Digital Privacy [*]

Itay P. Fainmesser[†] Andrea Galeotti[‡] Ruslan Momot[§]

September 2019

Abstract

This paper proposes a framework for studying data policy, information security, and privacy concerns in digital businesses. We offer a full characterization of the optimal design of data storage and data protection policies for a digital company and also of how those policies affect users' activity, privacy, and welfare. Our framework features a taxonomy that distinguishes between *advertisement-driven companies* (e.g., Facebook and Google) and *transaction-driven companies* (e.g., Amazon and Uber). This distinction reveals that advertisement-driven businesses store either all or none of the user-generated data whereas transaction-driven businesses exhibit a smoother pattern that may include an intermediate data storage policy. Comparing the amount of user information that these two types of companies store, we find that—contrary to public opinion—advertisement-driven companies do not invariably retain more of their users' data than do transaction-driven companies. Our study establishes that measuring the direct damage inflicted by adversaries on consumers significantly underestimates not only the welfare loss but also the loss of consumer surplus due to adversarial activity. Finally, we identify the conditions under which advertisement-driven businesses generate more consumer surplus than that generated by their transaction-driven counterparts.

> "The problem with data protection laws is that it presumes the data collection was ok."
>
> *Eduard Snowden, NSA whistleblower*
> September 19, 2019

[*] Preliminary and incomplete.
[†] Johns Hopkins University
[‡] London Business School
[§] HEC Paris

1 Introduction

An ever larger share of our social and economic activity is being conducted on and moderated by large digital businesses such as Facebook, Google, Amazon, Uber, and other "sharing economy" platforms. This activity—from the sharing of location data on Uber, to revealing personal information and thoughts posted on Facebook, to searching for medications or diagnoses on Google, to purchasing an array of goods on Amazon—generates extensive amounts of data on individual users. That information is then used to improve services provided by digital businesses. For example, Uber uses passengers' and drivers' location to find the best match, shorten wait times, and enhance safety; Facebook uses personal data to curate posts presented to users based on their tastes and to refine its targeting of users with advertisements; Amazon uses information to match users with the products they are likely to want; and a variety of online dating platforms use information to propose desirable matches between individuals. Digital businesses, in turn, extract some of the generated gains as increased profits.

However, recent news about Facebook and other digital businesses has publicly highlighted some undesirable consequences of the ubiquitous availability of massive data sets that include personal information. Some of the more notorious examples are the consulting firm Cambridge Analytica using Facebook data to sway election outcomes,[1] health insurance companies seeking to predict the health outcomes of potential insurers based on confidential private information,[2] identity theft from the databases of Equifax[3] and Marriott,[4] and criminal scammers seeking to extract money from vulnerable populations. One can only conclude that sharing personal information on the Web exposes individuals to a range of risks imposed by bad actors or *adversaries*.

A strong public reaction, which included movements such as #DeleteFacebook,[5] led consumer protection agencies and government authorities to seek regulation and also induced digital businesses to re-evaluate their business models. Examples include the EU's General Data Protection Regulation (GDPR), introduced in 2018 to protect personal data of users in the European Union,[6] as well as

[1] "How Trump Consultants Exploited the Facebook Data of Millions", *New York Times*, 17 March 2018; see also Papanastasiou [2018] and Candogan and Drakopoulos [2017].

[2] "Can a Facebook Post Make Your Insurance Cost More?", *Wall Street Journal*, 18 March 2019.

[3] " 'We've Been Breached': Inside the Equifax Hack", *Wall Street Journal*, 18 September 2017.

[4] "Marriott CEO Reveals New Details About Mega Breach", *Forbes*, 11 March 2019.

[5] "The #DeleteFacebook Movement Has Reached a Fever Pitch, as Former Facebook Insiders Turn on the Company", *Business Insider*, 21 March 2018.

[6] "The Birth of GDPR: What Is It and What You Need to Know", *Forbes*, 25 May 2018; `https://gdpr-info.eu`

this acknowledgment by Facebook's Mark Zuckerberg: "We don't have the strongest reputation on privacy right now, to put it lightly. But I'm committed to doing this well and starting a new chapter for our product."[7]

In this paper we develop an analytical framework to study data policy, information security, and privacy concerns in digital businesses. We focus on the following three research questions.

1. What are the incentives of a digital company to store and secure users' information?
2. How do those incentives depend on the company's revenue model? For instance, would an advertising-driven business such as Google have a different data policy than a transaction-driven business such as Uber?
3. What is the scope for policies to mitigate the exploitation of user information by adversaries?

Our model features agents of three strategic types: a digital business (we use the terms "digital business" and "company" interchangeably); users; and adversaries. A *digital business* offers services to users. *Users* have idiosyncratic valuations of the services offered, and they might benefit from interactions with other users (positive network effects). Users' activity (usage of the service) is endogenous and conveys information about their characteristics. This information may be used by the business to enhance its service and thus provide value to users. However, it could also be used maliciously by other economic agents, referred to as *adversaries*. These agents are heterogeneous in their ability to access the user information stored by the digital business.

The digital business benefits from user activity and also from its ownership of user information. To maintain these benefits, the digital business defines a data policy consisting of two instruments. The first is its *data protection* policy: How difficult is it for adversaries to access user information stored on the company's servers? A digital business can upgrade its firewall or antivirus protection, making it more costly for adversaries to access information. The company can also create more effective screening procedures and/or restrict access to user data through its application programming interface (API) in order to prevent excessive user data procurement by affiliated companies.[8] The second instrument is the company's *data storage* policy: How much of a user's information is actually stored in the company's database? Examples include the decision by Whatsapp to encrypt users'

[7] "Facebook's Zuckerberg Announces Privacy Overhaul: 'We Don't Have the Strongest Reputation' ", *The Guardian*, 30 April 2019.

[8] "Facebook Restricts APIs, Axes Old Instagram Platform Amidst Scandals", *TechCrunch*, 4 April 2018.

text messages (thus reducing the amount of information it stores)[9] and Facebook's practices (at least until August 2019) of transcribing audio chats (thereby increasing the amount of accessible information it stores)[10] and of retaining information from deleted accounts.[11]

Our first set of results takes the design of the digital business as given and then characterizes users' and adversaries' behaviors. Any improvement in the company's data *protection* policy leads to reduced adversarial activity and hence to increases in user activity and in the amount of information stored by the digital business. If the company's data *storage* policy becomes more stringent (i.e., if a smaller fraction of data is stored), then user activity and the resulting amount of information stored first increase and then decrease. The level of adversarial activity and consumer surplus follow a similar pattern when the business changes its data storage policy. This non-monotonicity reflects a natural tension between information's benefits and costs to consumers. On the one hand, consumers benefit from the better service that a business can provide when it retains more data. On the other hand, increased data availability makes it easier and more attractable for adversaries to harm users—a development that encourages bad agents to enter the market.

Our main result, which builds on the characterization of users' and adversaries' behavior, determines the optimal design of a digital business. If the costs of data protection are *low*, then a digital business will store and use all of the data generated by its users' activities; it will also provide enough protection that adversarial activity is sufficiently reduced that user activity is not compromised. For *intermediate* data protection costs, the digital business will choose an intermediate data storage policy and a less stringent data protection policy. Finally, if data protection costs are *high* then a digital business will save no user data and thus will have no need for data protection.

We underscore the economic content of this characterization by fitting it to the empirically relevant digital business models. One such model is adopted by digital companies, such as Facebook and Google, whose main revenue source is targeted advertising; we refer to these businesses as *advertisement-driven* (or simply *ad-driven*) companies. In adopting this business model, these companies are essentially trading in the user information they store. The second model applies to digital businesses whose main revenue source is user payments in the form of subscription fees (e.g., Dropbox) or of some fraction of user purchases (e.g., Amazon and eBay); these businesses

[9] "WhatsApp Introduces End-to-End Encryption", *New York Times*, 5 April 2016.

[10] "Facebook Paid Contractors to Transcribe Users' Audio Chats", *Bloomberg*, 13 August 2019.

[11] "OK, You've Deleted Facebook, but Is Your Data Still Out There?", *CBS News*, 23 March 2018.

are referred to as *transaction-driven* companies. Companies such as Amazon Marketplace and eBay—as well as online dating services and ride-hailing platforms like Uber and Lyft—all exploit user information in order to match users with services and products and to charge users either for the match itself (eBay, Uber), for access to the service (online dating platforms), or both (Amazon). In general: the more information these companies possess, the better the matches and thus the greater the value for a user who transacts through their platforms.

Our characterization reveals that an ad-driven business chooses *either* to store all data generated by user activity and provide sufficient data protection *or* to store no data and provide no protection—in effect, going out of business. In other words, the optimal design is discrete: there is no intermediate design in which the business provides some protection while adopting a selective data storage policy. In contrast, transaction-driven businesses exhibit a more continuous pattern whereby such an intermediate design may be optimal under certain conditions. Comparing the two types of companies in terms of how much user information each stores, we find that ad-driven companies do not always store more users' data; this outcome is at odds with the prevailing narrative.[12]

Using our final set of results, we evaluate the welfare loss due to adversarial activity. We show that measuring only the direct damage inflicted by adversaries on consumers significantly underestimates the welfare loss as well as the loss of consumer surplus. We define the *adversarial loss multiplier* calculated as a ratio: the consumer surplus loss due to adversaries *divided by* the direct damage to users from adversarial activity. We establish that the lower bound for this loss multiplier is always greater than 2. The lower bound increases with network strength, becoming unbounded from above even for finite network effects. Since also profits decline when adversaries are present, it follows that the total welfare loss is larger still. Finally, we compare advertisement- and transaction-driven companies with respect to the consumer surplus derived by their respective users and find that the relationship depends strongly on the cost of data protection.

This paper contributes to an active interdisciplinary area of research that studies the consequences of the unprecedented ability of digital institutions to amass large data sets consisting of user characteristics and behavior. The main issues discussed in the literature are consumer privacy, the management of consumer information, implications for managerial strategies, and the role of policy intervention.

[12] "The Facts About Facebook", *Wall Street Journal*, 24 January 2019.

With regard to privacy, the computer science literature has addressed the design of algorithmic mechanisms for anonymizing individual-level data (for reviews of this research stream, see e.g. Dwork and Roth 2014, Cummings et al. 2015, Ghosh and Roth 2015, Abowd and Schmutte 2019). Related studies on economics and marketing focus on understanding how users' information is revealed through their purchase decisions (Conitzer et al. 2012), their formation of social links (Acemoglu et al. 2017), and voluntary information disclosure; this literature explores how these mechanisms for extracting user information affect the design of targeted pricing (e.g., Candogan et al. 2012, Bloch and Quérou 2013, Fainmesser and Galeotti 2015, Montes et al. 2018, Fainmesser and Galeotti 2019, Ichihashi 2019b) and social image visibility (Ali and Bénabou 2016) as well as advertising strategies (Galeotti and Goyal 2009, Shen and Miguel Villas-Boas 2017), the extent of competition (Casadesus-Masanell and Hervas-Drane 2015), and overall consumer behavior (Goldfarb and Tucker 2011, Koh et al. 2015, Jann and Schottmüller 2016, Gradwohl 2017). Excellent surveys of this work are provided by Acquisti et al. [2016], Mayzlin [2016], and Bergemann and Bonatti [2019].

Our study complements the extant literature by focusing on the drivers and welfare implications of a digital company's data policy. For that purpose, we abstract from the specific mechanisms through which data can be collected and instead develop a full-fledged strategic model in which the activities of users and adversaries are endogenous to a company's data policy design. This model offers a clear framework within which to study the optimal data policy design of a digital business and to compare it across different domains—in particular, ad-driven and transaction-driven digital businesses.

