Dynamic Privacy Choices

Click here to download the latest version.

Shota Ichihashi*

May 27, 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 precise information about the consumer, thereby increasing platform profits. Although consumers value privacy, a platform is able to collect information by gradually lowering the level of privacy protection. In the long-run, consumers become "addicted" to the platform, whereby they lose privacy and receive low payoffs, but continue to choose high activity levels. Competition is unhelpful because consumers have a higher incentive to use a platform to which they have lower privacy.

*Bank of Canada, 234 Wellington Street West, Ottawa, ON K1A 0G9, Canada. Email: shotaichihashi@gmail.com. Website: http://shotaichi.weebly.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 Bank of Canada, Peking University, University of Rochester (Simon), Privacy Workshop at Princeton, and Canadian Economic Theory Conference 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. However, it also raises important concerns among consumers and policymakers (Crémer et al., 2019; Furman et al., 2019; Morton et al., 2019).

Consider a consumer (she) and a social media platform (it). The consumer writes posts, shares photos, 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, the consumer may value her privacy, or be concerned about the risk of data leakage, identity theft, and price or non-price discrimination.[1] The latter is a "privacy cost" of using the platform. If the consumer anticipates a high privacy cost, she may use the platform less actively, or may not join it in the first place. The platform can influence her decision through its privacy policy. For example, Facebook committed to not use cookies to track users.[2]

I formalize this story as a dynamic model: In each period, a 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 precision 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 per-period profit is increasing but the consumer's payoff is decreasing in the amount of information the platform has collected. Thus, the consumer chooses activity levels balancing the benefits from the service and the privacy costs. Anticipating the consumer's behavior, the platform sets privacy levels to maximize profits.

The key idea is that, when the consumer has less privacy, she faces a lower marginal privacy cost of using the same platform, repeatedly. For example, if Google already knows a lot about a consumer, she might not care about letting Google Maps track her location. In an extreme case, if the platform knows everything, then the marginal privacy cost is zero, because the consumer's

[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 May 19, 2020)

activity on a platform no longer affects what it knows about her.

The paper examines the dynamic implications of this idea. First, in equilibrium, the consumer chooses higher activity levels but receives lower payoffs over time. In the long-run, the consumer loses privacy and incurs a high privacy cost, but behaves as if there is no privacy cost. The privacy loss occurs even when the consumer values privacy highly. To induce such an outcome, the platform sets high privacy levels in early periods to incentivize the consumer to use the platform. However, the platform later reduces the privacy level to accelerate data collection. The baseline model assumes that the consumer is myopic and the platform can fully commit to future privacy levels. However, the result holds even when the consumer is forward-looking and the platform only has short-run commitment power (Section 6).

The decreasing marginal privacy cost implies that the consumer is more willing to use a platform to which she has less privacy. This consumer's tendency makes competition less effective. To see this, suppose that an incumbent (e.g., Google) has a lot of data on a consumer. Then, even if the entrant (e.g., DuckDuckGo) offers a better privacy protection, the consumer may stick to the incumbent since she incurs negligible marginal privacy cost of using the incumbent. As a result, the incumbent can keep collecting data without losing the consumer to the entrant.

I consider several privacy regulations. First, mandating the platform to adopt a strict privacy policy may perversely lower the privacy and welfare of consumers in the long-run. In contrast, the "right to be forgotten," which enables consumers to erase past information, may enhance consumer welfare and induce competition. Thus, ex ante and ex post privacy protections may have different impacts in a dynamic environment.

The paper has implications for consumer privacy. First, the consumer's long-run behavior seems consistent with the so-called privacy paradox: Consumers express concerns about data collection but actively share data with third parties (Acquisti et al., 2016). Second, the platform's equilibrium strategy rationalizes how online platforms, such as Facebook, seem to have relaxed its privacy policy over time. Third, the paper helps us understand why competition using better privacy policies might not be successful (Marthews and Tucker, 2019). Moreover, the results clarify how a privacy regulation such as the right to be forgotten can promote competition. The result also points to a challenge to firms that offer privacy-friendly alternatives to dominant platforms. Finally, the paper offers novel policy implications, such as restricting data collection hurting consumers.

This paper contributes to the literature on markets for data and the economics of privacy. First, the paper is 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 monetary compensation, and the data of some consumers reveal information about others. This "data externality" lowers the incentive of each consumer to protect privacy, leading to low compensation and excessive data sharing.[3] The economic force of this paper is similar to theirs: If a consumer provides data today, she has a lower incentive to protect privacy in the future. However, this paper considers a dynamic model, which enables me to study new issues. Specifically, I consider market entry, commitment to privacy policies, the right to be forgotten, data retention, consumer myopia, and the dynamics of privacy policies. My paper is also different from these papers in how data are generated. They assume that consumers hold data at the outset, and platforms buy data in exchange for monetary compensation. In contrast, in this paper, data arise as a byproduct of activity from which consumers derive utilities. This formulation directly applies to online platforms that do not pay consumers for data.

This paper also relates to the 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. My paper contributes to this literature by considering a less-explored setting in which data collection lowers consumer welfare. This assumption enables me to study issues related to consumer privacy. Hagiu and Wright (2020) allow price competition and consider a rich learning dynamics incorporating "within-user" and "across-user" learning. In contrast, I abstract from pricing, and focus on within-user learning.

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

[3] Relatedly, Easley et al. (2018) consider (positive) data externalities in a model where market transaction generates data.

are revealed, and they have lower incentives to engage in signal-jamming. Despite this connection, the two signal-jamming activities are quite 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 (i.e., 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 for not only increasing profit today but motivating the consumer to raise activity levels in the future. Many of my results come from this complementarity between data collection and higher activity levels. Finally, the analysis of competition and privacy regulations do not have counterparts in this literature.

Bonatti and Cisternas (2020) study consumer privacy in a continuous-time dynamic model. They consider a long-lived consumer with short-lived sellers. Sellers can learn about consumer preferences based on scores that aggregate purchase histories, and sellers use information for price discrimination. In contrast, I consider long-lived platforms and abstract away from how platforms use consumer information. Fainmesser et al. (2019) study the optimal design of a platform to store data and invest in information security. They consider a platform that cares about both the activity levels of consumers and the amount of data it can extract. They study how different objectives lead to different platform designs. I adopt simpler preferences for consumers and platforms, but consider a dynamic environment. Hörner and Skrzypacz (2016) considers a dynamic game of selling information. They show that the optimal selling strategy transmits information gradually. In my model, the benefit of gradualism appears in the platform's optimal privacy policy. Casadesus-Masanell and Hervas-Drane (2015) considers a static vertical differentiation framework in which firms compete in price and the level of privacy protection.

The rest of the paper is as follows. Section 2 presents a model of a monopoly platform and a myopic consumer. Section 3 presents the long-run equilibrium outcome and characterizes the equilibrium privacy policy of the platform. Section 4 considers platform competition. Section 5 studies the incentive of the consumer or platforms to erase past information. Section 6 considers a forward-looking consumer. Section 7 considers extensions, including consumer heterogeneity, time-varying types, and general payoff functions.

2 A Dynamic Model of Privacy Choices

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

In each period $t \in \mathbb{N}$, the consumer publicly chooses an *activity level* $a_t \geq 0$. After the consumer chooses a_t, the platform privately observes a signal $s_t = X + \varepsilon_t$, where $\varepsilon_t \sim \mathcal{N}\left(0, \frac{1}{a_t} + \gamma_t\right)$. A greater a_t reduces the variance of ε_t and makes s_t more informative about X, whereas a greater γ_t makes s_t less informative about X.[5] $\gamma_t \geq 0$ is the *privacy level* of the platform in period t. All the 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 $\boldsymbol{a}_t = (a_1, \ldots, a_t)$ and the platform has chosen privacy levels $\boldsymbol{\gamma}_t = (\gamma_1, \ldots, \gamma_t)$ up to period t. At the end of period t, the platform receives a payoff of $\sigma_0^2 - \sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t) \geq 0$. $\sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ is the variance of the conditional distribution of X given $(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$, derived from Bayes' rule.[6] I take $(\sigma_t^2(\cdot,\cdot))_{t\in\mathbb{N}}$ as a primitive. A lower $\sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ means that the platform has a more accurate estimate of X, which means that the consumer has less 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 a discount factor $\delta_P \in (0, 1)$.

The consumer's payoff in period t is $U(\boldsymbol{a}_t, \boldsymbol{\gamma}_t) := u(a_t) - v\left[\sigma_0^2 - \sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)\right]$. The first term $u(a_t)$ is the gross benefit of using the platform. $u : \mathbb{R}_+ \to \mathbb{R}$ is strictly concave, continuously differentiable, maximized at $a_{max} \in (0, \infty)$, and satisfies $u(0) = 0$. The second term $v\left[\sigma_0^2 - \sigma_t^2(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)\right]$ is a *privacy cost*, which captures the negative impact of data collection on the consumer. $v \in \mathbb{R}_+$ is the exogenous parameter that captures the value of privacy. The baseline model assumes that the consumer is myopic, i.e., she chooses a_t to maximize $U(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ in each period t. Section 6

[4]Even if the consumer privately observes X, all the results hold with respect to a pooling equilibrium in which the consumer of all types chooses the same activity level after any history. Unobservable X greatly simplifies exposition without changing the insights.

[5]If $a_t = 0$, then treat s_t as a pure noise that is independent of X.

[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 the further discussion.

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

shows that most of the results hold even if the consumer is forward-looking. However, the myopia assumption leads to cleaner results.

Payoffs are normalized so that if $a_t = 0$ for all t, then the platform and the consumer obtain zero payoffs in all periods. The primitives, σ_0^2, $u(\cdot)$, and v, are commonly known to the consumer and the platform (Section 7 assumes that the consumer privately knows v).

The timing of the game is as follows. Before $t = 1$, the platform commits to a *privacy policy* $\boldsymbol{\gamma} = (\gamma_1, \gamma_2, \dots) \in \mathbb{R}_+^\infty$. After observing $\boldsymbol{\gamma}$, the consumer chooses an activity level in each period. An *equilibrium* is a strategy profile such that (i) the consumer myopically chooses a_t to maximize $U(\boldsymbol{a}_t, \boldsymbol{\gamma}_t)$ following every history, breaking ties in favor of the highest activity level, and (ii) the platform, anticipating (i), optimally chooses a privacy policy $\boldsymbol{\gamma}$ before $t = 1$.

The model does not explicitly capture the consumer's decision to join the platform. However, we may say that the consumer joins in period t if t is the first period such that a_t is positive. We could extend the model so that the consumer incurs a one-time cost $\kappa > 0$ to join. The results continue to hold if κ is not too high.

2.1 Discussion of Modeling Assumptions

The model is not intended to capture all the details of how a platform collects and uses consumer data. Instead, I set up the model to highlight a set of conditions under which consumer privacy is difficult to sustain and data-driven platforms tend to be profitable. This subsection discusses which modeling assumptions are crucial for the main results.

2.1.1 Assumptions that are Crucial for the Results

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

Privacy cost function. The privacy cost $v(\sigma_0^2 - \sigma_t^2)$ captures monetary or non-monetary reasons for which a consumer wants a platform to have less information. For instance, a consumer may have intrinsic preferences for privacy (Kummer and Schulte, 2019; Lin, 2019; Tang, 2019). For another instance, a consumer may consider the risk of data breach, identity theft, and price or non-price discrimination by platforms and third parties. Section 7 shows that the main insight continues to hold for a privacy cost that is non-linear and non-monotonic in σ_t^2. In reality, the gross utility of using a platform may also depend on how much information the platform has. However, I do not consider such an extension.

Privacy cost is sunk. Even if the consumer chooses $a_s = 0$ for all $s \geq t$, she incurs a privacy cost of $-v(\sigma_0^2 - \sigma_t^2)$ in period s because $\sigma_s^2 = \sigma_t^2$. This observation is crucial: The consumer cannot delete data collected in the past, and thus she perceives the privacy cost from past data collection as sunk. This assumption reflects the difficulty of deleting digital data, which is referred to as "data persistence" (Tucker, 2018). For instance, suppose that a platform collects sensitive personal information and stores it even after users quit. Then, a consumer may face a risk of data leakage even when she is not active on the platform. For another instance, if a consumer inadvertently discloses some sensitive information to other users on a platform, then she may incur a psychological cost from the fact that other users know the information. Such a cost is likely to persist even if the consumer is inactive on the platform. Since the consumer regards the privacy cost as sunk, she chooses activity levels based on the *marginal* privacy cost rather than the level of privacy cost. This assumption also enables us to think what privacy regulation can improve consumer privacy by making the privacy cost non-sunk. I will later show that a regulation in line with "the right to be forgotten" achieves this goal. Finally, Section 7 considers an extension in which the consumer perceives only a fraction of the privacy cost as sunk.

Consumer's type is constant over time. This assumption implies that the consumer's activity today affects her welfare in the future though the privacy cost. The assumption is suitable if X is relatively persistent characteristics such as one's race and political affiliation. Section 7 presents a numerical analysis for an imperfectly persistent type and shows that the main insight extends.

2.1.2 Assumptions that are Not Crucial for the Results

Single consumer. I deliberately consider a single consumer to emphasize that the results do not rely on interactions among multiple consumers. Since the consumer's type is Gaussian, one could easily incorporate multiple consumers with "data externalities" as in Acemoglu et al. (2019) and

Bergemann et al. (2019) by assuming that their types are correlated. I conjecture that incorporating data externalities does not change the result. Indeed, the data externality reduces a consumer's incentive to protect privacy, but such an economic force already exists in the current model.

Consumer myopia. I consider a myopic consumer for two reasons. First, the myopia assumption captures the idea that, when a consumer uses a platform, she may not take into account how the data collected today may be misused against her in the future. This myopia assumption seems plausible because users may not necessarily foresee the future consequences of data sharing.[8] Second, the results under a monopoly platform continue to hold even if the consumer is arbitrarily patient (Section 6). Thus, I start with the myopia assumption to deliver clear intuitions.

