing and matched against the Bloomberg transaction file.

3.2 Sample and variable construction

The integration of the Bloomberg transaction file with the Orbis financial panel presents a canonical record linkage challenge, as neither database provides a common unique identifier, and company name representations diverge across the two sources.

Exact string matching would, therefore, produce a high rate of false non-matches and an attrition of the analytical sample. To address this, the two datasets were merged using a probabilistic fuzzy-matching procedure implemented with the Stata command *reclink* (Blasnik 2010), which applies the Jaro-Winkler string similarity metric to compute a pairwise match score for each candidate pair. This metric assigns a similarity score between 0 and 1, where 1 denotes an exact match, and 0 denotes dissimilarity. A match score threshold of 0.85 was adopted as the minimum criterion for a confirmed match, balancing precision against recall, consistent with the entity-resolution literature (Winkler 1990; Blasnik 2010).

Following the match, the dataset was restructured into long format using Stata's *reshape* command, yielding an initial panel of 1,044,384 firm-year observations spanning 1997 to 2025. This panel comprises 36,134 unique deal groups; the gap relative to the 44,189 Orbis-confirmed target matches (Section 3.1) reflects the requirement of a usable financial panel around each deal, as not every confirmed Orbis match yields at least one firm-year observation once the matched data are reshaped into long format. Stata's *xtreg* command performs automatic listwise deletion on any observation with a missing value in the dependent variable or any covariate. Given that Orbis financial variables are sparsely reported for private technology firms (Section 3.1), this reduces the estimation sample to 65,036 firm-year observations across 11,236 deal groups – the final sample on which all static regression results reported in Section 4 are estimated. Distribution analysis confirmed that all key financial-level variables exhibit right-skewed distributions congruent with a log-normal distribution, supporting the log transformation applied to operating revenues and total assets prior to estimation.

Industrial proximity is measured through two independent binary dummy variables constructed at the transaction level from Bloomberg's industry classification data. The sector-level proximity indicator, *overlap_sector*, equals 1 when both acquirer and target are in the same industry sector, and 0 otherwise. This variable captures broad competitive adjacency across the technology universe. The sub-sector-level indicator, *overlap_subsector*, equals 1 when the acquirer and target share both the same industry sector and industry sub-sector, and 0 otherwise. This variable captures direct product-market confrontation and is the primary identification variable for the Killer Acquisition test, serving as the technology-sector analogue of the therapeutic-area overlap criterion employed by Cunningham, Ederer and Ma (2021). The post-acquisition indicator, *post*, equals 1 for all calendar years on or after the year of acquisition announcement for deal i, and 0 for all pre-announcement years. The construction of the difference-in-differences interaction terms built from these dummies, and their application within each regression specification, are presented alongside the corresponding equations in Section 3.3.

The primary dependent variable is the natural logarithm of operating revenues, *log_Op_rev*. Operating revenue is selected for three reasons:

1. It is reported for both publicly listed and privately held technology firms in Orbis;
2. A post-acquisition deceleration in the target's revenue relative to control firms is an observable financial signature of suppressed commercial activity in a setting where trial-stage innovation data are unavailable;
3. Its distribution is confirmed to be log-normal, as evidenced by empirical density and normal quantile-quantile plots.

All financial variables were extracted from Orbis with EUR set as the batch-export reporting currency, so that cross-border currency heterogeneity is resolved at the data-extraction stage rather than through a subsequent conversion step. No deflation to constant prices was applied; reported values are therefore in nominal EUR terms, and the year fixed effects included in every specification are relied upon to absorb common time-varying price-level trends that would otherwise confound the post-acquisition comparison.

Two firm-level control variables are included in all specifications: the natural logarithm of total assets, *log_Tot_assets*, which controls for firm size and scale effects, and financial leverage, *Fin_leverage*, which accounts for differences in capital structure that may affect investment and revenue capacity. The natural logarithm transformation is applied where the underlying value is strictly positive. Observations recording zero or negative operating revenues are excluded via listwise deletion. This restriction is theoretically motivated rather than merely a data convenience, since in the technology sector, operating revenue constitutes the primary observable signal that a target has achieved product-market fit and is actively competing in a commercial market. A pre-revenue firm represents a speculative competitive threat; a revenue-generating firm, by contrast, has demonstrated the market viability that makes acquisition and subsequent suppression rational for an incumbent protecting monopoly rent. This framing is congruent with the technology-sector extension of Motta and Peitz (2021), who emphasise that KA incentives concentrate on targets whose commercial trajectory poses a credible near-term competitive threat, rather than on purely pre-commercial entities. The sample restriction therefore selects the population theoretically most relevant to the test. A Return on Invested Capital proxy is constructed in two steps. First, equity is derived from the DuPont identity, given that ROA and ROE share net income as their numerator:

$$\text{Equity}_{\text{derived}} = \text{Tot\_assets} \times \left(\frac{\text{ROA}}{\text{ROE}}\right).$$

Second, total debt is derived from *Fin_leverage*, which Orbis reports as a gearing ratio:

$$\text{TotalDebt}_{\text{derived}} = \text{Fin\_leverage} \times \text{Equity}_{\text{derived}}.$$

Invested capital is then:

$$\text{IC} = \text{Equity}_{\text{derived}} + \text{TotalDebt}_{\text{derived}} \quad \text{and} \quad \text{ROIC} = \frac{\text{EBIT}}{\text{IC}}, \text{if IC} > 0.$$

Given extreme kurtosis in the raw EBIT series, ROIC is winsorised at the first and ninety-ninth percentiles, yielding ROIC, which serves as the capital efficiency outcome variable in Section 4.5. *Fin_leverage* is accordingly excluded from the control set in Model M4 only, as it is embedded directly in the dependent variable's construction there.

A revenue-based Herfindahl-Hirschman Index (HHI) is constructed at the sub-sector × year level as the sum of squared revenue market shares scaled by 10,000:

$$HHI_{jt} = \sum_{i} s_{ijt}^{2} \times 10{,}000\ ,$$

where s_{ijt} denotes firm i's share of total observable operating revenue in sub-sector j in year t. A structural limitation of this measure is that it is computed over the target panel only: acquirer revenues, whose market share matters most, are absent from the Orbis-matched panel. The HHI therefore proxies target-pool concentration rather than true market concentration, an acknowledged limitation that motivates the alternative deal-frequency proxy described below.

An alternative concentration variable, deal-frequency (*n_ssa_deals*), counts the number of same-sub-sector acquisitions per sub-sector per calendar year. This variable measures incumbent acquisition intensity without requiring acquirer revenue data and is the primary moderation variable alongside the HHI specification. Both HHI and deal-frequency are standardised to mean zero and unit standard deviation prior to inclusion in interaction terms.

3.3 Empirical strategy

The empirical strategy is estimated throughout by fixed-effects panel regression, with heteroscedasticity-robust standard errors clustered at the deal level; deal fixed effects absorb all time-invariant unobserved heterogeneity at the transaction level, and year fixed effects absorb common macroeconomic and price-level shocks across the sample period.