The rest of our paper proceeds as follows. Section 2 presents the model and describes how it maps to different types of digital businesses. In Section 3 we characterize the behavior of users and adversaries, given the digital company's design, and then derive our main result: a characterization of the optimal business design in terms of its policies for data protection and data storage. Section 4 introduces the adversarial activity loss multiplier and presents a comprehensive welfare analysis. We conclude in Section 5 with summary remarks and a discussion of how the model can be used to answer recurring questions about privacy and the business models typically adopted by digital companies. Proofs for all results are given in Appendix B.

2 Model

A digital business (a company) chooses a *data policy* consisting of two components: a data protection policy and a data storage policy. Consumers decide how much to use the services provided by the digital business and those decisions, together with the company's data policy, determine how much consumer information is stored in the latter's database. Malicious agents or adversaries can, at a cost that depends on the data protection policy in place, attempt to access the database; if successful, they harm consumers. In this section we develop a parsimonious model that links all these elements before giving a few contemporary examples that showcase the model's applicability.

2.1 Users

A unit mass of users choose their respective levels of activity with a digital business, levels that define their usage of a service. Thus each user i chooses a costly action a_i, which represents i's usage level of the business's service. Let $\bar{a} = \int_j a_j \, dj$ denote average user activity. Usage of the service generates information about the user. The information that a digital business stores about its users—and that can be retrieved (either by the business itself or by third parties with access to the database)—is ξa_i. Here $\xi \in [0, 1]$ is the company's data storage policy, which ranges from storing none to all of the focal data. Suppose user i chooses a_i when she expects average activity to be $\bar{a}$;[13] then her utility is

$$U_i(a_i, \bar{a}) = a_i b_i - \tfrac{1}{2} a_i^2 + \beta a_i \bar{a} + a_i \xi \rho - a_i \xi \omega. \tag{1}$$

The first term is user i's stand-alone benefit b_i from the service; users are heterogeneous with respect to b_i, and $\bar{b} = \int_j b_j \, dj$. The second term is the private cost that user i pays for activity a_i—for example, the cost of time spent on a social platform like Facebook or the cost of purchases made on an e-commerce platform like Amazon. The third term, a classical expression for the positive network effect, depends on users' average activity and it is parameterized by $\beta \geq 0$.

The fourth and fifth terms in Eq. 1 determine (respectively) the user's benefits and costs that result from the company's storing her information $a_i \xi$ in its database. Those benefits, which are captured by the parameter $\rho \in [0, 1)$, reflect the digital company's services that are targeted to

[13] To ease the exposition, we use feminine and masculine pronouns for (respectively) the user and adversary.

particular users based on stored information about them. The user's cost, captured by $\omega a_i \xi$, depends on how often she expects her information to be misused by adversaries. In equilibrium, ω will be equal to the mass of active adversaries; in fact, it will become clear that ω can be interpreted as the mean number of adversaries' attacks per user. Hence we refer to ω as the *attack rate* or, more generally, as *adversarial activity.*

2.2 Adversaries

There is a large mass K of potential adversaries.[14] Adversaries are heterogeneous in their ability to access the information stored by a digital business. This heterogeneity is captured by the parameter γ, which we assume (for the sake of simplicity) to be uniformly distributed over $[0, K]$. An adversary knows his own γ and chooses to be active (action 1) or not (action 0). The gain to an inactive adversary is his outside offer, which we normalize to zero ($\pi(0) = 0$). If an adversary with ability γ is active, then he pays a fixed cost γC to access the stored information and attacks one user chosen uniformly at random. An adversary who attacks user i receives payoff $a_i \xi$, so his expected benefit is $\bar{a}\xi$. The parameter C is chosen by the digital business and reflects its data protection policy. Formally, the payoff expected by an active adversary of ability γ is[15]

$$\pi(1|\gamma) = \bar{a}\xi - \gamma C.$$

2.3 Digital business

The digital business chooses the data protection and data storage policies, or (C, ξ), that will maximize its strategic objectives. In particular, the company seeks to maximize users' average activity and the amount of user information that it can use. These objectives can be expressed as

$$\Pi(C, \xi, \bar{a}) = \alpha\bar{a} + \hat{\alpha}\bar{a}\xi - \psi C, \tag{2}$$

[14] We present results for the case in which K is large—strictly speaking, for $K \to \infty$. Doing so ensures the existence of at least some nonactive adversaries. The only role played by this restriction is in facilitating our presentation of the results, and there are no economic insights to be gained from the case of small and finite K.

[15] Another interpretation of this model of the adversary is that there is a single criminal who, upon gaining access to the digital company's data, attacks *all* users. Then γC represents the attack's cost; alternatively, C can denote this cost and $1/\gamma$ the likelihood of that attack being successful.

where ψ signifies the cost of protecting the database's stored information.[16] The parameters α and $\hat{\alpha}$ represent the extent to which the digital company's revenue depends on, respectively, average activity and stored information. The relation between α and $\hat{\alpha}$ depends primarily on the company's adopted business model. Drawing from our observations of existing companies and from our discussion in the Introduction, we classify digital businesses into the two broad categories described next.

Advertisement-driven companies: digital companies whose main source of revenue is offering targeted advertising (e.g., Facebook and Google). In selling targeted ad services, such a business capitalizes directly on the user information it stores. For these companies, then, $\alpha = 0$ and $\hat{\alpha} > 0$.

Transaction-driven companies: digital businesses whose main source of revenue is the payments made by users in the form of subscription fees or commissions. Leading examples of transaction-driven digital businesses are such companies as Amazon Marketplace and eBay, ride-hailing platforms like Uber and Lyft, and online dating services. Recall from the Introduction that such a business matches user information with services and products and then charges users either for the match itself (e.g., eBay) or for access to service (online dating platforms) or for both (Amazon). A digital business that collects more information can make better matches, in which case users gain more utility from its platform. For transaction-driven companies we have $\alpha > 0$ and $\hat{\alpha} = 0$.

2.4 Timing and equilibrium concept

We consider the following two-stage game. In the first stage, a digital business chooses its data policy design (C, ξ)—that is, its data protection policy and data storage policy. This choice is observed by users and adversaries. In the second stage, users choose their activity and adversaries choose whether or not to attack the digital business. Users and adversaries act simultaneously.[17]

In our context, the strategy of a digital business corresponds to a data policy choice $(C, \xi) \in \mathbb{R}_+ \times [0,1]$. The user's strategy is a function $a_i \colon \mathbb{R}_+ \times \mathbb{R}_+ \times [0,1] \to \mathbb{R}_+$ that specifies user i's activity for every possible b_i and (C, ξ). Finally, the strategy of an adversary is a function $v_j \colon [0,K] \times \mathbb{R}_+ \times [0,1] \to \{0,1\}$ that specifies, for every possible $\gamma \in [0,K]$, whether adversary j will attack the company. We use $\mathbf{a}$ and $\mathbf{v}$ to denote the *strategy profile* of (respectively) a user and

[16] If we put $\hat{\alpha} = \alpha_1 - \mu$, then μ is an explicit direct cost to the company for storing information and α_1 is that information's direct benefit to the company.

[17] Because there is a continuum of users and adversaries, no agent's action affects the best reply of others. Hence the analysis does not change if moves are sequential rather than simultaneous.

an adversary.

We shall characterize perfect Bayesian equilibria of the game. Here, an *equilibrium* is a design choice (C^*, ξ^*) and a strategy profile $(\mathbf{a}^*, \mathbf{v}^*)$ such that: (a) the digital business maximizes its profit given $(\mathbf{a}^*, \mathbf{v}^*)$; and (b) for every (C, ξ), $(\mathbf{a}^*, \mathbf{v}^*)$ is a Bayesian equilibrium in the ensuing subgame.

As is typical in the literature featuring models with positive network externalities, we assume that $\beta < 1$. This assumption guarantees that, for every design choice (C, ξ), there is a unique equilibrium in the game's second stage.

2.5 Examples

The heterogeneity of digital businesses reflects several factors: differences in their objectives, the various channels through which they create value for users, their interpretation of adversaries, and the instruments available for storing and protecting information. We provide several examples to show that our model can be mapped to real-world digital businesses. Our focus in this discussion is on the proposed taxonomy of advertisement- and transaction-driven companies.

Facebook (or any social networking platform)**.** For Facebook, a_i measures the content that user i creates and shares. Users exert effort and benefit from interactions with others (the network effect parameter $\beta > 0$), whereas stand-alone benefits are virtually zero ($b \to 0$). Facebook stores the information it collects from users and offers them a variety of targeted services, such as recalling past experiences ("remember X years ago") and promoting new apps and targeted offers that match the user's tastes. Therefore, ρ is positive even prior to our accounting for the potential benefit to users of targeted versus nontargeted advertising.

Companies affiliated with Facebook are granted access to its database (e.g., through Facebook's API). Many of these companies offer benefits to users and complement Facebook's service; thus their presence is reflected by a positive ρ. However, users view some affiliated companies as malicious because they misrepresent the value of their service and/or use the information from Facebook for other purposes. A recent example is Cambridge Analytica. We refer to such malicious companies as adversaries. Note that Cambridge Analytica did not directly hack Facebook. Instead, it obtained rightful access to the information by agreeing to follow Facebook's terms of service—which it later violated. Facebook has since changed its terms of service to reduce the amount and types of information that third parties can collect, a tacit admission that its prior terms were insufficient.

Our model represents such a change by an increase in data protection C.

Facebook is an ad-driven company and its revenue—via targeted advertising and charging affiliated companies for data access—is directly tied to the amount of data it stores. The platform acts as if $\alpha = 0$ and $\hat{\alpha}$ is large. Yet Facebook is typical of social networking platforms in having a different business model during its initial phase. In particular, it did not sell ads and focused instead on leveraging network effects while effectively subsidizing users. In this phase, the company's objective was to increase average users' activity. That metric was the key performance indicator; it was used to price the company's value and to raise capital, which had the benefit of accelerating growth. So during its initial phase, Facebook acted as if $\alpha > 0$ and $\hat{\alpha} = 0$.[18]

Amazon Marketplace. One of the main reasons for users to be active on Amazon Marketplace is finding suitable products. Whether this objective is achievable requires, inter alia, the presence of other users on the platform (i.e, peer buyers or sellers). It follows that network effects are positive.[19] Amazon uses sophisticated algorithms to recommend products to users as a function of such user characteristics as previous purchases. This process includes Amazon's use of user information to expand consumers' "consideration sets" by proposing products they may be unaware of but whose characteristics are correlated with the those of their respective previous purchases. Hence users may benefit from the information that is stored about them, so $\rho > 0$ in this case.

Amazon Marketplace is a transaction-driven company that earns commissions from successful matches between the platform's buyers and sellers. Transaction-driven companies attract adversaries like hackers who try to penetrate the platform's security and to gather, for example, credit card information.

This description of Amazon applies to most matching platforms, such as e-commerce platforms, dating platforms, and "sharing economy" platforms; it applies also to companies that rely on commissions or subscription fees. All of these businesses share the characteristics we have just

[18] This description can also be adapted to describe search engines like Google. Google began by offering free search services to users, and no paid advertising was allowed. In this initial phase, Google's main objective was to grow as rapidly as possible. Information stored about users was the main input used to improve the matches to user searches, but that information was not monetized by Google. Thus, during Google's initial phase, $\alpha > 0$ and $\hat{\alpha} = 0$. It was not until Google acquired a substantial share of the market for search engines that the company adopted an ad-driven revenue model.

[19] That users of Amazon Marketplace are of different types (i.e., buyers and sellers) leads to the possibility of distinguishing between across-side and within-side network effects. Such diversity allows also for richer modeling choices (e.g., multi-sided platform models). Although these considerations help explain, for example, how the business can "price" different users, they are not the focus of our study. Our aim is rather to understand the design and effects of a digital company's data policy.

described.

