A Model of Information Security and Competition*

Alexandre de Cornière† and Greg Taylor‡

September 6, 2024

Abstract

Cyberattacks are a pervasive threat in the digital economy, with the potential to harm firms and their customers. Larger firms constitute more valuable targets to hackers, thereby creating negative network effects. These can be mitigated by investments in security, which play both a deterrent and a protective role. We study equilibrium investment in information security under imperfect competition in a model where consumers differ in terms of security savviness. We show that the competitive implications of security depend on firms' business models: when firms compete in prices, security intensifies competition, which implies that it is always underprovided in equilibrium (unlike in the monopoly case). When firms are advertising-funded platforms, security plays a business-stealing role, and may be overprovided. Regarding policy, the structure of the optimal liability regime also depend on firms' business model.

*We thank Arrah-Marie Jo, Dann Arce and anonymous reviewers at the WEIS 2021 workshop for their comments, as well as the editor and reviewers at Marketing Science. Financial support from from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program) and the Digital Economics Research Network is gratefully acknowledged.

†Toulouse School of Economics, University of Toulouse Capitole, Toulouse, France; alexandre.decorniere@tse-fr.eu

‡Oxford Internet Institute, University of Oxford; greg.taylor@oii.ox.ac.uk

1 Introduction

Digital platforms—from computer operating systems to social media websites, cloud computing services to digital payment services—employ a broad range of business models to provide a stunning diversity of products and services. But a feature they share in common is that information security and data protection are at the core of how they run their businesses. Cybersecurity has shot to the top of the digital policy agenda on the back of a spate of major security breaches. Recently, the Solar Winds attack,[1] a major breach of Microsoft Exchange Server,[2] and the Colonial pipeline hack[3] have each affected millions of individuals, including the clients of many of the world's largest technology firms and those dependent on government services and infrastructure. Attacks are perpetrated by criminals or state actors who may seek to steal data, extort payments through so-called ransomware, or merely to cause damage to the victims. When a firm's system is targeted, its customers suffer through loss of personal data or loss of access to legitimate services. The Colonial hack, and others like it, make clear that cybersecurity is an issue that affects the physical world as well as the digital—a point that will become increasingly salient as connected devices such as autonomous vehicles or smart medical devices become widespread. Industry observers estimate the damage related to cybercrime at $1trn globally in 2020.[4] In response, firms are projected to spend more than $170bn per-year on cybersecurity by 2022.[5]

The fight against cybercrime is as much an economic as a technical one, with both attack-

[1]Remote IT monitoring software provided by Solar Winds to around 30,000 organizations, including governments and multinational businesses, was breached in 2020. A vulnerability in the software allowed hackers to expose the data and systems of around 18,000 organizations and their business partners to harm. See `https://en.wikipedia.org/wiki/2020_United_States_federal_government_data_breach`, accessed 25 July 2021.

[2]Microsoft Exchange Server is an email and calendaring software system used by organizations, including governments and large and small businesses, around the world. In 2021 a vulnerability allowed hackers to steal data or gain control of computer systems using this software. See `https://en.wikipedia.org/wiki/2021_Microsoft_Exchange_Server_data_breach`, accessed 20 July 2021.

[3]In 2021, hackers exploited a vulnerability in the network infrastructure of the Colonial Pipeline to shut down infrastructure that supplies almost half of the oil consumed on the US East Coast. See `https://en.wikipedia.org/wiki/Colonial_Pipeline_ransomware_attack`, accessed 25 July 2021.

[4]See `https://www.mcafee.com/enterprise/en-us/assets/reports/rp-hidden-costs-of-cybercrime.pdf`, accessed 17 February 2021.

[5]See `https://www.gartner.com/en/documents/3889055`, accessed 11 November 2020.

ers and defenders responding to incentives (Anderson and Moore, 2006). Unlike traditional analyses of the economics of crime, which have mostly focused on the incentive effects of criminal justice (e.g., Becker, 1968), a key issue in cybersecurity is that consumers rely on firms to invest in the security of their products, which thus becomes an integral element of market competition. Given that security is a "good" that results from choices made by various actors, a general question is how efficient is the market at providing it? In this paper, we address this question and investigate how firms' incentives to invest in cybersecurity are shaped by their competitive environment. We also investigate various policies aimed at correcting market failures. At a broad level, we show that both equilibrium investment in security and the design of the optimal regulatory regime depend on the interaction between market structure and business model.

We study a simple model of competition between platform firms that offer differentiated products to their customers (who may be businesses or consumers). Hackers are attracted by larger targets (e.g., because they have more data to steal, or because an attack there will generate more damage or publicity), and seek to exploit vulnerabilities in order to breach firms' IT systems. Because larger firms are more likely to be under attack, security concerns dampen the network effects generated by participation on the platforms. A successful attack causes harm to both the firm and its customers, and a key strategic decision for firms is how much to invest in preventing breaches by eliminating vulnerabilities.

The competitive environment is described by three parameters: the intensity of competition, consumers' awareness of security risks, and firms' business model. Competition is measured by market structure and the degree of substitutability among products. We allow for varying levels of consumer awareness, from the case where most consumers are naive about security to the opposite extreme where consumers are quite savvy (e.g., when IT departments of major corporations are procuring enterprise IT systems). Finally, we distinguish between firms whose business model consists in selling a product, and who thus have to choose a price (e.g., cloud service providers), and firms that have other means of monetizing users

(e.g., advertising-supported platforms or firms that sell consumers' data), and therefore seek to maximize demand. We call the former the *pricing regime* and the latter the *advertising regime*.

A principal theme of this paper is that the business model has a significant effect on investment incentives. We show that in the pricing regime firms invest less in security than the efficient solution (and, indeed, sometimes less than a monopolist). This is not only because of the externality security creates for consumers, but also because high levels of security reduce the negative network effects due to hacking, thereby intensifying price-competition. The underinvestment is especially pronounced when breaches cause a big loss for consumers because this amplifies the strength of the negative network effect. In the advertising regime, firms seek to maximize their market share and investing in security becomes a way to attract sophisticated consumers. If there are many such consumers, or if advertising revenue is large, this business stealing effect may even result in equilibrium over-investment and yields higher investment than does monopoly. Unlike the pricing regime, large consumer losses from hacks lead to higher investment as demand becomes more sensitive to security. In our baseline model only firms can exert protection efforts. We discuss the robustness of this assumption in Section 5, where we also allow sophisticated consumers to take preventive actions.

Since the competitive provision of security is generically sub-optimal, we also consider optimal liability regimes in which firms are fined for security breaches and consumers may be compensated for their loss. Here, too, the business model of firms plays an important role. Fines and compensation are strategic substitutes in the pricing regime but strategic complements in the advertising regime. Moreover, an optimal fine in the pricing regime is always punitive (i.e., it exceeds the loss incurred by consumers), whereas fines are non-punitive in the advertising regime.

Although our analysis is motivated by the market provision of security, it can also shed light on other situations. For example, we could let the 'hackers' be agents who post misleading or harmful content online and the firms be social media platforms that invest in

moderation processes to find and remove harmful content before users are exposed to it.[6] More broadly, since security reduces a firm's negative network effect, the analysis shares some features with environments where firms invest to reduce congestion in their service. To the extent that such services use different business models and firms also directly incur damages from congestion, our results can speak to such situations.[7]

To summarise, this paper makes several contributions to the literature on information security. Firstly, it provides a tractable analysis of the market provision of security under competition. Secondly, it studies the effect of firms' business models, which, we show, have significant implications for equilibrium security investment. Thirdly, it provides results on optimal cybersecurity policy under various competitive conditions. These contributions are developed over the coming sections, which are structured as follows: Section 2 outlines our baseline model of security and competition. Section 3 contains the main equilibrium analysis. We study regulatory interventions in Section 4, allow consumers to invest in mitigating the harms from an attack in Section 5, and conclude in Section 6. Most proofs are in the Appendix.

Related literature

Much early work on the economics of information security has its origins on the boundary between economics and computer science and focused on the role that economic forces (such as externalities or moral hazard) play in determining the overall security of a system; Anderson and Moore (2006) and Moore, Clayton, and Anderson (2009) provide an early overview. For a more recent survey of the theoretical literature, see Fedele and Roner (2020).

In monopoly environments, Gordon and Loeb (2002) introduce some of the basic economic trade-offs that face a firm when deciding how much to invest. August and Tunca (2006), Choi, Fershtman, and Gandal (2010) discuss the issue of users' incentives to patch (i.e., to install

[6] Liu, Yildirim, and Zhang (2022) study how a platform's moderation strategy depends on its revenue model.

[7] See, e.g., Johari, Weintraub, and Van Roy (2010) for an analysis of investment under congestion without these features.

security updates). August, Niculescu, and Shin (2014) study the question of versioning cloud versus on premises software. More recent contributions include Lam (2016) on the optimal liability regime, Jullien, Lefouili, and Riordan (2020) on incentives to screen malware-installing advertisers, and Toh (2017) on the role of reputation. Particularly related to our paper is Fainmesser, Galeotti, and Momot (2020), who discuss how business models can shape incentives to collect and protect consumer data for a monopolist. By contrast, our focus is on competitive behavior, which gives rise to a range of novel strategic implications.

A few papers study security investments in oligopolistic set-ups. In a policy paper, Geer, Jardine, and Leverett (2020) provide an overview of the relationship between concentration and cybersecurity risk. Formally, Garcia and Horowitz (2007) highlight that competition may not lead software vendors to invest more in security. Dey, Lahiri, and Zhang (2012) study competition among security providers. Gordon, Loeb, and Lucyshyn (2003) and Gal-Or and Ghose (2005) focus on the issue of information sharing among competitors. In Arce (2018) and Arce (2020), security concerns may offset positive network externalities and prevent tipping in an otherwise winner-take-all market. Indeed, negative network effects are known more generally to relax competition (Navon, Shy, and Thisse, 1995) and in this vein, Nagler (2011) considers a firm that can influence the negative spillovers its customers create for users of other products. In our model negative network effects arise endogenously as a consequence of hacker rationality. We distinguish ourselves from the literature by examining how these effects influence the competitive provision of security by firms and showing that the answer and welfare implications under an ad-funded business model differ from those under price competition. Formally, our model is one of a two-sided market where one side (hackers) creates a negative externality on the other (users), while externalities are positive the other way around. This situation is akin to one where advertisers exert a negative externality on consumers of a media platform (e.g. Anderson and Coate, 2005), except that in our situation the platform cannot charge the hackers and must invest to drive them away.

Empirical work about the link between security and market structure is relatively scant,

and has so far produced mixed evidence: using data about security patches in different software markets, Arora et al. (2010) finds a positive relation between competition and speed of patch releases, while Jo (2019) finds a negative one.

Finally, a large literature studies the question of security in networked environments, where attacks can propagate through connected nodes and where security becomes a public good (Hirschleifer, 1983; Varian, 2004; Goyal and Vigier, 2014; Acemoglu, Malekian, and Ozdaglar, 2016; Dziubiński and Goyal, 2017; Fabrizi, Lippert, and Rodrigues-Neto, 2019). We mostly abstract away from this dimension, even though our extension with protective investment by consumers introduces some externalities, as protection by consumers deters hackers from entering the market.

2 Model

The model consists of three types of agent: two firms indexed $i \in \{1, 2\}$, hackers, and consumers.

