Online Business Models, Digital Ads, and User Welfare*

Daron Acemoglu† Daniel Huttenlocher‡ Asuman Ozdaglar§ James Siderius¶

September 9, 2024

Abstract

We present a model where social media platforms offer plans that intermix entertaining content with digital advertising ("ads"). Consumers derive utility from entertainment and learn about their valuation for a product from ads. While some consumers are fully rational, others naïvely perceive digital ads as more informative than they actually are. We characterize the profit-maximizing business model of the platform and show that welfare is lower when the platform monetizes through advertising instead of subscription both for naïfs (because they are targeted by intense digital advertising, which makes them over-optimistic about product quality and over-purchase the product) and for sophisticates (because the demand from naïfs increases the firm's price). This negative welfare effect is intensified when the platform can offer mixed business models that separate the naïve and sophisticated consumers into different plans. Our results are robust to firm-level and platform-level competition, largely because digital ads soften competition between both firms and platforms. We also show how digital ad taxes can improve welfare.

JEL Classification: D83, D43, L13.

Keywords: digital advertising, welfare, online business models, digital ad taxation

*We are grateful to numerous participants at the INFORMS 2022 Annual Conference, UT Dallas Jindal School of Business, Santa Clara Leavey School of Business, Johns Hopkins Carey School of Business, University of Michigan Ross School of Business, Tuck School of Business at Dartmouth College, USC Marshall School of Business, Emory University, Chicago Booth School of Business, MIT Sloan School of Management, Columbia NYC Media Seminar, and AsuFest 2023 for their suggestions and feedback.

†Massachusetts Institute of Technology, NBER, and CEPR, daron@mit.edu

‡Massachusetts Institute of Technology, huttenlocher@mit.edu

§Massachusetts Institute of Technology, asuman@mit.edu

¶Tuck School of Business at Dartmouth College, james.siderius@tuck.dartmouth.edu

1 Introduction

Online platforms have become a dominant forum for entertainment and social interaction, overtaking traditional sources such as TV and print media (see Cinelli et al. (2020) and Sherman and Waterman (2016)). The average adult now spends over three hours a day on social media,[1] and even more time on streaming services such as Netflix, YouTube, and Hulu (Budzinski et al. (2021), Richter (2019), and Twenge et al. (2019)). There is an active debate on the costs and benefits of social media engagement, with many of the potential negative effects being related to targeted digital advertisement and its impacts on user beliefs and behavior (e.g., see Marwick and Lewis (2017), Allcott et al. (2020), and Allcott et al. (2022)).

The majority of social media platforms generate their revenues from digital advertising ("digital ads" for short).[2] Unlike traditional advertising, where the same product recommendation is broadcast to a large audience, digital advertising allows ads to be tailored and targeted to different users. While this may make such ads more informative about relevant products and services, it also opens the way to greater manipulation and enticement for other users (Bennett and Gordon (2020), Deng and Mela (2018), and De Jans et al. (2019)). Despite growing concerns on these topics, there is currently no framework in which digital ads have both informative and manipulative roles. There are also only very few analyses of online business models.

In this paper, we develop a parsimonious model where an online media platform offers both entertainment and digital ads and acts as a two-sided marketplace bringing together users that can learn from informative ads and a firm interested in advertising to users. The platform also has a choice of monetizing its services via advertising, subscription fees, or both. Digital ads are informative about the (user-specific) quality of the firm's product. Greater ad intensity is costly because it interrupts the entertaining content for users, but also beneficial because it provides informative signals about the product.

A distinguishing feature of our model is that there are both *sophisticated* users who have the correct model about the relationship between good signals from digital ads and product quality, and *naïve* users who have a misspecified model of the relationship between signals from digital ads and product quality. Specifically, naïve users underestimate the likelihood of "false positives", whereby a product that is low-quality for them may nonetheless generate a positive signal via ads. This may be because of their inherent naïveté or because they underestimate the degree to which the targeting of digital ads may exaggerate the appeal of the underlying product to them. This misspecification on the part of naïve agents opens the way to manipulation—it is profitable for the firm and hence for the platform to send more ads to naïve users to boost their demand for the product.

For expositional clarity, we first restrict the platform to two simple business models: a free-of-charge advertising-based plan or an ad-free plan with a subscription fee. Our first main result (Proposition 2) is a striking one: provided that naïve users do not have a model very close to that

[1] See https://www.forbes.com/sites/petersuciu/2021/06/24/americans-spent-more-than-1300-hours-on-social-media/ and https://whatagraph.com/blog/articles/how-much-time-do-people-spend-on-social-media.

[2] For example, digital ads made up 98% of Facebook's revenue from 2017-2019 (see https://www.nasdaq.com/articles/what-facebooks-revenue-breakdown-2019-03-28-0) and about 85% of YouTube's revenue in 2020 (even with a premium ad-free subscription plan offered, see https://spendmenot.com/blog/youtube-revenue-statistics/).

of sophisticated users, the unique equilibrium involve an advertising-based plan designed for naïfs, which consequently fully segments the market, with sophisticates not using the platform. This is because the platform chooses a very high level of ad intensity which is unattractive for sophisticates and fully extracts all surplus from naïfs. As a result, the *ex post* welfare of naïve users is even less than the benchmark without the platform, because they end up *de facto* manipulated into over-consuming the product. Interestingly, sophisticated users also have lower welfare than the scenario without the platform, because when digital ads inflate the demand from naïve agents, the firm prefers to charge a higher price, which sophisticates also pay.[3] An important comparative static is that the separating equilibrium with digital ads targeting naïve agents is more likely when the likelihood of false positive signals for naïfs is higher—implying that digital ads emerge precisely when they are more misleading. Intuitively, it is this misleading aspect of digital advertisement that makes them expand the demand for the product and create profitable monetization opportunities for the platform.

