Cyber Risk in General Equilibrium

Aniket Baksy* Daniele Caratelli†

March 26, 2026

Abstract

We study the implications of rising cyber risk for macroeconomic outcomes and the ability of policy to mitigate them. We develop a general equilibrium model that features strategic interactions between heterogeneous firms that invest in cybersecurity and attackers who direct their search toward the most attractive targets. Attackers trade-off the higher bounty from attacking a larger firm against the lower likelihood of overcoming their cybersecurity defenses. Cyberattacks generate a negative externality by reducing aggregate productivity, motivating a role for policy. We discipline the model using firm-level evidence on cybersecurity investment, documenting that cybersecurity employment as a share of total employment rises steeply with firm size. Introducing cyber risk to the economy reduces firm entry by 4 percent, aggregate productivity by 0.6 percent, and total output by 1.7 percent. Using the model, we evaluate two policy interventions, subsidies for cybersecurity investment and bailouts for attacked firms. While subsidies can raise aggregate output, bailouts reduce it, with effects that depend crucially on the severity and nature of cyber risk.

*aniket.baksy@unimelb.edu.au.

†danicaratelli@gmail.com.

We thank Chanelle Duley, Michael Junho Lee, Simon Loertscher, Mark Paddrik, Stacey Schreft, Lawrence Uren, Peyton Young and participants at the UniMelb faculty seminar for useful discussion and insightful comments. We thank the Privacy Rights Clearinghouse for generously sharing their data. Aniket gratefully acknowledges the hospitality of the ESRC-funded Digital Futures at Work Research Centre, where a portion of the research underlying this project was carried out.

1 Introduction

Over the past two decades, cyber risk has evolved from a narrow IT concern into a widespread threat that firms across the economy confront daily. As we document, data breaches, ransomware incidents, and disruptions to digital infrastructure now occur at an alarming frequency, with thousands of major cases recorded each year in the United States alone. These events inflict direct financial losses on firms[1] and ripple through the economy by disrupting supply chains, eroding customer trust, and weighing on productivity and investment. Accordingly, policymakers are starting to take notice[2].

Despite its growing importance, the macroeconomic consequences of rising cyber risk remain largely unexplored. Unlike other shocks to revenues faced by firms, which researchers often model as stochastic processes which firms take as given in the tradition of Hopenhayn (1992), exposure to cyber risk is an equilibrium outcome of interactions between firms and attackers. Thus, while existing research has made important progress in documenting firm-level consequences, such as declines in market valuation and investment (e.g., Kamiya et al. 2021; Akyildirim et al. 2024; Muktadir-Al-Mukit and Ali 2025), and in modeling cybersecurity decisions in static, partial-equilibrium settings (Moore, Clayton and Anderson 2009; Ramírez 2025; Ebel and Mitra 2024), these approaches fall short of providing a framework in which cyber risk can be evaluated as a systemic phenomenon and policy tradeoffs can be analyzed quantitatively.

In this paper, we propose such a framework. Our primary contribution is to develop a micro-founded, dynamic general equilibrium model of cyber risk. The model features a two-sided interaction between firms and attackers: firms' cybersecurity investments depend on their expected exposure to attack, while attackers direct their efforts toward firms whose size and defenses make them most attractive targets. These strategic interactions shape firms' entry and exit decisions and their exposure to attack, altering the equilibrium firm size distribution and amplifying the aggregate consequences of cyber risk through general equilibrium effects. To the best of our knowledge, ours is the first model that captures these interactions in an environment featuring rich firm heterogeneity and general equilibrium consequences, and in which serious policy counterfactuals can be conducted.

We now discuss our model's structure. The model integrates two classic building

[1] According to the Ponemon Institute and IBM (2025), the average cost of a breach was nearly $4.4 million in 2024.

[2] The financial sector is particularly prone to cyberattacks and also highly vulnerable given the interconnected nature of financial activities, leading to concerns from central banks and financial regulators. For instance, the Federal Reserve annually publishes a Cybersecurity and Financial System Resilience Report to Congress on actions taken to strengthen the security of the financial sector (Board of Governors of the Federal Reserve System 2025). Broader concerns have been raised around the security of critical infrastructure projects and national security.

blocks of modern macroeconomics: a search-and-matching framework à la Diamond (1982), and Mortensen and Pissarides (1994) and an endogenous firm entry–exit margin following Hopenhayn (1992). Firms, heterogeneous in their productivity, decide whether to incur a fixed cost to enter the market. Those that enter hire labor, produce output, and invest in cybersecurity to defend against potential attacks. Production is subject to fixed costs, which occasionally lead firms to exit. Attacks are launched by attackers, who choose the number and intensity of the attacks they launch against firms of a given size[3].

We now discuss the consequences of cyberattacks. A matching technology determines the probability that an attacker meets a firm, which therefore depends on the "tightness" of the cyber market, the ratio of the number of attackers to potential targets. Upon matching, attackers decide on the intensity of their attack subject to a capacity constraint; the highest possible attack intensity, which we call *attack capacity*, plays an important role in our analysis. Attacks succeed with a probability increasing in attack intensity and decreasing in the level of cybersecurity chosen by the targeted firm. Successful attacks reduce firm output over their duration and allow the attacker to extract rents increasing in the size of the targeted firm; the resulting loss of profits can incentivise attacked firms to exit. Finally, we allow for externalities from cyberattacks by having aggregate productivity decrease in the total share of attacked firms.

Our analysis is novel in that it captures the rich interactions between firms' entry, exit, and cybersecurity investment decisions and attackers' size-contingent targeting and attack intensity choices. On the one hand, firms anticipate that their specific size may make them more vulnerable to attacks and choose their cybersecurity investment accordingly. On the other hand, when attackers decide which firms to attack, they take these investment decisions into account. As a result, firms behave as Stackelberg leaders in deciding on the level of security they adopt. These interactions jointly determine firms' survival, resource allocation, and entry and exit dynamics.

The model delivers several insights. First, it predicts that the incidence of attacks follows an inverse-U shape in firm size. The smallest firms are unattractive targets because the payoff to attackers is limited, while the largest firms invest heavily in cybersecurity and are therefore difficult to breach. Mid-sized firms, in contrast, represent both a meaningful prize and a tractable target, making them disproportionately vulnerable. We provide empirical evidence for this claim using rich microdata on individual breaches[4].

[3]The role of attackers is analogous to that of vacancy-posting firms in search and matching models of the labor market, while firms play the role of unemployed workers meeting attackers at random. However, while it is typically assumed that unemployed workers desire matches with posted vacancies, firms in our model would rather not be contacted by cyberattackers.

[4]We note that empirical validation of this claim is severely complicated by the fact that cyberattacks remain among the most under-reported category of crimes, and that mandated disclosures are heavily

Second, the inclusion of cyber risk leads to a "missing middle" in the firm size distribution: relative to a counterfactual economy without cyberattacks, there are fewer medium-sized firms, whereas the mass of large and small firms is much less affected. Intuitively, our model features no permanent productivity heterogeneity; hence, the largest firms are ones which have drawn sequences of high productivity shocks, investing heavily in cybersecurity and in turn deterring attackers. Small firms represent unattractive targets, given that each cyber attack requires a fixed investment from an attacker. Mid-sized firms represent targets of sufficient value to be worth attacking.

Third, our model emphasizes the importance of general equilibrium effects operating via entry and exit, and the labor market, that are absent in partial-equilibrium reasoning. In particular, cyber risk leads to an adverse shift in the firm size distribution towards relatively small firms, as a result of which aggregate productivity falls. In our simulations, aggregate output falls by as much as 1.6% relative to an economy without cyber risk.

In counterfactual exercises, we consider two scenarios particularly likely to characterize the cyber landscape in the near future: larger average attack capacities, leading to larger attacks, and higher average match probabilities. We model the former by an increase in the attack capacity of the typical attacker, and the latter by an increase in the efficiency of the matching technology determining firm-attacker contact rates. Both of our scenarios can be interpreted as the results of technical change favouring cyberattackers and raising cyber risk. However, we use this experiment to demonstrate how evaluating the consequences of a rise in cyber risk requires understanding its source and a careful accounting of endogenous responses, with the attendant general equilibrium consequences. A doubling of attack capacity from our baseline leads to a 0.75% decline in output relative to our calibrated baseline, while a doubling of matching efficiency reduces it by close to 2%. In the former case, the decline is driven almost entirely by general equilibrium adjustments via entry, exit and the labour market, while in the latter case, it is driven largely by losses from spillovers due to an almost doubling of the share of firms under attack. In both cases, the measure of firms in the economy falls by almost 2% as a result.

Beyond positive analysis, our model allows us to consider normative questions taking strategic interactions between firms and attackers into account, a feature which to the best of our knowledge is novel to the literature. Cybersecurity investments are costly, but their benefits extend beyond the individual firm, since reduced vulnerability improves

skewed toward larger and public firms. Taking our model seriously would suggest that the under-reporting of cyberattacks is highest among medium-sized businesses, and that the collection of survey data on this category of firms is of critical importance. The Ponemon Institute's 2016 State of Cybersecurity in Small and Medium-Sized Business surveyed 598 businesses with between 100 and 1000 employees, finding that 55% of respondents claimed a cyber attack in the previous year, a number much higher than suggested by datasets tracking cyber attacks (Ponemon Institute 2016).

outcomes for trading partners and customers and lowers aggregate attack probabilities. This externality provides a rationale for government intervention. We explore two policies governments could adopt to mitigate cyber risk: firm subsidies for cybersecurity investments and ex-post bailouts to attacked firms. We find that the subsidy that maximizes total output in the economy would lead firms to invest almost 11% more in cybersecurity on average. Our results highlight that policy evaluation in the cyber domain must account for equilibrium spillovers. Policy interventions affect firm entry and exit margins and through this macroeconomic outcomes.

Related Literature. Our paper contributes to a nascent but fast-growing literature on cybersecurity. On the empirical side, several studies have measured the costs of breaches, including reductions in firm revenues, investment, and market valuations (Romanosky 2016; Anderson et al. 2013). A related strand documents the distribution and dynamics of breaches across industries and firm types (Edwards, Hofmeyr and Forrest 2016; Seh et al. 2020; Barati and Yankson 2022; Carfora and Orlando 2022), with recent work emphasizing the disproportionate vulnerability of medium-sized firms (Baksy, Caratelli and Olson 2025).

Another strand emphasizes the financial stability and macroprudential implications of cyber risk. Duffie and Younger (2019) point to the real economic effects cyber attacks can have through shared digital infrastructure and payment systems. Eisenbach, Kovner and Lee (2022) highlights how cyber shocks differ from traditional operational risks, arguing that their systemic nature warrants dedicated modeling frameworks and stresses the potential amplification channels through payment and settlement systems when a cyber shock hits the financial system. Using a natural experiment, Kotidis and Schreft (2025), provides causal evidence that cyber incidents propagate through financial networks, raising funding costs and weakening balance sheets of firms not directly attacked, highlighting an important channel through which cyber shocks generate systemic risk. Koo et al. (2022) develops a framework to integrate cyber events into macro-financial risk monitoring, demonstrating how breach data can inform stress testing exercises. Anand, Duley and Gai (2022) emphasise the importance of accounting for spillovers via shared software platforms. These contributions underscore the need for a fully micro-founded equilibrium framework that links micro-level cyber interactions to aggregate outcomes. Our paper advances this literature by embedding a game-theoretic treatment of cyberattacks into a dynamic general equilibrium model of firm entry and exit, thereby linking micro-level interactions to aggregate outcomes.

On the theory side, a smaller set of papers model cyberattacks as strategic interac-

tions between attackers and firms (Moore, Clayton and Anderson 2009), though typically in static and partial-equilibrium environments. Our framework extends this approach by embedding these interactions in a dynamic general equilibrium with endogenous firm entry and exit, allowing both firm behavior and attacker activity to adjust endogenously. To do so, we adapt a search-and-matching framework in which firms face random encounters with attackers and make ex-ante investment decisions that shape the consequences of these encounters, making cyber risk an equilibrium force governing firm dynamics and aggregate outcomes. The model allows us to evaluate policy interventions aimed at mitigating the macroeconomic consequences of rising cyber risk. Our results indicate that private incentives for cybersecurity investment can diverge from social incentives. This echoes the findings of Erol and Lee (2024), who show that competition among financial market infrastructures can lead to underinvestment in technological resiliency and system-wide "technological drag." We extend this insight to a general equilibrium, macroeconomic model, showing how cyber risk shapes firm entry, the size distribution, and aggregate productivity.

Outline. The remainder of the paper proceeds as follows. Section 2 documents new empirical evidence on the incidence of cyberattacks, focusing on the disproportionate risk borne by medium-sized firms. Section 3 develops a simple deterministic model to illustrate the attacker–firm interaction. Section 4 extends the framework to a dynamic environment with firm entry and exit. Section 5 calibrates the model to match features of the U.S. economy in 2000 and in the 2020s. Section 6 presents counterfactual exercises that highlight the importance of general equilibrium forces, and Section 7 discusses policy implications. Section 8 concludes.

2 Cyber Risk: What you need to know

In this section we document various features of increase in cyber risk over the last 20 years, relying both on the existing literature and several independent datasets. We document five facts. First, cyber risk is widespread and has risen dramatically over the past twenty years. Second, the incidence of cyberattacks occurs across the firm size distribution, with small and medium sized enterprises suffering a rising share of attacks. Third, we show that investments in cybersecurity, proxied for by the employment of cybersecurity professionals, scales rapidly with firm size, even after accounting for sector and firm-level variation. Fourth, we provide evidence that cyberattacks have dramatic and persistent firm-level effects. Finally, we review evidence on the indirect costs of cyber-

attacks and argue that these can dominate direct losses. The evidence presented in this section comes with an important caveat: cybercrime is selectively reported and hard to quantify, and that observed breach data substantially understate true attack incidence, particularly among smaller and private firms (Romanosky, Telang and Acquisti 2011), an issue that motivates our structural approach in the sections that follow.

2.1 Cyber Risk is Rising and All-Pervasive

In this section, we document the increase in cybersecurity threats in the US since 2005. To establish this we use the *Data Breach Chronology*, compiled by the Privacy Rights Clearinghouse (PRC). These data provide a broad but incomplete record of data breaches in the United States[5]. The dataset records information from a wide range of sources, including state attorneys-general, regulatory filings, media reports, and company press releases, and records details on each event. These include the identity and type of the affected organization, the date the breach was reported, the estimated number of individuals impacted, the type of breach, the sector of the affected entity, and the duration of the incident. The data contain 74,783 breaches; we focus on the 27,385 attacks which occurred between 2005 and 2024 inclusive, in the United States, and affected private business organisations.

Figure 1 shows that the total number of reported attacks has increased by several orders of magnitude over the past 20 years, rising from 101 attacks in 2008 to 4,603 in 2023. The rapid increase in the number of attacks is paralleled by a rise in their severity: the median attack compromised 168 records in 2008, but nearly 2,300 records in 2023.

2.2 Cyber attackers are highly likely to have financial motives

Economically damaging cyber attackers may a wide range of motives, including non-financial ones. However, we document that financially motivated attacks comprise a large share of recorded attacks.