Product market Firms' products are differentiated from consumers' point of view, with each firm located at opposite ends of a unit-length Hotelling segment along which consumers are uniformly distributed. Products may be subject to positive direct network effects, meaning they yield a benefit $b \geq 0$ for each consumer that uses them (we also allow the possibility of no direct network effects, in which case $b = 0$). The gross utility of a consumer who selects a product at a distance d from his ideal position and that is used by a mass n of consumers is $V - \tau d + nb$. The stand-alone value, V, is assumed large enough to ensure the market is covered in equilibrium, and we also assume that $\tau > b$ to rule-out cases where the direct network effect causes the market to tip.

We consider two kinds of business model for firms, depending on whether they generate revenues through *pricing* or *advertising*.[8] In the pricing regime, each firm sets a price p_i, which

[8]Throughout, we refer to the latter case as the advertising regime for concreteness. But what's important

enters consumers' utility negatively in a linear way. In the advertising regime, products are free; each firm faces a set of advertisers and runs an optimal auction to sell a slot. We denote by R the expected per-user advertising revenue, which, for most of the analysis, we assume to be the same for both firms. As we show in Section 3.3, allowing R to depend on a firm's individual security does not change the results.

Security Each firm's IT system has potential vulnerabilities that hackers seek to exploit. By fixing vulnerabilities, a firm can reduce the probability that an attack against it is successful. We denote this probability by $1-\sigma_i$, where σ_i is firm i's level of protection (e.g., the share of vulnerabilities that are fixed). Fixing vulnerabilities requires investing in security: achieving a given σ_i costs $\frac{k\sigma_i^2}{2}$.[9] These costs may include hiring software engineers to check for vulnerabilities in the code or to patch exposed vulnerabilities, or training of employees against phishing.

Each firm faces a continuum of mass 1 of hackers. Hackers observe the level of protection of their potential target, and must decide whether to launch an attack. The cost, c, of launching an attack, which includes the required effort as well as the risk of being caught, is independent across hackers and uniformly distributed on $[0,1]$. In case of a breach, a hacker gets an expected payoff of $h \in (0,1]$ per customer.[10]

Thus, if firm i serves n_i consumers, the payoff to attacking i is $h(1-\sigma_i)n_i - c$. Because there is a unit mass continuum of hackers, the number of attacks is equal to the probability that this payoff is positive. It follows that the expected number of *successful* attacks occurring against i is

$$\Pr\left[c < h(1-\sigma_i)n_i\right] \cdot (1-\sigma_i) = h(1-\sigma_i)^2 n_i, \tag{1}$$

where the two terms on the left correspond respectively to the number of attacks on firm i

for our analysis is that the firms rely on some means other than prices to generate revenue. Besides ads, this could include, for instance, selling consumers' data.

[9]Most of our results would be qualitatively unchanged if we allowed a more general convex cost function, $k(\sigma)$. We focus on the quadratic case as this allows us to give closed-form expressions which makes the intuition clearer in places.

[10]The model is also consistent with each attack only affecting a small share of the firm's customers.

and the probability that each attack succeeds.[11]

Consequences of a breach A successful attack imposes damage Δ on a firm. This may capture the administrative cost of responding to the attack and the IT costs of addressing any damage caused, the reputational damaged incurred, or even a fine imposed by regulators (see Section 4).[12] In the online appendix we discuss an alternative model in which this cost is proportional to the number of affected users, Δn_i. Although this makes the analysis less tractable, we show that the most important qualitative features of the model continue to hold in such an environment.

Writing r_i for firm i's per-consumer revenue (i.e., $r_i = p_i$ in the pricing regime, and $r_i = R$ in the advertising regime), firm i's payoff is then

$$\pi_i = \left[r_i - h(1-\sigma_i)^2\Delta\right] n_i - \frac{k\sigma_i^2}{2}, \tag{2}$$

The term $h(1-\sigma_i)^2\Delta$ functions like a marginal cost because the expected number of breaches, (1), depends on the firm's demand.

A successful attack on firm i also imposes an expected loss $L \geq 0$ on each of its customers, stemming from the corruption or fraudulent use of data, from privacy violations or identity fraud, or from interrupted access to compromised services.[13] The utility from choosing product i for a consumer located at a distance d from firm i, when a mass n_i of consumers do the same, is therefore $u_i = V - \tau d + n_i b - p_i - h(1-\sigma_i)^2 n_i L$ (where $p_i = 0$ in the advertising regime).

A fraction $\mu \in [0,1]$ of consumers are sufficiently sophisticated to be able to observe firm

[11] Note that the model is consistent with other assumptions about the security technology. It could be that attacks are always successful but investment in security make them more costly, e.g. with $C(\sigma_i) = c/(1-\sigma_i)^2$, or that both effects are at play, with a probability of success of $\sqrt{1-\sigma_i}$ and a cost $C(\sigma_i) = c/(1-\sigma_i)$. It could also be that security reduces the payoff in case of a successful attack to $h(1-\sigma_i)$.

[12] In an event study, Cavusoglu, Mishra, and Raghunathan (2004) estimate the cost of a revealed breach on publicly traded firms at \$1.6bn. Equifax reported that it incurred technology infrastructure costs (i.e., ignoring legal and liability costs) of \$82.8m after its 2017 breach—see `https://www.bankinfosecurity.com/equifaxs-data-breach-costs-hit-14-billion-a-12473`, accessed 11 November 2020.

[13] Alternatively, the model is consistent with each consumer being directly affected by each attack only with a small probability, which enters L.

i's security, σ_i, and incorporate the security risk into their decision-making. The remaining $1-\mu$ consumers are naive about the risk and ignore it when choosing a firm (formally, they behave as if $L=0$), but still suffer in the event of a breach for the purpose of evaluating welfare. Instead of being naive, we could assume these $1-\mu$ consumers simply don't care about security (meaning they don't incur any loss when a breach occurs). This alternative assumption leaves our results unchanged in the duopoly case—see the online Appendix for details.

Timing and equilibrium The timing is the following: in the first stage, firms simultaneously choose their investment level σ_i, observed by both firms and by savvy consumers. In the second stage, firms choose their prices (in the pricing regime). In the third stage, consumers choose a firm, after forming rational expectations about other consumers' participation.[14] In the fourth stage, hackers facing each firm observe its security and market share before deciding whether to attack. We look for symmetric subgame perfect equilibria. In order to focus on interior solutions throughout the paper, we assume that k is large enough, and that $4\Delta > L\mu$.

2.1 Discussion of assumptions

The literature on the competitive provision of security is nascent and there is still no universally agreed modelling approach. Let us therefore pause briefly to explain the institutional motivation for some of our key assumptions.

Many of the most high-profile cyber attacks exploit so-called *zero day vulnerabilities*. These are vulnerabilities in (software or hardware) systems that lie unknown to the developers or users of that system.[15] A single such vulnerability, if discovered by a hacker, potentially

[14] We do not interpret this as meaning that naive consumers correctly reason about the behaviour of savvy types, but rather require only that, upon observing the realized n_is, no consumer wants to switch to the other firm.

[15] Recent examples include the Solar Winds attack and the Microsoft Exchange Server breach, which each exposed millions of users to attacks.

exposes every user of the system to harm, meaning the firm's system as a whole is targeted rather than individual users. We focus on firms' investment in their systems' security for this reason. While a firm can close a zero day vulnerability at the system level, consumers may themselves be able to take actions to mitigate the harm from any breach. We introduce such consumer effort in Section 5.

Because a single zero day vulnerability can compromise all of a system's users, systems with many users are more attractive targets for hackers. This is reflected empirically in the market value of exploits against systems of various sizes.[16] Recently, U.S. policy makers have called for a halt in consolidation in the health care sector over fears that the creation of large providers increased cybersecurity risks, following an attack on UnitedHealth Group in 2024[17] More generally, hackers should be thought of as rational actors (see, e.g., Ransbotham and Mitra, 2009; Schechter and Smith, 2003). In other words, as Pierce (2016) notes, "Hackers may choose to target larger entities to obtain a large amount of information at once or look for the party with the most vulnerable system protocols."[18] Note that the quote also hints at hackers' ability to observe the security level of firms. We make this assumption in the model, but it is not critical. Indeed, assuming that hackers do not observe σ_i but form expectations σ_i^e about it would result in a risk of successful attack equal to $h(1-\sigma_i^e)(1-\sigma_i)n_i$ instead of $h(1-\sigma_i)^2 n_i$, thereby slightly changing the equilibrium values without affecting the fundamental logic of our arguments.

A typical individual consumer might have little understanding of technical security features, leaving them unable to compare the security of the products available. They correspond to our naive consumers. On the other hand, if the end customer is a business then its technology procurement is likely to be handled by a team of IT experts whose job is to carefully

[16]For example, a vulnerability in the Firefox desktop web browser has about 20% of the market value of a similar vulnerability in (the more popular) Chrome. See `https://zerodium.com/program.html` for example market values, accessed 3 May 2022.

[17]See `https://www.axios.com/2024/04/17/congress-change-healthcare-cyberattack`, accessed 2 July 2024.

[18]The Financial Times also reports a growing emphasis on attacks targeted at large firms—known as "big game hunting" in the hacker community. See, `https://www.ft.com/content/387eb604-4e72-11ea-95a0-43d18ec715f5`, accessed 8 September 2020.

evaluate the security of enterprise IT systems they purchase. We account for the spectrum of consumer savvyness through the parameter μ. Along with the different business models, this allows us to model different kinds of product market: enterprise tools are typically sold for a positive price to more sophisticated buyers, whereas consumer-facing social-media is more likely to be ad funded and targeted at a mostly naive user base.

Our main results do not hinge on the choice of the functional forms for the cost of investment ($k\sigma_i^2/2$) nor on the way investment reduces the expected number of breaches ($h(1-\sigma_i)^2 n_i$). For instance, we would obtain similar results with linear functional forms ($k\sigma_i$ and $h(1-\sigma_i)n_i$ respectively), as discussed in the online appendix.

Lastly, we discuss two features we have deliberately excluded from the model. First, we ignore the potential substitutability between firms on the hacker side. While it is possible that an increase in security by firm i could lead some hackers to target firm j instead of i, we believe this effect to be negligible because there are many firms in other markets, and hackers are not constrained to target a firm in the specific product market we model. Second, a firm's choice of business model is likely to depend on a wide range of market and business imperatives, of which security is only one small part. We therefore take firms' business model as given and focus on the implications of the prevailing business model for security investment.

3 Baseline analysis

3.1 Efficient investment

As a first benchmark, it is useful to compute the optimal decision of a social planner who could symmetrically impose a security investment of σ_w on firms (e.g., by directly regulating firms' security policies) and seeks to maximize total welfare excluding hackers' payoffs.

When firms have $\sigma_1 = \sigma_2 = \sigma_w$, the equilibrium of the ensuing subgame is symmetric

and each firm serves half of the market.[19] Given the assumptions of covered market and unit demand, prices are neutral from a welfare standpoint. It follows that the planner's choice of σ influences welfare only directly via the damage or loss from successful attacks and the firms' costs. Each successful attack generates a social cost of $\Delta + n^* L$, where n^* is the targeted firm's market share. The planner then optimally chooses σ to solve

$$\max_{\sigma \geq 0} \left\{ -n^* h (1-\sigma)^2 (\Delta + n^* L) - \frac{k\sigma^2}{2} \right\}. \tag{3}$$

The solution to this problem is found immediately by taking a first-order condition from the objective function:

$$2n^* h(1-\sigma)(\Delta + n^* L) - k\sigma = 0. \tag{4}$$

After setting $n^* = 1/2$ for the symmetric duopoly, this yields the following Lemma.

