

Information Security and Competition*

Alexandre de Cornière[†] and Greg Taylor[‡]

February 18, 2021

Abstract

Cyber-attacks 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

In the digital age, the issue of cybersecurity is of considerable importance. Some industry observers estimate that thirty million cyber attacks occurred in 2018,¹ and put the damage related to cybercrime at \$1trn globally in 2020.² Organizations face a stream of attempts to steal valuable data—from consumer data held by e-commerce firms to confidential busi-

*We are grateful for helpful comments and suggestions from Paul Bellflamme and David Ronayne.

[†]Toulouse School of Economics, University of Toulouse Capitole; alexandre.de-corniere@tse-fr.eu; <https://sites.google.com/site/adcorniere>

[‡]Oxford Internet Institute, University of Oxford; greg.taylor@oii.ox.ac.uk; <http://www.greg-taylor.co.uk>

¹See <https://purplesec.us/resources/cyber-security-statistics/>.

²See <https://www.mcafee.com/enterprise/en-us/assets/reports/rp-hidden-costs-of-cybercrime.pdf>, accessed 17 February 2021.

ness documents stored on cloud platforms.³ Other attacks are purely malicious, seeking to inflict the maximum damage possible. Denial of service attacks, for instance, overwhelm online infrastructure and services with network traffic, rendering them unavailable to legitimate users.⁴ Moreover, the so-called Internet of Things promises to bring these threats into the physical realm. As consumers and firms install “smart” devices into their homes and businesses, reports already abound of hacked baby monitors and security cameras.⁵ Increasingly connected devices—whether industrial machinery, autonomous vehicles, or connected medical devices—will expose yet more aspects of life to cyber threats.

The recent hack of Solar Winds, a major software provider, illustrates the risk to businesses and governments posed by hacking. During the early months of 2020, hackers suspected of working for the Russian government inserted malicious code into Solar Winds’ software system. As part of its regular updates, the company then unwittingly sent out this code to its customers, which include some large companies as well as several departments of the US government. “The code created a backdoor to customer’s information technology systems, which hackers then used to install even more malware that helped them spy on” Solar Winds’ customers,⁶ an estimated 18 000 of which having been left vulnerable to the hack. The attack was only uncovered at the end of 2020, leaving the hackers enough time to infiltrate a high number of very sensitive IT networks.

The fight against cybercrime is as much an economic as a technical one (Anderson and Moore, 2006). Indeed, attackers are generally thought of as rational agents following a cost-benefit analysis regarding their choice of targets and mode of attack, with an estimated 86% of breaches motivated by financial gain.⁷ Defenders also tend to respond to incentives, and the pervasiveness of externalities is one of the main challenges to overcome in order to achieve a secure environment. 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. More precisely, in this paper, we investigate how firms’ incentives to invest in cybersecurity

³For example, hackers exploited vulnerabilities at Adobe (2013), eBay (2014), Equifax (2017), LinkedIn (2012), Marriott (2014–18), Myspace (2013), and Sina Weibo (2020) to steal personal information of, in each case, hundreds of millions of users. See <https://www.csoonline.com/article/2130877/the-biggest-data-breaches-of-the-21st-century.html>, accessed 11 November 2020.

⁴One such attack in 2020 targeted Amazon Web Services, which supplies key infrastructure for, among other users, Netflix, LinkedIn, Facebook, the BBC, Baidu, and ESPN.

⁵See, e.g., <https://www.theguardian.com/technology/2019/dec/27/ring-camera-lawsuit-hackers-alabama>, <https://www.washingtonpost.com/technology/2018/12/20/nest-cam-baby-monitor-hacked-kidnap-threat-came-device-parents-say/>, accessed 7 October 2020.

⁶See, e.g., <https://www.businessinsider.fr/us/solarwinds-hack-explained-government-agencies-cyber-security>, accessed 14 January 2021.