The Privacy Rights Clearinghouse data do not contain a direct measure of attacker motives, making it difficult for us to distinguish between attacks primarily for financial gain and attacks that are driven by non-financial motives[6]. We therefore rely on a dataset of attacks maintained by the Center for International and Security Studies at Maryland

[5]The Privacy Rights Clearinghouse Data Breach Chronology has been widely employed in research on cybersecurity. Examples include and risk modeling (Carfora and Orlando 2022), breach impact analysis (Rosati 2021; Barati and Yankson 2022; Edwards, Hofmeyr and Forrest 2016), longitudinal trend analyses (**??**), and systematic literature reviews in breach classification (**?**).

[6]The Office for Financial Research's Annual Report to Congress (Office of Financial Research 2022) lists multiple possible alternative motivations that attackers may have, including revenge, reputation, geopolitics, and international security.

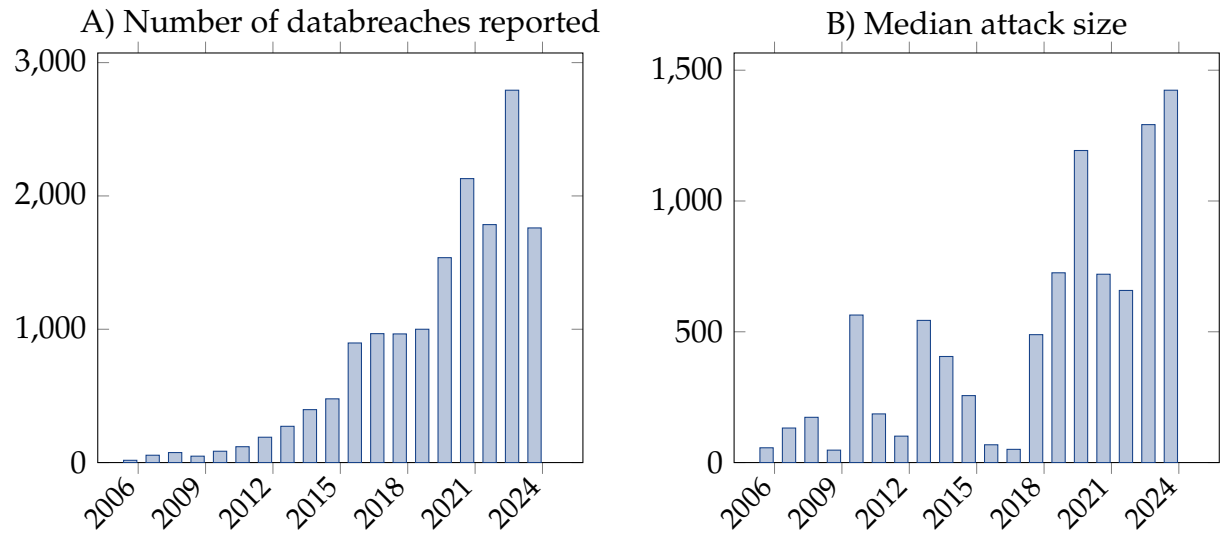


Figure 1: Number of attacks reported in the US (panel A) and the median number of records breached across breaches in a given year (panel B). All data from the Privacy Rights Clearinghouse.

(CISSM). Between 2014 and 2023, around 89% of all cyberattacks in the dataset[7] were explicitly motivated by financial motives. We note that the CISS data is known to pay more attention to attacks by state actors who are more likely to have geopolitical motives, likely biasing the true share of financially motivated attacks even higher.

Year	Financial	Espionage	Personal	Protest	Political	Sabotage
2014	103	14	13	19	9	1
2015	95	9	7	28	9	1
2016	146	5	8	33	5	10
2017	216	16	11	8	14	12
2018	293	29	5	8	12	1
2019	333	11	3	3	7	5
2020	486	26	5	6	20	3
2021	529	12	1	15	9	1
2022	747	6	0	32	6	1
2023	1052	8	0	23	8	6

Table 1: Cyber attack motives in the CISSM dataset. We pool espionage targeted against firms and governments into one category for parsimony.

[7] We focus on attacks on private firms in the United States.

2.3 Cyber Threats Increasingly Affect Small and Medium Enterprises

It is possible that the rise of large data-intensive corporations has incentivised cyberattacks against their databases, and that the rising number of attacks is concentrated against these large firms. We cannot verify this directly in our database, since it does not contain any information about the target of an attack other than the name[8]. In this section, we nonetheless provide two pieces of evidence that the rise in attacks increasingly affects even modest-sized firms.

First, while the DBC does not directly provide a measure of firm size, it does provide an estimate of the number of records breached in an incident; under the assumption that this number scales with firm revenue[9] and hence with firm size, the distribution of attacks across the number of breached records is informative about the distribution of attacks by firm size. Figure 2 depicts the distribution of attacks across bins of the number of records breached, year-on-year. The figure makes it clear that the rising mean number of breached records is not driven by headline-grabbing breaches compromising over a million records, which are likely to affect large corporations. Instead, the rise is driven by a stark increase in attacks compromising an above-average 1,000 to 10,000 records, the sizes of databases associated with mid-sized enterprises: such attacks have grown from comprising about 6.3% of all attacks to nearly 34% of them. Attacks involving the largest numbers of record breaches (over a million) account for only about 2% of all breaches, a share that has remained stable over our sample period.

Given that reporting standards in some cases require larger firms to report cyberattacks in greater detail, and that the coverage of our data is likelier to omit small and medium sized firms, this is likely a stark underestimate of the actual share of attacks that small and medium enterprises have faced. For instance, the Ponemon Institute's Review of Small and Medium-Sized Business cybersecurity in 2016 (Ponemon Institute 2016) found that 55% of all medium-sized firms (between 100-1000 employees) experienced a cyberattack in the year prior to the survey.

We provide additional evidence on the vulnerability of small and medium firms us-

[8]We have explored matching the dataset with Compustat to obtain firm sizes using fuzzy matching algorithms on firm name. However, such an exercise has not proved fruitful for two reasons: first, the match rate between Compustat and the DBC is persistently low, given that the former contains only public firms and the latter contains a large number of small, private firms, and second, because employment information in Compustat is only available annually and includes employees from non-US locations, making it difficult to measure firm sizes directly. We are working on matching our Revelio data with the DBC to provide additional information regarding domestic firm size.

[9]The assumption that the size of a firm's database is related positively to the firm's size features in recent theoretical and quantitative work on the role of data in determining the firm size distribution (Begenau, Farboodi and Veldkamp 2018; Farboodi et al. 2019).

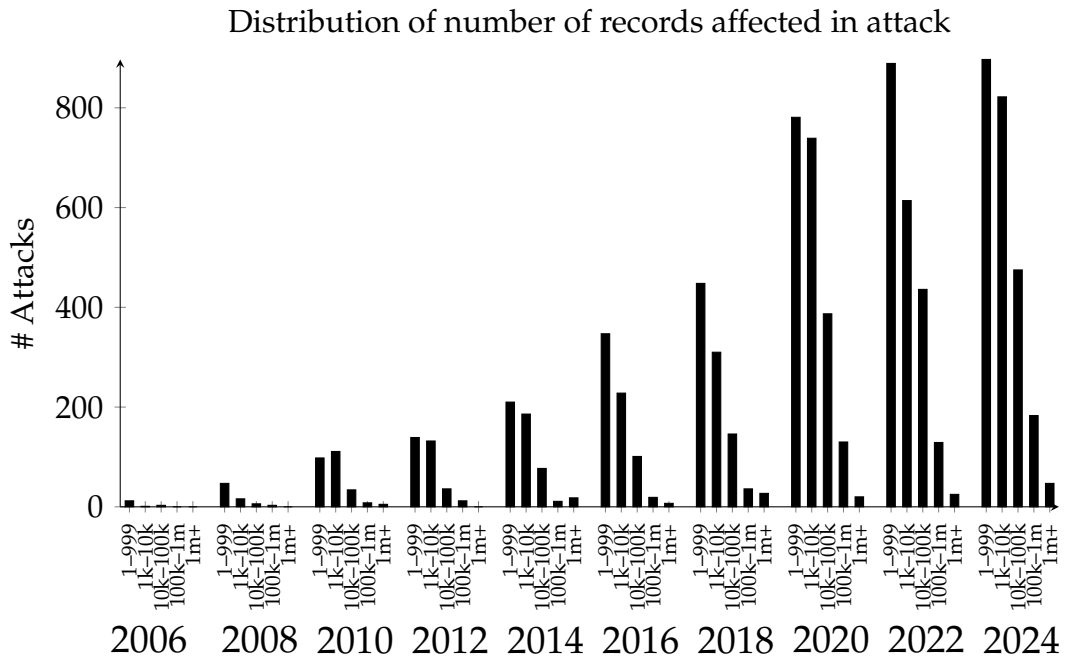


Figure 2: Distribution of attacks across record loss classes. All data from the Privacy Rights Clearinghouse.

ing data from the VERIS[10] Community Database, constructed by the Verizon RISK team using information crowdsourced from partners. The VERIS dataset has the advantage of including information on firm size categories and of capturing cyber attacks other than data breaches, including ransomware attacks. Figure 3 shows that for private-sector organisations primarily operating in the US, between 2010 and 2022[11], firms with employment between 101 and 1000 increased their share of reported incidents over two-fold from around 16% to nearly 36%.

2.4 The typical cyberattack lasts under a year, with notable outliers

Our data includes information on the starting and ending date of a data breach. However, this information is missing for a large number of observations, and we note that it is particularly likely to be missing for the largest attacks - which also have longer durations, on average. In figure 4, we construct the distribution of attacks after reweighting observations to match the overall distribution within year, firm organization type and bins of

[10]VERIS stands for the Vocabulary for Event Recording and Incident Sharing, and represents a unified framework within which organisations can anonymously share information about security incidents. The data is publicly available on GitHub and is used by Verizon to produce its annual Data Breach Investigations Reports, published since 2008.

[11]Prior to 2010, the number of reported incidents in the VCDB is too low for us to draw meaningful conclusions.

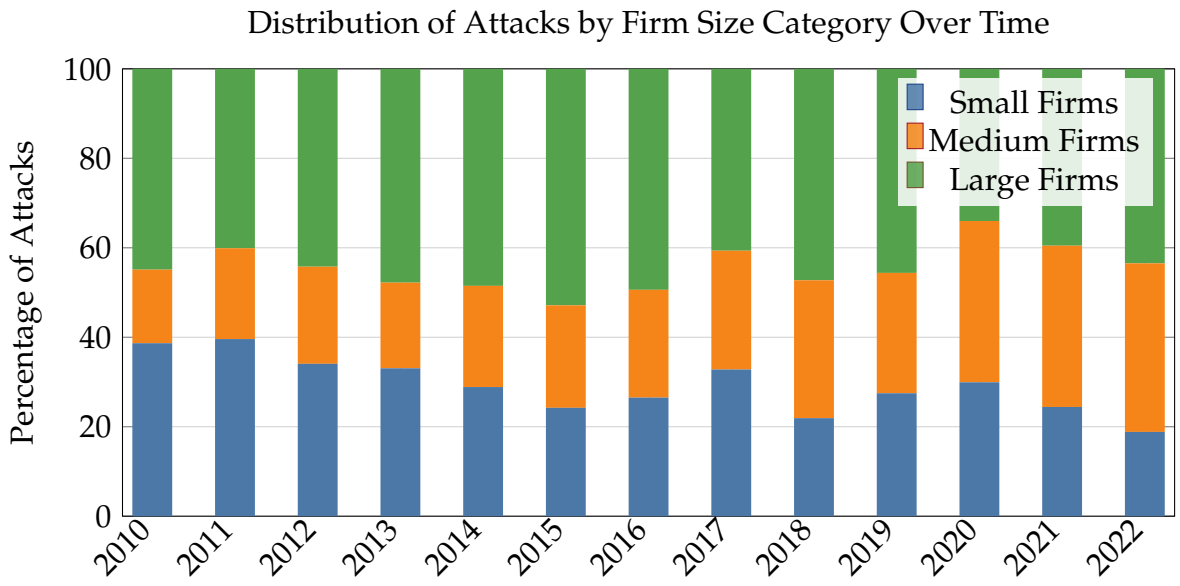


Figure 3: Distribution of attacks across firm size classes. Small firms: 1-99 employees, medium firms: 100-999 employees, large firms: 1000+ employees. All data from the VERIS database maintained by Verizon's RISK Team.

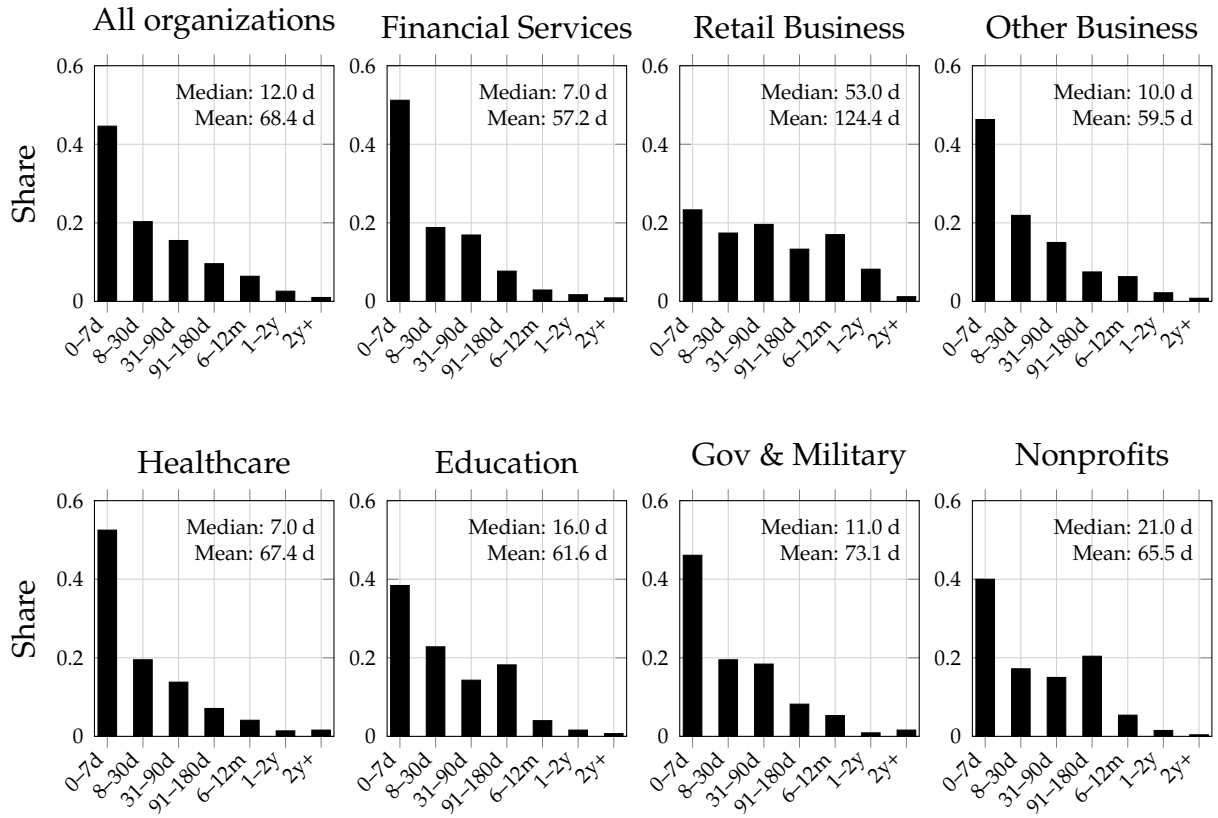


Figure 4: Distribution of attack durations, measured as the number of days between the start and end of a breach. All data from the Privacy Rights Clearinghouse.

attack size. We find that the mean attack duration is around 68 days. Our reweighting procedure is imperfect, and indeed, our estimated mean attack duration is much lower than other estimates, such as the 241-day typical time to identify and contain a breach estimated by the Ponemon Institute (Ponemon Institute and IBM 2025). However, our data also contains more small firms than the sampling frame of the Ponemon Institute's survey, and these smaller firms may be subject to lower complexity attacks which are quicker to recover from.

