

Dynamic Privacy Choices

by Shota Ichihashi



Banking and Payments Department
Bank of Canada
sichihashi@bankofcanada.ca

Bank of Canada staff working papers provide a forum for staff to publish work-in-progress research independently from the Bank's Governing Council. This research may support or challenge prevailing policy orthodoxy. Therefore, the views expressed in this paper are solely those of the authors and may differ from official Bank of Canada views. No responsibility for them should be attributed to the Bank.

Acknowledgements

For valuable comments, I thank Jason Allen, Byung-Cheol Kim, David Miller, Michael Sockin, Takuo Sugaya, Catherine Tucker, Lei Xu and Kai Hao Yang, as well as seminar and conference participants at the Bank of Canada, Duke University, the University of North Carolina, the Federal Reserve Bank of New York, New York University, Peking University, Queen's University, Simon Business School at the University of Rochester, the Privacy Workshop at Princeton University, the Canadian Economic Theory Conference 2020, the Twenty-First ACM Conference on Economics and Computation, SITE 2020, and the 2021 conference of the Society for the Advancement of Economic Theory. The opinions expressed in this article are the author's own and do not reflect the views of the Bank of Canada.

Abstract

I study a dynamic model of consumer privacy and platform data collection. In each period, consumers choose their level of platform activity. Greater activity generates more information about the consumer, thereby increasing platform profits. When the platform can commit to the future privacy policy, it collects information by committing to gradually decrease the level of privacy protection. In the long run, consumers lose privacy and receive low payoffs but choose high activity levels. In contrast, the platform with weaker commitment power may attain the commitment outcome or fail to collect any data, depending on consumer expectations regarding future privacy protection.

Topic: Economic models

JEL codes: D, D8, D82, D83

1 Introduction

Online platforms, such as Amazon, Facebook, Google, and Uber, analyze user activities and collect a large amount of data. The data collection may improve their services and benefit consumers. At the same time, it raises concerns for consumers and policymakers, highlighted by recent privacy scandals and data breaches, such as the Cambridge Analytica data scandal.¹

As an example, consider a consumer (she) and a social media platform (it). The consumer writes posts and reads news on the platform. The platform analyzes her activity and collects data such as her 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 service provided by the platform. On the other hand, she may value her privacy and be concerned about the risk of data leakage, identity theft, and price or non-price discrimination. Such risks are the “privacy costs” of using the platform. If the consumer anticipates a high privacy cost, she may use the platform less actively or may not join it. The platform can influence her decision through its privacy policy— e.g., Facebook committed to not use cookies to track users.²

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

The main idea is that the consumer has a decreasing marginal privacy cost—i.e., when the consumer has already lost some privacy, she faces even a lower marginal privacy cost of using the same platform. For example, if Google already knows a lot about a consumer, she might not care

¹In a survey conducted by Pew Research Center, “a majority of Americans report being concerned about the way their data is being used by companies.” See <https://www.pewresearch.org/internet/2019/11/15/americans-and-privacy-concerned-confused-and-feeling-lack-of-control-over-their-personal->

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

about letting Google Maps track her location today. In an extreme case, if the platform knows everything, the consumer faces a marginal privacy cost of zero, because her activity no longer affects what the platform knows about her type.

The main finding is that because of the decreasing marginal cost of losing privacy, a farsighted consumer eventually gives up all of her privacy, even when she anticipates future privacy choices of the platform. To induce such an outcome, in early periods the platform commits to not collect too much data. By doing so, it encourages the consumer to use the service and generate information when she has not yet lost her privacy. At the same time, the platform always collects full information in the long run. By committing to do so, the platform not only increases future profits but encourages the consumer to generate more information in early periods: Indeed, the consumer who anticipates future privacy loss faces a lower marginal privacy cost today and uses the platform more actively. As a result, in the long run, the consumer loses privacy and chooses a high activity level, regardless of discount factors and the consumer's value of privacy.

I first establish the above result when the platform's revenue depends only on information. However, the result is relevant even if revenue depends on the consumer's activity. We might think that if the platform's revenue comes mainly from consumer activity, it should refrain from collecting data to encourage the activity of consumers who have privacy concerns. On the contrary, the platform can collect data, reduce the consumer's marginal privacy cost, and increase her future activity levels. As a result even if the platform's revenue comes mainly from consumer activity, the equilibrium may still entail a low level of privacy in the long run.

I then study the role of the platform's commitment power regarding its future privacy choices. The baseline model assumes that the platform can commit to future privacy levels, in which case it collects information by committing to high privacy protection in early periods. Under a certain condition, the platform can attain the same outcome as long as it has one-period commitment power. However, weaker commitment power may create multiple equilibria: There can also be an equilibrium in which the platform offers the highest privacy protection and fails to collect any information. Such an equilibrium captures the platform's Coasian commitment problem.

The paper has implications for consumer privacy. First, the consumer's long-run behavior (in the equilibrium with data collection) is consistent with the so-called privacy paradox: Consumers express concern about their privacy, but actively share data with third parties ([Acquisti et al., 2016](#)).

The platform’s equilibrium strategy rationalizes how online platforms, such as Facebook, seem to have expanded the scope of data collection. Second, my results clarify the role of commitment and expectation in data collection: Depending on consumers’ expectation about their future privacy, the platform may collect data when consumers highly value their privacy, or it may fail to collect data when consumers do not much value their privacy.

The rest of the paper is as follows. [Section 2](#) discusses related literature, and [Section 3](#) presents the model. [Section 4](#) considers the platform with long-run commitment power, and [Section 5](#) considers the platform with one-period commitment power. [Section 6](#) considers extensions, including the time-varying type of the consumer. All proofs are in the Appendix.

2 Related Literature

The paper relates to three strands of literature. First, it relates to the growing work on consumer privacy and data collection. Several papers model a platform that collects data in exchange for money ([Acemoglu et al., 2019](#); [Bergemann et al., 2019](#); [Choi et al., 2019](#); [Easley et al., 2018](#); [Ichihashi, 2020](#)). In these papers the data on some consumers reveal information about others. This “data externality” could lower consumers’ private costs of providing data, leading to an inefficiently high level of data sharing at equilibrium.³ In my paper, the consumer’s cost of generating information is decreasing in the stock of data she provided in the past and the amount of data the platform will collect in the future. To isolate the main economic force of this paper from data externality, I study a model with one consumer. Several papers, such as [Fainmesser et al. \(2019\)](#), [Jullien et al. \(2018\)](#), and [Argenziano and Bonatti \(2020\)](#), study the design of privacy regulation or policy. I add to this literature by studying the dynamics of data collection and the role of the platform’s commitment power.

Second, the paper relates to the literature on strategic manipulation of information. In particular, my model relates to career concern models, which originated with [Holmström \(1999\)](#). In Holmström’s model, a young worker, whose ability has not yet been revealed, works hard to influence the market’s belief. In my model, a consumer who has not yet lost privacy uses the platform

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

less actively to generate less information. Over time, the information about the consumer and the worker are revealed, and they have lower incentives to engage in signal jamming. Despite this connection, the two signal jamming activities are different. In career concern models, a worker cares about the market’s belief about the expected quality. Also, the market wants the worker to engage in signal jamming (i.e., higher effort), which creates a trade-off between learning about ability and motivating high efforts (e.g., [Hörner and Lambert 2018](#)). In contrast, I abstract away from belief manipulation (e.g., strategically browsing websites to convince a platform that a consumer has certain characteristics) and assume that the consumer cares about the platform’s posterior variance about their type. Also, the platform wants the consumer to engage less in signal jamming. Thus, the platform prefers to collect information not only to increase profit today, but also to induce high activity levels in the future. Many of my results stem from this complementarity between data collection and consumer activity, which is absent in career concern models. Recent papers also study how the agent’s incentive to manipulate information distorts their behavior ([Frankel and Kartik, 2019b,a](#); [Bonatti and Cisternas, 2020](#); [Argenziano and Bonatti, 2020](#); [Ball, 2020b](#)). The common theme of these papers, which also appears in this paper, is that a certain information policy can mitigate the sender’s incentive to manipulate information.

Finally, the paper relates to the literature on dynamic information design (e.g., [Ely \(2017\)](#), [Smolin \(forthcoming\)](#), and [Ball \(2020a\)](#)).⁴ This literature typically considers a policy that provides information over time to influence the agent’s behavior. In particular, [Ball \(2020a\)](#) highlights the importance of making future information provision contingent on past actions. In contrast, I study the dynamic collection of information when the agent’s action generates a signal about their type. Decreasing marginal privacy cost, which could lower the value of the designer’s commitment power, has no counterpart in this literature.⁵

3 Model

I study a dynamic game between a consumer (she) and a platform (it). The consumer uses the platform’s service to receive benefits, but the usage generates information about her time-invariant

⁴See, e.g., [Kamenica \(2019\)](#) for a survey of static and dynamic information design.

⁵[Renault et al. \(2017\)](#) show that a greedy information disclosure policy can be optimal for the sender when the receivers of information are short-lived.

type. The platform chooses the level of privacy protection, which is the amount of noise added to the information generated. The platform provides the service for free and monetizes information. I model payoffs in a reduced-form way: In the baseline model, the platform prefers more information and the consumer prefers less information to be collected. [Appendix A](#) microfounds such preferences, assuming that the platform sells data to third-party sellers that price discriminate the consumer in the product market.

The formal description is as follows. Time is discrete and infinite, indexed by $t \in \mathbb{N}$. The consumer's type X is drawn from a normal distribution $\mathcal{N}(0, \sigma_0^2)$. The type is realized before $t = 1$ and fixed over time. The consumer does not observe X .⁶ The platform does not observe X either but receives signals about it.

In each period $t \in \mathbb{N}$, the consumer chooses an *activity level* a_t from a finite set $A \subset \mathbb{R}_+$ such that $\min A = 0$ and $a_{max} := \max A > 0$. The platform then observes a_t and a signal $s_t = X + \varepsilon_t$, where $\varepsilon_t \sim \mathcal{N}\left(0, \frac{1}{a_t} + \gamma_t\right)$. The consumer does not observe the signal.⁷ A higher a_t reduces the variance of ε_t and makes s_t more informative about X . Thus, a_t captures a consumer's online activity that generates data, such as browsing and posting content on social media. For a fixed a_t , the informativeness of the signal decreases in $\gamma_t \in \overline{\mathbb{R}}_+ := \mathbb{R}_+ \cup \{\infty\}$, which is the *privacy level* of the platform in period t . A higher γ_t implies the platform offers higher privacy protection. If $a_t = 0$ or $\gamma_t = \infty$, signal s_t is totally uninformative. Random variables X and $(\varepsilon_t)_{t \in \mathbb{N}}$ are mutually independent.

The payoffs are as follows. Suppose that the consumer has chosen activity levels $\mathbf{a}_t = (a_1, \dots, a_t) \in A^t$ and the platform has chosen privacy levels $\boldsymbol{\gamma}_t = (\gamma_1, \dots, \gamma_t) \in \overline{\mathbb{R}}_+^t$ up to period t . At the end of period t , the platform receives a payoff of $\sigma_0^2 - \sigma_t^2(\mathbf{a}_t, \boldsymbol{\gamma}_t) \geq 0$, where $\sigma_t^2(\mathbf{a}_t, \boldsymbol{\gamma}_t)$ is the posterior variance of X given $(\mathbf{a}_t, \boldsymbol{\gamma}_t)$ and Bayes' rule.⁸ A small $\sigma_t^2(\mathbf{a}_t, \boldsymbol{\gamma}_t)$ means that the platform can accurately estimate the consumer's type, or equivalently, the consumer has the low stock of privacy.

⁶Even if the consumer privately observes X , all results hold with respect to a pooling equilibrium in which consumers of all types choose the same activity level after any history. Such an equilibrium exists because the payoff of each player does not depend on a realization of X . Unobservable X simplifies exposition without changing the results.

⁷All the results continue to hold even if signals are public, because the payoff of each player does not depend on the realization of a signal.

⁸The equivalent formulation is that the platform observes (a_t, s_t) , chooses $b_t \in \mathbb{R}$, and obtains an ex post payoff of $-(X - b_t)^2$, which the platform does not observe. Writing the payoffs in terms of σ_t^2 simplifies exposition. See [Acemoglu et al. \(2019\)](#) for further discussion.

For any t and $\tau \leq t$, $\sigma_t^2(\mathbf{a}_t, \gamma_t)$ is decreasing in a_τ , increasing in γ_τ , and independent of s_τ .⁹ Where it does not cause confusion, I write $\sigma_t^2(\mathbf{a}_t, \gamma_t)$ as σ_t^2 . The platform discounts future payoffs with discount factor $\delta_P \in (0, 1)$.

The consumer's flow payoff in period t is $U(\mathbf{a}_t, \gamma_t) := u(a_t) - v \cdot [\sigma_0^2 - \sigma_t^2(\mathbf{a}_t, \gamma_t)]$. The first term $u(a_t)$ is her gross benefit of using the platform. We only assume $u(a)$ is strictly increasing in $a \in A$ and $u(0) = 0$.¹⁰ The second term $v \cdot [\sigma_0^2 - \sigma_t^2(\mathbf{a}_t, \gamma_t)]$ is a *privacy cost*, which captures the negative impact of data collection on the consumer. The parameter $v \in \mathbb{R}_{++}$ captures her value of privacy; it is exogenous and commonly known to the consumer and the platform. The consumer discounts future payoffs with discount factor $\delta_C \in [0, 1)$. I normalize the payoffs so that if $a_t = 0$ for all t , the platform and the consumer obtain zero payoffs in all periods.¹¹

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

I study two games that differ in the timing of moves. One is the game of *long-run commitment*. In this game, before $t = 1$, the platform commits to a *privacy policy* $\gamma = (\gamma_1, \gamma_2, \dots) \in \overline{\mathbb{R}}_+^\infty$, which is publicly observable. Then, in each period $t \in \mathbb{N}$ the consumer chooses a_t , and the platform learns about her type based on the realized signal. In this game, the platform moves only before $t = 1$. The other is the game of *one-period (or short-run) commitment*, in which the platform and the consumer move sequentially in every period: At the beginning of each period t , the platform sets γ_t . After observing γ_t , the consumer chooses a_t . Then, the platform observes the signal, and the game proceeds to period $t + 1$. In this case, the platform can commit to a privacy level only for one period.