3 Analysis

We solve the game as follows. In Section 3.1, we characterize the equilibrium—between adversaries' and users' actions—for every data policy design. Then, in Section 3.2, we present the optimal data policy design.

3.1 Equilibrium between users and adversaries

Proposition 1. *Fix the data policy design of the digital business characterized by data protection policy C and data storage policy ξ. Then the ensuing subgame has a unique equilibrium in which average users' activity is*

$$\bar{a}^*(C,\xi) = \frac{C(\bar{b}+\rho\xi)}{C(1-\beta)+\xi^2}, \tag{3}$$

and the equilibrium attack rate by adversaries is

$$\omega^*(C,\xi) = \frac{\xi\bar{a}^*(C,\xi)}{C}. \tag{4}$$

Many of the market's outcomes are determined by the average activity, $\bar{a}^*(C,\xi)$, and the amount of information stored, $\xi\bar{a}^*(C,\xi)$. These two factors play leading roles in determining a company's incentives when deciding on their design of a data policy. Note that the equilibrium attack rate $\omega^*(C,\xi)$ depends on the amount of information stored in the digital company's database, consumer surplus (CS), on the other hand, can be derived using average users' activity: CS $= \frac{1}{2}\int_i a_i^*(C,\xi,b_i)\,di = \frac{1}{2}[\sigma_b^2 + \bar{a}^*(C,\xi)]$, where σ_b^2 is the variance in the stand-alone benefits $\{b_i\}$. Given the importance of average activity and of how much information is stored, it is critical to understand the comparative statics of these quantities with respect to the choice of data policy described by C and ξ.

Proposition 2. *Average user activity ($\bar{a}^*(C,\xi)$) and amount of information stored ($\xi\bar{a}^*(C,\xi)$) each increase with a stronger data* protection *policy C; however, both of these metrics first increase and then decrease with increasing fraction of stored data—that is, with the company's* storage *policy ξ.*

Formally, there exist $\hat{\xi}(\bar{b}, C, \rho, \beta) \leq \tilde{\xi}(\bar{b}, C, \rho, \beta)$ *that are decreasing in* C *and such that:*[20]

(i) $\bar{a}^*(C, \xi)$ *increases with* ξ *for* $\xi \in [0, \min\{1, \hat{\xi}(\bar{b}, C, \rho, \beta)\}]$ *and decreases otherwise;*

(ii) $\xi\bar{a}^*(C, \xi)$ *increases with* ξ *for* $\xi \in [0, \min\{1, \tilde{\xi}(\bar{b}, C, \rho, \beta)\}]$ *and decreases otherwise.*

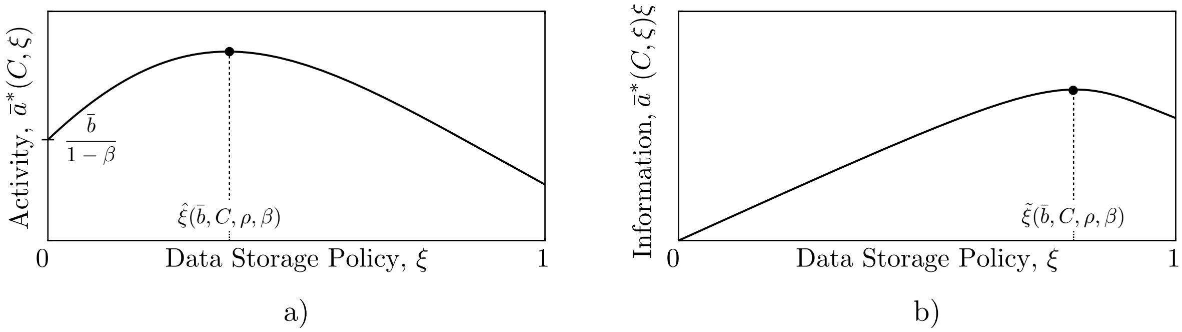

Figure 1: Users' average activity $\bar{a}^*(C, \xi)$ (panel a) and average stored information $\xi\bar{a}^*(C, \xi)$ (panel b) as a function of data storage policy ξ see text).

It is straightforward to intuit the comparative statics with respect to a data protection policy C. Yet the effect of data storage policy (ξ) is more subtle, as illustrated in Figure 1. When no information is stored (i.e., $\xi = 0$), the equilibrium attack rate is zero and user activity is determined solely by the interaction between a user's stand-alone benefit and positive network effects. By storing ever greater proportions of user data (increasing ξ), the digital business creates new benefits for users since it can then offer tailored services based on users' information (as reflected in the term $a_i\rho\xi$ in Eq. 1). At the same time, increasing ξ also boosts the attack rate. However, adversaries are unlikely to misuse information when ξ is still. So in this case, any increase in the amount of information stored will increase the extent of user activity.

With increasing ξ, adversaries have more to gain from every attack; this dynamic leads to a further increase not only in the attack rate but also in the loss that users suffer from any single attack. At some point, these negative effects outweigh the benefit to users from their information being used, which in turn reduces users' average activity. We remark that, even though users' average activity declines for every $\xi > \hat{\xi}(\bar{b}, C, \rho, \beta)$, total information stored declines only for $\xi > \tilde{\xi}(\bar{b}, C, \rho, \beta) \geq \hat{\xi}(\bar{b}, C, \rho, \beta)$. Yet for a sufficiently large ξ, any additional increase in ξ leads to a decline in user activity that is steep enough to reduce total information in ξ.

[20] In particular, $\hat{\xi}(\bar{b}, C, \rho, \beta) = -\frac{\bar{b}}{\rho} + \sqrt{\left(\frac{\bar{b}}{\rho}\right)^2 + C(1-\beta)}$ and $\tilde{\xi}(\bar{b}, C, \rho, \beta) = \frac{\rho C(1-\beta)}{\bar{b}} + \sqrt{C(1-\beta) + \left(\frac{\rho C(1-\beta)}{\bar{b}}\right)^2}$.

It is interesting that this increase in user activity with increased ξ (for sufficiently small ξ) is driven by the benefits that users obtain from tailored services (captured by $a_i\xi\rho$ in Eq. 1's utility function). In fact, if users do not benefit from their information being accessible ($\rho = 0$) then user activity decreases with increasing ξ for *any* data storage policy. Yet regardless of the value of ρ, accessible information $\bar{a}^*(C,\xi)\xi$ first increases and then decreases with ξ.

As suggested in our discussion preceding Proposition 2, the comparative statics for consumer surplus is identical to the comparative statics for average user activity. Similarly, the comparative statics for the attack rate (as a function of ξ) is identical to the corresponding comparative statics for the amount of information stored.

3.2 Optimal design

Having accounted for the reactions of users and adversaries, we can now express the digital company's data policy design problem in the following form:

$$\max_{(C,\xi)\in\mathbb{R}_+\times[0,1]} \Pi(\xi,C) = \alpha\bar{a}^*(C,\xi) + \hat{\alpha}\xi\bar{a}^*(C,\xi) - \psi C$$

$$\text{s.t. } \bar{a}^*(C,\xi) \text{ as given by Eq. (3).}$$

Let

$$\xi(\alpha,\hat{\alpha},\rho,\bar{b},\psi) = \frac{1}{2\hat{\alpha}\rho}\left(-\alpha\rho - \bar{b}\hat{\alpha} + |\bar{b}\hat{\alpha} - \alpha\rho|\sqrt{\frac{\psi}{\psi-\hat{\alpha}\rho}}\right); \tag{5}$$

$$C(\xi,\alpha,\hat{\alpha},\rho,\bar{b},\psi) = \frac{1}{1-\beta}\left(-\xi^2 + \xi\sqrt{\frac{1}{\psi}(\bar{b}+\rho\xi)(\alpha+\hat{\alpha}\xi)}\right). \tag{6}$$

Theorem 1 (Optimal Design)**.** *There exist cost-of-protection thresholds $0 < \psi_L < \psi_H$ such that the following statements hold.*

(i) *For $\psi \le \psi_L$, the digital business stores all user information, $\xi^* = 1$, and sets a protection level $C^* = C(1,\ldots)$.*

(ii) *For $\psi \in (\psi_L, \psi_H)$, the digital business stores only some user information, $\xi^* = \xi(\cdots) \in (0,1)$, and sets a protection level $C^* = C(\xi^*,\ldots)$.*

(iii) *For $\psi \ge \psi_H$, the digital business stores no information, $\xi^* = 0$, and sets a protection level*

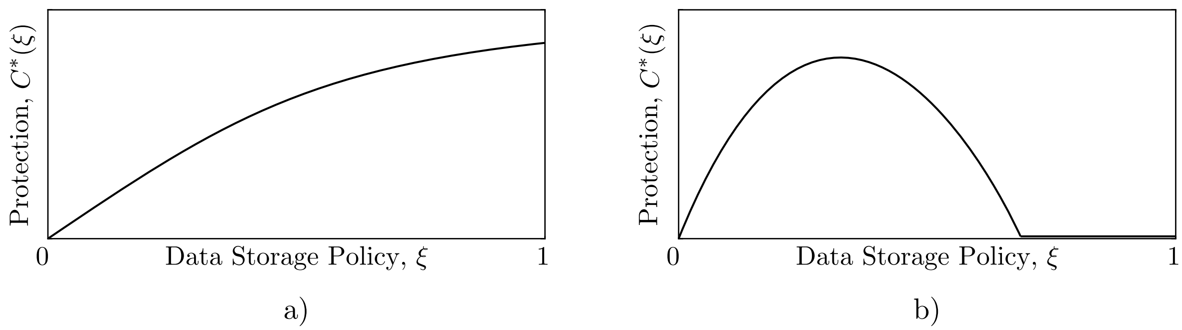

Figure 2: Typical behavior of optimal data protection level $C^*(\xi)$ for a given data storage level ξ in the case of low protection costs (panel a) and high protection costs (panel b).

$C^* = 0$.

Moreover, the optimal amount of data storage (ξ^*) and the optimal extent of data protection (C^*) are strategic complements in equilibrium.

To develop intuition, it is helpful to consider the optimal choice of protection for an arbitrary data storage policy ξ (see Lemma 1 in Appendix A). Figure 2 plots the digital company's optimal level of $C^*(\xi)$ for a given ξ. The figure's left-hand side (panel a) corresponds to the case where protection costs ψ are low; in this case, investing in protection increases with ξ. So for low values of ψ, our two instruments—the extent of protection and the amount of stored information—are complements.

The right-hand side of Figure 2 (panel b) corresponds to the case of intermediate/high protection costs. Much as in the low-cost case, if ξ is small then the optimal protection level $C^*(\xi)$ increases with ξ. Yet when ξ becomes large (i.e., a higher fraction of data is stored), the protection level actually declines with an increase in ξ. For large enough ξ, the digital company's best policy may even be to refrain from investing in any protection. In that case, the business's two instruments are substitutes. Two effects underlie this outcome. First, it is clear from Eq. 4 that a unit increase in the protection level C has the least influence on adversarial activity when there is little information stored. This fact, when coupled with the observation from Proposition 2 that the least information is stored when ξ is either low or high, implies that increasing the level of data protection is least effective when ξ is either low or high. Second, it follows from the user's utility (Eq. 1) that the direct damage from adversarial activity increases quadratically with ξ (to see this, plug the expression for $\omega^*(C,\xi)$ into the damage expression $\omega a_i \xi$ to obtain $a_i[\xi^2 \bar{a}^*(C,\xi)/C]$). In other words: if the business seeks to keep its data on user activity intact, then higher levels of information storage (ξ) require quadratically higher data protection, which is costly. That cost, when combined with the ineffectiveness of data protection at high levels of ξ, may lead the digital business to forgo protection altogether in this region of very high ξ.

