INCENTIVIZING INNOVATION

HOW PRIVATE EQUITY TIME HORIZONS DICTATE LONG-TERM VALUE CREATION

CARL MATTSSON

SENMARU HRNCIRIK-MARUYAMA

Bachelor Thesis

Stockholm School of Economics

2026

Incentivizing Innovation: How Private Equity Time Horizons Dictate Long-term Value Creation

Abstract:

This thesis aims to determine whether the holding period tendencies of Private Equity firms influence the level of innovation within their acquired portfolio companies. This hypothesis relies on the premise that ownership structures with a long-term view are better equipped to incentivize investments with delayed payoffs. Analyzing a sample of U.S. Private Equity buyouts from 1995 to 2015 and quantifying innovative output through both patenting volume and forward citation counts, our results show that while patenting quantity is not noticeably affected by investor time horizons, patenting quality is. This indicates that Private Equity investors with longer holding period strategies tend to push for a higher degree of innovation quality, prioritizing impactful advancements over sheer volume.

Keywords:

Private Equity, Innovation, Patents, Time Horizons, Investments

Authors:

Carl Mattsson (25931)
Senmaru Hrncirik-Maruyama (26005)

Tutor:

Alvin Chen, Assistant Professor, Department of Finance

Examiner:

Ramin Baghai, Professor, Department of Finance

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

1 Introduction

1.1 Motivation & Background

The Private Equity industry markets itself as a vital component of the financial system, one that not only provides private companies with capital but also drives growth and operational enhancements. Metrics such as IRR and MOIC are commonly used to measure the financial returns of Private Equity funds, allowing firms to quantify their value-add (Gompers, Kaplan and Mukharlyamov, 2016). However, these top-line metrics fail to distinguish whether this value creation stems from fundamental operational and governance improvements, or if it is primarily the result of financial engineering and macroeconomic tailwinds.

A robust measure of a company's long-term success and fundamental impact is its capacity for innovation and the subsequent commercial viability of those breakthroughs. Tracking this innovative output is an effective way to measure fundamental value creation, as it demonstrates that the firm is making meaningful contributions through new technologies. Understanding which governance and ownership structures most effectively foster this environment allows investors to make more informed decisions regarding long-term value. Because innovation is inherently a long-term endeavor, requiring immediate capital expenditures to capture benefits in the often distant future, companies that invest heavily in it must be guided by leadership with a sustained, forward-looking focus. Existing literature has extensively analyzed the differences between private and public companies in this regard. For instance, Lerner, Sørensen and Strömberg (2011) examined how taking public companies private through leveraged buyouts (LBOs) affected the acquired companies' innovation, concluding, based on patenting activity, that LBO-backed firms are more incentivized to promote innovation. Conversely, Bernstein (2015) examined the transition from private to public markets, finding that the process of an initial public offering (IPO) tends to reduce a firm's innovative output.

One argument made by Lerner et al. (2011) is that private company management teams are largely shielded from the short-term performance pressures of public markets, granting them the freedom to make decisions that benefit the firm in the long run. This dynamic is empirically supported by Guo, Pérez-Castrillo, and Toldrà-Simats (2019), who show that public companies under heavy coverage by financial analysts are more likely to reduce R&D spending and pivot away from innovation to meet short-term earnings forecasts.

The hypothesis that long-termism promotes innovation extends beyond the simple public-versus-private dichotomy. Because Private Equity funds do not share investment identical time horizons, their influence on a portfolio company's long-term strategic goals should naturally vary. This leads to the core research question of this thesis: How do the holding period strategies of Private Equity funds impact the extent to which they promote innovation within their portfolio companies?

Private Equity firms with longer holding periods have a greater stake in the sustained, long-term performance of their assets compared to firms seeking quick exits. Because the financial rewards of innovation take time to materialize, especially when compared to others, more immediate value-creation levers utilized by Private Equity investors, long-term focused funds should theoretically be more willing to invest in R&D, as they remain invested long enough to capture the eventual upside. Long-term Private Equity funds, in theory, have the right incentive structure to foster innovation in their portfolio companies relative to short-term funds (Manso, 2011). This paper will empirically evaluate this relationship.

1.2 Literature Review

A foundational study for this research is Lerner et al. (2011)'s analysis of Private Equity and long-run investment. The authors examine the impact of Private Equity-backed leveraged buyouts (LBOs) on innovation within acquired firms, arguing that shielding management from the short-term earnings pressures of public shareholders increases their willingness to make long-term strategic decisions, such as investing in R&D. Because R&D spending is often undisclosed or inconsistently reported for private companies, it cannot serve as a reliable comparative metric. Overcoming this data limitation, Lerner et al. (2011) utilizes patenting activity, specifically patent counts, forward citations, and original grants, as a proxy for innovation, as patent data is publicly available regardless of a company's ownership structure. They conclude that both the volume and quality of patenting activity tend to increase following an LBO.

Exploring a parallel dynamic, Bernstein (2015) operates on the same core hypothesis: that the short-term performance pressures of public markets negatively impact innovation. However, rather than analyzing Private Equity buyouts, Bernstein examines the reverse transition by observing private companies going public. Utilizing similar metrics of patent counts and forward citations, Bernstein (2015) concludes that transitioning to a public company reduces a firm's internal innovative output. A potential limitation of Bernstein's analysis, however, is its reliance on internal patenting activity pre- and post-IPO without fully accounting for the tendency of public firms to innovate through M&A. As Hovakimian and Hutton (2010) note, newly public companies often leverage their improved access to capital to acquire other innovative firms. This substitution of internal R&D for external acquisitions could partially explain the observed decline in post-IPO patenting activity.

Another highly relevant study is Barrot (2012), which investigates the relationship between investor time horizons and innovative behavior. While framed around "Private Equity," Barrot (2012)'s dataset primarily consists of early-stage Venture Capital funds. To proxy for time horizons, Barrot (2012) measures how close a fund is to its statutory 10-year liquidation date at the time of investment. The study finds that investments made closer to the fund's liquidation deadline are typically directed toward mature, less risky, and less innovative companies, whereas capital deployed early in the fund's lifecycle targets young, innovation-heavy firms. Like the aforementioned studies, Barrot (2012) quantifies innovation through patent volume and citations.

However, applying Barrot (2012)'s findings to traditional Private Equity presents several limitations. First, the study assumes a rigid 10-year lifespan for all funds and does not analyze idiosyncratic differences in holding period tendencies between different fund managers. Second, because Venture Capital funds typically acquire minority stakes, they lack the unilateral governance control necessary to drive post-investment innovation changes; Barrot (2012)'s findings therefore reflect investment selection rather than governance impact. Finally, analyzing the remaining time until a fund's liquidation is less relevant for limited partners (LPs) allocating capital, as LPs commit their funds entirely at the inception of the vehicle. Understanding the historical holding period tendencies of specific Private Equity funds would provide a much more actionable metric for LPs seeking long-term impact and returns.

This thesis builds directly upon this established body of literature. We adopt the foundational hypothesis that long-term oriented ownership fosters greater innovation compared to structures incentivizing short-term performance, and we utilize the established methodology of measuring innovation via patent counts and forward citations. However, we expand upon this framework by isolating a critical gap in the literature: examining how varying time horizons and holding period tendencies between different Private Equity funds dictate the extent to which they drive innovation within their portfolio companies.

1.3 Hypothesis Development

The Private Equity literature is divided across two opposing analytical frameworks, which we have summarized as the long-termism and short-termism views. This dichotomy not only serves as a helpful framework with which to understand the operating mechanisms behind holding periods, but forms the basis of our study.

Long-termism is the moniker used for the long-term investor. The long-term investor offers stability along with capital and long term leadership. Their investment structure is more willing to tolerate early failures in order to secure future payoffs built on business fundamentals. Conversely, short-termism, used to describe the short-term investor, represents the tendency to forego traditionally business fundamentals in pursuit, rapidly pushing financial metrics in pursuit of a profitable exit. As Manso (2011) laid out, innovation thrives in conditions where early failures are tolerated and long-term payoffs are rewarded. Thus, with the long-termism and short-termism frameworks in mind, our hypothesis is:

- **H0:** Longer time horizon Private Equity firms are not associated with higher post-acquisition innovation (an umbrella term measured by patent count and citations, a metric that will be justified in later sections)
- **H1:** Longer time horizon Private Equity firms are associated with higher post-acquisition innovation (an umbrella term measured by patent count and citations, a metric that will be justified in later sections)

We intend to test this hypothesis by (1) identifying a relevant dataset of companies acquired by Private Equity firms, (2) finding the patents associated with these companies, and testing the trends of this firm-level patent activity against the focal coefficient (time horizon of the Private Equity funds) before and after the transaction. The focal coefficient is identified from variation in average holding period across the Private Equity firms in our sample.

This is not an analysis on the comparative advantages of either of these (long-termism vs. short-termism), it is rather a dive into whether each of these clusters (and everything in between) is more conducive to meaningful innovation in their portfolio companies, and if so, how does a metric like time horizon relate to innovation? What is it that time horizon implies about the philosophy and the incentives undergirding a Private Equity fund's investments and leadership?

2 Data & Sample Construction

To conduct this analysis, our data collection process required three primary components: a definitive list of companies acquired by Private Equity firms, the corresponding patent data for each of those target companies, and the historical holding period data for the acquiring Private Equity firms.

2.1 Target Company Identification

To identify our initial sample of Private Equity-backed companies, we utilized the Capital IQ database. Because Capital IQ does not explicitly separate Private Equity from Venture Capital in its buyer categorization, we applied a series of rigorous screening criteria to isolate traditional Private Equity buyouts. We restricted our search to mergers and acquisitions (M&A) that met the following conditions:

- **Transaction Value:** Total transaction value exceeding 25 million USD, serving as a baseline to filter out smaller, early-stage Venture Capital investments. The decision to use 25 million USD specifically can be supported through the findings of Gompers, Kaplan, and Mukharlyamov(2016), stating that most Private Equity firms don't invest in companies with an enterprise value of less than that.
- **Ownership Stake:** Acquisitions representing 100% ownership of the target company. This further excluded Venture Capital minority stakes and ensured the acquiring Private Equity firms possessed full operational and strategic control over the targets.
- **Timeframe and Geography:** Transactions finalized between 1995 and 2015, where the target company was headquartered in the United States. The US limitation was imposed to ensure higher quality and consistency in patent reporting, similar to that of Lerner et al. (2011). The date range couldn't be too late back in time, as it could decrease the consistency of the patent data, and it couldn't be too close to the present as that could risk younger companies not having had enough time to properly accumulate forward citations on their patents.
- **Single Acquirer:** Transactions involving only one acquiring Private Equity firm. Club deals (multiple Private Equity acquirers) were removed to allow us to definitively isolate which firm was responsible for post-acquisition governance.
- **Initial Deal Only:** For target companies that underwent multiple secondary buyouts over time, only the initial was retained, ensuring a one-to-one match between target companies and Private Equity firms.
- **Relevant Sectors:** The target companies were restricted to patent-heavy industries: IT, Industrials, Health Care, Energy, Consumer Discretionary, and Communication Services. These sectors were decided based on similar sector filterings made by Ramzan, Lau, Bawazir, Belhaj (2025).