The solution concept is perfect Bayesian equilibrium in which the platform's posterior variance is given by $\sigma_t^2(\mathbf{a}_t, \gamma_t)$ after any history (\mathbf{a}_t, γ_t) .¹² Additionally, under long-run commitment I focus

⁹Throughout the paper, "increasing" means "non-decreasing." Similar conventions apply to "decreasing," "higher," "lower," and so on.

¹⁰As a result, we can more generally assume that the variance of noise ε_t is $\frac{1}{g(a_t)} + \gamma_t$ for any strictly increasing $g(\cdot)$ such that $g(0) = 0$; we can redefine $g(a)$ as an activity level.

¹¹While I do not explicitly model the consumer's participation decision, we may interpret $a_t = 0$ for all t as non-participation.

¹²This caveat pins down the payoff-relevant component of the platform's belief (i.e., σ_t^2) even off the equilibrium path. Because σ_t^2 is determined by the past observable outcome (\mathbf{a}_t, γ_t) , I omit the platform's belief from the description of equilibrium.

on equilibrium in which the consumer chooses the largest activity levels.¹³ For clarity, under long-run commitment, I use “optimal policy” for the platform’s equilibrium strategy.

3.1 Discussion of Assumptions

Modeling privacy. The model adopts the Gaussian signal structure and assumes that each player’s payoff depends on the platform’s belief only through residual variance σ_t^2 . This assumption excludes a consumer who cares about the first-order moment of the platform’s belief, e.g., consumers who want a lending platform to believe that their expected quality as a borrower is high.¹⁴ However, the assumption renders the model tractable: We can treat the game as that of perfect information, because σ_t^2 deterministically evolves as a function of past (a_t, γ_t) . The assumption also enables us to model the decreasing marginal privacy cost cleanly, because we can capture the stock of privacy or the amount of data as a scalar. [Section 6.3](#) maintains the Gaussian assumption and shows that all the results continue to hold if the variance of a noise ε_t is $\sigma_\varepsilon^2(a_t, \gamma_t)$ that satisfies certain conditions including $\frac{\partial^2 \sigma_\varepsilon^2}{\partial a \partial \gamma} \geq 0$, i.e., a higher privacy protection makes the variance of the noise less responsive to activity.

Consumer’s privacy preference. The privacy cost $v(\sigma_0^2 - \sigma_t^2)$ captures monetary or nonmonetary reasons why a consumer wants a platform to have less information—e.g., consumers may intrinsically value their privacy, or consider the risk of data breach and discrimination by third parties ([Kummer and Schulte, 2019](#); [Lin, 2019](#); [Tang, 2019](#)). The formulation would be relevant, for example, when consumers intrinsically want the platform to know less about their characteristics. As an example, [Appendix A](#) derives the privacy cost function in the context of third-degree price discrimination with linear demand.

Persistence of information. The baseline model assumes that the consumer’s type is persistent. This assumption enables the platform to learn about the consumer’s type over time. As a result, the consumer’s incentive depends on the stock of information collected in the past, which shapes the

¹³[Appendix C](#) shows that under any privacy policy, the consumer has an optimal sequence of activity levels that has greater a_t for all t than any other optimal sequence.

¹⁴To see how relaxing this assumption could complicate the analysis, suppose that the consumer has a type-dependent utility $u(a, X)$ and privately observes X . The model could then divert from the Gaussian setting because a_t could signal X . The model also requires us to keep track of the beliefs of the platform and the consumer separately after the consumer deviates, which further complicates the analysis of the consumer’s incentive to manipulate the platform’s belief.

equilibrium dynamics. From the consumer’s perspective, the assumption implies that the consumer cannot “delete” data collected in the past: Even if the consumer stops using the platform from period T on (i.e., $a_t = 0$ for all $t \geq T$), the consumer still incurs a privacy cost of $-v(\sigma_0^2 - \sigma_T^2)$ in any period $t \geq T$. The assumption reflects the difficulty of deleting data, which is referred to as “data persistence” (Tucker, 2018). For instance, suppose a platform collects personal information and shares it with third parties. The consumer may then face a risk of discrimination or malicious targeting even outside of the platform. In another example, if a consumer inadvertently discloses information to other users, she may incur a psychological cost because other users know the information. Such a cost would persist even when the consumer is not active on the platform. Section 6.2 studies an extension in which the consumer’s type varies over time.

4 Optimal Policy Under Long-Run Commitment

I begin by studying the game of long-run commitment. I first present a result under a stationary privacy policy, then study the outcome of the entire game. Given the platform’s information in the previous period and (a_t, γ_t) , the posterior variance evolves as follows.¹⁵

$$\sigma_t^2(\mathbf{a}_t, \gamma_t) = \frac{1}{\frac{1}{\sigma_{t-1}^2(\mathbf{a}_{t-1}, \gamma_{t-1})} + \frac{1}{a_t + \gamma_t}}. \quad (1)$$

Thus the consumer’s privacy cost in period t is

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

Define the privacy cost function as

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

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

¹⁵If $x|\mu \sim \mathcal{N}(\mu, \sigma^2)$ and $\mu \sim \mathcal{N}(\mu_0, \sigma_0^2)$, then $\mu|x \sim \mathcal{N}\left(\frac{\sigma_0^2}{\sigma^2 + \sigma_0^2}x + \frac{\sigma^2}{\sigma^2 + \sigma_0^2}\mu_0, \left(\frac{1}{\sigma_0^2} + \frac{1}{\sigma^2}\right)^{-1}\right)$.

Lemma 1 (Privacy Cost and Marginal Privacy Cost).

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

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

We now derive the consumer's problem. We can rewrite the evolution of posterior variances (1) as that of posterior precisions:

$$\frac{1}{\sigma_t^2(\mathbf{a}_t, \gamma_t)} = \frac{1}{\sigma_{t-1}^2(\mathbf{a}_{t-1}, \gamma_{t-1})} + \frac{1}{\frac{1}{a_t} + \gamma_t} = \dots = \frac{1}{\sigma_0^2} + \sum_{s=1}^t \frac{1}{\frac{1}{a_s} + \gamma_s}. \quad (2)$$

Once the platform commits to privacy policy $(\gamma_t)_{t \in \mathbb{N}}$, the consumer chooses activity levels over time, which gives the following problem:

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

The next result presents the consumer's response to a stationary privacy policy. Although the platform's optimal policy may not be stationary, this non-equilibrium analysis clarifies the intuition behind the consumer's dynamic incentive.¹⁶

¹⁶We may also view the consumer's behavior under a stationary policy as the long-run outcome of the model in which the platform sets a single privacy level to consumers who join the platform in different periods. To see this, suppose that a unit mass of consumers join the platform in each period, and each consumer stops using the platform in any period with probability $1 - \delta$. As $t \rightarrow \infty$ the mass of consumers who joined the platform k periods ago converges to δ^k , and every consumer solves (3) with $\delta_C = \delta$ and $\gamma_t = \gamma$ for all t .

Proposition 1. *Let $(a_t^*)_{t \in \mathbb{N}}$ denote the optimal choice of the consumer under a stationary privacy policy, i.e., $\gamma_t = \gamma$ for all $t \in \mathbb{N}$. There is a cutoff value $v^*(\gamma) \in \mathbb{R}_+$ with the following properties.*

1. *If $v < v^*(\gamma)$, then a_t^* increases in t , $\lim_{t \rightarrow \infty} a_t^* = a_{max}$, and $\lim_{t \rightarrow \infty} \sigma_t^2 = 0$. The consumer's continuation value decreases over time.*
2. *If $v > v^*(\gamma)$, then $a_t^* = 0$ and $\sigma_t^2 = \sigma_0^2$ for all $t \in \mathbb{N}$.*
3. *The cutoff $v^*(\gamma)$ is increasing in γ , and $\lim_{\gamma \rightarrow \infty} v^*(\gamma) = \infty$.*

The intuition is as follows. If the value of privacy is low, the consumer prefers a positive activity level $a_1^* > 0$ in $t = 1$. The consumer activity generates information, which reduces her payoff and marginal cost of using the platform. As a result, she chooses $a_2^* \geq a_1^*$ in $t = 2$. Repeating this argument, we can conclude that a_t^* increases over time. The platform can then observe the signals to perfectly learn the consumer's type as $t \rightarrow \infty$. Perfect learning in $t \rightarrow \infty$ implies that the marginal privacy cost goes to zero, and thus $a_t^* \rightarrow a_{max}$ (Point 1). In contrast, the consumer with a high v does not use the platform (Point 2). Finally, $v^*(\gamma)$ is increasing in γ because a higher privacy level reduces the cost of using the platform (Point 3).

Under a stationary policy, the consumer's continuation value decreases over time. In contrast, the flow payoffs may be non-monotone: Evaluated at the same activity level (i.e., $a_{t-1} = a_t$), the consumer incurs a higher privacy cost and earns a lower flow payoff in period t than $t - 1$, because $\sigma_t^2 \leq \sigma_{t-1}^2$. However, the optimal a_t is (weakly) higher than a_{t-1} , and the consumer may increase a_t so much that she enjoys a higher flow payoff than in the previous period at the expense of a low continuation value. [Online Appendix J](#) provides a numerical example.

If we interpret γ as a privacy regulation, [Proposition 1](#) implies that a stricter regulation (i.e., a higher γ) increases cutoff $v^*(\gamma)$ and expands the range of v 's under which the consumer loses privacy. Thus a stricter regulation can increase the platform's information by making consumers comfortable to use the platform. Because a higher γ decreases privacy costs for any given sequence of activity levels, it also increases the consumer's ex ante payoffs.

We now turn to the optimal policy of the platform that can commit to any (potentially nonstationary) privacy policy. The platform anticipates that the consumer solves (3) given any privacy policy. The following lemma presents a property of the consumer's optimal behavior that shapes the platform's choice.

Lemma 2. Let $(\bar{a}_t(\gamma))_{t \in \mathbb{N}}$ and $(\bar{a}_t(\gamma'))_{t \in \mathbb{N}}$ denote the optimal activity levels (i.e., solutions of (3)) under privacy policies $\gamma = (\gamma_t)_{t \in \mathbb{N}}$ and $\gamma' = (\gamma'_t)_{t \in \mathbb{N}}$, respectively. Let $\mathcal{T} = \{t \in \mathbb{N} : \gamma_t \neq \gamma'_t\}$ denote the set of all periods in which the two policies differ. Suppose $\frac{1}{\bar{a}_t(\gamma)} + \gamma_t \leq \frac{1}{\bar{a}_t(\gamma')} + \gamma'_t$ for all $t \in \mathcal{T}$. Then, $\bar{a}_t(\gamma) \geq \bar{a}_t(\gamma')$ for all $t \in \mathbb{N} \setminus \mathcal{T}$.

The lemma means that data collection in different periods are complements. For example, suppose that the platform changes a privacy level for period t from γ' to γ . The platform makes such a change at the ex ante stage, so the consumer will also adjust her activity levels over time. In particular, suppose her activity level in period t changes from a' to a . The lemma implies that if the new policy leads to better information in period t , i.e., $\frac{1}{a} + \gamma \leq \frac{1}{a'} + \gamma'$, then the consumer must also be choosing higher activity levels and the platform collects more information in any other period $s \neq t$, in which the privacy levels remain the same. The intuition comes from the decreasing marginal privacy cost: If the consumer anticipates more data collection in some periods, it reduces the cost of giving up her privacy in other periods.

Proposition 1 and Lemma 2 do not characterize how the consumer reacts to the general change of a privacy policy. However, we can use them to characterize the unique long-run outcome (see Appendix E for the proof; Appendix C proves the existence of an optimal policy).

Theorem 1. Any optimal policy of the platform induces the following outcome:

1. The consumer loses her privacy and chooses the highest activity level in the long run:

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

2. For any $T \in \mathbb{N}$, there is a $\underline{v} \in \mathbb{R}$ such that for any $v \geq \underline{v}$, we have $\gamma_t^* > 0$ for all $t \leq T$.

Point 1 says that a farsighted consumer eventually gives up all of her privacy, even when she anticipates future choices made by the platform. The privacy loss occurs for any discount factors, so a myopic platform that faces a patient consumer induces the long-run privacy loss. Similarly, the long-run outcome is independent of the value v of privacy, the shape of $u(\cdot)$, or which optimal policy we consider, in case the optimal policy is not unique.¹⁷ Point 2 implies that if the consumer highly values her privacy, the platform commits to garble signals to maximize information.

¹⁷The platform's ex ante payoff is uniquely determined, because we pin down the consumer's optimal choice by assuming the tie-breaking in favor of greater a_t 's, and the platform commits to a privacy policy ex ante.

We may think that the long-run privacy loss follows from the statistical argument—that the platform observes a signal in each period, so it will learn the true type as $t \rightarrow \infty$. The intuition is misleading for two reasons. First, the consumer can choose $a_t = 0$ to generate no information, and we do not impose any Inada-type condition on $u(\cdot)$ under which the consumer always prefers a positive a_t . Second, the platform’s objective is not the long-run payoff but the discounted payoffs. Depending on how data collection in one period affects the outcomes in other periods, the platform might strategically commit to not collect full information.¹⁸ The result stems from a combination of the decreasing marginal privacy cost and the platform’s ability to offer privacy protection.

The intuition for Point 1 is as follows. Suppose that the platform commits to collect no information after period $\tau \geq 2$, that is, it adopts $(\gamma_t)_{t \in \mathbb{N}}$ such that $\gamma_t = \infty$ for any $t \geq \tau$. We show that the platform can modify the policy to increase the precision of signals in all periods. First, [Proposition 1](#) says that there is a stationary policy $\gamma_t = \gamma^* < \infty$ under which the consumer chooses increasing activity levels. Given such a γ^* , suppose the platform replaces $\gamma_t = \infty$ with $\gamma_t = \gamma^*$ for all $t \geq \tau$. After the change, the consumer’s problem from period τ is as if she faces a lower initial variance (i.e., σ_t^2 instead of σ_0^2) and faces a stationary policy of γ^* . Because γ^* induces the consumer to choose positive activity levels when she starts with the initial variance σ_0^2 , the consumer chooses even higher activity levels when she faces γ^* from period τ onward. As a result the change of the policy increases the precision of signals after period τ . Second, the consumer also adjusts activity levels before period τ , but [Lemma 2](#) implies that the consumer will increase activity levels in any period $s < \tau$. Thus, from the platform’s perspective, any policy that stops data collection is dominated by another policy that induces privacy loss. The proof applies a similar argument to any policy such that σ_t^2 does not converge to 0.