The baseline specification is a single equation, estimated under two alternative definitions of one variable, Overlap:

$$\textbf{(1)}\ \ \mathrm{Rev}_{it} = \alpha + \beta_1 \mathrm{Post}_{it} + \beta_2 (\mathrm{Post} \times \mathrm{Overlap})_{it} + \gamma_1 \mathrm{Size}_{it} + \gamma_2 \mathrm{Leverage}_{it} + \delta_t + \\ + \eta_i + \varepsilon_{it}\ ,$$

where Rev is the log of operating revenue, Size is the log of total assets, and Leverage is financial leverage (Section 3.2). When Overlap is defined at the sector level, equation (1) tests whether targets acquired by firms in the same broad Bloomberg sector exhibit a differential post-acquisition revenue trajectory relative to targets acquired by firms in a different sector. When Overlap is defined at the sub-sector level, the primary identification variable for the KA test, the same equation restricts the treatment definition to targets whose acquirer operates in the identical sub-sector; a statistically significant and negative β_2 under this definition is the primary static criterion for consistency with the KA hypothesis. Both estimates of equation (1) reflect one specification under two definitions of Overlap, not two separate models, and are reported side by side in Table 1; β_2 in equation (1) is reported in the results and tables that follow as *did_sector* under the sector-level definition of Overlap, or *did_subsector* under the sub-sector-level definition, congruent with the table label “DiD × Post (T=0)”.

Every firm in the matched panel has been acquired by some incumbent, so there is no group of non-acquired firms available as a control. The implicit control group is therefore targets acquired by a non-competitor (Overlap = 0); β_2 captures the differential effect of being acquired by a same-sub-sector rival relative to being acquired by any other acquirer, not the effect of acquisition itself. The dynamic extension decomposes this effect across three temporal horizons, using the sub-sector definition of Overlap throughout:

$$\textbf{(2)}\ \mathrm{Rev}_{it} = \alpha + \beta_0(\mathrm{Post} \times \mathrm{Overlap})_{it} + \beta_1(\mathrm{Post} \times \mathrm{Overlap})_{i,t+1} + \beta_2(\mathrm{Post} \times \mathrm{Overlap})_{i,t+2} + \beta_3\mathrm{Post}_{it} + \gamma X_{it} + \delta_t + \eta_i + \varepsilon_{it}\ ,$$

where t+1 and t+2 denote one and two years after the acquisition year, respectively. *did_subsector* denotes the same contemporaneous coefficient as in equation (1), now labelled β_0; the lagged interactions β_1 and β_2 are reported as *did_sub_lag1* (T+1) and *did_sub_lag2* (T+2) respectively, and β_3, the unconditional post-acquisition indicator, is reported as *Post*. Under the KA hypothesis, β_0 may be positive (selection of high-growth targets at acquisition), β_1 should be negative and significant (suppression materialises within twelve months of completion), and β_2 should be negative or zero (the effect persists or decays). The sample restriction underlying this specification is set out in Section 3.4.

A further set of specifications extends this dynamic structure along three dimensions, without introducing additional equations. Market concentration is introduced as a moderator, first as an additive control using the standardised HHI, then as its interaction with the treatment indicator, and, on a separate sample, with the deal-frequency proxy in place of HHI. A second set of robustness checks replaces operating revenue with the log of intangible assets and the log of R&D expenditure. A third set applies the same dynamic structure to EBITDA margin, ROA, and ROIC, to test whether any revenue contraction is accompanied by the margin improvements a synergy explanation would predict.

Sample compositions, restrictions, and diagnostics for each of these extensions are addressed in Section 3.4; coefficients are reported in Section 4 and Appendix A.

3.4 Identification and robustness

The principal identification challenge is the potential endogeneity of the overlap treatment, as acquirers may select same-sub-sector targets based on unobservable characteristics that independently predict post-acquisition revenue trajectories. Deal fixed effects absorb all time-invariant unobserved heterogeneity at the transaction level, with identification derived from within-deal variation in revenue performance across the pre- and post-acquisition window. Year fixed effects absorb aggregate macroeconomic shocks and technology-cycle variation. The validity of the parallel trends assumption is assessed through an event study specification that augments equation (2) with two lead interaction terms (*did_sub_lead2*, *did_sub_lead1*), capturing the revenue differential in the two years preceding the acquisition announcement. Under the null of parallel trends, both lead coefficients should be statistically indistinguishable from zero. Their joint significance would indicate that same-sub-sector and non-overlapping targets were already on divergent revenue trajectories before treatment, precluding a fully causal interpretation of the dynamic DiD estimator. The estimated coefficients and the outcome of this test are reported in Section 4.3, together with the dynamic estimates to which they pertain.

Because a parallel trends violation is, by construction, a question about comparability of the treated and control groups on pre-treatment characteristics, a complementary diagnostic was implemented: entropy balancing. This procedure reweights the non-overlapping control group so that its first-moment distribution (mean) of pre-acquisition revenue growth, pre-acquisition size, and pre-acquisition leverage exactly matches that of the same-sub-sector treated group, without discarding observations through nearest-neighbour matching. Balancing weights were estimated on the cross-section of deal groups with a usable pre-acquisition observation. Deal groups for which no such observation exists cannot be assigned a weight and are excluded from this check, a further restriction of the identification sample distinct from the estimation samples used elsewhere. The event study and dynamic specifications were then re-estimated on the resulting entropy-weighted panel as a test of whether the pre-trend violation is attributable to observable selection on growth, size, or leverage, in which case balancing on these covariates should restore parallel trends, or to selection on unobservable characteristics that survive reweighting. The result is reported in Section 4.3 (Table 5, Appendix A).

Beyond the baseline and dynamic specifications, a battery of additional specifications tests the robustness of the central finding along four dimensions, namely a dose-response comparison across the two overlap definitions, alternative dependent variables (intangible assets, R&D expenditure), market-concentration moderators (HHI and deal-frequency), and operational efficiency outcomes (EBITDA margin, ROA, ROIC). The deal-frequency moderation is accompanied by a variance inflation factor diagnostic on the design matrix to rule out collinearity with the standardised level term. The HHI is computed over the target panel only, excluding acquirer market share, and is reported with this limitation in mind. The EBITDA margin and ROA specifications draw on different estimation samples, reflecting differential reporting coverage in Orbis; the R&D expenditure specification is restricted to the smaller sub-sample of firms reporting R&D data, a structurally different population skewed toward larger, more innovation-intensive firms. All dynamic models additionally impose a sample restriction to announcements on or before 31 December 2022, to ensure the second-lag (T+2) window falls within the Orbis coverage horizon; static specifications retain the full sample.

4. EMPIRICAL RESULTS

4.1 Sample and estimation frame

Tables 1 through 5 (Appendix A) report the full regression output underlying the findings below; the main text presents key coefficients only. The static specifications are estimated on 65,036 observations across 11,236 deal groups; the dynamic specifications, on 55,599 observations across 9,622 deal groups. The first empirical question is whether a post-acquisition revenue differential is detectable at the broad sector level and, if so, whether it concentrates at the sub-sector level, as the KA hypothesis predicts.

4.2 Sector-level vs. sub-sector-level baseline

The sector-level estimate of equation (1) yields a post-acquisition revenue differential for same-sector targets of $\beta = -0.035$ (SE = 0.030, $p = 0.243$). The coefficient is economically small and statistically indistinguishable from zero, as is the post indicator ($\beta = -0.013$, $p = 0.540$). The absence of a sector-level effect is itself informative, as a significant sector-level estimate would be expected if revenue underperformance were attributable to broad industry conditions rather than competitive overlap. Therefore, its absence is in line with the KA prediction that suppression incentives concentrate at the level of direct product-market threat.

