

How do cyber attacks impact corporate balance sheets and income statements? An empirical study

Teyyub Mutallimov* Dana Etgar Itzhaki[†] Tyler Moore[‡]

May 12, 2026

Abstract

Researchers have long sought to quantify what impact, if any, cyber attacks have on firms. The most common approach has been to conduct event studies using the stock prices of affected firms, which have found conflicting evidence of small negative, flat, or even positive impacts. Stock prices are not an ideal measure, however, as many other factors can influence their prices such as investor sentiment. In this paper, we study the economic consequences of cyber attacks by linking external incident records to a firm–quarter panel of SEC financial filings for 2014–2025. To estimate their impact, we implement a dynamic event-study difference-in-differences design that traces outcomes from four quarters before to four quarters after Incident, consistent with recent econometric work on staggered treatment timing. On average, cyber attacks induce persistent increases in operating and capital expenditures, financed through higher liabilities, and lead to sustained earnings declines. Disaggregating by incident type reveals sharp heterogeneity: Confidentiality breaches trigger a financing and investment response, with firms expanding their balance sheets to fund remediation and system upgrades. In contrast, ransomware attacks act as availability shocks that cause operational disruption and earnings losses without inducing capital restructuring. These findings show that cyber risk imposes material and persistent financial costs that depend on attack mechanism, and that these costs are only partially reflected in qualitative corporate disclosures.

Keywords: Economics of cybersecurity, Financial statement , SEC disclosures, Disclosure gap, Dynamic Difference-in-Differences, Data breaches vs. Ransomware.

*School of Cyber Studies, College of Engineering & Computer Science, The University of Tulsa. Email: tem9735@utulsa.edu

[†]Berglas School of Economics, Tel Aviv University. Email: etgaritzhki@mail.tau.ac.il

[‡]School of Cyber Studies, College of Engineering & Computer Science, University of Tulsa. Email: tyler-moore@utulsa.edu

1 Introduction

Cyber incidents represent a significant threat to modern firms, with annual losses now exceeding \$16 billion (Federal Bureau of Investigation (FBI), 2025), yet the financial impact of these events remains poorly understood. This gap is notable since, unlike many other operational risks, cyber risk is characterized by both high frequency and high heterogeneity across attack mechanisms. Publicly traded companies in the United States have long been required to disclose material cybersecurity incidents in filings with the Securities and Exchange Commission (SEC), following interpretive guidance issued in 2011 and expanded in 2018 (U.S. Securities and Exchange Commission, 2018). In practice, however, disclosure has remained incomplete and inconsistent. Although many firms eventually acknowledge an attack, few provide quantitative estimates of remediation costs, operational disruption, or investment responses, resulting in a fragmented understanding of the underlying economic consequences (Adams and Moore, 2025).

This absence of quantitative disclosure complicates empirical assessment for both researchers and policymakers, while obscuring the true risk exposure for investors. Existing empirical work largely relies on stock price reactions around disclosure events, yet recent evidence suggests these fluctuations often reflect transient investor sentiment rather than realized operational or accounting outcomes (Muktadir-Al-Mukit and Ali, 2025). Furthermore, these event studies typically focus on narrow windows of just a few days, which seminal work by Tetlock (2007) suggests may capture temporary price pressure and subsequent reversals rather than permanent changes in firm fundamentals. Whether cyber attacks materially alter firms' financial positions, their earnings capacity, or their capital structure therefore remains an open empirical question.

To address this gap, we collect data on publicly known cyber attacks and link them to quarterly financial statements for U.S. public firms, producing a firm–quarter panel that records financial outcomes before and after disclosure. This allows us to observe how cyber incidents affect balance sheet items such as total assets, total liabilities, and capital expenditures, as well as income statement components including operating costs and earnings, rather than inferring impacts from market reactions or qualitative disclosures. We estimate these effects using a dynamic event-study Difference-in-Differences design that traces outcomes from four quarters before to four quarters after disclosure.

Our results reveal three core findings. First, on average, cyber attacks cause sustained increases in operating and capital expenditures, financed primarily through higher liabilities, and lead to measurable declines in earnings. Second, the average effect masks stark heterogeneity by incident type. Data breaches trigger a financing and investment response in which firms expand their balance sheets to fund remediation and system upgrades. By contrast, ransomware attacks generate business interruption and earnings losses without inducing comparable capital restructuring. Third, these adjustments are not visible in qualitative corporate disclosures, suggesting a disclosure gap between what firms report and what they bear economically.

By analyzing incident-linked financial outcomes rather than market expectations, this study clarifies how cyber risk propagates through corporate financial statements and why existing market-based estimates may fail to capture heterogeneous underlying mechanisms.

Furthermore, our primary event-study results are robust to an alternative Year-over-Year

(YoY) Difference-in-Differences specification, which confirms that the documented financial adjustments are persistent at an annual scale.

The remainder of the paper proceeds as follows. Section 2 reviews related work on the financial and disclosure implications of cyber incidents, and also discusses recent economic literature on dynamic event-study Difference-in-Differences designs, which have become widely used in applied empirical research. Section 3 describes the construction of the firm–quarter panel and incident sample. Section 4 outlines the empirical design. Section 5 presents the main results and documents heterogeneity across incident types. Section 6 concludes.

2 Related Literature

This paper contributes to the empirical literature on cyber incidents and firm performance, as well as to recent methodological work on dynamic Difference-in-Differences designs. Our study bridges a critical gap between investor expectations, as measured by stock market reactions, and realized operational outcomes captured in corporate financial statements.

Early empirical studies primarily examined the impact of security breaches on firm value using stock market event-study techniques. These studies generally assume that capital markets efficiently assess the expected legal, operational, and reputational losses following a disclosure. Beginning with Campbell et al. (2003), a significant body of research documented negative abnormal returns at the time of breach disclosure, particularly when confidential customer information was compromised. Initial studies identified immediate drops in share price ranging from 2.7% within one day to 4.7% over a three-day window (Garg et al., 2003). These penalties are often moderated by firm characteristics such as size and industry type (Cavusoglu et al., 2004).

However, as the frequency of cyber attacks has increased, the market’s response has become increasingly inconsistent, leading to what we term a “consensus gap.” One camp of research identifies acute value destruction, with long-term abnormal returns falling as low as 15% to 18% in the year following an incident (Ali et al., 2021). Conversely, a second camp finds that many incidents produce statistically insignificant reactions. For instance, Hovav and D’Arcy (2003) found that availability-related attacks, such as Denial-of-Service events, often result in no significant market penalty. Similarly, Richardson et al. (2019) and Patsakis et al. (2018) suggest that the economic impact for many firms is “much ado about nothing,” with stock prices frequently recovering within days. Perhaps most surprising is a third camp documenting positive reactions; Michel et al. (2020) identify a “shareholder puzzle” where some firms experience positive returns after a breach, while Berkman et al. (2018) find that investors view post-breach investment announcements as value-adding activities.

Recent research by Gurjar et al. (2025) specifically tests this gap by examining whether the direct, internalized costs of ransomware trigger more definitive market penalties compared to traditional breaches. Despite the theoretical severity of these incidents, their analysis reveals that aggregate Cumulative Abnormal Returns (CARs) remain statistically weak and highly sector-dependent, with technology firms often showing rapid recovery or even positive corrections. This lack of a consistent market signal for such a tangible and quantifiable threat suggests that investor-based metrics may systematically fail to capture the full economic gravity of cyber incidents.

A critical limitation of the aforementioned literature is its reliance on investor sentiment and stock market reactions, which capture expectations and beliefs rather than realized economic outcomes. As shown by Baker and Wurgler (2006), investor sentiment systematically affects asset prices and expected returns, even in the absence of corresponding changes in firms’ fundamentals. Consequently, event-study based evidence around cyber incident disclosures primarily reflects how investors anticipate future losses, not whether these losses ultimately materialize in firms’ operational or accounting performance. To address this, a small number of studies have examined realized accounting metrics, though they too report mixed findings. Ko and Dorantes (2006) observed significant negative differences in Return on Assets (ROA) only in specific post-breach quarters. Gwebu et al. (2014) found that while direct costs (extraordinary items) are significant, the impact on core sales and market share often appears statistically insignificant. More recently, Frimpong and Chen (2021) documented significant revenue impacts following a breach, yet found no corresponding significant relationship with net income or future marketing expenses.

This study differs from prior work in three respects. First, we move beyond “vague” proxies by using a large firm-quarter panel linked to quarterly SEC financial statements, allowing us to trace realized adjustments in liabilities and expenditures. Second, we explicitly separate confidentiality attacks from availability attacks, showing that the mixed results in prior literature likely stem from grouping these distinct shocks together. Third, we provide quarter-by-quarter evidence on specific balance sheet items, thereby shedding light on the mechanisms through which cyber risk permanently reshapes corporate finance.