Value of privacy v is commonly known. Section 7 shows that the results under a monopoly platform continue to hold even if the consumer is privately informed of v. In other words, the results hold when a population of consumers have heterogeneous v's, and the platform has to set a privacy policy that is common to all consumers.

Privacy policies. A platform either chooses a privacy policy at the outset (as in the baseline model) or sets a privacy level at the beginning of each period. They are not fully general. For example, we might consider a mechanism that maps each history of activity levels $(a_s)_{s=1}^t$ to a privacy level γ_t in period t. However, such a mechanism seems impractical and prohibitively costly to communicate to consumers, and it would not change the long-run outcome.

Platform's payoffs. The platform's payoff can be any decreasing function of $(\sigma_t^2)_{t\in\mathbb{N}}$. All the results and proofs continue to hold without modification (see Section 7 for details).

3 Monopoly Platform

I begin with studying the consumer's behavior, taking the platform's strategy as given. Then, I present the long-run equilibrium outcome. After that, I characterize the platform's equilibrium privacy policy.

[8] For example, Tucker (2018) states that "Introducing the potential for myopia or hyperbolic discounting into the way we model privacy choices over the creation of data seems, therefore, an important step."

3.1 Consumer Behavior

Bayes' rule implies[9]

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

Thus, in each period t, the consumer chooses a_t to maximize

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

taking $\sigma_{t-1}^2(\boldsymbol{a}_{t-1}, \boldsymbol{\gamma}_{t-1})$ and γ_t as given. 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 summarizes the key properties of the privacy cost C and the marginal privacy cost $\frac{\partial C}{\partial a}$.

Lemma 1. *The following holds.*

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

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

□

Lemma 1 implies that the privacy cost C increases but the marginal cost $\frac{\partial C}{\partial a}$ decreases when the consumer has less privacy (i.e., σ^2 is small). Thus, if the consumer has less privacy, her utility is low but the marginal utility of using the platform is high. Intuitively, once a platform has

[9] The equation holds because if $x|\mu \sim N(\mu, \sigma^2)$ and $\mu \sim N(\mu_0, \sigma_0^2)$, then $\mu|x \sim N\left(\frac{\sigma_0^2}{\sigma^2 + \sigma_0^2} x + \frac{\sigma^2}{\sigma^2 + \sigma_0^2}\mu_0, \left(\frac{1}{\sigma_0^2} + \frac{1}{\sigma^2}\right)^{-1}\right)$.

collected a lot of information about a consumer, then the consumer's activity today does not affect the platform's learning much, leading to a lower marginal privacy cost.

Another important observation from Lemma 1 is that the marginal privacy cost is decreasing in the privacy level. Thus, the platform can incentivize the consumer to choose a higher activity level by committing to add a noise to the signal. Thus, the platform may be able to improve the quality of a signal by committing to underuse it.

To state the next result, let $a^*(\gamma, \sigma^2)$ denote the optimal activity level given a privacy level γ in the current period and the conditional variance σ^2 from the previous period:

$$a^*(\gamma, \sigma^2) := \max \left\{ \arg\max_{a \geq 0} u(a) - v \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma^2} + \frac{1}{\frac{1}{a} + \gamma}} \right) \right\}. \tag{3}$$

The following result follows from Lemma 1 (see Appendix A for the proof). Recall $a_{max} = \arg\max_{a \geq 0} u(a)$.

Lemma 2. *$a^*(\gamma, \sigma^2)$ is increasing in γ and decreasing in σ^2. For any (γ, σ^2), $\lim_{\hat{\gamma} \to \infty} a^*(\hat{\gamma}, \sigma^2) = \lim_{\hat{\sigma}^2 \to 0} a^*(\gamma, \hat{\sigma}^2) = a_{max}$.*

The next result presents the equilibrium of a subgame in which the platform has committed to a stationary privacy policy (see Appendix B for the proof).

Proposition 1. *Suppose that the platform chooses a stationary privacy policy, i.e., $\gamma_t = \gamma, \forall t \in \mathbb{N}$. Let $(a_t^*)_{t \in \mathbb{N}}$ denote the equilibrium activity levels of this subgame. There is a cutoff value $v^*(\gamma) \in \mathbb{R}_+$ such that:*

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 per-period payoff decreases over time.*

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

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

The intuition is as follows. If the value of privacy is low, then the consumer prefers a positive activity level $a_1^* > 0$ in $t = 1$. The consumer activity generates information, which reduces her payoff but increases the marginal net benefit of using the platform. Thus, in $t = 2$, the consumer chooses $a_2^* \geq a_1^*$. Repeating this argument, we can conclude that a_t^* increases over time. The platform can then use the signals and 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 behaves as if there is no privacy cost. 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.

Proposition 1 highlights 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 can increase the consumer's per-period payoffs in the short-run but decrease them in the long-run. If the regulator cares about the long-run consumer welfare, it may consider the regulation as detrimental.[10]

3.2 Equilibrium

I now present the equilibrium of the entire game (see Appendix C for the proof, which uses Proposition 3 in the next subsection).

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

$$\lim_{t \to \infty} a_t^* = a_{max}, \ \lim_{t \to \infty} \gamma_t^* = 0, \ \textit{and} \ \lim_{t \to \infty} \sigma_t^2 = 0. \tag{4}$$

[10]The caveat "if the regulator cares about the long-run consumer welfare" is important, because regardless of the consumer's discount factor, a higher γ increases the consumer's *ex ante* sum of discounted payoffs (which is equal to the first-period payoff for a myopic consumer). A higher privacy level is undesirable only if the regulator has a higher discount factor than the consumer.

Moreover, for any $T \in \mathbb{N}$, there is a $\underline{v} \in \mathbb{R}$ such that, for any $v \geq \underline{v}$, any equilibrium privacy policy satisfies $\gamma_t^* > 0$ for all $t \leq T$.

This result contrasts with Proposition 1 in that the privacy loss occurs for any $v \in \mathbb{R}$. The result also shows that if v is high, the equilibrium privacy policy is nonstationary: In early periods, the platform sets positive privacy levels. In the long-run, the platform offers a vanishing level of privacy. The long-run payoff of the consumer is $u(a_{max}) - v\sigma_0^2$, which can be arbitrarily low for a large v.

The intuition is as follows. In early periods, the platform does not know much about the consumer. Then, consumer activity has a large impact on what the platform learns about her type. Thus, the consumer faces a high marginal privacy cost, which discourages her from raising the activity level. To reduce the marginal cost, the platform commits to a high privacy level and slowly learns the consumer type. After a long period of interaction, the platform accurately knows the consumer's type, and thus she faces a low marginal cost. Thus, the platform can lower a privacy level to accelerate learning.

Proposition 2 illustrates how the combination of decreasing marginal privacy cost and the platform's commitment to underuse data makes consumer privacy difficult to sustain. Indeed, if the consumer faced an increasing convex loss of providing data, then the platform's learning could stop in the middle. If the platform had no commitment power, then the consumer might choose $a_t = 0$, anticipating that she incurs a high privacy cost.

Figure 1 depicts the equilibrium dynamics in a numerical example.[11] Figure 1(a) shows that 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 in Proposition 1.[12]

3.3 Implications of Proposition 2

First, Proposition 2 gives an economic explanation of the so-called *privacy paradox*: Consumers seem to casually share their data with online platforms, despite their demand for privacy and con-

[11] I compute the equilibrium strategy profile using Proposition 3.

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

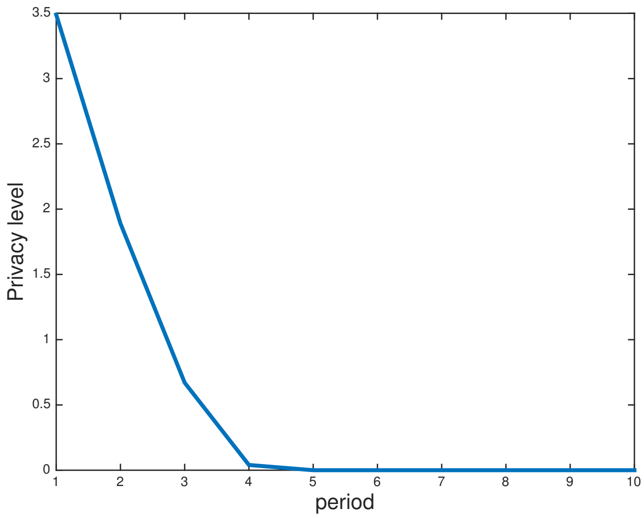

Figure 1(a): Privacy level γ_t

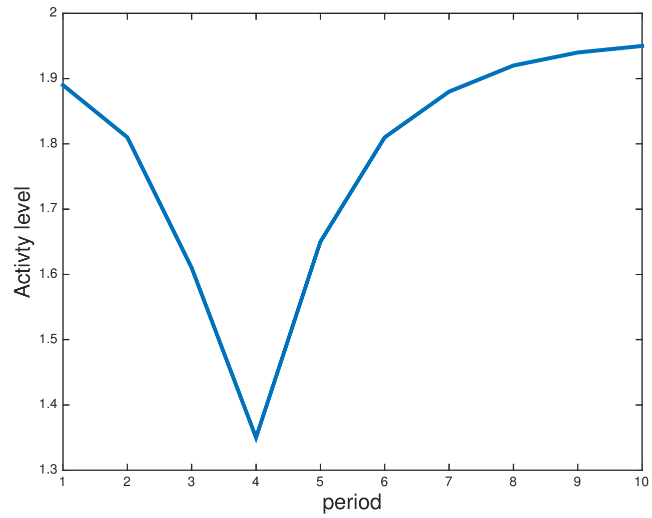

Figure 1(b): Activity level a_t

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

cerns about data collection.[13] One may view this puzzle as the long-run equilibrium outcome of this model, where the consumer faces a high privacy cost and zero marginal cost. In contrast to the explanation based on information externalities among consumers, the outcome in Proposition 2 occurs without a priori market imperfection. The result also points to the difficulty of applying the revealed preference argument to infer the value of privacy to a consumer. Indeed, a consumer's privacy choice in a single period may not tell much about his or her preferences for privacy (v), if the choice is made after the consumer revealed much information in the past. Moreover, $(a_t^*, \sigma_t^2) \to (a_{max}, 0)$ holds even if the consumer is forward-looking and arbitrarily patient (Proposition 9).

Second, the result establishes a novel connection between consumer privacy problem and rational harmful addiction (Becker and Murphy, 1988). The connection comes from the observation that the consumer's utility is decreasing but marginal utility is increasing in the amount of information collected in the past. One difference from a typical model of rational addiction is that the platform can dynamically adjust the degree of addiction through its privacy policy. As a result, even if the consumer values privacy arbitrarily highly ex ante, she becomes "addicted" to the platform.

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

At an anecdotal level, the equilibrium strategy of the platform seems consistent with how the privacy policies of online platforms have evolved. In 2004, Facebook's privacy policy stated that it would not use cookies to collect consumer information. In 2020, the privacy policy states that it uses cookies to track consumers on and possibly off the website.[14] Srinivasan (2019) describes how Facebook has acquired dominance in the social media market:

> "When Facebook entered the market, the consumer's privacy was paramount. The company prioritized privacy, as did its users—many of whom chose the platform over others due to Facebook's avowed commitment to preserving their privacy. Today, however, accepting Facebook's policies in order to use its service means accepting broad-scale commercial surveillance."

Fainmesser et al. (2019) describe how the business models of online platforms have changed from the initial phase where they expand a user base to the mature phase where they monetize collected information. The equilibrium dynamics Proposition 2 rationalize the described pattern.

3.4 Characterizing Equilibrium Privacy Policy

The following result characterizes the platform's equilibrium privacy policy. Recall that $a^*(\gamma, \sigma^2)$ is the activity level a_t chosen by the myopic consumer given $\gamma_t = \gamma$ and $\sigma^2_{t-1} = \sigma^2$.

Proposition 3. *The equilibrium privacy policy* $(\gamma_1^*, \gamma_2^*, \dots)$ *is recursively defined as follows.*

$$\gamma_t^* \in \arg\min_{\gamma \geq 0} \frac{1}{a^*(\gamma, \hat{\sigma}^2_{t-1})} + \gamma, \forall t \in \mathbb{N}, \tag{5}$$

$$\hat{\sigma}^2_0 = \sigma^2_0, \tag{6}$$

$$\hat{\sigma}^2_t = \frac{1}{\frac{1}{\hat{\sigma}^2_{t-1}} + \frac{1}{\frac{1}{a^*(\gamma_t^*, \hat{\sigma}^2_{t-1})} + \gamma_t^*}}, \forall t \in \mathbb{N}. \tag{7}$$

Moreover, given any privacy policy $\boldsymbol{\gamma}$*, let* $(\sigma^2_t)_{t \in \mathbb{N}}$ *denote the conditional variances induced by the consumer's optimal behavior. Then,* $\hat{\sigma}^2_t \leq \sigma^2_t$ *for all* $t \in \mathbb{N}$.

[14] 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

Proof. Take any privacy policy $(\gamma_t)_{t\in\mathbb{N}}$, and let $(\sigma_t^2)_{t\in\mathbb{N}}$ denote the sequence of the conditional variances induced by $a^*(\cdot,\cdot)$. I show $\hat{\sigma}_t^2 \leq \sigma_t^2$ for all $t \in \mathbb{N}$. The inequality holds with equality for $t = 0$. Take any $\tau \in \mathbb{N}$. Suppose $\hat{\sigma}_t^2 \leq \sigma_t^2$ for $t = 0, \ldots, \tau - 1$. It holds

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

The first inequality follows from the inductive hypothesis and decreasing $a^*(\gamma,\cdot)$. The second inequality follows from (5). $\forall t \in \mathbb{N}, \hat{\sigma}_t^2 \leq \sigma_t^2$ implies that the privacy policy described by (5), (6), and (7) is optimal, because the platform obtains a greater profit than any other privacy policy in all periods. □

The objective of the minimization problem (5), $\frac{1}{a^*(\gamma,\hat{\sigma}_{t-1}^2)} + \gamma$, is the variance of the noise ε_t in the signal $s_t = X + \varepsilon_t$ given the consumer's best response. The minimization problem captures the platform's trade-off. On the one hand, a higher privacy level γ leads to a higher activity level, which leads to a lower variance $\frac{1}{a^*(\gamma,\hat{\sigma}_{t-1}^2)}$ of ε_t. On the other hand, given any activity level, a higher γ lowers the informativeness of the signal. This cost is captured by the second term γ. The platform chooses γ_t^* by resolving this trade-off. As the platform solves (5) in each period, the conditional variance evolves according to (7) with the initial condition (6).