The sub-sector-level estimate of the same equation is positive and statistically insignificant ($\beta = +0.013$, SE = 0.033, $p = 0.695$). This static null does not contradict the KA hypothesis, since the static specification pools the acquisition year with all subsequent post-acquisition years within a single post indicator, attenuating the dynamic treatment effect reported in Section 4.3. The log total assets coefficient ($\beta = 0.836$, SE = 0.015, $p < 0.001$) is stable across both estimates, confirming a consistent revenue-to-size relationship.

4.3 The dynamic revenue profile: core evidence

The dynamic specification, decomposing the DiD treatment into contemporaneous (T=0), first-lag (T+1), and second-lag (T+2) components, is estimated on 55,599 observations across 9,622 deal groups (Table 1, Appendix A). Three findings warrant discussion.

First, the contemporaneous treatment coefficient is positive and statistically significant (*did_subsector*: $\beta = +0.084$, SE = 0.035, $p = 0.017$). Same-sub-sector targets exhibit operating revenues approximately 8.8 per cent higher than non-overlapping targets in the year of acquisition announcement. This premium is in alignment with acquirers selecting commercially high-trajectory targets – firms whose revenue growth simultaneously makes them attractive acquisition candidates and credible competitive threats – echoing the selection mechanism theorised by Cunningham, Ederer and Ma (2021) and confirmed in the event study below.

Second, the first-lag coefficient is negative and statistically significant (*did_sub_lag1*: $\beta = -0.062$, SE = 0.031, $p = 0.044$). One year post-acquisition, same-sub-sector targets underperform their non-overlapping counterparts by approximately 6.0 per cent in log operating revenues. This reversal of 0.146 log points relative to the T=0 premium constitutes the primary evidence of post-acquisition performance decay in the treated group, timed consistently with the operational-integration period over which a restriction on the target's independent commercial activity would first materialise in financial statements.

Third, the second-lag coefficient is negative but does not achieve conventional significance (*did_sub_lag2*: $\beta = -0.033$, SE = 0.030, $p = 0.278$). The directional persistence of the negative effect through T+2 is noted, though its statistical attenuation precludes inference. The overall post indicator ($\beta = -0.048$, SE = 0.021, $p = 0.021$) confirms that the post-acquisition period is, on average, characterised by revenue underperformance relative to the pre-acquisition benchmark embedded in the deal-fixed effect.

The event study specification, estimated on 51,410 observations across 9,387 deal groups, extends this profile to a full five-period window, T−2 to T+2 (Figure 1, Appendix B). At T−2, the sub-sector treatment indicator carries a coefficient of +0.112 (SE = 0.027, $p < 0.001$); at T−1, +0.048 (SE = 0.021, $p = 0.022$). A joint F-test of the null hypothesis that both pre-trend coefficients equal zero is rejected: $F(2, 9{,}386) = 14.26$, $p < 0.001$. At the moment of acquisition, the coefficient profile inflects sharply, as the contemporaneous effect drops to statistical indistinguishability from zero ($\beta = -0.027$, $p = 0.424$), before turning negative and marginally significant at T+1 ($\beta = -0.053$, $p = 0.079$) and negative and insignificant at T+2 ($\beta = -0.016$, $p = 0.608$). The positive, monotonically rising pre-trend at T−2 and T−1, followed by an inflection at the acquisition year and a reversal at T+1, is the central pattern of decline this thesis documents, namely a pre-acquisition revenue premium consistent with KA-predicted target selection (Jin, Leccese and Wagman 2023), followed by a post-acquisition reversal in alignment with commercial suppression.

As a further check on this rejection, the event study and dynamic specifications were re-estimated on the entropy-balanced sample described in Section 3.4 (Table 5, Appendix A; $N = 22{,}067$ and $N = 23{,}349$ respectively, reflecting the smaller pool of deal groups with a usable pre-acquisition observation). Balancing achieves exact equality of means on pre-acquisition size, growth, and leverage across treated and control groups (pre-acquisition size: 8.969 for both, against 8.969 versus 9.154 unweighted). Re-weighting on observables does not resolve the rejection, as the joint test of the two pre-trend leads continues to reject parallel trends, $F(2, 2{,}808) = 6.11$, $p = 0.0023$, though the F-statistic falls relative to the unweighted specification. The T−2 lead remains significant ($\beta = 0.074$, $p = 0.010$); the T−1 lead is attenuated to marginal significance ($\beta = 0.033$, $p = 0.054$). The post-acquisition decay is, if anything, amplified rather than attenuated, as the entropy-weighted dynamic estimate yields a T=0 premium of +0.127 log points ($p < 0.001$) and a T+1 reversal of −0.104 log points ($p = 0.001$), both larger in magnitude and more precisely estimated than their unweighted counterparts.

All findings in this and the following sections are characterised accordingly as a trajectory of post-acquisition performance decay compatible with the Killer Acquisition hypothesis, rather than as a causally identified treatment effect.

4.4 Market concentration and deal intensity

Two proxies for market concentration are tested as moderators of the primary revenue effect, motivated by the theoretical prediction that KA incentives intensify with the acquirer's market power (Cunningham, Ederer and Ma 2021; Motta and Peitz 2021).

The revenue-based HHI interaction (did_sub_HHI) carries a coefficient of +0.012 (SE = 0.096, $p = 0.901$) and yields no significant main or moderating effect in any specification (Table 2, Appendix A); in the dynamic moderation, the contemporaneous interaction is −0.037 (SE = 0.107, $p = 0.731$), with the lagged interactions similarly insignificant. Congruent with the construction limitation noted in Section 3.4, these null results are uninformative about market-power-driven amplification of the suppression effect, rather than evidence against it.

The interaction between deal-frequency and the treatment indicator (did_sub_nssa) carries a coefficient of −0.061 (SE = 0.029, $p = 0.037$): targets in sub-sectors with higher same-sub-sector acquisition intensity exhibit a more pronounced post-acquisition revenue contraction, estimated on 56,723 observations across 9,775 deal groups (Table 2b, Appendix A). This result controls for the direct level effect of deal-frequency itself (n_ssa_std: $\beta = 1.175$, $p = 0.194$, not significant) – an implausibly large coefficient on a standardised regressor in a log-revenue equation. The variance inflation factor diagnostic flagged in Section 3.4 returns a mean VIF of 1.15, with no individual VIF exceeding 1.26, ruling out collinearity with the interaction term; the magnitude is therefore reported as an unresolved feature of the level term that does not bear on did_sub_nssa, which is precisely estimated and correctly signed. The deal-frequency finding provides independent structural evidence that the revenue-suppression pattern intensifies in market contexts where the KA hypothesis predicts effects to be strongest, in contrast to the uninformative HHI result.

4.5 Operational efficiency: absence of synergy evidence

The operational efficiency analysis is the second pillar of the empirical strategy. Under a synergy-driven acquisition hypothesis, revenue reduction should be accompanied by margin expansion; under a suppression of commercial performance hypothesis, revenues decline without corresponding margin improvement. Models M2–M4 estimate the dynamic specification on EBITDA margin, ROA, and ROIC, on the $year\_announcement \leq 2022$ sub-sample (Table 4, Appendix A). Models M2 and M3 draw on different, though overlapping, estimation samples (M2: $N = 43{,}650$, 7,573 deal groups; M3: $N = 57{,}330$, 9,715 deal groups), for the data-coverage reasons set out in Section 3.4; the comparison of T=0 premia between the two models below should be read with this in mind.

