Dynamic Privacy Choices

Shota Ichihashi*

December 14, 2020

Abstract

I study a dynamic model of consumer privacy and platform data collection. In each period, consumers choose their level of platform activity. Greater activity generates more information about the consumer, thereby increasing platform profits. Although consumers value their privacy, the platform can collect information by gradually lowering the level of privacy protection. In the long run, consumers become "addicted" to the platform: They lose privacy and receive low payoffs, but choose high activity levels. I study the implications of these dynamics on the platform's business model, its commitment power regarding the future privacy policy, and competition between platforms.

Keywords: personal information; consumer privacy; privacy paradox; data policy; gradualism

JEL Classification: D44, D82, D83.

*Bank of Canada, 234 Wellington Street West, Ottawa, ON K1A 0G9, Canada. Email: shotaichihashi@gmail.com.
I thank Jason Allen, Byung-Cheol Kim, David Miller, Michael Sockin (discussant), Takuo Sugaya, Catherine Tucker, Lei Xu, and seminar and conference participants at the Bank of Canada, Peking University, University of Rochester (Simon), Privacy Workshop at Princeton, and Canadian Economic Theory Conference 2020, EC'20, and SITE 2020 for valuable comments. The opinions expressed in this article are the author's own and do not reflect the views of the Bank of Canada.

1 Introduction

Online platforms, such as Amazon, Facebook, Google, and Uber, analyze user activities and collect a large amount of data. This data collection may improve their services and benefit consumers, but it also raises concerns for consumers and policymakers (Crémer et al., 2019; Furman et al., 2019; Morton et al., 2019).

As an example, consider a consumer (she) and a social media platform (it). The consumer writes posts and reads news on the platform. The platform analyzes her activity and collects data such as her race, location, and political preferences. The platform can then generate revenue—e.g., via improved targeted advertising. The consumer faces a trade-off: On the one hand, she enjoys the services provided by the platform. On the other hand, she may value her privacy, or be concerned about the risk of data leakage, identity theft, and price or non-price discrimination.[1] Such risks are the "privacy costs" of using the platform. If the consumer anticipates a high privacy cost, she may use the platform less actively, or may not join it. The platform can influence her decision through its privacy policy. For example, Facebook committed to not use cookies to track users.[2]

I model such a situation as a dynamic game between a consumer and a platform. In each period, the consumer chooses her level of platform activity. Based on the level of activity, the platform observes a signal about the consumer's time-invariant type. The informativeness of the signal is increasing in the activity level, but decreasing in the platform's privacy level, which specifies the amount of noise added to the signal. The platform's profit is increasing, but the consumer's payoff is decreasing in the amount of information the platform has collected. As a result, the consumer chooses activity levels that balance the benefits of the service and the privacy costs. Anticipating the consumer's behavior, the platform chooses privacy levels.

The main idea is that the consumer has a decreasing marginal privacy cost—i.e., when the consumer has less privacy, she faces a lower marginal privacy cost of using the same platform. For example, if Google already knows a lot about a consumer, she might not care about letting Google Maps track her location today. In an extreme case, if the platform knows everything, the consumer faces a marginal privacy cost of zero, because her activity no longer affects what the

[1]Such concerns are highlighted by, for example, the Cambridge Analytica scandal.

[2]In 2004, Facebook's privacy policy stated that "we do not and will not use cookies to collect private information from any user." https://web.archive.org/web/20050107221705/http://www.thefacebook.com/policy.php (accessed on July 31, 2020)

platform knows about her type.

The main finding is that because of the decreasing marginal cost of losing privacy, a farsighted consumer eventually gives up all of her privacy, even when she anticipates future privacy choices made by the platform. To induce such an outcome, the platform commits to a high level of privacy protection in early periods: By doing so, the platform can encourage the consumer to use the service and generate information, when she has not yet lost her privacy. However, as the platform collects more data, the consumer faces lower marginal privacy costs. As a result, in later periods, the platform can decrease a privacy level to speed up data collection. In the long run, the consumer loses privacy and incurs a high privacy cost, but chooses a high activity level. This result does not depend on discount factors: For example, even if the platform is myopic and faces a patient consumer, it adopts a privacy policy that causes the long-run privacy loss.

I first show the above result in the baseline model in which the platform's revenue depends only on information it collects. However, the main insight holds even if the platform's revenue depends on the consumer's activity. Indeed, even if the platform earns revenue mainly from consumer activity, it may eventually collect as much information as a data-driven platform, because the consumer chooses the highest activity level when she has no privacy on the platform.

I then explore the implication and robustness of the main finding. First, I study the role of the platform's commitment power regarding its future privacy choices. If the platform can commit to future privacy levels, it can collect information by committing to offer high but finite privacy levels in early periods. Under a certain condition, the platform can implement the same policy as long as it has one-period commitment power. However, if the platform faces a high prior uncertainty about the consumer's type and cannot commit to a future privacy policy, there is also an equilibrium in which it fails to collect any information: The consumer refuses to provide data, because the platform, which fails to collect data today, will offer high privacy protection in the future. This equilibrium captures the platform's Coasian commitment problem.

Second, I examine the implication of the aforementioned decreasing marginal privacy cost on competition and regulations. The decreasing marginal cost implies that the consumer is more willing to use a platform on which she has less privacy. This consumer's tendency renders competition less effective. Also, ex ante and ex post privacy regulations have different impacts: Mandating that the platform pre-commit to a strict privacy policy may, perversely, lower the privacy and welfare of consumers in the long run. In contrast, enabling the consumer to delete information collected in the past may enhance welfare.

The paper has implications for consumer privacy. First, the consumer's long-run behavior (in the equilibrium with data collection) is consistent with the so-called privacy paradox: Consumers express concern about their privacy, but actively share data with third parties (Acquisti et al., 2016). The platform's equilibrium strategy rationalizes how online platforms, such as Facebook, seem to have expanded the scope of data collection. Second, my results clarify the role of commitment and expectation in data collection: Depending on consumers' expectation about their future privacy, the platform may collect data when consumers highly value their privacy, or it may fail to collect data when consumers do not much value their privacy.

The rest of the paper is as follows. Section 2 discusses related literature, and Section 3 presents the model. Section 4 considers the platform with long-run commitment power and presents the equilibrium. Section 5 assumes the platform has one-period commitment. In particular, assuming that the consumer has binary activity level, I characterize an equilibrium that is best or worst for the consumer. Section 6 studies competition between platforms. Section 7 considers extensions, including the impact of erasing past information.

2 Related Literature

This paper contributes to the literature on the economics of privacy and markets for data. This literature has studied several important questions, such as how to use consumer data to create market segmentation (Ali et al., 2020; Bonatti and Cisternas, 2020; Elliott and Galeotti, 2019; Haghpanah and Siegel, 2019; Loertscher and Marx, 2020; Yang, 2019; Ichihashi, 2020b); how to choose the optimal level of privacy protection and information security (Dwork et al., 2014; Fainmesser et al., 2019; Jullien et al., 2018); how information externalities create inefficiency or influence agents' behavior (Acemoglu et al., 2019; Bergemann et al., 2019; Choi et al., 2019; Easley et al., 2018; Liang and Madsen, 2020; Ichihashi, 2020a); how agents strategically manipulate data (Frankel and Kartik, 2019b,a; Argenziano and Bonatti, 2020; Ball, 2020); how consumer data and privacy interact with mechanism design (Brunnermeier et al., 2020; Calzolari and Pavan, 2006; Eilat et al., 2019; Ghosh and Roth, 2011); how to price and sell information (Agarwal et al., 2019;

Hörner and Skrzypacz, 2016; Bergemann et al., 2018); and how privacy and data affect competition (Casadesus-Masanell and Hervas-Drane, 2015; De Corniere and Taylor, 2020). I contribute to this literature by studying a firm's dynamic policy to acquire consumer data, and how the policy depends on the firm's commitment power and consumer expectation.

The paper is especially related to Acemoglu et al. (2019); Bergemann et al. (2019); and Choi et al. (2019). They consider static models in which a platform collects data in exchange for money. In their models, the data on some consumers reveal information about others. Under a certain information structure, this "data externality" lowers consumers' private costs of providing data relative to social costs. In this case, the equilibrium involves an inefficiently high level of data sharing.[3] In my paper, the consumer's cost of generating information is decreasing in the stock of data she provided in the past and the amount of data the platform will collect in the future. In this case, depending on the consumer's expectation, the equilibrium may involve an inefficiently high or low level of data collection. The dynamic model also enables me to study new issues, such as a platform's commitment and the impact of erasing past data.

This paper also relates to recent work on dynamic competition in digital markets. Hagiu and Wright (2020) study "data-enabled learning," whereby firms can improve their products and services through learning from the data they obtain from their customers. Prufer and Schottmüller (2017) assume that the cost of investing in quality is decreasing in the firm's past sales, and greater investment in quality leads to higher demand in the current period. In contrast to this literature, I assume data collection lowers consumer welfare. Such an assumption enables us to study issues related to consumer privacy. Hagiu and Wright (2020) allow price competition and study rich learning dynamics that incorporate "within-user" and "across-user" learning. In contrast, I abstract away from pricing, and focus on within-user learning and the design of a privacy policy.

How the consumer's incentive changes over time in my model is similar to that of career concern models, which originated with Holmström (1999). In career concern models, a young worker, whose ability has not yet been revealed to the market, works hard to influence the market's belief. In my model, a consumer who has not yet lost privacy uses the platform less actively to generate less information. Over time, the information about the consumer and the worker are revealed, and

[3] Bergemann et al. (2019) also consider an information structure under which the data externality renders the private cost greater than the social cost, which may lead to an inefficiently low level of data sharing.

they have lower incentives to engage in signal jamming. Despite this connection, the two signal jamming activities are different. In career concern models, the market wants the worker to engage in signal jamming, which corresponds to higher effort. Thus, there is a trade-off between learning the worker's ability and motivating high effort (e.g., Hörner and Lambert 2018). In my model, the platform wants the consumer to engage less in signal jamming. Thus, the platform prefers to collect information not only to increase profit today, but also to motivate the consumer to raise activity levels in the future. Many of my results stem from this complementarity between data collection and consumer activity, which is absent in career concern models.

3 Model

I study a dynamic game between a consumer (she) and a platform (it). The consumer uses the platform's service to receive benefits, but her use of the service generates information about her time-invariant (Gaussian) type. The platform chooses a privacy level, which is the amount of noise added to the information generated. The platform provides its service for free and monetizes the information. I model payoffs in a reduced-form way, so that, in the baseline model, the platform prefers more information and the consumer prefers less information to be collected. Appendix A microfounds such preferences, assuming that the platform sells data to third-party sellers that price discriminate the consumer.

The formal description is as follows. Time is discrete and infinite, indexed by $t \in \mathbb{N}$. The consumer's type X is drawn from a normal distribution $\mathcal{N}(0, \sigma_0^2)$. The type is realized before $t = 1$ and fixed over time. The consumer does not observe X.[4] The platform does not observe X either, but receives signals about it.

In each period $t \in \mathbb{N}$, the consumer chooses an *activity level* a_t from a finite set $A \subset \mathbb{R}_+$ such that $\min A = 0$ and $a_{max} := \max A > 0$. The platform then observes a_t and a signal $s_t = X + \varepsilon_t$, where $\varepsilon_t \sim \mathcal{N}\left(0, \frac{1}{a_t} + \gamma_t\right)$. The consumer does not observe the signal.[5] A higher a_t reduces the

[4] Even if the consumer privately observes X, all results hold with respect to a pooling equilibrium in which consumers of all types choose the same activity level after any history. Such an equilibrium exists because the payoff of each player does not depend on a realization of X. Unobservable X simplifies exposition without changing the results.

[5] All the results continue to hold even if signals are public, because the payoff of each player does not depend on the realization of a signal.

variance of ε_t and makes s_t more informative about X. For a fixed a_t, the informativeness of the signal decreases in $\gamma_t \in \overline{\mathbb{R}}_+ := \mathbb{R}_+ \cup \{\infty\}$, which is the *privacy level* of the platform in period t. A higher γ_t implies the platform offers higher privacy protection. If $a_t = 0$ or $\gamma_t = \infty$, signal s_t is totally uninformative. Random variables X and $(\varepsilon_t)_{t\in\mathbb{N}}$ are mutually independent.

The payoffs are as follows. Suppose that the consumer has chosen activity levels $\boldsymbol{a}_t = (a_1, \dots, a_t) \in A^t$ and the platform has chosen privacy levels $\boldsymbol{\gamma}_t = (\gamma_1, \dots, \gamma_t) \in \overline{\mathbb{R}}_+^t$ up to period t. At the end of period t, the platform receives a payoff of $\sigma_0^2 - \sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t) \geq 0$, where $\sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ is the posterior variance of X given $(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ and Bayes' rule.[6] A small $\sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ means the platform has an accurate estimate of X, or equivalently, the consumer has low privacy. For any t and $\tau \leq t$, $\sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ is decreasing in a_τ, increasing in γ_τ, and independent of s_τ.[7] Where it does not cause confusion, I write $\sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ as σ_t^2. The platform discounts future payoffs with discount factor $\delta_P \in (0, 1)$.

The consumer's flow payoff in period t is $U(\boldsymbol{a}_t, \boldsymbol{\gamma}_t) := u(a_t) - v \cdot [\sigma_0^2 - \sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)]$. The first term $u(a_t)$ is her gross benefit of using the platform, where $u(a)$ is strictly increasing in $a \in A$ and $u(0) = 0$. The second term $v \cdot [\sigma_0^2 - \sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)]$ is a *privacy cost*, which captures the negative impact of data collection on the consumer. The parameter $v \in \mathbb{R}_{++}$ captures her value of privacy; it is exogenous and commonly known to the consumer and the platform. The consumer discounts future payoffs with discount factor $\delta_C \in [0, 1)$. A special case is a *myopic consumer* (i.e., $\delta_C = 0$), who chooses $a_t \in A$ to maximize $U(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ in each period t. I normalize the payoffs so that if $a_t = 0$ for all t, the platform and the consumer obtain zero payoffs in all periods.

The informational assumptions are summarized as follows. The primitives, σ_0^2, A, $u(\cdot)$, and v, are commonly known. The past activity levels and privacy levels are publicly observable. The consumer's type is unobservable, and the signals are observable only to the platform.

I study two games that differ in the timing of moves. One is the game of *long-run commitment.* In this game, before $t = 1$, the platform commits to a *privacy policy* $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \dots) \in \overline{\mathbb{R}}_+^\infty$, which is publicly observable. Then, in each period $t \in \mathbb{N}$ the consumer chooses a_t, and the platform learns about her type based on a signal s_t. In this game, the platform moves only before $t = 1$.

[6]The equivalent formulation is that the platform observes (a_t, s_t), chooses $b_t \in \mathbb{R}$, and obtains an ex post payoff of $-(X - b_t)^2$, which the platform does not observe. Writing the payoffs in terms of σ_t^2 simplifies exposition. See Acemoglu et al. (2019) for further discussion.

[7]Throughout the paper, "increasing" means "non-decreasing." Similar conventions apply to "decreasing," "higher," "lower," and so on.

The other is the game of *one-period (or short-run) commitment,* in which the platform and the consumer move sequentially in every period: At the beginning of each period t, the platform sets γ_t. After observing γ_t, the consumer chooses a_t. Then, the platform observes the signal, and the game proceeds to period $t+1$. In this case, the platform can commit to a privacy level only for one period. In either game, the solution concept is subgame perfect equilibrium (SPE).[8]

Two remarks are in order. First, under the long-run commitment, the platform can commit to only a history-independent sequence of privacy levels. As a result, we cannot immediately conclude that the platform is better off under the long-run than short-run commitment.[9] Second, I do not model the consumer's participation, but we may interpret that the consumer joins the platform in period $t^* := \min\{t \in \mathbb{N} : a_t > 0\}$. The results continue to hold even if the consumer incurs a small one-time cost to join the platform.

3.1 Discussion of Assumptions

Data generation. In practice, consumer data are generated by their activity on a platform, such as browsing content and responding to posts. The model captures such activity by assuming that the precision of a signal is increasing in the activity level. To focus on the consumer's incentives to protect their privacy, I abstract away from belief manipulation, such as a consumer strategically browsing websites to influence a platform's inference.

Privacy cost function. The privacy cost $v(\sigma_0^2 - \sigma_t^2)$ captures the monetary or nonmonetary reasons why a consumer wants a platform to have less information—e.g., consumers may intrinsically value their privacy, or consider the risk of data breach and discrimination by third parties (Kummer and Schulte, 2019; Lin, 2019; Tang, 2019). Although I do not focus on a particular reason why the loss of privacy harms the consumer, Appendix A microfounds this privacy cost function in the context of third-degree price discrimination with linear demand. In practice, the consumer may

[8]The payoffs $u(a_t) - v \cdot \left[\sigma_0^2 - \sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)\right]$ and $\sigma_0^2 - \sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ incorporate the belief updating of the consumer and the platform. For a myopic consumer, an equilibrium refers to a strategy profile such that (i) the consumer chooses a_t to maximize $U(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ following every history, breaking ties in favor of higher activity levels, and (ii) the platform, anticipating (i), optimally chooses a privacy policy $\boldsymbol{\gamma}$ before $t = 1$ (under long-run commitment), or chooses γ_t at the beginning of each period t (under one-period commitment).

[9]If the platform had "full commitment power" and could commit to any history-contingent rule to set privacy levels, then it would weakly prefer such a regime to any other commitment regimes. However, Section 5 shows that if $|A| = 2$, the platform is indifferent between the full, one-period, and long-run commitment regimes.

prefer a certain level of data collection, which Section 4.1 takes into account. Section 7 studies the consumer who is privately informed of her value v of privacy.

The privacy cost is sunk. The consumer cannot delete past data. Thus she perceives the privacy cost from past data collection as sunk: Even if $a_t = 0$ for all $t \geq T$, the consumer incurs a privacy cost of $-v(\sigma_0^2 - \sigma_T^2)$ in any $t \geq T$. This assumption reflects the difficulty of deleting data, which is referred to as "data persistence" (Tucker, 2018). For instance, suppose a platform collects personal information and shares it with third parties. Then the consumer may face a risk of discrimination or malicious targeting even outside of the platform. In another example, if a consumer inadvertently discloses information to other users, she may incur a psychological cost because other users know the information. Such a cost may persist even when the consumer is not active on the platform. Because the consumer regards the privacy cost as sunk, she chooses activity levels based on the marginal privacy cost rather than the total privacy cost. Appendix L relaxes this assumption and studies extensions in which the consumer regards a part of or all of the privacy cost as non-sunk.

Single consumer. I consider a single consumer to emphasize that the results do not rely on interactions between multiple consumers. However, since the consumer's type is Gaussian, one could incorporate multiple consumers with "data externalities," by assuming that consumers' types are correlated (Acemoglu et al., 2019; Bergemann et al., 2019).

4 Equilibrium Under Long-Run Commitment

To examine how the platform designs its privacy policy, I begin by studying the game of long-run commitment. I first present a result under a stationary privacy policy, then study the equilibrium of the entire game. Given the platform's information in the previous period and (a_t, γ_t), the posterior variance evolves as follows.[10]

$$\sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t) = \frac{1}{\frac{1}{\sigma_{t-1}^2(\boldsymbol{a}_{t-1}, \boldsymbol{\gamma}_{t-1})} + \frac{1}{\frac{1}{a_t} + \gamma_t}}. \tag{1}$$

[10]If $x|\mu \sim N(\mu, \sigma^2)$ and $\mu \sim N(\mu_0, \sigma_0^2)$, then $\mu|x \sim N\left(\frac{\sigma_0^2}{\sigma^2+\sigma_0^2}x + \frac{\sigma^2}{\sigma^2+\sigma_0^2}\mu_0, \left(\frac{1}{\sigma_0^2} + \frac{1}{\sigma^2}\right)^{-1}\right)$.

Thus, the consumer's privacy cost in period t is

$$v\left[\sigma_0^2 - \sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)\right] = v\left[\sigma_0^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2(\boldsymbol{a}_{t-1}, \boldsymbol{\gamma}_{t-1})} + \frac{1}{\frac{1}{a_t} + \gamma_t}}\right].$$

Define the privacy cost function as

$$C(a, \gamma, \sigma^2) := v\left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma^2} + \frac{1}{\frac{1}{a} + \gamma}}\right).$$

The following lemma shows properties of privacy cost C and marginal privacy cost $\frac{\partial C}{\partial a}$.

Lemma 1 (Privacy Cost and Marginal Privacy Cost).

1. *$C(a, \gamma, \sigma^2)$ is decreasing in γ and σ^2, and increasing in a.*
2. *$\frac{\partial C}{\partial a}(a, \gamma, \sigma^2)$ is decreasing in γ and increasing in σ^2.*

Proof. Point 1 follows from equation (1). Point 2 follows from

$$\frac{\partial C}{\partial a} = v \cdot \frac{\frac{\frac{1}{a^2}}{\left(\frac{1}{a} + \gamma\right)^2}}{\left(\frac{1}{\sigma^2} + \frac{1}{\frac{1}{a} + \gamma}\right)^2} = \frac{v}{\left(\frac{1}{\sigma^2}(1 + \gamma a) + a\right)^2}.$$

□

Lemma 1 implies that if the consumer has less privacy (i.e., σ_t^2 is small), she faces a high privacy cost C but a low marginal privacy cost $\frac{\partial C}{\partial a}$. Intuitively, once a platform has collected a lot of information, the marginal privacy cost is low, because the consumer's activity today does not much affect the platform's learning. As a result, data collection harms the consumer, but incentivizes her to increase an activity level in the future. Also, the marginal privacy cost is decreasing in the level of privacy protection, γ. Thus, the platform can encourage the consumer's activity by committing to add a noise to the signal.