We note, further, that the disruptions from a cyberattack persist for significantly longer than the duration of the breach itself. These disruptions can include costs to business reputations, delays in reconstructing internal databases and time spent in reconfiguring security procedures. While the stock market reactions to cyberattacks are typically muted, suggesting that capital markets generally view attacks as short-term incidents with minimal impact on long-term firm performance, we note that estimation of the effects of large cyberattacks is complicated by the absence of data. A small literature finds that these effects can generally be detected for at least one year post-incident (Kamiya et al. 2021; Erkan-Barlow et al. 2023). In our model below, we report results assuming that the impact of a cyberattack lasts one year.

2.5 Indirect losses from a cyberattack can be massive

Estimates of the indirect losses from cyberattacks vary widely due to severe data challenges and issues around defining the exact nature of these losses[12] and yet it is generally accepted that the indirect losses dominate direct losses. Cybersecurity and Infrastructure Security Agency (2020) describe in detail the challenges associated with measuring the costs of a cyberattack and in reconciling the wide range of estimates available. Cobos and Cakir (2024) provide a rich review of the literature on estimating the costs of an attack, discussing evidence on losses driven by negative stock market reactions and reputational effects, the impacts of supply chain disruptions and spillover effects, and the delayed announcement of breaches. Relying on data from the NotPetya breach in 2017, Crosignani, Macchiavelli and Silva (2023) show that supply chain propagation of the effects of a major cyberattack can result in a four-fold amplification of the initial shock, which represents an empirical target for parameters governing spillover effects in our model below.

[12]Kamiya et al. (2021) show that the costs associated with investigation, remediation, legal and regulatory penalties associated with 75 major first-time cyberattacks were dominated by the loss in shareholder value by a factor of nearly 100. It is, however, unclear whether the entire loss in shareholder value, $104 billion, estimated by the authors should be attributed to the attack alone. For instance, susceptibility to a cyberattack may provide investors with adverse information about the resilience of a firm to alternative forms of risk, hence depressing its market value.

2.6 Cybersecurity investment rises with firm size

Estimates of the amount spent on cybersecurity investments are rare, and generally obtained by proprietary surveys which under-represent small firms[13]. To study how cybersecurity investments scale with firm size, we thus rely on data on the employment of cybersecurity professionals across firms. Our data comes from LinkedIn profiles compiled by Revelio labs, a provider of workforce intelligence. Revelio Labs' data covers nearly 90 million worker profiles a year over the 25-year period we study, including over 290,000 cybersecurity-related job spells, across nearly 8 million unique firms. Our sample is constructed using worker-firm job spells with a valid occupation, starting date and ending date between January 2000 and November 2024, when our data was obtained. We exclude all observations associated with firms that do not contain a valid Revelio identifier, but do not otherwise restrict our sample. We rely on Revelio's `k1500` occupation classification and identify cybersecurity related job spells by searching for titles containing the strings combining "{IT, information, cyber, network}" with "security", and "security" with "{analyst, architect and engineer}". Our results are robust to a wide range of alternative specifications, including one where we rely only on occupations mapped to the O*NET code 15.1212, Information security analysts.

Figure 5 shows a clear pattern: larger firms devote a higher share of employment to cybersecurity, with a pronounced rise among firms with more than 10,000 employees since the late 2010s. Table 2 shows that this correlation remains positive even when we consider variation within industries and firms over time, with a remarkably stable semi-elasticity - a firm that is twice as large has a cybersecurity share of its workforce that is almost a percentage point larger.

3 Macroeconomic Model of Cyber Attacks

There exists no established general equilibrium macroeconomic model to analyze cyber attacks. In this section we build one.

Model overview. In the model heterogeneous productivities are attacked by cyber agents and can choose to defend themselves through cybersecurity investments. The model nests a search and matching framework within a model of firm entry and exit à la Hopenhayn (1992). Time is discrete, and indexed by $t = 0, 1, \ldots$. There is one final good in

[13]For instance, Moody's Ratings (2023) estimates that corporations spent a median share of 8% of their IT budgets on cybersecurity, with a range of 5-10% across sectors. However, their sampling frame is firms whose debt issuances are rated by Moody's, which is likely to overrepresent large firms.

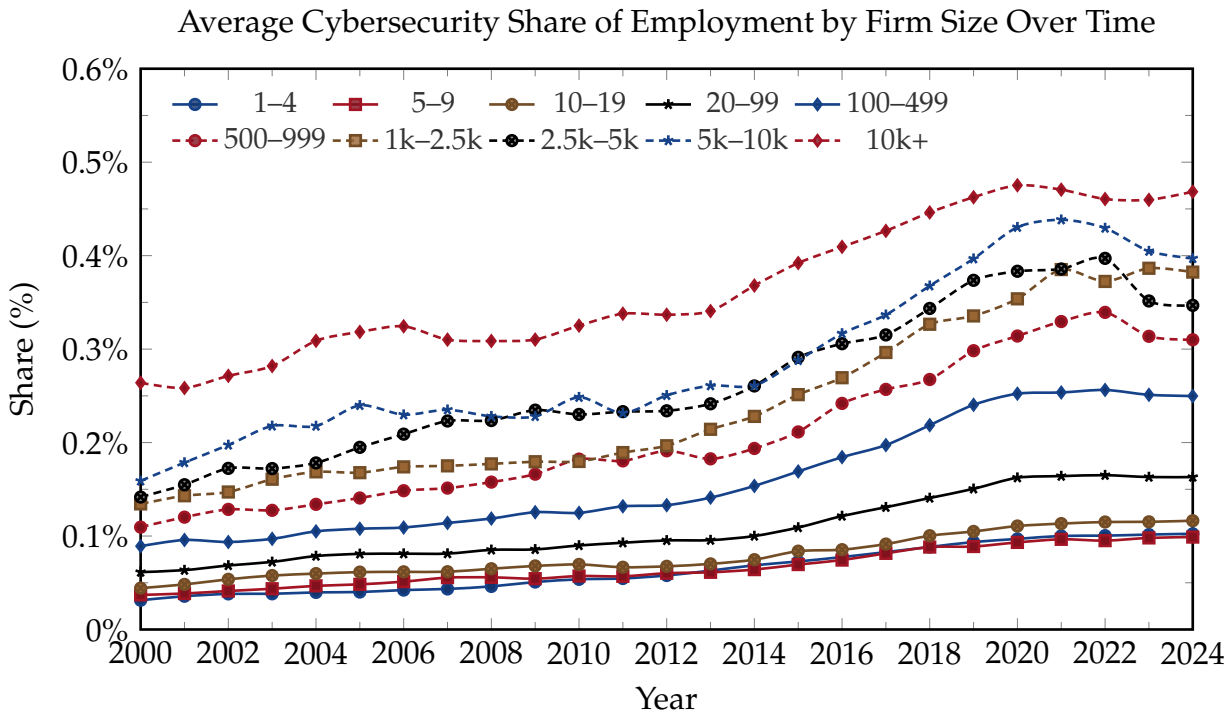


Figure 5: Share of cybersecurity-related employment across firm-size categories. Source: LinkedIn/Revelio Labs and authors' calculations.

Variable	Share of cybersecurity-related employment		
log(Total Workers)	0.0128***	0.0092***	0.0167***
	(0.0002)	(0.0002)	(0.0004)
Year FE	No	Yes	Yes
Industry FE	No	Yes	No
Firm FE	No	No	Yes

Table 2: Estimates of the semi-elasticity of the share of cybersecurity-related employees to firm size. All regressions are based on the set of firms in Revelio Labs data reporting at least one US employee with a valid industry code and firm identifier. Firm size is the number of employees of a firm recorded in the Revelio Labs dataset.

the economy which is used solely for consumption and is the numeraire. There are three sets of agents, households, firms and cyberattackers. Households supply labor to firms, which use labor to produce the final good using a technology featuring strict diminishing returns. All markets are competitive.

Firms are subject to occasional cyberattacks, which disrupt production and have spillover effects on aggregate productivity. To avoid the effects of attacks, firms make upfront investments in cybersecurity. In each period, attackers in our economy pay a cost to search

for firms of a given productivity level. Our setup is therefore analogous to a directed search model in which "markets" are segmented by productivity. Within each market, the probability of a firm and a cyberattacker meeting depends on the "tightness" of this market, the number of firms with that productivity level divided by the number of attackers targeting such firms.

When a firm and attacker meet, the latter decides on the quantity of resources it will devote to attacking the firm, subject to a limit. The outcome of a firm-attacker meeting depends on the level of cybersecurity and the chosen attack intensity; the higher the intensity of the attack relative to the cybersecurity level, the more likely an attack is to be *successful*. Successful attacks result in a loss of revenue for the firm and a positive share of revenue captured by the attacker.

In equilibrium, firms anticipate the likelihood of an attack and choose their ex-ante cybersecurity investments accordingly, trading off a convex cost of such investments against the value of a reduced likelihood of being attacked. Attackers anticipate the level of investment the typical firm of a given productivity level is likely to make, targeting their attacks against only those firm types whose bounty is sufficiently attractive and whose investment level is not excessively high. Fixed costs induce relatively low productivity firms to exit, and higher likelihoods of being attacked reduce the incentives of firms to enter, together shaping the equilibrium firm size distributions and hence determining aggregate productivity and output.

3.1 Firms

There is a mass of firms that produce a homogeneous good, which is the numeraire. Firms use labor as the only input into production. There is one unit of inelastically supplied labor whose price is w in the whole economy. A firm with productivity φ produces according to the production function $y(Z, \varphi, n) = Z\varphi n^{\alpha}$ where n is the firm's labor demand and $\alpha \in (0, 1)$ is the degree of returns to scale. Z is an aggregate productivity term all firms in the economy are subject to – we will return to Z later in detail. In addition to labor costs, firms must pay a per-period fixed cost f to operate. Firms discount the future with factor β and are hit with an exogenous exit shock with probability σ. Finally, firms' productivity follows a first-order Markov process with transition probability $\Pi(\varphi'|\varphi)$, which is a discretization of the following AR(1) process:

$$\log(\varphi') = \rho_z \log(\varphi) + \epsilon \tag{1}$$

where ρ_z is the persistence of the process and $\epsilon \sim \mathcal{N}(0, \sigma_z)$ is the innovation hitting productivity each period.

Safe firms. A *safe* firm is one that is not currently under a cyber attack. Safe firms choose the amount of labor to employ and the size of the cyber security investment, $x \geq 0$, which we assume is denominated in units of labour. We assume that cyber hiring x leads to cybersecurity level $C(x)$, where $C' > 0, C'' < 0$. The value of a safe firm with productivity φ is

$$\begin{aligned} V^s(Z,\varphi) &= \max_{n,x\geq 0} Z\varphi n^\alpha - wn - wx - f \\ &\quad +\beta(1-\sigma)\mathbb{E}_{\varphi'|\varphi}\left[q(Z',\varphi')V^m\left(Z',\varphi',x\right) + \left(1-q(Z',\varphi')\right)V^s(Z',\varphi')\right] \quad (2) \end{aligned}$$

What awaits the firm in the next period, if it survives, is one of two states: with probability $(1 - q(Z', \varphi'))$ the firm remains safe, and with probability $q(Z', \varphi')$ the firm meets a cyber attacker and derives value $V^m(Z', \varphi', x)$. We characterize V^m below.

Since there are no adjustment costs or search frictions in the labor market, the choice of labor is effectively a static one that solves the profit maximization problem

$$\pi(Z,\varphi;w) \equiv \max_{n\geq 0} Z\varphi n^\alpha - wn$$

which yields the labor demand curve for a firm, $n(Z, \varphi; w) = \left(\frac{w}{\alpha Z\varphi}\right)^{\frac{1}{\alpha-1}}$. The maximized profits for a firm are given by

$$\pi(Z,\varphi;w) = Z\varphi^{1-\alpha}\left(1-\alpha\right)\left(\frac{w}{\alpha}\right)^{\frac{\alpha}{\alpha-1}}$$

The firm's choice of cybersecurity investment x is given by the FOC equating the marginal benefit and marginal cost of cyber security,

$$w = \beta(1-\sigma)\mathbb{E}_{\varphi'|\varphi}\left[q(Z',\varphi')\frac{\partial V^m(Z',\varphi',x)}{\partial x}\right] \quad (3)$$

We can interpret the left hand side as the marginal cost of investing in cybersecurity, namely the wage paid to the labor invested in cyber security. The right-hand side is the benefit of this investment discounted by the firm's effective discount factor taking survival into account. We will describe this benefit next. We assume that cybersecurity investments must be made on an ongoing basis, and that the level of cybersecurity at date t is entirely determined by investments made at $t - 1$. Assuming instead that cybersecu-

rity levels were persistent would add a state variable to the model without substantially impacting the economics of our results. Further, given the continual emergence of new threats faced by firms, we expect that the level of cyber protection a firm has must be continuously refreshed in order to remain effective, which lends credence to our modelling approach.

Firms under attack. Meeting a cyber attacker does not imply a cyber attack will be successful. When attacker and firm meet, the value to the firm is

$$V^m(Z,\varphi,x) \;=\; \Lambda\left(a^*(x),x\right)V^a(Z,\varphi) + \left[1-\Lambda\left(a^*(x),x\right)\right]V^s(Z,\varphi) \tag{4}$$

With probability $\Lambda\left(a^*(x),x\right)$, where $a^*(x)$ is the optimal choice of attack intensity by the attacker, the firm is successfully attacked and gets value $V^a(Z,\varphi)$; otherwise the firm is safe. Note that the cyber security investment is *predetermined*, capturing the idea that cyber security investments must be ongoing and act as insurance that protects against attacks when they occur. We will describe the determination of Λ below.

