Distributional Consequences of Privacy Regulation on Platforms

Ben Casner* and Matthew Leisten†

Federal Trade Commission‡

This draft: January 10, 2026

Abstract

What are the welfare consequences of privacy regulation when a media platform offers both an ad-supported free version and an ad-free premium subscription? We present a model where privacy regulation lowers advertiser willingness to pay for advertising. The platform reacts by lowering the premium subscription price while increasing the advertising load imposed on free viewers. Low-income consumers—who consume the ad-supported version because they are more price sensitive and less advertising-averse—are harmed by the privacy regulation while high-income consumers are better off. In extensions we consider cases where wealthier consumers are more valuable to advertisers, where consumers have intrinsic preferences for privacy, and when the platform offers intermediate options with a lower price and less advertising, and provide conditions under which our main qualitative results still hold.

*bencasner.info, bcasner@ftc.gov

†mleisten.github.io, mleisten@ftc.gov

‡The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Trade Commission or any individual Commissioner. The authors thank Austin Drukker, Matthew Gentzkow, Ellie Prager, David Schmidt, Tat-How Teh, Tianle Zhang, and attendees at various conferences and seminars for valuable feedback. A special thank you to Tesary Lin both for her comments on the paper and for inspiring this project in the first place.

1 Introduction

Over the last decade, media platforms that were previously fully ad-funded, such as YouTube and Twitch.tv, have introduced subscriptions that allow consumers to pay a fee to watch content without ads. Similarly, previously fully subscription-based media services such as Netflix have introduced subscription tiers that offer a lower monthly fee in exchange for watching ads. This multi-tiered format acts as a form of price discrimination whereby price-insensitive or ad-averse consumers can access content without having to sit through advertising, while price-sensitive or ad-tolerant consumers can access content without having to pay a high monthly fee. However, the profitability of these advertising-supported tiers depends on how much the platforms can charge for ads, which in turn is a function of how much data the platform can share about the users being targeted [Johnson et al., 2020]. If a platform is subject to privacy regulations that inhibit its ability to sell advertising, the users of the platform will benefit from greater control over their data, but those regulations may then have the unintended effect of discouraging platforms from serving ad-tolerant consumers who are by and large low-income [Johnson et al., 2020]. Therefore, privacy regulation in this context is likely to have regressive welfare consequences.

To investigate this tradeoff, we model a monopoly media platform who serves consumers who vary in their income. The platform offers access to media through either an ad-free premium subscription or a free service with advertising, and its choice variables are the price of the premium subscription and the ad load on the ad-supported option. Similar to Casner and Teh [2025] we assume an exogenous advertising sector that has a constant willingness to pay for consumer impressions. Following Dworczak et al. [2021] and Leisten [2024], we model high-income consumers as having a low marginal value of money, and in accordance with findings in the empirical literature [Varian et al., 2005, Johnson et al., 2020] we assume higher-income consumers have a greater aversion to ads (or, equivalently, a high marginal value of time).

In our model, the platform optimally sets the price and the ad load. High-income consumers buy the premium service and low-income consumers buy the ad-supported service. So long as advertising is valuable but not too valuable, the marginal consumer between premium and ad-supported will indifferent between each of these and the outside option. If a regulator imposes privacy regulation, this reduces advertisers' willingness to pay for impressions, which reduces the profitability of the free version. The platform responds to this by shifting consumers off the ad-supported version and onto the paid subscription by reducing the price of the premium subscription and increasing the ad load. As a result, the low-income consumers are harmed because they have to sit through more advertising for the same content, while high-income consumers benefit because they pay a lower price. We establish that, while privacy increase it always performs worse in terms of the utilitarian, inequality-averse measure of consumer welfare proposed by Leisten [2024] than it does in terms of standard consumer surplus.

We consider three extensions to the model. In the first, we allow for a positive correlation between income and advertiser willingness-to-pay for eyeballs following Gentzkow et al. [2024]. In this extension, we find that our results from the base model carry through so long as privacy affects advertiser willingness to pay for high-income consumers more than for low-income consumers.

In the second extension, we allow consumers to have a taste for privacy. If the consumers' taste for privacy is simply a function of the current state of privacy regulation and independent of their own decisions, then the impact of this preference depends on the correlation between privacy preferences and wealth. If there is no correlation, then the increase in consumer welfare from increased privacy mitigates any welfare loss to low-income consumers from increased ad load. However, if privacy preferences are positively correlated with wealth—as is found in the empirical literature [Johnson et al., 2020, Lin and Strulov-Shlain, 2023]—then total inequality of outcome can *increase* relative to the base model,

exacerbating our results on the policy's distributional impact.

If consumers care about privacy only if their data is being used—i.e., only if they use the ad-supported version—then privacy regulation will increase the relative appeal of the ad-supported product and hence lead to higher ad-load and higher premium subscription prices relative to the base model. A marginal increase in privacy then has two effects:

1. The reduction in advertising price
2. An increase in appeal of the ad-supported subscription tier as regulation limits the extent to which consumer data can be used.

The second effect increases the value to consumers of watching ad-supported content and allows the platform to profitably raise the advertising level. If this effect dominates the reduction in ad price, then the value of ad-viewing consumers increases for the platform and the results of the base model are reversed. On the other hand, if the ad-price reduction dominates then the results of the base model carry through, but consumers who are the least ad-averse may have improved welfare as the increase in privacy utility will outweigh the reduced utility from increased ad-load.

Finally, we consider an extension where the platform offers an intermediate option with a positive price and a positive ad load. We find results analogous to the base model: when the advertising price decreases, the platform will increase the ad load on both the ad-supported and the intermediate option and decrease the price of the intermediate and premium options. Low-income consumers are hurt, and upper-middle-income consumers are also worse off, but interestingly lower-middle-income consumers are better off because they benefit more from the reduction in price of the middle option than they are hurt by the increase in advertising load.

1.1 Related Literature

Pay-TV and ad-supported media: Most articles on the co-existence of paid and ad-supported media business models posit that it is a market segmentation or price discrimination method: consumers with low time-value and high money-value will watch ad-supported content while those for whom the opposite is true prefer pay-tv [Armstrong and Weeds, 2007]. Alternatively, Calvano and Polo [2020] show that even with homogeneous consumers, the two business models can arise as a way for media platforms to differentiate themselves and hence soften competition.

Other papers examine the effects of pay-TV vs. ad-supported business models on the extent to which media providers differentiate themselves from each other through the choice of content provided [Peitz and Valletti [2008], Anderson et al. [2018]]. Anderson and Coate [2005] and Crampes et al. [2009] study the welfare implications of as compared to purely ad-supported vs. pay-tv business models. We take the quality and variety of media provided as given, and while it is the case that sorting consumers to ad-supported versus paid subscriptions is a key feature of the equilibrium in our model, we also take the price discrimination motive as given. Our main interest is exploring how privacy regulation has different impacts on the welfare of high- and low-income consumers through its impact on advertising prices.

Dietl et al. [2023] is much closer to our context, but they explore the competition between pay-tv and fully ad-supported businesses and the implications of the different business models on advertising levels rather than exploring differential impacts of privacy regulation.

Privacy: The classical literature on privacy frames an increase in the ability to maintain private info as a tradeoff between increased equality and reduced efficiency [Posner, 1981]. We contribute to a substantial literature that has developed since then on the unintentional effects of policy mandates and their interaction with privacy preferences. The economics

and marketing literature have demonstrated a positive correlation between age, income, and value for privacy [Varian et al., 2005, Lin and Strulov-Shlain, 2023] implying that firms which implement product designs to maximize data-gathering may unintentionally bias their sample by driving away older and wealthier consumers [Lin and Strulov-Shlain, 2023]. Much like Markovich and Yehezkel [2024] We show that a platform can instead use consumers' distaste for advertising to steer them toward a premium version of the service. However, this strand of the literature [Fainmesser et al., 2023, Markovich and Yehezkel, 2024] focuses on the interaction between privacy and a platform's choice of business model rather than the implications of privacy regulation on welfare given a business model.

When it comes to privacy regulations in particular, regulators must consider the impact of reduced advertising profitability and privacy regulations on the supply of content [Lefrere et al., 2024, Johnson et al., 2024], the ability to inform customers of a product's existence [Aridor et al., 2024], increased concentration in the digital advertising sector [Peukert et al., 2022], and increased search costs or reduced quality through lack of personalized service [Goldfarb and Tucker, 2019].[1] While we are not the first to note that these tradeoffs can fall particularly intensely on disadvantaged consumers, most of the previous literature has focused on the fact that alternative data sources are more robust for less marginalized groups [Tucker, 2023]. Our model instead finds that reduced profitability of serving low-income consumers leads to a more "damaged" product (in the sense of McAfee [2007]) for low-income consumers and lower prices for high-income consumers.

Inequality: In recent years there has been an increased awareness in the economics literature generally, and in the industrial organization and market design literature specifically, that excess inequality can lead to undesirable outcomes. Indeed, Dworczak et al. [2021] and Akbarpour et al. [2024] show that in an economy with significant inequality, policies

[1] We are focusing on the presence of these tradeoffs, but it is worth emphasizing that they must be balanced against the potential negative consequences of too much information sharing and benefits of privacy. See Dubé et al. [2024] for a comprehensive overview of this literature.

such as price controls can be welfare-enhancing in stark contrast to the intuition learned in introductory economics courses. Similarly, Leisten [2024] shows that an increase in inequality (as modeled by a mean preserving spread in consumers' shadow value of money) can lead to an increase in markups and hence reduced welfare for low-income consumers. Other papers consider policy responses to inequality such as potential contributions of antitrust policy toward creating more egalitarian economies [Baker and Salop, 2015], or measuring the impact that inequality has on macroeconomic outcomes [Hendren, 2020]. We contribute to this literature by connecting it to differential preferences for privacy, and by showing how privacy mandates may increase inequality even if the first order impact on total welfare is ambiguous.

Regulation and Product Quality: One of our contributions is to connect the economics of privacy to the literature that examines the tradeoff between regulation and the quality of products produced. Famously, regulation of the airline industry led to higher prices and a focus on luxury in-flight amenities [Borenstein, 1992]. Morrison and Winston [2010] report that deregulation starting in the Carter administration led to a significant drop in airline fares, making airline travel accessible to a much broader swathe of consumers, albeit at the cost of airlines no longer providing lobster dinners for in-flight meals. Similarly, the CAFE fuel efficiency standards introduced in 1975 have increased fuel economy across the American car fleet, but somewhat at the cost of quality [Ferrara, 2007]. We contribute by connecting this literature to advertising load choice and the implications of privacy regulations.

2 Model

A monopoly platform offers two products, ad-supported (A) and premium (P). For consumer i, utility for each product is

$$u_{iA} = V - \eta_i a$$
$$u_{iP} = V - \alpha_i p$$

where V is the baseline value of the platform, α_i is i's price sensitivity, η_i is i's aversion to advertising (thought of as the marginal value of time), a is ad load, and p is the premium price. There is a continuum of consumers with total mass 1, and we assume that wealthier people value their time more and poorer people value their dollar more. That is: consumer i's marginal value of money is $\alpha_i = \alpha(\eta_i)$, where $\alpha(.)$ is a smooth, decreasing function. Though this is not essential to the analysis, we assume that consumers with low marginal value of time and high marginal value of money are poor, while consumers with high marginal value of time and low marginal value of money are wealthy. The distribution of η is given by cdf $F(\eta)$, which is also smooth on support $[\eta_L, \eta_H]$, with $\eta_L > 0$. The pdf is $f(\eta)$, and it is positive everywhere on $[\eta_L, \eta_H]$. The function $\alpha(\eta)$ induces a distribution over α, and we assume $\alpha(\eta_H) > 0$ and $\alpha(\eta_L) < \infty$ so this distribution over α has a positive, bounded support.

The platform earns revenue p for each user who buys P and revenue $p_a a$ for each user who buys A. p_a converts "eyeballs" to dollars and is a function of the quality of information the platform has about its users, which can then be used to better target advertisements. We can think of privacy regulation as inducing a decrease in p_a as the platform is less able to leverage its information, or gather information, about consumers. The platform chooses an ad load, a, and a price for the premium product p, to maximize profits. Profits are given by

$$\pi = Q_A(a, p)p_a a + Q_P(a, p)p$$

where Q_A is the mass of consumers who consume A, and Q_P is the mass of consumers who consume P. We begin with the following lemma, which characterizes possible equilibria:

Lemma 1. *Fixing any $a > 0$ and $p > 0$, let $\bar{\eta}_A$ solve $V - \bar{\eta}_A a = 0$, let $\bar{\eta}_P$ solve $V - \alpha(\bar{\eta}_P)p = 0$, and let $\bar{\eta}$ solve $V - \bar{\eta}a = V - \alpha(\bar{\eta})p$. Then:*

1. *If $\bar{\eta}_A < \bar{\eta}_P$, then there is not full coverage of the market: consumers with $\eta \leq \bar{\eta}_A$ consume A, consumers with $\eta \in (\bar{\eta}_A, \bar{\eta}_P)$ consume 0, and consumers with $\eta \geq \bar{\eta}_P$ consume P.*
2. *If $\bar{\eta}_A \geq \bar{\eta}_P$, then there is full coverage of the market. $\bar{\eta} \in [\bar{\eta}_P, \bar{\eta}_A]$ and consumers with $\eta \leq \bar{\eta}$ consume A, and consumers with $\eta > \bar{\eta}$ consume P.*

All proofs are in the Appendix. The above lemma establishes that there are two possible scenarios: there may be an equilibrium with *full coverage* where everybody either consumes P or A, or an equilibrium without full coverage where some consumers do not engage with either of the platform's offerings. Furthermore, it establishes that the allocation of products to consumers can be characterized by thresholds: either $\bar{\eta}$ if there is full coverage of the market, or $\bar{\eta}_A$ and $\bar{\eta}_P$ if there is not.