Methodologically, this study connects to recent advances in Difference-in-Differences with staggered treatment timing. Traditional two-way fixed effects estimators may produce biased estimates when treatment effects vary across cohorts or over event time (Goodman-Bacon, 2021; Sun and Abraham, 2021; Callaway and Sant’Anna, 2021; de Chaisemartin and D’Haultfoeuille, 2020), motivating a shift toward dynamic event-study estimators that allow for heterogeneous treatment dynamics. We follow this approach by estimating event-time treatment effects from four quarters before to four quarters after disclosure, which aligns with current practice in applied empirical economics.

The analysis also relates to work that combines matching and Difference-in-Differences to improve pre-treatment comparability (Heckman et al., 1997; Abadie, 2005; Stuart, 2010). In our setting, treated firms are matched to similar controls based on pre-incident characteristics, and we adjust for remaining imbalance in profitability during estimation. This strategy is conceptually related to recent double-robust Difference-in-Differences estimators (Sant’Anna and Zhao, 2020), which combine outcome modeling with propensity weighting to improve identification. Taken together, our empirical implementation aligns with this design-based approach by applying modern event-study Difference-in-Differences methods to financial accounting outcomes in the context of cyber risk and by documenting heterogeneous effects across incident types.

3 Data and Sample Construction

This section describes the construction of the firm–quarter panel used to estimate the financial effects of cyber incidents¹. The dataset integrates information from three sources: (i) cyber incident records from the Temple University Cyber Incident Repository (Rege, 2025) and the Center for International & Security Studies at Maryland (CISSM) (Harry and Gallagher, 2018); (ii) quarterly financial statements for U.S. public firms obtained from the SEC EDGAR system, which aggregates mandatory Form 10-Q and Form 10-K filings and provides standardized accounting information; and (iii) SIC industry classifications. The EDGAR filings supply balance sheet and income statement variables at the firm–quarter level, allowing us to track changes in assets, liabilities, capital expenditures, operating costs, and earnings before and after an incident. Merging these components yields a firm–quarter event panel that spans the period 2014–2023.

3.1 Sample Construction and Data Assembly

To construct the analysis sample, we first link cyber incident reports to publicly traded U.S. firms in the SEC EDGAR database. Incident reports contain informal organization names, whereas SEC registrants use legal entity names. To reconcile these differences, we normalize all names from both sources by converting to lowercase, removing legal suffixes and punctuation, and alphabetically sorting the remaining tokens. A hybrid fuzzy matching procedure based on the `rapidfuzz` algorithm (Bachmann, 2021) is then applied to compute pairwise similarity scores. This matching procedure, combined with manual verification, identifies 484 distinct cyber incidents affecting 432 unique firms. To ensure clean estimation baselines, we exclude recurring incidents at the same firm that occur within a three-year window, preventing overlapping treatment effects or contamination from prior remediation costs. Multiple incidents for a single firm are only retained if separated by more than three years, providing a sufficient buffer for financial metrics to stabilize before a subsequent event.

To isolate the financial impact of these incidents, we construct a control group of matched pairs using one-to-one Propensity Score Matching (PSM). Our matching process begins with the 484 identified incidents, but we first impose a strict data availability requirement: both treated and potential control firms must have at least two quarterly financial filings in the pre-incident window and at least two in the post-incident window. This requirement ensures that valid before–after comparisons can be made and eliminates 116 incidents due to insufficient data. For the remaining 368 incidents, we build a pool of potential controls restricted to firms in the same two-digit SIC industry.

For each firm, we compute two pre-incident financial characteristics over the four quarters preceding the incident: (i) firm size, measured as the average of log-transformed total assets, and (ii) profitability, measured as the average of operating income scaled by assets. Using these measures, we estimate a propensity score for each firm via logistic regression. Each treated firm is matched to the nearest control firm within its SIC two-digit industry, provided the propensity score difference falls within a 0.2 standard deviation caliper. This matching procedure successfully pairs 338 of the 368 incidents; the remaining 30 incidents fail to find

¹We use the terms cyber incident and attack interchangeably.

suitable matches due either to a lack of available firms within the same industry (12 cases) or violations of the matching caliper (18 cases).

Post-matching balance diagnostics show that profitability is well balanced between treated and control firms (Standardized Mean Difference [SMD] = 0.0489), which lies well below the conventional 0.1 threshold. Firm size, however, exhibits a moderate imbalance (SMD = 0.1955). To address this discrepancy, we explicitly include pre-incident firm size as a control in the regression models. The matching process yields an analysis sample consisting of 676 firms (338 treated and 338 controls).

3.2 Event Window Construction and Stacked Panel

Cyber incidents occur at different calendar dates across firms, generating staggered treatment timing. To ensure comparability, we re-index quarterly financial reporting data relative to each incident date, defining the disclosure quarter as Q_0 and constructing a symmetric event window from Q_{-4} to Q_{+4} . Empirical analysis of our sample confirms that incident timing is approximately uniformly distributed throughout the quarterly window. For the 338 unique events, the mean disclosure date occurs at day 45.18 of the quarter (median = 43.0). This balanced distribution minimizes the risk of timing-related bias, as the estimated coefficients represent the average treatment effect across a heterogeneous set of within-quarter disclosure dates. The matched control firm inherits the same event-time structure. This event-time re-indexing is increasingly standard in empirical settings where treatments do not occur simultaneously.²

Since 22 firms are reused as control matches for multiple treated firms, we construct the panel using a *stacked* design. Under this approach, each treated–control match contributes a separate event window, and windows are stacked vertically to form the estimation sample. The stacked design prevents contamination of comparisons when control firms serve as matches for multiple events and ensures that identification relies on within-event contrasts rather than across-event aggregation. The final stacked panel covers the period 2014–2025 and contains 63,846 firm–quarter observations.

3.3 Incident Types, Timing, and Industry Distribution

Table 1 reports the distribution of the 338 cyber incidents used in estimation by year and incident type. Incident frequency increases substantially over time, particularly after 2019, driven by the rise of ransomware attacks. Data breaches dominate the earlier part of the sample, whereas ransomware becomes the modal attack type by 2023.

We classify incidents into three mutually exclusive categories: Data Breach (178 events), Ransomware (139 events), and a residual “Others” category (21 events). Table 2 reports the distribution of these incidents across two-digit SIC industries. Data breaches are concentrated in information-intensive sectors such as Services and FIRE, while ransomware attacks are concentrated in operationally intensive sectors such as Manufacturing and Transport/Utilities. These sectoral patterns are consistent with differences in confidentiality versus

²See Schmidheiny and Siegloch (2023) for a discussion of staggered adoption designs.

Table 1: Distribution of Cyber Incidents by Year

Year	Data Breach	Others	Ransomware	Total
2014	12	5	1	18
2015	9	3	0	12
2016	20	0	0	20
2017	7	0	4	11
2018	14	3	3	20
2019	12	0	8	20
2020	17	2	27	46
2021	13	1	16	30
2022	23	1	17	41
2023	30	3	43	76
2024	20	2	20	42
2025	1	1	0	2
Total	178	21	139	338

availability exposure across industries and motivate the heterogeneous treatment analysis in Section 5.2.

Table 2: Distribution of Cyber Incidents by Industry

Industry (SIC)	Data Breach	Others	Ransomware	Total
Manufacturing (20-39)	47	6	61	114
Services (70-89)	60	9	32	101
FIRE (60-67)	29	2	13	44
Transport/Util (40-49)	15	3	19	37
Retail (52-59)	22	1	5	28
Wholesale (50-51)	2	0	6	8
Mining (10-14)	2	0	1	3
Construction (15-17)	1	0	1	2
Agriculture (01-09)	0	0	1	1
Total	178	21	139	338

3.4 Pre-Attack Financial Characteristics

Before turning to the empirical analysis, we compare the financial characteristics of treated and matched control firms in the pre-incident period. These comparisons serve two purposes. First, they provide economic context regarding the scale and variability of firms in our sample. Second, they document the extent to which the matched treatment and control groups are comparable before the cyber incident, which is relevant for the validity of the Difference-in-Differences design.

Table 3 reports means, medians, and standard deviations for selected income statement and balance sheet variables, measured over the four quarters preceding the incident (Q_{-4} to Q_{-1}). Treated firms are slightly larger on average, as indicated by Total Assets and Total Liabilities, consistent with the moderate residual size imbalance discussed in Section 3.1. In contrast, profitability measures (Operating Income and Pre-Tax Income) and R&D expenditures are similar across groups. These patterns are consistent with the post-matching

Table 3: Pre-Attack Financial Characteristics (USD millions)

Financial Variable	Control Mean	Treated Mean	Control Median	Treated Median	Control StdDev	Treated StdDev
Operating Income	257.5	336.6	30.9	40.4	871.1	1,055.5
Pre-Tax Income	261.6	302.0	37.0	28.7	837.1	984.1
R&D Expenses	227.3	225.9	36.7	38.2	564.1	551.4
SG&A Expenses	530.1	690.6	192.7	207.2	1,115.8	1,550.8
Total Assets	23,493.0	27,077.1	4,956.1	5,755.1	56,790.3	66,070.9
Total Liabilities	18,878.8	19,437.1	2,347.9	2,573.7	54,089.9	60,247.1

balance diagnostics and suggest that, aside from scale, treated and control firms are comparable along key economic dimensions prior to the incident.