Applying these filters resulted in a preliminary sample of 776 target companies acquired by Private Equity firms.

2.2 Patent Data Acquisition

Unlike Lerner et al. (2011), who utilized a proprietary Harvard Business School database containing pre-cleaned company names, we sourced our patent data from PATSTAT (maintained by the European Patent Office), which comprehensively tracks both US and international patent filings. PATSTAT was selected primarily for its highly customizable SQL search capabilities, allowing for bulk querying.

A common challenge in patent research is the lack of standardization in applicant names, leading to numerous misspellings and variations for the same corporate entity, an issue especially prevalent among private companies. To address this, we employed a ”wildcard search” (Wang, Jia, Wang, Yang, Fu, and Xu, 2015), querying PATSTAT for all entities with names vaguely resembling our 776 target companies. This generated approximately 16,000 company entries.

To systematically match these 16,000 variations to our target list, we utilized a matching process using Claude, an AI by Anthropic (the exact methodology is detailed in Appendix C). This process successfully mapped 3,143 PATSTAT variations to 309 of our target companies from Capital IQ, rejecting the remainder. To validate the integrity of this matching process and screen for false positives, we conducted a manual audit on a $\sim$5% random sample (164 matches). The manual check confirmed 160 correct matches, yielding a 97.6% accuracy rate and increasing the reliability of the dataset.

Following the name-matching process, we extracted all associated patents, including their application filing dates and unique forward family citation counts, which are forward citations counted at the family level, preventing duplicate counts when the same patent is filed in multiple jurisdictions (Higham and Yoshioka, 2022). We then removed any target companies that exhibited zero patenting activity within a three-year window surrounding their respective Private Equity acquisition. This attrition left us with a refined sample of 218 patent-active target companies.

2.3 Estimating Private Equity Holding Periods

The 218 target companies in our refined sample were acquired by 140 distinct Private Equity firms. To calculate the historical holding period tendencies for these firms, we returned to Capital IQ to screen for historical buy and sell transactions.

We pulled all M&A transactions where the buyer was one of our 140 Private Equity firms, and subsequently pulled all M&A and IPO transactions where the seller was one of those same 140 firms, matching the entry and exit dates for historical portfolio companies between 1980 and 2020.

While calculating a firm’s average holding period using the full dataset includes some deals that occurred after our specific target acquisitions, we argue this approach is appropriate for our research question for three reasons. First, we view a firm’s holding period tendency as a structural, long-term proxy for its fundamental investment philosophy, which is best captured across the firm’s entire lifespan. Second, calculating a full 40-year average smooths out macroeconomic shocks, such as the 2008 financial crisis, which can inflate holding periods, as discussed by Joenväärä, Mäkiaho, and Torstila (2022). Finally, restricting the calculation strictly to historical, pre-acquisition deals would severely limit the sample size of deals for some of our firms, introducing significant small-sample bias into our independent variable. By utilizing the 1980–2020 panel, we ensure a robust, steady estimate of each firm’s true strategic timeline.

To avoid endogeneity, we explicitly excluded all 218 target companies in our primary sample from this calculation. This ensured that our proxy for a Private Equity firm’s average holding period was not derived from the exact same deals we were using to analyze innovation. To ensure statistical reliability in our time-horizon estimates, we also removed any Private Equity firm with a track record of fewer than four complete deals.

After calculating the average holding period for the remaining firms, we were left with robust time-horizon estimates for 98 private equity firms. Merging this final constraint with our patent dataset yielded a final, consolidated sample of 167 target companies associated with 3,435 patents, acquired by 98 distinct Private Equity firms. Our sample process can be visualized in Figure 1 below.

Figure 1: Sample Selection Process

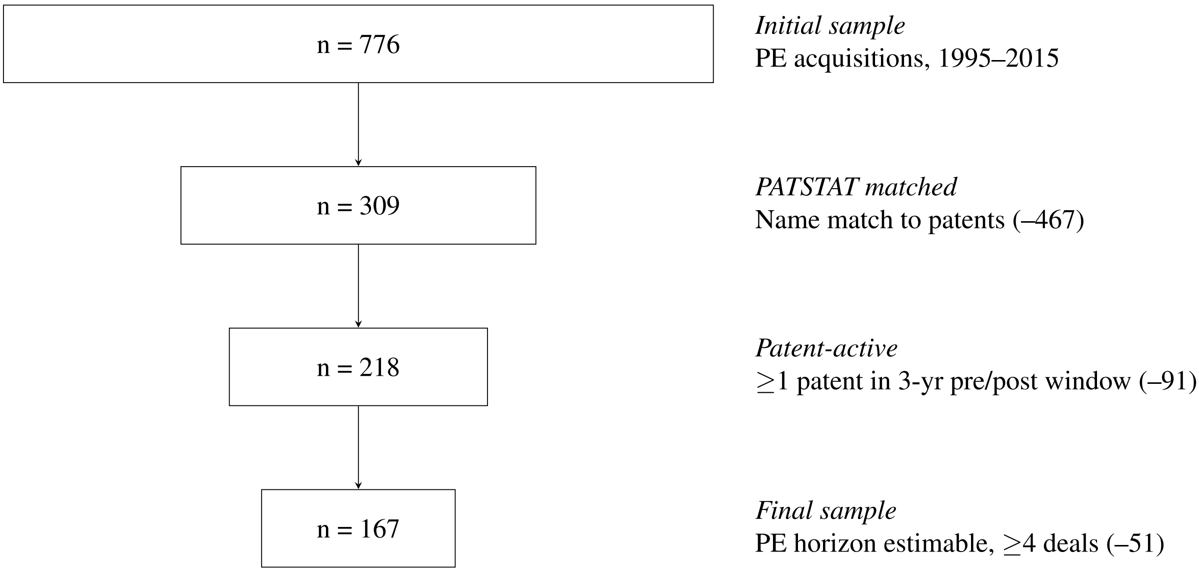

Notes: The flowchart illustrates the sequential data filtering and sample attrition process used to construct the final dataset. The initial sample consists of 776 US-headquartered target companies in patent-heavy sectors acquired via 100% Private Equity buyouts ($25m transaction value) between 1995 and 2015. The sample is reduced to 309 after matching target company names to PATSTAT patent records using algorithmic wildcard searches. Requiring at least one patent in the three-year window surrounding the acquisition further reduces the sample to 218 patent-active firms. Finally, isolating transactions where the acquiring Private Equity firm had a track record of ≥4 historical deals (between 1980 and 2020) to estimate investment horizons yields a final sample of 167 target companies associated with 3,435 patents across 98 distinct Private Equity firms.

2.4 Descriptive Statistics

Figure 2: Distribution of Private Equity Time Horizons (N = 167)

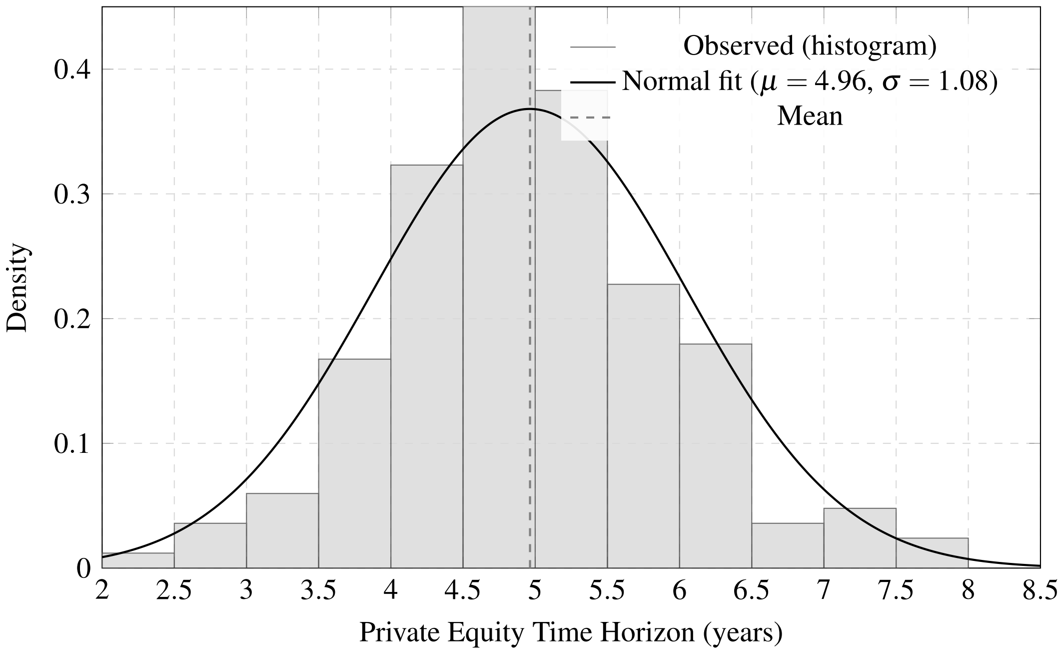

Notes: Distribution of Private Equity time horizons across the sample, overlaid with a fitted normal distribution. The empirical distribution is approximately bell-shaped but shows mild right skew, with a few longer-horizon firms in the upper tail (7–8 years).

The distribution of time horizons for the Private Equity firms acquiring the 167 target companies is visualized in figure 2. As shown, most firms average a time horizon of roughly 5 years, with the longest falling just under 8 years and the shortest just over 2 years. This spread is critical for our tests, as a lack of variance, where all firms shared an identical holding period, would make it difficult to identify meaningful statistical relationships in the data.

Figure 3: Distribution of Private Equity Deals by Year (N = 167)

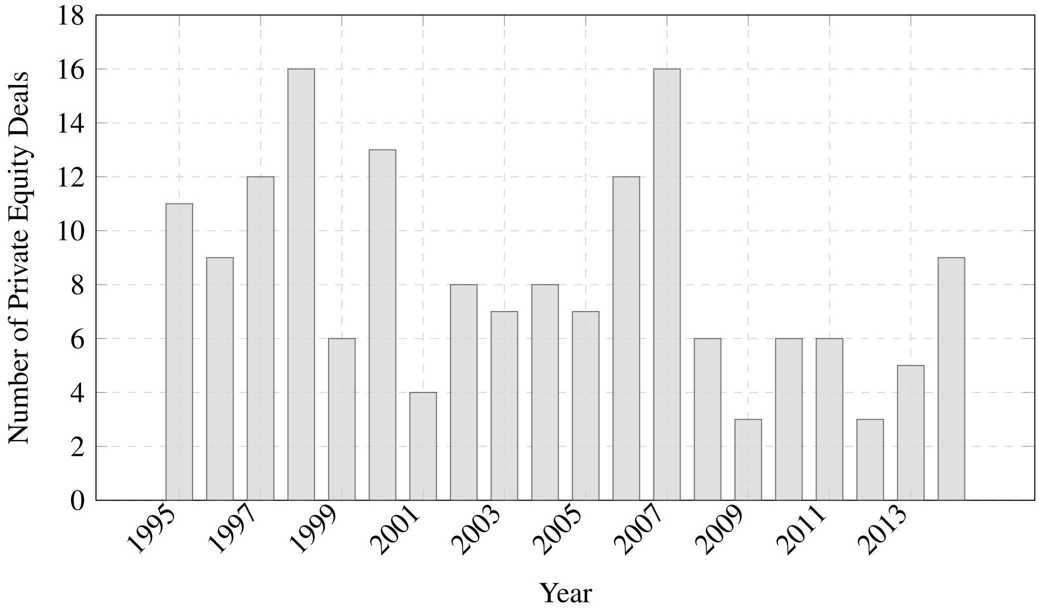

Notes: Distribution of Private Equity deals by year across the sample period (1995–2014). The sample shows two clear peaks of deal activity in 1998 and 2007, followed by sharp declines coinciding with the dot-com bust and the 2008 financial crisis respectively.