⁷See <https://enterprise.verizon.com/en-gb/resources/reports/dbir/>, accessed 11 November 2020.

are shaped by the competitive environment in which they operate. The levels of investment involved are substantial, with some projecting firms will spend more than \$170bn per-year by 2022.⁸ We also investigate various policies aimed at correcting market failures, such as the optimal liability regime or a certification scheme. 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 (themselves possibly being firms 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 into their IT system. A successful attack causes harm to both the firm and its customers, and firms' key strategic decision concerns 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 (we contrast monopoly and duopoly) and the degree of substitutability among products (i.e. the transportation cost in a Hotelling framework with exogenous locations). Consumers' awareness is captured by the share of consumers who can observe the security of each firm. In practice, consumer's awareness is likely to vary depending on which market we consider. In a B2C market, while final consumers may be aware of the use of encryption techniques for messaging or payment applications, they are probably less informed about the risks associated with different products than the IT departments of large companies in a B2B context (for instance when it comes to choosing a cloud provider). We introduce consumer heterogeneity by assuming that a share of consumers are naive and ignore security aspects. Finally, we distinguish between firms whose business model consists in selling a product, and who thus have to choose a price (e.g., cloud 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*.

Hackers choose whether or not to attack a firm based on three factors: the costs of launching an attack, the reward for a successful attack, and the likelihood of success. We assume the reward is increasing in the number of consumers a firm serves, which would be the case if hackers gain access to consumers' data for financial gain, if hacking a larger firm brings about more publicity, or if inflicting damage to many consumers is the objective of the attacker. The likelihood of succeeding in an attack is determined by firms' efforts in closing

⁸See <https://www.gartner.com/en/documents/3889055>, accessed 11 November 2020.

security vulnerabilities. If an attack is made successfully then both the target firm and each of its consumers suffer a loss. 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.

We first compare the first-best investment in security to the one chosen by a monopolist. In the pricing regime, we show that, even though consumers have a heterogeneous awareness of security risks, a monopolist's incentives to invest in security are aligned with the social planner's, provided there are enough sophisticated consumers. Because of negative network effects, sophisticated consumers' willingness to pay is lower than that of naive consumers. In order to sell to sophisticated consumers, the monopolist must therefore set the price of its product so as to compensate them for the expected security risk. Because the perceived security risk of the marginal consumer equals the actual risk of the average consumer, the firm's incentives to invest in order to raise the price lead to efficient investment (i.e. there is no distortion à la Spence, 1975). 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.

Since the competitive provision of security is generically sub-optimal, we also consider regulations that better align firms' incentives with those of the social planner. 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 (so larger fines are needed if consumers are not compensated), whereas the two instruments are strategic complements in the advertising regime. Moreover, the strategic effect noted above implies that an optimal fine in the pricing regime is always punitive (in the sense that it exceeds the loss incurred by consumers), whereas the presence of the business-stealing effect means that optimal fines are not 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.

In our baseline model only firms can exert protection efforts. We discuss the robustness of this assumption in Section 6, where we allow sophisticated consumers to take preventive actions.

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 overview. Among the issues considered in this literature is the incentive of a firm to invest in security. In an early contribution Gordon and Loeb (2002) introduce some of the basic economic trade-offs that face a firm when deciding how much to invest. More recently, economists have contributed, with emphasis on the novel economic and policy dimensions of the problem. For example, Lam (2016) studies the optimal design of liability policy (something we also consider) when a firm can invest in either preventing or mitigating the effects of an attack. August and Tunca (2006) and Choi, Fershtman, and Gandal (2010) emphasize the fact that security requires users to play an active role by installing updates and vulnerability patches, and that their potential failure to do so has policy as well as strategic implications for firms. More closely related to our paper, Fainmesser, Galeotti, and Momot (2020) study the incentives of a firm to protect consumer data as a function of its business model, and consider policies that can induce the socially optimal choice. Jullien, Lefouili, and Riordan (2020) consider a website's incentives to screen its advertisers when some of them install malware on users' devices. These papers focus on investment by a monopoly firm, whereas we study the strategic effects of security investments on market competition.