The EBITDA margin results (Model M2) show that same-sub-sector targets carry a significant profitability premium at the acquisition year (*did_subsector* = +2.01 percentage points, SE = 0.864, p = 0.020), consistent with the revenue premium documented at T=0 and confirming a high-quality selection profile. The T+1 and T+2 coefficients are statistically indistinguishable from zero (T+1: +0.10 pp, SE = 0.828, p = 0.905; T+2: +0.50 pp, SE = 0.739, p = 0.499): EBITDA margins neither improve nor deteriorate significantly following acquisition.

The pattern is replicated for return on assets (Model M3): a significant contemporaneous premium (+2.14 pp, SE = 0.782, p = 0.006), followed by insignificant and directionally negative T+1 (−0.33 pp, SE = 0.731, p = 0.647) and T+2 (−0.94 pp, SE = 0.579, p = 0.104) coefficients.

The ROIC proxy (Model M4) follows the same pattern, with a positive but only marginally significant contemporaneous differential (*did_subsector* = 0.032, SE = 0.017, p = 0.059, equivalent to a 3.2 percentage-point higher return on invested capital at T=0), followed by statistically indistinguishable T+1 (−0.002, SE = 0.016, p = 0.906) and T+2 (−0.014, SE = 0.013, p = 0.305) coefficients.

Taken together, same-sub-sector targets are selected for superior commercial and financial quality, evidenced by the T=0 revenue, margin, and capital-efficiency premia (Models C, M2, M3 and, at marginal significance, M4); their revenues underperform significantly at T+1; their margins and capital efficiency do not improve. This combination – revenue decline without margin improvement – is inconsistent with a synergy-motivated acquisition hypothesis and aligns with commercial suppression as the operative mechanism, characterised as a conditional empirical association rather than causal proof of Killer Acquisition behaviour.

Table 3 (Appendix A) reports results for three alternative dependent variables. Model G (static intangibles, full sample) yields *did_subsector* = −0.057 (SE = 0.071, p = 0.424); the dynamic extension (Model H) yields non-significant coefficients at all horizons. Model I (log R&D expenditure) yields *did_subsector* = −0.032 (SE = 0.081, p = 0.688), estimated on the sub-sample of 1,292 deal groups reporting R&D data, a structurally different population from the main panel as noted in Section 3.4. Neither innovation-asset measure exhibits a statistically significant differential trajectory for same-sub-sector targets. Within the current data environment, operating revenue remains the most reliably recoverable dependent variable for detecting post-acquisition performance suppression at scale.

5. CONCLUSIONS

5.1 Summary of findings

This thesis documents a profile of post-acquisition revenue decay among same-sub-sector technology targets that is consistent with the Killer Acquisition hypothesis, while stopping short of establishing causal proof of anti-competitive intent. The central finding is an inflection in the target revenue trajectory around the acquisition date, whereby same-sub-sector targets carry a significant revenue premium in the two years preceding announcement and at announcement itself, before this premium reverses into a significant 6.0 per cent revenue underperformance relative to non-overlapping targets one year post-acquisition. This reversal is accompanied by flat EBITDA margins and return on assets, inconsistent with the efficiency gains a synergy-driven acquisition rationale would predict. This pattern is broadly in line with comparable evidence from the wider digital-markets literature. Gautier and Lamesch (2021) documented that a substantial share of GAFAM-acquired targets are discontinued or absorbed without independent commercial continuation; Affeldt and Kesler (2021) find reduced post-acquisition innovation activity among apps acquired by dominant platforms. The deal-frequency moderation reported in Section 4.4 – suppression intensifying where acquisition activity is most concentrated – parallels the portfolio-effect concerns raised by Ivaldi, Petit and Unekbas (2023) in their assessment of EU digital merger enforcement. At the same time, Koski, Kässi and Braesemann (2026) find more mixed effects of acquisition on target-level innovation in a European technology sample, and Barnett (2023) cautions against treating structural performance patterns as sufficient grounds for inferring anti-competitive intent, without direct evidence of strategic suppression. In light of this more mixed evidence, the findings of this thesis are characterised as a structural pattern of performance decay in accordance with the Killer Acquisition hypothesis that warrants further antitrust scrutiny, rather than as definitive proof that Killer Acquisitions are occurring in the sampled transactions.

5.2 Policy implications

The findings carry several implications for merger policy in technology markets. Current merger control frameworks in both the European Union and the United States are calibrated primarily to transaction value and market share thresholds, structurally ill-suited to detect sub-threshold acquisitions of early-revenue technology targets, precisely the stratum where KA incentives are theoretically most concentrated (Crémer, De Montjoye and Schweitzer 2019; Cunningham, Ederer and Ma 2021). The revenue decay pattern documented here is detectable using commercially available financial data, suggesting that longitudinal outcome monitoring of completed acquisitions is technically feasible at regulatory scale.

The deal-frequency finding implies that sub-sectors with systematically high levels of same-sub-sector acquisition activity may merit proactive regulatory attention, even absent individual transaction notifications. A portfolio-based approach to merger scrutiny, assessing the cumulative competitive effect of repeated sub-threshold acquisitions within a defined competitive perimeter, would be responsive to this pattern and is consistent with the recommendations of Federico, Morton and Shapiro (2020). The null HHI result, alongside the significant deal-frequency moderation, suggests that traditional concentration indices may be inadequately calibrated for digital markets where dominant incumbents' revenues are not reflected in target financial data. Regulatory concentration assessments for technology mergers may benefit from supplementing HHI analysis with acquisition-intensity metrics that capture behavioural concentration directly.

Taken together, the main takeaway for policy institutions is that the absence of a notification obligation should not be read as the absence of a competitive concern: the segment of the market least visible to current thresholds is, on this evidence, also the segment where the revenue-decay pattern documented in this thesis is most pronounced. This is not evidence that Killer Acquisitions are occurring within this segment, nor does it justify a presumption against sub-threshold acquisitions as a class; the finding is structural and correlational, consistent with the KA hypothesis but not proof of anti-competitive intent. It does, however, argue for building a retrospective monitoring capability into merger oversight, so that patterns of this kind can be investigated as they emerge rather than reconstructed after the fact by academic research. The T+1/T+2 window used in this thesis offers a template: a mandatory post-completion reporting obligation for qualifying acquirers, timed to capture the horizon over which the observed pattern materialises, would let authorities generate exactly the kind of longitudinal evidence this thesis had to reconstruct from Orbis data.

5.3 Limitations

Several limitations constrain the inferential scope of these findings. First, and most consequential, the violation of parallel pre-trends (Section 4.3) formally precludes causal identification, as unobservable confounders driving both target selection and subsequent revenue trajectory cannot be excluded, and the entropy-balancing exercise indicates this is a genuine limit of selection-on-observables correction rather than an unaddressed robustness check. Second, the Orbis panel is subject to non-random survivorship and reporting bias toward larger, listed, European-domiciled targets, limiting generalisability to private, US-domiciled, and sub-threshold acquisitions. Third, the revenue-based HHI, computed over target-panel revenues only, is uninformative about true market-power amplification of the suppression effect and should not be read as evidence against the concentration hypothesis. Fourth, the pre-2022 restriction needed to observe the full T+2 window progressively censors recent acquisition cohorts, so the attenuated T+2 coefficient may reflect genuine decay, sample attrition, or both.