The platform chooses its strategy to maximize the sum of discounted profits. However, the equilibrium policy is as if the platform sets each γ_t to myopically maximize the precision of the signal. The reason is follows. In principle, the platform chooses (say) γ_1 to maximize the sum of period-1 profit and the continuation value. The period-1 profit is increasing in the precision of the signal in $t = 1$ by construction. As more information is generated in $t = 1$, the consumer faces lower marginal costs and chooses higher activity levels in the future. Thus, the continuation value is also increasing in the precision of the signal in $t = 1$. As a result, the platform can maximize the sum of discounted profits by maximizing the informativeness of signal in $t = 1$. A similar argument implies that the equilibrium privacy level in any period myopically maximizes the informativeness of the signal in that period.

Proposition 3 implies that the platform does not require as strong a commitment power as currently assumed.[15] I say that the platform has *short-run commitment power* if it chooses a privacy

[15]The proof of the following result also reveals that the equilibrium outcome is independent of the platform's

level γ_t at the beginning of each period t (before the consumer chooses a_t) without committing to future privacy levels.

Corollary 1. *Let* $(\boldsymbol{a}^*, \boldsymbol{\gamma}^*)$ *denote the equilibrium outcome of the baseline model in which the platform can commit to any privacy policy before* $t = 1$. *The same outcome* $(\boldsymbol{a}^*, \boldsymbol{\gamma}^*)$ *arises in an equilibrium when the platform has short-run commitment power.*

Proof. Suppose the platform only has short-run commitment power. Consider the following strategy profile: Following every history with the conditional variance σ^2, the platform sets $\gamma \in \arg\min_{\gamma \geq 0} \frac{1}{a^*(\gamma, \sigma^2)} + \gamma$. The consumer always acts according to $a^*(\cdot, \cdot)$. By construction, $(\boldsymbol{a}^*, \boldsymbol{\gamma}^*)$ arises on the path of play. Any deviation by the platform increases the conditional variances in all periods (Proposition 3). Thus, it has no profitable deviation. □

The result shows that long-run commitment has no value relative to short-run commitment. In contrast, the platform could be strictly worse off if it has no commitment power: If the platform sets γ_t after observing a_t, then in any equilibrium, the platform sets $\gamma_t = 0$ whenever $a_t > 0$. Anticipating this, the consumer chooses a lower activity level than under the short-run or long-run commitment. I argue that short-run commitment is a reasonable assumption in practice, because a platform could be sanctioned for the outright violation of its privacy policy.

4 Platform Competition

I now explore the effect of competition between two platforms, an incumbent (I) and an entrant (E). I is in the market from the beginning of $t = 1$. In period $t^* \geq 2$, E enters the market. t^* is exogenous, deterministic, and known to I at the outset.[16]

Before the entry ($t < t^*$), the consumer chooses the activity level $a_t^I \geq 0$ for I. After the entry ($t \geq t^*$), the consumer chooses $(a_t^I, a_t^E) \in \mathbb{R}_+^2$, where a_t^E is the activity level for E. I assume single-homing: The consumer can choose (a_t^I, a_t^E) if and only if $\min(a_t^I, a_t^E) = 0$. Single-homing is natural if platforms offer similar services such as search engines.

discount factor or time horizon.

[16]Exogenous entry is to simplify exposition. If E can choose to enter in any period $t \geq 2$ at a positive entry cost, then we obtain a similar result where E does not enter in equilibrium.

I consider two games that differ in the timing of moves. One is *competition with long-run commitment:* I publicly commits to $(\gamma_1^I, \gamma_2^I, \dots)$ at the beginning of $t = 1$, and E publicly commits to $(\gamma_{t^*}^E, \gamma_{t^*+1}^E, \dots)$ at the beginning of period t^*. The other is *competition with short-run commitment:* In each period, each platform chooses a privacy level, after which the consumer chooses an activity level. In particular, I and E set privacy levels γ_t^I and γ_t^E simultaneously in each period $t \geq t^*$.

As in the baseline model, activity on platform $k \in \{I, E\}$ generates a signal $s_t^k = X + \varepsilon_t^k$ with $\varepsilon_t^k \sim \mathcal{N}\left(0, \frac{1}{a_t^k} + \gamma_t^k\right)$. Each platform k privately observes s_t^k, that is, there is no information spillover. All 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 conditional variance of X given activity levels and privacy levels up to t on platform k. The consumer's payoff in period t is given by

$$u(a_t^I) - v\left(\sigma_0^2 - \sigma_{t,I}^2\right) + \mathbf{1}_{\{t \geq t^*\}} \cdot \left[u(a_t^E) - v\left(\sigma_0^2 - \sigma_{t,E}^2\right)\right], \tag{8}$$

where $\mathbf{1}_{\{t \geq t^*\}}$ is the indicator function that equals 1 or 0 if $t \geq t^*$ or $t < t^*$, respectively. Payoff (8) implies that even if the consumer switches to 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 (see the discussion in Section 2. Section 7 relaxes this assumption).

To ensure the existence of an equilibrium, I impose an upper bound on the feasible privacy levels. In practice, the bound might reflect the minimum amount of data that a platform needs to collect for offering services, or the maximum privacy level that a platform can credibly enforce.

Assumption 1. There is a $\bar{\gamma} \in \mathbb{R}_+$ satisfying $a^*\left(\bar{\gamma}, \sigma_0^2\right) > 0$ such that each platform can choose a privacy level of at most $\bar{\gamma}$.

$a^*\left(\bar{\gamma}, \sigma_0^2\right) > 0$ implies that if a platform chooses $\bar{\gamma}$, then the consumer chooses $a_t^I > 0$ or $a_t^E > 0$. This restriction on $\bar{\gamma}$ is necessary for a non-trivial equilibrium in which the consumer uses a platform.

4.1 Equilibrium under Competition

I present an equilibrium in which the consumer never switches to the entrant and the long-run outcome equals the monopoly outcome.[17] If the entry is sufficiently late, then the equilibrium out-

[17]I do not consider other equilibria, which are likely to exist because there are two patient players in the game.

come exactly equals the monopoly one. Moreover, there is no equilibrium in which the consumer permanently switches to the entrant.

Proposition 4. *Regardless of the commitment assumption:*

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 a* $\underline{t} \geq 2$ *such that, if the entry time* t^* *is greater than* $\underline{t}$, *then* I*'s privacy policy* $(\gamma_t^I)_{t\in\mathbb{N}}$ *in the above equilibrium coincides with the monopoly strategy.*

3. *Switching never occurs: There is no equilibrium in which for some* $\hat{t} \in \mathbb{N}$, $a_t^E > 0$ *and* $a_t^I = 0$ *for all* $t \geq \hat{t}$.

The intuition is as follows. Before the entry, I chooses privacy levels to induce the consumer to choose positive activity levels. Suppose that, upon entry, E sets the highest privacy level $\bar{\gamma}$. Since the privacy cost from collected data is sunk, the consumer chooses which platform to use based on the marginal (or more precisely, incremental) cost. Because the incumbent has collected data, the consumer faces a lower marginal cost of using I. Thus, if I also chooses $\bar{\gamma}$, then the consumer strictly prefers to use I. However, the equilibrium choice of I may not be $\bar{\gamma}$: I chooses a privacy level to maximize the precision of the signal subject to the constraint that the consumer prefers I. As time goes by, the constraint is relaxed, because the consumer's marginal cost of using I goes to zero. Thus, I can offer a vanishing privacy level over time.

The threat of future entry has no impact on I's strategy: Before the entry, I chooses the same privacy levels as monopoly, regardless of commitment assumption. This is because maximizing the amount of information makes consumer switching least likely.

Point 3 implies that, to poach the consumer, E needs some advantage in terms of the quality of service or the privacy level. I say that E can *successfully enter the market* if there exists an equilibrium in which the consumer switches to E upon entry, i.e., $a_t^E > 0$ and $a_t^I = 0$ for all $t \geq t^*$.[18]

[18] The result considers the entrant's advantage in terms of service quality. We obtain a similar result by considering the entrant that can choose a higher maximum privacy level $\bar{\gamma} + \Delta$ than the incumbent.

Proposition 5. *Suppose that the gross benefit of E's service is given by $u^E(\cdot) := u(\cdot) + \Delta$. Regardless of commitment assumption, there is a $\Delta^* > 0$ such that for any $\Delta \geq \Delta^*$, E can successfully enter the market. The minimum Δ^* satisfying this property is increasing in t^*.*

4.2 Antitrust Implication of Propositions 4 and 5

The results imply that data held by incumbents can be a barrier to entry. To see this, suppose that a platform has collected much data from consumers. Data collection lowers the welfare of consumers who value their privacy. However, consumers also face lower marginal privacy costs, because there is less for them to lose on the margin. If consumers regard collected data as sunk (for the reasons discussed in Section 2), they decide which platform to use based on marginal costs. Since the incumbent has an advantage in terms of lower marginal costs, switching and market entry become less likely to occur.

In the model, low marginal costs are associated with high privacy costs. Thus, switching and market entry are less likely when consumers suffer from a lack of privacy and receive low payoffs from the incumbent. This welfare implication contrasts with the existing idea of "data as an entry barrier," where dominant platforms use data to improve their services. For example, Furman et al. (2019) states 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" (italicized by the author).

As an application, consider the market for search engines: The incumbent is Google, and the entrant is a privacy-preserving alternative of Google, such as DuckDuckGo. My results suggest that, even if DuckDuckGo is as good a search engine as Google, DuckDuckGo may not be able to poach consumers. If consumers have no privacy to Google, then their marginal privacy cost of using Google is negligible. Thus, DuckDuckGo may have to offer much better a search engine than Google without collecting data, which seems unrealistic.

Proposition 4 relies on a strong assumption that the consumer incurs a privacy cost from one platform even after she migrates to the other platform. At the same time, such a model enables us to think what kind of privacy regulation alters this assumption and promotes competition. Indeed, the next section shows that if the consumer can erase past information, competition is more likely to occur.

5 Erasing Past Information

I now consider the incentive of the consumer or a platform to erase past information.

5.1 The Right to be Forgotten

Consider the "right to be forgotten," whereby the consumer can request a platform to erase past information. Below, I describe the model of competition but a similar extension applies to monopoly.

In each period, the consumer makes two decisions. First, the consumer chooses whether to erase past information of each platform in the market. Second, the consumer chooses a_t^I or (a_t^I, a_t^E), depending on whether t is before or after the entry. If she erases information of platform $k \in \{I, E\}$ in period t, then the conditional variance for platform k at the beginning of t becomes σ_0^2. At the end of the period, the consumer still incurs a privacy cost based on information generated in period t. The consumer does not incur any cost to erase information, but the results hold if the cost is not too high.

For example, suppose that the consumer erases information of both platforms in period t and uses platform E. Then, her payoff is

$$u(a_t^E) - v\left[\sigma_0^2 - \sigma_{1,E}^2(a_t^E, \gamma_t^E)\right], \tag{9}$$

where $\sigma_{1,E}^2(a_t^E, \gamma_t^E)$ is the conditional variance for E given one signal based on (a_t^E, γ_t^E). Thus, the privacy cost from E is only based on the signal of period t. Since the consumer has erased information and does not use I, she does not incur a privacy cost from I.

In contrast, suppose that the consumer has never erased information. If she uses platform E in period t, then her payoff in period t is

$$u(a_t^E) - v\left[\sigma_0^2 - \sigma_{t,E}^2(\boldsymbol{a}_t^E, \boldsymbol{\gamma}_t^E)\right] - v\left[\sigma_0^2 - \sigma_{t,I}^2(\boldsymbol{a}_t^I, \boldsymbol{\gamma}_t^I)\right], \tag{10}$$

where $\boldsymbol{a}_t^k = (a_1^k, \ldots, a_t^k)$ and $\boldsymbol{\gamma}_t^k = (\gamma_1^k, \ldots, \gamma_t^k)$ for each $k \in \{I, E\}$. Thus, the consumer incurs a privacy cost from both platforms based on past data collection.

As before, we can consider long-run commitment where a platform commits to the entire privacy policy, and short-run commitment where a platform sets a privacy level at the beginning of each period (before the consumer makes any decision). The following result summarizes the impact of the right to be forgotten.

Proposition 6 (The Right to be Forgotten). *If the consumer can costlessly erase past information, then regardless of the commitment assumption, the following equilibrium exists:*

1. *Under monopoly, in all periods, the consumer erases information and the platform sets a privacy level γ_1^* defined in (5).*

2. *Under competition, the consumer erases information in all periods, and both platforms set the highest privacy level $\bar{\gamma}$ in any period $t \geq t^*$ after the entry.*

3. *Under competition, suppose that the gross benefit of E's service is given by $u^E(\cdot) := u(\cdot) + \Delta$. For any $\Delta > 0$, E can successfully enter the market.*

The right to be forgotten benefits the consumer in three ways. First, it reduces privacy cost. Second, it incentivizes platforms to choose higher privacy levels: Once the consumer erases information, she incurs a high marginal privacy cost. Then, a monopoly platform always offers a period-1 privacy level in any period (Point 1). Under competition, erasing information makes platforms homogeneous and intensifies competition. As a result, platforms offer the highest privacy level (Point 2). Finally, erasing information eliminates the incumbency advantage and promotes the entry of a higher quality platform (Point 3).

