Digital Privacy*

Itay P. Fainmesser† Andrea Galeotti‡ Ruslan Momot§

First Draft: September 2019. This Draft: June 2020.

Abstract

We study the incentives of a digital business to collect and protect users' information. The information the business collects improves the service it provides to consumers, but it may also be accessed, at a cost, by third strategic parties in a way that harms users, imposing privacy costs. We characterise how the revenue model of the business shapes the equilibrium data policy. We compare the equilibrium data policy with the social optimum and show that a two-pronged policy, which combines a minimal data protection requirement with a tax proportional to the amount of data collected, restores efficiency.

"The problem with data protection laws is that it presumes the data collection was ok."

Eduard Snowden, NSA whistleblower

September 19, 2019

*We thank Susan Athey, Heski Bar-Issac, Otávio Bartalotti, Roland Bénabou, Devesh Raval, Douglas S. Smith, and participants of talks and seminars at the Retreat on Networks, Information, and Social Economics (RINSE), INFORMS, MSOM, London Business School, Paris School of Economics, Privacy Workshop at Princeton University, University of Pennsylvania, and the Federal Trade Commission for helpful comments and discussions. Andrea Galeotti gratefully acknowledges financial support from European Research Council through the ERC-consolidator grant (award no. 724356)

†The Johns Hopkins Carey Business School and The Department of Economics, The Johns Hopkins University, Baltimore, MD 21202 (e-mail: itay_fainmesser@jhu.edu)

‡Department of Economics, London Business School, London, NW1 4SA, United Kingdom (e-mail: agaleotti@london.edu)

§Operations Management Department, HEC Paris, 78350 Jouy-en-Josas, France (e-mail: momot@hec.fr)

1. Introduction

The growing social and economic activity conducted online – from the sharing of location data on Uber to searching for medications or diagnoses on Google – generates extensive amounts of data. The information collected is then often used to improve services offered to consumers. For example, Uber uses passengers' and drivers' location to shorten wait times and enhance safety, Facebook uses personal data to curate posts presented to users based on their tastes and to target users with advertisements, and online dating platforms use users' information to propose desirable matches. At the same time, there are also undesirable consequences to the ubiquitous availability of data. From firms like Cambridge Analytica using Facebook data to sway election outcomes to health insurance companies predicting health outcomes of potential insurers based on undisclosed personal information, and to private hackers preying on innocent users – opportunities for user data exploitation are rife. As a result, privacy has been singled out as one of the biggest challenges that the digital economy faces.[1]

This paper addresses four fundamental issues on the digital economy: (1) the trade-offs people face when using online services, (2) the determinants of harmful use of private information by third parties, (3) the incentives of digital businesses to collect and protect user data, and (4) actions that should be taken by regulatory authorities in order to protect consumer privacy and maximize social welfare.

In our model, users' activity on the platform of a digital business reveals information about some of their traits/preferences. This information allows the business to offer users a higher quality service (positive externality), but it also triggers users' concerns for privacy (negative externality): the risk that users perceive from their information being accessed and misused by the third parties – hereafter *adversaries*. We explicitly model the strategic behavior of the adversaries and, therefore, users' privacy costs are endogenous.

An important feature of our framework is that we allow for different business revenue models.

[1] A few articles that illustrate these undesirable consequences are "How Trump Consultants Exploited the Facebook Data of Millions", *New York Times*, 17 March 2018 (see also Papanastasiou 2020 and Candogan and Drakopoulos 2020), "Can a Facebook Post Make Your Insurance Cost More?", *Wall Street Journal*, 18 March 2019," 'We've Been Breached': Inside the Equifax Hack", *Wall Street Journal*, 18 September 2017, and "Marriott CEO Reveals New Details About Mega Breach", *Forbes*, 11 March 2019. For a common take on the top questions on the digital economy, see also https://www.youtube.com/watch?v=W32yoivhaL4

We postulate that the business' profit is an increasing function of the users' demand for service and of the information collected by the business. We introduce the following taxonomy. At the one extreme, we have *purely data-driven* digital businesses whose sole source of revenue stems from selling information or information-based services to third-parties , e.g., weather apps that are free and do not display advertisements, but rather collect user location data and sell it to data aggregators such as Acxiom and Infogroup. At the other extreme, we have *purely usage-driven* businesses that collect users' payments in the form of subscription fees or commissions, e.g., ride-hailing platforms like Uber and Lyft, and many online dating services. Between these two extremes, we have, among others, ad-driven companies like Facebook and Google, whose main source of revenue is offering targeted advertising. In selling targeted advertisement services, an ad-driven business capitalizes directly on the users' information it collects. At the same time, the business also wishes users to be active for them to view and click on the ads. Therefore, the business also attaches a direct value to user activity.

We first show that, as the business's data collection policy becomes more expansive (i.e., the fraction of user data that is collected increases), users' activity first increases and then decreases. The resulting amount of users' information, that is actually collected by the business and is stored on its servers, follows a similar pattern, but it starts decreasing at a higher data collection level. Such non-monotonicity reflects a natural tension between the positive and negative externalities that data collection by the business imposes on users. When the data collection policy is restricted, privacy costs are low because only a few adversaries are active. Hence, an increase in data collection improves the service to users at little privacy costs. As the business collects a larger fraction of users' data, adversaries activity increase, thereby imposing larger privacy costs on users, who, in turn, reduce their activity.

Our second result characterizes the equilibrium data collection policy: what part (or how much) of user's information is actually collected by the business.[2] Since the business's profit increases in user activity as well as in the amount of data collected, the optimal data collection policy lies

[2]Data collection policy examples include the decision by Whatsapp to encrypt users' text messages (thus reducing the amount of information it collects) – see also "WhatsApp Introduces End-to-End Encryption", *New York Times*, 5 April 2016., and Facebook's practices (at least until August 2019) of transcribing audio chats (thereby increasing the amount of accessible information it stores) as well as of retaining information from deleted accounts – see also "Facebook Paid Contractors to Transcribe Users' Audio Chats", *Bloomberg*, 13 August 2019 and "OK, You've Deleted Facebook, but Is Your Data Still Out There?", *CBS News*, 23 March 2018.

in between the cut-off levels that maximize these two quantities. In particular, a business with a revenue model that is more data-driven selects a more expansive data collection policy and, consequently, experiences lower user activity, and generates lower consumer surplus.

We then compare the equilibrium data collection policy with the data collection policy that maximizes overall welfare (hereafter socially optimal policy).[3] Purely usage-driven businesses select the socially optimal policy. All other businesses over-collect users' information, thus leading to larger privacy costs and lower users' activity. Such inefficiency in data collection can be corrected with a tax proportional to the amount of data collected, or, with a liability policy that imposes on businesses fines that are proportional to users' privacy costs.

In the second part of the paper, we expand the digital business's strategy to include a choice of a level of data protection: the business makes a costly investment that affects the ease with which collected information is accessible by adversaries. Examples include investments in firewall enhancements, API updates, and ethical hackers who help the business in detecting and repairing vulnerabilities. We provide conditions under which data collection and data protection are complementary instruments for the business. Such complementarity implies that, although data-driven businesses tend to over-collect user information, they may also invest more in data protection. In an empirically relevant case of the model, in which the digital business's profit is linear in user activity and in the information collected, we show that businesses with a revenue model that is intermediate between purely usage-driven and purely data-driven models generate better outcomes for consumers.

Finally, we show that when businesses choose both data collection and data protection, a regulatory policy that combines a requirement of a minimal level of data protection together with either a liability policy or a tax on information collected, induces a socially optimal decentralized equilibrium. We contrast such policy with the practice of the US Federal Trade Commission (FTC) which has a mandate to enforce a minimal level of data protection level.[4] In practice, the vast majority of the regulator's actions against firms come in response to documented data breaches, and once a data breach has been verified, the FTC issues heavy fines and often settlement is reached without a court ruling on whether the minimal protection levels were met. This introduces a, de

[3] The welfare objective is a weighted sum of consumer surplus and business profit.

[4] See, for example, the case of FTC vs D-Link: "D-Link agrees to 10 years of security audits to settle FTC case", *The Verge*, 4 Jul 2019; https://www.ftc.gov/enforcement/cases-proceedings/132-3157/d-link.

facto, liability element and, according to our normative analysis, it increases the effectiveness of FTC's practice. Noting that the implementation of this liability depends on a combination of detection, litigation, and negotiation, our result suggests that liability fines may be replaced by a tax on data collection.

This paper contributes to an active interdisciplinary area of research that studies the consequences for market outcomes of the ability of digital institutions to amass large data sets. The main issues discussed in the literature are the collection and management of consumer information, consumer privacy, and possible policy intervention.

Recent work has focused on understanding how user information is either voluntarily disclosed (Ali et al. 2019) or inferred from users' actions such as purchasing behavior over time (Conitzer et al. 2012, Fudenberg and Villas-Boas 2006), ratings (Bonatti and Cisternas 2017), formation of social links (Acemoglu et al. 2017), platform usage (Ichihashi 2019) or gathered through monetary transfers (Bergemann et al. 2020). We rely on this line of work and assume that there is a one-to-one mapping between any user's action (usage of the platform) and the information that is revealed about this user to the platform.[5]

Other work has explored how the mechanisms for extracting user information and possession of information itself affect the design of targeted/personalized pricing (e.g., Candogan et al. 2012, Bloch and Quérou 2013, Fainmesser and Galeotti 2016, 2020, Montes et al. 2018, Ichihashi 2020a, Valletti and Wu 2020), selective selling (Momot et al. 2019), service systems (Hu et al. 2020); what is the impact on social image visibility (Ali and Bénabou 2020), advertising strategies (Galeotti and Goyal 2009, Shen and Miguel Villas-Boas 2017), the extent of competition (Casadesus-Masanell and Hervas-Drane 2015) and mergers of digital businesses (Prat and Valletti 2019), overall consumer behavior (Goldfarb and Tucker 2011, Koh et al. 2015, Jann and Schottmüller 2020, Gradwohl 2017).[6]

Our contribution is to formulate a model in which privacy costs are endogenous and therefore change with the data policy of the business. This allows to assess positive and normative implications of data policy design and to compare conclusions across different domains—in particular, between

[5]The computer science literature has addressed the design of algorithmic mechanisms for anonymizing and protecting 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).

[6]Excellent surveys are provided by Acquisti et al. 2016, Mayzlin 2016, and Bergemann and Bonatti 2019. See also https://www.heinz.cmu.edu/~acquisti/economics-privacy.htm#Papers for a structured list of papers in this area by Alessandro Acquisti.

data-driven and usage-driven revenue models. In the last section, we demonstrate how our framework is useful to study broader policy issues such as the motivations and welfare effect of vertical integration in digital markets, the unique privacy risks posed by government adversaries, and the potential welfare benefits of introducing a carefully calibrated payment to users for their data.

The rest of the paper proceeds as follows. Section 2 presents the model and Section 3 characterizes equilibrium user and adversarial activities for given data collection and protection policies. Section 4 compares equilibrium and socially optimal data collection policies and offers policy recommendations. Section 5 extends the analysis to include data protection policies. Finally, Section 6 offers a discussion and broader policy implications. All proofs are given in Appendix A.

2. Model

A digital business chooses a data collection policy. Users decide how much to use the services provided by the business. Users' activity, together with the business' data collection policy, determine how much information is collected about the users by the business. Third parties (henceforth, *adversaries*) can, at a cost, attempt to access and use the data for purposes that are not in line with users' preferences. Thus, if successful, adversaries can harm users. We introduce formally these elements next.

2.1. Users

There is a unit mass of users of a digital business. Each user i chooses a costly action a_i that represents i's usage of the business' service – *user's activity*. Denote by $\bar{a} = \int_j a_j \, dj$ the average user activity. User i's activity reveals valuable information about her, and we assume that a_i also captures the amount of information revealed about user i.

Not all information revealed by users' activity is collected by the digital business. In particular, the amount of information that a digital business collects about user i—and that can be retrieved (either by the business itself or by adversaries)—is ξa_i. Here $\xi \in [0, 1]$ is the business' *data collection policy*, which ranges from collecting none to all of the data generated by usage. We provide specific examples of data collection technologies that micro-found this formulation in Section 2.5.

Suppose, user i chooses activity a_i and expects average activity to be $\bar{a}$. Then, her utility is

$$U_i(a_i, \bar{a}) = \underbrace{a_i b_i - \tfrac{1}{2}a_i^2}_{\text{private benefit and cost}} + \underbrace{\beta a_i \bar{a}}_{\text{network effects}} + \underbrace{a_i \xi(\rho - \omega).}_{\text{positive and negative information externalities}} \quad (1)$$

The first term summarizes user i's private benefits and costs for the service. Users are heterogeneous with respect to b_i; $\bar{b}$ is the average across users and σ_b^2 is the variance. The second term introduces classical positive network effects that are parameterized by $\beta \geq 0$.[7] The quadratic specification in 1 allows us to derive closed-form solutions of the second stage game, but the qualitative results generalize beyond the quadratic formulation, see the formal discussion in Online Appendix C.1.

The last term captures the positive and negative externalities to user i from the information that is collected by the business. On the one hand, every additional unit of information that the business collects, improves the service offered to users by $\rho \in [0, 1)$. On the other hand, as more information is collected, there is also a higher demand for information from adversaries. Adversarial activity, thus, imposes negative externality on users proportional to the amount of information the business has on them. This negative externality is captured by ω, which is defined as the number (i.e., mass) of adversarial activities using the data of a user. We refer to ω as to the *demand for user information from adversaries*.

