A Model of Information Security and Competition*

Alexandre de Cornière† and Greg Taylor‡

July 27, 2021

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, security plays a business-stealing role, and may be overprovided. In terms of policy, we show that both the structure of the optimal liability regime and the efficacy of certification schemes also depend on firms' business model.

1 Introduction

Cybersecurity has shot to the top of the digital policy agenda on the back of a spate of major security breaches. Within recent months, the Solar Winds attack,[1] a major breach of

*We thank Arrah-Marie Jo, Dann Arce and some anonymous reviewers at WEIS 2021 for their comments. De Cornière acknowledges funding from ANR under grant ANR-17-EURE-0010 (Investissements d'Avenir program).

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

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

[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.

Microsoft Exchange Server,[2] and the Colonial pipleline 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 attackers and defenders responding to incentives (Anderson and Moore, 2006). 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 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. A key feature of the model is therefore the presence of negative network effects, as larger firms are more likely to be under attack, which, if successful, would hurt consumers. 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 compe-

[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.

tition, consumers' awareness of security risks, and firms' business model. Competition is measured by market structure (we contrast monopoly and duopoly) 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*.

To briefly summarize the results, we find that in the pricing regime a monopolist's security investment incentives are aligned with the social planner's provided there are enough sophisticated consumers. In order to sell to sophisticated consumers, the monopolist must set the price of its product so as to compensate them for the expected security risk. This induces the firm to internalize the social value of security. In the advertising regime, on the other hand, a monopolist has no way of extracting the value that sophisticated consumers derive from security, and this leads to systematic under-investment. Things are quite different when there is competition on the market. In the pricing regime, high levels of security reduce the negative network effects due to hacking, thereby intensifying price-competition. Moreover, under competition the marginal consumer is not always sophisticated, as naive consumers compare prices and product (non-security) characteristics. These two forces lead firms to under-invest in equilibrium. 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. Otherwise, firms under-invest. In our baseline model only firms can exert protection efforts. We discuss the robustness of this assumption in Section 6, where we also allow sophisticated consumers to take preventive actions.

Since the competitive provision of security is generically sub-optimal, we also consider regulatiory interventions. Firstly, we 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. Lastly, our contrasting results on the effects of consumer information under different business models imply that certification schemes or other initiatives to increase transparency may be counter-productive

in the pricing regime.

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 provides two useful benchmarks: the social planner's solution and the monopolist's behavior. Section 4 contains the main equilibrium analysis under competition. We study regulatory interventions in Section 5, allow consumers to invest in mitigating the harms from an attack in Section 6, and conclude in Section 7. 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. 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, 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.

A few papers study security investments in an 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. Such a security-related negative

network effect (larger firms attract more hackers) also plays a role in our analysis.

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. Consumers are uniformly distributed along the segment, and the gross utility of a consumer who selects a product at a distance d from his ideal position is $V - td$. The stand-alone value, V, is assumed large enough to ensure the market is covered in equilibrium.

We consider two kinds of business model for firms, depending on whether they generate revenues through *pricing* or *advertising*.[6] In the pricing regime, each firm sets a price p_i, which enters consumers' utility negatively is 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 4.2, 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

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

share of vulnerabilities that are fixed). Fixing vulnerabilities requires investing in security: achieving a given σ_i costs $\frac{k\sigma_i^2}{2}$.[7] 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 per customer.[8]

Thus, if firm i serves n_i consumers, the payoff to attacking i is $h(1-\sigma_i)n_i - c$. It follows that the probability of a *successful* attack ocurring 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 and the probability that each attack succeeds.[9]

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 5).[10] Writing r_i for irm 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 breach probability, (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,

[7] 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.

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

[9] 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 $h(1-\sigma_i)$.

[10] 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.

stemming from the corruption or fraudulent use of data, from privacy violations or identity fraud, or from interrupted access to compromised services.[11] 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 $u_i = V - td - p_i - h(1-\sigma_i)^2 n_i L$ (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 i's security, σ_i, and incorporate the security risk into their decision-making. The remaining $1-\mu$ of 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.[12]

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. 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.[13] Recent examples include the Solar Winds attack[14] and the Microsoft Exchange Server breach,[15] which each exposed millions of users to attacks. A single such vulnerability, if discovered by a hacker, potentially 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

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

[12] 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 Appendix for details.

[13] See Buchanan (2020) and Perlroth (2021) for non-technical accounts of the history of cyber warfare.