5.2 Data Retention Policies

This section studies whether a platform has an incentive to voluntarily erase data collected in the past. This question relates to data retention policies, which have recently been paid attention by economists and legal scholars (Chiou and Tucker, 2017).

I consider a game with short-run commitment, where a platform sets a privacy level at the beginning of each period.[19] If there is one platform in the market, then the platform chooses whether to erase past information and sets a privacy level. Then, the consumer chooses an activity level. If there are two platforms, then they simultaneously choose whether to erase information. After observing this, they simultaneously choose privacy levels. Finally, the consumer chooses an activity level for each platform.

A platform's erasing information affects the conditional variances and payoffs in the same way as the consumer erasing information (see the previous subsection). The following result shows that a platform's optimal data retention policy is to not erase information at all.

Proposition 7. *A platform never erases information:*

1. *A monopoly platform never erases information in any period in any equilibrium. The equilibrium outcome equals Proposition 2.*

2. *Under competition, there is an equilibrium in which platforms never erase information in any period. Among the equilibria with this property, Proposition 4 holds.*

The result illustrates that a platform has different incentives to offer ex ante and ex post privacy protections. If consumer behavior is exogenous, then a platform has no incentive to raise a privacy level or erase past information, because it lowers the profit by reducing the amount of information. However, if consumer behavior is endogenous, a platform may have an incentive to increase a privacy level (at least in early periods) to reduce the consumer's marginal cost and increases her activity level. In contrast, a platform has no incentive to erase information because it increases the consumer's marginal privacy cost and decreases her activity level. Thus, we may not expect competition to encourage platforms to offer the right to be forgotten to consumers.

[19] The same result holds for other commitment assumptions. For example, if a platform can commit to a privacy policy and a data retention policy at the outset, then before $t = 1$, it commits to a privacy policy $(\gamma_1, \gamma_2, \dots)$ and the frequency $T \in \mathbb{N}$ with which the platform deletes information. Here, T means that the platform erases information in every T periods.

6 Forward-looking Consumer

The main results under a monopoly platform hold even when the consumer is forward-looking. Assume that the platform and the consumer have any discount factors $\delta_P \in (0,1)$ and $\delta_C \in (0,1)$, respectively. The solution concept is pure-strategy subgame perfect equilibrium. The proofs of the following results are in Appendix F.

6.1 Monopoly with Long-run Commitment Power

First, consider a monopoly platform that can commit to any privacy policy at the outset. The following result extends Proposition 1.

Proposition 8 (Stationary Policy with a Patient Consumer). *Take any privacy policy $\boldsymbol{\gamma}$ such that $\gamma_t = \gamma$ for all $t \in \mathbb{N}$. Let $(\bar{a}_t)_{t\in\mathbb{N}}$ denote the equilibrium strategy in the subgame following $\boldsymbol{\gamma}$. There is a $v^*(\gamma) > 0$ such that the following holds:*

1. *If $v < v^*(\gamma)$, then $\bar{a}_t$ is increasing in t, $\lim_{t\to\infty} \bar{a}_t = a_{max}$, and $\lim_{t\to\infty} \sigma_t^2 = 0$.*
2. *If $v > v^*(\gamma)$, then $\bar{a}_t = 0$ for all $t \in \mathbb{N}$.*

Moreover, $v^*(\gamma)$ is increasing and $\lim_{\gamma\to\infty} v^*(\gamma) = \infty$.

The following result extends Proposition 2. The fundamental features of the model are decreasing marginal privacy cost and the platform's ability to commit to offer privacy, and the myopia assumption is not necessary.

Proposition 9. *For any $v \in \mathbb{R}$ and $(\delta_P, \delta_C) \in (0,1)^2$, in any equilibrium:*

$$\lim_{t\to\infty} a_t^* = a_{max} \text{ and } \lim_{t\to\infty} \sigma_t^2 = 0. \tag{11}$$

A technical challenge is that if the platform chooses a non-stationary privacy policy, then the consumer faces a non-stationary dynamic programming, which is intractable. Thus, I first show that (11) holds under a stationary privacy policy $\boldsymbol{\gamma}^S = (\gamma^S, \gamma^S, \dots)$ for a large $\gamma^S \in \mathbb{R}$. To show that any equilibrium involves (11), suppose to the contrary that the equilibrium privacy policy $\boldsymbol{\gamma}^E$ involves imperfect learning (i.e., $\lim_{t\to\infty} \sigma_t^2 > 0$), which implies that the precision of signal s_t goes to zero as t grows large. Suppose that the platform modify $\boldsymbol{\gamma}^E$ by replacing $(\gamma_\tau^E, \gamma_{\tau+1}^E, \dots)$ with $(\gamma^S, \gamma^S, \dots)$ for a large $\tau \in \mathbb{N}$. This replacement increases the precisions of the signals after period τ. It also increases activity levels before period τ, because the consumer chooses higher activity levels when she anticipates more data collection in the future. The consumer's response to future data collection is specific to the forward-looking consumer. It comes from the supermodularity of her objective in activity levels and the informativeness of signals in different periods.

6.2 Monopoly with Short-run Commitment Power

Suppose now that the platform sets a privacy level γ_t at the beginning of each period t. We may think that the equilibrium outcome substantively changes once we jointly consider the patient consumer and the platform with weaker commitment. For example, the consumer might choose to not use the platform when it cannot commit to high privacy levels in the future. Contrary to this intuition, I show that any equilibrium involves privacy loss under a mild modification:

Assumption 2. The following holds.

1. In each period, the consumer chooses a_t from a finite set $A \subset \mathbb{R}_+$ such that $\min A = 0$ and $\max A = a_{max}$.
2. The platform can choose a privacy level of at most $\bar{\gamma}$ that satisfies $\bar{\gamma} > \left(\frac{v\sigma_0^2}{(1-\delta_C)u(a_{max})} - 1\right)\sigma_0^2 - \frac{1}{u(a_{max})}$.

Point 1 means that there are finitely many activity levels.[20] Point 2 means that there is a large finite upper bound on feasible privacy levels.

Proposition 10. *Even if the consumer is patient and the platform has short-run commitment power, the consumer eventually loses privacy: For any $v \in \mathbb{R}$ and $(\delta_P, \delta_C) \in (0,1)^2$, in any equilibrium, $\lim_{t\to\infty} \sigma_t^2 = 0$ and $\lim_{t\to\infty} a_t^* = a_{max}$.*

[20] To prove $\sigma_t^2 \to 0$ in Proposition 10, we can replace Point 1 with a weaker assumption that A is a compact set that has the smallest positive activity level, i.e., $\inf(A \setminus \{0\}) > 0$. I adopt Point 1 to simplify exposition.

7 Extensions

This section shows that the main insight holds in a more general environment. Also, I examine how the equilibrium dynamics interact with the platform's incentive to invest in quality. Unless otherwise noted, I consider a monopoly platform with long-run commitment power and a myopic consumer.

7.1 Consumers with Heterogeneous v

Proposition 2 holds when consumers have heterogeneous v's. In other words, the main insight does not depend on the platform knowing v at the outset. To see this, 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 some distribution with a finite support $V \subset \mathbb{R}_+$. Let $\alpha_v \in [0,1]$ denote the mass of consumers who have $v \in V$. 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. Thus, the platform learns about i's type only based on her past activity levels and privacy levels.

For each $i \in [0,1]$, let $\sigma_t^2(i)$ denote the conditional 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 $\int_{i \in [0,1]} \sigma_0^2 - \sigma_t^2(i) di$. If (almost) all consumers who have the same v choose the same activity level, then 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 conditional 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 could choose high activity levels even without privacy protection.[21]

However, there is no trade-off for the platform in the long-run: All consumers eventually lose privacy and choose the highest activity levels (see Appendix H.1 for the proof).

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

Proposition 11. *Let $(a_t^*(v), \sigma_t^2(v), \gamma_t^*)_{t\in\mathbb{N}, v\in V}$ denote the outcome of any equilibrium. Then,*

$$\forall v \in V, \lim_{t\to\infty}(a_t^*(v), \sigma_t^2(v)) = (a_{max}, 0) \text{ and } \lim_{t\to\infty} \gamma_t^* = 0. \tag{12}$$

To see the intuition, suppose that v is either $L = 0$ or $H > 0$, and the platform sets $\gamma_t = 0$ in early stages to obtain 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.

7.2 Endogenous Quality of Service

So far, the benefit $u(\cdot)$ from the platform's service has been exogenous. Suppose now that, before $t = 1$, the platform chooses a quality $q \geq 0$. Then, in each period, the consumer receives a gross benefit of $u_q(a) = qa - \frac{1}{2}a^2$, and the platform receives a payoff of $\sigma_0^2 - \sigma_t^2 - c(q)$ for some strictly increasing $c(\cdot)$. The platform chooses q once, but incurs $c(q)$ in every period.

Proposition 12. *A patient platform does not invest in quality: For any $\delta_P \in (0,1)$, let $q(\delta_P)$ denote the quality in an (arbitrarily chosen) equilibrium. Then, $\lim_{\delta_P\to 1} q(\delta_P) = 0$. Thus, regardless of the consumer's discount factor, as $\delta_P \to 1$, her ex ante sum of discounted payoffs converges to zero, and her long-run equilibrium payoff converges to $-v\sigma_0^2 < 0$.*

Proof. Given (δ_P, q), Let $\Pi(\delta_P, q)$ denote the platform's ex ante sum of discounted profits. For any $q > 0$, the platform's per-period payoff is at most $\sigma_0^2 - c(q)$. Thus, $(1-\delta_P)\Pi(\delta_P, q) \leq \sigma_0^2 - c(q)$. Suppose to the contrary that there is a sequence $\delta_n \to 1$ such that for some $q' > 0$, $q(\delta_n) \geq q'$ for infinitely many n's (for some selection of equilibria). Without loss of generality, assume $(1-\delta_n)\Pi(\delta_n, q(\delta_n)) \in [0, \sigma_0^2]$ has a limit. Then, $\lim_{n\to\infty}(1-\delta_n)\Pi(\delta_n, q(\delta_n)) \leq \sigma_0^2 - c(q') < \sigma_0^2 - c(q'/2)$. Proposition 9 implies that there is a $\boldsymbol{\gamma}$ under which $(a_t^*, \sigma_t^2) \to (a_{max}, 0)$ given quality $q'/2$. If the platform chooses $q'/2$ and $\boldsymbol{\gamma}$, then as $\delta_P \to 1$, its average payoff converges to $\sigma_0^2 - c(q'/2)$. Thus, the platform with a large δ_n strictly prefers $q'/2$ to $q(\delta_n)$, which is a contradiction. Thus, $\lim_{\delta_P\to 1} u_{q(\delta_P)}(a_{max}) - v\sigma_0^2 = -v\sigma_0^2$. Also, as the consumer's ex ante payoff is nonnegative but lower than $\frac{u_{q(\delta_P)}(a_{max})}{1-\delta_C}$, it converges to 0 as $\delta_P \to 1$. □

7.3 General Privacy Cost Function

This subsection generalizes consumer preferences in two ways. First, I relax the assumption that privacy cost is sunk. Second, I relax the assumption that privacy cost is linear in σ_t^2. Appendix H.2 contains the proof.

7.3.1 Relaxing "Privacy Cost is Sunk"

The baseline model assumes that 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 only a fraction $\alpha \in [0, 1)$ of the privacy cost when $a_t = 0$. 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.

Proposition 13. *Take any $v \in \mathbb{R}$. Let $(a_t^*, \gamma_t^*, \sigma_t^2)_{t\in\mathbb{N}}$ denote the outcome of any equilibrium. There is an $\alpha^* < 1$ such that for any $\alpha \geq \alpha^*$, $\lim_{t\to\infty} a_t^* = a_{max}$, $\lim_{t\to\infty} \gamma_t^* = 0$, and $\lim_{t\to\infty} \sigma_t^2 = 0$.*

Proposition 14. *There is an $\alpha^* < 1$ such that for any $\alpha \geq \alpha^*$, the result under competition (Proposition 4) holds.*

7.3.2 Relaxing Linearity

Suppose that the consumer's per period payoff is $u(a_t) - C\left(\sigma_t^2\right)$. Assume $C(\cdot) : \mathbb{R}_+ \to \mathbb{R}$ is continuously differentiable. In particular, the cost and the marginal cost at no privacy (i.e., $C(0)$ and $C'(0)$) are finite. $C(\cdot)$ can be non-monotone: For example, $C(\cdot)$ can be first decreasing and then increasing, which means that the consumer prefers some (but not too much) data collection. The following result shows that the long-run outcome remains the same.

Proposition 15. *In any equilibrium, $\lim_{t\to\infty} a_t^* = a_{max}$ and $\lim_{t\to\infty} \sigma_t^2 = 0$.*

7.4 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 $Var(X_t) = \sigma_0^2$

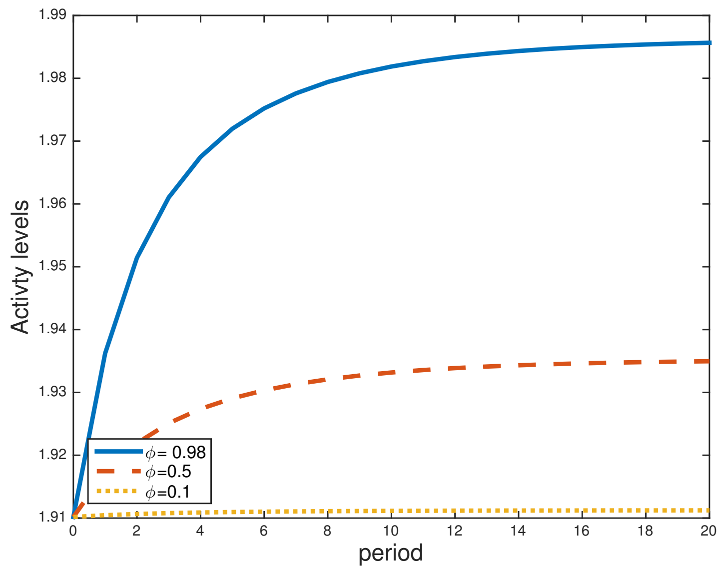

Figure 2: Activity levels $u(a) = 2a - \frac{1}{2}a^2$, $v = 10$, $\sigma_0^2 = 1$, $\phi \in \{0.1, 0.5, 0.98\}$, and $\gamma_t \equiv 4$.