The intuition for Point 2 is that in early periods, the platform knows little about the consumer, so the consumer’s activity has a large impact on what the platform can learn about her type. Thus the consumer faces a high marginal privacy cost, which discourages her from raising the activity level. The platform then commits to a high level of privacy protection to encourage consumer activity.

In contrast to [Proposition 1](#), under the optimal policy, a_t^* and γ_t^* may be non-monotone (see

¹⁸To see this, consider a situation in which the consumer’s activity level in the first period is decreasing in the precision of the signal generated in the future period. A myopic platform would then commit to $\gamma_t = \infty$ for all $t \geq 2$ and σ_t^2 does not converge to 0.

Online Appendix J for numerical examples). The activity level may decrease from period $t - 1$ to t , when the platform decreases γ_t^* to increase the signal quality at the expense of the activity level. The platform may increase γ_t^* if a lower σ_t^2 makes it easier for the platform to induce a high a_t through privacy protection.

Remark 1. I assume that the platform can only commit to a deterministic sequence of privacy levels. A natural question is whether the platform benefits from stronger commitment power—i.e., it can commit to an action-dependent policy that maps past actions (a_1, \dots, a_{t-1}) to the privacy level γ_t in any period $t \geq 2$. Online Appendix M provides a sufficient condition under which the platform earns a greater payoff from an action-dependent policy than under long-run commitment. Intuitively, the platform with the stronger commitment power may induce consumers to choose high activity levels by threatening them to set $\gamma_t = 0$ in future periods unless they choose a high activity level today. At the same time, Section 5 shows that if the consumer has binary activity levels and the players have the same discount factor, the platform does not benefit from action-dependent policies.

4.1 The Long-Run Privacy Loss Under a General Payoff Specification

We now ask to what extent Theorem 1 depends on assumptions on preferences. In particular, I have assumed that the consumer dislikes data collection and the platform cares only about information. I now relax these restrictions on preferences. The platform's per-period payoff is $\Pi(a_t, \sigma_t^2)$, which is strictly increasing in a_t and decreasing in σ_t^2 . For example, an advertising platform benefits from a high activity a_t and more data (i.e., low σ_t^2) because the consumer will then see many highly targeted ads. We impose no restrictions on the relative importance of activity and data for Π .

The consumer's per-period payoff is now $\hat{u}(a_t, \sigma_t^2)$ that satisfies the following conditions: There is some $\beta > 0$ such that for each $\sigma^2 \in [0, \sigma_0^2]$ and $a, a' \in A$ with $a > a'$, $\hat{u}(a, \sigma^2) - \hat{u}(a', \sigma^2) \geq \beta$; for each $a \in A$, $\hat{u}(a, \sigma^2)$ is differentiable in σ^2 ; and $\max_{a \in A, \sigma^2 \in [0, \sigma_0^2]} \left| \frac{\partial \hat{u}}{\partial \sigma^2}(a, \sigma^2) \right| < \infty$. The last inequality means that the marginal cost of losing privacy is uniformly bounded. For example, we can allow $\hat{u}(a, \sigma^2) = u(a) - C(\sigma^2)$ for any $C(\cdot)$ with bounded derivatives. The term $C(\cdot)$ can be first decreasing and then increasing, i.e., the consumer prefers some level of data collection. The original setting is $\hat{u}(a, \sigma^2) = u(a) - v(\sigma_0^2 - \sigma^2)$. Other parts of the game remain the same.

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

1. *There exists a privacy policy $(\gamma_t)_{t \in \mathbb{N}}$ under which the consumer eventually loses her privacy and chooses the highest activity level: Any optimal choice of the consumer induces $\lim_{t \rightarrow \infty} \sigma_t^2 = 0$ and $\lim_{t \rightarrow \infty} a_t^* = a_{max}$.*
2. *Fix the consumer's discount factor. Let $\sigma_\infty^2(\delta_P) := \lim_{t \rightarrow \infty} \sigma_t^2(\delta_P)$ denote the long-run posterior variance under an arbitrarily chosen optimal policy given the platform's discount factor δ_P . Then, $\lim_{\delta_P \rightarrow 1} \sigma_\infty^2(\delta_P) = 0$.*

Point 1 states that the platform can induce privacy loss and a high activity level for a broad class of consumer preferences. Whenever the consumer incurs bounded marginal costs $\frac{\partial \hat{u}}{\partial \sigma^2}$, the marginal cost with respect to an activity level goes to zero as $\sigma_t^2 \rightarrow 0$. Thus the platform can induce the highest activity level and full information collection in the long run.¹⁹

Point 1 alone does not imply that the platform collects all the information under an optimal policy. However, Point 2 states that a sufficiently patient platform does so. The consumer chooses the highest activity level when she has no privacy, and such an outcome maximizes the long-run profit of the platform. Thus for δ_P close to 1, even if a platform cares about user activity, it may eventually collect as much information as exclusively data-driven firms in the long run.

The result concerns only the long-run outcome induced by a patient platform. As a result, for a general δ_P and in a given period t , the platform may still induce a higher σ_t^2 and impose a lower privacy cost when it cares more about consumer activity. Also, the result fixes δ_C and takes $\delta_P \rightarrow 1$. Because a more patient consumer is less willing to raise an activity level (see [Section 6.1](#)), the long-run outcome in [Proposition 2](#) may differ from the one in which the consumer and the platform have the common discount factor that converges to 1.

4.2 Implications of [Theorem 1](#) and [Proposition 2](#)

First, [Theorem 1](#) potentially explains the privacy paradox: Consumers seem to casually share their data with online platforms, despite their concerns about data collection.²⁰ We may view this puz-

¹⁹Although in a different context, the intuition is similar to the idea that if a firm uses data for forecasts and the gain to a perfect forecast is finite, the returns to data must diminish at some point ([Farboodi and Veldkamp, 2020](#)).

²⁰[Acquisti et al. \(2016\)](#) conduct an insightful review of research on the economics of privacy, including the privacy paradox. Recent empirical work includes, for example, [Athey et al. \(2017\)](#).

zle as the long-run outcome of this model, in which the consumer faces a high privacy cost and negligible marginal cost. Such an outcome can arise even if firms adopt business models that do not much rely on data (Proposition 2), or the platform mainly cares about short-run profits. The result also points to the difficulty of applying the revealed preference argument to static privacy choices, because the consumer’s decision may depend on the stock of information they have already revealed.

Second, the result connects consumer privacy problem with rational addiction (Becker and Murphy, 1988). The connection stems from that a high activity level today decreases the consumer’s future utility, but increases her future marginal utility of using the platform. In contrast to models of rational addiction, the current model has a platform that can choose its privacy policy to influence the degree of “addiction.”²¹ As a result, even if consumers are patient and highly value their privacy, they keep using the platform and lose privacy.

Finally, at an anecdotal level, the optimal policy of the platform, which offers high privacy levels for early periods but not necessarily for later periods, seems consistent with how data policies of online platforms have evolved. In 2004, Facebook’s privacy policy stated that it would not use (first-party) cookies to collect user information. In 2020, the privacy policy states that it uses cookies to track users on and possibly off the website.²² Srinivasan (2019) describes how Facebook’s policy has changed from the one that preserves consumer privacy to “broad-scale commercial surveillance.” Also, Fainmesser et al. (2019) describe how online platforms’ business models have changed from the initial phase, in which they expand a user base, to the mature phase, in which they monetize the information collected. The dynamics presented in this paper is one way to rationalize the pattern described.²³

²¹Relatedly, Boone and Shapiro (2006), Zhang (2012), and Meng (2018) study dynamic contracts for addictive goods. A general framework in Pavan et al. (2014) could also apply to the design of a mechanism that provides addictive goods. In these papers, the principal uses monetary transfer and allocation rules to maximize profits, which differs from my model in which the platform can directly affect the degree of “addition” through a privacy policy. Methodologically, these papers use the first-order conditions to characterize the agent’s incentive; in my model, we cannot use the first-order conditions because the consumer faces a concave cost function.

²²In 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>

²³In addition to the force described in this paper, the platform may have an even stronger incentive to lower privacy protection if the expansion of user base strengthens a positive network externality and incentivizes users to increase activity levels even at lower privacy protection.

5 One-Period Commitment

I now study the game of one-period commitment, in which the two players move sequentially in every period. One-period commitment could be realistic, for example, if a platform may be sanctioned for the outright violation of its privacy policy, but it may still revise the policy over time. For some of the results, I study the following equilibria:

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

Proposition 3. Any pure-strategy MPE satisfies exactly one of the following two properties:

1. $\lim_{t \rightarrow \infty} \gamma_t = 0$, $\lim_{t \rightarrow \infty} a_t = a_{max}$, and $\lim_{t \rightarrow \infty} \sigma_t^2 = 0$; or
2. $\lim_{t \rightarrow \infty} \gamma_t = \infty$, $a_t > 0$ for all t , and $\lim_{t \rightarrow \infty} \sigma_t^2 > 0$.

Moreover, there is a $B > 0$ such that if $\sigma_0^2 \leq B$, the unique equilibrium involves $(\gamma_t, a_t) = (0, a_{max})$ for all $t \in \mathbb{N}$.

The first part of the result classifies any MPE into two types, depending on whether the long-run privacy loss (i.e., $\lim_{t \rightarrow \infty} \sigma_t^2 = 0$) occurs. In any equilibrium with $\lim_{t \rightarrow \infty} \sigma_t^2 = 0$, the platform eventually offers $\gamma_t = 0$ because if σ_t^2 is sufficiently low, the consumer will choose a high activity level even without privacy protection. In contrast, in any equilibrium with $\lim_{t \rightarrow \infty} \sigma_t^2 > 0$, the platform offers perfect privacy protection (i.e., $\lim_{t \rightarrow \infty} \gamma_t = \infty$) in the long run. In such an equilibrium, the weaker commitment power forces the platform to stop data collection. The second part of the result implies that an equilibrium that satisfies Point 2 exists only if the platform has not yet collected too much data on the consumer. Intuitively, if σ_0^2 is small, the marginal privacy cost is so small that the consumer prefers a_{max} regardless of the current and future privacy protection.

Proposition 3 does not say whether both types of equilibria exist. I now construct such equilibria to provide an intuition regarding how weaker commitment power affects data collection. To facilitate the analysis, we will use the following assumption:

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

In general, analyzing the consumer’s optimal choice can be complicated, because the consumer’s payoffs contain the privacy cost, which is concave in a . For example, even if $u(a)$ is concave, the consumer’s flow payoff may be neither concave nor convex in a . [Assumption 1](#) simplifies the consumer’s incentive and enables us to derive stronger results regarding the equilibrium dynamics. Moreover, the assumption is without loss of generality if $u(a)$ is weakly convex, in which case the consumer always finds it optimal to choose 0 or a_{max} from a general A (see the discussion at the end of [Appendix G](#)).

The following notions are useful for describing the results.

Definition 2. An equilibrium is *consumer-worst* if it minimizes the consumer’s ex ante sum of discounted payoffs across all equilibria (including those that are not MPE). We analogously define “*consumer-best*,” “*platform-worst*,” and “*platform-best*.”

The following result presents a consumer-worst equilibrium, which is also platform-best under a common discount factor. Note that in any MPE, the consumer’s strategy is written as $a(\sigma^2, \gamma)$, which specifies her activity level given residual variance σ^2 in the previous period and privacy protection γ in the current period. Similarly, the platform’s strategy is written as $\gamma(\sigma^2)$.

Theorem 2. *Under [Assumption 1](#), there is a consumer-worst MPE. This equilibrium is independent of the platform’s discount factor δ_P , and has the following properties:*

1. *The privacy level γ_t^* is decreasing in t and hits zero in a finite time. The consumer loses her privacy in the long run: $\lim_{t \rightarrow \infty} \sigma_t^2 = 0$.*
2. *In every period, the platform chooses the lowest privacy level γ_t that induces the consumer to take $a_t = a_{max}$. Formally, given her equilibrium Markov strategy $a(\sigma^2, \gamma)$, the platform’s equilibrium strategy is written as $\gamma(\sigma^2) = \min \{ \gamma \in \mathbb{R}_+ : a(\sigma^2, \gamma) = a_{max} \}$ for any $\sigma^2 \leq \sigma_0^2$.*

Point 1 extends the intuition for [Theorem 1](#). The consumer incurs high marginal privacy costs in early periods, so the platform initially chooses high γ_t to incentivize the consumer to generate information. As the platform collects more information, the consumer’s incentive to protect privacy declines; correspondingly, the platform sets a decreasing privacy level. [Lemma 1](#) alone

does not imply that the consumer faces a lower cost of choosing a_{max} when σ_t^2 is small, because her continuation value is endogenous. However, in this equilibrium, the consumer's continuation value $V(1/\sigma_t^2)$, as a function of the amount of information collected, is decreasing and convex in $1/\sigma_t^2$. As a result, the consumer's Markov decision problem exhibits a declining marginal loss of generating information.

Point 2 implies that the platform's optimal policy is greedy—i.e., in each period, the platform maximizes the precision of the signal for that period. The greedy policy is optimal for the platform because the decreasing marginal cost implies that the consumer's future activity level is increasing in the precision of the signal generated today. The optimality of greedy policy indicates that the platform may not benefit from stronger commitment power. The following result confirms this intuition. Recall that δ_C and δ_P denote the discount factors of the consumer and the platform, respectively.

Corollary 1. *Let γ_t^* denote the privacy level in the consumer-worst equilibrium. If $\delta_C = \delta_P$, the optimal policy under long-run commitment is $(\gamma_t^*)_{t \in \mathbb{N}}$. Moreover, the platform's payoff in the consumer-worst equilibrium is equal to the one when the platform can commit to any action-contingent privacy policy—i.e., it can pre-commit to any strategy that determines privacy level γ_t in every period $t \geq 2$ as a function of past activity levels (a_1, \dots, a_{t-1}) .*

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

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

In this equilibrium, the platform offers full privacy, because whenever it attempts to collect information by setting $\gamma_t < \infty$, the consumer chooses $a_t = 0$. The consumer prefers $a_t = 0$ following the platform's deviation, because the initial privacy loss, no matter how small, will lead to complete privacy loss and impose her a high cost in the future. Indeed, after any off-path event in which the platform collects some information (i.e., $\sigma_t^2 < \sigma_0^2$), the consumer-worst

equilibrium in [Theorem 2](#) is played. A grim trigger strategy—i.e., the platform’s deviation induces $a_t = 0$ forever— does not work, because the platform can set a large finite γ_t to render such a punishment suboptimal for the consumer. We may view [Theorem 3](#) as the platform’s Coasian commitment problem: The platform in period t competes with its future self, which offers the best privacy protection in any period $s \geq t + 1$. The result counters the idea that the platform’s weak commitment power causes insufficient privacy protection.