When a firm is successfully attacked, it loses a share ℓ of its output but must continue to pay the fixed cost of production and its labor bill. To capture the fact that cyber attacks have persistent consequences, the attacked firm escapes the attack with probability δ. If the firm chooses to, it can exit altogether. The value of a firm under attack is

$$\begin{aligned} V^a(Z,\varphi) = \max\big\{0,\; \max_n (1-\ell)\cdot Z\varphi n^\alpha - wn - f \\ + \beta(1-\sigma)\mathbb{E}_{\varphi'|\varphi}\left[(1-\delta)V^a(Z',\varphi') + \delta V^s(Z',\varphi')\right]\big\} \end{aligned} \tag{5}$$

We note that the profit maximization problem pinning down the choice of labor by the firm is isomorphic to that of a safe firm with a reduced productivity level of $(1-\ell)Z\varphi$, facilitating the computation of the labor demand, $n^a(Z,\varphi;w)$, and profit functions, $\pi^a(Z,\varphi;w)$. By the Envelope theorem, we have from equation 3 that

$$w = \beta(1-\sigma)\mathbb{E}_{\varphi'|\varphi}\left[q(Z',\varphi')\frac{\partial\Lambda\left(a^*(x),x\right)}{\partial x}\left(V^a(Z',\varphi') - V^s(Z',\varphi')\right)\right] \tag{6}$$

Firm entry and exit. There is a large pool of prospective entrant firms. To enter the market, firms must first pay an entry cost f_e and draw a productivity φ from a Pareto distribution with scale parameter ϖ whose (discretized) probability mass function is $g(\varphi)$ with (discretized) cumulative distribution function is $G(\varphi)$ and whose productivity states are $\{\varphi_1,\ldots\varphi_P\}$.

At this stage, firm φ can choose whether to become *operational*, or exit the market. The value of an entrant with productivity φ is

$$V^e(Z,\varphi) = \max\left\{0, V^s\left(Z,\varphi\right)\right\} \tag{7}$$

What emerges from this is that there exists a set $\Phi = \{\varphi : V^e\left(Z,\varphi\right) \geq 0\}$ that consists of the profitable productivities whereas firms with $\varphi \notin \Phi$ will immediately exit.[14]

Because there is an unbounded set of prospective entrants, the free entry condition for firms in this economy is

$$\sum_{p\in\Phi} V^e\left(Z,\varphi\right) g(\varphi) - f_e = 0 \tag{8}$$

Because productivity is stochastic, safe firms may *endogenously* choose to exit if their productivity state falls outside of Φ. An operational firm that is under attack, may also no longer find it worthwhile to stay in the market. Therefore, there is a set $\Phi^{'}$ of productivities where firms find it worthwhile to remain in the market, defined by $\Phi' = \{\varphi : V^a\left(Z,\varphi\right) > 0\} \subset \Phi$ whereas firms whose productivity $\varphi \in \Phi \setminus \Phi'$ will exit when under attacked.

If the mass of entrants in a given period is M, the total number of entrants who become operational equals $M\sum_{\varphi\in\Phi} g(\varphi)$. Denote the total number of firms in the market by $L = \sum_{\varphi\in\Phi} L(\varphi)$, where $L(\varphi)$ is the sum of firms that are safe or under attack in market φ. Denote by $s(Z,\varphi)$ the share of safe firms in market φ. A share σ of firms in the market in any period will exit exogenously. No safe firm in the market will voluntarily choose to exit. All firms under attack whose productivity $\varphi \in \Phi \setminus \Phi'$ will voluntarily exit the market. Additionally, the shares $s(Z,\varphi)\sum_{\varphi'\notin\Phi}\Pi(\varphi'|\varphi)$ and $(1-s(Z,\varphi))\sum_{\varphi'\notin\Phi'}\Pi(\varphi'|\varphi)$ will also voluntarily exit, the first are safe firms and the second are firms under attack that transition into non-profitable productivities. In a steady state equilibrium, the mass of firms in each operational market $\varphi \in \Phi'$ evolves according to

$$\begin{aligned} L\left(\varphi\right) \;=\; & \underbrace{(1-\sigma)L_{-1}(\varphi)}_{\text{no exogenous separation}} + \underbrace{Mg(\varphi)}_{\text{new entrants}} \\ & \underbrace{-(1-\sigma)L_{-1}(\varphi)\left(\sum_{\varphi'\neq\varphi}\Pi(\varphi|\varphi')\right)}_{\text{exit to }\varphi'} + \underbrace{(1-\sigma)L_{-1}(\varphi')\left(\sum_{\varphi'\neq\varphi}\Pi(\varphi'|\varphi)\right)}_{\text{entry from }\varphi'} \end{aligned}$$

[14] While in Hopenhayn (1992) there exists a lower bound φ^* above which all productivities are profitable, this is not necessarily the case here because mid-productivity firms may have lower expected profits than low-productivity firms if they fall victim to cyber attacks more frequently.

while in market $\varphi \in \Phi \setminus \Phi'$ the total mass of exiting firms (the RHS) must equal the total mass of entering operational firms (the LHS)

$$M \sum_{\varphi \in \Phi} g(\varphi) = \sigma L + (1-\sigma) \sum_{\varphi \in \Phi \setminus \Phi'} (1 - s(\theta))\, L(\varphi)$$

Breaking this down by productivity market, the total entrants and exiters in market φ are

$$entrants(Z,\varphi) = \begin{cases} Mg(\varphi) & \text{if } \varphi \in \Phi \\ 0 & \text{otherwise} \end{cases} \quad exiters(Z,\varphi) = \begin{cases} \sigma L(\varphi) & \text{if } \varphi \in \Phi' \\ \sigma L(\varphi) + (1 - s(Z,\varphi))\, L(\varphi) & \text{if } \varphi \in \Phi \setminus \Phi' \\ 0 & \text{otherwise} \end{cases}$$

Equality of entrants and exiters in a steady state implies the mass of operational firms in market φ is

$$L(\varphi) = \begin{cases} \frac{Mg(\varphi)}{\sigma} & \text{if } \varphi \in \Phi' \\ \frac{Mg(\varphi)}{\sigma + (1 - s(Z,\varphi))} & \text{if } \varphi \in \Phi \setminus \Phi' \\ 0 & \text{otherwise} \end{cases}$$

Additionally, the share of safe firms is governed by the following law of motion in steady state. If $\varphi \in \Phi \setminus \Phi'$ then $s(Z,\varphi) = 1$ since all firms under attack will voluntarily exit the market. If $\varphi \in \Phi'$ the share of safe firms evolves according to $(1 - s'(\varphi)) = (1 - \delta)(1 - \sigma)(1 - s(Z,\varphi)) + s(Z,\varphi)(1-\sigma) q(\varphi) \Lambda\left(a^*(\varphi), x^*(\varphi)\right)$ where s' is the share of safe firms in the next period. Because in steady state $s' = s$, we have

$$s\left(Z,\varphi\right)|_{\varphi \in \Phi'} = \frac{1 - (1-\delta)(1-\sigma)}{1 - (1-\delta)(1-\sigma) + (1-\sigma) q\left(\varphi\right) \Lambda\left(a^*\left(\varphi\right), x^*\left(\varphi\right)\right)}$$

3.2 Attackers

There is an unbounded set of cyber attackers who start out as *vacant*, that is they are not engaged in an attack. Vacant attackers can meet firms, which will lead to either a successful strike or a failed strike.

Vacant attackers. Vacant attackers must pay a fixed cost κ to acquire the necessary computational resources to strike a firm. Additionally, attackers must choose which firms φ to strike. An attacker who chooses firm φ will meet a firm of productivity φ in the next period with probability $\lambda(Z,\varphi)$. The value of a vacant attacker targeting firms φ is

$$W^v(Z,\varphi) = -\kappa + \left\{\lambda(Z,\varphi) W^m\left(Z,\varphi, x^*(Z,\varphi)\right) + \left(1 - \lambda(Z,\varphi)\right) W^v(Z,\varphi)\right\} \tag{9}$$

where $W^m(Z, \varphi, x^*(Z, \varphi))$ is the value derived from meeting a firm whose optimal cyber security investment is $x^*(Z, \varphi)$.

Active attackers. When an attacker meets a firm it chooses the attack intensity a with which to strike the firm. The value of meeting a firm of productivity φ and with cyber security investment x is

$$W^m(Z, \varphi, x) = \max_{a \in [0,A]} -a + \Lambda(a, x) W^a(Z, \varphi) + [1 - \Lambda(Z, a, x)] W^v(Z, \varphi) \tag{10}$$

where a translates into attack intensity $\mathcal{A}(a)$, which we assume is strictly increasing and concave. Note that the attacker chooses the resources to devote to the attack *after* observing the cyber security level of the firm. The FOC that pins down the attack intensity a upon meeting a firm with cyber security x is

$$\frac{\partial \Lambda(a, x)}{\partial a} [W^a(Z, \varphi) - W^v(Z, \varphi)] = 1 + \mu_A - \mu_0 \tag{11}$$

where μ_0 and μ_A are the Lagrange multipliers for the constraints $a \geq 0$ and $a \leq A$. Note that only one of these constraints can bind at a time.

If the attack is successful, the attacker gains value

$$W^a(Z, \varphi) = \chi \cdot Z\varphi \cdot (n^a)^\alpha + \beta \{(1 - \sigma) [(1 - \delta) W^a(Z, \varphi) + \delta W^v(Z, \varphi)] + \sigma W^v(Z, \varphi)\} \tag{12}$$

where χ is the share of output the cyber attacker is able to retain.

Successful Attacks. When a firm with cyber security x meets an attacker with attack intensity a, they respectively draw cyber shocks ϵ_f, ϵ_a from a Gumbell distribution with scale parameter ϱ. An attack is *successful* if $\mathcal{A}(a) + \epsilon_a \geq C(x) + \epsilon_f$ for a function $\mathcal{A}(\cdot)$ which represents the realized attack given intensity a and so the probability of a successful attack is

$$\Lambda(a, x) \equiv \Pr\left[C(x) + \epsilon_f \leq \mathcal{A}(a) + \epsilon_a\right] = \frac{\exp\left\{\frac{\mathcal{A}(a) - C(x)}{\varrho}\right\}}{1 + \exp\left\{\frac{\mathcal{A}(a) - C(x)}{\varrho}\right\}} \tag{13}$$

Note that the probability of an attack succeeding depends only on the relative levels of attack intensity a and cyber hiring x. However, in equilibrium, both x and a are contingent on firm productivity φ, implying that the probability of a successful attack varies across firm productivity levels.

3.3 Cyber Externalities

We introduce aggregate productivity spillovers arising from cybersecurity risk through the term Z, which affects firm-level production. This term captures the idea that cyberattacks generate negative externalities that extend beyond directly targeted firms. Through interconnected supply chains, shared payment systems, common software dependencies, and other forms of digital linkages, even firms that are not attacked can suffer disruptions or productivity losses.[15].

Analogous to the positive knowledge spillovers in Romer (1986, 1990), where aggregate productivity rises with the economy-wide stock of knowledge, we model cyber risk as generating a negative externality: aggregate productivity declines as a larger share of firms are compromised by a cyber attack. Formally, we specify

$$Z = \left(\int_{\varphi} s(Z, \varphi)\, d\varphi \right)^{\epsilon_{\text{att}}} \tag{14}$$

where $s(Z, \varphi)$ denotes the share of firms of type φ currently under attack. The higher the economy-wide prevalence of attacks, the greater the drag on overall productivity, even for firms that remain unbreached.

3.4 Household

There is a representative household that derives utility by consuming the homogeneous good, earns wage w by providing a unit of labor inelastically, and earns firm profits. Its welfare is defined by

$$\mathcal{W} = \log(C) \tag{15}$$

where the aggregate resource constraint it $C = w + T + \Pi - f_e M \int_{\varphi} g(\varphi)$. Consumption is financed via labor income net of transfers[16] T, and firm profits net of entry costs, where profits are defined as

$$\Pi = \int_{\varphi} \Big\{ s(Z, \varphi) \left[\pi(Z, \varphi; w) - w x(\varphi) \right] \tag{16}$$

$$+ (1 - s(Z, \varphi))(1 - \ell)\pi(Z, \varphi; w) \Big\} D(\varphi) d\varphi \tag{17}$$

[15] It is commonplace in the macro literature to have shocks to individual firms propagate through networks consistent with the network amplification mechanisms emphasized by Acemoglu et al. (2012)

[16] In the baseline, transfers are zero. We introduce them here anticipating our discussion of policies a government could undertake to reduce cyber risk in section 6.

We assume cyber agents are foreigners and therefore exclude the real resources paid to acquire attack intensity and the amount earned by attackers from the domestic resource constraint[17].

3.5 Equilibrium

Matching Technology. Let $s(Z,\varphi)$ be the share of firms of productivity φ that are safe from an attack and let $v(\varphi)$ be the vacancy rate of cyber attackers in market φ, defined as the number of vacant attackers per firm. If the mass of firms operating in market φ is $L(\varphi)$, only $s(Z,\varphi)L(\varphi)$ safe firms and $v(\varphi)L(\varphi)$ cyber vacancies give rise to new cyber attacks. The number of cyber attacks is $m\left(s(Z,\varphi)L(\varphi), v(\varphi)L(\varphi)\right)$, which is increasing in both arguments, concave, and homogeneous of degree 1. Defining cyber tightness as $\theta = \frac{s}{v}$, the vacancy-filling and the attack rates are

$$q\left(Z,\theta(Z,\varphi)\right) = m\left(\frac{s(Z,\varphi)}{v(Z,\varphi)},1\right) \tag{18}$$

$$\lambda\left(Z,\theta(Z,\varphi)\right) = m\left(1,\frac{v(Z,\varphi)}{s(Z,\varphi)}\right) = \theta(Z,\varphi)q\left(Z,\theta(Z,\varphi)\right) \tag{19}$$

Steady State Equilibrium. The steady state equilibrium is a collection

$$Z, w, M, \left\{x(\varphi), a(\varphi), n(\varphi), \theta(\varphi), s(Z,\varphi), \Lambda(\varphi), \theta(\varphi), \lambda(\varphi), q(\varphi)\right\}_{\varphi\in\Phi}$$

such that

1. Given w, firms choose $n(\varphi)$ to maximize their static profits and $x(\varphi)$ is chosen to solve their dynamic problem.
2. Taking the firm choices of security $x(\varphi)$ as given, attackers' choice of $a(\varphi)$ solve their dynamic problem.
3. Firm's free entry condition holds:

$$\sum_{\varphi\in\Phi} V^e\left(\varphi\right) g(\varphi) - f_e = 0$$

4. The probability of a successful attack is consistent with choices made by the firms

[17]The Council on Foreign Relations estimates that 77% of attacks on U.S. firms originated in five of its geopolitical rivals.

and attackers in each submarket,

$$\Lambda(\varphi) = \frac{\exp\left\{\frac{\mathcal{A}(a(\varphi))-C(x(\varphi))}{\varrho}\right\}}{1+\exp\left\{\frac{\mathcal{A}(a(\varphi))-C(x(\varphi))}{\varrho}\right\}}$$

5. Free entry for attackers targeting each productivity φ holds:

$$W^v(\varphi) = 0$$

6. The labor market clears

$$1 = \sum_{\varphi\in\Phi'} n^a(\varphi)(1-s(Z,\varphi))L(\varphi) + \sum_{\varphi\in\Phi} n(\varphi)s(Z,\varphi)L(\varphi)$$

   where the first summation is the labor demand coming from safe and attacked firms when endogenous exit does not occur and the second summation is the labor demand coming from safe firms only because attacked firms endogenously exit.

7. The externality in equation (14) holds.

8. Cyber tightness and the contact rates for firms and attackers are consistent with the matching function in the cyber market.

3.6 Discussion of Model Assumptions

One of this paper's main contributions is the development of a first macroeconomic, general equilibrium model to analyze cyber attacks and study policy related to attacks. In this subsection, we discuss some key modelling decisions and how they map to salient features that underpin the relationship between cyber attackers and the firms they strike.

