Advertising platforms and privacy

Ruizhi Zhu[1], Sridhar Moorthy[2], and Xianwen Shi[2]

[1]Northeastern University [2]University of Toronto

November 27, 2024
Keywords: Advertising, platform, privacy, consumer welfare

Abstract

We examine the implications of privacy-motivated targeting restrictions on consumer welfare, advertisers' customer acquisition costs, and platform revenue. In our model, competitive product firms reach consumers by placing informative ads on a monopoly advertising platform; consumers are horizontally differentiated in their product preferences and in their willingness to pay for a non-preferred product, but they don't have any intrinsic privacy preferences, nor an aversion to advertising *per se*.

In this context we show that both platform and consumers would be better off without any privacy restrictions if ad rates were exogenous to the privacy regime. However, in the more realistic scenario of endogenous ad rates, consumers with flexible product preferences are likely to be better off under privacy. This is because a platform with market power, selling informative ads to competitive product firms, will recognize the threat that cross-selling by mistargeted ads poses to ad volume. To compensate, the platform will lower ad rates, which end up benefiting flexible consumers in the form of lower product prices.

1 Introduction

The past decade has seen a surge in online advertising through third-party platforms such as Google and Meta. These platforms build profiles of users by tracking their online activities—browsing behavior, text and e-mail messages, photo and video posts, social media interactions—and sell them to advertisers. Advertisers value those profiles because they help them craft better ads and target consumers more efficiently. The multi-billion dollar valuations of companies like Alphabet and Meta are testimony to the value of targeted ads and the market power of advertising platforms.

Concurrent with these developments, consumer privacy concerns have also risen, and are now front and center in the debate about how to regulate the online platforms (O'Neil 2021; McKinnon 2022). Some governments have already taken measures to protect consumer privacy. For example, the European Union enacted the General Data Protection Regulation (GDPR) in 2016 to give users greater control over their personal information (Aridor et al. 2021; Johnson 2024); in the U.S., the State of California passed the California Consumer Privacy Act (CCPA) in 2018, and in China, the Personal Information Protection Law (PIPL) was adopted in 2021. Some companies, such as Apple, have embraced privacy as a marketing strategy.[1] Its App Tracking Transparency Initiative (ATT), announced in 2021, has required apps to seek users' explicit permission before tracking them across other apps and websites. Meta has described ATT as a "harmful policy" that makes "it harder and more expensive for businesses of all sizes to reach their customers" (Vranica et al. 2022).[2]

The academic discourse on privacy is a tale of two narratives with diametrically opposite conclusions for consumer welfare. On the one hand we have the standard textbook account that a monopolist, given the opportunity to target consumers based on their willingness to pay, would price discriminate, charging each consumer her reservation price (Tirole 1988, Chapter 3; Acquisti et al. 2016). Faced with such a firm, consumers would be better off guarding their privacy. On the other hand, in a competitive market, the same ability to target prices would intensify competition: each reservation price becomes its own market effectively, and firms compete vigorously knowing that their aggressive price offers in one market are shielded from other markets (Thisse and Vives 1988; Moorthy and Tehrani 2023). In this scenario, then, consumers might be better off without privacy. What remains elusive is to find an instrumental demand for privacy in a competitive market.

In this paper we show that it is possible to do so. A notable innovation—missing from the papers cited above—is the explicit modeling of the advertising platform as a third actor, besides consumers and product firms. In our model, consumers are horizontally differentiated in their product preferences and in their willingness to pay for a less-preferred product, but they don't have any intrinsic privacy preferences, nor an aversion to advertising *per se*. Competitive product firms reach these consumers by placing ads on a monopoly advertising platform. How targeted those ads

[1] https://www.apple.com/ca/privacy

[2] Wernerfelt et al.'s (2022) experiments at Meta suggest that this concern was well-founded. They find that the median advertising cost per incremental customer increases by 37% when Meta's targeting algorithms are prohibited from using "off-site data." Small scale advertisers and those in CPG, Retail, and E-commerce are especially affected, a result echoed in Aridor et al. (2024). Sun et al. (2024) report similar results based on experiments at Alibaba. Kraft et al. (2023) argue that ATT reduced publishers' ad revenue from Apple users by 21%.

are depends on the prevailing privacy regime. Under no privacy, the advertising platform can make it possible for advertisers to target specific consumer types and set type-specific prices and advertising schedules; under full privacy, neither of those things is possible. In each privacy regime, the platform chooses ad rates to maximize its ad revenue, anticipating the product-market equilibrium to follow;[3] advertisers take those ad rates and targeting constraints (if any) as a given when deciding how many ads to buy and what prices to charge.

Our principal result is that privacy can have positive value for some or all consumers even in competitive product markets. However, it is not easy to find such value: after all, our consumers are intrinsically privacy-neutral—to the extent they care about privacy only because it gives them better product-market outcomes—and our product markets are competitive. Moreover, giving up privacy has the attractive feature of making advertising more efficient; when advertisers are more productive, they advertise more, reaching more consumers and increasing consumption. The platform also benefits. Indeed, these arguments are powerful enough to ensure that consumers are better off giving up their privacy if ad rates are unable to respond to changes in the privacy regime, as might be the case if the ad market was perfectly competitive and ad rates were determined solely by the costs of running an ad platform. For consumers to prefer privacy in our model, it is necessary that ad rates be endogenous to the privacy regime.[4] This underscores the importance of recognizing the strategic role of the advertising platform when debating the value of privacy in competitive product markets; a partial-equilibrium analysis focusing on the product market alone is unlikely to reach the right conclusion.

When ad rates are endogenous to the privacy regime, several less obvious effects of privacy are exposed. First, when consumers are picky in their product preferences, the platform can easily compensate advertisers for their loss in advertising productivity by lowering ad rates while maintaining no-privacy consumer-level outcomes. In this scenario, platform revenue does not suffer, advertisers' customer acquisition costs do not go up, and consumers continue to be privacy-neutral. Second, when consumers are flexible in their product preferences, mistargeted ads are not entirely wasted: they generate sales among consumers who are not reached by ads for their first-best product. However, such sales dilute total surplus, leading to a lower demand for advertising. We show that the ad platform cannot find adequate compensation for this reduction in demand; not only do ad rates decline in equilibrium, ad volume does, too, resulting in a reduction in platform revenue. However, consumers gain, making them better off under privacy. Finally, when there is a mix of consumers,

[3]The mechanics of ad rate-setting in online markets differs, of course, from this idealization. Typically, an ad platform decides how many ad slots to offer, i.e., the ad load, and conducts an auction for each ad slot; ad rates arise as equilibrium outcomes of those auctions. For our purposes, however, it suffices to take a reduced-form approach and assume that ad rates are chosen directly. To see the correspondence between the two approaches, imagine that a platform is running a series of simultaneous single-ad-slot auctions based on the ad load it wants. Advertisers bid in those auctions, anticipating the effect of their bids on how they will compete in the product market. The resulting equilibrium generates an ad demand curve: a mapping from ad load to ad rates. By picking an ad load, then, the platform is effectively picking an ad rate.

[4]This is, of course, what we should expect in the real world given that it is platforms like Meta that are determining ad rates and they have market power. Indeed, Meta's sales have rebounded after a brief blip in 2022 following Apple's ATT announcement (https://www.macrotrends.net/stocks/charts/META/meta-platforms/revenue).

some picky and others flexible, we get asymmetric effects: firms producing the preferred products of the former see mistargeted ads generating sales among unreached flexible consumers, but not vice-versa. As a result, the former suffer less from mistargeted ads than the latter, and the platform's ad rate setting problem becomes more complicated. In this scenario, a variety of equilibria may arise, and consumers may differ in their attitude toward privacy. In particular, consumers with flexible product preferences are more likely to embrace it than picky consumers.

It is worth emphasizing that our results are obtained in a setting where product markets are competitive: all product firms make zero profits in equilibrium, both under privacy and under no privacy.[5] Furthermore, the ad platform always weakly prefers no privacy to privacy: it can replicate under no privacy whatever outcomes it achieves under privacy.[6] Indeed, our analysis shows that this preference becomes strict if the platform is unable to replicate under privacy the same consumer-level outcomes as under no privacy. These results are consistent with the findings in the empirical literature on privacy restrictions (Wernerfelt et al. 2022; Kraft et al. 2023; Aridor et al. 2024; Sun et al. 2024).

Our model puts the onus of consumers' privacy preferences on the market power of ad platforms. It suggests that stricter privacy regulations may induce ad platforms to cut advertising rates, possibly more for some products than others, which then has knock-on effects on how different advertisers compete in product markets, and hence consumer welfare. Our message to regulators is thus the following: When advertising is an important tool for providing information to consumers, the market power of ad platforms may be a more important determinant of how consumers fare under different privacy regulations than product market competitiveness itself.

The rest of the paper is organized as follows. In the next section we review the two literature most closely related to our work: the targeting literature and the privacy literature. Section 3 describes our model and defines what an equilibrium means in this paper. Section 4 describes the equilibrium under no privacy. In Section 5 we examine two symmetric cases of our model, one with all consumers equally very picky and the other with all consumers equally quite flexible, and show that both endogenous ad rates and flexible consumer preferences are necessary for consumers to be better off under privacy. Section 6 looks at the asymmetric counterpart, one consumer type very picky and the other quite flexible, and shows that the latter is more likely to embrace privacy. Section 7 puts the results of the preceding three sections in perspective by comparing results across the three versions of our model. Finally, Section 8 concludes the paper.[7]

[5]Even though privacy has no effect on advertisers' profits, it may have an effect on their behavior. As we document in Propositions 3 and 8, their advertising volume, customer acquisition costs, prices, and sales can all vary between no privacy and privacy.

[6]This is in contrast to papers such as Levin and Milgrom (2010), Amaldoss et al. (2015), Sayedi (2018), and Rafieian and Yoganarasimhan (2021), which have suggested that a platform might actually prefer some targeting restrictions to "thicken" markets; this motivation is absent in our model because our product firms are always competitive—under privacy and under no privacy.

[7]Due to space constraints, the main body of the paper only presents key ideas and intuitions. Proofs of the lemmas and propositions are divided between Appendix A and an Online Appendix.

2 Background

The issues explored in this paper are intimately related to two streams of literature: the literature on targeting and the literature on privacy. Before we discuss those connections, however, it would be useful to summarize here briefly the literature on informative advertising with probabilistic exposure, the advertising framework adopted in this paper (and many others: Butters (1977), Grossman and Shapiro (1984), Tirole (1988, Section 7.3.1), Stegeman (1991), Chen and Iyer (2002), Bagwell (2007), Esteves and Resende (2016), and Moorthy and Tehrani (2023)).

In this framework, consumers are unaware of the firms and uninformed about their products and product prices unless exposed to their advertising.[8] Advertising exposure is probabilistic: firms only control ads sent, not ads received. Consumers are not guaranteed to receive an ad—much less a specific ad; each consumer in a target market has an equal probability of receiving a targeted ad. Consumers who receive no ads can't buy; the rest choose, from all the ads received, the product price offer that maximizes their consumer surplus. Butters (1977) shows that this advertising technology implies a distribution of equilibrium product prices even in a homogeneous market with competitive firms; furthermore, the amount of advertising is socially optimal. The latter result, however, turns out to be fragile. In Stegeman's (1991) model with heterogeneous consumer valuations, the amount of advertising is less than socially optimal; in Grossman and Shapiro's (1984) model with horizontal product differentiation, the amount of advertising is more than socially optimal. Our model combines aspects of these two models: from Grossman and Shapiro (1984) we borrow horizontal product differentiation; from Stegeman (1991) we borrow a continuum of firms and consumers. But our model is also different from them: while these papers only focus on untargeted advertising and product prices, we compare targeted advertising and product prices with untargeted advertising and product prices; while ad rates are exogenous in these papers, we examine both exogenous and endogenous ad rates; finally, while there is no advertising platform in these papers, we have a monopoly ad platform intermediating between advertisers and consumers and setting ad rates.

As noted by Bergemann and Bonatti (2011, p. 435), "the Internet has introduced at least two technological innovations in advertising, namely (i) the ability to relate payments and performance (e.g., pay per click), and (ii) an improved ability to target advertisement messages to users."[9] Most of the targeting literature, including Chen et al. (2001), Iyer et al. (2005), Bergemann and Bonatti (2011), Johnson (2013), and Moorthy and Tehrani (2023), focus on effect (ii). Johnson (2013), for example, examines how improvements in targeting capability affect firms and consumers when the latter can use ad-avoidance technologies to avoid ads. He finds that while product firms are better off, consumer utility has a U-shaped relationship with targeting accuracy. In Chen et al. (2001), the advertising messages in question are price messages. They show that an improved ability to target price in Narasimhan's (1988) duopoly model may soften competition when targeting accuracy is low, leading to higher firm profits, but at high levels of targeting accuracy, further accuracy improvements lower profits. In the same model, Iyer et al. (2005) examine the implications of

[8] As discussed in footnote 17, this assumption is not as restrictive as it seems.

[9] See also Athey and Gans (2010) and Goldfarb (2014) for related discussion.

targeting both advertising spending and product prices ("full targeting") vis-a-vis each alone ("partial targeting"). They find that the marginal value of adding advertising spending targeting to price targeting is positive for both firms while the marginal value of adding price targeting to advertising spending targeting is zero. Moorthy and Tehrani (2023) suggest that these results may be special to Narasimhan's (1988) model; in Hotelling's model, where targeting would be on consumers' preferences rather than on consumers' loyalty and switching behaviors, either kind of partial targeting, as well as full targeting, may be collectively optimal, depending on the extent of product differentiation and the cost of advertising.

In all these papers, advertising costs are assumed to be fixed, exogenously, and uniform across firms.[10] In comparison, ad rates in our model are endogenously determined by an ad platform as a function of the privacy regime, and can differ across products. Bergemann and Bonatti (2011) develop a model with many firms and online and offline advertising media to study how the improved targetability afforded by online media affects equilibrium ad rates and the competition between online and offline media. Using Butters's (1977) advertising technology and allowing ad rates to be endogenously determined in equilibrium, they find that although targeting increases the social value of advertising, equilibrium ad rates first increases and then decreases in targeting ability. In their model, product prices are exogenous, ad rates are uniform across products, and the ad market is perfectly competitive. Our model uses the same advertising technology as them, but we allow both ad rates and product prices to be endogenously determined by the privacy regime in equilibrium. More importantly, in contrast with Bergemann and Bonatti (2011) and much of the literature,[11] we allow ad rates to differ across products (and potentially also across market segments). Thus, our model can capture not only effect (ii) above but also effect (i): the ad platform can relate ad rates to click-through rates and favour one type of product over another. Finally, by assuming a monopoly ad platform, our analysis comes closer to replicating the real-world context under which privacy regulations are most relevant—a context in which ad platforms such as Meta and Google hold considerable market power.

Since changes in targetability in our model are motivated by changes in privacy regulations, our paper is naturally linked to the growing literature on privacy (see Acquisti et al. (2016) for a survey). One strand of this literature, e.g., Choi et al. (2019), Acemoglu et al. (2022) and Bergemann et al. (2022), focuses on information externalities among users and shows how the social value of privacy may differ from personal values of privacy.[12] Another strand of literature recognizes that firms may offer different product prices and/or products to consumers based on their past purchasing behavior (Villas-Boas 1999; Fudenberg and Tirole 2000; Fudenberg and Villas-Boas 2012), or on

[10] See also Shin and Yu (2021). Their model also features fixed ad rates, but in their model consumers infer their preferences from the targeted ads they receive.

[11] A noticeable exception is Galeotti and Moraga-González (2008). In a related model with two consumer market segments where two homogeneous firms can each send targeted ads to one or both segments, they show that firms can earn strictly positive profit only if the per-consumer advertising costs differ significantly across segments. Their ad rates, however, are exogenous.

[12] Johnson (2013) and Garratt and Van Oordt (2021) suggest, for example, that such externalities may arise when using ad-avoidance technologies and electronic cash, respectively.

their voluntary disclosures (Ichihashi 2020). (Privacy regulations may determine whether sellers can use data on consumers' past purchasing behavior or must rely only on data voluntarily disclosed.) One way to model the value of privacy is in terms of privacy costs—costs incurred by consumers to stay anonymous. Conitzer et al. (2012) show that consumer welfare has an inverted-U shaped relationship with this cost, and Baye and Sappington (2020) argue that privacy protection benefits myopic consumers while imposing a cost on sophisticated ones. Similarly, in a duopoly setting where firms can use consumer information to price discriminate while consumers can pay a privacy cost to opt out of such discrimination, Montes et al. (2019) show that firms may be worse off and consumers be better off as privacy cost increases. Ke and Sudhir (2023) examine the full gamut of privacy rights embedded in privacy regulations such as GDPR, including the right to opt-in to share data and the right to be "forgotten." In a two-period model where the value of product personalization differs among consumers, they show that consumers may or may not be better off with privacy rights if the product market is a monopoly, but in an oligopoly, where the firms are ex-ante homogeneous, consumers will always be better off with privacy rights. By contrast, in our static model, privacy settings are exogenous (controlled by an authority outside the model), firms are horizontally differentiated, and the product market is competitive. Our goal is to examine how privacy regulations change ad prices and the ripple effect this has on ad volumes, product prices, and consumption outcomes. We show that different consumers may have different attitudes toward privacy even when the product firms themselves are neutral, making zero profits in each privacy regime.

Another stream of literature (Levin and Milgrom 2010; Amaldoss et al. 2015; Sayedi 2018; Rafieian and Yoganarasimhan 2021) argues that platforms may have an incentive to lower the level of ad targetability because imperfect targetability can improve market thickness and hence raise platform profits. In contrast, since the product market in our model is competitive, and hence thick both under privacy and under no privacy, this benefit of privacy is absent. Moreover, since the platform can always replicate privacy outcomes under no privacy—by simply ignoring the finer data at its disposal—it always weakly prefers no privacy to privacy.