The results from this simplified model with just two business plans generalize directly when the platform can offer multiple entertainment plans that intermix advertising and subscription. The equilibrium typically separates naïfs and sophisticates, but this time the platform can also extract surplus from sophisticates through a subscription fee. Welfare effects are similar to our baseline result, though with additional nuanced implications. Specifically, the separating equilibrium emerges when false positives from the ad technology are not too rare (otherwise the equilibrium is pooling). Equilibrium user welfare is then shown to be decreasing in the likelihood of false positives (which implies greater distortion in naïve agents' assessment of product quality) and is also decreasing in the overall informativeness of ads (because more informative ads make a separating equilibrium, which is worse for naïfs and features higher prices, more likely).

The insights from our analysis with a single platform and a single advertising firm also generalize to an environment with multiple platforms and multiple firms. The fundamental reason for this is that digital ads soften the competition between both firms and platforms. For example, two firms with identical products that would otherwise engage in Bertrand competition and earn zero profits, now gain market power, because users who obtain different information from the ads they see have different (derived) willingness to pay for the products of the two firms. We show that digital ads soften firm-level competition, because they enable endogenous differentiation of products based on the signals that users receive about the quality. As a result, provided that both false positive and true positive signals from the technology are sufficiently likely, the equilibrium is again separating and a high ad-intensity plan targets naïve users, while sophisticates are charged a subscription fee. As before, the separating equilibrium features higher markups and lower welfare for both types of users.

Platform-level competition also has nuanced effects on user welfare for similar reasons. All else equal, competition between platforms could reduce surplus extraction from users. Nevertheless, this offset is incomplete because naïve users overvalue ads, since they consider them to be more informative than they actually are. As a result, when ads are sufficiently informative, a separating

[3]This equilibrium allocation contrasts with a benchmark in which all users are fully rational, where the unique equilibrium is a subscription-based model with no ads. The first-best allocation is actually different than this fully-rational benchmark because it features some positive (but typically small) amount of digital advertisement, which is informative for users.

equilibrium emerges where naïve agents are targeted by frequent digital ads and welfare is low, despite between-platform competition.

The pervasive market failures and the equilibrium choice of online business models that are *de facto* manipulative for naïve agents raise the question of whether feasible regulatory policies might improve welfare. We show that the first best (where a social planner controls the full allocation without any incentive compatibility constraints in the self-selection of different types of users) is in general not achievable, but the second best (where the social planner can only influence the business models of platforms and has to respect the self-selection of different types of users) can be decentralized in equilibrium through non-linear taxation and subsidy policies on the platform and firm. We conclude by proving that simple policies that simply linearly tax digital ad revenue can improve consumer welfare over equilibrium welfare.

Related Literature. This paper relates to several strands of work. The first is the industrial organization literature on informational advertising (e.g., see Tirole (1988), Grossman and Shapiro (1984) and Dixit and Norman (1978)). Most closely related to our work is Meurer and Stahl (1994), which builds on the seminal paper of Butters (1977), to construct a model in which advertising is informative about partially substitutable, horizontally-differentiated products, and consumers use the information in the ads to decide which products to buy. We differ from this paper and from all others in this literature in three important ways. First, a platform is situated in-between the users and the ads, and users make active decisions influencing how many ads they will see (by their choice of plan). This type of platform intermediation makes us more closely related to the work on two-sided marketplaces (e.g., Rochet and Tirole (2006), De Reuver et al. (2018)). Second, because of the presence of naïve agents, ads in our setup are simultaneously informative and manipulative. This dual role of ads and the interaction between naïfs and sophisticates are at the root of all of our results. Lastly, there is no equivalent of choices over the business models in this literature, which is also critical for our results about separating naïve and sophisticated users and the resulting low welfare.

The second related literature focuses on deceptive and manipulative advertising (e.g., Danciu et al. (2014) and Eyal (2014)), and is tightly connected to the work on behavioral manipulation on platforms (see Acemoglu et al. (2023b) and Susser and Grimaldi (2021)). Within the strand, our paper is most similar to Piccolo et al. (2018), Hattori and Higashida (2012), and Gupta (2023), where ads can be potentially misleading and persuade users to take actions that benefit the advertiser but make the buyer worse off. In Piccolo et al. (2018), for example, products are vertically-differentiated and the focus is on the existence of pooling equilibria where advertising obfuscates true differences in quality. In Hattori and Higashida (2012), all consumers are gullible and take misleading advertising at face value rather than make inferences about product quality from the information contained in ads. In Gupta (2023), deceptive advertising is more persuasive to naïve consumers who do not internalize the possibility of false advertising in how they update their beliefs about product quality. Our work differs from this literature in three different dimensions again. The first is again the presence of a platform intermediating between firms and users, and choosing business models (which is the key vehicle for separating equilibria to emerge in our model). Second, ads in our framework are both informative and manipulative, and as noted above, this dual role of advertisement is critical for our results. Third,

required policy interventions are very different in our setup. While in the presence of purely deceptive advertising, it is optimal to ban or prevent advertising altogether or strictly regulate deception, in our setup the second-best includes positive level of advertisement and this can be achieved with the appropriate taxation scheme on digital ad quantities.

Finally, there is a nascent literature on online business models and monetization, to which we are also related. Sato (2019) characterize the optimal business model of a digital platform with users who have different demand elasticities for entertainment and advertising. This paper establishes that a two-item menu, comprising free ad-based plan and a paid-for premium plan with no ads, is profit maximizing for the platform. Building on this work, Zennyo (2020) studies a setup with multiple competing ad-based platforms. In contrast to our work, in these papers digital advertising does not play a crucial role (for example, are not informative and generate an exogenously-specified fixed revenue for the platform). More importantly, there is no notion of manipulative advertising, which is our key focus.