Remark 2 (Welfare Implications). Under a common discount factor and binary activity levels, the equilibria in [Theorems 2](#) and [3](#) are Pareto optimal in terms of the ex ante payoffs of the consumer and the platform. Indeed, once we write the platforms’ payoff as $v[\sigma_0^2 - \sigma_t^2]$, the sum of the payoffs of the consumer and the platform is $\sum_{t=1}^{\infty} \delta^{t-1} u(a_t)$, which is maximized at $a_t = a_{max}$ for all t . As a result, if the consumer chooses $a_t = a_{max}$ in all periods at two equilibria, both are Pareto optimal. Therefore we can view [Theorems 2](#) and [3](#) as two points on the Pareto frontier and the platform’s commitment as a way to select the consumer-worst outcome.

6 Extensions

6.1 Speed of Learning

We provide comparative statics about the speed of learning. The following results show that the platform collects information more slowly when consumers have a high value on privacy, a high discount factor, or a high prior variance for their types. A high v , δ_C , or σ_0^2 slows down data collection by increasing the cost for the consumer of using the platform and generating information.

Proposition 4. *Fix any privacy policy under long-run commitment. In any period, the optimal activity level and the precision of the signal are lower if the discount factor δ_C , the value v of privacy, or the prior variance σ_0^2 is higher.*

Proposition 5. *Suppose [Assumption 1](#) holds. Consider the consumer-worst equilibrium in [Theorem 2](#). In any period, the equilibrium privacy level is higher and the precision of the signal is lower if the consumer’s discount factor δ_C , the value v of privacy, or the prior variance σ_0^2 is higher.*

Because of [Corollary 1](#), [Proposition 5](#) also applies to long-run commitment: If the consumer has binary actions and the players have common discount factor $\delta_C = \delta_P = \delta$, the platform's optimal privacy levels increase in (v, δ, σ_0^2) under long-run commitment.

6.2 Time-Varying Type

This extension allows the consumer's type to change over time. The consumer's type in period t is X_t , and $(X_t)_{t \in \mathbb{N}}$ now follows an $AR(1)$ process: $X_0 \sim \mathcal{N}(0, \sigma_0^2)$, $X_{t+1} = \phi X_t + z_t$ with $\phi \in [0, 1]$, and $z_t \stackrel{iid}{\sim} \mathcal{N}(0, (1-\phi^2)\sigma_0^2)$. Note that $Var(X_t) = \sigma_0^2$ for all $t \in \mathbb{N}$. As in the baseline model (which corresponds to $\phi = 1$), the platform observes a signal $s_t = X_t + \varepsilon_t$ with $\varepsilon_t \sim \mathcal{N}\left(0, \frac{1}{a_t} + \gamma_t\right)$.

Suppose that at the beginning of period t , the platform holds posterior variance σ_{t-1}^2 . If the platform sets γ_t and the consumer chooses a_t , the platform observes a signal and holds posterior variance $\hat{\sigma}_t^2 = \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{a_t} + \gamma_t}$. Then the consumer receives $u(a_t) - v(\sigma_0^2 - \hat{\sigma}_t^2)$ and the platform receives $\sigma_0^2 - \hat{\sigma}_t^2$. After the state transition, the platform's posterior variance on X_{t+1} at the beginning of period $t + 1$ is

$$\sigma_t^2 = \phi^2 \hat{\sigma}_t^2 + (1 - \phi^2) \sigma_0^2 \geq \hat{\sigma}_t^2.$$

For example, if the state is independent across time (i.e., $\phi = 0$), the posterior variance at the beginning of each period is always σ_0^2 , but the platform may still collect information within the period (i.e., $\hat{\sigma}_t^2 < \sigma_0^2 = \sigma_t^2$). The following result generalizes [Proposition 1](#).

Proposition 6. *Let $(a_t^*)_{t \in \mathbb{N}}$ denote the optimal choice of the consumer under a stationary privacy policy $\gamma_t = \gamma$. There is a cutoff value $v^*(\gamma) \in \mathbb{R}_+$ with the following properties.*

1. *If $v < v^*(\gamma)$, then a_t^* increases in t , and σ_t^2 decreases in t and converges to a non-negative value as $t \rightarrow \infty$. The consumer's continuation value decreases over time.*
2. *If $v > v^*(\gamma)$, then $a_t^* = 0$ and $\sigma_t^2 = \sigma_0^2$ for all $t \in \mathbb{N}$.*
3. *The cutoff $v^*(\gamma)$ is increasing in γ , and $\lim_{\gamma \rightarrow \infty} v^*(\gamma) = \infty$.*

The next result generalizes [Theorem 2](#).

Proposition 7. *Assume $\delta_C = \delta_P$. Under [Assumption 1](#), there is a Markov perfect equilibrium with the following properties:*

1. The equilibrium is consumer-worst and platform-best, and the equilibrium outcome coincides with the one under long-run commitment.
2. The privacy level γ_t^* and posterior variance σ_t^2 are decreasing in t .
3. In each period, both on and off the equilibrium paths, the platform chooses the lowest privacy level γ_t that induces the consumer to take $a_t = a_{max}$ as opposed to $a_t = 0$.
4. The ex ante payoff of the consumer decreases and that of the platform increases in persistence ϕ .

When the consumer's type varies over time, the equilibrium no longer entails long-run privacy loss. However, these results show that much of the equilibrium dynamics driven by the decreasing marginal privacy cost continues to hold. Moreover, [Proposition 7](#) shows that greater persistence of the type could hurt the consumer and benefit the platform. From the platform's perspective, higher persistence could hinder data collection because the consumer is less willing to use the service. However, a higher ϕ also means that the platform can use the information generated today to predict future types. The result shows that under a certain setting, the beneficial impact dominates and a higher persistence facilitates data collection.

6.3 General Noise Structure

We have assumed that the noise term of a signal is distributed according to $\mathcal{N}\left(0, \frac{1}{a_t} + \gamma_t\right)$. However, all the results hold if the variance takes a more general form $\sigma_\varepsilon^2(a, \gamma)$ that satisfies several conditions. First, we assume that $\frac{\partial \sigma_\varepsilon^2(a, \gamma)}{\partial \gamma} > 0$, $\frac{\partial \sigma_\varepsilon^2(a, \gamma)}{\partial a} < 0$, $\frac{\partial^2 \sigma_\varepsilon^2}{\partial a \partial \gamma} \geq 0$, $\lim_{\gamma \rightarrow \infty} \sigma_\varepsilon^2(a, \gamma) = \infty$, and $\sigma_\varepsilon^2(0, \gamma) = \infty$. The posterior variance now evolves as follows:

$$\sigma_t^2 = \frac{1}{\frac{1}{\sigma_{t-1}^2} + \frac{1}{\sigma_\varepsilon^2(a_t, \gamma_t)}}.$$

Second, we assume that

$$\frac{\partial \sigma_t^2}{\partial a} = \frac{\partial \sigma_\varepsilon^2}{\partial a}(a_t, \gamma_t) \cdot \frac{1}{\left(\frac{\sigma_\varepsilon^2(a_t, \gamma_t)}{\sigma_{t-1}^2} + 1\right)^2}, \quad (4)$$

which is negative, converges to 0 as $\gamma \rightarrow \infty$.

Equation (4) implies that lower privacy (i.e., a lower σ_{t-1}^2) leads to a lower marginal cost of raising a_t , and the marginal cost approaches 0 as $\sigma_{t-1}^2 \rightarrow 0$. The condition $\frac{\partial^2 \sigma_\varepsilon^2}{\partial a \partial \gamma} \geq 0$ implies that the right-hand side of (4) increases in γ_t , i.e., the marginal privacy cost decreases in γ_t . Also, $\frac{\partial \sigma_t^2}{\partial a}$ approaches 0 as $\gamma_t \rightarrow \infty$, so the platform can induce positive activity levels with sufficiently high privacy levels. These conditions are sufficient to extend all the results of this paper. An example that satisfies all the conditions is $\sigma_\varepsilon^2(a, \gamma) = \frac{e^{\gamma-a}}{a}$. An example that satisfies all the conditions except $\frac{\partial^2 \sigma_\varepsilon^2}{\partial a \partial \gamma} \geq 0$ is $\sigma_\varepsilon^2(a, \gamma) = \frac{\gamma}{a}$. In this case, $\frac{\partial \sigma_t^2}{\partial a}$ could decrease in γ , so higher privacy protection decreases the privacy cost but may increase the marginal cost.

7 Conclusion

This paper studies a dynamic model of consumer privacy and platform data collection. The fundamental feature of the model is that consumers use a platform’s service more actively when they expect high privacy protection or have already lost their privacy. The platform can collect information over time by committing to not collect too much information in early periods. In equilibrium, the consumer eventually loses privacy but keeps choosing a high level of activity. This outcome can arise even if the platform prioritizes user activity over data collection. If the platform has weaker commitment, it may end up offering the highest privacy protection: The consumer refuses to provide information, anticipating that small privacy loss will lead to complete privacy loss.

References

- Acemoglu, Daron, Ali Makhdoui, 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.
- Argenziano, Rossella and Alessandro Bonatti (2020), “Information revelation and privacy protection.”

- Athey, Susan, Christian Catalini, and Catherine Tucker (2017), “The digital privacy paradox: Small money, small costs, small talk.” Technical report, National Bureau of Economic Research.
- Ball, Ian (2020a), “Dynamic information provision: Rewarding the past and guiding the future.” *Available at SSRN 3103127*.
- Ball, Ian (2020b), “Scoring strategic agents.” *arXiv preprint arXiv:1909.01888*.
- Becker, Gary S and Kevin M Murphy (1988), “A theory of rational addiction.” *Journal of political Economy*, 96, 675–700.
- Bergemann, Dirk, Alessandro Bonatti, and Tan Gan (2019), “The economics of social data.” Cowles foundation discussion paper 2203, Yale University.
- Bonatti, Alessandro and Gonzalo Cisternas (2020), “Consumer scores and price discrimination.” *The Review of Economic Studies*, 87, 750–791.
- Boone, Jan and Joel Shapiro (2006), “Selling to consumers with endogenous types.”
- 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.
- Easley, David, Shiyang Huang, Liyan Yang, and Zhuo Zhong (2018), “The economics of data.” *Available at SSRN 3252870*.
- Ely, Jeffrey C (2017), “Beeps.” *American Economic Review*, 107, 31–53.
- Fainmesser, Itay P, Andrea Galeotti, and Ruslan Momot (2019), “Digital privacy.” *Available at SSRN*.
- Farboodi, Maryam and Laura Veldkamp (2020), “A growth model of the data economy.” Technical report, MIT Sloan.
- Frankel, Alex and Navin Kartik (2019a), “Improving information from manipulable data.” *arXiv preprint arXiv:1908.10330*.
- Frankel, Alex and Navin Kartik (2019b), “Muddled information.” *Journal of Political Economy*, 127, 1739–1776.

- 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.”
- Ichihashi, Shota (2020), “The economics of data externalities.”
- Jullien, Bruno, Yassine Lefouili, and Michael Riordan (2018), “Privacy protection and consumer retention.”
- Kamenica, Emir (2019), “Bayesian persuasion and information design.” *Annual Review of Economics*, 11, 249–272.
- Kummer, Michael and Patrick Schulte (2019), “When private information settles the bill: Money and privacy in googles market for smartphone applications.” *Management Science*.
- Lin, Tesary (2019), “Valuing intrinsic and instrumental preferences for privacy.” *Available at SSRN*, 3406412.
- Meng, Dawen (2018), “Sequential nonlinear pricing of experience goods with network effects.” *Available at SSRN 3168438*.
- Milgrom, Paul, Chris Shannon, et al. (1994), “Monotone comparative statics.” *Econometrica*, 62, 157–157.
- Pavan, Alessandro, Ilya Segal, and Juuso Toikka (2014), “Dynamic mechanism design: A myersonian approach.” *Econometrica*, 82, 601–653.
- Renault, Jérôme, Eilon Solan, and Nicolas Vieille (2017), “Optimal dynamic information provision.” *Games and Economic Behavior*, 104, 329–349.
- Smolin, Alex (forthcoming), “Dynamic evaluation design.” *American Economic Journal: Microeconomics*.
- 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*.

Zhang, Wei (2012), “Endogenous preferences and dynamic contract design.” *The BE Journal of Theoretical Economics*, 12.

Appendix

A A Microfoundation of the Privacy Cost and the Platform Revenue

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

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

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

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

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

$$\frac{3}{8}\mathbb{E}[(\mu - \mu_0)^2] = \frac{3}{8} \cdot \frac{(\sigma_0^2)^2}{\sigma_0^2 + \sigma_\varepsilon^2}. \quad (5)$$

The posterior variance σ_t^2 of X and the variance σ_ε^2 of the noise ε satisfy the equation $\frac{1}{\sigma_t^2} = \frac{1}{\sigma_0^2} + \frac{1}{\sigma_\varepsilon^2}$. Solving this equation with respect to σ_ε^2 and plugging it into (5), we obtain

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

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

$$\frac{1}{4} (\sigma_0^2 - \sigma_t^2), \quad (6)$$

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

We obtain the original model by assuming that the platform sells information to sellers who use it to price discriminate the consumer. The detail is as follows. Outside of the platform, the consumer interacts with seller t in period t , and her willingness to pay for seller t 's product is $X_t \sim \mathcal{N}(\mu_0, \sigma_0^2)$, which is now *IID across t* .²⁴ The consumer's activity on the platform in each period yields her utilities and generates signals for future sellers. Specifically, each period t consists of the following events: (i) the consumer chooses a_t , (ii) the platform collects and sells information (by posting a price) to seller t , and (iii) the seller sets the price and the consumer chooses quantity. Precisely, the platform can sell seller t the signal $s_t = X_t + \varepsilon_t$ with

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

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

²⁴I impose the IID assumption so that the prior distribution of the consumer's willingness to pay in the product market is the same across t .

according to (1). The consumer's payoff is $u(a_t) - \frac{3}{8}(\sigma_0^2 - \sigma_t^2)$.