Using this lemma, we can write the profit function of the platform. When there is full coverage of the market, platform profits are

$$p_a a F(\bar{\eta}) + p(1 - F(\bar{\eta})).$$

When there is not full coverage of the market, platform profits are

$$p_a a F(\bar{\eta}_A) + p(1 - F(\bar{\eta}_P)).$$

Conditional on there being full coverage of the market, the platform can always do strictly better by increasing a and p, such that $\bar{\eta}$ remains fixed and all consumers still earn more

than 0 utility (thereby ensuring there remains full coverage of the market). This means that we can write $a = \frac{V}{\bar{\eta}}$ and $p = \frac{V}{\alpha(\bar{\eta})}$. When there is not full coverage of the market, $a = \frac{V}{\bar{\eta}_A}$ and $p = \frac{V}{\alpha(\bar{\eta}_P)}$. Profits are then

$$\begin{aligned} V\left(\frac{p_a F(\bar{\eta})}{\bar{\eta}} + \frac{1 - F(\bar{\eta})}{\alpha(\bar{\eta})}\right) & \quad \text{if full coverage} \\ V\left(\frac{p_a F(\bar{\eta}_A)}{\bar{\eta}_A} + \frac{1 - F(\bar{\eta}_P)}{\alpha(\bar{\eta}_P)}\right) & \quad \text{if not.} \end{aligned}$$

We introduce the following assumption throughout the rest of this section:

Assumption-conc. $\frac{F(\eta)}{\eta}$ is concave in η, and $\frac{(1-F(\eta))}{\alpha(\eta)}$ is concave in η.

This assumption guarantees concavity of the profit functions by assuming concavity of its two components (profits from A and P, respectively). This assumption is slightly stronger than log concavity. The next lemma provides the conditions for p_a to neither be so small that the platform only wants to sell P, or so large that the platform only wants to offer A to consumers.

Lemma 2. *There exists $\bar{p}_a$ and $\bar{\bar{p}}_a$ such that, if and only if $p_a > \bar{p}_a$ and $p_a < \bar{\bar{p}}_a$, there is a separating equilibrium. That is, some consumers consume A and some consumers consume P. Specifically $\bar{p}_a = \frac{\eta_L(f(\eta_L)\alpha(\eta_L)+\alpha'(\eta_L))}{f(\eta_L)\alpha(\eta_L)^2}$ and $\bar{\bar{p}}_a = \frac{f(\eta_H)-\eta_H^2}{\alpha(\eta_H)(f(\eta_H)\eta_H-1)}$. It can be shown $\bar{p}_a < \bar{\bar{p}}_a$*

The next assumption ensures full coverage— that is, that no customer chooses the outside good in equilibrium. It is easy to verify that this assumption ensures that either the marginal revenue of A is positive at η_H or that the marginal revenue of P is positive at η_L. Combined with our assumption on concavity, this ensures that the platform either prefers selling A to all customers over selling O to any of them or that the platform prefers selling P to all customers over selling O to any of them.[2]

[2]It is easy to see that, in a non-full coverage equilibrium, marginal changes to p_a have no effect on a and p. This is because the marginal consumer for A is indifferent between A and O. This means that the platform

Assumption-fc. Either or both of the following hold:

1. $\eta_H f(\eta_H) > 1$
2. $\alpha(\eta_L) f(\eta_L) > -\alpha'(\eta_L)$

Under this assumption, the platform's problem reduces to choosing a threshold $\bar{\eta}$, where consumers with $\eta < \bar{\eta}$ prefer A, consumers with $\eta > \bar{\eta}$ prefer P, and consumers with $\eta = \bar{\eta}$ are indifferent between A, P, and O. Therefore, the platform's problem becomes

$$\bar{\eta} = \text{arg max}_{\eta} \; V\left(\frac{p_a F(\eta)}{\eta} + \frac{1 - F(\eta)}{\alpha(\eta)}\right).$$

With this result we are finally to state the main proposition of this section.

Proposition 1. *Suppose there is a marginal increase in privacy regulation. Then, if $p_a \in (\bar{p}_a, \bar{\bar{p}}_a]$, then a increases and p decreases, some consumers switch from A to P, and the poorest consumers are harmed while the wealthiest consumers benefit.*

If $p_a \notin (\bar{p}_a, \bar{\bar{p}}_a]$, then an increase in privacy regulation does not change a, p, or the choices of any consumer.

Figure 1 illustrates some of the main results concerning welfare. As η grows, for any a and p, consumer surplus increases for P and decreases for A. In equilibrium we can plug the platform's choice of a and p into consumer utility to get

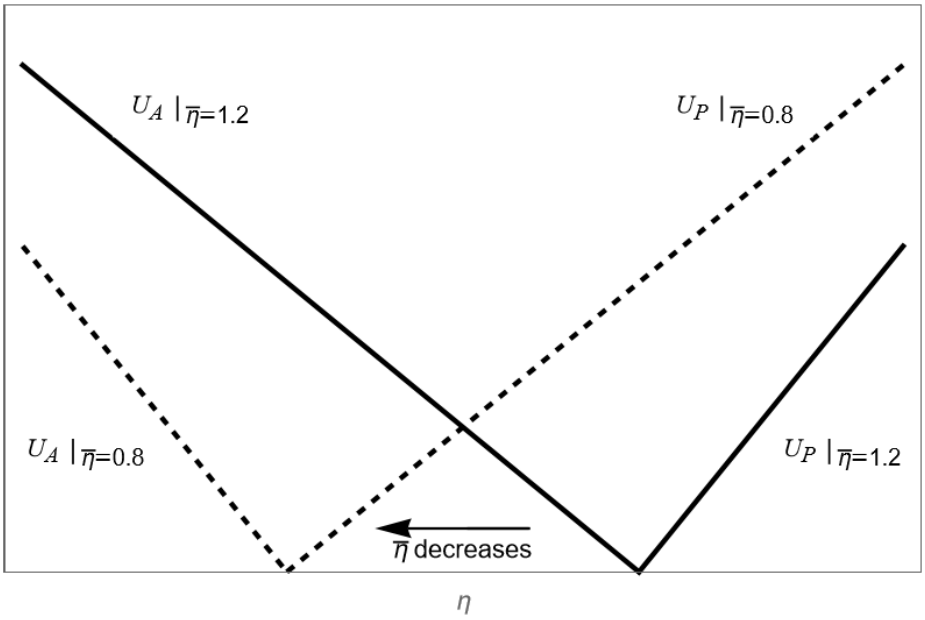
$$U_A(\eta) = V\left(1 - \frac{\eta}{\bar{\eta}}\right)$$
$$U_P(\eta) = V\left(1 - \frac{\alpha(\eta)}{\alpha(\bar{\eta})}\right)$$

Consumers choose between A, P, and O, so therefore overall surplus is represented by the area between the upper envelope of $U_A(\eta)$ and $U_P(\eta)$, and the x-axis. In any full-coverage

has set a to make the marginal revenue of this consumer equal to 0, and changing p_a—which enters profits multiplicatively—would not change this fact.

separating equilibrium, the platform increases a and p until the marginal consumer (the one with $\eta = \bar{\eta}$, who is indifferent between A and P) has zero consumer surplus. If the platform increases a and decreases p, then $U_A(\eta)$ shifts and rotates downward while $U_P(\eta)$ shifts up and becomes more horizontal. For a local change in $\bar{\eta}$ this means that consumers to for whom $\eta_i < \bar{\eta}$ are harmed while those who have $\eta_i > \bar{\eta}$ have their utility increase. Additionally, because the utility curve is rotating, the magnitude of harm is greatest for those consumers with η close to $\bar{\eta}$. While the high-income consumers whose welfare increases the most are also those closest to the cutoff. Intuitively: the consumers who are most affected by a change in a are those with the greatest aversion to advertising, while those who most benefit from a drop in the premium price are those with the most marginal value for money.

Figure 1: Impact on consumer welfare



More formally, we consider consumer surplus as a measure of welfare:

$$CS = \int_{\eta_L}^{\bar{\eta}} \frac{1}{\alpha(\eta)} U_A(\eta)\, f(\eta)\, d\eta + \int_{\bar{\eta}}^{\eta_H} \frac{1}{\alpha(\eta)} U_P(\eta)\, f(\eta)\, d\eta.$$

While consumer surplus is the most common measure of consumer welfare in partial equilibrium applications, Leisten [2024] highlights some of its shortcomings when differences in income lead to different marginal utility of consumption and proposes an alternative measure, inequality-adjusted consumer surplus:

$$\begin{aligned} IACS = & \frac{\int_{\eta_L}^{\bar{\eta}} U_A(\eta)\ f(\eta)\ d(\eta) + \int_{\bar{\eta}}^{\eta_H} U_P(\eta)\ f(\eta)\ d\eta}{\int_{\eta_L}^{\eta_H} \alpha(\eta)\ f(\eta)\ d\eta} \\ = & \int_{\eta_L}^{\bar{\eta}} \frac{1}{\bar{\alpha}} U_A(\eta)\ f(\eta)\ d(\eta) + \int_{\bar{\eta}}^{\eta_H} \frac{1}{\bar{\alpha}} U_P(\eta)\ f(\eta)\ d\eta \end{aligned}$$

where $\bar{\alpha} \equiv \int_{\eta_L}^{\eta_H} \alpha(\eta) f(\eta)\ d\eta$ is teh average price coefficient across all consumers.

The following proposition characterizes welfare changes across these two measures and formallizes the notion that, because the increase in privacy is regressive, inequality adjusted consumer surplus is more adversely affected by privacy than traditional consumer surplus.

Proposition 2. *A marginal increase in privacy always leads to a more positive (or, equivalently, less negative) marginal change in CS than IACS.*

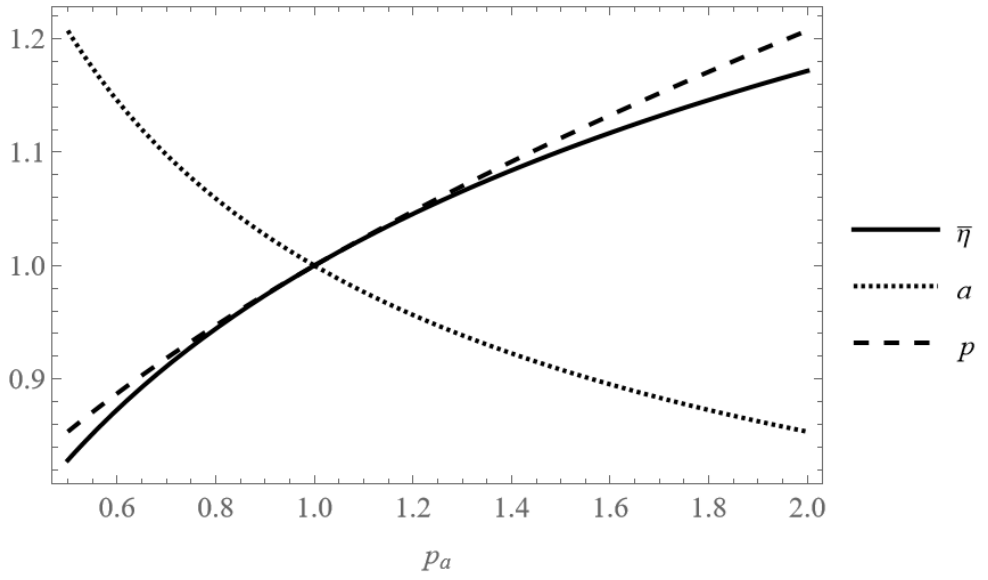
Proposition 2 implies that if both surplus measures increase (decrease), then CS will increase more (decrease less) than IACS, and if the signs are opposite then the change in CS will be positive while that of IACS negative. We demonstrate this result in the following numerical example:

2.1 Numerical example

Here, we let $\eta \sim U[0.5, 1.5]$ and $\alpha(\eta) = 2 - \eta$, which leads α to also be distributed $U[0.5, 1.5]$. Let $V = 1$. Profits are

$$\frac{p_a F(\bar{\eta}_A)}{\bar{\eta_A}} + \frac{1 - F(\bar{\eta}_P)}{2 - \bar{\eta}_P}$$

Figure 2: Numerical example: Ad load decreases and price increases in p_a



if there is not full coverage, and

$$\frac{p_a F(\bar{\eta})}{\bar{\eta}} + \frac{1 - F(\bar{\eta})}{2 - \bar{\eta}}$$

if there is. It is easy to see that $\frac{p_a F(\bar{\eta}_A)}{\bar{\eta}_A}$ is strictly increasing in $\bar{\eta}_A \in [0.5, 1.5]$, so there must be full coverage. Fixing any p_a, $\bar{\eta}(p_a)$ solves

$$\bar{\eta}(p_a) = \arg\max_{\eta} \left(\frac{p_a F(\eta)}{\eta} + \frac{1 - F(\eta)}{2 - \eta} \right).$$

Using $F(\eta) = \eta - 0.5$ for $\eta \in [0.5, 1.5]$ then taking first order conditions and some algebra yields $\bar{\eta}(p_a) = \frac{2p_a - 2\sqrt{p_a}}{p_a - 1}$ whenever this number is between 0.5 and 1.5. We plot $\bar{\eta}(p_a)$ and the implied a and p in Figure 2. Note that, as p_a decreases, $\bar{\eta}$ decreases, a increases and p decreases.