4 Empirical Methodology

To estimate the financial impact of cyber attacks, we employ a dynamic Difference-in Differences (DiD) event-study model on the matched sample described in Section 3. The goal is to compare how financial outcomes evolve for firms that experience a cyber incident relative to similar firms that do not, before and after the incident, while allowing the treatment effect to vary across quarters. This dynamic structure is appropriate in our setting because the consequences of cyber incidents unfold gradually over time through remediation expenditures, legal and regulatory processes, operational disruptions, and financing decisions, rather than producing a single one-period shock.

For each firm i and calendar quarter t , let Y_{it} denote the transformed financial outcome of interest. We define $\text{Treated}_i = 1$ for firms that experience a cyber incident in our sample period, and $\text{Treated}_i = 0$ for matched control firms. We construct an event-time index $\tau \in \{-4, -3, -2, -1, 0, 1, 2, 3, 4\}$ relative to the disclosure quarter of the incident, where $\tau = 0$ denotes the disclosure quarter and $\tau = -1$ is the final pre-incident quarter. Matched control firms are assigned the same event-time index as their corresponding treated firm. Using this framework, we estimate:

$$Y_{it} = \sum_{k=-4, k \neq -1}^4 \beta_k (\text{Treated}_i \cdot 1[\tau = k]) + \alpha_i + \delta_j + \gamma_t + \lambda(X_{i,pre} \times \gamma_t) + \epsilon_{it}. \quad (1)$$

In this specification, α_i denotes firm (CIK) fixed effects, γ_t denotes calendar-quarter fixed effects, and δ_j denotes event-stack fixed effects. Firm fixed effects absorb all time-invariant heterogeneity across firms, such as industry affiliation, long-run firm size, business model, or cybersecurity posture, while quarter fixed effects absorb macroeconomic conditions, seasonal financial reporting cycles, and regulatory shocks common to all firms in a given calendar quarter. The event-stack fixed effects δ_j are required since 22 firms in the matched control pool are reused as matches for multiple treated firms. We therefore construct stacked event windows, where each treated-control pair contributes a separate nine-quarter window spanning $\tau = -4$ to $\tau = +4$. Without stacking, the reuse of control firms would create contamination across unrelated incidents by mixing pre and post-treatment periods from

distinct events. The use in δ_j ensure that identification of the β_k coefficients is driven solely by within-event deviations and not by comparisons across stacked windows.

The term $\text{Treated}_i \cdot 1[\tau = k]$ identifies the relative-quarter treatment effect. Because $\tau = -1$ is omitted from the summation, the coefficients β_k are interpreted relative to the quarter immediately preceding the incident. For $k < 0$, the coefficients allow us to test for differential pre-trends between treated and control firms, and for $k \geq 0$, they trace out the post-incident trajectory of the financial outcome. In particular, β_0 captures the contemporaneous effect in the disclosure quarter, and β_1, \dots, β_4 capture the dynamic evolution of the effect over the subsequent four quarters.

The vector $X_{i,pre}$ consists of pre-incident firm characteristics, computed over the pre-treatment window $\tau \in \{-4, -3, -2, -1\}$. In our setting, $X_{i,pre}$ includes the average of log-transformed assets, which remains moderately imbalanced between treated and control firms after matching. Interacting $X_{i,pre}$ with calendar-quarter fixed effects allows the influence of this baseline characteristic to vary across time and improves the credibility of the comparison between treated and control firms.

Since financial variables in SEC filings vary widely in scale and often include negative values, we transform all outcomes before estimation. The choice of transformation depends on whether a variable can take zero or negative values. For variables that are always positive, such as *Total Assets*, *Total Liabilities*, *SG&A Expense*, *Accounts Payable*, and *Cost of Goods Sold*, we apply the $\log(1 + y)$ transformation. This stabilises the distribution and allows the estimated coefficients to be interpreted as approximate percentage changes. We use $\log(1 + y)$ rather than $\log(y)$ to guard against zero-valued edge cases; at the scale of the financial magnitudes in our sample, the two are numerically indistinguishable.

For variables that can be zero or negative, $\log(1 + y)$ is not applicable. We therefore apply the inverse hyperbolic sine (IHS) transformation, $\text{asinh}(y) = \log(y + \sqrt{y^2 + 1})$, which is defined for all real values and behaves like a logarithm for large positive numbers while remaining well-defined at zero and for negative values. Distributional diagnostics show that this applies to most income-statement and cash-flow variables: for example, *Operating Income* is negative in 25.3% of firm-quarters, *Non-operating Income/Expense* in 66.5%, and *Operating Cash Flow* in 30.3%. Balance-sheet accounts such as *Long-Term Debt (Current)*, *PP&E (Net)*, *Accounts Receivable*, *Inventory*, and *R&D Expense* are also assigned to this group because they contain zero observations in the data.

Standard errors are two-way clustered at the `event_id` and firm (`cik`) levels. This clustering structure accounts for arbitrary serial correlation within the incident window as well as correlation within firms that appear in multiple event stacks. Under the identifying assumption that, absent a cyber incident, treated and control firms would have followed parallel trajectories conditional on fixed effects and baseline controls, the post-incident coefficients β_0, \dots, β_4 represent the causal effect of cyber attacks on the financial outcome Y_{it} .

5 Results

This section presents the results from the dynamic event-study Difference-in-Differences specification in Equation (1). We first validate the parallel-trends assumption and then report the estimated post-incident effects, both in the aggregate and by incident type. Full

sets of pre-incident coefficients validating the parallel-trends assumption are provided in the Appendix; Table A3 reports these estimates for the pooled sample, while Table A4 provides the validation for all significant heterogeneous outcomes.

To begin, we assess whether treated firms and their matched controls exhibit similar pre-incident dynamics. The identifying requirement for the event-study design is that, absent a cyber incident, the two groups would have evolved along comparable trajectories in the relevant financial outcomes.

Table 4 reports the pre-incident coefficients for event-time quarters Q_{-4} , Q_{-3} , and Q_{-2} , with Q_{-1} omitted as the reference period. For Total Assets, Total Liabilities, Operating Income, and Pre-Tax Income, the estimated pre-treatment coefficients are small in magnitude and statistically indistinguishable from zero. The absence of systematic pre-trends supports the validity of the Difference-in-Differences strategy and justifies interpreting the post-incident coefficients as causal effects of cyber incidents.

Table 4: Validation of Parallel Trends (Pre-Attack Coefficients, Model 1)

	Assets	Liabilities	Operating Income	Pre-Tax Income
Q_{-4}	-0.085 (0.064)	-0.040 (0.045)	1.502 (1.336)	-0.480 (1.646)
Q_{-3}	-0.014 (0.017)	0.049 (0.037)	0.925 (1.272)	1.462 (1.536)
Q_{-2}	-0.005 (0.009)	-0.009 (0.015)	-0.542 (1.147)	-1.922 (1.447)
Transformation	log1p	log1p	asinh	asinh
Observations	5,795	4,471	3,976	3,201
FE (cik)	601	477	504	454
FE (event_id)	338	318	314	312
FE (period_qidx)	47	47	47	47

Notes: Coefficients report pre-attack differences between treated firms and matched controls at Q_{-4} , Q_{-3} , and Q_{-2} , relative to the omitted quarter Q_{-1} . All models are estimated by OLS and include firm, event, and calendar-quarter fixed effects. Standard errors are two-way clustered by `event_id` and `cik` and reported in parentheses. Significance codes: . $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

5.1 Average Effects of Cyber Incidents

We next estimate the average dynamic impact of cyber incidents on firm financials using the full sample of 338 matched events. Table 5 summarize the post-incident coefficients Q_0 through Q_4 for key balance sheet and income statement variables. All regressions include firm, event, and calendar-quarter fixed effects, as well as interactions between quarter fixed effects and pre-incident firm size, with standard errors two-way clustered by `event_id` and firm (`cik`).

Three main patterns emerge.