The rest of the paper is organized as follows. The next section introduces our model, describes agent payoffs, and defines user welfare. Section 3 characterizes the unique (Berk-Nash) equilibrium of the model and provides comparative statics. Section 4 generalizes the baseline model to allow the platform to adopt richer (mixed) business models. Section 5 studies the effects of introducing firm-level and/or platform-level competition. Section 6 concludes welfare-increasing policy interventions, while Section 7 concludes and the Appendix contains the proofs and some additional analysis omitted from the text.

2 A Model of Content Platforms

There are three types of agents: firms, platforms, and users. Our baseline model consists of a single firm and a single platform. The firm is a monopolist who sells a single horizontally-differentiated product. The media platform supplies entertainment and (digital) ads to its users, but can intermix advertisements (from now on, simply ads) that market the product by providing information to the user about her idiosyncratic preferences for that product. Users consume the content offered by the platform and are potential consumers for the product of the firm.

Users. There is a continuum of users who each have a two-dimensional type $(\tau_i, \theta_i) \in \{S, N\} \times \{0, 1\}$. The first dimension corresponds to the user's sophistication level; each user i is either sophisticated ($\tau_i = S$, with probability λ) or naïve ($\tau_i = N$, with probability $1 - \lambda$). The second dimension of the user's type, $\theta_i \in \{0, 1\}$, represents whether the product offer by the firm is high or low quality for her. Specifically, the product is high-quality for user i ($\theta_i = 1$), with prior probability q. All events are independent across users and other random variables. Users derive utility from the products they purchase and from entertaining content on the platform, as we will describe below.

Firm. There is a monopolist who sells a single product and chooses a unit price p. This implies that any user i who purchases z_i pays price pz_i. The firm's marginal cost for production is constant equal to c.

Platform. The platform operates as a two-sided marketplace, connecting firms with users. It offers engaging content, such as videos and music, while also displaying digital ads on behalf of firms. Each user spends a total time $T > 0$ on the platform, during which the platform determines the proportion of time allocated to ads versus entertaining content. Specifically, we assume there is a single ad that may be shown multiple times to the user. The appearance of ads follows a Poisson process with a rate of α, meaning the total number of ad displays is distributed according to Poisson(αT). Each ad lasts for a normalized duration of 1. The platform selects the parameter α.

The probability that a user sees the ad at least once is $1 - e^{-\alpha T}$. When an ad is viewed, it provides the user with an informative signal about her type (i.e., the product's quality for her), denoted by θ_i. However, because advertising reduces the time a user spends consuming content she enjoys, more frequent advertising comes at a cost. Under our Poisson assumption, the expected time the user spends viewing entertaining content is $(1-\alpha)T$, with the remaining αT allocated to ads.

Information Structure. If user i views the ad, it provides a binary signal $s_i \in \{G, B\}$ about the product, which is independent across users. However, if the same ad is shown multiple times to the same user, she does not obtain additional information from this. The signal distribution for the ad is given by

$$\begin{cases} s_i = G, & \text{with probability } \phi_1 \text{ if } \theta_i = 1\,, \\ s_i = B, & \text{with probability } 1-\phi_1 \text{ if } \theta_i = 1\,, \\ s_i = G, & \text{with probability } \phi_0 \text{ if } \theta_i = 0\,, \\ s_i = B, & \text{with probability } 1-\phi_0 \text{ if } \theta_i = 0\,, \end{cases} \tag{1}$$

where we assume that $\phi_1 > \phi_0$, which means that a positive ("good") signal provides information to the user that the product is *more* likely to be high quality for her ($\theta_i = 1$), while a negative ("bad") signal is bad news about the match quality. This implies that there are both type-I and type-II errors. We assume throughout that the signal distribution of (1) is the objective model or the "ground truth".[4]

Users evaluate signals not necessarily according to the objective model given in (1), but according to their *subjective model*. This is where the type of the user matters. Specifically, we assume that the subjective model of user i of type τ_i on signal distribution is

$$\begin{cases} s_i = G, & \text{with probability } \phi_1 \text{ if } \theta_i = 1\,, \\ s_i = B, & \text{with probability } 1-\phi_1 \text{ if } \theta_i = 1\,, \\ s_i = G, & \text{with probability } \phi_{0,\tau_i} \text{ if } \theta_i = 0\,, \\ s_i = B, & \text{with probability } 1-\phi_{0,\tau_i} \text{ if } \theta_i = 1\,. \end{cases} \tag{2}$$

[4]While we assume that only a single ad is shown on the platform, our analysis generalizes to the case of multiple different ads whereby each ad provides additional informative signals about θ_i. In this case, the number of ads seen by each user would still be given by $k \sim$ Poisson(αT), but now the ads generate incremental information about the user's preferences. This implies, in particular, that the number of ads with signal $s_i = G$ would be drawn as a binomial distribution with k trials and success probability ϕ_{θ_i}. Our results go through identically under this alternative formulation, with the exception that the (post-ad) conditional probability of user i that $\theta_i = 1$ becomes $\pi_i = (\phi_1^{k_+}(1-\phi_1)^{k_-}q)(\phi_1^{k_+}(1-\phi_1)^{k_-}q + \phi_{0,\tau_i}^{k_+}(1-\phi_{0,\tau_i})^{k_-}(1-q))^{-1}$, which depends on the number of positive signals k_+ and the number of negative signals k_-, with $k_+ + k_- = k$.

The subjective model summarizes the extent to which agents understand how ads are targeted and customized to their specific circumstances. For instance, Facebook utilizes browsing data, clicks, shares, and likes to classify users into "custom audiences" for various advertisers (see Tran (2017), Kruikemeier et al. (2016), Galán et al. (2019) for details on these marketing strategies). New generative AI tools has further enabled microtargeted advertising techniques, where hyper-personalized content can be generated individually in real-time (see Simchon et al. (2024) and Golab-Andrzejak (2023)).