2.2. Adversaries

There is a large mass K of potential adversaries.[8] Adversaries are heterogeneous in their ability to access information collected by the digital business. This heterogeneity is captured by the parameter γ, which we assume (for the sake of simplicity) to be drawn for each adversary from a uniform distribution over $[0, K]$.[9] An adversary knows his own γ and chooses whether to be active (action 1) or not (action 0). The gain to an inactive adversary is his outside option, which we normalize to zero ($\pi(0|\gamma) = 0$). If an adversary with ability γ is active, then he pays a fixed cost γC to access the collected information and targets one user, chosen uniformly at random. An

[7] The specification of network effects is widely used in network economics literature (see, e.g., Bloch and Quérou 2013, Candogan et al. 2012, Fainmesser and Galeotti 2016, to name a few).

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

[9] The draw of γ for different adversaries does not need to be independent and our analysis holds for any correlation structure. See also in Online Appendix C.1 for a generalization to non-uniform distributions.

adversary who targets user i receives payoff $a_i\xi$, so his expected benefit is $\bar{a}\xi$. Formally, the payoff expected by an active adversary of ability γ is[10]

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

The parameter C describes how well data are protected against misuse by third parties. In the first part of our analysis, this parameter is exogenous. In Section 5 we allow the digital business to invest in data protection.

2.3. Digital Business

The digital business chooses a data collection policy ξ to maximize its objective: a function which is increasing in users' average activity and in the amount of user information that the business collects. Formally, the digital business's profit function has the following form:

$$\Pi(\xi) = \Phi(\bar{a}, \xi\bar{a}). \tag{2}$$

We denote by $\Phi'_{\bar{a}}$ and $\Phi'_{\xi\bar{a}}$ the partial derivatives of Φ with respect to the first and the second arguments, respectively. That is, $\Phi'_{\bar{a}}$ is the additional profit to the business for a marginal increase in users' average activity, ceteris-paribus, and $\Phi'_{\xi\bar{a}}$ is the additional profit to the business for a marginal increase in users' information collected by the business. We assume that $\Phi'_{\bar{a}} \geq 0$, $\Phi'_{\xi\bar{a}} \geq 0$ and that the function Φ is concave in its two arguments.

The function Φ captures the revenue model of the business as well as the level of market competition that the business faces in markets in which it extracts revenue (e.g., the advertising market, markets for user data, etc.). With respect to the revenue model, we distinguish the following two polar cases (see the Introduction for specific examples of different business models).

A *purely data-driven* business's sole source of revenue stems from selling information or information-based services to third parties. For such businesses, $\Phi'_{\bar{a}} = 0$ and $\Phi'_{\xi\bar{a}} > 0$. On the contrary, a *purely usage-driven* business's source of revenue are the payments made by users in the form of subscription

[10] Another interpretation of this model of the adversary is that there is a single adversary, who upon gaining access to the digital business's data, attacks *all* users. If we choose this latter interpretation, we can interpret γC as the cost of accessing the data or interpret C as the cost of accessing the data and $1/\gamma$ as the likelihood that the adversary manages to access the data after investing C.

fees or commissions. For such companies, $\Phi'_{\bar{a}} > 0$ and $\Phi'_{\xi\bar{a}} = 0$. Between these two extremes, lie, among others, advertisement-driven companies, like Facebook and Google, that capitalize directly on the users' information they collect, but that also needs users to be active to view and click on the ads, and, thus, it also attaches a direct value to user activity. The exact way that an advertisement-driven company weighs users' activity vs. information collected may depend, among other things, on the life cycle of the business: at the startup phase, the objective is, generally, to maximize activity as this is the metric that allows raising capital through investors. As the platform matures, the weight is often shifted towards monetizing collected information.

2.4. Timeline and Equilibrium Concept

We consider the following sequential game. In the first stage, the digital business chooses its data collection policy ξ – the choice is observed by users and adversaries. Then, users choose their activity levels and, simultaneously, adversaries decide whether or not to be active.[11] Notably, because there is a continuum of users and adversaries, no one agent's action affects the best reply of others. Hence the analysis does not change if users' and adversaries' moves are sequential rather than simultaneous.

The strategy of a digital business corresponds to a data collection policy $\xi \in [0, 1]$. The user's strategy is a function $a_i \colon \mathbb{R}_+ \times [0, 1] \to \mathbb{R}_+$ that specifies user i's activity for every possible b_i and ξ. The strategy of an adversary is a function $v_j \colon [0, K] \times [0, 1] \to \{0, 1\}$ that specifies, for every possible $\gamma \in [0, K]$ and ξ, whether adversary j is active and will attack the business's database. We use $\mathbf{a}$ and $\mathbf{v}$ to denote the *strategy profiles* of users and adversaries respectively.

We characterize perfect Bayesian equilibria of the game: an information collection choice 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 ξ, $(\mathbf{a}^*, \mathbf{v}^*)$ is a Bayesian equilibrium in the ensuing subgame.[12]

Hereafter, we maintain the following assumption:

[11] The assumption that users know the data collection policy of the digital business is consistent with common practice where users accept *terms and conditions* when opening an account in a digital business. However, in reality, some users may not read those terms. This could be because they have no privacy concerns or because they lack awareness and do not internalize those costs. The model can easily accommodate heterogeneity across users in their sophistication or awareness levels and the qualitative results do not change. We also note that, in recent years, there is an increase in user awareness regarding the true extent of data collection and the potential privacy costs. In particular, many users can now request their information from digital businesses under the European GDPR law (see Art. 15 GDPR - Right of access by the data subject). The GDPR, therefore, allows users to have accurate knowledge of what data is collected about them.

[12] In each subgame, users and adversaries have a common prior that γ is uniformly distributed between 0 and K.

Assumption 1 *Assume that* $\beta < 1$ *and that* $C \leq \frac{\bar{b}}{(1-\beta)(2\rho+\bar{b})}$.

As is typical in the literature featuring models with positive network externalities, the assumption that $\beta < 1$ guarantees that, for every ξ, there is a unique equilibrium in the game's second stage. The restriction on C guarantees that the equilibrium ξ^* is interior.

2.5. Information Externalities and Revenue Models: Simple Examples

The following two examples demonstrate the workings behind our reduced-form formulation of information externalities and how they play out in businesses with different revenue models. The examples are not intended to capture the full complexity of a digital ecosystem but rather to provide some intuitions for our taxonomy of revenue models and information externalities.

Example 1 *Consider a digital business with* ***a purely data-driven revenue model****, say a weather app. Downloading and using the app is free. Private benefit and cost of using the service as well as any potential positive network effects are captured by the first and second terms of expression* (1).

The weather app sells information to a data aggregator, who, among other things, uses the data to target users with products (or sells access to the data to marketers who seek to do so). Now, suppose that each user i *has a taste characteristic* $\theta_i \in \{0, 1, ..., n\}$*, which is initially unknown to the aggregator. If the aggregator gets to know* θ_i*, it can target user* i *with a product that creates a value* V *to her.*

Ex ante, the aggregator knows that for any $m \in \{0, 1, ..., n\}$, $\theta_i = m$ *with probability* $1/n$*. User* i*'s activity provides signals to the aggregator about her tastes parameter* θ_i*, but only if data generated by this activity is collected. In particular, if the digital business sets data collection policy* ξ *and the user's activity level is* a_i*, then* ξa_i *is the probability that the aggregator learns the true realization of* θ_i*, whereas with the remaining probability the aggregator learns nothing.*[13] *Therefore, the expected probability that the aggregator creates a value* V *to the consumer is:*[14]

$$\frac{1}{n}\left[1 + (n-1)\, a_i \xi\right].$$

[13] For the purposes of this example, assume that b_i is sufficiently low so that $a_i < 1$

[14] An alternative, mathematically equivalent, model considers each user i as having a large set of attributes, each of which can receive a 0/1 value. Then, activity level a_i generates data that can reveal the value of a_i of a subset of the attributes and a business' data collection strategy ξ implies that the business collects data on a fraction ξ of the attributes and therefore ends up with accurate information on $a_i\xi$ of the attributes of user i and for any of the other attributes it only knows that they are 0 with probability 1/2. Knowing more of the attributes allows the business to match the user with a more suitable product which increases the expected value created.

Depending on the competitiveness of the market for products and for data, this value is shared between the data aggregator, the digital business, and the user. Let S_b be the share that is extracted by the digital business and S_u be the share that is extracted by the user. Then the profit to the digital business from user i is $S_b V\left[1+(n-1)\, a_i \xi\right]/n$. The positive externality to the user is $S_u V\left[1+(n-1)\, a_i \xi\right]/n$ (that is, $\rho = S_u\,(n-1)\, V/n$).

The negative externalities of the adversaries can be modelled similarly by considering adversaries who want to choose actions that match the least favorable user's taste characteristic.

Example 2 *Consider a digital business with* ***a purely usage-driven revenue model****, say a dating application such as OkCupid or Clover, which offers in-app purchases that enhance functionality. Much of the formulation from Example 1 carries over with the following differences: In this example, the business doesn't sell data to an aggregator. Instead, if the business knows the taste characteristic θ_i, it can offer matches suitable to user i. A suitable match creates a value V to the user and an unsuitable match creates zero value.*

The user receives this value (as well as the private benefits and network effects benefits) minus the price which is paid to the platform for in-app purchases, which are assumed to be proportional to usage, or $a_i p$ where p is the price per unit of usage. Then the profit to the business from user i is $a_i p$. In the user's utility, $a_i p$ can be included in $b_i a_i$ as a shift parameter, and the positive information externality to the consumer is $V\left[1+(n-1)\, a_i \xi\right]/n$. Hence, $\rho = V\,(n-1)\,/n$.

In Online Appendix B, we provide an additional example of a business with a hybrid revenue model that has both data- and usage-driven components (e.g., Facebook or Amazon).

3. The Users-Adversaries Game

We start by analyzing the equilibrium in the second stage of the game for any data collection policy ξ. Given ξ, each user trades off the benefit of using the services of the business with the associated costs. Part of these costs is the negative externalities that are imposed by the demand for information from adversaries. Adversaries' demand is, in turn, endogenously determined by their cost and benefit considerations.

Proposition 1 *Fix a data collection policy ξ. The ensuing subgame has a unique equilibrium*

in which average users' activity is

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

adversaries' demand for information is

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

and consumer surplus is equal to

$$\mathrm{CS}(\xi) = \frac{1}{2}\left[\sigma_b^2 + \bar{a}^*(\xi)^2\right]. \tag{5}$$

Because the business' objective is a function of average activity and of information collected, it is critical to understand the comparative statics of these quantities with respect to the choice of data collection policy ξ.

Corollary 1 *Average users' activity* $(\bar{a}^*(\xi))$ *and the amount of information stored* $(\xi\bar{a}^*(\xi))$ *both first increase and then decrease in the fraction of information collected, i.e. the business's data policy* ξ*. That is, there exist* $0 < \underline{\xi} < \bar{\xi} < 1$ *such that:*[15]

(i) $\bar{a}^*(\xi)$ *increases with* ξ *for* $\xi \in [0, \underline{\xi}]$ *and decreases otherwise;*

(ii) $\xi\bar{a}^*(\xi)$ *increases with* ξ *for* $\xi \in [0, \bar{\xi}]$ *and decreases otherwise.*

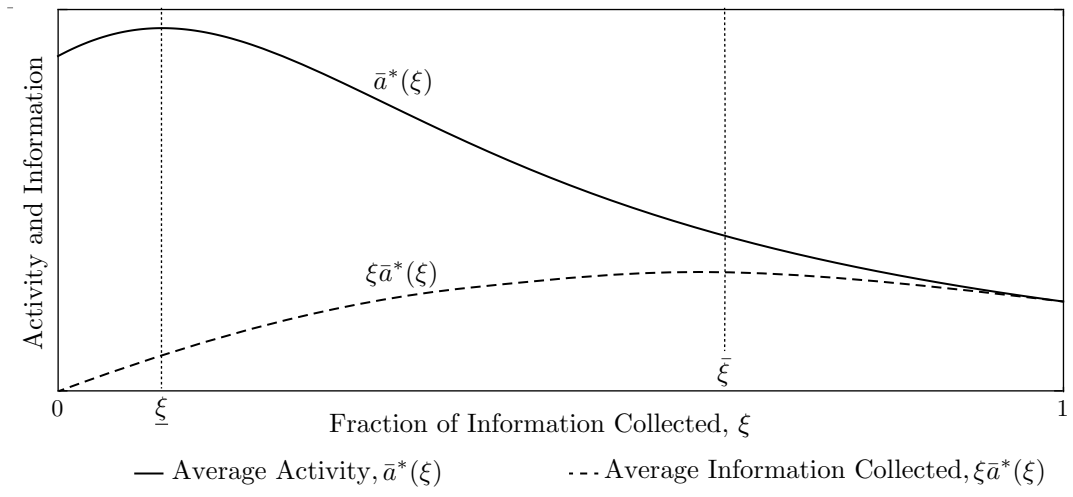

Figure 1: Average activity of users $\bar{a}^*(\xi)$ and information collected $\xi\bar{a}^*(\xi)$ as a function of data collection policy ξ.

The effect of a change in data collection policy (ξ) on average activity and average information

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

collected is illustrated in Figure 1. When no information is collected (i.e., $\xi = 0$), the equilibrium adversaries' demand for information $\omega^*(0)$ is zero (since there is no information to take) and user activity is determined solely by the interaction between a user's stand-alone benefit and cost, and positive network effects.