Figure 3 represents the 167 companies we are analyzing, distributed by the years they were acquired. As seen, there are more companies acquired earlier in the time period than later, with an especially low amount post 2008, which could be explained by the financial crash. The tendency for the deals to be earlier in the time period is likely because we've only kept companies that we've been able to find patent data on and there being a higher degree of patent data in the 90s and early 2000s than in later periods.

Figure 4: Distribution of Target Companies by Sector (N = 167)

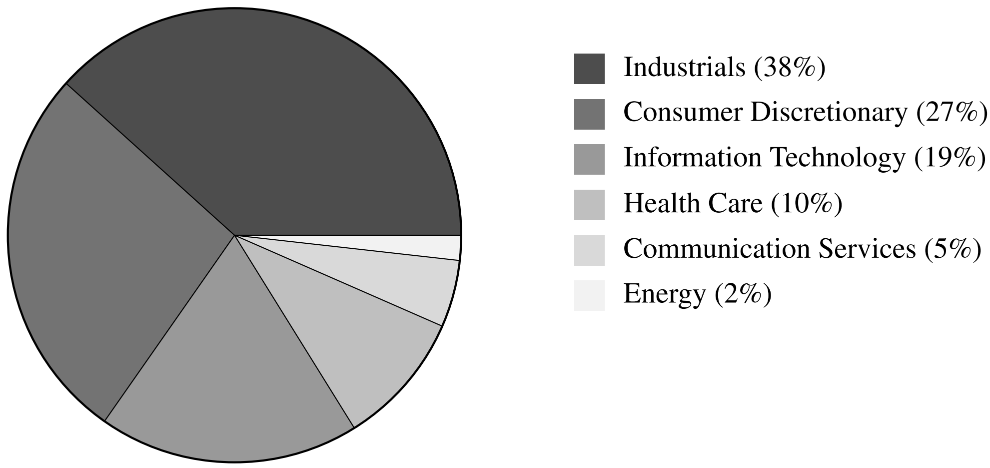

Notes: Distribution of target companies by sector (N = 167). Industrials is the most common sector in the sample (38%), followed by Consumer Discretionary (27%) and Information Technology (19%).

The distribution of companies within our final could either suggest that innovation most commonly occurs within the Industrials and Consumer Discretionary sectors, that they are the sectors most likely to be bought out by Private Equity firms, or that there are more firms within those sectors in general.

Figure 5: Patent Counts by Sector: Pre- vs. Post-Acquisition

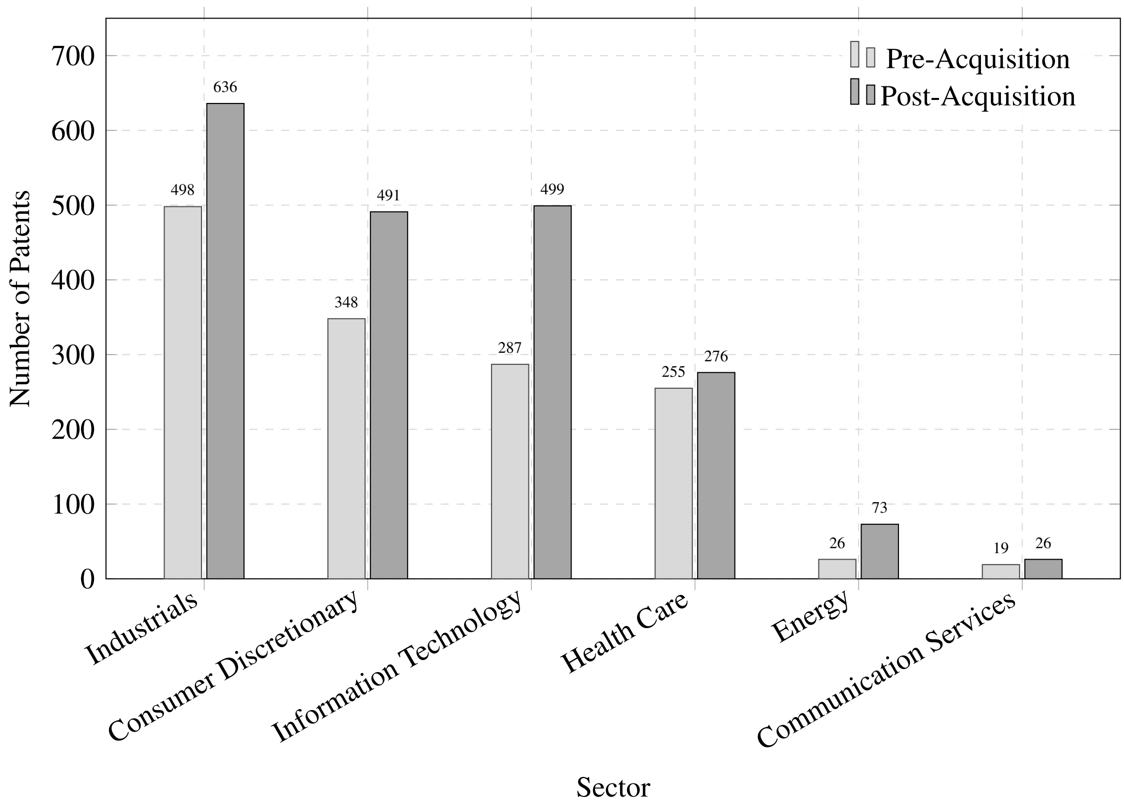

Notes: Patent counts by sector before and after acquisition. Every sector shows an increase in patenting activity post-acquisition, with the largest absolute gains in Industrials (+138), Information Technology (+212), and Consumer Discretionary (+143). Information Technology shows the largest relative increase (+74%).

The patenting activity has increased post Private Equity acquisition for every sector, which can be seen through the increase of patenting application counts. This is in line with Lerner et al. (2011) as it suggests that Private Equity firms have a positive effect on the innovation activity within their portfolio companies, and is a good starting point for us as we can then take it further and examine the extent to which this increase in innovation activity is affected by Private Equity time horizons.

Figure 6: Average Patents per Company Around the Acquisition Year

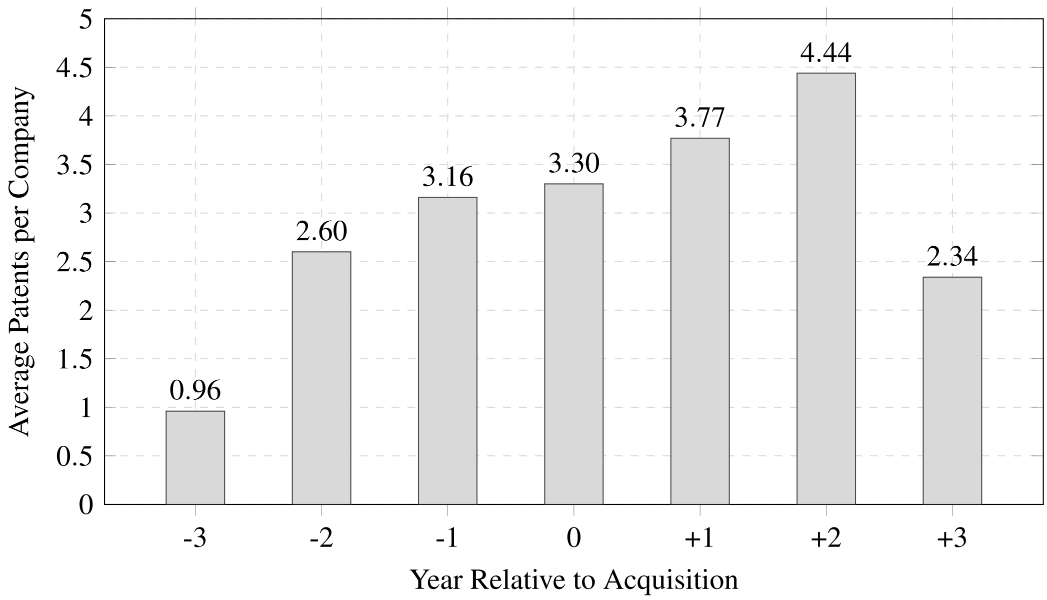

Notes: Average patents per company in the seven-year window centered on the acquisition year (year 0). Patent activity rises through the pre-acquisition period, continues climbing through the first two post-acquisition years (peaking at 4.44 patents per company in year +2), and then drops sharply in year +3.

When deconstructing the patent counts into the years pre and post acquisition, we get a more detailed view over the patenting trends. The increase in patenting activity up until acquisition could be firms looking to sell to Private Equity firms, actively filing for patents to increase their perceived value. Patenting does tend to increase in the 2 years following acquisitions, which could either be Private Equity firms pushing for more innovation within the companies or simply just pushing for more patent applications on innovations already made to make it look like more innovation is being made. Patents decreasing in year 3 post acquisition could be an indicator of this, as if this is what they are doing they would after a short while run out of old innovative efforts to patent. The possibility that the patents following acquisition are simply patents applied for pre acquisition that took time to be granted is not a risk, as the dates we have assigned to each patent are for when the patents are initially applied for, not for when they are granted.