We assume that sophisticated agents are aware of these marketing strategies, so for them, $\phi_{0,S} = \phi_0$. On the other hand, naïve agents are not fully aware, and we represent their parameter as $\phi_{0,N} = \omega_N \omega_P \phi_0$. Here, $\omega_N \leq 1$ reflects their naïveté (which would apply even without personalization), while $\omega_P \leq 1$ accounts for the personalized tailoring and targeting of ads, which may not be fully understood by naïve agents. We assume $\omega_N \omega_P < 1$, although the specific manner in which this bias is introduced is not crucial to our results, it may be relevant for certain informational interventions. For simplicity, throughout the rest of the paper we will treat ϕ_0 and $\phi_{0,N}$ as model primitives, suppressing the dependence on ω_N and ω_P.

The case where $\omega_N = \omega_P = 1$ represents the *fully-rational benchmark*, where naïveté does not affect how ads are perceived, and digital ads are either not tailored and targeted, or if they are, their targeting is fully understood by naïve agents. Alternatively, the case where $\lambda = 1$, meaning all agents are sophisticated, also corresponds to this benchmark. Our primary focus is on the situation where $\phi_{0,N} < \phi_0$, which we interpret as reflecting a common real-world scenario in which targeted *digital advertising* can mislead at least some agents. In this context, the difference between $\phi_{0,N}$ and ϕ_0 can be seen as the extent to which the platform's technology can *de facto* manipulate naïve agents.

2.1 Actions and Timing

Next, we define the exact strategic game played by our agents. The game will consist of five stages, denoted $t = 1, 2, 3, 4$, and 5, as depicted in Figure 1.

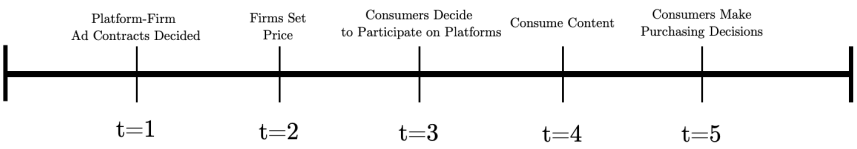


Figure 1. Timing of the Advertising Model.

(i) At $t = 1$, the platform and firm negotiate a contract that specifies an ad intensity α and a monetary transfer m from the firm to the platform (the "advertising revenue"). For simplicity, we assume this takes the form of a take-it-or-leave-it offer (α, m) from the platform to the firm, which is then either accepted or rejected by the firm. If the firm rejects the contract (or the platform does not offer one), the platform can advertise at whatever rate it desires. The platform can also set a subscription fee P for users to join the platform, and the acceptance decision of the firm can be conditioned on P (since this determines participation in the platform).

(ii) At $t = 2$, the firm sets its price p^* for its product.

(iii) At $t = 3$, the platform produces content and advertises at rate α. Each user i makes a binary decision $x_i \in \{0, 1\}$ about whether to spend time T ($x_i = 1$) or no time ($x_i = 0$) on the platform, given α and subscription price P.[5] A decision $x_i = 0$ to not participate gives user i an outside option $v > 0$.

(iv) At $t = 4$, the user digests the platform content (including any ads offered at the Poisson rate α). She receives utility value of entertainment equal to $(1 - \alpha)T$. Any user who does not engage with the platform ($x_i = 0$) views no content whatsoever (ads or entertainment) and just receives her outside option v.

(v) At $t = 5$, each user i decides how much of the product to purchase, z_i, at the price p, based on her updated posterior about θ_i.

2.2 Payoffs and Solution Concept

Platform. Recall that the platform can generate revenue by charging the firm for advertising and/or by charging users a subscription fee. When the firm accepts a contract (α^*, m^*) and the platform charges P^* to users, its payoff is

$$m^* + \int_0^1 P^* x_i \, di \,,$$

with the convention that if the firm rejects the contract, then the monetary transfer is $m^* = 0$.

Firm. The firm generates a profit by selling its product, but pays the platform for advertising. That is, the firm receives a payoff

$$\int_0^1 (p^* - c) z_i^* \, di - m^*,$$

where z_i^* is the consumption decision of agent i and m^* is the agreed upon transfer for an accepted contract (and zero for a rejected contract).

Users. Each user i receives utility both from product consumption and from content consumption on the platform. As specified above, the utility from the platform is $(1 - \alpha)T$, when the user is on the platform or v when she chooses her outside option. In addition, given her type θ_i and consumption level z_i, she receives a consumption utility $U(z_i; \theta_i) = \beta \theta_i z_i - z_i^2/2$. This implies that her expected utility from consumption is

$$\max_{z_i \geq 0} \mathbb{E}^{\tau_i} \left[U(z_i; \theta_i) - p z_i \right] = \max_{z_i \geq 0} \mathbb{E}^{\tau_i} \left[\beta \theta_i z_i - z_i^2/2 - p z_i \right],$$

where $\mathbb{E}^{\tau_i}$ is the expectation according to type τ_i's subjective probably distribution.[6] Given the linear-quadratic utility, the parameter β is the slope of the demand for the product and thus determines the

[5] Our results are not sensitive to a continuous time allocation decision $x_i \in [0, 1]$ for each user, since consuming platform content (potentially with ads) has diminishing marginal utility and the outside option is constant, and thus there will be two potential candidates for consumption $x_i \in \{0, \bar{x}\}$. The only additional complication in this case would be $\bar{x}$ may change with some parameters, rather than always being equal to T.