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 conditional 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. Since we cannot use the first-order condition to solve the consumer's problem, we are not able to obtain a simple set of equations to characterize the steady state.

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

Figure 3 presents equilibria taking into account the platform's optimization. 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 relatively a 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

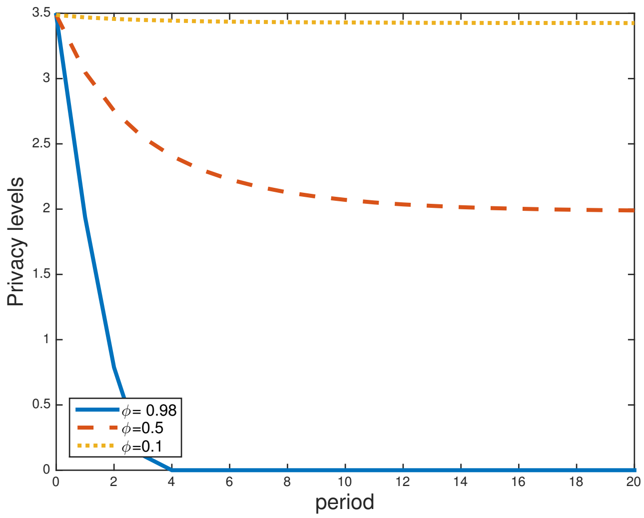

Figure 3(a): Privacy level γ_t

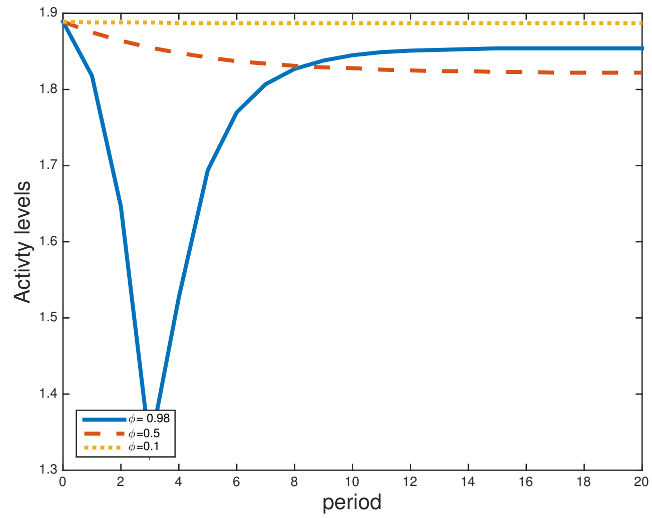

Figure 3(b): Activity level a_t

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

to obtain non-trivial amount of information in the steady state even if v is larger.[22] 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 is not necessarily decreasing in ϕ. 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 $\phi = 0.1$. Thus, the activity level is no longer monotone in ϕ once we consider the platform's optimal privacy policy.

7.5 General Payoffs for the Platform

All the results of this paper continue to hold if the platform's final payoff from any sequence of conditional 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 affect the analysis because, in the equilibrium under monopoly or competition, any deviation by a platform weakly increases σ_t^2 for all $t \in \mathbb{N}$.

One natural specification of $\Pi(\cdot)$ is as follows: Suppose that the platform sells information to a sequence of short-lived data buyers. Any information sold in period t is freely replicable later and

[22]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 $a_t = a_{max} = 2$.

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 newly generated in period t. Thus, the ex ante payoff is $\sum_{t=1}^{\infty} \delta_P^{t-1}(\sigma_{t-1}^2 - \sigma_t^2)$. This objective is strictly decreasing in each σ_t^2, because the coefficient of each σ_t^2 is $-\delta_P^{t-1} + \delta_P^t < 0$.

8 Conclusion

This paper studies a dynamic model of consumer privacy and platform data collection. The fundamental feature of the model is that a consumer faces a lower marginal loss of giving up privacy as a platform collects more data. The paper explores dynamic implications of this idea. First, a monopoly platform is able to collect much information over time by committing to underuse data in the early stages. In equilibrium, the consumer eventually loses privacy but keeps choosing a high level of activity. The result is robust to extensions such as consumer heterogeneity and a platform's weaker commitment power. Second, decreasing marginal privacy cost makes competition unhelpful, because a consumer is more likely to stick to a platform to which she has less privacy. The results rely on "data persistence," where data generated today can adversely affect consumers in the future. I show how regulating data collection could perversely lower the long-run privacy and consumer welfare, and how a regulation such as the right to be forgotten promotes competition and benefits consumers.

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.

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.

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.

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

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*.

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

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

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*.

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

Harris, Christopher (1985), “Existence and characterization of perfect equilibrium in games of perfect information.” *Econometrica: Journal of the Econometric Society*, 613–628.

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.

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).

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

Marthews, Alex and Catherine Tucker (2019), "Privacy policy and competition." *Brookings Paper*.

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*.

Appendix

A Properties of the Consumer’s Best Response: Proof of Lemma 2

Proof. Define $U(a,\gamma,\sigma^2) := u(a) - v\left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma^2}+\frac{1}{\frac{1}{a}+\gamma}}\right)$. Lemma 1 implies that $\frac{\partial U}{\partial a}$ is increasing in γ and decreasing in σ^2. The standard argument of monotone comparative statics implies that $a^*(\gamma,\sigma^2)$ is increasing in γ and decreasing in σ^2 (e.g., Milgrom et al. 1994).

Suppose to the contrary that $\lim_{\sigma^2\to 0} a^*(\gamma,\sigma^2) = a_{max}$ fails. As $a^*(\gamma,\sigma^2) \leq a_{max}$ for all (γ,σ^2), there are $\varepsilon > 0$ and $(\sigma_n^2)_{n\in\mathbb{N}}$ such that $\lim_{n\to\infty} \sigma_n^2 = 0$, $a^*(\gamma,\sigma_n^2) \leq a_{max} - \varepsilon$ for all $n \in \mathbb{N}$, and $\lim_{n\to\infty} a^*(\gamma,\sigma_n^2) = a_{max} - \varepsilon$. Suppose that the consumer chooses a_{max} instead of $a < a_{max}$. Then, the payoff difference is

$$\Delta(a,\sigma^2) := u(a_{max}) - u(a) - v\left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma^2}+\frac{1}{\frac{1}{a_{max}}+\gamma}}\right) + v\left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma^2}+\frac{1}{\frac{1}{a}+\gamma}}\right). \qquad (13)$$

Note that $\lim_{n\to\infty} \Delta(a^*(\gamma,\sigma_n^2),\sigma_n^2) = u(a_{max}) - u(a_{max}-\varepsilon) > 0$. This implies that $\Delta(a^*(\gamma,\sigma_n^2),\sigma_n^2) > 0$ for a large n. Thus, for a large n, the consumer strictly prefers a_{max} to $a^*(\gamma,\sigma_n^2)$, which is a contradiction. A similar argument implies that $\lim_{\hat{\gamma}\to\infty} a^*(\hat{\gamma},\sigma^2) = a_{max}$. □

B The Long-run Outcome Under a Stationary Privacy Policy: Proof of Proposition 1

Proof. To clarify that the optimal activity level depends on v, I write $a^*(\gamma,\sigma^2)$ as $a^*(v,\gamma,\sigma^2)$, which is decreasing in v. Define $v^*(\gamma)$ as follows:

$$v^*(\gamma) = \sup\left\{v \in \mathbb{R} : a^*\left(v,\gamma,\sigma_0^2\right) > 0\right\}. \qquad (14)$$

Note that $\frac{\partial U}{\partial a} = u'(a) - v\frac{1}{\left(\frac{1}{\sigma_0^2}(1+\gamma a)+a\right)^2}$. This implies that

$$\left.\frac{\partial U}{\partial a}\right|_{a=0} = u'(0) - v \cdot (\sigma_0^2)^2,$$
$$\left.\frac{\partial U}{\partial a}\right|_{a=a'} \leq u'(0) - v \cdot \frac{1}{\left(\frac{1}{\sigma_0^2}\left(1+\gamma a_{max}\right)+a_{max}\right)^2}, \forall a' \in [0, a_{max}].$$

The second inequality holds because $u(\cdot)$ and the privacy cost function are increasing and concave. For a sufficiently small v, the right hand side of the first equality is positive. Thus, $v^*(\gamma)$ is well-defined and positive. For a sufficiently large v, the right hand side of the second inequality is negative. Thus, $v^*(\gamma)$ is finite.

Suppose $v < v^*(\gamma)$. By the construction of $v^*(\gamma)$, $a^*(v, \gamma, \sigma_0^2) > 0$. σ_t^2 decreases in t for any sequence of activity levels. Thus, if $\gamma_t = \gamma$ for any $t \in \mathbb{N}$, then $a^*(v, \gamma, \sigma_t^2)$ is increasing in t and greater than $a_1^* > 0$ for all t. This implies that $\lim_{t\to\infty} \sigma_t^2 = 0$, because

$$0 \leq \sigma_t^2 \leq \frac{1}{\frac{1}{\sigma_0^2} + \frac{t}{\left(\frac{1}{a_1^*}+\gamma\right)}} \to 0 \quad \text{as} \quad t \to \infty.$$

Lemma 2 implies $\lim_{t\to\infty} a_t^* \to a_{max}$. For $v > v^*(\gamma)$, note that $a^*(v, \gamma, \sigma_0^2) = 0$, which implies that $a_t^* = 0$ for all t. $v^*(\gamma)$ is increasing in γ, because $a^*(v, \gamma, \sigma_0^2)$ is increasing in γ. Finally, $\lim_{\gamma\to\infty} a^*(\gamma, \sigma^2) = a_{max}$ in Lemma 1 implies $\lim_{\gamma\to\infty} v^*(\gamma) = \infty$. □

C Properties of Equilibrium: Proof of Proposition 2

Proof. The proof relies on Proposition 3, which I prove in the main text. Let $(a_t^*, \hat{\sigma}_t^2)_{t\in\mathbb{N}}$ denote the equilibrium activity levels and conditional variances. First, I prove $\lim_{t\to\infty} \hat{\sigma}_t^2 = 0$. Proposition 1 implies that there is a stationary privacy policy $\gamma_t \equiv \gamma$ such that $\lim_{t\to\infty} \sigma_t^2 = 0$. Proposition 3 implies that $\hat{\sigma}_t^2 \leq \sigma_t^2$ for all $t \in \mathbb{N}$, which implies $\lim_{t\to\infty} \hat{\sigma}_t^2 = 0$.

To show $\lim_{t\to\infty} a_t^* = a_{max}$, suppose to the contrary that there is an $\varepsilon > 0$ such that $a_t^* \leq a_{max} - \varepsilon$ for infinitely many t's. Without loss of generality, suppose $a_t^* \leq a_{max} - \varepsilon$ for all t. Following the proof of Lemma 2, we can conclude that, for a large t, the consumer strictly prefers a_{max} to a_t^*. Indeed, if the consumer chooses a_{max} instead of a_t^*, $u(\cdot)$ increases by at least $u(a_{max})-$

$u(a_{max} - \varepsilon) > 0$ whereas the increment of privacy cost goes to zero. This is a contradiction.

To prove $\lim_{t\to\infty} \gamma_t^* = 0$, suppose to the contrary that there is a $\underline{\gamma} > 0$ such that $\gamma_t^* \geq \underline{\gamma}$ for infinitely many t's. To simplify exposition, suppose $\gamma_t^* \geq \underline{\gamma}$ for all $t \in \mathbb{N}$. Take any $\varepsilon \in (0, \underline{\gamma})$. Then, since $\lim_{t\to\infty} a_t^* = a_{max} > 0$, for a sufficiently large t, the minimized value in (5) is weakly greater than $\frac{1}{a_{max}} - \varepsilon + \underline{\gamma}$. To show a contradiction, let T denote the first period such that $a^*(0, \hat{\sigma}_{T-1}^2) > 0$. Then, $a^*(0, \hat{\sigma}_{t-1}^2) > 0$ for any $t \geq T$. If the platform chooses $\gamma_t = 0$ instead of γ_t^* in period $t \geq T$, then the minimand in (5) equals $\frac{1}{a^*(0,\hat{\sigma}_t^2)}$, which converges to $\frac{1}{a_{max}} < \frac{1}{a_{max}} - \varepsilon + \underline{\gamma}$ for a sufficiently large t. This implies that for a sufficiently large t, the platform can strictly increase its payoff in period t by setting $\gamma_t = 0$, which is a contradiction.