Finally, there is a growing literature on the interaction between platforms and consumer privacy. In Casadesus-Masanell and Hervas-Drane (2015), consumers disclose information to firms to receive better service but such disclosure entails a disutility. Firms compete for consumers in both privacy and product prices and can generate revenue from consumer purchases and/or from sales of consumer data to third parties. They show that it is optimal for firms to focus on one revenue source, and more intense competition does not necessarily lead to a higher level of privacy. In Campbell et al. (2015), firms differ in the scope of products they are offering, and a more specialized firm offers a narrower range of products but with higher product quality. Consumer data can help firms optimize their product offerings, but consumers suffer an intrinsic loss when their data are collected and used by these firms. They show that privacy regulation has a differential impact on firms selling different ranges of products. In comparison, consumers in our model do not have intrinsic preferences for privacy and their attitude toward privacy is completely determined by the consumer surplus they receive under different privacy regimes. In this respect we are similar to De Corniere and De Nijs (2016), who also assume that consumers do not have intrinsic preferences for privacy. However, in their model firms compete in an auction for the monopoly right to serve a consumer; the ad platform (auctioneer) can choose whether to disclose consumer information to allow bidders to have a better estimate of the value of serving a particular consumer. The privacy regime is determined by the platform, and it prefers no privacy (i.e., full disclosure) only when the number of firms is sufficient large. By contrast, in our model, privacy settings are exogenous—controlled by a third party outside the model—and the platform always weakly prefers no privacy. We focus on the differential impact of privacy on different types of consumers in a competitive as opposed to a monopoly product market, in contrast to De Corniere and De Nijs's focus on the platform's incentive to disclose. Thus the two papers are complementary.

3 Model

We build a model which, while minimalistic, is still capable of delivering on three objectives: (i) capture key strategic features of the online advertising environment in which privacy issues are playing out in the real world, (ii) generate rich competition patterns in the product market, and (iii) provide a micro foundation for the differential effects of privacy regulations on different types of consumers.

As noted in the Introduction, the online advertising environment is characterized by advertising platforms like Meta with a lot of market power, while the advertisers advertising in them have hardly any.[13] Accordingly, in our model, the advertising platform is a monopoly and advertisers are a continuum of infinitesimal product firms with no market power. See Figure 1.

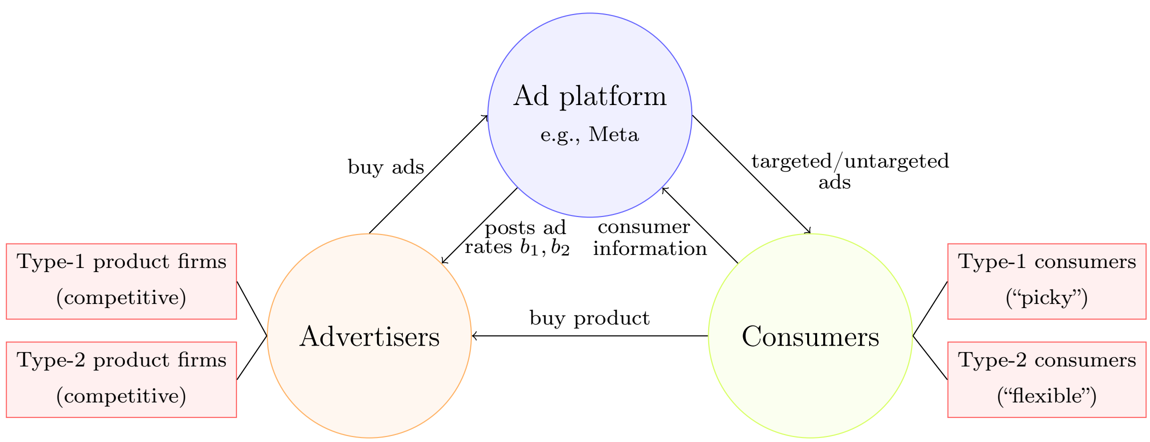

Figure 1: Model framework

On the consumer side, we have a continuum of infinitesimal consumers, horizontally differentiated

[13] As Wernerfelt et al. (2022) note: "Small businesses rely heavily on digital and social media advertising, whereas larger businesses generally advertise on more traditional channels (Moorman 2022, Peck 2022)."

into two types, $i \in \{1, 2\}$, based on the product type they prefer. For example, in the cereal category we can identify two consumer types: people who prefer gluten-free cereals and people who prefer "regular" cereals. As this example suggests, corresponding to the two consumer types, two product types are defined: product type 1, the product type preferred by type-1 consumers, and product type 2, the product type preferred by type-2 consumers. Denote the fraction of type-1 consumers by γ; then the fraction of type-2 consumers is $1 - \gamma$.[14]

Each consumer is assumed to have a demand of up to one unit of a product. Horizontal product differentiation is captured in the following utility function:

$$u_{ij} = \begin{cases} u & \text{if} \quad i = j \\ \beta_i u & \text{if} \quad i \neq j \end{cases}$$

with $u > 0$ and $0 \leq \beta_i < 1$. In other words, a type-1 consumer is willing to pay up to u for a unit of product-type 1, but is only willing to pay up to $\beta_1 u < u$ for a unit of product-type 2; similarly, a type-2 consumer is willing to pay u for a unit of product type 2, but is only willing to pay $\beta_2 u$ for a unit of product type 1. β_i thus represents how "picky" or "flexible" a type-i consumer is; the higher the β, the less picky (more flexible) she is. By placing a type index on β we are allowing for the possibility of asymmetric pickiness—the two consumer types differing in their pickiness, making this a variant of the familiar Hotelling model.[15]

Turning to the supply side, the market for each product type consists of a continuum of product firms with no market power. We call the firms making type-1 product, type-1 firms, and the firms making type-2 product, type-2 firms.[16] Their marginal cost of production is $c \in (0, u)$ and they have no fixed costs (other than advertising costs). Thus, when a type-j firm sells to a consumer of type i with $j = i$ (product type matches consumer type), the total surplus per unit is $u - c$, but when a type-j firm sells to a consumer of type i with $j \neq i$ (product type mismatched with consumer type), total surplus per unit is $\beta_i u - c$. Denote the ratio of these total surpluses by ρ_i:

$$\rho_i \equiv \frac{\beta_i u - c}{u - c}, \quad i = 1, 2. \tag{1}$$

When $\rho_i < 0$, i.e., $\beta_i u < c$, type-i consumers are so picky that it is impossible to sell unmatched products to them while covering variable costs; then the only equilibrium possible under privacy is one where type-i consumers buy matched products. By contrast, when $\rho_i \geq 0$, type-i consumers are flexible enough that it is possible to have a privacy equilibrium where they buy an unmatched product. Later we will argue that if $\rho_i < 0$ for both types of consumer, then it is impossible for

[14] Sometimes, for ease of exposition, we refer to the fraction of type-i consumers as $\gamma_i \in (0, 1)$ ($\gamma_1 + \gamma_2 = 1$).

[15] To see the correspondence, imagine two firms and two groups of consumers located at either end of a unit line (but none in between). In the standard Hotelling model, unit "transportation costs" are assumed to be the same for all consumers. Translated to our context, this is as if consumers are homogeneous in their reservation prices for their preferred product *and* homogeneous in their (lower) reservation prices for the non-preferred product. Since we are not assuming the latter homogeneity, we are allowing for the possibility that the two groups of consumers have different unit transportation costs.

[16] Since the market is competitive for each product type, it doesn't make a difference whether a firm makes only one type of product or both types of product.

consumers to be better off under privacy. In other words, for consumers to benefit from privacy, at least some consumers must be flexible enough to make $\rho_i > 0$.

Let us turn now to the advertising technology used by the product firms. Following Butters (1977), ads in our model are informative and have probabilistic exposure within the markets to which they are targeted. As discussed in Section 2, this is the advertising technology used in virtually the entire informative advertising literature, including: Grossman and Shapiro (1984), Tirole (1988, Section 7.3.2), Stegeman (1991), Chen and Iyer (2002), Bagwell (2007), Bergemann and Bonatti (2011), Esteves and Resende (2016), and Moorthy and Tehrani (2023). In this framework, exposure to a firm's ads is necessary to become informed about its product attributes and price, and ad exposure, within targeted groups, is random.[17,18] If a block A of ads is sent to a unit measure of consumers, the fraction of consumers who are exposed to an ad is $1 - e^{-A}$.[19] Consumers decide whether and which product to buy based on the ads they see. If a consumer sees ads from only one product, she can buy only that product, which she will, if it yields a positive consumer surplus; on the other hand, if she sees ads from multiple products (of the same or different types), she will compare their respective offers and choose the one offering the highest positive consumer surplus.

Product firms compete via the number of ads they send and the prices they advertise. The targetability of ads depends on the privacy regime prevailing, presumably dictated by regulations. We consider two privacy regimes: no privacy (NP) and full privacy (P). Under NP, the platform observes consumers' types and has the capability to target ads to particular consumer types; advertisers, taking advantage of this capability, can decide whether to target one or both types, and if the latter, whether to customize their offers by consumer type or not. Under P, the platform cannot identify consumers' types, making it impossible for advertisers to target specific types; in that case their product-price offers must necessarily be the same for all consumer types. We assume that consumers do not have any intrinsic preference for privacy (Becker 1980; Lin 2022), nor an aversion to ads *per se*, or an intrinsic preference for well-targeted over mis-targeted ads.[20] In other words, going in, all

[17] The second part of this assumption is simply an empirical fact of online advertising: firms only control "intend to expose," not actual exposure. For example, even if a consumer is in a targeted group, she will not see any ads if she is using an ad blocker (Johnson 2013). Even if a consumer is not using an ad blocker, there is no guarantee that she will see an ad just because it appears in her "feed." Lacking visibility into which among their targeted consumers will see their ads and which won't, advertisers only control probability of exposure, not actual exposure. (The only thing advertisers can guarantee is that non-targeted consumers will *not* see their ads.) As for the first part of the assumption, it should not be interpreted literally as requiring that consumers start from ground zero—that they have never heard of the firms and know nothing about their products and prices until exposed to their advertising. A more realistic interpretation would allow for the possibility that consumers have heard of the advertisers before, and possibly also know something about their products, but that those memories have faded—faded enough that they can't act on them without further advertising exposure. The function of advertising exposure, then, is to rekindle and refresh those memories (Iyer et al. 2005). As Moorthy and Tehrani (2023) note, "this extends the scope of informative advertising to virtually any kind of advertising—–even to so-called 'uninformative advertising.'"

[18] This advertising technology, in conjunction with infinitesimal firms, delivers competitive product markets, i.e., markets in which firms have no market power. It is more challenging to justify the benefits of privacy for consumers in such a setting than, say, a setting where the product market is a monopoly or an oligopoly.

[19] To see this, consider a finite market with n consumers. An ad falls randomly on one of these n consumers, so each consumer has $1/n$ chance of being hit by that ad. If a firm sends A units of ads per consumer, i.e., An total number of ads, then the probability that any given consumer observes none of those ads is $(1 - 1/n)^{An}$. This probability converges to e^{-A} as n tends to infinity because $(1 - 1/n)^{An} = (1 + 1/(n-1))^{-An}$ and $\lim_{n\to\infty} (1 + 1/n)^n = e$.

[20] This is not because we don't think these preferences exist. Rather, it is because we want to create an environment

consumers are privacy-neutral. To the extent they prefer privacy, then, it is because their consumer surplus is higher in the privacy equilibrium than in the no-privacy equilibrium.

We assume that the ad platform sets ad rates by product type and target market to maximize ad revenues. Under NP, there are potentially three target markets for each type of firm—type-1 consumers, type-2 consumers, all consumers—so the platform may quote six ad rates; under P, since the only target market is the set of all consumers, there are only two ad rates.[21] The interaction between platform and advertisers is a two-stage game: the platform sets ad rates first; advertisers, taking those ad rates as given, then decide how many ads to buy and what product prices to advertise.[22] We call the equilibrium in the second stage of the game the "product-market equilibrium." In a subgame perfect equilibrium, the platform will recognize that the product-market equilibrium depends on the ad rates it sets, and will set those rates anticipating the advertising revenue to come from the product-market equilibrium.

Following Stegeman (1991), we represent the product-market equilibrium in a target market as a pair of non-decreasing continuous advertising distribution functions, $(A_1(\cdot), A_2(\cdot))$, where $A_i(p)$ $(i = 1, 2)$ is the total number of ads with product prices less than or equal to p, sent by type-i firms per unit measure of the target market.[23] Note that these advertising distribution functions represent the collective advertising of all type-i firms, not the advertising of an individual firm—the latter is infinitesimal by definition (because the firm itself is infinitesimal). Furthermore, by writing these advertising quantities as "per unit measure of the target market," we are really capturing advertising *intensities,* which is the right metric to take to the privacy-no privacy comparison. For the size of the target market will vary by privacy regime: under privacy the target market will have to be the entire market; under no privacy, it can be a specific consumer type or the entire market.[24]

Let the ad rate for product i in a particular target market be b_i per ad unit; then, if a product-i advertiser buys n ad units targeting that market it will pay nb_i to the ad platform. Denote the marginal profit of a type-i firm in a target market from an additional ad with product price p, when type-i firms are collectively sending A_i ads, by $\pi_i(p; A_1, A_2)$.

Definition (Equilibrium in a target market) *An equilibrium in a target market is a pair of ad rates* (b_1, b_2) *and a pair of non-decreasing continuous advertising distribution functions* $(A_1(\cdot), A_2(\cdot))$, $A_i(p) : [c + b_i, u] \to \mathbb{R}_+$, $i = 1, 2$, *such that:*

that is privacy-neutral—neither stacked in favor of privacy nor stacked against it. Whatever privacy preferences develop in our model will be driven purely by instrumental considerations.

[21] Later we will argue that the additional capability of targeting by target market under NP is redundant because in equilibrium, advertisers with type-i products will target type-i consumers only.

[22] As noted in footnote 3, this model of ad platform choosing ad rates directly may be seen as a reduced-form representation of ad auctions where the platform sets the ad load and ad rates arise as auction equilibria.

[23] Why are these functions and not a single (A, p) pair? The reason has to do with the probabilistic nature of advertising exposure. No single price can be an equilibrium price because a firm can gain either by advertising a lower price—hoping for additional sales from consumers who see this offer and no better offer—or by advertising a higher price—hoping for additional margin from consumers who see only this offer. The logic is akin to the logic of mixed-strategy equilibria in Narasimhan's (1988) duopoly model.

[24] In Stegeman's (1991) model, products are homogeneous while consumers are vertically differentiated. Furthermore, consumers' types cannot be identified or separated, so all consumers are equally likely to receive each ad. In comparison, in our model under no privacy, the platform can identify and target individual consumer types, so the advertising market can operate independently for each consumer type.

1. *(b_1, b_2) maximizes ad revenue for the monopoly platform, anticipating the product-market equilibrium to follow in the target market.*

2. *The product-market equilibrium in a target market is characterized by a pair of advertising distribution functions $(A_1(\cdot), A_2(\cdot))$ satisfying:*

   (i) $A_i\,(c + b_i) = 0$,

   (ii) $\pi_i\,(p; A_1, A_2) \leq 0$ *for all product prices* $p \in [c + b_i, u]$,

   (iii) $A_i(p') = A_i(p'')$ *if* $\pi_i\,(p; A_1, A_2) < 0$ *for all* $p \in [p', p'']$,

where $A_i\,(p)$ $(i = 1, 2)$ is the number of ads with prices less than or equal to p sent by type-i firms per unit measure of the target market.

Equilibrium advertising distribution functions have a domain $[c + b_i, u]$ because a price below $c + b_i$ will produce a sure loss for a type-i firm, while consumers will not accept a price above u. Condition (i) requires that type-i firms do not advertise product price $c + b_i$, for if a firm did, while it breaks even on the ads seen by a consumer, it incurs a loss on the ads not seen (which always occurs with positive probability). Condition (ii) says that the marginal profit from an additional ad is non-positive at all feasible product prices $p \in [c + b_i, u]$. Condition (iii) says that firms will not advertise product prices that generate a negative marginal profit. We will rely on condition (ii) to construct the product-market equilibrium in each privacy regime. As in Stegeman (1991), because each target market is competitive, all firms will have zero expected profits in equilibrium in each privacy regime.

The following lemma follows immediately from our definition of equilibrium.

Lemma 1 (Necessary condition for equilibrium ad rates). *Equilibrium ad rates in each target market must satisfy $b_{-i}/b_i \geq \rho_i$ for $i = 1, 2$.* □

In other words, it will never be in the platform's interest to set ad rates in a target market so lopsidedly in favor of one type of firm that the other type is unable to compete "in its own backyard"—the consumers who prefer it. Ad rates that satisfy this lemma guarantee that mismatched firms cannot out-compete matched firms.[25]

Since the game between platform and advertisers is two-stage, a subgame-perfect equilibrium calls for computing the product-market equilibrium given ad rates first, and then optimizing the platform's ad rates to maximize ad revenue. It turns out that this exercise is relatively straightforward under no privacy and relatively complicated under privacy. Therefore, to simplify the exposition, we will proceed as follows:

1. First, in Section 4, we will derive the no-privacy equilibrium analytically.

2. Then, in Section 5, we will begin our analysis of the privacy equilibrium by examining two symmetric cases of our model: both consumer types equally very picky ($\rho_1 = \rho_2 < 0$) and both

[25] Equal ad rates automatically satisfy this lemma.

consumer types equally somewhat flexible ($\rho_1 = \rho_2 > 0$). These examples will illustrate how consumer and firm behavior change under privacy, and the impact this has on how the platform sets ad rates. Two conditions will emerge as necessary for consumers to benefit from privacy: (1) existence of consumers with flexible preferences, and (2) ad rates that are endogenous to the privacy regime.

3. Finally, in Section 6, we will examine the asymmetric case where one type of consumer (say type-1) is very picky ($\rho_1 < 0$) while the other is somewhat flexible ($\rho_2 > 0$). This case will show that different consumer types may develop different attitudes toward privacy. In particular, it will confirm the intuition coming from the symmetric examples that flexible consumers stand to gain more from privacy regulations than picky consumers.

4 Equilibrium under no privacy

In this section we will execute Step 1 of the agenda outlined above by first computing the product-market equilibrium for given ad rates (Section 4.1), and then optimizing the ad rates to maximize the platform's ad revenues (Section 4.2).

4.1 Product-market equilibrium for given ad rates

Under no privacy, the ad platform can classify consumers by type, and advertisers can use this classification to target consumers with matched or unmatched product preferences (or both). However, it is easy to see that the latter cannot happen in a no-privacy equilibrium. It will never be in the platform's interest to allow mismatched sales. This is because total surplus is maximized when type-1 firms sell to type-1 consumers and type-2 firms sell to type-2 consumers, and the platform can also maximally extract surplus by restricting targeting to matched ads only. The platform can guarantee this outcome simply by prohibiting mismatched targeting or, more naturally, by setting (prohibitively) high ad rates for firms seeking mismatched consumers (compared to firms seeking matched consumers).[26] Effectively, then, the ad rates that matter are the ad rates that apply to matched consumers.