We now derive the consumer's problem. We can rewrite the evolution of posterior variances

(1) as that of posterior precisions:

$$\frac{1}{\sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)} = \frac{1}{\sigma_{t-1}^2(\boldsymbol{a}_{t-1}, \boldsymbol{\gamma}_{t-1})} + \frac{1}{\frac{1}{a_t} + \gamma_t} = \cdots = \frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{a_s} + \gamma_s}. \tag{2}$$

As a result, if the platform commits to a privacy policy $(\gamma_t)_{t\in\mathbb{N}}$, the consumer solves the following maximization problem:

$$\max_{(a_t)_{t\in\mathbb{N}}\in A^\infty} \sum_{t=1}^{\infty} \delta_C^{t-1} \left[u(a_t) - v \cdot \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{a_s}+\gamma_s}} \right) \right]. \tag{3}$$

The next result presents the consumer's response to a stationary privacy policy. Although the platform's equilibrium policy may not be stationary, this non-equilibrium analysis clarifies the intuition behind the consumer's dynamic incentive (see Appendix C for the proof).

Proposition 1. *Suppose the platform commits to a stationary privacy policy, i.e., $\gamma_t = \gamma$ for all $t \in \mathbb{N}$. Let $(a_t^*)_{t\in\mathbb{N}}$ denote the equilibrium activity levels of this subgame. There is a cutoff value $v^*(\gamma) \in \mathbb{R}_+$ with the following properties.*

1. *If $v < v^*(\gamma)$, then a_t^* increases in t, $\lim_{t\to\infty} a_t^* = a_{max}$, and $\lim_{t\to\infty} \sigma_t^2 = 0$. The consumer's continuation value decreases over time.*

2. *If $v > v^*(\gamma)$, then $a_t^* = 0$ and $\sigma_t^2 = \sigma_0^2$ for all $t \in \mathbb{N}$.*

3. *The cutoff $v^*(\gamma)$ is increasing in γ, and $\lim_{\gamma\to\infty} v^*(\gamma) = \infty$.*

The intuition is as follows. If the value of privacy is low, the consumer prefers a positive activity level $a_1^* > 0$ in $t = 1$. The consumer activity generates information, which reduces her payoff and marginal cost of using the platform. As a result, she chooses $a_2^* \geq a_1^*$ in $t = 2$. Repeating this argument, we can conclude that a_t^* increases over time. The platform can then observe the signals to perfectly learn the consumer's type as $t \to \infty$. Perfect learning in $t \to \infty$ implies that the marginal privacy cost goes to zero, and thus $a_t^* \to a_{max}$. To sum up, if v is below the cutoff, the consumer eventually loses her privacy, but acts as if there is no privacy cost (Point 1). In contrast, the consumer with a high v does not use the platform (Point 2). Finally, $v^*(\gamma)$ is increasing in γ because a higher privacy level reduces the cost of using the platform (Point 3).

Proposition 1 implies a perverse effect of privacy regulation: Suppose that a regulator, who cares about consumer privacy, mandates a stricter privacy policy—i.e., $\gamma_t = \gamma$ becomes $\gamma_t = \gamma' > \gamma$ for all $t \in \mathbb{N}$. The result implies that this regulation increases the cutoff from $v^*(\gamma)$ to $v^*(\gamma')$, and expands the range of v's under which the consumer loses privacy (Point 1). To see the welfare implication, suppose $v > \frac{u(a_{max})}{\sigma_0^2}$ holds. For a small γ, the consumer may choose $a_t^* = 0$ and obtain a payoff of zero in all periods. If the regulator enforces a large γ', then the consumer chooses $a_1^* > 0$. However, $a_1^* > 0$ implies $(a_t^*, \sigma_t^2) \to (a_{max}, 0)$, and thus the consumer's per-period payoff converges to $u(a_{max}) - v\sigma_0^2 < 0$. Thus, the regulation may increase the consumer's payoffs in the short run but decrease them in the long run. If the regulator cares about long-run consumer welfare, it may consider a higher γ to be detrimental.[11]

The next result shows the equilibrium dynamics of the entire game, in which the platform can commit to any (potentially nonstationary) privacy policy. In equilibrium, the platform anticipates that the consumer solves (3) given any privacy policy (see Appendix D for the proof; Appendix B proves the existence of an equilibrium).

Theorem 1. *In any equilibrium under the long-run commitment, the following holds.*

1. *The consumer eventually loses her privacy and chooses the highest activity level:* $\lim_{t\to\infty} \sigma_t^2 = 0$ *and* $\lim_{t\to\infty} a_t^* = a_{max}$.
2. *For any $T \in \mathbb{N}$, there is a $\underline{v} \in \mathbb{R}$ such that for any $v \geq \underline{v}$, we have $\gamma_t^* > 0$ for all $t \leq T$.*
3. *If the consumer is myopic, there is a $T' \in \mathbb{N}$ such that for all $t \geq T'$, $\gamma_t^* = 0$.*

Because of the decreasing marginal cost of losing privacy, a farsighted consumer eventually gives up all of her privacy, even when she anticipates future choices made by the platform. The privacy loss occurs even if δ_C is close to 1 and the service utility $u(\cdot)$ is small relative to the privacy cost. Point 2 implies that if the consumer highly values her privacy, the platform commits to high privacy protection in early periods to induce the long-run privacy loss.

The intuition is as follows. In early periods, the platform knows little about the consumer, so the consumer's activity has a large impact on what the platform can learn about her type. Thus the

[11]The caveat "if the regulator cares about the long-run consumer welfare" is important, because a higher γ increases the consumer's ex ante sum of discounted payoffs calculated based on δ_C. A higher privacy level is undesirable for the regulator only if the regulator's discount factor is different from the consumer's.

consumer faces a high marginal privacy cost, which discourages her from raising the activity level. The platform then commits to a high level of privacy protection to encourage consumer activity. As a result, in early periods the platform slowly learns her type.

The above argument does not exclude the possibility that data collection stops in the middle—e.g., the platform might need to provide a growingly high privacy protection over time, so that the precision of the signal eventually goes to zero. However, the result implies that data collection never stops because of the decreasing marginal cost of privacy loss: As time goes by, the platform accurately knows the consumer's type, which renders it cheaper for the consumer to generate additional information. Therefore, the platform is able to collect full information over time.

The consumer can be forward-looking. Thus, we may think that the platform could benefit from pre-committing to high privacy levels for future periods, if it encourages the consumer to generate more information in early periods. This intuition is inaccurate: The consumer's objective in (3) is supermodular in today's activity a_t and the precision of future signals $(\frac{1}{(a_{t+s})^{-1}+\gamma_{t+s}})_{s\in\mathbb{N}}$, so the consumer chooses a higher activity level when she anticipates to lose her privacy in the future. The result also implies that the platform collects full information for any discount factors: For example, even a myopic platform that faces a patient consumer adopts a privacy policy that causes the long-run privacy loss.

Figure 1 depicts the equilibrium dynamics for a myopic consumer in a numerical example.[12] Figure 1(a) shows the platform offers a decreasing privacy level, hitting zero in $t = 5$. Figure 1(b) shows that the equilibrium activity level first decreases but eventually approaches $a_{max} = 2$. The non-monotonicity of a_t^* contrasts with the case of a stationary privacy policy.[13]

A natural question is to what extent Theorem 1 depends on assumptions on preferences. In particular, the consumer and the platform hold the opposite preferences over data collection, and the platform earns revenue only from data. I now relax these assumptions and extend the main result.

[12] I assume $A = \{0, 0.01, 0.02, \ldots, 2\}$ and use Claim 1 to compute an equilibrium.

[13] I have not managed to prove the non-monotonicity of $(a_t^*)_{t\in\mathbb{N}}$. A numerical exercise suggests that the non-monotonicity occurs for a wide range of parameters (v, σ_0^2) with a myopic consumer, when the equilibrium privacy level is strictly decreasing in early periods.

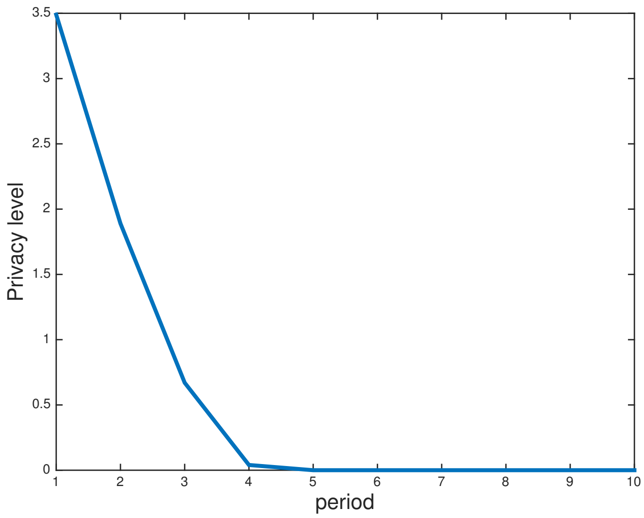

Figure 1(a): Privacy level γ_t

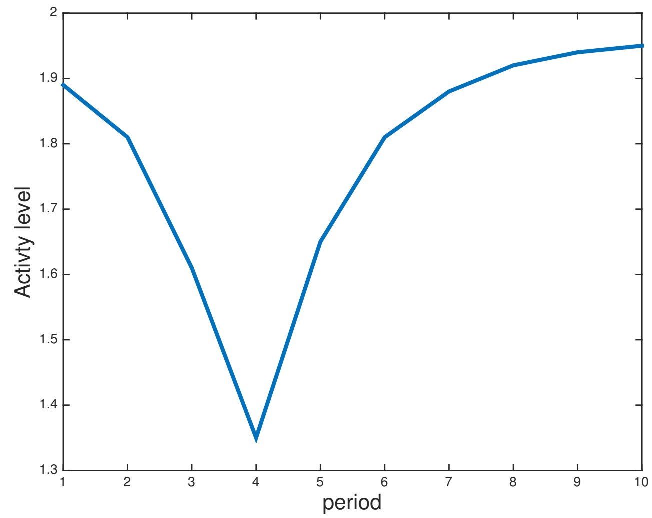

Figure 1(b): Activity level a_t

Figure 1: Equilibrium under $u(a) = 2a - \frac{1}{2}a^2$, $v = 10$, and $\sigma_0^2 = 1$.

4.1 The Implication on Business Models

So far, I have assumed that the consumer dislikes data collection, and the platform's revenue depends only on information. While these assumptions help us highlight the intuition, the main insight holds in a more general setting. To see this, suppose now that the platform's per-period payoff is $\Pi(a_t, \sigma_t^2)$, which is strictly increasing in a_t and decreasing in σ_t^2. For example, an advertising platform benefits from a high activity a_t and more data (i.e., low σ_t^2) because the consumer will then see many highly targeted ads. We impose no restrictions on the relative importance of activity and data for $\Pi(\cdot,\cdot)$. Also, the consumer's per-period payoff is now $u(a_t) - C(\sigma_t^2)$. The original setting is $C(\sigma_t^2) = v(\sigma_0^2 - \sigma_t^2)$, but $C(\cdot)$ is now any differentiable function such that $\sup_{x \in [0,\sigma_0^2]} |C'(x)| < \infty$ and $C(\sigma_0^2) = 0$. The cost function can be non-monotone. For example, $C(\cdot)$ is first decreasing and then increasing, i.e., the consumer prefers some level of data collection.

The following result extends Theorem 1: The platform can always induce the consumer to give up her privacy and choose the highest activity level in the long run, and a sufficiently patient platform chooses to do so (see Appendix E for the proof).

Proposition 2. *In the above setting, the following holds:*

1. *There is a privacy policy under which the consumer loses her privacy and chooses the highest*

activity level in the long run: $\lim_{t\to\infty} \sigma_t^2 = 0$ *and* $\lim_{t\to\infty} a_t^* = a_{max}$.

2. *Given the platform's discount factor* δ_P, *let* $\sigma_\infty^2(\delta_P) := \lim_{t\to\infty} \sigma_t^2(\delta_P)$ *denote the long-run posterior variance in an arbitrarily chosen equilibrium. Then,* $\lim_{\delta_P\to 1} \sigma_\infty^2(\delta_P) = 0$. *Similarly, the long-run activity level converges to* a_{max} *as* $\delta_P \to 1$, *provided the limits exist—i.e.,* $\lim_{\delta_P\to 1} \lim_{t\to\infty} a_t = a_{max}$ *for arbitrarily chosen equilibria.*

Even if a platform's profit places a small weight on information, it may collect as much information as exclusively data-driven firms, because high user activity is consistent with no privacy. Intuitively, so long as $C(\sigma_t^2)$ has a bounded derivative, the marginal cost of using the platform approaches zero when the consumer loses all of her privacy.[14] As a result, the platform can induce the highest activity and the lowest privacy at the same time, which is the best outcome in the long-run regardless of the relative importance of data and activity for the firm's revenue.

4.2 Implications of Theorem 1 and Proposition 2

First, Theorem 1 potentially explains the privacy paradox: Consumers seem to casually share their data with online platforms, despite their concerns about data collection.[15] We may view this puzzle as the long-run equilibrium outcome of this model, in which the consumer faces a high privacy cost and negligible marginal cost. Such an outcome can arise even if firms adopt business models that do not much rely on data (Proposition 2). The result also points to the difficulty of applying the revealed preference argument to static privacy choices, because the consumer's decision may depend on the stock of information they have already revealed.

Second, the result connects consumer privacy problem with rational addiction (Becker and Murphy, 1988). The connection stems from that a high activity level today decreases the consumer's future utility, but increases her future marginal utility of using the platform. In contrast to models of rational addiction, the current model has a platform that can choose its privacy policy to influence the degree of addiction. As a result, even if consumers are patient and highly value their privacy, they become "addicted" to the platform.

[14] While in a different context, the intuition is similar to the idea that if a firm uses data for forecasts and the gain to a perfect forecast is finite, the returns to data must diminish at some point (Farboodi and Veldkamp, 2020).

[15] Acquisti et al. (2016) conduct an insightful review of research on the economics of privacy, including the privacy paradox. Recent empirical work includes, for example, Athey et al. (2017).

Finally, at an anecdotal level, the equilibrium strategy of the platform, which offers high privacy levels for early periods but not necessarily for later periods, seems consistent with how the data collection strategies of online platforms have evolved. In 2004, Facebook's privacy policy stated that it would not use (first-party) cookies to collect consumer information. In 2020, the privacy policy states that it uses cookies to track consumers on and possibly off the website.[16] Srinivasan (2019) describes how Facebook's policy has changed from the one that preserves consumer privacy to "broad-scale commercial surveillance." Also, Fainmesser et al. (2019) describe how online platforms' business models have changed from the initial phase, in which they expand a user base, to the mature phase, in which they monetize the information collected. The equilibrium dynamics in Theorem 1 rationalize the pattern described, and Proposition 2 implies that such dynamics could occur even if a platform earns revenue mainly from user activity.

4.3 Myopic Consumer

If the consumer is myopic, we can characterize the equilibrium. Let $a^*(\gamma, \sigma^2) \in A$ denote the best response of a myopic consumer, given a privacy level γ in the current period and the posterior variance σ^2 from the previous period:

$$a^*(\gamma, \sigma^2) := \max \left\{ \arg\max_{a \in A} u(a) - v\left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma^2} + \frac{1}{\frac{1}{a} + \gamma}}\right) \right\}. \tag{4}$$

The following result characterizes the equilibrium (see Appendix F for the proof; it also provides the recursive characterization of the equilibrium privacy policy).

Claim 1. *Suppose the consumer is myopic. Consider the greedy policy of the platform—i.e., given the consumer's best response* (4)*, it chooses privacy level γ_t to myopically maximize the precision of the signal in each period t. The sequence $(\gamma_t^*)_{t \in \mathbb{N}}$ obtained in this way is the platform's equilibrium strategy. Also, the same outcome arises under one-period commitment.*

Although the platform is patient, the optimal policy is greedy. The result comes from the

[16] In 2020, Facebook's privacy policy states that "we use cookies if you have a Facebook account, use the Facebook Products, including our website and apps, or visit other websites and apps that use the Facebook Products (including the Like button or other Facebook Technologies)." https://www.facebook.com/policies/cookies

decreasing marginal privacy cost (and not necessarily from the consumer myopia, as we see in the next section). To see this, consider the platform's problem of choosing γ_t. The platform faces a trade-off: Higher privacy protection reduces the consumer's marginal cost and encourages her activity. However, privacy protection implies that the platform learns less about the consumer's type from the signal. Balancing this trade-off, the platform chooses γ_t to maximize the sum of period-t profit and its continuation value. However, the platform does not face a dynamic trade-off: If it collects more information today, the consumer will face lower marginal privacy costs and choose higher activity levels in the future. As a result, both the platform's period-t profit and its continuation value are increasing in the precision of the signal in period t. The platform can then maximize the sum of discounted profits by myopically maximizing the precision of the signal in each period t. Since the platform's optimal policy is greedy, it is time consistent. The next section shows that the equilibrium privacy policy can be greedy even if the consumer is patient.

5 One-Period Commitment

I now consider one-period commitment to study how the lack of commitment affects the platform's ability to collect data. One-period commitment could be realistic in some contexts. For example, a platform may be sanctioned for the outright violation of its privacy policy, but it may still revise its policy over time. Below, I first show that the equilibrium is unique if the prior uncertainty of the consumer's type (i.e., σ_0^2) is small. I then consider a general σ_0^2.

5.1 Unique Equilibrium Under a Small σ_0^2

Despite the dynamic nature of the game, the equilibrium is unique if the initial uncertainty about the consumer' type is small.

Proposition 3. *There is a $B > 0$ such that if $\sigma_0^2 \leq B$, then any equilibrium involves $(\gamma_t, a_t) = (0, a_{max})$ for all $t \in \mathbb{N}$, and $\lim_{t\to\infty} \sigma_t^2 = 0$.*

Proof. Let a' denote the second highest activity level in A. Take any B that satisfies $u(a_{max}) - u(a') - \frac{v}{1-\delta}B > 0$. In any period, if the consumer chooses $a_t = a_{max}$ instead of $a_t \in A \setminus \{a_{max}\}$, her gross payoff increases by at least $u(a_{max}) - u(a') > 0$, and her privacy cost increases by at most $\frac{v}{1-\delta}B$. For $\sigma_0^2 \leq B$, the consumer chooses a_{max}, and the platform optimally chooses $\gamma_t = 0$. $\square$

Intuitively, if σ_0^2 is small, the marginal privacy cost is so small that the consumer prefers a_{max} regardless of the current and future privacy protection. Thus, the long-run privacy could arise in equilibrium only if the platform has not yet learned much about the consumer.

5.2 Binary Activity Level

I now consider a general σ_0^2. To facilitate the analysis, I impose the following assumption and definitions.

Assumption 1. The consumer has a binary activity level: $A = \{0, a_{max}\}$.

Definition 1. An equilibrium is *platform-best* if it maximizes the platform's ex ante sum of discounted payoffs across all subgame perfect equilibria. We analogously define "*platform-worst*," "*consumer-best*," and "*consumer-worst*."

Definition 2. A *Markov perfect equilibrium (MPE)* is an equilibrium in which after any history, the platform's choice γ_t depends only on σ_{t-1}^2, and the consumer's choice a_t depends only on $(\sigma_{t-1}^2, \gamma_t)$.

The following result presents a consumer-worst equilibrium, which is also platform-best under a common discount factor. Recall that δ_C and δ_P denote the discount factors of the consumer and the platform, respectively (see Appendix G for the proof).

Theorem 2. *Under Assumption 1, there is a consumer-worst Markov perfect equilibrium. This equilibrium is independent of δ_P, and has the following properties:*

1. *If $\delta_C = \delta_P$, the equilibrium is platform-best. The privacy levels $(\gamma_t^*)_{t \in \mathbb{N}}$ coincide with an equilibrium policy under long-run commitment.*

2. *The privacy level γ_t^* is decreasing in t and hits zero in a finite time. Also, the consumer loses her privacy in the long run: $\lim_{t \to \infty} \sigma_t^2 \to 0$.*

3. *The platform's strategy is greedy: Given the consumer's strategy, after any history, the platform sets a privacy level γ_t to maximize the informativeness of the signal in period t.*

Point 1 implies that one-period commitment power can be enough for the platform to attain its best outcome. Indeed, we can strengthen Point 1 as follows (see Appendix G for the proof). Suppose the platform has the strongest commitment power—i.e., it can commit to any rule that determines privacy levels as a function of past and future outcomes. Even so, the platform's payoff cannot exceed the payoff from the equilibrium in Theorem 2. Intuitively, this consumer-worst equilibrium attains the highest discounted privacy cost across all outcomes such that the consumer's ex ante payoff exceeds a certain lower bound. We can show that this lower bound applies even if the platform has a stronger commitment power. Also, under a common discount factor, the consumer's discounted privacy cost is proportional to the platform's discounted profit. Therefore, the platform cannot increase its profit even if it has the strongest commitment power.

Points 2 extend the intuition in Theorem 1: The platform initially chooses high privacy levels to incentivize the consumer to generate information. As the platform collects more information, her incentive to protect privacy declines; correspondingly, the platform sets a decreasing privacy level, which hits zero in a finite period. Lemma 1 alone does not imply that the consumer faces a lower cost of choosing a_{max} when she has less privacy, because her continuation value is endogenous. However, in this equilibrium, the consumer's continuation value $V(1/\sigma_t^2)$, as a function of the amount of information collected, is decreasing and convex in $1/\sigma_t^2$. As a result, the consumer's Markov decision problem exhibits a declining marginal loss of generating information.

Point 3 states that the platform adopts a greedy policy, given the consumer equilibrium strategy. The proof also reveals that any deviation by the platform reduces the precision of the signal in any future period. Thus we obtain the same consumer-worst equilibrium as long as the platform prefers to have more information—e.g., the platform's objective does not need to be additively separable over time (see Section 7.4).