To show the final part, 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 $\gamma_{t^*}^*(v_n) = 0$. This follows from the proof of Proposition 1, where I show that the consumer with a sufficiently large v chooses $a_1^* = 0$ for a fixed γ_1. This contradicts the optimality of $\boldsymbol{\gamma}^*$ because if the platform sets a sufficiently large γ_{t^*}, then the consumer chooses a positive activity level and the minimand in (5) becomes finite. □

D Equilibrium under Competition: Proofs for Section 4

D.1 Proof of Proposition 4

Proof. For each $k \in \{I, E\}$, I use $-k$ to mean E or I if $k = I$ or $k = E$, respectively. Suppose that, at the beginning of period $t \geq t^*$, the conditional variance for platform k is $\sigma_{t-1,k}^2$. Let γ_t^k denote the privacy level of platform k in period t. The consumer weakly prefers to use platform k (i.e. $a_t^{-k} = 0$ maximizes her period-t payoff) if

$$\arg\max_{a\geq 0} u(a) - v[\sigma_0^2 - \sigma_{t,k}^2(\gamma_t^k, a|\sigma_{t-1,k}^2)] - v[\sigma_0^2 - \sigma_{t-1,-k}^2]$$
$$\geq \arg\max_{a\geq 0} u(a) - v[\sigma_0^2 - \sigma_{t,-k}^2(\gamma_t^{-k}, a|\sigma_{t-1,-k}^2)] - v[\sigma_0^2 - \sigma_{t-1,k}^2],$$

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

$$\arg\max_{a\geq 0} u(a) - v[\sigma^2_{t-1,k} - \sigma^2_{t,k}(\gamma^k_t, a|\sigma^2_{t-1,k})] \geq \arg\max_{a\geq 0} u(a) - v[\sigma^2_{t-1,-k} - \sigma^2_{t,-k}(\gamma^{-k}_t, a|\sigma^2_{t-1,-k})]. \tag{15}$$

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

First, I consider competition with short-run 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 conditional variance (if k^* is not unique, then set $k^* = I$). Then, platform $-k^*$ chooses the highest privacy level $\bar{\gamma}$. Platform k^* chooses a privacy level $\gamma^{k^*}_t$ that solves

$$\begin{aligned} &\min_{\gamma\in[0,\bar{\gamma}]} \frac{1}{a^*(\gamma, \sigma^2_{t-1,k^*})} + \gamma \\ \text{s.t.} \quad &\arg\max_{a\geq 0} u(a) - v[\sigma^2_{t-1,k^*} - \sigma^2_{t,k^*}(\gamma, a|\sigma^2_{t-1,k^*})] \\ &\geq \arg\max_{a\geq 0} u(a) - v[\sigma^2_{t-1,-k^*} - \sigma^2_{t,-k^*}(\bar{\gamma}, a|\sigma^2_{t-1,-k^*})]. \end{aligned} \tag{16}$$

In each period, the consumer myopically chooses a^I_t (if $t < t^*$) or (a^I_t, a^E_t) (if $t \geq t^*$) to maximize her per-period payoff. If indifferent, then the consumer 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 $a^k_t = 0$ for one of $k \in \{I, E\}$ with equal probability, and chooses a^{-k}_t to maximize her period-t payoff.)

I show that 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: 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.

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

$$\min_{\gamma} \frac{1}{a^*(\gamma, \sigma^2_{s-1,I})} + \gamma$$
$$\text{s.t.} \quad \arg\max_{a\geq 0} u(a) - v[\sigma^2_{s-1,I} - \sigma^2_{s,I}(\gamma, a|\sigma^2_{s-1,I})] \tag{17}$$
$$\geq \arg\max_{a\geq 0} u(a) - v[\sigma^2_{s-1,E} - \sigma^2_{s,E}(\bar{\gamma}, a|\sigma^2_{s-1,E})].$$

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 that the set of γ that satisfies the constraint is smaller. Thus, the minimized value in (17) 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. This is 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 (17) by increasing $\sigma^2_{s-1,I} - \sigma^2_{s,I}(\gamma, a|\sigma^2_{s-1,I})$.

On the equilibrium path, $a^E_t = 0$ for all $t \in \mathbb{N}$. $\lim_{t\to\infty} \sigma^2_{I,t} = 0$ holds because it holds even if I adopts $\gamma_t = \bar{\gamma}$ for all t, and I chooses each γ^I_t to achieve even lower conditional variances. Given this result, $\lim_{t\to\infty} a^I_t = a^*$ follows the same proof as monopoly.

Suppose that γ^I_t does not converge to 0. Then, there is a convergent subsequence $\gamma^I_{t(n)}$ such that $\lim_{n\to\infty} \gamma^I_{t(n)} = \gamma' > 0$. For a sufficiently large n, both $\gamma = 0$ and $\gamma = \gamma^I_{t(n)}$ satisfy the constraint in (17), because $\sigma^2_{s-1,E} - \sigma^2_{s,E}(\bar{\gamma}, a|\sigma^2_{s-1,E}) = \sigma^2_0 - \sigma^2_{1,E}(\bar{\gamma}, a^*(\bar{\gamma}, \sigma^2_0)|\sigma^2_0) > 0$, but $\lim_{s\to\infty} \sigma^2_{s-1,I} - \sigma^2_{s,I}(0, a^*(0, \sigma^2_{s-1,I})|\sigma^2_{s-1,I}) \leq \lim_{s\to\infty} \sigma^2_{s-1,I} = 0$. 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 (17) than $\gamma = \gamma'$. This is a contradiction and thus $\lim_{t\to\infty} \gamma^I_t \to 0$ in the equilibrium.

For a sufficiently large t^*, $\sigma^2_{t^*-1,I} \leq \sigma^2_0 - \sigma^2_{t^*,E}(\bar{\gamma}, a^*(\sigma^2_0, \bar{\gamma})|\sigma^2_0)$. Then, for any period $t \geq t^*$, the constraint (17) holds for any $\gamma \leq \bar{\gamma}$. This implies that I's problem is equal to the monopolist's problem after the entry. Combined with the above result, I's choice equals the monopolist's.

Finally, there is no equilibrium in which $a^E_t > 0$ and $a^I_t = 0$ for all $t \geq t^*$. This is because I can then choose $\bar{\gamma}$ for all periods. Given this, the consumer strictly prefers to use I for any period $t \geq t^* \geq 2$, because the consumer has generated information on I in periods $t < t^*$, which leads to a strictly lower marginal privacy cost.

A similar proof applies to competition with long-run commitment. In particular, 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 (17) after t^*.

□

D.2 Successful Entry: Proof of Proposition 5

Proof. Consider the following strategy profile: In any period $t \leq t^*$, I chooses a monopoly strategy. In any period $t \geq t^*$, I chooses $\bar{\gamma}$, whereas E solves

$$\begin{aligned} &\min_{\gamma\in[0,\bar{\gamma}]} \frac{1}{a^*(\gamma, \sigma^2_{t-1,E})} + \gamma \\ \text{s.t.}\quad &\arg\max_{a\geq 0} u(a) + \Delta - v[\sigma^2_{t-1,E} - \sigma^2_{t,E}(\gamma, a|\sigma^2_{t-1,E})] \\ &\geq \arg\max_{a\geq 0} u(a) - v[\sigma^2_{t-1,I} - \sigma^2_{t,I}(\bar{\gamma}, a|\sigma^2_{t-1,I})]. \end{aligned} \tag{18}$$

Let Δ^* denote the lowest Δ such that the set of γ's that satisfy (18) is nonempty given $t = t^*$, $\sigma^2_{t^*-1,E} = \sigma^2_0$, and the monopoly outcome $\sigma^2_{t-1,I}$. Δ^* is well-defined because the set of all γ's satisfying the constraint is non-empty for a large Δ, and the set is upper hemicontinuous in Δ. The rest of the strategy profile is specified analogously to Proposition 4. The same argument as Proposition 4 confirms that this is an equilibrium. Δ^* is increasing in t^*, because a larger t^* decreases $\sigma^2_{t^*-1,I} - \sigma^2_{t^*,I}(\bar{\gamma}, a|\sigma^2_{t-1,I})$.

Finally, suppose $\Delta < \Delta^*$ but there is an equilibrium in which the consumer only uses E in any period $t \geq t^*$. If I adopts a monopoly strategy for any $t < t^*$ and chooses $\gamma_t^I = \bar{\gamma}$ in period t^*, then the consumer strictly prefers to use I in t^*. This weakly increases I's payoff for any period $t < t^*$ and strictly increases I's payoff in period t^*. This is a contradiction. □

E Erasing Past Information: Omitted Proofs from Section 5

E.1 The Right to be Forgotten: Proof of Proposition 6

Proof. Consider monopoly with long-run commitment. 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 (5) with $t = 1$. If the platform has short-run commitment power, then the platform sets γ_t to maximize the amount of information in each period. Because $\frac{1}{a^*(\sigma^2,\gamma)} + \gamma$ is increasing in σ^2, it is optimal for the consumer to erase information, which leads to a weakly lower amount of information generated. In either case, the platform's problem leads to $\gamma_t = \gamma_1^*$ for all t.

For competition, consider the following strategy profile: Before entry, the consumer erases information in all periods, and chooses the activity level according to $a^*(\cdot,\cdot)$. After entry, the consumer erases information in all periods, and chooses the platform that offers $\gamma_t = \bar{\gamma}$ for all $t \geq t^*$ if there is such a platform (for a node in which both platform have deviated, I assign any equilibrium of that subgame). On the path of play, I sets $\gamma_t^I = \gamma_1^*$ for all $t < t^*$ and $\gamma_t^I = \bar{\gamma}$ for all $t \geq t^*$. E sets $\gamma_t^E = \bar{\gamma}$ for all $t \geq t^*$ upon entry. I can pick any equilibrium in any subgame in which the consumer deviates and chooses to not erase information, because the consumer is worse off relative to no deviation.

Finally, for any $\Delta > 0$, we can construct an equilibrium in which (i) the consumer erases in formation in all periods and sets $a_t^I = 0$ for any $t \geq t^*$, (ii) I sets $\bar{\gamma}$ in any period $t \geq t^*$, and (iii) E sets the lowest γ_t^E such that the consumer is indifferent between I and E and therefore chooses E. This is an equilibrium in which $a_t^E > 0$ and $a_t^I = 0$ for all $t \geq t^*$. □

E.2 Data Retention Policies: Proof of Proposition 7

Proof. A monopolists' problem is to solve (5) by choosing a privacy level and whether to erase information. Whenever $\sigma_{t-1}^2 < \sigma_0^2$, erasing information strictly increases the conditional variance, increases the consumer's marginal cost, and shifts $a^*(\cdot, \sigma^2)$ downward. Thus, erasing information strictly lowers the platform's payoff.

In the model of competition, consider the strategy profile in which platforms never erase information on the path of play, and all players behave in the same way as the strategy profile constructed for Proposition 4. The action of each player straightforwardly extends to nodes in which a platform has deleted information, because the relevant state variable in that strategy profile is $(\sigma_{I,t-1}^2, \sigma_{E,t-1}^2)$. If a platform erases information, it lowers the payoff and increases the consumer's cost of using the platform. Thus, it is optimal for each platform to not erase information. □

F Forward-looking Consumer and Platform with Long-run Commitment: Proof of Proposition 9

This appendix consists of three steps. First, I prove the existence of an equilibrium in which the consumer breaks ties and chooses the "greatest" sequence of activity levels. Second, I prove useful properties of the consumer's value function in her dynamic optimization. Finally, I use these results to prove Propositions 8 and 9.

I prepare notations. Let $\mathcal{A} := [0, a_{max}]^{\mathbb{N}}$ denote the set of all sequences of activity levels between 0 and a_{max}. It is without loss of generality to exclude an activity level strictly above a_{max}. Let $\boldsymbol{a}$ denote a generic element of $\mathcal{A}$, with the t-th coordinate denoted by a_t. Let $\boldsymbol{\Gamma}$ denote the set of all sequences of non-negative real numbers. Let $\boldsymbol{\gamma}$ denote a generic element of $\boldsymbol{\Gamma}$, with the t-th coordinate denoted by γ_t. I consider product topology for $\mathcal{A}$ and $\boldsymbol{\Gamma}$.

F.1 Existence of an Equilibrium

Take any privacy policy $\boldsymbol{\gamma} \in \boldsymbol{\Gamma}$. The consumer's problem is

$$\max_{\boldsymbol{a} \in \mathcal{A}} \sum_{t=1}^{\infty} \delta_C^{t-1} \left[u(a_t) - v \cdot \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{a_s} + \gamma_s}} \right) \right]. \tag{19}$$

For any $\boldsymbol{\gamma} \in \boldsymbol{\Gamma}$, let $\mathcal{A}^*(\boldsymbol{\gamma}) \subset \mathcal{A}$ denote the set of all maximizers in (19).

