Erasmus University Rotterdam

Erasmus School of Economics

Master Thesis

Financial Economics

Cybersecurity risk and stock returns

March 2023

Miguel Vaz

Supervisor: Professor Esad Smajlbegovic

Co-reader: Professor Laurens Swinkels

The views stated in this thesis are those of the author and not necessarily those of the supervisor, second assessor, Erasmus School of Economics or Erasmus University Rotterdam.

Abstract

Using textual analysis of earnings conference calls, I construct a measure of cybersecurity risk. I compile a dictionary of cybersecurity-related terms to that end. I then study the relation between this risk and US stock returns over 16 years. I find that investors require a premium to buy high cyber risk stocks. The results are valid across different industries and over the whole sample period.

Acknowledgements

I would like to thank my supervisor, Prof. Esad Smajlbegovic, for his feedback. It was essential to overcome obstacles and move forward with my thesis. I also thank my parents, who always invested in my education.

Contents

1. Introduction 1

2. Literature review 3

2.1. Cyber risk 3

2.2. Textual analysis 4

2.3. Cybersecurity risk derived from textual analysis 4

2.4. Asset pricing 6

3. Hypotheses 8

4. Data 9

4.1. Using earnings calls as data 10

5. Methodology 10

5.1. Textual analysis 10

5.2. Cybersecurity risk measure 12

5.3. Asset pricing tests 13

6. Validation 17

7. Summary statistics 19

8. Results 22

8.1. Portfolio sorts 22

8.2. Fama-Macbeth regressions 26

8.3. Time-series 29

9. Limitations 30

10. Conclusion 32

References 34

Appendix 37

1. Introduction

In recent years, the rapid advancement of technology and increased reliance on digital systems have led to an unprecedented rise in cybersecurity risk. Particularly now due to the pandemic, corporates have a bigger digital footprint, thus are more vulnerable to these online threats. According to the World Economic Forum, the number of cyberattacks increased by 22%, substantially more than in the years before (Greenberg, 2021).

Policymakers and practitioners fear a major cyberattack in the next years. Relatively to a global financial crisis, it is now more likely to happen and can have similar systemic effects (Fung, 2021). Another development has been cyberattacking from nation-state actors. Since Russia invaded Ukraine in February 2021, we have seen cyber weapons intensifying geopolitical tensions and vice-versa (Moschetta et al., 2023).

Besides, on a more micro level, cyber threats have become a major concern for companies. There are many types of cybercrimes committed against corporations. The most common is phishing, a type of attack that attempts to trick individuals into revealing sensitive information. The list continues, from hacking that shuts down information systems, unauthorized access via poor passwords, data stealing or ransomware, when the intruders demand money to prevent an attack.

Motivated by these insights, my thesis seeks to answer the following research question: h*ow does cybersecurity risk, derived from textual analysis, affect US stock returns.*

I first face the challenge of constructing a measure for this risk. Since this is considered a latent risk, not immediately visible, it is difficult to predict and quantify. Sometimes companies just realize they were attacked weeks later. It is hard to estimate the total costs that may only be fully apparent months or years later[1].

My strategy is to establish a dictionary of terms related to cybersecurity risk and count the times they appear in the firms' conference calls each quarter. If companies' directors and analysts frequently use these terms, it indicates they are worried about

[1] Equifax, for example, experienced a massive cyberattack between mid-May and July 2017, but did not publicly disclose it until September of that year.

the potential damage a security breach or unauthorized access can cause to the firm. Hence, they are more exposed to cybersecurity risk. The majority of firms listed on a US stock exchange hold an earnings conference call each quarter. The management provides future-looking insights about financial performance, business initiatives, and their associated risks. In the end, there is a question-and-answer period, which is important to clear any doubts of the analysts and other interested parties in the call.

I confirm the validity of the measure by analysing how it evolves over time and behaves across industries. Next, I investigate the relation between stock returns and cybersecurity risk, i.e. to what extent my proxy predicts variation in the cross-section of stock returns. I use for this purpose portfolio sorts and Fama-Macbeth regressions.

I find that firms more exposed to cyber risk[2] earn higher returns. Particularly, a one-unit increase in risk increases returns by 0.1% per month, and a high minus low cyber risk portfolio earns a monthly 0.55% abnormal return.

Lastly, I further analyse my measure. A highly positive correlation with cybersecurity ETFs might indicate it captures other aspects than risk.

My thesis contributes to the literature by introducing a simple cybersecurity risk measure captured from companies' conference calls. Without the need for machine learning techniques, my proxy shows the relevance of cybersecurity risk in the stock market. Moreover, my research offers a detailed description of the parsing methods which makes it replicable. This is not always the case in finance textual studies.

The remaining of the thesis is organized as follows: section 2 reviews the relevant literature around this topic; section 3 presents the hypothesis; section 4 describes the data I employ; section 5 elaborates on the methods I use for the textual analysis and asset pricing tests; section 6 validates my cybersecurity risk measure; section 7 shows descriptive statistics; section 8 documents the results; section 9 discusses limitations of my measure; and section 10 concludes.

[2] I use interchangeably the terms "cyber risk" and "cybersecurity risk" throughout the thesis.

2. Literature review

In this section, I first look at recent literature on cyber risk in a broader context and on textual analysis in finance. Then, I cover the studies more closely related to mine which combine the latter two fields. Lastly, I review the asset pricing theory.

2.1. Cyber risk

Tosun (2021) focus on the effects on firm characteristics following data breach disclosures, using event studies. The abnormal returns around 58 first official reports on intentional attacks significantly drop. In the long term, while the market value is not affected, firm policies are. The reputational damage an attack entails makes companies decrease dividend payments and R&D expenses. Furthermore, a difference-in-difference analysis to disentangle the causal effects of a breach shows that larger firms, with higher Tobins' Q, leverage and operating profits face more negative returns. This finding is consistent with Kamiya et al. (2021).

This paper documents that hackers indeed cherry-pick their targets. Relatively to non-attacked firms, they are more visible, and more represented among Fortune 500 companies. Event study analysis indicates that only attacks involving personal data theft cause significant announcement returns. Although different results than Tosun (2021), the authors use a larger and so less restricted sample of cyberattacks.

To understand the impacts on the supply chain, Crosignani et al. (2023) investigate a particular cyberattack named NotPetya, considered the most damaging in history. The Russian military targeted an accounting software to try to paralyze the Ukranian economy[3]. Affected firms, both customers and suppliers, recorded significantly lower profits relative to similar non-affect firms.

Gambacorta et al. (2022) draw attention to the difficulty of estimating cyber risk, naming it a "known unknow tail risk". Therefore, the insurance market cannot be sufficiently solid. Firm size is again positively correlated with cyber costs; using cloud services and IT spending are negatively correlated. The financial sector, an attractive

[3] Experts believe that it was a politically motivated cyberattack with ties to the conflict in Crimea.

target for hackers due to their critical importance, invests more in its IT systems and might be more experienced in facing cyber risk.

2.2. Textual analysis

I add to the literature in finance using text as data, still a recent and developing field. Loughran & McDonald (2016) provide an updated literature review on this topic. While it is clear that investors not only incorporate quantitative data in the stock market, using qualitative information poses various challenges. First, extracting text should not be extracting a collection of characters, but the actual information which depends on context, word sequence or type of document. Another source of imprecision comes from the lack of consistency in the text formatting. It is likely that the structure, of XML files in my case, is correlated with firm size and time period.

My thesis uses "bag-of-words" techniques, which involve creating a list of words that share sentiment and then counting those words in each document. This is also the method used by Jamilov et al. (2021). I implement a slight variation, by using a different term-weighting scheme, the term frequency-inverse document frequency. Florackis et al. (2022) in turn, compute the cosine similarity between two vectors of words, to measure the semantic similarity between two texts collapsed into these vectors.

2.3. Cybersecurity risk derived from textual analysis

Florackis et al. (2022), Jiang et al. (2020), Jamilov et al. (2021), and Lhuissier & Tripier (2021) derive cybersecurity risk using textual analysis.

Florackis et al. (2022) analyse the item 1A Risk Factors from the 10-K disclosure of US-listed companies, looking to extract keywords and sentences associated with cybersecurity. The goal is to construct a measure of this risk. To that end, the authors use a training sample. It isolates firms that ex-ante had a high exposure to cybersecurity risk and were subject to a major attack. Then they calculate the similarity between each firm's cyber risk disclosure and past disclosures in the training sample.

As Jiang et al. (2020), the paper finds that the stock market prices the cybersecurity risk in the cross-section of returns, also in the long term. Both studies document a price of risk and alphas not explained by the traditional asset pricing models. Specifically, the portfolios that invest in high-risk firms and shorts low-risk firms earn a statistically significant return, of around 9% over the following year.

This premium may be due to compensation for hedging against the consequences of a catastrophic hack, but it likely captures other risks besides cybersecurity. However, Florackis et al. (2022) show the premium is not driven simply by firms in the technology sector, which during the sample period face higher cyber risk and outperform the market, at the same time.