Thus we have shown that, for an exogenous information storage level ξ, the design parameters C and ξ might be complements or substitutes. However, in an optimal design, the data protection policy C and data storage policy ξ are always complements. It is intuitive that, if increasing ξ results in the platform finding it too costly to protect information and thus deciding to lower the protection level, then ξ is not optimal. To see why, first note that C and ξ are substitutes when ξ is large. But in that case, both the total quantity of stored information and average user activity are decreasing in ξ (see Proposition 2 and Figure 1); the implication is that the company will benefit from reducing ξ even while holding the protection level fixed.

Theorem 1 is in line with this intuition. When the cost of adequate protection is too high, the digital business does not find it profitable to protect information and therefore sets $C^* = 0$. Because C and ξ are complements at the optimum, the business also chooses not to store any data ($\xi^* = 0$); hence it resorts to maintaining user activity at the level $\bar{b}/(1-\beta)$ and relying solely on transactions for revenue. It is easy to see that, in the absence of revenues from user activity ($\alpha = 0$), the threshold ψ_H approaches infinity and so a business will never move to the regime of no information storage and no data protection.[21]

Yet when the cost of protection is intermediate, the business finds a middle ground: some level of information storage and sufficient protection to generate returns from the information stored. Finally, if the cost of protection is low then a digital business can store all information and effectively cap adversarial activity by increasing data protection. Figure 3 plots the company's optimal design choice as a function of ψ.

To gain further intuition, we apply Theorem 1 to the two simple cases of advertisement- and transaction-driven companies—which, it may be recalled, exemplify the two most commonly observed business models. For ease of presentation, we examine the limiting case of $\bar{b} \to 0$, or the case in which users' stand-alone benefits are negligible. This approach should not affect comparisons

[21] This statement holds unless, simultaneously with $\alpha = 0$, also $\hat{\alpha} = 0$ (in which case the business makes zero profits) or $\bar{b} \to 0$ (in which case users receive no stand-alone benefits; see Corollary 2 to follow). In either of these cases, the business may prefer to set $\xi^* = 0$ and $C^* = 0$.

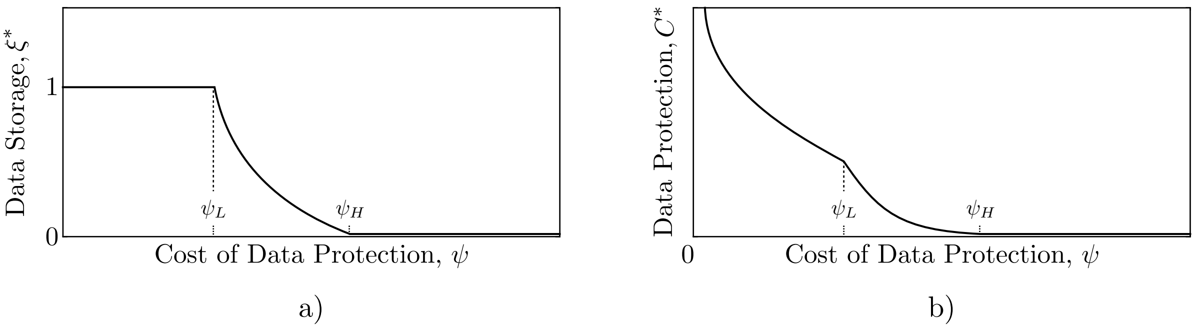

Figure 3: Optimal data storage policy ξ^* (panel a) and optimal data protection policy C^* (panel b) as a function of the cost ψ of protection; the points ψ_L and ψ_H are defined in Theorem 1.

between these two types of businesses because we have no reason to believe that they are inherently different in terms of stand-alone benefits.

Corollary 1 (Transaction-Driven Companies, $\alpha > 0$ and $\hat{\alpha} = 0$)**.** *If $\bar{b} \to 0$ then the optimal data policy design of a transaction-driven digital business, and the subsequent equilibrium, take one of the following forms.*

(i) *If $\psi \leq \alpha\rho/4$ then the business stores all information, $\xi^* = 1$, and chooses protection level $C^* = (1-\beta)^{-1}\left[\sqrt{\rho\alpha/\psi} - 1\right]$; user activity is then $\bar{a}^*(C^*, \xi^*) = C^*\sqrt{\rho\psi/\alpha}$ and the attack rate is $\omega^*(C^*, \xi^*) = \sqrt{\rho\psi/\alpha}$.*

(ii) *If $\psi > \alpha\rho/4$ then the business stores some user information, $\xi^* = \alpha\rho/4\psi$ and chooses protection level $C^* = (1-\beta)^{-1}(\xi^*)^2$; user activity is then $\bar{a}^*(C^*, \xi^*) = (\rho/2(1-\beta))\xi^*$ and the attack rate is $\omega^*(C^*, \xi^*) = \rho/2$.*

A transaction-driven digital business has incentives to store information because the platform can use it to offer services that create value for users, which in turn stimulates user activity. The level of information stored is a weakly increasing function of ρ and a weakly decreasing function of ψ. For example, if ρ increases—reflecting perhaps a new algorithm that allows for more efficient matches—then the digital business will relax its storage policy (i.e., increase ξ^*) and also increase the level of protection C^*. As a result, average users' activity increases, the attack rate declines, and the amount of stored information increases. The company does not change its storage policy when the strength of network effects changes (e.g., when β increases) but it does then increase its level of protection, which (as before) leads to increased average user activity.

Corollary 2 (Advertisement-Driven Companies, $\alpha = 0$ and $\hat{\alpha} > 0$)**.** *If $\bar{b} \to 0$ then the optimal data policy design of an ad-driven digital business, and the subsequent equilibrium, take one of the following forms.*

(i) *If $\psi < \hat{\alpha}\rho$ then the company stores all information, $\xi^* = 1$, and chooses protection level $C^* = (1-\beta)^{-1}\big[\sqrt{\rho\hat{\alpha}/\psi} - 1\big]$; user activity is then $\bar{a}^*(C^*, \xi^*) = C^*\sqrt{\rho\psi/\hat{\alpha}}$ and the attack rate is $\omega^*(C^*, \xi^*) = \sqrt{\rho\psi/\hat{\alpha}}$.*

(ii) *If $\psi \geq \hat{\alpha}\rho$ then the company is inactive; hence it stores no information ($\xi^* = 0$) and chooses a zero level of protection ($C^* = 0$), which leads to a zero user activity ($\bar{a}^*(C^*, \xi^*) = 0$).*

The policy design choice of an advertisement-driven digital business is extreme: as long as the company is active, it stores all users' information.[22] Therefore, the company will need to secure that information and so invests a positive amount in protecting it. We posit that the company is more incentivized to invest in protection when users benefit more from the information stored (higher ρ) and/or when network effects are stronger (higher β). So for companies with large ρ and β, we predict larger investments in data protection, lower attack rates, and higher average activity levels.

As a final observation, we consider the implication of these two corollaries for the evolution of the data policy of an ad-driven digital business over its life cycle. The initial phase of most such companies is delicate because their survival depends on building a solid and active customer base. For that reason, the initial focus of an ad-driven company is on increasing user activity. As mentioned in Section 2.5, this performance indicator is used to determine the company's valuation, to raise capital, and thus to accelerate growth. So a digital business in its initial phase acts as if α is large and $\hat{\alpha} = 0$. After establishing a user base and consolidating users' network effects, the company enters a more mature phase and its business model changes to monetizing the information it stores from users' activity. For an advertisement-driven company, this means offering targeted ad services. The company then acts as if $\alpha = 0$ and $\hat{\alpha}$ is large.

The foregoing characterization suggests that a move from the initial phase to the mature phase comes with a discrete change in the company's data policy. In the case of an intermediate value of ψ, for instance, a digital business will move from partial information storage to complete information storage—accompanied by a discrete increase in the data protection level C.

[22] Note that an ad-driven company stores a nonzero (but strictly smaller than 1) fraction of its users' information if $\psi \geq \hat{\alpha}\rho$ and users' stand-alone benefits are nonzero (i.e., if $\bar{b} > 0$.

4 Welfare

Although digital businesses account for a significant share of all transactions, the impact of digital privacy issues on consumers is not yet fully understood. One way of measuring such effects is by evaluating the direct loss due to adversarial uses of digitally stored data; in our model, that loss is expressed as $\mathcal{D} = \int \omega^* a_i^* \xi \, di = \omega^* \bar{a}^* \xi$. A January 2019 BBC report states that UK victims of cybercrime lose £190,000 *per day*.[23] Since the UK population is about 66 million, each individual suffers (on average) a daily loss of about £0.003, or 3.5 US pennies.

However, these numbers may underestimate the welfare loss induced by the adversaries. The reason is that an adversary, apart from inflicting direct losses on users, also induces them to change their sharing behaviors and forces digital businesses to alter the design of their data policy. Note that, for a fixed data policy design (C, ξ), an absence of adversaries will lead users to increase their activity levels. Moreover, if the risk of adversaries is removed then the digital company's optimal design becomes $(C, \xi) = (0, 1)$, which allows for full data utilization with no adverse effects and also saves on data protection costs.

In order to evaluate the full effect of adversarial activity on consumers, we define the *adversarial loss multiplier* $\mathcal{M}$ as the ratio of total adversary-caused consumer surplus loss to the direct loss that adversaries inflict:

$$\mathcal{M} = \frac{\text{CS}_{\text{no adversaries}} - \text{CS}^*}{\mathcal{D}}, \tag{7}$$

where $\text{CS}_{\text{no adversaries}}$ is the equilibrium consumer surplus in a game that excludes adversaries. Our next proposition characterizes the loss multiplier and establishes its significance.[24]

Proposition 3 (Adversarial Loss Multiplier)**.** *The total decrease in consumer surplus due to the presence of adversaries is equal to $\mathcal{M} \cdot \mathcal{D}$, where $\mathcal{M}$ is the adversarial loss multiplier:*

$$\mathcal{M} = \frac{C^*}{\xi^*}\left(\left(\frac{\bar{b}+\rho}{1-\beta}\right)^2 \frac{1}{(\bar{a}^*)^2} - 1\right) \geq \frac{2}{1-\beta}.$$

Thus the total loss to users from the presence of adversaries is more than double users' equilibrium

[23] "UK Cyber-crime Victims Lose £190,000 a Day", *BBC News*, 27 January 2019.

[24] We remark that $\text{CS}_{\text{no adversaries}}$ is equal to the consumer surplus when a digital business sets $C \to \infty$, which fully protects users' data from adversaries and maximizes CS. However, that approach can never be a part of a digital company's optimal design when $\psi > 0$.

direct loss due to adversarial use of stored data. Furthermore, the loss multiplier increases with the strength of network effects and is unbounded even when those effects are finite. In fact, total welfare loss is greater even than the loss of consumer surplus. The additional loss is due to the effect of adversaries on the digital company's profit by altering users' activity and increasing its protection costs.

Proposition 3 shows that the direct damage $\mathcal{D}$ from adversarial activity severely underestimates the effects that adversaries have on consumer surplus and total welfare. We next turn to evaluating the overall consumer surplus and profits attainable by advertisement- and transaction-driven companies. As discussed in Section 5, this examination sheds light on an ongoing public debate about which revenue model is more aligned with consumer preferences.