Figure 7: Average Patent Citations by Sector: Pre- vs. Post-Acquisition

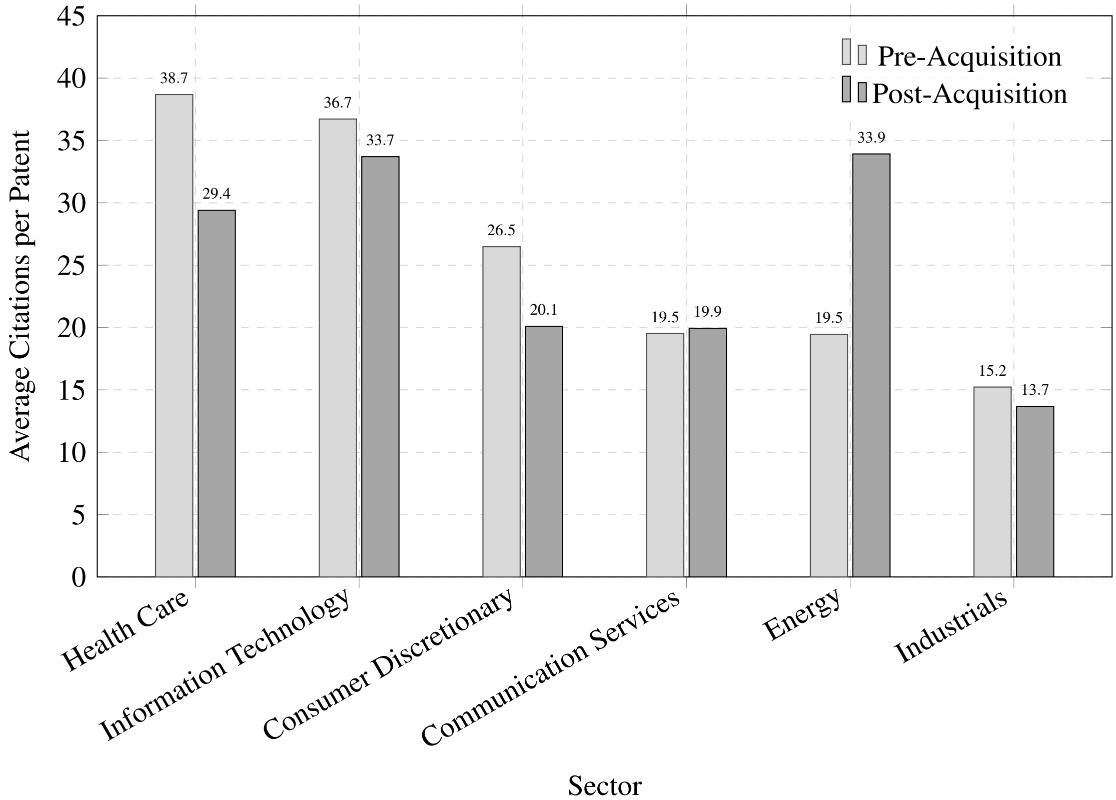

Notes: Average patent citations by sector before and after acquisition. Unlike patent counts, citation quality declined post-acquisition in four of six sectors — Health Care (−24%), Consumer Discretionary (−24%), Industrials (−10%), and Information Technology (−8%). Energy is a striking exception, with average citations rising 74%, while Communication Services remained essentially unchanged.

Average citation amounts per patent don't have a clear increase between pre and post acquisition as patent counts did. The only sector that displays an increase is energy, but the sample size for the energy sector is only 3 which is too small to derive a significant result. Average citation amounts decreasing post acquisition does not have to mean that the Private Equity firms reduce the quality of innovation within their portfolio companies, but can be explained by the fact that older patents have more time to accumulate citations than newer ones. To compare whether Private Equity firms have an effect on innovation quality or not, we'd have to compare a sample of Private Equity acquired companies with a sample of non Private Equity acquired ones. This has already been done by Lerner et al. (2011) and as we are comparing the difference in effect between different Private Equity firms, we do not have to do this in our analysis.

Figure 8: Distribution of PE Deal Transaction Values (N = 167, log scale)

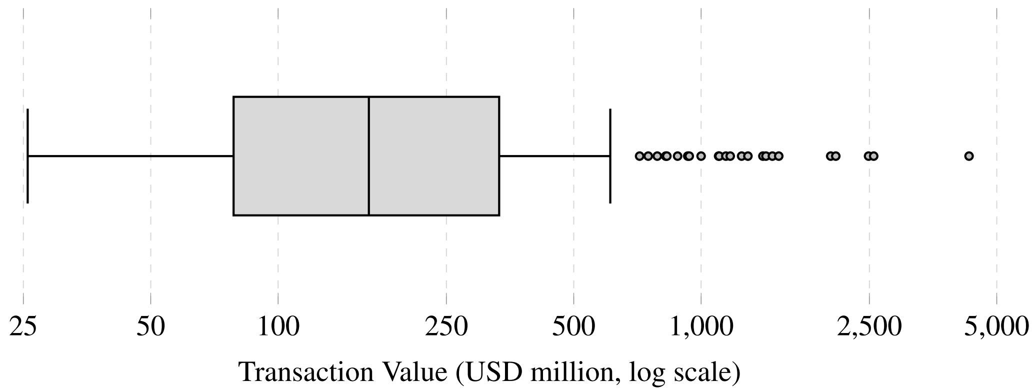

Notes: Box plot of transaction values across the sample (N = 167) on a logarithmic x-axis. The box represents the interquartile range (IQR): Q1 = USD 79M, median = USD 164M, Q3 = USD 333M. Whiskers extend to the most extreme observations within 1.5×IQR of the quartiles. The 24 markers beyond the upper whisker represent outlier deals — transactions of USD 716M or more, including a single transaction of USD 4.3 billion. The log scale visually compresses the heavy right tail and makes the structure of the distribution easier to read.

As seen from the distribution of purchase values of the target companies in our list, there is a lot of variance in the size of the companies that we are analyzing, although most sit around the ∼80 to ∼330 million USD mark.

3 Methodology

3.1 Variable and Outcomes

To measure "innovation" in our portfolio companies, we chose to focus on patent count and patent citation. Patent count refers to the frequency with which a company has filed for patents. Patent citation, however, refers to the number of times a patent has been cited by other patents, which is the standard for measuring innovation quality and economic impact (Squicciarini, Dernis, and Criscuolo, 2013). Patent citation output of a company is also a strong predictor of future earnings (Gu, 2005), which further strengthens its usefulness as a measurement relevant to investors looking for future returns. To prevent analyzing general patenting activity again in our citations analysis, as we've already done so in the patent count tests, we'll be analyzing the change in average citation amount per patent, giving us a measurement of the average quality of a company's patents, unaffected by the quantity of patents they output. In both cases, we measure whether patent count and/or citations change before and after the transaction.

To do this, we had to establish the number of years before and after the acquisition we would consider in our analysis. Patent count generally reflects a strong innovation culture, thus there would be considerable serial correlation between the number of patents issued today and in the future. Furthermore, patent citations similarly exhibit serial correlation, with patents that are cited frequently in the first few years being heavily cited throughout their lifetime. Using a similar analysis timeline as Lerner et al. (2011), we will be analyzing patents 3 years prior to and 3 years post Private Equity acquisition. Having too short of a timeframe would risk not capturing the effect of the Private Equity firms on patenting activity, as changes in innovative behavior and measurable outcomes can take time. Having too long of a timeframe would add the risk of us analyzing patenting activity occurring after some companies have already been sold off by their respective Private Equity firms, as some of the Private Equity firms in our data only hold their companies for about 3 years.

A central issue we tackled was the level of innovation to consider, whether to include all patent applications or just the ones that were granted. Many patents receive citations before they are ever granted, and many granted patents don't receive citations. Moreover, it is highly relevant when a firm increases the number of patent applications for our study even if those applications don't result in a grant. Thus we decided to delineate between innovation activity (a broader scope of analysis using the amount of patent applications and the subsequent citations on the patents-in-progress) and innovation outcome (a narrower scope of analysis using the amount of patents granted and the citations on those patents specifically). Since we had to remove patents that hadn't been granted for our innovation outcome tests, we ended up with 157 companies with granted patent data, compared to the 167 companies we had for the innovation activity tests.

A natural concern regarding our choice to anoint activity with the same or similar legitimacy as outcome would be the economic value of ungranted patents. Farre-Mensa, Hegde and Ljungqvist (2020) argue that patent applications or ungranted patents still have economic value because a patent application is often used as collateral for bank loans, and that ungranted patents still become prior art (evidence that an invention is already known or publicly available) and describe the holder's intellectual assets. Thus, regardless of whether patents are eventually granted, they still serve as a defensible measure of innovation output.

For the innovation quality tests we had to conduct additional filtering of the data. Since we are measuring average patent citation amounts, we couldn't include companies that had 0 patents either before or after acquisition, as we can't calculate average strength of patents if there are no patents. Including these cases would inappropriately affect the data, as companies that go from 0 patents to patents or vice versa, would appear to have a very strong change in patent quality. To counter this issue, we only included companies that had at least one patent before acquisition and at least one patent after acquisition. For the innovation activity tests we went from 167 companies to 98 companies, and for the innovation outcome tests we went from 157 to 84 companies.

Using both of these dichotomies, activity/outcome and quantity/quality, we get four different outcome variables. For innovation activity, post acquisition patent count (granted + ungranted) and average forward citations per patent (granted + ungranted) are the variables we are measuring. For innovation outcome, we measure post acquisition granted patent count and average forward citations per granted patent. Throughout all of these models, we use a control variable (the pre-acquisition equivalent) for these outcome variables. We will expand on how these 4 outcome variables are designed into their regression specifications in the section below.

3.2 Empirical Specification

Each of the four outcome variables defined in the section above enters a parallel regression. Counts (patent counts) use negative binomial regression; per-patent averages (citations) use OLS. The estimating equations are:

$$\mathbb{E}[Post\_Count_i|X_i] = \exp(\alpha + \beta(PE\_Time\_Horizon_i) + \gamma(Pre\_Count_i) + \delta(Z_i)) \qquad (1)$$

estimated by negative binomial maximum likelihood, and

$$Post\_Citation\_AVG_i = \alpha + \beta(PE\_Time\_Horizon_i) + \gamma(Pre\_Citation\_AVG_i) + \delta(Z_i) + \varepsilon_i \quad (2)$$

estimated by OLS. Z_i includes industry fixed effects and (for the count model) deal value. The focal coefficient β is identified from variation in Private Equity holding period across portfolio companies. The error term contains everything correlated with both holding period and innovation outcomes that varies across deals, most importantly, unobserved firm quality at the time of acquisition, which Private Equity firms may select. This unobserved firm quality remains one of the main limitations of our analysis.

The choice of negative binomial for our count model follows Lerner et al. (2011). Patent counts are heavily overdispersed: a small number of portfolio companies produce hundreds of patents while most produce none (or close to none), so the variance far exceeds the mean. Poisson is inappropriate given the variance-to-mean ratio in our sample.

Table 1: Descriptive Statistics: VMR on Patent Counts and Granted Patents

Variable	**Mean**	**Variance**	**VMR**	**Zeros**	**Zero %**
Panel A: Activity — Full sample (N = 167)					
Pre Count	8.5808	191.9076	22.3647	41	24.6%
Post Count	11.9820	338.5238	28.2526	28	16.8%
Panel B: Activity — Zero-patent firms excluded (N = 98)					
Pre Count	12.9694	265.7207	20.4883	0	0.0%
Post Count	18.3673	450.8121	24.5442	0	0.0%
Panel C: Outcome (granted patents) — Full sample (N = 157)					
Pre Patent Count	6.6369	106.8097	16.0932	38	24.2%
Post Patent Count	7.9936	161.0833	20.1515	35	22.3%
Panel D: Outcome (granted patents) — Zero-patent firms excluded (N = 84)					
Pre Patent Count	10.5238	158.6380	15.0742	0	0.0%
Post Patent Count	13.4762	229.1922	17.0072	0	0.0%

Notes: VMR = Variance-to-Mean Ratio. A VMR substantially greater than 1 indicates overdispersion, suggesting a negative binomial model is more appropriate than Poisson. Panels A and B refer to the Activity dataset (patent applications); Panels C and D refer to the Outcome dataset (granted patents). Each "zero-patent firms excluded" panel drops firms with zero patents in either the pre- or post-acquisition period.