They conclude with an out-of-sample event study of the Solar Winds hack, in 2020. The results show that firms with higher ex-ante exposure to cybersecurity risk exhibited negative cumulative abnormal returns (CARs) around the event, with the negative association being stronger for more highly exposed firms. They further find that among companies that have higher ex-ante exposure to cybersecurity risk, SolarWinds customers earn negative CARs, as opposed to non-affected companies which earn positive CARs.

Jamilov et al. (2021) take the quarterly earnings conference calls to identify relevant terms related to cyber risk. The authors argue they contain more and richer information than the 10-K disclosures used by Florackis et al. (2022). The firm-level exposure is measured as the number of mentions of over 30 unique terms, normalized by the number of words in each transcript. The method I use for the measure construction closely follows this. Yet, their sample is larger, covering 85 countries over 18 years. They find evidence that this firm-level risk can spread through companies in the same country and industry, causing spillover effects. Even if a company has zero cyber risk exposure, others' higher exposure is associated with negative returns.

Their approach to study the asset pricing implications is somewhat different from mine, and the ones of Jiang et al. (2020) and Florackis et al. (2022). The authors consider that firm-level cyber risk has a factor structure. So, they construct a novel pricing factor, compute its betas and use them to sort stocks into portfolios.

In addition, they report that cybersecurity ETFs reflect more market risk and the conventional size factor rather than exposure to cyber risk. As such, they cannot be used as a hedge against cybersecurity risk.

Jiang et al. (2020) have a different view on cybersecurity ETFs. The authors assume they should act as a hedge. The returns of their cyber risk high-minus-low portfolio are negatively correlated with the ETFs' returns. This validates their measure, showing that it is actually capturing the market's view on cyber threats. Unlike Jamilov et al. (2021), but similar to Florackis et al. (2022), the focus is to examine the cross-section of stock returns. In turn, the methodology differs. To obtain a cyber risk proxy, they apply a variety of machine learning techniques to estimate the ex-ante probability of a firm being attacked. Further, they analyse institutional investors' demand and conclude they buy low cyber risk stocks and sell those which score high.

Lastly, Lhuissier & Tripier (2021) rely on *tweets* containing a combination of the term "cyber" with "risk", "attack" or "threat" to build a cyber risk index. This relatively new way of data gathering[4] does not follow the bottom-up approach of the previous studies to identify cybersecurity risk at the firm-level. Nevertheless, it has advantages. It provides daily and real-time tracking while covering the entire economy rather than listed companies. They focus on the benefits for cybersecurity companies that may arise from higher cyber risk. Although the impact on returns is not significant, it is for the COVID-19 period. As Lallie et al. (2021) document, the lockdown and working from home seem to have increased not only the number of cyberattacks but also, according to Lhuisser et al. (2021), the demand for the services provided by the cybersecurity industry. It represents the so-called *hack effect* (MarketInsite, 2019).

2.4. Asset pricing

My thesis lies on the traditional asset pricing theory, i.e. on the risk-return relationship. I use the models presented in this subsection for the asset pricing tests. The more risk investors take the more they should be compensated by the returns on their portfolios. This idea is present in the CAPM, a model developed by Sharpe (1964), which assumes a single risk factor, the market risk. This factor captures the

[4] Twitter offers a product track of Twitter API for researchers since 2021.

common movements in prices not specific to individual companies. According to the CAPM, this factor is enough to capture all systematic risk.

Later, Fama & French (1993) extend the CAPM with two more risk factors, size and value. The Fama-French 3-factor model (FF3) offers a better description of the stock market. It is able to capture anomalies found in the CAPM. The first, *SMB* (small minus big market capitalization stocks), aims to capture the superior performance of small companies relative to large ones. One explanation is that the size factor reflects liquidity differences. Small caps are more difficult to quickly buy and sell without affecting the price. This may create opportunities for patient investors to earn higher returns. However, the size effect seems to have disappeared since the mid-1980s. Black (1993) argues it actually never existed. It was found due to data mining. Schwert (2003) in turn, say investors eliminate the premium by investing in small stocks, bidding up their prices.

On the other hand, the value effect keeps producing gains. It is named HML, which stands for high minus low book-to-market stocks. Value stocks tend to outperform growth stocks and the explanations are either rational or behavioural. First, value is riskier because it is associated with companies that are in distress or facing challenges, hence more vulnerable to economic downturns. Alternatively, during such periods, these firms have more trouble shifting their activities to new profitable ones (Zhang, 2005). As for behavioural reasons, Lakonishok et al. (1994) document that investors overreact to positive news about growth stocks, driving up their prices. At the same time, they underestimate value stocks' growth opportunities.

Jegadeesh & Titman (1993) find the momentum factor. Its abbreviation, UMD, stands for stocks that have gone up minus stocks that have gone down. Later, Carhart (1997) added it to FF3, a model known as FFC (Fama-French Carhart). So, the investment strategy consists of buying winners from the past 6 to 12 months, and short losers over the same period. While it is one of the most robust and persistent anomalies, it is difficult to explain economically. Researchers mostly cite behavioural reasons. Investors may overreact slowly to news and persistently increase prices. Or, another possible explanation, they may underreact and keep rising prices as they learn.

More recently, Fama & French (2015) include a profitability and an investment factors, making it a five-factor model (FF5). The authors argue they might assume the role of the earlier FF factors. In my analyses, I use both this model and the FFC.

RMW is robust minus weak profitability stocks. More profitable companies tend to outperform less profitable companies, even after accounting for size and value.

CMA stands for conservative minus aggressive stocks, meaning high and low investment firms. Firms that invest more tend to face a higher risk of financial distress, which leads to higher expected returns. The results of this paper indicate these two additions take the role of HML.

3. Hypotheses

My hypotheses are as follows:

H1: cybersecurity risk is priced; in other words, expected returns increase from low to high-cybersecurity-risk stocks.

The literature on the relationship between cybersecurity risk and returns using a textual measure finds that investors require a premium to invest in highly exposed stocks. The returns of portfolios sorted according to this variable increase with the risk in Florackis et al. (2022) and Jiang et al. (2020). This result is confirmed as well for individual stocks in Fama & MacBeth (1973) regressions, which control for other factors that influence returns.

Highly exposed stocks may incur losses because of cyber incidents. Some of the consequences are costly remediation expenses, reputational damage, loss of intellectual property, and regulatory fines. Additionally, they can disrupt business operations, leading to lower profitability and reduced investor confidence. These risks cause uncertainty and volatility in the stock price, which leads investors to demand compensation.

My results should not be only driven by specific industries or time periods. Companies that belong to the technology sector inevitably mention more cyber-related terms in their conference calls. So, they should have a relatively high cyber risk score

and may be the ones responsible for the cyber risk-returns relation I find. In addition, the emergence of the COVID-19 pandemic has led to a historic acceleration in cyber risks. Thus, I might observe a regime shift in my time-series data which can affect my main relation of interest.

H2: the cybersecurity risk premium cannot be explained by other common risk factors.

Both Florackis et al. (2022) and Jiang et al. (2020) find annualized alphas of around 9%, depending on the risk model, in a portfolio long on high cybersecurity stocks and short by the same amount on low cybersecurity stocks. Thus, I expect that a zero-investment strategy in this portfolio, with stocks sorted by my proxy, generates abnormal returns.

This suggests that the portfolio has unique characteristics that are not captured by the traditional risk factors. These unique characteristics should not be idiosyncratic cybersecurity risk or other company-specific factors.

4. Data

I retrieve data from CRSP for US stock returns, Compustat for financial information and Yahoo Finance for ETF prices. The pricing factors are from Kenneth's French website. I download the conference calls transcripts from Thomson Street Events.

My sample period goes from 2005 to 2021 and covers 5,150 unique companies listed on the NYSE, NASDAQ and AMEX. Prior to 2005, the transcripts available are limited. In addition, the files' structure lacks consistency (Loughran & McDonald, 2016).

I compile a list of words related to cybersecurity based on three different sources, for the sake of robustness. These are: Financial Stability Board (international association), National Cyber Security Centre (UK) and Cybersecurity and Infrastructure Security Agency (US). These credible institutions have their dictionaries publicly available. My idea is to start with a broad pre-defined dictionary of words and then refine it based on certain conditions.

4.1. Using earnings calls as data

For the textual analysis, there are advantages in using earnings calls over, for example, 10-K disclosures. Earnings conference calls typically happen four times per year, while 10-K reports are released annually. So, they provide more up-to-date information on a company's performance and risk factors.

During an earnings call, analysts can ask questions and provide additional insights into a company's performance and risk factors that may not be captured in a 10-K. Cyberattacks are underreported, so these Q&As may force managers to reveal more information.

In fact, cyberattack disclosures by public companies, about 300 from 2010 to 2015 seem limited compared to the thousands of reports from outside sources (Amir et al., 2018). The more severe the attacks the more likely underreporting is.