By collecting 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) and outlined in the Examples 1 and 2 in Section 2.5). This effect increases users' demand for the business's service. At the same time, increasing ξ also boosts the adversaries' demand for information. This creates a negative externality on users' participation and offsets the increase in users' demand for the business. Which effect dominates depends on the level of ξ.

Adversaries are unlikely to misuse information when ξ is small. In this case, the negative externalities that adversaries impose on users are small and an increase in ξ will increase the extent of user activity. With increasing ξ further, adversaries have more to gain from every attack; this leads to a further increase not only in adversaries' demand for user information but also in the loss that users suffer from adversaries. At some point, these negative effects outweigh the benefits to users from their information being used for tailored services. When this happens, users' average activity starts declining in ξ.

Even though users' average activity declines for every $\xi > \underline{\xi}$, total information collected by the business keeps increasing when $\xi \in [\underline{\xi}, \bar{\xi}]$. In this region, negative externalities imposed by adversaries on users are sufficiently significant to lead to a decrease in usage, but not severe enough to make this decrease large. Consequently, in this region, even though average activity decreases in ξ, information collected still increases in ξ. It is only when $\xi > \bar{\xi}$ that any additional increase in ξ leads to a decrease in user activity that is steep enough to reduce total information collected, notwithstanding an increase in the fraction of information stored.

4. Data Collection

Our first result in this section characterizes the business's equilibrium data collection policy and, as a result, how the business's revenue model affects user activity, the amount of user information

collected, adversaries' demand for information, and consumer surplus. To this end, we begin by defining a notion of a *more data-driven* revenue model.

Definition 1 *We say that a business with objective function $\Phi(\bar{a}, \xi\bar{a})$ has a revenue model that is more data-driven than a business with objective function $\tilde{\Phi}(\bar{a}, \xi\bar{a})$ if, for all $\bar{a}$ and ξ:*

$$\frac{\Phi'_{\xi\bar{a}}(\bar{a}, \xi\bar{a})}{\Phi'_{\bar{a}}(\bar{a}, \xi\bar{a})} > \frac{\tilde{\Phi}'_{\xi\bar{a}}(\bar{a}, \xi\bar{a})}{\tilde{\Phi}'_{\bar{a}}(\bar{a}, \xi\bar{a})} \tag{6}$$

Next, recall that $\bar{a}^*(\xi)$ is the average user activity in response to data collection policy ξ, and that $\underline{\xi}$ and $\bar{\xi}$ are the data collection policies that maximize average usage and information collection, respectively. Let

$$r(\xi) = -\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi} \Big/ \frac{\mathrm{d}(\bar{a}^*(\xi))}{\mathrm{d}\xi}. \tag{7}$$

The function $r(\xi)$ measures how the total amount of information collected, $\xi\bar{a}^*(\xi)$, changes in response to ξ relative to the change in average activity, $\bar{a}^*(\xi)$. At $\xi = \underline{\xi}$ this function is equal to $+\infty$, it decreases in ξ and it equals zero at $\xi = \bar{\xi}$. We can then prove the following result:

Proposition 2 *Let $\xi^\star$ be the digital business's equilibrium data collection policy. Then, $\xi^\star \in [\underline{\xi}, \bar{\xi}]$ and*

$$r(\xi^\star) = \frac{\Phi'_{\bar{a}}(\bar{a}^*(\xi^\star), \xi^\star\bar{a}^*(\xi^\star))}{\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi^\star), \xi^\star\bar{a}^*(\xi^\star))}. \tag{8}$$

Furthermore,

- *For a purely usage-driven business (i.e., $\Phi'_{\xi\bar{a}} = 0$), $\xi^\star$ maximizes equilibrium average activity, $\bar{a}^*(\xi)$, i.e., $\xi^\star = \underline{\xi}$.*
- *For a purely data-driven business (i.e., $\Phi'_{\bar{a}} = 0$), $\xi^\star$ maximizes information collected, $\xi\bar{a}^*(\xi)$, i.e., $\xi^\star = \bar{\xi}$.*
- *A business with a revenue model that is more data-driven chooses a higher $\xi^\star$, collects overall more user data, $\xi^\star\bar{a}^*(\xi^\star)$, and faces lower average user activity, $\bar{a}^*(\xi^\star)$. A more data-driven revenue model also leads to lower consumer surplus, $\mathrm{CS}(\xi^\star)$ and to higher adversaries' demand for user information, $\omega^*(\xi^\star)$.*

To gain some intuition for Proposition 2, we first note the following: (1) the business's objective increases in average activity and in total information collected, (2) both of those quantities increase in ξ for $\xi < \underline{\xi}$ and decrease in ξ for $\xi > \bar{\xi}$. Therefore, $\xi^\star \in [\underline{\xi}, \bar{\xi}]$. That is, the business' information

collection policy must be at least as high as the level at which average activity is maximized and no higher than the level at which information collection is maximized. Second, we note that when $\xi \in [\underline{\xi}, \bar{\xi}]$, the marginal profit with respect to ξ reads as follows:

$$\frac{\mathrm{d}\Pi(\xi)}{\mathrm{d}\xi} = \underbrace{\Phi'_{\bar{a}}(\bar{a}^*(\xi), \xi\bar{a}^*(\xi)) \cdot \frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}}_{\text{Marginal cost}} + \underbrace{\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi), \xi\bar{a}^*(\xi)) \cdot \frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}}_{\text{Marginal benefit}}$$

The first term represents the implicit marginal cost for the business to increase ξ: an increase in ξ reduces business's profit because users will demand less of the business's services (users' average activity $\bar{a}^*(\xi)$ decreases in ξ for $\xi > \underline{\xi}$). The magnitude of the cost depends on the price the business obtains from average usage, i.e., $\Phi'_{\bar{a}}(\bar{a}^*(\xi), \xi\bar{a}^*(\xi))$. The second term is the marginal benefit of the business to increase ξ: an increase in ξ increases the business's profits because the business is able to collect and use more of users' information and this information is priced positively at $\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi), \xi\bar{a}^*(\xi))$.

The business's choice of data collection balances marginal cost and marginal benefit. By substituting $r(\xi)$ into the business's first-order condition we obtain (8), where the right-hand side is the ratio of the per-unit price that the business obtains from average usage and from information collected. The comparative statics in Proposition 2 follow from (8) and from the connection between usage and consumer surplus, as well as from the connection between information collected and adversaries' demand for such.

4.1. Socially Optimal Data Collection

Consider a benevolent planner who seeks to choose a data collection policy to maximize social welfare: a weighted average of consumer surplus and business's profit, i.e.,

$$W(\xi) = \alpha \mathrm{CS}(\xi) + (1 - \alpha)\Pi(\xi),$$

where $\alpha \in [0, 1]$. Our next result characterizes the socially optimal data collection policy and shows how equilibrium inefficiencies depend on the business's revenue model.

Proposition 3 *Let ξ^W be the socially optimal data collection policy. Then $\xi^W \in [\underline{\xi}, \bar{\xi}]$ and*

$$r(\xi^W) = \frac{\alpha \bar{a}^*(\xi^W) + (1-\alpha)\Phi'_{\bar{a}}(\bar{a}^*(\xi^W), \xi^W \bar{a}^*(\xi^W))}{(1-\alpha)\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi^W), \xi^W \bar{a}^*(\xi^W))}. \tag{9}$$

Furthermore,

- *If the business is purely usage-driven, i.e., $\Phi'_{\xi\bar{a}} = 0$, then the equilibrium data collection policy of the business is socially optimal, i.e, $\xi^W = \xi^\star$;*
- *Otherwise, relative to the socially optimal outcome, the equilibrium business's data collection policy and total users' information collected are too large ($\xi^\star > \xi^W$ and $\xi^\star \bar{a}^*(\xi^\star) > \xi^W \bar{a}^*(\xi^W)$), whereas average activity is too low ($\bar{a}^*(\xi^\star) < \bar{a}^*(\xi^W)$).*

Relative to the planner's function, the business generally does not fully internalize that an increase in data collection policy, ξ, hurts consumers by decreasing their average activity. The welfare-maximizing data collection policy ξ^W is, therefore, lower than the equilibrium business policy, $\xi^\star$. The only exception is a purely usage-driven business that, in equilibrium, chooses the welfare-maximizing data collection policy, ξ^W.

To gain some intuition, note that using the equilibrium characterization of consumer surplus (Proposition 1), we obtain that the socially optimal information collection policy ξ^W is defined as follows:

$$\xi^W = \arg\max_{\xi \in [0,1]} \ \alpha \cdot \frac{1}{2}[\sigma_b^2 + (\bar{a}^*(\xi))^2] + (1-\alpha) \cdot \Phi(\bar{a}^*(\xi), \xi \bar{a}^*(\xi)).$$

Much like consumer surplus and the business's profit, welfare is an increasing function of average activity and total information collected. Therefore, $\xi^W \in [\underline{\xi}, \bar{\xi}]$. Furthermore, when $\xi \in [\underline{\xi}, \bar{\xi}]$ we can decompose marginal welfare with respect to ξ into marginal cost and marginal benefit as follows:

$$\frac{\mathrm{d}W(\xi)}{\mathrm{d}\xi} = \underbrace{[\alpha \bar{a}^*(\xi) + (1-\alpha)\Phi'_{\bar{a}}(\bar{a}^*(\xi), \xi \bar{a}^*(\xi))] \cdot \frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}}_{\text{Marginal cost}} + \underbrace{(1-\alpha)\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi), \xi \bar{a}^*(\xi)) \cdot \frac{\mathrm{d}(\xi \bar{a}^*(\xi))}{\mathrm{d}\xi}}_{\text{Marginal benefit}}.$$

Solving for the first-order condition and recalling the definition of $r(\xi)$ in (7), we get that welfare is maximized for ξ^W that solves (9). We further note that, with the exception of purely usage-driven businesses, at the business's optimal data collection policy $\xi^\star$, marginal welfare is negative, i.e., $\frac{\mathrm{d}W(\xi)}{\mathrm{d}\xi}|_{\xi^\star} = \alpha \bar{a}^*(\xi^\star) \frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}|_{\xi^\star} < 0$.

4.2. Regulation

The US Federal Trade Commission (FTC) has a mandate to enforce a minimal level of data protection. In our model, such requirement maps into guaranteeing a level of data protection C not lower than a certain threshold. However, since Proposition 3 has been derived for an arbitrary level of C, it follows that enforcing a minimal level of data protection alone is insufficient to align the data collection incentives of the business with those of the social planner.

In practice, the vast majority of the FTC's actions against firms come in response to documented data breaches. Once such breaches have been exposed and verified, the FTC imposes heavy fines on the businesses involved. In this section, we characterize two policies that induce socially optimal data collection by businesses and argue that one of those policies is akin to the latter practice by FTC.

We first define the damage $D_i(\xi)$ that a user i expects incur to due to adversarial activity:

Definition 2 *User i expects to face a damage $D_i(\xi) = a_i \xi \omega^*(\xi)$ caused by adversarial activity. Also denote by $D(\xi) = \int_i D_i(\xi) \mathrm{d}i$ the average damage incurred by a user.*

We next show that imposing fines that are proportional to the damage inflicted on users (as prescribed in the Definition 3), will, in fact, provide businesses with the correct incentives to set a socially optimal data collection policy ξ^W.

Definition 3 *In the event a user suffers damage D_i from adversarial activity, a liability policy imposes a fine on the business which equals $\ell \times D_i$.*

Building on Proposition 3, which derives the socially optimal data collection policy ξ^W, we obtain the following result.

Proposition 4 *Consider a business which is not purely usage-driven and let $\alpha < 1$. Under liability policy*

$$\ell^\star = \frac{\alpha}{1-\alpha} \frac{C}{2r(\xi^W)\xi^W} \tag{10}$$

the business chooses a data collection policy which is socially optimal, i.e., $\xi^\star = \xi^W$. Furthermore, if the welfare objective corresponds to consumer surplus, i.e., $\alpha = 1$, then

$$\ell^\star = \frac{\Phi'_{\xi\bar{a}}(\bar{a}^*(\underline{\xi}), \underline{\xi}\bar{a}^*(\underline{\xi}))}{\rho}.$$

To see the mechanism behind such policy in more details, note that under arbitrary liability policy ℓ and arbitrary data collection policy ξ, the average fine that the business expects to pay is

$$F(\xi, \ell) = \ell \times D(\xi) = \ell \times \omega^*(\xi) \times \xi \bar{a}^*(\xi) = \frac{\ell}{C}[\xi \bar{a}^*(\xi)]^2,$$

where the second equality follows by noticing that, from Proposition 1, $\omega^*(\xi) = \xi \bar{a}^*(\xi)/C$. Hence, the objective function of the business becomes

$$\Pi(\xi, \ell) = \Phi(\bar{a}^*(\xi), \xi \bar{a}^*(\xi)) - F(\xi, \ell).$$

The fine $F(\xi, \ell)$ is an increasing function of total information collected, and therefore it reduces the benefit that the business obtains by capitalizing on the information. Such a reduction increases with the level of ℓ. Since it is exactly the direct capitalization on user information that drives the over-collection of information by the business in the first place, the introduction of such policy, if rightly calibrated, eliminates the misalignment with the social objective.

