On Equilibrium Cyber Risk*

Carlos A. Ramírez

March 11, 2025

Abstract

I develop a simple model to study how the interplay between institutions' and hackers' incentives can alter cyber risk within an equilibrium context. By formalizing the strategic interaction between institutions and hackers, the model characterizes how changes in heterogeneity across institutions, cybersecurity technologies, and hacker competition can lead to material shifts in cyber risk.

JEL Codes: D39, D22, M15.
Keywords: equilibrium cyber risk, cybersecurity, cyberattacks, hacks.

Highlights:

- The model establishes an equilibrium link between the distribution of cyberattacks and the size distribution of institutions.
- When hackers' rewards are increasing in the size of their targets, my results are consistent with the observed concentration of cyberattacks among larger institutions.
- As cybersecurity technologies improve or hacker competition intensifies, institutions invest less on cybersecurity while hackers face weaker incentives to attack and become less selective in their targeting.

*Federal Reserve Board. These are my views and do not reflect those of the Federal Reserve Board or other members of its staff. Email: carlos.ramirez@frb.gov. There are no competing interests to declare.

1 Introduction

In recent years, cyber risk has become a pressing concern for investors, businesses, regulators, and academics alike.[1] Yet much of the existing literature overlooks equilibrium considerations in its assessments. This oversight likely stems from the challenges of observing the actions of both hackers and institutions, coupled with the lack of consensus on how their motivations interact in equilibrium.[2] To fill this gap, I propose a model that links the motivations of hackers and institutions within an equilibrium framework, offering a clearer understanding of how basic economic factors can reshape cyber risk.

My model incorporates two key features. First, institutions and hackers would like to outguess one another before making decisions. Second, hackers' rewards are influenced by both the size of their targets and the effectiveness of cybersecurity technologies. With this model in hand, I establish a mapping between the distribution of cyberattacks and the size distribution of institutions, and explore how changes in heterogeneity across institutions, cybersecurity technologies, and hacker competition can alter cyber risk.

Within my model, the link between hackers' rewards and the size of their targets plays a key role in shaping equilibrium behavior. I show that when hackers' rewards increase (decrease) with the size of their targets, larger (smaller) institutions invest more in cybersecurity as they are targeted more frequently. In this case, larger (smaller) institutions also become less (more) attractive targets as cybersecurity technologies improve. Additionally, I show that institutions invest less on cybersecurity while hackers face weaker incentives to attack and become less selective in their targeting when cybersecurity technologies improve or hacker competition intensifies. Although both heterogeneity across institutions and cybersecurity technologies influence equilibrium behavior, their impact on cyber risk diminishes as hacker competition intensifies.

My findings contribute to the rapidly expanding literature on the drivers and consequences of cyber risk. Within this literature, Ahnert et al. (2024) and Anand et al. (2024) are the most closely related to my work. Ahnert et al. (2024) study how regulators can correct inefficiencies engendered by the lack of observability in firms'

[1] World Economic Forum Report and Federal Reserve Board Report underscore concerns of businesses and regulators about the consequences of cyberattacks. Duffie and Younger (2019) and Kashyap and Wetherilt (2019) highlight the potential systemwide implications of cyberattacks and the policy challenges that stem from them.

[2] See Ablon (2018) and Chng et al. (2022) for a description of hackers' motivations.

decisions. Anand et al. (2024) explore how cyberattacks influence banks' decision-making and their likelihood of runs. While our work shares an interest in the interaction between institutions and hackers, my paper is the first to explore how the interplay between cybersecurity technologies, heterogeneity across institutions, and hacker competition can reshape cyber risk in equilibrium.[3]

2 The Hacking Game

Though stylized, the baseline model conveys the main intuition for how institutions' and hackers' motivations jointly shape cyber risk in equilibrium. The Online Appendix shows that this model can be extended to incorporate: heterogeneity in bargaining power, alternative costs structures, and uncertainty about the precise impact of cyberattacks.

Consider an economy consisting of a unit continuum of institutions (firms, for short), each varying in size, alongside a single hacker motivated purely by financial gain. I focus on games wherein firms and the hacker choose actions simultaneously, and their payoffs (which are common knowledge) depend on a combination of their selected actions.

Consider a firm of size $s \in (0, 1)$ and the hacker. The proposed game—referred to as the Hacking Game—shares a distinguishing feature with the game of Matching Pennies: both players would like to outguess one another before selecting their actions.[4] The firm would like to anticipate whether the hacker will attack before investing in cybersecurity, while the hacker would like to know how much the firm invests in cybersecurity before attacking.