Furthermore, there is often a considerable time lag between the occurrence, its acknowledgement, and its disclosure. The massive data breach at Equifax was first reported in September 2017, but the hackers had access to sensitive information, such as Social Security numbers and birth dates, from mid-May to July. The SolarWinds supply chain attack was discovered in December 2020, but it is believed to have started as early as March. The Uber data breach is another example of a cyberattack not being reported on time. It occurred in October 2016 and involved the personal information theft of 57 million customers and drivers. However, the company paid the attackers $100,000 to destroy the stolen data and kept the breach a secret for over a year. The breach was finally made public in November 2017.

5. Methodology

5.1. Textual analysis

In this subsection, I provide a detailed description of how I define my cybersecurity risk dictionary and parse the corporate calls' transcripts.

First, I need to avoid homograph words, i.e. words with other meanings besides the one cybersecurity-related. For example, "confidentiality" or "compromise". This is a

challenge faced by textual analyses. While for humans distinguishing meanings in context is an easy task, for a computer it is not (Loughran & McDonald, 2016).

Yet, I keep terms like "cloud" or "cookie". Although those may possess alternative connotations, their usage in the context of a conference call in any other sense would be odd. As such, the amount of noise they produce is negligible. I manually double-check those transcripts where "cloud" is most frequently mentioned. Its usage soars, from 232 times in 2005 to 16,516 in 2020. It would not make sense then that the meaning conveyed concerns the masses in the atmosphere.

Sometimes the meaning can very likely be either the one I am interested in or another. But if it is a relevant term that should not be left out, another word must immediately follow. For instance, "data" has a broad definition, so I only count bigrams such as "data security", "data loss" or "data leak"[5].

Next, I must search for word roots instead of full words, although the English language has little inflection, i.e. a tendency to have words identical to their roots. This ensures I count any inflections of the word roots. For example, instead of extracting individually "hacker", "hacked", and then probably forgetting some inflections, I look for "hack" followed by a space, to allow for any inflections[6].

In addition, there are companies whose names include terms from my list. Company names are often written in the transcripts, so not taking this into account leads to biases. For instance, firms like "Check Point *Software* Technologies", "*Cloud*flare" or "*Cyber*Ark Software".

To extract data from the XML[7] files I use a pre-packaged solution, namely a Phyton library. For each transcript, I record how many times each word in my dictionary is mentioned. The final list of 126 terms is in the Appendix.

[5] "Cyber", "digital", "identity" and "threat" are other examples.

[6] The parse is not case sensitive, so the distinction between proper and common nouns is not a concern.

[7] XML is a type of file format used to store structured data. It uses a set of tags, specified by the file creator, to define the structure and meaning of the data.

5.2. Cybersecurity risk measure

Next, I define three possible measures of cybersecurity risk.

For my baseline measure, Cyber1, I use a common term weighting scheme from the literature on textual analysis, the term frequency-inverse document frequency (tf-idf). Instead of treating all words equally, tf-idf assigns more weight to words that are relatively rare in the overall corpus.

Let N be the total number of transcripts in my sample, df_t the number of transcripts containing term t; $tf_{t,d}$ represents the count of term t in document d, whereas a_d is the average word count in document d.

If $tf_{t,d} > 0$, tf-idf is defined as:

$$tf\text{-}idf = \frac{\left(1 + log\left(tf_{t,d}\right)\right)}{(1 + log(a_d))} \log\frac{N}{df_t}$$

Next, Cyber2, is simply the sum of the counted terms, at the transcript-level. Cyber3 is equal to the sum divided by the total number of words. This proxy thus accounts for the length of discussions. To count the total number of words in the transcripts, I exclude irrelevant bits like intervenient names and titles.

Formally, for transcript i, quarter t and where N_{it} is the total words:

$$Cyber = \frac{1}{N_{it}} \sum_{c} CyberTerm_{it}^{c}$$

For the non-normalized version, Cyber2, N_{it} is zero.

Loughran & McDonald (2011) argue this technique produces regressions with better-fit results than the simple sum of terms or the proportion. For portfolio sorting using one or another is not likely to have an impact, since the variable is just used to rank stocks. The scenario changes for the Fama & MacBeth (1973) regressions, as one of the three is used as the dependent variable.

The maximum weight assigned to a term represented by $\log\frac{N}{df_t}$, the idf component of the equation above, is 5.24. The minimum is 0.68.

To smooth out any fluctuations in the data and provide a more stable estimate of the risk level, it is reasonable to use a moving average for my cybersecurity risk measure. I use a rolling average of the most recent 4 periods.

To match the quarterly data with monthly returns data, I extend my initial dataset. The procedure I employ is analogous to the conventional practice in finance for annual variables, such as the book value. For instance, if *Cyber1* is 1 in month t for firm i, the entries for the months up until the next conference call, say t+4, also take the value of 1. Sometimes transcripts are missing, and the gap is larger than the expected 4 months. In such cases, I still take the value of my proxy for months t+5, t+6, until t+12 at most. Therefore, my dataset is now at the monthly-firm level, ready for asset pricing analysis.

The distinct company identifier codes challenge the merging of the two datasets While the transcripts data employs RIC codes (Refinitiv Identifier Code), CRSP uses PERMNO. To address this issue, I use CUSIP as the linking code.

5.3. Asset pricing tests

Portfolio sorts

I start with univariate portfolio sorts to examine whether exposure to cybersecurity risk can predict future excess returns. Fama & French (2015) use this approach. I rank individual stocks according to one characteristic, cybersecurity risk. At each month I group stocks into five portfolios. I track the performance of each portfolio until the next month. Then, the portfolios are rebalanced. Portfolio 1 has the lowest exposure to this risk, while 5 has the highest.

The flexibility of this procedure allows me to choose non-evenly spaced percentiles for the breakpoint calculation (Bali et al.,2014). This choice is due to three reasons.

First, if I split the sample into just three portfolios, the dispersion of the sort variable is lower across the portfolios. It is more difficult to detect the cross-sectional relation between cybersecurity risk and returns. Second, and crucially, the sort variable is heavily skewed, as seen in the descriptive statistics. For some time periods the breakpoints for both the 20^{th} and 40^{th} percentiles are zero, so no firms would fall under portfolio 2 in the five even breakpoints scenario. Third, in any case, I am more

interested in the upper tail of the statistical distribution of my cybersecurity risk proxy. This is where cyber risk resides. As Crosignani et al. (2023) highlight, it is a tail risk. It is unlikely to occur but, if it does, has a significant impact on the value of an investment. With uneven breakpoints, I gain a deeper understanding of how extreme values of cybersecurity risk affect stock returns.

At the end of the month, I compute equal and value-weighted returns of these 5 portfolios[8] and the difference portfolio. The latter consists of going long on portfolio 5 and short portfolio 1 by the same amount, representing a zero net investment strategy.

Value-weighted returns are more appropriate when it comes to stock analysis. Since this approach uses market capitalization to weight returns, it reflects better the actual composition of the market and the returns from the investor perspective. Equal-weighted returns may assign large weights to low-cap stocks, which are less liquid, hence more costly to trade.

To test hypothesis 1, I examine whether the time-series means of the 6 portfolios, especially the high-minus-low, is significantly different from zero. In other words, whether they generate alphas. This means there is a cross-sectional relation between my proxy and future returns. More specifically, I want to see if this relation persists after adjusting for systematic risk factors.

So, I run time-series regressions of these portfolio returns against common models of risk adjustment, the CAPM, FF3, FFC and FF5. The three factors are the market risk, size and value. Secondly, I use the Fama-French (2015) model, also an extension of the FF3 model. It additionally includes a profitability and investment policy factors.

Next, I analyse bivariate independent-sort portfolios. In addition to ranking stocks by their cybersecurity risk score, I use other three sort variables, namely, size, Tobin's Q and tangibility. For the second variable, I calculate a single breakpoint, the median, which then divides the sample into two groups, at each time period. Each portfolio is the intersection of the sets formed with both sorts. Thus, I compute the returns for 10 portfolios, plus two zero-cost portfolios[9]. The rest of the exercise is identical to the univariate sorts.

[8] These portfolios consist of long positions in each of the stocks included.

[9] Those represent the difference between the high and low cybersecurity risk groups, within the two groups of the second sort variable.

I use these two variables to investigate their relation with cybersecurity risk, as it seems evident in the literature. By comparing portfolios' returns based on different levels of the second variable, I can gain insight into the nature of that relationship and how it affects stock market returns.

Fama-Macbeth procedure

Portfolio sorting can throw away relevant information in the cross-section of returns. To get a more complete picture of the relation between stock returns and the risk factors driving them, in particular cybersecurity risk, I perform the Fama-Macbeth (FMB) regression analysis. In comparison with portfolio sorts, it allows the inclusion of multiple variables as controls. One disadvantage is, however, assuming a functional linear form of the cross-sectional relation between the variables analysed.

Before starting with the actual method, I run rolling-window time-series regressions to find market betas. This involves estimating a regression model over a rolling window and then updating the regression coefficients at each time period. Following common practice, I use a window of 24 monthly excess returns for this estimation. This approach allows for the beta estimates to vary over time, which helps to capture changes in the risk exposures of the individual assets.