Our analysis thus suggests that the FTC's practice of imposing fines on businesses, based on the documented data breaches, could be optimal. We do note, however, that imposing liability fines requires a litigation process to establish and quantify damages. An alternative policy that takes a more legislative and bureaucratic path is imposing a *tax on collected user information.* Let t be the tax rate imposed by the regulator on each unit of collected information. That is, the expected amount of tax that the business pays is $t \times \xi \bar{a}^*(\xi)$. We can then say the following:

Corollary 2 *Consider a business which is not purely usage-driven and let $\alpha < 1$. Under a tax rate*

$$t^\star = \frac{\alpha}{1-\alpha} \frac{\bar{a}^*(\xi^W)}{r(\xi^W)}, \tag{11}$$

the business chooses a data collection policy which is socially optimal, i.e., $\xi^\star = \xi^W$. Furthermore, if the welfare objective corresponds to consumer surplus, i.e., $\alpha = 1$, then $t^\star = \Phi'_{\xi\bar{a}}(\bar{a}^*(\underline{\xi}), \underline{\xi}\bar{a}^*(\underline{\xi}))$.

We hypothesize that imposing such tax on the amount of collected data could be a viable option for a regulator since the amount of collected data could be easily quantified.

Finally, it is important to note that even under the socially optimal level of data collection, the

adversarial activity causes some direct disutility, or damage, to users (see Definition 2). While this direct damage may be small, the following result suggests that for any data protection level, and even with the socially optimal data collection policy, the total welfare loss from the presence of adversaries can be arbitrarily large.

Proposition 5 *Let* $\mathrm{CS}_{no\ adversaries}(\xi)$ *be consumer surplus in the absence of adversaries.*[16] *For any data collection policy* ξ*, the total decrease in consumer surplus due to the presence of adversaries,* $\mathrm{CS}_{no\ adversaries}(\xi) - \mathrm{CS}(\xi)$*, equals* $\mathcal{M}(\xi) \cdot D(\xi)$*, where* $D(\xi)$ *is the average damage caused to a user as described in Definition 2 and* $\mathcal{M}(\xi)$ *is the adversarial loss multiplier. Moreover,*

$$\mathcal{M}(\xi) \geq \frac{1}{1-\beta}.$$

Proposition 5 shows that for businesses with significant network effects (i.e., high β and, in particular, $\beta \to 1$), the welfare loss from adversarial activity could be arbitrarily large, even if the direct damage $D(\xi)$ is relatively small. This is the case because part of the change in welfare is explained by the change in users' behavior (decrease in activity) due to their concern that their data may become available to the adversaries. Network effects amplify this effect, as any individual user decreasing activity has a strong impact on the incentives of other users to be active.

More generally, there are two reasons for the total welfare loss to be larger than the direct damage from adversarial activity. First, in expectation of an adversarial activity, users reduce their usage of a service that a digital business provides. Second, a digital business itself commits to a more restricted data collection policy (lower ξ) in order to minimize the reduction in user activity. Of these two effects, only the first one is captured by the lower bound in Proposition 5, suggesting that the total welfare loss from adversarial activity is even larger.

From a policy perspective, the lesson from Proposition 5 is that even with optimal regulation of data collection in place, there are still high returns to averting data misuse in the digital economy. The next session considers the case where adversaries' demand for users information can be reduced by making a costly investment that increases data protection.

[16] In particular, $\bar{a}^*(\xi)_{\text{no adversaries}} = \frac{\bar{b}+\rho\xi}{1-\beta}$ and $\mathrm{CS}_{\text{no adversaries}}(\xi) = \frac{1}{2}\left[\sigma_b^2 + (\bar{a}^*(\xi)_{\text{no adversaries}})^2\right]$.

5. Data Protection

In this section, we allow the digital business to choose also a data protection policy: a costly investment that increases adversaries' costs to access the collected data. For example, a business could invest in firewall upgrades, modify its API to add additional restrictions on data access or hire ethical hackers to uncover database vulnerabilities. Formally, we assume that the business can choose a protection level C at a cost $K(C)$, where $K(C)$ is an increasing and convex function. That is, the business's profit function becomes $\Pi(\xi, C) = \Phi(\bar{a}, \bar{a}\xi) - K(C)$. To simplify the exposition, we make a technical assumptions that $K'(0) = 0$ and $K'(\bar{C}) = \infty$, where $\bar{C} = \frac{\bar{b}}{(1-\beta)(2\rho+\bar{b})}$. These two assumptions assure that the business's optimal choice of data protection C satisfies Assumption 1 and, therefore, equilibrium usage level is unique and ξ is interior.

To stress the dependency of equilibrium average activity on *both* data protection and data collection levels, we extend the notation for average activity to $\bar{a}^*(\xi, C)$. It is immediate from the characterisation in Proposition 1 that the equilibrium average activity, $\bar{a}^*(\xi, C)$, and the total equilibrium information collected, $\xi\bar{a}^*(\xi, C)$, are both strictly increasing in C. The business's optimal data protection and data collection policies solve the following problem:

$$(\xi^\star, C^\star) = \arg \max_{\substack{\xi \in [0,1] \\ C \geq 0}} \Phi(\bar{a}^*(\xi, C), \xi\bar{a}^*(\xi, C)) - K(C). \tag{12}$$

Let $r(\xi, C)$ be the counterpart of $r(\xi)$ (see expression 7) for the case when C is endogenous. Deriving the first-order conditions we obtain that $(\xi^\star, C^\star)$ solves

$$r(\xi^\star, C^\star) = \frac{\Phi'_{\bar{a}}(\bar{a}^*(\xi^\star, C^\star),\, \xi^\star\bar{a}^*(\xi^\star, C^\star))}{\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi^\star, C^\star),\, \xi^\star\bar{a}^*(\xi^\star, C^\star))}, \tag{13}$$

$$\frac{\partial \bar{a}^*(\xi^\star, C^\star)}{\partial C} \left[\Phi'_{\bar{a}}(\bar{a}^*(\xi^\star, C^\star), \xi^*\bar{a}^*(\xi^\star, C^\star)) + \xi\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi^\star, C^\star), \xi\bar{a}^*(\xi^\star, C^\star))\right] = K'(C^\star). \tag{14}$$

Part (iii) of Proposition 2 states that for the same level of data protection, businesses with more data-driven revenue models (see Definition 1) collect more information as compared to businesses with the less data-driven revenue models. In turn, this induces lower users' average activity in the platform and consequently leads to lower consumer surplus. However, the equilibrium condition (14) shows that the marginal benefit of data protection (the left-hand side of expression 14) also

depends on the revenue model of the business. Indeed, an increase in the marginal value to the business of a unit of users' activity, i.e., an increase in $\Phi'_{\bar{a}}$ or in $\xi\Phi'_{\xi\bar{a}}$, increases the marginal returns to investment in data protection, ceteris paribus. Furthermore, any direct change in data collection and data protection due to a change in the revenue model will create additional feedback depending on whether the level of collected information and the protection level are substitute or complement instruments for the digital business. We can say the following:

Proposition 6 *Suppose that for any $\bar{a}$ and ξ, $\Phi''_{\xi\bar{a},\bar{a}}(\bar{a},\xi\bar{a}) \geq -\xi\Phi''_{\xi\bar{a},\xi\bar{a}}(\bar{a},\xi\bar{a})$, then at the business's equilibrium policy, data collection ξ and data protection C are complements. That is, $\frac{\partial^2\Pi(\xi,C)}{\partial\xi\partial C}|_{\xi^\star,C^\star} > 0$.*

The condition of Proposition 6 guarantees that, for the business's profit, the complementary between users' activity and users' information is sufficiently strong relative to the second-order effects of users' information. This condition is satisfied if Φ is, for example, a linear function, a Cobb-Douglas function, or a CES function. It is also satisfied for a business that is purely usage-driven, regardless of the specific functional form of Φ.

An implication of Proposition 6 is that a change to the revenue model of the digital business or a regulatory change, which creates additional incentives for the businesses to invest in data protection, will also lead to an increase in data collection by the business. The former effect (increase in data protection) will increase consumer surplus, whereas the latter effect (increase in data collection) will lead to a decrease in consumer surplus. The result presents a challenge to policymakers: what is the optimal intervention in a situation in which correcting one aspect of the problem amplifies inefficiencies on another margin? Before addressing this regulatory challenge, we provide a simple example that lends itself to an analytical and numerical analysis of the dependency of the business's data collection and data protection policies and the way in which both are affected by the business's revenue model. In the subsequent section, we return to the general model and derive a general welfare-maximizing regulation recommendation that is surprisingly simple given the complexity of the business' incentives.

Example: Linear Model. Suppose that $\Phi(\bar{a},\xi\bar{a}) = P_u \cdot \bar{a} + P_d \cdot \xi\bar{a}$. That is, $\Phi'_{\bar{a}} = P_u$ is the "price" the business extracts for each unit of users' average activity (usage) and $\Phi'_{\xi\bar{a}} = P_d$ is the respective price the business extracts for each unit of data collected. The ratio P_d/P_u defines the

extent to which the business is data-driven (based on our Definition 1). Notably, this example satisfies the property of Proposition 6 that, for any $\bar{a}$ and ξ, $\Phi''_{\xi\bar{a},\bar{a}}(\bar{a},\xi\bar{a}) \geq -\xi\Phi''_{\xi\bar{a},\xi\bar{a}}(\bar{a},\xi\bar{a})$. Therefore, at the business's equilibrium action, data collection ξ and data protection C are complements, and the marginal benefit of data collection increases in data protection.

Now consider two businesses, one of which is more data-driven (higher P_d/P_u). From equation (13), the incentives to collect information of such more data-driven business are strictly higher. Then, complementarity between data protection and data collection also implies an indirect effect that also leads to higher incentives of this business to protect information. This can be observed from equation (14), which can be rewritten as $P_u + \xi P_d = K'(C)/\bar{a}'_C$ to highlight that an increase in ξ requires higher protection C (notice, that the right-hand side increases in C). At the same time, a business can be more data-driven with higher or lower $P_u + \xi P_d$. If we know that the more-data driven business also has a higher $P_u + \xi P_d$, we could conclude that this business has higher incentives to protect the information. However, if the reverse holds, the lower $P_u + \xi P_d$ implies lower incentives to protect information, countering the effects of the complementarity between C and ξ.

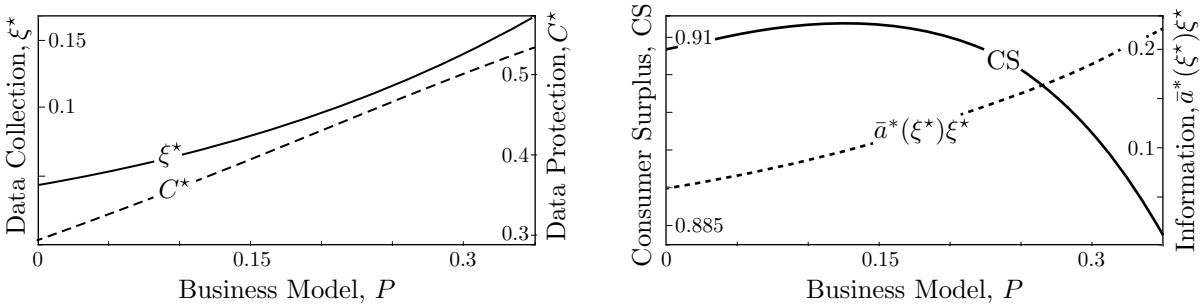

Figure 2: A numerical example for linear model, with $\Phi(\bar{a},\xi\bar{a}) = (1-P)\cdot\bar{a} + P\cdot\xi\bar{a}$ and $K(C) = \psi C \ln\frac{\bar{C}}{\bar{C}-C}$. (a) equilibrium data collection and data protection policies, $\xi^\star, C^\star$; (b) equilibrium consumer surplus and information collected, $\text{CS}, \xi^\star\bar{a}^*(\xi^\star)$ as functions of the extent to which a business is data-driven, P. Numerical parameters: $\bar{b} = 0.8, \rho = 0.4, \beta = 0.4, \psi = 0.05$ for which $\bar{C} = 0.83$.

Figure 2 illustrates the case in which $P_u = 1 - P$ and $P_d = P$ with $P \in [0,1]$. An increase in P thus represents a shift to a more data-driven revenue model. At the same time, an increase in P decreases $P_u + \xi P_d = 1 - (1-\xi)P$, thus leading to lower incentives to protect information. In this particular example, the aggregate effect is that data collection and data protection levels are higher for the more data-driven business. However, even in this simple example, user activity and

consumer surplus first increase and then decrease with the shift to a more data-driven revenue model. This is in contrast with the case of exogenous data-protection policy (see Proposition 2) in which a shift to a more data-driven revenue model unambiguously leads to a decrease in user activity and consumer surplus. When data protection is also endogenous, businesses with hybrid revenue models generate the highest levels of consumer surplus.

5.1. Equilibrium Inefficiencies and Optimal Regulation

In this section, we show that businesses can find it profitable to over- or under-collect information and over- or under-protect their data relative to the socially optimal levels. Nevertheless, regardless of the deviation from the social optimum, a simple policy consisting of a minimal protection requirement combined with liability fines or tax on data collection achieves the social optimum.