[14] `https://en.wikipedia.org/wiki/2020_United_States_federal_government_data_breach`, accessed 20 July 2021.

[15] `https://en.wikipedia.org/wiki/2021_Microsoft_Exchange_Server_data_breach`, accessed 20 July 2021.

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 6.

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, with hackers able to sell knowledge of vulnerabilities in more popular systems for higher prices.[16] More generally, hackers should be thought of as rational actors (see, e.g., Ransbotham and Mitra, 2009; Schechter and Smith, 2003). The likelihood of a firm being targeted increases in the size of the reward to a successful attack and decreases in the barriers to success (i.e., the security countermeasures implemented by the target).[17] 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 or of the latest emerging cyber threats. On the other hand, if the end customer is a business then its IT procurement is likely to be handled by a team of IT experts and to involve site visits by vendors who brief the customer on a product's security features (a key selling point for enterprise IT systems).[19] 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 sophisticated buyers, whereas consumer-facing social-media is more likely to be ad funded and targeted at a less savvy user base.

Lastly, we discuss two features we have deliberately excluded from the model. First, we

[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 25 July 2021.

[17]According to an industry report, around 86% of breaches are motivated by financial gain. See `https://enterprise.verizon.com/en-gb/resources/reports/dbir/`, accessed 11 November 2020.

[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.

[19]Additionally, major providers of IT platforms, such as Amazon Web Services, publish white papers that describe the underlying security measures for the firm's business clients. See, e.g., `https://docs.aws.amazon.com/whitepapers/latest/introduction-aws-security/introduction-aws-security.pdf`, accessed 23 July 2021.

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 Benchmarks: social planner and monopoly

3.1 Efficient investment under duopoly

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.[20] 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*

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

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

[20] We establish this formally in Section 4.1 below.

k, the cost of providing security.

3.2 Monopoly

Before proceeding to the main analysis, it is useful to also consider the benchmark of monopoly, which will help highlight the effect competition has on investment in security. For tractability, we focus on the case where $t = 0$, i.e. where the only dimension of heterogeneity is consumers' awareness of the security risks.[21]

When the monopolist serves all consumers, the efficient level of investment is found from (4) after substituting $n^* = 1$:

$$2h(1-\sigma)(\Delta + L) - k\sigma = 0 \tag{6}$$

Pricing regime A savvy consumer who expects the monopolist to serve n consumers estimates the security risk to be equal to $hL(1-\sigma)^2 n$, and his willingness to pay (given that $t = 0$) is $V - hL(1-\sigma)^2 n$: demand by savvy consumers exhibits negative network effects, a feature we discuss at length in the following sections. In contrast, naive consumers' willingness to pay is simply V. The monopolist therefore has two available strategies: serving all consumers, or only the naive ones.

For a given security level σ, the highest price resulting in full market coverage is $p = V - hL(1-\sigma)^2$. Substituting this price into the firm's profit, (2), we obtain a first-order condition, $\frac{\partial \pi}{\partial \sigma} = 0$, that coincides exactly with (6). Thus, the monopolist implements the efficient level of investment. Intuitively, if the firm reduces the probability of a successful breach by ϵ, savvy consumers' willingness to pay increases by $L\epsilon$. The firm can increase the price by $L\epsilon$ and the naive will pay even though their willingness to pay has not moved (because they are inframarginal). This means that the firm fully internalizes consumers' losses.

Alternatively, if the monopolists decides to price the μ savvy consumers out of the market, its optimal price is $p = V$. Substituting this along with $n = 1 - \mu$ into the firm's profit, (2), the optimal investment solves

$$\frac{\partial \pi}{\partial \sigma} = 2(1-\mu)h(1-\sigma)\Delta - k\sigma = 0.$$

Comparing this with (6) reveals that the monopolist under-invests from a social perspective because (naive) consumers' willingness to pay does not respond to security investment,

[21] When $t > 0$ the demand function exhibits at least two and up to four kinks depending on the parameter values, which makes the analysis very cumbersome for relatively little economic insight. This problem does not occur under duopoly, provided that V is large enough.

meaning the firm does not internalize consumers' losses when choosing σ.

Profit is independent of μ if the firm serves savvy consumers, but is decreasing in μ if it does not. Thus, there exists a $\bar{\mu} \in [0,1]$ such that savvy consumers are served if $\mu > \bar{\mu}$ but not if the inequality is reversed.