Lemma 1. *A social planner that can control $\sigma_1 = \sigma_2 \equiv \sigma$ to maximize total welfare optimally selects*

$$\sigma_w^* = \frac{h(L + 2\Delta)}{2k + h(L + 2\Delta)}. \tag{5}$$

The comparative statics are rather intuitive: the efficient investment level is increasing in h (hackers' gains), Δ (damage to firms) and L (damage to consumers), and decreasing in k, the cost of providing security.

3.2 Pricing regime

Turning to the case of duopoly competition, we know that the number of successful attacks against firm i in the last stage of the game is $h(1-\sigma_i)^2 n_i$. We now proceed by backward induction, starting from consumers' decisions.

Demand Because only sophisticated consumers can observe security levels and take them into account, their behavior differs from naive consumers.

[19] We establish this formally in Section 3.2 below.

The sophisticated consumer who is indifferent between firm 1 and 2 is located at $x \in [0,1]$ solving

$$V - x\tau + n_1 b - p_1 - h(1-\sigma_1)^2 n_1 L = V - (1-x)\tau + n_2 b - p_2 - h(1-\sigma_2)^2 n_2 L,$$

i.e.,

$$x = \frac{\tau - p_1 + p_2 + b(n_1 - n_2)}{2\tau} - \frac{hL}{2\tau}\left(n_1(1-\sigma_1)^2 - n_2(1-\sigma_2)^2\right). \tag{6}$$

The first term on the right-hand side of (6) is the demand in a standard Hotelling model with direct network effects. The second term, on the other hand, shows the endogenous presence of negative indirect network externalities in the model: as more consumers choose firm i, it is targeted more often (for a given σ_i), which makes i less attractive to other consumers. Let us stress that, while security concerns create a negative network effect, it need not be the case that consumers prefer smaller firms. Indeed, if b is large relative to hL then consumers will enjoy a net benefit when more consumers patronize the same firm, even though it is more likely to be breached (for a given security level).

The indifferent naive consumer is located at $y \in [0,1]$ such that

$$V - y\tau + n_1 b - p_1 = V - (1-y)\tau + n_2 b - p_2,$$

i.e.,

$$y = \frac{\tau - p_1 + p_2 + b(n_1 - n_2)}{2\tau}. \tag{7}$$

Naive consumers do not perceive the greater risk of attack as n_i increases, and there is thus no security-based negative network externality term in their demand.

For equilibrium consistency we must have $n_1 = \mu x + (1-\mu)y$ and $n_2 = \mu(1-x) + (1-\mu)(1-y)$. Solving this system of equations yields the demand functions

$$n_1^* = \frac{[\tau - b] + p_2 - p_1 + hL\mu(1-\sigma_2)^2}{2[\tau - b] + hL\mu[(1-\sigma_1)^2 + (1-\sigma_2)^2]}, \quad n_2^* = 1 - n_1^*. \tag{8}$$

Observe that a one-unit increase in the strength of network effects, b, has the same effect on demand as a one-unit decrease in the transport cost, τ.

Pricing stage Given σ_1 and σ_2, firms choose prices to maximize (2), with $r_i = p_i$ and demand given by (8). Firm i's first-order condition is $\frac{\partial \pi_i}{\partial p_i} = 0$ and solving this system yields the equilibrium prices:

$$p_i^* = \tau - b + \frac{1}{3}h\Big\{ (\Delta + L\mu)\left[3 - (2-\sigma_i)\,\sigma_i - (2-\sigma_j)\,\sigma_j\right] - \Delta\,(2-\sigma_i)\,\sigma_i - L\mu\,(2-\sigma_j)\,\sigma_j \Big\}. \quad (9)$$

In a standard Hotelling game with direct network effects only we would have $p_i^* = \tau - b$; adding security concerns introduces the second term. The next result will play an important role in the subsequent analysis.

Lemma 2. *In the pricing subgame, prices are a decreasing function of the level of security:* $\frac{\partial p_i^*}{\partial \sigma_i} < 0, \frac{\partial p_i^*}{\partial \sigma_j} < 0$.

Investment in security has the strategic effect of intensifying subsequent price competition, which plays an important role in the analysis to follow. This effect operates through two channels, a cost and a demand one. Regarding the cost channel, a firm that has invested a lot in security faces a lower effective marginal cost. Indeed, whenever an extra consumer chooses firm i, the expected damage incurred by i increases by $(1-\sigma_i)^2\Delta$, which is decreasing in σ_i. This lower effective marginal cost leads firm i (and j, by strategic complementarity) to reduce prices.[20]

For the demand channel, notice that the price-elasticity of firm i's demand is

$$\eta_i = \frac{p_i}{\tau - b - p_i + p_j + hL\mu(1-\sigma_j)^2}. \quad (10)$$

[20] Anecdotally, IBM Security presents raising the price charged to end-users as one possible strategic response to increased security costs: `https://securityintelligence.com/articles/going-up-how-to-handle-rising-cybersecurity-costs/`.

As firm j increases σ_j, firm i's demand becomes more price-elastic. Indeed, inspection of (8) reveals that an increase in σ_j reduces firm i's demand and increases its sensitivity ($\partial^2 n_i/\partial p_i \partial \sigma_j < 0$), as the negative network effects become smaller. Because of this increased price-elasticity, a rise in σ_j leads firm i to charge a lower price. By strategic complementarity of prices, firm j also lowers its price following an increase in σ_j. In other words, more investment reduces the strength of the negative indirect network effects due to security concerns, thereby intensifying price competition.

The fact that a firm's security induces a downward shift in that firm's reaction function highlights a fundamental difference between security and a more traditional notion of "quality", which often leads to an upward shift in the firm's own reaction function. Below we will elaborate on the implications of this observation in terms of comparative statics.

Investment stage In the first stage of the game, each firm's problem is

$$\max_{\sigma_i \geq 0} \left\{ \left[p_i^* - h(1-\sigma_i)^2 \Delta \right] n_i^* - \frac{k\sigma_i^2}{2} \right\}, \tag{11}$$

where p_i^* and n_i^* are respectively given in (9) and (8). Best-response functions are decreasing and cross only once if $\tau \geq \hat{\tau}_1 \equiv \frac{12hk^2\left(h\left(16\Delta^2+7\mu^2L^2+4\Delta L\mu\right)-18kL\mu\right)}{(h(4\Delta-L\mu)+6k)^3} + b$. In that case there is a unique equilibrium, which is symmetric. We assume that this is the case for the time being, and later come back to the case where $\tau < \hat{\tau}_1$.

Computing $\frac{\partial \pi_i}{\partial \sigma_i}$ and imposing symmetry ($\sigma_i = \sigma_j$) yields $\frac{1}{6}(h(4\Delta - L\mu)(1-\sigma) - 6k\sigma) = 0$. This is solved by the symmetric equilibrium level of investment, σ_p^*:

$$\sigma_p^* = \frac{h(4\Delta - L\mu)}{6k + h(4\Delta - L\mu)}. \tag{12}$$

As one might expect, a firm's equilibrium investment in security is increasing in the gains from hacking, h, and in the damages from a successful attack, Δ, while it is decreasing in the cost of investing, k.

The effect of the parameters L and μ on σ_p^* are more novel: the equilibrium investment is decreasing in the share of sophisticated consumers (μ) and in the damage consumers incur in case of a breach (L). This is due to the strategic effect mentioned above: looking at (10), we see that the effect of σ_j on η_i is stronger for larger values of μ and L. Thus, as L and μ increase, incentives to invest in security are weakened by the competition-intensifying strategic effect (the reason h does not play the same role is that it also enters the expected cost).

One can also notice that the intensity of competition, captured by the inverse of $\tau - b$, does not affect the equilibrium investment in security (provided that $\tau \geq \hat{\tau}_1$). This is because of two opposite effects. On one hand, demand for firm i is less sensitive to σ_i as $\tau - b$ increases, by (8). On the other hand, the equilibrium price increases with $\tau - b$ (see below), which means that each additional customer attracted by improved security is worth more. In the current specification with linear transportation costs, these two effects exactly cancel one another.

In terms of efficiency, comparing (5) and (12), we find that

$$\sigma_w^* - \sigma_p^* = \frac{2hk(2\Delta + L(3+\mu))}{(2k + h(L+2\Delta))(6k + 4h\Delta - hL\mu)} > 0, \tag{13}$$

so firms under-invest in security in equilibrium. This happens for two reasons. Firstly, unlike the social planner, firms do not fully-internalize consumers' losses when choosing the optimal investment. Secondly, the aforementioned strategic effect gives firms an incentive to under-invest in order to soften price competition from their rival.

We summarise these results in the following proposition (the proof is immediate from (12) and (13)).

Proposition 1. *Suppose that $\tau \geq \hat{\tau}_1$. In the pricing regime, firms under-invest in security compared to the socially optimal solution.*

Firms' investment in security is decreasing in μ, L, and k; increasing in h and Δ; and

independent of τ and b.

As for equilibrium prices, substituting σ_p^* into (9) we obtain:

$$p^* = \tau - b + \frac{36hk^2(\Delta + L\mu)}{(6k + 4h\Delta - hL\mu)^2} = \tau - b + h\Delta(1 - \sigma_p^*)^2 + \frac{36hk^2 L\mu}{(6k + 4h\Delta - hL\mu)^2}. \qquad (14)$$

Recall that a standard Hotelling model with direct network effects and marginal costs $h\Delta(1-\sigma_p)^2$ would yield an equilibrium price of $\tau - b + h\Delta(1 - \sigma_p)^2$. Because of the presence of endogenous negative network effects discussed above, the price-elasticity of demand is lower than in the standard Hotelling model, leading to higher prices in equilibrium.

The equilibrium price is an increasing function of μ and L: these parameters amplify the negative indirect network effects, and make firms less willing to cut prices to attract new consumers. Similarly, an increase in the cost of security k leads to higher prices, as less security means stronger negative network effects. The effect of an increase in the hacking activity h is more ambiguous. Indeed, we have $\partial p^*/\partial h > 0$ if and only if $h < k/(4\Delta - L\mu)$, meaning that there is an inverted-U relationship between h and p^*. Two opposite effects are at play here. On the one hand, an increase in the prevalence of hacking induces firms to invest more in security, which intensifies competition and pushes prices down. On the other hand, more hacking means that the negative network effects are larger, which softens competition. When the cost of providing security k is large, the second effect dominates (σ is not very responsive to h), and prices go up.

By affecting firms' pricing and investment decisions, any consumer who becomes savvy exerts an externality on their peers. Indeed, substituting the equilibrium values of σ_i, p_i and $n_i = 1/2$ into u_i, we have

$$\frac{\partial u_i}{\partial \mu} = -\frac{36hk^2 L[6k + h(6\Delta + L(1 + \mu))]}{(6k + 4h\Delta - hL\mu)^3} < 0.$$

The externality is therefore negative. In particular, this means that the presence of savvy consumers harms naive ones. Additionally, notice that, using the envelope theorem, an

increase in μ affects i's equilibrium profits only via its effect on σ_j and p_j. Since a higher μ causes the rival to be a softer competitor (σ_j decreases and p_j increases), firms' profits must increase as more consumers become savvy.