Balance sheet expansion. Total Assets (log1p) increase significantly beginning in the incident quarter Q_0 , with the largest effect at Q_2 . The Q_2 coefficient of approximately

Table 5: Summary of Post-Treatment Regression Results (Model 1)

Panel A: Balance Sheet and Cash-Flow Outcomes					
	Total Assets	PPE (Net)	SG&A Expense	Accounts Rec.	Op. Cash Flow
Q0	0.018** (0.007)	0.006 (0.010)	0.035 . (0.018)	0.040 (0.031)	-12.031 (8.427)
Q1	0.031* (0.013)	0.026 (0.017)	0.044* (0.021)	0.128 . (0.066)	-1.509 (9.820)
Q2	0.053** (0.017)	0.044* (0.019)	0.071* (0.029)	0.054 (0.043)	-6.543 (9.546)
Q3	0.042* (0.019)	0.037 (0.022)	0.068* (0.029)	0.044 (0.041)	-1.370 (2.770)
Q4	0.043 . (0.022)	0.028 (0.032)	0.056 . (0.033)	0.123 (0.084)	-9.770 (8.295)
Transformation	log1p	asinh	log1p	asinh	asinh
Observations	5,795	4,752	2,225	3,829	1,305
FE (cik)	601	520	305	401	521
FE (event_id)	338	320	227	282	326
FE (period_qidx)	47	47	47	47	47

Panel B: Income-Statement and Liabilities Outcomes					
	Liabilities	Gross Profit	Cont. Op. Inc.	Non-op. Inc.	LTD (Curr.)
Q0	0.023 (0.016)	-0.128* (0.056)	0.536 (1.322)	0.507 (2.205)	0.360 (0.791)
Q1	0.051* (0.022)	0.039 (0.127)	-1.028 (1.497)	-1.534 (2.222)	0.265 (0.831)
Q2	0.069* (0.031)	-0.019 (0.072)	-1.795 (1.567)	-8.358** (2.851)	1.817 . (1.032)
Q3	0.056 . (0.033)	0.314 (0.384)	-3.774** (1.405)	0.253 (2.340)	1.469 (1.060)
Q4	0.036 (0.031)	-0.119 . (0.071)	-2.749 . (1.534)	-0.828 (2.124)	1.086 (1.163)
Transformation	log1p	asinh	asinh	asinh	asinh
Observations	4,471	2,022	3,201	1,457	1,347
FE (cik)	477	261	454	202	175
FE (event_id)	318	183	312	176	155
FE (period_qidx)	47	47	47	47	47

Notes: All models are OLS estimations on the matched sample. Coefficients report post-incident effects relative to the omitted quarter Q_{-1} . Standard errors (in parentheses) are clustered by `event_id` and `cik`. Significance codes: . $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

0.053 corresponds to a 5.3% increase relative to Q_{-1} ; for a firm with \$2.5 billion in assets, this implies roughly \$132 million in additional assets. Total Liabilities (log1p) also increase significantly from Q_1 , with a peak effect of about 6.9% at Q_2 (around \$138 million for a firm with \$2 billion in liabilities). Firms thus respond to cyber incidents by expanding their balance sheets, primarily via higher liabilities rather than asset liquidation.

Increases in operating and capital spending. The balance sheet expansion is mirrored in expenditure accounts. SG&A Expense (log1p) rises significantly from Q_1 through Q_3 , with the largest effect at Q_2 (about 7.1% relative to Q_{-1}), corresponding to roughly \$35 million in additional spending for a firm with \$500 million in quarterly SG&A. PP&E Net (asinh) also increases significantly at Q_2 , indicating tangible investments in hardware, infrastructure, and related remediation projects.

Profitability declines. On the income statement, firms experience a deterioration in profitability. Gross Profit (asinh) declines significantly in the incident quarter Q_0 , and in later quarters we observe large, negative and significant effects on Non-operating Income/Expense and Income from Continuing Operations. These patterns are consistent with a combination of direct remediation costs, legal and regulatory expenses, and temporary disruptions to operations.

Overall, the pooled estimates indicate that cyber incidents lead to sustained increases in operating and capital expenditures, financed through higher liabilities, and to a persistent decline in profitability over the following four quarters.

5.2 Heterogeneous Effects by Incident Type

We now examine whether financial impacts differ across incident types. We focus on the two major categories defined in Section 3.3: Data Breaches (178 events), which primarily compromise confidentiality, and Ransomware attacks (139 events), which primarily disrupt availability. The residual “Others” category (21 events) is small and not central to our interpretation.

Table 6 reports heterogeneous post-incident effects for Data Breaches and Ransomware attacks. Panel A presents balance-sheet outcomes (Total Assets and Total Liabilities), and Panel B presents profitability outcomes (Operating Income and Income before Noncontrolling Interest). Additional account-level specifications (e.g., Accounts Receivable, COGS, Inventory, PP&E, R&D) are provided in Appendix Tables A1–A2.

Data breaches: financing and investment response. Data breaches exhibit a pronounced balance sheet response. As illustrated in the left column of Figure 1, Total Liabilities (log1p) increase significantly beginning in Q_0 and remain elevated through Q_4 . The Q_3 coefficient of about 0.129 implies a 12.9% increase relative to Q_{-1} , corresponding to roughly \$258 million in additional obligations for a firm with \$2 billion in pre-incident liabilities. Total Assets rise in parallel, indicating that the additional liabilities are used to finance asset accumulation rather than simply covering cash shortfalls. Panel (e) confirms that this

Table 6: Heterogeneous Effects of Cyberattacks by Incident Type

Panel A: Balance Sheet Effects				
	Assets		Liabilities	
	Breach	Ransomware	Breach	Ransomware
Q0	0.029** (0.011)	0.001 (0.009)	0.063** (0.024)	-0.029 (0.024)
Q1	0.050** (0.018)	0.014 (0.020)	0.106*** (0.032)	-0.011 (0.034)
Q2	0.065** (0.021)	0.043 (0.030)	0.127*** (0.037)	0.006 (0.055)
Q3	0.067* (0.028)	0.016 (0.029)	0.129** (0.038)	-0.023 (0.059)
Q4	0.074* (0.033)	0.008 (0.032)	0.126** (0.043)	-0.067 (0.046)
Transformation	log1p	log1p	log1p	log1p
Observations	3,058	2,401	2,379	1,870
FE (cik)	327	268	260	212
FE (event_id)	178	139	167	132
FE (period_qidx)	47	46	47	46
Panel B: Profitability Effects				
	Operating Income		Noncont. Income	
	Breach	Ransomware	Breach	Ransomware
Q0	1.906 (1.616)	-1.568 (1.955)	2.232 (1.912)	-2.149 (2.415)
Q1	1.052 (1.803)	-0.493 (2.126)	-0.614 (1.983)	-1.946 (2.692)
Q2	4.193* (1.640)	-2.162 (2.025)	1.318 (1.912)	-7.111** (2.486)
Q3	0.800 (1.681)	-5.120* (2.196)	-2.145 (1.722)	-5.398* (2.575)
Q4	-0.261 (1.778)	-4.779* (2.211)	-0.766 (1.761)	-5.075 (2.875)
Transformation	asinh	asinh	asinh	asinh
Observations	2,052	1,689	1,663	1,369
FE (cik)	265	235	240	210
FE (event_id)	162	131	164	130
FE (period_qidx)	47	37	47	36

Notes: Standard errors clustered by `event_id` and `cik` are reported in parentheses. Noncont. Income refers to Income (Loss) from Noncontrolling Interest. Significance codes: . $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

financing is partly used for tangible investments in Property, Plant, and Equipment. Appendix results show that working-capital accounts such as Accounts Receivable and Accounts

Payable also increase, and that COGS, Inventory, PP&E, and R&D all rise in the quarters following a breach. Importantly, Operating Income does not experience large, persistent declines; in some quarters (notably Q_2), it increases, which is consistent with remediation and investment being financed in a way that preserves overall profitability.

Ransomware: operational disruption and earnings losses. Ransomware incidents display a very different pattern. Consistent with the visual evidence in the right column of Figure 1, neither Total Assets nor Total Liabilities show robust, sustained increases, suggesting that firms do not respond by materially expanding their balance sheets. Instead, the primary effects are on operations and earnings. Appendix evidence indicates that Accounts Receivable decline significantly in later quarters, consistent with business interruption and delayed invoicing. Table 6 shows that, for ransomware events, Gross Profit declines significantly at Q_2 , Operating Income drops sharply at Q_3 , and Income before Noncontrolling Interest is significantly negative at Q_2 and Q_3 . Unlike data breaches, there is little sign of systematic increases in COGS, Inventory, PP&E, or R&D, reinforcing the view that ransomware acts mainly as an availability shock that temporarily halts production and revenue generation.