We can compute consumer surplus as

$$CS(p_a) = \int_{0.5}^{\bar{\eta}(p_a)} \frac{1}{\alpha(\eta)}(1 - \eta a(p_a))\, d\eta + \int_{\bar{\eta}(p_a)}^{1.5} \frac{1}{\alpha(\eta)}(1 - \alpha(\eta)p(p_a))\, d\eta$$

and the inequality-adjusted consumer surplus

$$IACS(p_a) = \frac{\int_{0.5}^{\bar{\eta}(p_a)}(1 - \eta a(p_a))\, d\eta + \int_{\bar{\eta}(p_a)}^{1.5}(1 - \alpha(\eta)p(p_a))\, d\eta}{\int_{0.5}^{1.5} \alpha(\eta)\, d\eta}.$$

We define the terms $\Delta CS(\eta)$ and $\Delta IACS(\eta)$ as the change in *individual consumer welfare* at a given value of η for a change in p_a, weighted by the unit adjustment that the respective surplus measures use to be stated in dollar terms. In other words it is the change in the interior of the integral at a given value of the variable being integrated over. Formally, for a change from p_a^1 to p_a^2

$$\Delta CS(\eta) \equiv \begin{cases} \frac{1}{\alpha(\eta)}[(1 - \eta a(p_a^2)) - (1 - \eta a(p_a^1))], & \eta \leq \bar{\eta}(p_a^2) \\ \frac{1}{\alpha(\eta)}[(1 - \alpha(\eta)p(p_a^2)) - (1 - \eta a(p_a^1)], & \bar{\eta}(p_a^2) < \eta \leq \bar{\eta}(p_a^1) \\ \frac{1}{\alpha(\eta)}[(1 - \alpha(\eta)p(p_a^2)) - (1 - \alpha(\eta)p(p_a^1))], & \eta > \bar{\eta}(p_a^1) \end{cases}$$

And $\Delta IACS(\eta)$ is defined analogously. The total change in welfare is found by integrating $\Delta IACS(\eta)$ across the range of η. These terms show the change in welfare resulting from privacy regulation at each point in the income spectrum. We plot these effects of a change in p_a from 1.75 to 0.9 in Figure 3. Consistent with the illustration in Figure 1, poorer consumers are harmed and wealthier consumers benefit, with the largest harms accrued by poorer consumers who are "lower-middle class." Furthermore, the inequality adjustments in IACS place greater weight on the harms imposed on poorer consumers and less on the benefits accrued by wealthier consumers. The total change in consumer surplus, represented by the integral under this curve, is positive without the inequality adjustments

(an increase in privacy is beneficial) but negative with the inequality adjustments.

Figure 3: Numerical example: Heterogeneous impacts of privacy on consumers

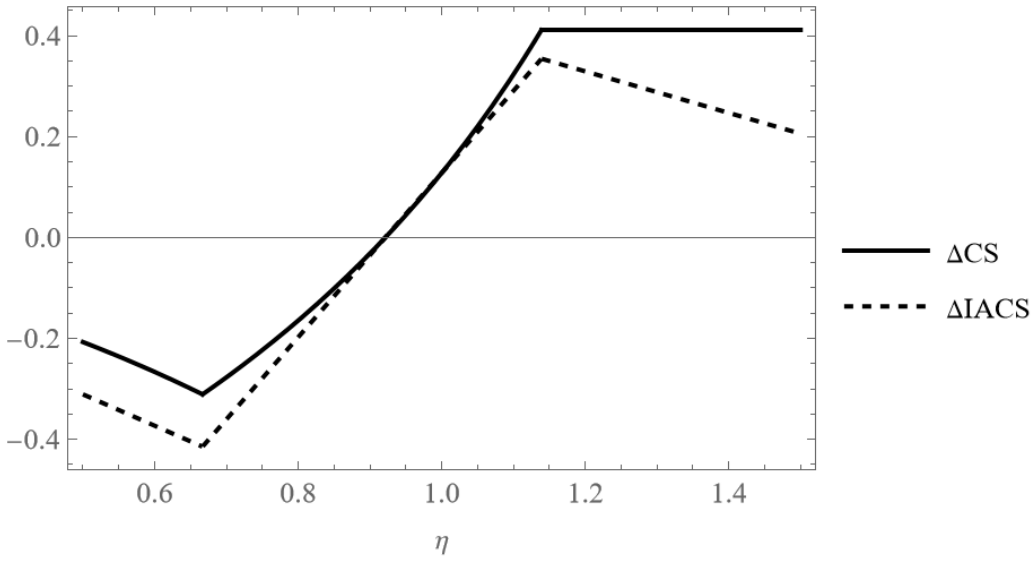
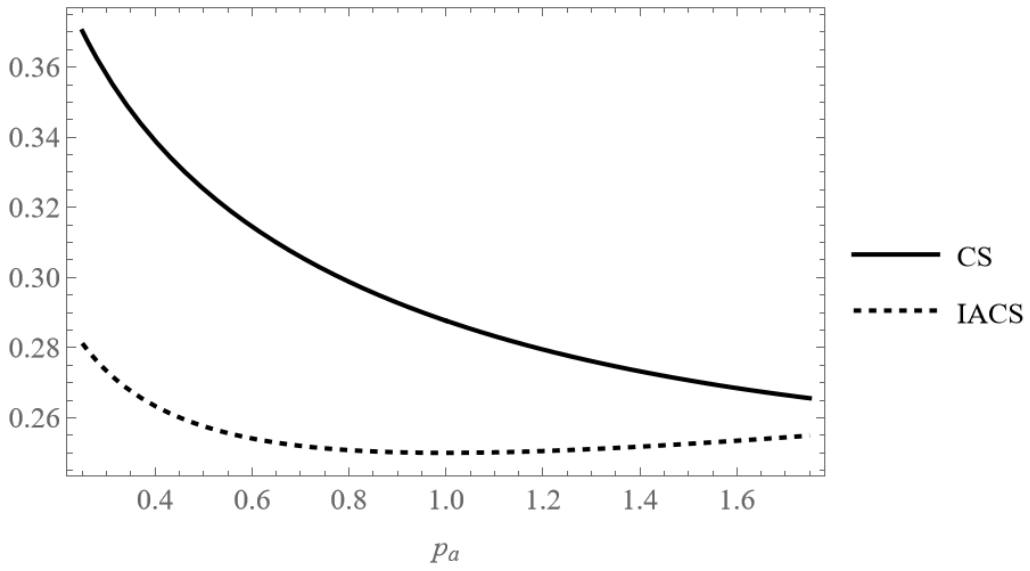


Figure 4 plots aggregate consumer surplus and inequality-adjusted consumer surplus against p_a. While, in this example, consumer surplus unambiguously increases with more privacy (decreases in p_a), inequality-adjusted consumer surplus first declines and then increases with more privacy.

Figure 4: Numerical example: Inequality-adjusted consumer surplus reflects the harm to low-income consumers



3 Extensions

3.1 Heterogeneous value of attention

In our baseline specification, p_a converts attention (i.e., a unit of consumer-advertisement time) to dollars. However, consumers may be heterogeneous in the marginal profitability of their attention [Gentzkow et al., 2024]. Here, we suppose the conversion from ad load to revenues is a function $p_a(\eta;\theta)$, where $\theta \in \Theta$ indexes the stringency of privacy regulations and $\frac{\partial p_a}{\partial \eta} > 0$ so wealthy consumers' attention is more profitable than poor consumers' attention. We also assume $\frac{\partial p_a}{\partial \theta} < 0$ and that $p_a > 0 \; \forall \eta \in [\eta_L, \eta_H]$ and $\forall \theta \in \Theta$.

Whenever there is a full-coverage separating equilibrium, the platform's profits are

$$\max_{\bar{\eta}} \int_{\eta_L}^{\bar{\eta}} \frac{V p_a(\eta) f(\eta)}{\bar{\eta}} \, d\eta + \frac{V(1 - F(\bar{\eta}))}{\alpha(\bar{\eta})} \tag{1}$$

Proposition 3. *Suppose we have a separating, full coverage equilibrium. Then if* $\frac{\partial^2 p_a}{\partial \eta \partial \theta} < 0$*, an increase in privacy regulations increases* a*, decreases* p*, and causes some consumers on the margin to switch from* A *to* P*.*

When θ increases, two things happen to the platform's tradeoff when setting the advertising level:

1. The ad revenue gained from attracting new users to A from P decreases.
2. The relative profitability of showing ads to high- and low-income consumers changes and may drive the platform to serve more high-income consumers or more low-income consumers.

The first factor always pushes $\bar{\eta}$ down and a up. The second factor may point in the same or opposite direction from the first factor. The condition $\frac{\partial^2 p_a}{\partial \eta \partial \theta} < 0$ ensures that high-income consumers become relatively less profitable as ad viewers, which in turn ensures that both the first and second effects point in the same direction. Therefore, a increases, p decreases, and consumers on the margin switch from A to P.[3]

3.2 Taste for privacy

We have thus far focused on the market outcomes resulting from privacy regulation and have ignored the possibility that consumers experience direct benefits from privacy regulation. In this subsection, we address this in a by introducing a direct utility of privacy regulation $z(\theta)$, with $z(0) = 0$ and $z'(\theta) > 0$. We consider two different forms of privacy taste:

1. *Intrinsic taste*: By which we mean a general preference for having less information

[3] Tucker [2023] finds that there are more alternative data sources for wealthier groups, so we might expect that to mitigate the reduction in willingness to pay from loss of ability to target consumers and the condition $\frac{\partial^2 p_a}{\partial \eta \partial \theta} < 0$ would not hold. However, recall that this condition is referring to the absolute value of the cross derivative, so even if the proportional change in advertiser willingness to pay is smaller for high-income consumers it is entirely likely that the absolute change in advertising price is increasing in income.

gathered by external parties. $z(\theta)$ in this model is independent of both consumers' participation decisions and which method they use to access content. In this case consumers get benefit $z(\theta)$ regardless of the choices they make.

2. *Instrumental taste*: This taste for privacy represents a concern for the economic impact of information sharing. For example, a firm with a lot of information about its customers could use personalized pricing to extract a higher proportion of surplus than a naive firm. Here we follow Markovich and Yehezkel [2024] in assuming that consumers in a media context only have this concern for privacy when their information is actively being used by advertisers. Accordingly $z(\theta)$ in the instrumental taste paradigm enters only the utility of viewing the ad-supported content.[4]

Intrinsic tastes for privacy: In an intrinsic privacy taste paradigm, $z(\theta)$ is always a part of consumer utility and it is independent of market participant decisions. With this additional term, utility becomes

$$u_{iA} = V - \eta_i a + z(\theta)$$
$$u_{iP} = V - \alpha_i p + z(\theta)$$

and utility is $z(\theta)$ if consumers choose the outside option. Because $z(\theta)$ enters the consumer utility additively and is independent of consumer decisions (including participation), all market outcomes in the model are unaffected. Therefore we can leap immediately to

[4]This definition of instrumental privacy is narrower than that elsewhere in the literature [Lin, 2022]. One could easily imagine an instrumental concern that is largely independent of one's consumption decisions. For example the more firms which possess one's data, the more likely it is that one's data will be exposed in a security breach, leaving one more vulnerable to identity theft. Similarly, ad-avoidance subscriptions often also come with less data collection (Markovich and Yehezkel [2024] discuss pay-for-privacy programs). Nevertheless we maintain this narrow definition as it provides a useful distinction for the different effects privacy preferences could have on market outcomes, and it seems sensible to think that instumental preferences would have an outsize impact on utility of products where data is more likely to be collected while intrinsic preferences would be more associated with the general state of privacy regulation. To the extent that these forms of privacy preference cross over into having both market impact and utility independent of market outcomes, we could simply model this as $z(\theta)$ being a function of both intrinsic and instrumental utility in either of the paradigms we outline.

considering the implications for consumer surplus. When there is a full coverage, separating equilibrium, consumer surplus is:

$$\begin{aligned} CS &= \int_{\eta_L}^{\bar{\eta}(p_a)} \frac{1}{\alpha(\eta)}(V - \eta a + z(\theta))f(\eta)d\eta + \int_{\bar{\eta}(p_a)}^{\eta^H} \frac{1}{\alpha(\eta)}(V - \alpha(\eta)p + z(\theta))f(\eta)d\eta \\ &= \int_{\eta_L}^{\bar{\eta}(p_a)} \frac{1}{\alpha(\eta)}(V - \eta a)f(\eta)d\eta + \int_{\bar{\eta}(p_a)}^{\eta^H} \frac{1}{\alpha(\eta)}(V - \alpha(\eta)p)f(\eta)d\eta \\ &\quad + z(\theta)\int_{\eta_L}^{\eta^H} \frac{1}{\alpha(\eta)}f(\eta)d\eta \end{aligned}$$

and inequality-adjusted surplus is

$$\begin{aligned} IACS &= \frac{\int_{\eta_L}^{\bar{\eta}(p_a)}(V - \eta a + z(\theta))d\eta + \int_{\bar{\eta}(p_a)}^{\eta^H}(V - \alpha(\eta)p + z(\theta))f(\eta)d\eta}{\int_{\eta_L}^{\eta^H}\alpha(\eta)f(\eta)d\eta} \\ &= \frac{\int_{\eta_L}^{\bar{\eta}(p_a)}(V - \eta a)f(\eta)d\eta + \int_{\bar{\eta}(p_a)}^{\eta^H}(V - \alpha(\eta)p)f(\eta)d\eta}{\int_{\eta_L}^{\eta^H}\alpha(\eta)f(\eta)d\eta} + \frac{z(\theta)}{\int_{\eta_L}^{\eta^H}\alpha(\eta)f(\eta)d\eta} \end{aligned}$$

Note that in both cases $z(\theta)$ enters consumer surplus as a constant multiplied by a weight ($E[\frac{1}{\alpha(\eta)}]$ for CS and $1/E[\alpha(\eta)]$ for IACS). This pushes consumer surplus up and somewhat mitigates the negative impact of the privacy policy on low income consumers, reducing disparity in utility outcomes and increasing the likelihood that $\Delta IACS$ and ΔCS will both be positive after an increase in privacy regulation.