Non-covered market If we relax our assumption that the market is covered then competition between firms can expand total demand. The standard Hotelling model is not well-suited to such an exercise,[21] but we can proceed by modifying the model in the following way: Suppose that the two firms are, as before, located at points 0 and 1 on the Hotelling line. "Interior" consumers are distributed uniformly with density λ over $[0, 1]$. Additionally, each firm has a "hinterland" of captive consumers who consider only its product and whose distance from the firm is distributed uniformly with density $1 - \lambda$ over $[0, \infty)$. Thus, even if the interior segment is covered and subject to competition, a price cut leads to market expansion by increasing the number of hinterland consumers who buy from the firm. We focus on situations where the interior market segment is covered so that there is meaningful competition between the firms. Letting $\lambda = 1$ returns us to the baseline model.

This setup yields an alternative demand function but the analysis otherwise follows the same steps described above.[22] The model with $\lambda < 1$ is less tractable than the case of $\lambda = 1$ studied above, but can be solved numerically. The strategic effect identified previously is still in place: investing in security strengthens price competition. But the role played by savvy consumers is now more nuanced. Adding more savvy consumers strengthens the strategic effect as before, but also increases the rate at which security investments attract new hinterland consumers to the market. Thus, the relationship between μ and σ_p^* depends on which of these two forces is strongest. When λ is large, most of a firm's marginal consumers

[21] If the market in a Hotelling model is not covered then firms do not compete at all and act as local monopolists.

[22] Specifically, the demand for firm 1 from savvy consumers is $D_1^{\text{savvy}} = \lambda x + (1-\lambda)x_1^H$ where x is its share of consumers located in $[0, 1]$, given by (6), and x_1^H is the indifferent hinterland consumer, solving $v - p - \tau x_1^H + [b - h(1-\sigma_1)^2 L]n_1 = 0$. Demand from naive consumers is similarly given by $D_1^{\text{naive}} = \lambda y + (1-\lambda)y_1^H$ where y is found in (7) and y_1^H solves $v - p - \tau y_2^H + bn_1 = 0$. Lastly, we have $n_1 = \mu D_1^{\text{savvy}} + (1-\mu)D_1^{\text{naive}}$. This system of equations can be solved for n_1, (and analogously for n_2) after which we can follow the same steps described above to compute the equilibrium values of p_i and σ_i.

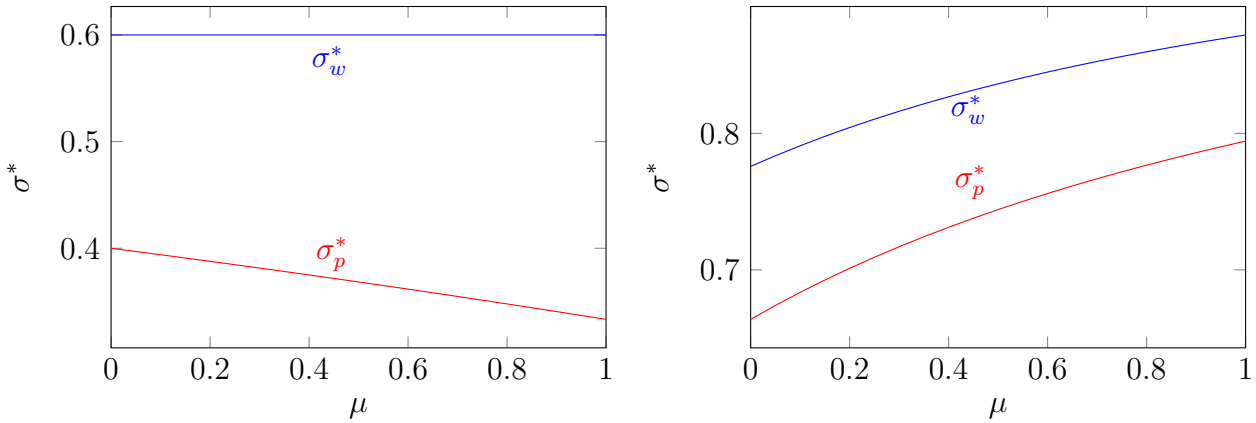

(a) $\lambda = 0.99$ (b) $\lambda = 1/2$

Figure 1: When λ is large the strategic effect dominates and investment decreases in μ. For smaller λ the market-expansion effect leads investment to be increasing in μ. Investment is inefficiently low in both cases. The plot is drawn for $\Delta = h = L = k = \tau - b = 1$ and $v = 3$.

come from its rival, meaning the strategic effect is strongest and σ_p^* is decreasing in μ. Conversely, when λ is small most of the marginal consumers come from market expansion and σ_p^* is increasing in μ. In either case, the strategic effect leads to underinvestment. An example can be seen in Figure 1.

We emphasise that our main point is not to argue that the strategic effect we identify is the only factor in firms' decision-making, but rather to highlight this is a novel strategic force at play in markets for information security with the potential to distort security away from its optimum level.

Market structure Other factors that can influence the strength of the strategic effect are market structure and concentration. In Appendix B we investigate how these influence investment in security. There we show that, when V is large enough,[23] a monopolist chooses the efficient level of investment, fully internalizing consumers' losses via its price. Going from one to two firms therefore results in worse investment incentives for firms because the strategic effect described in Lemma 2 comes into effect. Starting from duopoly, on the other hand, an additional firm weakens the strategic effect because a change in each firm's investment has

[23]Large V simplifies the analysis by ensuring the market is covered. But even if we allow for a non-covered market we still sometimes find that monopoly yields better investment incentives than duopoly.

a small impact on rivals' pricing when it is just one of many competitors. Overall, then, we find that the number of competitors can have a non-monotonic effect on investment in the pricing regime.

We can extend this intuition to think about what happens when the number of firms is held fixed at three but the market concentration is varied. A firm's investment decision exerts a stronger strategic effect if it is an important competitor for its rivals. The strategic effect is therefore strongest for the one or two firms that account for the largest share of highly concentrated markets. On the other hand, markets with symmetric firms tend to minimize the size of the strategic effect and yield higher investment in security. Details can be found in Appendix B.

Asymmetric equilibrium When $\tau < \hat{\tau}_1$, the equilibrium described above is no longer stable. Instead, there are two stable asymmetric equilibria, mirror images of one another. One firm (say firm 1) invests more than the other, charges a lower price, and attracts a higher market share. Interestingly, even though firm 1 invests more, in equilibrium it is more likely to suffer a breach because it attracts more hackers due to its larger market share. Firm 2 invests less, charges a higher price and serves fewer consumers. Firm 2 serves relatively more savvy consumers because it suffers fewer breaches. The analytical solution is complex and not illuminating, but Figure 2 shows some equilibrium values of interest. As τ approaches $\hat{\tau}_1$, the asymmetric equilibria smoothly converge to the stable one studied above.

3.3 Advertising regime

In the advertising regime, demand is given by (8) with $p_i = p_j = 0$. Firms' security investment is chosen to solve

$$\max_{\sigma_i > 0} \left\{ \left[R - h(1-\sigma_i)^2 \Delta \right] n_i^* - \frac{k\sigma_i^2}{2} \right\}, \tag{15}$$

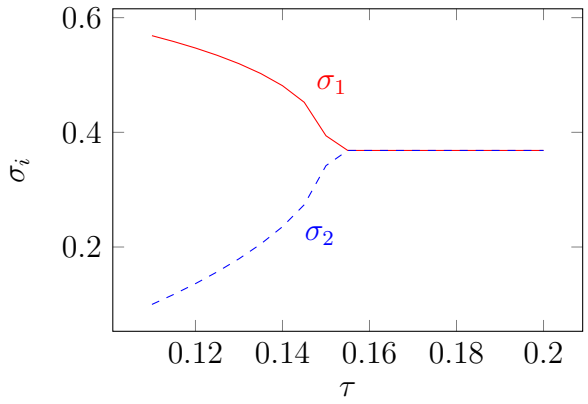

(a) Security investments.

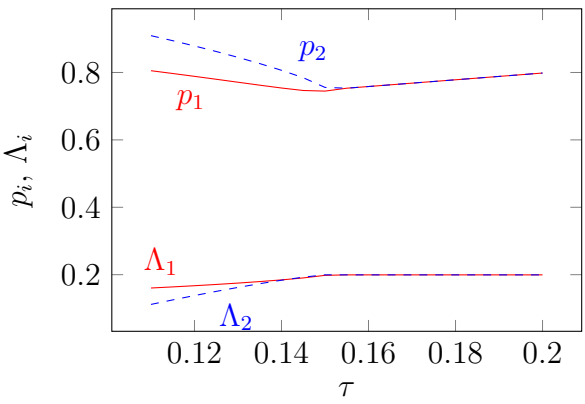

(b) Prices and consumer losses from breaches.

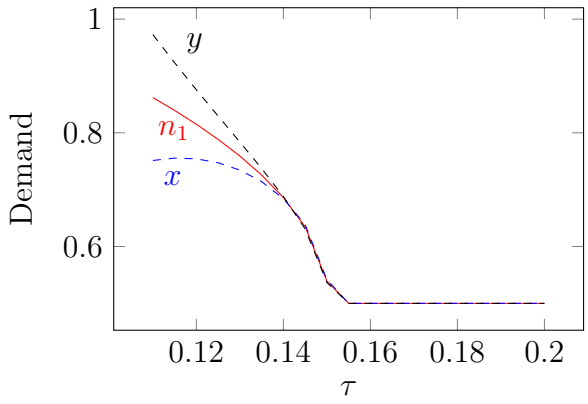

(c) Demands.

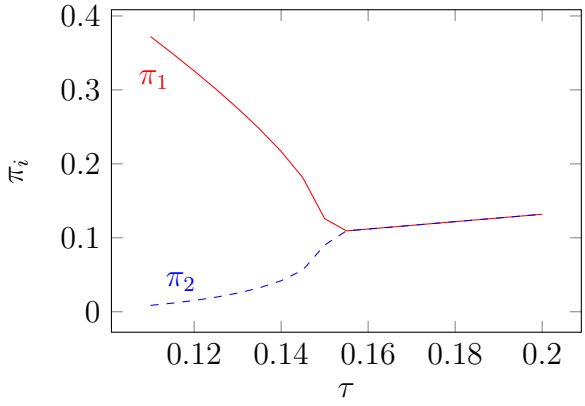

(d) Profits.

Figure 2: Asymmetric equilibria for τ small, $b = 0$ and $\Delta = h = L = k = 1, \mu = \frac{1}{2}$. The consumer losses from breaches at firm i is defined as $\Lambda_i = h(1-\sigma_i)^2 L n_i$.

where n_i^* is given in (8) (again, with $p_1 = p_2 = 0$). As in the pricing regime, the symmetric equilibrium is stable if and only if τ is large enough, which we assume from now on.[24] The symmetric equilibrium σ_a^* is implicitly given by evaluating firms' first-order conditions at $\sigma_i = \sigma_j = \sigma$:

$$\frac{h\left\{2(\tau-b)\Delta + L\mu\left[R + h\Delta(1-\sigma)^2\right]\right\}(1-\sigma)}{2\left(\tau - b + hL\mu(1-\sigma)^2\right)} - k\sigma = 0. \tag{16}$$

Applying the implicit function theorem to (16) allows us to study how the equilibrium in-

[24]The cut-off value for τ, $\hat{\tau}_2$, is different from $\hat{\tau}_1$, and its analytical expression is complicated. When $\tau < \hat{\tau}_2$, the asymmetric equilibria yield negative profits. In that case we would expect one firm to drop out, resulting in a monopoly, but we do not model such an entry game.

vestment level responds to the model's parameters. The following result summarizes and also compares equilibrium investment to the socially optimal solution. Its proof is in Appendix A.