Advertising regime The firm's profit is $\pi = R - h(1-\sigma)^2\Delta - \frac{k\sigma^2}{2}$. The associated first-order condition is $\frac{\partial \pi}{\partial \sigma} = 2h(1-\sigma)\Delta - k\sigma = 0$. Comparison with (6) 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.

In summary, we have the following result:

Proposition 1. *Suppose that $t = 0$.*

1. *In the pricing regime there exists a $\bar{\mu} \in [0,1]$ such that a monopolist implements the efficient level of security if $\mu \geq \bar{\mu}$, and under-invests if $\mu < \bar{\mu}$.*

2. *In the advertising regime, a monopolist always invests less than the efficient level.*

4 Equilibrium under duopoly

4.1 Pricing regime

Turning to the case of duopoly, we know that the probability of a successful attack 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.

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

$$V - xt - p_1 - h(1-\sigma_1)^2 n_1 L = V - (1-x)t - p_2 - h(1-\sigma_2)^2 n_2 L.$$

i.e.

$$x = \frac{t - p_1 + p_2}{2t} - \frac{hL}{2t}\left(n_1(1-\sigma_1)^2 - n_2(1-\sigma_2)^2\right) \tag{7}$$

The first term on the right-hand side of (7) is the demand in a standard Hotelling model. The second term shows the presence of negative network externalities in the model: as more

consumers choose firm i, the probability that it becomes a target (for a given σ_i) increases, which makes i less attractive to other consumers.

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

$$V - yt - p_1 = V - (1-y)t - p_2.$$

i.e.

$$y = \frac{t - p_1 + p_2}{2t} \tag{8}$$

Naive consumers do not perceive the greater risk of attack as n_i increases, and there is thus no 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

$$n_1^* = \frac{p_2 - p_1 + t + hL\mu(1-\sigma_2)^2}{2t + hL\mu[(1-\sigma_1)^2 + (1-\sigma_2)^2]}, \quad n_2^* = 1 - n_1^*. \tag{9}$$

Pricing stage Given σ_1 and σ_2, firms choose prices to maximize (2), with $r_i = p_i$ and demand given by (9). Firm i's first-order condition is $\frac{\partial \pi_i}{\partial p_i} = 0$ and solving this system yields the equilibrium prices:

$$p_i^* = t + \frac{1}{3}h\Big\{\Delta\left[3 - 2\left(2-\sigma_i\right)\sigma_i - \left(2-\sigma_j\right)\sigma_j\right] + L\mu\left[3 - \left(2-\sigma_i\right)\sigma_i - 2\left(2-\sigma_j\right)\sigma_j\right]\Big\}. \tag{10}$$

In a standard Hotelling game we would have $p_i^* = t$; 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 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.

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

$$\eta_i = \frac{p_i}{t - p_i + p_j + hL\mu(1-\sigma_j)^2} \tag{11}$$

As firm j increases σ_j, firm i's demand becomes more price-elastic. Indeed, inspection of equation 9 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 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

$$\max_{\sigma_i \geq 0} \left\{ \left[p_i^* - h(1-\sigma_i)^2 \Delta \right] n_i^* - \frac{k\sigma_i^2}{2} \right\}, \tag{12}$$

where p_i^* and n_i^* are respectively given in (10) and (9). Making this substitution, 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^*:

$$\sigma_p^* = \frac{h(4\Delta - L\mu)}{6k + h(4\Delta - L\mu)}. \tag{13}$$

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 (11), 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) the parameter t, does not affect the equilibrium investment in security. This is because of two opposite effects. On one hand, demand for firm i is less sensitive to σ_i as t increases, by (9). On the other hand, the equilibrium price increases with t (see below), which means that

each additional customer attracted by an 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 (13), we find that

$$\sigma_w^* - \sigma_p^* = \frac{2hk(2\Delta + L(3+\mu))}{(2k + h(L+2\Delta))(6k + 4h\Delta - hL\mu)} > 0, \tag{14}$$

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 (13) and (14)).

Proposition 2. *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 t.

As for equilibrium prices, substituting σ_p^* into (10) we obtain:

$$p^* = t + \frac{36hk^2(\Delta + L\mu)}{(6k + 4h\Delta - hL\mu)^2} = t + h\Delta(1-\sigma_p^*)^2 + \frac{36hk^2 L\mu}{(6k + 4h\Delta - hL\mu)^2}. \tag{15}$$

Recall that a standard Hotelling model with marginal costs $h\Delta(1-\sigma_p)^2$ would yield an equilibrium price of $t + h\Delta(1-\sigma_p)^2$. Because of the presence of 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 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.