However, it is important to note that this result depends critically on the assumption that intrinsic privacy preferences are independent of socioeconomic status. The literature on privacy preferences does not support this assumption, instead finding that there is a significant positive correlation between taste for privacy and income [Turow et al., 2009, Lin and Strulov-Shlain, 2023]. To account for this, suppose that privacy utility is given by $z(\theta;\eta)$

with $z(0;\eta) = 0\ \forall\eta$, and $\frac{\partial^2 z}{\partial\theta\partial\eta} > 0$. In this case consumer surplus is

$$CS = \int_{\eta_L}^{\bar{\eta}(p_a)} \frac{1}{\alpha(\eta)}(V - \eta a + z(\theta;\eta))f(\eta)d\eta + \int_{\bar{\eta}(p_a)}^{\eta^H} \frac{1}{\alpha(\eta)}(V - \alpha(\eta)p + z(\theta;\eta))f(\eta)d\eta$$

and inequality-adjusted surplus is

$$IACS = \frac{\int_{\eta_L}^{\bar{\eta}(p_a)}(V - \eta a + z(\theta,\eta))d\eta + \int_{\bar{\eta}(p_a)}^{\eta^H}(V - \alpha_i p + z(\theta,\eta))f(\eta)d\eta}{\int_{\eta_L}^{\eta^H}\alpha(\eta)f(\eta)d\eta}$$

It is obvious from inspection that we can no longer pull out privacy preferences as a term independent of the α normalizations in CS, meaning that CS will place lower weight on the privacy gains of low-income consumers and higher weight on high-income consumers relative to IACS. This, combined with the fact that higher income consumers gain more from a privacy increase overall, means that intrinsic privacy preferences *strengthen* the increase in inequality when they are positively correlated with wealth.[5]

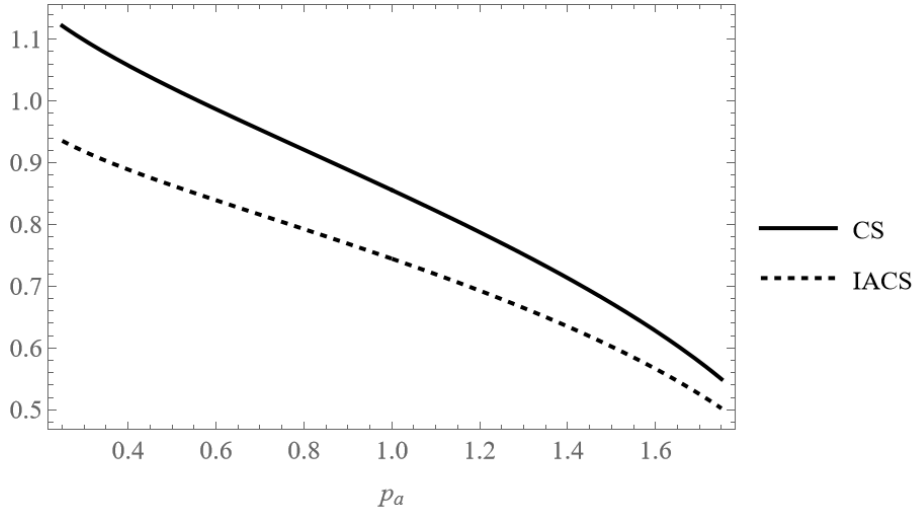
The impact of intrinsic privacy preferences on welfare analysis of a privacy enhancement policy depend on the degree of correlation with income. If preferences are only weakly correlated with income, then they suggest that any second order inequality enhancing effects from the policy change need to be balanced against an inequality reducing first order privacy gain. However, if preference for privacy is positively correlated with income, then intrinsic privacy preferences and these second order effects both push in the same direction when it comes to inequality of outcome.

To illustrate, we add to the numerical example in Figure 4 an intrinsic preference for privacy

[5] We can spell this out more explicitly using the Leibniz Rule. Specifically, focusing on $z(p_a,\eta)$'s contribution to CS and IACS, the change from a privacy enhancing policy action is $\int_{\eta_L}^{\bar{\eta}(p_a)} \frac{1}{\alpha(\eta)}(D_1 z(\theta,\eta))f(\eta)d\eta + \int_{\bar{\eta}(p_a)}^{\eta^H} \frac{1}{\alpha(\eta)}(D_1 z(\theta,\eta))f(\eta)d\eta$, while for IACS the change is $\frac{\int_{\eta_L}^{\bar{\eta}(p_a)}(D_1 z(\theta,\eta))f(\eta)d\eta + \int_{\bar{\eta}(p_a)}^{\eta^H}(D_1 z(\theta,\eta))f(\eta)d\eta}{\int_{\eta_L}^{\eta^H}\alpha(\eta)f(\eta)d\eta}$. $D_1 z(\theta,\eta)$ is positive and increasing in η by assumption, meaning that the underweighting of low-income consumers in CS will give that measure a more positive evaluation of the policy action when privacy preferences are taken into account than will IACS.

equal to $z(\theta, \eta) = \frac{1}{2}\sqrt{\theta\eta}$ and let $p_a = 2 - \theta$ for $\theta \in [0, 2]$. We plot out resulting consumer surplus and inequality-adjusted consumer surplus in Figure 5. The addition of intirnsic preferences for privacy mean that increases in privacy increase both consumer surplus and inequality-adjusted consumer surplus, but because privacy preferences are strongest for the wealthiest consumers, the gap between CS and IACS is especially pronounced when p_a is low.

Figure 5: Income correlated intrinsic preferences for privacy lead to an increased gap between CS and IACS as privacy increases



Instrumental tastes for privacy: Following Markovich and Yehezkel [2024], in this section we model privacy as a concern about one's data being *used*, rather than a general concern about it being collected in the first place. Consumers therefore only benefit from privacy regulation if they consume A. Utility in this paradigm is the

$$u_{iA} = V - \eta_i a + z(\theta)$$
$$u_{iP} = V - \alpha_i p$$

and 0 for the outside option. $z(\theta)$ acts as a demand shifter for A, making it relatively more

appealing for all consumers by the same amount when privacy regulation is positive.

An immediate consequence is that for any fixed a and p, $\bar{\eta}_A$ and $\bar{\eta}$ are higher in this extension than in the base model, since they now solve $V+\bar{\eta}_A a+z(\theta)=0$ and $V-\bar{\eta}a+z(\theta)=V-\alpha(\bar{\eta})p$, respectively. $\bar{\eta}_P$ is unchanged as it solves $V-\alpha(\bar{\eta})p=0$, so the privacy utility does not come into play. This then implies that the platform will set $a=\frac{V+z(\theta)}{\bar{\eta}_A}$ when the market is not fully covered and $a=\frac{V+z(\theta)}{\bar{\eta}}$ when it is. The expressions determining p remain $\frac{V}{\alpha(\bar{\eta}_P)}$ and $\frac{V}{\alpha(\bar{\eta})}$. Profits are then

$$\Pi=\begin{cases}(V+z(\theta))\left(\frac{p_a F(\bar{\eta})}{\bar{\eta}}\right)+V\left(\frac{1-F(\bar{\eta})}{\alpha(\bar{\eta})}\right) & \text{if full coverage}\\ (V+z(\theta))\left(\frac{p_a F(\bar{\eta}_A)}{\bar{\eta}_A}\right)+V\left(\frac{1-F(\bar{\eta}_P)}{\alpha(\bar{\eta}_P)}\right) & \text{if not.}\end{cases}$$

Consider the following assumption

assumption-sePriv. $p_a \in \left(\frac{V}{(V+z(\theta))}\frac{\eta_L(f(\eta_L)\alpha(\eta_L)+\alpha'(\eta_L))}{f(\eta_L)\alpha(\eta_L)^2}, \frac{V}{(V+z(\theta))}\frac{f(\eta_H)-\eta_H^2}{\alpha(\eta_H)(f(\eta_H)\eta_H-1)}\right)$.

We can then derive the following analogue to Lemma 2

Lemma 3. *Under Assumption-conc and Assumption-sePriv, there is a separating equilibrium. That is, some consumers consume A and some consumers consume P.*

The logic of Lemma 3 is essentially the same as for Lemma 2, but the increased appeal of A stemming from privacy preferences means that the condition on p_a which ensures the platform will always want to have some users consuming the free version has loosened, and that ensuring the platform wants some users buying the premium subscription has tightened.

Assumption-fc continues to be sufficient to ensure a full coverage equilibrium, so we are ready to prove the following analogue to Proposition 1:

Proposition 4. *Suppose a regulator increases θ*

- *If $\frac{d}{d\theta}\big(p_a(V+z(\theta))\big)>0$ then a marginal increase in privacy causes $\bar{\eta}$ and p to*

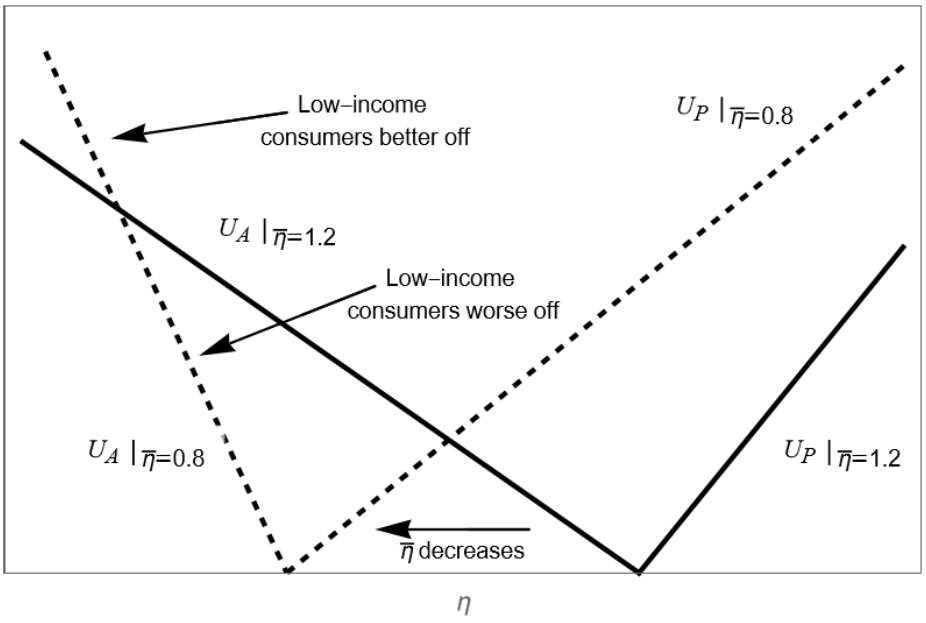
increase. The change in a is ambiguous but the poorest consumers' welfare improves while the richest consumers are worse off.

- *If $\frac{d}{d\theta}\left(p_a(V + z(\theta))\right) < 0$ then a marginal increase in privacy causes $\bar{\eta}$ and p to decrease while a increases.*
  - *Consumers for whom $\eta \in \left[\frac{z'(\theta)\bar{\eta}}{z'(\theta) - \frac{\frac{d\bar{\eta}}{d\theta}}{\bar{\eta}}(V+z(\theta))}, \bar{\eta}\right]$ are worse off*
  - *Consumers for whom $\eta \notin \left[\frac{z'(\theta)\bar{\eta}}{z'(\theta) - \frac{\frac{d\bar{\eta}}{d\theta}}{\bar{\eta}}(V+z(\theta))}, \bar{\eta}\right]$ are better off.*

To see the intuition behind Proposition 4, first note that $\frac{d}{d\theta}\left(p_a(V + z(\theta))\right)$ represents the change in value of ad-viewing consumers to the platform as θ increases. The advertising price decreases as privacy increases as discussed above, but consumers' valuation for A also increases, and the platform is able to capture some of that increase in value. If that increase in value dominates the change in advertising price, then the first part of Proposition 4 is essentially Proposition 1, but with the signs of changes reversed and with roughly analogous intuition. The change in ad load is ambiguous because the logic of Proposition 1 says it should decrease, but also the increase in valuation by consumers means the platform can increase the ad load as θ increases without driving away consumers.

For the second part of Proposition 4, the logic in terms of equilibrium market outcomes is identical to that of Proposition 1, but the effect on low-income consumers' welfare is ambiguous because there is a direct increase in utility from increased privacy in addition to the indirect effects of the platform increasing the ad load. We illustrate this graphically using Figure 6. As in the base model, a decrease in $\bar{\eta}$ is rotating U_A downward while rotating U_P upward, but the increase in privacy utility leads to a countervailing upward shift in the vertical intercept of U_A. This upward shift has the greatest impact on consumers with the least aversion to advertising, so there is a region of particularly low-income consumers whose overall welfare improves with the increase in privacy regulation.

Figure 6: With instrumental privacy preferences some low-income consumers' welfare is improving in θ



For an intuitive explanation: suppose for the sake of illustration that $\eta_L = 0$, then trivially there must be some neighborhood around η_L such that $|z'(\theta)| > \eta \left|\frac{da}{d\theta}\right|$. Consumers who are not particularly bothered by advertising are not significantly impacted by a change in the advertising level. For these consumers the increase in privacy outweighs any loss in utility from increased ad-load. On the other hand, let $\bar{\eta}'$ be the new value of $\bar{\eta}$ after the privacy increase. Consumers for whom $\eta = \bar{\eta}'$ had positive utility from consuming A before the privacy increase, but because they have relatively high value for time they are more affected by an increase in advertising and hence now have 0 utility due to the increased ad load, so these consumers are worse off. The lower bound of the interval in the proposition is the cutoff η at which the impact of changing at load and benefits of increased privacy are equal, while for $\eta > \bar{\eta}$ consumers purchase P and neither ad load nor $z(\theta)$ are relevant to their utility.

3.3 Richer Menus

In the base model the platform is limited to offering a purely ad-supported service and a pure subscription service. One might wonder what would happen were the platform to offer a more complex menu including intermediate products with lower, but positive, ad loads and cheaper subscription prices. We explore this question by introducing a third product M to the platform's menu. This product has both a subscription fee and ads and offers utility

$$U_M = V - \eta a_M - \alpha(\eta) p_M$$

The platform still offers A and P with ad load a_A and price p_P respectively. The assumptions from the base model are sufficient to ensure full coverage of the market—Assumption-fc guarantees that both A and P are more profitable than providing an outside option to any consumer—and that the platform will not sell just A or just P to all consumers. For this section we introduce the additional assumption that $\alpha''(\eta) > 0$, which together with the other assumptions guarantees that it is profit maximizing for the platform to provide 0 utility to consumers who are indifferent between products.[6]

Let $\bar{\eta}_{AM}$ denote the η such that $U_M(\bar{\eta}_{AM}) = U_A(\bar{\eta}_{AM})$, and define $\bar{\eta}_{PM}$ analogously for the indifference point between P and M. Like in the base model, the platform will maximize profits by setting the utility at these indifference points equal to 0 (otherwise the platform could increase ad-load and/or prices while not losing consumers). Its problem then reduces to setting $\bar{\eta}_{AM}$ and $\bar{\eta}_{PM}$. Note: Assumption-se does not guarantee that it is profit maximizing for the platform to offer all 3 products. It is possible to have $\bar{\eta}_{AM} = \eta_L$, $\bar{\eta}_{PM} = \eta_H$, or $\bar{\eta}_{PM} \leq \bar{\eta}_{AM}$ in which case the platform is effectively only offering P and M, A and M, or

[6]This restriction eliminates uninteresting equilibria: If $\alpha(\cdot)$ were linear (e.g. $1 - \alpha\eta$) then the platform could set $a_M = \alpha p_M$, in which case utility for the middle option cancels to $U_M = V - p_M$ and the platform can extract all surplus by setting $p_M = V$. The restriction also aids tractability: if $\alpha''(\eta) < 0$, then U_M would be convex in η and so the platform could only offer M to a positive mass of consumers in a separating equilibrium with all three products if it offered positive utility to all consumers of at least one of A or P, which creates an intractable model.