B Proof of Lemma 1

Proof. Point 2 follows from

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

□

C The Existence of an Optimal Policy Under Long-Run Commitment

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

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

$$\max_{\mathbf{a} \in \mathcal{A}} \sum_{t=1}^{\infty} \delta_C^{t-1} \left[u(a_t) - v \cdot \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^t \frac{1}{\frac{1}{a_s} + \gamma_s}} \right) \right]. \quad (7)$$

For any $\gamma \in \Gamma$, let $\mathcal{A}^*(\gamma) \subset \mathcal{A}$ denote the set of all maximizers of (7).

Lemma 3. *The correspondence $\mathcal{A}^*(\gamma)$ is non-empty, compact, and upper hemicontinuous in γ .*

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 $(\mathbf{a}^n)_{n=1}^\infty$ such that $\mathbf{a}^n \rightarrow \mathbf{a}^*$. This implies that for each $t \in \mathbb{N}$, $\lim_{n \rightarrow \infty} a_t^n \rightarrow a_t^*$. The consumer's period- t payoff $U_t(\mathbf{a}, \gamma) :=$

$u(a_t) - v \cdot \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^t \frac{1}{\frac{1}{a_s} + \gamma_s}} \right)$ is bounded from above and below by $u(a_{max})$ and $-v\sigma_0^2$, respectively. Define $B := \max(u(a_{max}), v\sigma_0^2) > 0$. Take any $\varepsilon > 0$, and let T^* satisfy $\frac{\delta_C^{T^*}}{1-\delta_C} B < \frac{\varepsilon}{4}$. Take a sufficiently large n so that for each $t \leq T^*$, $\delta_C^{t-1} |U_t(\mathbf{a}^n, \gamma) - U_t(\mathbf{a}^*, \gamma)| < \frac{\varepsilon}{2T^*}$. These inequalities imply that

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

Thus the objective function in (7) is continuous in \mathbf{a} . Berge maximum theorem implies that $\mathcal{A}^*(\gamma)$ is non-empty, compact, and upper hemicontinuous in γ . \square

Next, I show properties of the consumer's objective $U(\mathbf{a}, \gamma) := \sum_{t=1}^{\infty} \delta_C^{t-1} U_t(\mathbf{a}, \gamma)$. Abusing notation, for any $\mathbf{a}, \mathbf{a}' \in \mathcal{A}$, write $\mathbf{a} \geq \mathbf{a}'$ 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 γ , $U(\mathbf{a}, \gamma)$ is supermodular in \mathbf{a} .*

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

$$(\mathbf{a} \vee \mathbf{b})^n = \begin{cases} \max(a_t, b_t) & \text{if } t \leq n, \\ a_t & \text{if } t > n. \end{cases} \quad (8)$$

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

$$(\mathbf{a} \wedge \mathbf{b})^n = \begin{cases} \min(a_t, b_t) & \text{if } t \leq n, \\ a_t & \text{if } t > n. \end{cases} \quad (9)$$

Also, define \mathbf{b}^n as

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

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

The supermodularity implies the consumer has the “greatest” optimal choice. This is the consumer’s optimal choice I select using the tie-breaking rule.

Lemma 5. *For each γ , the set $\mathcal{A}^*(\gamma)$ of optimal choices is a sublattice of \mathcal{A} . There is an $\bar{\mathbf{a}} \in \mathcal{A}^*(\gamma)$ such that for any $\mathbf{a} \in \mathcal{A}^*(\gamma)$, $\bar{\mathbf{a}} \geq \mathbf{a}$.*

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

$$\mathcal{A}_t^*(\gamma) := \{a \in A : \exists \mathbf{a}^* \in \mathcal{A}^*(\gamma) \text{ s.t. } \mathbf{a}_t^* = a\}. \quad (11)$$

For each $k \in \mathbb{N}$, let \mathbf{a}^k denote an optimal policy such that the consumer chooses $a_k = \max \mathcal{A}_k^*(\gamma)$ in period k . Define $\bar{\mathbf{a}}^k := \mathbf{a}^1 \vee \dots \vee \mathbf{a}^k$. Because $\mathcal{A}^*(\gamma)$ is sublattice, for any $k \in \mathbb{N}$, $\bar{\mathbf{a}}^k$ maximizes (7). We also have $\bar{\mathbf{a}}^k \rightarrow \bar{\mathbf{a}}$, where $\bar{a}_t = \max \mathcal{A}_t^*(\gamma)$ for any $k \in \mathbb{N}$. Because $\mathcal{A}^*(\gamma)$ is compact, $\bar{\mathbf{a}} \in \mathcal{A}^*(\gamma)$. By construction, for any $\mathbf{a} \in \mathcal{A}^*(\gamma)$, $\bar{\mathbf{a}} \geq \mathbf{a}$. \square

For each $\gamma \in \Gamma$, let $\bar{\mathbf{a}}(\gamma) := (\bar{a}_t(\gamma))_{t \in \mathbb{N}}$ denote the greatest strategy of the consumer defined in [Lemma 5](#).

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

Proof. By [Lemma 3](#), $\mathcal{A}^*(\gamma)$ is upper hemicontinuous, so the set $\mathcal{A}_t^*(\gamma)$ of all activity levels that can be chosen in period t is upper hemicontinuous in γ . Thus, it is enough to show that for any upper hemicontinuous and compact-valued correspondence $\phi : X \rightarrow \mathbb{R}$, $f(x) := \max \phi(x)$ is upper semicontinuous. To show this, take any $x_n \rightarrow x$. For each n , define $y_n = f(x_n)$. Because there is a subsequence $y_{n(k)}$ of y_n that converges to $\limsup y_n$, it holds that $\limsup y_n = \lim y_{n(k)} = \lim f(x_{n(k)}) \leq f(\lim x_{n(k)}) = f(x)$. The inequality holds because ϕ has a closed graph. Connecting the left and right sides, we establish that $f(\cdot)$ is upper semicontinuous. \square

Lemma 7. *There exists an equilibrium in the game of long-run commitment power.*

Proof. The platform's objective is

$$\sum_{t=1}^{\infty} \delta_P^{t-1} \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^t \frac{1}{\bar{a}_s(\gamma) + \gamma_s}} \right). \quad (12)$$

To show it is upper semicontinuous, take $\gamma^n \rightarrow \gamma$. Then,

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

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

D Consumer Behavior Under a Stationary Privacy Policy:

Proof of Proposition 1

This Appendix uses notations introduced at the beginning of [Appendix C](#).

D.1 Properties of Consumer Value Function

First, I prove useful properties of the consumer's value function that hold for any privacy policy.

Let $(\bar{a}_t(\gamma))_{t \in \mathbb{N}}$ denote the greatest best response of the consumer constructed in [Lemma 5](#). For each privacy policy $\gamma \in \Gamma$, define

$$V_\gamma(\rho) := \sum_{t=1}^{\infty} \delta_C^{t-1} \left[u(\bar{a}_t(\gamma)) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho + \sum_{s=1}^t \frac{1}{\frac{1}{\bar{a}_s(\gamma)} + \gamma_s}} \right) \right]. \quad (13)$$

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

Lemma 8. *For any $\gamma \in \Gamma$, $V_\gamma(\cdot) : \mathbb{R}_{++} \rightarrow \mathbb{R}$ is decreasing and convex. For any $\rho > 0$ and $\Delta > 0$, $\lim_{\rho \rightarrow \infty} V_\gamma(\rho) - V_\gamma(\rho + \Delta) = 0$.*

Proof. Fix any privacy policy γ . Hereafter, I omit γ 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 $\frac{1}{\sigma_{t-1}^2} = \rho$:

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

Here, $\rho_{t-1} = \rho$, and $(\rho_t, \dots, \rho_{T-1})$ are recursively defined by Bayes' rule given (a_t, \dots, a_{T-1}) .

The standard argument of dynamic programming implies that for each $t = 1, \dots, T$,

$$V_t^T(\rho) = \max_{a \in A} u(a) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a} + \gamma_t}} \right) + \delta_C V_{t+1}^T \left(\rho + \frac{1}{\frac{1}{a} + \gamma_t} \right), \quad (14)$$

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

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

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

$$V^T(\rho) \geq \sum_{s=1}^T \delta_C^{s-1} \left(u(\bar{a}_s) - v \left(\sigma_0^2 - \frac{1}{\rho_{s-1} + \frac{1}{\frac{1}{\bar{a}_s} + \gamma_s}} \right) \right). \quad (15)$$

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

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

D.2 Proof of Proposition 1

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

$$V(\rho) = \max_{a \in A} u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a} + \gamma}} \right) + \delta_C V \left(\rho + \frac{1}{\frac{1}{a} + \gamma} \right). \quad (16)$$

Again, I suppress the dependence of $V(\cdot)$ on γ . Lemma 8 implies that $V(\cdot)$ is decreasing and convex. Thus, the maximand in (16) has the increasing differences in (a, ρ) . Thus, $\bar{a}(v, \gamma, \rho)$, the greatest maximizer, is increasing in ρ . Note that $\rho_t \leq \rho_{t+1}$, and the inequality is strict if and only

if $a_t > 0$. As a result, the consumer's optimal behavior is either (i) $a_t = 0$ for all t , or (ii) $a_1 > 0$ and a_t is increasing in t . Now, define

$$v^*(\gamma) := \sup \{v \in \mathbb{R} : \bar{a}(v, \gamma, \rho_0) > 0\}, \quad \text{where } \rho_0 = \frac{1}{\sigma_0^2}. \quad (17)$$

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

If $v < v^*(\gamma)$, then $a_t \geq a_1 > 0$ for all t . Since $\gamma < \infty$, we obtain $\lim_{t \rightarrow \infty} \sigma_t^2 \rightarrow 0$, or equivalently, $\lim_{t \rightarrow \infty} \rho_t = \infty$ with $\rho_t := \frac{1}{\sigma_t^2}$. By Lemma 8, for any $\rho > 0$ and $\Delta > 0$, $\lim_{\rho \rightarrow \infty} V(\rho) - V(\rho + \Delta) = 0$. This, combined with $\lim_{t \rightarrow \infty} \rho_t = \infty$, implies $\lim_{t \rightarrow \infty} \bar{a}_t(v, \gamma, \rho_t) = a_{max}$. \square

E The Optimal Policy Under Long-Run Commitment: Proof of Theorem 1

E.1 Lemmas

Proof of Lemma 2. Let β be any one of γ and γ' . I decompose the consumer's problem (7) into two steps. First, given any $(a_t)_{t \in \mathcal{T}}$, the consumer chooses $(a_t)_{t \notin \mathcal{T}}$ to maximize the following hypothetical objective function:

$$\sum_{t=1}^{\infty} \delta_C^{t-1} \left[\mathbf{1}_{\{t \notin \mathcal{T}\}} u(a_t) - v \cdot \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \sum_{s=1}^t \frac{1}{a_s + \beta_s}} \right) \right]. \quad (18)$$

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

For any $t \in \mathcal{T}$, a_t affects (18) only through $\frac{1}{a_t} + \gamma_t$, because $\mathbf{1}_{\{t \notin \mathcal{T}\}} = 0$. Also the same argument as in the proof of Lemma 4 implies that (18) is supermodular in $\left((a_t)_{t \notin \mathcal{T}}, \left\{ \left(\frac{1}{a_s} + \gamma_s \right)^{-1} \right\}_{s \in \mathcal{T}} \right)$.

This implies that if $\frac{1}{\bar{a}_t(\gamma)} + \gamma_t \leq \frac{1}{\bar{a}_t(\gamma')} + \gamma'_t$ for all $t \in \mathcal{T}$, then $\bar{a}_t(\gamma) \geq \bar{a}_t(\gamma')$ for all $t \in \mathbb{N} \setminus \mathcal{T}$. \square

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

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

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

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

First, we have $\lim_{\gamma_{max} \rightarrow \infty} D(\sigma_0^2, \gamma_{max}) = 0$, and $D(\sigma_t^2, \gamma_{max}) \leq D(\sigma_0^2, \gamma_{max})$ for any $\sigma_t^2 \leq \sigma_0^2$. Thus, for any γ_{max} such that $D(\sigma_0^2, \gamma_{max}) < u(a_{max}) - u(a')$, the consumer's optimal action is a_{max} in period t . Also, even for $\gamma_{max} = 0$, $\lim_{\sigma_{t-1}^2 \rightarrow 0} D(\sigma_{t-1}^2, 0) = 0$. Thus for a sufficiently small σ_{t-1}^2 , the consumer chooses $a_\tau = a_{max}$ for all $\tau \geq t$ under any (continuation) privacy policy. \square

E.2 Proof of Theorem 1

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

Let $\gamma_{max} \in [0, +\infty)$ denote the privacy level defined in Lemma 9—i.e., the consumer chooses $a_t = a_{max}$ if $\gamma_t = \gamma_{max}$. If the platform commits to $\gamma_t = \gamma_{max}$, the variance of the noise of the signal in period t is $B := \frac{1}{a_{max}} + \gamma_{max}$. Take T^* such that for all $t \geq T^*$, $\frac{1}{a_t^*} + \gamma_t^* > B$. If the

platform replaces γ_t^* with γ_{max} for all $t \geq T^*$ and commits to such a new policy ex ante, then the precision of the signal increases from $\frac{1}{\frac{1}{a_t^*} + \gamma_t^*}$ to B^{-1} in any period $t \geq T^*$. [Lemma 2](#) implies that after the policy change, the consumer also chooses a weakly greater a_t for all $t < T^*$. To sum up, the platform can strictly increase its profit by replacing γ_t^* with γ_{max} for all $t \geq T^*$, which is a contradiction. The second part of [Lemma 9](#) then implies that there is some T such that for all $t \geq T$, $a_t^* = a_{max}$.