Let b_i denote the ad rate for type-i firms targeting type-i consumers ($i = 1, 2$) and $\left(A_1^{NP}(\cdot), A_2^{NP}(\cdot)\right)$ the advertising distribution functions of the two types of firms in the resulting product-market equilibrium. Then, by property (ii) of the definition of equilibrium, for all prices in the support of these functions, the marginal profit of a type-i firm from sending an additional ad to type-i consumers must be non-positive. To calculate this marginal profit, suppose a type-i firm tries to send Z additional ad units (per unit measure of type-i consumers) with product price p to these consumers. These new ads will generate additional sales in the amount of $\left(1 - e^{-Z}\right) e^{-A_i^{NP}}$ among those type-i consumers who do not receive type-i product ads with product prices less than p (because $e^{-A_i^{NP}(p)}$ is the fraction of type-i consumers who do not receive any type-i product ads with prices less than p

[26] See https://www.facebook.com/business/help/430291176997542?id=561906377587030 for how Facebook does this in an ad auction by setting "estimated action rates" and "ad quality."

in the putative equilibrium and $1 - e^{-Z}$ is the fraction of type-i consumers who receive at least one of the Z additional ads). The margin on each such sale is $(p - c)$, and the firm incurs an advertising cost of $b_i Z$. Hence its incremental profit will be

$$\Pi_i^{NP}\left(Z; p, A_1^{NP}(\cdot), A_2^{NP}(\cdot)\right) = \left(1 - e^{-Z}\right) e^{-A_i^{NP}(p)} (p - c) - b_i Z.$$

and its marginal profit will be

$$\pi_i^{NP}\left(p; A_1^{NP}(\cdot), A_2^{NP}(\cdot)\right) \equiv \left.\frac{\partial \Pi_i^{NP}\left(Z; p, A_1^{NP}(\cdot), A_2^{NP}(\cdot)\right)}{\partial Z}\right|_{Z=0} = e^{-A_i^{NP}(p)} (p - c) - b_i.$$

To obtain the product-market equilibrium we set this marginal profit equal to zero and solve for the advertising distribution function.

Lemma 2 (product-market equilibrium under no privacy). *Given ad rates $\{b_i\}$, $i = 1, 2$, for type-i firms targeting matched consumers, the unique product-market equilibrium under no privacy $(A_1^{NP}(\cdot), A_2^{NP}(\cdot))$ is given by*

$$A_i^{NP}(p) = \ln\left(\frac{p - c}{b_i}\right) \tag{2}$$

for $p \in [c + b_i, u]$ and $i = 1, 2$. □

Equilibrium advertising distribution functions, being cumulative, are increasing in price by definition. The more interesting observation is that the advertising density function, $1/(p - c)$, is decreasing in price, i.e., advertising intensity is decreasing in price. This is because of competitive pressure: as price increases, the chances of being undercut increases. Total advertising volume is $\gamma_i A_i^{NP}(u)$. Each advertising distribution function implies a corresponding sales distribution function (total sales of type-i firms with price $\leq p$)

$$S_i^{NP}(p) = \gamma_i \left[1 - e^{-A_i^{NP}(p)}\right] = \gamma_i \left(1 - \frac{b_i}{p - c}\right)$$

and total sales volume $S_i^{NP}(u) = \gamma_i (1 - b_i/(u - c))$.

Lemma 3 (Comparative statics of the product-market equilibrium under no privacy). *As b_i increases, type-i firms reduce advertising volume and increase average product price; there are no cross-product effects.* □

Under no privacy, advertisers of each type target their natural market only; there is no advertising spill-over to the "other" market. Type-i firms' total advertising is $\gamma_i A_i^{NP}(u) = \gamma_i(\ln(u - c) - \ln b_i)$, which is decreasing in b_i. Since the advertising density function is independent of b_i, and the lower bound of the support is increasing in b_i, average product price is increasing in b_i.

4.2 Optimal ad rates

Under no privacy, the platform may set ad rates that vary by product type and target market. However, as discussed earlier, the latter variation does not have any bite; it will never be in the platform's interest to allow mismatched sales. So, effectively, the platform has to optimize only two ad rates: b_1 for type-1 firms targeting type-1 consumers and b_2 for type-2 firms targeting type-2 consumers.

This optimization is straightforward. All that the ad platform has to do is maximize its ad revenue

$$R^{NP} = \gamma_1 b_1 A_1^{NP}(u) + \gamma_2 b_2 A_2^{NP}(u),$$

with respect to b_1 and b_2. This yields the following equilibrium under no privacy.

Proposition 1 (Equilibrium under no privacy). *In the equilibrium under no privacy, the optimal ad rate for type-i firms targeting type-i consumers is given by*

$$b_i^{NP} = \frac{u-c}{e} \quad \textit{for } i = 1, 2,$$

which implies the following product-market equilibrium:

$$A_i^{NP}(p) = 1 + \ln\left(\frac{p-c}{u-c}\right) \quad \textit{for } p \in [c + (u-c)/e, u] \textit{ and } i = 1, 2. \tag{3}$$

□

The platform sets equal ad rates for the two types of firm under no privacy because it wants both types to target their natural markets; this maximizes total surplus. (Under privacy, as we will see, this will not always be possible.) Both types of firms end up with the same sales and customer acquisition costs; both types of consumers enjoy the same consumer surplus.

4.3 Properties of the no-privacy equilibrium

- Sales distribution function

$$S_i^{NP}(p) = \gamma_i \left[1 - \frac{(u-c)}{e(p-c)}\right] \quad \text{for } p \in [c + (u-c)/e, u]$$

- Total sales of type-i product

$$S_i^{NP}(u) = \gamma_i \left(\frac{e-1}{e}\right) \quad \text{for } p \in [c + (u-c)/e, u]$$

- Total ad volume of type-i products

$$\gamma_i A_i^{NP}(u) = \gamma_i \left(\ln(u-c) - \ln b_i^{NP}\right) = \gamma_i$$

- Per-capita customer acquisition cost for type-i firms

$$\frac{b_i^{NP}\gamma_i A_i^{NP}(u)}{S_i^{NP}(u)} = \frac{u-c}{e-1}$$

- Platform revenue

$$R^{NP} = \frac{u-c}{e}$$

- Aggregate consumer surplus of type-i consumers

$$\int_{c+(u-c)/e}^{u} (u-p)\, d\left(\frac{S_i^{NP}(p)}{\gamma_i}\right) = \left(1-\frac{2}{e}\right)(u-c)\,.$$

5 Two symmetric cases

We will begin our analysis of the privacy equilibrium by examining two symmetric cases of our model: (i) both consumer types equally (very) picky in their product preferences ($\beta_1 = \beta_2 \equiv \beta \leq 0$) and (ii) both consumer types equally quite flexible in their product preferences ($\beta_1 = \beta_2 \equiv \beta = .5$). Both cases will share the additional symmetric assumption, $\gamma = 1/2$, and the normalization, $u = 1$, $c = 0$; then $\rho_1 = \rho_2 \equiv \rho = \beta$.

Since Proposition 1 doesn't depend on β_i or γ, the no-privacy equilibrium is the same in the two cases: $b_i^{NP} = 1/e$ and $A_i^{NP}(p) = 1 + \ln(p)$ for $p \in [1/e, 1]$, $i = 1, 2$, platform revenue $= 1/e$, each type's consumer surplus $= (1 - 2/e)$.

Under privacy the platform cannot identify consumers' types. Hence advertisers cannot target specific consumer types: each ad falls randomly on the entire market. Similar to the case of no privacy, the equilibrium under privacy is described by a pair of advertising distribution functions, $(A_1^P(\cdot), A_2^P(\cdot))$, but now $A_i^P(p)$ is the total number of ads with product prices less than or equal to p sent by type-i firms per unit measure of *all consumers*. With that definition, the number of *matched* type-i ads—type-i ads received by type-i consumers—with product prices no higher than p is $\gamma_i A_i^P(p)$, i.e., $A_i^P(p)$ type-i ads per unit measure of type-i consumers.

5.1 Case 1: Both consumer types equally (very) picky

When $\beta \leq 0$, consumers are so picky that, even though they get exposed to ads for all products under privacy, they only buy matched products. We call such an equilibrium "a W equilibrium"—"W" indicating *within* product type competition only.

If a type-i firm sends ads at product price p, those ads will make a sale only with type-i consumers. The fraction of such consumers is $(1/2)e^{-A_1^P(p)}$. Hence the marginal profit of a type-i firm sending an additional ad at product price p is simply

$$\pi_i^P\left(p; A_1(\cdot), A_2(\cdot)\right) = \frac{1}{2} e^{-A_i^P(p)}\left(p\right) - b_i$$

Setting this equal to zero and solving for the advertising distribution functions yields

$$A_i^P(p) = \ln\left(\frac{p}{2b_i}\right) \quad \text{for} \quad p \in [2b_i, 1],\ i = 1, 2. \tag{4}$$

For given ad rates, this product-market equilibrium has a higher lower bound in the support than under no privacy—$2b_i$ instead of b_i—and $A_i^P(p) < A_i^{NP}(p)$, $dA_i^P/dp = dA_i^{NP}/dp$ for $p \in [2b_i, 1]$. Prices that are advertised are advertised at the same intensity as under no privacy, but average advertised price is higher and total advertising volume is lower. Hence we get the following result.

Proposition 2. *When ad rates are exogenous to the privacy regime and both consumer types are equally very picky, both platform and consumers are worse off under privacy.* □

When both consumer types are picky, consumer behavior doesn't differ substantively between the privacy and no-privacy product-market equilibria: in both cases, consumers buy matched products only. So, why the worse welfare outcome under privacy? It reflects a fundamental, inevitable effect of privacy: mistargeted ads—type-1 consumers receiving type-2 product ads, and vice-versa. When those mismatched ads reach picky consumers, they are wasted. This reduction in advertising productivity leads firms to pull back on their advertising, especially advertising featuring low prices. The reduction in advertising volume in conjunction with an increase in average advertised price makes both platform and consumers worse off.

Of course, ad rates are not exogenous to the privacy regime: the platform chooses them, so it can certainly optimize them for the privacy regime it is operating under. This means solving the following maximization problem:

$$\max_{b_1, b_2} R^P = b_1 \ln\left(\frac{1}{2b_1}\right) + b_2 \ln\left(\frac{1}{2b_2}\right), \tag{5}$$

where $\ln(1/2b_i) = A_i^P(1)$ is the total ad demand of firms of type i in the privacy product-market equilibrium.[27] The solution is:

$$b_i^P = \frac{1}{2e} = \frac{b_i^{NP}}{2} \qquad \text{for } i = 1, 2. \tag{6}$$

which yields the product-market equilibrium,

$$A_i^P(p) = 1 + \ln p = A_i^{NP}(p) \qquad \text{for } p \in [1/e, 1],\ i = 1, 2.$$

Equilibrium advertising rates under privacy are half of what they are under no privacy, to

[27] There are no γ_i terms in this revenue expression because the advertising functions have been defined per unit measure of the target market, which is the entire market now.

compensate for the fact that half of each firm's advertising is wasted. With that fix, however, everything remains the same for consumers: they see the same prices, advertised at the same intensities. Hence they are just as well off under privacy as under no privacy. The platform, too, is unaffected, because whatever it gives up in ad rates it makes up on ad volume.[28] Hence our welfare result changes.

Proposition 3. *When ad rates are endogenous to the privacy regime and both consumer types are equally very picky, both platform and consumers are equally well off under privacy and no privacy.* $\square$

In Proposition 2, with exogenous ad rates, consumers are strictly worse off under privacy, and in Proposition 3, with endogenous ad rates, they are equally well off under privacy and no privacy. But we are yet to see a situation where they are better off under privacy. This simply reflects a natural predisposition of our model towards consumers favoring no privacy, which makes our subsequent results all the more interesting. We will show next (and in Section 6) that when consumers are flexible in their product preferences, it is possible for them to be better off under privacy.

5.2 Case 2: Both consumer types equally flexible with $\beta = .5$

Since privacy entails ads randomly distributed across the entire market, consumers may be in one of three states: (i) see ads for their preferred product only, (ii) see ads for their non-preferred product only, or (iii) see ads for both preferred and non-preferred products. With flexible product preferences, they may consider purchasing a non-preferred product if it is well-priced. For instance, a type-2 consumer will buy a type-1 product if she is in state (ii) and $p_1 \leq .5$ or she is in state (iii) and $.5 - p_1 \geq \max\{1 - p_2, 0\}$, where p_1 and p_2 are the lowest prices of type-1 and type-2 products seen by this consumer. In other words, mistargeted ads may not be entirely wasted and there is a possibility of cross-product competition.

Each advertiser faces a 2×2 matrix of considerations: along one dimension, whether to advertise high prices that appeal to matched consumers only or to advertise low prices that appeal to all consumers; and along the other dimension, the nature of competition for each consumer type—whether it comes from the same type of firm or from both types of firm. In general, three types of product-market equilibria may emerge under privacy: (1) a W equilibrium where each product type sells only to matched consumers (as in the picky-consumers case just considered), (2) a C equilibrium where each product type sells both to matched and mismatched consumers (here "C" refers to the presence of *cross-product* competition), and (3) an E equilibrium where one type of firm sells to both types of consumers and the other type is *excluded*; see Figure 2 below. As might be guessed, an E equilibrium, which is asymmetric by definition, will not arise in this symmetric model under symmetric ad rates: an excluded firm can surely offer matching consumers better deals than the mismatched firms. A W product-market equilibrium is theoretically possible: if ad rates are relatively high, then firms will only send ads with high product prices to appeal to matched

[28] As we note in Section 6, this property may not hold in all instances of a W equilibrium. When it does, as here, we call it a W1 equilibrium; when it does not, i.e., the platform is worse off under privacy, we call it a W2 equilibrium.

consumers as in Section 5.1. However, as we shall see, in this symmetric model, equilibrium ad rates will not be so high, and only a C product-market equilibrium survives.

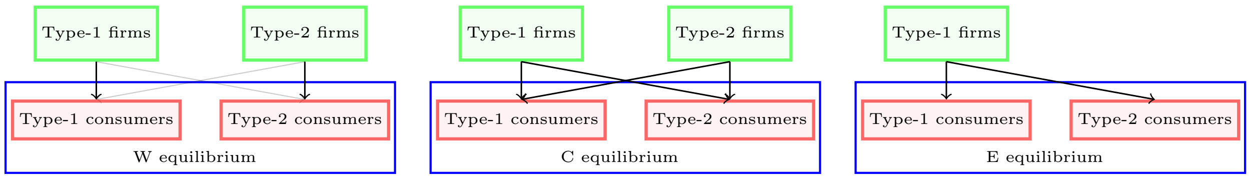

Note: Dark arrows indicate ad flow with product sales; grey arrows indicate ad flow without product sales.

Figure 2: Different types of product-market equilibria under privacy

Suppose ad rates satisfy Lemma 1 and we are in a C situation. Then a firm may sell to matched consumers only if its price is greater than .5 or it may sell to all consumers if its price is less than .5. See Figure 3 below.

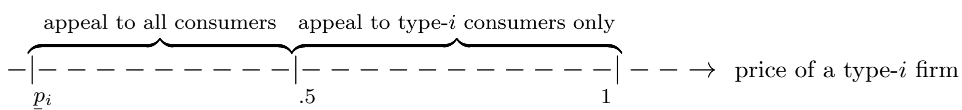

Figure 3: How a type-i firm's market varies by price in a C equilibrium

Therefore, if a type-1 firm sends Z additional units of ads at price $p \in [\underline{p}_1, .5]$, those ads will make a sale not only with type-1 consumers who do not receive product-1 ads with prices lower than p, but also with type-2 consumers who receive neither ads of product 1 with product prices lower than p nor ads of product 2 with product prices lower than $p + .5$. The fraction of such consumers is $.5e^{-A_1^P(p)}\left(1 + e^{-A_2^P(p+.5)}\right)$. Since the fraction of all consumers receiving at least one of those additional ads is $1 - e^{-Z}$, its net profit is

$$\Pi_1^P(Z; p, A_1^P(\cdot), A_2^P(\cdot)) = .5e^{-A_1^P(p)}\left(1 + e^{-A_2^P(p+.5)}\right)\left(1 - e^{-Z}\right)p - b_1 Z.$$

On the other hand, if a type-1 firm sends Z additional units of ads at product price $p \in (.5, 1]$, those ads will make a sale only with type-1 consumers who receive neither product-1 ads with prices lower than p nor product-2 ads with prices lower than $p - .5$. The fraction of such consumers is $.5e^{-A_1^P(p)}e^{-A_2^P(p-.5)}$. Hence, its net profit will be

$$\Pi_1^P(Z; p, A_1^P(\cdot), A_2^P(\cdot)) = .5e^{-A_1^P(p)}e^{-A_2^P(p-.5)}\left(1 - e^{-Z}\right)p - b_1 Z.$$

Setting the marginal profits equal to zero for each type of firm and solving for the advertising distribution functions, we show in Appendix A.6 that the product-market equilibrium for given ad rates (b_1, b_2) is:

$$A_i^P(p) = \begin{cases} \ln p - \ln \underline{p}_i & \text{if} \quad p \in [\underline{p}_i, .5] \\ -2\ln 2 - \ln b_i & \text{if} \quad p \in (.5, \tilde{u}_i) \\ -\ln 2 + \ln p - \ln b_i & \text{if} \quad p \in [\tilde{u}_i, \bar{p}_i], \end{cases} \tag{7}$$

where

$$\bar{p}_1(b_1, b_2) = b_2 - b_1 + \frac{1}{4} + \sqrt{b_1 + \left(b_2 - b_1 + \frac{1}{4}\right)^2}, \tag{8a}$$

$$\bar{p}_2(b_1, b_2) = b_1 - b_2 + \frac{1}{4} + \sqrt{b_2 + \left(b_1 - b_2 + \frac{1}{4}\right)^2}, \tag{8b}$$

and $\underline{p}_i = \bar{p}_{-i} - .5$ for $i = 1, 2$. The lower prices $[\underline{p}_i, .5]$ appeal to both types of consumers while the higher prices $[\tilde{u}_i, \bar{p}_i]$ appeal to matched consumers only; there may be a gap in the support to reflect the discrete drop in sales when transitioning from the entire market to matched consumers only. Note that $\underline{p}_i = \bar{p}_{-i} - .5$ implies that firms make mismatched sales only among consumers who do not receive any matched ads. This reflects the role of Lemma 1.

Notice that if $b_2 > b_1$, then $\bar{p}_1 > \bar{p}_2$, $\underline{p}_1 < \underline{p}_2$, and $A_1^P(\bar{p}_1) > A_2^P(\bar{p}_2)$. Firms favored by the platform in the form of a lower ad rate advertise more, and over a wider price range. When ad rates are symmetric, we can compare these advertising distribution functions with the corresponding functions under no privacy in Lemma 2 and see that total advertising volume will be lower under privacy, with the volume of matched advertising lower still. Sales go down under privacy as a result, reducing consumer welfare.