Assumption #1: Stackelberg Structure. In our model, firms pre-commit to a level of defense, acting as Stackelberg leaders. Our modeling approach is standard in the literature on security games[18] and is closest in spirit to Ebel and Mitra (2024), and builds on the intuition that firm-level security investments are ongoing investments in generally maintaining a level of security which is broadly observable. For instance, firms may actively disclose investments in cyber security in annual reports to investors (Al Amosh

[18]This assumption is also standard in the game theory literature on security investments more broadly (Korzhyk et al. 2011), which is based on the idea that attackers constitute rational agents able to observe ongoing security investments and activities. In the cybersecurity context, a recent piece extending the analysis of Stackelberg security games is Huang, Naghizadeh and Liu (2024).

and Khatib 2025). The key question in interpreting this assumption is whether the typical attacker can infer the cybersecurity level at a prospective target based on observables like firm performance and size with a high degree of precision; we argue that this is eminently plausible. Direct evidence for such a mechanism would require a systematic forensic analysis of cyberattackers and their information sets.

Assumption #2: Search, Matching and Randomness in Attack Success. We model two layers of randomness in the rate at which attackers and firms meet: first, there is uncertainty due to the search framework in which attacks on firms that are more likely to be targeted are less likely to meet the firm at all, and second, conditional on meeting, attacks may fail even if the realised intensity of an attack exceeds the security level of the firm. We now discuss both layers.

The first layer of uncertainty, driven by imperfect matching, reflects the fact that targeting the most attractive targets is subject to congestion. Intuitively, we wish to model the idea that attackers face congestion frictions in attacking the largest firms. This congestion arises because at any given point in time, attackers who successfully breach a firm crowd out other attackers in three possible ways. First, ransoming a firm's data may physically prevent access to other potential attackers by encrypting files and rendering systems or data inaccessible (Microsoft 2023; National Cyber Security Centre 2023). Second, data compromised by a first attacker lowers the marginal return to additional attackers because the target's systems are already disrupted or partially monetized (though repeat attacks can still occur Federal Bureau of Investigation and CISA 2023). Third, attacked firms typically respond by isolating systems, reducing exposure, and patching vulnerabilities, making them less likely targets and hence less attractive targets for subsequent attackers (Cybersecurity and Infrastructure Security Agency 2025).

The second layer of uncertainty is that conditional on matching, attackers may fail to breach a firm's defenses stochastically. We do this because whether or not an attack succeeds depends both on the decisions of the parties at play – firms choose how much to invest on their defenses and attackers choose which firms and how intensely to attack them – but also from randomness such as individuals clicking malicious links (Verizon 2024; National Institute of Standards and Technology 2023). Further, we note that attacks on firm systems often fail to breach all the defenses they seek to overcome; rather, attackers are often caught and blocked rapidly, a factor that Ponemon Institute and IBM (2025) argues is responsible for the minor decline in losses from a typical attack.

Assumption #3: Spillovers. We model spillovers as productivity losses which are a concave function of the share of firms under attack: as the share under attack grows, the rate at which spillover-driven losses rise increases. Our motivation for this is an extensive literature documenting that cyberattacks generate losses via network effects through supply chains (Acemoglu, Malekian and Ozdaglar 2016) or through common choices for security vendors (Anand, Duley and Gai 2022), and that losses from densely connected networks are initially convex in the share of affected nodes. Extending our work to include explicit network-driven spillovers and subsequent losses is an important area for future research, and will require much more data to adequately discipline.

4 Calibration

We calibrate the model in two stages. First, we discipline it to a "no-cyber" steady state designed to approximate the U.S. economy around the year 2000, when firms were largely unaffected by cyber threats. Second, we re-calibrate the model to match key features of the U.S. economy in the 2020s, a period characterized by the pervasive presence of cyberattacks.

4.1 "No cyber" calibration

We choose the 2000 baseline to represent a period in which firms were essentially not subject to cyber attacks[19]. To do so, we set the share of output lost by firms and the corresponding gain accruing to attackers during a successful cyber incident to zero, ensuring that cyber activity is absent in equilibrium. With these restrictions, no firm has an incentive to incur the cost of cybersecurity investment, and no attacker has an incentive to pay the fixed cost of entry.

The core technology and discounting parameters are set externally. The degree of decreasing returns is $\alpha = 0.85$, in line with the estimate in Ottonello and Winberry (2020), and the discount factor is $\beta = 0.96$, targeting a 4 percent annual interest rate. Total labor supply is normalized to one. The remaining parameters are estimated via simulated method of moments to match three salient features of the U.S. firm distribution in 2000: (i) the distribution of entrants, (ii) the overall distribution of firm sizes, and (iii) the average

[19] In 2000, digital penetration in the economy was still in its nascent stages, with internet access being highly limited. Only 6% of the world population was online in 2000, compared to 68% by 2025. As of March 2001, only 6.3% of US firms were procuring inputs online (OECD 2001), and in 1999, sales over the internet amounted to 0.6% of total sales (while Canada was at 0.2%) (StatCan 2000). Accordingly, cyberattacks in this period were largely conducted for notoriety, and limited digitalisation also limited losses from such attacks(ICAEW Insights 2023).

exit probability of firms. This yields estimates for the shape of the Pareto productivity distribution at entry, ϖ, the exogenous exit probability σ, the fixed costs of entry and operation (f_e and f), and the parameters governing the log-productivity process, $(\mu_z, \rho_z, \sigma_z)$. Table 3 reports the resulting parameter values.

The implied firm and entrant size distributions are shown in Figure 6. The model closely matches the empirical distributions and generates an average exit probability of 6.0%, compared to the observed 6.2% from Lee and Mukoyama (2015).

4.2 Cyber calibration

Holding fixed the parameters estimated in the "no cyber" economy, we calibrate the additional parameters that govern cyber risk. The escape probability, δ, is set to match the half-life of the decline in firm sales following a cyber attack. We normalize the EV scale parameter to $\varrho = 1$, the maximum attack capacity to $A = 1$, and the cost-scaling parameter of the attack function to $\kappa_A = 1$.

The remaining parameters are estimated via simulated method of moments to match four sets of moments of the cyber economy in the 2020s: (i) the average cybersecurity employment share by firm size in 2024, (ii) total cyber losses equivalent to 1.16 percent of GDP, (iii) the share of firms experiencing a cyber incident (0.67 percent), and (iv) an amplification factor linking direct to total economic losses from cyber shocks equal to 4.

The empirical targets underlying these moments are drawn from a combination of micro and macro sources. The first moment, the average cybersecurity employment share by firm size in 2024, has been discussed in detail earlier in the paper. The second moment, total cyber losses equal to 1.16 percent of GDP, is based on macroeconomic estimates of cyber losses reported in the International Monetary Fund's Global Financial Stability Report (International Monetary Fund 2023). The third moment, the share of firms experiencing a cyber incident of 0.67 percent, is computed using incident-level loss data from Coveware as reported via Bloomberg, combined with the firm size distribution from the U.S. Business Dynamics Statistics to map observed incidents into an economy-wide firm-level attack probability. Finally, the amplification factor mapping direct to total losses is taken from Crosignani, Macchiavelli and Silva (2023), who document that propagation through supply chains raises total economic losses to roughly four times the initial direct loss.

The estimated parameters are shown in table 3. The loss and gain shares, $\ell = \chi = 0.275$, are assumed to be equal so that the resource constraint is satisfied without any further adjustments.[20] The matching function is constant returns to scale, $m(\theta) = \psi\theta^{\eta}$,

[20]This assumption is conservative: evidence suggests that firm losses can exceed attacker gains (Anderson et al. 2013).

with $\psi = 0.01$ and $\eta = 0.35$. The attack function is $K(a) = \kappa_A \frac{a^{1-\gamma_A}}{1-\gamma_A}$ with $\gamma_A = 0.2$, and the cybersecurity function is $C(x) = \kappa_C \frac{x^{1-\gamma_C}}{1-\gamma_C}$ with $\kappa_C = 0.25$ and $\gamma_C = 0.825$. The attacker vacancy posting cost is $\kappa_v = 0.0002$, and the externality parameter is 1.0001.

	Parameter (quarterly frequency)	Value
β	Discount factor	0.96
α	Returns to scale	0.85
ϖ	Pareto distribution scale	0.7
σ	Exit probability (exogenous)	0.00
f	Per-period fixed cost of production	0.183
f_e	Fixed cost of firm entry	3.83
ρ_z	Persistence in productivity	0.925
σ_z	Stdv. of productivity innovations	0.5
ℓ	Firm loss share	0.275
χ	Attacker gain share	0.275
δ	Escape probability	1
η	Elasticity of matching function	0.35
ψ	Scaling matching function	0.01
$K(a)$	Cost of attack function	$\frac{a^{1-0.2}}{1-0.2}$
$C(x)$	Cost of security function	$0.25 \cdot \frac{x^{1-0.825}}{1-0.825}$
κ_v	Elasticity of matching function	0.0002
ϵ_{att}	Romer externality	1.0001

Table 3: Model parameters. The first set of parameters is set by the authors to match an annual interest rate of 4% and to match the decreasing returns in Ottonello and Winberry (2020). The second set of parameters is estimated via SMM to match key moments from the 2000 ("no cyber") U.S. economy. The third set of parameters is estimated via SMM to match cyber-related moments from the 2020 ("cyber") U.S. economy.

The calibrated model delivers a loss-to-output ratio of 0.92% (against the empirical 1.16%), a share of firms under attack of 0.62% (vs. 0.67% in the data), and a ratio of indirect to direct losses of 3.94 (vs. 4.0 taken from Crosignani, Macchiavelli and Silva 2023). Panel A of figure 7 compares the model-predicted cybersecurity employment shares by firm size with data from Revelio Labs (via LinkedIn). The model replicates the increasing investment share in firm size, with large firms allocating 0.44% of labor to cybersecurity compared to 0.46% in the data. Panel B of figure 7 shows that the model also qualitatively captures the varying annual probability of attack of firms of different sizes – an untargeted moment. The probability of a firm being attacked peaks for mid-sized firms both in the model and in the available data from Coveware. These two last pieces of evidence establish that larger firms exhibit stronger cyber preparedness than smaller firms and that medium-sized are those most susceptible to attack.

The aggregate consequences of introducing cyber risk into the economy are shown in Table 4. Through the propagation externality, cyber attacks reduce aggregate total factor

productivity by 0.62%. Lower productivity, together with the risk of being attacked and losing revenue, reduces the value of entry for firms, leading to a 3.61% decline in the mass of entrants M. As a result of these forces, the introduction of cyber risk lowers aggregate output by 1.76%.

	Productivity (Z)	Entry Mass (M)	Output (y)
2000	1.0000	0.00370	0.4862
2020	0.9938	0.00357	0.4776
Percent change (%)	−0.62	−3.61	−1.76

Table 4: Aggregate Outcomes in 2000 and 2020

These results resonate with the mechanism discussed in Baksy, Caratelli and Olson (2025). Small firms face few attacks because they offer little reward to attackers. Large firms are attractive targets but offset this risk through extensive cybersecurity defense. In contrast, mid-sized firms occupy a vulnerable middle ground: valuable enough to attract attackers but too resource-constrained to sustain comprehensive defenses.

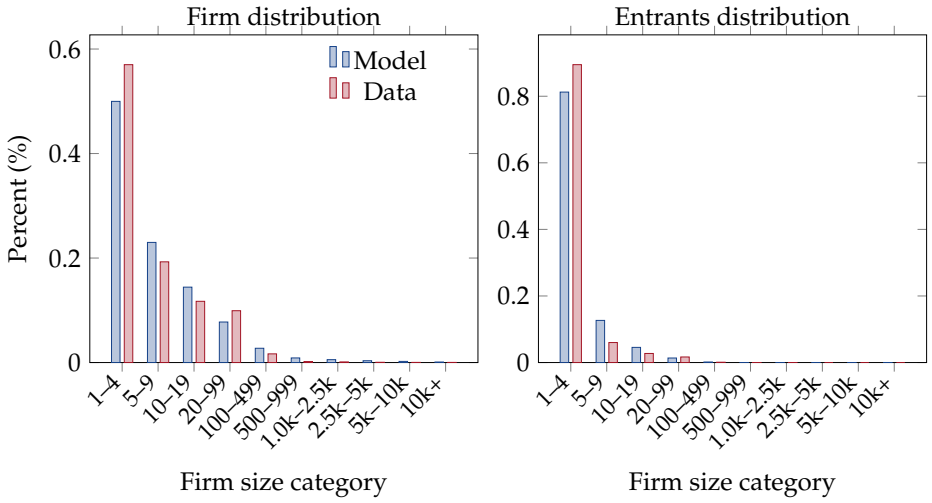


Figure 6: Panel A shows the firm distribution by firm size in the data (red) and the model (blue). Panel B shows the distribution of entering firms by firm size in the data (red) and the model (blue). Empirical distribution reflects the year 2000. Model represent the calibrated model without cyber-attacks. Source: Compustat and authors' analysis.

5 Counterfactual Exercises

Cyber risk has grown over the past two decades and is likely to intensify as firms accumulate sensitive data, become more digitally dependent, and as technologies such as



Figure 7: Distributional moments. Panel A shows the average cybersecurity share of employment by firm size in the data (2024 data) and the model. Panel B displays the annual probability of attack in model and data across the firm size distribution. In panel B, the data refers specifically to "ransomware", model refers to a generic cyber attack.

artificial intelligence enhance the effectiveness of cyber attackers. To quantify how these structural shifts may affect equilibrium outcomes, we study two counterfactual scenarios. The first increases attackers' capacity, A, while the second raises the probability that firms are successfully targeted. Both exercises capture plausible forms of technological progress that favor cyber agents. Together, they show that evaluating cyber risk requires accounting for endogenous firm responses and general equilibrium feedbacks.

For each counterfactual, we compare the steady state of the economy to the benchmark equilibrium and express the resulting change in aggregate output as a percent deviation from the benchmark. We decompose the change in aggregate output into three components. Let Y denote equilibrium aggregate output. Direct losses are measured as

$$L^{\text{direct}} \equiv \ell Z \sum_{\varphi \in \Phi'} \varphi \left(n^a(\varphi)\right)^{\alpha} L(\varphi)$$

which captures the loss in output from successful cyber attacks. Indirect losses are defined as

$$L^{\text{indirect}} \equiv W \sum_{\varphi \in \Phi} x(\varphi) L(\varphi) + \frac{1-Z}{Z} Y,$$

and capture the added costs due to cyber security investment and the lost output due to the cyber externality. The residual is given by

$$L^{\text{residual}} \equiv Y - L^{\text{direct}} - L^{\text{indirect}}.$$

The residual component reflects all endogenous general equilibrium responses, including adjustments in firm entry and in attack incidence.

Panels (A) and (C) report the output decomposition for changes in attack capacity and matching efficiency, respectively, while Panels (B) and (D) display the associated endogenous equilibrium responses underlying the residual component.