Proposition 2. *In the advertising regime, firms over-invest compared to the socially optimal solution if*

$$\mu\left(R-\frac{4hk^2(L+\Delta)}{(2k+h(L+2\Delta))^2}\right)>\tau-b \tag{17}$$

and under-invest if the inequality is reversed.

Firms' investment in security is decreasing in τ and k; and increasing in μ, L, h, b, and Δ.

The model with ad-funded firms delivers different predictions from the one where firms compete in prices. First, the comparative statics with respect to several key parameters are different. Equilibrium investment increases in L and μ: as consumers become more sensitive to security differences, firms invest more. There is no strategic effect through which security would intensify price-competition. Security is also greater when competition is more intense (τ is small) or network effects strong (b is large) because the mark-up is independent of τ and b and therefore does not offset the effect on demand sensitivity, as under price-competition.

Second, there can be over-investment in equilibrium compared to the social planner's solution, σ_w^*. This can happen because of a business-stealing effect: when R or μ are large, or when $\tau - b$ is small, the private payoffs from increasing security are larger than the social one, resulting in over-investment. The ad-funded business model is typically used in B2C markets, where the consumers are less likely to be savvy about security risks (low μ). We would therefore expect over-investment to arise only when products exhibit little differentiation.

We show in Appendix B that, unlike the pricing regime, investment in security is higher than that chosen by a monopolist (holding the scale of operations fixed). This is because the strategic effect of Lemma 2 is not active in the advertising regime, leaving duopolists to

compete for savvy consumers via security investments.

Endogenous ad revenue We have assumed that R is exogenously fixed. But one might expect R to be determined by a firm's choices in a way that interacts with security. For example, suppose advertisers do not wish to be associated with insecure firms—their inverse demand for ad impressions is $P(A, \sigma)$ when a firm shows A ads to each consumer, with $\partial P/\partial \sigma > 0$. The firms choose $A^*(\sigma_i)$ as the solution to $\max_A P(A, \sigma_i)A$. We then have $R(\sigma_i) = P(A^*(\sigma_i), \sigma_i)A^*(\sigma_i)$ and $R'(\sigma_i) > 0$.

We can easily allow $R'(\sigma_i) \neq 0$ in the model. Then (16) becomes

$$\frac{h\left\{2(\tau - b)\Delta + L\mu\left[R + h\Delta(1-\sigma)^2\right]\right\}(1-\sigma)}{2\left(\tau - b + hL\mu(1-\sigma)^2\right)} = k\sigma - \frac{1}{2}R'(\sigma).$$

It is immediate that, given basic regularity conditions on $R(\cdot)$, this is equivalent to a transformation of the marginal cost of investment and our results go through.

3.4 Comparison to a standard model of investment in quality

Let us compare our results to a standard model of Hotelling competition with investment in quality to highlight how security is different. To that end, suppose savvy and non-savvy consumers respectively perceive utility $V + q_i - \tau x - p_i + bn_i$ and $V - \tau x - p_i + bn_i$ when they are located at distance x from firm i, where q_i is i's quality. Profit is $\pi = r_i n_i - kq_i^2/2$. The timing is as above (but with firms choosing q_i instead of σ_i and no role for hackers).

Pricing regime Computing the location of the indifferent consumers, demand is $n_i = \frac{1}{2(\tau - b)}\left[p_j - p_i + \tau - b + \mu(q_i - q_j)\right]$. Substituting this into the profit function, and solving the system of best responses, $\{\frac{\partial \pi_i}{\partial p_i} = 0\}_{i=1,2}$, yields equilibrium price: $p_i = \tau - b + \frac{\mu}{3}(q_i - q_j)$. Unlike (9), we see that the strategic effect that led a firm's price to be decreasing in its own security level is absent when firms instead choose quality. This difference arises for two reasons. Firstly, security weakens the negative network effect and strengthens competition,

whereas quality does not.[25] Secondly, quality investment only increases the value of the product (analogous to a product innovation), whereas security also reduces firms' damages, which enter profit like a marginal cost, and therefore corresponds to both a process and product innovation.

Substituting the price into the profit function, we can compute the symmetric equilibrium quality,

$$\frac{\partial \pi_i}{\partial q_i}|_{q_1=q_2=q} = 0 \iff q_p^* = \mu/3k.$$

Because of the diminished strategic effect of investment, equilibrium quality is now an increasing rather than a decreasing function of μ. As consumers become more responsive to quality, firms invest more in order to create value that it can extract through a higher price.

Advertising regime We substitute n_i into π_i, let $r_i = R$, and solve $\frac{\partial \pi_i}{\partial q_i}|_{q_1=q_2=q} = 0$ to yield the equilibrium quality,

$$q_a^* = \frac{R\mu}{2k(\tau - b)}.$$

Here the results share much in common with security investment: quality is increasing in μ and can be excessive if R is large.[26] Thus, the important distinction between business models in the security investment game vanishes when firms make more standard quality investments.

4 Regulation

The previous analysis suggests that equilibrium investment in security is unlikely to be socially optimal, and that there is therefore scope for policy interventions aimed at correcting distortions. Broadly speaking, there are three main policy approaches: transparency initiatives such as notification requirements or certification schemes, regulated minimum security

[25] Formally, this means that quality translates into a parallel shift in demand, whereas security produces a rotation.

[26] The social planner in the quality investment game solves $\max \frac{1}{2}q - \frac{k}{2}q^2 \implies q_w^* = 1/2k$.

standards, and financial penalties or liability for breaches. However, there does not yet exist a globally consistent approach to policy in this area. In the United States, few laws exist at the federal level, except with respect to specific industries such as health. States have moved to fill this vacuum, with the main focus being on obligations to disclose security breaches (e.g., 2003 California Notice of Security Breach Act) and the requirement for minimum security standards (e.g., as imposed in the 2004 California Assembly Bill 1950). Firms can also be held accountable for security breaches under civil litigation if they can be shown to have been negligent. The European Union has been more active in policy-making. As well as obligations to disclose breaches, the EU Cybersecurity Act created a certification scheme aimed to increase the transparency of firms' security arrangements, while the GDPR more recently introduced significant statutory fines for firms that suffer a breach. For example, in 2020, British Airways and hotel chain Marriott were respectively fined £20m and £18.4m for data breaches affecting hundreds of thousands or millions of customers.[27]

4.1 Optimal liability regime

Suppose that the regulator can impose a fine $f \geq 0$ on a firm in case of a breach, and can award a compensation $g \in [0, L]$ to consumers. Such instruments are, for instance, available under the EU GDPR (Articles 82 and 83). The actual loss for the firm is now $\Delta + f$, while the harm to consumers is $L - g$. We say that a pair $\{f, g\}$ is optimal if the equilibrium choice of σ under this liability regime coincides with the efficient level, σ_w^*.

Pricing regime The condition for equilibrium investment to be at the socially optimal level is

[27] See https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2020/10/ico-fines-british-airways-20m-for-data-breach-affecting-more-than-400-000-customers/ and https://ico.org.uk/about-the-ico/news-and-events/news-and-blogs/2020/10/ico-fines-marriott-international-inc-184million-for-failing-to-keep-customers-personal-data-secure/, accessed 10 November 2020.

$$\sigma_w^* = \sigma_p^* \iff \frac{h(L+2\Delta)}{2k+hL+2h\Delta} = \frac{h[4(\Delta+f)-(L-g)\mu]}{6k+4h(\Delta+f)-h(L-g)\mu}.$$

There is therefore a continuum of (f, g) pairs that implement the planner's solution, with the optimal fine being $f_p^*(g) = \frac{1}{4}[2\Delta - g\mu + L(3+\mu)]$.

Several observations are in order. First, the optimal fine is a decreasing function of the compensation awarded to consumers. In other words, f and g are strategic substitutes. The reason for this is that as the amount of compensation rises, the price elasticity of demand for firm i, which equals $p_i / [(\tau - b) - p_i + p_j + h(L-g)(1-\sigma_j)^2]$, becomes less sensitive to σ_j, so that the strategic effect leading to under-investment weakens.

Second, as long as $g \leq L$, we have $f_p^*(g) > L/2$. In words, the fine exceeds the loss incurred by consumers. While a fine of $L/2$ in case of a breach would lead firms to internalize consumers' losses, it would not be enough to correct the strategic effect leading to under-investment. The optimal fine therefore needs to be punitive in order to induce efficient levels of investment.

Advertising regime The first-order condition determining equilibrium investment is given in (16). Suppose we implement a budget-balanced policy that fully-compensates consumers ($g = L$ and $f = L/2$). Making this substitution, (16) simplifies to $\frac{1}{2}h(L+2\Delta)(1-\sigma) - k\sigma = 0$, which is precisely the condition solved by the social planner (cf. equation 4). We therefore observe that this budget-balanced full-compensation policy exactly implements the planner's solution. Intuitively, setting $f = L/2$ causes firms to completely internalize consumers' losses so that there is no externality distortion. Moreover, if $g = L$ then there's no business-stealing effect because consumers become insensitive to firms' investments. Both of the effects that might cause equilibrium to depart from the efficient level of investment are therefore neutralized.

More generally, the optimal (f, g) solve

$$\sigma_w^* = \sigma_a^* \iff \frac{h\left\{2(\tau - b)(\Delta + f) + (L - g)\mu\left[R + h(\Delta + f)(1 - \sigma_w^*)^2\right]\right\}(1 - \sigma_w^*)}{2\left((\tau - b) + h(L - g)\mu(1 - \sigma_w^*)^2\right)} = k\sigma_w^*. \tag{18}$$

This implies an optimal fine, $f_a^*(g)$. Unlike the pricing regime, f and g are strategic complements. Compensation reduces consumers' losses and makes them less responsive to firms' investments in security; a bigger fine is then needed to restore firms' incentives to invest.

Because $f_a^*(g)$ is increasing, one can also remark that, unlike the previous case, the optimal fine is never punitive. Indeed, for any $g \leq L$, $f_a^*(g) \leq L/2$.

Results regarding the optimal liability regime are summarized in the next proposition:

Proposition 3. *Under both the pricing and the advertising regimes, there exists $G \subseteq [0, L]$ such that:*

$$\forall g \in G, \quad \exists f^*(g) \geq 0 \text{ s.t. } \{f^*(g), g\} \text{ implements } \sigma = \sigma_w^*.$$

In pricing regime, f and g are strategic substitutes ($f^{*'}(g) < 0$). The optimal fine is always punitive, i.e. $f^*(g) > L/2$ for all $g \in G$.

In the advertising regime, f^* and g^* are strategic complements ($f^{*'}(g) > 0$). The optimal fine is non-punitive: $f^*(g) \leq L/2$ for all $g \in G$.

In the pricing regime, $G = [0, L]$, which means that the social optimum can be achieved using fines only. This may be relevant in contexts where a compensation scheme might be costly to administer. In the advertising regime, on the other hand, G may take the form $[\underline{g}, L]$, with $\underline{g} > 0$, depending on the parameters of the model. This implies that fines alone may be insufficient to achieve the socially optimal investment level. In particular, when there is over-investment in equilibrium, setting $g = 0$ would require a negative fine in case of a breach in order to achieve the efficient outcome.