A and P respectively. Nonetheless in numerical simulation we find the platform offering all products in most of the parameter spaces we explore, so for the remainder of this section we present results assuming $\eta_L < \bar{\eta}_{AM} < \bar{\eta}_{PM} < \eta_H$.

We can use the 0 utility at indifference condition to find ad load and prices as a function of the indifference points

Lemma 4. *When it is profitable for the platform to offer all three products, the profit maximizing price and ad load for M are given by:*

$$a_M = V\frac{\alpha(\bar{\eta}_{AM}) - \alpha(\bar{\eta}_{PM})}{\bar{\eta}_{PM}\alpha(\bar{\eta}_{AM}) - \bar{\eta}_{AM}\alpha(\bar{\eta}_{PM})} \tag{2}$$

$$p_M = V\frac{\bar{\eta}_{PM} - \bar{\eta}_{AM}}{\bar{\eta}_{PM}\alpha(\bar{\eta}_{AM}) - \bar{\eta}_{AM}\alpha(\bar{\eta}_{PM})} \tag{3}$$

Plugging these into platform profit we get

$$\begin{aligned}\Pi = V\Bigg[&p_a\frac{F(\bar{\eta}_{AM})}{\bar{\eta}_{AM}} \\ &+ (F(\bar{\eta}_{PM}) - F(\bar{\eta}_{AM}))\left[p_a\frac{\alpha(\bar{\eta}_{AM}) - \alpha(\bar{\eta}_{PM})}{\bar{\eta}_{PM}\alpha(\bar{\eta}_{AM}) - \bar{\eta}_{AM}\alpha(\bar{\eta}_{PM})} + \frac{\bar{\eta}_{PM} - \bar{\eta}_{AM}}{\bar{\eta}_{PM}\alpha(\bar{\eta}_{AM}) - \bar{\eta}_{AM}\alpha(\bar{\eta}_{PM})}\right] \\ &+ \frac{1 - F(\bar{\eta}_{PM})}{\alpha(\bar{\eta}_{PM})}\Bigg]\end{aligned}$$

To guarantee concavity and sufficiency of the first order conditions as the platform sets these cutoffs we introduce

assumption-concMenu. $(F(\bar{\eta}_{PM}) - F(\bar{\eta}_{AM}))\left[\frac{\bar{\eta}_{PM} - \bar{\eta}_{AM}}{\bar{\eta}_{PM}\alpha(\bar{\eta}_{AM}) - \bar{\eta}_{AM}\alpha(\bar{\eta}_{PM})}\right]$ and $(F(\bar{\eta}_{PM}) - F(\bar{\eta}_{AM}))\left[p_a\frac{\alpha(\bar{\eta}_{AM}) - \alpha(\bar{\eta}_{PM})}{\bar{\eta}_{PM}\alpha(\bar{\eta}_{AM}) - \bar{\eta}_{AM}\alpha(\bar{\eta}_{PM})}\right]$ are both concave in $\bar{\eta}_{AM}$ and $\bar{\eta}_{AM}$.

With this assumption we can state the main proposition of this extension

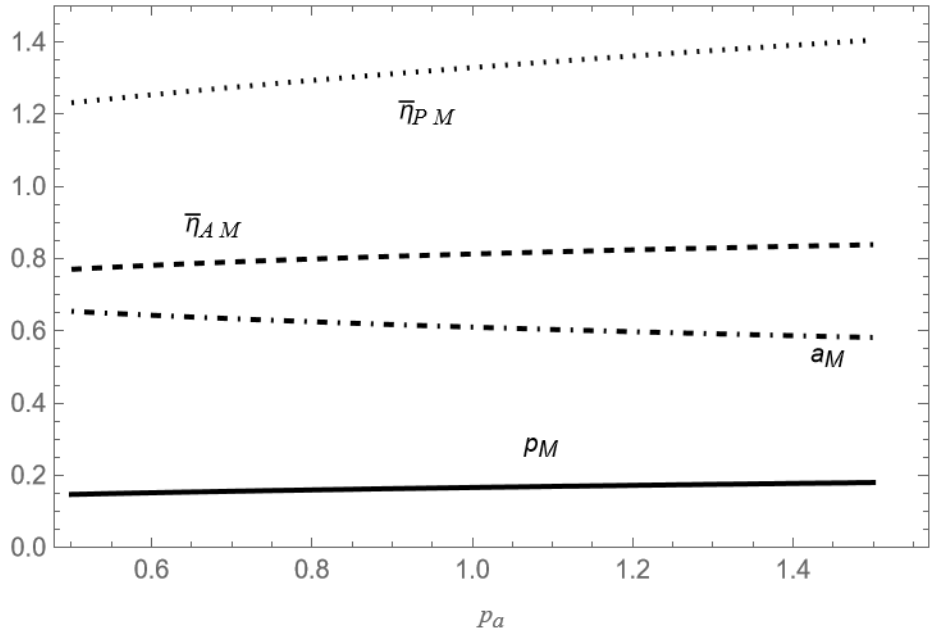
Proposition 5. *A marginal increase in privacy (reduction in p_a) reduces both $\bar{\eta}_{AM}$ and $\bar{\eta}_{PM}$. Consumers in $[\eta_L, \bar{\eta}_{AM})$ are worse off, consumers in $(\bar{\eta}_{PM}, \eta_H]$ are better off, and*

there exist cutoffs $\{\eta_{M1}, \eta_{M2}\} \in [\bar{\eta}_{AM}, \bar{\eta}_{PM}]$ *such that consumers in* $[\bar{\eta}_{AM}, \eta_{M1})$ *are better off and consumers in* $(\eta_{M2}, \bar{\eta}_{PM}]$ *are worse off.*

The platform decreases provision of A and increases provision of P for the same reasons as in the base model: ad impressions are less valuable relative to consumer payments, so the platform is willing to sacrifice the former to get more of the latter. It does so by increasing advertising and reducing fees for both M and P. Consumers close to $\bar{\eta}_{AM}$ are relatively insensitive to ads, so they benefit more from the reduction in price than they are hurt by the increase in advertising, while the reverse is true for consumers close to $\bar{\eta}_{PM}$.

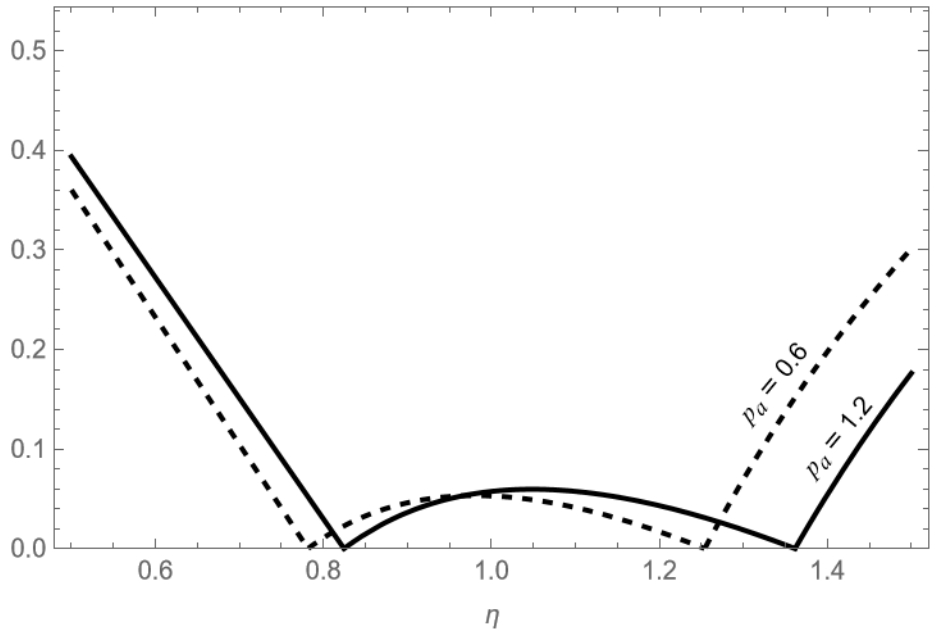
We extend the numerical example from previous sections to the menu, with the small change that we assume $\alpha(\eta) = \frac{2}{\eta^2}$. Figure 7 shows how the two cutoffs and the price and ad-load for the intermediate menu option vary with p_a. As outlined in proposition 5, a

Figure 7: Both indifference points and p_M decrease as p_a decreases, while a_M increases.



decrease in ad price increases the advertising load and decreases the other three variables. We can see the impact this has on utility: From Figure 8, we can see that a decrease in the advertising price causes consumer utility curves (as a fucntion of η to shift to the left,

Figure 8: A decrease in ad price shifts all parts of the utility curve to the left



which—broadly speaking—benefits consumers whose utility is increasing in η but hurts those whose utility is decreasing. Considering this as a welfare change as broken down by η instead of looking at utility directly, we can see from Figure 9 that the positive impact on high-income consumers is much greater than the positive impact on M consumers near $\bar{\eta}_{AM}$.

Thus, similar to the base model, if the platform offers a menu of products which vary in price and ad-load, the reduced targeting from privacy regulation will cause an across the board increase in ad-load and reduction in price, which benefits high-income consumers much more than middle-income consumers, and harms low-income consumers. The overall impact on consumer surplus may be positive, but this depends in large part on the high weight that standard CS calculations put on consumers with low marginal value of a dollar.

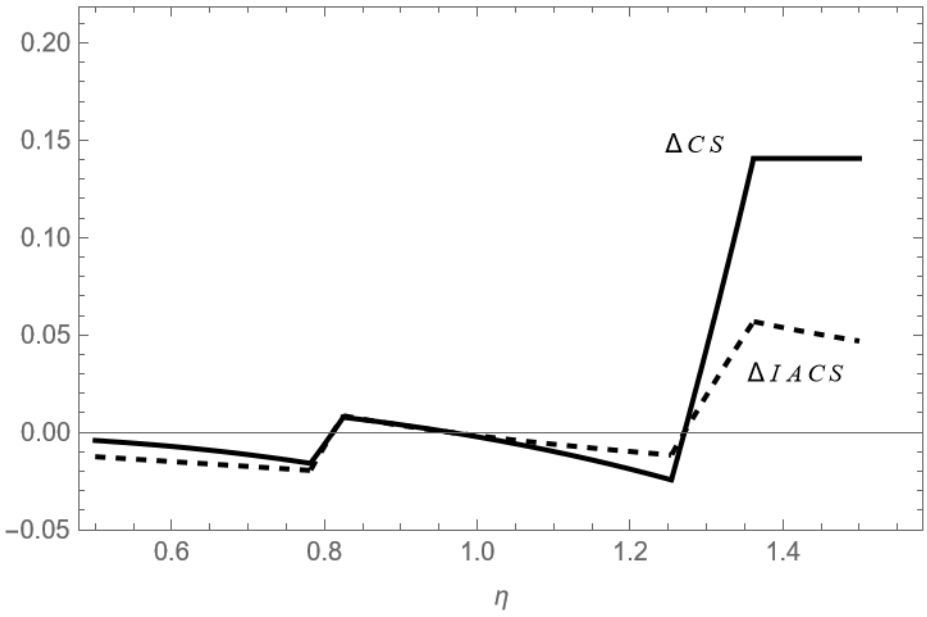


Figure 9: The increase in welfare for high income consumers is much larger than the increase for middle-income consumers.

4 Discussion

4.1 Policy Implications

What takeaways should policy-makers have from these results besides "be careful of unintended distributional consequences of privacy regulation?" If a policymaker has a preference for privacy regulation that accounts for distributional concerns, then Proposition 4 implies that they would likely want to focus their policy on assuaging consumer concerns while allowing for targeting. Part 1 of Proposition 4 implies that a policy with this intent that is designed such that the impact on consumer welfare (e.g. eliminating targeting of vulnerable consumers by malicious advertisers) is significant, while still allowing for ad targeting will likely *reduce* inequality.[7]

[7]Whether policies should be concerned with mitigating inequality is a more complex topic than we have room to go over in this article. We take the concern as given and leave a more nuanced discussion of normative distributional concerns to other works.

For example, Adjerid et al. [2016b] find that consumers value privacy more when they feel the party potentially gaining their information is less trustworthy; therefore a policy that allows consumers to blacklist certain firms or categories of firms from getting their data would allow firms which have good reputations to access data while those which are less consumer friendly would be shut out.[8] Alternatively, a data storage exchange could enforce a code of conduct and require minimum behavior standards for access, which would significantly alleviate consumer privacy concerns as shown in a health context by Adjerid et al. [2016a].

Lin [2022] shows that one of the main reasons many people are concerned about privacy is that they worry about the potential harms from data breaches. Therefore, a policy requiring strict data security and safeguards on advertiser behavior while still allowing for tools that permit advertisers to target their ads to specific demographics would—according to our results—be more beneficial for low-income consumers than forbidding the gathering of the data in the first place.