Using the equilibrium characterization of consumer surplus, social welfare can be captured by

$$W(\xi, C) = \frac{1}{2}\alpha[\sigma_b^2 + (\bar{a}^*(\xi, C))^2] + (1-\alpha)[\Phi(\bar{a}^*(\xi, C), \xi\bar{a}^*(\xi, C)) - K(C)],$$

with a first-order condition for data collection policy ξ that is structurally equivalent to condition (9). The first-order condition for data protection, C, can be written as follows:

$$\frac{\alpha}{1-\alpha}\bar{a}(\xi, C) + \Phi'_{\bar{a}}(\bar{a}^*(\xi, C), \xi\bar{a}^*(\xi, C)) + \xi\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi, C), \xi\bar{a}^*(\xi, C)) = \frac{K'(C)}{\frac{\partial\bar{a}(\xi, C)}{\partial C}}. \tag{15}$$

The left-hand side of (15) is increasing in α and the case of $\alpha = 0$ corresponds to the first-order condition of the profit-maximizing business. We thus obtain that for an exogenously given data collection level ξ, the optimal protection level of the business is always lower than the socially optimal protection level. This observation carries over for purely usage-driven businesses even when the business can choose both the levels of data collection and protection. Moreover, because for usage-driven businesses, C and ξ always exhibit complementarities at the optimum, usage-driven businesses will always under-collect and under-protect information relative to the social optimum.

A business with a revenue model that includes a significant data-driven component may, however, choose data collection and protection levels that are too high or too low relative to the social

optimum.[17] Figure 3 illustrates the three possible scenarios using the linear model example developed earlier.

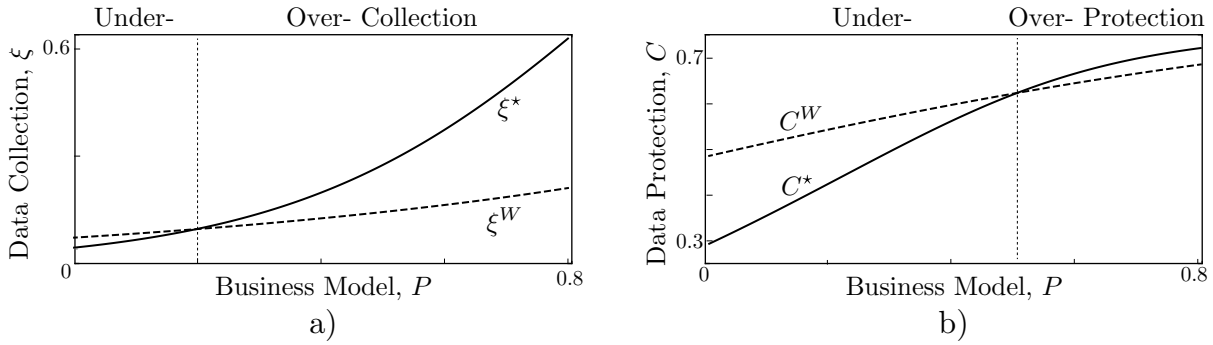

Figure 3: A numerical example for the linear model with the same set of parameters as Figure 2. (a) equilibrium and welfare-maximizing data collection policies ($\xi^\star$ and ξ^W respectively); (b) equilibrium and welfare-maximizing data protection policies ($C^\star$ and C^W respectively) as a function of the extent to which a business is data-driven, P.

The next proposition shows that regardless of the direction of the deviations from the social optimum, a simple two-pronged policy restores efficiency.

Proposition 7 *The following two-pronged policy induces an equilibrium in which the business chooses the socially optimal data collection and data protection policies:*

a) a required minimum level of data protection $C_{\min} = C^W$ *(where* C^W *is the socially-optimal data protection level)* ***combined with***

b) Either a liability fine proportional to expected damages from adversarial activity (see Definition 3) with the rate $\ell^\star$ *defined as in Proposition 4, or a tax rate* $t^\star$ *on information collected as defined in Corollary 2.*

Furthermore, if a business is purely usage-driven (i.e., if $\Phi'_{\xi\bar{a}} = 0$*), then only imposing a minimum level of data protection requirement is sufficient.*

Proposition 7 prescribes a minimal data protection requirement, $C_{\min}$, in-line with a similar policy used by FTC, combined with a liability fine proportional to inflicted damage (as described in Definition 3) or a tax rate on information collected.

Intuitively, imposing a liability fine or a tax on information collected guarantees that socially-optimal data collection level ξ^W solves the business' first-order condition with respect to data

[17] In fact, the only combination of deviation from the social optimum that is *not possible in equilibrium* is over-protection and under-collection.

collection under socially-optimal data protection level C^W. At the same time, by reducing the level of data collection, a liability fine also eliminates any incentive the business may have to overprotect the data. A minimal protection level is then sufficient to guarantee that the business protects the collected data appropriately from a welfare perspective.

Recall, as per the discussion in Section 4.2 (and Footnote 4), that the current policy of the Federal Trade Commission requires a minimal level of data protection. This is sufficient to achieve the socially optimal behavior of a purely usage-driven business. On the other hand, to achieve the socially optimal behavior of a business with a considerable data-driven component, Proposition 7 suggests that the FTC's policy should be complemented by liability fines proportional to the damage to the users from adversarial activity or a tax rate on information collected. Whether the FTC's practice of responding to data breaches by seeking large fines is calibrated correctly, as to mimic our liability recommendation and induce a socially optimal data collection level, are open empirical questions.

6. Discussion and Broader Policy Implications

We conclude by elaborating on other broad policy issues that can be studied by expanding the framework that we have developed.

6.1. Government Adversaries

As we discuss in the introduction and thereafter, we collect under the title "adversary" any agent or entity whose usage of data harms users and thus reduces their utilities. That is, an adversary can be a hacker aiming at identity thefts, as well as a government agency seeking to use the data to track down and arrest users or to crack down on dissent.[18] Our main analysis thus focuses on the important commonalities across these different types of adversaries and distinguishes them only by their cost structures.

However, some correlations between adversary types and adversaries' cost structures are interest-

[18]For an example of users' arrests in the US, see "How ICE Picks Its Targets in the Surveillance Age", *The New York Times*, 17 April 2019. For a recent discussion on how COVID-19 contact-tracing apps can turn governments into surveillance states, see "Europe Rolls Out Contact Tracing Apps, With Hope and Trepidation", *The New York Times*, 16 June 2020 and "Coronavirus apps: the risk of slipping into a surveillance state", *Financial Times*, 28 April 2020.

ing. For example, a government agency, such as ICE or the FBI in the US, may find it much less costly to obtain data from a digital business, because they may be able to receive a court order. This may be even easier in authoritarian political regimes, in which the so-called adversary may have direct control of the courts. More importantly, a digital business may find it prohibitively expensive or even impossible to protect against a government adversary.[19] Finally, the ability of a government to act as an adversary is scalable, and may thus alter the distribution of adversaries costs altogether.

It is immediate from our model that both profits and consumer surplus will decline with a decrease in the cost adversaries incur to obtain data. Thus suggesting that government overreach leads to potentially harsher consequences than private hackers, and that to secure the future of welfare-enhancing digital services, appropriate checks and balances should be put in place and conducted by an organization that is, credibly, independent from the government.

6.2. Vertical Integration

The last two decades have seen significant consolidation of digital businesses, with many mergers and acquisitions being led by the Big Five Tech Giants or GAFAM (Google, Amazon, Facebook, Apple, and Microsoft).[20] Many of those acquisitions are vertical. That is, the acquired business operates in a separate market and/or provides a distinct service from the acquiring business. However, in privacy terms, even seemingly unrelated mergers may have an important effect. This is so because data may be shared between different subsidiaries of the same parent company.[21] A few natural questions arise: How would data collection practices change with data sharing? How will the merger affect privacy costs? Overall, how will data mergers affect profits, consumer surplus, and aggregate welfare?

Much in-line with public discourse, database mergers often involve businesses that can collect different types of information and have heterogeneous abilities to capitalize on the different types

[19] For example, a law approved in 2014 by Russia's president, Vladimir Putin, requires domestic and foreign companies to store the personal data of Russian citizens on servers in Russia. For details, see "Facebook and Twitter could be blocked in Russia in data storage row", *The Guardian*, 2 October 2019.

[20] See also https://www.visualcapitalist.com/the-big-five-largest-acquisitions-by-tech-company/

[21] For one of many examples, see https://www.eff.org/deeplinks/2020/04/google-fitbit-merger-would-cement-googles-data-empire about the Google-Fitbit merger. For a recent paper that analyzes how combining datasets could feed back into user behavior see Liang and Madsen 2020.

of data.[22] The following simple example demonstrates how our model can incorporate such considerations and shows that whether users benefit or suffer from a merger depends heavily on the extent to which the merger affects the businesses' decision making process with respect to data collection.

Consider the linear model that we introduce in Section 5 and two digital businesses $j \in \{1, 2\}$. The profit function of business j has now an added term $\delta_j \mathbb{1}_{merger} \xi_{-j} \bar{a}^*_{-j}$, where $\delta_j > 0$ and $\mathbb{1}_{merger}$ is the indicator function, taking a value of 1 if the businesses allow for reciprocal access to their databases and 0 otherwise (i.e., $\mathbb{1}_{merger} = 1$ if business 1 can access database of business 2 and vice versa). Users' utilities and adversaries' cost-benefit structure remain the same as before.

Now consider the following three scenarios: Scenario 0 is the pre-merger scenario, the two businesses act independently without sharing information. In Scenario DM (Data Merger), the businesses allow for reciprocal access to their databases but maintain autonomy over setting their data collection policies. Finally, in Scenario SM (Strategy Merger), in addition to the reciprocal access to data, businesses decide jointly on their data collection policies.

It is easy to see that in Scenario DM, businesses' data collection policies, user activity, privacy costs, and consumer surplus will remain the same as in Scenario 0, but profits will increase. In contrast, in Scenario SM the businesses will internalize the positive externalities they impose on each other by collecting information. As a result, data collection, profits, and privacy costs will increase, whereas user activity and consumer surplus will decrease relative to Scenarios 0 and DM. Intuitively, the move to Scenario SM is akin to a business becoming more data-driven, leading to increased data collection and a decrease in user activity.

Notably, public concerns with regards to mergers and acquisitions of digital businesses are not limited to privacy concerns. Another consideration is whether merging the data will affect market power in one of the markets in which the seemingly unrelated businesses operate in. Such discussions are clearly interesting. Methodologically, considering market power requires extending our model to allow for competition between digital businesses, which we defer to future work.

[22] See, e.g., https://www.propublica.org/article/google-has-quietly-dropped-ban-on-personally-identifiable-web-tracking

6.3. Pay For Data

Whether users should be paid for their data is an interesting question that pertains to the property rights over individual-level data and transcends the analysis of this paper.[23] However, our framework could be useful in evaluating the effect that pay-for-data schemes may have on data collection and data protection, and subsequently on consumer surplus and total welfare.

On the business's side, the payment for user data can be captured in our model in a way similar to the analysis tied to Corollary 2. However, in contrast to the tax per data stored, the payment goes to users and is proportional to their own usage and thus increases their activity.

In Online Appendix C.2, we illustrate, using an example of the linear model, that when a regulator requires businesses to pay users for their data, users' direct incentives to exert activity increase (since they are now paid for their data) while business's marginal returns from each unit of user information decrease. An interplay between the two defines whether the business collects a higher fraction of information or not and whether users benefit from such "pay for data" policy. We show that there exist regimes in which incentives of users to increase activity are so strong that the business can gain from an imposed payment by increasing data collection, ξ, in spite of the lowered returns from user information. When the price per data is sufficiently large, we find that consumer surplus can be higher than without the "pay for data" policy, and it can actually reach the socially optimal level.

Notably, information generated by different users might be interdependent, so that information generated by one user can reveal information on another. Such considerations could be introduced into our model by: (1) replacing $a_i\xi$ in user's i utility with a function of other users' usage levels, and (2) replacing everywhere $\bar{a}$ with a different function aggregating the usage levels of all users. We defer this analysis for future work and refer the interested reader to work on the topic by Acemoglu et al. 2019, Bergemann and Bonatti 2019, Gradwohl 2017, and Ichihashi 2020b.

[23] In related work, Arrieta-Ibarra et al. 2018 make the case that users should be paid for their data as if that data were labor, whereas Ichihashi 2019 explores a scenario in which competing data brokers compensate users for their data, and Bergemann and Bonatti 2019 and Acemoglu et al. 2019 study a setting where a data intermediary extracts users' information by offering monetary transfers. Emerging work in the marketing literature seeks to evaluate users' valuation of privacy via empirical and experimental approaches (see Lin 2019 and references therein).

6.4. Annoying Ads

In our narrative throughout the paper, we included the targeting of ads as a positive factor in users' utility functions. That is, users prefer relevant ads and information revealed by users through their activity improves the match between them and the ads they observe. However, in practice, targeting of certain ads may also lead to a reduction in users' utility. Consider, for instance, targeting adds of addictive products to vulnerable users (e.g., an AA member can be targeted with an ad for alcoholic beverages), similarly some ads may be misleading or even manipulative.[24]

At first glance, such annoying or harmful targeting does not seem to fit with our formulation of adversaries because the digital business sells the ads to the marketers. However, harmful targeting relies on data collected and quality targeting technology in order to be effective. Moreover, if, as digital businesses often argue, they prefer not to advertise harmful ads, they may invest in the better screening of the advertising content that they post. Circumventing the additional screening is costly to marketers who seek to post harmful ads, and thus the extra screening can be thought of as a form of data protection (or protection of the ability to target using the data). Our analysis of optimal regulation thus follows.

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.

D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar. Too much data: Prices and inefficiencies in data markets. Technical report, National Bureau of Economic Research, 2019.