In order to diagnose overdispersion in a dataset, the VMR must be greater than 1. As displayed in the table above, the variance is approximately 15-28 times greater than the mean in our datasets. Negative binomial generalizes Poisson by allowing the conditional variance to exceed the conditional mean, with dispersion parameter alpha estimated from the data. Lerner et al. (2011) work at the patent level and estimate negative binomial models for citation counts directly, observing at the individual patent level. Our dataset is constructed at the portfolio-company level, with the citation outcome organized as average forward citations per patent. Thus the citations variable becomes a continuous and non-negative ratio rather than a count, with OLS becoming the proper method of modeling. The issue with this approach is that a few outliers (foundational patents) may have an outsized effect on the end result. We have Winsorized our Post Citation AVG variable at the 99th percentile to eliminate the leverage that data points existing several standard deviations above the sample mean may have on the tests.

We included sector fixed effects to account for systematic differences in innovation activity between disparate sectors. We have restricted our sample to include only the industries that had sufficient patent activity for estimation: Consumer Discretionary, Information Technology, Industrials, Health Care, Energy and Communication Services.

When beginning to run tests, we took some statistical measures to ensure the explanatory value of our models. We clustered our standard errors at the PE firm level. Multiple portfolio companies sharing the same acquiring firm and thus fund-level strategy would remain unobserved and likely correlated if not for the clustering. Ignoring this would overstate the effective sample size and likely produce standard errors that are too small.

The Pre_Count and Pre_Citation_AVG variables make sure that we account for the patenting volume and strength of companies prior to acquisition, as the independent variables are measured in pure amount differences, not in ratios. Companies with higher patenting volume and strength prior to acquisition should have a higher difference between pre and post acquisition, all else equal, so this makes sure that the time horizon variable isn't affected by this.

Table 2: Sample Size: 4 Outcome Variables

	Patent Count	Average Patent Citations
All	167	98
Granted	157	84

Within each estimation, we ran a test on three different models:

1. M1: only taking into account time horizon and pre-count/pre-citation AVG
2. M2: M1 + industry fixed effects
3. M3: M2 + transaction value + year fixed effects

We report three models for each outcome. The first model (M1) only takes into account the Private Equity time horizons, as well as the pre acquisition patent counts / citations. The second model (M2) also includes industry fixed effects. The third model (M3) takes the second model and adds a control variable for transaction value as well as year fixed effects. Since we don't have that many data points in our tests, adding a dummy variable for each year would significantly decrease the statistical power. To reduce the amount of fixed effects in our models while still accounting for macroeconomic changes between time periods, we've split the years into groups. We have one fixed effect for companies acquired between 1995-2000, one for 2000-2005, one for 2005-2010, and one for 2010-2015. The last year in each period is viewed as "less than" in the calculations, so there are no overlaps between any of the periods.

4 Analysis

4.1 Measuring Innovation Quantity

4.1.1 Quantity of Innovation Activity

Table 3: PE Time Horizon Coefficient Estimates: Innovation Activity (Patent Count)

Statistic	**M1A**	**M2A**	**M3A**
Coefficient (β)	0.0846	0.0668	0.0164
Std. Error (clustered)	0.1011	0.1021	0.1067
z-statistic	0.8372	0.6545	0.1539
p-value (clustered)	0.4025	0.5128	0.8777
95% CI lower	-0.1135	-0.1332	-0.1926
95% CI upper	0.2827	0.2668	0.2255
IRR = $\exp(\beta)$	1.0883	1.0691	1.0166
IRR CI lower	0.8927	0.8753	0.8248
IRR CI upper	1.3267	1.3058	1.2529

Notes: Standard errors clustered at the PE firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. None of the coefficient estimates reach conventional levels of statistical significance.

The Private Equity Time Horizon coefficients are positive in all three models so there is some positive association between a longer time horizon and more innovation activity. A one-year increase in Private Equity holding period is indicative of an estimated 1.6-8.5% increase in expected post-acquisition patent applications. However, none of the coefficients pass the statistically significant threshold. In fact, our p-value indicates that in M2 and M3, the point estimate is more consistent with the null hypothesis than the alternative.

Table 4: Regression Results (Extended): Innovation Activity (Patent Count)

	M1A				M2A				M3A			
Variable	Coef	(SE)	z	[IRR]	Coef	(SE)	z	[IRR]	Coef	(SE)	z	[IRR]
const	1.3555**	(0.5464)	2.481	[3.8789]	1.2793**	(0.5774)	2.216	[3.5941]	1.6810***	(0.6462)	2.602	[5.3711]
PE_Time_Horizon	0.0846	(0.1011)	0.837	[1.0883]	0.0668	(0.1021)	0.655	[1.0691]	0.0164	(0.1067)	0.154	[1.0166]
Pre_Count	0.0499***	(0.0079)	6.352	[1.0512]	0.0481***	(0.0068)	7.054	[1.0492]	0.0497***	(0.0079)	6.281	[1.0509]
Transaction_Value	—	—	—	—	—	—	—	—	-0.0003	(0.0006)	-0.423	[0.9997]
Consumer_Discretionary	—	—	—	—	0.2847	(0.3081)	0.924	[1.3294]	0.4575	(0.3137)	1.458	[1.5801]
Information_Technology	—	—	—	—	0.3432	(0.3036)	1.131	[1.4095]	0.5062	(0.3201)	1.581	[1.6590]
Health_Care	—	—	—	—	0.1052	(0.3600)	0.292	[1.1110]	0.1031	(0.3494)	0.295	[1.1087]
Energy	—	—	—	—	1.2197*	(0.7221)	1.689	[3.3860]	1.7152*	(0.9690)	1.770	[5.5580]
Communication_Services	—	—	—	—	-0.5650	(0.3985)	-1.418	[0.5684]	-0.1153	(0.3884)	-0.297	[0.8911]
yr_2000_2005	—	—	—	—	—	—	—	—	-0.2062	(0.2838)	-0.727	[0.8137]
yr_2005_2010	—	—	—	—	—	—	—	—	-0.5971*	(0.3498)	-1.707	[0.5504]
yr_2010_2015	—	—	—	—	—	—	—	—	-0.1446	(0.3837)	-0.377	[0.8654]
alpha	1.5503***	(0.1836)	8.445	[4.7131]	1.4781***	(0.1743)	8.480	[4.3846]	1.4856***	(0.1952)	7.610	[4.4177]

Notes: Cluster-robust standard errors by PE firm in parentheses. IRR = exp(coef). Reference industry: Industrials. Reference year block: yr_1995_2000. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

To get a clearer view on how the mechanisms within the test are working relative to each other, we can turn to our extended results in Table 4. In it we see that the patent count prior to acquisition is significant at the 1% level in all three models. Effectively, the amount of patent applications the firm files prior to the acquisition is more indicative regarding the patent applications filed after the acquisition than all other factors. This confirms that there is strong serial correlation at the firm level when it comes to patenting. Within the industry dummies in M2 and M3, Energy stands out as distinguishable from our reference industry with marginal statistical significance. The same is true of the year block 2005-2010 against our grouped year reference block.

4.1.2 Quantity of Innovation Outcome

Table 5: PE Time Horizon Coefficient Estimates: Innovation Outcome (Patent Count)

Statistic	M1A	M2A	M3A
Coefficient (β)	0.1549	0.1417	0.1309
Std. Error (clustered)	0.1104	0.0999	0.1017
z-statistic	1.4026	1.4185	1.2871
p-value (clustered)	0.1607	0.1561	0.1980
95% CI lower	-0.0615	-0.0541	-0.0684
95% CI upper	0.3712	0.3375	0.3303
IRR = exp(β)	1.1675	1.1523	1.1399
IRR CI lower	0.9403	0.9473	0.9338
IRR CI upper	1.4495	1.4015	1.3914

Notes: Standard errors clustered at the PE firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. None of the coefficient estimates reach conventional levels of statistical significance.

The Private Equity Time Horizon coefficients are once again positive but insignificant. However, our estimates here are significantly larger than in our previous test: a one-year increase in Private Equity holding period is associated with an estimated 13.1-15.5% increase in post-acquisition grants. This indicates that while there is still no definitive association between time horizon and patent count, there is a stronger tie between holding period and patent grants than holding period and applications alone.

Table 6: Regression Results (Extended): Innovation Outcome (Patent Count)

	M1A				M2A				M3A			
Variable	Coef	(SE)	z	[IRR]	Coef	(SE)	z	[IRR]	Coef	(SE)	z	[IRR]
const	0.6372	(0.5896)	1.081	[1.8912]	0.4299	(0.5587)	0.770	[1.5371]	0.5694	(0.5844)	0.974	[1.7673]
PE_Time_Horizon	0.1548	(0.1104)	1.403	[1.1675]	0.1417	(0.0999)	1.418	[1.1523]	0.1309	(0.1017)	1.287	[1.1399]
Pre_Count	0.0608^{***}	(0.0095)	6.434	[1.0627]	0.0616^{***}	(0.0088)	7.007	[1.0635]	0.0608^{***}	(0.0093)	6.527	[1.0626]
Transaction_Value	—	—	—	—	—	—	—	—	-0.0002	(0.0002)	-0.999	[0.9998]
Consumer_Discretionary	—	—	—	—	0.5851^{*}	(0.3051)	1.918	[1.7951]	0.6875^{**}	(0.3078)	2.233	[1.9887]
Information_Technology	—	—	—	—	0.0669	(0.2934)	0.228	[1.0692]	0.1531	(0.3010)	0.509	[1.1655]
Health_Care	—	—	—	—	0.3838	(0.3780)	1.015	[1.4679]	0.4450	(0.3960)	1.124	[1.5604]
Energy	—	—	—	—	1.3690^{**}	(0.6903)	1.983	[3.9313]	1.7485^{**}	(0.8479)	2.062	[5.7460]
Communication_Services	—	—	—	—	-0.4127	(0.4666)	-0.884	[0.6619]	-0.1549	(0.4807)	-0.322	[0.8565]
yr_2000_2005	—	—	—	—	—	—	—	—	0.1095	(0.2771)	0.395	[1.1157]
yr_2005_2010	—	—	—	—	—	—	—	—	-0.4782	(0.3051)	-1.567	[0.6199]
yr_2010_2015	—	—	—	—	—	—	—	—	-0.1212	(0.2778)	-0.436	[0.8859]
alpha	1.4569^{***}	(0.2002)	7.277	[4.2925]	1.3253^{***}	(0.1694)	7.825	[3.7632]	1.2711^{***}	(0.1761)	7.218	[3.5647]

Notes: Cluster-robust standard errors by PE firm in parentheses. IRR = exp(coef). Reference industry: Industrials. Reference year block: yr_1995_2000. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In Table 6, we see once again that Pre_Count is highly significant, but we also see a change in the industry FEs: energy reaches significance at the 5% level in M2 and M3, and so does consumer discretionary, if only marginally in M2. This means that energy-sector portfolio companies grant nearly four times the patents of industrial companies (our reference industry). The larger effect of our industry dummies in the grant regression (relative to the application regression) suggests differences in industry have larger total effects on grants than on just applications. This implies some kind of sector-specific difference in grant rates.

4.2 Measuring Innovation Quality

Now we move into our analysis on patent citations, marking a shift in our analysis from frequency in innovative behavior to the quality of that behavior.

4.2.1 Quality of Innovation Activity

In this section, we look into how Private Equity Time Horizon affects the citations on all patents (those that are eventually granted and those that were not).