4.2 Modeling Decisions

Perfect negative correlation between η and $\alpha(\eta)$:

One of the implications of our assumption that $\alpha'(\eta) < 0$ is that the interesting results of our model occur only when there is full coverage of the market. If the market is not fully covered then there are no consumers whose marginal decision is between the advertising-supported and premium products offered by the platform. In this case, the marginal revenue of this consumer is zero regardless of p_a, and many of the results of the model go away.

While we feel assuming a strong negative correlation is justified by the empirical literature as discussed earlier in this article, in the real world many consumers do not watch YouTube

[8] Such choice screens would impose their own costs on consumers, but Farronato et al. [2025] show that this problem is surmountable.

or Twitch.tv, so this full coverage assumption is not necessarily fully realistic. However, we do not believe that it makes a significant difference for our results. For example, suppose consumers have a taste for the platform ϵ which is independently distributed from η. In this case, there would be some consumers—those with high ϵ and intermediate η—who would be indifferent between A and P. An increase in p_a would make the marginal revenue of these consumers greater for A than for P and would therefore induce some incentive for the platform to shift these consumers to consuming A by lowering a and/or increasing p. By this reasoning, the main intuition of Proposition 1 would still go through.

Correlation between instrumental privacy tastes and η:

We do not extend our analysis of instrumental privacy tastes to allow for a correlation between $z(\theta)$ and η, largely because we believe doing so would not add any interesting results. As we have established in the discussion of intrinsic tastes, the most realistic assumption would be for the correlation to be positive. However, if $z(\theta)$ were increasing in wealth, this would push the results closer to the base model. An increase in θ would increase the appeal of A for high-income consumers more than for low-income, meaning that the platform would have more incentive to reduce p as privacy regulation increases. Meanwhile, the relatively small benefits for low-income consumers mean that the increasing utility scenario in part 1 of Proposition 4 is less likely to be relevant for consumers below $\bar{\eta}$ and we would expect the second part—which is roughly equivalent to Proposition 1—would be the more relevant of the two scenarios outlined in Proposition 4.

5 Conclusion

We have created a model of media provision that accounts for the differing impact of privacy regulation on consumers across the income spectrum. We model a monopoly platform that offers two options for viewing content: ad-supported or a premium subscription. The platform's revenue comes from a payment per ad impression which is set by an exoge-

nous advertising sector, and the premium subscription price paid by consumers. It sets the advertising level for the ad-supported option and the access price for the premium subscription. Consumers vary in their income; following results from the empirical literature, higher-income consumers are more averse to advertising and less price sensitive. In equilibrium, relatively low-income consumers will choose the advertising supported option while higher-income consumers choose the premium subscription, and the platform will set the price and ad load such that the marginal consumer is just indifferent between both the two options and not viewing content at all.

If a regulator strengthens privacy regulations, the platform earns less per advertising impression. The platform responds by reducing the price of the premium subscription and increasing the ad load to drive more consumers to pay for content directly. This improves the welfare of high-income consumers who are infra-marginal consumers of the premium option but lowers the welfare of low-income who are inframarginal ad-viewing consumers. Overall, in our baseline model the privacy regulation has adverse distributional consequences.

We consider three extensions. In the first we allow for a positive correlation between advertiser willingness to pay for ad impressions and income. All of the results of the base model go through so long as the marginal impact of privacy regulation on advertising price increases with income.[9] In the second extension, we explicitly model the benefits of privacy regulation. If privacy is an intrinsic benefit that consumers feel regardless of their actions, then the impact of this extension depends on the correlation between taste for privacy and income. If privacy benefits are negatively (positively) correlated with income, the distributional effects are mitigated (exacerbated) relative to the baseline outcome.

Next, we allow for privacy to be an instrumental consideration where consumers only care about their data being used by advertisers and are less concerned about whether it is col-

[9]This is a fairly weak condition, as it can be satisfied even if the *percentage* change in advertising price is decreasing in income.

lected. In this case, privacy regulation increases the relative appeal of the ad-supported option, and the results depend on the relative impact of the privacy benefit and the reduced ad price. If the privacy benefit dominates, then the platform can capture some (but not all) of this benefit by increasing the ad-load, and the overall size of the privacy benefit is sufficient to make ad-viewing consumers more valuable to the platform. If this is the case, then the platform will *increase* the subscription fee, and the welfare effects of the base model are reversed. If the effect of ad prices dominates, then the results of the base model go through mostly unchanged, but there may be a set of low-income consumers who have such a low value of time that they benefit from the privacy increase more than they are hurt by the increase in ad-load. In this case, middle-income consumers are hurt, while the poorest and wealthiest benefit.

Finally, we allow for a richer menu of product offerings by the platform by introducing an intermediate option with positive ad load and price. When privacy regulation reduces the advertising price, the ad load goes up for both the ad-supported and intermediate options, while the subscription price goes down for both the intermediate and premium options. Low-income consumers and upper-middle-income consumers are harmed, while lower-middle-income and high-income consumers benefit.

Our results suggest that regulators who care about both privacy and distributional consequences should carefully consider the equilibrium effects of privacy regulation, as it may hurt low-income consumers. We show that one way to mitigate the negative distributional impact of privacy regulation is to focus on safeguarding consumer welfare rather than banning data collection altogether. If the policy can be designed such that the benefit to consumers outweighs the reduction in advertiser willingness to pay, then the adverse distributional consequences of the policy within our model are reversed.

References

Idris Adjerid, Alessandro Acquisti, Rahul Telang, Rema Padman, and Julia Adler-Milstein. The impact of privacy regulation and technology incentives: The case of health information exchanges. *Management Science*, 62(4):1042–1063, 2016a.

Idris Adjerid, Sonam Samat, and Alessandro Acquisti. A query-theory perspective of privacy decision making. *The Journal of Legal Studies*, 45(S2):S97–S121, 2016b.

Mohammad Akbarpour, Piotr Dworczak, and Scott Duke Kominers. Redistributive allocation mechanisms. *Journal of Political Economy*, 132(6):1831–1875, 2024.

Simon P. Anderson and Stephen Coate. Market provision of broadcasting: A welfare analysis. *The Review of Economic Studies*, 72(4):947–972, 2005. ISSN 00346527, 1467937X. URL http://www.jstor.org/stable/3700696.

Simon P Anderson, Øystein Foros, and Hans Jarle Kind. Competition for advertisers and for viewers in media markets. *The Economic Journal*, 128(608):34–54, 2018.

Guy Aridor, Yeon-Koo Che, Brett Hollenbeck, Maximilian Kaiser, and Daniel McCarthy. Evaluating the impact of privacy regulation on e-commerce firms: Evidence from apple's app tracking transparency. *Marketing Science Institute Working Paper*, pages 24–124, 2024.

Mark Armstrong and Helen Weeds. Programme quality in subscription and advertising-funded television. *Unpublished working paper*, 2007.

Jonathan B Baker and Steven C Salop. Antitrust, competition policy, and inequality. *Geo. LJ Online*, 104:1, 2015.

Severin Borenstein. The evolution of us airline competition. *Journal of Economic perspectives*, 6(2):45–73, 1992.

Emilio Calvano and Michele Polo. Strategic differentiation by business models: Free-to-air and pay-tv. *Economic Journal*, 130(625):50–64, 2020.

Ben Casner and Tat-How Teh. Content-hosting platforms: discovery, membership, or both? *RAND Journal of Economics*, 2025.

Claude Crampes, Carole Haritchabalet, and Bruno Jullien. Advertising, competition and entry in media industries. *The Journal of Industrial Economics*, 57(1):7–31, 2009. ISSN 1467-6451. doi: 10.1111/j.1467-6451.2009.00368.x. URL http://dx.doi.org/10.1111/j.1467-6451.2009.00368.x.

Helmut Dietl, Markus Lang, and Panlang Lin. The effects of introducing advertising in pay tv: a model of asymmetric competition between pay tv and free tv. *The BE Journal of Theoretical Economics*, 23(1):291–326, 2023.

Jean-Pierre Dubé, Dirk Bergemann, Mert Demirer, Avi Goldfarb, Garrett Johnson, Anja Lambrecht, Tesary Lin, Anna Tuchman, Catherine E Tucker, and John G Lynch. The intended and unintended consequences of privacy regulation for consumer marketing: A marketing science institute report. *Available at SSRN 4847653*, 2024.

Piotr Dworczak, Scott Duke Kominers, and Mohammad Akbarpour. Redistribution through markets. *Econometrica*, 89(4):1665–1698, 2021.

Itay P Fainmesser, Andrea Galeotti, and Ruslan Momot. Digital privacy. *Management Science*, 69(6):3157–3173, 2023.

Chiara Farronato, Andrey Fradkin, and Tesary Lin. Designing consent: Choice architecture and consumer welfare in data sharing. *Working Paper*, 2025.

Ida Ferrara. Automobile quality choice under pollution control regulation. *Environmental and Resource Economics*, 38:353–372, 2007.

Matthew Gentzkow, Jesse M Shapiro, Frank Yang, and Ali Yurukoglu. Pricing power in advertising markets: Theory and evidence. *American Economic Review*, 114(2):500–533, 2024.

Avi Goldfarb and Catherine Tucker. Digital economics. *Journal of economic literature*, 57 (1):3–43, 2019.

Nathaniel Hendren. Measuring economic efficiency using inverse-optimum weights. *Journal of public Economics*, 187:104198, 2020.

Garrett Johnson, Tesary Lin, James C Cooper, and Liang Zhong. Coppacalypse? the youtube settlement's impact on kids content. *SSRN Working Paper No. 4430334*, 2024.

Garrett A Johnson, Scott K Shriver, and Shaoyin Du. Consumer privacy choice in online advertising: Who opts out and at what cost to industry? *Marketing Science*, 39(1):33–51, 2020.

Vincent Lefrere, Logan Warberg, Cristobal Cheyre, Veronica Marotta, and Alessandro Acquisti. Does privacy regulation harm content providers? a longitudinal analysis of the impact of the gdpr. *A Longitudinal Analysis of the Impact of the GDPR (April 1, 2024)*, 2024.

Matthew Leisten. Economic inequality and market power. *Available at SSRN 4634316*, 2024.

Tesary Lin. Valuing intrinsic and instrumental preferences for privacy. *Marketing Science*, 41(4):663–681, 2022.

Tesary Lin and Avner Strulov-Shlain. Choice architecture, privacy valuations, and selection

bias in consumer data. *University of Chicago, Becker Friedman Institute for Economics Working Paper*, (2023-58), 2023.

Sarit Markovich and Yaron Yehezkel. Competing for cookies: Platforms' business models in data markets with network effects. *Available at SSRN 4770577*, 2024.

R Preston McAfee. Pricing damaged goods. *Economics*, 1(1):20070001, 2007.

Steven Morrison and Clifford Winston. *The evolution of the airline industry*. Brookings Institution Press, 2010.

Martin Peitz and Tommaso Valletti. Content and advertising in the media: Pay-tv versus free-to-air. *International Journal of industrial organization*, 26(4):949–965, 2008.

Christian Peukert, Stefan Bechtold, Michail Batikas, and Tobias Kretschmer. Regulatory spillovers and data governance: Evidence from the gdpr. *Marketing Science*, 41(4): 746–768, 2022.

Richard A Posner. The economics of privacy. *The American economic review*, 71(2): 405–409, 1981.

Catherine Tucker. Algorithmic exclusion: The fragility of algorithms to sparse and missing data. *Center on Regulation and Markets at Brookings, working paper (February), https://www. brookings. edu/wp-content/uploads/2023/02/Algorithmicexclusion-FINAL. pdf*, 2023.

Joseph Turow, Jennifer King, Chris Jay Hoofnagle, Amy Bleakley, and Michael Hennessy. Americans reject tailored advertising and three activities that enable it. *Available at SSRN 1478214*, 2009.

Hal Varian, Fredrik Wallenberg, and Glenn Woroch. The demographics of the do-not-call list [security of data]. *IEEE Security & Privacy*, 3(1):34–39, 2005.

A Proofs

Proof of Lemma 1

Proof. Suppose $\bar{\eta}_A < \bar{\eta}_P$. Note that, because $V - \eta a$ is decreasing in η, if $\eta_i < \bar{\eta}_A$, then i prefers A to 0. Furthermore, note that because $V - \alpha(\eta)p$ is increasing in η, if $\eta_i < \bar{\eta}_A$ then $\eta_i < \bar{\eta}_P$, and i prefers 0 to P. Therefore, i consumes A. Symmetric arguments ensure that if $\eta_i > \bar{\eta}_p$, then i prefers P to 0 and i prefers 0 to A, so i consumes P. Finally, if $\eta_i \in [\bar{\eta}_A, \bar{\eta}_P]$, then the monotonicity of $V - \eta a$ and $V - \alpha(\eta)p$ ensures i prefers 0 to A and prefers 0 to P, so i consumes 0. This establishes the partial coverage equilibrium described in the statement of the lemma.