Theorem 2 indicates that the lack of long-run commitment may not prevent the platform from collecting consumer data. At the same time, the result does not imply the uniqueness of the equilibrium when σ_0^2 is not small. Indeed, the platform with only one-period commitment power may fail to collect any information (see Appendix H for the proof).

Theorem 3. *Suppose Assumption 1 and $\delta_C \geq \frac{1}{2}$ hold. There is a (v-dependent) $\sigma^2 < \infty$ such that if $\sigma_0^2 \geq \sigma^2$, there is a consumer-best and platform-worst Markov perfect equilibrium, in which the platform sets $\gamma_t = \infty$ and the consumer chooses $a_t = a_{max}$ in all periods.*

In this equilibrium, the platform offers full privacy, because whenever it attempts to collect information by setting $\gamma_t < \infty$, the consumer chooses $a_t = 0$. The consumer prefers $a = 0$ following the platform's deviation, because the initial privacy loss, no matter how small, will lead to the complete privacy loss and impose her a high cost in the future. Indeed, after any off-path event in which the platform collects some information (i.e., $\sigma_t^2 < \sigma_0^2$), the consumer-worst equilibrium in Theorem 2 is played. A grim trigger strategy—i.e., the platform's deviation induces $a_t = 0$ forever— does not work, because the platform can set a large finite γ_t to render such a punishment suboptimal for the consumer (i.e., Lemma 11 in Appendix I). We may view Theorem 3 as the plaform's Coasian commitment problem: The platform in period t competes with its future self, which offers the best privacy protection in any period $s \geq t + 1$.

Implication on Introducing New Digital Services. Consider a new smart speaker through which a firm collects signals about consumer characteristics. Also, consider company A that operates other digital services and already holds consumer data (i.e., a low σ_0^2), and a new company B with little data (i.e., a high σ_0^2). We may think that company B has a stronger incentive to introduce the smart speaker, because it faces a higher marginal value of information. However, Proposition 3 and Theorem 3 suggest that only company A may successfully introduce the new product and collect consumer data. Indeed, under the platform-worst outcome, company B does not value the smart speaker, because consumers will refuse to use it. Intuitively, consumers will find it less costly to have their information collected by company A, which already knows a lot about consumers.

Remark 1 (Welfare Implications). Under a common discount factor δ, the total surplus is

$$\sum_{t=1}^{\infty} \delta^{t-1} \left[u(a_t) + (1-v)(\sigma_0^2 - \sigma_t^2) \right].$$

If $v > 1$, the efficient outcome is $(\gamma_t, a_t) = (\infty, a_{max})$ for all $t \in \mathbb{N}$. Thus, the consumer-best equilibrium in Theorem 3 is efficient. If $v < 1$, the efficient outcome is $(\gamma_t, a_t) = (0, a_{max})$ for all $t \in \mathbb{N}$. This outcome may not arise in any equilibrium, because it may give the consumer a negative payoff. However, the inefficiency will disappear in the platform-best equilibrium, in that $(\gamma_t, a_t) = (0, a_{max})$ is played after some finite period. Finally, if the platform has long-run commitment power, it implements the platform-best outcome in any equilibrium. In this case, if $v > 1$, any equilibrium leads to an inefficiently high level of data collection.

5.3 General Set of Activity Levels

In this subsection, the consumer can choose activity levels from any finite set A. The general characterization of the set of equilibria is beyond the scope of the paper. However, I can construct an equilibrium that has similar properties to Theorem 2, under the following technical assumption:

Assumption 2. The platform chooses a privacy level from a finite set $\Gamma \subset \overline{\mathbb{R}}_+$ that contains some finite $\bar{\gamma} > \frac{v(\sigma_0^2)^2}{(1-\delta_C)u(a_{max})} - \frac{1}{\sigma_0^2} - \frac{1}{a_{max}}$.

This assumption allows $\min \Gamma = 0$ and $\max \Gamma = \infty$. Proposition 3 implies that any subgame that starts from $\sigma_t^2 \leq B$ has a unique equilibrium, in which $(\gamma_s, a_s) = (0, a_{max})$ for all $s \geq t + 1$. We can then use the backward induction with respect to σ_t^2 to construct an MPE, starting from any σ_0^2 (see Appendix I for the proof).

Proposition 4. *Under Assumption 2, for any σ_0^2, there is a Markov perfect equilibrium in which (i) there is a $T \in \mathbb{N}$ such that for all $t \geq T$, $(\gamma_t^*, a_t^*) = (0, a_{max})$, and thus (ii) $\lim_{t\to\infty} \sigma_t^2 = 0$.*

6 Platform Competition with a Myopic Consumer

This section shows that the decreasing marginal privacy cost renders competition ineffective in increasing consumer privacy. Specifically, I study a model with an incumbent (I), an entrant (E), and a myopic consumer. Platform I is in the market from the beginning of $t = 1$. In period $t^* \geq 2$, E enters the market. The entry period t^* is exogenous, deterministic, and commonly known.[17] Let γ_t^k denote the privacy level of platform k in period t.

Before the entry ($t < t^*$), the consumer chooses an activity level $a_t^I \in A$ for I. After the entry ($t \geq t^*$), the consumer chooses $(a_t^I, a_t^E) \in A^2$, where a_t^E is the activity level for E. The consumer

[17] I obtain qualitatively the same result when the entry is endogenous and costly for the entrant.

can choose $(a_t^I, a_t^E) \in A^2$ if and only if $\min(a_t^I, a_t^E) = 0$. This restriction captures single-homing, which is natural if platforms offer similar services.

Since the result does not depend on the commitment regime, I examine competition with one-period commitment (Appendix J considers both commitment regimes). At the beginning of each period, a platform chooses a privacy level, after which the consumer chooses an activity level. In particular, I and E simultaneously set privacy levels γ_t^I and γ_t^E in each period $t \geq t^*$, without making any commitment to future privacy levels.

As before, platform $k \in \{I, E\}$ receives a signal $s_t^k = X + \varepsilon_t^k$ with $\varepsilon_t^k \sim \mathcal{N}\left(0, \frac{1}{a_t^k} + \gamma_t^k\right)$ in period t. Each platform k privately observes s_t^k, and all of the noise terms $(\varepsilon_t^k)_{k,t}$ are independent across $(k,t) \in \{I, E\} \times \mathbb{N}$. The payoff of platform $k \in \{I, E\}$ in period t is $\sigma_0^2 - \sigma_{t,k}^2$, where $\sigma_{t,k}^2$ is the posterior variance of the consumer's type, given activity levels and privacy levels. The consumer's payoff in period t is

$$u(a_t^I) - v\left(\sigma_0^2 - \sigma_{t,I}^2\right) + \mathbf{1}_{\{t \geq t^*\}} \cdot \left[u(a_t^E) - v\left(\sigma_0^2 - \sigma_{t,E}^2\right)\right], \tag{5}$$

where $\mathbf{1}_{\{t \geq t^*\}}$ is the indicator function that equals 1 or 0 if $t \geq t^*$ or $t < t^*$, respectively. Payoff (5) implies that even if the consumer switches to (say) E and never uses I from some period on, she continues to incur a privacy cost based on the information collected by I in the past (Appendix L relaxes this assumption).

To obtain a non-trivial result, I impose an upper bound on the feasible privacy levels. The bound might capture the minimum amount of data a platform needs to collect in order to maintain services, or the maximum privacy protection a platform can credibly enforce. Recall that $a^*\left(\bar{\gamma}, \sigma_0^2\right)$ is the optimal activity level of a myopic consumer, defined in (4).

Assumption 3. There is a $\bar{\gamma} \in \mathbb{R}_+$ satisfying $a^*\left(\bar{\gamma}, \sigma_0^2\right) > 0$ such that platforms I and E can choose a privacy level of at most $\bar{\gamma}$.

I present an equilibrium that involves a monopolistic outcome (see Appendix J for the proof).

Proposition 5. *Under Assumption 3, the following holds.*

1. *There is an equilibrium in which $a_t^E = 0$ for all $t \in \mathbb{N}$, $\lim_{t\to\infty} a_t^I = a_{max}$, $\lim_{t\to\infty} \sigma_{I,t}^2 = 0$, and $\lim_{t\to\infty} \gamma_t^I = 0$.*

2. *There is $\underline{t} \geq 2$ such that if the entry time t^* is greater than $\underline{t}$, any equilibrium outcome for the consumer and the incumbent, $(a_t^I, \gamma_t^I)_{t\in\mathbb{N}}$, coincides with the monopoly outcome in Claim 1.*

The intuition is as follows. Suppose that upon entry, the entrant sets the highest privacy level $\bar{\gamma}$. Since the privacy cost from collected data is sunk, the consumer decides which platform to use based on her marginal (or, more precisely, incremental) costs. Now, the consumer faces a lower marginal cost of using the incumbent, which has already collected some data. Thus if the incumbent also chooses $\bar{\gamma}$, the consumer prefers to use it. However, the equilibrium choice of the incumbent may not be $\bar{\gamma}$: The incumbent chooses a privacy level that maximizes the precision of the signal, subject to the constraint that the consumer does not switch to the entrant. As time goes by, the constraint is relaxed, because the consumer's marginal cost for the incumbent goes to zero. As a result, the incumbent offers a vanishing privacy level over time. Finally, the threat of future entry does not affect the incumbent's strategy: Before the entry, it chooses the same privacy levels as a monopoly, because collecting more information renders consumer switching less likely.

Under the decreasing marginal privacy cost, switching could be less likely when consumers have low privacy and receive low payoffs from the incumbent. This observation may contrast with the existing idea of "data as an entry barrier," in which dominant platforms use data to improve their services and attract users.[18]

As an example, consider search engines: The incumbent is Google, and the entrant is a privacy-preserving alternative of Google, such as DuckDuckGo. If consumers have no privacy on Google, they face negligible marginal privacy costs of using it. Then even if DuckDuckGo is as good a search engine as Google and offers better privacy protection, it may not be able to poach consumers. The result also implies that the effect of such competition may depend on whether consumers regard data collected by Google as sunk.

7 Extensions

This section examines several extensions. For simplicity, I focus on a monopoly with long-run commitment and a myopic consumer. Appendix K contains omitted proofs.

[18] For example, Furman et al. (2019) state that "data can act as a barrier to entry in digital markets. A data-rich incumbent is able to cement its position by *improving its service and making it more targeted for users*, as well as making more money by better targeting its advertising." (italics added)

7.1 Erasing Past Information

This extension studies the incentive of the consumer or the platform to erase past information.

7.1.1 The Right to be Forgotten

First, I consider the right to be forgotten, whereby the consumer can request a platform to delete past information. At the beginning of each period, the consumer chooses whether to erase past information, then chooses an activity level. If she erases information in period t, the posterior variance at the beginning of t becomes the prior variance σ_0^2. At the end of the period, the consumer still incurs a privacy cost based on information generated in that period. For example, if the consumer erases information in period t, her payoff is $u(a_t) - v\left[\sigma_0^2 - \sigma_1^2(a_t, \gamma_t)\right]$, where $\sigma_1^2(a_t, \gamma_t)$ is the posterior variance given one signal based on (a_t, γ_t). Thus, the privacy cost is only based on the signal of period t. In contrast, if the consumer has never erased information, her payoff in period t is $u(a_t) - v\left[\sigma_0^2 - \sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)\right]$.

Claim 2. *If the consumer can costlessly erase past information, there is an equilibrium in which the platform commits to a stationary privacy policy $\gamma_t \equiv \gamma_1^*$, where γ_1^* is defined in (21). In this equilibrium, the consumer erases information in every period.*

Once the consumer erases information, she incurs a high marginal privacy cost. Then the platform offers a period-1 privacy level in any period. As a result, the equilibrium involves neither privacy loss nor vanishing privacy protection.

When there are multiple platforms, erasing past information promotes competition and further benefits consumers: Once the consumer deletes information, the incumbent and the entrant become identical in terms of the amounts of data they hold. As a result, they offer the highest privacy protection to attract consumers.

7.1.2 Data Retention Policies

Does the platform have an incentive to voluntarily erase past data? This question relates to data retention policies, which have recently drawn the attention of economists and legal scholars (Chiou and Tucker, 2017). Here, at the beginning of this game, the platform commits to a privacy policy $(\gamma_t)_{t\in\mathbb{N}}$ and the set $\mathcal{T} \subset \mathbb{N}$ of periods to delete information. The platform erases past information at the beginning of each period $t \in \mathcal{T}$. The platform's erasing information affects the posterior variance and payoffs in the same way as the consumer erasing information (see the previous subsection). As a result, erasing information increases σ_t^2 to σ_0^2, and decreases the myopic consumer's activity level. Thus we obtain the following result.

Claim 3. *In any equilibrium, the platform never erases information: $\mathcal{T} = \emptyset$.*

The result implies that the platform has different incentives to offer ex ante and ex post privacy protections: It may voluntarily offer high privacy levels in early periods, because committing to collect less information encourages the consumer's activity. However, the platform has no incentive to delete past information, because it increases the consumer's marginal cost and decreases her activity level.

7.2 Consumers with Heterogeneous Values of Privacy

The main insight does not depend on whether the platform knows v at the outset. To see this, I extend the model as follows: There is a unit mass of consumers. Each consumer $i \in [0,1]$ has v_i, which is distributed according to a distribution with a finite support $V \subset \mathbb{R}_+$. Let $\alpha_v \in [0,1]$ denote the mass of consumers with $v \in V$. Each consumer i is privately informed of v_i, and the platform knows V and $(\alpha_v)_{v \in V}$.

The game is a natural extension of the baseline model. Before $t = 1$, the monopoly platform chooses a privacy policy $(\gamma_t)_{t \in \mathbb{N}}$, which is common across all consumers. Then each consumer i myopically chooses activity levels $(a_t(i))_{t \in \mathbb{N}}$. The types and signals are independent across consumers.

For each $i \in [0,1]$, let $\sigma_t^2(i)$ denote the posterior variance for consumer i at the end of period t. Then i's payoff is $u(a_t(i)) - v_i[\sigma_0^2 - \sigma_t^2(i)]$, and the platform's payoff is $\int_{i \in [0,1]} \sigma_0^2 - \sigma_t^2(i) di$. In equilibrium, consumers who have the same v choose the same sequence of activity levels. As a result, we can write the platform's profit as $\sum_{v \in V} \alpha_v \left[\sigma_0^2 - \sigma_t^2(v)\right]$, where $\sigma_t^2(v)$ is the posterior variance of consumers with v.

The platform faces a new trade-off: A high privacy level encourages consumers with high v to choose positive activity levels. However, the platform obtains less information from consumers with low v, who would choose high activity levels without privacy protection.[19] This static trade-off also creates a dynamic trade-off: For example, a more myopic platform may set a low privacy level to quickly collect data from consumers with low v, whereas a patient platform may set high privacy levels to collect information from all consumers over time.

However, there is no trade-off for the platform in the long run—i.e., all consumers eventually lose privacy and choose the highest activity levels.

Proposition 6. *Let $(a_t^*(v), \sigma_t^2(v), \gamma_t^*)_{t\in\mathbb{N}, v\in V}$ denote the outcome of any equilibrium. Then,*

$$\forall v \in V, \lim_{t\to\infty}(a_t^*(v), \sigma_t^2(v)) = (a_{max}, 0) \text{ and } \lim_{t\to\infty}\gamma_t^* = 0. \tag{6}$$

To see the intuition, suppose that v is either $L = 0$ or $H > 0$, and the platform sets $\gamma_t = 0$ in early periods to collect information only from L-consumers. During this period, only $\sigma_t^2(L)$ decreases over time. However, once $\sigma_t^2(L)$ gets close to zero, the platform finds it more profitable to increase a privacy level to encourage H-consumers to use the platform. Thus, the platform eventually obtains information from all consumers.

7.3 Time-Varying Type of the Consumer

The baseline model assumes that the consumer's type X is constant over time. However, we can conceptually extend the model so that her type is some stochastic process $(X_t)_{t\in\mathbb{N}}$. One possibility, which I adopt for a numerical analysis, is as follows: $X_{t+1} = \phi X_t + \zeta_t$ with $\phi \in [0,1]$, $X_0 \sim \mathcal{N}(0, \sigma_0^2)$, and $\zeta_t \overset{iid}{\sim} \mathcal{N}(0, (1-\phi^2)\sigma_0^2)$. The variance of each ζ_t is normalized so that $Var(X_t) = \sigma_0^2$ for all $t \in \mathbb{N}$. As in the baseline model, given an activity level a_t and a privacy level γ_t in period t, the platform observes a signal $s_t = X_t + \varepsilon_t$ with $\varepsilon_t \sim \mathcal{N}\left(0, \frac{1}{a_t} + \gamma_t\right)$. The posterior variance evolves according to $\sigma_t^2 = \frac{1}{\frac{1}{\phi^2\sigma_{t-1}^2 + (1-\phi^2)\sigma_0^2} + \frac{1}{\frac{1}{a_t}+\gamma_t}}$.

A natural question is how the equilibrium converges to the steady state. However, such an analysis is difficult, partly because the consumer's objective is neither concave nor convex in a_t. Thus I present a numerical analysis to examine the convergence to the steady state, and how the equilibrium responds to the persistence of the consumer's type. Intuitively, if the type is less persistent (i.e., ϕ is small), a larger amount of new information arrives in each period. Then, she

[19] A similar trade-off arises in Lefouili and Toh (2019).

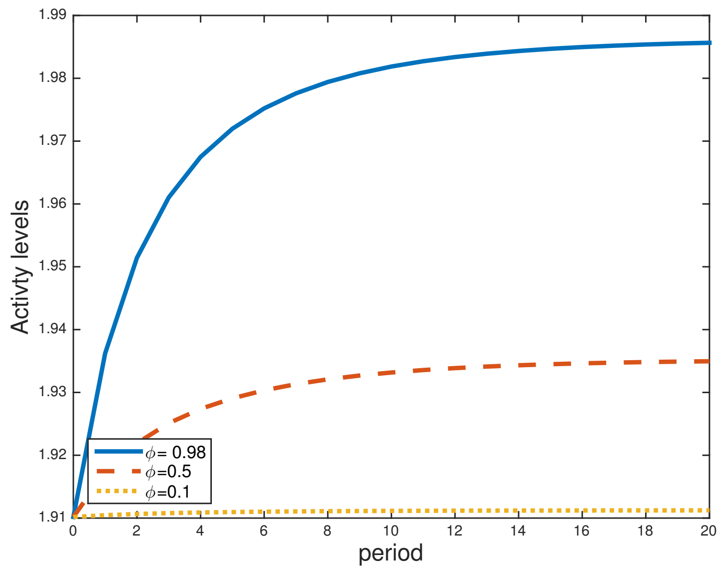

Figure 2: Activity levels under a stationary policy, given $u(a) = 2a - \frac{1}{2}a^2$, $v = 10$, $\sigma_0^2 = 1$, $\phi \in \{0.1, 0.5, 0.98\}$, and $\gamma_t \equiv 4$.

faces a higher marginal cost and chooses a lower activity level. Figure 2 confirms this intuition: Given a stationary privacy policy, the optimal activity levels converge to the steady states, which seem to increase in ϕ.

Figure 3 presents equilibria, taking the platform's optimization into account. First, the numerical analysis suggests that the main insight of this paper is not specific to the baseline specification $\phi = 1$. Namely, the platform offers a relatively high privacy level in early periods, but later reduces it (Figure 3(a)). While Figure 3 fixes v, a similar numerical exercise shows that the platform is able to obtain a nontrivial amount of information in the steady state even if v is larger.[20] Second, the platform offers a higher privacy level when the consumer's type is less persistent. This observation is consistent with the intuition that the consumer faces a higher privacy cost when her type is less persistent. Finally, the equilibrium activity level can be non-monotone in ϕ, when the platform chooses a privacy policy. Indeed, the steady-state activity level at $\phi = 0.98$ is higher than the one at $\phi = 0.5$, but lower than the one at $\phi = 0.1$.

[20]For example, if $\phi = 0.5$ and $v = 200$, then in the steady state the platform offers $\gamma_t \approx 90$ and the consumer chooses $a_{max} = 2$.

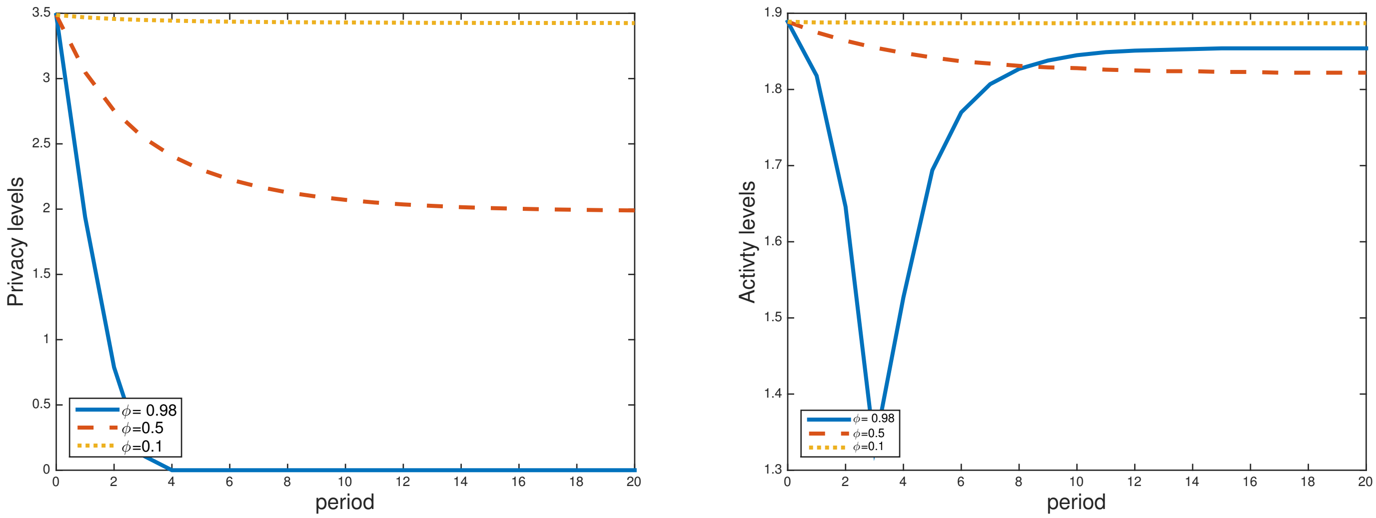

Figure 3(a): Privacy level γ_t

Figure 3(b): Activity level a_t

Figure 3: Equilibrium under $u(a) = 2a - \frac{1}{2}a^2$, $v = 10$, $\sigma_0^2 = 1$, $\phi \in \{0.1, 0.5, 0.98\}$.