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

$$v \left(\frac{1}{\frac{1}{\sigma_0^2} + (t^* - 1)a_{max}} - \frac{1}{\frac{1}{\sigma_0^2} + (t^* - 1)a_{max} + a_{min}} \right) > 0,$$

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

F General Payoffs: Proof of [Proposition 2](#)

Proof of Proposition 2. First, we show Point 1. Note that Point 1 is a “non-equilibrium” result and only shows the existence of the platform’s policy that induces the long-run privacy loss. Define $v := \sup_{a, x \in [0, \sigma_0^2]} \left| \frac{\partial \hat{u}}{\partial \sigma^2}(a, x) \right|$ and $C(x) = v(\sigma_0^2 - x)$. Define $u(a) = \beta a$ for $\beta > 0$ that satisfies $\hat{u}(a, \sigma^2) - \hat{u}(a', \sigma^2) \geq \beta$ for any σ^2 and $a > a'$. We have $C'(\sigma^2) = -v \leq -\frac{\partial \hat{u}}{\partial \sigma^2}(a, \sigma^2)$ and $\hat{u}(a', \sigma^2) - \hat{u}(a, \sigma^2) \geq u(a') - u(a) > 0$ for any $a' > a$. As a result, the consumer incurs a lower marginal privacy cost and a higher marginal gross benefit under \hat{u} than under (u, C) , uniformly across all σ^2 . [Proposition 1](#) implies that there is a γ^* such that if the platform commits to $\gamma_t = \gamma^*$ for all $t \in \mathbb{N}$, the consumer’s optimal behavior induces $a_t \rightarrow a_{max}$ and $\sigma_t^2 \rightarrow 0$ when she has

(u, C) . Suppose to the contrary that σ_t^2 does not converge to 0 when the consumer with \hat{u} acts optimally. Then for some $T \in \mathbb{N}$, $a_t = 0$ for all $t \geq T$. In any period $t \geq T$, the consumer with (u, C) strictly prefers some $a_t > 0$ to $a_t = 0$ (i.e., Point 1 of [Proposition 1](#)). The consumer with \hat{u} can mimic this strategy to strictly increase her continuation value relative to taking zero activity levels forever. This is a contradiction. Finally, $\sigma_t^2 \rightarrow 0$ implies $a_t \rightarrow a_{max}$ by the same argument as [Lemma 9](#).

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

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

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

G Omitted Proofs for [Section 5](#)

Proof of [Proposition 3](#)

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

To show the first part, take any pure-strategy MPE. Because σ_t^2 is non-negative and decreasing, it converges to a non-negative value. If $\sigma_t^2 \rightarrow 0$, then $\sigma_s^2 < B$ for a sufficiently large s , so the second part of the result implies Point 1.

Next, suppose that σ_t^2 converges to a positive value. We show $a_t > 0$ for all t . Let t^* denote the smallest t such that $a_t = 0$ in equilibrium. Then $\sigma_{t^*+1}^2 = \sigma_{t^*}^2$. Because we consider an MPE, we

can inductively show that $a_s = 0$ for all $s \geq t^*$. However, the platform will then have a profitable deviation: In period t^* , the platform can choose a sufficiently high but finite γ_t . If the consumer chooses $a_{t^*} = 0$, then $a_s = 0$ for all $s \geq t^*$ according to the equilibrium strategy (following the platform's deviation). Compared to this outcome, the consumer is strictly better off if she chooses $a_{t^*} > 0$ and $a_s = 0$ for all $s \geq t^* + 1$. This is a contradiction, so $a_t > 0$ for all t . Now, suppose to the contrary that γ_t does not diverge to ∞ . It means that there is some $\bar{\gamma} < \infty$ and the platform chooses $\gamma_t \leq \bar{\gamma}$ for infinitely many t 's. Because A is finite, the lowest positive activity level is well-defined. Therefore the platform can use the signals in these periods to perfectly learn the consumer's type in the long run, i.e., $\sigma_t^2 \rightarrow 0$. This is a contradiction, so we have $\gamma_t \rightarrow \infty$. \square

Consumer-Worst Equilibrium: Proof of Theorem 2

Proof. We use precision $\rho_t = \frac{1}{\sigma_t^2}$ as a state variable of MPE. Let $\gamma(\rho)$ denote the platform's choice of γ_t given $\rho_{t-1} = \rho$, and let $a(\rho, \gamma)$ denote the consumer's choice of a_t given $(\rho_{t-1}, \gamma_t) = (\rho, \gamma)$. We adopt notations $\frac{1}{0} = \infty$ and $\frac{1}{\infty} = 0$. We construct $(W(\cdot), V(\cdot, \cdot), \Pi(\cdot), a(\cdot, \cdot), \gamma(\cdot))$ that satisfies the following functional equations: For all $\rho \in [\frac{1}{\sigma_0^2}, \infty)$ and $\gamma \geq 0$,

$$W(\rho, \gamma) = \max_{a \in \{0, a_{max}\}} \left\{ u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a} + \gamma}} \right) + \delta V \left(\rho + \frac{1}{\frac{1}{a} + \gamma} \right) \right\}, \quad (19)$$

$a(\rho, \gamma)$ is the largest maximizer of the right-hand side of (19),

$$V(\rho) = W(\rho, \gamma(\rho)), \quad (20)$$

$$\Pi(\rho) = \max_{\gamma \geq 0} \left\{ \sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a(\rho, \gamma)} + \gamma}} + \delta \Pi \left(\rho + \frac{1}{\frac{1}{a(\rho, \gamma)} + \gamma} \right) \right\}, \quad \text{and} \quad (21)$$

$$\gamma(\rho) \text{ is the largest maximizer of the right-hand side of (21)}. \quad (22)$$

Let V^* denote the solution of

$$V(\rho) = \max_{a \in \{0, a_{max}\}} \left\{ u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + a} \right) + \delta V(\rho + a) \right\}. \quad (23)$$

This Bellman equation describes the consumer's choice when the platform always sets $\gamma_t = 0$. Blackwell's condition for a contraction implies that the equation has a unique solution V^* . The

relevant contraction mapping maps any decreasing convex function to a decreasing convex function. Thus V^* is decreasing and convex. Let $W(\rho, \gamma)$ satisfy (19) with $V = V^*$. If the consumer increases a from 0 to a_{max} , she incurs a flow privacy cost and a loss from a lower continuation value. These losses are smaller when ρ and γ are higher, so the maximizer $a(\rho, \gamma)$ is increasing in ρ and γ . Given $a(\rho, \gamma)$, equation (21) defines a Bellman equation. The similar argument as equation (19) implies that we have a unique solution $\Pi(\cdot)$ to equation (21), and it is increasing. Equation (22) then determines $\gamma(\rho)$.

The remaining equation is (20). Consider the platform's problem in the right-hand side of (21). Note that $\Pi(\cdot)$ is increasing. If $a(\rho, 0) = a_{max}$, the platform chooses $\gamma(\rho) = 0$. If $a(\rho, 0) = 0$, the platform chooses the lowest γ such that $a(\rho, \gamma) = a_{max}$, so the consumer is indifferent between a_{max} and 0. In either case, the right-hand side of (19) evaluated at $\gamma = \gamma(\rho)$ is equal to the one such that she chooses a optimally against $\gamma = 0$. As a result we have

$$W(\rho, \gamma(\rho)) = \max_{a \in \{0, a_{max}\}} \left\{ u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + a} \right) + \delta V^*(\rho + a) \right\} = V^*(\rho).$$

Thus we have constructed $(W(\cdot), V(\cdot, \cdot), \Pi(\cdot), a(\cdot, \cdot), \gamma(\cdot))$ that satisfies equations (19)-(22).

We show that $(\gamma(\cdot), a(\cdot, \cdot))$ consists of an equilibrium. Equations (19) and (20) imply that the consumer has no profitable one-shot deviation after any history. Equation (21) implies the platform has no profitable one-shot deviation. The one-shot deviation principle implies $(\gamma(\cdot), a(\cdot, \cdot))$ consists of an equilibrium.

As shown above, in this equilibrium the consumer's payoff is the same as when she faces the platform that sets $\gamma_t = 0$ after any history. Thus the equilibrium is consumer-worst. Note that the platform's strategy is $\gamma(\rho) = \min \{ \gamma \in \mathbb{R}_+ : a(\rho, \gamma) = a_{max} \}$, so Point 2 holds. Both $a(\rho, \gamma)$ and $\gamma(\rho)$ are independent of δ_P ; therefore, the equilibrium strategy is independent of δ_P .

Finally, we show Point 1. Because $a(\rho, \gamma)$ is increasing in ρ , $\gamma(\rho)$ is decreasing in ρ . Because ρ_t is increasing on-path, γ_t decreases over time. Moreover, we have $\rho_{t+1} \geq \rho_0 + t \cdot \frac{1}{\frac{1}{a_{max}} + \gamma(\rho_1)}$, and thus $\lim_{t \rightarrow \infty} \rho_t = \infty$. Proposition 3 implies $\gamma_t = 0$ after some finite t . \square

Proof of Corollary 1

We describe the game of action-contingent commitment. Before $t = 1$, the platform publicly commits to a mapping $\gamma(\cdot) : \{\phi\} \cup (\cup_{s=1}^{\infty} A^s) \rightarrow \mathbb{R}_+$, which determines $\gamma_1 = \gamma(\phi)$ and maps past actions $(a_1, \dots, a_{t-1}) \in A^{t-1}$ to the privacy level γ_t in every period $t \geq 2$. The consumer observes $\gamma(\cdot)$ and chooses activity levels over time.

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

We now show that $(\gamma_t^*)_{t \in \mathbb{N}}$ is indeed optimal. We begin with two observations. First, multiplying a positive constant to the platform's payoff does not change the optimal policy. Thus we can assume that the platform's payoff equals the consumer's privacy cost, $v[\sigma_0^2 - \sigma_t^2]$. Second, recall that in the consumer-worst equilibrium, the consumer's ex ante payoff is $V(\rho_0)$, which is her ex ante payoff when she acts optimally against a (hypothetical) platform that always sets $\gamma_t = 0$ in any period. The consumer can secure a payoff of at least $V(\rho_0)$ regardless of the platform's behavior. Thus under long-run commitment and action-contingent commitment, the consumer's ex ante payoff satisfies

$$\sum_{t=1}^{\infty} \delta^{t-1} u(a_t) - \sum_{t=1}^{\infty} \delta^{t-1} \pi_t \geq V(\rho_0),$$

where π_t is the privacy cost (which equals the platform's payoff) in period t . As a result, the

platform's ex ante payoff satisfies

$$\sum_{t=1}^{\infty} \delta^{t-1} u(a_{max}) - V(\rho_0) \geq \sum_{t=1}^{\infty} \delta^{t-1} \pi_t. \quad (24)$$

In the consumer-worst equilibrium, the consumer chooses a_{max} in all periods and obtains ex ante payoff $V(\rho_0)$. Thus the platform's payoff is exactly the upper bound $\sum_{t=1}^{\infty} \delta^{t-1} u(a_{max}) - V(\rho_0)$ in (24). Since the platform with long-run commitment can attain the same outcome with $(\gamma_t^*)_{t \in \mathbb{N}}$, it is an optimal policy under long-run commitment. The same upper bound $\sum_{t=1}^{\infty} \delta^{t-1} u(a_{max}) - V(\rho_0)$ applies to the case of action-contingent commitment. Therefore the consumer-worst equilibrium gives the platform the best possible payoff under action-contingent commitment. \square

Consumer-Best Equilibrium: Proof of Theorem 3

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

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

$$u(a_{max}) - v \left(\frac{1}{\rho_0} - \frac{1}{\rho_0 + \frac{1}{\frac{1}{a_{max}} + \gamma}} \right) = u(a_{max}) - v \frac{\frac{1}{\frac{1}{a_{max}} + \gamma}}{\rho_0 \left(\rho_0 + \frac{1}{\frac{1}{a_{max}} + \gamma} \right)} \geq 0. \quad (25)$$

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

Recall from the proof of Theorem 2 that in the consumer-worst equilibrium, there is a cutoff state $\rho(0)$ such that if $\rho_t \geq \rho(0)$, $(\gamma, a) = (0, a_{max})$ is chosen in the continuation game, and if

$\rho_t < \rho(0)$, the consumer's continuation value is the same as taking $a_t = 0$ for all t .

Take any $\bar{\rho} > 0$ such that $\bar{\rho} + \frac{1}{\frac{1}{a_{max}} + \hat{\gamma}(\bar{\rho})} \leq \rho(0)$. For any initial state $\rho_0 \leq \bar{\rho}$, the above strategy profile is an equilibrium. First, it is an equilibrium at any (off-path) state $\rho > \rho_0$ by construction. At $\rho = \rho_0$, the consumer has no profitable deviation when the platform offers $\gamma = \infty$, because she can receive the best payoff of $u(a_{max})$ in the current and any future periods. Suppose that the platform deviates for the first time in period t and chooses $\gamma_t < \infty$. Suppose to the contrary that the consumer strictly benefits from the one-shot deviation to $a = a_{max}$. Then $\rho_0 + \frac{1}{\frac{1}{a_{max}} + \gamma_t} \leq \rho(0)$ must hold, because $\rho_0 \leq \bar{\rho}$ and the platform chooses $\gamma_t \geq \hat{\gamma}(\rho_0) \geq \hat{\gamma}(\bar{\rho})$. Thus her payoff in period t is at most $u(a_{max})$, whereas her continuation value from period $t + 1$ is equal to the one from choosing $a_s = 0$ for all $s \geq t + 1$. In contrast, if the consumer chooses $a_t = 0$ and follows her strategy thereafter, her payoff is $\frac{\delta}{1-\delta}u(a_{max})$, because she sets $a_t = 0$ in period t and the state remains ρ_0 . Thus, the consumer has a profitable deviation only if $\frac{\delta}{1-\delta}u(a_{max}) < u(a_{max})$, which contradicts $\delta \geq 1/2$. \square

Discussion on Assumption 1

The assumption of binary actions is without loss of generality if $u(a)$ is weakly convex (or precisely, $u(\cdot) : A \rightarrow \mathbb{R}$ is a restriction of a weakly convex function $\hat{u} : \mathbb{R}_+ \rightarrow \mathbb{R}$ to A). Under this assumption, the consumer chooses 0 or a_{max} as an optimal activity level, so we can solve the game as if she has binary actions. To see this, consider the proof of [Theorem 2](#). We can extend equations (19)-(23) by replacing “ $a \in \{0, a_{max}\}$ ” with “ $a \in A$.” Without any restriction on A or $u(a)$, the value function V that solves (23) is decreasing and convex, $a(\rho, \gamma)$ is increasing in ρ and γ , and Π is increasing. However, we need [Assumption 1](#) to verify (20), which implies that V is indeed the consumer's value function when the platform follows $\gamma(\cdot)$. For a general A and $u(a)$, there is no guarantee that (20) holds, i.e., the consumer's value function $W(\rho, \gamma(\rho))$ is not necessarily equal to $V(\rho)$, which is her optimal payoff facing $\gamma_t \equiv 0$. However, if $u(a)$ is weakly convex, the maximands of the consumer's Bellman equation (19) is strictly convex in a . Thus for any ρ , we have $a(\rho, \gamma) = 0$ if γ is below some threshold, and $a(\rho, \gamma) = a_{max}$ if γ is above the threshold. The platform's Markov strategy $\gamma(\rho)$ is then defined as the lowest γ that induces the consumer to choose a_{max} . Therefore we obtain (20) by the same logic as in the original proof.