The finding that profitability declines become most pronounced in Q_2 and Q_3 rather than the incident quarter (Q_0) aligns with the "occurrence-to-recognition" cycle documented in banking and operational risk literature. Specifically, Aldasoro et al. (2020) find an average lag of 184 days between the discovery of an incident and its formal recognition in a firm's books. This "recognition lag" represents the time required for organizations to internalize the costs of an event, which often involves assessing the full financial impact and resolving legal or regulatory proceedings before the loss is officially recorded in the accounts. Furthermore, our results are influenced by statutory reporting lags documented by Easton and Zmijewski (1993), who find that firms typically file mandatory SEC reports very close to their legal deadlines, specifically 45 days for Form 10-Q and 90 days for Form 10-K following the end of a fiscal period. These cumulative operational and regulatory delays ensure that the financial consequences of an incident occurring in Q_0 may not be fully realized or disclosed until Q_2 or Q_3 .

Summary. The heterogeneous effects make clear that the pooled average impact of cyber incidents is a composite of two distinct mechanisms. Data breaches trigger a financing and investment response, with firms raising additional resources and investing in remediation and system upgrades while largely preserving profitability. Ransomware incidents, by contrast, cause acute operational disruptions and substantial earnings losses, with little evidence of strategic balance-sheet adjustment. Distinguishing between these mechanisms is essential for understanding the true economic consequences of cyber risk.

5.3 Validation Model

Our second model, Model 2, is a Year-over-Year (YoY) Difference-in-Differences (DiD) used for validation. This model compares the change in a financial variable from one year to the next, which implicitly removes firm-level fixed effects.

Table 7: Year-over-Year (YoY) Impact: Data Breach vs. Ransomware

Panel A: Data Breach (Confidentiality Attacks)				
	Total Assets	Total Liabilities	Inventory (Net)	Accounts Rec.
Q_0 vs Q_{-4}	0.163 (0.110)	0.120* (0.056)	0.042 (0.072)	0.157 (0.114)
Q_1 vs Q_{-3}	0.082** (0.026)	0.084. (0.049)	0.139* (0.065)	0.491 (0.328)
Q_2 vs Q_{-2}	0.075** (0.024)	0.121* (0.051)	0.138* (0.060)	0.268* (0.133)
Q_3 vs Q_{-1}	0.070* (0.030)	0.106* (0.044)	0.162* (0.063)	0.193. (0.109)
Transformation	log1p	log1p	asinh	asinh
Observations	316	198	100	140

Panel B: Ransomware (Availability Attacks)				
	Total Assets	Total Liabilities	Total Revenue	Accounts Rec.
Q_0 vs Q_{-4}	0.009 (0.038)	-0.105 (0.110)	-0.121 (0.115)	-0.032 (0.048)
Q_1 vs Q_{-3}	-0.010 (0.044)	-0.204 (0.137)	0.346 (0.464)	-0.072 (0.055)
Q_2 vs Q_{-2}	-0.007 (0.013)	0.053 (0.071)	-0.084 (0.092)	-0.097 (0.065)
Q_3 vs Q_{-1}	0.013 (0.030)	0.011 (0.097)	-0.101 (0.067)	-0.119* (0.055)
Q_4 vs Q_0	-0.008 (0.033)	-0.009 (0.058)	-0.135* (0.058)	-0.117* (0.052)
Transformation	log1p	log1p	asinh	asinh
Observations	252	148	58	130

Notes: Estimates are from YoY OLS regressions with clustered standard errors in parentheses. Panel A shows confidentiality attacks (Data Breaches) and Panel B shows availability attacks (Ransomware). Total Revenue for Ransomware is measured using Revenue from Contract with Customer. Significance codes: . $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

$$\Delta Y_{it} = \beta_0 + \beta_1 \text{Treated}_i + \gamma_t + \beta_3 \text{mean_log_assets}_i + \epsilon_{it} \quad (2)$$

In this model, the main variable (Δy_{it}) is the year-over-year change in a financial outcome. We retain the pre-attack size variable and calendar time fixed effects (γ_t) to maintain consistency with the dynamic specification.

Table 7 presents these validation results, disaggregated by incident type. The findings in Panel A and Panel B strongly mirror the primary dynamic results. For Data Breaches, we observe significant annual increases in Total Assets and Total Liabilities, particularly in the year-over-year windows corresponding to the mid-remediation period (Q_2 vs Q_{-2}).

Conversely, Panel B confirms the operational disruption mechanism for Ransomware, where Total Revenue and Accounts Receivable show significant year-over-year declines in the later quarters of the post-incident year. By yielding consistent estimates across an alternative functional form, these results reinforce the causal interpretation of our main event-study coefficients.

6 Conclusion

This paper examines how cyber attacks affect the financial performance and balance sheet decisions of publicly traded U.S. firms. Leveraging a matched sample of 338 incidents from 2014–2025 and a stacked dynamic Difference-in-Differences event-study design, we document how both the composition and the consequences of cyber incidents have evolved over time. The empirical strategy combines Propensity Score Matching with rich fixed effects and event-time re-indexing, allowing us to isolate the financial impact of cyber incidents from confounding industry and macroeconomic trends.

The consistency of these findings across both dynamic event-study and Year-over-Year validation models suggests that the financial impacts of cyber risk are not transient artifacts of high-frequency quarterly data but represent substantial annual shifts in corporate fundamentals.

Three findings emerge. First, the average cyber incident induces a multi-quarter adjustment in corporate finances. Firms increase operating and capital expenditures beginning in the disclosure quarter and finance these increases primarily through expanded liabilities. These systematic increases in obligations reveal a persistent financial burden that is frequently omitted from the initial “materiality” assessments found in official corporate disclosures. These adjustments are accompanied by a contemporaneous decline in profitability, particularly in the quarters immediately following the incident. This pattern suggests that cyber incidents impose persistent remediation and compliance costs that are not fully offset by operating revenues in the short run.

Second, the average effect masks substantial heterogeneity across incident types. Data breaches, which compromise confidentiality, trigger a financing and investment response: firms expand their balance sheets, increase working capital, and undertake remediation investments while preserving core profitability. In contrast, ransomware attacks, which compromise availability, lead to operational disruption. These attacks depress revenue generation and receivables and generate sizable earnings losses without inducing capital raising or balance sheet expansion. The contrast between these two attack types underscores the importance of distinguishing between disruption-driven and compliance-driven cyber events.

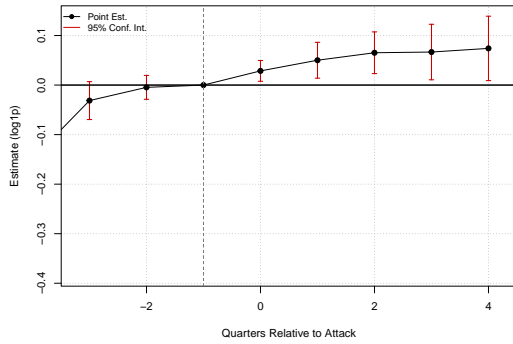
Third, the absence of differential pre-trends and the successful matching of treated and control firms provide empirical support for the identifying assumptions that underpin the causal interpretation of the estimates. The results therefore not only characterize the financial consequences of cyber incidents but also help reconcile seemingly conflicting evidence in the existing literature regarding whether cyber attacks destroy economic value, shift expenditure, or simply reconfigure financial reporting outcomes. Specifically, our results suggest that the disclosure gap hides significant, measurable impacts on a firm’s internal financial structure that market-based studies often fail to capture.

Taken together, the findings imply that cyber risk is a material financial risk for public firms and that the nature of the risk depends critically on the attack mechanism. To the extent that regulatory reporting and insurance coverage continue to expand, the distinction between confidentiality and availability attacks will become increasingly important for investors, regulators, and insurers. Understanding this heterogeneity may also help guide firms' cyber defense investments, insurance pricing, and post-attack recovery strategies.


Future work may extend this analysis in several directions. Linking financial outcomes to market valuation or credit spreads would allow an assessment of how investors price cyber risk. Incorporating time-to-resolution measures or insurance recovery data would help identify the duration of remediation and the role of risk transfer mechanisms. Finally, analyzing non-public firms or international settings would improve the external validity of the estimates. As cyber threats continue to grow in frequency and sophistication, understanding how firms absorb, finance, and recover from these incidents will remain central to both academic research and policy design.

Acknowledgments


Moore gratefully acknowledges support from the US National Science Foundation (NSF) Award No. 2452738.