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

[24] See, e.g., Liu et al. 2020 and references therein; see also "Facebook Says It Won't Back Down From Allowing Lies in Political Ads", *The New York Times*, 9 Jan 2020, on Facebook's policy regarding misleading campaign ads.

S. N. Ali and R. Bénabou. Image versus information: Changing societal norms and optimal privacy. *American Economic Journal: Microeconomics*, forthcoming, 2020.

S. N. Ali, G. Lewis, and S. Vasserman. Voluntary disclosure and personalized pricing. Technical report, National Bureau of Economic Research, 2019.

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*, 11:85–107, 2019.

D. Bergemann, A. Bonatti, and T. Gan. The economics of social data. *SSRN 3548336*, 2020.

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

A. Bonatti and G. Cisternas. Ratings-based price discrimination. *Working Paper*, 2017.

O. Candogan and K. Drakopoulos. Optimal signaling of content accuracy: Engagement vs. misinformation. *Operations Research*, 68(2):497–515, 2020.

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: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, 2016.

I. P. Fainmesser and A. Galeotti. Pricing network effects: Competition. *American Economic Journal*, forthcoming, 2020.

D. Fudenberg and J. M. Villas-Boas. Behavior-based price discrimination and customer recognition. *Handbook on economics and information systems*, 1:377–436, 2006.

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.

M. Hu, R. Momot, and W. Jianfu. Privacy management in service systems. *SSRN 3628751*, 2020.

S. Ichihashi. Dynamic privacy choices. *SSRN 3472151*, 2019.

S. Ichihashi. Online privacy and information disclosure by consumers. *American Economic Review*, 110(2):569–95, 2020a.

S. Ichihashi. The economics of data externalities. 2020b. URL https://shota2.github.io/research/externality.pdf.

O. Jann and C. Schottmüller. An informational theory of privacy. *The Economic Journal*, 130(625): 93–124, 2020.

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

M. A. Lariviere. A note on probability distributions with increasing generalized failure rates. *Operations Research*, 54(3):602–604, 2006.

A. Liang and E. Madsen. Data and incentives. 2020. URL https://bit.ly/31pJxfK.

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

J. Z. Liu, M. Sockin, and W. Xiong. Data privacy and temptation. 2020. URL http://wxiong.mycpanel.princeton.edu/papers/Privacy.pdf.

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

R. Momot, E. Belavina, and K. Girotra. The use and value of social information in selective selling of exclusive products. *Management Science*, forthcoming, 2019.

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. *Management Science*, 66(5):1826–1846, 2020.

A. Prat and T. M. Valletti. Attention oligopoly. *SSRN 3197930*, 2019.

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

T. Valletti and J. Wu. Consumer profiling with data requirements: Structure and policy implications. *Production and Operations Management*, 29(2):309–329, 2020.

A. Appendix: Proofs of the Main Results

A.1. Proof of Proposition 1 (Page 11)

Given users' expectation of the adversaries' demand for information 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 $\rho,\xi\in[0,1]$ and $\omega\geq 0$ 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 adversaries' demand for information is $\bar{a}(\omega)\xi/C$ which should be consistent with the initial belief ω, hence ω should solve $\omega=\frac{\bar{a}(\omega)\xi}{C}$. The expressions for $\bar{a}^*(\xi)$ and $\omega^*(\xi)$ then follow by substitution.

Consumer surplus can be derived as follows: $\mathrm{CS}(\xi)=\int U_i(a_i^*(\xi),\bar{a}^*(\xi))\mathrm{d}b_i$ where $a_i^*(\xi)=b_i+\beta\bar{a}^*(\xi)-\omega^*(\xi)\xi+\rho\xi$.

A.2. Proof of Corollary 1 (Page 12)

Derivative of $\bar{a}^*(\xi)$ 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 $\underline{\xi}$ which can be found as the largest solution to the corresponding quadratic equation. Similarly, derivative of $\xi\bar{a}^*(\xi)$ 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 $\bar{\xi}$ which is the largest solution to the corresponding quadratic equation. Finally, $\bar{\xi}>\underline{\xi}$ holds trivially when $\frac{\bar{b}}{\rho}<\frac{\rho C(1-\beta)}{\bar{b}}$, 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.

It always holds that $\underline{\xi} \geq 0$, for $\underline{\xi} \leq 1$ we need to require $C \leq \frac{2\bar{b}+\rho}{\rho(1-\beta)}$. Similarly, $\bar{\xi} \geq 0$ always, while for $\bar{\xi} \leq 1$ we need to require $C \leq \frac{\bar{b}}{(1-\beta)(\bar{b}+2\rho)}$. The latter inequality is binding and is guaranteed by Assumption 1.

A.3. Proof of Proposition 2 (Page 14)

First-order condition (FOC) is given by Eq. (8):

$$-\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}\Big/\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi} \equiv r(\xi) = \frac{\Phi'_{\bar{a}}(\bar{a}^*(\xi), \xi\bar{a}^*(\xi))}{\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi), \xi\bar{a}^*(\xi))}$$

We will first show that solution to FOC exists and is unique. $r'(\xi) = -\frac{2(C(1-\beta)+\xi^2)(\bar{b}^2+\rho^2C(1-\beta))}{(2\bar{b}\xi+\rho(\xi^2-C(1-\beta)))^2} \leq 0$. Also, $r(\underline{\xi}) = +\infty$ and $r(\bar{\xi}) = 0$. Furthermore, the RHS of the FOC is non-negative on $[\underline{\xi}, \bar{\xi}]$. Assuming that the RHS of the FOC is finite at $\underline{\xi}$ (i.e. that $\frac{\Phi'_{\bar{a}}(\bar{a}^*(\underline{\xi}),\underline{\xi}\bar{a}^*(\underline{\xi}))}{\Phi'_{\xi\bar{a}}(\bar{a}^*(\underline{\xi}),\underline{\xi}\bar{a}^*(\underline{\xi}))}$ is finite), existence follows. Denote $\xi^\star$ – a solution to FOC on $[\underline{\xi}, \bar{\xi}]$. We will show that $\frac{\partial}{\partial\xi}\left(\frac{\Phi'_{\bar{a}}(\bar{a}^*(\xi),\xi\bar{a}^*(\xi))}{\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi),\xi\bar{a}^*(\xi))}\right)\Big|_{\xi=\xi^\star} \geq 0$ from where uniqueness follows. Taking this derivative, we obtain:

$$\frac{1}{(\Phi'_{\xi\bar{a}})^2}\cdot\Big[\Phi'_{\xi\bar{a}}\cdot(\Phi''_{\bar{a},\bar{a}}\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi} + \Phi''_{\xi\bar{a},\bar{a}}\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}) - \Phi'_{\bar{a}}\cdot(\Phi''_{\xi\bar{a},\bar{a}}\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi} + \Phi''_{\xi\bar{a},\xi\bar{a}}\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi})\Big]$$

where we omitted repeated arguments of the functions for brevity. In order to evaluate this expression at $\xi^\star$, we substitute from FOC: $\Phi'_{\xi\bar{a}} = \Phi'_{\bar{a}}/r(\xi^\star)$ (derivatives are evaluated at $\xi^\star$). We obtain that the sign of the derivative of the RHS is defined by:

$$-\Phi'_{\bar{a}}\frac{(\mathrm{d}\bar{a}^*(\xi)/\mathrm{d}\xi)^2}{\mathrm{d}(\xi\bar{a}^*(\xi))/\mathrm{d}\xi}\Phi''_{\bar{a},\bar{a}} - 2\Phi'_{\bar{a}}\Phi''_{\xi\bar{a},\bar{a}}\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi} - \Phi'_{\bar{a}}\Phi''_{\xi\bar{a},\xi\bar{a}}\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi} \tag{16}$$

all evaluated at $\xi^\star$. Note that on $[\underline{\xi}, \bar{\xi}]$, we have $\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi} \leq 0$ and $\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi} \geq 0$, also we have $\Phi'_{\bar{a}} \geq 0$. From concavity of Φ, we have $\Phi''_{\bar{a},\bar{a}} \leq 0, \Phi''_{\xi\bar{a},\xi\bar{a}} \leq 0$ and $|\Phi''_{\xi\bar{a},\bar{a}}| \leq \sqrt{\Phi''_{\bar{a},\bar{a}}\Phi''_{\xi\bar{a},\xi\bar{a}}}$. If at $\xi^\star$, $\Phi''_{\xi\bar{a},\bar{a}} \geq 0$ then the expression (16) is positive. Similarly, if at $\xi^\star$, $\Phi''_{\xi\bar{a},\bar{a}} \leq 0$, then substitute $\Phi''_{\xi\bar{a},\bar{a}} = -\sqrt{\Phi''_{\bar{a},\bar{a}}\Phi''_{\xi\bar{a},\xi\bar{a}}}$ to expression (16) to find its lower bound which can be shown to be non-negative. Hence evaluated at $\xi^\star$, derivative $\frac{\partial}{\partial\xi}\left(\frac{\Phi'_{\bar{a}}}{\Phi'_{\xi\bar{a}}}\right)$ is positive, from where uniqueness of $\xi^\star$ on $[\underline{\xi}, \bar{\xi}]$ follows.

We will derive $\Pi''(\xi)$ to check for second-order conditions at $\xi^\star$:

$$\Phi''_{\bar{a},\bar{a}}\Big(\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}\Big)^2 + 2\Phi''_{\xi\bar{a},\bar{a}}\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi} + \Phi''_{\xi\bar{a},\xi\bar{a}}\Big(\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}\Big)^2 + \Phi'_{\bar{a}}\frac{\mathrm{d}^2\bar{a}^*(\xi)}{\mathrm{d}\xi^2} + \Phi'_{\xi\bar{a}}\frac{\mathrm{d}^2(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi^2}$$

The sign of the first three terms can be shown to be negative using concavity of Φ as above. The sign of the last two terms can be determined by substituting $\Phi'_{\xi\bar{a}} = \Phi'_{\bar{a}}/r(\xi^\star)$ from the FOC and calculating the derivatives explicitly, which leads to $\Phi'_{\bar{a}}\frac{2C(\bar{b}^2+C(1-\beta)\rho^2)}{(C(1-\beta)+\xi^2)(\bar{b}(\xi^2-C(1-\beta))-2\rho\xi C(1-\beta))}$ which can be shown to be negative for any $\xi < \bar{\xi}$. Hence, $\Pi''(\hat{\xi}) \leq 0$.

If $\Phi'_{\bar{a}} = 0$ then FOC is $r(\xi) = 0$ which is satisfied at $\bar{\xi}$. If $\Phi'_{\xi\bar{a}} = 0$ then FOC is $r(\xi) = \infty$ which is satisfied at $\underline{\xi}$. A decrease in the RHS of FOC (i.e., increase in $\Phi'_{\xi\bar{a}}/\Phi'_{\bar{a}}$) leads to higher solution $\xi^\star$ to FOC ($r(\xi)$ is decreasing, RHS is increasing at $\xi^\star$). The rest of the results follows from Corollary 1.

A.4. Proof of Proposition 3 (Page 15)

First-order condition has the following form for the welfare maximization:

$$r(\xi) = \frac{\alpha\bar{a}^*(\xi) + (1-\alpha)\Phi'_{\bar{a}}(\bar{a}^*(\xi), \xi\bar{a}^*(\xi))}{(1-\alpha)\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi), \xi\bar{a}^*(\xi))}$$

Existence of the solution can be shown in the same way as in the proof of Proposition 2. $\alpha = 0$ corresponds to the case of profit maximization. Rewrite the RHS of the FOC as $\frac{\alpha\bar{a}^*(\xi)}{(1-\alpha)\Phi'_{\xi\bar{a}}} + \frac{\Phi'_{\bar{a}}}{\Phi'_{\xi\bar{a}}}$, the second term is the same as in the FOC of Proposition 2 (profit maximization). The first term is positive. Given that $r(\xi)$ is decreasing, convex, conclude that $\forall\alpha$: $\xi^W < \xi^\star$ (unless $\Phi'_{\xi\bar{a}} = 0$) and properties of information and activity follow from the fact that $\xi^W \in [\underline{\xi}, \bar{\xi}]$. Notice that there might be several solutions ξ^W to FOC. Also, it is not guaranteed anymore that derivative of the RHS is positive at all such solutions ξ^W. Indeed, for the RHS, the derivative wrt ξ is

$$\frac{\left(\alpha\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi} + (1-\alpha)\left(\Phi''_{\bar{a},\bar{a}}\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi} + \Phi''_{\bar{a},\xi\bar{a}}\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}\right)\right)\Phi'_{\xi\bar{a}} - (\alpha\bar{a}^*(\xi) + (1-\alpha)\Phi'_{\bar{a}})\left(\Phi''_{\bar{a},\xi\bar{a}}\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi} + \Phi''_{\xi\bar{a},\xi\bar{a}}\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}\right)}{(1-\alpha)(\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi), \xi\bar{a}^*(\xi)))^2}$$

Where we omitted the arguments of derivatives of Φ for the sake of brevity. At ξ^W – solution of the FOC, substitute the numerator of the RHS of the FOC into the derivative. We thus get that the

sign of the derivative of the RHS at ξ^W is defined by:

$$\left(\alpha\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}+(1-\alpha)\left(\Phi''_{\bar{a},\bar{a}}\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}+\Phi''_{\bar{a},\xi\bar{a}}\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}\right)\right)\Phi'_{\xi\bar{a}}-(1-\alpha)\Phi'_{\xi\bar{a}}r(\xi)(\Phi''_{\bar{a},\xi\bar{a}}\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}+\Phi''_{\xi\bar{a},\xi\bar{a}}\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi})$$

Evaluated at ξ^W. After further simplification (and multiplication by $-\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}$) we obtain that the sign of this expression is defined by:

$$-\alpha\left(\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}\right)^2-(1-\alpha)\left[\Phi''_{\bar{a},\bar{a}}\left(\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}\right)^2+2\Phi''_{\bar{a},\xi\bar{a}}\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}+\Phi''_{\xi\bar{a},\xi\bar{a}}\left(\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}\right)^2\right]$$

all evaluated at ξ^W. The first term is negative. The expression in the brackets is shown to be negative using concavity of Φ in the proof of Proposition 2.