5.4 Directions for future research

Several extensions would strengthen the evidence base. Patent, R&D, and intangible-asset trajectories, benchmarked against synthetic control firms, would provide a direct innovation-discontinuation test analogous to the pharmaceutical KA benchmark, circumventing the financial-reporting heterogeneity that limits revenue-based proxies. Applying the Callaway and Sant'Anna (2021) heterogeneity-robust estimator to a single sub-sector or geography would yield cleaner identification of the staggered treatment effect, addressing the two-way fixed-effects limitation noted in Section 3.4. Moreover, a synthetic control or propensity-matched control group on pre-acquisition growth, profitability, and sector characteristics would partially restore parallel trends and improve causal credibility, using the event study pre-trend as a matching-quality diagnostic. Product-level discontinuation data from app stores or company-profiling services would permit a technology-sector analogue of the trial-discontinuation variable in Cunningham, Ederer and Ma (2021). Finally, restricting the sample to venture-capital-backed targets at acquisition would test whether the KA consistent decay is more pronounced where the entry-for-acquisition dynamic of Kamepalli, Rajan and Zingales (2020) is theoretically most salient.

REFERENCES

Affeldt, P. and Kesler, R. (2021) 'Competitors' Reactions to Big Tech Acquisitions: Evidence from Mobile Apps', *SSRN Electronic Journal* [Preprint].

Arrow, K.J. (1972) 'Economic Welfare and the Allocation of Resources for Invention', in C.K. Rowley (ed.) *Readings in Industrial Economics*. London: Macmillan Education UK, pp. 219–236.

Barnett, J. (2023) '"Killer Acquisitions" Reexamined: Economic Hyperbole in the Age of Populist Antitrust', *SSRN Electronic Journal* [Preprint].

Blasnik, M. (2010) 'RECLINK: Stata module to probabilistically match records'.

Cabral, L. (2024) 'Big Tech Acquisitions', *International Journal of Industrial Organization*, 103, p. 103127.

Callaway, B. and Sant'Anna, P.H.C. (2021) 'Difference-in-Differences with multiple time periods', *Journal of Econometrics*, 225(2), pp. 200–230.

Crémer, J., De Montjoye, Y.-A. and Schweitzer, H. (2019) *Competition policy for the digital era*. Publications Office of the European Union.

Cunningham, C., Ederer, F. and Ma, S. (2021) 'Killer Acquisitions', *Journal of Political Economy*, 129(3), pp. 649–702.

Federico, G., Morton, F.S. and Shapiro, C. (2020) 'Antitrust and Innovation: Welcoming and Protecting Disruption', *Innovation Policy and the Economy*, 20, pp. 125–190.

Gans, J.S. and Stern, S. (2000) 'Incumbency and R&D Incentives: Licensing the Gale of Creative Destruction', *Journal of Economics & Management Strategy*, 9(4), pp. 485–511.

Gautier, A. and Lamesch, J. (2021) 'Mergers in the digital economy', *Information Economics and Policy*, 54, p. 100890.

Gilbert, R.J. and Newbery, D.M. (1982) 'Preemptive patenting and the persistence of monopoly', *The American Economic Review*, pp. 514–526.

Hainmueller, J. (2011) 'Entropy Balancing for Causal Effects: A Multivariate Reweighting Method to Produce Balanced Samples in Observational Studies', *SSRN Electronic Journal* [Preprint].

Ivaldi, M., Petit, N. and Unekbas, S. (2023) 'Killer Acquisitions: Evidence from EC Merger Cases in Digital Industries', *SSRN Electronic Journal* [Preprint].

Jin, G.Z., Leccese, M. and Wagman, L. (2023) 'How Do Top Acquirers Compare in Technology Mergers? New Evidence from an S&P Taxonomy', *International Journal of Industrial Organization*, 89, p. 102891.

Kalemli-Özcan, Ş. *et al.* (2024) 'How to Construct Nationally Representative Firm-Level Data from the Orbis Global Database: New Facts on SMEs and Aggregate Implications for Industry Concentration', *American Economic Journal: Macroeconomics*, 16(2), pp. 353–374.

Kamepalli, S.K., Rajan, R. and Zingales, L. (2020) *Kill Zone*. w27146. Cambridge, MA: National Bureau of Economic Research, p. w27146.

Katz, M.L. (2021) 'Big Tech mergers: Innovation, competition for the market, and the acquisition of emerging competitors', *Information Economics and Policy*, 54, p. 100883.

Kim, J.D. (2022) 'Startup acquisitions, relocation, and employee entrepreneurship', *Strategic Management Journal*, 43(11), pp. 2189–2216.

Koski, H., Kässi, O. and Braesemann, F. (2026) 'Killers on the road of emerging start-ups—implications for market entry and venture capital financing', *Industrial and Corporate Change*, p. dtaf061.

Letina, I., Schmutzler, A. and Seibel, R. (2024) 'KILLER ACQUISITIONS AND BEYOND: POLICY EFFECTS ON INNOVATION STRATEGIES', *International Economic Review*, 65(2), pp. 591–622.

Motta, M. and Peitz, M. (2021) 'Big tech mergers', *Information Economics and Policy*, 54, p. 100868.

Nadler, J. (no date) 'SUBCOMMITTEE ON ANTITRUST, COMMERCIAL, AND ADMINISTRATIVE LAW OF THE COMMITTEE ON THE JUDICIARY OF THE HOUSE OF REPRESENTATIVES'.

Reinganum, J.F. (1983) 'Uncertain innovation and the persistence of monopoly', *The American Economic Review*, 73(4), pp. 741–748.

Renneboog, L. and Vansteenkiste, C. (2019) 'Failure and success in mergers and acquisitions', *Journal of Corporate Finance*, 58, pp. 650–699.

Schumpeter, J.A. (2013) *Capitalism, socialism and democracy*. Routledge.

Seru, A. (2014) 'Firm boundaries matter: Evidence from conglomerates and R&D activity', *Journal of Financial Economics*, 111(2), pp. 381–405.

Winkler, W.E. (1990) 'String comparator metrics and enhanced decision rules in the fellegi-sunter model of record linkage.'