[6] For sufficiently high values of β, we always have $z_i^* > 0$ in equilibrium, and for simplicity, we focus on such cases and drop the non-negativity constraint from z_i. We view this as the empirically relevant configuration, since estimates in the

elasticity of demand. We also note that linear-quadratic utility is a simplifying assumption and implies that, for Bayesian agents with the correct probability distribution, additional information does not change the expected quantity consumed, which is a convenient benchmark.[7]

Solution Concept. We use the notion of (perfect) Berk-Nash equilibrium (Esponda and Pouzo (2016)) to model agents' beliefs under misspecified signal structures. Perfection here simply means that we impose sequential rationality at each information set, given beliefs, and when this causes no confusion, we refer to our equilibrium notion as Berk-Nash equilibrium or simply as "equilibrium". This implies in particular that all agents are Bayesian, but only given their subjective model. Because the subjective model of sophisticates is the objective model, a sophisticated agent will have a standard Bayesian belief π^S about $\theta_i = 1$ conditional on viewing an ad:

$$\pi^S(s_i) = \begin{cases} \frac{\phi_1 q}{\phi_1 q + \phi_0(1-q)}, & \text{if } s_i = G, \\ \frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_0)(1-q)}, & \text{if } s_i = B. \end{cases}$$

In particular, $\pi^S \mid \theta_i$ will be distributed as a multinomial with

$$\pi^S \mid \theta_i \sim \begin{cases} q, & \text{with probability } e^{-\alpha T}, \\ \frac{\phi_1 q}{\phi_1 q + \phi_0(1-q)}, & \text{with probability } \phi_{\theta_i}(1 - e^{-\alpha T}), \\ \frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_0)(1-q)}, & \text{with probability } (1-\phi_{\theta_i})(1 - e^{-\alpha T}), \end{cases}$$

where $\alpha \in [0,1]$ is the advertising intensity of the platform and $T \in \mathbb{R}$ is the time allocated on the platform by the user. We refer to F_{S0} and F_{S1} as the distributions over $\pi^S \mid \theta_i$ for $\theta_i = 0$ and $\theta_i = 1$, respectively.

Naïfs, on the other hand, update according to their subjective model and have:

$$\pi^N(s_i) = \begin{cases} \frac{\phi_1 q}{\phi_1 q + \phi_{0,N}(1-q)}, & \text{if } s_i = G, \\ \frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_{0,N})(1-q)}, & \text{if } s_i = B, \end{cases}$$

where $\pi^N \mid \theta_i$ is distributed as a multinomial with

$$\pi^N \mid \theta_i \sim \begin{cases} q, & \text{with probability } e^{-\alpha T}, \\ \frac{\phi_1 q}{\phi_1 q + \phi_{0,N}(1-q)}, & \text{with probability } \phi_{\theta_i}(1 - e^{-\alpha T}), \\ \frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_{0,N})(1-q)}, & \text{with probability } (1-\phi_{\theta_i})(1 - e^{-\alpha T}). \end{cases}$$

Importantly, because $\phi_{0,N} < \phi_0$, we have $\pi^S \mid \theta_i \preceq_{FOSD} \pi^N \mid \theta_i$ for both $\theta_i \in \{0,1\}$. In other words,

literature suggest relatively lower demand elasticities (high β) for products as compared to entertainment (see Chyi (2005), Vock et al. (2013), and Chyi and Ng (2020)). See also Berger et al. (2015), Sherman and Waterman (2016), and Flew (2021). Appendix B relaxes the assumption that $z_i^* > 0$ and shows that with the kinked demand curves that arise in this case, our results are essentially identical.

[7] Beyond the linear-quadratic case, expected consumption may increase or decrease depending on whether the implied demand curve is concave or convex.

the beliefs of naïve agents that $\theta_i = 1$ are more favorable than the beliefs of sophisticated agents given their ad viewership.[8]

Note also that the distributions over $\pi^N \,|\, \theta_i$ (for $\theta_i = 0$ and $\theta_i = 1$), denoted by F_{N0} and F_{N1}, are governed by the objective probability distribution over signals—rather than the naïves' subjective model. That is, while the *interim beliefs* $\pi^N(G)$ and $\pi^N(B)$ are updated using $\phi_{0,N}$ for the naïve agents, the induced probabilities over $\{q, \pi^N(G), \pi^N(B)\}$ are governed by the objective model using ϕ_0 when $\theta_i = 0$. For *(ex post)* welfare it will be these distributions based on objective measures that are relevant.

Equilibrium. We can determine the unique (perfect) Berk-Nash equilibrium of this sequential game via backward induction.

(a) At $t = 5$, each user holds belief π_i that she values the product ($\theta_i = 1$) and choosess her optimal consumption z_i^* to solve $C_i(\pi_i, p) \equiv \max_{z_i} \pi_i U(z_i; \theta_i = 1) + (1 - \pi_i)U(z_i; \theta_i = 0) - pz_i$ given π_i, where $C_i(\pi_i, p)$ represents the expected consumption utility *given* a belief π_i about $\theta_i = 1$ and a given product price p.

(b) At $t = 4$, user i's belief π_i is determined by Bayes rule given the realization of the signal for this user (conditional on her participation decision $x_i \in \{0, 1\}$ and the ad intensity α) and her subjective model, varying according to her sophistication type τ_i.

(c) At $t = 3$, each user decides whether to participate on the platform by solving $\max_{x_i \in \{0,1\}} x_i(\mathbb{E}^{\tau_i}_{\pi_i}[U(C_i(\pi_i, p); \theta_i = 1) - pC_i(\pi_i, p) \,|\, \alpha] + (1 - \alpha)T) + (1 - x_i)(U(C_i(q, p); \theta_i = q) - pC_i(q, p) + v)$, where recall that $\mathbb{E}^{\tau_i}$ is the expectation with respect to the subjective model of a user with type τ_i. When participating on the platform and observing digital ads ($x_i = 1$), the user obtains an informational value, since they believe that these ads lead to better decisions. Specifically, the perceived (interim) informational value from digital ads α for users of type τ_i is:[9]

$$\begin{aligned} I_{\tau_i}(\alpha) =& q\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=1}[U(C_i(\pi_i, p^*);\, \theta_i = 1) - p^* C_i(\pi_i, p^*) \,|\, \alpha] \\ &+ (1 - q)\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=0}[U(C_i(\pi_i, p^*);\, \theta_i = 0) - p^* C_i(\pi_i, p^*) \,|\, \alpha] - (U(C_i(q, p^*);\, \theta_i = q) - p^* C_i(q, p^*)). \end{aligned}$$

(d) At $t = 2$, given ad intensity α, the firm sets price p by solving $\Pi(\alpha) \equiv \max_{p \geq 0} \int_0^1 \left[x_i^*(\alpha, p)(p - c)\mathbb{E}_{\tau_i}[\mathbb{E}^{\tau_i}_{\pi_i}[z_i^*(\pi_i, p) \,|\, \alpha]] + (1 - x_i^*(\alpha, p))(p - c)z_i^*(q, p)\right] di$, where $\Pi(\alpha)$ is the expected profit of the firm given an advertising intensity α on the platform.