The facts that $\frac{\partial \sigma_p^*}{\partial \mu} < 0$ and $\frac{\partial \sigma_p^*}{\partial t} = 0$ reinforce our earlier observation that what we call "security" is different from a mere investment in the quality of product i. Indeed, to take a specific example, a Hotelling model with endogenous qualities such that $u_i = q_i - p_i - td$ would have $\frac{\partial q^*}{\partial \mu} > 0$ and $\frac{\partial q^*}{\partial t} < 0$. In the present model, security is best thought of as an investment in reducing the strength of the negative network effect.

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.

Recall from Section 3.2 that whenever a monopolist serves the savvy consumers it chooses the first-best σ.[22] Introducing competition can therefore reduce investment in security (because of the strategic effect described in Lemma 2, which is not active for a monopolist).[23] Starting from duopoly, on the other hand, an additional firm weakens the strategic effect because a change in each firm's investment has a small impact on rivals' pricing when it is just one of many competitors (we verify this intuition in Appendix B). 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. Details can be found in Appendix B.

4.2 Advertising regime

In the advertising regime, demand is given by (9) with $p_i = p_j = 0$. Firms' security investment is chosen to solve

$$\max_{\sigma_i > 0} \left\{ \left[R - h(1-\sigma_i)^2 \Delta \right] n_i^* - \frac{k\sigma_i^2}{2} \right\}, \tag{16}$$

[22] If μ is small then the monopolist does not serve the savvy consumers and invests less than does a duopolist.

[23] It may seem that the monopolist only invests more than a duopolist because of its larger scale. But we can eliminate this scale effect by setting $n = 1/2$ in Section 3.2, and thereby isolate the strategic effect. We still find that the monopolist implements the first-best in this reduced-size market (hence, with higher investment than in duopoly).

where n_i^* is given in (9) (with $p_1 = p_2 = 0$). The symmetric equilibrium σ is implicitly given by evaluating firms' first-order conditions at $\sigma_i = \sigma_j = \sigma$:

$$\frac{h\left\{2t\Delta + L\mu\left[R + h\Delta(1-\sigma)^2\right]\right\}(1-\sigma)}{2\left(t + hL\mu(1-\sigma)^2\right)} - k\sigma = 0. \tag{17}$$

Applying the implicit function theorem to (17) allows us to study how the equilibrium investment 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 3. *In the advertising regime, firms over-invest compared to the socially optimal solution if*

$$\frac{t}{\mu} < R - h(L+\Delta)(1-\sigma_w^*)^2, \tag{18}$$

where σ_w^* is given in (5)*, and under-invest if the inequality is reversed.*

Firms' investment in security is decreasing in t and k; and increasing in μ, L, h, 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 (t is small) because the mark-up is independent of t 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 t 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.

Third, holding the scale of operations fixed, investment is higher than under monopoly. Indeed, if we normalize the size of the market served by the monopolist to $n = 1/2$, the monopolist's marginal return to investing (from Section 3.2) is $\frac{\partial \pi}{\partial \sigma} = h(1-\sigma)\Delta - k\sigma$. This is less than the left-hand side of (17): competition forces firms to invest more to avoid losing savvy consumers to a rival.[24]

[24]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.

We have assumed that R is exogenously fixed. But one might expect R to depend on a firm's investment (e.g., because advertisers prefer to be associated we secure firms). We could easily incorporate this into the model by letting $R'(\sigma_i) \neq 0$. Then (17) becomes

$$\frac{h\left\{2t\Delta + L\mu\left[R + h\Delta(1-\sigma)^2\right]\right\}(1-\sigma)}{2\left(t + hL\mu(1-\sigma)^2\right)} = k\sigma - \frac{1}{2}R'(\sigma).$$

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.

A remark on positive network effects Many IT products exhibit positive network effects, and our framework could easily be amended to include them, provided they are small enough so as to not lead to complete market tipping. Suppose indeed that consumers enjoy network benefits b for each other user of the product they choose. Then one can show that, provided that $b < t$, the equilibrium in both the pricing and the advertising regimes is the same as one in which there are no network benefits but the transportation parameter is $\tilde{t} \equiv t - b$. On implication of this observation is that comparative statics results with respect to b are of the opposite sign to those related to t.