A.5. Proof of Proposition 4 (Page 17)

Given liability policy ℓ, business is choosing ξ to maximize $\Pi(\xi,\ell)=\Phi(\bar{a}^*(\xi),\xi\bar{a}^*(\xi))-\frac{\ell}{C}(\xi\bar{a}^*(\xi))^2$. The FOC is $\frac{\mathrm{d}}{\mathrm{d}\xi}(\Phi(\bar{a}^*(\xi),\xi\bar{a}^*(\xi)))=2\frac{\ell}{C}\xi\bar{a}^*(\xi)\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}$. From Proposition 3, socially optimal data collection policy ξ^W solves $\frac{\mathrm{d}}{\mathrm{d}\xi}(\Phi(\bar{a}^*(\xi),\xi\bar{a}^*(\xi)))=-\frac{\alpha}{1-\alpha}\bar{a}^*(\xi)\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}$. In order for liability $\ell^\star$ to correct inefficiency in data collection, it must be:

$$2\frac{\ell^\star}{C}\xi^W\bar{a}^*(\xi^W)\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}\Big|_{\xi=\xi^W}=-\frac{\alpha}{1-\alpha}\bar{a}^*(\xi^W)\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}\Big|_{\xi=\xi^W}$$

Recall that $r(\xi)=-\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}\Big/\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}$. Then the expression for $\ell^\star$ follows. If $\alpha=1$ then FOC for welfare maximization is $\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}=0$ which is satisfied at $\xi^W=\underline{\xi}$. Hence, optimal liability $\ell^\star$ follows from $\Phi'_{\bar{a}}\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}+\Phi'_{\xi\bar{a}}\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}=2\frac{\ell^\star}{C}\xi\bar{a}^*(\xi)\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}$ evaluated at $\xi=\underline{\xi}$. Notice that $\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}|_{\underline{\xi}}=\bar{a}^*(\underline{\xi})$. Then $\ell^\star=\Phi'_{\xi\bar{a}}\frac{C}{2\underline{\xi}\bar{a}^*(\underline{\xi})}$. Substituting $\bar{a}^*(\underline{\xi})$ and simplifying we get the formula for $\ell^\star$.

A.6. Proof of Corollary 2 (Page 18)

Proof is similar to the one of Proposition 4. Given tax t, the business is choosing ξ to maximize $\Pi(\xi,t)=\Phi(\bar{a}^*(\xi),\xi\bar{a}^*(\xi))-t\xi\bar{a}^*(\xi)$. The FOC is $\frac{\mathrm{d}}{\mathrm{d}\xi}(\Phi(\bar{a}^*(\xi),\xi\bar{a}^*(\xi)))=t\frac{\mathrm{d}(\xi\bar{a}^*(\xi))}{\mathrm{d}\xi}$. From Proposition 3, socially optimal data collection policy ξ^W solves $\frac{\mathrm{d}}{\mathrm{d}\xi}(\Phi(\bar{a}^*(\xi),\xi\bar{a}^*(\xi)))=-\frac{\alpha}{1-\alpha}\bar{a}^*(\xi)\frac{\mathrm{d}\bar{a}^*(\xi)}{\mathrm{d}\xi}$. In order for tax $t^\star$ to correct inefficiency in data collection, it must be that: $t^\star=\frac{\alpha}{1-\alpha}\frac{\bar{a}^*(\xi^W)}{r(\xi^W)}$. Case of $\alpha=1$ follows from the fact that $\xi^W=\underline{\xi}$ when $\alpha=1$.

A.7. Proof of Proposition 5 (Page 19)

Average activity in the case with no adversaries can be obtained by setting $\omega = 0$ in the proof of Proposition 1. Average damage from Definition 2 is $D(\xi) = \frac{(\xi \bar{a}^*(\xi))^2}{C}$. Then we have:

$$\mathcal{M}(\xi) = \frac{\text{CS}_{\text{no adversaries}}(\xi) - \text{CS}(\xi)}{D(\xi)} = \frac{C}{2\xi^2}\left(\left(\frac{\bar{b}+\rho\xi}{1-\beta}\right)^2 - \left(\frac{C(\bar{b}+\rho\xi)}{C(1-\beta)+\xi^2}\right)^2\right) \Big/ \left(\frac{C(\bar{b}+\rho\xi)}{C(1-\beta)+\xi^2}\right)^2$$

Simplifying we obtain $\mathcal{M}(\xi) = \frac{\xi^2+2C(1-\beta)}{2C(1-\beta)^2} = \frac{\xi^2}{2C(1-\beta)^2} + \frac{1}{1-\beta} \geq \frac{1}{1-\beta}$.

A.8. Second-Order Conditions for Maximization Problem (12) (Page 20)

First, solve first-order condition (13) and (14) wrt $\Phi'_{\bar{a}}$ and $\Phi'_{\xi\bar{a}}$. We get:

$$\Phi'_{\bar{a}}(\bar{a}^*(\xi, C)) = -K'(C)\frac{(\xi^2+\tilde{C})(\bar{b}(\xi^2-\tilde{C}) - 2\rho\xi\tilde{C})}{\xi^2(\bar{b}+\rho\xi)^2}$$
$$\Phi'_{\xi\bar{a}}(\bar{a}^*(\xi, C)) = K'(C)\frac{(\xi^2+\tilde{C})(2\bar{b}\xi + \rho\xi^2 - \tilde{C}\rho)}{\xi^2(\bar{b}+\rho\xi)^2}$$

where $\tilde{C} = C(1-\beta)$. As before, we will use shorthand notation $\Phi''_{\bar{a},\bar{a}}, \Phi''_{\xi\bar{a},\xi\bar{a}}, \Phi''_{\xi\bar{a},\bar{a}}$ to denote second-order derivatives of function Φ. For the rest of the proof, all functions are evaluated at optimal $\xi^\star, C^\star$, yet we will write ξ, C for simplicity. We will also use the following notation: $A \equiv \frac{\partial \bar{a}^*(\xi,C)}{\partial \xi}, B \equiv \frac{\partial(\xi\bar{a}^*(\xi,C))}{\partial \xi}$ and $D \equiv \frac{\partial \bar{a}^*(\xi,C)}{\partial C}, E \equiv \frac{\partial(\xi\bar{a}^*(\xi,C))}{\partial C}$. Then second-order derivatives of the function $\tilde{\Pi}(\xi, C) = \Phi(\bar{a}^*(\xi, C), \xi\bar{a}^*(\xi, C))$:

$$\frac{\partial^2\tilde{\Pi}(\xi,C)}{\partial\xi^2} = \Phi''_{\bar{a},\bar{a}}A^2 + 2\Phi''_{\xi\bar{a},\bar{a}}AB + \Phi''_{\xi\bar{a},\xi\bar{a}}B^2 + \Phi'_{\bar{a}}\frac{\partial^2\bar{a}^*(\xi,C)}{\partial\xi^2} + \Phi'_{\xi\bar{a}}\frac{\partial^2(\xi\bar{a}^*(\xi,C))}{\partial\xi^2}$$
$$\frac{\partial^2\tilde{\Pi}(\xi,C)}{\partial\xi\partial C} = \Phi''_{\bar{a},\bar{a}}AD + \Phi''_{\xi\bar{a},\bar{a}}(AE+BD) + \Phi''_{\xi\bar{a},\xi\bar{a}}BE + \Phi'_{\bar{a}}\frac{\partial^2\bar{a}^*(\xi,C)}{\partial\xi\partial C} + \Phi'_{\xi\bar{a}}\frac{\partial^2(\xi\bar{a}^*(\xi,C))}{\partial\xi\partial C}$$
$$\frac{\partial^2\tilde{\Pi}(\xi,C)}{\partial C^2} = \Phi''_{\bar{a},\bar{a}}D^2 + 2\Phi''_{\xi\bar{a},\bar{a}}DE + \Phi''_{\xi\bar{a},\xi\bar{a}}E^2 + \Phi'_{\bar{a}}\frac{\partial^2\bar{a}^*(\xi,C)}{\partial C^2} + \Phi'_{\xi\bar{a}}\frac{\partial^2(\xi\bar{a}^*(\xi,C))}{\partial C^2}$$

Substituting $\Phi'_{\bar{a}}, \Phi'_{\xi\bar{a}}$ found above, we get:

$$\tilde{\Pi}''_{\xi\xi}(\xi, C) = \Phi''_{\bar{a},\bar{a}}A^2 + 2\Phi''_{\xi\bar{a},\bar{a}}AB + \Phi''_{\xi\bar{a},\xi\bar{a}}B^2 + \lambda_{\xi\xi}, \text{ where } \lambda_{\xi\xi} = -2K'(C)\frac{C(\bar{b}^2 + \tilde{C}\rho)}{\xi^2(\bar{b} + \rho\xi)^2} < 0$$

$$\tilde{\Pi}''_{\xi C}(\xi, C) = \Phi''_{\bar{a},\bar{a}}AD + \Phi''_{\xi\bar{a},\bar{a}}(AE + BD) + \Phi''_{\xi\bar{a},\xi\bar{a}}BE + \lambda_{\xi C}, \text{ where } \lambda_{\xi C} = 2K'(C)\frac{\tilde{C}}{\xi(\xi^2 + \tilde{C})} > 0$$

$$\tilde{\Pi}''_{\xi\xi}(\xi, C) = \Phi''_{\bar{a},\bar{a}}D^2 + 2\Phi''_{\xi\bar{a},\bar{a}}DE + \Phi''_{\xi\bar{a},\xi\bar{a}}E^2 + \lambda_{CC}, \text{ where } \lambda_{CC} = -2K'(C)\frac{(1-\beta)}{\xi^2 + \tilde{C}} < 0$$

It can be easily shown that $\tilde{\Pi}''_{\xi\xi}(\xi, C) \leq 0$ and $\tilde{\Pi}''_{CC}(\xi, C) \leq 0$ using concavity of Φ (e.g., see, proof of Proposition 2). Finally, $\tilde{\Pi}_{\xi\xi}(\xi, C)\tilde{\Pi}_{CC}(\xi, C) - \tilde{\Pi}^2_{\xi C}(\xi, C)$ after simple algebraic manipulations can be rewritten as:

$$\left(\Phi''_{\bar{a},\bar{a}}\Phi''_{\xi\bar{a},\xi\bar{a}} - (\Phi''_{\xi\bar{a},\bar{a}})^2\right)(BD - AE)^2 + \lambda_{CC}\lambda_{\xi\xi} - \lambda^2_{\xi C} + 2\Phi''_{\xi\bar{a},\bar{a}}(AB\lambda_{CC} - (BD + AE)\lambda_{\xi C} + DE\lambda_{\xi\xi}) +$$
$$\Phi''_{\bar{a},\bar{a}}(A^2\lambda_{CC} - 2AD\lambda_{\xi C} + D^2\lambda_{\xi\xi}) + \Phi''_{\xi\bar{a},\xi\bar{a}}(B^2\lambda_{CC} - 2BE\lambda_{\xi C} + E^2\lambda_{\xi\xi})$$

The first term is non-negative due to concavity of Φ. The second and the third terms can be combined and simplified (substitute the expressions for λ) to $\frac{4\tilde{C}(\bar{b}\xi - \rho\tilde{C})^2K'(C)^2}{\xi^2(\xi^2+\tilde{C})^2(\bar{b}+\rho\xi)^2}$ which is non-negative. Thus, we are left to determine the sign of the last three terms. We can rewrite those as follows:

$$\frac{2CK'(C)}{(\xi^2 + \tilde{C})^4}\left(-\Phi''_{\bar{a},\bar{a}}\mu^2 - 2\Phi''_{\xi\bar{a},\bar{a}}\mu\eta - \Phi''_{\xi\bar{a},\xi\bar{a}}\eta^2 - \Phi''_{\xi\bar{a},\xi\bar{a}}\tilde{C}\xi^2(\bar{b} + \rho\xi)^2\right) \tag{17}$$

where $\mu = \bar{b}\xi - \tilde{C}\rho$ and $\eta = \bar{b}(\xi^2 - \tilde{C}) - 2\rho\xi\tilde{C}$. Notice that sign of μ can be either positive or negative, while $\eta < 0$ since $\xi < \bar{\xi}$ (by Assumption 1), where $\bar{\xi}$ solves $\eta = 0$. The multiplier of expression (17) is positive. The last term is also positive since $\Phi''_{\xi\bar{a},\xi\bar{a}} \leq 0$. Finally, irrespective of the sign of μ The first three terms can also be shown positive using concavity of Φ (applying the same technique as in the proof of Proposition 2). Hence, we showed that at optimal $\xi^\star, C^\star$ Hessian of $\tilde{\Pi}(\xi, C)$ is negative semi-definite. Given that $K''(C) \geq 0$, we have Hessian of the profit function $\Pi(\xi, C)$ is also negative semi-definite at $\xi^\star, C^\star$.