Relatively few papers study security in oligopolistic frameworks. On the theory side, Gordon, Loeb, and Lucyshyn (2003) and Gal-Or and Ghose (2005) focus on the issue of information sharing among competitors, highlighting its strategic use and the associated free-riding problem. Empirical work about the link between security and market structure is also 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.

Broadly speaking, there are two categories of cyber-attacks: targeted and non-targeted ones. The former are aimed at a specific target based on its value to the attacker and its level of protection, while the latter exploit contagion mechanisms in networked environments, and

seek to infect as many devices as possible. Fabrizi, Lippert, and Rodrigues-Neto (2019) focus on firms that sell protection to interconnected individuals subject to untargeted attacks. A kind of positive network effect arises because firms’ costs of providing protection are assumed to be lower when more individuals protect themselves. In Dey, Lahiri, and Zhang (2012), individuals may be subject to either untargeted or targeted attacks. In the former case, there are positive network effects because protection prevents contagion between users, whereas the latter case leads to negative spillovers because each new protected user leaves a smaller group of unprotected individuals to bear the brunt of attacks. August, Niculescu, and Shin (2014) looks at the choice of software vendor of whether to offer a cloud version on top of its on-premises version, when the two versions differ in terms of their quality and the type of attacks (targeted or non-targeted) they face. In contrast to these papers, we focus on targeted attacks against firms that invest in protecting themselves (on their consumers’ behalf). This generates network effects that are negative because larger firms are more attractive targets. Our analysis therefore applies to different kinds of market situation, addresses new questions about security and competition, and yields quite different results. We also study a range of policy responses to the market failures we identify.

A literature (e.g., Acemoglu, Malekian, and Ozdaglar, 2016; Dziubiński and Goyal, 2017; Goyal and Vigier, 2014) studies how the effects of an attack are mediated by the structure of a network that connects potential target nodes. The focus is on issues such as the optimal defensive network structure and the spillovers between connected nodes. This is also related to work on the optimal allocation of defensive resources among potential targets (e.g., Bier, Oliveros, and Samuelson, 2007; Dziubiński and Goyal, 2017; Kovenock and Roberson, 2012). If the effects of an attack can spill-over between interconnected agents then security investment becomes a public good whose equilibrium properties depend upon the underlying security technology (Hirschleifer, 1983; Varian, 2004). Rather than focusing on the security decisions of populations of individual (but interdependent) agents, we study markets where firms are responsible for protecting the security of all their customers’ data. This allows us to examine how security investment fits in to the wider picture of market competition under different business models.

2 Model and preliminaries

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

Product market The firms offer different varieties of a product, for which consumers have unit demand. We adopt the Hotelling formulation, and assume that the two firms are located at opposite ends of a unit-length segment. Consumers are uniformly distributed along the segment, and the gross utility of a consumer who selects a product located at a distance d from his ideal position is $V - td$, where the stand-alone value, V , is assumed large enough to ensure the market is always covered in equilibrium.

We consider two kinds of business models for firms, depending on whether they generate revenues through *pricing* or *advertising*.⁹ In the pricing regime, each firm sets a price p_i , which enters consumers' utility negatively in a linear way. In the advertising regime, revenues are exogenously given,¹⁰ and proportional to the number of consumers served. We denote by R the per-user advertising revenue. Note that we do not study the choice of business model by firms, but rather how the type of business model affects firms' cybersecurity strategies.

Security In the field of information security it is recognised that (social or technical) systems may have *vulnerabilities* that, when confronted with a security threat, lead to a security *breach* (see Anderson, 2008, ch1). We focus on threats that take the form of deliberate, targeted attacks by *hackers*.¹¹ Ransbotham and Mitra (2009) and Schechter and Smith (2003) emphasise the rational basis for choice of targets by hackers. 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. Ransbotham and Mitra (2009) find that the reward may take the form of material gain (e.g., sale or use of stolen data) and status associated with successfully compromising a high-profile target, while a principal barrier to success is the security countermeasures implemented by the target. 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."¹²

Formally, we capture these elements as follows. We assume that each firm's IT system is

⁹Throughout, 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.