I use the FMB two-step procedure to test both H1 and H2. In the first stage, I run cross-sectional regressions at each time t in my sample. I use four different specifications, by adding variables to the previous specification.

In the second stage, I compute the average of the cross-sectional slope coefficients, i.e. individual risk premiums, estimated in these individual monthly regressions to obtain a final estimate of the factor risk premium.

The goal is to test whether these coefficients are statistically distinguishable from zero. If so, the exposure to cybersecurity risk predicts a cross-sectional variation in stock returns. To that end, I calculate standard errors and t-statistics. To account for the potential presence of autocorrelation and heteroskedasticity in the error terms, I use Newey & West (1987) standard errors.

Further tests

I conduct a couple more tests, for the sake of robustness and exploring other possibilities.

First, I suspect that there is a structural break in my data which corresponds to the start of the COVID-19 pandemic. This can be visually observed in Figure 1 and confirmed by various papers and news, as I previously stated.

The question is whether I can improve the FMB regression model by splitting the model and fitting two separate subsamples. If the intercept and slope parameters change at some time in my sample, the model suffers from a structural break. Essentially, I run a comparison between two models, an unrestricted model, which contains all variables, versus a restricted model where I impose the null hypothesis, setting the coefficients equal to zero.

To split my data into two subsamples, I create a dummy variable D_1 equal to one if t is after January 2020. Then I interact it with all X variables and I test the joint significance of all these variables, dummy and slope-dummy terms. If I reject the null that all these coefficients are zero, there is a significant structural break. This is a method known as the Chow test.

Next, my cybersecurity risk measure perhaps suffers from biases due to the influence of technological companies. To test this, I create another dummy D_2 equal to 1 if the firm belongs either to the Software Services or the Technology Hardware Equipment GGROUPs.

As a concluding remark, I use the one-month-ahead excess returns in all the time-series regressions. This is a common practice in the finance literature. Investors take time to incorporate news about the market or about a company's financials.

6. Validation

In this section, I run different tests to validate my measure, in the fashion of Florackis et al. (2022). First, I check the time-series trend. In line with Jamilov et al. (2021)., cyber risk over the years increases by roughly five times from its lowest point in 2008 to the highest in 2020, as Figure 1 below indicates. The largest jumps in the data occur in February, March, and April 2020, coinciding with the pandemic period.

According to recent news and reports this is indeed what happened (Lallie et al., 2021). It created a perfect storm of new vulnerabilities for cybercriminals. Online activity surged as it had never before because people had to stay at home and work remotely. The IT systems' security could not follow this abrupt change. Moreover, there were several scams with COVID-19 as a theme, such as phishing and malicious emails, targeting companies' employees (Zammost & Schlesinger, 2020).

Figure 1: Cybersecurity Risk by month

This figure shows the average estimate of cybersecurity risk in each month of my sample, from January 2005 to December 2022.

Second, in Figure 2 I observe how my proxy behaves across the 24 GICS industry groups (GGROUPs)[10]. Again, in line with Jamilov et al. (2021) and Florackis et al. (2022), the groups that have the highest average scores are Technology Hardware and Equipment and Software Services. One expects companies that specialize in a particular area to have more detailed discussions about risks related to their field of expertise. Still, these alone cannot account for the high overall score.

Figure 2: Cybersecurity risk across industries
This figure displays the average estimate of my proxy across GGROUPS.

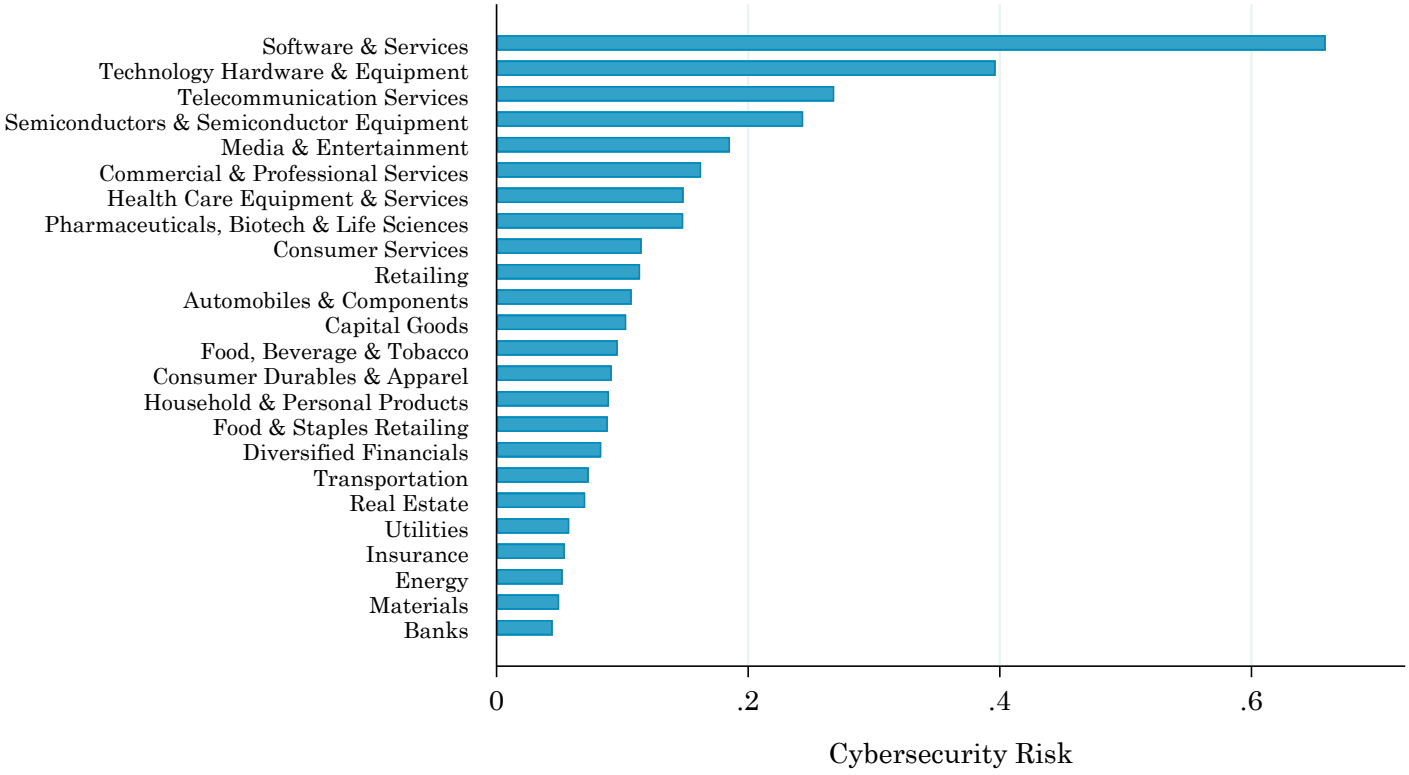

Industries that rely more on technology for their operations and business processes are more susceptible to cyberattacks that disrupt systems and compromise sensitive data. Technology companies typically have complex IT infrastructures, with multiple interconnected systems, applications, and devices. These are more difficult to secure and more vulnerable to cyber threats. Besides, technology companies are often at the

[10] The Global Industry Classification Standard (GICS), created by MSCI Inc. and S&P Dow Jones Indices, is a common global classification standard used by market participants.

forefront of innovation, working with cutting-edge technologies that introduce new security risks. Figure 2 shows this is indeed what happens with my proxy.

With the measure constructed and validated, I can move on to study the effect of this risk on the stock market.

7. Summary statistics

I start this section by documenting statistics of my cybersecurity dictionary, in Table 1. "Software" is the most mentioned term in my list; 184,413 times over 37,013 transcripts. "Cloud", "hardware", "protocol", "app" and "virus" come after accounting for over 20 thousand mentions each. Among the least said, just once over the whole sample, are "dictionary attack", "key bundle" or "malvertising".

Table 1: Top-mentioned terms

This table documents the 20 most mentioned terms from my dictionary in the companies' conference calls, the number of times they appear in the transcripts, and the resulting weight assigned, represented by the inverse term frequency. The "cyber" count excludes all other terms that contain the word "cyber", such as "cybersecurity".

Terms	Total	Transcripts	ITF
Software	184,413	37,013	0.68
Cloud	137,760	18,346	0.99
Hardware	54,134	16,824	1.02
Protocol	27,180	12,796	1.14
App	22,345	8,640	1.31
Virus	15,371	6,325	1.45
Cyber	15,046	4,243	1.62
Domain	13,030	6,205	1.46
Hack	9,149	1,532	2.06
Operating system	8,250	4,716	1.58
Digital transformation	7,691	3,025	1.77
Firewall	4,831	1,444	2.09
Router	4,654	1,936	1.96
Confidentiality	3,406	2,870	1.79
Internet of things	3,096	1,719	2.01
Cybersecurity	2,987	1,290	2.14
Cookie	2,885	1,339	2.12
Encryption	2,847	1,013	2.24
Credentials	1,995	1,258	2.15
Trojan	1,793	342	2.71

It is also interesting to note substantial upward trends of some terms. Besides "cloud" already stated in the previous section, "digital transformation" goes from 5 in 2005 to 2,222 in 2020. "App", "critical infrastructure" and "ransomware" follow similar paths.