(e) At $t = 1$, given a contract (α, m) and subscription fee P, the firm accepts the contract if and only if $\Pi(\alpha) - m \geq \max_{p \geq 0}(p - c)z_i^*(q, p)$. The platform then selects the contract (α, m) that maximizes m conditional on the acceptance rule by the firm.

We assume the outside option satisfies $v > I_N(1)$, so that users only participate on the platform when at least some entertainment content is shown.

[8] A particularly transparent special case has $\phi_{0,N} = 0$, where a positive ad is always interpreted by naïve agents as evidence that the product is high quality.

[9] That this quantity does not depend on price p^* is established in Lemma A.1.

2.3 User Welfare

When a user abstains from participating on the platform, her utility is

$$W(\tau_i, x_i = 0) = \underbrace{v}_{\text{Outside Option}} + \underbrace{qU(z_i^*(q,p^*);\,\theta_i = 1) + (1-q)U(z_i^*(q,p^*);\,\theta_i = 0) - p^* z_i^*(q,p^*)}_{\text{Expected Product Consumption Utility}},$$

where p^* is the price chosen by the firm in equilibrium.

On the other hand, suppose that a user engages with the platform, where (α^*, m^*) is the contract accepted (with the convention $m^* = 0$ for a rejected contract), p^* is the price chosen by the firm in equilibrium, and P^* is the subscription price charged by the platform. Recall that we let F_{S0}, F_{S1}, F_{N0} and F_{N1} denote the distributions over π_i for the respective types, $(S,0)$, $(S,1)$, $(N,0)$, and $(N,1)$, using the objective model (which depend on α^*). We are using the objective model here because, as already noted, all of our welfare results focus on *ex post* utility, which depends on the objective model, and not on the subjective interim beliefs of naïve types. Average user welfare by type (conditional on platform engagement) is given directly by

$$\begin{aligned} W(\tau_i, x_i = 1) = \underbrace{(1-\alpha^*)T - P^*}_{\text{Content Consumption Surplus}} &+ q\underbrace{\mathbb{E}_{\pi\sim F_{\tau_i 1}}\left[U(z^*(\pi,p^*);\,\theta_i = 1) - p^* z^*(\pi,p^*)\right]}_{\text{Product Consumption Utility Conditional on } \theta_i = 1} \\ &+ (1-q)\underbrace{\mathbb{E}_{\pi\sim F_{\tau_i 0}}\left[U(z^*(\pi,p^*);\,\theta_i = 0) - p^* z^*(\pi,p^*)\right]}_{\text{Product Consumption Utility Conditional on } \theta_i = 0}. \end{aligned}$$

The use of the distributions F_{S0}, F_{S1}, F_{N0}, and F_{N1} in computing welfare—rather than the agent's subjective model—reiterates that we are evaluating *ex post* welfare. We denote the welfare of sophisticated and naïve agents, respectively, by $W^*(S) = W(S, x_S^*)$ and $W^*(N) = W(N, x_N^*)$.

2.4 First Best

We start by characterizing the first-best allocation, which clarifies how the utility of users can be maximized when a planner has complete control over all aspects of the allocation.

Proposition 1. *The first-best user welfare occurs when the platform advertises at the rate α_S^{FB} to sophisticates and at rate α_N^{FB} to naïfs, where $\alpha_N^{FB} \leq \alpha_S^{FB}$. Moreover, for type τ user, $W_{FB}(\tau) > W_{base}(\tau)$, where $W_{base}(\tau)$ is the base case welfare with no platform.*

Proposition 1 shows that the first-best allocation involves a small amount of advertising on the platform to both types of agents—but crucially different amounts for different types. For sophisticated agents, advertising improves their decision-making and is socially valuable, and in addition, they enjoy the content on the platform, and hence $W_{FB}(S) > W_{\text{base}}(S)$. The same forces are present for naïfs, but the social planner prefers to send them fewer ads because they tend to misinterpret the information in the ads and thus they derive less (*ex post*) utility from ads. The fact that the social planner can prefer strictly positive ads underscores the informative nature of digital ads in our model.

We can also already see another important point, which plays a central role in our analysis below. If naïve agents were given a choice between the two levels of advertising, α_S^{FB} and $\alpha_N^{FB} \leq \alpha_S^{FB}$, they

would choose the higher one, α_S^{FB}. In fact, given the option, they would prefer even higher levels of the advertising than α_S^{FB}. This is because, at the interim stage, they erroneously think the ads are more informative than they truly are. This is one of the reasons why the first best will never be implementable in a decentralized equilibrium in this model.

3 Baseline Equilibrium Characterization

To build intuition, we start with a simplified version of our model, where we allow the platform to either charge a subscription fee or use digital ads, but *not both.* This, as we will see, has no major effect on the insights our model generates, but will simplify our initial characterization.

3.1 Equilibrium Business Models and Digital Ads

We start by presenting a number of lemmas, which then leads to the main characterization results (for behavior and welfare) in this baseline environment.