¹⁰We treat R as exogenous in order to cleanly emphasize the main difference between the price and ad regimes—namely, a strategic effect that works via consumer prices. But, as we show in Section 4.2, we can easily allow R to depend on a firm's security without changing the results.

¹¹The term hacking is often colloquially used to refer specifically to technical attacks against a computer system, but we use it more broadly to refer to any attacks aimed at exploiting a vulnerability to breach the security of a system. Thus, a hack could encompass either a technical attack against a computer security system or a non-technical attack such as persuading a worker to reveal their password or a security guard to let you into the computer room.

¹²The 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.

exposed to potential vulnerabilities, which hackers seek to exploit. By fixing vulnerabilities, each 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. Fixing vulnerabilities requires investing in security: in order to achieve a level of protection σ_i , firm i needs to pay $\frac{k\sigma_i^2}{2}$.¹³ Related 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 hacker with probability h . 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 idiosyncratic to the hacker and uniformly distributed on $[0, 1]$. In case of a breach, the hacker gets a payoff of 1 per customer.

Thus, if firm i serves n_i consumers, the payoff to attacking firm i is $(1 - \sigma_i)n_i - c$. It follows that the probability of a *successful* attack against i is

$$h \cdot \Pr [c < (1 - \sigma_i)n_i] \cdot (1 - \sigma_i) = h(1 - \sigma_i)^2 n_i,$$

where the three terms on the left correspond respectively to the probability a hacker is present, the probability they find it worthwhile to attack this particular firm, and the probability that the attack succeeds.

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, we believe this effect to be negligible, as there are many firms in other markets, and hackers are not constrained to target a firm in the specific product market we model.

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 damage incurred, or even a fine imposed by regulators (see Section 5).¹⁴ Firm i 's payoff is then

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

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

A successful attack on firm i also imposes a loss $L \geq 0$ on each of its customers, either

¹³Most 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.

¹⁴In 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/equifax-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 from interrupted access to compromised services. 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 is replaced by zero 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 fully 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.¹⁵ This parameterization allows us to apply the model to various kinds of market. In business-to-business (B2B) markets, where the “consumers” are themselves firms, technology purchasing decisions will often be made by a dedicated IT team that understands and prioritizes security (μ is relatively large). Moreover, B2B markets are generally organized around price competition rather than the ad-funded business model. In contrast, business-to-consumer (B2C) markets are likely to have less savvy customers (μ small, or even zero) and may involve either price competition or ad-funded business models.

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 which firm to patronize. In the fourth stage, hackers facing each firm observe its security and market share before deciding whether to attack. We look for subgame perfect equilibria.

For quick reference, Table 1 summarises the main parameters and variables of the model. In order to focus on interior solutions throughout the paper, we assume that k is large enough, and that $4\Delta > L\mu$.

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.

¹⁵Instead 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.

Table 1: Summary of important variables and parameters in the model.

Variable	Meaning
σ_i	Firm i 's security level.
π_i	Firm i 's profit.
n_i	Firm i 's demand.
p_i	Firm i 's price (under price competition).
R	Firms' per-consumer ad revenue (in the ad-funded business model).
t	Transport cost—a measure of product differentiation.
μ	Fraction of sophisticated consumers (who observe firms' security).
Δ	Direct damages incurred by a firm that suffers a breach.
L	Loss incurred by consumers whose chosen firm suffers a breach.
h	Overall prevalence of hacking in the market.
k	Parameter scaling the cost of investment in security.

When firms have $\sigma_1 = \sigma_2 = \sigma$, the equilibrium of the ensuing subgame is symmetric and each firm serves half of the market.¹⁶ 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

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

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

$$2n^*h(1 - \sigma)(\Delta + n^*L) - k\sigma = 0. \quad (3)$$

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*

$$\sigma_w^* = \frac{h(L + 2\Delta)}{2k + hL + 2h\Delta}. \quad (4)$$

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

¹⁶We establish this formally in Section 4.1 below.