Next suppose $\bar{\eta}_A \geq \bar{\eta}_P$. Then for $\eta_i \in [\bar{\eta}_P, \bar{\eta}_A]$, by our monotonicity arguments, i prefers both A and P to 0. Because the functions $V - \eta a$ and $V - \alpha(\eta)p$ are continuous, by the intermediate value theorem $\bar{\eta} \in [\bar{\eta}_P, \bar{\eta}_A]$, and therefore if $\eta_i < \bar{\eta}$, i prefers A to P, and therefore consumes A and if $\eta_i > \bar{\eta}$, i consumes P. □ □

Proof of Lemma 2

Proof. The marginal revenue from increasing the share of customers who consume A by increasing $\bar{\eta}_A$ is

$$\frac{\partial}{\partial \bar{\eta}_A}\left[\frac{Vp_aF(\bar{\eta}_A)}{\bar{\eta}_A}\right] = \frac{Vp_a(\bar{\eta}_A f(\bar{\eta}_A) - F(\bar{\eta}_A))}{\bar{\eta}_A^2}. \tag{4}$$

The marginal revenue from increasing the share of consumers who buy P by decreasing $\bar{\eta}_P$ is

$$-\frac{\partial}{\partial \bar{\eta}_P}\left[\frac{V(1 - F(\bar{\eta}_P))}{\alpha(\bar{\eta})}\right] = \frac{V(\alpha(\bar{\eta}_P)f(\bar{\eta}_P) + (1 - F(\bar{\eta}_P))\alpha'(\bar{\eta}_P)}{\alpha(\bar{\eta}_P)^2}. \tag{5}$$

First, the marginal revenue of A is positive at $\bar{\eta}_A = \eta_L$. To see this, note that $F(\eta_L) = 0$. Plugging in to equation 4 obtains $\frac{Vp_a\eta_L f(\eta_L)}{\eta_L^2}$, and $\eta_L > 0$ so this is positive. Similarly, marginal revenue of P is positive at $\bar{\eta}_P = \eta_H$. Note $1 - F(\eta_H) = 0$, so marginal revenue in equation 5 is $\frac{V\alpha(\eta_H)f(\eta_H)}{\alpha(\eta_H)^2}$. $\alpha(\eta_H) > 0$,

The next step is to show that the marginal revenue of A is greater than that of P at η_L and that the marginal revenue of A is less than that of P at η_H if and only if $p_a \in (\bar{p}_a, \bar{\bar{p}}_a)$. Plugging in η_L to our expressions for marginal revenue, the inequality for η_L can be stated as

$$\frac{Vp_a\eta_L f(\eta_L)}{\eta_L^2} > \frac{V\alpha(\eta_L)f(\eta_L) + \alpha'(\eta_L)}{\alpha(\eta_L)^2}.$$

Some algebra on this inequality yields the expression $p_a > \frac{\eta_L(f(\eta_L\alpha(\eta_L) + \alpha'(\eta_L))}{f(\eta_L)\alpha(\eta_L)^2} = \bar{p}_a$. Similarly, the requisite inequality for η_H can be stated

$$\frac{Vp_a(\eta_H f(\eta_H) - 1)}{\eta_H^2} < \frac{V\alpha(\eta_H)f(\eta_H)}{\alpha(\eta_H)^2}.$$

Simple algebra can restate this inequality as $p_a < \frac{f(\eta_H)\eta_H^2}{\alpha(\eta_H)(f(\eta_H)\eta_H - 1)} = \bar{\bar{p}}_a$.

The final step of the proof is to combine these inequalities to show that there is a separating equilibrium. Because the marginal revenue of A is greater than that of both P and O at η_L the platform will optimally have some consumers in a neighborhood of η_L consume A. A symmetric argument, the platform will optimally have some consumers in a neighborhood of η_H consume P. If p_a is not in the requisite range, then either the marginal revenue A will be higher than for P at η_H, and only A will be sold, or the marginal revenue for P will

be weakly higher than A at η_L, and only P will be sold. □ □

Proof of Proposition 1

Proof. Suppose $p_a \in (\bar{p}_a, \bar{\bar{p}}_a)$. Then there is a full-coverage separating equilibrium. This means that there is some $\bar{\eta}$ where consumers are indifferent between A and P, and all consumers with $\eta < \bar{\eta}$ will consume A and all consumers with $\eta \geq \bar{\eta}$ will consume P. It is a dominant strategy for the platform to choose a and p so that consumers at $\bar{\eta}$ obtain zero utility. Therefore, $a = \frac{V}{\bar{\eta}}$ and $p = \frac{V}{\alpha(\bar{\eta})}$. Then $\bar{\eta}$ solves

$$\bar{\eta} = \arg\max_{\eta} \; V\left(\frac{p_a F(\eta)}{\eta} + \frac{1 - F(\eta)}{\alpha(\eta)}\right).$$

The first order condition is

$$p_a \frac{\bar{\eta} f(\bar{\eta}) - F(\bar{\eta})}{\bar{\eta}^2} = \frac{\alpha(\bar{\eta}) f(\bar{\eta}) + (1 - F(\bar{\eta}))\alpha'(\bar{\eta})}{\alpha(\bar{\eta})^2} \tag{6}$$

By Assumption-conc, the left hand side of the above equation is decreasing in $\bar{\eta}$ and the right hand side is increasing in $\bar{\eta}$. Because there is full coverage, both sides of the above equation are positive, i.e., there is positive marginal revenue for each product (otherwise, it would not be profitable to serve the consumers located at $\bar{\eta}$ and there would not be full coverage). Therefore, decreasing p_a decreases the left hand side of this equation, and therefore the right hand side must also decrease. This means $\bar{\eta}$ must decrease. Since there is a separating equilibrium by Assumption-se, $\bar{\eta} \in [\eta_L, \eta_H]$, and this decrease in $\bar{\eta}$ must lead to an increase in a, a decrease in p, and some consumers near $\bar{\eta}$ will switch from A to P.

Accordingly, some marginal consumers (those with $\eta_i = \bar{\eta}$) switch from A to P, and $a = \frac{V}{\bar{\eta}}$ increases while $p = \frac{V}{\alpha(\bar{\eta})}$ must decrease. It follows that inframarginal consumers of A are made worse off while inframarginal consumers of P are made better off.

If $p_a \leq \bar{p}_a$, then in equilibrium all consumers consume P. The platform sets p and a such that the consumer with η_L will be indifferent between A, P and O. A marginal reduction in p_a still ensures that $p_a \leq \bar{p}$ and we have an equilibrium in which all consumers consume P and the consumer at η_L is still indifferent between A, P, and O. Because consumer tastes have not changed, prices do not change.

If $p_a > \bar{\bar{p}}_a$, we have an equilibrium where all consumers consume A, and the platform will optimally make the consumer at η_H indifferent between A, P, and O. A marginal decrease in p_a does not change that, and we still have an equilibrium in which all consumers consume

A and the consumer at η_H is indifferent between A, P, and O.[10] □ □

Proof of Proposition 2

Proof. Consider a marginal increase in privacy, i.e. a marginal decrease in p_A. Under our assumptions, the $\bar{\eta}$ that solves equation 2 is increasing and continuous in p_A, so a marginal increase in privacy corresponds to a marginal decrease in $\bar{\eta}$. Therefore, it suffices to show that $\frac{\partial IACS}{\partial \bar{\eta}} > \frac{\partial CS}{\partial \bar{\eta}}$.

The expressions for these derivatives are

$$\frac{\partial IACS}{\partial \bar{\eta}} = \frac{1}{\bar{\alpha}} \left(\int_{\eta_L}^{\bar{\eta}} \frac{\partial U_A(\eta)}{\partial \bar{\eta}} \, f(\eta) \, d\eta + \int_{\bar{\eta}}^{\eta_H} \frac{\partial U_P(\eta)}{\partial \bar{\eta}} \, f(\eta) \, d\eta \right) \tag{7}$$

and

$$\frac{\partial CS}{\partial \bar{\eta}} = \int_{\eta_L}^{\bar{\eta}} \frac{1}{\alpha(\eta)} \frac{\partial U_A(\eta)}{\partial \bar{\eta}} \, f(\eta) \, d\eta + \int_{\bar{\eta}}^{\eta_H} \frac{1}{\alpha(\eta)} \frac{\partial U_P(\eta)}{\partial \bar{\eta}} \, f(\eta) \, d\eta.$$

Noting that $\frac{\partial U_A(\eta)}{\partial \bar{\eta}}$ is always positive and $\frac{\partial U_P(\eta)}{\partial \bar{\eta}}$ is always negative, the mean value theorem for definite integrals states that there exists $\eta_1 \in (\eta_L, \bar{\eta})$ and $\eta_2 \in (\bar{\eta}, \eta_H)$ such that

$$\frac{\partial CS}{\partial \bar{\eta}} = \frac{1}{\alpha(\eta_1)} \int_{\eta_L}^{\bar{\eta}} \frac{\partial U_A(\eta)}{\partial \bar{\eta}} \, f(\eta) \, d\eta + \frac{1}{\alpha(\eta_2)} \int_{\bar{\eta}}^{\eta_H} \frac{\partial U_P(\eta)}{\partial \bar{\eta}} \, f(\eta) \, d\eta.$$

Rearranging equation 7 and plugging in to the above, we obtain

$$\frac{\partial CS}{\partial \bar{\eta}} = \frac{1}{\alpha(\eta_1)} \int_{\eta_L}^{\bar{\eta}} \frac{\partial U_A(\eta)}{\partial \bar{\eta}} \, f(\eta) \, d\eta + \frac{1}{\alpha(\eta_2)} \left(\bar{\alpha} \frac{\partial IACS}{\partial \bar{\eta}} - \int_{\eta_L}^{\bar{\eta}} \frac{\partial U_A(\eta)}{\partial \bar{\eta}} f(\eta) \, d\eta \right).$$

Some algebra yields

$$\frac{\partial CS}{\partial \bar{\eta}} = \underbrace{\frac{1}{\bar{\alpha}} \left(\frac{\alpha(\eta_2)}{\alpha(\eta_1)} - 1 \right) \int_{\eta_L}^{\bar{\eta}} \frac{\partial U_A(\eta)}{\partial \bar{\eta}} f(\eta) \, d\eta}_{(A)} + \frac{\partial IACS}{\partial \bar{\eta}}$$

Inspecting term (A) in the above equation: $\frac{1}{\bar{\alpha}}$ is positive; $\int_{\eta_L}^{\bar{\eta}} \frac{\partial U_A(\eta)}{\partial \bar{\eta}} f(\eta) \, d\eta$ is positive

[10] If $p_a = \bar{p}_a$, then a marginal decrease in p_a induces a switch to a separating equilibrium and the above reasoning applies.

because $\frac{\partial U_A(\eta)}{\partial \bar{\eta}}$ is always positive, and $\left(\frac{\alpha(\eta_2)}{\alpha(\eta_1)} - 1\right)$ is negative because $\eta_1 < \eta_2$ and $\alpha(.)$ is a decreasing function. Therefore, term (A) is negative and $\frac{\partial CS}{\partial \bar{\eta}} < \frac{\partial IACS}{\partial \bar{\eta}}$. □ □

Proof of Proposition 3

Proof. Taking the first order condition of equation 1 and applying the Leibniz rule, we obtain

$$\frac{p_a(\bar{\eta})f(\bar{\eta})}{\bar{\eta}} - \int_{\eta_L}^{\bar{\eta}} \frac{p_a(\eta)f(\eta)}{\bar{\eta}^2}\, d\eta = \frac{\alpha(\bar{\eta})f(\bar{\eta}) + \alpha'(\bar{\eta})(1 - F(\bar{\eta}))}{\alpha(\bar{\eta})^2}. \tag{8}$$

Differentiating the left hand side of this equation with respect to θ, we obtain

$$\frac{\partial p_a(\bar{\eta})}{\partial \theta}\frac{f(\bar{\eta})}{\bar{\eta}} - \int_{\eta_L}^{\bar{\eta}} \frac{\partial p_a(\eta)}{\partial \theta}\frac{f(\eta)}{\bar{\eta}^2}\, d\eta.$$

which can be expressed as

$$\frac{1}{\bar{\eta}^2}\left(\frac{\partial p_a(\bar{\eta})}{\partial \theta} f(\bar{\eta})\bar{\eta} - F(\bar{\eta})\mathbb{E}\left[\frac{\partial p_a(\eta)}{\partial \theta}\middle| \eta \leq \bar{\eta}\right]\right).$$

This expression is negative if and only if

$$\frac{\partial p_a(\bar{\eta})}{\partial \theta} f(\bar{\eta})\bar{\eta} < F(\bar{\eta})\mathbb{E}\left[\frac{\partial p_a(\eta)}{\partial \theta}\middle| \eta \leq \bar{\eta}\right] \tag{9}$$

By our assumption that $\frac{\partial p_a}{\partial \theta} < 0$ and $\frac{\partial^2 p_a}{\partial \eta \partial \theta} < 0$, we have $\frac{\partial p_a(\bar{\eta})}{\partial \bar{\eta}} < \mathbb{E}\left[\frac{\partial p_a(\eta)}{\partial \theta}\middle| \eta \leq \bar{\eta}\right] < 0$. This means

$$F(\bar{\eta})\mathbb{E}\left[\frac{\partial p_a(\eta)}{\partial \theta}\middle| \eta \leq \bar{\eta}\right] > F(\bar{\eta})\frac{\partial p_a(\bar{\eta})}{\partial \theta}$$

Finally, our assumption that marginal revenues are positive ensures $f(\bar{\eta})\eta > F(\bar{\eta})$, so $F(\bar{\eta})\frac{\partial p_a(\bar{\eta})}{\partial \theta} > \bar{\eta} f(\bar{\eta})\frac{\partial p_a(\bar{\eta})}{\partial \theta}$. Stringing together inequalities, this means $\bar{\eta} f(\bar{\eta})\frac{\partial p_a(\bar{\eta})}{\partial \theta} < [F(\bar{\eta})\mathbb{E}\left[\frac{\partial p_a(\eta)}{\partial \theta}\middle| \eta \leq \bar{\eta}\right]$, which is the condition in inequality 9. Therefore, the left hand side of Equation 8 decreases when θ increases, and so must the right hand side. By our assumption on concavity, this means $\bar{\eta}$ must decrease. It immediately follows that a increases, p decreases, and some consumers on the margin switch from A to P. □ □

Proof of Lemma 3

Proof. Following the logic of the proof of Lemma 2, the marginal profit from increasing $\bar{\eta}_A$ is

$$\frac{\partial}{\partial \bar{\eta}_A}\left[\frac{(V+z(\theta))p_a F(\bar{\eta}_A)}{\bar{\eta}_A}\right] = \frac{(V+z(\theta))p_a(\bar{\eta}_A f(\bar{\eta}_A) - F(\bar{\eta}_A))}{\bar{\eta}_A^2}. \tag{10}$$

And the marginal profit of decreasing $\bar{\eta}_P$ is unchanged

$$-\frac{\partial}{\partial \bar{\eta}_P}\left[\frac{V(1-F(\bar{\eta}_P))}{\alpha(\bar{\eta})}\right] = V\frac{(\alpha(\bar{\eta}_P)f(\bar{\eta}_P) + (1-F(\bar{\eta}_P))\alpha'(\bar{\eta}_P)}{\alpha(\bar{\eta}_P)^2}. \tag{11}$$

By the exact same logic as Lemma 2, the former is positive at $\bar{\eta}_A$ and the latter at $\bar{\eta}_P$, so it remains to find a range of p_a such that marginal revenue of A is greater than P at η_L and that of P is greater at η_H.