Because the solution of this game involves uncertainty about what players will do, let q_s denote the probability that a firm of size s chooses to defend itself and p_s denote the probability that the hacker chooses to target such a firm. Define $v(s)$ as the value of

[3]The Online Appendix provides a more detailed discussion of the differences between my model and Ahnert et al. (2024) and Anand et al. (2024). By theoretically studying the interaction between institutions and hackers, my paper also complements (1) the literature that explores the connection between firms' characteristics and the likelihood of cyberattacks and (2) the literature that studies the impact of cyberattacks. Aldasoro et al. (2020) and Chang et al. (2024) document relationships between the size of institutions and the likelihood of cyberattacks. Jamilov et al. (2021), Kamiya et al. (2021), Florackis et al. (2023), and Jiang et al. (2024) underscore the impact of cyber risk on stock returns, while Curti et al. (2023) show that hacks can increase the financing costs of state and local governments. Eisenbach et al. (2022), Kotidis and Schreft (2022), Crosignani et al. (2023) emphasize how hacks can affect different industries via propagation along supply chains, payment systems, or technology providers.

[4]See (Gibbons, 1992, chap 1.3).

that firm vulnerable to cyberattacks—where $v(\cdot)$ is an arbitrary continuous function of s satisfying $0 \leq v(s) < 1$, $\forall s$. Let $\alpha \in (0,1)$. For a given tuple of choice variables (q_s, p_s), Table 1 reports the payoffs of both players.[5]

HACKER | FIRM OF SIZE s	`Defend` with prob. q_s	`Do not defend` with prob. $(1-q_s)$
`Attack` with prob. p_s	Hacker: $\frac{(1-\alpha)v(s)}{2} - p_s$ Firm: $\alpha v(s) + \frac{(1-\alpha)v(s)}{2} - q_s$	Hacker: $(1-\alpha)v(s) - p_s$ Firm: $\alpha v(s)$
`Do not attack` with prob. $(1-p_s)$	Hacker : 0 Firm: $v(s) - q_s$	Hacker : 0 Firm: $v(s)$

Table 1: Matrix of payoffs

Explanation of payoffs.—To appreciate how motivations manifest in the Hacking Game, it is useful to analyze Table 1. When the hacker attacks and the firm does not defend itself, the firm's payoff is $\alpha v(s)$ while the hacker's payoff is $(1-\alpha)v(s) - p_s$. That is, the firm gets a fraction α of $v(s)$ while the hacker gets the remaining fraction—net of the costs of implementing the hack, p_s. The hacker's rewards, $(1-\alpha)v(s)$, serve as a metaphor for the expected value that is lost to cyberattacks. Thus, parameter α can be thought of as firms' recovery rate in the face of hacks. Intuitively, α reflects the effectiveness of cybersecurity technologies relative to cyberattack technologies: more effective cybersecurity technologies are associated with a higher α. As $\alpha \to 1$ hacks are expected to be harmless, while the opposite happens as $\alpha \to 0$.

When the hacker attacks and firm defends itself, the firm's payoff is $\alpha v(s) + \frac{(1-\alpha)v(s)}{2} - q_s$. The first term, $\alpha v(s)$, represents the payoff of a firm that chooses not to defend itself. The second term, $\frac{(1-\alpha)v(s)}{2}$, relates to the expected impact of cyberattacks, $(1-\alpha)v(s)$. For simplicity, I assume that $(1-\alpha)v(s)$ is divided equally among players when both exert effort, and, thus, the hacker's payoff is $\frac{(1-\alpha)v(s)}{2} - p_s$.[6] The third term, $-q_s$, captures the firm's cost of investing in cybersecurity. Consequently, q_s can also be interpreted as how aggressively the firm invests in cybersecurity, while p_s can be interpreted as the hacking intensity on such a firm.[7]

[5] In any game in which players would like to outguess one other, there is no Nash equilibrium in pure strategies, as the solution involves uncertainty about what players will do.

[6] Players are implicitly assumed to have equal bargaining power. The Online Appendix shows that introducing differential bargaining power between players only complicates notation without materially changing the qualitative nature of my results.

[7] Incorporating (q_s, p_s) into payoffs as in Table 1 not only enriches the model, as it helps it to account for the costs associated with cyberattacks and cybersecurity investments, but also ensures the uniqueness of the equilibrium by enforcing concavity in the expected payoffs of both players.

Intuitively, firms would opt to defend if the hacker attacks. And firms prefer not to defend themselves if the hacker does not attack. While firms obtain higher payoffs when there is no attack, the hacker gains nothing from abstaining.