Lemma 3. *$\mathcal{A}^*(\boldsymbol{\gamma})$ 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 $u(a^*) > 0$ and $-v\sigma_0^2 < 0$, respectively. Define $B := \max(u(a^*), v\sigma_0^2) > 0$. 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, equation (19) is continuous in $\boldsymbol{a}$. For each privacy policy $\boldsymbol{\gamma}$, let $\mathcal{A}^*(\boldsymbol{\gamma}) \subset \mathcal{A}$ denote the set of all maximizers. Berge maximum theorem implies that $\mathcal{A}^*(\boldsymbol{\gamma})$ is non-empty, compact, and upper hemicontinuous. □

Next, I prove some 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 4. *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

$$(\boldsymbol{a} \vee \boldsymbol{b})^n = \begin{cases} \max(a_t, b_t) & \text{if} \quad t \leq n, \\ a_t & \text{if} \quad t > n. \end{cases} \tag{20}$$

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

$$(\boldsymbol{a} \wedge \boldsymbol{b})^n = \begin{cases} \min(a_t, b_t) & \text{if} \quad t \leq n, \\ a_t & \text{if} \quad t > n. \end{cases} \tag{21}$$

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

$$\boldsymbol{b}^n = \begin{cases} b_t & \text{if} \quad t \leq n, \\ a_t & \text{if} \quad t > n. \end{cases} \tag{22}$$

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 $n \in \mathbb{N}$, $U(\boldsymbol{a})$ is supermodular in the first n activity levels, $(a_1, \ldots, a_n) \in \mathbb{R}^n_+$. Thus, $U((\boldsymbol{a} \vee \boldsymbol{b})^n) + U((\boldsymbol{a} \wedge \boldsymbol{b})^n) \geq U(\boldsymbol{a}) + U(\boldsymbol{b}^n)$. Since $U(\cdot)$ is continuous, we can take $n \to \infty$ and obtain $U(\boldsymbol{a} \vee \boldsymbol{b})) + U(\boldsymbol{a} \wedge \boldsymbol{b}) \geq U(\boldsymbol{a}) + U(\boldsymbol{b})$. □

Lemma 5. *There is an $\bar{\boldsymbol{a}} \in \mathcal{A}^*(\gamma)$ such that, for any $\boldsymbol{a} \in \mathcal{A}^*(\gamma)$, $\bar{\boldsymbol{a}} \geq \boldsymbol{a}$.*

Proof. First, Corollary 2 of Milgrom et al. (1994) implies that $\mathcal{A}^*(\boldsymbol{\gamma})$ is a sublattice of $\mathcal{A}$. Since $\mathcal{A}^*(\boldsymbol{\gamma})$ is compact, for each $t \in \mathbb{N}$, the projection of $\mathcal{A}^*(\boldsymbol{\gamma})$ on the t-th coordinate, i.e.,

$$\mathcal{A}_t^*(\boldsymbol{\gamma}) := \left\{a_t \in [0, a^*] : \exists \boldsymbol{a}_{-t} = (a_s)_{s\in\mathbb{N}\setminus\{t\}} \in \mathcal{A}^*(\boldsymbol{\gamma}) \text{ s.t. } (a_t, \boldsymbol{a}_{-t}) \in \mathcal{A}^*(\boldsymbol{\gamma})\right\}, \tag{23}$$

is compact (here, $(a_t, \boldsymbol{a}_{-t})$ is a sequence of activity levels such that the consumer takes a_t in period t and acts according to $\boldsymbol{a}_{-t}$ in other periods). For each $k \in \mathbb{N}$, let $\boldsymbol{a}^k$ denote an optimal policy such that $\boldsymbol{a}^k = \max \mathcal{A}_k^*(\boldsymbol{\gamma})$. Define $\bar{\boldsymbol{a}}^k := \boldsymbol{a}^1 \vee \cdots \vee \boldsymbol{a}^k$. Since $\mathcal{A}^*(\boldsymbol{\gamma})$ is sublattice, for any $k \in \mathbb{N}$, $\bar{\boldsymbol{a}}^k$ maximizes (19). Also, $\bar{\boldsymbol{a}}^k \to \bar{\boldsymbol{a}}$, where $\bar{a}_t = \max \mathcal{A}_k^*(\boldsymbol{\gamma})$ for any $k \in \mathbb{N}$. Since $\mathcal{A}^*(\boldsymbol{\gamma})$ is compact, $\bar{\boldsymbol{a}} \in \mathcal{A}^*(\boldsymbol{\gamma})$. By construction, for any $\boldsymbol{a} \in \mathcal{A}^*(\boldsymbol{\gamma})$, $\bar{\boldsymbol{a}} \geq \boldsymbol{a}$. □

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 5.

Lemma 6. *For each $t \in \mathbb{N}$, $\bar{\boldsymbol{a}}_t(\boldsymbol{\gamma})$ is upper semicontinuous in $\boldsymbol{\gamma}$.*

Proof. By Lemma 3, $\mathcal{A}^*(\boldsymbol{\gamma})$ is upper hemicontinuous. Thus, the set $\mathcal{A}_t^*(\boldsymbol{\gamma})$ of all activity levels for 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. □

Lemma 7. *There exists an equilibrium.*

Proof. The platform's objective is

$$\sum_{t=1}^{\infty} \delta_P^{t-1} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{\bar{a}_s(\boldsymbol{\gamma})} + \gamma_s}}\right). \tag{24}$$

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. $(\boldsymbol{\gamma}^*, \bar{\boldsymbol{a}}(\cdot))$ is an equilibrium. □

F.2 Properties of Consumer Value Function

For each privacy policy $\boldsymbol{\gamma} \in \boldsymbol{\Gamma}$, define

$$V_{\boldsymbol{\gamma}}(\rho) := \sum_{t=1}^{\infty} \delta_C^{t-1} \left[u(\bar{a}_t(\boldsymbol{\gamma})) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho + \sum_{s=1}^t \frac{1}{\frac{1}{\bar{a}_s(\boldsymbol{\gamma})} + \gamma_s}} \right) \right]. \tag{25}$$

$V_{\boldsymbol{\gamma}}(\rho)$ is the consumer's maximum value of the objective starting from the conditional variance $\sigma^2 = \frac{1}{\rho}$.

Lemma 8. *For any* $\boldsymbol{\gamma} \in \boldsymbol{\Gamma}$, $V_{\boldsymbol{\gamma}}(\cdot)$ *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 $\frac{1}{\sigma_{t-1}^2} = \rho$:

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

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

$$V_t^T(\rho) = \max_{a\geq 0} u(a) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a}+\gamma_t}}\right) + \delta_C V_{t+1}^T\left(\rho + \frac{1}{\frac{1}{a} + \gamma_t}\right), \tag{26}$$

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 is trivially decreasing and convex. Suppose that V_{t+1}^T is decreasing and convex. Since $-v \cdot \left(\sigma_0^2 - \frac{1}{\rho+\frac{1}{\frac{1}{a}+\gamma_t}}\right)$ has the same property, $V_t^T(\cdot)$ is also decreasing and convex. Thus, $V_1^T(\cdot)$ is decreasing and convex.

Define $V^\infty(\rho) := \lim_{T\to\infty} V_1^T(\rho)$. $V^\infty(\rho)$ 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 $(\bar{a}_1, \bar{a}_2, \ldots) \in \mathcal{A}^*(\boldsymbol{\gamma})$ denote the optimal policy. For any finite T,

$$V_1^T(\rho) \geq \sum_{s=1}^{T} \delta_C^{s-1} \left(u(\bar{a}_s) - v\left(\sigma_0^2 - \frac{1}{\rho_{s-1} + \frac{1}{\frac{1}{\bar{a}_s}+\gamma_s}}\right)\right). \tag{27}$$

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 $V_1^T(\rho) - \frac{\delta_C^T}{1-\delta_C} v\sigma_0^2 > V(\rho)$. If the consumer in the original infinite horizon problem adopts the T-optimal policy that gives $V_1^T(\rho)$ up to period t, then she can obtain a strictly greater payoff than $V(\rho)$, which is a contradiction. Thus, $V^\infty(\rho) = V(\rho)$.

Suppose that 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, $\lim_{\rho\to\infty} V(\rho) - V(\rho+\Delta) = 0$. □

F.3 Consequences of Previous Lemmas

The following is the proof of Proposition 8.

Proof. Since $\gamma_t = \gamma$ for all t, the value function $V(\cdot)$ satisfies the Bellman equation

$$V(\rho) = \max_{a\geq 0} u(a) - v\left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a}+\gamma}}\right) + \delta_C V\left(\rho + \frac{1}{\frac{1}{a}+\gamma}\right). \tag{28}$$

Lemma 8 implies that $V(\cdot)$ is decreasing and convex. Thus, the maximand in (28) has the increasing differences in (a, ρ). Thus, $\bar{a}(v, \gamma, \rho)$, the greatest maximizer, is increasing in ρ.

Define

$$v^*(\gamma) := \sup\left\{v \in \mathbb{R} : \bar{a}_1(v, \gamma, \rho_0) > 0\right\}, \quad \text{where} \quad \rho_0 = \frac{1}{\sigma_0^2}. \tag{29}$$

$v^*(\gamma)$ is increasing because $\bar{a}_1(v, \gamma, \rho_0)$ is. Suppose to the contrary that there is a sequence $\gamma^n \to \infty$ such that $v^*(\gamma^n) \leq \bar{v}$ for some $\bar{v} < \infty$. Take the consumer with $v > \bar{v}$. Suppose that the consumer takes $a_t = a_{max} = \arg\max_{a\geq 0} u(a)$ for all $t \in \mathbb{N}$. As, $\gamma^n \to \infty$, the consumer's period-t payoff converges to $u(a_{max})$ for each $t \in \mathbb{N}$. As the consumer's objective is continuous in per-period payoffs with product topology, the sum of discounted payoffs converges to $\frac{u(a_{max})}{1-\delta_C} > 0$. This contradicts that, for all n, the consumer with $v > \bar{v}$ should choose $\bar{a}_1(v, \gamma_n, \rho_0) = 0$ and thus $a_t = 0$ for all t. Thus, $\lim_{\gamma\to\infty} v^*(\gamma) = \infty$.

By the identical argument with the case of the myopic consumer, we can conclude that the consumer's activity level is positive and increasing in t if $v < v^*(\gamma)$. This implies $\lim_{t\to\infty} \sigma_t^2 \to 0$, or equivalently, $\lim_{t\to\infty} \rho_t = \infty$ with $\rho_t := \frac{1}{\sigma_t^2}$. By Lemma 8, 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 $\lim_{t\to\infty} \bar{a}_t(v, \gamma, \rho_t) = a_{max}$. Finally, $v > v^*(\gamma)$ implies $\bar{a}_1 = 0$. This implies $\bar{a}_t = 0$ for all $t \in \mathbb{N}$ because the conditional variance does not change. □

The following result states that whenever the change in a privacy policy increases the precisions of signals in some periods, the consumer chooses greater activity levels in other periods.

Lemma 9. *Take any $\boldsymbol{\gamma}, \boldsymbol{\gamma}' \in \boldsymbol{\Gamma}$. Define $\mathcal{T} = \{t \in \mathbb{N} : \gamma_t = \gamma'_t\}$. Suppose that $\frac{1}{\bar{a}_t(\boldsymbol{\gamma})} + \gamma_t \leq \frac{1}{\bar{a}_t(\boldsymbol{\gamma}')} + \gamma'_t$ for all $t \in \mathbb{N} \setminus \mathcal{T}$. Then, $\bar{a}_t(\boldsymbol{\gamma}) \geq \bar{a}_t(\boldsymbol{\gamma}')$ for all $t \in \mathcal{T}$.*

Proof. Let $\boldsymbol{\beta}$ be any one of $\boldsymbol{\gamma}$ and $\boldsymbol{\gamma}'$. I decompose the consumer's problem (19) into two steps. First, given any $(a_t)_{t \notin \mathcal{T}}$, the consumer chooses $(a_t)_{t \in \mathcal{T}}$ to maximize the following hypothetical objective function:

$$\sum_{t=1}^{\infty} \delta_C^{t-1} \left[\mathbf{1}_{\{t \in \mathcal{T}\}} u(a_t) - v \cdot \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^{t} \frac{1}{\frac{1}{a_s} + \beta_s}} \right) \right]. \tag{30}$$

Note that the consumer does not receive a benefit of $u(a_t)$ in period $t \notin \mathcal{T}$. This leads to a mapping that maps any $(a_t)_{t \notin \mathcal{T}}$ to the (greatest) optimal choice of $(a_t)_{t \in \mathcal{T}}$. In the second step, the consumer chooses $(a_t)_{t \notin \mathcal{T}}$ to maximize her original objective, taking the mapping $(a_t)_{t \notin \mathcal{T}} \mapsto (a_t)_{t \in \mathcal{T}}$ as given.

For any $t \notin \mathcal{T}$, a_t affects (30) only through $\frac{1}{a_t} + \gamma_t$, because $\mathbf{1}_{\{t \in \mathcal{T}\}} = 0$. Moreover, the same argument as in the proof of Lemma 4 implies that (30) is supermodular in $\left((a_t)_{t \in \mathcal{T}}, \left\{ \left(\frac{1}{a_s} + \gamma_s \right)^{-1} \right\}_{s \notin \mathcal{T}} \right)$. 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 $t \in \mathbb{N} \setminus \mathcal{T}$, then $\bar{a}_t(\boldsymbol{\gamma}) \geq \bar{a}_t(\boldsymbol{\gamma}')$ for all $t \in \mathcal{T}$. □

F.4 $\lim_{t\to\infty} a_t^* = a_{max}$ and $\lim_{t\to\infty} \sigma_t^2 = 0$: Proof of Proposition 9

Proof. Let $\boldsymbol{\gamma}^*$ denote the equilibrium privacy policy, and let $\boldsymbol{a}^*$ denote the equilibrium activity levels. First, I show $\lim_{t\to\infty} \sigma_t^2 = 0$. Suppose to the contrary that $\lim_{t\to\infty} \sigma_t^2 \neq 0$. As σ_t^2 is decreasing, $\lim_{t\to\infty} \sigma_t^2 > 0$ exists. This implies $\frac{1}{a_t^*} + \gamma_t^* \to \infty$, which I prove to be a contradiction.