Table 7: PE Time Horizon Coefficient Estimates: Innovation Activity (Citations)

Statistic	**M1B**	**M2B**	**M3B**
Coefficient (β)	4.1117	3.7795	3.5435
Std. Error (clustered)	2.2792	2.2748	2.1273
t-statistic	1.8040	1.6615	1.6657
p-value (clustered)	0.0712*	0.0966*	0.0958*
95% CI lower	−0.3555	−0.6789	−0.6260
95% CI upper	8.5789	8.2379	7.7130

Notes: Standard errors clustered at the PE firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. PE_Time_Horizon is marginally significant at the 10% level across all three specifications.

In Table 7, all three Time Horizon coefficients are positive, even reaching marginal significance. This is a meaningful increase in statistical significance. While this is not quite definitive proof of an association, we can preliminarily report a stronger link between time horizon and quality of innovation than between time horizon and quantity of innovation. At this level, each additional year in time horizon equates to an estimated 3.5-4.1 increase in average citations.

Table 8: Regression Results (Extended): Innovation Activity (Citations)

	M1B			**M2B**			**M3B**		
Variable	Coef	(SE)	t	Coef	(SE)	t	Coef	(SE)	t
const	-2.8327	(11.5191)	-0.246	-4.7702	(12.4845)	-0.382	3.9074	(12.9552)	0.302
PE_Time_Horizon	4.1117*	(2.2792)	1.804	3.7795*	(2.2748)	1.661	3.5435*	(2.1273)	1.666
Pre_Citation_AVG	0.2191***	(0.0773)	2.832	0.2077**	(0.0824)	2.521	0.2375***	(0.0852)	2.787
Transaction_Value	—	—	—	—	—	—	-0.0047	(0.0038)	-1.260
Consumer_Discretionary	—	—	—	4.8728	(4.9547)	0.983	4.2407	(5.3811)	0.788
Information_Technology	—	—	—	5.7795	(5.1546)	1.121	7.3568	(5.0586)	1.454
Health_Care	—	—	—	6.1244	(6.1788)	0.991	-0.1746	(8.4017)	-0.021
Energy	—	—	—	3.7119	(4.9764)	0.746	1.1429	(5.6085)	0.204
Communication_Services	—	—	—	23.9812	(24.0581)	0.997	30.8040	(25.7552)	1.196
yr_2000_2005	—	—	—	—	—	—	-8.3296	(5.7482)	-1.449
yr_2005_2010	—	—	—	—	—	—	-9.5221*	(5.5431)	-1.718
yr_2010_2015	—	—	—	—	—	—	-12.3244*	(7.2833)	-1.692

Notes: OLS regression with cluster-robust standard errors by PE firm in parentheses. Reference industry: Industrials. Reference year block: yr_1995_2000. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Taking a deeper look into our test on Time Horizon and quality on innovative activity, we can report that the relationship of a firm's pre-acquisition patent citations on post-acquisition citations is statistically significant in all three models, with this relationship being significant at the 1% level in M3. This points to citation quality, like patent volume, being serially correlated within the firm. Unlike previous models on patent counts, none of the industry dummies are significantly distinguishable from our reference industry. The same is true of the grouped year dummies, aside from a marginally significant result for 2005-2010. This indicates that patent citations were more or less a similar challenge between industries and year groups.

4.2.2 Quality of Innovation Outcome

In this section, we look into how Private Equity Time Horizon affects the citations on granted patents.

Table 9: PE Time Horizon Coefficient Estimates: Innovation Outcome (Citations)

Statistic	**M1B**	**M2B**	**M3B**
Coefficient (β)	3.4550	3.2514	2.6429
Std. Error (clustered)	2.2620	2.2289	2.0401
t-statistic	1.5274	1.4587	1.2955
p-value (clustered)	0.1267	0.1446	0.1952
95% CI lower	-0.9783	-1.1173	-1.3557
95% CI upper	7.8884	7.6201	6.6414

Notes: Standard errors clustered at the PE firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. None of the coefficient estimates reach conventional levels of statistical significance.

Table 9 informs our understanding of the focal relation further. While Time Horizon coefficients are positive, it sits just outside the realm of marginal significance. This, in tandem with the findings from our previous discussion on quality of innovation activity implies that many of the patents that were not ultimately granted had a meaningful amount of citations. Interestingly enough, this finding indicates that the ungranted patents had higher average citations than those that ended up getting granted.

Table 10: Regression Results (Extended): Innovation Outcome (Citations)

	M1B			**M2B**			**M3B**		
Variable	Coef	(SE)	t	Coef	(SE)	t	Coef	(SE)	t
const	-3.7982	(11.6908)	-0.325	-4.3156	(12.3709)	-0.349	4.6694	(11.7198)	0.398
PE_Time_Horizon	3.4550	(2.2620)	1.527	3.2514	(2.2289)	1.459	2.6429	(2.0401)	1.295
Pre_Citation_AVG	0.4216^{**}	(0.1738)	2.426	0.3928^{**}	(0.1867)	2.105	0.5377^{***}	(0.1719)	3.127
Transaction_Value	—	—	—	—	—	—	-0.0157^{***}	(0.0049)	-3.202
Consumer_Discretionary	—	—	—	2.5003	(5.5026)	0.454	1.4530	(6.3246)	0.230
Information_Technology	—	—	—	6.8383	(7.0061)	0.976	6.8861	(5.4225)	1.270
Health_Care	—	—	—	2.1422	(6.3337)	0.338	-5.1532	(7.8501)	-0.656
Energy	—	—	—	-1.5014	(6.5205)	-0.230	-3.7528	(7.5494)	-0.497
yr_2000_2005	—	—	—	—	—	—	-5.2627	(5.3881)	-0.977
yr_2005_2010	—	—	—	—	—	—	-10.1914^{*}	(6.1854)	-1.648
yr_2010_2015	—	—	—	—	—	—	-6.9497	(4.9711)	-1.398

Notes: OLS regression with cluster-robust standard errors by Private Equity firm in parentheses. Reference industry: Industrials. Reference year block: yr_1995_2000. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In Table 10, we see that, for the first time in our analyses, transaction value is statistically significant. One might expect that the transaction value would have a positive correlation with patents and likely citations (argument being that large firms have more disposable capital and talent with which to innovate at a higher level than smaller firms). However, our findings suggest the opposite, that insofar as citations on granted patents are concerned, transaction value has a decidedly negative effect. This makes sense if you view the primary innovation vehicle of larger firms to be acquisitions. Since transaction value has been negative throughout all of our analyses (though typically it has been statistically indistinguishable from zero), the acquisition-first view of larger enterprise innovation strategy makes sense, however this relation seems specifically true in citations on granted patents.

5 Robustness Checks

5.1 Standard Errors

Table 11: Robustness Check: Activity Quality (Clustered SEs vs. HC3)

Model	**SE Method**	p **(clustered)**	p **(HC3)**	**Significance Change?**
M1B	cluster-robust	0.0712*	0.1121	Yes ($^* \rightarrow$ n.s.)
M2B	cluster-robust	0.0966*	0.1408	Yes ($^* \rightarrow$ n.s.)
M3B	cluster-robust	0.0958*	0.1423	Yes ($^* \rightarrow$ n.s.)

Notes: Sensitivity check on PE_Time_Horizon's statistical significance under alternative standard error specifications. Cluster-robust standard errors are clustered at the Private Equity firm level (the preferred specification given firm-level nesting). HC3 standard errors are heteroscedasticity-consistent but do not account for within-firm clustering. In all three model specifications, the marginal significance under cluster-robust SEs (* at $p < 0.10$) does not survive under HC3 SEs. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 12: Robustness Check: Outcome Quality (Clustered SEs vs. HC3)

Model	**SE Method**	p **(clustered)**	p **(HC3)**	**Significance Change?**
M1B	cluster-robust	0.1267	0.2024	No (n.s. under both)
M2B	cluster-robust	0.1446	0.2250	No (n.s. under both)
M3B	cluster-robust	0.1952	0.2935	No (n.s. under both)

Notes: Sensitivity check on PE_Time_Horizon's statistical significance in the winsorized sample under alternative standard error specifications. Cluster-robust standard errors are clustered at the Private Equity firm level (the preferred specification given firm-level nesting). HC3 standard errors are heteroscedasticity-consistent but do not account for within-firm clustering. In contrast to the unwinsorized sample (Table 11), the coefficient on PE_Time_Horizon is not statistically significant under either standard error method in any of the three specifications.

In Table 11, we can see that going from clustered standard errors to HC3 standard errors reduces the statistical significance of our findings to be just outside the bounds of marginal significance. The trend is the same in Table 12 but as the results were not marginally significant in the first place, the change is less striking. However, this trend implies one of two things:

1. The marginal significance under clustered SEs was just an artifact of cluster correction, not a feature of the data. The cluster correction effectively estimates a smaller standard

error because it accounts for negative within-firm correlation. However, in a sample size of 98 firms, it's difficult to say that the cluster correction is itself estimated precisely.

2. The cluster robust SE approach is theoretically the correct one. We know that there is likely within Private Equity firm dependence because the same Private Equity firms can undergo multiple transactions each within our dataset. Since the HC3 structure ignores this, it is inherently irrelevant for our data.

We believe that the latter is true. Cameron and Miller (2015) make the explicit warning regarding cluster counts from less than 20 to less than 50, which puts our 98 in a "safe" range. However, it is still a gray zone and left up to the sensibilities of the reader and of future research to decide conclusively.

5.2 Sensitivity to Winsorization

One clear robustness check we wanted to do is to review the effect that winsorizing our citation data at the 99th percentile has on the results of our test. To do this, we ran an analogous regression without winsorizing our data and these were our findings:

Table 13: Robustness Check: Activity Quality without Winsorization

Model	**Coefficient (β)**	**SE**	t	p **(clustered)**	p **(HC3)**
M1B	4.1657	2.3187	1.7966	0.0724*	0.1126
M2B	3.8325	2.3132	1.6568	0.0976*	0.1411
M3B	3.5936	2.1624	1.6619	0.0965*	0.1424

Notes: Robustness check using the unwinsorized sample. Cluster-robust standard errors at the Private Equity firm level; HC3 standard errors are heteroscedasticity-consistent. PE_Time_Horizon retains marginal significance under clustered SEs but loses significance under HC3 SEs in all three specifications — a pattern that suggests the marginal significance is partly dependent on the SE method used. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 14: Robustness Check: Outcome Quality without Winsorization

Model	**Coefficient (β)**	**SE**	t	p **(clustered)**	p **(HC3)**
M1B	3.5018	2.2937	1.5268	0.1268	0.2011
M2B	3.2966	2.2600	1.4587	0.1447	0.2236
M3B	2.6835	2.0672	1.2981	0.1942	0.2908

Notes: Robustness check using the unwinsorized sample. Cluster-robust standard errors at the Private Equity firm level; HC3 standard errors are heteroscedasticity-consistent. None of the coefficient estimates reach conventional levels of statistical significance under either standard error method. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

In the tables above, you can see that our p-values with clustered standard errors remain virtually the same as in our original regression, indicating that winsorizing our data had little to no effect on the descriptive power of the tests. This goes for both clustered and HC3 models. This means that the marginal significance was not driven by tail observations. When taking a closer look, we found that only one option in the Activity Quality dataset was eventually winsorized.