5 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 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.[25]

We studied the relationship of a regulated security standard in the previous section. In this section we use our model to shed light on two policy tools, namely the optimal liability regime and a certification scheme.

5.1 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

$$\sigma_w^* = \sigma_p^* \iff \frac{h(L+2\Delta)}{2k+hL+2h\Delta} = \frac{h[4(\Delta+f)-(L-g)\mu]}{6k+4h(\Delta+f)-h(L-g)\mu}. \tag{19}$$

There is therefore a continuum of (f, g) pairs that implement the planner's solution, with the optimal fine being

$$f_p^*(g) = \frac{1}{4}[2\Delta - g\mu + L(3+\mu)]. \tag{20}$$

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 $p_i/\left(t - p_i + p_j + h(L-g)(1-\sigma_j)^2\right)$, 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. Notice that $f_p^*(g) > L/2$ implies that the optimal liability regime generates

[25] 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.

a budget surplus, which can be redistributed through lump-sum payments to consumers or firms.

Replacing Δ by $\Delta + f^*(g)$ and L by $L - g$ in (15), we find that the equilibrium price is a decreasing function of g. Among all the efficient policies $\{f^*(g), g\}$, the lowest price is therefore achieved when $g = L$. Full insurance along with a punitive fine maximizes consumer surplus over all efficient policies.

Advertising regime The first order condition determining equilibrium investment is given in (17). Suppose we implement a budget-balanced policy that fully-compensates consumers ($g = L$ and $f = L/2$). Making this substitution, (17) 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.

As in the case of price competition, there are multiple ways to implement the planner's desired level of investment. Indeed, any (f, g) such that (17) holds at $\sigma = \sigma_w^*$ works. In other words, the optimal (f, g) solve

$$\frac{h\left\{2t(\Delta + f) + (L - g)\mu\left[R + h(\Delta + f)(1 - \sigma_w^*)^2\right]\right\}(1 - \sigma_w^*)}{2\left(t + h(L - g)\mu(1 - \sigma_w^*)^2\right)} - k\sigma_w^* = 0. \tag{21}$$

Unlike the previous case, though, the relationship between f and g is one of strategic complementarity. To see this simply, consider a pair (f, g) that implements σ_w^*. Suppose that we increase g. By the comparative statics of Proposition 3, where we replace L by $L - g$, we know that σ_a^* is a decreasing function of g, so that firms would react to $dg > 0$ by reducing their investment. In order to offset this and stay on the efficient investment locus, the regulator needs to increase the fine f (because, still by Proposition 3 where we replace Δ by $\Delta + f$, σ_a^* is an increasing function of f).

Because $f_a^*(g)$ is increasing, one can also remark that, unlike the previous case, the optimal fine is never punitive in the sense that it is never larger than the damage suffered by consumers. Indeed, for any $g \leq L$, $f_a^*(g) \leq f_a^*(L) = L/2$.

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

Proposition 4. *Under both the pricing and the advertising regimes, there is a continuum of $\{f, g\}$ pairs that implement the efficient level of investment. Formally, there exists $G \subseteq [0, L]$*

such that:

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

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 $g \in [0, L]$.

In the advertising regime, f^* and g^* are strategic complements ($f^{*'}(g) > 0$). The optimal fine is not punitive: $f^*(g) \leq L/2$ for all $g \in [0, L]$.

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 these assumptions. Indeed, the important feature of the model is the existence of a strategic effect, whereby under-investment in security softens price-competition. For this effect to matter, the industry needs to be concentrated enough, and firms need to be able to observe (or infer) the level of security offered by their rivals. If these conditions do not hold, we should expect the optimal liability regime to involve non-punitive fines and to exhibit strategic complementarity between fines and the level of compensation.

5.2 Certification

Another policy instrument at the disposal of regulators is the use of a certification scheme, whereby an independent entity would evaluate the security level of firms, and publicize the results. In the EU, for example, the Cybersecurity Act of 2019 established a cybersecurity certification framework, where the requirements are tailored to specific products or businesses, and where several levels of security can be certified (basic, substantial, high). One important consequence of a certification scheme is that it allows consumers to observe the security level of firms more easily. A natural way to incorporate this policy into our model is to model it as an increase in the share of sophisticated consumers μ.