The inequality for η_L can be stated as

$$\frac{(V+z(\theta))p_a\eta_L f(\eta_L)}{\eta_L^2} > V\frac{\alpha(\eta_L)f(\eta_L) + \alpha'(\eta_L)}{\alpha(\eta_L)^2}.$$

Solving for p_a yields the expression $p_a > \frac{V}{(V+z(\theta))}\frac{\eta_L(f(\eta_L\alpha(\eta_L)+\alpha'(\eta_L))}{f(\eta_L)\alpha(\eta_L)^2}$. Similarly, the inequality for η_H can be stated

$$\frac{(V+z(\theta))p_a(\eta_H f(\eta_H) - 1)}{\eta_H^2} < \frac{V\alpha(\eta_H)f(\eta_H)}{\alpha(\eta_H)^2}.$$

Which yields $p_a < \frac{V}{(V+z(\theta))}\frac{f(\eta_H)\eta_H^2}{\alpha(\eta_H)(f(\eta_H)\eta_H - 1)}$.

The rest of the proof follows the same steps as the proof of Lemma 2. □

Proof of Proposition 4

Proof. By the same logic as for the original proof of the proposition $a = \frac{V+z(\theta)}{\bar{\eta}}$ and $p = \frac{V}{\alpha(\bar{\eta})}$, so $\bar{\eta}$ solves:

$$\bar{\eta} = \arg\max_{\eta} \left((V+z(\theta))\frac{p_a F(\eta)}{\eta}\right) + \left(V\frac{1-F(\eta)}{\alpha(\eta)}\right).$$

We find the first order condition:

$$p_a(V + z(\theta))\frac{\bar{\eta}f(\bar{\eta}) - F(\bar{\eta})}{\bar{\eta}^2} = V\frac{\alpha(\bar{\eta})f(\bar{\eta}) + (1 - F(\bar{\eta}))\alpha'(\bar{\eta})}{\alpha(\bar{\eta})^2} \tag{12}$$

Assumption-conc is still sufficient to ensure sufficiency of this first order condition for profit maximization, so the effect of privacy regulation will depend on $\frac{d}{d\theta}\big(p_a(V + z(\theta))\big)$. By assumption $\frac{dp_a}{d\theta} < 0$, and $z'(\theta) > 0$ implies $\frac{d(V+z(\theta))}{d\theta} > 0$, so the sign of $\frac{d}{d\theta}\big(p_a(V + z(\theta))\big)$ is ambiguous. If it is positive, then the left side of Equation 12 is increasing in θ, which from the same logic as in the proof of Proposition 1 means $\bar{\eta}$ increases which implies that $p = \frac{V}{\alpha(\bar{\eta})}$ increases, reducing the welfare of inframarginal high-income consumers.

Recall $a = \frac{V+z(\theta)}{\bar{\eta}}$. The top and bottom of this fraction are both increasing in θ in this case, so the change in a is ambiguous. However we can substitute this identity for a into consumer utility to re-write low-income consumers' utility as

$$V - \eta\frac{V + z(\theta)}{\bar{\eta}} + z(\theta) = (V + z(\theta))(1 - \frac{\eta}{\bar{\eta}}) \tag{13}$$

Since $z(\theta)$ and $\bar{\eta}$ are increasing in θ, the inframarginal consumers' welfare is increasing in θ.

If $\frac{d}{d\theta}\big(p_a(V + z(\theta))\big) < 0$ then $\bar{\eta}$ decreases, and p decreases. Because $z(\theta)$ and $\bar{\eta}$ decrease, $a = \frac{V+z(\theta)}{\bar{\eta}}$ must increase. Inframarginal consumers of the premium subscription are better off by analogous logic to their being worse off above, but the effect on inframarginal consumers of the ad-supported product is ambiguous. The first term in the right hand side of Equation 13 is increasing, but because $\bar{\eta}$ is decreasing, so is the second term. We can resolve this ambiguity by taking the derivative of the right hand side of Equation 13 with regard to θ, and finding conditions such that it is negative:

$$z'(\theta)(1 - \frac{\eta}{\bar{\eta}}) + (V + z(\theta))\frac{\frac{d\bar{\eta}}{d\theta}\eta}{\bar{\eta}^2} < 0$$

Solving for η

$$\eta > \frac{z'(\theta)\bar{\eta}}{z'(\theta) - \frac{\frac{d\bar{\eta}}{d\theta}}{\bar{\eta}}(V + z(\theta))}$$

We have already shown that $\frac{d}{d\theta}\big(p_a(V + z(\theta))\big) < 0$ implies $\frac{d\bar{\eta}}{d\theta} < 0$, so we conclude that $0 < \frac{z'(\theta)}{z'(\theta) - \frac{\frac{d\bar{\eta}}{d\theta}}{\bar{\eta}}(V+z(\theta))} < 1$, and the interval from the second part of the proposition is non-empty. For consumers with $\eta \in \left[\frac{z'(\theta)\bar{\eta}}{z'(\theta) - \frac{\frac{d\bar{\eta}}{d\theta}}{\bar{\eta}}(V+z(\theta))}, \bar{\eta}\right]$, utility is decreasing as θ increases. If η is below the lower end of this interval, then we have shown that the increase in privacy

dominates the change in ad load and utility increases for those consumers, while if $\eta > \bar{\eta}$ consumers will choose P and their welfare is increasing. □

Proof of Lemma 4

Proof. From the 0 utility at indifference condition

$$V - \bar{\eta}_{AM} a_A = 0 \Rightarrow a_A = \frac{V}{\bar{\eta}_{AM}} \tag{14}$$

Substitute this into the indifference between A and M

$$V - \bar{\eta}_{AM} \frac{V}{\bar{\eta}_{AM}} = 0 = V - \bar{\eta}_{AM} a_M - \alpha(\bar{\eta}_{AM}) p_M$$
$$\Rightarrow a_M = \frac{V - \alpha(\bar{\eta}_{AM}) p_M}{\bar{\eta}_{AM}}$$

By similar logic

$$p_P = \frac{V}{\alpha(\bar{\eta}_{PM})} \tag{15}$$

which gives

$$\bar{\eta}_{PM} a_M = V - \alpha(\bar{\eta}_{PM}) p_M$$

Solving for price, then substituting into a_M gives

$$a_M = V \frac{\alpha(\bar{\eta}_{AM}) - \alpha(\bar{\eta}_{PM})}{\bar{\eta}_{PM}\alpha(\bar{\eta}_{AM}) - \bar{\eta}_{AM}\alpha(\bar{\eta}_{PM})} \tag{16}$$

$$p_M = V \frac{\bar{\eta}_{PM} - \bar{\eta}_{AM}}{\bar{\eta}_{PM}\alpha(\bar{\eta}_{AM}) - \bar{\eta}_{AM}\alpha(\bar{\eta}_{PM})} \tag{17}$$

□

Proof of Proposition 5

Proof. The first order conditions are $[\bar{\eta}_{AM}]$:

$$\begin{aligned}
p_a\Bigg(&\frac{\bar{\eta}_{AM}f(\bar{\eta}_{AM})-F(\bar{\eta}_{AM})}{\bar{\eta}_{AM}}-f(\bar{\eta}_{AM})\frac{\alpha(\bar{\eta}_{AM})-\alpha(\bar{\eta}_{PM})}{\bar{\eta}_{PM}\alpha(\bar{\eta}_{AM})-\bar{\eta}_{AM}\alpha(\bar{\eta}_{PM})}\\
&(F(\bar{\eta}_{PM})-F(\bar{\eta}_{AM}))\Bigg[\frac{\alpha'(\bar{\eta}_{AM})}{\bar{\eta}_{PM}\alpha(\bar{\eta}_{AM})-\bar{\eta}_{AM}\alpha(\bar{\eta}_{PM})}\\
&\qquad-\frac{[\alpha(\bar{\eta}_{AM})-\alpha(\bar{\eta}_{PM})](\bar{\eta}_{PM}\alpha'(\bar{\eta}_{AM})-\alpha(\bar{\eta}_{PM}))}{(\bar{\eta}_{PM}\alpha(\bar{\eta}_{AM})-\bar{\eta}_{AM}\alpha(\bar{\eta}_{PM}))^2}\Bigg]\Bigg)\\
=&f\left(\bar{\eta}_{AM}\right)\frac{\bar{\eta}_{PM}-\bar{\eta}_{AM}}{\bar{\eta}_{PM}\alpha\left(\bar{\eta}_{AM}\right)-\bar{\eta}_{AM}\alpha\left(\bar{\eta}_{PM}\right)}\\
&+\left(F\left(\bar{\eta}_{PM}\right)-F\left(\bar{\eta}_{AM}\right)\right)\left[\frac{1}{\bar{\eta}_{PM}\alpha\left(\bar{\eta}_{AM}\right)-\bar{\eta}_{AM}\alpha\left(\bar{\eta}_{PM}\right)}\right.\\
&\qquad\left.+\frac{\left(\bar{\eta}_{PM}-\bar{\eta}_{AM}\right)\left(\bar{\eta}_{PM}\alpha'\left(\bar{\eta}_{AM}\right)-\alpha\left(\bar{\eta}_{PM}\right)\right)}{\left(\bar{\eta}_{PM}\alpha\left(\bar{\eta}_{AM}\right)-\bar{\eta}_{AM}\alpha\left(\bar{\eta}_{PM}\right)\right)^2}\right]
\end{aligned}$$

and $[\bar{\eta}_{PM}]$:

$$\begin{aligned}
p_a\Bigg[&f\left(\bar{\eta}_{PM}\right)\frac{\left[\alpha\left(\bar{\eta}_{AM}\right)-\alpha\left(\bar{\eta}_{PM}\right)\right]}{\bar{\eta}_{PM}\alpha\left(\bar{\eta}_{AM}\right)-\bar{\eta}_{AM}\alpha\left(\bar{\eta}_{PM}\right)}\\
&+\left(F\left(\bar{\eta}_{PM}\right)-F\left(\bar{\eta}_{AM}\right)\right)\left(\frac{-\alpha'\left(\bar{\eta}_{PM}\right)}{\bar{\eta}_{PM}\alpha\left(\bar{\eta}_{AM}\right)-\bar{\eta}_{AM}\alpha\left(\bar{\eta}_{PM}\right)}\right.\\
&\qquad\left.+\frac{\left(\alpha\left(\bar{\eta}_{AM}\right)-\alpha\left(\bar{\eta}_{PM}\right)\right)\left[\bar{\eta}_{AM}\alpha'\left(\bar{\eta}_{PM}\right)-\alpha\left(\bar{\eta}_{AM}\right)\right]}{\left(\bar{\eta}_{PM}\alpha\left(\bar{\eta}_{AM}\right)-\bar{\eta}_{AM}\alpha\left(\bar{\eta}_{PM}\right)\right)^2}\right)\Bigg]\\
=&\frac{f\left(\bar{\eta}_{PM}\right)}{\alpha\left(\bar{\eta}_{PM}\right)}+\frac{\alpha'\left(\bar{\eta}_{PM}\right)\left(1-F\left(\bar{\eta}_{PM}\right)\right)}{\alpha\left(\bar{\eta}_{PM}\right)^2}-f\left(\bar{\eta}_{PM}\right)\frac{\bar{\eta}_{PM}-\bar{\eta}_{AM}}{\bar{\eta}_{PM}\alpha\left(\bar{\eta}_{AM}\right)-\bar{\eta}_{AM}\alpha\left(\bar{\eta}_{PM}\right)}\\
&-\left(F\left(\bar{\eta}_{PM}\right)-F\left(\bar{\eta}_{AM}\right)\right)\left[\frac{1}{\bar{\eta}_{PM}\alpha\left(\bar{\eta}_{AM}\right)-\bar{\eta}_{AM}\alpha\left(\bar{\eta}_{PM}\right)}\right.\\
&\qquad\left.-\frac{\left(\bar{\eta}_{PM}-\bar{\eta}_{AM}\right)\left[\alpha\left(\bar{\eta}_{AM}\right)-\bar{\eta}_{AM}\alpha'\left(\bar{\eta}_{PM}\right)\right]}{\left(\bar{\eta}_{PM}\alpha\left(\bar{\eta}_{AM}\right)-\bar{\eta}_{AM}\alpha\left(\bar{\eta}_{PM}\right)\right)^2}\right]
\end{aligned}$$

Assumption-conc and assumption-concMenu together imply that the left hand side of each FOC is decreasing in the relevant cutoff while the right hand side is increasing. The left hand side decreases as p_a gets smaller, and so by the same logic as for Proposition 1, the solution to these FOCs decreases in as p_a decreases.

$\bar{\eta}_{AM}$ decreasing implies a_A increasing, while $\bar{\eta}_{PM}$ decreasing implies p_P decreasing, so the welfare results for the inframarginal A and P consumers follow immediately. For

consumers near $\bar{\eta}_{AM}$, first note that $U(\bar{\eta}_{AM}) = 0$ while equilibrium utility is positive for consumers not indifferent between two products is positive, so if $\bar{\eta}_{AM}$ decreases, the utility for consumers at the (former) $\bar{\eta}_{AM}$ must increase. Given continuity of U_M and the fact that $\frac{dU_M(\bar{\eta}_{AM})}{dp_a} > 0$, there must be some neighborhood of $\bar{\eta}_{AM}$ such that $\frac{dU_M(\eta)}{dp_a} > 0$ for η in that neighborhood. Call the boundary of this neighborhood η_{M1}. The proof of existence of a neighborhood where welfare is decreasing around $\bar{\eta}_{PM}$ follows analogous logic, except noting that for the consumers at the *post* change $\bar{\eta}_{PM}$ utility is 0 and so for an instantaneous change in p_a there must be a neighborhood below $\bar{\eta}_{PM}$ such that utility is decreasing. $\square$