Lemma 1. *Let $\Pi_S^*(\alpha)$ denote the firm's profit from the sophisticated agents, $\int_{\tau_i=S}(p^* - c)z_i^* \, di$, under an advertising scheme with intensity $\alpha > 0$. Then $\Pi_S^*(\alpha)$ is independent of α. In other words, the firm extracts no surplus from advertising to the sophisticated agents.*

Lemma 1 demonstrates that the firm cannot extract advertising rents from sophisticated agents' participation on the platform. This outcome arises from the linear-quadratic utility and the resulting linear demand curves. Agents who like the product receive more positive ads on average, while those who dislike the product encounter more negative ads. Although the former group is willing to purchase more and pay higher prices, the latter group's lower willingness-to-pay and reduced consumption perfectly offset these gains. This balance is a direct result of the Martingale property of Bayesian beliefs combined with linear demand.

Put differently, the user's demand curve without advertising is $z_i^*(p) = \beta q - p$, whereas after advertising it becomes $z_i^*(p) = \beta\pi_i - p$. Since sophisticated agents are fully Bayesian, their expected posterior equals the prior, $\mathbb{E}^S[\pi_i] = q$. This implies that the expected demand after advertising remains the same as the demand before advertising. This doesn't imply that users do not benefit from information—they do make more informed decisions. However, the firm cannot capture any of this surplus because the users' expected demand remains unchanged. As a result, there's no surplus for the platform to capture by charging the firm to display digital ads to users. Therefore, the platform also does not profit from showing digital ads to sophisticated users.

Although this result does not hold exactly with concave demand curves, it transparently illustrates why the main source of profits for the firm (and thus ad revenue for the platform) is from naïve agents—ads will generally have a small effect on the expected purchases of sophisticated agents, but potentially much larger effects on the expected purchases of naïve agents who overestimate the likelihood of a high-quality product given a positive signal.

Lemma 2. *Let $\Pi_N^*(\alpha)$ denote the firm's profit from the naïve agents, $\int_{\tau_i=N}(p^* - c)z_i^* \, di$, under an advertising scheme with intensity $\alpha > 0$. Then $\Pi_N^*(\alpha)$ is positive and increasing in α. In other words, the*

firm extracts positive surplus from advertising to the naïve agents and this surplus is greater when there is more advertising.

In contrast to sophisticated agents, naïfs' average demand curve drifts upward as the advertising intensity increases, even though they also have a linear demand curve. This result is rooted in the fact that naïve agents have the wrong model and update their beliefs under the perception that low-quality products generate positive signals with probability $\phi_{0,N}$, while in truth they generate such signals with probability $\phi_0 > \phi_{0,N}$. Then using our notation of F_{N0} and F_{N1} for the distribution of beliefs π_i for naïfs, we can write their demand curve as $z_i^*(p) = \beta(q\mathbb{E}_{\pi_i \sim F_{N1}}[\pi_i \mid \alpha^*] + (1-q)\mathbb{E}_{\pi_i \sim F_{N0}}[\pi_i \mid \alpha^*]) - p$, which is strictly greater than their expected purchases without ads, $\beta q - p$, and also increasing in α. This allows the firm to charge higher prices and secure higher profits in the product market when there are naïve agents receiving digital ads. This surplus by the firm is then extracted by the platform via the monetary transfer, m^*.

The same forces underlying Lemma 2 also leads to our next result about the (interim) informational value that sophisticated and naïve agents derive from digital ads.

Lemma 3. *For any α, $I_N(\alpha) > I_S(\alpha) > 0$ and $\arg\max_{\alpha\in[0,1]} I_N(\alpha) + (1-\alpha)T > \arg\max_{\alpha\in[0,1]} I_S(\alpha) + (1-\alpha)T$. Moreover, $I_S(\alpha)$ and $I_N(\alpha)$ are concave and monotonically increasing in α.*

Because naïves mistakenly believe that digital ads are more informative than they truly are, their subjective (interim) value from participating in the platform is greater. This implies in particular, that naïve agents are more tolerant of digital ads and in fact would choose a higher level of digital ads than sophisticated agents, as we also noted in our discussion of why the first-best allocation cannot be implemented. It also implies that they are more willing to take part in an ad-based platform. The utility of an agent of type τ_i is from platform participation is $I_{\tau_i}(\alpha) + (1-\alpha)T - v \geq 0$, and thus Lemma 3 implies that the platform participation constraint will always bind for sophisticates before it binds for naïfs. Put differently, whenever sophisticates participate in the platform, so do naïfs, but not vice-versa.

These three lemmas together with platform maximization yield our next result:

Lemma 4. *If the platform adopts an advertising-based business model, it sets α^* such that $I_N(\alpha^*) + (1-\alpha^*)T - v = 0$ and $I_S(\alpha^*) + (1-\alpha^*)T - v < 0$. In other words, the platform extracts all surplus from naïve agents, while sophisticates do not participate on an advertising-based platform.*

Intuitively, from Lemmas 1 and 2, the total advertising revenue the platform can generate is increasing in the ad intensity α. Moreover, from Lemma 3, the participation constraint will bind first for the sophisticates, and from Lemma 1, the platform does not collect any additional advertising revenue from serving ads to sophisticates. Therefore, the platform finds a simple form of separation profitable: sophisticates are excluded and naïfs receive relatively high intensity of digital advertising. Whether the platform will actually choose this separating allocation depends on how much it can collect as a subscription fee from both types. This calculation leads to our main characterization result in this baseline environment:

Proposition 2. *There exists $\hat{\phi}_0(\lambda, \phi_1, \phi_{0,N}) > \phi_{0,N}$ such that:*

(a) If $\phi_0 < \hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$*, the platform chooses a subscription model with* $P^* = T - v$ *and the firm sets a price* $p^* = \bar{p}^* \equiv (\beta q + c)/2$*;*

(b) If $\phi_0 > \hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$*, the platform chooses an advertising model with ad intensity* $\hat{\alpha}^* = \arg\max\{\alpha \in [0,1] : I_N(\alpha) + (1-\alpha)T - v = 0\}$ *and the firm chooses a price* $\hat{p}^* > \bar{p}^*$*.*