APPENDICES

Appendix A: regression tables

Table A.0 — Descriptive statistics

Variable	N (obs)	Mean	Std. Dev.	p25	Median (p50)	p75
Panel A: Financial Variables (EUR thousands, raw values)						
Operating Revenue	108,185	220,540	1,253,531	1,260	5,540	23,459
Total Assets	125,626	222,249	1,303,906	1,087	4,300	18,780
Intangible Assets	115,386	58,804	413,296	0	11	571
R&D Expenditure	14,937	322,039	1,097,292	0	671	9,016
Panel B: Efficiency and Leverage Ratios (%)						
EBITDA Margin (%)	43,650	8.30	22.54	1.83	8.68	18.46
Return on Assets (%)	57,330	3.34	22.39	−1.80	4.50	12.88
Financial Leverage (%)	65,036	65.99	120.39	0.61	17.12	72.40
ROIC (winsorised, ratio)	65,686	0.151	0.403	0.005	0.114	0.283
Panel C: Treatment and Identification Variables						
Sub-Sector Overlap (0/1)	957,110	0.347	0.476	0	0	1
Sector Overlap (0/1)	957,110	0.507	0.500	0	1	1
HHI (sub-sector × year)	1,044,384	454.5	808.3	209.8	299.3	483.4

Notes: Financial variables from Moody's Orbis Bureau van Dijk, in EUR thousands. Full panel: 1,044,384 firm-year observations across 36,134 deal groups, 1997–2025 (of the 44,189 Orbis-confirmed target matches reported in Section 3.1, the balance corresponds to confirmed matches with no usable firm-year observation once reshaped into the long-format panel, Section 3.2). R&D sub-sample: 1,292 deal groups with non-missing R&D data. ROIC proxy = EBIT / IC, where IC = Equity_derived + Total Debt; Equity_derived = Tot_assets × (ROA/ROE) (DuPont identity); Total Debt = Fin_leverage × Equity_derived (Fin_leverage read as a debt-to-equity gearing ratio); winsorised at 1st–99th percentiles. Sub-sector and sector overlap indicators: proportion of firm-year observations flagged as horizontal deals. HHI constructed at sub-sector × year level over target-panel revenues only (acquirer revenues excluded). Log transformations applied to Operating Revenue, Total Assets, EBITDA, Intangible Assets, and R&D prior to estimation (conditional on strictly positive values). Regression sample after listwise deletion: 65,036 obs / 11,236 groups (static models A–B); 55,599 obs / 9,622 groups (dynamic models C and related). N for EBITDA Margin, Return on Assets, Financial Leverage and ROIC reflects each variable's own regression-estimation sample (consistent with the analytic sample emphasis above), not full-panel data availability; the corresponding Mean, SD, and percentiles are computed on the full available data for that variable.

Table 1 — Main DiD results: post-acquisition revenue effects

	(A) *Sector-Level*	(B) *Sub-Sector*	(C) *Dynamic DiD*
DiD × Post (T=0)	-0.035 *(0.030)*	0.013 *(0.033)*	0.084** *(0.035)*
DiD × Post (T+1)	—	—	-0.062** *(0.031)*
DiD × Post (T+2)	—	—	-0.033 *(0.030)*
Post-Acquisition Indicator	-0.013 *(0.021)*	-0.033* *(0.019)*	-0.048** *(0.021)*
Log Total Assets	0.836*** *(0.015)*	0.836*** *(0.015)*	0.821*** *(0.018)*
Financial Leverage	-0.000* *(0.000)*	-0.000* *(0.000)*	-0.000*** *(0.000)*
Deal & Year FE	Yes	Yes	Yes
Observations	65,036	65,036	55,599
Deal Groups	11,236	11,236	9,622
R^2 (within)	0.393	0.393	0.372
Sample	Full	Full	≤ 2022

*DV: log operating revenues (EUR thousands). In Model (A), DiD T=0 refers to did_sector (sector-level overlap × post); in Models (B)–(C), to did_subsector (sub-sector overlap × post). Models (A) and (B): full sample. Model (C): announcements ≤ 2022. Robust SEs in parentheses, clustered at deal level. Deal and year FE absorbed. * p<0.10, ** p<0.05, *** p<0.01.*

Table 2 — HHI robustness and market concentration moderation

	(B) *Baseline*	**(D)** *HHI Control*	**(E)** *HHI Moderation*	**(F)** *Dynamic + HHI*
DiD × Post (T=0)	0.013 *(0.033)*	0.013 *(0.033)*	0.016 *(0.036)*	0.076* *(0.040)*
DiD × Post (T+1)	—	—	—	-0.045 *(0.043)*
DiD × Post (T+2)	—	—	—	-0.050 *(0.045)*
HHI (standardised)	—	-0.007 *(0.020)*	-0.007 *(0.020)*	-0.008 *(0.017)*
DiD × HHI (T=0)	—	—	0.012 *(0.096)*	-0.037 *(0.107)*
DiD × HHI (T+1)	—	—	—	0.069 *(0.129)*
DiD × HHI (T+2)	—	—	—	-0.070 *(0.133)*
Post-Acquisition Indicator	-0.033* *(0.019)*	-0.033* *(0.019)*	-0.033* *(0.018)*	-0.047** *(0.021)*
Log Total Assets	0.836*** *(0.015)*	0.836*** *(0.015)*	0.836*** *(0.015)*	0.821*** *(0.018)*
Financial Leverage	-0.000* *(0.000)*	-0.000* *(0.000)*	-0.000* *(0.000)*	-0.000*** *(0.000)*
Deal & Year FE	Yes	Yes	Yes	Yes
Observations	65,036	65,036	65,036	55,599
Deal Groups	11,236	11,236	11,236	9,622
R^2 (within)	0.393	0.393	0.393	0.372
Sample	Full	Full	Full	≤ 2022

*DV: log operating revenues (EUR thousands). HHI standardised (mean 0, SD 1). Models (B), (D), (E): full sample. Model (F): announcements ≤ 2022. Robust SEs in parentheses, clustered at deal level. Deal and year FE absorbed. * p<0.10, ** p<0.05, *** p<0.01.*

Table 2b — Alternative concentration proxy: deal frequency moderation

	(E) *HHI Moderation*	**(E2)** *Deal Frequency*
DiD × Post (T=0)	0.016 *(0.036)*	0.008 *(0.037)*
HHI (standardised)	-0.007 *(0.020)*	—
DiD × HHI (T=0)	0.012 *(0.096)*	—
Deal Frequency (standardised)	—	1.175 *(0.905)*
DiD × Deal Frequency (T=0)	—	-0.061** *(0.029)*
Post-Acquisition Indicator	-0.033* *(0.018)*	-0.039* *(0.021)*
Log Total Assets	0.836*** *(0.015)*	0.828*** *(0.017)*
Financial Leverage	-0.000* *(0.000)*	-0.000** *(0.000)*
Deal & Year FE	Yes	Yes
Observations	65,036	56,723
Deal Groups	11,236	9,775
R^2 (within)	0.393	0.374
Sample	Full	≤ 2022

*DV: log operating revenues (EUR thousands). Deal frequency = count of same-sub-sector acquisitions per sub-sector × year (standardised). Model (E): full sample. Model (E2): announcements ≤ 2022. Robust SEs in parentheses, clustered at deal level. Deal and year FE absorbed. * p<0.10, ** p<0.05, *** p<0.01.*

Table 3 — Robustness: alternative dependent variables

	(B) *Log Revenue*	**(G)** *Log Intangibles*	**(H)** *Log Intg. (Dyn.)*	**(I)** *Log R&D*
DiD × Post (T=0)	0.013 *(0.033)*	-0.057 *(0.071)*	-0.049 *(0.077)*	-0.032 *(0.081)*
DiD × Post (T+1)	—	—	0.052 *(0.058)*	—
DiD × Post (T+2)	—	—	-0.099 *(0.061)*	—
Post-Acquisition Indicator	-0.033* *(0.019)*	0.011 *(0.042)*	-0.008 *(0.048)*	-0.016 *(0.055)*
Log Total Assets	0.836*** *(0.015)*	0.903*** *(0.027)*	0.907*** *(0.031)*	0.578*** *(0.021)*
Financial Leverage	-0.000* *(0.000)*	0.001*** *(0.000)*	0.001*** *(0.000)*	0.000 *(0.000)*
Deal & Year FE	Yes	Yes	Yes	Yes
Observations	65,036	54,988	46,810	8,877
Deal Groups	11,236	9,818	8,361	1,292
R^2 (within)	0.393	0.162	0.150	0.315
Sample	Full	Full	≤ 2022	Full