Best Response Functions.—Solving the first-order conditions of both players yields:

$$q^*(v, p_s) = \left(\frac{1-\alpha}{4}\right) p_s v \quad \text{and} \quad p^*(v, q_s) = \left(\frac{1-\alpha}{2}\right)\left(1 - \frac{q_s}{2}\right) v \tag{1}$$

That is, a firm's best response, $q^*(v, p_s)$, increases with both its value vulnerable to hacks, v, and the likelihood/intensity of being targeted, p_s. And it decreases with the effectiveness of cybersecurity technologies, α. In turn, the hacker's best response, $p^*(v, p_s)$, increases with v and decreases with both α and q_s. Intuitively, the system of equations (1) reflects both (a) the hacker's understanding that firms invest more in cybersecurity when they are more likely to be targeted and (b) firms' understanding that increasing cybersecurity discourages the hacker from targeting them.

Equilibrium.—In equilibrium, players have no unilateral incentive to deviate as their strategies are best responses to one another. The next proposition demonstrates that the Hacking Game has a unique equilibrium.

PROPOSITION 1 *The simultaneous move game between the hacker and firms has a unique equilibrium. In such equilibrium, a firm of size s faces a targeting probability of p_s^e while investing q_s^e in cybersecurity, where*

$$p_s^e = \frac{(1-\alpha)v(s)}{2\left(1 + \frac{(1-\alpha)^2}{16} v^2(s)\right)} \quad \textit{and} \quad q_s^e = \frac{(1-\alpha)^2 v^2(s)}{8\left(1 + \frac{(1-\alpha)^2}{16} v^2(s)\right)}. \tag{2}$$

Proposition 1 shows that the equilibrium distribution of cyberattacks, captured by p_s^e, and firms' cybersecurity investments, captured by q_s^e, are intimately linked to firms' size distribution—via function $v(\cdot)$—and reshaped by the effectiveness of cybersecurity technologies, α.

Equilibrium Characteristics.—The next lemma underscores the role of cybersecurity technologies and firm heterogeneity on equilibrium cyber risk.

LEMMA 1 *p_s^e and q_s^e decrease with α. In addition, $sign\left(\frac{\partial p_s^e}{\partial s}\right) = sign\left(\frac{\partial q_s^e}{\partial s}\right) = sign\left(\frac{\partial v}{\partial s}\right)$ and $sign\left(\frac{\partial^2 p_s^e}{\partial s \partial \alpha}\right) = sign\left(\frac{\partial^2 q_s^e}{\partial s \partial \alpha}\right) = -sign\left(\frac{\partial v}{\partial s}\right)$, where $sign(\cdot)$ denotes the sign function.*

Lemma 1 highlights three important observations about equilibrium behavior. The first observation relates to how changes in cybersecurity technologies directly alter cyber risk. As cybersecurity technologies become more effective, the hacker faces weaker incentives to attack. In response, firms invest less in cybersecurity.

The next two observations relate to the role of firm heterogeneity on cyber risk. First, the direction in which p_s^e and q_s^e change with s is entirely determined by $v(\cdot)$. Specifically, p_s^e and q_s^e increase with s only if $v(\cdot)$ is an increasing function. When $v(\cdot)$ is increasing (decreasing), hackers opt to target larger (smaller) firms as they obtain higher rewards. As larger firms are expected to be targeted more (less) often, they invest more (less) on cybersecurity. Thus, when $v(\cdot)$ is an increasing function, the model's results are consistent with the observed concentration of cyberattacks among larger institutions—see, Chang et al. (2024).

The second observation relates to the interplay between firm heterogeneity and cybersecurity technologies, which manifest itself in the cross-derivatives $\frac{\partial^2 p_s^e}{\partial s \partial \alpha}$ and $\frac{\partial^2 q_s^e}{\partial s \partial \alpha}$. In particular, whether $\frac{\partial p_s^e}{\partial \alpha}$ and $\frac{\partial q_s^e}{\partial \alpha}$ become more or less negative as s increases depends on $v(\cdot)$. To illustrate this point, suppose $v(\cdot)$ increases (decreases) with s. As cybersecurity technologies become less effective, larger firms experience a relatively more (less) pronounced increase in their hacking intensity when compared to smaller firms. This is because the hacker obtains a larger (smaller) gain per unit of investment when targeting larger firms. In response, larger firms invest relatively more (less) in cybersecurity.