Proposition 4. *When ad rates are symmetric and exogenous to the privacy regime, and both consumer types are equally flexible with $\beta = .5$, both platform and consumers are worse off under privacy.* □

In other words, going to flexible consumers is not sufficient to make them better off under privacy. Will that change if ad rates were endogenous to the privacy regime? To find out, let us examine the full equilibrium under flexible preferences.

Using (7), the platform's optimization problem is

$$\max_{b_1, b_2} R^P = b_1 \ln\left(\frac{\bar{p}_1(b_1, b_2)}{2b_1}\right) + b_2 \ln\left(\frac{\bar{p}_2(b_1, b_2)}{2b_2}\right), \tag{9}$$

subject to the constraints $A_i^P(1) > 0$, $A_i^P(.5) > 0$ and $A_i^P(1) \geq A_i^P(.5)$. The first constraint ensures that firms actually send ads in equilibrium; the second constraint guarantees that firms advertise low prices to appeal to mismatched consumers, which is a defining characteristic of a C equilibrium. The third constraint verifies the non-decreasing nature of the advertising distribution functions. If the second constraint is violated, firms would only advertise high prices, resulting in a W equilibrium.

The revenue function with flexible consumers is similar to the revenue function with picky consumers (5), with one crucial difference: the numerators of the ad demand functions are functions of (b_1, b_2), instead of 1. It is easy to check from (8) that as b_i goes up, $\bar{p}_i$ goes down and $\bar{p}_{-i}$ goes up; there are cross-product effects. In particular, by increasing the ad rate for a particular type of product, the platform can raise the ad demand curve of the other type of product.[29]

Unconstrained optimization of (9) yields $b_i^P = .172021 < .5b_i^{NP}$ for $i = 1, 2$, which implies the advertising distribution functions

$$A_i^P(p) = 1.45 + \ln p \quad \text{ for } p \in [.234, .5],$$

for $i = 1, 2$; $\underline{p}_i = .234$. These advertising distribution functions satisfy the three constraints, so we do indeed have a C equilibrium (and can eliminate the possibility of a W equilibrium). Notice that there is no gap in the support of these functions: firms sell to matched and mismatched consumers throughout the price range. The platform's ad rates under privacy are lower than its ad rates under no privacy, as in the first example, but there is a critical difference. Because consumers are flexible in their product preferences, mismatched ads are not entirely wasted: they generate sales among consumers who are not reached by matched firms. However, mismatched sales generate lower total surplus than matched sales, and the demand for advertising shrinks. It is no longer enough to simply compensate advertisers for the lower productivity of their ads; the platform has to reduce ad rates below $.5b_i^{NP}$. But even with this price concession, total advertising volume cannot be restored to no-privacy levels: $A_i^P(.5) = .7569 < 1 = A_i^{NP}(1)$. This has the predictable effect of lowering platform revenue below no privacy levels:

$$R^P = .261 < 1/e = R^{NP}.$$

However, what is bad for the platform is good for consumers. Type-1 consumers' (as well as that of a type-2 consumers') aggregate consumer surplus under privacy is

$$\int_{\bar{p}_1 - .5}^{.5} (1-p) d\left(\frac{S_{1\to 1}^P(p)}{\gamma_1}\right) + \int_{\bar{p}_2 - .5}^{.5} (.5-p) d\left(\frac{S_{2\to 1}^P(p)}{\gamma_1}\right) = .395 > 1 - \frac{2}{e},$$

where $S_{1\to 1}^P(p)$ is the (cumulative) matched sales distribution function and $S_{2\to 1}^P(p)$ is the (cumulative) mismatched sales distribution function. Thus we get:

Proposition 5. *When ad rates are endogenous to the privacy regime, and both consumer types are equally flexible with $\beta = .5$, consumers benefit from privacy but the platform is worse off.* □

So what do we learn from these symmetric examples? First, it is hard to find consumer benefits from privacy in competitive markets when ad rates are viewed as exogenous parameters, as much of the privacy literature has done. The two symmetric examples that we examined couldn't be more different: in one the consumers were very picky and in the other they were quite flexible. Yet the two

[29]This comparative statics will play an important role in the next section where, with asymmetric consumer types, the platform will have reason to manipulate relative advertising rates to increase or decrease advertising by particular types of firms.

examples produced the same welfare result under exogenous ad rates. The reason is, no matter how low ad rates are, as long as they are fixed, the platform can't compensate advertisers for the natural reduction in advertising volume that occurs under privacy due to reduced ad productivity. Only by going to the more realistic assumption that ad rates are endogenous—that the platform should be expected to adjust those ad rates when the privacy regime changes—were we able to generate a demand for privacy among consumers. Second, even with endogenous ad rates, consumers will not benefit from privacy unless they are somewhat flexible in their product preferences. For when they are very picky, the platform has the tools to compensate advertisers for their lower ad productivity under privacy while maintaining ad revenue; once compensated, competitive product firms will deliver sufficient advertising for consumers to maintain their no-privacy consumption levels even if some of those ads are wasted. Flexible consumer preferences changes the calculus of ad productivity under privacy: mistargeted ads are not necessarily wasted. However, mismatched sales generate lower total surplus than matched sales, reducing ad demand. As the second example demonstrates, the platform can partially ameliorate the situation by lowering ad rates, but ad volume will suffer. Together, these effects end up taking a toll on platform revenue. Consumers benefit via lower product prices and market expansion.

To summarize, flexibility in consumer preferences, coupled with the platform's ability to adjust ad rates in response to changes in the privacy regime, is crucial for realizing the benefits of privacy for consumers in competitive product markets.

6 The asymmetric case

In this section we will examine the equilibrium under privacy in an asymmetric version of our model that combines the features of the two cases just considered. One consumer type will be very picky, like the consumers in Section 5.1, and the other type will be flexible, like the consumers in Section 5.2. Specifically, we will assume that $\beta_1 \leq c/u < \beta_2$. That is, type-2 consumers are less picky than type-1 consumers, who are picky enough that type-2 firms will not find it cost-effective to serve them. Harking back to our cereal example, a type-1 consumer is one who will consume gluten-free cereals only, whereas a type-2 consumer, while she prefers cereals with gluten, is willing to consider gluten-free products.[30]

Thus type-1 firms may sell to type-2 consumers. However, they are still at a competitive disadvantage doing so: type-2 consumers prefer type-2 products. We will represent this disadvantage via $\rho \equiv \rho_2 = (\beta_2 u - c)/(u - c) \in (0, 1)$. As β_2 increases, ρ approaches one, and the competitive disadvantage of type-1 firms with respect to type-2 consumers decreases.

As discussed in Section 5.2, several types of product-market equilibria are possible under privacy when consumers are flexible in their product preferences:

1. **W equilibrium**. Firms of a given type compete among themselves only and there is no

[30] A useful mnemonic, that might help identify the two types, is that type-*2* consumers find *both* product types acceptable whereas type-*1* consumers find only *one* product type acceptable.

cross-selling. Such an equilibrium can happen in two ways. First, type-2 consumers are so picky—β_2 is so small—that $\underline{p}_1 > \beta_2 u$. (This is what happened in Section 5.1.) Second, type-2 consumers are picky but not picky enough; in this case, $\underline{p}_1 = \beta_2 u$. We call the first kind of equilibrium W1 equilibrium and the second W2 equilibrium.

2. **C equilibrium**. Some type-2 consumers buy type-1 products, i.e., some cross-selling occurs. Since type-1 firms can attract both types of consumers with $p < \beta_2 u$ and only type-1 consumers with $p > \beta_2 u$, they experience a discrete drop in demand at $p = \beta_2 u$. Due to this discontinuity, they may not advertise any price immediately above $\beta_2 u$, but resume advertising at some price $\tilde{u} > \beta_2 u$. That is, the prices advertised by type-1 firms may take the form $[\underline{p}_1, \beta_2 u] \cup [\tilde{u}, u]$. If $\tilde{u} < u$, both intervals are being used, and we call this kind of equilibrium "C1 equilibrium." If $\tilde{u} \geq u$, only the lower interval is being used, and we call it a "C2 equilibrium." (Our equilibrium in Section 5.2 is a C2 equilibrium.)

3. **E equilibrium**. Only one product type is sold to both types of consumers. If this product type is type-1, we call it an E1 equilibrium, otherwise, an E2 equilibrium.

Thus we have potentially six types of product-market equilibrium under privacy in the asymmetric case: W1, W2, C1, C2, E1, and E2. However, not all equilibria are equally likely. In fact, as we show in Proposition 10, an E2 equilibrium will never arise, and E1 and W2 rarely. W1, C1 and C2 are all arguably equally likely *a priori*. Since we have already seen a W1 equilibrium in Section 5.1 and a C2 equilibrium in Section 5.2, here we will focus on the C1 equilibrium; discussion of the other types of privacy equilibria is relegated to the Online Appendix.

6.1 C1 product-market equilibrium under privacy for given ad rates

The following lemma describes the advertising distribution functions in the C1 product-market equilibrium for given ad rates (b_1, b_2) satisfying Lemma 1.

Lemma 4 (C1 product-market equilibrium under privacy). *For ad rates (b_1, b_2) satisfying Lemma 1, the C1 product-market equilibrium under privacy is given by*

$$A_1^P(p) = \begin{cases} \ln(p-c) - \ln(\underline{p}_1 - c) & \text{if} \quad p \in [\underline{p}_1, \beta_2 u] \\ \ln \gamma_1 + \ln(\beta_2 u - c) - \ln b_1 & \text{if} \quad p \in (\beta_2 u, \tilde{u}) \\ \ln \gamma_1 + \ln(p-c) - \ln b_1 & \text{if} \quad p \in [\tilde{u}, u] \end{cases} \tag{10a}$$

$$A_2^P(p) = \ln \gamma_2 + \ln(p-c) - \ln b_2 \quad \text{for} \quad p \in [c + b_2/\gamma_2, \underline{p}_1 + (1-\beta_2)u] \tag{10b}$$

where $\underline{p}_1$ solves $b_1/(\underline{p}_1 - c) = \gamma_1 + b_2/(\underline{p}_1 + (1-\beta_2)u - c)$ and $\tilde{u}$ solves $A_1^P(\beta_2 u) = A_1^P(\tilde{u})$. □

Type-2 firms only sell to type-2 consumers in a C1 equilibrium, hence their advertising distribution functions are similar to the advertising distribution functions in the W1 equilibrium in Section 5.1 (4), the key difference being a lower upper bound in the support. The advertising distribution function for type-1 firms is more complex, reflecting the possibility of cross-selling to type-2 consumers. Still, the two types of firms do not compete "head-to-head": type-1 firms only sell to type-2 consumers when they receive no type-2 offers. To see this, note that the *maximum* price difference between type-2 and type-1 products is $\underline{p}_1 + (1-\beta_2)u - \underline{p}_1 = (1-\beta_2)u$, which is just the difference in reservation prices of type-2 consumers for type-2 versus type-1 products. In other words, in a C1 equilibrium, all type-2 consumers who see ads for both type-1 and type-2 products still buy type-2 products.

Because there is cross-selling in a C1 product-market equilibrium, its comparative statics with respect to ad rates show the possibility of cross-product effects.

Lemma 5 (Comparative statics of the C1 product-market equilibrium under privacy).

1. *Own-product effects: For $i = 1, 2$, as b_i increases, type-i firms advertise less, and their average product price increases.*

2. *Cross-product effects: as b_1 increases, type-2 firms advertise more, and their average product price increases; as b_2 increases, type-1 firms advertise in the same amount, but their average product price decreases.* □

The own-product effects in Lemma 5 are directionally similar to those in the W1 equilibrium of Section 5.1 (and in the no-privacy equilibrium): each type of firm, when facing higher advertising costs, advertises less and increases its average price. While this much is similar between the two types of firms, not everything is the same: from (10) and (20), it is easily verified that type-2 firms' ad demand is more ad rate-sensitive than type-1 firms' ad demand.

To interpret the cross-product effects, note first that type-1 firms advertise both low prices and high prices—the former to appeal to type-2 consumers, the latter to exploit type-1 consumers. Their advertising volume depends only on their own advertising cost b_1, but the distribution of advertised prices depends also on how type-2 firms behave, which is a function of b_2. When b_2 increases, type-2 firms advertise less, which softens the competition for type-2 consumers. As a result, type-1 firms send more ads with low prices and fewer ads with high prices. (This shift in the distribution of prices is characteristic of a C1 equilibrium, distinguishing it from a C2 equilibrium.) Therefore, as b_2 increases, the average price of type-1 products goes down. Type-2 firms appeal to type-2 consumers only, but they face competition from firms of their own type and from type-1 firms. When b_1 increases, type-1 firms advertise less. This gives type-2 firms more room to advertise high prices, increasing their average price.

6.2 Privacy versus no privacy under exogenous ad rates

We offer two propositions comparing the privacy and no-privacy product-market equilibria under exogenous ad rates. The first is specific to the C1 equilibrium, while the second applies to all types of privacy equilibria.

Proposition 6 (C1 versus no-privacy equilibrium under exogenous ad rates). *When ad rates satisfying Lemma 1 are exogenous to the privacy regime, a C1 product-market equilibrium has the following properties vis-a-vis the no-privacy equilibrium:*

1. *Matched advertising volume: lower for both types of products*

2. *Average price: lower for type-1 products, possibly higher for type-2 products*

3. *Sales: lower for type-2 products, possibly higher for type-1 products* □

The first part of Proposition 6 reflects the common effect of privacy on both types of firms: loss in ad productivity from mistargeted ads. However, type-1 firms suffer less than type-2 firms because some of their mistargeted ads—those carrying low prices—are actually productive in generating sales among type-2 consumers. This explains the second and third parts of the proposition.

Proposition 7 (Welfare: privacy versus no privacy under exogenous ad rates). *When ad rates satisfying Lemma 1 are exogenous to the privacy regime, both platform and consumers are better off under no privacy.* □

Why are consumers better off under no privacy when ad rates are exogenous? First, advertising is more productive under no privacy, which implies more advertising, lower average prices, and more consumption. Second, with ad rates fixed exogenously, higher ad productivity does not translate to pricier ads (or, stated differently, lower ad productivity does not imply cheaper ads). With the platform's hands tied in this way, the bounties conferred by greater ad productivity flow unhindered to consumers.

6.3 Optimal ad rates in a C1 privacy equilibrium

Determining optimal ad rates under privacy is significantly more complicated than determining optimal ad rates under no privacy. For one thing, several types of product-market equilibria may arise, and within each type of equilibrium different (b_1, b_2) may be optimal. For example, it is easy to see that the optimal rates in a W1 equilibrium will necessarily be different than the optimal ad rates in a C1 equilibrium: in the former the platform has to contend with own-product effects only whereas in the latter it has to contend with own-product and cross-product effects. Furthermore, an overarching complication is that the choice of ad rates itself may alter the type of product-market equilibrium that prevails. In other words, there is an "outer loop" optimization (selecting a product-market equilibrium) and on an "inner loop" optimization (selecting ad rates within a product-market equilibrium). Said differently, the ad rates that are optimal within a particular type of product-market equilibrium must also induce that type of product-market equilibrium

Since the outer-loop optimization depends on the inner-loop optimization, we will begin with the latter. At the end of this step we will be able to make conditional comparisons of the C1 equilibrium under privacy with the no-privacy equilibrium of the form: "If the platform chooses to induce a C1 equilibrium under privacy, then ..." They will become unconditional when we work out the outer-loop optimization in Proposition 10. That proposition identifies the market conditions under which the platform will indeed induce different types of product-market equilibria.

In a C1 product-market equilibrium, the platform's optimization problem is to choose b_1 and b_2 to maximize ad revenue

$$R^P = b_1 A_1^P(u) + b_2 A_2^P(u),$$

where $A_1^P(u)$ and $A_2^P(u)$ are, respectively, the total ad demands for products 1 and 2 in the product-market equilibrium (from (10)). Define $b_1^* = \gamma_1 b_1^{NP}$ and $b_2^* = \gamma_2 b_2^{NP}$; these ad rates are *welfare-neutral* because they compensate advertisers directly for the lower productivity of their ads under privacy. In Section 5.1, where a W1 privacy equilibrium was induced, these were sufficient to make both platform and consumers indifferent between privacy and no privacy. In a C1 product-market equilibrium, however, they are no longer sufficient.

Lemma 6 (Optimal ad rates in a C1 privacy equilibrium). *In a C1 privacy equilibrium, optimal ad rates satisfy:* $b_1^P > b_1^*$, $b_2^P < b_2^*$.

There are two things about the optimal ad rates in Lemma 6 worth noting. First, that they are not the welfare-neutral rates that were optimal in the W1 equilibrium of the first symmetric case (Section 5.1). Second, that they are not symmetric unidirectional distortions from those rates as in the C2 equilibrium of the second symmetric case (Section 5.2). Instead, the platform is now making asymmetric distortions, clearly tilting the playing field in favor of type-2 firms. Why? As an ad-revenue maximizing firm, the platform is concerned about two things: ad rates and ad volumes. If it relied on only one type of advertiser, this would simply be a matter of asking whether ad demand from that type of firm is price-elastic or not. But in this case, the platform obtains ad revenue from two types of advertisers. Not only own-price elasticities matter, cross-price elasticities matter, too. Lemma 5 says that increasing b_1 will increase ad demand from type-2 firms, but increasing b_2 will not increase ad demand from type-1 firms. Furthermore, as noted above, the slope of type-2 firms' ad demand function is larger than the slope of type-1 firms' ad demand function (with respect to their respective ad rates). In short, both own-price and cross-price elasticities point in the same direction: increase b_1, reduce b_2.

Lemma 6 provides the clue as to why the two types of consumers may differ in their preference for privacy, and, in particular, why type-2 consumers may view it more favorably than type-1 consumers.

6.4 C1 privacy equilibrium versus no-privacy equilibrium under endogenous ad rates

We are now ready to compare the outcomes in a C1 privacy equilibrium under the optimal ad rates identified in Lemma 6 with the outcomes in a no-privacy equilibrium (assuming the platform will indeed want to induce a C1 equilibrium with those ad rates given model parameters). (Similar comparisons for the other types of privacy equilibria are in the Online Appendix.)

Proposition 8 (C1 versus no-privacy equilibrium under endogenous ad rates). *When ad rates are endogenous to the privacy regime and a C1 privacy equilibrium is induced, it has the following properties vis-a-vis the no-privacy equilibrium:*

1. *Matched advertising volume: lower for type-1 products, higher for type-2 products,*

2. *Average price: higher for type-1 products, lower for type-2 products,*

3. *Sales: higher for type-2 products; for type-1 products, higher if and only if $\rho > g_1(\gamma)$,*

4. *Per-capita customer acquisition cost: lower for type-2 products; for type-1 products, higher if and only if $\rho > g_2(\gamma)$,* □

where $g_1(\cdot)$ and $g_2(\cdot)$ are defined in Appendix A.14 and illustrated in Figure 4.