5.3 Poisson Test

As a robustness check on the count regressions, we re-estimate our patent counts using a Poisson regression.

Table 15: Robustness Check: Poisson Test on Activity Quantity

Statistic	**M1A**	**M2A**	**M3A**
Coefficient (β)	0.1557	0.1780	0.1611
Std. Error (clustered)	0.1042	0.1092	0.1170
z-statistic	1.4947	1.6304	1.3764
p-value (clustered)	0.1350	0.1030	0.1687
IRR = $\exp(\beta)$	1.1685	1.1948	1.1748
95% CI IRR lower	0.9527	0.9647	0.9340
95% CI IRR upper	1.4332	1.4800	1.4777

Notes: Standard errors clustered at the PE firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. None of the coefficient estimates reach conventional levels of statistical significance.

Table 16: Robustness Check: Poisson Test on Outcome Quantity

Statistic	**M1A**	**M2A**	**M3A**
Coefficient (β)	0.2328	0.2409	0.2407
Std. Error (clustered)	0.1198	0.1113	0.1145
z-statistic	1.9426	2.1639	2.1028
p-value (clustered)	0.0521^{*}	0.0305^{**}	0.0355^{**}
IRR = $\exp(\beta)$	1.2621	1.2724	1.2721
95% CI IRR lower	0.9979	1.0230	1.0165
95% CI IRR upper	1.5962	1.5826	1.5921

Notes: Standard errors clustered at the PE firm level. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. PE_Time_Horizon is significant at the 5% level in both M2A and M3A.

From our results, we gather that the two emerging patterns both point to Negative Binomial being the correct estimator for our data. First, the Poisson point estimates are declaratively larger than those in the Negative Binomial. This is largely due to the sensitivity of the Poisson model to overdispersion. The Negative Binomial model's variance function appropriately weights extreme data points whereas Poisson tends to underweight them, allowing the high count observations to dominate the likelihood, resulting in an outcome quantity p-value that is between 3-5%, a shocking distinction from the 16-19% with the Negative Binomial model. The significant results displayed in our outcome quantity model is not evidence of a robust effect, but an artifact of an irrelevant model. Thus, this robustness check serves as a confirmation of our methodological choice.

6 Broader Discussion

We initially hypothesized that longer Private Equity holding periods are associated with greater innovation in portfolio companies. The logic is essentially that a fund holding assets for longer faces a higher opportunity cost of capital, so they have to generate real operational value to justify the extended horizon. Short term financial engineering or unsustainable cost-cutting are generally sufficient to boost returns over a two-year deal, but they are unlikely to sustain robust performance over five or more years. Thus, long-term funds have stronger incentives to invest in innovation, which provides a competitive moat for portfolio companies but takes time to pay off. As mentioned in our hypothesis section, this is in line with the work done by Manso (2011). He theorized that innovation flourishes under conditions that tolerate early failures and reward long-term success. He presented innovation as a real option and the tolerance for early failures from long-term investors gives them the right to exercise that option. Short-horizon owners can push their companies to file patents but they avoid underwriting slow, uncertain bets that result in groundbreaking work that future inventors build on.

Across the four outcomes in our analysis so far, application volume, grant volume, application citation quality, grant citation quality, the coefficient on Private Equity Time Horizon is positive in every model. None of the coefficients reach significance, with only the activity-citation regression crosses the 10% threshold, and the outcome-citation regression sits just short of it. This means that longer Private Equity holding periods are associated with more, and better, portfolio-company innovation, but the only dimension where the effect is statistically distinguishable from zero is citation quality of inventive activity.

The natural reading is that Private Equity holding period shapes innovation quality more than quantity, a finding that is heavily substantiated by Lerner et al. (2011)'s work. Longer-horizon Private Equity firms are not associated with portfolio companies that produce noticeably more patents, applications or grants, but they are associated with portfolio companies whose patents prove more impactful, accumulating roughly three to four additional citations per patent for each additional year of Private Equity ownership. The effect is detected with greatest precision in the application-citation regression, where the sample is largest and the citation distribution is least filtered by the patent approval process. The grant-citation regression points in the same direction with similar magnitude but loses statistical power on the smaller granted-patent sample, which negates citations of patents pending approval (where they end up getting rejected).

Where most associations between Private Equity and Innovation usually run into the same dichotomy - the long-termism or the short-termism hypothesis, this pattern aligns more closely with the long-termism hypothesis, but in a specific and limited way. Long-termism, in our analysis, does not raise the rate of inventive output; portfolio companies with long-horizon owners do not patent more frequently than those with short-horizon owners. What long-termism seems to enable is depth: longer ownership horizons appear to be associated with portfolio companies innovating in such a way that the inventions value manifests later, in the form of citations from subsequent inventors. Private Equity firms with short time horizons may be pushing portfolio companies toward inventions that pay off within the holding period or create a cosmetic band-aid to increase value for an exit; Private Equity firms with longer horizons may be more willing to back inventions whose impact grows over years.

Bloom, Jones, Reenen and Webb (2020) point out that while historical studies from the 1920s to 1990s used patent count as a reliable, stable proxy for ideas, the conditions allowing for that association have changed. Starting in the 1980s, USPTO patent grants surged faster than ever before, driven by changes in regulation, lowering the bar for a patent to be granted. Guntersdorfer (2003) argued that this trend culminated in the US abandoning both the business method exception (the rule stating that a process or a method cannot be patented) and physical requirements (that in order for a patent to be granted, the invention must have a tangible real world manifestation). As a result, a patent today represents something entirely different when compared to a patent from the 1970s. The dubious nature of a patent grant's relationship to an idea and its quality over time makes patent count a more questionable metric for innovation than we initially considered.

That being said, there are many diagnostic shortcomings to address. First of all, the citation-quality effect is marginal, it crosses the 10% line in two of our four specifications and falls short in the other two. Furthermore, the design of the experiment is inherently conditional, not causal. Private Equity firms that stay longer might just be picking firms whose research happens to mature slowly. Moreover, sector FEs strip out the obvious sector-level confounders but not the unobserved differences between firms that Private Equity firms have access to. And the count results are flatly null across the board. Taking our story at face value means treating the citation result as a real signal and the count result as uninformative, and at our sample size, neither call can be made with confidence.

Our analysis does not settle the long-termism vs short-termism question as it relates to innovation. What it does suggest is that the question may be framed wrong if it's only about volume. If Private Equity holding period matters for portfolio-company innovation at all, our evidence says it shows up in what gets invented, not how much.

7 Fit Tests

The four models fit the data as one might expect, given the reputation of patent data being overdispersed and serially correlated. The dispersion parameter α is positive and significant at the 1% level in every count regression ($\alpha = 1.27 - 1.55$), meaning patent counts are overdispersed at the firm level and that negative binomial is the right method to use. Pre_Count in the count models and Pre_Citation_AVG in the citation models are highly significant throughout, accounting for inside-firm differences in patenting and citation and building the coefficient on PE_Time_Horizon from the acquisition-specific changes. Likelihood-ratio and F-tests testing the inclusion of industry fixed effects yield mixed results. We find that industry effects significantly improve model fit only for the innovation outcome count model ($p = 0.043$), whereas they fail to provide significant explanatory power over the baseline for the other three outcomes. Furthermore, expanding the models to include transaction value and year fixed effects (M3) provides a statistically significant improvement to the citation outcome model ($p = 0.004$). However, similar improvements in the activity models should be interpreted with caution due to the reduction in sample size caused by missing transaction data. The fit test tables can be found in Appendix D.

8 Conclusion

The goal of this paper was to determine whether the time horizon tendencies of Private Equity firms have an effect on the innovation within their acquired portfolio companies. By using patent counts and citations to quantify innovation output, our analysis revealed that while a Private Equity firm's historical time horizon does not have a noticeable effect on patenting quantity, it does have a meaningful, positive effect on patenting quality. This suggests that Private Equity firms with longer time horizons are more willing to invest the necessary capital and runway to ensure strong, foundational innovative outcomes, rather than simply pushing for a higher volume of less significant patents. Prioritizing high-quality patents is a strong indicator of a focus on long-term value creation over short-term gains, especially since the financial rewards of such innovation take time to materialize, if they are realized at all.

8.1 Implications

The implications of these findings are highly relevant for LP investors choosing where to allocate their capital. LPs prioritizing long-term value and structural growth need empirical ways to evaluate which private equity firms provide the best vehicle for such investments. Furthermore, understanding which ownership structures best promote innovation benefits the target companies themselves. Just as the findings of Lerner et al. (2011) and Bernstein (2015) assist companies in weighing the innovative benefits of private versus public ownership, our findings provide LPs and company founders with a framework for determining which Private Equity profiles are best suited for long-term, innovation-centric goals. Misaligned ownership that pushes for short-term gains can be highly detrimental to a company's long-term trajectory.

8.2 Limitations

Our final sample of 167 portfolio companies (98 for the activity-quality regression and 84 for the outcome-quality regression) limits statistical and descriptive power. With such a limited sample size, we can't definitively determine that the patent count nulls are true nulls or if the marginal significance of the activity-quality regression isn't just noise. Thus, replication with a larger sample size and more statistical power would alleviate this ambiguity.

Additionally, our lack of access to a database where all patent names are standardized and all subsidiaries linked adds limitations to the extent to which we can ensure that all patents filed by our analyzed companies have actually been included in our data. Researchers replicating our tests using a more advanced database (like the Harvard Business School database used by Lerner et al. (2011)) would likely be able to produce more accurate results.

While our sample is reliable to the phenomenon it seeks to describe, US-headquartered firms in patent-intensive industries acquired by PE firms between 1995 and 2015, this also introduces a heavy skew in our data due to the 1990s LBO wave. We recognize an overrepresentation of older companies in our sample, a problem faced by many researchers including Lerner et al. (2011).

Furthermore, the nature of our study is more conditional than causal. Our sector and year FEs can isolate between sector and year differences but are inherently unable to isolate within sector and year distinctions. This goes back to the perennial shortcoming of this study, the selection into treatment. Private Equity firms with longer-time horizons may select companies with some unobserved firm quality that is conducive to slowly maturing innovation cycles.

8.3 Work for Future Researchers

Our research is limited to identifying a correlation between longer Private Equity time horizons and innovation quality. The specific mechanisms through which these longer time horizons convert into stronger innovation are not yet definitively established, as our methodology analyzes innovation outcomes rather than operational inputs. This presents a clear avenue for future research to investigate the specific actions taken by long-term oriented Private Equity firms to drive this innovation, whether through changes to R&D spending, shifting executive compensation structures, or other levers.

Additionally, our analysis is somewhat limited by the availability of data, in order to address this question with definitive statistical power, the sample size available to researchers must be more extensive. In order to correct for the unobserved firm quality conundrum, a future study would have to rely not only on quantitative data, which is fundamentally blind to the nature of this phenomenon. Instead, future researchers would have to explore a qualitative dimension, through which the behavioral and more intangible stewardship qualities of the Private Equity firms may be revealed.