Illustrative Examples.—To illustrate the relevance of the above observations, assume that s follows a Beta distribution with long right tails, making it comparable to the size distribution across U.S. firms. To fix ideas, consider $s \overset{d}{\sim} \beta(2, 8)$. Figure 1a depicts the distributions of both of s and p_s^e under various parameter configurations while considering two distinct functions for $v(s)$. The first function, $v(s) = s$, is increasing in s, while the second one, $v(s) = e^{-s}$, is decreasing in s. For each of these functions, figure 1b depicts the size of the average target, $\mathbb{E}[s_t]$, as a function of α.

Figure 1 shows that, irrespective of the functional form for $v(\cdot)$, the average targeted firm becomes smaller as cybersecurity technologies become more effective. This is because p_s^e decreases linearly with α. Notably, changes in cybersecurity technologies also influence the hacker's targeting strategy. An increase in α not only increases the frequency of smaller values of p_s^e, but also causes the entire probability density function of p_s^e to

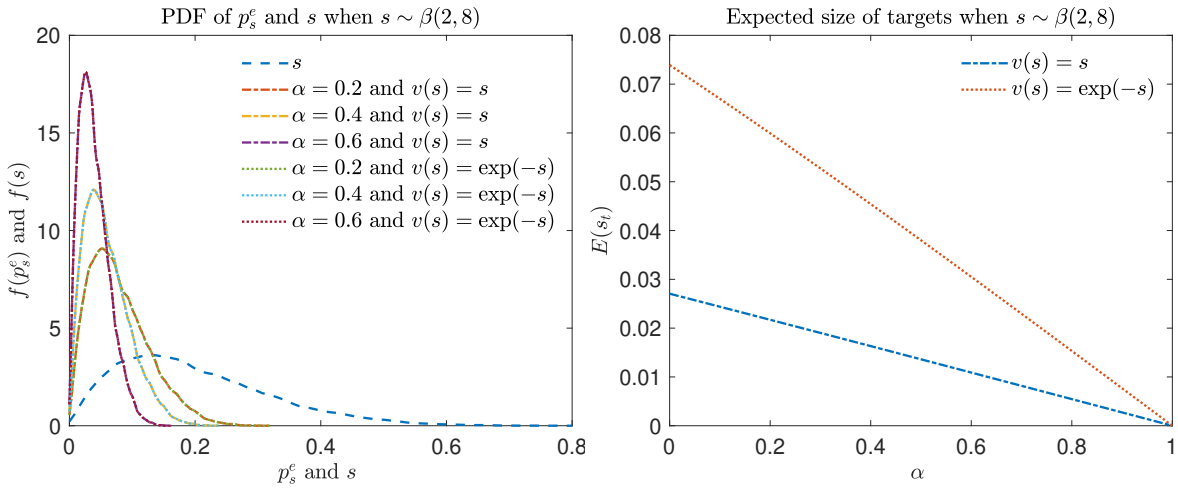

(a) $f(s)$ and $f(p_s^e)$ when $s \overset{d}{\sim} \beta(2,8)$, **(b)** Size of average target, $\mathbb{E}[s_t]$

Figure 1: Impact of function $v(\cdot)$ and α on the size of the average target.

become more concentrated around its mean. That is, targeting becomes less tailored as cybersecurity technologies become more effective.[8]

3 The Hacking Game with Competition

This section explores how competition among hackers can modify firms' and hackers' incentives, ultimately reshaping cyber risk. Besides a unit continuum of firms, consider an economy with $n \geq 2$ hackers that compete when targeting firms. For ease of exposition, I consider economies wherein hacking rewards are split equally among hackers and firms' payoffs are independent of n.[9] For tractability, I focus on equilibria wherein hackers choose the same strategy.

Table 2 reports players' payoffs in an extended version of the Hacking Game that accounts for hacker competition. Reported values consider a firm of size s and an individual hacker—where q_s and p_s are defined as before. Panel A assumes that the remaining $(n-1)$ hackers attack, while panel B assumes that they do not attack.

Before characterizing the equilibrium, it is useful to highlight two key observations that

[8] The Online Appendix also shows that when $v(\cdot)$ is an increasing (decreasing) function, the size of the average targeted firm increases (decreases) as firms become more heterogeneous in size. Because it is more profitable to target larger (smaller) firms, the size of average target increases (decreases) as larger (smaller) firms become more common.

[9] Although assuming that payoffs are divided equally among hackers is not essential, it helps removing equilibrium considerations associated with heterogeneity in bargaining power among hackers. The Online Appendix shows that my results hold even when firms' payoffs directly depend on n.