7.4 General Payoffs of the Platform

Most of the results continue to hold if the platform's final payoff from a sequence of posterior variances is $\Pi((\sigma_t^2)_{t\in\mathbb{N}})$, where $\Pi : \mathbb{R}_+^\infty \to \mathbb{R}$ is bounded and coordinate-wise strictly decreasing. This generalization does not change the analysis, because in the equilibrium under monopoly or competition, a deviation by the platform increases σ_t^2 for all $t \in \mathbb{N}$.[21] For example, suppose he platform sells information to a sequence of short-lived data buyers. Any information sold in period t is freely replicable later and thus has a price of zero in any period $s \geq t+1$. Then, the platform's payoff in period t equals the value of information generated in period t—i.e., the platform's ex ante payoff is $\sum_{t=1}^{\infty} \delta_P^{t-1}(\sigma_{t-1}^2 - \sigma_t^2)$, which is decreasing in each σ_t^2.

8 Conclusion

This paper studies a dynamic model of consumer privacy and platform data collection. The fundamental feature of the model is that data collection today reduces a consumer's marginal loss of

[21] The platform's deviation may not uniformly increase posterior variances in Theorem 1. However, the proof of this theorem rests on the argument that if the equilibrium fails to meet certain conditions such as $\sigma_t^2 \to 0$, the platform can deviate and uniformly decrease posterior variances. Thus, Theorem 1 continues to hold with the same proof under this general $\Pi(\cdot)$.

giving up privacy in the future. I examine dynamic implications of this idea. First, a monopoly platform is able to collect information over time by committing to not collect too much information in early periods. In equilibrium, the consumer eventually loses privacy but keeps choosing a high level of activity. This outcome can arise even if a platform prioritizes user activity over data collection. Second, if the platform has weaker commitment power and does not initially hold much consumer information, it may fail to collect information in some equilibrium: The consumer refuses to provide information, anticipating that small privacy loss will lead to the complete privacy loss. This result implies that a company that already holds consumer data is more capable of collecting additional data. Finally, a decreasing marginal privacy cost could render competition unhelpful, because a consumer is more likely to stick with a platform on which they have less privacy.

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar (2019), "Too much data: Prices and inefficiencies in data markets." Working Paper 26296, National Bureau of Economic Research.

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman (2016), "The economics of privacy." *Journal of economic Literature*, 54, 442–92.

Agarwal, Anish, Munther Dahleh, and Tuhin Sarkar (2019), "A marketplace for data: An algorithmic solution." In *Proceedings of the 2019 ACM Conference on Economics and Computation*, 701–726.

Ali, S Nageeb, Greg Lewis, and Shoshana Vasserman (2020), "Voluntary disclosure and personalized pricing." In *Proceedings of the 21st ACM Conference on Economics and Computation*, 537–538.

Argenziano, Rossella and Alessandro Bonatti (2020), "Information revelation and privacy protection."

Athey, Susan, Christian Catalini, and Catherine Tucker (2017), "The digital privacy paradox: Small money, small costs, small talk." Technical report, National Bureau of Economic Research.

Ball, Ian (2020), "Scoring strategic agents." *arXiv preprint arXiv:1909.01888*.

Becker, Gary S and Kevin M Murphy (1988), "A theory of rational addiction." *Journal of political Economy*, 96, 675–700.

Bergemann, Dirk, Alessandro Bonatti, and Tan Gan (2019), "The economics of social data." Cowles foundation discussion paper 2203, Yale University.

Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin (2018), "The design and price of information." *American Economic Review*, 108, 1–48.

Bonatti, Alessandro and Gonzalo Cisternas (2020), "Consumer scores and price discrimination." *The Review of Economic Studies*, 87, 750–791.

Brunnermeier, Markus K, Rohit Lamba, and Carlos Segura-Rodriguez (2020), "Inverse selection." *Available at SSRN*.

Calzolari, Giacomo and Alessandro Pavan (2006), "On the optimality of privacy in sequential contracting." *Journal of Economic theory*, 130, 168–204.

Casadesus-Masanell, Ramon and Andres Hervas-Drane (2015), "Competing with privacy." *Management Science*, 61, 229–246.

Chiou, Lesley and Catherine Tucker (2017), "Search engines and data retention: Implications for privacy and antitrust." Technical report, National Bureau of Economic Research.

Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim (2019), "Privacy and personal data collection with information externalities." *Journal of Public Economics*, 173, 113–124.

Crémer, Jacques, Yves-Alexandre de Montjoye, and Heike Schweitzer (2019), "Competition policy for the digital era." *Report for the European Commission*.

De Corniere, Alexandre and Greg Taylor (2020), "Data and competition: a general framework with applications to mergers, market structure, and privacy policy."

Dwork, Cynthia, Aaron Roth, et al. (2014), "The algorithmic foundations of differential privacy." *Foundations and Trends in Theoretical Computer Science*, 9, 211–407.

Easley, David, Shiyang Huang, Liyan Yang, and Zhuo Zhong (2018), "The economics of data." *Available at SSRN 3252870*.

Eilat, Ran, Kfir Eliaz, Xiaosheng Mu, et al. (2019), *Optimal Privacy-Constrained Mechanisms*. Centre for Economic Policy Research.

Elliott, Matthew and Andrea Galeotti (2019), "Market segmentation through information." *Available at SSRN 3432315*.

Fainmesser, Itay P, Andrea Galeotti, and Ruslan Momot (2019), "Digital privacy." *Available at SSRN*.

Farboodi, Maryam and Laura Veldkamp (2020), "A growth model of the data economy." Technical report, MIT Sloan.

Frankel, Alex and Navin Kartik (2019a), "Improving information from manipulable data." *arXiv preprint arXiv:1908.10330*.

Frankel, Alex and Navin Kartik (2019b), "Muddled information." *Journal of Political Economy*, 127, 1739–1776.

Furman, Jason, D Coyle, A Fletcher, D McAules, and P Marsden (2019), "Unlocking digital competition: Report of the digital competition expert panel." *HM Treasury, United Kingdom*.

Ghosh, Arpita and Aaron Roth (2011), "Selling privacy at auction." In *Proceedings of the 12th ACM conference on Electronic commerce*, 199–208.

Haghpanah, Nima and Ron Siegel (2019), "Pareto improving segmentation of multi-product markets." Technical report, Working Paper.

Hagiu, Andrei and Julian Wright (2020), "Data-enabled learning, network effects and competitive advantage." Technical report, Working paper.

Holmström, Bengt (1999), “Managerial incentive problems: A dynamic perspective.” *The Review of Economic Studies*, 66, 169–182.

Hörner, Johannes and Nicolas S Lambert (2018), “Motivational ratings.”

Hörner, Johannes and Andrzej Skrzypacz (2016), “Selling information.” *Journal of Political Economy*, 124, 1515–1562.

Ichihashi, Shota (2020a), “The economics of data externalities.”

Ichihashi, Shota (2020b), “Online privacy and information disclosure by consumers.” *American Economic Review*, 110, 569–95.

Jullien, Bruno, Yassine Lefouili, and Michael Riordan (2018), “Privacy protection and consumer retention.”

Kummer, Michael and Patrick Schulte (2019), “When private information settles the bill: Money and privacy in googles market for smartphone applications.” *Management Science*.

Lefouili, Yassine and Ying Lei Toh (2019), “Privacy regulation and quality investment.” Technical report, Toulouse School of Economics (TSE).

Liang, Annie and Erik Madsen (2020), “Data and incentives.” *arXiv preprint arXiv:2006.06543*.

Lin, Tesary (2019), “Valuing intrinsic and instrumental preferences for privacy.” *Available at SSRN*, 3406412.

Loertscher, Simon and Leslie M Marx (2020), “Digital monopolies: privacy protection or price regulation?” *International Journal of Industrial Organization*, 102623.

Milgrom, Paul, Chris Shannon, et al. (1994), “Monotone comparative statics.” *Econometrica*, 62, 157–157.

Morton, Fiona Scott, Theodore Nierenberg, Pascal Bouvier, Ariel Ezrachi, Bruno Jullien, Roberta Katz, Gene Kimmelman, A Douglas Melamed, and Jamie Morgenstern (2019), “Report: Committee for the study of digital platforms-market structure and antitrust subcommittee.” *George*

J. Stigler Center for the Study of the Economy and the State, The University of Chicago Booth School of Business.

Prufer, Jens and Christoph Schottmüller (2017), "Competing with big data." *Tilburg Law School Research Paper*.

Srinivasan, Dina (2019), "The antitrust case against facebook: A monopolist's journey towards pervasive surveillance in spite of consumers' preference for privacy." *Berkeley Business Law Journal*, 16, 39.

Tang, Huan (2019), "The value of privacy: Evidence from online borrowers." Technical report, Working Paper, HEC Paris.

Tucker, Catherine (2018), "Privacy, algorithms, and artificial intelligence." *The Economics of Artificial Intelligence: An Agenda*.

Yang, Kai Hao (2019), "Selling consumer data for profit: Optimal market-segmentation design and its consequences." Technical report, Working Paper.

Appendix

A A Microfoundation of the Privacy Cost and the Platform Revenue

This appendix microfounds the payoff functions. We first describe a static interaction between a consumer and a "seller" in a product market, without considering the platform. We then embed it into the original model.

Consider a market that consists of a consumer and a seller. The consumer chooses a quantity $q \in \mathbb{R}$ to maximize her utility $Xq - \frac{1}{2}q^2 - pq$ given a unit price $p \in \mathbb{R}$. Her type $X \sim \mathcal{N}(\mu_0, \sigma_0^2)$ is now the willingness to pay for the product. The seller knows the prior distribution (μ_0, σ_0^2) and receives a signal $s = X + \varepsilon$ with $\varepsilon \sim \mathcal{N}(0, \sigma_\varepsilon^2)$. Only the consumer observes the true X. The seller first sets p to maximize its revenue, then the consumer chooses q to maximize her utility.

Although a more general setup appears in the literature, I provide the analysis for completeness. Suppose the seller observes a signal and holds a posterior mean μ of X. The consumer's demand is $q = X - p$, and thus the seller's expected revenue is $p(\mu - p)$, so the optimal price is $p^* = \frac{\mu}{2}$ with the quantity $X - \frac{\mu}{2}$. The consumer's expected payoff is

$$\begin{aligned}
&\mathbb{E}\left[(X - p^*)q^* - \frac{1}{2}(q^*)^2\right] \\
=&\frac{1}{2}\mathbb{E}\left[\left(X - \frac{1}{2}\mu\right)^2\right] \\
=&\frac{1}{2}\mathbb{E}\left[(X-\mu)^2\right] + \frac{1}{8}\mathbb{E}\left[(\mu-\mu_0)^2\right] + \frac{1}{8}\mu_0^2 \\
=&\frac{1}{2}\mathbb{E}\left[(X-\mu_0)^2\right] - \frac{3}{8}\mathbb{E}\left[(\mu-\mu_0)^2\right] + \frac{1}{8}\mu_0^2.
\end{aligned}$$

The expectation $\mathbb{E}$ is with respect to the joint distribution of (X, μ). The first and the last terms do not depend on the signal structure. As a result, the signal decreases consumer surplus from this transaction by

$$\frac{3}{8}\mathbb{E}[(\mu - \mu_0)^2] = \frac{3}{8} \cdot \frac{(\sigma_0^2)^2}{\sigma_0^2 + \sigma_\varepsilon^2}. \tag{7}$$

The posterior variance σ_t^2 of X and the variance σ_ε^2 of the noise ε satisfy the equation $\frac{1}{\sigma_t^2} = \frac{1}{\sigma_0^2} + \frac{1}{\sigma_\varepsilon^2}$. Solving this equation with respect to σ_ε^2 and plugging it into (7), we obtain

$$\frac{3}{8} \cdot \frac{(\sigma_0^2)^2}{\sigma_0^2 + \frac{1}{\frac{1}{\sigma_t^2} - \frac{1}{\sigma_0^2}}} = \frac{3}{8} \cdot \frac{(\sigma_0^2)^2 \cdot \left(\frac{1}{\sigma_t^2} - \frac{1}{\sigma_0^2}\right)}{\sigma_0^2 \cdot \left(\frac{1}{\sigma_t^2} - \frac{1}{\sigma_0^2}\right) + 1} = \frac{3}{8} \cdot \frac{(\sigma_0^2)^2 \cdot \left(\frac{1}{\sigma_t^2} - \frac{1}{\sigma_0^2}\right)}{\frac{\sigma_0^2}{\sigma_t^2}} = \frac{3}{8}\left(\sigma_0^2 - \sigma_t^2\right),$$

which is the original privacy cost function $v(\sigma_0^2 - \sigma_t^2)$ with $v = \frac{3}{8}$. Similarly, the information increases the seller's revenue by

$$\frac{1}{4}\left(\sigma_0^2 - \sigma_t^2\right), \tag{8}$$

which is equivalent to the platform's payoff in the original model.

We obtain the original model by assuming that the platform sells information to sellers who use it to price discriminate the consumer. The detail is as follows. Outside of the platform, the consumer interacts with seller t in period t, and her willingness to pay for seller t's product is $X_t \sim \mathcal{N}(\mu_0, \sigma_0^2)$, which is now *IID across* t.[22] The consumer's activity on the platform in each period yields her utilities and generates signals for future sellers. Specifically, each period t consists of

[22] I impose the IID assumption so that the prior distribution of the consumer's willingness to pay in the product market is the same across t.

the following events: (i) the consumer chooses a_t, (ii) the platform collects and sells information (by posting a price) to seller t, and (iii) the seller sets the price and the consumer chooses quantity. Precisely, the platform can sell seller t the signal $s_t = X_t + \varepsilon_t$ with

$$\varepsilon_t \sim \mathcal{N}\left(0, \frac{1}{\sum_{s=1}^{t} \frac{1}{\frac{1}{a_s}+\gamma_s}}\right).$$

That is, the platform collects information about the consumer's willingness to pay for product t based on her past activities. We now obtain the baseline model: The platform's per-period payoff, which is the revenue it can earn by selling the signal to seller t, is $\frac{1}{4}\left(\sigma_0^2 - \sigma_t^2\right)$, where σ_t^2 evolves according to (1). The consumer's payoff is $u(a_t) - \frac{3}{8}\left(\sigma_0^2 - \sigma_t^2\right)$.

B Existence of Equilibrium Under Long-Run Commitment

I prove the existence of an equilibrium under long-run commitment with $\delta_C > 0$ (for a myopic consumer, Claim 1 constructs an equilibrium). I introduce some notations. Let $\mathcal{A} := A^\infty$ denote the set of all sequences of activity levels. Because $A \subset \mathbb{R}_+$ is finite, it is compact, so $\mathcal{A}$ is compact with respect to product topology. Let $\boldsymbol{a}$ denote a generic element of $\mathcal{A}$, with the t-th coordinate denoted by a_t. Let $\boldsymbol{\Gamma} := [0,\infty]^{\mathbb{N}}$ denote the set of all privacy policies. Let $\boldsymbol{\gamma}$ denote a generic element of $\boldsymbol{\Gamma}$, with the t-th coordinate denoted by γ_t. I consider the ordered topology for $\overline{\mathbb{R}}_+$ and the product topology for $\boldsymbol{\Gamma}$. Finally, let $U_t(\boldsymbol{a}, \boldsymbol{\gamma})$ denote the consumer's flow payoff in period t, given an outcome $(\boldsymbol{a}, \boldsymbol{\gamma})$. Note that $U_t(\boldsymbol{a}, \boldsymbol{\gamma})$ depends only on $(a_1, \ldots, a_t)$ and $(\gamma_1, \ldots, \gamma_t)$.

Given any privacy policy $\boldsymbol{\gamma} \in \boldsymbol{\Gamma}$, the consumer's problem is

$$\max_{\boldsymbol{a}\in\mathcal{A}} \sum_{t=1}^{\infty} \delta_C^{t-1} \left[u(a_t) - v \cdot \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{a_s}+\gamma_s}}\right)\right]. \tag{9}$$

For any $\boldsymbol{\gamma} \in \boldsymbol{\Gamma}$, let $\mathcal{A}^*(\boldsymbol{\gamma}) \subset \mathcal{A}$ denote the set of all maximizers of (9).

Lemma 2. *The correspondence $\mathcal{A}^*(\boldsymbol{\gamma})$ is non-empty, compact, and upper hemicontinuous in $\boldsymbol{\gamma}$.*

Proof. First, $\mathcal{A}$ is compact with respect to product topology. Second, the objective function is continuous: To see this, take any sequence of the consumer's choices $(\boldsymbol{a}^n)_{n=1}^{\infty}$ such that $\boldsymbol{a}^n \to \boldsymbol{a}^*$.

This implies that for each $t \in \mathbb{N}$, $\lim_{n\to\infty} a_t^n \to a_t^*$. The consumer's period-$t$ payoff $U_t(\boldsymbol{a}, \boldsymbol{\gamma}) := u(a_t) - v \cdot \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^t \frac{1}{\frac{1}{a_s}+\gamma_s}}\right)$ is bounded from above and below by $u(a_{max})$ and $-v\sigma_0^2$, respectively. Define $B := \max(u(a_{max}), v\sigma_0^2) > 0$. Take any $\varepsilon > 0$, and let T^* satisfy $\frac{\delta_C^{T^*}}{1-\delta_C} B < \frac{\varepsilon}{4}$. Take a sufficiently large n so that for each $t \leq T^*$, $\delta_C^{t-1} |U_t(\boldsymbol{a}^n, \boldsymbol{\gamma}) - U_t(\boldsymbol{a}^*, \boldsymbol{\gamma})| < \frac{\varepsilon}{2T^*}$. These inequalities imply that

$$\left| \sum_{t=1}^{\infty} \delta_C^{t-1} U_t(\boldsymbol{a}^n, \boldsymbol{\gamma}) - \sum_{t=1}^{\infty} \delta_C^{t-1} U_t(\boldsymbol{a}^*, \boldsymbol{\gamma}) \right| < \varepsilon.$$

Thus the objective function in (9) is continuous in $\boldsymbol{a}$. Berge maximum theorem implies that $\mathcal{A}^*(\boldsymbol{\gamma})$ is non-empty, compact, and upper hemicontinuous in $\boldsymbol{\gamma}$. □

Next, I show properties of the consumer's objective $U(\boldsymbol{a}, \boldsymbol{\gamma}) := \sum_{t=1}^{\infty} \delta_C^{t-1} U_t(\boldsymbol{a}, \boldsymbol{\gamma})$. Abusing notation, for any $\boldsymbol{a}, \boldsymbol{a}' \in \mathcal{A}$, write $\boldsymbol{a} \geq \boldsymbol{a}'$ if and only if $a_t \geq a_t'$ for all $t \in \mathbb{N}$. $\geq$ is a partial order on $\mathcal{A}$, and $(\mathcal{A}, \geq)$ is a lattice.

Lemma 3. *For any $\boldsymbol{\gamma}$, $U(\boldsymbol{a}, \boldsymbol{\gamma})$ is supermodular in $\boldsymbol{a}$.*

Proof. Take any $\boldsymbol{\gamma}$. Below, I omit $\boldsymbol{\gamma}$ and write $U(\cdot, \boldsymbol{\gamma})$ as $U(\cdot)$. Take any $\boldsymbol{a}, \boldsymbol{b} \in \mathcal{A}$. For each $n \in \mathbb{N}$, define $(\boldsymbol{a} \vee \boldsymbol{b})^n$ as

$$(\boldsymbol{a} \vee \boldsymbol{b})^n = \begin{cases} \max(a_t, b_t) & \text{if} \quad t \leq n, \\ a_t & \text{if} \quad t > n. \end{cases} \tag{10}$$

Similarly, define $(\boldsymbol{a} \wedge \boldsymbol{b})^n$ as

$$(\boldsymbol{a} \wedge \boldsymbol{b})^n = \begin{cases} \min(a_t, b_t) & \text{if} \quad t \leq n, \\ a_t & \text{if} \quad t > n. \end{cases} \tag{11}$$

Also, define $\boldsymbol{b}^n$ as

$$\boldsymbol{b}^n = \begin{cases} b_t & \text{if} \quad t \leq n, \\ a_t & \text{if} \quad t > n. \end{cases} \tag{12}$$

In product topology, $(\boldsymbol{a} \vee \boldsymbol{b})^n \to \boldsymbol{a} \vee \boldsymbol{b}$, $(\boldsymbol{a} \wedge \boldsymbol{b})^n \to \boldsymbol{a} \wedge \boldsymbol{b}$, and $\boldsymbol{b}^n \to \boldsymbol{b}$ as $n \to \infty$. For each $t \in \mathbb{N}$ and $n \in \mathbb{N}$, $U_t(\boldsymbol{a}, \boldsymbol{\gamma})$ is supermodular in $(a_1, \ldots, a_n)$, because it has increasing differences in each pair (a_t, a_s). Thus for each $n \in \mathbb{N}$, $U(\boldsymbol{a})$ is supermodular in the first n activity levels, $(a_1, \ldots, a_n) \in \mathbb{R}^n_+$. We then have $U((\boldsymbol{a} \vee \boldsymbol{b})^n) + U((\boldsymbol{a} \wedge \boldsymbol{b})^n) \geq U(\boldsymbol{a}) + U(\boldsymbol{b}^n)$. Because $U(\cdot)$ is continuous, we can take $n \to \infty$ and obtain $U(\boldsymbol{a} \vee \boldsymbol{b})) + U(\boldsymbol{a} \wedge \boldsymbol{b}) \geq U(\boldsymbol{a}) + U(\boldsymbol{b})$. $\square$

The supermodularity implies the consumer has the "greatest" optimal choice.