Higher attack capacity A. Panels (A) and (B) report the effects of increasing attacker capacity from $A = 1$ to $A = 5$. Aggregate output declines monotonically with attack capacity, falling by about 2.5 percent relative to the benchmark economy at $A = 5$. Panel (A) shows that only a limited share of this decline is attributed to cyber losses. Direct losses from successful attacks account for roughly 0.4 percentage points (about 16 percent of the total decline), while indirect losses associated with the propagation externality contribute an additional 0.07 percentage points. The remaining share of the output contraction—over four-fifths—is driven by endogenous equilibrium responses.

Panel (B) illustrates the mechanisms behind this residual component. As attack capacity increases, the mass of active firms declines sharply, reflecting weaker entry incentives in a riskier cyber environment. At the same time, surviving firms raise cybersecurity investment in response to heightened risk. Despite these defensive efforts, the share of firms that are successfully attacked rises substantially, highlighting the limits of individual firm responses when attack capacity increases. Together, reduced firm entry, higher defensive costs, and increased attack incidence generate large general equilibrium effects that dominate the direct and indirect loss channels.

Higher matching efficiency ψ. Panels (C) and (D) present the analogous decomposition for increases in cyber matching function efficiency ψ. As ψ rises, aggregate output declines sharply and is almost 7 percent lower when $\psi = 0.05$ compared to the benchmark of $\psi = 0.01$. Panel (C) shows that this decline is driven primarily by indirect losses associated with the propagation externality, which account for about 4.87 percentage points of the output reduction ($\sim$ 70% of the total). Direct losses from cyber attacks contribute an additional 0.245 percentage points ($\sim$ 3.6% of the total), while the remaining 1.78 percentage points ($\sim$ 26% of the total) are attributable to endogenous general equilibrium responses captured by the residual component.

Panel (D) highlights the underlying equilibrium adjustments. An increase in matching efficiency ψ has qualitatively similar effects to an increase in attack capacity in the previous counterfactual, but the quantitative implications are markedly different. In particular, the response of the share of firms under attack is substantially stronger when ψ

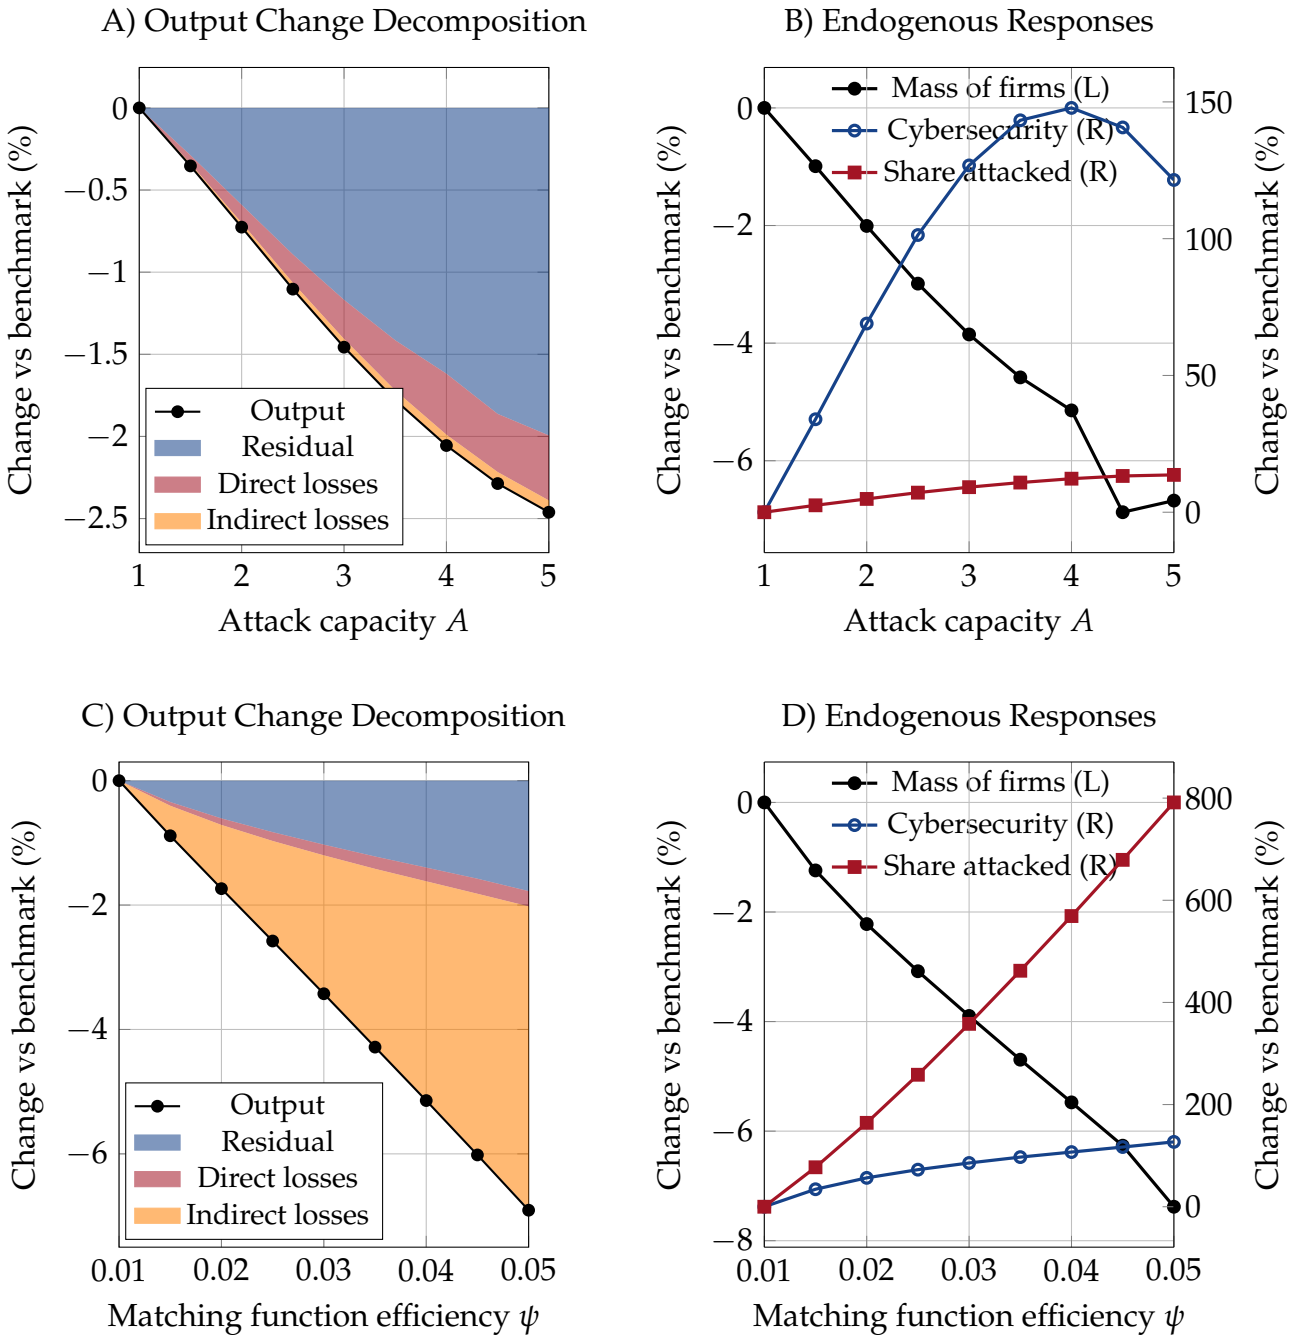


Figure 8: Output decomposition and endogenous responses to cyber risk and labor market efficiency. Panels (A) and (B) report the effects of higher attack capacity A. Panel (A) decomposes the percent change in aggregate output relative to the benchmark economy into direct losses from cyber attacks, indirect losses associated with the propagation externality, and a residual. Panel (B) shows the corresponding endogenous equilibrium responses: firm entry, cybersecurity investment, and the share of firms under attack. Panels (C) and (D) present the same objects for changes in matching function efficiency ψ.

rises. As Panel (D) shows, the fraction of firms experiencing a successful attack increases by almost a factor of eight as ψ moves from 0.01 to 0.05. By contrast, the same object rises by only about 13 percent when attack capacity increases from $A = 1$ to $A = 5$.

This stark difference explains why indirect losses play a much larger role in the ψ counterfactual. The sharp increase in attack incidence strongly depresses aggregate productivity through the cyber propagation externality, generating a large negative contribution of indirect losses to output. By comparison, changes in cybersecurity investment and in the mass of active firms are of similar magnitude across the two counterfactuals. Next, we look into why this channel is so much stronger when ψ rises than when A rises.

A versus ψ. The stronger role of indirect losses under increases in matching efficiency ψ reflects how attackers reallocate their effort across firm sizes. When attack capacity A rises, cyber attackers optimally concentrate attacks on large, high-output firms, which are more attractive targets. These large firms have a harder time keeping up their cyber investments when cyber attackers' A increases. As the panels in figure 9 show, attack success probabilities increase primarily at the top of the firm size distribution. Because large firms represent only a small mass of firms, this shift in target by attackers limits the aggregate propagation externality, even though these firms are economically important.

By contrast, higher matching efficiency ψ raises the effectiveness of attacks across a much broader range of firm sizes. Attack success probabilities increase substantially not only for large firms but also for small and medium-sized firms, which are far more numerous. This expansion sharply amplifies the propagation externality and explains why indirect losses dominate the output response in the ψ counterfactual, despite similar aggregate changes in cybersecurity investment and firm entry.

These shifts across the firm size distribution translate into changes in market structure. Figure 10 reports the share of aggregate output produced by each firm size bin. Panel A varies attacker capacity A, while panel B varies matching efficiency ψ. The two adverse cyber scenarios have opposite implications for output concentration. As A rises, large firms, being more exposed to successful attacks, lose output share, leading to a decline in concentration at the top of the size distribution. In contrast, higher ψ benefits, in relative terms, larger firms, which are better able to defend against attacks. As a result, output shifts away from small and mid-sized firms toward the largest firms. These patterns highlight how different dimensions of cyber risk can push market structure in opposite directions, an important consideration in economies where production is increasingly concentrated and potentially more vulnerable to systemic cyber events.

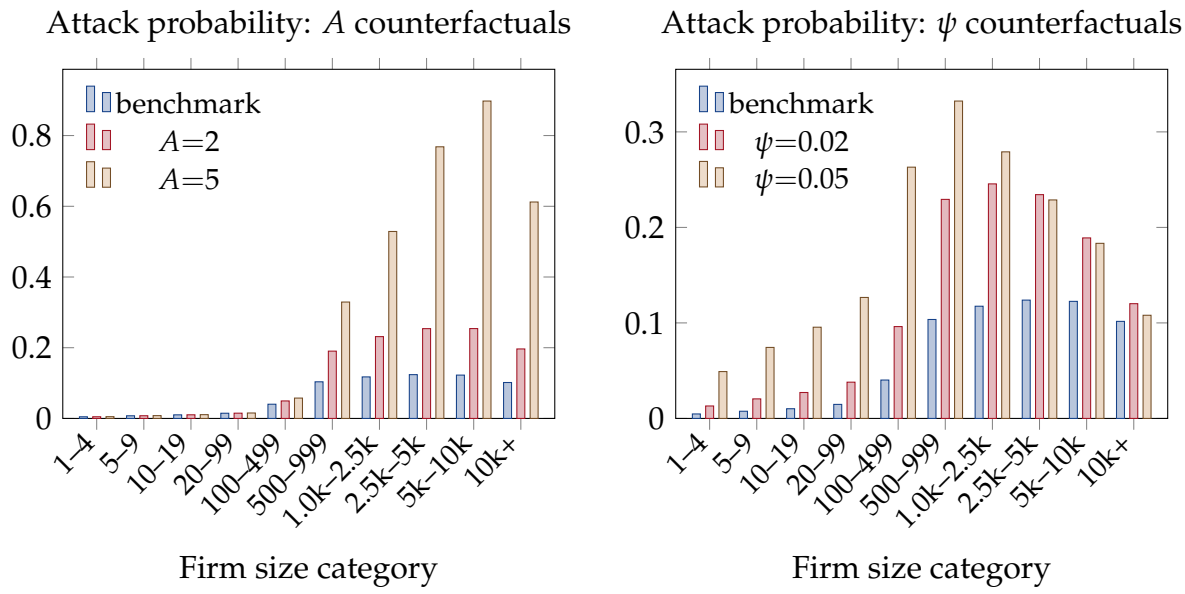


Figure 9: Attack success probability by firm size under changes in attacker capacity A (left) and matching efficiency ψ (right). Bars show the benchmark and two counterfactual values for each experiment.

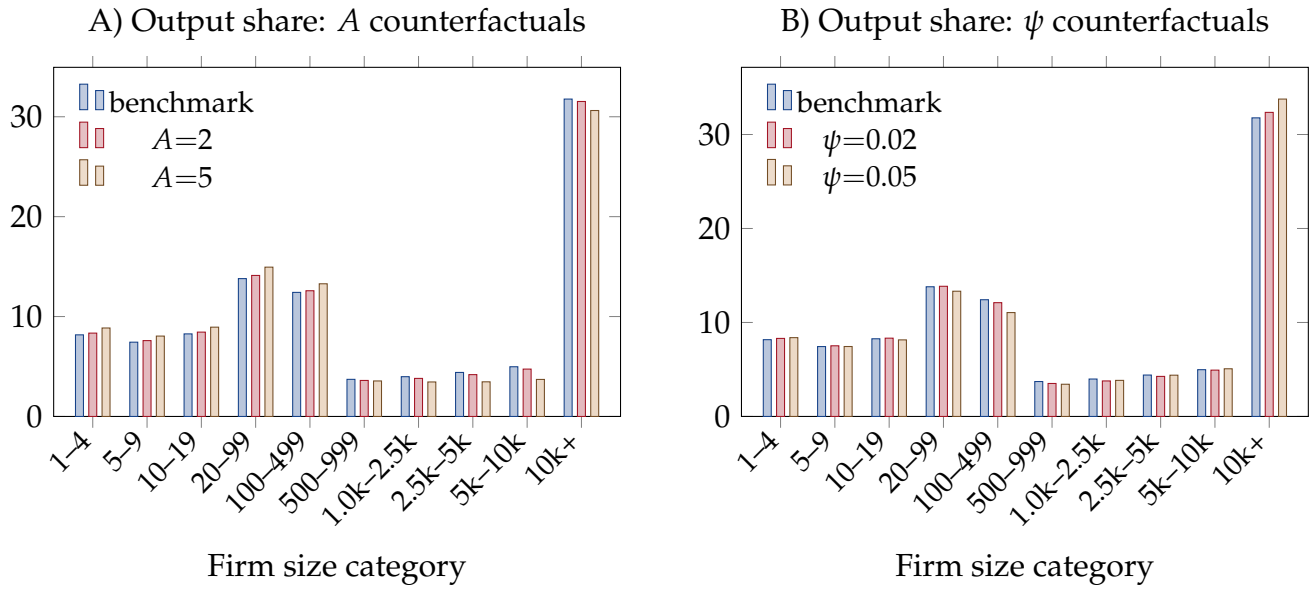


Figure 10: Share of output produced across the firm size distribution under changes in attacker capacity A (panel A) and matching efficiency ψ (panel B). Bars show the benchmark and two counterfactual values for each experiment.