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.¹⁷

When the monopolist serves all consumers, the efficient level of investment is found from (3) after substituting $n^* = 1$:

$$2h(1 - \sigma)(\Delta + L) - k\sigma = 0 \quad (5)$$

Pricing regime A savvy consumer who expects the monopolist to serve n consumers estimates the security risk to be equal to $hL(1 - \sigma)^2n$, and his willingness to pay (given that $t = 0$) is $V - hL(1 - \sigma)^2n$: 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, (1), we obtain a first-order condition, $\frac{\partial \pi}{\partial \sigma} = 0$, that coincides exactly with (5). 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 monopolist 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, (1), the optimal investment solves

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

Comparing this with (5) reveals that the monopolist under-invests from a social perspective because (naive) consumers' willingness to pay does not respond to security investment, 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

¹⁷When $t > 0$ 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.

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 (5) 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:

Lemma 2. *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

We know that, given investment levels and consumers' choices, the probability of a successful attack against firm i in the last stage of the game is $h(1 - \sigma_i)^2n_i$. We then 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

$$V - xt - p_1 - h(1 - \sigma_1)^2n_1L = V - (1 - x)t - p_2 - h(1 - \sigma_2)^2n_2L.$$

i.e.

$$x = \frac{t - p_1 + p_2}{2t} - \frac{hL}{2t} (n_1(1 - \sigma_1)^2 - n_2(1 - \sigma_2)^2) \quad (6)$$

The first term on the right-hand side of (6) 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.

$$y = \frac{t - p_1 + p_2}{2t} \quad (7)$$

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

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

when firms compete in prices. Demand in the ad-funded business model is found by setting $p_1 = p_2 = 0$ in (8).

Pricing stage Given σ_1 and σ_2 , firms choose prices to maximize (1), 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:

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

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

Proof. We have

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

Investment in security has the strategic effect of softening subsequent price competition, which plays an important role in the analysis to follow. From the second term of (9), we see that the strategic effect of security works through two channels: firms' damages (the part multiplying Δ) and consumers' losses (the part multiplying L). Regarding firm damages, a firm that has invested a lot in security faces a lower effective marginal cost because the extra attacks attracted by an increase in demand are less likely to damage the firm. The lower

marginal cost leads a firm (and its rival, by strategic complementarity) to reduce prices. To understand the effect via consumers' losses, notice that the price-elasticity of firm i 's demand is

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

As firm j increases σ_j , firm i 's demand becomes more price-elastic. Indeed, inspection of equation 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 network effects due to security concerns, thereby intensifying price competition

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

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

where p_i^* and n_i^* are respectively given in (9) and (8). 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^* :

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

As one might expect, a firm's equilibrium investment in security is increasing in the rate of hacking attacks, 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 surprising: 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) the parameter t , does not affect the equilibrium investment in security. This is because of two opposite effects. On the one hand, demand for firm i is less sensitive to σ_i as t increases, by (8). 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 (4) and (12), we find that

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

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.

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

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

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.

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

Proposition 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 t .

Note that the facts that $\frac{\partial \sigma_p^*}{\partial \mu} < 0$ and $\frac{\partial \sigma_p^*}{\partial t} = 0$ demonstrate that what we call “security”

is different from a mere investment in the quality of product i , as 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 .¹⁸ Introducing competition can therefore reduce investment in security (because of the strategic effect described in Lemma 3, which is not active for a monopolist).¹⁹ 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 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 this regime, demand is given by (8) with $p_i = p_j = 0$. Firms' security investment is chosen to solve

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

¹⁸If μ is small then the monopolist does not serve the savvy consumers and invests less than does a duopolist.