We can also prove [Theorem 3](#) for a general A if $u(a)$ is weakly convex. The only part of the

proof we need to modify is that we redefine $\hat{\gamma}(\rho_0)$ in the proof as the lowest γ under which the consumer's flow payoff from an optimal $a \in A$ is non-negative when she faces (ρ_0, γ) .

H Speed of Learning: Appendix for Section 6.1

Proof of Proposition 4. Fix any privacy policy of the platform and all parameters except the consumer's discount factor. Take two discount factors of the consumer, δ_V and $\delta_W > \delta_V$. As in the proof of Lemma 8, we begin with 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 $\frac{1}{\sigma_{t-1}^2} = \rho$:

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

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

$$V_t^T(\rho) = \max_{a \in A} u(a) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a} + \gamma_t}} \right) + \delta_V V_{t+1}^T \left(\rho + \frac{1}{\frac{1}{a} + \gamma_t} \right), \quad (26)$$

where $V_{T+1}^T(\cdot) \equiv 0$. Similarly if the consumer has discount factor δ_W , the relevant value function W_t^T satisfies

$$W_t^T(\rho) = \max_{a \in A} u(a) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a} + \gamma_t}} \right) + \delta_W W_{t+1}^T \left(\rho + \frac{1}{\frac{1}{a} + \gamma_t} \right), \quad (27)$$

where $W_{T+1}^T(\cdot) \equiv 0$. We can recursively show that V_t^T and W_t^T are decreasing and convex for any t . We show that for any $x, y \geq 0$ such that $y > x$, $W_k^T(y) - W_k^T(x) \leq V_k^T(y) - V_k^T(x)$ for all $k = 1, \dots, T + 1$. The inequality holds with equality for $k = T + 1$. Suppose it holds for $k = t + 1$. Let $a_t^V(\rho)$ and $a_t^W(\rho)$ denote the maximizers for (26) and (27), respectively. The envelope theorem

implies

$$V_t^T(y) - V_t^T(x) = - \int_x^y \frac{v}{\left(\rho + \frac{1}{\frac{1}{a_t^V(\rho)} + \gamma_t}\right)^2} + \delta_V \frac{dV_{t+1}^T}{d\rho} \left(\rho + \frac{1}{\frac{1}{a_t^V(\rho)} + \gamma_t}\right) d\rho \quad (28)$$

and

$$W_t^T(y) - W_t^T(x) = - \int_x^y \frac{v}{\left(\rho + \frac{1}{\frac{1}{a_t^W(\rho)} + \gamma_t}\right)^2} + \delta_W \frac{dW_{t+1}^T}{d\rho} \left(\rho + \frac{1}{\frac{1}{a_t^W(\rho)} + \gamma_t}\right) d\rho. \quad (29)$$

The induction hypothesis states that W_{t+1}^T decreases more rapidly than V_{t+1}^T , which implies $a_t^W(\rho) \leq a_t^V(\rho)$. Thus we have $W_t^T(y) - W_t^T(x) \leq V_t^T(y) - V_t^T(x)$. By induction, we have $V_1^T(y) - V_1^T(x) \geq W_1^T(y) - W_1^T(x)$. Taking $T \rightarrow \infty$, we have $W_1(y) - W_1(x) \leq V_1(y) - V_1(x)$. We can apply the same argument to the consumer's problem that starts from period t and conclude that $W_t(y) - W_t(x) \leq V_t(y) - V_t(x)$ for any $t \geq 1$.

In $t = 1$, the consumer's problems under δ_V and δ_W are respectively

$$\max_{a \in A} \left[u(a) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho_0 + \frac{1}{\frac{1}{a} + \gamma_1}} \right) + \delta_W W_2 \left(\rho_0 + \frac{1}{\frac{1}{a} + \gamma_1} \right) \right] \quad (30)$$

and

$$\max_{a \in A} \left[u(a) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho_0 + \frac{1}{\frac{1}{a} + \gamma_1}} \right) + \delta_V V_2 \left(\rho_0 + \frac{1}{\frac{1}{a} + \gamma_1} \right) \right]. \quad (31)$$

Because $W_2(\cdot)$ decreases more rapidly than $V_2(\cdot)$, we have $a_1^W(\rho) \leq a_1^V(\rho)$ and $\rho_1^W \leq \rho_1^V$. Suppose now that inequalities $a_k^W(\rho) \leq a_k^V(\rho)$ and $\rho_k^W \leq \rho_k^V$ hold up to period $k = t - 1$. In period t , the consumer solves

$$\max_{a \in A} \left[u(a) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho_t^W + \frac{1}{\frac{1}{a} + \gamma_t}} \right) + \delta_W W_{t+1} \left(\rho_{t-1}^W + \frac{1}{\frac{1}{a} + \gamma_t} \right) \right] \quad (32)$$

and

$$\max_{a \in A} \left[u(a) - v \cdot \left(\sigma_0^2 - \frac{1}{\rho_t^V + \frac{1}{\frac{1}{a} + \gamma_t}} \right) + \delta_V V_{t+1} \left(\rho_{t-1}^V + \frac{1}{\frac{1}{a} + \gamma_t} \right) \right]. \quad (33)$$

Thus we have $a_{t+1}^W(\rho) \leq a_{t+1}^V(\rho)$ and $\rho_{t+1}^W \leq \rho_{t+1}^V$.

The same argument implies that if v is higher, the consumer chooses a lower activity level in every period.

Finally, suppose the prior variance σ_0^2 of the consumer's type increases, or equivalently, ρ_0 decreases. We use V for the continuation value of the consumer. In $t = 1$, the consumer solves

$$\max_{a \in A} u(a) + \frac{v}{\rho_0 + \frac{1}{\frac{1}{a} + \gamma_1}} + \delta V_2 \left(\rho_0 + \frac{1}{\frac{1}{a} + \gamma_1} \right). \quad (34)$$

Because V is decreasing and convex, the consumer chooses a lower activity level when ρ_0 is small. Thus if ρ_0 is smaller, ρ_1 becomes smaller and the consumer chooses a lower activity level in $t = 2$. Repeating the same argument, we can show that the consumer chooses a lower activity level in every period when she starts from a greater σ_0^2 . \square

Proof of Proposition 5. Take two discount factors of the consumer, δ_V and $\delta_W > \delta_V$. Let V and W satisfy [equation \(23\)](#) in the proof of [Theorem 2](#) given δ_V and δ_W , respectively:

$$V(\rho) = \max_{a \in \{0, a_{max}\}} \left\{ u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + a} \right) + \delta_V V(\rho + a) \right\}, \quad (35)$$

$$W(\rho) = \max_{a \in \{0, a_{max}\}} \left\{ u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + a} \right) + \delta_W W(\rho + a) \right\}. \quad (36)$$

We show $\delta_W > \delta_V$ implies W decreases more rapidly than V , i.e., $W(y) - W(x) \leq V(y) - V(x)$ for all $y > x \geq 0$. For $\delta \in \{\delta_W, \delta_V\}$, define an operator \mathcal{F}_δ as follows:

$$\mathcal{F}_\delta h(\rho) = \max_{a \in \{0, a_{max}\}} \left\{ u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + a} \right) + \delta h(\rho + a) \right\}. \quad (37)$$

The operator \mathcal{F}_δ maps any decreasing and convex h to a function with the same property. We show that if h decreases more rapidly than g , i.e., if $h(y) - h(x) \leq g(y) - g(x)$ for all $y > x \geq 0$, then $\mathcal{F}_{\delta_W} h$ decreases more rapidly than $\mathcal{F}_{\delta_V} g$. The envelope theorem implies

$$\begin{aligned} \mathcal{F}_{\delta_W} h(y) - \mathcal{F}_{\delta_W} h(x) &= \int_x^y \frac{-v}{(\rho + a_{W,h}(\rho))^2} + \delta_W h'(\rho + a_{W,h}(\rho)) d\rho, \quad \text{and} \\ \mathcal{F}_{\delta_V} g(y) - \mathcal{F}_{\delta_V} g(x) &= \int_x^y \frac{-v}{(\rho + a_{V,g}(\rho))^2} + \delta_V g'(\rho + a_{V,g}(\rho)) d\rho, \end{aligned}$$

where

$$a_{W,h}(\rho) \in \arg \max_{a \in \{0, a_{max}\}} \left\{ u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + a} \right) + \delta_W h(\rho + a) \right\} \quad (38)$$

$$a_{V,g}(\rho) \in \arg \max_{a \in \{0, a_{max}\}} \left\{ u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + a} \right) + \delta_V g(\rho + a) \right\}. \quad (39)$$

We have $a_{W,h}(\rho) \leq a_{V,g}(\rho)$ because $\delta_W h(\rho + a)$ decreases more rapidly than $\delta_V g(\rho + a)$. Thus we have $\mathcal{F}_{\delta_W} h(y) - \mathcal{F}_{\delta_W} h(x) \leq \mathcal{F}_{\delta_V} g(y) - \mathcal{F}_{\delta_V} g(x)$. For any $n \in \mathbb{N}$, we have $\mathcal{F}_{\delta_W}^n V(y) - \mathcal{F}_{\delta_W}^n V(x) \leq \mathcal{F}_{\delta_V}^n V(y) - \mathcal{F}_{\delta_V}^n V(x) = V(y) - V(x)$. Taking $n \rightarrow \infty$, we have $W(y) - W(x) \leq V(y) - V(x)$.

In the consumer-worst equilibrium, the platform sets the lowest privacy level that makes the consumer willing to choose a_{max} as opposed to 0. Given δ_W , the platform sets

$$\gamma_W(\rho) = \min \left\{ \gamma \in \mathbb{R}_+ : u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a} + \gamma}} \right) + \delta_W W \left(\rho + \frac{1}{\frac{1}{a} + \gamma} \right) \geq -v \left(\sigma_0^2 - \frac{1}{\rho} \right) + \delta_W W(\rho) \right\}.$$

Given δ_V , the platform sets

$$\gamma_V(\rho) = \min \left\{ \gamma \in \mathbb{R}_+ : u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a} + \gamma}} \right) + \delta_V V \left(\rho + \frac{1}{\frac{1}{a} + \gamma} \right) \geq -v \left(\sigma_0^2 - \frac{1}{\rho} \right) + \delta_V V(\rho) \right\}.$$

The consumer incurs a greater loss from a lower continuation value given δ_W than δ_V , so we have $\gamma_W(\rho) \geq \gamma_V(\rho)$. We now show that the platform chooses a greater privacy level when the consumer is more patient. In period 1, the statement holds because $\gamma_W(\rho_0) \geq \gamma_V(\rho_0)$. Suppose the statement holds up to period $t - 1$. Let $\rho_{t-1}(W)$ and $\rho_{t-1}(V)$ denote the states at the beginning of period t under δ_W and δ_V , respectively. We have $\rho_{t-1}(W) \leq \rho_{t-1}(V)$, because the consumer chooses a_{max} and the platform chooses higher privacy levels before period t . Because γ_W and γ_V are decreasing, we have $\gamma_W(\rho_{t-1}(W)) \geq \gamma_V(\rho_{t-1}(W)) \geq \gamma_V(\rho_{t-1}(V))$. Therefore the platform sets a higher privacy level and obtains a less accurate signal in every period if the consumer's more patient.

Applying the same argument as above, we can show that the platform offers a higher privacy level if v is higher or σ_0^2 is higher. In either case, the platform has to choose a higher privacy level to induce the consumer to choose a_{max} , because she faces a higher incremental cost of increasing

an activity level. □

I Time-Varying Type: Appendix for Section 6.2

Proof of Proposition 6. The consumer's problem is a dynamic programming with state variable $\rho_t := \frac{1}{\sigma_t^2}$. The Bellman equation is

$$V(\rho) = \max_{a \in A} \left\{ u(a) - v \left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{\frac{1}{a} + \gamma}} \right) + \delta_C V \left(\frac{1}{\phi^2 \left(\frac{1}{\rho + \frac{1}{\frac{1}{a} + \gamma}} \right) + (1 - \phi^2)\sigma_0^2} \right) \right\}. \quad (40)$$

Blackwell's sufficient condition ensures that the equation has a unique solution, V^* . We show V is decreasing and convex. The second term of the maximand is decreasing and convex in ρ . The last term has the same property, because it is a composite of a decreasing convex function V and increasing concave function $\phi^2 \left(\frac{1}{\rho + \frac{1}{\frac{1}{a} + \gamma}} \right) + (1 - \phi^2)\sigma_0^2$. The upper envelope of decreasing and convex functions is decreasing and convex, so the right-hand side of the Bellman equation is decreasing and convex in ρ . As a result V^* is decreasing and convex. A similar fixed point argument implies that for any ρ , $V^*(\rho)$ decreases in v .