6 Subsidies versus Bailouts

In this section we examine the role that government policy can play in countering the rising tide of cyberattacks. We focus on two interventions: a subsidy to cyber security investment and bailouts that compensate firms when attacks occur. While such measures may currently appear remote, large-scale and systemic cyber incidents may increasingly force governments to intervene, as illustrated by the UK government's £1.5 billion loan guarantee to Jaguar following a major cyberattack (UK Government, Department for Business and Trade and UK Export Finance and HM Treasury 2025).

The two instruments operate through different incentive channels. Subsidies act on firms' ex-ante cyber security choices, while bailouts provide ex-post insurance against realized losses. As a result, the two policies differ sharply in how they affect firm entry, cybersecurity investment incentives, and the strength of the aggregate cyber externality. The remainder of the section makes these differences explicit both analytically and quantitatively.

Subsidy. We first consider a government that subsidizes firms' cyber security expenditures. The subsidy affects firm behavior only through the marginal cost of cyber security investment x. Equations (2) and (5) become

$$\begin{aligned} V^s(\varphi) &= \max_{n,x\geq 0} \ \varphi n^\alpha - wn - wx(1-\tau) - f \\ &\quad +\beta(1-\sigma)\mathbb{E}_{\varphi'|\varphi}\left[q(\varphi')V^m(\varphi',x) + (1-q(\varphi'))V^s(\varphi')\right], \end{aligned} \tag{20}$$

$$\begin{aligned} V^a(\varphi) = \max\Big\{0, \quad &\max_n (1-\ell)\varphi n^\alpha - wn - f \\ &+\beta(1-\sigma)\mathbb{E}_{\varphi'|\varphi}\left[(1-\delta)V^a(\varphi') + \delta V^s(\varphi')\right]\Big\}. \end{aligned} \tag{21}$$

The government taxes the representative household with a lump sum tax T and subsidizes a fraction τ of each firm's cyber security wage bill wx. Its budget constraint equates total tax receipts and subsidy expenditures so that the tax rate T adjusts to cover the government's cost of providing the subsidy.

$$T = \underbrace{\tau w \int_j x(j) D^s(j)\, dj}_{\text{subsidy outlays}}. \tag{22}$$

Bailout. Next we analyze a bailout policy in which the government reimburses a share ι of losses suffered by attacked firms. In contrast to the subsidy scenario, bailouts leave the cost of cyber security unchanged but partially insure attacked firms by reducing the effective revenue loss from ℓ to $\ell(1-\iota)$. The Bellman equations become

$$\begin{aligned} V^s(\varphi) &= \max_{n,x\geq 0} \varphi n^\alpha - wn - wx - f \\ &\quad +\beta(1-\sigma)\mathbb{E}_{\varphi'|\varphi}\left[q(\varphi')V^m(\varphi',x) + (1-q(\varphi'))V^s(\varphi')\right], \end{aligned} \tag{23}$$

$$\begin{aligned} V^a(\varphi) = \max\Big\{0, \quad & \max_n \left(1-\ell\cdot(1-\iota)\right)\varphi n^\alpha - wn - f \\ & +\beta(1-\sigma)\mathbb{E}_{\varphi'|\varphi}[(1-\delta)V^a(\varphi') + \delta V^s(\varphi')]\Big\}. \end{aligned} \tag{24}$$

The government finances these transfers with a lump sum tax T on the representative household.

$$T = \underbrace{\ell\cdot\iota\int_j \varphi n^\alpha(j)D^a(j)\,dj}_{\text{bailout payments}}. \tag{25}$$

Comparing policies. We now quantify how these distinct incentive channels play out in general equilibrium. A partial equilibrium perspective would suggest that the two policies generate opposite investment incentives: bailouts reduce incentives to invest in cybersecurity, while subsidies increase them. While this intuition is broadly correct, it is only partly reflected once general equilibrium adjustments are taken into account.

We highlight three intertwined general-equilibrium channels: firm entry, the endogenous tax rate required to finance the policies, and the strength of the cyber propagation externality. Figure 11 maps these channels by showing the effect that different subsidy and bailout rates have on aggregate economic indicators. In each of the six panels, the x-axis shows the policy rate – subsidy rate τ or bailout rate ι – the blue line shows the indicator under the different subsidy rate scenarios, and the red line shows the indicator under the different bailout rate scenarios. In each of these, the 0 rate case corresponds to the benchmark economy.

The two policies have starkly different consequences once their general equilibrium mechanisms are taken into account. The subsidy directly reduces the marginal cost of cyber security, inducing entering firms to invest more in cyber security (panel B). Higher investment improves resilience to cyberattacks, reducing the share of firms under attack (panel C) and weakening the cyber externality, thereby raising economy-wide produc-

tivity (panel D). For moderate subsidy rates, the reduction in attack risk and the associated weakening of the cyber externality dominate the distortionary effects of taxation, encouraging entry (panel A) despite higher taxes. However, as the subsidy becomes too generous ($\tau > 0.7$), the tax burden becomes increasingly onerous. While aggregate productivity continues to rise, firm entry declines, and the resulting contraction in the mass of firms ultimately reduces aggregate output (panel E) and consumption (panel F). The model finds that, with respect to the current cyber environment, a subsidy of 0.7% is output and welfare-maximizing.

The effects of bailouts operate through a different channel. Even at relatively low bailout rates ι, bailouts weaken incentives to invest in cybersecurity through moral hazard (panel B). Lower investment raises attack incidence (panel C), and strengthens the propagation externality (panel D). The lower productivity coupled with the increased tax burden discourages firm entry (panel A). As a result, aggregate output and consumption fall monotonically as the bailout rate increases (panels E and F).

Taken together, these results highlight that policies targeting ex-ante prevention and ex-post insurance differ sharply once general equilibrium forces are taken into account.

State-dependent policies. The policy evaluation above suggests that subsidies can be effective at combating cyber attacks and increasing output. However, as cyber risk grows, the policy prescription may change. In figure 12 we consider the scenarios from the previous section. Panel A shows the tax to output ratio, and panel C aggregate output, as the subsidy rate τ varies for both the benchmark calibration (solid blue), the economy in which attack capacity increases ($A = 5$), and the economy in which attacker matching efficiency increases ($\psi = 5$). In these two counterfactuals, the same policy has very different implications.

Unlike the benchmark economy, in which output increases up to a subsidy rate $\tau = 0.7$, when cyber agents have higher attack capacity (red dotted line), output is flat at first, and then quickly decreases for mid and high values of τ. This is coupled with a larger incentive to invest in cybersecurity by firms, which drives up the total tax bill much more than in the benchmark economy with $A = 1$. In contrast, with $\psi = 0.05$ the economy has output increasing for an even larger subsidy rate ($\tau = 0.8$) than the benchmark despite there too the tax to output ratio increasing considerably relative to the benchmark.

The reason for the stark difference in the aggregate output response as τ changes is that, when cyber risk intensifies through higher attack capacity, subsidies disproportionately benefit large firms who are most targeted when attack capacity is high and therefore become fiscally costly, whereas when risk rises through broader attack effectiveness,

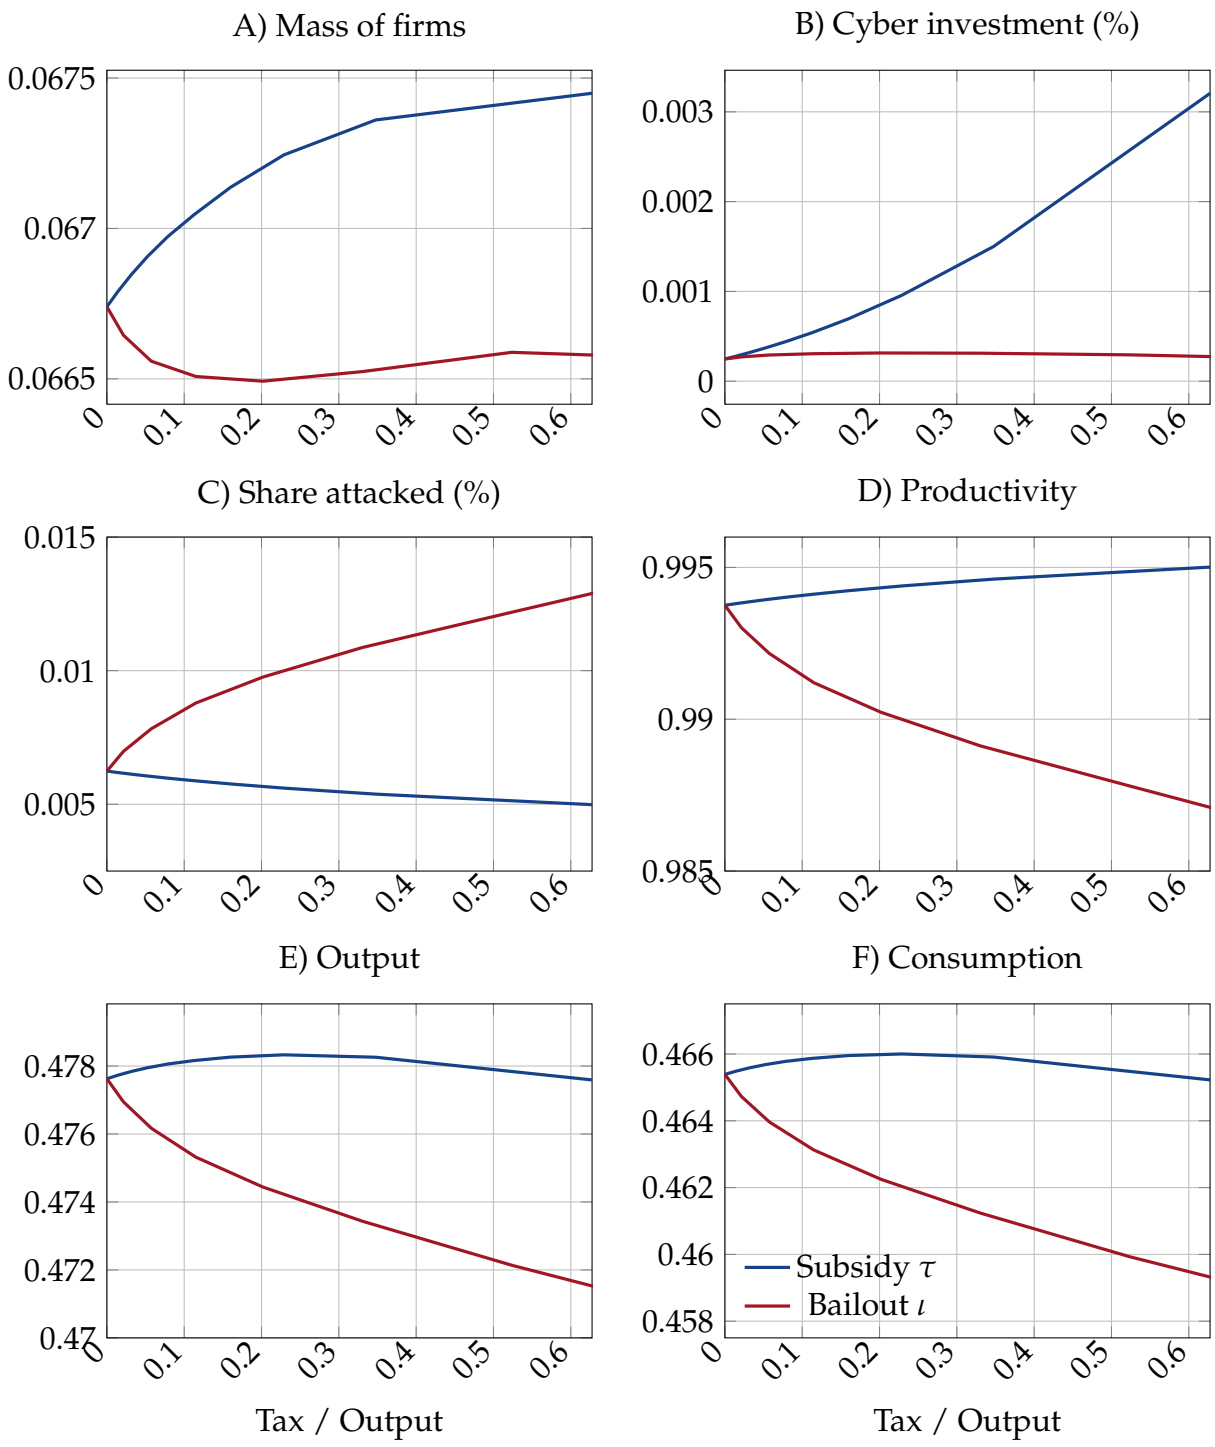


Figure 11: Effect of cyber security subsidies (τ) and bailouts (ι) rates on (A) the mass of active firms, (B) the cyber investment share, (C) the share of firms under attack, (D) aggregate productivity defined as 1 minus the cyber externality term, (E) aggregate output, and (F) total consumption.

subsidies improve resilience across the firm distribution and remain effective. As figure 12 shows, the investment increase as the subsidy policy becomes more generous is overwhelmingly driven by the largest firms when $A = 5$. In contrast, in the $\psi = 0.05$ economy, the changes are much more homogeneous across the firm size distribution with all firms increasing their investment in cybersecurity. Therefore, unlike the $\psi = 0.05$ scenario, when $A = 5$, the subsidy acts as a net transfer from small to large firms when $A = 5$. This extra fiscal burden without a commensurate benefit, dissuades firms from entering when $A = 5$ (figure 12) and lowers the total mass of firms in the economy, depressing output as a consequences.

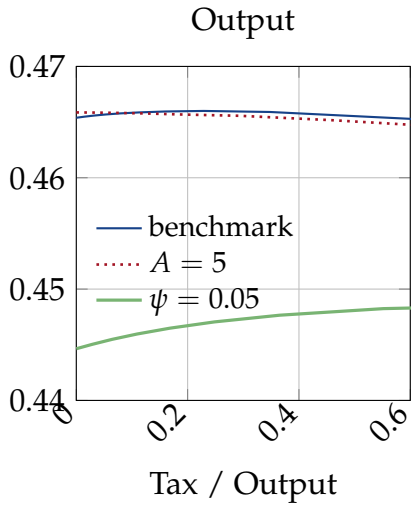


Figure 12: Equilibrium responses to changes in the policy rate. Panel (A) reports tax revenue as a share of output; Panel (C) reports aggregate output. Solid blue lines correspond to the benchmark economy, red dotted lines correspond to the economy with high attack capacity ($A = 5$), and green dashed lines correspond to the economy with higher matching efficiency ($\psi = 0.05$).

Takeaway. In sum, this quantitative analysis highlights a sharp distinction between the two policy instruments. Subsidies operate by lowering the marginal cost of preventive investment, strengthening resilience, and attenuating the cyber externality, which supports entry and raises output over a wide range of policy intensities. Bailouts, by contrast, primarily provide insurance ex-post, but weaken investment incentives and amplify the cyber externality, causing aggregate productivity and output to decline as policy generosity increases.