We have the following result:

Proposition 5. *(i) In the pricing regime, a certification scheme lowers the equilibrium security level. When coupled with an optimal liability regime, a certification scheme requires a larger fine in case of a breach.*

(ii) In the advertising regime, a certification scheme increases the equilibrium security level. When coupled with an optimal liability regime, a certification scheme requires a smaller fine in case of a breach.

Proposition 5 is a corollary of Propositions 2 and 3 (regarding the effect of an increase in μ), and of Equations 20 and 21 (regarding the link with the optimal liability regime).

Part (ii) of Proposition 5 is probably the result that corresponds to the common intuition regarding certification: by making security more transparent, certification enables consumers to compare offers along this dimension, which leads firms to invest more. Notice though that, even in this regime, certification is not necessarily optimal if we start from a situation where μ is already large enough so that there is over-investment in equilibrium. Interestingly, in such a regime, certification is a substitute to a fine: regulators can therefore focus on one instrument and achieve a large part of the gains from regulation.

Part (i), however, is a cautionary tale, as it highlights a potential drawback from more transparency. Indeed, the existence of the strategic effect implies that firms under-invest as μ increases, so as to soften competition. In order to offset this effect, the regulator would need to increase the fine imposed on firms in case of a breach.

6 Consumer self-protection

Beyond relying on firms to invest in sufficient security, consumers may take some protecting 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 consumers to play $\hat{e}$.

His surplus if he plays e equals

$$S(e,\hat{e},\sigma) = V - \frac{t}{4} - p - L(e)\frac{h(\hat{e}\mu)}{2}(1-\sigma)^2 - e, \tag{22}$$

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 (22) reveals two features of consumer investment in this model. First, investment 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 who plays σ while its rival plays $\hat{\sigma}$ and consumers play e. This profit is obtained from the analysis of Section 4. 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 (13) and (17) 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 2). 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 1a). Changing a parameter causes one or both curves (and hence the equilibrium point) to shift as in Figure 1b. Applying standard comparative statics techniques to this equilibrium system yields, for any parameter $z \in \{\mu,\Delta,k,t\}$,

$$\frac{\partial \sigma^*}{\partial z} = \frac{\frac{\partial \hat{\sigma}}{\partial z} + \frac{\partial \hat{\sigma}}{\partial e}\frac{\partial \hat{e}}{\partial z}}{1 - \frac{\partial \hat{\sigma}}{\partial e}\frac{\partial \hat{e}}{\partial \sigma}}, \tag{23}$$

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 (23) is positive; the sign is then given by that of the numerator. The following proposition describes the comparative statics.

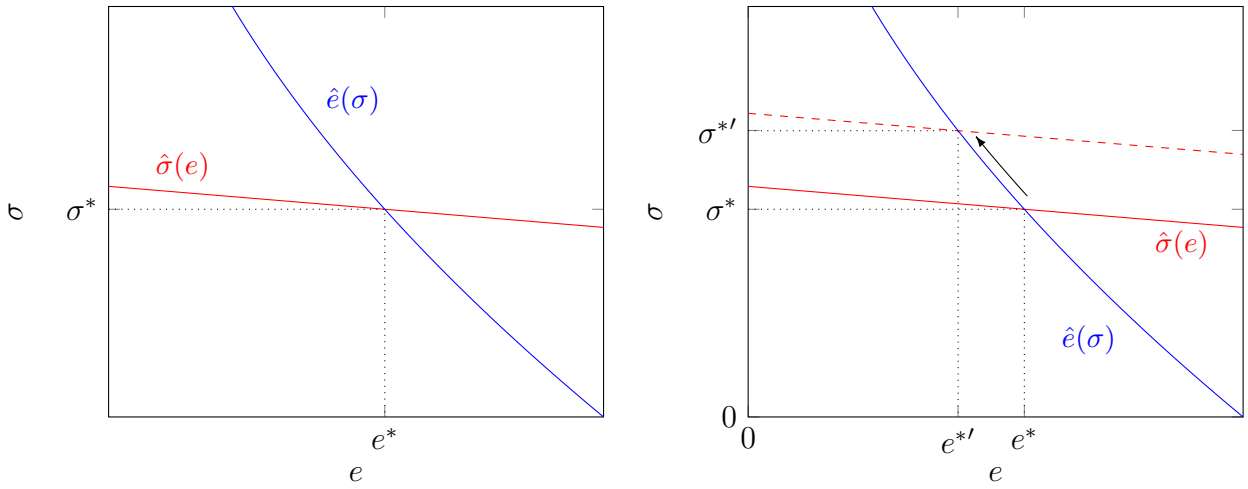

Figure 1: (a) Equilibrium is found where $\hat{\sigma}(e)$ and $\hat{e}(\sigma)$ intersect. (b) Effect of an increase in Δ (which causes $\hat{\sigma}(e)$ to increase).