Panel A: remaining $(n-1)$ hackers attack

HACKER | FIRM OF SIZE s	`Defend` with prob. q_s	`Do not defend` with prob. $(1-q_s)$
`Attack` with prob. p_s	Hacker: $\frac{(1-\alpha)v(s)}{2n} - p_s$ Firm: $\alpha v(s) + \frac{(1-\alpha)v(s)}{2} - q_s$	Hacker: $\frac{(1-\alpha)v(s)}{n} - p_s$ Firm: $\alpha v(s)$
`Do not attack` with prob. $(1-p_s)$	Hacker : 0 Firm: $\alpha v(s) + \frac{(1-\alpha)v(s)}{2} - q_s$	Hacker : 0 Firm: $\alpha v(s)$

Panel B: remaining $(n-1)$ hackers do not attack

HACKER | FIRM OF SIZE s	`Defend` with prob. q_s	`Do not defend` with prob. $(1-q_s)$
`Attack` with prob. p_s	Hacker: $\frac{(1-\alpha)v(s)}{2} - p_s$ Firm: $\alpha v(s) + \frac{(1-\alpha)v(s)}{2} - q_s$	Hacker: $(1-\alpha)v(s) - p_s$ Firm: $\alpha v(s)$
`Do not attack` with prob. $(1-p_s)$	Hacker : 0 Firm: $v(s) - q_s$	Hacker : 0 Firm: $v(s)$

Table 2: Matrix of payoffs with competition among $n \geq 2$ hackers

help distill the role of competition on equilibrium outcomes. First, a hacker's incentives to attack diminish rapidly as n grows—which is interpreted as intensified competition hereinafter. To illustrate this observation, assume $p_s \leq 1/2$. Given (p_s, q_s), a hacker's expected payoffs equals

$$\begin{aligned}
\mathbb{E}[\pi^{\text{hacker}}(p_s)|q_s] &= p_s^n \underbrace{\left\{ q_s \left(\frac{(1-\alpha)v}{2n} - p_s \right) + (1-q_s) \left(\frac{(1-\alpha)v}{n} - p_s \right) \right\}}_{\text{other hackers attack}} \\
&+ p_s(1-p_s)^{n-1} \underbrace{\left\{ q_s \left(\frac{(1-\alpha)v}{2} - p_s \right) + (1-q_s) \left((1-\alpha)v - p_s \right) \right\}}_{\text{other hackers do not attack}} \\
&= \mathcal{O}((1-p_s)^n) \text{ as } n \to \infty.
\end{aligned} \tag{3}$$

That is, $\mathbb{E}[\pi^{\text{hacker}}(p_s)|q_s]$ decline at a rate of $(1-p_s)^n$ when n grows large. Hence, an individual hacker's incentives to attack decrease quickly as competition intensifies.[10]

Second, hackers' actions become less responsive to firms' behavior as competition intensifies. To illustrate this observation, let us consider how n affects a hacker's rewards when other hackers attack. When the firm defends itself, $(1-\alpha)v$ is divided between the firm and hackers—where the firm gets half of $(1-\alpha)v$ while hackers collectively get

[10]Let $f(\cdot)$ and $g(\cdot)$ denote two functions of n. I write $f(n) = \mathcal{O}(g(n))$ as $n \to \infty$ if there exist $\lambda > 0$ and $n_0 \in \mathcal{N}$ such that $|f(n)| \leq \lambda |g(n)|$ for all $n \geq n_0$. For ease of exposition, equation 3 assumes $p_s \leq 1/2$. If $p_s > 1/2$, then $\mathbb{E}[\pi^{\text{hacker}}(p_s)|q_s] = \mathcal{O}(p_s^n)$.

the other half. When the firm opts not to defend itself, $(1-\alpha)v$ is divided only among hackers. Therefore, from an individual hacker's perspective, the difference in payoffs between these cases becomes relevant as it captures the extent to which firms' behavior affects her expected payoffs.

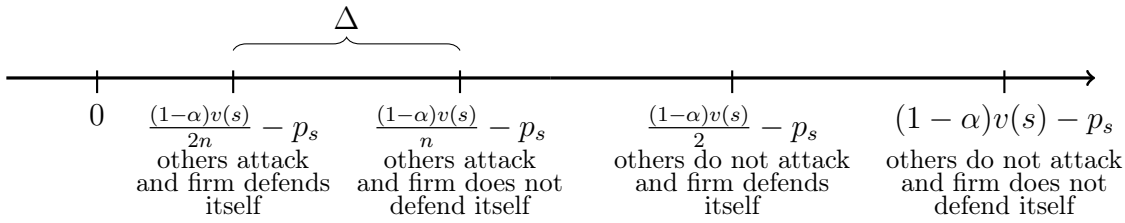

Figure 2: A hacker's payoffs from choosing p_s when $n \geq 3$.