Proposition 4. *Consider an advertisement-driven company* $(\Pi = \hat{\alpha}\bar{a}\xi - \psi C)$ *and a transaction-driven company* $(\Pi = \alpha\bar{a} - psiC)$, *and assume that* $\alpha = \hat{\alpha}$. *Then the following statements hold.*

(i) *Profit is always weakly lower for the ad-driven company and is strictly lower when data storage costs are sufficiently high.*

(ii) *Consumer surplus is lower (resp. higher) for the ad-driven company when data protection costs are high (resp. intermediate), and it is equal across business models when data protection costs are low.*

It is clear that digital businesses differ in terms of more than their revenue model. Therefore, Proposition 4(i) should not be interpreted as a claim that digital businesses will always do better if they choose a transaction-based revenue model. Instead, it reflects that advertisement-driven companies are more vulnerable to adversarial activities than are transaction-driven companies. One can see this by observing that, in the environment of Proposition 4, the values of $\Pi_{\text{no adversaries}}$ and $\text{CS}_{\text{no adversaries}}$ are the same for ad-driven and transaction-driven companies. It follows that the difference in profits and in consumer surplus can be fully attributed to the presence of adversaries. In Section 5 we argue that a successful digital business's data protection costs are likely to be in the intermediate range. Part (ii) of Proposition 4 suggests that, in a world with adversaries, an advertisement-driven business model may be more aligned with consumer needs than is a transaction-driven model.

The logic underlying Proposition 4 is as follows. When data protection costs are *low*, advertisement- and transaction-driven companies can each afford to store all of the available data ($\xi = 1$). At that point, their respective data storage policies are such that both choose their data protection level to maximize $\alpha\bar{a} - \psi C$; hence both types of businesses choose the same data protection level C, have identical profits, and provide the same utilities to users. For *intermediate* data protection costs, both types of companies must scale down their data protection because it is now more expensive. To discourage increased adversarial activity—a natural consequence of reduced protection—both companies also choose to reduce the fraction of data they store. Yet this reduction in ξ is more damaging to an ad-driven company's profits, which are *directly* affected by the reduced average user activity and also by the reduced fraction of data stored. Hence the ad-driven company chooses to store more data and is willing to spend more on protecting it. This strategy mitigates the ad-driven company's relatively greater loss in consumer surplus. The effect that reducing ξ has on profits carries over to the case of higher data protection costs.

When data protection costs become sufficiently *high*, the advertisement-driven company essentially "gives up" and sets $(C, \xi) = (0, 0)$. At that point the ad-driven company makes zero profits, and its users are affected because they no longer receive any of the benefits associated with their data being used.[25] Figure 4 (panel a) illustrates how consumer surplus differs between the two types of firms.

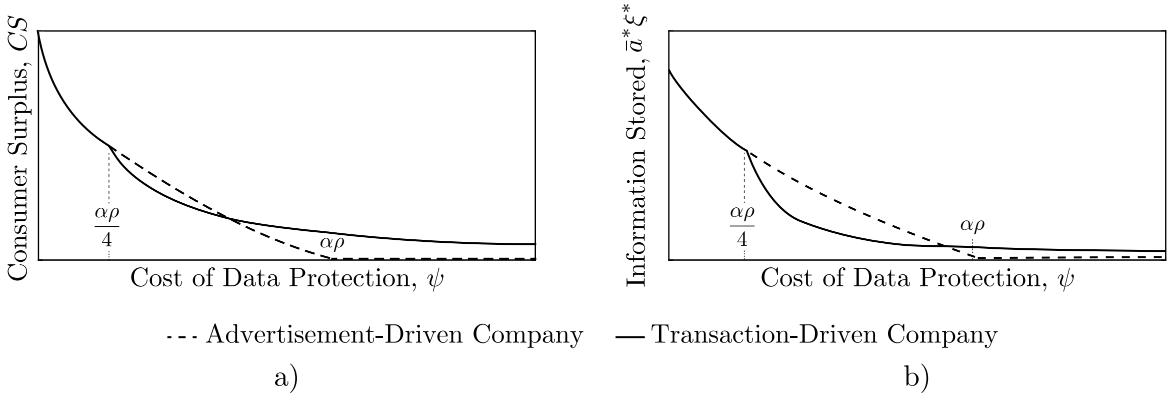

Figure 4: Typical behavior of consumer surplus CS (panel a) and total amount of information stored $\bar{a}^*(C^*, \xi^*)\xi^*$ (panel b)—as a function of the cost ψ of protection—for advertisement-driven (dashed curves) and transaction-driven (solid curves) companies.

[25] The results for high data protection costs become even more pronounced if we assume, as seems reasonable, that a company making no profits goes out of business.

5 Discussion and Conclusions

In this paper we develop a framework to capture the major factors that determine a digital company's data policies. This framework provides a useful perspective and terminology for discussing questions and issues faced by users, digital businesses, and regulatory authorities. Many of these questions have been raised by the media, public officials, firms, and users of digital businesses in the public domain. In a recent *Wall Street Journal* op-ed, for example, Facebook CEO Mark Zuckerberg acknowledges and responds to several such questions.[26] We begin this section by demonstrating how our model facilitates expressing and analyzing these questions (see Appendix B.8 for our formal analysis). In particular,

> *"there's the important question of whether the advertising model encourages companies like ours [Facebook] to use and store more information than we otherwise would."*

After comparing Corollaries 1 and 2 while assuming $\alpha = \hat{\alpha}$, we conclude that—when the cost of data protection is intermediate—advertisement-driven companies design less stringent data storage policies (i.e., they set a higher ξ) than do their transaction-driven counterparts. Figure 4 (panel b) illustrates that this design choice shapes users' and adversaries' incentives in such a way that the amount of available information stored on the company's servers ($\bar{a}\xi$) is greater for the ad-driven company in this intermediate range of protection costs. When data protection costs are high, however, this relationship is reversed. In that case, the ad-driven company's cost of keeping adversaries at bay outweighs the profits derivable from users' stored data; hence the digital business may simply abandon the market.

> *"Some worry that ads create a misalignment of interests between us [Facebook] and people who use our services."*

Because advertisement-driven companies have more permissive data storage policies, it is tempting to conclude that such business models lead to the company's and its users' interests not being aligned. However, users also extract benefits from the digital business's use of the data: they may well enjoy better matches with content, products, and peers as well as more relevant ads. The merit of Proposition 4 is to point out that, in the case of intermediate data protection costs,

[26] "The Facts About Facebook", *Wall Street Journal*, 24 January 2019.

consumer surplus is higher for users of an advertisement-driven firm than for users of a comparable transaction-driven firm.

However, our results do show that, regardless of the specific revenue model, there is some basic misalignment of interest between businesses and users. For example, holding all else equal, users always prefer a stronger data protection policy. Governments around the world responded with added regulation, with the most notable being the EU General Data Protection Regulation (GDPR).[27] While regulation in different countries take different forms, a common component is the requirement that digital businesses improve the ways they protect users' data. Our analysis shows that when businesses are strategic, data protection and data collection are complements. Therefore, in equilibrium, increasing the data protection requirements (i.e., requiring C to be above some binding threshold) may lead businesses to collect and store more data. Whether users benefit from this change or not depends on how much added value users have for the business' use of their data and how effective data protection is in deterring adversaries, as well as on the revenue model of the digital business. Either way, as our front matter quote of Eduard Snowden suggests, data protection and data collection/storage policies cannot be treated as independent when designing regulations to enhance consumer surplus and aggregate welfare.

Interestingly, the basic idea that users extract benefits from a digital company's use of their data has been the focus of increased recent attention. Turow et al. (2015, 2018), as well as the *New York Times*,[28] reference a large-scale survey when making two claims. First, user decisions about sharing their data are not entirely rational. Many users do not understand privacy policies and data practices, sharing their data merely because it's seemingly too difficult to "opt out". Second, nearly six in ten users would prefer *not* to see targeted ads, news, and/or discounts—and that proportion increases when users learn just how they are tracked by websites. Turow et al. [2018] respond to the claim, frequently made by marketers and tech C-suite executives, that users prefer relevant ads to random ads.[29]

Our framework allows us to explore further the meaning of Turow et al. (2018) results and to pose new questions for empirical work. A possible interpretation is that users find no benefit

[27] https://eugdpr.org/

[28] See "Sharing Data for Deals? More Like Watching It Go with a Sigh", *New York Times*, 28 December 2018; and "Mark Zuckerbeg's Delusion of Consumer Consent", *New York Times*, 31 January 2019.

[29] For example, Zuckerberg argues (in his *Wall Street Journal* op-ed cited previously) that "[p]eople consistently tell us that if they're going to see ads, they want them to be relevant."

from their information being stored; in the language of our model, this amounts to $\rho = 0$ or even $\rho < 0$. Alternatively, users may wrongly suppose that data storage is for the purposes of targeted advertising only; that misconception is suggested by the difficulty (alluded to previously) of understanding privacy practices and also by the observation that targeted ads are rejected at higher rates when users are fully informed of the data storage practices that enable such targeting. In that case, it is possible that the fraction of ρ attributed to targeted advertising is, in itself, too small to justify the digital company's "surveillance" of users' Web activity. Another possibility is that users dislike ad targeting under any circumstances: in particular, regardless of whether eliminating it would affect the data storage policies of digital businesses such as Facebook. Finally, consumers may be sophisticated enough to prefer the equilibrium that would result from a ban on targeted advertising—given that, as emphasized by Corollaries 1 and 2 and by the foregoing discussion, such a ban would likely result in a lower fraction of data being stored. Analyzing equilibria that include naïve or irrational users is beyond the scope of this paper; even so, our framework is flexible enough that it could be modified to incorporate different types of users, including some forms of naïveté.

Another interesting question that transcends our analysis pertains to property rights over individual-level data. Arrieta-Ibarra et al. [2018] make the case that users should be paid for their data as if that data were labor, and Ichihashi [2019a] explores a scenario in which competing data brokers compensate users for their data. Emerging work in the marketing literature seeks to evaluate users' valuation of privacy via empirical and experimental approaches (see Lin 2019 and the references therein). We defer to future work those extensions of our framework that accommodate payment for data usage.

References

J. M. Abowd and I. M. Schmutte. An economic analysis of privacy protection and statistical accuracy as social choices. *American Economic Review*, 109(1):171–202, 2019.

D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar. Privacy-constrained network formation. *Games and Economic Behavior*, 105:255–275, 2017.

A. Acquisti, C. Taylor, and L. Wagman. The economics of privacy. *Journal of Economic Literature*, 54(2):442–92, 2016.

S. N. Ali and R. Bénabou. Image versus information: Changing societal norms and optimal privacy. Technical report, National Bureau of Economic Research, 2016.

I. Arrieta-Ibarra, L. Goff, D. Jiménez-Hernández, J. Lanier, and E. G. Weyl. Should we treat data as labor? moving beyond" free". In *aea Papers and Proceedings*, volume 108, pages 38–42, 2018.

D. Bergemann and A. Bonatti. Markets for information: An introduction. *Annual Review of Economics*, 11, 2019.

F. Bloch and N. Quérou. Pricing in social networks. *Games and economic behavior*, 80:243–261, 2013.

O. Candogan and K. Drakopoulos. Optimal signaling of content accuracy: Engagement vs. misinformation. *Misinformation (October 11, 2017)*, 2017.

O. Candogan, K. Bimpikis, and A. Ozdaglar. Optimal pricing in networks with externalities. *Operations Research*, 60(4):883–905, 2012.

R. Casadesus-Masanell and A. Hervas-Drane. Competing with privacy. *Management Science*, 61(1): 229–246, 2015.

V. Conitzer, C. R. Taylor, and L. Wagman. Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science*, 31(2):277–292, 2012.

R. Cummings, K. Ligett, M. M. Pai, and A. Roth. The strange case of privacy in equilibrium models. *arXiv preprint arXiv:1508.03080*, 2015.