Proposition 6. *In a stable interior equilibrium of the game with consumer investment: (i) The signs of* $\frac{d\sigma^*}{d\Delta}$, $\frac{d\sigma^*}{dk}$, *and* $\frac{d\sigma^*}{dt}$ *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 arerobust except as detailed in part (iii) of Proposition 6. 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

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

where the first set of square brackets enclose the damage per-breach and the second two terms measure the number of successful attacks. Figure 2 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.

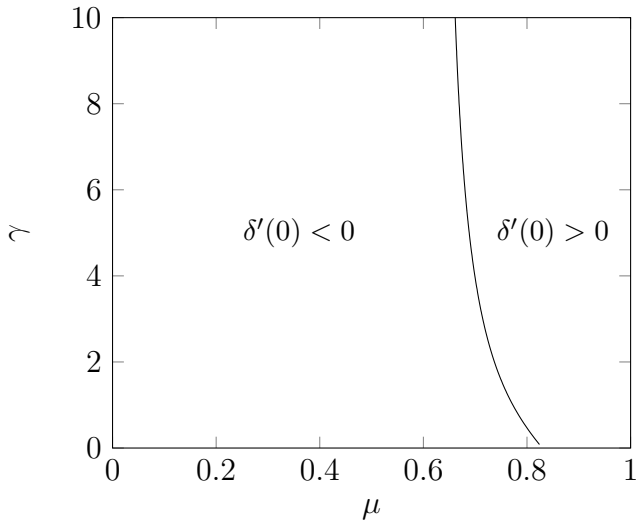

Figure 2: Effect of introducing a small fine on total security breach damages in the pricing regime when $k = 1$, $\Delta = 1/2$, and $L(e) = 2(1 - e)^2$.

7 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 a monopolist will often

choose the socially efficient level of security—fully internalizing the value of this investment to consumers—because its marginal customers are likely to be those most aware of security risks. Introducing competition leads to a fall in investment below the efficient level. 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. We contrast this to an alternative business model where firms are ad funded. Here we find that a monopolist under-invests because it can no longer use prices to capture the incremental value of security to consumers. This problem is mitigated by competition, which induces firms to invest more as they compete for security-savvy customers. Indeed, we may even witness over-investment when firms are ad funded because of the business-stealing effect of investment.

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 extert 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 high social damages from security breaches.

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.

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.

Buchanan, Ben (2020). *The Hacker and the State. Cyber Attacks and the New Normal of Geopolitics*. Harvard University Press.

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*.

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.

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

Perlroth, Nicole (2021). *This Is How They Tell Me the World Ends. The Cyberweapons Arms Race*. Bloomsbury Publishing PLC.

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 3. Let $\psi \equiv \frac{\partial \pi_i}{\partial \sigma_i}\Big|_{\sigma_i=\sigma_j=\sigma}$ (the left-hand side of (17)). 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

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

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 (17) is positive at $\sigma = \sigma_w^*$. Substituting $\sigma = \sigma_w^*$ into (17) and noting that $k\sigma_w^* = h\left(\frac{L}{2} + \Delta\right)(1 - \sigma_w^*)$ (from equation 4), the left-hand side of (17) becomes

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

This is positive when (18) is satisfied. ∎

Proof of Proposition 6. (i) For $z \in \{\Delta, k, t\}$ we have from (22) that $\frac{\partial \hat{e}}{\partial z} = 0$. From (23), therefore, $\frac{d\sigma^*}{dz}$ has the same sign as $\frac{\partial \hat{\sigma}}{\partial z}$, which is just the equilibrium effect of Section 4.

(ii) Applying standard comparative statics methods to (22) 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 (13) (and suppressing the arguments for readability),

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

With these ingredients, calculating the numerator of (23) 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 4. ∎

B Market structure and the strategic effect

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.