Figure 2 illustrates this difference, labeled as Δ. Intuitively, $\Delta = \frac{(1-\alpha)v(s)}{2n}$ represents the reward an individual hacker forgoes when the firm defends itself. As Δ diminishes, the hacker becomes less sensitive to firms' actions. Consequently, as competition intensifies, hackers' actions become less responsive to changes in firms' behavior, and, as a result, also less responsive to changes in both α and $v(s)$.

Equilibrium.— The next proposition demonstrates that there is a unique equilibrium wherein hackers select the same strategy.

PROPOSITION 2 *The simultaneous move game between $n \geq 2$ hackers and firms has a unique equilibrium wherein hackers choose the same strategy. In such an equilibrium, a firm of size s faces a targeting probability of p_s^e while investing q_s^e in cybersecurity, where p_s^e and q_s^e solve*

$$\begin{aligned} q_s^e &= \left(\frac{1-\alpha}{4}\right)\left(\frac{(p_s^e)^{n-1} + p_s^e(1-p_s^e)^{n-1}}{(p_s^e)^{n-1} + p_s^e(1-p_s^e)^{n-1} + (1-p_s^e)^n}\right) v(s), \\ p_s^e &= \arg\max \mathbb{E}[\pi^{hacker}(p_s)|q_s^e]. \end{aligned} \tag{4}$$

Consistent with previous results, a firm's cybersecurity strategy, q_s^e, increases with its value vulnerable to hacks, $v(s)$, and decreases with the effectiveness of cybersecurity technologies, α. While q_s^e also increases with p_s^e, the precise relationship between q_s^e and p_s^e is reshaped by the intensity of hacker competition, n.

Because of the lack of closed form solutions of (4), I solve the model numerically and present comparative statics results in figures. Figure 3 depicts p_s^e and q_s^e as a function of v and α for different values of n. As the first observation highlights, hackers' incentives