9 References

Barrot, Jean-Noël, 2012, Investor horizon and innovation: Evidence from private equity funds, working paper.

Bernstein, Shai, 2015, Does going public affect innovation?, Journal of Finance 70, 1365-1403.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb, 2020, Are ideas getting harder to find?, American Economic Review 110, 1104-1144.

Cameron, A. Colin, and Douglas L. Miller, 2015, A practitioner's guide to cluster-robust inference, Journal of Human Resources 50, 317-372.

Farre-Mensa, Joan, Deepak Hegde, and Alexander Ljungqvist, 2020, What is a patent worth? Evidence from the U.S. patent "lottery", Journal of Finance 75, 639-682.

Gompers, Paul A., Steven N. Kaplan, and Vladimir Mukharlyamov, 2016, What do private equity firms say they do?, Journal of Financial Economics 121, 449-476.

Gu, Feng, 2005, Innovation, future earnings, and market efficiency, Journal of Accounting, Auditing Finance 20, 385-418.

Guntersdorfer, Michael, 2003, Software patent law: United States and Europe compared, Duke Law Technology Review 2, 1-12.

Guo, Bibo, David Pérez-Castrillo, and Anna Toldrà-Simats, 2019, Firms' innovation strategy under the shadow of analyst coverage, Journal of Financial Economics 131, 456-483.

Higham, Kyle W., and Tohru Yoshioka, 2022, Patent citation generation at the triadic offices: Mechanisms and implications for analysis, SSRN working paper.

Hovakimian, Armen, and Irena Hutton, 2010, Merger-motivated IPOs, Financial Management 39, 1547-1573.

Joenväärä, Juha, Juho Mäkiaho, and Sami Torstila, 2022, Prolonged private equity holding periods: Six years is the new normal, The Journal of Alternative Investments 25, 65-93.

Lerner, Josh, Morten Sørensen, and Per Strömberg, 2011, Private Equity and long-run investment: The case of innovation, Journal of Finance 66, 445-477.

Manso, Gustavo, 2011, Motivating innovation, Journal of Finance 66, 1823-1860.

Ramzan, Muhammad, Wee-Yeap Lau, Abdullah Abdulaziz Bawazir, and Fozi Ali Belhaj, 2025, Does innovation impact the financial performance of high-tech firms during different life cycle stages? Evidence from SP500 firms, working paper.

Squicciarini, Mariagrazia, Hélène Dernis, and Chiara Criscuolo, 2013, Measuring patent quality: Indicators of technological and economic value, OECD Science, Technology and Industry Working Papers 2013/03.

Wang, Daoshun, Xiaohua Jia, Cong Wang, Kan Yang, Sheng Fu, and Minqiang Xu, 2015, Generalized pattern matching string search on encrypted data in cloud systems, IEEE Transactions on Cloud Computing 3, 329-342.

10 Appendix

10.1 Appendix A: Summary Statistics by Sample

Table 17: Summary Statistics by Sample

Variable	**N**	**Mean**	**Median**	**SD**	**Min**	**Max**
Panel A: Activity — Quantity (Raw N = 167)						
PE Time Horizon (years)	167	4.964	4.870	1.084	2.360	7.660
Pre-Acq. Patent Count	167	8.581	3.000	13.853	0.000	87.000
Post-Acq. Patent Count	167	11.982	4.000	18.399	0.000	108.000
Transaction Value (USD m)	151	207.289	127.500	198.869	25.590	937.580
Panel B: Activity — Quality (Raw N = 98)						
PE Time Horizon (years)	98	5.078	4.979	1.034	2.759	7.593
Pre-Acq. Avg. Citations	98	32.586	22.548	40.877	1.000	335.500
Post-Acq. Avg. Citations	98	25.213	19.431	23.302	0.000	127.857
Transaction Value (USD m)	98	399.075	172.766	611.786	26.900	4300.000
Panel C: Outcome — Quantity (Raw N = 167)						
PE Time Horizon (years)	157	4.944	4.873	1.072	2.356	7.658
Pre-Acq. Patent Count	157	6.637	3.000	10.335	0.000	64.000
Post-Acq. Patent Count	157	7.994	3.000	12.692	0.000	86.000
Transaction Value (USD m)	157	362.340	163.925	559.792	25.589	4300.000
Panel D: Outcome — Quality (Raw N = 84)						
PE Time Horizon (years)	84	5.103	4.998	1.085	2.759	7.593
Pre-Acq. Avg. Citations	84	32.413	26.111	27.632	1.000	153.000
Post-Acq. Avg. Citations	84	27.527	20.988	24.870	0.000	127.857
Transaction Value (USD m)	84	333.510	153.000	552.276	26.900	4300.000

Notes: Activity samples include all PE-acquired firms in the dataset (with patent applications as the dependent variable). Outcome samples restrict to firms with eventually-granted patents. Quality panels (B and D) further restrict to firms with non-zero patents in both pre- and post-acquisition periods, allowing for citation averages to be computed. Transaction Value is missing for some Activity-Quantity observations (N = 151).

10.2 Appendix B: AI Disclosure

We used Claude AI (by Anthropic) to assist us with some parts of our thesis. In the data collection process, we used it to match the PATSTAT names with the Capital IQ names. We also used it to help us write the SQL code used to pull data from PATSTAT, the Python Code we used to run the tests, and the Latex code we used to format the tables, graphs, and final thesis document.

10.3 Appendix C: Claude Name Matching Methodology

The following overview of the methodology used by Claude AI to match the PATSTAT names with the Capital IQ names was generated by Claude following its match:

1. Overview and Core Challenge

This appendix outlines the methodology used to match patent assignee names from the EPO PATSTAT database to authoritative company names from S&P Capital IQ. The matching process evaluated approximately 16,000 raw PATSTAT entries against 776 Capital IQ reference names. The primary challenge in this process was overcoming significant search noise within PATSTAT; querying a target company's name often returned thousands of unrelated entities sharing a common keyword (e.g., searching for "Power Holdings" yielded over 6,000 unrelated companies or individuals with the name "Power").

2. Algorithmic Matching Approach

To ensure high data integrity and avoid the false positives common in automated fuzzy matching, a deterministic, three-layer rule-based system was utilized:

- **Layer 1: Exclusion Rules:** Prioritized regular expressions designed to explicitly block known false positives (e.g., similarly named but demonstrably separate entities) before any matching was attempted.
- **Layer 2: Match Rules:** Ordered regular expressions mapping specific PATSTAT patterns to a single Capital IQ entity. These rules were designed to capture the target company, its known legal variants, acquired entities, and verified subsidiaries.
- **Layer 3: Default Rejection:** Any PATSTAT entry that passed through the exclusion rules without triggering a definitive match rule was conservatively classified as a "NO MATCH."

3. Verification and Match Confidence

To construct accurate Match Rules, corporate relationships were verified using SEC Exhibit 21 filings (subsidiary lists), historical M&A data, public corporate profiles (Bloomberg, PitchBook), and patent assignee records. Matches were subsequently categorized by confidence:

- **High Confidence:** Unambiguous relationships, including exact name matches with minor legal suffix variations, confirmed subsidiaries, verified name changes/acquisitions, and obvious typographical errors.
- **Medium Confidence:** Plausible relationships that lacked explicit public verification, such as inferred holding/operating company pairs, unconfirmed subsidiaries, or partial name matches requiring domain knowledge.

4. Rejection Criteria

To minimize false positives, entries were aggressively rejected under the following conditions:

- **Shared Word, Different Company:** Entities sharing a keyword but operating in entirely different industries.
- **Individual Inventors:** Entries representing personal names rather than corporate assignees.
- **Similar Prefix / Generic Terms:** Separate organizations sharing a prefix, or names consisting solely of generic business terms (e.g., "National," "Global," "Power").
- **Historical / Post-Acquisition Operations:** Patent records predating the CapIQ entity's lineage or entities that operated independently post-acquisition.

5. Summary Statistics

The guiding principle throughout the matching process was that false negatives (missing a true match) were preferable to false positives (incorrectly attributing patents to the wrong firm).

When in doubt, entries were rejected.

10.4 Appendix D: Fit Test Tables

Table 18: Model Fit Statistics: Innovation Outcome (Citations, OLS)

Model	SE Method	Clusters	R^2	Adj. R^2	Log-Lik	AIC	BIC	N	K
M1B	cluster-robust	60	0.266	0.248	−375.24	756.48	763.77	84	3
M2B	cluster-robust	60	0.276	0.220	−374.66	763.32	780.34	84	8
M3B	cluster-robust	60	0.411	0.330	−366.01	754.02	780.76	84	12

Notes: Best AIC: M3B. Best BIC: M1B. Nested F-Tests indicate that adding industry fixed effects (M1B vs. M2B) does not significantly improve model fit ($p = 0.898$), whereas adding year blocks and transaction value (M2B vs. M3B) provides a significant improvement ($p = 0.004$).

Table 19: Model Fit Statistics: Innovation Outcome (Patent Count, NegBin)

Model	SE Method	Clusters	Log-Lik	AIC	BIC	N	K	α
M1A	cluster-robust	93	−450.45	908.90	921.13	157	3	1.457
M2A	cluster-robust	93	−444.71	907.43	934.93	157	8	1.325
M3A	cluster-robust	93	−441.99	909.98	949.72	157	12	1.271

Notes: Best AIC: M2A. Best BIC: M1A. All models successfully converged. Likelihood Ratio Tests indicate that adding industry fixed effects (M1A vs. M2A) significantly improves fit ($p = 0.043$), but subsequently adding year blocks and transaction value (M2A vs. M3A) does not yield a significant improvement ($p = 0.245$).

Table 20: Model Fit Statistics: Innovation Activity (Citations, OLS)

Model	SE Method	Clusters	R^2	Adj. R^2	Log-Lik	AIC	BIC	N	K
M1B	cluster-robust	70	0.192	0.175	−436.11	878.23	885.98	98	3
M2B	cluster-robust	70	0.228	0.168	−433.91	883.82	904.50	98	8
M3B	cluster-robust	70	0.288	0.196	−429.98	883.95	914.97	98	12

Notes: Best AIC: M1B. Best BIC: M1B. Nested F-Tests indicate that adding industry fixed effects (M1B vs. M2B) does not significantly improve fit ($p = 0.533$). The nested F-Test for M2B vs. M3B is omitted due to a reduction in sample size resulting from missing transaction value data.

Table 21: Model Fit Statistics: Innovation Activity (Patent Count, NegBin)

Model	SE Method	Clusters	Log-Lik	AIC	BIC	N	K	α
M1A	cluster-robust	97	−538.14	1084.29	1096.76	167	3	1.550
M2A	cluster-robust	97	−534.65	1087.29	1115.36	167	8	1.478
M3A	cluster-robust	92	−483.08	992.16	1031.38	151	12	1.486

Notes: Best AIC: M3A. Best BIC: M3A. All models successfully converged. Likelihood Ratio Tests indicate that adding industry fixed effects (M1A vs. M2A) does not significantly improve fit ($p = 0.221$). The statistically significant improvement from adding year blocks and transaction value (M2A vs. M3A, $p < 0.001$) should be interpreted cautiously due to the reduction in the sample size (from N=167 to N=151).