In Table 2 I present summary statistics of my baseline cybersecurity risk measure, Cyber1, but also the other two possible ones, Cyber2 and Cyber3, after merging my data with CRSP and Compustat. To address issues of non-normality and extreme values, I apply logarithmic transformations to size, book-to-market and R&D expenses, which otherwise show highly skewed distributions.

The percentiles and the positive skewness of the cybersecurity variables indicate that most of its values are zero or close. So, most of the variation occurs in the upper tail of the distribution.

Table 2: Descriptive statistics

Panel A shows the descriptive statistics of the cybersecurity risk variables after the 4-month moving average is applied. Cyber2 is multiplied by 100. Panel B presents statistics of several financial variables. Size, book-to-market and R&D expenditure are logarithmized.

Panel A: Descriptive statistics of cybersecurity risk variables

	Mean	*Stdev*	*P50*	*P70*	*P80*	*P90*	*Max*	*Skew*	*Kurt*
Cyber1	**0.17**	**0.34**	**0.04**	**0.26**	**0.47**	**1.59**	**7.89**	**4.97**	**42.85**
Cyber2	3.00	8.57	0.5	1.5	3	7.25	261.5	6.74	71.23
Cyber3	0.38	1.02	0.048	0.19	0.39	0.96	20	5.84	50.73

Panel B: Descriptive statistics of financial variables

	Mean	*Stdev*	*P1*	*P25*	*P50*	*P75*	*P90*	*Skew*	*Kurt*
Size	6.99	1.99	2.51	5.66	6.98	8.29	11.82	0.7	2.94
Book-to-market	-0.86	0.88	-3.53	-1.33	-0.77	-0.28	0.35	-0.85	5.94
Tobin's Q	2.05	1.79	0.68	1.11	1.49	2.27	9.16	6.19	126.9
Tangibility	0.21	0.23	0.00	0.04	0.12	0.30	0.61	1.39	3.99
ROA	-0.01	0.21	-0.86	-0.01	0.02	0.69	0.12	-7.05	215.1
ROE	-0.89	4.66	-2.93	-0.02	0.81	0.15	0.25	-114.4	16621
R&D expenditure	3.37	1.89	-1.14	2.15	3.37	4.46	5.79	0.13	3.73

Palo Alto Networks in 2019 Q4 records the maximum cyber risk score. This is a natural outcome given that it is a leading provider of cybersecurity products, such as firewall technology, cloud security and endpoint protection. If I zoom in at the 99th percentile of the distribution I mostly find firms belonging either to the Software and Services or Technology Hardware and Equipment GGROUPs, as Figure 2 suggests. One exception is Equifax Inc., a consumer credit reporting agency victim of a major data breach in 2017.

Regarding correlations, reported in Table 3, I highlight the almost zero correlation between my measure and size. It is an unexpected result, as evidence suggests that larger companies are more exposed to cyber risks. Hackers tend to view them as high-profile targets, and they likely hold more sensitive data such as customer information, employee records or proprietary business information (Kamiya et al., 2021). Size can also be correlated because smaller companies' transcripts tend to lack structure. So, when parsing, I potentially lose records of these firms because crucial information, like date or company identifier, is missing. It does not seem that is the case.

Table 3: Correlations

This table shows the correlation coefficients between my baseline measure (Cyber1 in table 2) and financial characteristics. * indicates statistical significance at the 1% level.

		(i)	*(ii)*	*(iii)*	*(iv)*	*(v)*	*(vi)*	*(vii)*
(i)	*Cybersecurity Risk*	1.00						
(ii)	*Size*	0.08 *	1.00					
(iii)	*Book-to-market*	-0.19 *	-0.28 *	1.00				
(iv)	*Tobin's Q*	0.18 *	0.13 *	-0.69 *	1.00			
(v)	*Tangibility*	-0.17 *	0.10 *	0.09 *	-0.13 *	1.00		
(vi)	*ROA*	-0.03 *	0.32 *	0.04 *	-0.14 *	0.09 *	1.00	
(vii)	*ROE*	-0.00	0.06 *	0.06 *	-0.02 *	0.01 *	0.17 *	1.00
(viii)	*R&D expenditure*	0.19 *	0.76 *	-0.21*	0.06 *	-0.06 *	0.14 *	0.02 *

Companies that hold more intangible assets, like intellectual property or software, are more prone to cyber threats, as the negative correlation for tangibility suggests.

Moreover, there is a positive correlation with R&D expenses. Intuitively, R&D investments can help cyber-riskier companies develop new solutions that detect and prevent cyberattacks. This result should be interpreted with caution nevertheless, as only approximately half of the observations of the total sample have non-missing values.

8. Results

Before performing the asset pricing tests, I drop observations whose price is below 5 and share code is not 10 or 11. Penny stocks are more speculative and less liquid than higher-priced stocks, hence more susceptible to manipulation and other distortions. Additionally, low-priced stocks are often associated with small firms, more volatile and less well-established than larger firms.

Other share codes than 10 and 11 are used for different types of securities, such as preferred stock, warrants, or options. By dropping observations with those, I analyse a more homogenous set of assets and reduce the impact of outliers or unusual securities. To this end as well, I winsorize the upper 1% of the cybersecurity variable distribution, over all firms and time periods.

8.1. Portfolio sorts

Since it is constructed at the firm-level, it is possible that my cybersecurity risk measure, is capturing idiosyncratic risk which is not priced by investors. In other words, risks that are specific to individual firms and not related to systematic risk factors, priced in financial markets.

To alleviate this concern, it is important to control for other firm-specific factors that can be driving the relationship between cybersecurity risk and stock returns. I include control variables such as firm size, profitability, and leverage in the Fama-Macbeth regressions to account for these factors. As for portfolio sorts, the analysis is at the portfolio-level. Most portfolios include 150+ stocks from different sectors. As for

the bivariate sorts, a few exceptions stay above 60 assets. Thus, idiosyncratic risk is diversified away and this is not a concern.

Univariate sorts

Table 4 presents the results of the analysis conducted on the relationship between my cybersecurity risk proxy and future (1-month) portfolio returns, as well as some firm-level characteristics.

Panel A shows the future portfolio returns sorted by my cybersecurity risk proxy, with the portfolios ranging from low-cyber risk stocks (P1) to high-cyber risk stocks (P5), and the zero-cost strategy, long on P5 and short on P1 (P5-P1). The results are presented for four different asset pricing models: CAPM, Fama-French 3-factor model (FF3), Fama-French-Carhart 4-factor model (FFC), and Fama-French 5-factor model (FF5).

The table indicates that for all models, returns significantly increase from the low-cyber risk portfolio (P1) to the high-cyber risk portfolio (P5), supporting hypothesis 1. For the value-weighted portfolios, the CAPM alpha increases from 0.64% to 1.21% per month. Controlling for other risk factors does not affect the significance of the findings. The spread portfolio also earns positive abnormal returns, confirming hypothesis 2. It indicates that investors are willing to pay a premium for holding stocks highly exposed to cybersecurity risk. It also implies a cross-sectional relation between my measure and returns.

If stocks are equally weighted, the zero-net investment portfolio does not earn statistically significant returns. The results of value-weighted portfolios are more reliable though. This approach considers the importance of each stock in the overall market. It is closer to a portfolio of a real investor, who would not invest relatively as much in low market-cap stocks, usually less liquid and more costly to trade.

Panel B presents the number of stocks for each average portfolio and firm-level characteristics, namely cybersecurity risk, book-to-market, size, return on assets (ROA), return on equity (ROE), tangibility, Tobin's Q, and R&D expenditure. Except for cyber risk, which increases from P1 to P5, other characteristics are rather stable. P1 has a substantially higher number of stocks because of the uneven breakpoints I choose and explain in subsection 5.3.

Table 4: Cybersecurity risk sorted portfolios

Panel A shows average excess returns (return in excess of the risk-free rate) adjusted by the CAPM, Fama-French 3 and 5-factor models and Carhart's 4-factor. Starting in January 2005 I rank stocks at the end of each month and allocate them into 5 portfolios. I compute their returns weighted by each stock's market capitalization (vw) and weighted equally (ew). Newey & West (1987) t-statistics are reported in brackets. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

Panel B reports the average number of firms in each portfolio, as well as firm characteristics equally weighted.