C. Dwork and A. Roth. The algorithmic foundations of differential privacy. *Foundations and Trends® in Theoretical Computer Science*, 9(3–4):211–407, 2014.

I. P. Fainmesser and A. Galeotti. Pricing network effects. *The Review of Economic Studies*, 83(1): 165–198, 2015.

I. P. Fainmesser and A. Galeotti. Pricing network effects: Competition. *American Economic Journal*, (Forthcoming), 2019.

A. Galeotti and S. Goyal. Influencing the influencers: a theory of strategic diffusion. *The RAND Journal of Economics*, 40(3):509–532, 2009.

A. Ghosh and A. Roth. Selling privacy at auction. *Games and Economic Behavior*, 91:334–346, 2015.

E. L. Glaeser and J. Scheinkman. Non-market interactions. Technical report, National Bureau of Economic Research, 2000.

A. Goldfarb and C. E. Tucker. Privacy regulation and online advertising. *Management science*, 57 (1):57–71, 2011.

R. Gradwohl. Information sharing and privacy in networks. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pages 349–350. ACM, 2017.

S. Ichihashi. Non-competing data intermediaries. 2019a. URL `https://shota2.github.io/research/data.pdf`.

S. Ichihashi. Online privacy and information disclosure by consumers. *Forthcoming at American Economic Review*, 2019b.

O. Jann and C. Schottmüller. An informational theory of privacy. 2016.

B. Koh, S. Raghunathan, and B. R. Nault. Is voluntary profiling welfare enhancing? *Management Information Systems Quarterly, Forthcoming*, 2015.

T. Lin. Valuing intrinsic and instrumental preferences for privacy. *Available at SSRN 3406412*, 2019.

D. Mayzlin. Managing social interactions. In *The Oxford Handbook of the Economics of Networks*. 2016.

R. Montes, W. Sand-Zantman, and T. Valletti. The value of personal information in online markets with endogenous privacy. *Management Science*, 65(3):1342–1362, 2018.

Y. Papanastasiou. Fake news propagation and detection: A sequential model. *Available at SSRN 3028354*, 2018.

Q. Shen and J. Miguel Villas-Boas. Behavior-based advertising. *Management Science*, 64(5): 2047–2064, 2017.

J. Turow, M. Hennessy, and N. Draper. The tradeoff fallacy: How marketers are misrepresenting american consumers and opening them up to exploitation. *Available at SSRN 2820060*, 2015.

J. Turow, M. Hennessy, and N. Draper. Persistent misperceptions: Americans' misplaced confidence in privacy policies, 2003–2015. *Journal of Broadcasting & Electronic Media*, 62(3):461–478, 2018.

Appendix

A Additional Results

Lemma 1 (Optimal Data Protection Policy $C^*(\xi)$ as a Function of Data Storage Policy ξ)**.** *Given company's data storage policy ξ, the company sets $C^*(\xi) = \max\{C_h(\xi), 0\}$. Optimal protection policy is non-zero if either protection is cheap ($\psi < \hat{\alpha}\rho$) or company doesn't store much information: $\xi < \min\{1, \frac{\hat{\alpha}\bar{b}+\rho\alpha+\sqrt{(\hat{\alpha}\bar{b}-\alpha\rho)^2+4\alpha\bar{b}\psi}}{2(\psi-\hat{\alpha}\rho)}\}$. Here $C_h(\xi) = \frac{1}{1-\beta}\left(-\xi^2 + \xi\sqrt{\frac{1}{\psi}(\bar{b}+\rho\xi)(\alpha+\hat{\alpha}\xi)}\right)$.*

Lemma 2 (Optimal Data Storage Policy $\xi^*(C)$ as a Function of Data Protection Policy C)**.** *Given company's data protection policy C, the company stores:*

- *part of the information $\xi^*(C) = -\kappa(C) + \sqrt{\kappa(C)^2 + C(1-\beta)} < 1$ (increasing in C) if data protection policy level is low $C < C_l$;*
- *all the information ($\xi^*(C) = 1$) if data protection policy level is high $C > C_l$.*

Here $\kappa(C) = \frac{\bar{b}\alpha-\hat{\alpha}\rho C(1-\beta)}{\hat{\alpha}\bar{b}+\alpha\rho}$ and $C_l = \frac{1}{1-\beta}\frac{2\alpha\bar{b}+\alpha\rho+\hat{\alpha}\bar{b}}{2\hat{\alpha}\rho+\alpha\rho+\hat{\alpha}\bar{b}}$.

Lemma 3. *Digital business's data storage policy (ξ) and data protection policy (C) are complements for sufficiently high data protection levels C and substitutes otherwise. Mathematically, $\frac{\partial^2\Pi(C,\xi)}{\partial\xi\partial C} > 0$ if $C(1-\beta) > \frac{2\alpha\bar{b}+\hat{\alpha}\bar{b}\xi+\alpha\rho\xi}{\alpha(2\bar{b}+3\rho\xi)+\hat{\alpha}\xi(3\bar{b}+4\rho\xi)}\xi^2$ (an increasing convex function of ξ).*

B Proofs

B.1 Proof of Proposition 1 (Page 12).

Given users' expectation of the attack rate ω, there exists a unique response of the users. Existense (sufficient condition) follows from Glaeser and Scheinkman [2000]: $\exists_{\widetilde{a}\geq 0}\forall_{\bar{a}\leq\widetilde{a}}\frac{\partial U_i(a_i,\bar{a})}{\partial a_i}\Big|_{a_i=\widetilde{a}} < 0$ or $\exists_{\widetilde{a}\geq 0}\forall_{\bar{a}\leq\widetilde{a}}\ b_i + \beta\bar{a} + \xi[\rho-\omega] - \widetilde{a} < 0$. Because $\rho, \omega, \xi \in [0,1]$ this is satisfied whenever $\exists_{\widetilde{a}\geq 0}\ b_i + \beta\widetilde{a} + 1 - \widetilde{a} < 0$ or $\exists_{\widetilde{a}\geq 0}\ b_i + 1 < \widetilde{a}\,(1-\beta)$, which can only be satisfied if $\beta < 1$.

Condition for the uniqueness of the users' response also follows from Glaeser and Scheinkman [2000]: $\forall_i \left|\frac{\partial^2 U_i}{\partial a_i \partial\bar{a}} / \frac{\partial^2 U_i}{\partial a_i^2}\right| < 1$ which is satisfied if $\beta < 1$.

Next, we derive our characterization of the unique response. Denote $\mathbf{a}_{-i}$ the activity choice that i conjecture about the other users. Then user i's best reply is $a_i = b_i + \beta\bar{a} - \omega\xi + \rho\xi$ where $\bar{a} = \int_j a_j dj$. In equilibrium users' expectation are correct and so $\int_i a_i di = \bar{b} + \beta\bar{a} - \omega\xi + \rho\xi = \bar{a}$ or

$\bar{a}(\omega) = \frac{\bar{b}+\rho\xi-\omega\xi}{1-\beta}$. Such response induces adversaries with $\bar{a}(\omega)\xi \geq \gamma C$ to be active. Therefore, the induced attack rate is $\bar{a}(\omega)\xi/C$ which should be consistent with the initial belief ω. The expressions for $\bar{a}(C,\xi)$ and $\omega^*(C,\xi)$ then follow.

B.2 Proof of Proposition 2 (Page 12).

Derivative of $\bar{a}^*(C,\xi)$ wrt C is $\frac{\xi^2(\bar{b}+\rho\xi)}{(C(1-\beta)+\xi^2)^2} \geq 0$; wrt ξ it is $\frac{C(\rho C(1-\beta)-2\bar{b}\xi-\rho\xi^2)}{(C(1-\beta)+\xi^2)^2}$ the sign of which is defined by the sign of $\rho C(1-\beta) - 2\bar{b}\xi - \rho\xi^2$. At $\xi = 0$ the latter expression is positive and has negative derivative. It changes sign to negative only once for $\xi > 0$ at $\hat{\xi}(\bar{b}, C, \rho, \beta)$ which can be found as the largest solution to the corresponding quadratic equation. Similarly, derivative of $\bar{a}^*(C,\xi)\xi$ wrt ξ is $\frac{C(\bar{b}(C(1-\beta)-\xi^2)+2C(1-\beta)\rho\xi)}{(C(1-\beta)+\xi^2)^2}$ which has sign of $-\bar{b}\xi^2 + 2\rho\xi C(1-\beta) + \bar{b}C(1-\beta)$. The latter expression is positive and has positive derivative at $\xi = 0$, it changes sign only once at $\tilde{\xi}(\bar{b}, C, \rho, \beta)$ which is the largest solution to the corresponding quadratic equation. Finally, $\tilde{\xi}(\ldots) > \hat{\xi}(\ldots)$ holds trivially when $\bar{b}/\rho < \rho C(1-\beta)/\bar{b}$, otherwise, rewrite it as $\frac{\rho C(1-\beta)}{\bar{b}} + \frac{\bar{b}}{\rho} > \sqrt{\left(\frac{\bar{b}}{\rho}\right)^2 + C(1-\beta)} - \sqrt{\left(\frac{\rho C(1-\beta)}{\bar{b}}\right)^2 + C(1-\beta)}$, RHS is positive when $\bar{b}/\rho \geq \rho C(1-\beta)/\bar{b}$, taking square of the both sides and rearranging, we can show that this inequality holds.

B.3 Proof of Lemma 1 (Page 28).

Plug expression for $\bar{a}^*(C,\xi)$ into Eq. 2. It is easy to verify that profit function is concave in C. Function $C_h(\xi) = \frac{1}{1-\beta}\left(-\xi^2 + \xi\sqrt{\frac{1}{\psi}(\bar{b}+\rho\xi)(\alpha+\hat{\alpha}\xi)}\right)$ solves the FOC. Differentiating $C_h(\xi)$ wrt ξ, we get:

$$(1-\beta)\frac{dC_h(\xi)}{d\xi} = -2\xi + \sqrt{\frac{1}{\psi}(\bar{b}+\rho\xi)(\alpha+\hat{\alpha}\xi)} + \frac{\xi}{2\psi}\frac{\rho(\alpha+\hat{\alpha}\xi)+\hat{\alpha}(\bar{b}+\rho\xi)}{\sqrt{\frac{1}{\psi}(\bar{b}+\rho\xi)(\alpha+\hat{\alpha}\xi)}}$$

It is positive at $\xi = 0$. Developing equality $\frac{dC_h(\xi)}{d\xi} = 0$, we get that it is equivalent to:

$$\xi^4 16\hat{\alpha}\rho(\hat{\alpha}\rho-\psi) + 8\xi^3(\hat{\alpha}\bar{b}+\alpha\rho)(3\hat{\alpha}\rho-2\psi)+$$
$$\xi^2(9\hat{\alpha}^2\bar{b}^2 + 34\alpha\hat{\alpha}\bar{b}\rho + 9\alpha^2\rho^2 - 16\alpha\bar{b}\psi) + 12\xi\alpha\bar{b}(\hat{\alpha}\bar{b}+\alpha\rho) + 4\alpha^2\bar{b}^2 = 0$$

Consider two cases. First, let $\hat{\alpha}\rho > \psi$ then also $3\hat{\alpha}\rho > 2\psi$. Note also that $34\alpha\hat{\alpha}\bar{b}\rho - 16\alpha\bar{b}\psi > 0$. Then coefficients of the polynomial don't change sign, hence there are no positive roots and hence $\frac{dC_h(\xi)}{d\xi} > 0$, $\forall\xi \geq 0$. Since, $C_h(0) = 0$, then $C_h(\xi) > 0, \forall\xi > 0$ in this case.