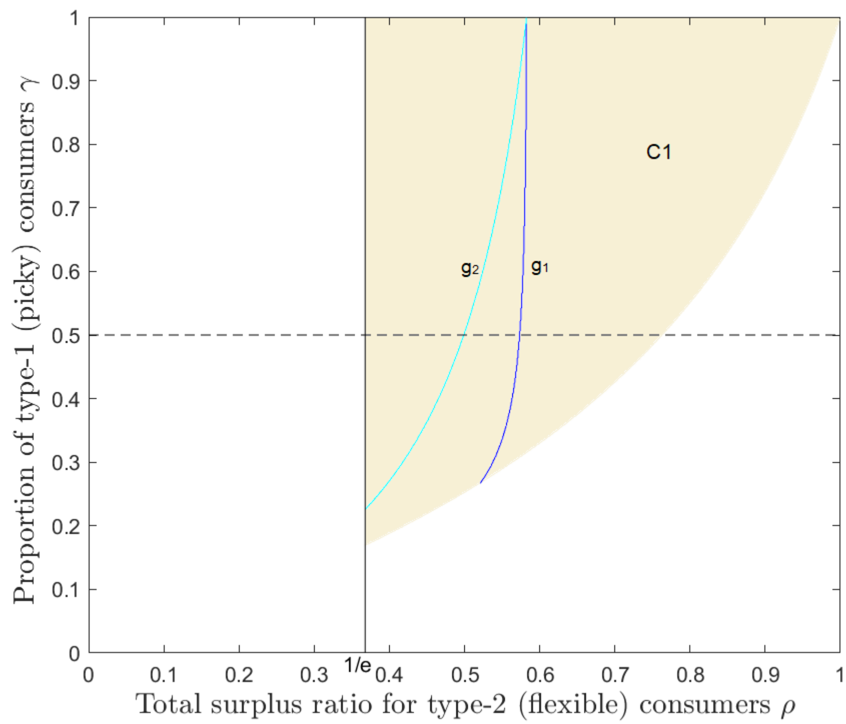

Figure 4: Illustrating the g_1 and g_2 functions of Proposition 8

Proposition 9 (Welfare: C1 versus no-privacy equilibrium under endogenous ad rates). *When ad rates are endogenous to the privacy regime and a C1 equilibrium is induced under privacy:*

1. *The platform will be worse off.*

2. *Type-2 consumers will be better off, but type-1 consumers will be worse off.* □

In a C1 privacy equilibrium, the platform recognizes that it has to be judicious about how it compensates advertisers for their lower advertising productivity. Since type-2 products have narrower appeal than type-1 products, they need to be compensated more. In addition, the boost in demand for type-1 firms from cross-selling means that their ad demand is less sensitive to ad rates. So the platform's optimal solution involves lowering b_2^P below b_2^* and raising b_1^P above b_1^*. The lower ad rate for type-2 firms induces them to advertise more than under no privacy, resulting in greater sales and lower per-capita customer acquisition costs. But for type-1 firms, sales volume and per-capita customer acquisition costs are not entirely driven by higher ad rates; they are tempered by the cross-selling opportunity. For these firms, sales and per-capita customer acquisition costs may be higher when type-2 consumers are sufficiently flexible in their preferences (large ρ).

Type-1 consumers see fewer matched ads and higher average product prices in a C1 privacy equilibrium than in no-privacy equilibrium, making them worse off under privacy. Type-2 consumers, however, are better off—in two ways. First, via lower type-2 product prices: type-2 firms lower their product prices both because of a lower ad rate b_2^P and because of cross-selling pressure from type-1 firms. Second, some type-2 consumers who would not have consumed at all under no privacy, now consume type-1 products.

As far as the platform is concerned, since the platform cannot achieve the same consumer behavior in a C1 privacy equilibrium as under no privacy, its revenue is lower. The distortion in ad rates—$b_1^P > b_1^*$, $b_2^P < b_2^*$—is simply a reflection of the platform's straitened circumstances under privacy.

6.5 Equilibrium selection under privacy

The propositions above have compared a C1 privacy equilibrium with the no-privacy equilibrium assuming that the platform will want to induce a C1 equilibrium under privacy. Now we ask: Under what circumstances will the platform do so? Proposition 10 provides a comprehensive answer to this question. It identifies specific regions of the parameter space under which a W1, W2, C1, C2, or E1 equilibrium will be induced (see Figure 5). (An E2 equilibrium is never induced.) This proposition relies on the property that for given parameter values, for each pair of ad rates (b_1, b_2), the product-market equilibrium is unique. The platform, by choosing ad rates to its advantage, can thus induce a particular product-market equilibrium.

Proposition 10 (Equilibrium selection under privacy). *Under privacy, the platform will induce a:*

1. *W1 equilibrium if* $\rho \leq 1/e$,
2. *W2 equilibrium if* $\rho > 1/e$ *and* $\gamma \in (h_2(\rho), h_1(\rho))$,
3. *C1 equilibrium if* $\rho > 1/e$ *and* $\gamma > \max\{h_1(\rho), h_3(\rho)\}$,
4. *C2 equilibrium if* $\rho > 1/e$ *and* $\gamma < \min\{h_1(\rho), h_2(\rho), h_4(\rho)\}$,
5. *E1 equilibrium if* $\gamma \in (h_4(\rho), h_3(\rho))$,

where $h_1(\cdot), \ldots, h_4(\cdot)$ *are defined in Appendix A.15.* □

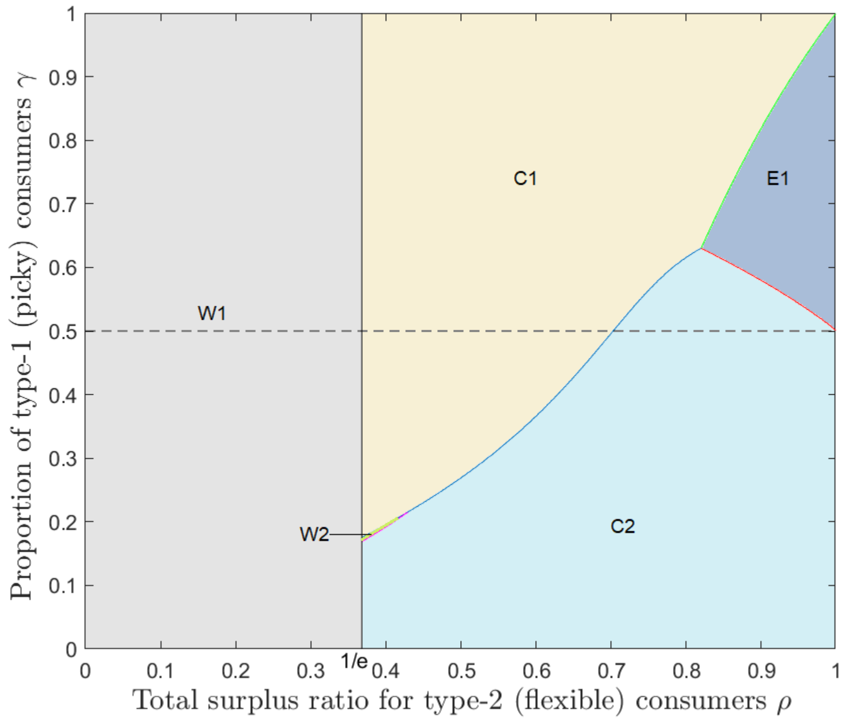

Figure 5: Platform's choice of product-market equilibrium in different regions of the parameter space

The key to understanding the platform's equilibrium selection problem is to think of it as a quest to replicate under privacy the same consumer-level outcomes as those that occur under no privacy, i.e., type-i consumers buying type-i products only. The question is, what costs must be borne in order to do so.

It turns out that when $\rho \leq 1/e$, i.e., when type-1 consumers do not consider type-2 products to be good substitutes, the platform has to bear no costs. (This was the situation in Section 5.1.) A W1 privacy equilibrium with the same consumer-level outcomes as the no-privacy equilibrium can be induced simply by charging the welfare-neutral ad rates. Although these rates are lower than the ad rates under no privacy, the platform doesn't suffer because it makes up in ad volume what it loses in price. In the small adjacent region where a W2 equilibrium is induced, the argument is similar, except that now the platform has to depart from the welfare-neutral ad rates in order to replicate no privacy-like consumer behavior. Now there is a cost.

As ρ increases beyond $1/e$, type-2 consumers are increasingly accepting of type-1 products, and those consumers are now an attractive market expansion opportunity for type-1 firms. It is no longer cost-effective for the platform to fight this tendency; it would rather accommodate cross-selling. This is how we get the C1 and C2 privacy equilibria. The former occurs when the proportion of picky consumers is relatively large; the latter occurs when the proportion of picky consumers is relatively small (as in Section 5.2, where there were none). In inducing a C1 equilibrium, the platform is still showing some ambivalence toward cross-selling; it is trying to straddle the fence between tolerating cross-selling and preventing it. However, the ambivalence disappears when inducing a C2 equilibrium; now the platform fully embraces cross-selling.

For large values of ρ and γ, an E1 equilibrium emerges, reflecting the dominant presence of type-1 consumers and the increased flexibility of type-2 consumers. Now the platform dispenses with type-2 product ads altogether; both types of consumers consume type-1 products only. For the platform, the loss in ad revenue from type-2 products is compensated for by gains in revenue from type-1 products. Finally, an E2 equilibrium is never induced. This is because, from the platform's perspective, such an equilibrium needlessly leaves money on the table from all those type-1 consumers who don't see ads for a product they can buy: it can do better simply by selling type-1 product ads at a price that induces those firms to advertise product prices between $\beta_2 u$ and u (selling to type-1 consumers while avoiding type-2 cannibalization).

7 Discussion

The picture that emerges from our traversal over three versions of our model—two symmetric, one asymmetric—is that two conditions are necessary for consumers to benefit from privacy in competitive product markets: preference flexibility and market power for the ad platform. The latter may seem ironic: we usually associate firm market power with consumer harm! What aligns the interests of the platform and consumers is that, when ads are informative, both consumers and platform share an interest in promoting ad volume. Without ads there will be no consumption, and there will be no ad revenue. Privacy regulations threaten ad volume by reducing ad targetability. Mistargeted ads are either completely wasted or must carry low prices in order to appeal. Competitive product firms react to the loss of ad productivity by cutting back on advertising.

In this setting it is the market power of the platform that comes to the consumer's rescue. For it confers on the platform the ability to respond to reductions in ad-targetability by reducing ad rates. Competitive product markets dutifully pass on those savings to consumers. This, then, provides a path for consumers to benefit from privacy. By contrast, an ad platform playing in a perfectly competitive ad market would simply lack the tools to adapt to changes in privacy conditions: its ad rates would be based on the costs of running an ad platform rather than ad productivity. This was the point of our "exogenous ad rates" propositions—Propositions 2, 4 and 7—which all pointed in one direction: consumers better off giving up their privacy.

A nuance to this line of reasoning is that consumers must also show flexibility in product preferences for the market power of the ad platform to come to their aid under privacy. This was the point of Section 5.1 where we examined a model in which consumers were so picky that they wouldn't consider a non-preferred product. Under those conditions, while the platform must still discount ad rates under privacy, those discounts need not be excessive; they can be calibrated to be just enough to compensate advertisers for their wasted ads while preserving the consumer behavior under no privacy. With no changes in consumer behavior, consumers stay privacy-neutral. Flexible product preferences changes this calculus. Now there is a possibility of cross-selling under privacy:

consumers buying mismatched products when the only ads they see are mismatched ads. Such sales are, by definition, total-surplus-diluting, causing a further reduction in ad volume and necessitating a further reduction in ad rates. Consumers benefit from privacy under these conditions.

Finally, our analysis of the asymmetric model shows that when some consumers are picky and others are flexible, we may see a diversity of views on privacy: flexible consumers favorable, picky consumers unfavorable. Again, the source of this divergence can be traced to the divergent responses of the platform under privacy toward the firms catering to those groups of consumers. The platform favors firms who will "stay in their lane" over firms who will be tempted to cross-sell.

It is important to note that we made some simplifying assumptions along the way to making the above points. In particular, we assumed that consumers don't have intrinsic privacy preferences and don't mind seeing mistargeted ads. In reality, neither is true. Incorporating the first factor will increase the demand for privacy, while incorporating the second will decrease it. Additionally, if consumers have limited attention, mistargeted ads could reduce market efficiency by crowding out matched ads, making privacy more likely to harm consumers.

8 Conclusion

In this paper we examined the impact of privacy-motivated targeting restrictions on consumers in a model where competitive product firms reach consumers by placing informative ads on a monopoly advertising platform. We asked two main questions: Is it possible for consumers who are otherwise privacy-neutral to be better off under privacy? And if so, is it possible for their welfare outcomes to differ? We have answered both questions in the affirmative.

The demand side of our model recognizes a natural heterogeneity that exists in most markets, namely, that not only are consumers different in their product preferences, they also differ in how picky they are: some consumers are more willing to try non-preferred products than others. On the supply side, we emphasize the market power of the advertising platform while deemphasizing the market power of product firms. Our results suggest that both preference flexibility and platform market power are crucial to finding a demand for privacy in competitive product markets. If an ad platform lacks market power to set ad rates—and by implication, is unable to respond to reductions in ad productivity by reducing ad rates—then the greater productivity of ads under no privacy carries the day: better targeted ads induce product firms to advertise more, increasing consumption. However, ad-platform market power by itself is not enough. It must be married to consumer preference flexibility for consumers to be better off under privacy. When consumers have flexible preferences, competitive product firms can't avoid cross-selling to mismatched consumers under privacy, diluting total surplus and inducing them to cut back on advertising; a platform with market power will react by cutting ad rates, which, when passed on to flexible consumers by competitive product markets, end up benefiting those consumers.

In the academic discourse on privacy, it is commonly assumed that the only way to defend privacy regulations is by appealing to the "market thickening" effects of privacy on advertising platforms. In this paper, the advertising platform does not experience any market thickening under privacy: in fact, it is always worse off under privacy. What this paper highlights instead is that when advertising is informative, both consumers and platform share an interest in preserving ad volume, which, when threatened by privacy, can make an ad platform with market power act in some consumers' interests.

Declarations

Funding and Competing Interests

This work was supported by the Social Sciences and Humanities Research Council of Canada via Grant #435–2017–0121 and #435-2020-0073. All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript.

References

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asu Ozdaglar (2022). Too Much Data: Prices and Inefficiencies in Data Markets. *American Economic Journal: Microeconomics* **14**(4), 218–256.

Acquisti, Alessandro, Curtis Taylor, and Liad Wagman (2016). The Economics of Privacy. *Journal of Economic Literature* **54**(2), 442–492.

Amaldoss, Wilfred, Preyas S Desai, and Woochoel Shin (2015). Keyword search advertising and first-page bid estimates: A strategic analysis. *Management Science* **61**(3), 507–519.

Aridor, Guy, Yeon-Koo Che, Brett Hollenbeck, Daniel McCarthy, and Maximilian Kaiser (2024). *Evaluating the Impact of Privacy Regulation on E-Commerce Firms: Evidence from Apple's App Tracking Transparency*. URL: https://papers.ssrn.com/abstract=4717752 (visited on 08/02/2024). Pre-published.

Aridor, Guy, Yeon-Koo Che, and Tobias Salz (2021). "The Effect of Privacy Regulation on the Data Industry: Empirical Evidence from GDPR." *Proceedings of the 22nd ACM Conference on Economics and Computation*. EC '21: The 22nd ACM Conference on Economics and Computation. Budapest Hungary: ACM, pp. 93–94.

Athey, Susan and Joshua S Gans (2010). The Impact of Targeting Technology on Advertising Markets and Media Competition. *American Economic Review* **100**(2), 608–13.

Bagwell, Kyle (2007). The Economic Analysis of Advertising. *Handbook of Industrial Organization* **3**, 1701–1844.

Baye, Michael R and David EM Sappington (2020). Revealing Transactions Data to Third Parties: Implications of Privacy Regimes for Welfare in Online Markets. *Journal of Economics & Management Strategy* **29**(2), 260–275.

Becker, Gary S. (1980). Privacy and Malfeasance: A Comment. *The Journal of Legal Studies* **9**(4), 823–826.

Bergemann, Dirk and Alessandro Bonatti (2011). Targeting in Advertising Markets: Implications for Offline versus Online Media. *The RAND Journal of Economics* **42**(3), 417–443.

Bergemann, Dirk, Alessandro Bonatti, and Tan Gan (2022). The Economics of Social Data. *The RAND Journal of Economics* **53**(2), 263–296.

Butters, Gerard R. (1977). Equilibrium Distributions of Sales and Advertising Prices. *The Review of Economic Studies* **44**(3), 465–491.

Campbell, James, Avi Goldfarb, and Catherine Tucker (2015). Privacy Regulation and Market Structure. *Journal of Economics & Management Strategy* **24**(1), 47–73.

Casadesus-Masanell, Ramon and Andres Hervas-Drane (2015). Competing with Privacy. *Management Science* **61**(1), 229–246.

Chen, Yuxin and Ganesh Iyer (2002). Research Note Consumer Addressability and Customized Pricing. *Marketing Science* **21**(2), 197–208.

Chen, Yuxin, Chakravarthi Narasimhan, and Z John Zhang (2001). Individual Marketing with Imperfect Targetability. *Marketing Science* **20**(1), 23–41.

Choi, Jay Pil, Doh-Shin Jeon, and Byung-Cheol Kim (2019). Privacy and Personal Data Collection with Information Externalities. *Journal of Public Economics* **173**, 113–124.

Conitzer, Vincent, Curtis R Taylor, and Liad Wagman (2012). Hide and Seek: Costly Consumer Privacy in a Market with Repeat Purchases. *Marketing Science* **31**(2), 277–292.

De Corniere, Alexandre and Romain De Nijs (2016). Online Advertising and Privacy. *The RAND Journal of Economics* **47**(1), 48–72.

Esteves, Rosa-Branca and Joana Resende (2016). Competitive Targeted Advertising with Price Discrimination. *Marketing Science* **35**(4), 539–692.

Fudenberg, Drew and Jean Tirole (2000). Customer Poaching and Brand Switching. *RAND Journal of Economics* **31**(4), 634–657.