Given the assumptions of symmetry and perfect information, it is natural that we can find

$\{f, g\}$ pairs that induce efficient levels of investment. Interestingly, the qualitative features of the optimal schedules differ across the two classes of business models, a property that do not seem to hinge on our specific functional-form assumptions. Indeed, the important feature of the model is the existence of a strategic effect, whereby under-investment in security softens price-competition.

5 Consumer self-protection

Beyond relying on firms to invest in sufficient security, consumers may take some protective measures themselves. Such measures may include storing more sensitive data elsewhere, encrypting data, checking regularly for breaches, or insuring against loss. In this section we study equilibrium in which both consumers and firms can invest in security.

To incorporate this possibility in the model, we assume that, in the first stage (i.e. at the same time firms choose σ), savvy consumers can incur effort e to reduce the loss they incur to $L(e)$, such that $L'(e) < 0$, $L''(e) > 0$, and $\lim_{e\to\infty} L(e) \geq 0$. Some kinds of protection (e.g., insurance) may leave hackers' incentives relatively unchanged, while others (e.g., encrypting stored data) reduce the payoff to a successful breach by preventing the hackers from using some of the stolen data. Formally, we assume that, if a firm's savvy consumers choose e on average, the expected gain from hacking is $h(e\mu) \equiv 1-\gamma e\mu$, where $\gamma \geq 0$ measures the extent to which consumers' effort reduces hackers' payoff as well as their own loss.

Suppose that a savvy consumer expects firms to play σ and other savvy consumers to play $\hat{e}$. His surplus if he plays e equals

$$S(e, \hat{e}, \sigma) = V - \frac{t}{4} + \frac{b}{2} - p - L(e)\frac{h(\hat{e}\mu)}{2}(1-\sigma)^2 - e, \tag{19}$$

where $p = 0$ in the advertising regime. Notice that a single consumer cannot affect the average payoff to hacking, which is why h depends on $\hat{e}$ and not on e.

Expression (19) reveals two features of consumer investment in this model. First, invest-

ment exerts a positive externality on other consumers, as the security risk decreases with the level of consumer self-protection $\hat{e}$. Second, consumers' efforts are strategic substitutes: as other consumers invest more, a consumer faces less risk, and thus has a lower incentive to invest himself. In order to focus on equilibria with a positive level of consumer protection, we assume that $L'(0) = -\infty$, which ensures that $\frac{\partial S(0,0,\sigma)}{\partial e} > 0$ for any $\sigma < 1$. Because $\frac{\partial^2 S(e,\hat{e},\sigma)}{\partial e \partial \hat{e}} < 0$, there exists a unique fixed point $\hat{e}(\sigma)$ which maximizes $S(e, \hat{e}(\sigma), \sigma)$. One can readily check that $\hat{e}(\sigma)$ is downward sloping: investment in security by the firm crowds-out consumer effort.

Let $\pi(\sigma, \hat{\sigma}, e)$ be the profit of a firm that plays σ while its rival plays $\hat{\sigma}$ and consumers play e. This profit is obtained from the analysis of Section 3. Let $\hat{\sigma}_p(e)$ and $\hat{\sigma}_a(e)$ be respectively the equilibrium choice of firms in the pricing and advertising regimes when consumers' effort is e (given by (12) and (16) where we replace L by $L(e)$ and h by $h(e\mu)$). Whereas investment by firms unambiguously reduces consumers' incentive to invest, the slope of $\hat{\sigma}_p(e)$ is ambiguous: an increase in e leads to a simultaneous decrease in L and h, which have opposite effects on σ (by Proposition 1). In the advertising regime, the slope of $\hat{\sigma}_a(e)$ is negative, as both L and h induce firms to invest more (Proposition 2).

An interior equilibrium is then given by a pair, (e^*, σ^*), such that

$$e^* = \hat{e}(\sigma^*) \text{ and } \sigma^* = \hat{\sigma}(e^*)$$

(see Figure 3a). Changing a parameter causes one or both curves (and hence the equilibrium point) to shift as in Figure 3b. Applying standard comparative statics techniques to this equilibrium system yields, for any parameter $z \in \{\mu, \Delta, k, t, b\}$,

$$\frac{\partial \sigma^*}{\partial z} = \frac{\frac{\partial \hat{\sigma}}{\partial z} + \frac{\partial \hat{\sigma}}{\partial e}\frac{\partial \hat{e}}{\partial z}}{1 - \frac{\partial \hat{\sigma}}{\partial e}\frac{\partial \hat{e}}{\partial \sigma}}, \tag{20}$$

with a symmetric expression for $\partial e^*/\partial z$. Moreover, a necessary condition for the equilibrium to be stable is $|\hat{\sigma}'(e)||\hat{e}'(\sigma)| < 1$, implying the denominator of (20) is positive; the sign is then

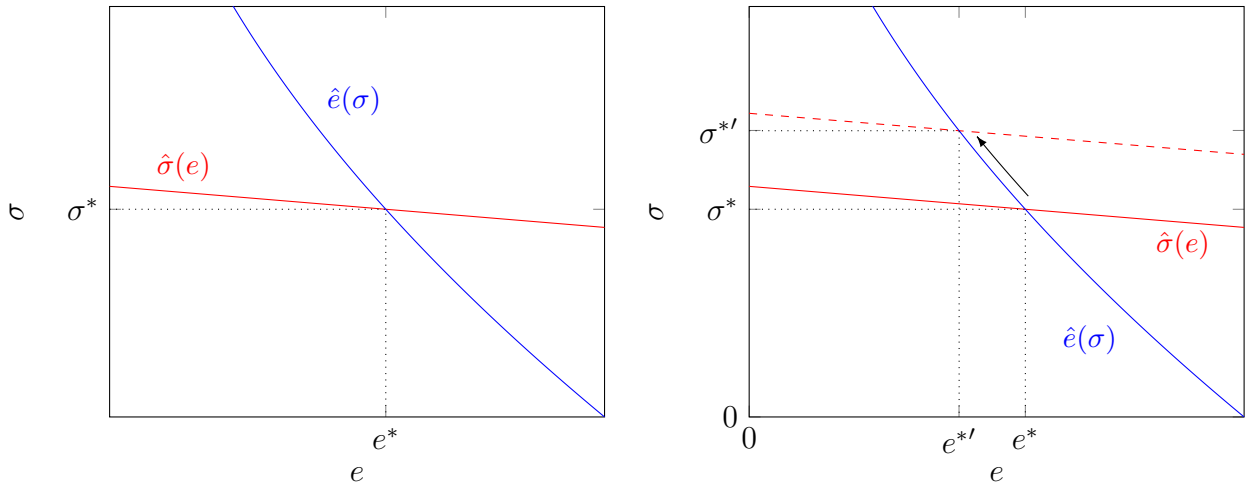

Figure 3: (a) Equilibrium is found where $\hat{\sigma}(e)$ and $\hat{e}(\sigma)$ intersect. (b) Effect of an increase in Δ (which causes $\hat{\sigma}(e)$ to increase).

given by that of the numerator. The following proposition describes the comparative statics.

Proposition 4. *In a stable interior equilibrium of the game with consumer investment: (i) The signs of* $\frac{d\sigma^*}{d\Delta}$, $\frac{d\sigma^*}{dk}$, $\frac{d\sigma^*}{dt}$, *and* $\frac{d\sigma^*}{db}$ *are the same as in the baseline model. (ii) The sign of* $\frac{d\sigma^*}{d\mu}$ *is the same as in the baseline model in the pricing regime. (iii) The sign of* $\frac{d\sigma^*}{d\mu}$ *is the same as in the baseline model in the advertising regime if* γ *is sufficiently small.*

The baseline comparative statics results are robust except as detailed in part (iii) of Proposition 4. In the advertising regime, adding savvy consumers (who invest) may reduce firm effort simply because fewer hackers are active and the firm feels less need to protect itself. This is especially true if γ is large (i.e., when consumer investment in security quickly reduces the payoff to hacking).

The fact that $\hat{e}'(\sigma) < 0$ also has some policy implications. Although a fine increases $\hat{\sigma}$, it also crowds-out consumer effort, blunting the effect on overall security. Write $e^*(f)$ and $\sigma^*(f)$ for the equilibrium when a fine increases firms' perceived damages from Δ to $\Delta + f$. The aggregate damage from all breaches is then

$$\delta(f) = [\Delta + (1-\mu)L(0) + \mu L(e^*(f))] \times [1 - \gamma\mu e^*(f)] \times \frac{1}{2}[1 - \sigma^*(f)]^2,$$

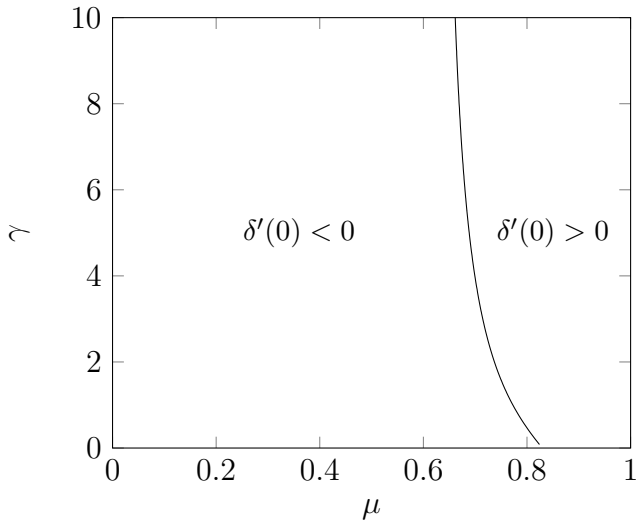

Figure 4: Effect of introducing a small fine on total security breach damages in the pricing regime when $k = 1$, $\Delta = 1/2$, and $L(e) = 2(1 - e)^2$.

where the first set of square brackets enclose the damage per-breach and the second two terms measure the number of successful attacks. Figure 4 shows, for a particular $L(e)$, how the introduction of a small fine affects these damages in the pricing regime. If μ is not too large then the dominant effect of a fine is to increase firm investment and, much like our baseline model, damage is reduced. If, on the other hand, most consumers are sophisticated then a fine crowds-out consumer investment to such an extent that total damages increase.

6 Conclusion

The issue of information security has rapidly climbed the strategic and public policy agenda as digitization not only expands the technological frontier, but also creates new kinds of security threat for homes and businesses. Consumers entrust firms with their personal data and financial affairs, while emerging technologies such as the Internet of things expose their physical environment to cybersecurity threats. When a consumer hands their credit card number to an e-commerce firm, a parent installs a "smart" baby monitor, or a business stores its research data in the cloud, they all depend on the firms providing this technology to have invested sufficient effort in ensuring its security. Those investment decisions take

place in the context of a market and this paper has addressed the natural question of how market competition affects the strategic incentives to undertake such effort.

In order to do this, we constructed a model with the following key features: (i) firms invest in the security of their products; (ii) hackers choose whether to attack, based on the rewards to a successful attack (which depend on the number of users compromised), and the likelihood of success; and (iii) some users are more informed than others about the security of different products. Externalities are pervasive in this kind of market and the socially optimal level of security investment accounts for the harm that attacks impose on both firms and their customers. A recurring theme throughout the paper is that the prevailing business model significantly and qualitatively affects the level of investment, how investment strategically responds to changes in the environment, and the relevant policy prescription. In markets where firms compete in prices, we find that investment is below the efficient level, especially when attacks cause a lot of damage for consumers. This happens not only because firms fail to internalize consumers' losses, but also because of a novel strategic effect whereby investment in security intensifies price competition. This effect can be so strong that an increase in competition leads to lower investment. We contrast this to an alternative business model where firms are ad funded. Here we find that firms use security as an instrument to compete for security-savvy customers so competition increases investment. Indeed, we may even witness over-investment when firms are ad funded because of the business-stealing effect of investment.