Panel A: Future (1-month) portfolio returns sorted by my cybersecurity proxy

		Portfolios					
		[P1]	[P2]	[P3]	[P4]	[P5]	[P5]-[P1]
CAPM alpha	vw	**0.64****	**0.69****	**0.71****	**0.98*****	**1.21*****	**0.56*****
		(2.02)	(2.28)	(2.13)	(3.10)	(3.43)	(2.65)
	ew	**0.76***	**0.78***	**0.74***	**0.86****	**1.00****	**0.24**
		(1.82)	(1.91)	(1.86)	(2.16)	(2.50)	(1.41)
FF3 alpha	vw	**0.65****	**0.69****	**0.71****	**0.96*****	**1.16*****	**0.51****
		(2.03)	(2.24)	(2.09)	(3.00)	(3.29)	(2.41)
	ew	**0.75***	**0.76***	**0.71***	**0.82****	**0.97****	**0.22**
		(1.79)	(1.86)	(1.77)	(2.03)	(2.39)	(1.25)
FFC alpha	vw	**0.68****	**0.71****	**0.74****	**0.98*****	**1.20*****	**0.52****
		(2.11)	(2.31)	(2.19)	(3.07)	(3.41)	(2.47)
	ew	**0.77***	**0.79***	**0.74***	**0.85****	**0.102****	**0.25**
		(1.83)	(1.93)	(1.85)	(2.12)	(2.55)	(1.46)
FF5 alpha	vw	**0.70****	**0.74****	**0.75****	**1.06*****	**1.24*****	**0.55****
		(2.12)	(2.35)	(2.16)	(3.25)	(3.45)	(2.50)
	ew	**0.86****	**0.89****	**0.83****	**0.93****	**1.10*****	**0.25**
		(1.99)	(2.12)	(2.01)	(2.26)	(2.67)	(1.38)
Panel B: Firm characteristics							
Number of firms		911	356	181	181	180	-
CyberE3		0.01	0.10	0.21	0.36	0.90	-
Book-to-market (ln)		-0.80	-0.99	-1.08	-1.16	-1.27	-
Size (ln)		7.34	7.50	7.48	7.44	7.77	-
ROA		0.03	0.02	0.01	0.01	0.02	-
ROE		0.07	0.02	0.01	-0.05	0.02	-
Tangibility		0.26	0.21	0.19	0.16	0.12	-
Tobin's Q		1.83	2.16	2.38	2.56	2.80	-
R&D (ln)		3.28	3.65	3.68	3.83	4.38	-

Bivariate sorts

I now rank stocks according to three characteristics and place them into two groups, low and high. As before, I also form 5 groups of stocks ranked according to their cyber risk score. This results in 10 portfolios. Table 5 presents only the average returns of the two difference portfolios, i.e. high minus low cyber risk stocks within the two groups of the second sort variable. The results are presented in terms of the FFC and FF% factor alphas.

Overall, panel A shows that the cross-sectional relation reported between returns and cyber risk reported in the univariate sorts remains after controlling for the second variable. More specifically, it suggests that high-size, high-Tobin's Q, and low-tangibility stocks have higher FFC and 5-factor alphas than low-size, low-Tobin's Q, and high-tangibility stocks, respectively.

For example, the first row of Panel A reports the results for size. For equally weighted portfolios, the FFC alpha is 0.20% for low-size stocks and 0.33% for high-size stocks. The corresponding FF5 alpha is 0.25% for low-size stocks and 0.36% for high-size stocks. The numbers in parentheses represent the t-statistics for the estimates, adjusted by Newey & West (1987). The average returns are statistically significant only within the high size subsample.

Similarly, the second row of Panel A shows the results for Tobin's Q. The table shows that for the equal-weighted portfolios, the FFC alpha is 0.19% for low-Q stocks and 0.29% for high-Q stocks. The corresponding FF5 alpha is 0.23% for low-Q stocks and 0.28% for high-Q stocks.

Finally, the third row shows the results for tangibility. The table shows that for the equal-weighted portfolios, the FFC alpha is 0.35% for low-tangibility stocks and 0.18% for high-tangibility stocks. The corresponding FF5 alpha is 0.38% for low-tangibility stocks and 0.20% for high-tangibility stocks.

These results further corroborate hypothesis 2.

Table 5: Double-sorted portfolios

This table reports the alphas adjusted by Carhart's 4-factor and Fama-French 5-factor models. Starting in January 2005 I rank stocks at the end of each month and allocate them into 5 groups. Independently, I also sort stocks into 2 groups according to the firm characteristics displayed. 10 portfolios are then formed, plus 2 spread portfolios, whose alphas I report. I compute their returns weighted by each stock's market capitalization (vw) and weighted equally (ew). Newey & West (1987) t-statistics are reported in brackets. ***, ** and * indicate statistical significance at 1%, 5% and 10% levels, respectively.

		Equal-weighted portfolios High-low Cyber Risk Stocks		Value-weighted portfolios High-low Cyber Risk Stocks	
		FFC alpha	*FF5 alpha*	*FFC alpha*	*FF5 alpha*
Panel A: Firm characteristics					
	LOW	**0.20**	**0.25**	**0.30**	**0.30**
Size		(0.90)	(1.08)	(1.05)	(0.99)
	HIGH	**0.33***	**0.36****	**0.53****	**0.56****
		(1.93)	(2.00)	(2.51)	(2.55)
	LOW	**0.19**	**0.23**	**0.38**	**0.34**
Tobin's Q		(0.80)	(0.91)	(1.52)	(1.31)
	HIGH	**0.29***	**0.28***	**0.43***	**0.46****
		(1.80)	(1.72)	(1.93)	(2.01)
	LOW	**0.35***	**0.38***	**0.58****	**0.56****
Tangibility		(1.82)	(1.93)	(2.16)	(2.00)
	HIGH	**0.18**	**0.20**	**0.50****	**0.54****
		(1.01)	(1.05)	(2.37)	(2.46)

8.2. Fama-Macbeth regressions

To further explore my main cross-sectional relation of interest I perform the Fama & MacBeth (1973) procedure. By running individual cross-sectional regressions, I can capture information that may get lost in portfolio analysis.

In all model specifications presented in Table 6, the dependent variable is excess returns in period t+1. The independent variable of interest, cybersecurity risk, is statistically significant at the 5% level in model specifications [3] and [4] but not in [1] or [2]. So, to find the relation I am primarily interested in, it is necessary to control for additional effects.

In general, the coefficients suggest that firms with higher levels of cybersecurity risk have higher excess returns. All variables are standardized, i.e. subtracted by the

mean and divided by the standard deviation to allow for a more straightforward interpretation of the slopes.

In model 4, the time-series average of the cross-sectional slope implies that a one standard deviation increase in cybersecurity risk increases returns by 0.10% per month. This validates hypothesis 1.

The other independent variables included in the regression are beta, size, book-to-market, momentum, and short-term reversal. Beta is statistically significant at the 5% level in all four columns, indicating a positive relation between beta and excess returns. None of the other variables is statistically significant at conventional levels.

Table 6: Fama-Macbeth regressions

This table presents the results from Fama-MacBeth regressions on the relation between my cybersecurity risk proxy and subsequent 1-month stock returns. For each month of the sample, I run cross-sectional regressions of excess stock returns on lagged cybersecurity risk and a set of firm characteristics that are also lagged. These are beta, size, book-to-market, momentum and short-term. All the variables are standardized. The coefficients represent the time-series averages of the estimates from the cross-sectional regressions. The t-statistics in brackets are based on the Newey & West (1987) standard errors. *, ** and *** denote statistical significance at 10%, 5% and 1% levels, respectively.

	Excess Returns$_{t+1}$			
	[1]	[2]	[3]	[4]
Cybersecurity risk	**0.08**	**0.09**	**0.09****	**0.10****
	(1.54)	(1.65)	(1.97)	(2.30)
Beta		**0.34****	**0.33****	**0.26***
		(2.07)	(2.09)	(1.82)
Size (ln)			**-0.002**	**-0.01**
			(-0.03)	(-0.07)
Book-to-market (ln)			**-0.04**	**-0.03**
			(-0.46)	(-0.40)
Momentum				**-0.06**
				(-0.31)
Short-term reversal				**-0.15****
				(-2.41)
Constant	**0.90****	**0.90****	**0.88****	**0.71***
	(2.13)	(2.13)	(2.07)	(1.76)
Observations	360230	360230	360230	360230

Next, I examine whether the relationship between cybersecurity risk and returns is stronger for technology companies. To do so, I create a dummy equal to one for technology companies and interact it with my measure.

As before, Table 7 reports the coefficient estimates for each independent variable for the four model specifications.

Table 7: Fama-Macbeth regressions with an interaction effect

This table presents the results from Fama-MacBeth regressions on the relation between my cybersecurity risk proxy and subsequent 1-month stock returns, including the interaction term *Tech * Cyber risk*. For each month of the sample, I run cross-sectional regressions of excess returns on lagged cybersecurity risk and a set of firm characteristics also lagged. I control for the same variables as previously. All the variables are standardized. The coefficients represent the time-series averages of the estimates from the cross-sectional regressions. The t-statistics in brackets are based on the Newey-West standard errors. *, ** and *** denote statistical significance at 10%, 5% and 1% levels, respectively.