Fudenberg, Drew and J Miguel Villas-Boas (2012). In the Digital Economy. *The Oxford Handbook of the Digital Economy*, 254–272.

Galeotti, Andrea and José Luis Moraga-González (2008). Segmentation, Advertising and Prices. *International Journal of Industrial Organization* **26**(5), 1106–1119.

Garratt, Rodney J and Maarten RC Van Oordt (2021). Privacy as a Public Good: A Case for Electronic Cash. *Journal of Political Economy* **129**(7), 000–000.

Goldfarb, Avi (2014). What is Different About Online Advertising? *Review of Industrial Organization* **44**(2), 115–129.

Grossman, Gene M. and Carl Shapiro (1984). Informative Advertising with Differentiated Products. *The Review of Economic Studies* **51**(1), 63–81.

Ichihashi, Shota (2020). Online Privacy and Information Disclosure by Consumers. *American Economic Review* **110**(2), 569–595.

Iyer, Ganesh, David Soberman, and J. Miguel Villas-Boas (2005). The Targeting of Advertising. *Marketing Science* **24**(3), 461–476.

Johnson, Garrett A. (2024). “Economic Research on Privacy Regulation: Lessons from the GDPR and Beyond.” Goldfarb, Avi and Catherine Tucker. *The Economics of Privacy.* National Bureau of Economic Research Conference Report. Chicago: University of Chicago Press, pp. 97–126.

Johnson, Justin P (2013). Targeted Advertising and Advertising Avoidance. *The RAND Journal of Economics* **44**(1), 128–144.

Ke, T. Tony and K. Sudhir (2023). Privacy Rights and Data Security: GDPR and Personal Data Markets. *Management Science* **69**(8), 4389–4412.

Kraft, Lennart, Bernd Skiera, and Tim Koschella (2023). *Economic Impact of Opt-in versus Opt-out Requirements for Personal Data Usage: The Case of Apple's App Tracking Transparency (ATT)*. URL: https://papers.ssrn.com/abstract=4598472 (visited on 05/21/2024). Pre-published.

Levin, Jonathan and Paul Milgrom (2010). Online Advertising: Heterogeneity and Conflation in Market Design. *The American Economic Review* **100**(2), 603–607.

Lin, Tesary (2022). Valuing Intrinsic and Instrumental Preferences for Privacy. *Marketing Science* **41**(4), 663–681.

McKinnon, John D. (2022). Online Privacy Protections Gain Traction With Lawmakers, Tech Industry. *The Wall Street Journal*.

Montes, Rodrigo, Wilfried Sand-Zantman, and Tommaso Valletti (2019). The Value of Personal Information in Online Markets with Endogenous Privacy. *Management Science* **65**(3), 1342–1362.

Moorthy, Sridhar and Shervin S. Tehrani (2023). Targeting Advertising Spending and Price on the Hotelling Line. *Marketing Science*.

Narasimhan, Chakravarthi (1988). Competitive Promotional Strategies. *The Journal of Business* **61**(4), 427–449.

O'Neil, Cathy (2021). Comment: How to Make Google and Facebook Care About Privacy. *HeraldNet*.

Rafieian, Omid and Hema Yoganarasimhan (2021). Targeting and privacy in mobile advertising. *Marketing Science* **40**(2), 193–218.

Sayedi, Amin (2018). Real-time bidding in online display advertising. *Marketing Science* **37**(4), 553–568.

Shin, Jiwoong and Jungju Yu (2021). Targeted Advertising and Consumer Inference. *Marketing Science* **40**(5), 900–922.

Stegeman, Mark (1991). Advertising in Competitive Markets. *The American Economic Review* **81**(1), 210–223.

Sun, Tianshu, Zhe Yuan, Chunxiao Li, Kaifu Zhang, and Jun Xu (2024). The value of personal data in internet commerce: A high-stakes field experiment on data regulation policy. *Management Science* **70**(4), 2645–2660.

Thisse, Jacques-Francois and Xavier Vives (1988). On The Strategic Choice of Spatial Price Policy. *The American Economic Review* **78**(1), 122–137.

Tirole, Jean (1988). *The Theory of Industrial Organization*. Cambridge, Massachusetts: The MIT Press.

Villas-Boas, J Miguel (1999). Dynamic competition with Customer Recognition. *The Rand Journal of Economics*, 604–631.

Vranica, Suzanne, Patience Haggin, and Salvador Rodriguez (2022). Inside Facebook's $10 Billion Breakup With Advertisers. *The Wall Street Journal*.

Wernerfelt, Nils, Anna Tuchman, Bradley T Shapiro, and Robert Moakler (2022). "Estimating the Value of Offsite Data to Advertisers on Meta." Available at SSRN: https://ssrn.com/abstract=4198438.

A Appendix

A.1 Proof of Lemma 1

Suppose, to the contrary, that $(u-c)/b_2 < (\beta_2 u - c)/b_1$. Then type-2 firms are at a competitive disadvantage in their own natural market, as well as in type-1 firms' natural market. Therefore, if type-1 firms are making non-positive profits in equilibrium (as per condition (ii)), type-2 firms must be incurring a loss should they advertise. Hence they will not advertise. But then the platform can lower b_2 so that $(u-c)/b_2 = (\beta_2 u - c)/b_1$ without affecting its ad revenue. Therefore, it is without loss to require $b_{-i}/b_i \geq \rho_i$. ∎

A.2 Proof of Lemma 2

To see that this is indeed the unique equilibrium, note that by condition (ii), for all $p \in [c+b_i, u]$, $\pi_i^{NP}\left(p; A_1^{NP}, A_2^{NP}\right) = e^{-A_i^{NP}(p)}(p-c) - b_i \leq 0$. Hence, $A_i^{NP}(p) \geq \ln(p-c) - \ln b_i$ for all i. It suffices to argue that this inequality must hold with equality in equilibrium. Suppose by contradiction that $A_i^{NP}(\hat{p}) > \ln(\hat{p}-c) - \ln b_i$ for some $\hat{p} \in [c+b_i, u]$. Since $A_i^{NP}(p)$ is continuous, there must exist an interval containing $\hat{p}$ such that $A_i^{NP}(p) > \ln(p-c) - \ln b_i$ for all p in that interval. Let (p', p'') denote the maximal such interval. Then, for all $p \in (p', p'')$, $\pi_i^{NP}\left(p; A_1^{NP}, A_2^{NP}\right) < 0$; hence by equilibrium condition (iii) and continuity, $A_i^{NP}(p') = A_i^{NP}(p'')$. Since $A_i^{NP}(c+b_i) = 0 = \ln(c+b_i-c) - \ln b_i$ and interval (p', p'') is maximal, we must have $A_i^{NP}(p') = A_i^{NP}(p'') = \ln(p'-c) - \ln b_i$. This observation, together with $\pi_i^{NP}\left(p'; A_1^{NP}, A_2^{NP}\right) = 0$, implies that

$$\begin{aligned}
\pi_i^{NP}\left(p''; A_1^{NP}, A_2^{NP}\right) &= e^{-A_i^{NP}(p'')}\left(p''-c\right) - b_i \\
&= e^{-A_i^{NP}(p')}\left(p''-p'\right) \\
&> 0,
\end{aligned}$$

which contradicts condition (ii). ∎

A.3 Proof of Lemma 3

Recall from Lemma 2 that $A_i^{NP}(p) = \ln[(p-c)/b_i]$. Hence, the total amount of advertising by type-i firms is $A_i^{NP}(u) = \ln[(u-c)/b_i]$, which is decreasing in b_i. The density of the advertising distribution function, $\partial A_i^{NP}(p)/\partial p = 1/(p-c)$, while decreasing in product price, is independent of ad rate. Therefore, an increase in ad rate b_i, by raising the lower bound $\underline{p}_i$ of the advertising distribution function's support, will increase the average advertised product price for product i. ∎

A.4 Proof of Proposition 2

Under privacy, the marginal profit of type-i firms by sending one additional ad at price $p \in [\underline{p}_i, u]$ is

$$\pi_i^P\left(p; A_1^P(\cdot), A_2^P(\cdot)\right) = \gamma_i e^{-A_i^P(p)}(p-c) - b_i^P.$$

The zero-profit condition $\pi_i^P\left(p; A_1^P, A_2^P\right) = 0$ for all $p \in [\underline{p}_i, u]$ with $\underline{p}_1 > \beta_2 u$ implies that

$$A_i^P\left(p\right) = \ln \gamma_i + \ln\left(p - c\right) - \ln b_i^P.$$

It follows from $A_i^P(\underline{p}_i) = 0$ that

$$\underline{p}_i = c + b_i^P/\gamma_i.$$

Define the sales function $S_{j\to i}^P(p)$ as the total sales of product j to type-i consumers at product prices less than or equal to p.Matched ad volume is $\gamma_1 A_1^\lambda(u) + \gamma_2 A_1^\lambda(\beta_2 u)$ for type-1 product and $\gamma_2 A_2^\lambda(u)$ for type-2 product.Total sales of type-i product is given by $S_{i\to 1}^P(u) + S_{i\to 2}^P(u)$.

Under **no privacy**, total sales of type-i product is

$$S_i^{NP}(u) = \gamma_i\left(1 - e^{-A_i^{NP}(u)}\right) = \gamma_i\left(1 - \frac{b_i}{u-c}\right).$$

Under **privacy** and **W equilibrium**, both types of consumers only buy their preferred product and $S_{2\to 1}^P(p) = S_{1\to 2}^P(p) = 0$ for all equilibrium product prices p. Total sales of product i is

$$S_{i\to 1}^P(u) + S_{i\to 2}^P(u) = \gamma_i\left(1 - e^{-A_i^P(u)}\right) = \gamma_i\left(1 - \frac{b_i}{\gamma_i(u-c)}\right) < S_i^{NP}(u).$$

Therefore, compared to no privacy environment, the market sizes of both products in both markets would shrink, while the average product prices go up.

For consumer welfare, consider first the welfare of type-1 consumers. The gain of a type-i consumers from having privacy is

$$\int_{c+b_i/\gamma_i}^{u}(u-p)d\frac{S_{i\to i}^P(p)}{\gamma_i} - \int_{c+b_i}^{u}(u-p)d\frac{S_i^{NP}(p)}{\gamma_i} = \frac{i}{\gamma_i}\left(\int_{c+b_i/\gamma_i}^{u} S_{i\to i}^P\left(p\right)dp - \int_{c+b_i}^{u} S_i^{NP}\left(p\right)dp\right) < 0.$$

where the equality follows from integration by parts, and the inequality follows from that $S_i^{NP}\left(p\right) > S_{i\to i}^P\left(p\right)$ for all $p \in [c + b_i/\gamma_i, u]$. It follows immediately that type-i consumers benefit from no privacy. ∎

A.5 Proof of Proposition 3

The total demand for ads of product i is

$$A_i^P\left(u\right) = \ln \gamma_i + \ln\left(u - c\right) - \ln b_i^P.$$

Given ad rates b_1^P and b_2^P, the platform's ad revenue $b_1^P A_1^P\left(u\right) + b_2^P A_2^P\left(u\right)$ can be rewritten as

$$R_W^P = b_1^P\left(\ln \gamma_1 + \ln\left(u-c\right) - \ln b_1^P\right) + b_2^P\left(\ln \gamma_2 + \ln\left(u-c\right) - \ln b_2^P\right) \tag{11}$$

The optimization problem of the monopoly platform is to choose b_1^P and b_2^P to maximize R_W^P subject to $A_1^P(u) \geq 0$, $A_2^P(u) \geq 0$ and

$$\underline{p}_1 \geq \beta_2 u. \tag{12}$$

Since consumers are very picky, constraint (12) is not binding, and the solutions to the optimization problem without any constraints are

$$b_i^P = b_i^* \equiv \frac{\gamma_i (u-c)}{e}.$$

The equilibrium advertising distribution functions are

$$A_i^P(p) = \ln(p-c) - \ln(u-c) + 1.$$

It is easy to verify that constraints $A_1^P(u) \geq 0$ and $A_2^P(u) \geq 0$ are satisfied. Constraint (12) is equivalent to

$$c + \frac{u-c}{e} \geq \beta_2 u \Longleftrightarrow \rho \leq \frac{1}{e}.$$

The equilibrium sales functions are

$$S_{i \to i}^P(p) = \gamma_i - \frac{b_i^P}{p-c} = \gamma_i - \frac{\gamma_i}{e}\frac{u-c}{p-c}.$$

The total number of matched ads of type-i product in advertising market is

$$\gamma_i A_i^P(u) = \gamma_i \left(\ln \gamma_i + \ln(u-c) - \ln b_i^P\right) = \gamma_i,$$

and the market size of type-i product in product market is

$$S_{i \to i}^P(u) = \gamma_i \left(1 - \frac{b_i^P}{\gamma_i (u-c)}\right) = \gamma_i \frac{e-1}{e}.$$

The consumer acquisition cost per consumer for type-i firms is

$$\frac{b_i^P A_i^P(u)}{S_i^P(u)} = \frac{u-c}{e-1}.$$

Hence, by comparing to what we obtain in Section 4.3, we conclude that the two privacy modes generate identical equilibrium advertising distribution functions, sales functions, market sizes, and customer acquisition costs for every product. It simply follows that consumers and the platform are also equally well off. ∎

A.6 Proof of Proposition 4

To show that that $\underline{p}_i = \bar{p}_{-i} - .5$ for $i = 1, 2$, note that for type-2 firms $\pi_2^P(p') = \frac{1}{2} e^{-A_2^P(p')} e^{-A_1^P(p'-.5)} p' - b_2$ for $p' \in (\tilde{u}_2, \bar{p}_2]$ (these are the marginal profits from an additional type-2 ad with price p' reaching type-2 consumers who not receive a lower price from type-2 firms nor a price lower than $p' - .5$ from type-1 firms). If a type-1 firm sends an additional ad with product price $p' - .5 \in (\underline{p}_1, .5)$, its marginal expected profit (from an additional type-1 ad reaching type-1 consumers whose only offers from type-1 firms are higher-priced and type-2 consumers whose offers from type-1 firms are higher-priced and whose offers from type-2 firms are higher-priced than p') is

$$\begin{aligned}\pi_1^P(p' - .5) &= \left[.5 + .5e^{-A_2^P(p')}\right] e^{-A_1^P(p'-.5)}(p' - .5) - b_1 \\ &= \left(\frac{1 + e^{-A_2^P(p')}}{e^{-A_2^P(p')}}\right)\left(\frac{p' - .5}{p'}\right)\left(\pi_2^P(p') + b_2\right) - b_1,\end{aligned}$$

where both ratios in the expression are increasing in p'. Therefore, there cannot be an interval $[a_1, a_2] \subset (\tilde{u}_2, \bar{p}_2)$ such that $\pi_1^P(p - .5) = 0$ and $\pi_2^P(p) = 0$ for all $p \in [a_1, a_2]$. It follows that $\bar{p}_2 - .5 \leq \underline{p}_1$. However, $\bar{p}_2 - .5 < \underline{p}_1$ is impossible because then type-2 firms can earn strictly positive profit by sending ads with product price $\bar{p}_2 + \epsilon$ for some small $\epsilon > 0$. Hence, we must have $\underline{p}_1 = \bar{p}_2 - .5$. The argument for $\underline{p}_2 = \bar{p}_1 - .5$ is identical.

It follows that the marginal profit of a type-i firm sending an additional ad at product price p is

$$\pi_i^P(p; A_1(\cdot), A_2(\cdot)) = \begin{cases} .5\left(1 + e^{-A_{-i}(\bar{p}_{-i})}\right) e^{-A_i(p)}p - b_i & \text{if} \quad p \in [\bar{p}_{-i} - .5, .5] \\ .5e^{-A_i(p)}p - b_i & \text{if} \quad p \in (.5, \bar{p}_i] \end{cases}$$

From the zero marginal-profit conditions for type-1 and type-2 firms, we can get Equation (8a) and Equation (8b). Hence, under privacy, the total demand for type-i ads is

$$A_i^P(\bar{p}_i(b_1, b_2)) = \ln\left(\frac{\bar{p}_i(b_1, b_2)}{2 + b_i}\right), \quad i = 1, 2. \tag{13}$$

Under equal ad rates $b_1 = b_2$, it is easy to see that $A_i^P(p) \leq A_i^{NP}(p)$ for all p. Then:

$$S_{1\to 1}^P(p) = \gamma_1\left(1 - e^{-A_1^P(p)}\right) = \frac{1}{2}\left(1 - \frac{\bar{p}_1 - .5}{p}\right),$$

and

$$S_{2\to 1}^P(p) = \gamma_1 e^{-A_1^P(.5)}\left(1 - e^{-A_2^P(p)}\right) = \frac{1}{2}\left(\frac{\bar{p}_1 - .5}{.5}\right)\left(1 - \frac{\bar{p}_2 - .5}{p}\right),$$

are such that $S_{1\to 1}^P(p) \leq S_1^{NP}(p)$, and $S_{2\to 1}^P(p - .5) \leq S_1^{NP}(p)$. Hence the aggregate consumer surplus of type-1 consumers (as well as that of a type-2 consumers) under privacy is

$$\begin{aligned} &\int_{\bar{p}_1 - .5}^{\bar{p}_1}(1 - p)d\left(\frac{S_{1\to 1}^P(p)}{\gamma_1}\right) + \int_{\bar{p}_1 - .5}^{.5}(.5 - p)d\left(\frac{S_{2\to 1}^P(p)}{\gamma_1}\right) \\ < \ &\int_{\bar{p}_1 - .5}^{\bar{p}_1}(1 - p)d\left(\frac{S_1^{NP}(p)}{\gamma_1}\right) + \int_{\bar{p}_1 - .5}^{.5}(.5 - p)d\left(\frac{S_1^{NP}(p + .5)}{\gamma_1}\right) \\ < \ &\int_{b_1}^{1}(1 - p)d\left(\frac{S_1^{NP}(p)}{\gamma_1}\right), \end{aligned}$$

the consumer surplus under no privacy. ■

A.7 Proof of Proposition 5

Consider the relaxed problem without any constraints. The first-order conditions are

$$\ln \bar{p}_i - \ln 2 - \ln b_i^P + \frac{b_i^P \frac{\partial \bar{p}_i}{\partial b_i^P}}{\bar{p}_i} + \frac{b_{-i}^P \frac{\partial \bar{p}_{-i}}{\partial b_i^P}}{\bar{p}_{-i}} - 1 = 0, \tag{14}$$

where

$$\frac{\partial \bar{p}_i}{\partial b_i^P} = \frac{-\bar{p}_i + .5}{\sqrt{\left(b_{-i}^P - b_i^P + \frac{1}{4}\right)^2 + b_i^P}}, \quad \frac{\partial \bar{p}_i}{\partial b_{-i}^P} = \frac{\bar{p}_i}{\sqrt{\left(b_{-i}^P - b_i^P + \frac{1}{4}\right)^2 + b_i^P}}.$$

Applying these partial derivatives to (13), it is easy to see that each product type's ad volume increases as its ad rate decreases. Thus, given the opportunity to reset ad rates under privacy, the platform will compensate for the reduction in ad demand by reducing ad rates.