The model also yields implications for other settings. For example, consider competition between social media platform firms. We can reinterpret the hackers as agents who post misleading or harmful content to the platform (such as political disinformation) in the hope of reaching as many users as possible, and firms' investment in security as their effort to identify and remove such harmful content. The model can then be used to study firms' equilibrium incentives to invest in content moderation. For example, competition does not generally lead to socially optimal levels of investment and, moreover, a switch to a subscription business

model might make the situation worse rather than better. In this interpretation, the model is related to the work of Liu, Yildirim, and Zhang (2022), albeit with a focus on competitive strategy rather than the situations of monopoly considered by those authors.

Given that externalities, business stealing, and the strategic effect via prices all generically lead to market failure, we investigate the potential for regulatory interventions to restore efficiency. In the pricing regime, the planner's solution can be achieved with an appropriately chosen fine. This fine must be punitive in order to offset the strategic effect as well as inducing firms to internalize consumers' losses. In the ad-funded regime, fines along may not suffice to align incentives, meaning the optimal policy mix sometimes includes a degree of insurance for consumers.

Lastly, we study how consumers' efforts to mitigate the losses from any attack interact with firms' investments. We observe a crowding-out effect whereby consumers exert less effort if firms invest more in security. This blunts the efficacy of policy interventions designed to reduce the damages from cybercrime by inducing firm investment. Indeed, consumers' response can sometimes be so strong that a policy intervention like a fine would lead to higher social damages from security breaches.

7 Funding and Competing Interests

Author A received financial support from from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program). Author B received financial support from the Digital Economics Research Network.

References

Acemoglu, Daron, Azarakhsh Malekian, and Asu Ozdaglar (2016). "Network Security and Contagion". *Journal of Economic Theory* 166, pp. 536–585.

Anderson, Ross and Tyler Moore (2006). "The Economics of Information Security". *Science* 314.5799, pp. 610–613.

Anderson, Simon P. and Stephen Coate (2005). "Market Provision of Broadcasting: A Welfare Analysis". *The Review of Economic Studies* 72.4, pp. 947–972.

Arce, Daniel G (2018). "Malware and market share". *Journal of Cybersecurity* 4.1, tyy010.

— (2020). "Cybersecurity and platform competition in the cloud". *Computers & Security* 93, p. 101774.

Arora, Ashish et al. (2010). "Competition and patching of security vulnerabilities: An empirical analysis". *Information Economics and Policy* 22.2, pp. 164–177.

August, Terrence, Marius Florin Niculescu, and Hyoduk Shin (2014). "Cloud implications on software network structure and security risks". *Information Systems Research* 25.3, pp. 489–510.

August, Terrence and Tunay I Tunca (2006). "Network software security and user incentives". *Management Science* 52.11, pp. 1703–1720.

Becker, Gary S. (1968). "Crime and Punishment: An Economic Approach". *Journal of Political Economy* 76.2, pp. 169–217.

Cavusoglu, Huseyin, Birendra Mishra, and Srinivasan Raghunathan (2004). "The effect of internet security breach announcements on market value: Capital market reactions for breached firms and internet security developers". *International Journal of Electronic Commerce* 9.1, pp. 70–104.

Choi, Jay Pil, Chaim Fershtman, and Neil Gandal (2010). "Network security: Vulnerabilities and disclosure policy". *The Journal of Industrial Economics* 58.4, pp. 868–894.

Dey, Debabrata, Atanu Lahiri, and Guoying Zhang (2012). "Hacker Behavior, Network Effects, and the Security Software Market". *Journal of Management Information Systems* 29.2, pp. 77–108.

Dziubiński, Marcin and Sanjeev Goyal (2017). "How do you defend a network?" *Theoretical Economics* 12.1, pp. 331–376.

Fabrizi, Simona, Steffen Lippert, and José A. Rodrigues-Neto (2019). "Attack, Defense, and the Market for Protection". *Working Paper*.

Fainmesser, Itay P., Andrea Galeotti, and Ruslan Momot (2020). "Digital Privacy". *Working Paper*.

Fedele, Alessandro and Cristian Roner (2020). "Dangerous Games: A Literature Review on Cybersecurity Investments". *BEMPS-Bozen Economics & Management Paper Series* BEMPS75.

Gal-Or, Esther and Anindya Ghose (2005). "The economic incentives for sharing security information". *Information Systems Research* 16.2, pp. 186–208.

Garcia, Alfredo and Barry Horowitz (2007). "The potential for underinvestment in internet security: implications for regulatory policy". *Journal of Regulatory Economics* 31.1, pp. 37–55.

Geer, Dan, Eric Jardine, and Eireann Leverett (2020). "On market concentration and cybersecurity risk". *Journal of Cyber Policy* 5.1, pp. 9–29.

Gordon, Lawrence A. and Martin P. Loeb (2002). "The Economics of Information Security Investment". *ACM Transactions on Information and System Security* 5.4, pp. 438–457.

Gordon, Lawrence A, Martin P Loeb, and William Lucyshyn (2003). "Sharing information on computer systems security: An economic analysis". *Journal of Accounting and Public Policy* 22.6, pp. 461–485.

Goyal, Sanjeev and Adrien Vigier (2014). "Attack, Defence, and Contagion in Networks". *The Review of Economic Studies* 81.4, pp. 1518–1542.

Hirschleifer, Jack (1983). "From Weakest-Link to Best-Shot: The Voluntary Provision of Public Goods". *Public Choice* 41.3, pp. 371–386.

Jo, Arrah-Marie (2019). "The effect of competition intensity on software security – An empirical analysis of security patch release on the web browser market". *Working Paper*.

Johari, Ramesh, Gabriel Y. Weintraub, and Benjamin Van Roy (2010). "Investment and Market Structure in Industries with Congestion". *Operations Research* 58.5, pp. 1303–1317.

Jullien, Bruno, Yassine Lefouili, and Michael H Riordan (2020). "Privacy Protection, Security, and Consumer Retention". *TSE Working Paper*.

Lam, Wing Man Wynne (2016). "Attack-Prevention and Damage-Control Investments in Cybersecurity". *Information Economics and Policy* 37, pp. 42–51.

Liu, Yi, Pinar Yildirim, and Z John Zhang (2022). "Implications of revenue models and technology for content moderation strategies". *Marketing Science*.

Moore, Tyler, Richard Clayton, and Ross Anderson (2009). "The Economics of Online Crime". *Journal of Economic Perspectives* 23.3, pp. 3–20.

Nagler, Matthew G. (2011). "Negative Externalities, Competition AND Consumer Choice". *The Journal of Industrial Economics* 59.3, pp. 396–421. URL: `http://www.jstor.org/stable/41289460` (visited on 10/30/2023).

Navon, Ami, Oz Shy, and Jacques-François Thisse (1995). "Product Differentiation in the Presence of Positive and Negative Network Effects". *CEPR Discussion Paper No.1306*.

Pierce, Justin C. (2016). "Shifting Data Breach Liability: A Congressional Approach". *William & Mary Law Review* 53.3, pp. 975–1017.

Ransbotham, Sam and Sabyasachi Mitra (2009). "Choice and chance: A conceptual model of paths to information security compromise". *Information Systems Research* 20.1, pp. 121–139.

Schechter, Stuart E. and Michael D. Smith (2003). "How Much Security is Enough to Stop a Thief? The Economics of Outsider Theft via Computer Systems and Networks". *in Financial Cryptography*. Springer-Verlag, pp. 122–137.

Toh, Ying Lei (2017). "Incentivizing Firms to Protect Consumer Data: Can Reputation Play a (Bigger) Role?" *Working Paper*.

Varian, Hal (2004). "System Reliability and Free Riding". *Working Paper*.

A Proofs

Proof of Lemma 2. We have

$$\frac{\partial p_i^*}{\partial \sigma_i} = -\frac{2}{3}h(2\Delta + L\mu)(1-\sigma_i) < 0, \ \frac{\partial p_i^*}{\partial \sigma_j} = -\frac{2}{3}h(\Delta + 2L\mu)(1-\sigma_j) < 0. \qquad \blacksquare$$

Proof of Proposition 2. Let $\psi \equiv \frac{\partial \pi_i}{\partial \sigma_i}\Big|_{\sigma_i=\sigma_j=\sigma}$ (the left-hand side of (16)). It is easily checked that $\frac{\partial^2 \pi_i}{\partial \sigma_i^2} > \frac{\partial \psi}{\partial \sigma}$, meaning $\frac{\partial^2 \pi_i}{\partial \sigma_i^2} < 0 \implies \frac{\partial \psi}{\partial \sigma} < 0$. Now, using standard comparative statics methods along with $\frac{\partial \psi}{\partial \sigma} < 0$, we have

$$\operatorname{sgn}\frac{\partial \sigma}{\partial t} = -\operatorname{sgn}\frac{\frac{\partial \psi}{\partial t}}{\frac{\partial \psi}{\partial \sigma}} = \operatorname{sgn}\frac{\partial \psi}{\partial t} = \operatorname{sgn}\left(-\frac{hL\mu\left[R - h\Delta(1-\sigma)^2\right](1-\sigma)}{2\left[\tau - b + hL\mu(1-\sigma)^2\right]^2}\right) < 0.$$

Comparative statics with respect to the other parameters are obtained analogously.

Since $\frac{\partial \psi}{\partial \sigma} < 0$, we have over-investment in equilibrium if the left-hand side of (16) is positive at $\sigma = \sigma_w^*$. Substituting $\sigma = \sigma_w^*$ into (16) and noting that $k\sigma_w^* = h\left(\frac{L}{2} + \Delta\right)(1-\sigma_w^*)$ (from equation 4), the left-hand side of (16) becomes

$$\frac{hL\left[\mu\left(R - h(L+\Delta)(1-\sigma_w^*)^2\right) - \tau + b\right](1-\sigma_w^*)}{2\left(\tau - b + hL\mu(1-\sigma_w^*)^2\right)}.$$

This is positive when (17) is satisfied. $\blacksquare$

Proof of Proposition 4. (i) For $z \in \{\Delta, k, t, b\}$ we have from (19) that $\frac{\partial \hat{e}}{\partial z} = 0$. From (20), therefore, $\frac{d\sigma^*}{dz}$ has the same sign as $\frac{\partial \hat{\sigma}}{\partial z}$, which is just the equilibrium effect of Section 3.

(ii) Applying standard comparative statics methods to (19) yields

$$\frac{\partial \hat{e}}{\partial \mu} = -\frac{\hat{e}h'(\hat{e}\mu)L'(\hat{e})}{\mu h'(\hat{e}\mu)L'(\hat{e}) + h(\hat{e}\mu)L''(\hat{e})}.$$

Moreover, from (12) (and suppressing the arguments for readability),

$$\frac{\partial \hat{\sigma}}{\partial \mu} = -\frac{6k\left[hL - e(4\Delta - \mu L)h'\right]}{\left[6k + h(4\Delta - \mu L)\right]^2}, \ \frac{\partial \hat{\sigma}}{\partial e} = \frac{6k\mu\left[(4\Delta - \mu L)h' - hL'\right]}{\left[6k + h(4\Delta - \mu L)\right]^2}.$$

With these ingredients, calculating the numerator of (20) yields

$$\left(\frac{\partial\hat{\sigma}}{\partial z}+\frac{\partial\hat{\sigma}}{\partial e}\frac{\partial\hat{e}}{\partial z}\right)=\frac{6kh\left[\mu h'L'\left(eL'-L\right)-\left(hL-e(4\Delta-\mu L)h'\right)L''\right]}{\left(6k+h(4\Delta-\mu L)\right)^2\left(\mu h'L'+hL''\right)}.$$

Given $L'(e)<0$, $L''(e)>0$, and $h'(e\mu)\leq 0$, the denominator of the right-hand side is positive and the numerator is negative.