*DV by column: (B) log operating revenues; (G) log intangible assets (static); (H) log intangible assets (dynamic); (I) log R&D expenditure. Models (B), (G), (I): full sample. Model (H): announcements ≤ 2022. Model (I) estimated on sub-sample of deal groups reporting R&D data (N=8,877; 1,292 groups). Robust SEs in parentheses, clustered at deal level. Deal and year FE absorbed. * p<0.10, ** p<0.05, *** p<0.01.*

Table 4 — Margin and capital efficiency analysis (dynamic DiD)

	(C) *Log Revenue*	(M2) *EBITDA Margin*	(M3) *ROA*	(M4) *ROIC*
DiD × Post (T=0)	0.084** *(0.035)*	2.013** *(0.864)*	2.140*** *(0.782)*	0.032* *(0.017)*
DiD × Post (T+1)	-0.062** *(0.031)*	0.099 *(0.828)*	-0.334 *(0.731)*	-0.002 *(0.016)*
DiD × Post (T+2)	-0.033 *(0.030)*	0.499 *(0.739)*	-0.942 *(0.579)*	-0.014 *(0.013)*
Post-Acquisition Indicator	-0.048** *(0.021)*	-1.432*** *(0.417)*	-1.089*** *(0.376)*	-0.023*** *(0.008)*
Log Total Assets	0.821*** *(0.018)*	3.862*** *(0.292)*	2.459*** *(0.235)*	0.025*** *(0.004)*
Financial Leverage	-0.000*** *(0.000)*	-0.017*** *(0.001)*	-0.027*** *(0.001)*	— (omitted)
Deal & Year FE	Yes	Yes	Yes	Yes
Observations	55,599	43,650	57,330	56,125
Deal Groups	9,622	7,573	9,715	9,566
R^2 (within)	0.372	0.039	0.038	0.008
Sample	≤ 2022	≤ 2022	≤ 2022	≤ 2022

*DV by column: (C) log operating revenues; (M2) EBITDA margin (%); (M3) return on assets (%); (M4) ROIC = EBIT / Invested Capital (Equity_derived + Total Debt; Equity_derived = Tot_assets × ROA/ROE by the DuPont identity; Total Debt = Fin_leverage × Equity_derived). Fin_leverage is excluded from Model M4's control set, as it enters the dependent variable's construction directly. All models: announcements ≤ 2022 (T+2 fully observed). Margin and efficiency models estimated on sub-samples with available financial ratio data. Robust SEs in parentheses, clustered at deal level. Deal and year FE absorbed. * p<0.10, ** p<0.05, *** p<0.01.*

Table 5 — Parallel-trends remediation: entropy-weighted estimates

	(1) *Event Study (leads+logs)*	(2) *Model C (T=0, T+1, T+2)*
DiD Lead: T−2 (pre-trend)	0.074**	
	(0.029)	
DiD Lead: T−1 (pre-trend)	0.033*	
	(0.017)	
DiD: Sub-sector × Post (T=0)	0.042	0.127***
	(0.033)	*(0.035)*
DiD: Sub-sector × T+1	-0.097***	-0.104***
	(0.031)	*(0.032)*
DiD: Sub-sector × T+2	-0.024	-0.046
	(0.034)	*(0.034)*
Post-Acquisition Indicator	-0.119***	-0.150***
	(0.024)	*(0.024)*
Log Total Assets	0.830***	0.834***
	(0.024)	*(0.024)*
Financial Leverage	-0.000	-0.000
	(0.000)	*(0.000)*
Observations	22,067	23,349
Deal Groups	2,809	2,818
R^2 (within)	0.482	0.460

*Notes: Dependent variable is log operating revenues (EUR thousands; nominal, not deflated). Robust standard errors clustered at the deal level in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Sample is entropy-balanced (Hainmueller 2011) on pre-acquisition revenue growth, size, and leverage (targets(1), first moment); deal groups lacking a usable pre-acquisition observation are excluded, which is why N differs from the unweighted specifications in Table 1. Column (1) re-estimates the event study specification of Section 3.4 on the entropy-weighted sample; column (2) re-estimates Model C (the main dynamic DiD). Compare against the unweighted estimates in Table 1, Model C.*

Appendix B: figures

Figure 1 — Event study: revenue effect around acquisition

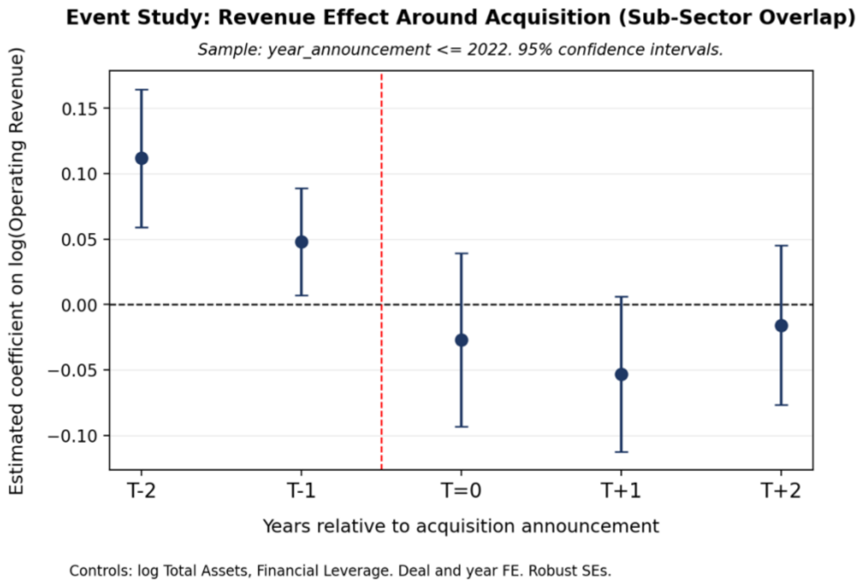

*Note: Coefficient plot from the event study specification augmenting Model C with two lead interaction terms (did_sub_lead2, did_sub_lead1). Dependent variable: log operating revenues. Sample: announcements ≤ 2022 (N = 51,410; 9,387 deal groups). 95% confidence intervals shown. Controls: log total assets, financial leverage. Deal and year FE absorbed. Dashed vertical line at T=0 (acquisition announcement year). Coefficients: T−2 = +0.112*** (0.027); T−1 = +0.048** (0.021); T=0 = −0.027 (0.034); T+1 = −0.053* (0.030); T+2 = −0.016 (0.031).*