From an output perspective, the model therefore favors subsidies over bailouts as a tool to counter rising cyber risk: subsidies directly strengthen preventive investment and

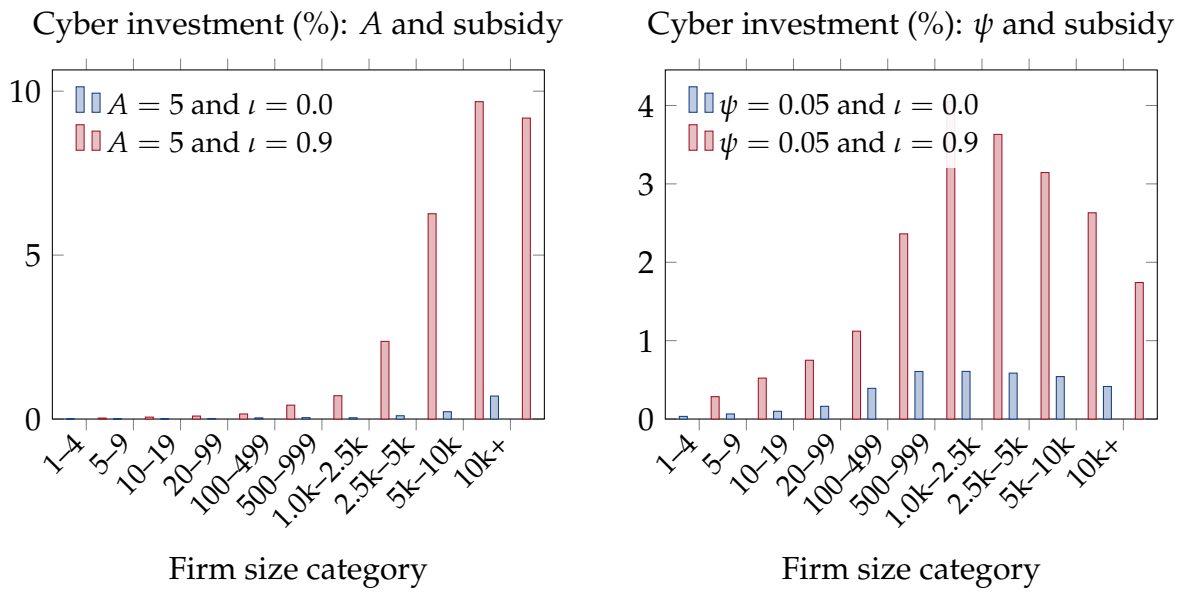


Figure 13: Cyber investment as a share of labor by firm size under adverse cyber scenarios for different subsidy policies. Panel A shows cyber investment for high attack capacity ($A = 5$) for no (blue) and high (red) subsidy rates. Panel A shows cyber investment for high attack matching efficiency ($\psi = 0.05$) for no (blue) and high (red) subsidy rates.

attenuate the cyber externality, while bailouts primarily insure losses ex post and exacerbate moral hazard. However, at what level the subsidy should be set or even whether it is beneficial, will depend on the state of the economy and may very well change as cyber risk evolves.

7 Conclusion

Cyber risk has emerged as a first-order economic concern, yet its macroeconomic implications and appropriate policy responses remain poorly understood. Our first and main contribution in this paper is to construct a dynamic general equilibrium model in which these strategic cybersecurity investments interact with strategic choices by attackers regarding which firms to target and what attack intensity to deploy against them.

Second, to interpret this empirical pattern and study its aggregate consequences, we develop a dynamic general equilibrium model in which heterogeneous firms invest in cybersecurity, attackers optimally target firms, and successful cyberattacks propagate through the economy by reducing aggregate productivity. The model is disciplined to match the observed size–investment gradient alongside key moments on firm dynamics, attack incidence, and loss amplification, allowing us to connect micro-level investment

behavior to macroeconomic outcomes.

The analysis highlights the central role of general equilibrium forces in shaping both the aggregate impact of cyber risk and the effectiveness of policy interventions. Cyber-attacks generate losses not only through direct disruptions at affected firms, but also through a propagation externality that depresses economy-wide productivity and feeds back into firms' entry and investment decisions. As a result, the macroeconomic cost of cyber risk substantially exceeds the direct losses typically measured in firm-level data.

We use the model to evaluate two policy instruments that operate through distinct channels: subsidies to cybersecurity investment and bailouts to firms that suffer cyber losses. While partial-equilibrium intuition suggests opposing incentive effects, the general equilibrium comparison reveals a sharper contrast. Subsidies directly strengthen preventive investment, reduce attack incidence, and mitigate the propagation externality, raising aggregate productivity and, at moderate levels, encouraging firm entry despite higher taxes. Bailouts, by contrast, primarily provide ex post insurance, weakening incentives to invest through moral hazard, amplifying cyber risk, and discouraging entry. As a result, aggregate output declines as bailouts become more generous, even when fiscal costs remain limited.

The effectiveness of preventive policy is state-dependent. When cyber risk intensifies through higher attack capacity, subsidies become increasingly costly and yield diminishing aggregate gains, particularly for smaller firms. When risk rises through broader attack effectiveness, however, subsidies remain an effective tool for containing the externality and supporting aggregate activity. These results underscore that policy design must account not only for fiscal considerations, but also for the heterogeneity in firms' exposure to cyber risk and their capacity to invest in resilience, heterogeneity that is clearly reflected in the data.

Overall, the paper shows that the empirical size gradient in cybersecurity investment is a key organizing feature of cyber risk in the economy, with important implications for both aggregate outcomes and optimal policy design. Policies that target ex ante resilience are more effective at containing the macroeconomic consequences of cyber risk than those that insure losses ex post. More broadly, cyber risk should be viewed not merely as an operational or firm-level concern, but as a macroeconomic externality with meaningful implications for productivity, entry, and aggregate output.

As cyber threats continue to grow in scale and sophistication, understanding their aggregate consequences and grounding policy analysis in empirically observed patterns of firm behavior will remain an increasingly important challenge for researchers and policymakers alike.

References

Acemoglu, Daron, Azarakhsh Malekian, and Asuman Ozdaglar. 2016. "Network security and contagion." *Journal of Economic Theory*, 166: 536–585.

Acemoglu, Daron, Vasco M Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi. 2012. "The network origins of aggregate fluctuations." *Econometrica*, 80(5): 1977–2016.

Akyildirim, Erdinc, Thomas Conlon, Shaen Corbet, and Yang (Greg) Hou. 2024. "HACKED: Understanding the stock market response to cyberattacks." *Journal of International Financial Markets, Institutions and Money*, 97(C): None.

Al Amosh, Hamzeh, and Saleh F. A. Khatib. 2025. "Cybersecurity Transparency and Firm Success: Insights From the Australian Landscape." *Australian Economic Papers*, 64(2): 189–204.

Anand, Kartik, Chanelle Duley, and Prasanna Gai. 2022. "Cybersecurity and financial stability."

Anderson, Ross, Chris Barton, Rainer Böhme, Richard Clayton, Michel JG Van Eeten, Michael Levi, Tyler Moore, and Stefan Savage. 2013. "Measuring the cost of cybercrime." *The economics of information security and privacy*, 265–300.

Baksy, Aniket, Daniele Caratelli, and Luke M. Olson. 2025. "Cyberattacks and Firm Size: The Vulnerability of Mid-Size Firms." *The OFR Blog, Office of Financial Research*.

Barati, Mehdi, and Benjamin Yankson. 2022. "Predicting the Occurrence of a Data Breach." *International Journal of Information Management Data Insights*, 2(2): 100128.

Begenau, Juliane, Maryam Farboodi, and Laura Veldkamp. 2018. "Big data in finance and the growth of large firms." *Journal of Monetary Economics*, 97: 71–87.

Board of Governors of the Federal Reserve System. 2025. "Cybersecurity and Financial System Resilience Report." Report to Congress.

Carfora, Maria Francesca, and Albina Orlando. 2022. "Some Remarks on Malicious and Negligent Data Breach Distribution Estimates." *Computation*, 10(12).

Cobos, Estefania Vergara, and Selcen Cakir. 2024. "A Review of the Economic Costs of Cyber Incidents." World Bank Working Paper.

Crosignani, Matteo, Marco Macchiavelli, and André F. Silva. 2023. "Pirates without Borders: The Propagation of Cyberattacks through Firms' Supply Chains." *Journal of Financial Economics*, 147(2): 432–448.

Cybersecurity and Infrastructure Security Agency. 2020. "Cost of a Cyber Incident: Systematic Review and Cross-Validation." Cybersecurity and Infrastructure Security Agency.

Cybersecurity and Infrastructure Security Agency. 2025. "StopRansomware Guide." CISA. Accessed: 2026-03-17.

Diamond, Peter A. 1982. "Aggregate Demand Management in Search Equilibrium." *Journal of Political Economy*, 90(5): 881–894.

Duffie, Darrell, and Joshua Younger. 2019. *Cyber runs.* Brookings.

Ebel, Austin, and Debasis Mitra. 2024. "Economics and optimal investment policies of Attackers and Defenders in cybersecurity." *Journal of Cybersecurity*, 10(1): tyae019.

Edwards, Benjamin, Steven Hofmeyr, and Stephanie Forrest. 2016. "Hype and heavy tails: A closer look at data breaches." *Journal of Cybersecurity*, 2(1): 3–14.

Eisenbach, Thomas M, Anna Kovner, and Michael Junho Lee. 2022. "Cyber risk and the US financial system: A pre-mortem analysis." *Journal of Financial Economics*, 145(3): 802–826.

Erkan-Barlow, Asligul, Thanh Ngo, Rajni Goel, and Denise W. Streeter. 2023. "An in-depth analysis of the impact of cyberattacks on the profitability of commercial banks in the United States." *Journal of Global Business Insights*, 8: 120–135.

Erol, Selman, and Michael Junho Lee. 2024. "Financial System Architecture and Technological Vulnerability." Federal Reserve Bank of New York 1122. Revised October 2024.

Farboodi, Maryam, Roxana Mihet, Thomas Philippon, and Laura Veldkamp. 2019. "Big Data and Firm Dynamics." *AEA Papers and Proceedings*, 109: 38–42.

Federal Bureau of Investigation and CISA. 2023. "StopRansomware: Dual Ransomware Attacks." FBI Internet Crime Complaint Center (IC3). Accessed: 2026-03-17.

Hopenhayn, Hugo A. 1992. "Entry, exit, and firm dynamics in long run equilibrium." *Econometrica*, 60(5): 1127–1150.

Huang, Ziyuan, Parinaz Naghizadeh, and Mingyan Liu. 2024. "Interdependent security games in the Stackelberg style: how first-mover advantage impacts free riding and security (under-)investment." *Journal of Cybersecurity*, 10(1): tyae009.

ICAEW Insights. 2023. "The 21st-century evolution of cyber security."

International Monetary Fund. 2023. *Global Financial Stability Report, October 2023: Cyber Risk: A Growing Concern for Macrofinancial Stability.* Washington, DC:International Monetary Fund.

Kamiya, Shinichi, Jun-Koo Kang, Jungmin Kim, Andreas Milidonis, and René M. Stulz. 2021. "Risk management, firm reputation, and the impact of successful cyberattacks on target firms." *Journal of Financial Economics*, 139(3): 719–749.

Koo, Helga, Remco van der Molen, Robert Vermeulen, Ralph Verhoeks, and Alessandro Pollastri. 2022. "A macroprudential perspective on cyber risk." DNB.

Korzhyk, Dmytro, Zhengyu Yin, Christopher Kiekintveld, Vincent Conitzer, and Milind Tambe. 2011. "Stackelberg vs. Nash in Security Games: An Extended Investigation of Interchangeability, Equivalence, and Uniqueness." *Journal of Artificial Intelligence Research*, 297–327.

Kotidis, Antonis, and Stacey L. Schreft. 2025. "The Propagation of Cyberattacks through the Financial System: Evidence from an Actual Event." *The Journal of Finance*, 80(6): 3313–3358.

Lee, Yoonsoo, and Toshihiko Mukoyama. 2015. "Entry and Exit of Manufacturing Plants over the Business Cycle." *European Economic Review*, 77: 20–27.

Microsoft. 2023. "Ransom:Win32/BlackCat.B." Microsoft Security Intelligence. Accessed: 2026-03-17.

Moody's Ratings. 2023. "Cyber budgets increase, executive overview improves, but challenges lurk under the surface." Moody's.

Moore, Tyler, Richard Clayton, and Ross Anderson. 2009. "The Economics of Online Crime." *Journal of Economic Perspectives*, 23(3): 3–20.

Mortensen, Dale T., and Christopher A. Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies*, 61(3): 397–415.

Muktadir-Al-Mukit, Dewan, and Md Hakim Ali. 2025. "The Dynamics of Stock Market Responses Following Cyber-Attacks News: Evidence from Event Study." *Information Systems Frontiers*.

National Cyber Security Centre. 2023. "Mitigating Malware and Ransomware Attacks." UK National Cyber Security Centre. Accessed: 2026-03-17.

National Institute of Standards and Technology. 2023. "Phish Scale: The Impact of Increasing Message Volume on Phishing Susceptibility." National Institute of Standards and Technology (NIST) NIST Special Publication 800-203 (Pre-publication draft). Accessed: 2026-03-17.

OECD. 2001. "The Internet and Business Performance." OECD OECD Digital Economy Papers 57, OECD Publishing.

Office of Financial Research. 2022. "Annual Report to Congress." U.S. Department of the Treasury. Office of Financial Research Annual Report.

Ottonello, Pablo, and Thomas Winberry. 2020. "Financial Heterogeneity and the Investment Channel of Monetary Policy." *Econometrica*, 88(6): 2473–2502.

Ponemon Institute. 2016. "SMBs are vulnerable to cyber attacks." *https://www.ponemon.org/research/ponemon-library/security/smbs-are-vulnerable-to-cyber-attacks.html*, Report Release.

Ponemon Institute and IBM. 2025. "2025 Cost of a Data Breach Report." Ponemon Institute.

Ramírez, Carlos A. 2025. "On Equilibrium Cyber Risk." *Economics Letters*, 251.

Romanosky, Sasha. 2016. "Examining the costs and causes of cyber incidents." *Journal of Cybersecurity*, 2(2): 121–135.

Romanosky, Sasha, Rahul Telang, and Alessandro Acquisti. 2011. "Do data breach disclosure laws reduce identity theft?" *Journal of Policy Analysis and Management*, 30(2): 256–286.

Romer, Paul M. 1986. "Increasing Returns and Long-Run Growth." *Journal of Political Economy*, 94(5): 1002–1037.

Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy*, 98(5, Part 2): S71–S102.

Rosati, Pierre. 2021. "Linking data breaches and financial markets: An augmented Privacy Rights Clearinghouse dataset." *Data in Brief*, 34: 106719.

Seh, Adil Hussain, Mohammad Zarour, Mamdouh Alenezi, Amal Krishna Sarkar, Alka Agrawal, Rajeev Kumar, and Raees Ahmad Khan. 2020. "Healthcare Data Breaches: Insights and Implications." *Healthcare*, 8(2): 133.

StatCan. 2000. "E-commerce and business use of the Internet." *The Daily*.

UK Government, Department for Business and Trade and UK Export Finance and HM Treasury. 2025. "Government backs Jaguar Land Rover with £1.5 billion loan guarantee." *https://www.gov.uk/government/news/government-backs-jaguar-land-rover-with-15-billion-loan-guarantee*, Press release.

Verizon. 2024. "2024 Data Breach Investigations Report." Verizon Business. Accessed: 2026-03-17.