(a) Data Breach: Total Assets



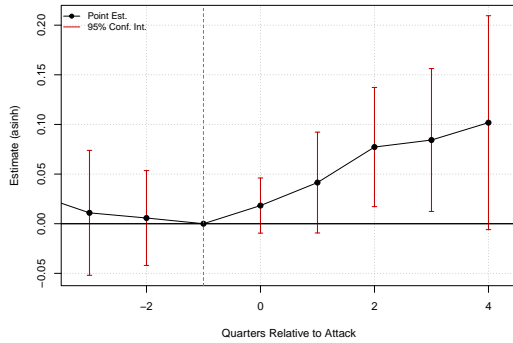
(b) Ransomware: Total Assets



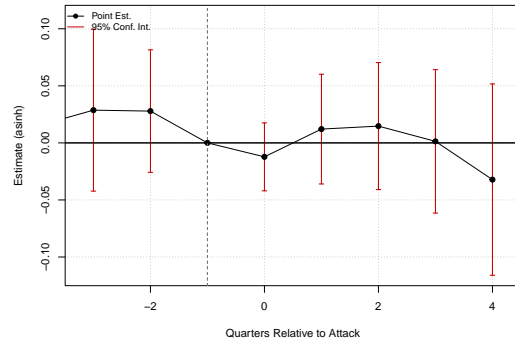
(c) Data Breach: Total Liabilities



(d) Ransomware: Total Liabilities



(e) Data Breach: PP&E (Net)



(f) Ransomware: PP&E (Net)

Figure 1: Heterogeneous balance sheet response by incident type. Event-study coefficients are estimated relative to Q_{-1} using a specification that includes leads from Q_{-4} onward. For clarity, the Q_{-4} coefficient is omitted from the figure due to imprecision; estimates from Q_{-3} to Q_4 are shown with 95% confidence intervals for Data Breaches (left column) and Ransomware attacks (right column).

References

- Abadie, A. (2005). Semiparametric difference-in-differences estimators. *Review of Economic Studies*, 72(1):1–19.
- Adams, M. and Moore, T. (2025). How informative are cybersecurity risk disclosures? empirical analysis of firms targeted by ransomware. *Computers & Security*, 159:104626.
- Aldasoro, I., Gambacorta, L., Giudici, P., and Leach, T. (2020). Operational and cyber risks in the financial sector.
- Ali, S. E. A., Lai, F.-W., Hassan, R., and Shad, M. K. (2021). The long-run impact of information security breach announcements on investors’ confidence: The context of efficient market hypothesis. *Sustainability*, 13(3):1066.
- Bachmann, M. (2021). Rapidfuzz: A fast string matching library for python.
- Baker, M. and Wurgler, J. (2006). Investor sentiment and the cross-section of stock returns. *The Journal of Finance*, 61(4):1645–1680.
- Berkman, H., Jona, J., Lee, G., and Soderstrom, N. (2018). Cybersecurity awareness and market valuations. *Journal of Accounting and Public Policy*, 37(6):508–526.
- Callaway, B. and Sant’Anna, P. H. (2021). Difference-in-differences with multiple time periods. *Journal of Econometrics*, 225(2):200–230.
- Campbell, K., Gordon, L. A., Loeb, M. P., and Zhou, L. (2003). The economic cost of publicly announced information security breaches: empirical evidence from the stock market. *Journal of Computer Security*, 11(3):431–448.
- Cavusoglu, H., Mishra, B., and Raghunathan, S. (2004). The effect of internet security breach announcements on market value: capital market reactions for breached firms and internet security developers. *International Journal of Electronic Commerce*, 9(1):70–104.
- de Chaisemartin, C. and D’Haultfoeuille, X. (2020). Two-way fixed effects estimators with heterogeneous treatment effects. *American Economic Review*, 110(9):2964–2996.
- Easton, P. D. and Zmijewski, M. E. (1993). Sec form 10k/10q reports and annual reports to shareholders: Reporting lags and squared market model prediction errors. *Journal of Accounting Research*, 31(1):113–129.
- Federal Bureau of Investigation (FBI) (2025). Internet crime report 2024. Technical report, Internet Crime Complaint Center (IC3). Accessed: 2025-05-12.
- Frimpong, B. and Chen, L. (2021). The effects of data breaches on public companies: A mirage or reality? *Advances in Information and Communication*, pages 381–393.
- Garg, A., Curtis, J., and Halper, H. (2003). The financial impact of it security breaches: what do investors think? *Information Management & Computer Security*, 12(1):22–33.

- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.
- Gurjar, A., Manatova, D., Staples, B., Chambers, S., and Camp, L. J. (2025). Is Ransomware an Economically Distinct Attack Type? An Event Study of Market Reactions . In *2025 APWG Symposium on Electronic Crime Research (eCrime)*, pages 1–10, Los Alamitos, CA, USA. IEEE Computer Society.
- Gwebu, K. L., Wang, J., and Xie, W. (2014). Understanding the cost associated with data security breaches. *Pacific Asia Conference on Information Systems (PACIS)*.
- Harry, C. and Gallagher, N. (2018). Classifying cyber events. *Journal of Information Warfare*, 17(3):17–31.
- Heckman, J. J., Ichimura, H., and Todd, P. E. (1997). Matching as an econometric evaluation estimator: Evidence from evaluating a job training programme. *Review of Economic Studies*, 64(4):605–654.
- Hovav, A. and D’Arcy, J. (2003). The impact of denial-of-service attack announcements on the market value of firms. *Risk Management and Insurance Review*, 6(2):97–121.
- Ko, M. and Dorantes, C. (2006). The impact of information security breaches on financial performance of the breached firms: an empirical investigation. *Journal of Information Technology Management*, 17(2):13–22.
- Michel, A., Oded, J., and Shaked, I. (2020). Do security breaches matter? the shareholder puzzle. *European Financial Management*, 26(2):288–315.
- Muktadir-Al-Mukit, D. and Ali, M. H. (2025). The dynamics of stock market responses following the cyber-attacks news: Evidence from event study. *Information Systems Frontiers*, 27(5):1741–1788.
- Patsakis, C., Charemis, A., Papageorgiou, A., Mermigas, D., and Pirounias, S. (2018). The market’s response toward privacy and mass surveillance: The snowden aftermath. *computers & security*, 73:194–206.
- Rege, A. (2025). Critical infrastructure ransomware attacks (CIRA) dataset. Temple University. Version 12.15.
- Richardson, V. J., Smith, R. E., and Watson, M. W. (2019). Much ado about nothing: The (lack of) economic impact of data privacy breaches. *Journal of Information Systems*, 33(3):227–265.
- Sant’Anna, P. H. and Zhao, J. (2020). Doubly robust difference-in-differences estimators. *Journal of Econometrics*, 219(1):101–122.
- Schmidheiny, K. and Sieglöcher, S. (2023). On event-study designs and distributed-lags in two-way fixed effects models. *Journal of Econometrics*.

- Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical Science*, 25(1):1–21.
- Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.
- Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of finance*, 62(3):1139–1168.
- U.S. Securities and Exchange Commission (2018). Commission statement and guidance on public company cybersecurity disclosures. Release No. 33-10459.

Appendix

A.1 Balance Sheet Decomposition

This section reports additional account-level estimates from the dynamic Difference-in-Differences specification in Equation (1). These results support the heterogeneous response patterns documented in Section 5.2.

Table A1 disaggregates the balance sheet response for the two major incident types—Data Breaches and Ransomware—into three accounts: Accounts Receivable (Current), Long-Term Debt (Current), and Accounts Payable (Current). The results show that data breaches induce substantial working-capital adjustments, reflected in significant increases in Accounts Receivable and Accounts Payable beginning in Q_0 and Q_2 , respectively. In contrast, ransomware incidents produce no systematic changes in these accounts, and, in later quarters, Accounts Receivable decline significantly, consistent with business interruption.

Table A1: Impact of Cyberattacks on Balance Sheet (by Account)

	Acc. Receivable (N)		LT Debt (C)		Acc. Payable (C)	
	Breach	Ransom.	Breach	Ransom.	Breach	Ransom.
Q0	0.109* (0.053)	-0.051 (0.032)	0.952 (1.032)	-0.296 (1.090)	0.015 (0.044)	0.014 (0.032)
Q1	0.259* (0.130)	-0.012 (0.046)	1.329 (0.694)	-0.631 (1.251)	0.064 (0.055)	0.019 (0.050)
Q2	0.208** (0.074)	-0.093 (0.048)	1.270 (1.513)	2.152 (1.607)	0.107* (0.053)	-0.019 (0.040)
Q3	0.189** (0.070)	-0.114* (0.048)	0.085 (1.920)	2.225 (1.562)	0.108 (0.056)	-0.036 (0.041)
Q4	0.389* (0.169)	-0.162** (0.055)	-0.478 (1.913)	1.662 (1.545)	0.117 (0.065)	-0.011 (0.066)
Transformation	asinh	asinh	asinh	asinh	log1p	log1p
Observations	1,851	1,726	563	698	1,964	1,702
FE (cik)	201	196	79	92	212	200
FE (event_id)	141	123	71	73	143	120
FE (period_qidx)	47	44	47	34	47	44

Notes: Standard errors clustered by `event_id` and `cik` are reported in parentheses. Standard abbreviations used for account names (N = Net, C = Current). Significance codes: . $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

A.2 Operational Expense and Investment Decomposition

To further examine how firms adjust operationally following cyber incidents, Table A2 reports estimates for four accounts: Cost of Goods and Services Sold (COGS), Inventory, Property, Plant and Equipment (PP&E), and Research and Development (R&D) Expense.