Figure 2 — Distribution of HHI across sub-sectors

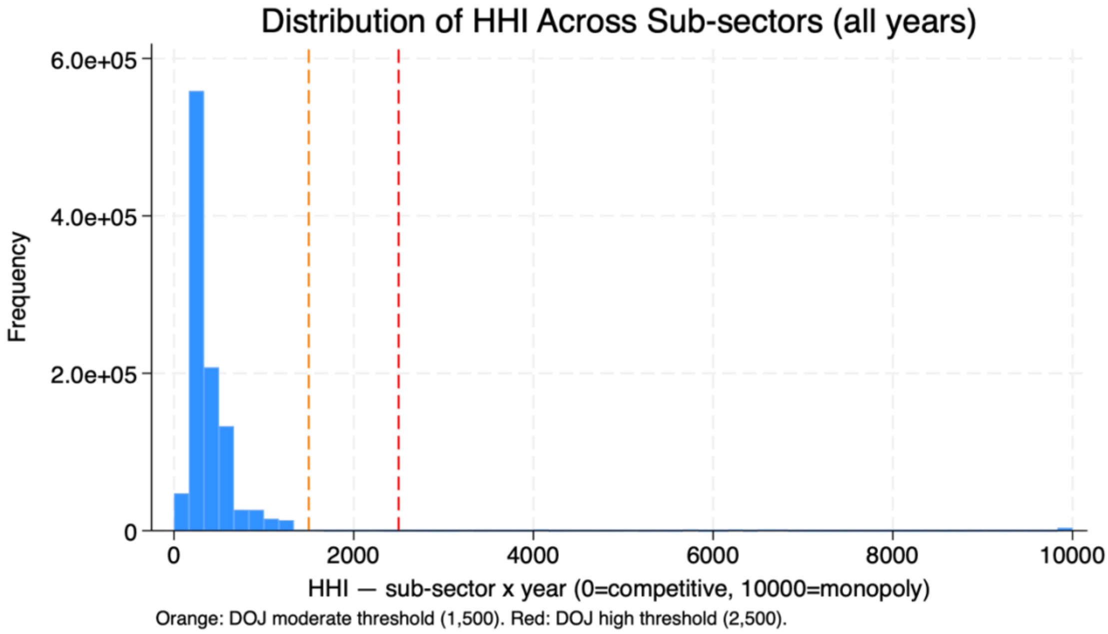

Note: Herfindahl-Hirschman Index (HHI) computed at the sub-sector × year level as the sum of squared revenue market shares scaled by 10,000. The HHI is constructed over the target panel only; acquirer revenues are absent from the denominator. Orange dashed line: DOJ moderate concentration threshold (1,500). Red dashed line: DOJ high concentration threshold (2,500). The distribution is right-skewed, with the median sub-sector × year cell at HHI ≈ 299, reflecting moderate-to-low target-pool concentration.

Figure 3 — Distribution of log operating revenue

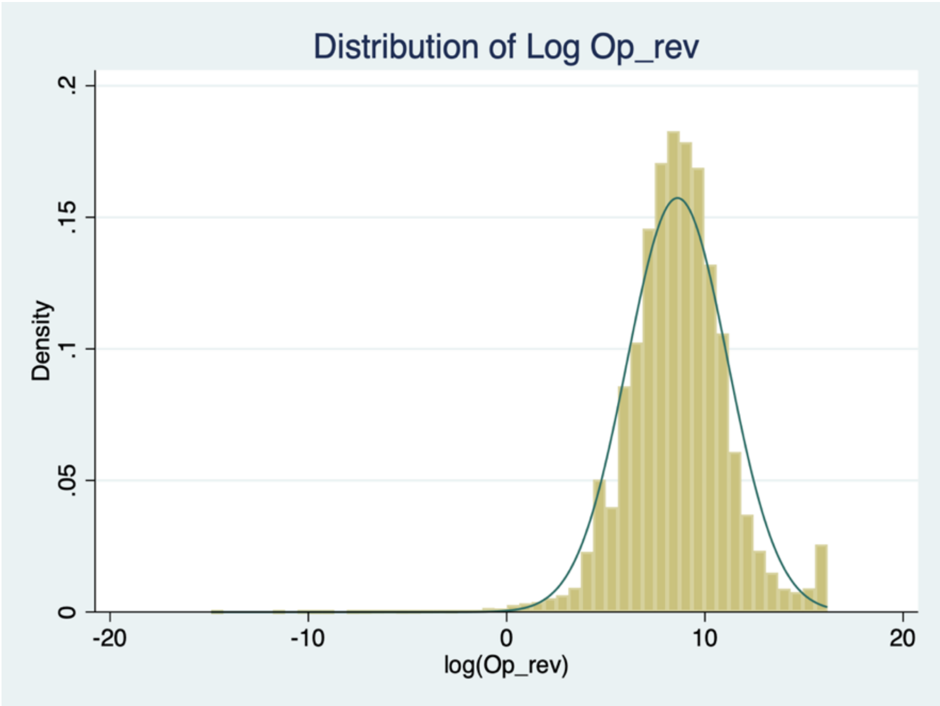

Note: Histogram of log operating revenues with normal density overlay. The distribution is right-skewed in levels and approximately normal in logs, supporting the log transformation applied in all specifications. Sample: all firm-year observations with Op_rev > 0 (N = 108,185).

Figure 4 — Distribution of log total assets

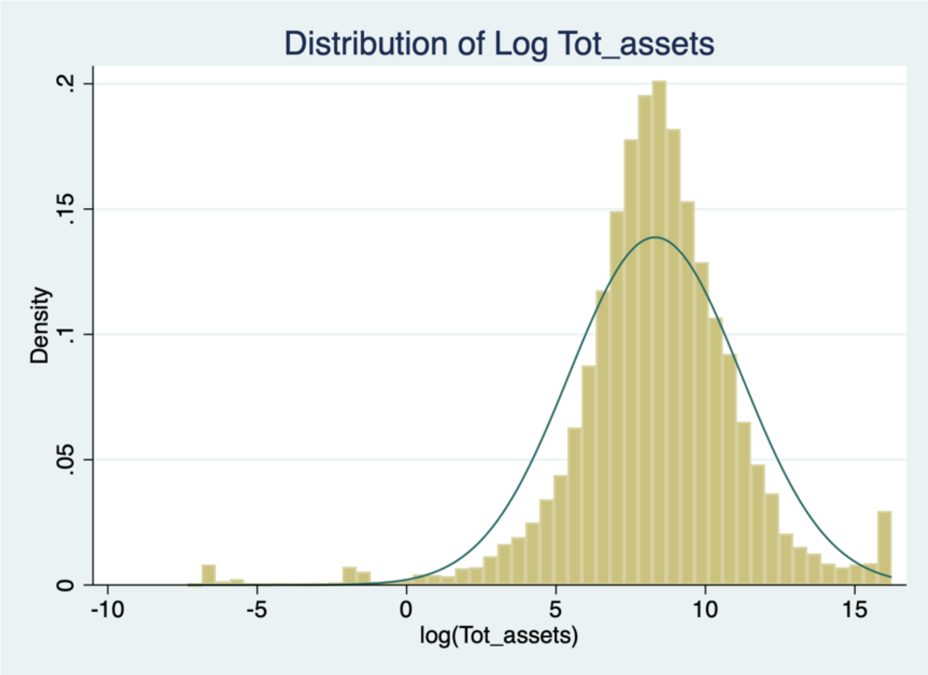

Note: Histogram of log total assets with normal density overlay. The distribution is heavily right-skewed in levels (SD/mean ≈ 5.9×) and approximately normal in logs, supporting the log transformation used as a control across all specifications. Sample: all firm-year observations with Tot_assets > 0 (N = 125,626).