Imposing symmetry on the first-order conditions and solving for b_1^P we get

$$b_1^P = \left(\bar{p}_1 - \frac{1}{4}\right)^2 - \frac{1}{16} = \bar{p}_1 \left(\bar{p}_1 - \frac{1}{2}\right)$$

where $\bar{p}_1$ solves (14), or equivalently

$$-\ln(2\bar{p}_1 - 1) = \frac{\bar{p}_1}{2\bar{p}_1 - \frac{1}{2}}$$

Thus, $\bar{p}_1 = 0.734274$, $b_1^P = 0.172021$, and $A_1^P = 0.758117$. We can now pin down $\tilde{u}_1$ by appealing to the indifference condition of type-1 firms:

$$\ln\left(\frac{1}{2} + \frac{1}{2}e^{-A_2^P(\bar{p}_2)}\right) + \ln\frac{1}{2} - \ln b_1^P = \ln\frac{1}{2} + \ln\tilde{u}_1 - \ln b_1^P$$

which yields

$$\tilde{u}_1 = \frac{1}{2} + \frac{1}{2}e^{-A_2^P(\bar{p}_2)} = \frac{1}{2} + \frac{b_1^P}{\bar{p}_1} = \bar{p}_1$$

In short, the support of a type-i firm in a C product-market equilibrium will be $[\bar{p}_{-i} - .5, .5]$. ■

A.8 Product-market equilibrium under privacy

To find the product-market equilibria under privacy, we first characterize the support of advertising distribution functions for different privacy equilibria.

Lemma A.1 (Support of advertising distribution functions). *The supports of the advertising distribution functions take the following forms in the different privacy equilibria:*

- ***W* equilibrium:** *The support of* $A_2^P(p)$ *is* $[\underline{p}_2, u]$ *and the support of* $A_1^P(p)$ *is* $[\underline{p}_1, u]$ *with* $\underline{p}_1 \geq \beta_2 u$.

- ***C* equilibrium:** *The support of* $A_2^P(p)$ *is* $[\underline{p}_2, \underline{p}_1 + (1-\beta_2)u]$. *The support of* $A_1^P(p)$ *is either* $[\underline{p}_1, \beta_2 u] \cup [\tilde{u}, u]$ *(C1) or* $[\underline{p}_1, \beta_2 u]$ *(C2) with* $\underline{p}_1 < \beta_2 u$ *and* $\beta_2 u < \tilde{u} < u$.

- ***E* equilibrium:** *In E1 equilibrium, the support of* $A_2^P(p)$ *is empty and the support of* $A_1^P(p)$ *is either* $[\underline{p}_1, u]$ *(same as W),* $[\underline{p}_1, \beta_2 u] \cup [\tilde{u}, u]$ *(same as C1) or* $[\underline{p}_1, \beta_2 u]$ *(same as C2). In E2 equilibrium, the support of* $A_1^P(p)$ *is empty and the support of* $A_2^P(p)$ *is* $[\underline{p}_2, u]$.

Proof. See Lemmas 1-5 in the Online Appendix. ■

Next we present the proof of Lemma 4 for the advertising distribution functions under W1 and C1 equilibria.

A.9 Proof of Lemma 4

C equilibrium features cross-product competition for type-2 consumers. Such an equilibrium must satisfy $A_1^P(u) \geq 0$, $A_2^P(u) \geq 0$ and constraints (15) and (16):

$$A_1^P(\beta_2 u) \geq 0, \tag{15}$$

and

$$A_1^P(u) \geq A_1^P(\beta_2 u). \tag{16}$$

The zero-profit condition $\pi_2^P\left(p; A_1^P, A_2^P\right) = 0$ for all $p \in [\underline{p}_2, \underline{p}_1 + (1-\beta_2)u]$ implies that

$$\gamma_2 e^{-A_2^P(p)}(p-c) - b_2^P = 0 \tag{17}$$

which implies that

$$A_2^P(p) = \ln \gamma_2 + \ln(p-c) - \ln b_2^P.$$

There are two sub-cases: either $\bar{p}_1 = u$ (C1) or $\bar{p}_1 = \beta_2 u$ (C2, discussed in Online Appendix).

In the C1 equilibrium, $\bar{p}_1 = u$ and the set of product prices that type-1 firms advertise in equilibrium takes the form of $[\underline{p}_1, \beta_2 u] \cup [\tilde{u}, u]$ with $\underline{p}_1 < \beta_2 u < \tilde{u} < u$. The marginal profit of type-1 firms by sending one additional ad at product price $p \leq \beta_2 u$ is

$$\pi_1^P\left(p; A_1^P, A_2^P\right) = \left(\gamma_1 + \gamma_2 e^{-A_2^P\left(\underline{p}_1 + (1-\beta_2)u\right)}\right) e^{-A_1^P(p)}(p-c) - b_1^P.$$

For any equilibrium product price $p \leq \beta_2 u$ advertised by type-1 firms, zero-profit condition implies that

$$\left[\gamma_1 + \gamma_2 e^{-A_2^P\left(\underline{p}_1 + (1-\beta_2)u\right)}\right] e^{-A_1^P(p)}(p-c) - b_1^P = 0. \tag{18}$$

By setting $p = \underline{p}_1 + (1-\beta_2)u$ in (17) and $p = \underline{p}_1$ in (18) and canceling out $e^{-A_2^P\left(\underline{p}_1 + (1-\beta_2)u\right)}$, we obtain $\underline{p}_1$ implicitly as a solution to

$$\frac{b_1^P}{\underline{p}_1 - c} = \gamma_1 + \frac{b_2^P}{\underline{p}_1 + (1-\beta_2)u - c}. \tag{19}$$

We can solve $\underline{p}_1$ explicitly as

$$\underline{p}_1 = c + \frac{b_1^P - b_2^P - \gamma_1(1-\beta_2)u + \sqrt{\Delta}}{2\gamma_1} \tag{20}$$

where

$$\Delta = \left(b_1^P - b_2^P - \gamma_1(1-\beta_2)u\right)^2 + 4\gamma_1 b_1^P (1-\beta_2)u. \tag{21}$$

The marginal profit of type-1 firms by sending one additional ad at product price $p \in [\tilde{u}, u]$ is

$$\pi_1^P\left(p; A_1^P, A_2^P\right) = \gamma_1 e^{-A_1^P(p)}(p-c) - b_1^P.$$

We can now derive $A_i^P(p)$ for all equilibrium product prices p. For $p \in [c + b_2^p/\gamma_2, \underline{p}_1 + (1-\beta_2)u]$, the zero profit condition for type-2 firms is

$$\gamma_2 e^{-A_2^P(p)}(p-c) - b_2^P = 0 \implies A_2^P(p) = \ln\gamma_2 + \ln(p-c) - \ln b_2^P.$$

For $p \in [\underline{p}_1, \beta_2 u]$, we obtain

$$A_1^P(p) = \ln\left(\gamma_1 + \gamma_2 e^{-A_2^P\left(\underline{p}_1 + (1-\beta_2)u\right)}\right) + \ln(p-c) - \ln b_1^P = \ln(p-c) - \ln(\underline{p}_1 - c),$$

where the last equality follows from (19). For $p \in [\tilde{u}, u]$, the zero-profit condition for type-1 firms is

$$\gamma_1 e^{-A_1^P(p)}(p-c) - b_1^P = 0 \implies A_1^P(p) = \ln\gamma_1 + \ln(p-c) - \ln b_1^P,$$

where $\tilde{u}$ is given by type-1 firms' indifference condition between advertising $\beta_2 u$ and $\tilde{u}$:

$$\ln(\beta_2 u - c) - \ln(\underline{p}_1 - c) = \ln\gamma_1 + \ln(\tilde{u} - c) - \ln b_1^P. \tag{22}$$

■

A.10 Proof of Lemma 5

The "own-product" effects of ad rate change simply follow the same logic as under no privacy. For "cross-product" effects in C1 equilibrium, there is an indirect effect through $\underline{p}_1$. It is easy to check from (20) that $\underline{p}_1$ is increasing in b_1 and decreasing in b_2. To verify the "cross-product" effects in the lemma, consider first the effect of an increase in b_2 on type-1 product. The total amount of type-1 ads is

$$A_1^P(u) = \ln \gamma_1 + \ln(u-c) - \ln b_1,$$

which is unaffected by b_2. An increase in b_2 would decrease $\underline{p}_1$ and increase $\tilde{u}$. Once again, it has no effect on the density $1/(p-c)$ at each advertised product price. More ads with lowest prices are sent while fewer ads with higher prices are sent. Therefore, the average price of type-1 product will decrease.

If b_1 increases, $\underline{p}_1$ and the upper bound of the support of $A_2^P(u)$ will increase. Therefore, type-2 firms will send more ads with higher prices, and the average type-2 product price will go up. ■

A.11 Proof of Proposition 6

The matched ad volume of type-2 product is

$$\gamma_2 A_2^P(u) = \gamma_2 \left(\ln \gamma_2 + \ln\left(\underline{p}_1 + (1-\beta_2)u - c\right) - \ln b_2\right),$$

which is smaller than under no privacy. Total sales of type-2 product is

$$S_{2\to1}^P(u) + S_{2\to2}^P(u) = \gamma_2\left(1 - e^{-A_2^P(u)}\right) = \gamma_2\left(1 - \frac{b_2}{\gamma_2\left(\underline{p}_1 + (1-\beta_2)u - c\right)}\right),$$

which is also smaller than under no privacy.

For type-1 product, matched ad volume of type-1 product is

$$\gamma_1 A_1^P(u) = \gamma_1 \left(\ln \gamma_1 + \ln(u-c) - \ln b_1\right).$$

which decreases under privacy. Total sales of type-1 product is

$$S_{1\to1}^P(u) + S_{1\to2}^P(u) = \gamma_1\left(1 - \frac{b_1}{\gamma_1(u-c)}\right) + \frac{b_2}{\underline{p}_1 + (1-\beta_2)u - c}\frac{\beta_2 u - \underline{p}_1}{\beta_2 u - c}$$

which is smaller than under no privacy if and only if

$$-\frac{\gamma_2 b_1}{u-c} + \frac{b_2}{\underline{p}_1 + (1-\beta_2)u - c}\frac{\beta_2 u - \underline{p}_1}{\beta_2 u - c} < 0$$

The left hand side is decreasing in b_1 and increasing in b_2. When we decrease b_1 or increase b_2, the equilibrium continues to be C1 equilibrium until $\underline{p}_1 = c + b_2/\gamma_2$ where the equilibrium turns to E1 equilibrium. Therefore, as shown for E1 equilibrium in Online Appendix, total sales may go up under privacy for small b_1.

Average advertised product price of type-2 product can be expressed as

$$\frac{\int_{c+b_2/\gamma_2}^{\underline{p}_1+(1-\beta_2)u} p dA_2^P(p)}{A_2^P(\underline{p}_1+(1-\beta_2)u)} = c + \frac{\underline{p}_1+(1-\beta_2)u-c-b_2/\gamma_2}{\ln\gamma_2+\ln\left(\underline{p}_1+(1-\beta_2)u-c\right)-\ln b_2}$$

The average product price rises compared to no privacy case if

$$\frac{\ln\left(\underline{p}_1+(1-\beta_2)u-c\right)-\ln\left(b_2/\gamma_2\right)}{\left(\underline{p}_1+(1-\beta_2)u-c\right)-b_2/\gamma_2} < \frac{\ln(u-c)-\ln b_2}{(u-c)-b_2}$$

As the changes in b_1, b_2 are bound by the conditions that $c+b_1<\underline{p}_1<\beta_2 u$, when b_1 is large such that $\underline{p}_1$ is close to $\beta_2 u$, the average advertised product price is higher than under no privacy, as analyzed in W equilibrium. When b_1 is small and γ_2 is large, the opposite of the inequality will be true. The average advertised product price of type-1 product is

$$\frac{\int_{\underline{p}_1}^{u} p dA_1^P(p)}{A_1^P(u)} = c + \frac{\beta_2 u-\underline{p}_1+u-\tilde{u}}{\ln\gamma_1+\ln(u-c)-\ln b_1} < c + \frac{\beta_2 u-c-b_1+u-\tilde{u}}{\ln\gamma_1+\ln(u-c)-\ln b_1} < c + \frac{u-c-b_1}{\ln(u-c)-\ln b_1}$$

where $\tilde{u}$ is given by $\gamma_1(\tilde{u}-c)=\beta_2 u-c$ and the second inequality is derived in the discussion of E1 equilibrium in the Online Appendix. Therefore, the average advertised price is lower than under no privacy.

For consumer welfare, type-1 consumers have been shown to be worse off in the proof of Proposition 2. The change in consumer welfare for a type-2 consumer is

$$\int_{c+b_2}^{u}(u-p)d\frac{S_{2\to 2}^{NP}(p)}{\gamma_2} - \left(\int_{c+b_2/\gamma_2}^{\underline{p}_1+(1-\beta_2)u}(u-p)d\frac{S_{2\to 2}^{P}(p)}{\gamma_2} + \int_{\underline{p}_1}^{\beta_2 u}(\beta_2 u-p)d\frac{S_{1\to 2}^{P}(p)}{\gamma_2}\right) = \frac{b_2}{\gamma_2}g(b_1,b_2)$$

where

$$G(b_1,b_2)=\gamma_1-\gamma_2\ln\frac{u-c}{b_2}+\ln\frac{\gamma_2(\underline{p}_1+(1-\beta_2)u-c)}{b_2}+\frac{\underline{p}_1-c}{\underline{p}_1+(1-\beta_2)u-c}\ln\frac{\beta_2 u-c}{\underline{p}_1-c}.$$

It can be verified that $G(b_1,b_2)$ is increasing in b_1 and decreasing in b_2. Therefore, $G(b_1,b_2)>0$ if $G(b_1',b_2')>0$ for some $b_1'\le b_1$ and $b_2'\ge b_2$. However, the decrease of b_1 and increase of b_2 cannot be arbitrary. For C equilibrium to exist, we need $A_1^P(u)\ge 0$, $A_2^P(u)\ge 0$ and $\underline{p}_1\le\beta_2 u$, $c+b_2/\gamma_2\le\underline{p}_1+(1-\beta_2)u$. When b_1 continues to decrease, the first three constraints still hold, but the last one may be violated. When we decrease b_1 to the point where $c+b_2/\gamma_2=\underline{p}_1+(1-\beta_2)u$, this will be E1 equilibrium where no type-1 product ads are sent. Therefore, in order to show that type-2 consumers prefer no privacy in C equilibrium, i.e. $G(b_1,b_2)>0$, it suffices to show that type-2 consumers prefer no privacy in any E1 equilibrium, which we do in the Online Appendix. ■

A.12 Proof of Proposition 7

Type-1 consumers have been shown to be worse off in all types of privacy equilibria compared to no-privacy equilibrium in the proof of Proposition 2. For type-2 consumers, it remains to be proven for W2 and E equilibria. (The proof for C2 equilibrium is included in the proof of Proposition 6.)

For W2 and E2 equilibria, it follows that $A_2^{NP}(p) > A_2^P(p)$ and hence $S_{2\to 2}^{NP}(p) > S_{2\to 2}^P(p)$ for all $p \in [c + b_2/\gamma_2, u]$. Type-2 consumers would prefer no privacy as only type-2 products will be considered.

In E1 equilibrium, the advertising distribution function $A_1^P(p)$ takes the form as in either W or C equilibrium. For W form, all the advertised prices of product 1 are not accepted by type-2 consumers and hence type-2 consumers are obviously worse off. For C form, the change in consumer welfare for a type-2 consumer is

$$\int_{c+b_2}^{u} (u-p)d\frac{S_{2\to 2}^{NP}(p)}{\gamma_2} - \int_{c+b_1}^{\beta_2 u} (\beta_2 u - p)d\frac{S_{1\to 2}^{P}(p)}{\gamma_2} = b_1 + (1-\beta_2)u - b_2 - b_2 \ln\frac{u-c}{b_2} + b_1 \ln\frac{\beta_2 u - c}{b_1}$$

which is increasing in b_1. Therefore, if $b_1/b_2 \geq \rho$, we have

$$\begin{aligned} & \int_{c+b_2}^{u} (u-p)d\frac{S_{2\to 2}^{NP}(p)}{\gamma_2} - \int_{c+b_1}^{\beta_2 u} (\beta_2 u - p)d\frac{S_{1\to 2}^{P}(p)}{\gamma_2} \\ \geq\ & \frac{\beta_2 u - c}{u-c} b_2 + (1-\beta_2)u - b_2 - b_2 \ln\frac{u-c}{b_2} + \frac{\beta_2 u - c}{u-c} b_2 \ln\frac{u-c}{b_2} > 0 \end{aligned}$$

where the last inequality follows because, for $x < 1$, $x(1-\ln x)$ is increasing in x and hence $x(1-\ln x) < 1$. Hence, type-2 consumers are worse in E1 equilibrium. ■

A.13 Proof of Lemma 6

The platform's optimization problem is to choose b_1^P and b_2^P to maximize its ad revenue

$$R_{C1}^P = b_1^P \left[\ln\gamma_1 + \ln(u-c) - \ln b_1^P\right] + b_2^P \left[\ln\gamma_2 + \ln\left(\underline{p}_1 - c + (1-\beta_2)u\right) - \ln b_2^P\right]$$

subject to $A_1^P(u) \geq 0$, $A_2^P(u) \geq 0$ and constraints (15) and (16). Consider the relaxed problem without any constraints. The first-order conditions are

$$\ln\gamma_1 + \ln(u-c) - \ln b_1^P + \frac{b_2^P \frac{\partial \underline{p}_1}{\partial b_1^P}}{\underline{p}_1 - c + (1-\beta_2)u} - 1 = 0, \tag{23}$$

$$\ln\gamma_2 + \ln\left(\underline{p}_1 + (1-\beta_2)u - c\right) - \ln b_2^P + \frac{b_2^P \frac{\partial \underline{p}_1}{\partial b_2^P}}{\underline{p}_1 - c + (1-\beta_2)u} - 1 = 0. \tag{24}$$

It follows that

$$\begin{aligned}\ln \gamma_1 + \ln(u-c) - \ln b_1^P - 1 &< 0,\\ \ln \gamma_2 + \ln\left(\underline{p}_1 + (1-\beta_2)u - c\right) - \ln b_2^P - 1 &> 0,\end{aligned}$$

and hence

$$b_1^P > b_1^* \text{ and } b_2^P < \frac{\gamma_2(\underline{p}_1 - c + (1-\beta_2)u)}{e}. \tag{25}$$

Since $\underline{p}_1 \leq \beta_2 u$, the second part of (25) immediately implies that $b_2^P < b_2^*$. Therefore, b_2^P is lower than the welfare-neutral ad rate b_2^* while b_1^P is higher than the welfare-neutral ad rate b_1^*. ■

A.14 Proofs of Proposition 8 and Proposition 9

We need to derive conditions under which the solution to the relaxed problem satisfies all dropped constraints. Constraint (16) of $A_1^P(u) \geq A_1^P(\beta_2 u)$ can be rewritten as

$$\ln \gamma_1 + \ln(u-c) - \ln b_1^P \geq \ln\left(\gamma_1 + \frac{b_2^P}{\underline{p}_1 - c + (1-\beta_2)u}\right) + \ln(\beta_2 u - c) - \ln b_1^P$$

which is equivalent to

$$b_2^P \leq \frac{\gamma_1(1-\rho)}{\rho}\left(\underline{p}_1 - c + (1-\rho)(u-c)\right). \tag{26}$$

Constraint (15) of $A_1^P(\beta_2 u) \geq 0$ is equivalent to

$$\underline{p}_1 \leq \beta_2 u, \tag{27}$$

which is also equivalent to

$$\left(\gamma_1 + \frac{b_2^P}{\underline{p}_1 - c + (1-\beta_2)u}\right)(\beta_2 u - c) \geq b_1^P \Longleftrightarrow \gamma_1 \geq \frac{b_1^P}{\rho(u-c)} - \frac{b_2^P}{\underline{p}_1 - c + (1-\rho)(u-c)}. \tag{28}$$

Note that the constraints $A_i^P(u) \geq 0$ are equivalent to

$$\begin{aligned}A_1^P(u) &\geq 0 \Longleftrightarrow b_1^P \leq \gamma_1(u-c),\\ A_2^P(u) &\geq 0 \Longleftrightarrow b_2^P \leq \gamma_2\left(\underline{p}_1 - c + (1-\beta_2)u\right).\end{aligned}$$

The first inequality is implied by (26) and (28), and the second is implied by (25). Therefore, if γ_1 and ρ satisfy conditions (26) and (27), then the ad rates given by (23) and (24) are indeed optimal for the platform.