Lemma 4. *For each $\boldsymbol{\gamma}$, the set $\mathcal{A}^*(\boldsymbol{\gamma})$ of optimal choices is a sublattice of $\mathcal{A}$. There is an $\bar{\boldsymbol{a}} \in \mathcal{A}^*(\boldsymbol{\gamma})$ such that for any $\boldsymbol{a} \in \mathcal{A}^*(\boldsymbol{\gamma})$, $\bar{\boldsymbol{a}} \geq \boldsymbol{a}$.*

Proof. First, Corollary 2 of Milgrom et al. (1994) implies that $\mathcal{A}^*(\boldsymbol{\gamma})$ is a sublattice of $\mathcal{A}$. Let $\mathcal{A}^*_t(\boldsymbol{\gamma})$ denote the projection of $\mathcal{A}^*(\boldsymbol{\gamma})$ on the t-th coordinate, i.e.,

$$\mathcal{A}^*_t(\boldsymbol{\gamma}) := \left\{a \in A : \exists \boldsymbol{a}^* \in \mathcal{A}^*(\boldsymbol{\gamma}) \text{ s.t. } \boldsymbol{a}^*_t = a\right\}. \tag{13}$$

For each $k \in \mathbb{N}$, let $\boldsymbol{a}^k$ denote an optimal policy such that the consumer chooses $a_k = \max \mathcal{A}^*_k(\boldsymbol{\gamma})$ in period k. Define $\bar{\boldsymbol{a}}^k := \boldsymbol{a}^1 \vee \cdots \vee \boldsymbol{a}^k$. Because $\mathcal{A}^*(\boldsymbol{\gamma})$ is sublattice, for any $k \in \mathbb{N}$, $\bar{\boldsymbol{a}}^k$ maximizes (9). We also have $\bar{\boldsymbol{a}}^k \to \bar{\boldsymbol{a}}$, where $\bar{a}_t = \max \mathcal{A}^*_k(\boldsymbol{\gamma})$ for any $k \in \mathbb{N}$. Because $\mathcal{A}^*(\boldsymbol{\gamma})$ is compact, $\bar{\boldsymbol{a}} \in \mathcal{A}^*(\boldsymbol{\gamma})$. By construction, for any $\boldsymbol{a} \in \mathcal{A}^*(\boldsymbol{\gamma})$, $\bar{\boldsymbol{a}} \geq \boldsymbol{a}$. $\square$

For each $\boldsymbol{\gamma} \in \boldsymbol{\Gamma}$, let $\bar{\boldsymbol{a}}(\boldsymbol{\gamma}) := (\bar{a}_t(\boldsymbol{\gamma}))_{t \in \mathbb{N}}$ denote the greatest strategy of the consumer defined in Lemma 4.

Lemma 5. *For each $t \in \mathbb{N}$, $\bar{a}_t(\boldsymbol{\gamma})$ is upper semicontinuous in $\boldsymbol{\gamma} \in \boldsymbol{\Gamma}$.*

Proof. By Lemma 2, $\mathcal{A}^*(\boldsymbol{\gamma})$ is upper hemicontinuous, so the set $\mathcal{A}^*_t(\boldsymbol{\gamma})$ of all activity levels that can be chosen in period t is upper hemicontinuous in $\boldsymbol{\gamma}$. Thus, it is enough to show that for any upper hemicontinuous and compact-valued correspondence $\phi : X \twoheadrightarrow \mathbb{R}$, $f(x) := \max \phi(x)$ is upper semicontinuous. To show this, take any $x_n \to x$. For each n, define $y_n = f(x_n)$. Because there is a subsequence $y_{n(k)}$ of y_n that converges to $\limsup y_n$, it holds that $\limsup y_n = \lim y_{n(k)} = \lim f(x_{n(k)}) \leq f(\lim x_{n(k)}) = f(x)$. The inequality holds because ϕ has a closed graph. Connecting the left and right sides, we establish that $f(\cdot)$ is upper semicontinuous. $\square$

Lemma 6. *There exists an equilibrium in the game of long-run commitment power.*

Proof. The platform's objective is

$$\sum_{t=1}^{\infty} \delta_P^{t-1} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{\bar{a}_s(\boldsymbol{\gamma})} + \gamma_s}} \right). \tag{14}$$

To show it is upper semicontinuous, take $\boldsymbol{\gamma}^n \to \boldsymbol{\gamma}$. Then,

$$\begin{aligned}
&\limsup_{n\to\infty} \sum_{t=1}^{\infty} \delta_P^{t-1} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{\bar{a}_s(\boldsymbol{\gamma}^n)} + \gamma_s^n}} \right) \\
&= \lim_{k\to\infty} \sup_{n\geq k} \sum_{t=1}^{\infty} \delta_P^{t-1} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{\bar{a}_s(\boldsymbol{\gamma}^n)} + \gamma_s^n}} \right) \\
&\leq \lim_{k\to\infty} \sum_{t=1}^{\infty} \delta_P^{t-1} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \sup_{n\geq k} \frac{1}{\frac{1}{\bar{a}_s(\boldsymbol{\gamma}^n)} + \gamma_s^n}} \right) \\
&= \sum_{t=1}^{\infty} \delta_P^{t-1} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \lim_{k\to\infty} \sup_{n\geq k} \frac{1}{\frac{1}{\bar{a}_s(\boldsymbol{\gamma}^n)} + \gamma_s^n}} \right) \\
&= \sum_{t=1}^{\infty} \delta_P^{t-1} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\liminf_{n\to\infty} \frac{1}{\bar{a}_s(\boldsymbol{\gamma}^n)} + \gamma_s^n}} \right) \\
&\leq \sum_{t=1}^{\infty} \delta_P^{t-1} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{\limsup_{n\to\infty} \bar{a}_s(\boldsymbol{\gamma}^n)} + \lim_{k\to\infty} \inf_{n\geq k} \gamma_s^n}} \right) \\
&\leq \sum_{t=1}^{\infty} \delta_P^{t-1} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{\bar{a}_s(\boldsymbol{\gamma})} + \gamma}} \right).
\end{aligned}$$

The second equality comes from the dominated convergence theorem, and the last inequality uses the upper semicontinuity of $\bar{a}_s(\boldsymbol{\gamma})$. Thus, given the consumer's optimal behavior, the platform's objective is upper semicontinuous. Since $\boldsymbol{\Gamma}$ is compact, there is a privacy policy $\boldsymbol{\gamma}^*$ that maximizes the platform's objective. The strategy profile $(\boldsymbol{\gamma}^*, \bar{\boldsymbol{a}}(\cdot))$ is an equilibrium. $\square$

C Consumer Behavior Under a Stationary Privacy Policy: Proof of Proposition 1

This Appendix uses notations introduced at the beginning of Appendix B.

C.1 Properties of Consumer Value Function

First, I prove useful properties of the consumer's value function that hold for any privacy policy. Let $(\bar{a}_t(\boldsymbol{\gamma}))_{t\in\mathbb{N}}$ denote the greatest best response of the consumer constructed in Lemma 4. For each privacy policy $\boldsymbol{\gamma} \in \boldsymbol{\Gamma}$, define

$$V_{\boldsymbol{\gamma}}(\rho) := \sum_{t=1}^{\infty} \delta_C^{t-1} \left[u(\bar{a}_t(\boldsymbol{\gamma})) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho + \sum_{s=1}^{t} \frac{1}{\frac{1}{\bar{a}_s(\boldsymbol{\gamma})} + \gamma_s}} \right) \right]. \tag{15}$$

$V_{\boldsymbol{\gamma}}(\rho)$ is the consumer's continuation value, starting from the posterior variance $\sigma^2 = \frac{1}{\rho}$.

Lemma 7. *For any $\boldsymbol{\gamma} \in \boldsymbol{\Gamma}$, $V_{\boldsymbol{\gamma}}(\cdot) : \mathbb{R}_{++} \to \mathbb{R}$ is decreasing and convex. For any $\rho > 0$ and $\Delta > 0$,* $\lim_{\rho\to\infty} V_{\boldsymbol{\gamma}}(\rho) - V_{\boldsymbol{\gamma}}(\rho + \Delta) = 0$.

Proof. Fix any privacy policy $\boldsymbol{\gamma}$. Hereafter, I omit $\boldsymbol{\gamma}$ from subscripts (thus, the consumer value function is $V(\cdot)$). Consider the "T-period problem, " in which the consumer's payoff in any period $s \geq T+1$ is exogenously set to zero. For any $t \leq T$, let $V_t^T(\rho)$ denote the consumer's continuation value in the T-period problem starting from period t given $\frac{1}{\sigma_{t-1}^2} = \rho$:

$$V_t^T(\rho) = \max_{(a_t,\ldots,a_T)\in A^{T-t+1}} \sum_{s=t}^{T} \delta_C^{s-t} \left(u(a_s) - v \left(\sigma_0^2 - \frac{1}{\rho_{s-1} + \frac{1}{\frac{1}{a_s} + \gamma_s}} \right) \right).$$

Here, $\rho_{t-1} = \rho$, and $(\rho_t, \ldots, \rho_{T-1})$ are recursively defined by Bayes' rule given $(a_t, \ldots, a_{T-1})$. The standard argument of dynamic programming implies that for each $t = 1, \ldots, T$,

$$V_t^T(\rho) = \max_{a\in A} u(a) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a} + \gamma_t}} \right) + \delta_C V_{t+1}^T \left(\rho + \frac{1}{\frac{1}{a} + \gamma_t} \right), \tag{16}$$

where $V_{T+1}^T(\cdot) \equiv 0$. I use induction to show that $V_1^T(\rho)$ is decreasing and convex. First, $V_{T+1}^T \equiv 0$ is trivially decreasing and convex. Suppose V_{t+1}^T is decreasing and convex. Because $-v \cdot$

$\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a}+\gamma_t}}\right)$ has the same property and the upper envelope of decreasing convex functions are decreasing and convex, so does $V_t^T(\cdot)$. This induction argument implies that for each T, $V^T(\rho) = V_1^T(\cdot)$ is decreasing and convex. Also, for any ρ and $\Delta > 0$, $\lim_{\rho\to\infty} V^T(\rho) - V^T(\rho+\Delta) \to 0$.

Define $V^\infty(\rho) := \lim_{T\to\infty} V^T(\rho)$. $V^\infty(\rho)$ is decreasing and convex, because these properties are preserved under pointwise convergence. I show that $V^\infty(\rho)$ is the value function of the original problem, i.e., $V^\infty(\cdot) = V(\cdot)$. Take any ρ, and let $(\bar{a}_1, \bar{a}_2, \dots) \in \mathcal{A}^*(\boldsymbol{\gamma})$ denote the optimal policy. For any finite T,

$$V^T(\rho) \geq \sum_{s=1}^{T} \delta_C^{s-1}\left(u(\bar{a}_s) - v\left(\sigma_0^2 - \frac{1}{\rho_{s-1} + \frac{1}{\frac{1}{\bar{a}_s}+\gamma_s}}\right)\right). \tag{17}$$

By taking $T \to \infty$, we obtain $V^\infty(\rho) \geq V(\rho)$. Suppose to the contrary that $V^\infty(\rho) > V(\rho)$. Then, there is a sufficiently large $T \in \mathbb{N}$ such that $V^T(\rho) - \frac{\delta_C^T}{1-\delta_C} v\sigma_0^2 > V(\rho)$. If the consumer in the original infinite horizon problem adopts the T-optimal policy that gives $V^T(\rho)$ up to period t, then she can obtain a strictly greater payoff than $V(\rho)$, which is a contradiction. Thus, $V^\infty(\rho) = V(\rho)$.

Finally, I show that for any ρ and $\Delta > 0$, $\lim_{\rho\to\infty} V(\rho) - V(\rho+\Delta) \to 0$. Suppose the consumer starting from $\rho + \Delta$ chooses the policy $(\bar{a}_t^\rho)_{t\in\mathbb{N}}$ that is optimal for ρ. Let $(\hat{\rho}_t)_{t=1}^\infty$ denote the induced sequence of the precisions after $\rho + \Delta$, i.e., $\hat{\rho}_t = \rho + \Delta + \sum_{s=1}^{t} \frac{1}{\frac{1}{\bar{a}_s^\rho}+\gamma_s}$. Note that $\hat{\rho}_t \geq \rho_t$ for all $t \in \mathbb{N}$. Then, it holds that $0 \leq V(\rho) - V(\rho+\Delta) \leq \sum_{t=1}^\infty \delta_C^{t-1}\left(\frac{1}{\rho} - \frac{1}{\rho+\Delta}\right) = \frac{1}{1-\delta_C}\left(\frac{1}{\rho} - \frac{1}{\rho+\Delta}\right)$. Thus, $\lim_{\rho\to\infty} V(\rho) - V(\rho+\Delta) = 0$. □

C.2 Proof of Proposition 1

Proof. If γ_t is constant across t, the consumer problem is a stationary dynamic programming. Suppose $\gamma_t = \gamma \in \mathbb{R}_+$ for all t. The value function $V(\cdot)$ satisfies the Bellman equation

$$V(\rho) = \max_{a\in A} u(a) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a}+\gamma}}\right) + \delta_C V\left(\rho + \frac{1}{\frac{1}{a}+\gamma}\right). \tag{18}$$

Again, I suppress the dependence of $V(\cdot)$ on γ. Lemma 7 implies that $V(\cdot)$ is decreasing and convex. Thus, the maximand in (18) has the increasing differences in (a, ρ). Thus, $\bar{a}(v, \gamma, \rho)$, the greatest maximizer, is increasing in ρ. Note that $\rho_t \leq \rho_{t+1}$, and the inequality is strict if and only if $a_t > 0$. As a result, the consumer's optimal behavior is either (i) $a_t = 0$ for all t, or (ii) $a_1 > 0$ and a_t is increasing in t. Now, define

$$v^*(\gamma) := \sup\{v \in \mathbb{R} : \bar{a}(v, \gamma, \rho_0) > 0\}, \quad \text{where} \quad \rho_0 = \frac{1}{\sigma_0^2}. \tag{19}$$

The consumer's payoff from any strategy with $a_1 > 0$ is strictly decreasing in v and strictly increasing in γ, whereas her payoff from $a_t \equiv 0$ is independent of (v, γ). As a result, if $\bar{a}(v, \gamma, \rho_0) > 0$, then $\bar{a}(v', \gamma', \rho_0) > 0$ for any $v' < v$ and $\gamma' > \gamma$. Therefore, the consumer's behavior follows (i) and (ii) above if $v > v^*(\gamma)$ and $v < v^*(\gamma)$, respectively, and $v^*(\gamma)$ is increasing in γ. For any given v, as $\gamma \to \infty$, the consumer's ex ante payoff from (say) $a_t = a_{max} > 0$ for all t becomes positive. Thus, $\lim_{\gamma \to \infty} v^*(\gamma) = \infty$.

If $v < v^*(\gamma)$, then $a_t \geq a_1 > 0$ for all t. Since $\gamma < \infty$, we obtain $\lim_{t \to \infty} \sigma_t^2 \to 0$, or equivalently, $\lim_{t \to \infty} \rho_t = \infty$ with $\rho_t := \frac{1}{\sigma_t^2}$. By Lemma 7, for any $\rho > 0$ and $\Delta > 0$, $\lim_{\rho \to \infty} V(\rho) - V(\rho + \Delta) = 0$. This, combined with $\lim_{t \to \infty} \rho_t = \infty$, implies $\lim_{t \to \infty} \bar{a}_t(v, \gamma, \rho_t) = a_{max}$. □

D Equilibrium Under Long-Run Commitment: Proof of Theorem 1

D.1 Lemmas

I begin with two lemmas. First, suppose the platform changes privacy levels in any period t that belongs to a set $\mathcal{T} \subset \mathbb{N}$. If the change affects the consumer behavior and increases the precisions of signals of all periods in $\mathcal{T}$, she chooses higher activity levels in all other periods. Recall that $\bar{a}(\boldsymbol{\gamma}) \in \mathcal{A}$ denote the greatest best response of the consumer constructed in Lemma 4.

Lemma 8. *Take any $\boldsymbol{\gamma}, \boldsymbol{\gamma}' \in \boldsymbol{\Gamma}$. Define $\mathcal{T} = \{t \in \mathbb{N} : \gamma_t \neq \gamma'_t\}$. Suppose $\frac{1}{\bar{a}_t(\boldsymbol{\gamma})} + \gamma_t \leq \frac{1}{\bar{a}_t(\boldsymbol{\gamma}')} + \gamma'_t$ for all $t \in \mathcal{T}$. Then, $\bar{a}_t(\boldsymbol{\gamma}) \geq \bar{a}_t(\boldsymbol{\gamma}')$ for all $t \in \mathbb{N} \setminus \mathcal{T}$.*

Proof. Let $\boldsymbol{\beta}$ be any one of $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}'$. I decompose the consumer's problem (9) into two steps. First, given any $(a_t)_{t \in \mathcal{T}}$, the consumer chooses $(a_t)_{t \notin \mathcal{T}}$ to maximize the following hypothetical objective function:

$$\sum_{t=1}^{\infty} \delta_C^{t-1} \left[\mathbf{1}_{\{t \notin \mathcal{T}\}} u(a_t) - v \cdot \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{a_s} + \beta_s}} \right) \right]. \tag{20}$$

The consumer receives a benefit of $u(a_t)$ only in period $t \notin \mathcal{T}$. This leads to a mapping that maps any $(a_t)_{t \in \mathcal{T}}$ to the (greatest) optimal choice of $(a_t)_{t \notin \mathcal{T}}$. In the second step, the consumer chooses $(a_t)_{t \in \mathcal{T}}$ to maximize her original objective, taking the mapping $(a_t)_{t \in \mathcal{T}} \mapsto (a_t)_{t \notin \mathcal{T}}$ as given.

For any $t \in \mathcal{T}$, a_t affects (20) only through $\frac{1}{a_t} + \gamma_t$, because $\mathbf{1}_{\{t \notin \mathcal{T}\}} = 0$. Also the same argument as in the proof of Lemma 3 implies that (20) is supermodular in $\left((a_t)_{t \notin \mathcal{T}}, \left\{ \left(\frac{1}{a_s} + \gamma_s \right)^{-1} \right\}_{s \in \mathcal{T}} \right)$. This implies that if $\frac{1}{\bar{a}_t(\boldsymbol{\gamma})} + \gamma_t \leq \frac{1}{\bar{a}_t(\boldsymbol{\gamma}')} + \gamma_t'$ for all $t \in \mathcal{T}$, then $\bar{a}_t(\boldsymbol{\gamma}) \geq \bar{a}_t(\boldsymbol{\gamma}')$ for all $t \in \mathbb{N} \setminus \mathcal{T}$. $\square$

Next, the platform can commit to a high privacy level to induce a_{max} in any period.

Lemma 9. *There is a $\gamma_{max} < +\infty$ such that if the platform commits to $\gamma_t = \gamma_{max}$, then regardless of the privacy levels in other periods, the consumer chooses $a_t = a_{max}$. Also, there is a $\bar{\sigma}^2$ such that if $\sigma_{T-1}^2 \leq \bar{\sigma}^2$, then the consumer chooses $a_t = a_{max}$ for all $t \geq T$ for any $(\gamma_\tau)_{\tau \geq T}$.*

Proof. Let a' denote the second highest activity level in A. Take any $(a_t)_{t \in \mathbb{N}} \in \mathcal{A}$ such that $a_t < a_{max}$. Suppose the consumer changes her action in period t from a_t to a_{max}. This change increases her period-t benefit from $u(\cdot)$ by at least $u(a_{max}) - u(a') > 0$. The change also increases the sum of discounted privacy costs (from the perspective of period t) by

$$\begin{aligned}
&\sum_{s=t}^{\infty} \delta^{s-t} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a_{max}} + \gamma_{max}} + \sum_{\tau=t+1}^{s} \frac{1}{\frac{1}{a_\tau} + \gamma_\tau}} \right) - \sum_{s=t}^{\infty} \delta^{s-t} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a_t} + \gamma_{max}} + \sum_{\tau=t+1}^{s} \frac{1}{\frac{1}{a_\tau} + \gamma_\tau}} \right) \\
\leq &\sum_{s=t}^{\infty} \delta^{s-t} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a_{max}} + \gamma_{max}}} \right) - \sum_{s=t}^{\infty} \delta^{s-t} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a_t} + \gamma_{max}}} \right) \\
= &\frac{1}{1-\delta} \left(\frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a_t} + \gamma_{max}}} - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a_{max}} + \gamma_{max}}} \right) =: D(\sigma_{t-1}^2, \gamma_{max}).
\end{aligned}$$

First, we have $\lim_{\gamma_{max} \to \infty} D(\sigma_0^2, \gamma_{max}) = 0$, and $D(\sigma_t^2, \gamma_{max}) \leq D(\sigma_0^2, \gamma_{max})$ for any $\sigma_t^2 \leq \sigma_0^2$. Thus, for any γ_{max} such that $D(\sigma_0^2, \gamma_{max}) < u(a_{max}) - u(a')$, the consumer's optimal action is

a_{max} in period t. Also, even for $\gamma_{max} = 0$, $\lim_{\sigma^2_{t-1} \to 0} D(\sigma^2_{t-1}, 0) = 0$. Thus for a sufficiently small σ^2_{t-1}, the consumer chooses $a_\tau = a_{max}$ for all $\tau \geq t$ under any (continuation) privacy policy. $\square$

D.2 Proof of Theorem 1

Proof. First, I show $\lim_{t\to\infty} \sigma^2_t = 0$. Let $\boldsymbol{\gamma}^*$ denote the equilibrium privacy policy, and let $\boldsymbol{a}^*$ denote the equilibrium activity levels. Suppose to the contrary that $\lim_{t\to\infty} \sigma^2_t \neq 0$. Because σ^2_t is decreasing, $\lim_{t\to\infty} \sigma^2_t > 0$ exists. This implies $\frac{1}{a^*_t} + \gamma^*_t \to \infty$. I derive a contradiction.