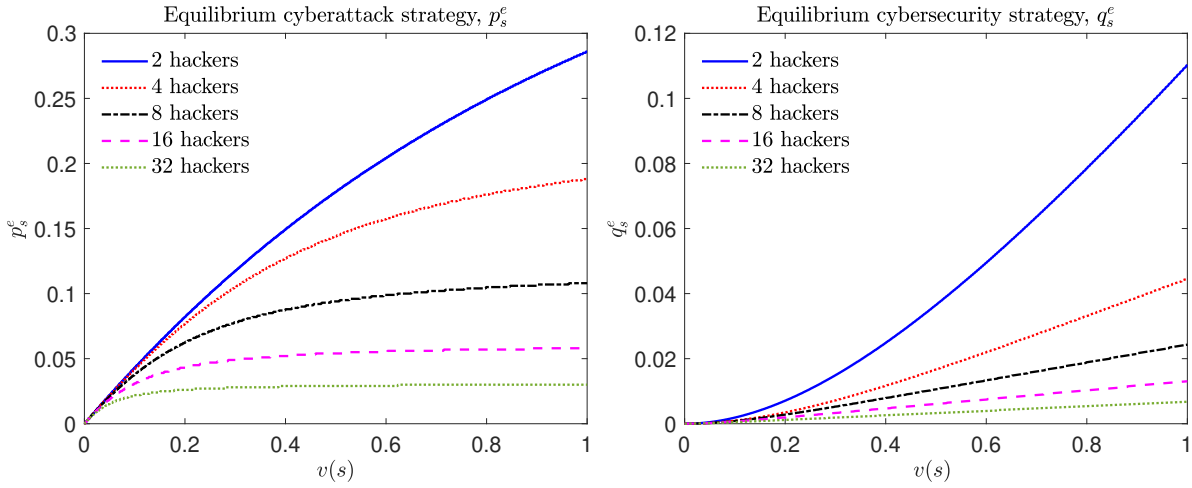

(a) $p^e(v)$ with $\alpha = 0.1$

(b) $q^e(v)$ with $\alpha = 0.1$

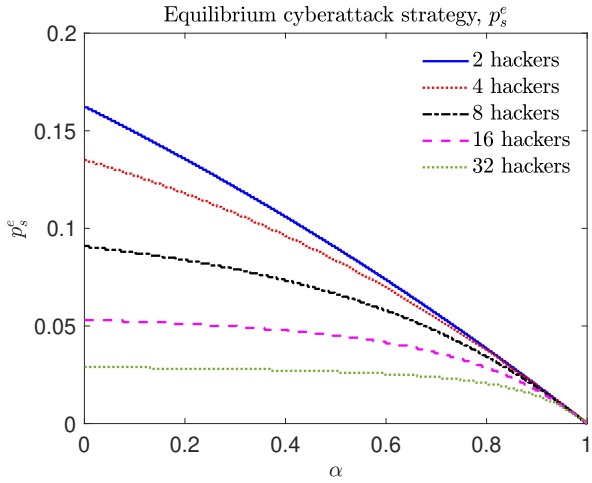

(c) $p^e(\alpha)$ with $v = 0.4$

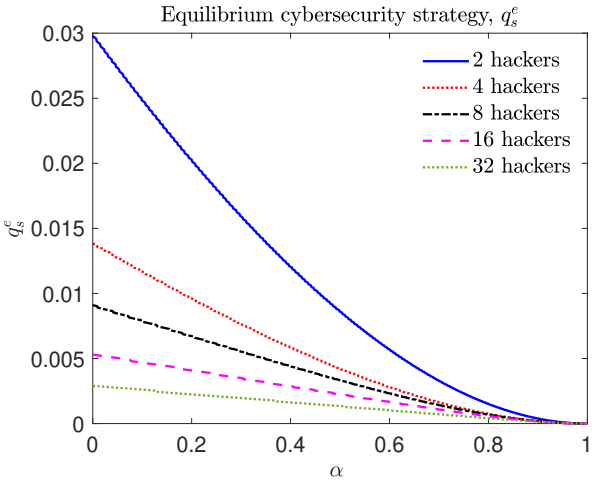

(d) $q^e(\alpha)$ with $v = 0.4$

Figure 3: p^e and q^e as a function of v, α, and n.

to attack diminish rapidly as competition intensifies. And, as the second observation highlights, p^e becomes less sensitive to changes in v and α when n grows—all of which is illustrated by figs. 3a and 3c. Importantly, because firms anticipate this behavior, they also decrease their cybersecurity investments, q_s^e, in a manner that reflects hackers' behavior—as illustrated by figs. 3b and 3d. In sum, increased competition not only decreases hacking incentives and cybersecurity investments, but also reduces the influence of both cybersecurity technologies, α, and firms' vulnerability to cyberattacks, $v(s)$, on equilibrium outcomes, p_s^e and q_s^e.

Competition and the distribution of cyberattacks.—The previous observations have

also an implication for the distribution of cyberattack intensities. Because hackers become less concerned with the precise value of $v(s)$ as n grows, they also become less concerned about firm heterogeneity as n grows. Consequently, the probability density function (pdf) of p_s^e not only moves to the left but also becomes more concentrated around its mean as competition increases. In other words, targeting becomes less tailored as competition intensifies. Figure 4a illustrates this point. To fix ideas, figure 4a assumes that $s \overset{d}{\sim} \beta(2,8)$, $v(s) \in \{s, e^{-s}\}$, and $\alpha = 0.1$.

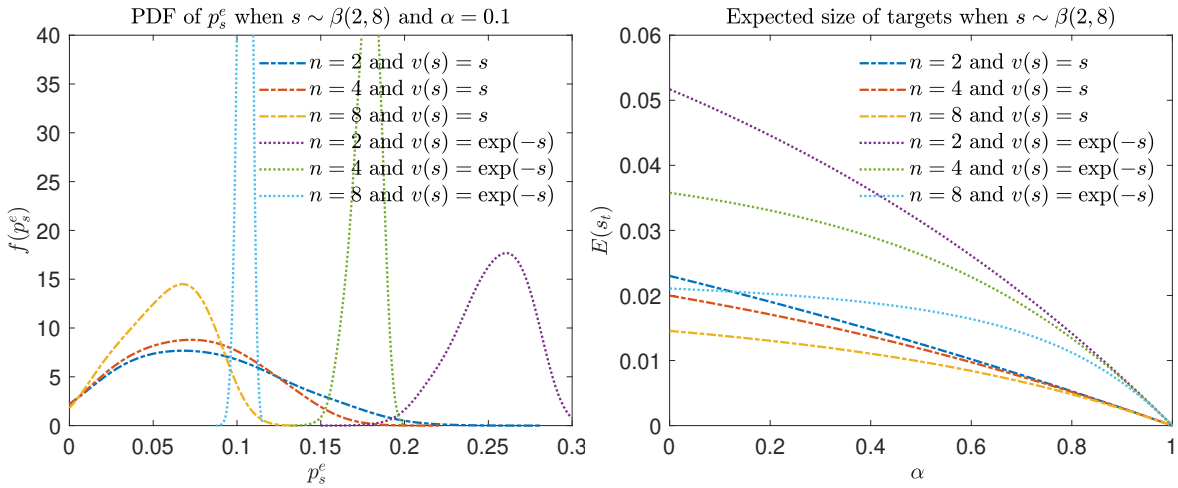

(a) $f(p_s^e)$ when $s \overset{d}{\sim} \beta(2,8)$.