Next we would like to argue that in equilibrium with cross-product competition and $\bar{p}_1 = u$, type-2 consumers are better off while type-1 consumers are worse off with privacy. Consider first the welfare of type-2 consumers. Under privacy, they receive product 2 price offers distributed according to $A_2^P(p)$ with $p \in [\underline{p}_2, \underline{p}_1 + (1-\beta_2)u]$ as well as product 1 price offers distributed according to $A_1^P(p)$

with $p \in [\underline{p}_1, \beta_2 u]$. A product 1 offer at price $p-(1-\beta_2)u$ generates the same surplus for type-2 consumers as a product 2 offer at price p. Under no privacy, they receive only product 2 offers with prices distributed according to $A_2^{NP}(p)$ for $p \in [b_2^*/\gamma_2 + c, u]$. Note that $\underline{p}_2 = b_2^P/\gamma_2 + c < b_2^*/\gamma_2 + c$, so a sufficient condition for type-2 consumers to prefer privacy is

$$\begin{cases} A_2^P(p) > A_2^{NP}(p) & \text{for} \quad p \in [\underline{p}_2, \underline{p}_1 + (1-\beta_2)u] \\ \frac{\partial A_1^P(p-(1-\beta_2)u)}{\partial p} \geq \frac{\partial A_2^{NP}(p)}{\partial p} & \text{for} \quad p \in [\underline{p}_1 + (1-\beta_2)u, u] \end{cases}$$

The first part of the condition says that type-2 consumers receive more ads from type-2 firms under privacy with product prices no higher than p for every p in the range of prices advertised by type-2 firms in equilibrium under privacy. The second part of the condition implies that type-2 consumers receive more ads from type-1 firms under privacy which generate the same surplus as ads from type-2 firms of product prices p for every p that is advertised by type-2 firms under no privacy but not under privacy. The first part is implied by $b_2^P < b_2^*$ and the second part is always true. Therefore, type-2 consumers are better off with privacy.

For type-1 consumers who buy only product 1, the sales distribution function $S_{1\to 1}^P(p)$ is

$$S_{1\to 1}^P(p) = \gamma_1\left(1 - e^{-A_1^P(p)}\right) = \begin{cases} \gamma_1\left(1 - \frac{\underline{p}_1 - c}{p-c}\right) & \text{if} \quad p \in [\underline{p}_1, \beta_2 u] \\ \gamma_1\left(1 - \frac{\underline{p}_1 - c}{\rho(u-c)}\right) & \text{if} \quad p \in (\beta_2 u, \tilde{u}] \\ \gamma_1\left(1 - \frac{b_1^P}{\gamma_1(p-c)}\right) & \text{if} \quad p \in (\tilde{u}, u] \end{cases}$$

The consumer surplus for a type-1 consumer is given by

$$\int_{\underline{p}_1}^{u}(u-p)d\left(\frac{S_{1\to 1}^P(p)}{\gamma_1}\right) = (u-c) - \frac{b_1^P}{\gamma_1} + \frac{b_1^P}{\gamma_1}\ln\frac{b_1^P}{\gamma_1(u-c)} - \left(\frac{b_1^P}{\gamma_1} - (\underline{p}_1 - c)\right)\ln\frac{\underline{p}_1 - c}{\rho(u-c)}$$

Hence, type-1 consumers would prefer no privacy if and only if

$$\frac{b_1^P}{\gamma_1(u-c)} - \frac{b_1^P}{\gamma_1(u-c)}\ln\frac{b_1^P}{\gamma_1(u-c)} + \left(\frac{b_1^P}{\gamma_1(u-c)} - \frac{\underline{p}_1 - c}{u-c}\right)\ln\frac{\underline{p}_1 - c}{\rho(u-c)} \geq \frac{2}{e}.$$

It can be numerically verified that this inequality holds under conditions (26) and (27).

The total number of matched ads of type-1 product in advertising market is

$$\gamma_1 A_1^P(u) = \gamma_1\left(\ln\gamma_1 + \ln(u-c) - \ln b_1^P\right) = \gamma_1 \ln\frac{b_1^*}{b_1^P} < \gamma_1,$$

and the total number of matched ads of type-2 product in advertising market is

$$\gamma_2 A_2^P(u) = \gamma_2\left(\ln\gamma_2 + \ln\left(\underline{p}_1 + (1-\beta_2)u - c\right) - \ln b_2^P\right) > \gamma_2.$$

The market size of type-1 product in product market is

$$S^P_{1\to1}(u)+S^P_{1\to2}(\beta_2 u)=\frac{b^P_1}{\underline{p}_1-c}-\frac{b^P_1}{u-c}-\frac{b^P_1}{\beta_2 u-c}+\gamma_1\frac{\underline{p}_1-c}{\beta_2 u-c}.$$

Plugging in b^P_1 and b^P_2, it can be numerically shown that it is smaller than the total number of sales under no privacy $\gamma_1(e-1)/e$ if and only if $\rho<g_1(\gamma_1)$ for some weakly increasing function g_1. The market size of type-2 product in product market is

$$S^P_{2\to2}(\beta_2 u)=\gamma_2\left(1-\frac{b^P_2}{\gamma_2(\underline{p}_1+(1-\beta_2)u-c)}\right)>\gamma_2\frac{e-1}{e}.$$

The consumer acquisition cost per consumer for type-1 firms is

$$\frac{b^P_1 A^P_1(u)}{S^P_{1\to1}(u)+S^P_{1\to2}(\beta_2 u)} = \frac{b^P_1\ln\frac{\gamma_1(u-c)}{b^P_1}}{\frac{b_1}{\underline{p}_1-c}-\frac{b_1}{u-c}-\frac{b_1}{\beta_2 u-c}+\gamma_1\frac{\underline{p}_1-c}{\beta_2 u-c}},$$

which can be shown numerically smaller than the consumer acquisition cost under no privacy $(u-c)/(e-1)$ if and only if $\rho<g_2(\gamma_1)$ for some weakly increasing function g_2. The consumer acquisition cost per consumer for type-2 firms is

$$\frac{b^P_2 A^P_2(\underline{p}_1+(1-\beta_2)u-c)}{S^P_{2\to2}(\underline{p}_1+(1-\beta_2)u-c)}=(\underline{p}_1+(1-\beta_2)u-c)\frac{\ln\frac{\gamma_2(\underline{p}_1+(1-\beta_2)u-c)}{b^P_2}}{\frac{\gamma_2(\underline{p}_1+(1-\beta_2)u-c)}{b^P_2}-1}<\frac{\underline{p}_1+(1-\beta_2)u-c}{e-1}<\frac{u-c}{e-1}$$

from condition (25). ■

A.15 Proof of Proposition 10

First, W1 is the optimal choice for the platform whenever possible (i.e., $\rho\le 1/e$), as it achieves the profit under no privacy which is an upper-bound for the platform's profit under full privacy. Second, C1 is better than W2 as it solves a relaxed problem of W2 by dropping the binding equality constraint of (15). As for the rest, they are induced in regions of the parameter space demarcated by four functions, $h_1(\rho)$, …, $h_4(\rho)$; see Figures 6–10. These functions are implicitly defined by (numerically) comparing platform's ad revenue across different equilibria: C1 equilibrium generates higher ad revenue than C2 if $\gamma>h_1(\rho)$ (Figure 6); W2 generates higher ad revenue than C2 if $\gamma>h_2(\rho)$ (Figure 7); C1 generates higher ad revenue than E1 if $\gamma>h_3(\rho)$ (Figure 8); E1 generates higher ad revenue than C2 if $\gamma>h_4(\rho)$ (Figure 9); and finally the comparison between W2 and E1 is shown in Figure 10.

Proposition 10 follows directly from these comparisons. ■

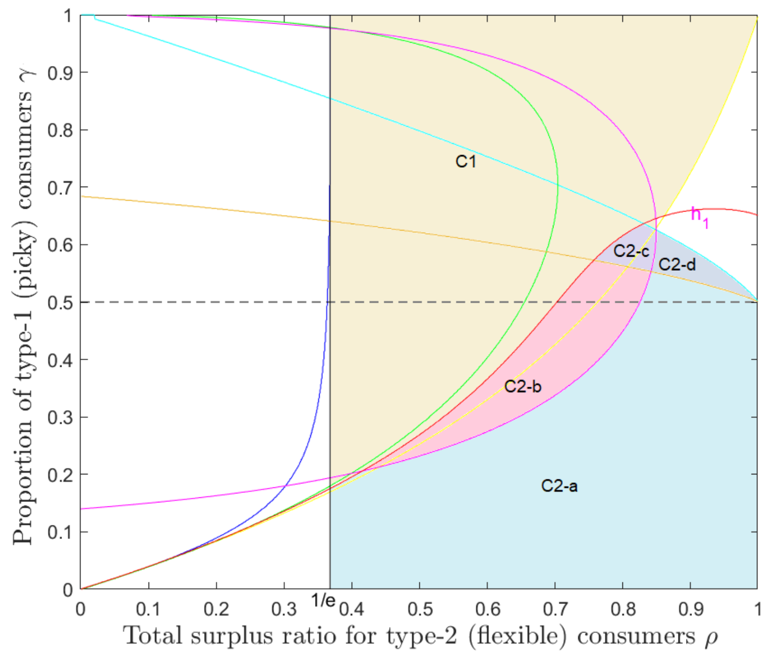

Figure 6: Platform's choice between C1 and C2 equilibrium

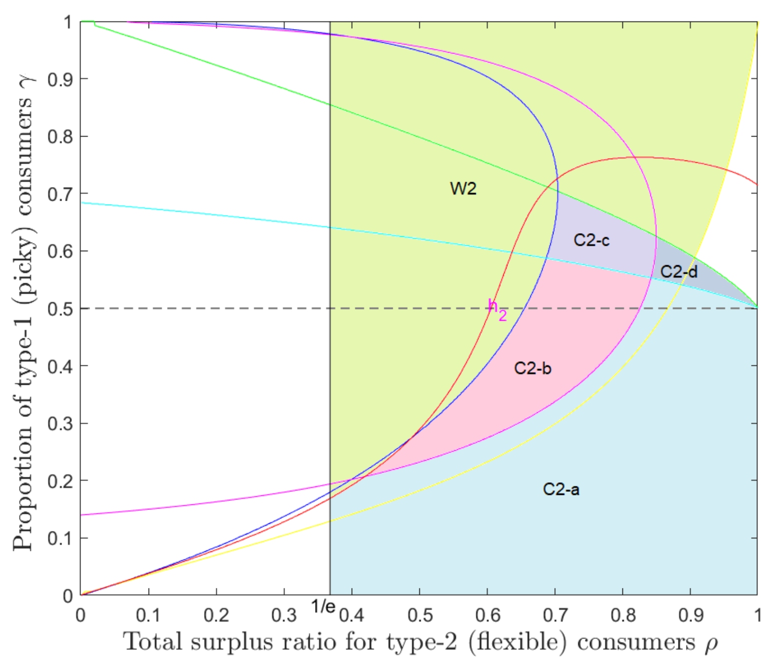

Figure 7: Platform's choice between C2 and W2 equilibrium

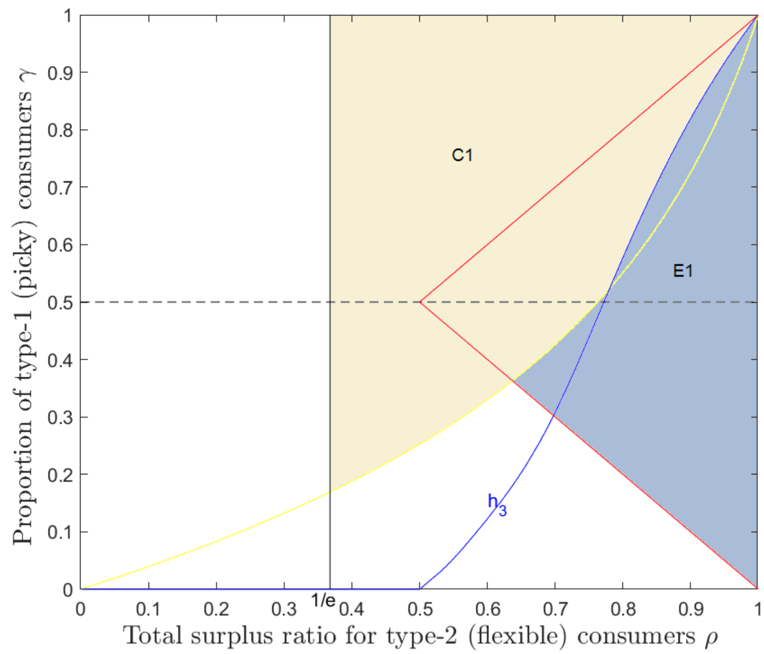

Figure 8: Platform's choice between C1 and E1 equilibrium

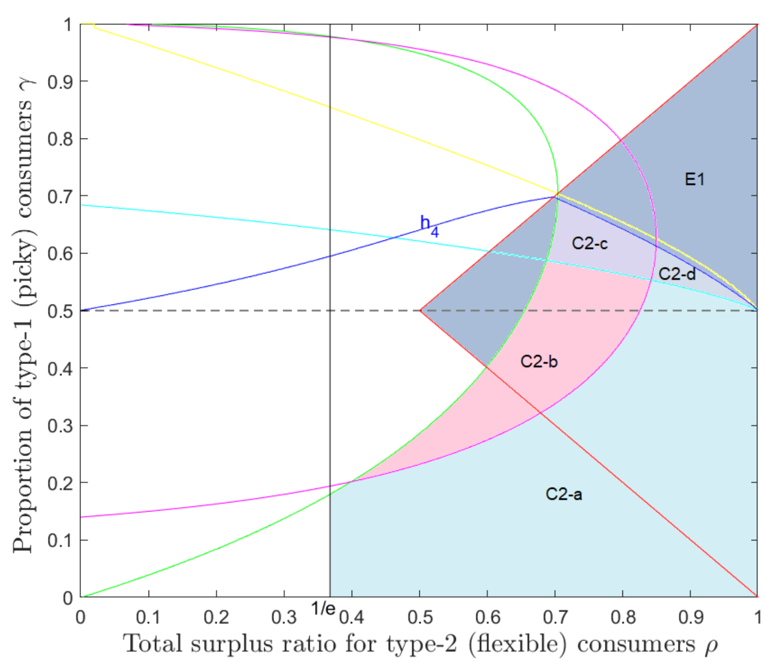

Figure 9: Platform's choice between C2 and E1 equilibrium

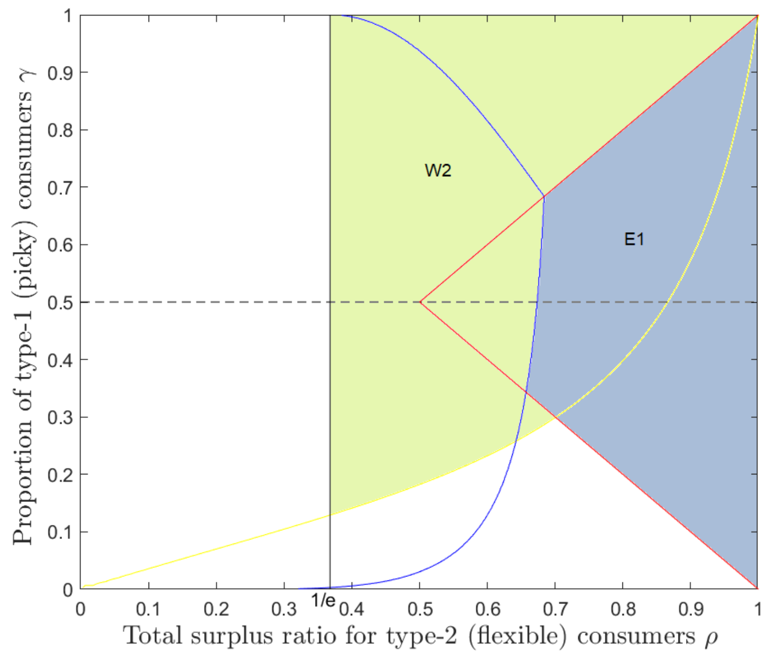

Figure 10: Platform's choice between W2 and E1 equilibrium