¹⁹It 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 (8) (with $p_1 = p_2 = 0$). The symmetric equilibrium σ is implicitly given by evaluating firms' first-order conditions at $\sigma_i = \sigma_j = \sigma$:

$$\frac{h \{2t\Delta + L\mu [R + h\Delta(1 - \sigma)^2]\} (1 - \sigma)}{2(t + hL\mu(1 - \sigma)^2)} - k\sigma = 0. \quad (16)$$

Applying the implicit function theorem to (16) 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 2. *In the advertising regime, firms over-invest compared to the socially optimal solution if*

$$\frac{t}{\mu} < R - h(L + \Delta)(1 - \sigma_w^*)^2, \quad (17)$$

where σ_w^* is given in (4), 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 .

Discussion 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 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 (16): competition forces firms to invest more to avoid losing savvy consumers to a rival.²⁰

²⁰If 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 with secure firms). We could easily incorporate this into the model by letting $R'(\sigma_i) \neq 0$. Then (16) becomes

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

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.

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.²¹

In this section we use our model to shed light on two policy tools, namely the optimal liability regime and a certification scheme.

²¹See <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.

5.1 Optimal liability regime

Suppose that the regulator can impose a fine $f \geq 0$ to 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

$$\sigma_w^* = \sigma_p^* \iff \frac{h(L + 2\Delta)}{2k + hL + 2h\Delta} = \frac{h[4(\Delta + f) - (L - g)\mu]}{6k + 4h(\Delta + f) - h(L - g)\mu}. \quad (18)$$

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

$$f_p^*(g) = \frac{1}{4}[2\Delta - g\mu + L(3 + \mu)]. \quad (19)$$

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 $p_i / (t - p_i + p_j + h(L - g)(1 - \sigma_j)^2)$, 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 $f_p^*(g) > L/2$. 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 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 (14), 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 (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 3). 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 (16) holds at $\sigma = \sigma_w^*$ works:

$$\frac{h \{2t(\Delta + f) + (L - g)\mu [R + h(\Delta + f)(1 - \sigma_w^*)^2]\} (1 - \sigma_w^*)}{2(t + h(L - g)\mu(1 - \sigma_w^*)^2)} - k\sigma_w^* = 0. \quad (20)$$

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 2, 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 2 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 3. *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:*

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

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 $g \in [0, L]$.*

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 $[g, L]$, with $\underline{g} > 0$, depending on the parameters of the model. This implies that fines 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 4. *(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 4 is a corollary of Propositions 1 and 2 (regarding the effect of an increase in μ), and of Equations 19 and 20 (regarding the link with the optimal liability regime).

Part (ii) of Proposition 4 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 $\lim_{e \rightarrow \infty} L(e) \geq 0$. 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 prevalence of 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

$$S(e, \hat{e}, \sigma) = V - \frac{t}{4} - p - L(e) \frac{h(\hat{e}\mu)}{2} (1 - \sigma)^2 - e, \quad (21)$$

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

Expression (21) 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 has thus a lower incentive to invest himself. We further assume that the functions L and h are such that $\frac{\partial S(e, 0, \sigma)}{\partial e} > 0$, i.e. that a consumer who expects other consumers not to protect themselves will exert a positive effort. Because $\frac{\partial^2 S(e, \hat{e}, \sigma)}{\partial e \partial \hat{e}} < 0$, this ensures that 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: as savvy consumers expect firms to invest more in security, they reduce their own 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 (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 in general: 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.

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 $z \in \{\mu, \Delta, k, t\}$,

$$\frac{d\sigma^*}{dz} = \frac{\frac{\partial \hat{\sigma}}{\partial z} + \frac{\partial \hat{\sigma}}{\partial e} \frac{\partial \hat{e}}{\partial z}}{1 - \frac{\partial \hat{\sigma}}{\partial e} \frac{\partial \hat{e}}{\partial \sigma}}, \quad (22)$$

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

Proposition 5. *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.*

Proof. (i) For $z \in \{\Delta, k, t\}$ we have from (21) that $\frac{\partial \hat{e}}{\partial z} = 0$. From (22), 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 (21) yields

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




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).