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

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

i.e.,

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

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

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

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

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

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 4.1: the system of first-order conditions

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

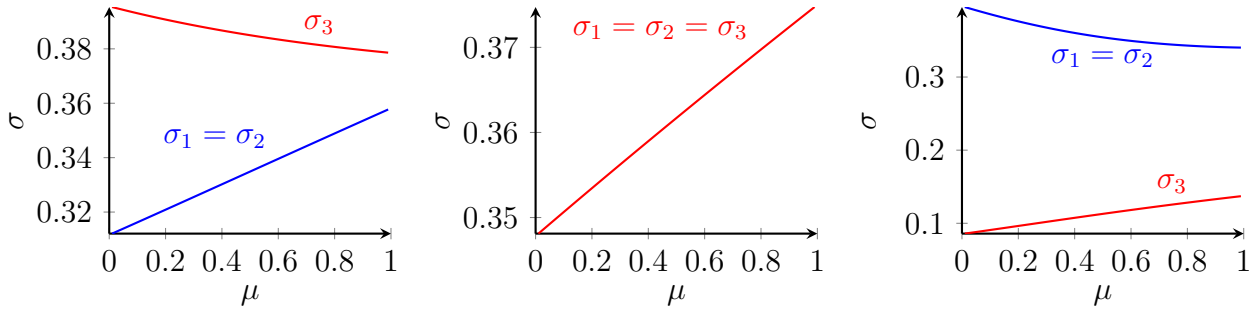

(a) $m = 1/8$ (b) $m = 1/3$ (c) $m = 7/8$

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

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

In Figure 3a 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 3c 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 HHI $= (D_1)^2 + (D_2)^2 + (D_3)^2$. Figure 4 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

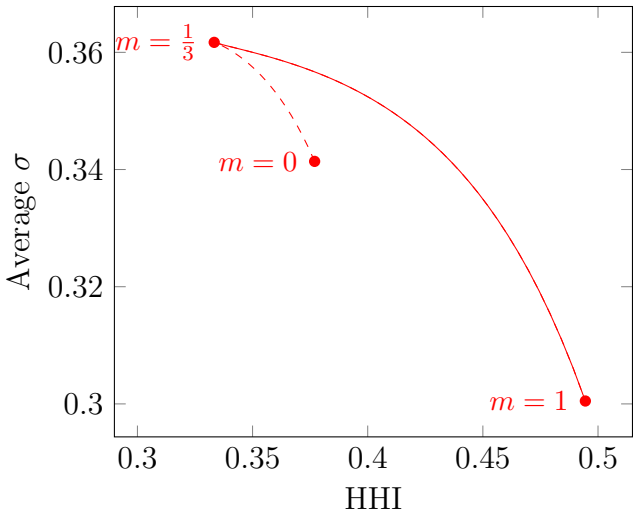

Figure 4: Relationship between market concentration (measured by HHI) and average security investment.

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.

C Non-naive consumers

Instead of assuming that a mass $1 - \mu$ consumers are naive (don't account for security when choosing a firm, but do suffer a loss from a breach), we could suppose that they simply don't care about security or have nothing to lose in an attack. Formally, this means that the social planner now wishes to maximize $2[-n^* h(1-\sigma)^2(\Delta + n^* \mu L) - k\sigma^2/2]$, where $n^* = 1/2$ is the equilibrium market share. This implies

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

Firms' decisions problem (and hence equilibrium outcomes) are unchanged.

Comparison of this new value of σ_w^* with (13), we find that there is still under-investment under price competition, even though the socially optimal investment is now lower. Following Section 5, we can replace Δ with $\Delta + f$ and L with $L - g$ to study the optimal liability regime. Efficient investment is achieved in the price competition model by $f^*(g) = \frac{1}{4}(2\Delta - g\mu + 4L\mu)$. The preserves the same properties we observed in Section 5: f and g are strategic substitutes

and the optimal fine is always punitive (i.e. $f^*(g) > L\mu/2$).

Turning to the case of ad-funded business models, we must compare the planner's first-order condition, $k\sigma = h\left(\Delta + \frac{L\mu}{2}\right)(1-\sigma)$, with (17). We then find that there is over-investment if and only if $t < R - h(\Delta + L\mu)\left(1-\sigma_w\right)^2$, which is analogous to the threshold in Proposition 3. The arguments regarding optimal liability under ad-funded business models from Section 5 continue to hold; in particular, fines and consumer compensation are still strategic complements. Moreover, if we let $g = L$ and $f = L\mu/2$ we still find that the budget-balanced, full-insurance scheme implements first-best.