Let $\gamma_{max} \in [0, +\infty)$ denote the privacy level defined in Lemma 9—i.e., the consumer chooses $a_t = a_{max}$ if $\gamma_t = \gamma_{max}$. If the platform commits to $\gamma_t = \gamma_{max}$, the variance of the noise of the signal in period t is $B := \frac{1}{a_{max}} + \gamma_{max}$. Take T^* such that for all $t \geq T^*$, $\frac{1}{a^*_t} + \gamma^*_t > B$. If the platform replaces γ^*_t with γ_{max} for all $t \geq T^*$ and commits to such a new policy ex ante, then the precision of the signal increases from $\frac{1}{\frac{1}{a^*_t}+\gamma^*_t}$ to B^{-1} in any period $t \geq T^*$. Lemma 8 implies that after the policy change, the consumer also chooses a weakly greater a_t for all $t < T^*$. To sum up, the platform can strictly increase its profit by replacing γ^*_t with γ_{max} for all $t \geq T^*$, which is a contradiction. The second part of Lemma 9 then implies that there is some T such that for all $t \geq T$, $a^*_t = a_{max}$.

Next, I write $\gamma^*_t(v)$ to clarify the dependence of the equilibrium privacy level on v. Suppose to the contrary that there is a T such that, for any $\underline{v}$, there is some $v \geq \underline{v}$ such that $\gamma^*_t(v) = 0$ for some $t \leq T$. Then we can find $v_n \to \infty$ and $t^* \leq T$ such that $\gamma^*_{t^*}(v_n) = 0$ for all n. However, for a sufficiently large v_n, $a^*_{t^*} = 0$ if $\gamma^*_{t^*}(v_n) = 0$. The reason is as follows. If the consumer changes her activity level from 0 to some $a > 0$, her gross payoff from $u(\cdot)$ increases by at at most $u(a_{max})$. In contrast, her privacy cost increases by at least

$$v\left(\frac{1}{\frac{1}{\sigma^2_0} + (t^*-1)a_{max}} - \frac{1}{\frac{1}{\sigma^2_0} + (t^*-1)a_{max} + a_{min}}\right) > 0,$$

where a_{min} is the smallest positive activity level in A. This expression is independent of the history and diverges to ∞ as $v \to \infty$. Thus for a large v, the consumer prefers $a = 0$. However, the platform can then commit to a high privacy level for period t^* to induce $a_{t^*} > 0$. By the same argument as the previous paragraph, this change also weakly increases the activity levels in all other periods. This is a contradiction.

Finally, we show that γ_t^* is zero after finite periods for a myopic consumer. If this is not the case, then there is a sequence of positive numbers $(g_t)_{t=1}^{\infty}$ such that $\gamma_t^* > g_t$ for infinitely many t's. Take T such that $\frac{1}{\frac{1}{\sigma_0^2}+\frac{T-1}{B}} < \bar{\sigma}^2$ and $\gamma_T^* > g_T$. Because the variance of the noise ε_t is at most B in each t, $\sigma_{T-1}^2 < \bar{\sigma}^2$ in equilibrium. The second part of Lemma 9 implies that if the platform sets $\gamma_T = 0$, the consumer still chooses a_{max}, which strictly decreases σ_T^2. As a result, the change of γ_T increases the platform payoff in any period $t \geq T$. Also, this change of γ_T does not affect the consumer's choice in $t < T$ because she is myopic. Thus, the platform benefits from changing γ_T from γ_T^* to 0, which is a contradiction. $\square$

E General Payoffs: Proof of Proposition 2

Proof of Proposition 2. First, we show Point 1. Define $v := \sup_{x\in[0,\sigma_0^2]} |C'(x)|$ and $C_v(x) = v(\sigma_0^2 - x)$. We have $C_v(\sigma_0^2) = C(\sigma_0^2) = 0$, $C_v'(\sigma^2) = -v \leq C'(\sigma^2)$, and thus $C_v(x) \geq C(x)$ for all $x \in [0, \sigma_0^2]$. As a result, the consumer under $C_v(\cdot)$ incurs greater privacy costs and marginal privacy costs than under $C(\cdot)$. Because the cost function $C_v(\cdot)$ satisfies the assumption of the original model, Proposition 1 implies that there is a γ^* such that if the platform chooses a privacy policy with $\gamma_t = \gamma^*$ for all $t \in \mathbb{N}$, the consumer's optimal behavior induces $a_t \to a_{max}$ and $\sigma_t^2 \to 0$. We show that the same long-run outcome arises when the consumer faces $C(\cdot)$. Suppose to the contrary that under $\gamma_t \equiv \gamma^*$, σ_t^2 does not converge to 0 when the consumer with $C(\cdot)$ acts optimally. Then for some $T \in \mathbb{N}$, the consumer chooses $a_t = 0$ for all $t \geq T$. However, in such a period T, the consumer facing $C_v(\cdot)$ strictly prefers some $a_t > 0$ to $a_t = 0$ (i.e., Point 1 of Proposition 1). As a result, the consumer facing $C(\cdot)$ can mimic this strategy to strictly increase her continuation value relative to taking zero activity levels forever, because she faces a uniformly lower marginal privacy cost. This is a contradiction. Finally, $\sigma_t^2 \to 0$ implies $a_t \to a_{max}$ by the same argument as Lemma 9.

Second, we show Point 2. For each $\delta_P \in [0,1]$, let $(\sigma_t^2(\delta_P))_{t\in\mathbb{N}}$ denote the equilibrium sequence of posterior variances, and define $\sigma_\infty^2(\delta_P) = \lim_{t\to\infty} \sigma_t^2(\delta_P)$. Suppose to the contrary that we can find a sequence $\delta_n \to 1$ and $\varepsilon > 0$ such that $\sigma_\infty^2(\delta_n) \geq \varepsilon$ for all n. Then, the platform's average revenue satisfies

$$(1-\delta_n)\sum_{t=1}^{\infty}\Pi(a_t(\delta_n),\sigma_t^2(\delta_n)) \le \Pi(a_{max},\varepsilon).$$

Point 1 implies that the platform has a privacy policy $\boldsymbol{\gamma}^*$ that induces $\lim_{t\to\infty}(a_t,\sigma_t^2) = (a_{max},0)$. As $\delta_P \to 1$, the platform's average payoff converges to $\Pi(a_{max},0)$. Thus, a sufficiently patient platform strictly prefers $\boldsymbol{\gamma}^*$ to the equilibrium policy, which is a contradiction.

Finally, suppose the long-run equilibrium activity $a_\infty(\delta_P) := \lim_{t\to\infty} a_t^*(\delta_P)$ has a well-defined limit $\lim_{\delta_P\to 1} a_\infty(\delta_P)$. Then, it must be a_{max}. Otherwise, the same argument as above implies that a patient platform prefers to deviate to $\boldsymbol{\gamma}^*$ that induces $\lim_{t\to\infty}(a_t,\sigma_t^2) = (a_{max},0)$. □

F Myopic Consumer: Proof of Claim 1

I prove two results that lead to Claim 1.

Lemma 10. *If the consumer is myopic, the platform adopts a greedy policy that myopically maximizes the precision of the signal in each period. Formally, the equilibrium policy* $(\gamma_1^*,\gamma_2^*,\dots)$ *is recursively defined as follows:*

$$\gamma_t^* \in \arg\min_{\gamma\ge 0} \frac{1}{a^*(\gamma,\hat{\sigma}_{t-1}^2)} + \gamma, \forall t\in\mathbb{N}, \tag{21}$$

$$\hat{\sigma}_0^2 = \sigma_0^2, \tag{22}$$

$$\hat{\sigma}_t^2 = \frac{1}{\frac{1}{\hat{\sigma}_{t-1}^2} + \frac{1}{\frac{1}{a^*(\gamma_t^*,\hat{\sigma}_{t-1}^2)}+\gamma_t^*}}, \forall t\in\mathbb{N}. \tag{23}$$

Also, given any privacy policy $\boldsymbol{\gamma}$*, let* $(\sigma_t^2)_{t\in\mathbb{N}}$ *denote the posterior variances induced by the consumer's optimal behavior. Then the optimal policy attains uniformly lower posterior variances:* $\hat{\sigma}_t^2 \le \sigma_t^2$ *for all* $t\in\mathbb{N}$.

Proof. Lemma 1 implies $a^*(\gamma,\sigma^2)$ is increasing in γ and decreasing in σ^2. Take any privacy policy $(\gamma_t)_{t\in\mathbb{N}}$ and let $(\sigma_t^2)_{t\in\mathbb{N}}$ denote the sequence of posterior variances induced by $a^*(\cdot,\cdot)$. I show $\hat{\sigma}_t^2 \le \sigma_t^2$ for all $t\in\mathbb{N}$. The inequality holds with equality for $t=0$. Take any $\tau\in\mathbb{N}$. Suppose

$\hat{\sigma}_t^2 \le \sigma_t^2$ for $t = 0, \ldots, \tau - 1$. It holds that

$$\sigma_\tau^2 = \frac{1}{\frac{1}{\sigma_{\tau-1}^2} + \frac{1}{\frac{1}{a^*(\gamma_\tau, \sigma_{\tau-1}^2)} + \gamma_\tau}} \ge \frac{1}{\frac{1}{\hat{\sigma}_{\tau-1}^2} + \frac{1}{\frac{1}{a^*(\gamma_\tau, \hat{\sigma}_{\tau-1}^2)} + \gamma_\tau}} \ge \frac{1}{\frac{1}{\hat{\sigma}_{\tau-1}^2} + \frac{1}{\frac{1}{a^*(\gamma_\tau^*, \hat{\sigma}_{\tau-1}^2)} + \gamma_\tau^*}} = \hat{\sigma}_\tau^2.$$

The first inequality follows from the inductive hypothesis and decreasing $a^*(\gamma, \cdot)$. The second inequality follows from (21). We now have $\hat{\sigma}_t^2 \le \sigma_t^2$ for all t, which implies the privacy policy described by (21), (22), and (23) is optimal. □

Corollary 1. *Assume the consumer is myopic, and let $(\boldsymbol{a}^*, \boldsymbol{\gamma}^*)$ denote the equilibrium outcome under long-run commitment in Lemma 10. The same outcome $(\boldsymbol{a}^*, \boldsymbol{\gamma}^*)$ arises in an equilibrium under one-period commitment.*

Proof. Suppose the platform has one-period commitment power. Consider the following strategy profile: At any node with posterior variance σ^2, the platform sets $\gamma \in \arg\min_{\gamma \ge 0} \frac{1}{a^*(\gamma, \sigma^2)} + \gamma$, and the consumer acts according to $a^*(\cdot, \sigma^2)$. Then, $(\boldsymbol{a}^*, \boldsymbol{\gamma}^*)$ arises on the path of play. The platform's deviation increases the posterior variances in all periods, an decreases its profit (Lemma 10). □

G Consumer-Worst Equilibrium: Proof of Theorem 2

Proof. **Step 1: Construction of MPE.** We write the consumer's discount factor δ_C as δ, and use a precision $\rho_t = \frac{1}{\sigma_t^2}$ as a state variable of MPE. Along any path of play, ρ_t is non-decreasing in t. Let $\gamma^*(\rho)$ denote the platform's choice of γ_t given $\rho_{t-1} = \rho$, and let $a^*(\rho, \gamma)$ denote the consumer's choice of a_t given $(\rho_{t-1}, \gamma_t) = (\rho, \gamma)$. Also, let $V_0(\rho)$ denote the consumer's continuation value when the initial state is ρ and $(\gamma_t, a_t) = (0, a_{max})$ for all $t \in \mathbb{N}$:

$$V_0(\rho) := \sum_{t=1}^{\infty} \delta^{t-1} \left[u(a_{max}) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho + t a_{max}} \right) \right]. \tag{24}$$

$V_0(\rho)$ is continuous, strictly decreasing, and strictly convex in $\rho \ge 0$.

First, we show that there is a $\rho(0) \in \mathbb{R}_{++}$ such that any strategy that satisfies $\gamma^*(\rho) = 0$ and $a^*(\rho, \gamma) = a_{max}$ for any $\rho \ge \rho(0)$ and any γ is an MPE in the game that starts from any $\rho \ge \rho(0)$. Given such $a^*(\cdot, \cdot)$ and the initial state $\rho \ge \rho(0)$, the platform prefers $\gamma = 0$ after any history, because the subsequent outcome is $(\gamma, a) = (0, a_{max})$ in all future periods, which maximizes the platform's payoff across all outcomes. It suffices to show that the consumer does not strictly benefit from a one-shot deviation to $a = 0$. At $\gamma = 0$, this condition is written as

$$u(a_{max}) - v\left(\sigma_0^2 - \frac{1}{\rho + a_{max}}\right) + \delta V_0\left(\rho + a_{max}\right) \geq -v\left(\sigma_0^2 - \frac{1}{\rho}\right) + \delta V_0(\rho)$$
$$\iff u(a_{max}) + \frac{v}{\rho + a_{max}} - \frac{v}{\rho} + \delta\left[V_0\left(\rho + a_{max}\right) - V_0(\rho)\right] \geq 0.$$

Both $\frac{v}{\rho+a_{max}} - \frac{v}{\rho}$ and $V_0\left(\rho + a_{max}\right) - V_0(\rho)$ are continuous and strictly increasing in ρ, and converge to 0 as $\rho \to \infty$. Since $u(a_{max}) > 0$, there is a unique $\rho(0) < \infty$ such that the inequality holds if and only if $\rho \geq \rho(0)$. Also, for any $\rho \geq \rho(0)$, the consumer does not strictly benefit from a one-shot deviation to $a = 0$ after the platform's deviation to $\gamma > 0$.

We have constructed an MPE with $\gamma_t^* \equiv 0$ and $a_t^* \equiv a_{max}$ for any initial state $\rho \geq \rho(0)$. Next, we construct an MPE for any initial state $\rho \in [\rho(1), \rho(0)]$, where $\rho(1) < \rho(0)$. Define

$$V(\rho, \gamma) := u(a_{max}) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a_{max}} + \gamma}}\right) + \delta V_0\left(\rho + \frac{1}{\frac{1}{a_{max}} + \gamma}\right).$$

$V(\rho, \gamma)$ is the consumer's continuation value when (i) the initial state is ρ, (ii) the platform sets γ and the consumer chooses a_{max} in the first period, and (iii) from the next period on, they always choose $(\gamma_t, a_t) = (0, a_{max})$. Consider the inequality

$$V(\rho, \gamma) \geq -v\left(\sigma_0^2 - \frac{1}{\rho}\right) + \delta \cdot V(\rho, \gamma),$$

which is written as

$$u(a_{max}) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a_{max}} + \gamma}}\right) + \delta V_0\left(\rho + \frac{1}{\frac{1}{a_{max}} + \gamma}\right)$$
$$\geq - v\left(\sigma_0^2 - \frac{1}{\rho}\right) + \delta \cdot \left[u(a_{max}) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a_{max}} + \gamma}}\right) + \delta V_0\left(\rho + \frac{1}{\frac{1}{a_{max}} + \gamma}\right)\right], \quad (25)$$

or equivalently,

$$(1-\delta)u(a_{max}) + \frac{v}{\rho + \frac{1}{\frac{1}{a_{max}}+\gamma}} - \frac{v}{\rho}$$
$$+\delta V_0\left(\rho + \frac{1}{\frac{1}{a_{max}}+\gamma}\right) + \delta\left[v\left(\sigma_0^2 - \frac{1}{\rho+\frac{1}{\frac{1}{a_{max}}+\gamma}}\right) - \delta V_0\left(\rho + \frac{1}{\frac{1}{a_{max}}+\gamma}\right)\right] \geq 0. \tag{26}$$

We show several properties of the left-hand side of (26). First, both sides of (25) are continuous in γ, and the left-hand side increases more than the right-hand side if γ increases (because of discounting). Thus, the left-hand side of (26) is continuous and strictly increasing in γ. It is also continuous and strictly increasing in ρ. In particular,

$$\begin{aligned}
&V_0\left(\rho + \frac{1}{\frac{1}{a_{max}}+\gamma}\right) + \left[v\left(\sigma_0^2 - \frac{1}{\rho+\frac{1}{\frac{1}{a_{max}}+\gamma}}\right) - \delta V_0\left(\rho + \frac{1}{\frac{1}{a_{max}}+\gamma}\right)\right]\\
=&K + v\sum_{t=1}^{\infty}\delta^{t-1}\frac{1}{\rho+\frac{1}{\frac{1}{a_{max}}+\gamma}+ta_{max}} - v\sum_{t=1}^{\infty}\delta^{t-1}\frac{1}{\rho+\frac{1}{\frac{1}{a_{max}}+\gamma}+(t-1)a_{max}}\\
=&K + v\sum_{t=1}^{\infty}\delta^{t-1}\left[\frac{1}{\rho+\frac{1}{\frac{1}{a_{max}}+\gamma}+ta_{max}} - \frac{1}{\rho+\frac{1}{\frac{1}{a_{max}}+\gamma}+(t-1)a_{max}}\right]
\end{aligned}$$

is strictly increasing in ρ, where K is a term that does not depend on ρ.

Because (25) holds with equality at $(\rho,\gamma)=(\rho(0),0)$, it holds with strict inequality at $\rho=\rho(0)$ for any $\gamma>0$. Then for any ρ that is smaller than but sufficiently close to $\rho(0)$, we can find a unique $\gamma(\rho)>0$ that satisfies (25) with equality. The left-hand side of (26) is increasing in ρ and γ. Thus, if $\gamma(\rho)$ exists for some $\rho<\rho(0)$, $\gamma(\rho')$ exists for any $\rho'\in[\rho,\rho(0))$. If ρ is such that no γ satisfies (25), then define $\gamma(\rho)=\infty$.

Because $\gamma(\rho)$ is decreasing in $\rho\leq\rho(0)$, for a ρ smaller than but close to $\rho(0)$, we obtain $\rho+\frac{1}{\frac{1}{a_{max}}+\gamma(\rho)}\geq\rho(0)$. As a result, $\rho(1)=\min\left\{\rho\in[\frac{1}{\sigma_0^2},\infty):\rho+\frac{1}{\frac{1}{a_{max}}+\gamma(\rho)}\geq\rho(0)\right\}$ is well-defined. If $\rho(1)>\frac{1}{\sigma_0^2}$, we have $\rho(1)+\frac{1}{\frac{1}{a_{max}}+\gamma(\rho(1))}=\rho(0)$.

We now construct an MPE starting from any $\rho\in[\rho(1),\infty)$. Consider the following strategy profile: For any $\rho\in[\rho(1),\rho(0)]$, the platform sets $\gamma(\rho)$ that solves (25) with equality. The consumer chooses a_{max} if $\gamma\geq\gamma(\rho)$ and $\rho+\frac{1}{\frac{1}{a_{max}}+\gamma}\geq\rho(0)$. If $\gamma<\gamma(\rho)$, she chooses $a=0$.

If $\gamma \geq \gamma(\rho)$ but $\rho' = \rho + \frac{1}{\frac{1}{a_{max}}+\gamma} < \rho(0)$, she chooses some optimal activity level, taking the continuation values $V(\rho, \gamma(\rho))$ (after $a = 0$) and $V(\rho', \gamma(\rho'))$ (after $a = a_{max}$) as given. Once the state reaches $\rho \geq \rho(0)$, the MPE for $\rho \geq \rho(0)$ is played—i.e., the platform sets $\gamma = 0$ and the consumer chooses a_{max} after any history. This strategy profile is an MPE: First, by construction, the consumer has no profitable one-shot deviation after any history. Second, the platform does not benefit from any one-shot deviation: If it increases γ, the deviation decreases the precision in the current and any future periods compared to without deviation. If it decreases γ, the consumer chooses $a = 0$ and the deviation decreases the precision in the current and any future periods, compared to without deviation.

For each $\rho \geq \rho_0$, define

$$V(\rho) = \begin{cases} -\frac{v}{1-\delta}\left(\sigma_0^2 - \frac{1}{\rho}\right) & \text{if } \rho \leq \rho(0), \\ \sum_{t=1}^{\infty} \delta^{t-1}\left[u(a_{max}) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho + t a_{max}}\right)\right] & \text{if } \rho(0) \leq \rho. \end{cases} \tag{27}$$

For $\rho \geq \rho(1)$, $V(\rho)$ is the consumer's continuation value in the above MPE. The value function $V(\rho)$ is decreasing, convex, and continuous (but not differentiable at $\rho = \rho(0)$). We now construct an MPE starting from any $\rho \in [\rho(2), \infty)$, where $\rho(2) < \rho(1)$. For each (ρ, γ) such that $\rho \leq \rho(1)$, define

$$V_2(\rho, \gamma) := u(a_{max}) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a_{max}}+\gamma}}\right) + \delta V\left(\rho + \frac{1}{\frac{1}{a_{max}} + \gamma}\right).$$

Consider the inequality

$$V_2(\rho, \gamma) \geq -v\left(\sigma_0^2 - \frac{1}{\rho}\right) + \delta V_2(\rho, \gamma). \tag{28}$$

For each $\rho < \rho(1)$, we consider the smallest γ that satisfies (28). Note that if we fix ρ and take γ that satisfies (28), $\rho + \frac{1}{\frac{1}{a_{max}}+\gamma} \leq \rho(0)$ holds; otherwise, it contradicts the definition of $\rho(1)$. As a result, (28) is equivalent to

$$u(a_{max}) - \frac{v}{1-\delta}\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a_{max}}+\gamma}}\right) \geq -v\left(\sigma_0^2 - \frac{1}{\rho}\right) + \delta\left[u(a_{max}) - \frac{v}{1-\delta}\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a_{max}}+\gamma}}\right)\right] \tag{29}$$

$$\iff (1-\delta)u(a_{max}) + \frac{v}{\rho + \frac{1}{\frac{1}{a_{max}}+\gamma}} - \frac{v}{\rho} \geq 0.$$

The left-hand side is continuous and strictly increasing in γ, and it is positive for $\gamma = \infty$. It is also continuous and strictly increasing in ρ. As a result, for each $\rho < \rho(1)$, we can find a unique $\gamma(\rho) > 0$ such that (29) holds with equality. By construction, $\gamma(\rho)$ is decreasing. Define $\rho(2) = \min\left\{\rho \in [\frac{1}{\sigma_0^2}, \infty) : \rho + \frac{1}{\frac{1}{a_{max}}+\gamma(\rho)} \geq \rho(1)\right\}$. If $\rho(2) > \frac{1}{\sigma_0^2}$, then $\rho(2) + \frac{1}{\frac{1}{a_{max}}+\gamma(\rho(2))} = \rho(1)$.