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

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

With these ingredients, calculating the numerator of (22) yields

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

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. ■

The baseline comparative statics results are, for the most part, robust. The main exception is in the advertising regime when γ is large (i.e., when consumer investment in security quickly reduces the payoff to hacking). In that case, adding savvy consumers (who invest) may reduce firm effort simply because fewer hackers are active and the firm feels less need to protect itself.

The results in this section also have some implications for policy. Because $\hat{e}'(\sigma) < 0$, the effects of any policy (e.g., a fine) that increases $\hat{\sigma}$ will be partly offset by a fall in consumers' equilibrium effort. In words: a policy that induces firm investment also crowds-out consumer effort, blunting the effect on overall security.

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.

We identify a variety of forces that cause a divergence of equilibrium investment from the socially optimal level. These include externalities and business-stealing, as well as a novel strategic effect that arises from the interaction of security investment and price competition. We also studied how liability regulation can be used to remedy these distortions. 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, security is under-provided; improved consumer information exacerbates this problem in concentrated markets; and an optimal liability scheme involves punitive fines. When firms are ad-funded, security can be under- or over-provided; improving consumer information or intensifying competition increases investment; and the optimal liability regime involves modest (or even zero) fines.

References

- Acemoglu, Daron, Azarakhsh Malekian, and Asu Ozdaglar (2016). “Network Security and Contagion”. *Journal of Economic Theory* 166, pp. 536–585.
- Anderson, Ross (2008). *Security Engineering, Second Edition*. 2nd ed. Indianapolis, IN: Wiley.
- Anderson, Ross and Tyler Moore (2006). “The Economics of Information Security”. *Science* 314.5799, pp. 610–613.
- 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.
- Bier, Vicki, Santiago Oliveros, and Larry Samuelson (2007). “Choosing What to Protect: Strategic Defensive Allocation against an Unknown Attacker”. *Journal of Public Economic Theory* 9.4, pp. 563–587.
- 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*.
- Gal-Or, Esther and Anindya Ghose (2005). “The economic incentives for sharing security information”. *Information Systems Research* 16.2, pp. 186–208.
- 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*.
- Kovenock, Dan and Brian Roberson (2012). “Conflicts with Multiple Battlefields”. *The Oxford Handbook of the Economics of Peace and Conflict*. Ed. by Michelle R. Garfinkel and Stergios Skaperdas. Oxford: Oxford University Press, pp. 1–30.
- 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.
- 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.
- Spence, A Michael (1975). “Monopoly, quality, and regulation”. *The Bell Journal of Economics*, pp. 417–429.
- Varian, Hal (2004). “System Reliability and Free Riding”. *Working Paper*.

A Proof of Proposition 2

Let $\psi \equiv \left. \frac{\partial \pi_i}{\partial \sigma_i} \right|_{\sigma_i = \sigma_j = \sigma}$ (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

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

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 3), the left-hand side of (16) becomes

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

This is positive when (17) is satisfied.

B Market structure and the strategic effect

This section extends the model to incorporate a third firm and thereby study the role of the strategic effect (Lemma 3) 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.,

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

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




Figure 2: 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 .

the analogous steps to those found in Section 4.1: the system of first-order conditions

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

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

In Figure 2a 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 2c 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




Figure 3: Relationship between market concentration (measured by HHI) and average security investment.

using the Herfindahl-Hirschman Index (HHI). The equilibrium σ s 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 3 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 of one half ($\text{HHI} = (1/2)^2 + (1/2)^2 = 1/2$). 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 (12), 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 $f^*(g) = \frac{1}{4}(2\Delta - g\mu + 4L\mu)$. This 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, $k\sigma = h(\Delta + \frac{L\mu}{2})(1 - \sigma)$, with (16). We then find that there is over-investment if and only if $t < R - h(\Delta + L\mu)(1 - \sigma_w)^2$, which is analogous to the threshold in Proposition 2. 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.