Given precision ρ in the current period, let $a(\rho)$ denote the consumer's optimal choice derived from the Bellman equation. The maximand of the Bellman equation has increasing differences in (a, ρ) . Thus $a(\rho)$ is increasing. Let $g(\rho)$ denote the precision in the following period:

$$g(\rho) = \frac{1}{\phi^2 \left(\frac{1}{\rho + \frac{1}{\frac{1}{a(\rho)} + \gamma}} \right) + (1 - \phi^2)\sigma_0^2}. \quad (41)$$

Because $a(\rho)$ is increasing, so is $g(\rho)$. If and only if $a(\rho) > 0$, we have $g(\rho_0) > \rho_0$. Thus we have one of the following dynamics. If $a(\rho_0) = 0$, the consumer chooses $a_t = 0$ in any period t . If $a(\rho_0) > 0$, the consumer chooses $a(\rho_t)$ with $\rho_t = g^t(\rho_0)$ in period t . Precision ρ_t and the consumer's activity level $a(\rho_t)$ increase in t . The existence of cutoff $v^*(\gamma)$ and the comparative statics follow from the same proof as [Proposition 1](#). □

Proof of Proposition 7. Let $\delta_C = \delta_P = \delta$. The construction of the equilibrium follows that of

Theorem 2: Consider the following set of functional equations:

$$W(\rho, \gamma) = \max \left\{ u(a_{max}) - v \left(\sigma_0^2 - \frac{1}{\rho + \frac{1}{a_{max} + \gamma}} \right) + \delta V \left(\frac{1}{\phi^2 \left(\frac{1}{\rho + \frac{1}{a_{max} + \gamma}} \right) + (1 - \phi^2)\sigma_0^2} \right), \right. \quad (42)$$

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

$$V(\rho) = W(\rho, \gamma(\rho)), \quad \text{and} \quad (43)$$

$$\Pi(\rho) = \max_{\gamma \geq 0} \left[\sigma_0^2 - \frac{1}{\rho + \frac{1}{a(\rho, \gamma) + \gamma}} + \delta \Pi \left(\frac{1}{\phi^2 \left(\frac{1}{\rho + \frac{1}{a(\rho, \gamma) + \gamma}} \right) + (1 - \phi^2)\sigma_0^2} \right) \right], \quad (44)$$

where $a(\rho, \gamma)$ is the maximizer of (42) and $\gamma(\rho)$ is the maximizer of (44). We show that there exists $(W(\cdot), V(\cdot, \cdot), \Pi(\cdot), a(\cdot, \cdot), \gamma(\cdot))$ that satisfies the above equations. Let V^* denote the solution of

$$V(\rho) = \max \left\{ u(a_{max}) - v \left(\sigma_0^2 - \frac{1}{\rho + a_{max}} \right) + \delta V \left(\frac{1}{\phi^2 \left(\frac{1}{\rho + a_{max}} \right) + (1 - \phi^2)\sigma_0^2} \right), \right. \quad (45)$$

$$\left. -v \left(\sigma_0^2 - \frac{1}{\rho} \right) + \delta V \left(\frac{1}{\phi^2 \left(\frac{1}{\rho} \right) + (1 - \phi^2)\sigma_0^2} \right) \right\}$$

The standard argument of dynamic programming implies that the equation has a unique solution V^* , which is decreasing and convex. Let $V = V^*$ in (42), which also determines the left-hand side $W(\rho, \gamma)$. The maximizer $a(\rho, \gamma)$ is increasing in ρ and γ . Equation (44) uniquely determines $\Pi(\cdot)$. Because the relevant contraction mapping for (44) maps an increasing function to an increasing function, $\Pi(\cdot)$ is increasing. To solve the right-hand side of (44), the platform chooses $\gamma(\rho) = \min \{ \gamma \in \mathbb{R}_+ : a(\rho, \gamma) = a_{max} \}$. Specifically, if $a(\rho, 0) = a_{max}$, the platform chooses $\gamma = 0$. If $a(\rho, 0) = 0$, the platform choose γ that makes the consumer indifferent between a_{max} and 0. As a

result we have

$$\begin{aligned}
W(\rho, \gamma(\rho)) &= \max \left\{ u(a_{max}) - v \left(\sigma_0^2 - \frac{1}{\rho + a_{max}} \right) + \delta V \left(\frac{1}{\phi^2 \left(\frac{1}{\rho + a_{max}} \right) + (1 - \phi^2) \sigma_0^2} \right), \right. \\
&\quad \left. -v \left(\sigma_0^2 - \frac{1}{\rho} \right) + \delta V \left(\frac{1}{\phi^2 \left(\frac{1}{\rho} \right) + (1 - \phi^2) \sigma_0^2} \right) \right\} \\
&= V(\rho).
\end{aligned}$$

We can now construct an equilibrium in which the platform follows $\gamma(\rho)$ and the consumer follows $a(\rho, \gamma)$. Equations (42) and (43) imply that the consumer has no profitable one-shot deviation after any history. Equation (44) implies the platform has no profitable one-shot deviation. The one-shot deviation principle implies $(\gamma(\cdot), a(\cdot, \cdot))$ consists of an equilibrium.

The consumer's ex ante payoff $V(\rho)$ is what she can secure by acting optimally against zero privacy levels. Thus the equilibrium is consumer-worst. It is also platform-best because the sum of the players' payoffs is maximized (i.e., it is $\frac{u(a_{max})}{1-\rho}$) and the consumer's payoff is minimized. The same logic as Corollary 1 implies that the equilibrium outcome is the same as the outcome under long-run commitment, which proves Point 1. Point 2 follow from the same proof as Proposition 6, and Point 3 follows from the derivation of $\gamma(\cdot)$ above.

We show Point 4. Take ϕ and $\hat{\phi} > \phi$. Let \mathcal{F}_ϕ and $\mathcal{F}_{\hat{\phi}}$ denote the relevant contracting mappings for the Bellman equation (45) given persistence ϕ and $\hat{\phi}$, respectively. Also let V_ϕ and $V_{\hat{\phi}}$ denote the corresponding solutions. We show $V_\phi(\rho) \geq V_{\hat{\phi}}(\rho)$ for all ρ . We have $V_\phi(\rho) = \mathcal{F}_\phi[V_\phi](\rho) \geq \mathcal{F}_{\hat{\phi}}[V_\phi](\rho)$. Because \mathcal{F}_ϕ and $\mathcal{F}_{\hat{\phi}}$ are monotone, for any n we have $V_\phi = \mathcal{F}_\phi^n[V_\phi] \geq \mathcal{F}_{\hat{\phi}}^n[V_\phi]$. Taking $n \rightarrow \infty$, we have $V_\phi \geq V_{\hat{\phi}}$, i.e., the consumer is worse off when ϕ is high. The platform is better off when ϕ is high, because the profit is proportional to the consumer's privacy cost given $\delta_C = \delta_P$. □

Online Appendix: Not for Publication

J Numerical Examples for Section 4

J.1 Non-monotonicity of flow payoffs under a stationary privacy policy

Consider the following parametrization: $A = \{0, 1, 3\}$, $u(1) = 1$, $u(3) = 1.04$, $v = 2$, $\delta = 0.9$, $\sigma_0^2 = 0.1$, and $\gamma_t = 0$ for all $t \in \mathbb{N}$. Define $U(a) := \sum_{t=1}^{\infty} \delta^{t-1} \left[u(a) - v \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + ta} \right) \right]$. First, under the consumer's optimal policy, there is some period t^* such that $a_{t^*-1} = 1$ and $a_{t^*} = 3$ if $U(1) > U(3) > U(0)$. The reason is as follows. [Proposition 1](#) states that the optimal policy under a stationary privacy policy is either $a_t = 0$ for all t , or a_t is positive and weakly increasing in t . $U(3) > U(0)$ implies that the consumer chooses the latter, and $U(1) > U(3)$ implies that $a_1 = 1$. Because $a_1 = 1$ and $a_t = 3$ for some finite t , there is a t^* such that $a_{t^*-1} = 1$ and $a_{t^*} = 3$. The flow payoff increases from $t^* - 1$ to t^* if $u(3) - v \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + 1 + 3} \right) > u(1) - v \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + 1} \right)$. The inequality holds if $u(3) - u(1) > B := \frac{3v}{\left(\frac{1}{\sigma_0^2} + 4\right)\left(\frac{1}{\sigma_0^2} + 1\right)}$. We can numerically show that $U(1) \approx 9.17$, $U(3) = 9.13$, $u(3) - u(1) = 0.04$, and $B = 0.039$. Thus we have $U(1) > U(3) > U(0)$ and $u(3) - u(1) > B$, so the consumer receives a higher flow payoff in period t^* than in $t^* - 1$. This example shows that the consumer's flow payoffs are non-monotone, because once a_t hits a_{max} , the flow payoffs strictly decrease in t .

J.2 Non-monotonicity of a_t in equilibrium

[Figure 1](#) depicts the equilibrium dynamics for a myopic consumer. I assume $A = \{0, 0.01, 0.02, \dots, 2\}$ and use [Claim 1](#) in [Appendix K](#) to compute an equilibrium. ([Claim 1](#) also implies that long-run commitment and one-period commitment lead to the same outcome given a myopic consumer.) [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.




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 $\sigma_0^2 = 1$.

J.3 Non-monotonicity of γ_t in equilibrium

Under different parameters, [Figure 2](#) depicts another equilibrium dynamics for a myopic consumer. [Figure 2\(a\)](#) shows that γ_t can be non-monotone. In particular, the platform increases a privacy level from $t = 1$ to $t = 2$ because it becomes less costly to induce the highest activity level through privacy protection.

K Myopic Consumer

I characterize the equilibrium under a myopic consumer, which facilitates numerical analysis. Let $a^*(\gamma, \sigma^2) \in A$ denote the best response of a myopic consumer, given a privacy level γ in the current period and the posterior variance σ^2 from the previous period:

$$a^*(\gamma, \sigma^2) := \max \left\{ \arg \max_{a \in A} \left[u(a) - v \left(\sigma_0^2 - \frac{1}{\frac{1}{\sigma^2} + \frac{1}{a} + \gamma} \right) \right] \right\}. \quad (46)$$

The following result characterizes the equilibrium.

Claim 1. *Consider the game with long-run commitment. If the consumer is myopic, the plat-*




Figure 2(a): Privacy level γ_t




Figure 2(b): Activity level a_t

Figure 2: Equilibrium under $A = \{0, 1, 2\}$, $u(1) = 10$, $u(2) = 11$, $v = 20$, and $\sigma_0^2 = 1$.

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

$$\gamma_t^* \in \arg \min_{\gamma \geq 0} \frac{1}{a^*(\gamma, \hat{\sigma}_{t-1}^2)} + \gamma, \forall t \in \mathbb{N}, \quad (47)$$

$$\hat{\sigma}_0^2 = \sigma_0^2, \quad (48)$$

$$\hat{\sigma}_t^2 = \frac{1}{\frac{1}{\hat{\sigma}_{t-1}^2} + \frac{1}{\frac{1}{a^*(\gamma_t^*, \hat{\sigma}_{t-1}^2)} + \gamma_t^*}}, \forall t \in \mathbb{N}. \quad (49)$$

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

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

The first inequality follows from the inductive hypothesis and decreasing $a^*(\gamma, \cdot)$. The second inequality follows from (47). We now have $\hat{\sigma}_t^2 \leq \sigma_t^2$ for all t , which implies the privacy policy

described by (47), (48), and (49) is optimal. □

L General Payoffs of the Platform

Most of the results continue to hold if the platform's final payoff from a sequence of posterior variances is $\Pi((\sigma_t^2)_{t \in \mathbb{N}})$, where $\Pi : \mathbb{R}_+^\infty \rightarrow \mathbb{R}$ is coordinate-wise strictly decreasing. This generalization does not change the analysis, because in equilibrium a deviation by the platform increases σ_t^2 for all $t \in \mathbb{N}$. An exception is [Theorem 1](#), where the platform's deviation may not uniformly increase posterior variances. However, the proof of this theorem rests on the argument that if the equilibrium fails to meet certain conditions such as $\sigma_t^2 \rightarrow 0$, the platform can deviate and uniformly decrease posterior variances. Thus, [Theorem 1](#) continues to hold with the same proof under this general $\Pi(\cdot)$.

For example, suppose the platform sells information to a sequence of short-lived data buyers. Any information sold in period t is freely replicable later and thus has a price of zero in any period $s \geq t + 1$. The profit in period t equals the value of information generated in period t —i.e., the platform's ex ante payoff is $\sum_{t=1}^{\infty} \delta_P^{t-1} (\sigma_{t-1}^2 - \sigma_t^2)$, which is decreasing in each σ_t^2 .

M Full Commitment

This appendix considers the platform with action-contingent commitment power: Before $t = 1$, the platform publicly commits to a mapping $\gamma(\cdot) : \{\phi\} \cup (\cup_{s=1}^{\infty} A^s) \rightarrow \mathbb{R}_+$, which determines $\gamma_1 = \gamma(\phi)$ and maps past actions $(a_1, \dots, a_{t-1}) \in A^{t-1}$ to the privacy level γ_t in every period $t \geq 2$.

To provide a condition under which action-contingent commitment benefits the platform, we prepare some notations. First, take any equilibrium under long-run commitment. Let $(\hat{a}_t)_t$, $(\hat{\gamma}_t)_t$, and $(\hat{\sigma}_t^2)_t$ denote the activity levels, privacy levels, and posterior variances at the equilibrium, respectively. Let $\hat{U}_2(\sigma^2)$ denote the consumer's continuation value starting from $t = 2$ when the posterior variance at the beginning of $t = 2$ is σ^2 and the consumer faces $(\hat{\gamma}_t)_{t \geq 2}$. Also, let $U^0(\sigma^2)$ denote the consumer's sum of discounted payoffs when the platform always set $\gamma_t = 0$ and the posterior variance is σ^2 .

Claim 2. *Suppose $\hat{a}_1 < a_{max}$. The platform's payoff under action-contingent commitment is*

strictly greater than the one under long-run commitment if

$$u(a_{max}) - v \left[\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \frac{1}{a_{max} + \hat{\gamma}_1}} \right] + \delta_C \hat{U}_2 \left(\frac{1}{\frac{1}{\sigma_0^2} + \frac{1}{a_{max} + \hat{\gamma}_1}} \right) \quad (50)$$

$$\geq \max_{a \in A} \left\{ u(a) - v \left[\sigma_0^2 - \frac{1}{\frac{1}{\sigma_0^2} + \frac{1}{a + \hat{\gamma}_1}} \right] + \delta_C U^0 \left(\frac{1}{\frac{1}{\sigma_0^2} + \frac{1}{a + \hat{\gamma}_1}} \right) \right\}. \quad (51)$$

Proof. Given the deterministic policy $(\hat{\gamma}_t)_t$ under long-run commitment, we create an action-dependent policy that is strictly better for the platform. Consider the following policy $\gamma^*(\cdot)$. If the consumer chooses $a_1 < a_{max}$ in $t = 1$, the platform sets $\gamma_t = 0$ in any period $t \geq 2$. If the consumer chooses a_{max} in $t = 1$, the platform sets $\hat{\gamma}_t$ in any period $t \geq 2$, i.e., it adopts a deterministic policy from $t = 2$ on. The left-hand side (50) is the consumer's payoff when she chooses a_{max} in $t = 1$ and behave optimally from $t = 2$ on. The right-hand side (51) is the consumer's payoff from the best possible deviation in $t = 1$. Thus the display inequality means that the consumer chooses $a_{max} > \hat{a}_1$ in $t = 1$. Note that the consumer's behavior after $t = 2$ under $\gamma^*(\cdot)$ is different from that under long-run commitment. However, the consumer faces a lower posterior variance in $t = 2$ under the former. Proposition 4 implies that the consumer's activity level under $\gamma^*(\cdot)$ is greater than the one under long-run commitment in any period $t \geq 2$. Thus, $\gamma^*(\cdot)$ gives the platform a higher payoff in any period than under long-run commitment. \square