We can then construct a Markov perfect equilibrium for any initial state in $[\rho(2), \infty)$. For any $\rho \in [\rho(2), \rho(1)]$, the platform sets $\gamma(\rho)$ that solves (29) with equality. The consumer chooses a_{max} if $\gamma \geq \gamma(\rho)$ and $\rho + \frac{1}{\frac{1}{a_{max}}+\gamma} \geq \rho(1)$. If $\gamma < \gamma(\rho)$, she chooses $a = 0$. If $\gamma \geq \gamma(\rho)$ but $\rho + \frac{1}{\frac{1}{a_{max}}+\gamma} < \rho(1)$, she chooses some optimal activity level, taking the relevant continuation values as given. Once the state reaches $\rho \geq \rho(1)$, the MPE for $\rho \geq \rho(1)$ is played. We can show that this is an MPE by the same argument as the case of $\rho \in [\rho(1), \rho(0)]$. In particular, the platform's deviation to $\gamma > \gamma(\rho)$ will uniformly increase the current and future ρ_t's.

Given the initial state $\rho \in [\rho(2), \rho(1)]$, the consumer's continuation value is $V(\rho) = -\frac{v}{1-\delta}\left(\sigma_0^2 - \frac{1}{\rho}\right)$, which is the same as that in the previous step. As a result, we can use the incentive constraint (29) to recursively construct a sequence $\rho(3), \rho(4), \ldots$ and an MPE for any $k \in \mathbb{N}$ and the initial state $\rho \in [\rho(k), \rho(k-1)]$. The smallest ρ we consider is $\rho_0 = \frac{1}{\sigma_0^2}$. Thus, $\frac{1}{\frac{1}{a_{max}}+\gamma(\rho)} \geq \frac{1}{\frac{1}{a_{max}}+\gamma(\rho_0)}$ for any $\rho \geq \rho_0$. As a result, $\rho(k) - \rho(k+1) \geq \frac{1}{\frac{1}{a_{max}}+\gamma(\rho_0)} > 0$ for any $\rho \geq \rho_0$, whenever $\rho(k) > \rho_0$. Thus there is a smallest finite $K^* \in \mathbb{N}$ such that $\rho(K^*) < \rho_0$. Redefine $\rho(K^*)$ as ρ_0. We now have an MPE starting from $\rho = \rho_0$.

Step 2: Consumer-worst and platform-best. Let U_0 denote the hypothetical payoff of the consumer, when she acts optimally against the platform that commits to zero privacy levels in all periods. I show that the consumer's payoff is U_0 in the above MPE. If $\rho_0 \geq \rho(0)$, the platform sets $\gamma_t = 0$ for all t, and thus the consumer obtains U_0. If $\rho_0 < \rho(0)$, the consumer is indifferent between following the equilibrium strategy and choosing $a_t = 0$ for all t, because of the binding incentive constraint (25) or (29). Now, if the platform committed to zero privacy levels for $\rho < \rho(0)$, the consumer would choose $a_t = 0$ for all t. As a result, the consumer's ex ante payoff is $U_0 = 0$. In any equilibrium, the consumer's payoff cannot be strictly lower than U_0. Therefore, the above equilibrium is consumer-worst.

To show the equilibrium is platform-best under a common discount factor, let Π denote the platform's ex ante sum of discounted payoffs. If there is another equilibrium in which the platform obtains $\Pi' > \Pi$, the consumer's payoff is at most $\frac{u(a_{max})}{1-\delta} - v\Pi' < \frac{u(a_{max})}{1-\delta} - v\Pi = U_0$. This is a contradiction. Thus, we have shown the first part of Point 1 (we will show the second part at the end). Finally, even if the platform can commit to any rule to set privacy levels, the consumer can secure U_0 by acting as if privacy levels are zero. Thus, the platform cannot attain a strictly greater payoff even if it has a stronger commitment power.

Step 3: Other properties of the equilibrium. We show that $\gamma(\rho)$ is decreasing in ρ. First, $\gamma(\rho)$ is decreasing on $\rho \leq \rho(1)$, because $\gamma(\rho)$ is determined by the binding (29). Second, $\gamma(\rho)$ is decreasing on $[\rho(1), \rho(0)]$, because it is determined by the binding (25). Third, $\gamma(\rho) = 0$ for all $\rho \geq \rho(0)$. These observations, combined with the continuity of $\gamma(\rho)$, imply $\gamma(\rho)$ is decreasing. From period t to $t+1$, the state increases by $\rho_{t+1} - \rho_t = \frac{1}{\frac{1}{a_{max}} + \gamma(\rho_t)} \geq \frac{1}{\frac{1}{a_{max}} + \gamma(\rho_0)}$. Thus ρ_t is strictly increasing in t and diverges to $+\infty$ (or equivalently, $\sigma_t^2 \to 0$ in equilibrium). As a result, γ_t is strictly decreasing in equilibrium and hits zero in period T, which is the smallest T with $\rho_{T-1} \geq \rho(0)$. We now have Points 2. Also, I constructed the above MPE so that for any σ_{t-1}^2, the platform chooses the lowest γ_t that induces a_{max}. Such behavior is equivalent to a greedy policy. Thus, Point 3 holds.

Finally we prove the second part of Point 1. Let $(\gamma_t^*)_{t\in\mathbb{N}}$ denote the (on-path) equilibrium privacy levels in the above MPE. I show that if the platform commits to $(\gamma_t^*)_{t\in\mathbb{N}}$ ex ante, the consumer chooses a_{max} in all periods. To see this, we compare (i) the consumer's (single-agent) decision problem given $(\gamma_t^*)_{t\in\mathbb{N}}$ under long-run commitment to (ii) her problem given the platform's Markov strategy under one-period commitment. Take any strategy of the consumer, and consider the privacy level in period t. In (i), the consumer faces γ_t^*. In (ii), the consumer faces γ_{n+1}^*, where n is how many times the consumer chose $a = a_{max}$ instead of $a = 0$ before (and including) period $t-1$. We have $\gamma_{n+1}^* \geq \gamma_t^*$ after any history. Thus for any strategy, the consumer faces lower privacy levels in all periods under long-run commitment than one-period commitment. As a result, the consumer's optimal payoff under the former cannot exceed the one under the latter. Now, the consumer's optimal strategy under one-period commitment is $a_t = a_{max}$ for all $t \in \mathbb{N}$. She can achieve the same outcome under long-run commitment by choosing $a_t = a_{max}$ for all $t \in \mathbb{N}$. As a result, the consumer prefers $a_t = a_{max}$ for all t under long-run commitment.

We have shown that if the platform commits to $(\gamma_t^*)_{t\in\mathbb{N}}$ under long-run commitment, the consumer chooses a_{max} in all periods and obtains U_0 defined in Step 2. The same argument as Step 2 implies that the platform's optimal policy is $(\gamma_t^*)_{t\in\mathbb{N}}$ even under long-run commitment. $\square$

H Consumer-Best Equilibrium: Proof of Theorem 3

Proof. We write $\delta_C = \delta \geq 1/2$. Following the proof of Theorem 2, we write a Markov strategy of each player as a function of a precision $\rho_t = \frac{1}{\sigma_t^2}$. Let $\rho_0 = \frac{1}{\sigma_0^2}$. Define the strategy profile as follows: Let $\gamma(\rho_0) = \infty$. For any $\rho > \rho_0$, let $\gamma(\rho)$ be the strategy in the consumer-worst MPE in Theorem 2. Let $a(\rho_0, \infty) = a_{max}$, and $a(\rho_0, \gamma) = 0$ for any $\gamma < \infty$. For any $\rho > \rho_0$, let $a(\rho, \gamma)$ be her strategy in the MPE in Theorem 2. On the path of play, $(\gamma_t, a_t) = (\infty, a_{max})$ is chosen in all periods. This outcome is best for the consumer and worst for the platform.

Given the above strategy profile, suppose the platform deviates and offers $\gamma < \infty$ at $\rho = \rho_0$. If the consumer chooses $a = 0$, her future continuation value is $\frac{1}{1-\delta}u(a_{max})$, which is her best possible outcome. As a result, a necessary condition for the consumer to choose a_{max} following the platform's deviation at ρ_0 is that she obtains a nonnegative payoff in the current period:

$$u(a_{max}) - v\left(\frac{1}{\rho_0} - \frac{1}{\rho_0 + \frac{1}{\frac{1}{a_{max}}+\gamma}}\right) = u(a_{max}) - v\frac{\frac{1}{\frac{1}{a_{max}}+\gamma}}{\rho_0\left(\rho_0 + \frac{1}{\frac{1}{a_{max}}+\gamma}\right)} \geq 0. \tag{30}$$

Let $\hat{\gamma}(\rho_0)$ denote the minimum γ that satisfies this constraint. $\hat{\gamma}(\rho_0)$ is decreasing in ρ_0, positive for a small ρ_0, and $\lim_{\rho_0 \to 0} \hat{\gamma}(\rho_0) = \infty$.

Take any $\bar{\rho} > 0$ such that $\bar{\rho} + \frac{1}{\frac{1}{a_{max}}+\hat{\gamma}(\bar{\rho})} \leq \rho(0)$, where $\rho(0)$ is the cutoff constructed for Theorem 2, above which $(\gamma, a) = (0, a_{max})$ is chosen. For any initial state $\rho_0 \leq \bar{\rho}$, the above strategy profile is an equilibrium. First, it is an equilibrium at any (off-path) state $\rho > \rho_0$ by construction. At $\rho = \rho_0$, the consumer has no profitable deviation when the platform offers $\gamma = \infty$, because she can receive the best payoff of $u(a_{max})$ in the current and any future periods. Suppose that the platform deviates and chooses $\gamma_t < \infty$. Suppose to the contrary that the consumer strictly benefits from the one-shot deviation to $a = a_{max}$. Then, $\rho_0 + \frac{1}{\frac{1}{a_{max}}+\gamma} \le \rho(0)$ must hold. Thus her payoff in period t is at most $u(a_{max})$, whereas her continuation value from period $t+1$ is nonpositive (recall that in the consumer-worst equilibrium, the consumer's continuation payoff starting from $\rho \le \rho(0)$ is nonpositive). In contrast, if the consumer chooses $a_t = 0$ and follows her strategy thereafter, her payoff is $\frac{\delta}{1-\delta}u(a_{max})$, because she sets $a_t = 0$ in period t and the state remains ρ_0. Thus, the consumer has a profitable deviation only if $\frac{\delta}{1-\delta}u(a_{max}) < u(a_{max})$, which contradicts $\delta \ge 1/2$. □

I An MPE for a General A: Proof of Proposition 4

For simplicity we write δ_C as δ.

Lemma 11. *If the platform sets $\bar{\gamma}$ in Assumption 2 in period t, the consumer strictly prefers (i) $a_t = a_{max}$ and $a_s = 0$ for all $s \ge t+1$ to (ii) $a_s = 0$ for all $s \ge t$, regardless of the platform's continuation strategy.*

Proof. Define $\rho_{t-1} = \frac{1}{\sigma^2_{t-1}}$. The consumer prefers (i) to (ii) if and only if

$$u(a_{max}) - \frac{v}{1-\delta}\left(\sigma_0^2 - \frac{1}{\rho_{t-1} + \frac{1}{\frac{1}{a_{max}}+\bar{\gamma}}}\right) \ge -\frac{v}{1-\delta}\left(\sigma_0^2 - \frac{1}{\rho_{t-1}}\right)$$

$$\iff u(a_{max}) - \frac{v}{1-\delta}\left[\frac{1}{\rho_{t-1}\left(\rho_{t-1}\left(\frac{1}{a_{max}}+\bar{\gamma}\right)+1\right)}\right] \ge 0. \tag{31}$$

The left-hand side of the last inequality is at least

$$H := u(a_{max}) - \frac{v}{1-\delta}\left[\frac{1}{\rho_0\left(\rho_0\left(\frac{1}{a_{max}}+\bar{\gamma}\right)+1\right)}\right].$$

The inequality $H > 0$ is equivalent to the one for $\bar{\gamma}$ in Assumption 2. □

Proof of Proposition 4. Let a^+ denote the smallest positive activity level in A, and let γ^+ denote the highest finite privacy level in Γ. Define $\Delta^* := \frac{1}{\frac{1}{a_+}+\gamma+}$. Proposition 3 implies that there is $\rho(0)$

such that if the initial state is above $\rho(0)$, then $(\gamma_t, a_t) = (0, a_{max})$ for all $t \in \mathbb{N}$ is an MPE. Let $V_0(\cdot) : [\rho(0), \infty) \to \mathbb{R}$ and $\Pi_0(\cdot) : [\rho(0), \infty) \to \mathbb{R}$ respectively denote the consumer's and the platform's continuation values in that MPE. We extend these functions so that $V_0(\rho) = \Pi_0(\rho) = -\infty$ for $\rho < \rho(0)$. Note that $\Pi_0(\cdot)$ is increasing. Also, define $\rho(1) := \rho(0) - \Delta^*$. Finally, let $A_+ = A \setminus \{0\}$ denote the set of all positive activity levels. For any $\rho \in [\rho(1), \rho(0)]$, consider the optimization problem

$$\Pi_1(\rho) := \max_{\gamma \in \Gamma, a(\rho,\gamma) \in A} \sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a(\rho,\gamma)} + \gamma}} + \Pi_0\left(\rho + \frac{1}{\frac{1}{a(\rho,\gamma)} + \gamma}\right) \tag{32}$$

$$s.t. \quad a(\rho,\gamma) \in \arg\max_{a \in A_+} u(a) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a}+\gamma}}\right) + \delta V_0\left(\rho + \frac{1}{\frac{1}{a} + \gamma}\right), \quad \text{and} \tag{33}$$

$$\begin{aligned} &u(a(\rho,\gamma)) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a(\rho,\gamma)}+\gamma}}\right) + \delta V_0\left(\rho + \frac{1}{\frac{1}{a(\rho,\gamma)} + \gamma}\right) \\ &\geq -v\left(\sigma_0^2 - \frac{1}{\rho}\right) + \delta \cdot \left[u(a(\rho,\gamma)) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a(\rho,\gamma)}+\gamma}}\right) + \delta V_0\left(\rho + \frac{1}{\frac{1}{a(\rho,\gamma)} + \gamma}\right)\right]. \end{aligned} \tag{34}$$

Let $\bar{\gamma}$ denote the privacy level in Lemma 11. First, we show that there is $(\gamma, a(\rho,\gamma)) = (\gamma^*, a^*)$ that satisfies the constraints. Take $\gamma^* = \bar{\gamma}$, and let $a(\rho, \gamma^*) = a^*$ denote the solution of (33). Suppose, to the contrary, that (34) fails, i.e., we obtain

$$\begin{aligned} &u(a^*) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a^*}+\gamma^*}}\right) + \delta V_0\left(\rho + \frac{1}{\frac{1}{a^*} + \gamma^*}\right) \\ &< -v\left(\sigma_0^2 - \frac{1}{\rho}\right) + \delta\left[u(a^*) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a^*}+\gamma^*}}\right) + \delta V_0\left(\rho + \frac{1}{\frac{1}{a(\rho,\gamma)} + \gamma}\right)\right]. \end{aligned}$$

This inequality implies

$$
\begin{aligned}
-\frac{v}{1-\delta}\left(\sigma_0^2-\frac{1}{\rho}\right) &> u(a^*)-v\left(\sigma_0^2-\frac{1}{\rho+\frac{1}{\frac{1}{a^*}+\gamma^*}}\right)+\delta V_0\left(\rho+\frac{1}{\frac{1}{a^*}+\gamma^*}\right)\\
&\geq u(a_{max})-v\left(\sigma_0^2-\frac{1}{\rho+\frac{1}{\frac{1}{a_{max}}+\gamma^*}}\right)+\delta V_0\left(\rho+\frac{1}{\frac{1}{a_{max}}+\gamma^*}\right)\\
&\geq u(a_{max})-\frac{v}{1-\delta}\left(\sigma_0^2-\frac{1}{\rho+\frac{1}{\frac{1}{a_{max}}+\gamma^*}}\right),
\end{aligned}
$$

which contradicts the definition of $\gamma^*=\bar{\gamma}$ in Lemma 11. Let $(\gamma(\rho), a(\rho,\gamma(\rho)))$ denote the solution of the above problem. Note that $\gamma(\rho)<\infty$ and $a(\rho,\gamma(\rho))>0$. Let $V_1(\cdot):[\rho(1),\infty)\to\mathbb{R}$ denote the extension of $V_0(\cdot)$ such that for all $\rho\in[\rho(1),\rho(0)]$,

$$
V_1(\rho)=u(a(\rho,\gamma(\rho)))-v\left(\sigma_0^2-\frac{1}{\rho+\frac{1}{\frac{1}{a(\rho,\gamma(\rho))}+\gamma(\rho)}}\right)+\delta V_0\left(\rho+\frac{1}{\frac{1}{a(\rho,\gamma(\rho))}+\gamma(\rho)}\right). \tag{35}
$$

Let $\Pi_1(\cdot):[\rho(1),\infty)\to\mathbb{R}$ denote the extension of $\Pi_0(\cdot)$ such that for all $\rho\in[\rho(1),\rho(0)]$,

$$
\Pi_1(\rho)=\sigma_0^2-\frac{1}{\rho+\frac{1}{\frac{1}{a(\rho,\gamma(\rho))}+\gamma(\rho)}}+\delta\Pi_0\left(\rho+\frac{1}{\frac{1}{a(\rho,\gamma(\rho))}+\gamma(\rho)}\right). \tag{36}
$$

Now, suppose that we have constructed $(\Pi_{n-1}(\cdot), V_{n-1}(\cdot))$ defined on $[\rho(n-1),\infty)$. Extend these functions by setting $\Pi_{n-1}(\rho)=V_{n-1}(\rho)=-\infty$ for any $\rho<\rho(n-1)$, then consider the following problem:

$$\Pi_n(\rho) := \max_{\gamma\in\Gamma, a(\rho,\gamma)\in A} \sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a(\rho,\gamma)}+\gamma}} + \Pi_{n-1}\left(\rho + \frac{1}{\frac{1}{a(\rho,\gamma)}+\gamma}\right) \tag{37}$$

$$s.t. \quad a(\rho,\gamma) \in \arg\max_{a\in A_+} u(a) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a}+\gamma}}\right) + \delta V_{n-1}\left(\rho + \frac{1}{\frac{1}{a}+\gamma}\right), \quad \text{and} \tag{38}$$

$$u(a(\rho,\gamma)) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a(\rho,\gamma)}+\gamma}}\right) + \delta V_{n-1}\left(\rho + \frac{1}{\frac{1}{a(\rho,\gamma)}+\gamma}\right)$$
$$\geq - v\left(\sigma_0^2 - \frac{1}{\rho}\right) + \delta \cdot \left[u(a(\rho,\gamma)) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a(\rho,\gamma)}+\gamma}}\right) + \delta V_{n-1}\left(\rho + \frac{1}{\frac{1}{a(\rho,\gamma)}+\gamma}\right)\right]. \tag{39}$$

By the same argument for $n = 1$, we can find a solution $(\gamma(\rho), a(\rho, \gamma(\rho)))$ for each $\rho \in [\rho(n), \rho(n-1)]$, where $\rho(n) = \rho(n-1) - \Delta^*$. We can then construct $V_n(\cdot)$ and $\Pi_n(\cdot)$, which are respectively the extensions of $V_{n-1}(\cdot)$ and $\Pi_{n-1}(\cdot)$ to $[\rho(n), \infty)$. Repeating this, we can find a finite n such that $\rho(n) \leq \rho_0 = \frac{1}{\sigma_0^2}$. We now have a function $\gamma(\rho)$ defined on $[\rho_0, \infty)$. Also, for any $n \in \mathbb{N}$, $\rho \in [\rho(n), \rho(n-1)]$, and $\gamma \in \Gamma$, let $a(\rho, \gamma)$ denote the solution of (38).

We use $(\gamma(\rho), a(\rho, \gamma))$ to construct an MPE, $(\gamma^*(\rho), a^*(\rho, \gamma))$. First, set $\gamma^*(\cdot) \equiv \gamma(\cdot)$. Second, we define $a^*(\rho, \gamma)$. Take any $\rho < \rho(0)$ and $\gamma < \infty$, and let $n \in \mathbb{N}$ satisfy $\rho \in [\rho(n), \rho(n-1)]$. Let $a^*(\rho, \gamma) = a(\rho, \gamma)$ if

$$u(a(\rho,\gamma)) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a(\rho,\gamma)}+\gamma}}\right) + \delta V_{n-1}\left(\rho + \frac{1}{\frac{1}{a(\rho,\gamma)}+\gamma}\right) \geq -v\left(\sigma_0^2 - \frac{1}{\rho}\right) + \delta V_n(\rho).$$

If this inequality fails, then $a^*(\rho, \gamma) = 0$. If $\gamma = \infty$, then $a^*(\rho, \gamma) = a_{max}$. For $\rho \geq \rho(0)$, define $(\gamma^*(\rho), a^*(\rho, \cdot)) = (0, a_{max})$.