These results clarify that data breaches generate an investment-oriented response: COGS increases beginning in *Q1*, Inventory rises beginning in *Q2*, PP&E rises beginning in *Q2*, and R&D rises beginning in *Q1*. Conversely, ransomware exhibits neither systematic increases in operational inputs nor increases in investment, reinforcing the operational-disruption mechanism highlighted in the main text.

Table A2: Impact of Cyberattacks on Operations (Four Accounts)

	COGS		Inventory (N)		PPE (N)		R&D Exp.	
	Breach	Ransom.	Breach	Ransom.	Breach	Ransom.	Breach	Ransom.
Q0	0.059 (0.038)	-0.032 (0.041)	0.028 (0.027)	0.033 (0.026)	0.018 (0.014)	-0.012 (0.015)	0.011 (0.062)	0.105. (0.062)
Q1	0.170* (0.071)	0.005 (0.045)	0.037 (0.051)	-0.002 (0.033)	0.041 (0.026)	0.012 (0.024)	0.122* (0.047)	0.001 (0.045)
Q2	0.134. (0.076)	-0.003 (0.045)	0.084* (0.039)	0.023 (0.040)	0.077* (0.030)	0.015 (0.028)	0.143* (0.057)	-0.025 (0.056)
Q3	0.188. (0.100)	0.018 (0.034)	0.139** (0.046)	0.022 (0.045)	0.084* (0.036)	0.001 (0.032)	0.129* (0.063)	-0.010 (0.055)
Q4	0.271* (0.112)	-0.070 (0.052)	0.188* (0.078)	0.016 (0.051)	0.102. (0.055)	-0.032 (0.042)	0.262. (0.139)	-0.071 (0.088)
Transformation	log1p	log1p	asinh	asinh	asinh	asinh	asinh	asinh
Observations	726	881	1,339	1,363	2,485	1,981	704	576
FE (cik)	111	130	148	156	285	228	104	89
FE (event_id)	91	91	98	97	166	133	80	63
FE (period_qidx)	46	32	47	37	47	44	43	31

Notes: Standard errors clustered by `event_id` and `cik` are reported in parentheses. Standard abbreviations: COGS (Cost of Goods and Services Sold), N (Net), R&D (Research and Development). Significance codes: . $p < 0.1$, * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

A.3 Interpretation

Taken together, the appendix results confirm that:

- Data breaches induce a *financing and investment* response through working-capital expansion and increased operational inputs.
- Ransomware induces an *availability disruption* without financial restructuring, ultimately reducing billable output and earnings.

These patterns provide the micro-level accounting basis for the aggregate patterns observed in Table 6 it, and reinforce the conclusion that pooled cyber effects are a statistical mixture of two fundamentally distinct economic mechanisms.

A.4 Validation of Parallel Trends

A fundamental requirement for the causal interpretation of our dynamic Difference-in-Differences estimates is the parallel trends assumption. This assumption requires that, in the absence of a cyber incident, treated and control firms would have followed statistically indistinguishable trajectories in their financial outcomes. We formally test this by examining the pre-treatment coefficients (β_{-4} to β_{-2}), using Q_{-1} as the omitted reference period.

Table A3 reports these pre-attack differences for the primary variables used in the aggregate pooled analysis. Across both balance sheet and income statement items, the estimated coefficients are small in magnitude and statistically insignificant at the conventional 5% level. While a marginal difference is noted in Accounts Receivable at Q_{-4} , this gap disappears in the quarters immediately preceding the incident, confirming a stable baseline for our estimation.

Table A3: Validation of Parallel Trends: Pooled Sample Pre-Attack Coefficients

Panel A: Balance Sheet and Cash-Flow Outcomes					
	Total Assets	PPE (Net)	SG&A Expense	Accounts Rec.	Op. Cash Flow
Q-4	-0.085 (0.064)	0.018 (0.029)	0.001 (0.034)	-0.077 (0.044)	1.502 (1.336)
Q-3	-0.014 (0.017)	0.018 (0.023)	-0.006 (0.029)	-0.028 (0.043)	0.925 (1.272)
Q-2	-0.005 (0.009)	0.015 (0.017)	0.020 (0.023)	-0.035 (0.050)	-0.542 (1.147)
Panel B: Income-Statement and Liabilities Outcomes					
	Liabilities	Gross Profit	Cont. Op. Inc.	Non-op. Inc.	LTD (Curr.)
Q-4	-0.040 (0.045)	0.296 (0.407)	-0.480 (1.646)	-0.462 (2.726)	0.392 (0.752)
Q-3	0.049 (0.037)	0.013 (0.136)	1.462 (1.536)	0.872 (2.176)	-1.066 (0.798)
Q-2	-0.009 (0.015)	0.659 (0.481)	-1.922 (1.447)	-1.682 (2.169)	-0.998 (0.677)

Notes: Table reports pre-attack differences relative to Q_{-1} . All models include firm, event, and calendar-quarter fixed effects. Standard errors are two-way clustered by `event_id` and `cik`.

To further ensure that our heterogeneous results are not driven by pre-existing trends, Table A4 provides validation specifically for the variables that exhibited significant post-attack adjustments, sorted alphabetically for comparison. In Panel A (Data Breach), variables driving the capital expansion narrative show flat trends leading up to the breach. In Panel B (Ransomware), variables driving operational paralysis are statistically indistinguishable from zero prior to the attack.

Table A4: Validation of Parallel Trends: Pre-Attack Coefficients for Significant Heterogeneous Outcomes (Alphabetical)

Panel A: Data Breach (Confidentiality Attacks)				
	Acc. Payable	Accounts Rec.	Assets (log1p)	COGS (log1p)
Q_{-4}	-0.054 (0.061)	-0.082 (0.070)	-0.149 (0.119)	-0.036 (0.059)
Q_{-3}	0.005 (0.057)	-0.048 (0.074)	-0.031 (0.019)	0.001 (0.068)
Q_{-2}	-0.028 (0.037)	-0.079 (0.096)	-0.005 (0.012)	0.048 (0.051)
	Inventory (asinh)	Liabilities (log1p)	LTD (asinh)	Op. Inc. (asinh)
Q_{-4}	0.034 (0.042)	-0.089 (0.068)	1.359 (0.876)	0.307 (1.797)
Q_{-3}	0.056 (0.040)	0.013 (0.024)	-0.264 (0.699)	0.439 (1.775)
Q_{-2}	0.032 (0.027)	0.001 (0.021)	-1.484 (0.892)	-0.609 (1.391)
	PPE (Net)	Pre-Tax Inc.	R&D Exp.	
Q_{-4}	0.031 (0.034)	0.772 (2.069)	0.001 (0.054)	
Q_{-3}	0.011 (0.032)	1.847 (2.273)	0.078 (0.060)	
Q_{-2}	0.006 (0.014)	-0.128 (1.798)	0.039 (0.047)	
Panel B: Ransomware (Availability Attacks)				
	Acc. Payable	Accounts Rec.	Assets (log1p)	COGS (log1p)
Q_{-4}	-0.026 (0.056)	-0.084 (0.063)	0.001 (0.032)	0.025 (0.058)
Q_{-3}	-0.012 (0.047)	-0.021 (0.057)	0.019 (0.032)	0.042 (0.056)
Q_{-2}	-0.012 (0.037)	0.020 (0.030)	0.002 (0.015)	0.061 (0.049)
	Inventory (asinh)	Liabilities (log1p)	LTD (asinh)	Op. Inc. (asinh)
Q_{-4}	0.079 (0.046)	0.047 (0.062)	0.538 (1.190)	-1.215 (2.143)
Q_{-3}	0.075 (0.041)	0.123 (0.085)	-0.845 (1.297)	1.095 (1.886)
Q_{-2}	0.037 (0.035)	-0.009 (0.025)	-0.319 (1.091)	-0.042 (2.048)
	PPE (Net)	Pre-Tax Inc.	R&D Exp.	
Q_{-4}	0.015 (0.049)	-2.429 (2.973)	0.051 (0.079)	
Q_{-3}	0.029 (0.054)	0.562 (2.314)	-0.007 (0.064)	
Q_{-2}	0.003 (0.027)	-4.681 (2.497)	-0.049 (0.074)	

Notes: Table reports pre-attack differences relative to Q_{-1} . All models include firm, event, and calendar-quarter fixed effects. Standard errors are two-way clustered by `event_id` and `cik`. Significance codes: . $p < 0.1$, * $p < 0.05$, ** $p < 0.01$.

A.5 Sample Selection and Differential Reporting

The number of observations varies substantially across outcome variables in our main analysis, from approximately 1,345 firm-quarter observations for *Net Cash from Operations* to over 5,700 for *Total Assets*. This appendix investigates whether that variation reflects a selection problem that could bias our estimates. We address three specific concerns raised in the literature on event studies with unbalanced panels: (i) whether treated and control firms differ systematically in which outcomes they file, (ii) whether the cyber attack itself causes firms to exit the estimation sample, and (iii) whether either pattern differs across attack types.