A.9. Proof of Proposition 6 (Page 21)

Result follows from the proof A.8. Evaluate $\tilde{\Pi}_{\xi C}(\xi, C)$ at $\xi^\star, C^\star$. We know that $\lambda_{\xi C} > 0$. The rest can be rewritten as follows:

$$\frac{\partial \bar{a}^*(\xi,C)}{\partial C}\frac{\partial \bar{a}^*(\xi,C)}{\partial \xi}\left(\Phi''_{\bar{a},\bar{a}} + 2\Phi''_{\xi\bar{a},\bar{a}}\xi + \xi^2\Phi''_{\xi\bar{a},\xi\bar{a}}\right) + \bar{a}^*(\xi,C)\frac{\partial \bar{a}^*(\xi,C)}{\partial C}\left(\Phi''_{\xi\bar{a},\bar{a}} + \xi\Phi''_{\xi\bar{a},\bar{a}}\right)$$

The first term is non-negative due to concavity of function Φ (see proof of Proposition 2). The second term is non-negative by assumption of the Proposition.

A.10. Proof of Proposition 7 (Page 24)

We will prove statement of the proposition for the case when liability policy $\ell^\star$ is used. Proof when tax rate $t^\star$ is used is similar. The profit function of the business under liability policy is $\Pi(\xi, C, \ell) = \Phi(\bar{a}^*(\xi, C), \xi\bar{a}^*(\xi, C)) - K(C) - \frac{\ell}{C}[\xi\bar{a}^*(\xi, C)]^2$. The first-order conditions are:

$$\begin{aligned}
\frac{\partial \Pi(\xi,C,\ell)}{\partial \xi} &= \frac{\partial \Phi(\bar{a}^*(\xi,C),\xi\bar{a}^*(\xi,C))}{\partial \xi} - 2\frac{\ell}{C}\bar{a}^*(\xi,C)\xi\frac{\partial(\xi\bar{a}^*(\xi,C))}{\partial \xi} \\
\frac{\partial \Pi(\xi,C,\ell)}{\partial C} &= \frac{\partial \Phi(\bar{a}^*(\xi,C),\xi\bar{a}^*(\xi,C))}{\partial C} + \frac{\ell}{C}\bar{a}^*(\xi,C)\xi^2\left(\frac{\bar{a}^*(\xi,C)\xi}{C} - 2\frac{\partial(\xi\bar{a}^*(\xi,C))}{\partial C}\right)
\end{aligned}$$

Now, evaluate these conditions at socially-optimal ξ^W, C^W. We will suppress subscript W from ξ^W, C^W for the rest of the proof. Substitute $\frac{\partial \Phi}{\partial \xi}, \frac{\partial \Phi}{\partial C}$ from the first-order conditions for welfare function (expression (15) and its counterpart for ξ), also substitute $\ell^\star$ from the expression (10). After simplification, we obtain:

$$\begin{aligned}
\frac{\partial \Pi(\xi,C,\ell^\star)}{\partial \xi}\Big|_{\xi^W,C^W} &= 0 \\
\frac{\partial \Pi(\xi,C,\ell^\star)}{\partial C}\Big|_{\xi^W,C^W} &= \frac{C^W\rho\xi^W(\bar{b}+\rho\xi^W)^2}{2((\xi^W)^2 + C^W(1-\beta))(\bar{b}((\xi^W)^2 - C^W(1-\beta)) - 2\rho\xi^W C^W(1-\beta))} < 0
\end{aligned}$$

Where the second inequality follows from the fact that ξ^W, C^W are such that $\xi^W < \bar{\xi}$ (based on Assumption 1) – see Corollary 1. Thus, business under minimum data protection $C \geq C_{\min}$ requirement and liability policy $\ell^\star$, sets ξ^W, C^W. In case if $\Phi'_{\xi\bar{a}} = 0$, welfare is maximized when activity $\bar{a}^*(\xi, C)$ is maximized, hence condition $\frac{\partial \Pi(\xi,C)}{\partial \xi}|_{\xi^W,C^W,\ell} = 0$ is satisfied with $\ell = 0$.

Online Appendices

B. Hybrid Revenue Model – an Additional Example for Section 2.5 (Page 10)

Consider a digital business with the revenue model that has both an information-driven component and a usage driven component, say Facebook. Using Facebook is free. Private benefit and cost of using the service as well as positive network effects are captured by the first and the second terms of expression 1. Each user i has a characteristic $\theta_i \in \{0, 1, \ldots, n\}$, which captures a unit demand to one of a class of products. θ_i is ex-ante unknown to the business. However, if business knew this characteristic, then it could target user i with a product of value V to her. Note that, in contrast with a pure data-driven business model, Facebook wishes that users are active and use its service in order to be able to show them ads.

The business knows that for any $m \in \{0, 1, \ldots, n\}$, $\theta_i = m$ with probability $1/n$. The user's activity provides signals to the business about θ_i, but only if this activity is registered and processed. In particular, if the business sets ξ and the user exerts activity level a_i, then ξa_i is the probability that the business observes the true realization of θ_i, whereas with the remaining probability the business learns nothing. In this case, the ex-ante expected probability that the platform creates value V to the user is:

$$\frac{1}{n}\left[1 + (n-1)\, a_i \xi\right].$$

Depending on the competitiveness of the market, this value will be shared between the advertiser and the user, and then the platform is able to extract payments from the advertiser via advertising fees (described below). Let S be the share that is extracted by the advertiser. Then, the positive externality to the user is $V(1-S)\left[1 + (n-1)\, a_i \xi\right]/n$ and $\rho = V(1-S)\,(n-1)/n$.

In addition to revenues from selling a product (in expectation $SV\left[1 + (n-1)\, a_i \xi\right]/n$), advertisers benefit from views that contribute to brand recognition. That is, the total value to an advertiser (or advertisers) from targeting a user is $SV\left[1 + (n-1)\, a_i \xi\right]/n + v a_i$ where v is the value in terms of brand recognition of each ad-view. Let s be the share of this value extracted by the platform in the form of per view fees. Then the profit to the platform from user i is $s\left(SV\left[1 + (n-1)\, a_i \xi\right]/n + v a_i\right)$.

C. Additional Analysis

C.1. General User Utility and Adversary Specification

Assume that user i's utility is generalized to the following form: $U_i(a_i) = U[a_i, \xi(\rho - \omega)]$.[25] The first argument is the user's activity, while the second argument is the user's expectation on the benefit/downside of exerting each unit of activity. We assume that function U is concave in its two arguments and the arguments are complements. We will denote U'_x, U'_y – partial derivatives of U wrt the first and the second arguments respectively. The first-order condition for the user is thus $U'_x(a, \xi(\rho - \omega)) = 0$.

We assume that adversaries' abilities γ are distributed with cdf G and pdf g. We also assume that distribution of γ satisfies increasing generalized failure rate (IGFR) property (see Lariviere 2006). In other words, we have that $xg(x)/(1 - G(x))$ increases in x or alternatively $g'(z) \geq -\frac{g(z)^2}{1-G(z)}$.

The main driving force of the results of the paper was the fact that the equilibrium average activity was non-monotone in the level of information collection policy ξ. In particular, activity increases with ξ for low ξ and it decreases in ξ when ξ is large. This non-monotonicity in user activity is the result of the interplay between positive and negative externalities that information imposes on users along with the adversaries' endogenous demand for information. When ξ is low, adversaries' demand for information is small, hence information produces more positive than negative externality to users – thus, an increase in ξ incentivises users to be more active. When ξ is large, the adversaries' demand for information is strong, and so negative externalities dominate, and users decrease their activity as ξ grows. We can show that the behavior of user activity inherits such traits in this more general model.

Similar to the result of Proposition 1 we can solve for user's equilibrium activity. In particular, it is such $a^*(\xi)$ solves: $U'_x(a, \xi\rho - \xi G(a\xi/C)) = 0$. Derivative of user activity wrt ξ:

$$\frac{\mathrm{d}a^*(\xi)}{\mathrm{d}\xi} = -\frac{U''_{x,y}(a, \xi\rho - \xi G(a\xi/C))[\rho - G(a\xi/C) - g(a\xi/C)a\xi/C]}{U''_{x,x}(a, \xi\rho - \xi G(a\xi/C)) - U''_{x,y}(a, \xi\rho - \xi G(a\xi/C))g(a\xi/C)\xi^2/C}|_{a=a^*(\xi)}$$

Thus, the sign of this expression is defined by $\mu(\xi) = \rho - G(a^*(\xi)\xi/C) - g(a^*(\xi)\xi/C)a^*(\xi)\xi/C$. At $\xi = 0$ this expression is positive. Its derivative wrt ξ is $\mu'(\xi) = (a^*(\xi)'\xi + a^*(\xi))(-2g(a^*(\xi)\xi/C) -$

[25] For simplicity, we omit network effects.

$g'(a^*(\xi)\xi/C))$. In order to show that $a^*(\xi)$ first increases and then decreases with ξ, we need to show that $\mu(\xi)$ either always remains positive (then $a^*(\xi)$ always increases) or crosses 0 only once. Assume that there exists $\hat{\xi}$, such that $\mu(\hat{\xi}) = 0$, then also $\frac{\mathrm{d}a^*(\xi)}{\mathrm{d}\xi}|_{\hat{\xi}} = 0$. Using IGFR property of distribution of γ, we conclude that sign of $\mu'(\hat{\xi})$ is defined by $-2 + 2G(a^*(\hat{\xi})\hat{\xi}/C) + g(a^*(\hat{\xi})\hat{\xi}/C)$ which is negative given that we know that at $\hat{\xi}$: $\rho = G(a^*(\hat{\xi})\hat{\xi}/C) + g(a^*(\hat{\xi})\hat{\xi}/C)$. Clearly $a^*(\xi)\xi$ increases with ξ when $a^*(\xi)$ increases. Behavior of $a^*(\xi)\xi$ when $a^*(\xi)$ decreases is defined by the sign of $\xi U''_{x,y}(\ldots)(\rho - G(a^*(\xi)\xi/C)) - U''_{x,x}(\ldots)a^*(\xi)$ and depends on the higher-order derivatives of function U. Notice that for $a^*(\xi)\xi$ to decrease, necessarily it must be that $\rho < \omega = G(a^*(\xi)\xi/C)$ or $a^*(\xi)\xi$ is high enough. $a^*(\xi)\xi$ increases first and then decreases if $\xi(\rho - \omega)/a^*(\xi)$ single-crosses $U''_{x,x}(a^*(\xi), \xi(\rho - \omega))/U''_{x,y}(a^*(\xi), \xi(\rho - \omega))$.

C.2. Paying Users for Data (Section 6.3, Page 28)

Users' utility is modified by adding $t\bar{a}\xi$, which, in the case of exogenous t is equivalent to setting $\tilde{\rho} = \rho + t$. Thus, $\bar{a}^*(\xi)$ is modified by increasing ρ to $\tilde{\rho}$. Denote modified users' response as $\tilde{a}^*(\xi)$. Then business's profit without paying for data is $\Pi(\xi) = \Phi(\bar{a}^*(\xi), \xi\bar{a}^*(\xi))$ and when paying for data it is $\tilde{\Pi} = \Phi(\tilde{a}^*(\xi), \xi\tilde{a}^*(\xi)) - t\tilde{a}^*(\xi)\xi$. We will further assume linear function form for Φ, such that $\Phi(\bar{a}, \xi\bar{a}) = P_u\bar{a} + P_d\xi\bar{a}$. Also wlog let $\beta = 0$.

Equilibrium data collection policy of the business which doesn't pay for data is $\xi^\star$ such that solves (8). Equilibrium data collection policy of the business which pays for data $\xi^\star(t)$ has the following form: $P_u\frac{\mathrm{d}\tilde{a}^*(\xi)}{\mathrm{d}\xi} + (P_d - t)\frac{\mathrm{d}(\xi\tilde{a}^*(\xi))}{\mathrm{d}\xi} = 0$. Solving for equilibrium $\xi^\star(t)$ we obtain: $\xi^\star(t) = -\kappa(t) + \sqrt{\kappa(t)^2 + C}$ – decreasing in $\kappa(t)$, where $\kappa(t) = \frac{\bar{b}P_u - (P_d - t)(\rho + t)C}{(P_d - t)\bar{b} + P_u(\rho + t)}$. The sign of the derivative of $\kappa(t)$ is defined by $t^2C(P_u - \bar{b}) + 2tC(\bar{b}P_d + \rho P_u) + \bar{b}P_u(\bar{b} - P_u) + C(P_u\rho^2 - \bar{b}P_d^2)$. Notice that if $\bar{b} = 0$ then this expression is positive for any t and hence, $\xi^\star(t) \leq \xi^\star$ for any t. If $\bar{b} > 0$, it may happen that $\xi^\star(t) > \xi^\star$. Consider, for instance, $t \to 0$, then sufficient condition for $\xi^\star(t)'_t > 0$ is $\bar{b}P_u(\bar{b} - P_u) + C(P_u\rho^2 - \bar{b}P_d^2) < 0$. If this condition holds, then at least for small t business collects higher fraction of user information than with $t = 0$. Now, consider $t = P_d$, then business sets $\underline{\xi}(t = P_d)$ (s.t. $\frac{\mathrm{d}\tilde{a}^*(\xi)}{\mathrm{d}\xi} = 0$) – data collection policy maximizing consumer surplus.