(iii) γ (and hence $h'(e\mu)$) small implies $\frac{\partial\hat{e}}{\partial\mu}$ is small, so that $\frac{d\sigma^*}{d\mu}$ has the same sign as $\frac{\partial\hat{\sigma}}{\partial\mu}$. Moreover, because $h'(e\mu)$ is small, the sign of $\frac{\partial\hat{\sigma}}{\partial\mu}$ is the same as in Section 3. ■

B Market structure and the strategic effect

B.1 Monopoly

B.1.1 Pricing regime

Suppose only a single firm was active, located at one end of the Hotelling line and setting price p and security level σ. The firm serves every consumer if $p\leq\underline{p}=V-\tau+b-h(1-\sigma)^2L$ and never has a reason to charge a price less than this. To keep the analysis simple, assume that V is large enough that the market is covered.[28]

Substituting price $\underline{p}$ into the firm's profit, (2), we obtain its first-order condition, $\frac{\partial\pi}{\partial\sigma}=0$, which coincides exactly with (4) (evaluated at $n^*=1$). Thus, the monopolist implements the efficient level of investment. This is in contrast to Proposition 1, where investment was below the efficient level.

B.1.2 Advertising regime

Now consider a monopolist in the advertising regime. The firm's profit, if serving n consumers, is $\pi=R-h(1-\sigma)^2\Delta n-\frac{k\sigma^2}{2}$. The associated first-order condition is $\frac{\partial\pi}{\partial\sigma}=$

[28] As usual, the marginal effect of an increase in p on profit is $n(p)+[p-c]n'(p)$, where n is demand, $c=\Delta(1-\sigma)h$ is the firm's effective marginal cost, and n' is the right-hand derivative if n exhibits a kink. If V (and hence $\underline{p}$) is large then $[p-c]n'(p)$ is large and non-positive for all $p\geq\underline{p}$, meaning the firm does not wish to increase price above $\underline{p}$.

$2h(1-\sigma)\Delta n - k\sigma = 0$. Comparison with (4) reveals that the monopolist under-provides security relative to the efficient level. This is because the firm has no way to extract the value of security to consumers and therefore fails to internalize consumers' losses from security breaches.

Moreover, holding the scale of operations fixed, the monopolist invests less than a duopolist. Indeed, if we normalize the size of the market served by the monopolist to $n = 1/2$,[29] the monopolist's marginal return to investing is $\frac{\partial \pi}{\partial \sigma} = h(1-\sigma)\Delta - k\sigma$. This is less than the left-hand side of (16): competition forces firms to invest more to avoid losing savvy consumers to a rival.

B.2 Oligopoly concentration

This section extends the model to incorporates a third firm and thereby study the role of the strategic effect (Lemma 2) under a wider variety of market structures. Recall that the strategic effect leads firms to invest less as μ increases (because investments then more strongly intensify price competition).

We revise the model as follows: suppose there are three firms, $i \in \{1, 2, 3\}$. Between each pair of firms is a Hotelling segment of length 1. The segment between firms 1 and 2 has uniformly distributed mass $m \in [0, 1]$ of consumers, while the two segments between 3 and its rivals each have mass $(1-m)/2$. If $m = 1/3$ then all three firms are ex ante symmetric. If $m < 1/3$ then there is a single dominant firm (firm 3), whereas $m > 1/3$ corresponds to a market structure where firm 3 is smaller than its two rivals. Each firm chooses a single security level, σ_i, followed by a single price, p_i. To keep things simple, we set $b = 0$.

Let n_{ij} be the share of consumers on the segment that connects firms i and j who choose firm i, and $M_{ij} \in \{m, \frac{1-m}{2}\}$ be the total mass of consumers on that segment. A sophisticated

[29] If we let $n = 1$ then the monopolist has an extra incentive to invest compared to duopolists because serving twice as many consumers makes it a more attractive target for hackers. It is then possible that the monopolist might invest more.

consumer is indifferent if they are located at x_{ij} solving

$$V - x_{ij}\tau - p_i - h(1-\sigma_i)^2 M_{ij} n_{ij} L = V - (1-x_{ij})\tau - p_j - h(1-\sigma_j)^2 M_{ij}(1-n_{ij})L,$$

i.e.,

$$x_{ij} = \frac{\tau - p_i + p_j}{2\tau} - \frac{hLM_{ij}}{2\tau}\left(n_{ij}(1-\sigma_i)^2 - (1-n_{ij})(1-\sigma_j)^2\right).$$

A consumer with $x < x_{ij}$ prefers i. Unsophisticated consumers are indifferent if located at y_{ij} solving $V - y_{ij}\tau - p_i = V - (1-y_{ij})\tau - p_j$, i.e.,

$$y_{ij} = \frac{\tau - p_i + p_j}{2\tau}.$$

A consumer with $y < y_{ij}$ prefers i. We then have

$$n_{ij} = \mu x_{ij} + (1-\mu) y_{ij} = \frac{p_j - p_i + \tau + hLM_{ij}\mu\left(1-\sigma_j\right)^2}{2\tau + hLM_{ij}\mu\left(2-(2-\sigma_i)\sigma_i - (2-\sigma_j)\sigma_j\right)}.$$

Lastly, firm 1's demand can be found as $N_1 = mn_{12} + \frac{1-m}{2}n_{13}$, firm 2's demand is $N_2 = mn_{21} + \frac{1-m}{2}n_{23}$, and firm 3's demand is $N_3 = \frac{1-m}{2}(n_{31} + n_{32})$. Given these demands, we are in a position to write firm i's profits as $\pi_i = [p_i - h(1-\sigma_i)^2\Delta]N_i - \frac{k\sigma_i^2}{2}$. From here we follow the analogous steps to those found in Section 3.2: the system of first-order conditions

$$\left\{\frac{\partial \pi_i}{\partial p_i} = 0\right\}_{i\in\{1,2,3\}}$$

can be solved analytically for the equilibrium prices, $p_i^*(\sigma_1, \sigma_2, \sigma_3)$. Substituting these prices into π_i, we can then solve the system

$$\left\{\frac{\partial \pi_i}{\partial \sigma_i} = 0\right\}_{i\in\{1,2,3\}}$$

for σ_i^*. Because of the asymmetry when $m \neq 1/3$, σ_i^* must be computed numerically. Figure

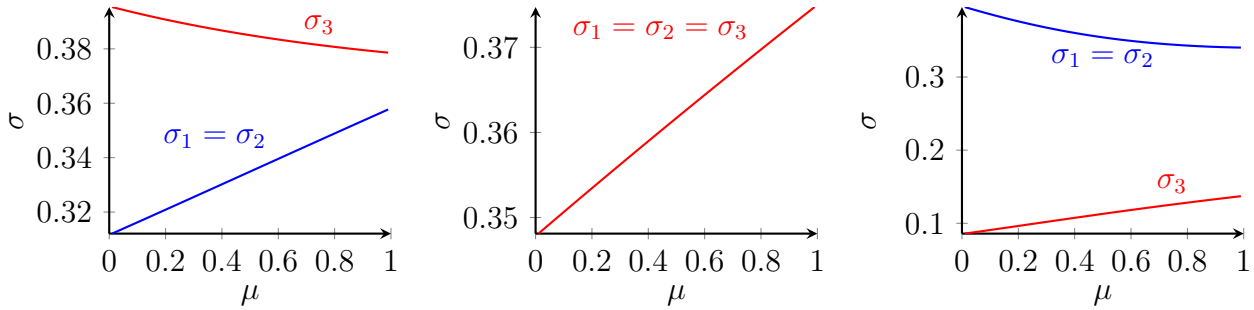

(a) $m = 1/8$ (b) $m = 1/3$ (c) $m = 7/8$

Figure 5: Equilibrium security as a function of μ for three different market structures (with $L = 3$, $h = 1$, $\tau = 1$, $\Delta = 1$, and $k = 1$). The strategic effect dominates for a firm if its σ is decreasing in μ.

5 shows these equilibrium security investment levels as a function of μ for three different market structures.

In Figure 5a m is small (firm 3 is dominant). Here we see that the strategic effect dominates for firm 3 (i.e., σ_3 is decreasing in μ). Intuitively, 3 is the most important competitor for both its rivals, so firm 3 is particularly sensitive to the fact that its investment will distort its rivals' pricing incentives. On the other hand, the strategic effect does not dominate for firms 1 and 2 ($\sigma_1 = \sigma_2$ is increasing in μ). This is because firm $i \in \{1, 2\}$ is only half the competition faced by 3. Firm i's investments therefore have a smaller effect on the pricing of its main competitor.

In Figure 5c the roles are reversed and firm 3 is smaller than its rivals. It is now firms 1 and 2 for whom the strategic effect dominates. Firm 3 does not experience a strong strategic effect because its two rivals are too busy competing with each other to be much influenced by the investment of such a small actor in the market.

We can make more explicit the relationship between concentration and security investment using the Herfindahl-Hirschman Index (HHI). The equilibrium σs imply demand $D_i(\sigma_1, \sigma_2, \sigma_3)$. We can then compute the HHI as $\text{HHI} = (D_1)^2 + (D_2)^2 + (D_3)^2$. Figure 6 shows how the average security experienced by a consumer varies with the level of market concentration. Beginning at $m = 1$, consumers consider only firms 1 and 2, which each enjoy a market share

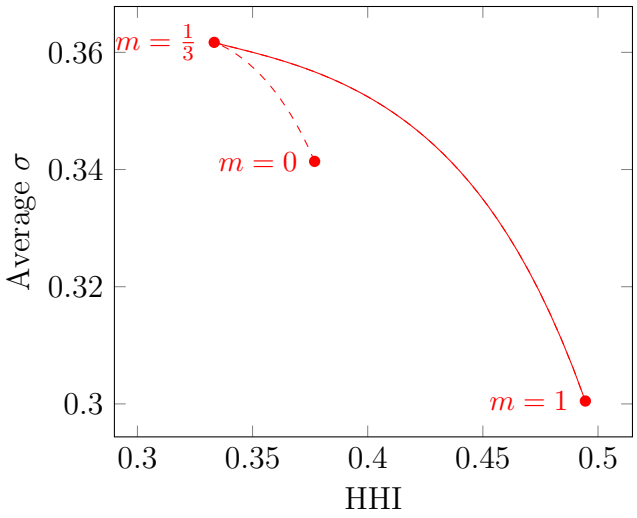

Figure 6: Relationship between market concentration (measured by HHI) and average security investment.

of one half (HHI $= (1/2)^2 + (1/2)^2 = 1/2$). Lowering m reduces market concentration and causes security to increase along the top curve until the HHI reaches its minimum of $1/3$ (at $m = 1/3$). Thereafter, further lowering m causes concentration to increase as firm 3 becomes dominant and security falls along the bottom (dashed) curve. Overall, then, we indeed find that security is decreasing in the level of concentration.