Approach. We define the estimation universe as all 676 unique firm-events in the propensity-score matched sample (338 treated, 338 control). For each outcome variable, *coverage* is the fraction of those 676 firm-events that have at least one observed value within the ± 4 quarter window around the attack date. A treated-control gap in coverage indicates that the two groups differ in which financial items they file, a structural feature of SEC disclosure, not necessarily a threat to identification. A post-attack decline in coverage among treated firms, relative to controls, would indicate that the attack itself causes firms to drop out of the estimation sample, which could bias estimates of the financial impact.

A.5.1 Coverage by Treatment Group

Table A5 reports, for each main outcome, the fraction of treated and control firm-events for whom the outcome is observed and the treated-minus-control gap.

Table A5: Sample Coverage by Outcome Variable and Treatment Group

Outcome Variable	Treated	Control	Gap (pp)
Total Assets	100.0%	100.0%	0.0
PP&E, Net	88.2%	85.8%	+2.4
Net Cash from Operations	90.8%	95.3%	-4.5
Total Liabilities	79.3%	77.8%	+1.5
Pre-tax Income	79.6%	74.9%	+4.7
Accounts Receivable	71.6%	63.6%	+8.0
SG&A Expense	50.3%	52.1%	-1.8
Gross Profit	46.2%	39.3%	+6.9
Non-operating Income	36.7%	33.7%	+3.0
Current Long-Term Debt	30.2%	30.5%	-0.3

Notes: Coverage is the share of the 676 firm-events in the matched sample (338 treated, 338 control) with at least one observed value for the outcome over the ± 4 quarter window. Gap = Treated – Control (percentage points).

Two patterns emerge. First, coverage is far from universal: it ranges from 30% (*Current Long-Term Debt*) to 100% (*Total Assets*), confirming that the variation in N across outcomes is structural—most firms do not file every financial line item in every quarter, and the dataset

contains only rows where a value was disclosed. Second, treated firms have slightly higher coverage than controls for most outcomes (gaps of 2–8 pp), but the direction is not uniform: treated firms have *lower* coverage for *Net Cash from Operations* (–4.5 pp) and *SG&A* (–1.8 pp). There is no systematic pattern indicating that treatment status drives selection into filing.

A.5.2 Pre- versus Post-Attack Coverage Change

Table A6 examines whether coverage changes around the attack date. For each outcome we report average quarterly coverage in the pre-period (quarters –4 to –1) and the post-period (quarters 0 to +4) for treated and control firms separately. We then compute the within-group change and form the Difference-in-Differences (DiD). We also report the p -value of the treated \times post interaction from a firm-quarter logistic regression of sample inclusion on `post`, `treated`, and their interaction.

Table A6: Pre/Post Coverage Change Around the Attack—Difference-in-Differences

Outcome Variable	Treated		Control		Treated Δ	Control Δ	DiD	p (interaction)
	Pre	Post	Pre	Post				
Total Assets	97.0%	90.9%	97.3%	96.6%	–6.1	–0.7	–5.4	0.001***
PP&E, Net	81.9%	73.5%	80.4%	78.3%	–8.4	–2.1	–6.3	0.004***
Total Liabilities	76.0%	70.9%	73.9%	73.9%	–5.1	0.0	–5.1	0.024**
Pre-tax Income	53.6%	52.3%	52.4%	53.1%	–1.3	+0.7	–2.0	0.418
Accounts Receivable	67.6%	64.1%	60.8%	60.0%	–3.5	–0.8	–2.7	0.251
SG&A Expense	36.4%	34.5%	38.0%	37.8%	–1.9	–0.2	–1.7	0.505
Gross Profit	37.7%	34.4%	31.7%	29.9%	–3.3	–1.8	–1.5	0.595
Non-operating Income	25.4%	23.9%	23.7%	23.4%	–1.5	–0.3	–1.2	0.564
Current Long-Term Debt	22.9%	20.5%	23.7%	23.0%	–2.4	–0.7	–1.7	0.414
Net Cash from Operations	21.5%	20.9%	22.9%	23.1%	–0.6	+0.2	–0.8	0.731

Notes: Pre = quarters –4 to –1; Post = quarters 0 to +4. Δ = Post – Pre (percentage points). DiD = Treated Δ – Control Δ . The p -value is from a firm-quarter logistic regression of sample inclusion on `post`, `treated`, and `post \times treated`. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

For three balance sheet outcomes—*Total Assets*, *Total Liabilities*, and *PP&E*—treated firms exhibit a statistically significant post-attack decline in coverage of 5–6 percentage points, while control firms show no comparable decline ($p < 0.05$ or better in each case). This indicates that a small share of attacked firms miss or delay quarterly filings following the incident, causing them to exit the estimation sample in the post-attack period. For the remaining seven outcomes, including all income statement and cash flow variables, the treated \times post interaction is not statistically significant and coverage changes are small and symmetric across groups.

A.5.3 Heterogeneity by Attack Type

We re-run the coverage analysis separately for the 115 data breach events and the 96 ransomware events, excluding the 15 events classified as “Other.” Ransomware attacks directly

encrypt firm systems and could plausibly delay filings more than data breaches, which involve data theft without necessarily impairing operations. Table A7 reports the difference-in-differences in coverage for each attack type.

Table A7: Post-Attack Coverage Change by Attack Type—Difference-in-Differences

Outcome Variable	Data Breach			Ransomware			p (DB)	p (RW)
	Treated Δ	Control Δ	DiD	Treated Δ	Control Δ	DiD		
Total Assets	-6.7	-1.3	-5.4	-5.4	-0.8	-4.6	0.051*	0.028**
PP&E, Net	-8.9	-2.2	-6.7	-8.4	-2.6	-5.8	0.020**	0.159
Total Liabilities	-6.5	-1.0	-5.5	-3.5	+0.4	-3.9	0.075*	0.292
Accounts Receivable	-3.5	-1.3	-2.2	-3.9	-0.5	-3.4	0.522	0.324
Pre-tax Income	-1.1	+0.5	-1.6	-2.0	+1.4	-3.4	0.652	0.419
SG&A Expense	-0.6	+0.6	-1.2	-3.6	-1.4	-2.2	0.708	0.572
Gross Profit	-3.7	-1.5	-2.2	-3.1	-2.4	-0.7	0.541	0.878
Non-operating Income	-1.2	-1.6	+0.4	-2.0	+0.8	-2.8	0.881	0.410
Current Long-Term Debt	-3.3	-0.4	-2.9	-1.5	-0.3	-1.2	0.212	0.740
Net Cash from Operations	-1.3	-0.2	-1.1	+0.7	-0.4	+1.1	0.686	0.768

Notes: Δ = Post coverage – Pre coverage (percentage points). DiD = Treated Δ – Control Δ . p (DB) and p (RW) are p -values of the treated \times post interaction from separate logistic regressions estimated within each attack type sub-sample. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Two findings stand out. First, the post-attack balance sheet attrition is present under both attack types and of comparable magnitude: the DiD for *PP&E* is -6.7 pp for data breaches and -5.8 pp for ransomware; for *Total Assets* the figures are -5.4 pp and -4.6 pp. The result documented in Section A.5.2 is therefore not driven by one attack type alone. Second, contrary to the operational disruption hypothesis, data breach events produce modestly larger balance sheet attrition than ransomware events. This is consistent with the more protracted disclosure obligations and legal processes triggered by data breaches—including mandatory notification deadlines and regulatory investigations—which may create administrative burdens that delay filings independent of any direct system disruption. For income statement and cash flow outcomes, the DiDs are small under both attack types and statistically indistinguishable from zero.

A.5.4 Implications for the Main Analysis

Taken together, the three analyses above point to a clear and bounded conclusion. The variation in N across outcomes in our main results is structural: it reflects differences in the prevalence of each financial line item across SEC filings, not differential propensity to report correlated with treatment status. For income statement and cash flow outcomes, the variables that exhibit the largest economic effects in our main analysis, there is no evidence of selective attrition, and the corresponding estimates are unaffected by the concerns investigated here.

For three balance sheet outcomes (*Total Assets*, *Total Liabilities*, and *PP&E*), a small but statistically significant share of treated firms drops out of the estimation sample in the post-attack period. This attrition affects approximately 5–6% of treated firms and is present under both data breach and ransomware events. Because the firms most likely to exit the sample post-attack are those most severely disrupted by the incident, this implies that our main estimates for balance sheet outcomes *understate* the true financial impact. Our

reported coefficients for these three variables should therefore be interpreted as conservative lower bounds on the average effect of a cyber attack on firm balance sheets.