By Proposition 8, there exists a $\hat{\gamma}$ such that $v^*(\hat{\gamma}) > v$. That is, if the platform commits to $\gamma_t = \hat{\gamma}$ for all t, then the consumer chooses some $\bar{a} > 0$ in $t = 1$. Define $B := \frac{1}{\bar{a}} + \hat{\gamma}$. Consider T^* such that, for all $t \geq T^*$, $\frac{1}{a_t^*} + \gamma_t^* > B$. Suppose that the platform replaces γ_t^* with $\hat{\gamma}$ for all $t \geq T^*$, and commits to such a new policy ex ante. Take any period $t \geq T^*$. Since the consumer's activity levels after T^* solve the Bellman equation with the "initial state" of $\rho = \frac{1}{\sigma_{T^*-1}^2} \geq \frac{1}{\sigma_0^2}$, the consumer chooses an activity level greater than $\bar{a} > 0$ after period T^*. Thus, the variance of the noise ε_t in the signal s_t is at most $\frac{1}{\bar{a}} + \hat{\gamma} < \frac{1}{a_t^*} + \gamma_t^*$. Thus, this change in the privacy policy strictly increases the platform's profit in any period $t \geq T^*$. By Lemma 9, this change also increases the consumer's activity level for any period $t < T^*$. Thus, as a result of the deviation, the platform's payoffs increase in all periods and strictly increase in some periods, which contradicts $\boldsymbol{\gamma}^*$ being optimal. Thus, in equilibrium, $\lim_{t\to\infty} \sigma_t^2 = 0$ holds. Finally, $\lim_{t\to\infty} \sigma_t^2 = 0$ implies $a_t^* \to a_{max}$: Otherwise, there is a convergent subsequence $a_{t(n)}^* \to a' < a_{max}$, however, the consumer could be strictly better off by choosing a_{max}, due to Lemma 8. □

G Forward-looking Consumer and Platform with Short-run Commitment: Proof of Proposition 10

Proof. The existence of an equilibrium follows from Harris (1985). To prove the claim, first, suppose to the contrary that, in some equilibrium, $\lim\limits_{t\to\infty} \sigma_t^2 > 0$. The inequality implies $\lim_{t\to\infty} \frac{1}{a_t^*} + \gamma_t^* = \infty$, which implies $\lim_{t\to\infty} a_t^* = 0$. Point 1 of Assumption 2 implies that there is a $T \in \mathbb{N}$ such that for all $t \geq T$, $a_t^* = 0$. Suppose that the platform sets a privacy level $\bar{\gamma}$ in period T. If the consumer chooses $a_t = 0$ for all $t \geq T$, her continuation payoff is $-\frac{v}{1-\delta_C}\left(\sigma_0^2 - \sigma_{T-1}^2\right)$. If she chooses $a_T = a_{max} > 0$ and $a_s = 0$ for all $s \geq T+1$, then her continuation payoff is $u(a_{max}) - \frac{v}{1-\delta_C}\left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_{T-1}^2} + \frac{1}{\frac{1}{a_{max}} + \bar{\gamma}}}\right)$. Thus, the consumer prefers the latter if and only if

$$
\begin{aligned}
& u(a_{max}) - \frac{v}{1-\delta_C}\left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_{T-1}^2} + \frac{1}{\frac{1}{a_{max}} + \bar{\gamma}}}\right) > -\frac{v}{1-\delta_C}\left(\sigma_0^2 - \sigma_{T-1}^2\right) \\
\iff & u(a_{max}) - \frac{v}{1-\delta_C}\left[\frac{1}{\frac{1}{\sigma_{T-1}^2}\left(\frac{1}{\sigma_{T-1}^2}\left(\frac{1}{a_{max}} + \bar{\gamma}\right) + 1\right)}\right] > 0. \qquad (31)
\end{aligned}
$$

Moreover,

$$
u(a_{max}) - \frac{v}{1-\delta_C}\left[\frac{1}{\frac{1}{\sigma_{T-1}^2}\left(\frac{1}{\sigma_{T-1}^2}\left(\frac{1}{a_{max}} + \bar{\gamma}\right) + 1\right)}\right] \geq u(a_{max}) - \frac{v}{1-\delta_C}\left[\frac{1}{\frac{1}{\sigma_0^2}\left(\frac{1}{\sigma_0^2}\left(\frac{1}{a_{max}} + \bar{\gamma}\right) + 1\right)}\right] > 0,
$$

where the last inequality comes from Point 2 of Assumption 2. Since (31) holds, $a_t = 0\ \forall t \geq T$ cannot be an optimal continuation strategy following $\gamma_T = \bar{\gamma}$. In other words, the consumer chooses $a_s > 0$ for at least one $s \geq T$. Thus, in period T, the platform must be choosing γ_T such that $a_s > 0$ for some $s \geq T$, which contradicts $a_t^* = 0$ for all $t \geq T$.

Second, suppose to the contrary that $\lim_{t\to\infty} a_t^* = a_{max}$ fails, which implies that there is a strictly increasing sequence $(t(k))_{k\in\mathbb{N}}$ of natural numbers such that $a_{t(k)}^* < a_{max}$ for all k. Take any k. Suppose that the consumer deviates in period k so that $a_t = a_{max}$ for all $t \geq t(k)$. Then, her continuation payoff is at least

$$\frac{u(a_{max})}{1-\delta_C} - v\sum_{n=1}^{\infty}\delta_C^{n-1}\left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_{t(k)-1}^2} + na_{max}}\right), \tag{32}$$

which occurs if the platform sets $\gamma_t = 0$ for all $t \geq t(k)$. If the consumer follows the equilibrium strategy, her continuation payoff is at most

$$\sum_{t=t(k)}^{\infty}\delta^{t-t(k)}u(a_t^*) - \frac{v}{1-\delta_C}\left(\sigma_0^2 - \sigma_{t(k)}^2\right). \tag{33}$$

Consider the difference

$$\underbrace{\frac{u(a_{max})}{1-\delta_C} - \sum_{t=t(k)}^{\infty}\delta^{t-t(k)}u(a_t^*)}_{=X_k} - \underbrace{\left\{v\sum_{n=1}^{\infty}\delta_C^{n-1}\left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_{t(k)-1}^2} + na^*}\right) - \frac{v}{1-\delta_C}\left(\sigma_0^2 - \sigma_{t(k)}^2\right)\right\}}_{=Y_k}. \tag{34}$$

First, $X_k \geq u(a_{max}) - u(a') > 0$, where a' is the second largest element of A. Second, $\lim_{k\to\infty} Y_k = 0$, because both the first and second terms in Y_k converge to $\frac{v}{1-\delta_C}\sigma_0^2$ because of $\lim_{k\to\infty}\sigma_{t(k)}^2 = 0$. Thus, (34) is positive for a sufficiently large k. In other words, the consumer has a profitable deviation in period $t(k)$ for a large k, which is a contradiction. □

H Omitted Proofs for Section 7

H.1 Heterogeneous Consumers: Proof of Proposition 11

Take any equilibrium with $(a_t^*(v), \sigma_t^2(v), \gamma_t^*)_{t\in\mathbb{N}, v\in V}$. 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

$$\Delta_t := \frac{1}{1-\delta_P}\sum_{v\in V}\alpha_v\left[\sigma_0^2 - \sigma_\infty^2(v)\right] - \frac{1}{1-\delta_P}\sum_{v\in V}\alpha_v\left[\sigma_0^2 - \sigma_{t-1}^2(v)\right]. \tag{35}$$

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

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

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 $\frac{\alpha_v B}{1-\delta_P} > \Delta_t$, or equivalently,

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

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 Proposition 2.

H.2 General Privacy Cost: Proofs of Propositions 13, 14, and 15

Proof of Proposition 13. Consider any equilibrium. In period t, the consumer chooses a positive activity level if

$$\max_{a\geq 0} u(a) - v\left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a}+\gamma_t^*}}\right) \geq -\alpha v\left(\sigma_0^2 - \sigma_{t-1}^2\right)$$

$$\iff \max_{a\geq 0} u(a) - v\left(\alpha\sigma_{t-1}^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a}+\gamma_t^*}}\right) \geq (1-\alpha)v\sigma_0^2.$$

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., $\alpha = 1$). Then, let $y_1 = \frac{1}{\hat{a}} + \hat{\gamma}$. Define $f(\alpha, x, y) := \alpha x - \frac{1}{\frac{1}{x}+\frac{1}{y}}$. 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 $\alpha\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2}+\frac{1}{y_1}} > 0$. Take any $\varepsilon > 0$ and let $\alpha' \in [0, 1)$ satisfy

$$u(\hat{a}_1) - v\left(\alpha'\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \frac{1}{\frac{1}{\hat{a}_1}+\hat{\gamma}_1+\varepsilon}}\right) \geq (1-\alpha')v\sigma_0^2. \tag{36}$$

Such an α' exists. Indeed, if $\varepsilon = 0$, then the above inequality holds for $\alpha' = 1$. Thus, if $\varepsilon > 0$, then the inequality holds for some $\alpha' < 1$. Also, let $\hat{\alpha} < 1$ satisfy $\hat{\alpha}\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2}+\frac{1}{\frac{1}{\hat{a}_1}+\hat{\gamma}_1+\varepsilon}} > 0$. Define $\alpha^* = \max(\alpha', \hat{\alpha}) < 1$. Now, take any $\alpha \in [\alpha^*, 1]$. Then,

$$\begin{aligned}
(36) &\Rightarrow u(\hat{a}_1) - v\left(\alpha^*\sigma_{t-1}^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{\hat{a}_1}+\hat{\gamma}_1+\varepsilon}}\right) \geq (1-\alpha^*)v\sigma_0^2 \\
&\Rightarrow u(\hat{a}_1) - v\left(\alpha\sigma_{t-1}^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{\hat{a}_1}+\hat{\gamma}_1+\varepsilon}}\right) \geq (1-\alpha)v\sigma_0^2 \\
&\Rightarrow \max_{a\geq 0} u(a) - v\left(\alpha\sigma_{t-1}^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\frac{1}{a}+\hat{\gamma}_1+\varepsilon}}\right) \geq (1-\alpha)v\sigma_0^2 \qquad (37)
\end{aligned}$$

(37) implies that, in any period, if the platform sets $\gamma_t = \hat{\gamma}_1 + \varepsilon$, then the consumer chooses $a_t > 0$. Recall that $\hat{a}_1 > 0$ is the optimal positive activity level given $\hat{\gamma}_1$ at σ_0^2 in $t = 1$. Thus, $a_t \geq \hat{a}_1$ 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

$$\max_{a\geq 0} u(a) - v\left(\alpha\sigma_{t-1}^2 - \frac{1}{\frac{1}{\sigma_{t-1}^2}+\frac{1}{\frac{1}{a}+\gamma_t}}\right) \geq (1-\alpha)v\sigma_0^2$$

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 Proposition 2, $\sigma_t^2 \to 0$ implies $a_t^* \to a_{max}$ and $\gamma_t^* \to 0$. □

Proof of Proposition 14. I adopt the notations in the proof of Proposition 4. 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:

$$\arg\max_{a\geq 0} u(a) - v[\sigma_0^2 - \sigma_{t,k}^2(\gamma_t^k, a|\sigma_{t-1,k}^2)] - \alpha v[\sigma_0^2 - \sigma_{t-1,-k}^2]$$
$$\geq \arg\max_{a\geq 0} u(a) - v[\sigma_0^2 - \sigma_{t,-k}^2(\gamma_t^{-k}, a|\sigma_{t-1,-k}^2)] - \alpha v[\sigma_0^2 - \sigma_{t-1,k}^2],$$

and

$$\arg\max_{a\geq 0} u(a) - v[\sigma_0^2 - \sigma_{t,k}^2(\gamma_t^k, a|\sigma_{t-1,k}^2)] - \alpha v[\sigma_0^2 - \sigma_{t-1,-k}^2] \geq -\alpha v[\sigma_0^2 - \sigma_{t-1,k}^2] - \alpha v[\sigma_0^2 - \sigma_{t-1,-k}^2].$$

These inequalities are equivalent to

$$\arg\max_{a\geq 0} u(a) - \underbrace{v\left[\alpha\sigma_{t-1,k}^2 - \frac{1}{\frac{1}{\sigma_{t-1,k}^2}+\frac{1}{\frac{1}{a}+\gamma_t^k}}\right]}_{\text{(A)}} \geq \arg\max_{a\geq 0} u(a) - \underbrace{v\left[\alpha\sigma_{t-1,-k}^2 - \frac{1}{\frac{1}{\sigma_{t-1,-k}^2}+\frac{1}{\frac{1}{a}+\gamma_t^{-k}}}\right]}_{\text{(B)}} \tag{38}$$

and

$$\arg\max_{a\geq 0} u(a) - \underbrace{v\left[\alpha\sigma_{t-1,k}^2 - \frac{1}{\frac{1}{\sigma_{t-1,k}^2}+\frac{1}{\frac{1}{a}+\gamma_t^k}}\right]}_{\text{(A)}} \geq (1-\alpha)v\sigma_0^2. \tag{39}$$

Let $a(\bar{\gamma}) := a^*(\bar{\gamma}, \sigma_0^2)$ in (3). By the same argument as Proposition 13, 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 that 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 4 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 (38) and (39). The rest of the proof follows the proof of Proposition 4. $\square$

Proof of Proposition 15. Consider the consumer's problem in period t. Given the conditional variance σ^2 at the end of period $t-1$ and the privacy level γ in period t, the consumer chooses a to maximize $U(a, \gamma, \sigma^2) := u(a) - C\left(\frac{1}{\frac{1}{\sigma^2} + \frac{1}{\frac{1}{a}+\gamma}}\right)$. It holds that

$$\frac{\partial U}{\partial a} = u'(a) + C'\left(\frac{1}{\frac{1}{\sigma^2} + \frac{1}{\frac{1}{a}+\gamma}}\right) \cdot \frac{1}{\left(\frac{1}{\sigma^2}(1+\gamma a) + a\right)} \geq u'(a) - B \cdot \frac{1}{\left(\frac{1}{\sigma^2}(1+\gamma a) + a\right)}, \quad (40)$$

where $B := \sup_{x \in [0, \sigma_0^2]} |C'(x)| < \infty$. If $\lim_{t\to\infty} \sigma_t^2 > 0$, then $\lim_{t\to\infty} \frac{1}{a_t^*} + \gamma_t^* = \infty$. Consider a hypothetical payoff function

$$U_B(a, \gamma, \sigma^2) = u(a) - B \cdot \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma^2} + \frac{1}{\frac{1}{a}+\gamma}}\right).$$

(40) implies $\frac{\partial U}{\partial a} \geq \frac{\partial U_B}{\partial a}$. Take any γ' such that $a_B^*(\gamma', \sigma^2) := \max\{\arg\max_{a \geq 0} U_B(a, \gamma', \sigma_0^2)\} > 0$. Then, for any $\sigma^2 \leq \sigma_0^2$, $a^*(\gamma', \sigma^2) \geq a_B^*(\gamma', \sigma^2) \geq a_B^*(\gamma', \sigma_0^2) > 0$. Take T such that for all $t \geq T$, $\frac{1}{a_t^*} + \gamma_t^* \geq \frac{1}{a_B^*(\gamma', \sigma_0^2)} + \gamma'$ Then, the platform can achieve a lower $\frac{1}{a_t} + \gamma_t$ for any $t \geq T$ by replacing γ_t^* with γ', which is a contradiction. A similar argument implies that $\lim_{t\to\infty} a_t^* = a_{max}$. $\square$