Now let $\psi > \hat{\alpha}\rho$. In case $3\hat{\alpha}\rho < 2\psi$, irrespective of the sign of the coefficient in front of ξ^2, coefficients of the polynomial change sign only once, hence there is one positive root and thus $\frac{dC_h(\xi)}{d\xi}$ changes sign only once for $\xi > 0$. If $3\hat{\alpha}\rho > 2\psi$ then $34\alpha\hat{\alpha}\bar{b}\rho > 16\alpha\bar{b}\psi$ even in the worst case of $\psi = \frac{3}{2}\hat{\alpha}\rho$, hence coefficient in front of ξ^2 is positive. Thus, there is only one change of sign and hence also one root in this case. Thus, we conclude that when $\psi > \hat{\alpha}\rho$, derivative $\frac{dC_h(\xi)}{d\xi}$ changes sign only once for $\xi > 0$.

It is easy to see that $C_h(0) = 0$. Also, solving $C_h(\xi) = 0$ for $\xi > 0$ we obtain: $\xi = \frac{\hat{\alpha}\bar{b}+\rho\alpha+\sqrt{(\hat{\alpha}\bar{b}-\alpha\rho)^2+4\alpha\bar{b}\psi}}{2(\psi-\hat{\alpha}\rho)} > 0$ when $\psi > \hat{\alpha}\rho$. Knowing that $C_h'(\xi)$ changes sign only once if $\psi > \hat{\alpha}\rho$ and that it increases at $\xi = 0$, we conclude that $C_h(\xi) > 0$ for low enough ξ if $\psi > \hat{\alpha}\rho$. The result of lemma now follows.

B.4 Proof of Lemma 2 (Page 28).

Plug expression for $\bar{a}^*(C,\xi)$ into Eq. 2. FOC is: $-\xi^2(\hat{\alpha}\bar{b} + \alpha\rho) - 2\xi(\alpha\bar{b} - \hat{\alpha}\rho C(1-\beta)) + C(1-\beta)(\hat{\alpha}\bar{b} + \alpha\rho) = 0$. It has one positive root. Divide FOC by $\hat{\alpha}\bar{b} + \alpha\rho$, then we get the positive root $-\kappa(C) + \sqrt{\kappa(C)^2 + C(1-\beta)}$ with the notation for $\kappa(C)$ from Lemma's formulation. $-\kappa(C) + \sqrt{\kappa(C)^2 + C(1-\beta)}$ is increasing and concave in C. Indeed, it's derivative is wrt C:

$$\frac{\hat{\alpha}(1-\beta)\rho}{\hat{\alpha}\bar{b} + \alpha\rho} + \frac{(1-\beta)(\alpha^2\rho^2 + \hat{\alpha}^2(\bar{b}^2 + 2\rho^2 C(1-\beta)))}{2(\hat{\alpha}\bar{b} + \alpha\rho)^2\sqrt{C(1-\beta) + \kappa(C)^2}} > 0$$

Second derivative is negative. Solving inequality $-\kappa(C) + \sqrt{\kappa(C)^2 + C(1-\beta)} < 1$ we derive C_l. Finally, second-derivative of the profit function has the sign of $\xi^3 + 3\xi^2\kappa(C) - 3\xi C(1-\beta) - C(1-\beta)\kappa(C)$. Plugging $-\kappa(C) + \sqrt{\kappa(C)^2 + C(1-\beta)}$ into this expression, we obtain $-2(C(1-\beta) + \kappa(C)^2)(-\kappa(C) + \sqrt{\kappa(C)^2 + C(1-\beta)}) < 0$.

B.5 Proof of Lemma 3 (Page 28).

Cross-partial derivative of the profit function with activity level $\bar{a}^*(C,\xi)$ has the following form:

$$\frac{\partial^2\Pi(C,\xi)}{\partial\xi\partial C} = -\frac{\xi}{(\xi^2 + C')^3}\left(2\alpha\bar{b}(\xi^2 - C') + \rho\alpha\xi(\xi^2 - 3C') + \hat{\alpha}\bar{b}\xi(\xi^2 - 3C') - 4\hat{\alpha}\xi^2\rho C'\right)$$

Where $C' = C(1-\beta)$. The expression in the brackets is a decreasing linear function of C'. It is positive when $C = 0$. Cross-partial derivative is positive when $C' > \frac{2\alpha\bar{b}+\hat{\alpha}\bar{b}\xi+\alpha\rho\xi}{\alpha(2\bar{b}+3\rho\xi)+\hat{\alpha}\xi(3\bar{b}+4\rho\xi)}\xi^2$. The RHS is 0 at $\xi = 0$ and it is easy to verify that it is increasing and convex in ξ.

B.6 Proof of Theorem 1 (Page 14).

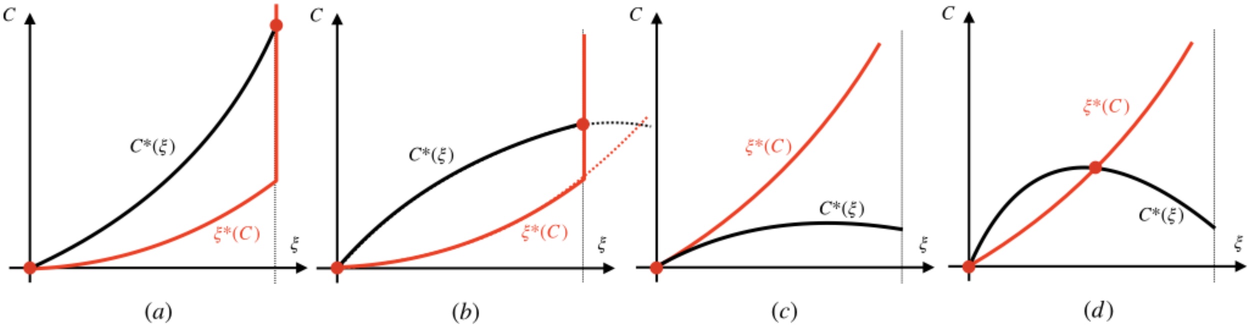

Figure 5: Functions $C^*(\xi)$ and $\xi^*(C)$ depending on ψ: (a) $\psi < \hat{\alpha}\rho$; (b) $\psi \in [\hat{\alpha}\rho, \psi_L]$; (c) $\psi > \psi_H$; (d) $\psi \in [\psi_L, \psi_H]$. Points ψ_L, ψ_H are defined in the proof of Theorem 1.

Lemmas 1 and 2 give us optimal data protection and data storage policies as a response to when one policy is being fixed. We know that $-\kappa(C) + \sqrt{\kappa(C)^2 + C(1-\beta)}$ is increasing and concave in C. Denote $C_{inv}(\xi)$ solution to $-\kappa(C) + \sqrt{\kappa(C)^2 + C(1-\beta)} = \xi$ wrt C (here $\kappa(C)$ is defined in Lemma 2):

$$C_{inv}(\xi) = \frac{1}{1-\beta}\frac{\xi[(\hat{\alpha}\bar{b}+\alpha\rho)\xi + 2\bar{b}\alpha]}{\hat{\alpha}\bar{b}+\alpha\rho+2\xi\hat{\alpha}\rho}$$

<u>Case $\psi < \hat{\alpha}\rho$</u>: We will show that $C_{inv}(\xi) < C_h(\xi)$ (where $C_h(\xi)$ is defined in Lemma 1) and they only intersect at $\xi = 0$ under this condition. We need to show:

$$\xi + \frac{(\hat{\alpha}\bar{b}+\alpha\rho)\xi + 2\bar{b}\alpha}{\hat{\alpha}\bar{b}+\alpha\rho+2\xi\hat{\alpha}\rho} < \sqrt{\frac{1}{\psi}(\bar{b}+\rho\xi)(\alpha+\hat{\alpha}\xi)}$$

Notice that LHS is $\frac{2(\bar{b}+\rho\xi)(\alpha+\hat{\alpha}\xi)}{\hat{\alpha}\bar{b}+\alpha\rho+2\xi\hat{\alpha}\rho}$. Developing, we need to show that $4\psi(\bar{b}+\rho\xi)(\alpha+\hat{\alpha}\xi) < (\hat{\alpha}\bar{b}+\alpha\rho+2\xi\hat{\alpha}\rho)^2$. LHS is the highest under $\psi = \hat{\alpha}\rho$. It is easy to check that inequality holds in this case and hence $C_{inv}(\xi) < C_h(\xi), \forall\xi > 0$. Thus, the only two points of intersection of best-responses $C^*(\xi)$ and $\xi^*(C)$ are $\xi = 0, C = 0$ and $\xi = 1, C = C_h(1)$. We evaluate company's profit at both points. The first point gives $\pi_0 = \frac{\alpha\bar{b}}{1-\beta}$, while the second gives $\pi_1(\psi) = \frac{1}{1-\beta}((\alpha+\hat{\alpha})(\bar{b}+\rho)+\psi-2\sqrt{\psi(\alpha+\hat{\alpha})(\bar{b}+\rho)})$ - decreasing in ψ on the interval $\psi < \hat{\alpha}\rho$ (derivative wrt ψ is $1-\sqrt{(\alpha+\hat{\alpha})(\bar{b}+\rho)/\psi}$). $\pi_1(0) > \pi_0$, also it is easy to verify that $\pi_1(\hat{\alpha}\rho) > \pi_0$, hence platform is choosing $\xi^* = 1, C^* = C_h(1)$. Figure 5 (a) illustrates this case.

Case $\psi > \hat{\alpha}\rho$: Solve $-\kappa(C_h(\xi)) + \sqrt{\kappa(C_h(\xi))^2 + C_h(\xi)(1-\beta)} = \xi$ for ξ. We get two non-negative roots: $\xi = 0$ and $\xi_i = \xi(\alpha, \hat{\alpha}, \rho, \bar{b}, \psi)$ defined in Theorem 1. We will establish first conditions when $\xi_i \in [0,1]$. We will consider case of $\bar{b}\hat{\alpha} > \alpha\rho$ (case $\bar{b}\hat{\alpha} < \alpha\rho$ is considered similarly and leads to the same result). $\xi_i > 0$ when $\psi \in [\hat{\alpha}\rho, \psi_H]$, where $\psi_H = \frac{(\hat{\alpha}\bar{b}+\alpha\rho)^2}{4\alpha\bar{b}}$ is a solution to $\xi_i = 0$ wrt ψ. Similarly, $\xi_i < 1$ when $\psi > \psi_L$, where $\psi_L = \frac{(\hat{\alpha}\bar{b}+\alpha\rho+2\hat{\alpha}\rho)^2}{4(\alpha+\hat{\alpha})(\bar{b}+\rho)}$ is a solution to $\xi_i = 1$ wrt ψ. Both $\psi_L > \hat{\alpha}\rho$ and $\psi_H > \hat{\alpha}\rho$, also $\psi_H > \psi_L$ (ξ_i decreases with ψ).

On the interval $\psi \in [\hat{\alpha}\rho, \psi_L]$, $\xi_i > 1$, hence on $\xi \in [0,1]$ we have $C_{inv}(\xi) < C_h(\xi)$. Thus, the two points of intersection of the best-response functions are $\xi = 0, C = 0$ and $\xi = 1, C = C_h(1)$. $\pi_1(\psi)$ decreases with ψ on $[\hat{\alpha}\rho, \psi_L]$ (derivative $\pi_1'(\psi)$ changes sign only once and at ψ_L it is negative). Plugging ψ_L into $\pi_1(\psi)$ defined above, we get $\pi_1(\psi_l) > \pi_0$. We thus conclude that $\pi_1(\psi) > \pi_0$ on $[\hat{\alpha}\rho, \psi_L]$ and thus $\xi^* = 1, C^* = C_h(1)$. Figure 5 (b) illustrates this case.