	Excess Returns$_{t+1}$			
	[1]	[2]	[3]	[4]
Cybersecurity risk	**0.09**	**0.11**	**0.10**	**0.12****
	(1.16)	(1.40)	(1.51)	(1.98)
Tech dummy	**0.08**	**0.06**	**0.00**	**0.08**
	(0.63)	(0.54)	(0.61)	(0.76)
*Tech * Cyber risk*	**-0.04**	**-0.05**	**-0.00**	**-0.06**
	(-0.54)	(-0.69)	(-0.63)	(-0.93)
Beta		**0.34****	**0.33****	**0.26***
		(2.08)	(2.09)	(1.83)
Size (ln)			**-0.00**	**-0.01**
			(-0.05)	(-0.07)
Book-to-market (ln)			**-0.04**	**-0.03**
			(-0.48)	(-0.41)
Momentum				**-0.06**
				(-0.32)
Short-term reversal				**-0.16****
				(-2.44)
Constant	**0.88****	**0.89****	**0.87****	**0.70***
	(2.13)	(2.14)	(2.07)	(1.74)
Observations	360,230	360,230	360,230	360,230

It seems that the relationship between cybersecurity risk and excess returns is not significantly stronger for tech companies relative to non-tech companies.

The coefficient for the interaction term between cybersecurity risk and the tech dummy if anything is negative, so the effect of cybersecurity risk on returns might be weaker for technology assets.

Therefore, it seems that the relation found in Table 6 between cybersecurity risk and excess returns is not driven solely by tech companies, and it applies to companies in other sectors as well.

8.3. Time-series

To test for structural changes in the relationship between variables over time I use the Chow test. I divide the sample T into two subsamples: T1, before, and T2, after the suspected breakpoints. I test 3 breakpoints, January, February, and March 2020. These points coincide with the beginning of the pandemic in the U.S. when cybersecurity risk surged.

The p-values associated with the F-statistics are 0.32, 0.28 and 0.56, greater than any significance level. Therefore, I do not reject the null hypothesis. This means I cannot conclude that there are statistically significant differences in the coefficients of the regression models. So, I assume that the coefficients stay the same across the different sub-samples of data.

The last two results show that my cybersecurity risk-returns relation is not affected by the technology sector or by a regime change in the time-series data. This offers additional support to hypothesis 1.

To conclude section 8, the results are robust to my alternative cyber risk measures, both the simple unweighted sum of terms (Cyber2) and the sum of terms normalized by the total number of words in each transcript (Cyber3).

9. Limitations

In this section, I discuss the limitations of my measure and ways to address them.

My findings point out that investors are willing to pay a premium for holding stocks highly exposed to cybersecurity risk. Also, the results are not driven by technological companies that inevitably have longer discussions about cybersecurity and hence score higher.

However, my measure might have limitations. In the fashion of Jiang et al. (2020) and Jamilov et al. (2021), In Figure3, I look at two cybersecurity ETFs, First Trust Nasdaq Cybersecurity ETF (CIBR) and ETFMG Prime Cyber Security ETF (HACK). There are other similar ETFs, but these were established first. So, more data on prices are available, since January and August 2015, respectively.

Figure 3: Cyber long-short portfolio and ETF returns

This figure illustrates the correlation between returns of the spread portfolio (value-weighted) sorted by my cyber risk measure, and CIBR ETF, from 2015-08 to 2021-12. The correlation between CIBR and HACK is +96% so it is redundant to plot the latter.

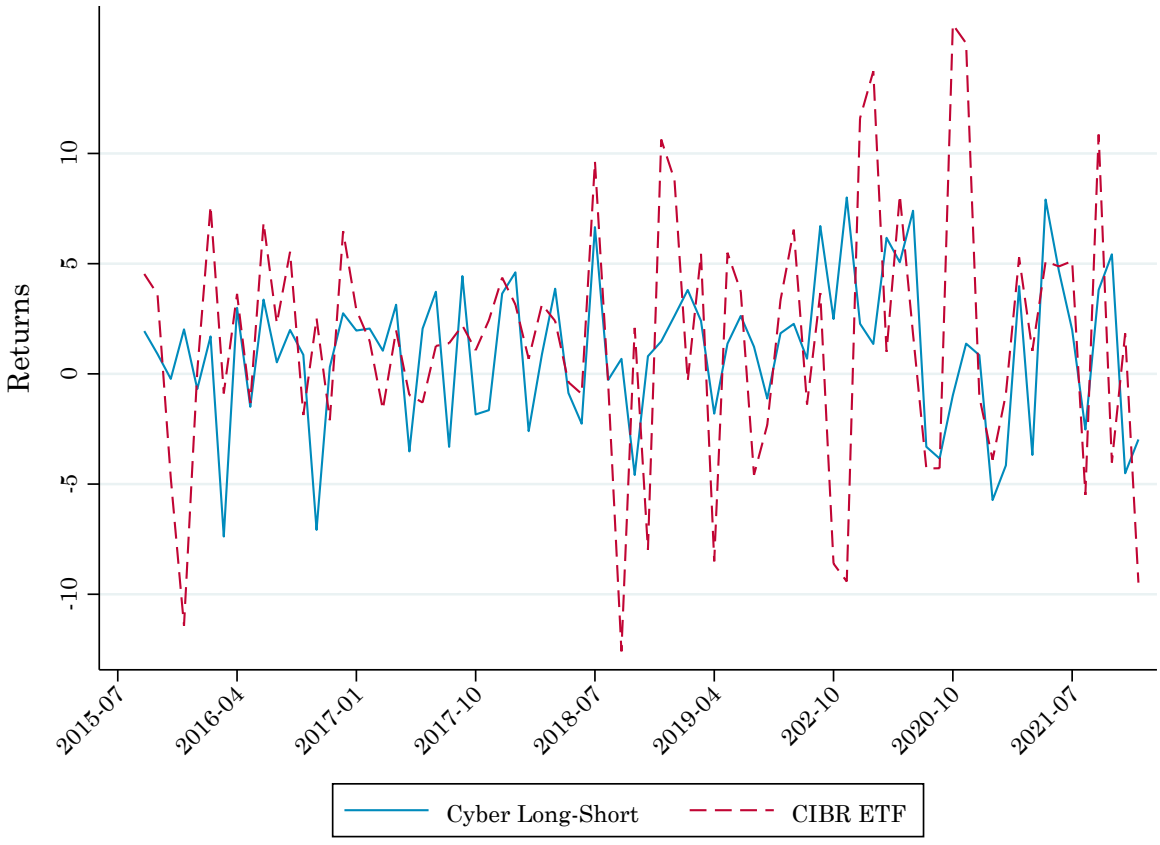

I test whether these ETFs' returns are correlated with my high-minus-low cybersecurity risk portfolio. It turns out the correlations are significantly positive, both around +30%, as Figure 3 illustrates.

This result contrasts with that of Jiang et al. (2020), who report a -30% correlation. The authors argue this result is expected, as such ETFs should provide a hedge against cybersecurity risk, thus having minimal exposure to it. Moreover, it documents that the highest number of data breaches occurs in the Financial Activities industry. In both measures of Jamilov et al. (2021) and Florackis et al. (2022), the finance sector scores high too. With my measure it happens somewhat the opposite, as Figure 2 in section 6 indicates.

Yet, Jamilov et al. (2021) argue these ETFs reflect mainly market risk. I find a +80% correlation with the S&P 500. Further, I am not convinced that these cybersecurity firms are in the low range of cyber risk, as they have expertise in the field. At the same time, by holding such intellectual property and sensitive data, hackers can see them as attractive targets, as suggested by the following examples.

In December 2020, a highly sophisticated attack targeted FireEye, a top US cybersecurity firm. It used techniques never seen before. The probably state-sponsored hackers stole sensitive information about FireEye's clients, as well as hacking tools used by the company to test the vulnerability of its clients' systems, both corporate and government. Hackers can then exploit these tools to attack high-profile targets (Sanger & Perlroth, 2020). With the same goal, a group known as ShadowBrokers hacked into the US National Security Agency, in 2016.

These results about the ETFs are not conclusive but lead me to further evaluate my proxy. It appears that my measure may not always capture the same aspects of cybersecurity risk. Jiang et al. (2020) approach is focused on estimating the probability of a cyberattack occurring and the loss severity, given firm characteristics and textual analysis of their 10-K fillings. They primarily attempt to tackle the latent feature of this risk. My proxy does not aim at this nor at predicting cyberattacks.

Instead, it is likely a combination of actual cybersecurity risk and the extent to which firms discuss cybersecurity issues during their calls. One way to disentangle the two is to see if firms in the high cybersecurity risk quantile experience more cybersecurity breaches and/or cyber insurance claims than those in the low cybersecurity risk quantile. Or else, as the papers I just mentioned, use more

sophisticated methods to construct the measure, such as machine learning techniques or sentiment analysis.

Another possibility to address this issue, but without changing the way my proxy is constructed, is using a more restrictive dictionary. So, of the 126 terms I list, I could keep only the most directly associated with cybersecurity risk, which would account for less than 30. But then my measure is scarce, and I cannot conduct the asset pricing tests as I did. For instance, I can form only two instead of five portfolios. One where the terms count is zero, and a second where is different from zero.

10. Conclusion

Overall, my research is a good starting point for understanding the relationship between cybersecurity risk and stock returns.