(b) $\mathbb{E}[s_t](\alpha)$ when $s \overset{d}{\sim} \beta(2,8)$.

Figure 4: Impact of $v(\cdot)$ and n on the distribution of p_s^e and $\mathbb{E}[s_t]$.

Random targeting. Figure 4a also helps uncovering a limiting result. As n grows large, the pdf of p_s^e approaches a distribution which assigns positive probability to a single value. Notably, this outcome would be equivalent to an economy wherein hackers target firms uniformly at random as every firm faces the same targeting probability. That is, random targeting becomes observationally equivalent to an economy wherein infinitely many hackers compete.[11]

Competition and the size of the average target.—Figure 4b highlights that increased competition can alter how $\mathbb{E}[s_t]$ decreases with α. As competition intensifies, hacking rewards become smaller, causing hackers to pay less attention to changes in α. As a result, the magnitude of $\left|\frac{\partial \mathbb{E}[s_t]}{\partial \alpha}\right|$ decreases as n increases. That is, the size of the average target becomes less sensitive to variation in cybersecurity technologies as competition intensifies.

[11] The Online Appendix provides further support for this result.

4 Conclusion

I propose a simple model to study how cyber risk can be influenced by the motivations of both institutions and hackers. By establishing a flexible mapping that connects the distribution of cyberattacks to the size distribution of institutions, my model provides insights into how changes in cybersecurity technologies, heterogeneity across institutions, and hacker competition can reshape cyber risk in equilibrium. My findings underscore the importance of understanding the strategic interaction between institutions and hackers when assessing cyber risk.

References

Ablon, Lillian, 2018, Data thieves: The motivations of cyber threat actors and their use and monetization of stolen data, Testimony presented before the House Financial Services Committee, Subcommittee on Terrorism, and Illicit Finance, on March 15, 2018.

Ahnert, Toni, Michael Brolley, David A. Cimon, and Ryan Riordan, 2024, Cyber risk and security investment, *Mimeo* .

Aldasoro, Iñaki, Leonardo Gambacorta, Paolo Giudici, and Thomas Leach, 2020, The drivers of cyber risk, *BIS Working Paper* .

Anand, Kartik, Chanelle Dulley, and Prasanna Gai, 2024, Cybersecurity and financial stability, *Mimeo* .

Chang, Jin-Wook, Kartik Jayachandran, Carlos A. Ramírez, and Ali Tintera, 2024, On the anatomy of cyberattacks, *Economics Letters* 238, 111676.

Chng, Samuel, Han Yu Lu, Ayush Kumar, and David Yau, 2022, Hacker types, motivations and strategies: A comprehensive framework, *Computers in Human Behavior Reports* 5, 100167.

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, 432–448.

Curti, Filippo, Ivan Ivanov, Marco Macchiavelli, and Tom Zimmermann, 2023, City hall has been hacked! the financial costs of lax cybersecurity, *Mimeo* .

Duffie, Darrell, and Joshua Younger, 2019, Cyber runs: How a cyber attack could affect u.s. financial institutions, *Hutchins Center Working Paper* 51.

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

Florackis, Chris, Christodoulos Louca, Roni Michaely, and Michael Weber, 2023, Cybersecurity risk, *Review of Financial Studies* 36, 351–407.

Gibbons, Robert, 1992, *Game Theory for Applied Economists* (Princeton University Press).

Jamilov, Rustam, Hélene Rey, and Ahmed Tahoun, 2021, The anatomy of cyber risk, *NBER Working Paper Series* .

Jiang, Hao, Naveen Khanna, Qian Yang, and Jiayu Zhou, 2024, The cyber risk premium, *Management Science* .

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

Kashyap, Anil K., and Anne Wetherilt, 2019, Some principles for regulating cyber risk, *AEA Papers and Proceedings* 109, 482–487.

Kotidis, Antonis, and Stacey L. Schreft, 2022, Cyberattacks and financial stability: Evidence from a natural experiment, *Finance and Economics Discussion Series (FEDS)* .