On the interval $\psi > \psi_H$, $C_{inv}(\xi) > C_h(\xi)$, $\forall \xi \in [0,1]$. Hence, the only point of intersection of the best-responses is $\xi^* = 0, C^* = 0$. Figure 5 (c) illustrates this case.

Finally, on the interval $\psi \in [\psi_L, \psi_H]$, there are two points of intersection of the best-responses: $\xi = 0, C = 0$ (delivering profit π_0) and $\xi = \xi_i \in [0,1], C = C_h(\xi_i)$ with profit:

$$\pi_{\xi_i}(\psi) = \frac{\psi}{4\hat{\alpha}^2\rho^2}\left[\hat{\alpha}\bar{b} + \alpha\rho + \hat{\alpha}\rho(\hat{\alpha}\bar{b} - \alpha\rho)\frac{z(\psi)}{\psi} - \hat{\alpha}\bar{b}z(\psi) + \alpha\rho z(\psi)\right]^2$$

Where $z(\psi) = \sqrt{\frac{\psi}{\psi - \hat{\alpha}\rho}}$ - decreasing function of ψ. Taking derivative $\pi'_{\xi_i}(\psi)$ we obtain:

$$\pi'_{\xi_i}(\psi) = \frac{1}{4\psi\hat{\alpha}^2\rho^2}\frac{\hat{\alpha}\rho z(\psi)}{z(\psi)^2 - 1}\left[\alpha^2\rho^2(z(\psi)+1)^2 - \hat{\alpha}^2\bar{b}^2(z(\psi)-1)^2\right]$$

Where we substituted $\psi = \frac{z(\psi)^2\hat{\alpha}\rho}{z(\psi)^2-1}$ and also $z(\psi) > 1$ on $\psi > \hat{\alpha}\rho$. Notice that $\xi_i = \frac{1}{2\hat{\alpha}\rho}[\bar{b}\hat{\alpha}(z(\psi) - 1) - \alpha\rho(z(\psi)+1)]$ and $\xi_i > 0$ on $\psi < \psi_H$. Developing $\pi'_{\xi_i}(\psi)$ and comparing to ξ_i we obtain:

$$\pi'_{\xi_i}(\psi) = \frac{1}{4\psi\hat{\alpha}^2\rho^2}\frac{\hat{\alpha}\rho z(\psi)}{z(\psi)^2 - 1}\left[\alpha\rho(z(\psi)+1) - \hat{\alpha}\bar{b}(z(\psi)-1)\right] \cdot \left[\alpha\rho(z(\psi)+1) + \hat{\alpha}\bar{b}(z(\psi)-1)\right] < 0$$

To check that $\pi_{\xi_i}(\psi) > \pi_0$ on $\psi \in [\psi_L, \psi_H]$ we thus need to check this inequality at ψ_H. It is easy to verify that $\pi_{\xi_i}(\psi_H) = \frac{\alpha\bar{b}}{1-\beta} = \pi_0$. Thus, on $\psi \in [\psi_L, \psi_H]$ digital business optimally chooses

$\xi^* = \xi_i \in [0,1]$ and $C^* = C_h(\xi_i)$. Figure 5 (d) illustrates this case.

In order to check that ξ, C are complements, we need to evaluate cross-partial derivative at ξ^*, C^*. The sign of the derivative is defined by (see proof of Lemma 3):

$$2\alpha\bar{b}(C' - \xi^2) + \rho\alpha\xi(3C' - \xi^2) + \hat{\alpha}\bar{b}\xi(3C' - \xi^2) + 4\hat{\alpha}\xi^2\rho C' \tag{8}$$

Where $C' = C(1-\beta)$. First, consider $\psi < \psi_L$, then $\xi^* = 1$ and $C^* = C_h(1) = -1+\sqrt{\frac{1}{\psi}(\bar{b}+\rho)(\alpha+\hat{\alpha})}$. Evaluating expression 8 at ξ^*, C^* we get: $-4(\bar{b}+\rho)(\alpha+\hat{\alpha}) + \sqrt{\frac{(\alpha+\hat{\alpha})(\bar{b}+\rho)}{\psi}}\left(2\alpha\bar{b} + 3\hat{\alpha}\bar{b} + 3\alpha\rho + 4\hat{\alpha}\rho\right)$. It is easy to verify that this expression is positive even at $\psi = \psi_L$. Hence, ξ, C are complements on $\psi < \psi_L$.

Consider now region of $\psi \in [\psi_L, \psi_H]$. Recall $z = \sqrt{\frac{\psi}{\psi - \hat{\alpha}\rho}}$ (decreasing with ψ and trivially $z > 1$). $\psi = \frac{z^2\hat{\alpha}\rho}{z^2-1}$ and the range of z is: $z \in \left[z_L = \frac{\hat{\alpha}\bar{b}+\alpha\rho}{|\hat{\alpha}\bar{b}-\alpha\rho|}, z_H = \frac{\hat{\alpha}\bar{b}+\alpha\rho+2\hat{\alpha}\rho}{|\hat{\alpha}\bar{b}-\alpha\rho|}\right]$. We will consider case $\bar{b}\hat{\alpha} > \alpha\rho$, the opposite case is considered similarly. Evaluating expressions 8 in this region we get that its sign is defined by the following expression: $\hat{\alpha}\bar{b}^2(z-1)^2 - \alpha_0^2\rho^2(1+z)^2$, which is increasing in z and is positive on the entire interval $z \in [z_L, z_H]$. Hence, ξ, C are complements at the optimum.

Finally, notice that when $\hat{\alpha}\bar{b} = \alpha\rho$ we have $\psi_L = \psi_H$, thus the moderate cost of protection region doesn't exist.

B.7 Proof of Proposition 3 (Page 19).

We first note that when there are no adversaries $\omega = 0$ and the digital business sets $(C, \xi) = (0, 1)$. As a result, user i's maximization becomes $U_i(a_i, \bar{a}) = a_i b_i - \frac{1}{2}a_i^2 + \beta a_i \bar{a} + a_i\rho$, and average users' activity level is $\bar{a} = \frac{\bar{b}+\rho}{1-\beta}$. Substituting this into the expression for $CS_{no-criminals}$ and substituting the expression for CS^* and $\mathcal{D}$ and reorganizing we get $\mathcal{M} = \frac{C^*}{\xi^*}\left(\left(\frac{\bar{b}+\rho}{1-\beta}\right)^2 \frac{1}{(\bar{a}^*)^2} - 1\right)$ as required. To see that $\mathcal{M} \geq \frac{2}{(1-\beta)}$ we first note that holding C and ξ fixed, the restricted multiplier $\tilde{\mathcal{M}}$ is $\frac{\xi^2+2C(1-\beta)}{C(1-\beta)^2}$ which is greater than $\frac{2}{(1-\beta)}$ because $\frac{\xi^2+2C(1-\beta)}{C(1-\beta)^2} = \frac{\xi^2}{C(1-\beta)^2} + \frac{2C(1-\beta)}{C(1-\beta)^2} = \frac{\xi^2}{C(1-\beta)^2} + \frac{2}{(1-\beta)}$. We next argue that $\mathcal{M} \geq \tilde{\mathcal{M}}$. That is, allowing the business to adjust C and ξ when there are no adversaries enhances CS. To show that, it is sufficient to show that when there are no adversaries, CS is maximized at $(C, \xi) = (0, 1)$, which we show next. When there are no adversaries user i's maximization becomes $U_i(a_i, \bar{a}) = a_i b_i - \frac{1}{2}a_i^2 + \beta a_i \bar{a} + a_i\xi\rho$, and average users' activity level is $\bar{a} = \frac{\bar{b}+\xi\rho}{1-\beta}$. That is, average users' activity is independent of C and increases in ξ. Recalling that, in equilibrium, $CS = \bar{a}^2 + \sigma_b^2$ we get that when there are no adversaries, CS is independent of C and maximized at $\xi = 1$ as required.

B.8 Proof of Proposition 4 (Page 20) and Claims in Section 5

Let $\alpha = \hat{\alpha}$. Using Corollary 1, when $\psi < \alpha\rho/4$: $\Pi = \frac{\psi}{1-\beta}\left(\sqrt{\frac{\rho\alpha}{\psi}} - 1\right)^2$, $CS \sim \frac{1}{1-\beta}\sqrt{\frac{\rho\psi}{\alpha}}\left(\sqrt{\frac{\rho\alpha}{\psi}} - 1\right)$ and stored information $\bar{a}^*\xi^* = \frac{1}{1-\beta}\sqrt{\frac{\rho\psi}{\alpha}}\left(\sqrt{\frac{\rho\alpha}{\psi}} - 1\right)$. When $\psi > \alpha\rho/4$: $\Pi = \frac{\alpha^2\rho^2}{16\psi(1-\beta)}$, $CS \sim \frac{\alpha\rho^2}{8\psi(1-\beta)}$, $\bar{a}^*\xi^* = \frac{\alpha^2\rho^3}{32\psi^2(1-\beta)}$. Similarly, for advertisement-driven company, using Corollary 2, we obtain, when $\psi < \alpha\rho$: $\Pi = \frac{\psi}{1-\beta}\left(\sqrt{\frac{\rho\alpha}{\psi}} - 1\right)^2$, $CS \sim \frac{1}{1-\beta}\sqrt{\frac{\rho\psi}{\alpha}}\left(\sqrt{\frac{\rho\alpha}{\psi}} - 1\right)$ and stored information $\bar{a}^*\xi^* = \frac{1}{1-\beta}\sqrt{\frac{\rho\psi}{\alpha}}\left(\sqrt{\frac{\rho\alpha}{\psi}} - 1\right)$, o/w all zeros. Clearly, when $\psi < \alpha\rho/4$, all measures are equal; when $\psi > \alpha\rho$, transaction-driven company has non-zero profit, CS, and stores non-zero amount of information as compared to advertisement-driven company, for which these quantities are zero. Consider now $\psi \in [\alpha\rho/4, \alpha\rho]$. Difference in profits $\Pi_t - \Pi_a \sim \frac{\alpha^2\rho^2}{16\psi} - \psi\left(\sqrt{\frac{\rho\alpha}{\psi}} - 1\right)^2$ which is positive when $\alpha\rho + 4\psi > 4\sqrt{\rho\alpha\psi}$ which is always true except for $\psi = \alpha\rho/4$ when it is equality. Difference in CS: $CS_a - CS_t \sim \rho - \sqrt{\frac{\rho\psi}{\alpha}} - \frac{\alpha\rho^2}{8\psi}$, substituting $\phi = \sqrt{\frac{\psi}{\alpha\rho}} - \frac{1}{2}$, we get that $CS_a > CS_t$ when $2\phi^2 + \phi - \frac{1}{2} < 0$ which is true on $\phi \in [0, \phi_1]$, where ϕ_1 is the positive solution to the quadratic equation and it is easy to check that ϕ_1 is such that corresponding $\psi < \alpha\rho$. Finally, for stored information $(\bar{a}^*\xi^*)_a - (\bar{a}^*\xi^*)_t \sim 1 - \sqrt{\frac{\psi}{\alpha\rho}} - \frac{\alpha^2\rho^2}{32\psi^2}$, again substituting $\phi = \sqrt{\frac{\psi}{\alpha\rho}} - \frac{1}{2}$, we get that the difference is positive when $32(\phi+1/2)^4(1/2-\phi) > 1$, which is equivalent to $-32\phi^4 - 48\phi^3 - 16\phi^2 + 8\phi + 6 > 0, \forall z > 0$. The latter expression has only one positive root (coefficients of the polynomial change sign once). It is easy to verify that this root is s.t. corresponding $\psi < \alpha\rho$.