We show that $(\gamma^*(\rho), a^*(\rho, \gamma))$ is an MPE by showing that there is no profitable one-shot deviation. The optimality of $a^*(\rho, \gamma)$ holds by construction. The optimality of $\gamma^*(\rho)$ holds for the following reason. First, it is not optimal for the platform to set γ such that $a(\rho, \gamma) = 0$. Thus, facing $a^*(\rho, \gamma)$, any optimal strategy of the platform induces a positive activity level, i.e., it chooses γ such that $a(\rho, \gamma) \in A_+$. By (38), among such privacy levels, $\gamma(\rho)$ is optimal by (38). Finally, in each period, ρ_t increases by at least $\Delta^* > 0$ defined at the beginning. Once ρ_{T-1} exceeds $\rho(0)$, we have $(\gamma_t^*, a_t^*) = (0, a_{max})$ for all $t \geq T$. $\square$

J Equilibrium Under Competition: Proof of Proposition 5

Proof. First, we construct an equilibrium that satisfies Point 1. Suppose that, at the beginning of period $t \geq t^*$, the conditional variance for platform k is $\sigma^2_{t-1,k}$. Let γ_t^k denote the privacy level of platform k in period t. The (myopic) consumer weakly prefers to use platform k (i.e. $a_t^{-k} = 0$ maximizes her period-t payoff) if

$$\begin{aligned}
&\arg\max_{a \in A} u(a) - v[\sigma_0^2 - \sigma^2_{t,k}(\gamma_t^k, a | \sigma^2_{t-1,k})] - v[\sigma_0^2 - \sigma^2_{t-1,-k}] \\
\geq &\arg\max_{a \in A} u(a) - v[\sigma_0^2 - \sigma^2_{t,-k}(\gamma_t^{-k}, a | \sigma^2_{t-1,-k})] - v[\sigma_0^2 - \sigma^2_{t-1,k}],
\end{aligned}$$

where $\sigma^2_{t,k}(\gamma, a|\sigma^2_{t-1,k})$ is the posterior variance at the end of period t when platform k chooses γ, the consumer chooses a, and the posterior variance from the previous period is $\sigma^2_{t-1,k}$. Arranging this inequality, we obtain

$$\arg\max_{a \in A} u(a) - v[\sigma^2_{t-1,k} - \sigma^2_{t,k}(\gamma_t^k, a|\sigma^2_{t-1,k})] \geq \arg\max_{a \in A}\ u(a) - v[\sigma^2_{t-1,-k} - \sigma^2_{t,-k}(\gamma_t^{-k}, a|\sigma^2_{t-1,-k})].$$

This inequality implies that the consumer prefers to use k if and only if the gross benefit from the service minus the incremental privacy cost is greater for k than $-k$.

First, I consider competition with one-period commitment. Consider the following strategy profile. For each period $t < t^*$, I chooses a monopoly privacy level γ_t^*. Take any period $t \geq t^*$. Let $k^* \in \arg\min_{k=I,E} \sigma^2_{t-1,k}$ denote the platform that has the lower posterior variance (if k^* is not unique, we set $k^* = I$). Then platform $-k^*$ chooses the highest privacy level $\bar{\gamma}$. Platform k^* chooses a privacy level $\gamma_t^{k^*}$ that solves

$$\begin{aligned}
&\min_{\gamma \in [0,\bar{\gamma}]} \frac{1}{a^*(\gamma, \sigma^2_{t-1,k^*})} + \gamma \\
\text{s.t.} \quad &\arg\max_{a \in A} u(a) - v[\sigma^2_{t-1,k^*} - \sigma^2_{t,k^*}(\gamma, a|\sigma^2_{t-1,k^*})] \\
&\geq \arg\max_{a \in A} u(a) - v[\sigma^2_{t-1,-k^*} - \sigma^2_{t,-k^*}(\bar{\gamma}, a|\sigma^2_{t-1,-k^*})].
\end{aligned} \tag{40}$$

In each period, the consumer myopically chooses a_t^I (if $t < t^*$) or (a_t^I, a_t^E) (if $t \geq t^*$) to maximize her per-period payoff. If indifferent, she uses the platform for which she chose a positive activity level in the most recent period. (If she chose zero activity levels up to period $t-1$, then she sets $a_t^k = 0$ for one of $k \in \{I, E\}$ with equal probability, and chooses a_t^{-k} to maximize her period-t payoff.)

I show the above strategy profile is an equilibrium. First, the consumer's behavior is optimal by construction. Second, I verify that platforms have no profitable deviation. Without loss of generality, consider a node in period t in which $I = k^*$ and $E = -k^*$. The strategy of E is optimal: Suppose the consumer uses I in period t (i.e. $\sigma^2_{t-1,I} \leq \sigma^2_{t-1,E}$). By construction, even if E chooses $\bar{\gamma}$ in all periods $s \geq t$, the consumer uses I in any future periods as long as I and the consumer follow the above strategy. Thus, E's payoff does not change if E lowers privacy levels. Thus, E has no profitable deviation.

Suppose now that I chooses a privacy level such that the consumer chooses E in period t. If $\sigma^2_{t,E} \leq \sigma^2_{t,I}$, then the consumer uses E in any period $s \geq t+1$. In this case, I's deviation is not profitable. Otherwise, $\sigma^2_{t,E} > \sigma^2_{t,I}$ hold. Note that I obtains a lower payoff in period t, because it is not maximizing the informativeness of the signal. Moreover, at any future period s, I faces an optimization problem

$$\begin{aligned} &\min_{\gamma} \frac{1}{a^*(\gamma, \sigma^2_{s-1,I})} + \gamma \\ \text{s.t.} \quad &\arg\max_{a \in A} u(a) - v[\sigma^2_{s-1,I} - \sigma^2_{s,I}(\gamma, a | \sigma^2_{s-1,I})] \\ &\geq \arg\max_{a \in A} u(a) - v[\sigma^2_{s-1,E} - \sigma^2_{s,E}(\bar{\gamma}, a | \sigma^2_{s-1,E})]. \end{aligned} \tag{41}$$

After deviation, I faces a strictly lower $\sigma^2_{s-1,E} - \sigma^2_{s,E}(\bar{\gamma}, a|\sigma^2_{s-1,E}) > 0$ because the consumer generated information on E in period t. This means the set of γ satisfying the constraint shrinks. Thus, the minimized value in (41) becomes greater for any period $s \geq t+1$ after deviation. This implies that I's payoff is weakly lower for any period $s \geq t$ after the deviation. A similar argument implies that it is not profitable for I to deviate from a monopoly strategy before entry, because the deviation lowers I's payoff before and after entry. In particular, the deviation shrinks the set of γ's satisfying the constraint in (41) by increasing $\sigma^2_{s-1,I} - \sigma^2_{s,I}(\gamma, a|\sigma^2_{s-1,I})$.

On the equilibrium path, $a_t^E = 0$ for all $t \in \mathbb{N}$. $\lim_{t\to\infty} \sigma_{I,t}^2 = 0$ holds because it holds even if I adopts $\gamma_t = \bar{\gamma}$ for all t, and I chooses each γ_t^I to achieve even lower posterior variances. Given this result, $\lim_{t\to\infty} a_t^I = a_{max}$ follows the same proof as monopoly.

Suppose γ_t^I does not converge to 0. Then, there is a convergent subsequence $\gamma_{t(n)}^I$ such that $\lim_{n\to\infty} \gamma_{t(n)}^I = \gamma' > 0$. For a sufficiently large n, both $\gamma = 0$ and $\gamma = \gamma_{t(n)}^I$ satisfy the constraint in (41), because $\sigma_{s-1,E}^2 - \sigma_{s,E}^2(\bar{\gamma}, a|\sigma_{s-1,E}^2) = \sigma_0^2 - \sigma_{1,E}^2(\bar{\gamma}, a^*(\bar{\gamma}, \sigma_0^2)|\sigma_0^2) > 0$, but $\lim_{s\to\infty} \sigma_{s-1,I}^2 - \sigma_{s,I}^2(0, a^*(0, \sigma_{s-1,I}^2)|\sigma_{s-1,I}^2) \leq \lim_{s\to\infty} \sigma_{s-1,I}^2 = 0$. As $n \to \infty$, the value of the objective converges to $\frac{1}{a_{max}}$ and $\frac{1}{a_{max}} + \gamma'$ for $\gamma = 0$ and $\gamma = \gamma'$, respectively. Thus, for a large n, $\gamma = 0$ achieves a strictly lower value in (41) than $\gamma = \gamma'$. This is a contradiction and thus $\lim_{t\to\infty} \gamma_t^I \to 0$ in the equilibrium.

Next, we show Point 2. For a sufficiently large t^*, $\sigma_{t^*-1,I}^2 \leq \sigma_0^2 - \sigma_{t^*,E}^2(\bar{\gamma}, a^*(\sigma_0^2, \bar{\gamma})|\sigma_0^2)$. Then, for any period $t \geq t^*$, the constraint (41) holds for any $\gamma \leq \bar{\gamma}$. This implies that in any equilibrium, I's problem is equal to the monopolist's problem, which proves Point 2.

A similar proof applies to competition with long-run commitment. In this game, I commits to $(\gamma_1^I, \gamma_2^I, \dots)$ before $t = 1$, then the consumer (myopically) chooses a_t^I for each $t < t^*$. At the beginning of t^*, E publicly commits to $(\gamma_{t^*}^E, \gamma_{t^*+1}^E, \dots)$, after which the consumer chooses (a_t^I, a_t^E) in each period $t \geq t^*$. Here, I consider an equilibrium in which E commits to $\gamma_t^E = \bar{\gamma}\ \forall t \geq t^*$, and I commits to monopoly privacy levels before t^* and sets privacy levels by recursively solving (41) after t^*. □

K Omitted Proofs for Section 7

K.1 Erasing Past Information: Proofs for Section 7.1

Proof of Claim 2. Since the consumer's action does not affect a privacy policy, it is optimal for the consumer to erase information in all periods. Anticipating this, the platform maximizes the amount of information generated in each period, by solving the problem (21) with $t = 1$. Thus the platform sets $\gamma_t = \gamma_1^*$ for all t. □

Proof of Claim 3. The platform's problem is to solve (21) by choosing a privacy level and whether to erase information. Whenever $\sigma_{t-1}^2 < \sigma_0^2$, erasing information strictly increases the posterior variance, increases the consumer's marginal cost, and shifts $a^*(\cdot, \sigma^2)$ downward. Because erasing information strictly lowers the platform's payoff, it chooses $\mathcal{T} = \emptyset$ in equilibrium. $\square$

K.2 Heterogeneous Consumers: Proof of Proposition 6

Proof. Take any equilibrium $(a_t^*(v), \sigma_t^2(v), \gamma_t^*)_{t\in\mathbb{N}, v\in V}$. For each $v \in V$, define $\sigma_\infty^2(v) := \lim_{t\to\infty} \sigma_t^2(v)$. First, suppose, to the contrary, that there is some $v^* \in V$ such that $\sigma_\infty^2(v^*) > 0$. Define

$$\Delta_t := \frac{1}{1-\delta_P} \sum_{v\in V} \alpha_v \left[\sigma_0^2 - \sigma_\infty^2(v)\right] - \frac{1}{1-\delta_P} \sum_{v\in V} \alpha_v \left[\sigma_0^2 - \sigma_{t-1}^2(v)\right]. \tag{42}$$

It holds $\lim_{t\to\infty} \Delta_t = 0$. Now, take any $\gamma_v^* \in \arg\min_\gamma \frac{1}{a^*(v^*,\gamma,\sigma_0^2)} + \gamma$. It holds that for any $\sigma^2 \in [\sigma_\infty^2(v^*), \sigma_0^2]$,

$$\sigma^2 - \frac{1}{\frac{1}{\sigma^2} + \frac{1}{\frac{1}{a^*(v^*,\gamma_v^*,\sigma^2)}+\gamma_v^*}} \geq \sigma^2 - \frac{1}{\frac{1}{\sigma^2} + \frac{1}{\frac{1}{a^*(v^*,\gamma_v^*,\sigma_0^2)}+\gamma_v^*}} \geq M := \min_{\sigma^2\in[\sigma_\infty^2(v^*),\sigma_0^2]} \sigma^2 - \frac{1}{\frac{1}{\sigma^2} + \frac{1}{\frac{1}{a^*(v^*,\gamma_v^*,\sigma_0^2)}+\gamma_v^*}} > 0.$$

The first inequality follows from $a^*(v^*, \gamma, \sigma_0^2) \leq a^*(v, \gamma, \sigma^2)$ for $\sigma^2 \leq \sigma_0^2$. The last inequality holds because the minimand is continuous and positive on $[\sigma_\infty^2(v^*), \sigma_0^2]$. For a sufficiently large t, we obtain $\frac{\alpha_v M}{1-\delta_P} > \Delta_t$, or equivalently,

$$\frac{\alpha_v M}{1-\delta_P} + \frac{1}{1-\delta_P} \sum_{v\in V} \alpha_v \left[\sigma_0^2 - \sigma_{t-1}^2(v)\right] > \frac{1}{1-\delta_P} \sum_{v\in V} \alpha_v \left[\sigma_0^2 - \sigma_\infty^2(v)\right].$$

The left hand side is the lower bound of the time-t continuation value that the platform can get by deviating to the privacy level γ_v^* from time t on. The right hand side is the upper bound of the time-t continuation value without deviation. Thus, the platform is strictly better off by committing to a privacy policy that sets γ_v^* from time t on. This is a contradiction. $\lim_{t\to\infty} a_t^*(v) = 0$ and $\lim_{\to\infty} \gamma_t^* = 0$ follow the proof of Theorem 1. $\square$

L Relaxing "The Privacy Cost is Sunk"

In the baseline model, the consumer incurs a privacy cost of $-v(\sigma_0^2 - \sigma_{t-1}^2)$ even if she chooses $a_t = 0$. Suppose now that the consumer incurs a fraction $\alpha \in [0, 1)$ of the privacy cost when $a_t = 0$. Namely, if $a_t > 0$, her payoff is $u(a_t) - v\left(\sigma_0^2 - \sigma_t^2\right)$. If $a_t = 0$, it is $-\alpha v \left(\sigma_0^2 - \sigma_{t-1}^2\right)$.

The main results under monopoly and competition continue to hold for α close to 1. The following considers monopoly.

Proposition 7. *There is an $\alpha^* < 1$ such that for any $\alpha \in [\alpha^*, 1]$, any equilibrium outcome $(a_t^*, \gamma_t^*, \sigma_t^2)_{t\in\mathbb{N}}$ satisfies $\lim_{t\to\infty} a_t^* = a_{max}$, $\lim_{t\to\infty} \gamma_t^* = 0$, and $\lim_{t\to\infty} \sigma_t^2 = 0$.*

Proof. Consider any equilibrium. In period t, the consumer chooses a positive activity level if

$$\max_{a\in A} u(a) - v\left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a}+\gamma_t^*}}\right) \geq -\alpha v\left(\sigma_0^2 - \sigma_{t-1}^2\right)$$
$$\iff \max_{a\in A} u(a) - v\left(\alpha\sigma_{t-1}^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a}+\gamma_t^*}}\right) \geq (1-\alpha)v\sigma_0^2.$$

Let $\hat{a}_1$ and $\hat{\gamma}_1$ denote the equilibrium activity level and privacy level, respectively, in $t = 1$ of the baseline model (i.e., $\alpha = 1$). Define $y_1 := \frac{1}{\hat{a}} + \hat{\gamma}$ and $f(\alpha, x, y) := \alpha x - \frac{1}{\frac{1}{x}+\frac{1}{y}}$. The function f is strictly convex in x. Thus, on the interval $[0, \sigma_0^2]$, $f(\alpha, \cdot, y)$ is maximized at $x = \sigma_0^2$ if $f(\alpha, \sigma_0^2, y) > f(\sigma, 0, y)$, or equivalently, $\alpha\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2}+\frac{1}{y}} > 0$. Moreover, the left hand side is decreasing in y. Thus, this inequality holds for all $y \leq y_1$ if and only if $\alpha\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2}+\frac{1}{y_1}} > 0$. Let $\alpha^* < 1$ satisfy $\alpha^*\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2}+\frac{1}{\frac{1}{\hat{a}_1}+\hat{\gamma}_1}} > 0$. For any $\alpha \in [\alpha^*, 1]$, we have

$$u(\hat{a}_1) - v\left(\alpha\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \frac{1}{\frac{1}{\hat{a}_1}+\hat{\gamma}_1}}\right) \geq (1-\alpha)v\sigma_0^2$$
$$\Rightarrow u(\hat{a}_1) - v\left(\alpha\sigma_{t-1}^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{\hat{a}_1}+\hat{\gamma}_1}}\right) \geq (1-\alpha)v\sigma_0^2$$
$$\Rightarrow \max_{a\in A} u(a) - v\left(\alpha\sigma_{t-1}^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a}+\hat{\gamma}_1}}\right) \geq (1-\alpha)v\sigma_0^2.$$

The first inequality holds because it is independent of α' and holds for $\alpha' = 1$. The last inequality implies that in any period, if the platform sets $\gamma_t = \hat{\gamma}_1$, then the consumer chooses $a_t > 0$. Also $a_t \geq \hat{a}_1$ holds because $\gamma_t > \hat{\gamma}$ and $\sigma_{t-1}^2 \leq \sigma_0^2$. In equilibrium, the platform sets γ_t to minimize the variance of the noise in s_t subject to the constraint that

$$\max_{a\in A} u(a) - v\left(\alpha\sigma_{t-1}^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a}+\gamma_t}}\right) \geq (1-\alpha)v\sigma_0^2.$$

The above argument implies that the variance of the noise in s_t is at most $\frac{1}{\hat{a}_1} + \hat{\gamma} + \varepsilon$, which implies $\sigma_t^2 \to 0$ in equilibrium. By the same proof as Theorem 1, $\sigma_t^2 \to 0$ implies $a_t^* \to a_{max}$ and $\gamma_t^* \to 0$. □

The following considers competition.

Proposition 8. *There is an $\alpha^{**} < 1$ such that for any $\alpha \in [\alpha^{**}, 1]$, the result under competition (Proposition 5) holds.*

Proof. I adopt the notations in the proof of Proposition 5. In any period, the consumer weakly prefers to use platform k (i.e., $a_t^k > 0$ and $a_t^{-k} = 0$) if the following two conditions hold:

$$\begin{aligned}&\arg\max_{a\in A} u(a) - v[\sigma_0^2 - \sigma_{t,k}^2(\gamma_t^k, a|\sigma_{t-1,k}^2)] - \alpha v[\sigma_0^2 - \sigma_{t-1,-k}^2]\\ \geq &\arg\max_{a\in A} u(a) - v[\sigma_0^2 - \sigma_{t,-k}^2(\gamma_t^{-k}, a|\sigma_{t-1,-k}^2)] - \alpha v[\sigma_0^2 - \sigma_{t-1,k}^2],\end{aligned}$$

and

$$\arg\max_{a\in A} u(a) - v[\sigma_0^2 - \sigma_{t,k}^2(\gamma_t^k, a|\sigma_{t-1,k}^2)] - \alpha v[\sigma_0^2 - \sigma_{t-1,-k}^2] \geq -\alpha v[\sigma_0^2 - \sigma_{t-1,k}^2] - \alpha v[\sigma_0^2 - \sigma_{t-1,-k}^2].$$

These inequalities are respectively equivalent to

$$\arg\max_{a\in A} u(a) - \underbrace{v\left[\alpha\sigma_{t-1,k}^2 - \frac{1}{\frac{1}{\sigma_{t-1,k}^2} + \frac{1}{\frac{1}{a}+\gamma_t^k}}\right]}_{\text{(A)}} \geq \arg\max_{a\in A} u(a) - \underbrace{v\left[\alpha\sigma_{t-1,-k}^2 - \frac{1}{\frac{1}{\sigma_{t-1,-k}^2} + \frac{1}{\frac{1}{a}+\gamma_t^{-k}}}\right]}_{\text{(B)}} \tag{43}$$

and

$$\underset{a \in A}{\arg\max}\, u(a) - v \underbrace{\left[\alpha \sigma_{t-1,k}^2 - \frac{1}{\frac{1}{\sigma_{t-1,k}^2} + \frac{1}{\frac{1}{a} + \gamma_t^k}} \right]}_{\text{(A)}} \geq (1-\alpha) v \sigma_0^2. \tag{44}$$

By the same argument as Proposition 7, there is $\alpha^{**} < 1$ such that for any $\alpha \geq \alpha^{**}$, the following holds: For any $\frac{1}{a} + \gamma_t^k \leq \frac{1}{a(\bar{\gamma})} + \bar{\gamma}$, (A) is maximized at $\sigma_{t-1,k}^2 = \sigma_0^2$; for any $\frac{1}{a} + \gamma_t^{-k} \leq \frac{1}{a(\bar{\gamma})} + \bar{\gamma}$, (B) is maximized at $\sigma_{t-1,-k}^2 = \sigma_0^2$. These observations imply the following. First, I can induce $a_t^I > 0$ before the entry, by setting $\gamma_t = \bar{\gamma}$. Second, after I collects some information, if I and E set the same privacy level $\bar{\gamma}$, then the consumer optimally sets $a_t^I > 0 = a_t^E$. We can then apply the proof of Proposition 5 to construct an equilibrium such that (i) E sets $\gamma_t^E = \bar{\gamma}$ for all $t \in \mathbb{N}$, (ii) I sets γ_t^I to minimize the variance of the noise of s_t subject to constraints (43) and (44). The rest of the proof follows the proof of Proposition 5. $\square$