To study this cross-sectional relation, I first build a firm-level measure by parsing the firm's earnings conference calls, counting cybersecurity-related terms. The more terms mentioned the higher the cybersecurity risk score. I check the time-series evolution and its behaviour across industries to show my proxy is capturing cybersecurity risk.

Using portfolio sorts and Fama-Macbeth regressions, I find that investors require higher returns to compensate for additional units of cybersecurity risk. A zero-investment strategy earns significant abnormal returns, net of common risk factors.

However, a simple measure like mine cannot always assure that I am capturing the risk of adverse impacts caused by a cyber incident. Sometimes high cyber risk firms according to my proxy are just firms which discuss cybersecurity in more detail in their conference calls. So perhaps, they keep high levels of cyber-hygiene, and not the opposite as the high score suggests.

Ultimately, the precise factors driving the relation between cybersecurity risk and risk-adjusted returns require further investigation to determine. It may be a combination of factors, including differences in the quality of cybersecurity practices, susceptibility to cyber threats, and market perceptions of risk.

To conclude, cybersecurity risk is clearly not slowing down, as the world relies more on technology. Hackers will continue to improve their skills and develop increasingly

sophisticated weapons. Companies need to take active steps to protect their digital assets and reduce their vulnerability to cyberattacks.

References

Amir, E., Levi, S., & Livne, T. (2018). Do firms underreport information on cyber-attacks? Evidence from capital markets. *Review of Accounting Studies, 23*(3).

Ang, A. (2014). *Asset Management: A Systematic Approach to Factor Investing*. Oxford University Press.

Bali, T., Engle, R., & Murray, S. (2014). *Empirical Asset Pricing: The Cross Section of Stock Returns: An Overview*. John Wiley & Sons, Ltd.

Black, F. (1993). Beta and Return. *The Journal of Portfolio Management, 20*(1).

Carhart, M. M. (1997). On Persistence in Mutual Fund Performance. *The Journal of Finance (New York), 52*(1), 57-82.

Crosignani, M., Macchiavelli, M., & Silva, A. F. (2023). Pirates without borders: The propagation of cyberattacks through firms' supply chains. *Journal of Financial Economics, 147*(2), 432-448.

Fama, E. F., & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics, 33*(1), 3-56.

Fama, E. F., & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics, 116*(1), 1-22. 10.1016/j.jfineco.2014.10.010

Fama, E. F., & MacBeth, J. D. (1973). Risk, Return, and Equilibrium: Empirical Tests. *Journal of Political Economy, 81*(3), 607-636. 10.1086/260061

Florackis, C., Louca, C., Michaely, R., & Weber, M. (2022). Cybersecurity Risk. *The Review of Financial Studies, 36*(1), 351-407. 10.1093/rfs/hhac024

Fung, B. (2021, Apr 12,). Cyberattacks are the number-one threat to the global financial system, Fed chair says.*CNN Wire Service*

Gambacorta, L., Giudici, P., Leach, T., & Aldasoro, I. (2022). The drivers of cyber risk. *Journal of Financial Stability, 60*, 100989.

Greenberg, I. (2021). *Fifth-generation cyberattacks are here. How can the IT industry adapt?* . World Economic Forum.

Hassan, T. A., Hollander, S., van Lent, L. A. G. M, & Tahoun, A. (2019). Firm-level political risk: Measurement and effects. *The Quarterly Journal of Economics, 134*(4).

Jamilov, R., Tahoun, A., & Rey, H. (2021). *The Anatomy of Cyber Risk.* (). Cambridge, Mass: National Bureau of Economic Research.

Jegadeesh, N., & Titman, S. (1993). Returns to Buying Winners and Selling Losers: Implications for Stock Market Efficiency. *The Journal of Finance (New York), 48*(1), 65-91.

Jiang, H., Khanna, N., Yang, Q., & Zhou, J. (2020). The Cyber Risk Premium. *SSRN Electronic Journal.*

Kamiya, S., Kang, J., Kim, J., Milidonis, A., & Stulz, R. M. (2021). Risk management, firm reputation, and the impact of successful cyberattacks on target firms. *Journal of Financial Economics, 139*(3), 719-749.

Lakonishok, J., Shleifer, A., & Vishny, R. W. (1994). Contrarian Investment, Extrapolation, and Risk. *The Journal of Finance, 49*(5), 1541-1578.

Lallie, H. S., Shepherd, L. A., Nurse, J. R. C., Erola, A., Epiphaniou, G., Maple, C., & Bellekens, X. (2021). Cyber security in the age of COVID-19: A timeline and analysis of cyber-crime and cyber-attacks during the pandemic.*Computers & Security, 105*, 102248.

Lhuissier, S., & Tripier, F. (2021). Measuring Cyber Risk.

Loughran, T., & McDonald, B. (2011). When Is a Liability Not a Liability? Textual Analysis, Dictionaries, and 10-Ks.*The Journal of Finance, 66*(1), 35-65.

Loughran, T., & McDonald, B. (2016). Textual Analysis in Accounting and Finance: A Survey. *Journal of Accounting Research, 54*(4), 1187-1230.

MarketInsite. (2019). *The Hack Effect: The Effect of Data Breaches on the Nasdaq CTA Cybersecurity Index.*

Moschetta, G., Beato, F. & Joshi, A. (2023). *Cybersecurity must be tightened up in this era of polycrisis.*

Newey, W. K., & West, K. D. (1987). A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix. *Econometrica, 55*(3), 703-708.

Novy-Marx, R. (2013). The other side of value: The gross profitability premium. *Journal of Financial Economics, 108*(1), 1-28.

Sanger, D., & Perlroth, N. (2020, Dec 8,). FireEye, a Top Cybersecurity Firm, Says It Was Hacked by a Nation-State. *New York Times*

Sautner, Z. (2021). *Pricing Climate Change Exposure*. SSRN.

Schwert, G. W. (2003). Anomalies and market efficiency. *Handbook of the Economics of Finance* (pp. 939-974). Elsevier.

Sharpe, W. F. (1964). Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk. *The Journal of Finance (New York), 19*(3), 425.

Tosun, O. K. (2021). Cyber-attacks and stock market activity. *International Review of Financial Analysis, 76*.

Venkataramakrishnan, S. (2022). *In charts: Cyber security risks and companies' readiness*. The Financial Times Limited.

Zammost, S., & Schlesinger, J. (2020). *US Secret Service warns that coronavirus email scams are on the rise* . CNBC.

Zhang, L. (2005). The Value Premium. *The Journal of Finance, 60*(1), 67-103.

Appendix

Term List

access management
adware
antivirus
app
baseline security
blacklist
botnet
bring-your-own-device
brute force attack
card breach
card fraud
certificate management
chief information officer
cloud
confidentiality
cookie
credentials
critical infrastructure
cryptographic
cyber advisory
cyberattack
cyber
cyber event
cyber incident
cyber infrastructure
cyber risk
cybersecurity
cyber threat
dark web
data at rest
data breach
data compromise
data fraud
data integrity
data leak
data loss
data loss prevention
denial-of-service
dictionary attack
digital footprint
digital signature
digital transformation
domain count
electronic evidence
electronic key
email compromise
encryption
end user device
extranet
file security
firewall
firmware
hack
hardware
hash
honeypot
ict
identity management
identity proofing
identity token
identity validation
identity verification
information theft
internetof things
intranet
intrusion detection system
ip security
it risk
it security
kerberos
key bundle
major information system
malicious code
malvertising
operating system
password
patch management
penetration testing / pentest
personal identification number
personal identity verification
pharming
phishing
private key
protective technology
protocol
public key
ransomware
remote access
router
secret key
sensitive information
social engineering
software
spam
spear phishing
spoofing
spyware
supply-chain attack
system integrity
system interconnection
system outage
threat actor
threat assessment
trojan
unauthorized access
unauthorized disclosure
virtual private network
virus
vpn
vulnerability assessment
vulnerability management

data security	malware	whaling
data theft	man-in-the-middle attack	whitelist
ddos attack	multifactor authentication	wireless local area network
decrypt	network security	worm
deleted file	online attack	zero-day

Variable definitions

Variable	Description	Source
Beta	The market beta of individual stocks estimated using monthly returns of the past 24 months.	CRSP
Book-to-market	Book value of common equity (ceq) / market value of common equity (prcc_f*csho)	Compustat
Firm size	|Total number of shares outstanding (shrout) * price of one share (altprc) | / 1000 The absolute value accounts for the fact that CRSP reports a negative price when the reported value is calculated as the average of a bid and ask price.	CRSP
Momentum	The cumulative return of a stock over a period of 11 months ending one day prior to month t.	CRSP
R&D expenditure	R&D expenditures (xrd). Missing values are replaced with zero.	Compustat
ROA	Net income (ni) / total assets (at)	Compustat
ROE	Net income (ni) / common equity (ceq)	Compustat
Short-term reversal	The return of a stock during month t.	CRSP
Tangibility	Total property, plant and equipment (ppent) / total assets (at)	Compustat
Tobin's Q	(Total assets (at) – common equity (ceq) + market value of equity (prcc_f * csho)) / total assets (at)	Compustat