Moreover, $\hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$ *is increasing in* λ *and* $\phi_{0,N}$*, and decreasing in* ϕ_1*.*[10]

Proposition 2 follows from Lemmas 1-4, which collectively indicate that the profit-maximizing business model depends on the extent to which the platform can extract digital ad revenue from the firm, which, in turn, extracts surplus from naïve agents. When ϕ_0 is low (relative to $\phi_{0,N}$)—more precisely, when it is lower than the threshold $\hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$—there exists a small gap between the actual likelihood of false positives (positive signals from low-quality products) and the perceived likelihood of false positives by naïve agents. In this scenario, referred to as regime (a), the amount of surplus that can be extracted from naïve agents is small. Consequently, it is more profitable to include all agents on the platform and charge a subscription fee equal to the additional utility they derive from the entertaining content, $P^* = T - v$, rather than attempting to extract informational surplus from naïve agents. In this regime, the profit-maximizing monopoly price for the firm is $\bar{p}^* \equiv (\beta q + c)/2$, determined by the *ex ante* linear demand curves (as no further information acquisition occurs without digital ads).

Conversely, when ϕ_0 is higher than the threshold $\hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$, it is more profitable to opt for an ad-based business model and exclude sophisticates from the platform. We refer to this case as regime (b). In this regime, the price charged by the firm for its product is greater than in regime (a), $\hat{p}^* > \bar{p}^*$. This is because with digital ads, naïve agents have an overinflated perception of the likelihood that they will like the product and this raises the monopoly price of the firm.

With higher values of ϕ_1, ads are more informative because they are more likely to give good signals when the product is of high-quality for the agent. Increasing the informativeness of ads makes them more appealing for the users, and thus allows the platform to increase the intensity of ads while still retaining naïve users on the platform. This means that the platform can still extract sufficient surplus from the naïve agents when ϕ_0 is lower, compensating a smaller false positive rate with a higher advertising intensity. This lowers the cutoff $\hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$ and increases the range of ϕ_0 where the platform adopts an advertising-based business model.

Three other points are important to note. First, $\hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$ is increasing in λ for obvious reasons. When there are more sophisticated agents, excluding them is more costly, and thus a subscription-based business model becomes more likely. Second and more importantly, digital ads targeted at naïve agents creates negative spillover to sophisticates: not only are they excluded from the platform, but they face a higher price for the product, $\hat{p}^* > \bar{p}^*$, than they would have done in regime (a). Third, the role of the incorrect model that naïve agents rely on for evaluating the meaning of positive signals is

[10] In all of the comparative static results we present, we will vary one parameter while holding the others fixed. In doing so, we treat ϕ_0 and $\phi_{0,N}$ as independent (free) parameters. One could alternatively define ω_N and ω_P as independent and let $\phi_{0,N} = \omega_N \omega_P \phi_0$, so that varying ϕ_0 simultaneously varies $\phi_{0,N}$. This alternative formulation has no substantial impact on our results (the equilibrium would retain the same cutoff structure and comparative statics with respect to λ and ϕ_1 would remain unchanged) but would increase the cutoff $\hat{\phi}_0$ in Proposition 2 (because we are increasing both ϕ_0 and $\phi_{0,N}$).

critical for this result, which can be seen from the fact that the equilibrium business model depends on the gap between the true likelihood of positive signals for low-quality products, ϕ_0, and their perceived likelihood, $\phi_{0,N}$. This can be seen also from considering the fully-rational benchmark, which we derive in the next subsection.

3.2 Fully-Rational Benchmark

It is useful to also consider the fully-rational benchmark where $\phi_{0,N} = \phi_0$ and thus there is no misperception on the part of naïve agents.

Proposition 3. *If $\phi_{0,N} = \phi_0$, the profit-maximizing platform business model of the platform is subscription-based with $P^* = T - v$ and the firm sets price $p^* = \bar{p}^*$. The welfare of agents of type $\tau \in \{S, N\}$ of the fully-rational benchmark, $W_{\text{fully-rational}}(\tau)$, is equal to their base case welfare, $W_{\text{base}}(\tau)$, with no platform at all.*

In this scenario, naïve agents have an accurate understanding of how digital ads generate signals, resulting in no distinction between sophisticated and naïve agents. This implies that neither the firm nor the platform can extract any informational surplus from naïve agents. Consequently, the profit-maximizing business model is subscription-based, with the same subscription fee as in Proposition 2, $P^* = T - v$, which captures all the surplus users would derive from consuming entertainment. The firm also sets the same price as before, $\bar{p}^*$. Since the platform is extracting all the surplus it helps create, the utilities of both types of agents are equivalent to what they would have been in the hypothetical case where the platform did not exist.

3.3 Digital Advertising: Welfare Analysis

Armed with the characterization of equilibrium in Proposition 2, we next determine user welfare in the benchmark equilibrium, separately in regimes (a) and (b).

Proposition 4.

(a) When $\phi_0 < \hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$, user welfare is $\hat{W}^*(\tau) = W_{\text{fully-rational}}(\tau) = W_{\text{base}}(\tau)$ for both $\tau \in \{S, N\}$.

(b) When $\phi_0 > \hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$, user welfare is $\hat{W}^*(\tau) < W_{\text{fully-rational}}(\tau) = W_{\text{base}}(\tau)$ for both $\tau \in \{S, N\}$.

The first part of this proposition is not surprising. When $\phi_0 < \hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$, the profit maximizing business model involves a pure subscription fee and no user sees any digital ads. The platform then captures the full surplus it creates and the equilibrium monopoly price is the same as the case without the platform as well, $\bar{p}^*$. Consequently, user welfare, for both sophisticated and naïve agents, is exactly the same as it would have been without the platform (which also coincides with the fully-rational welfare levels as we saw in Proposition 3).

The second part of the proposition is our main result and is quite a striking one. It shows that both sophisticates and naïfs have lower utility than the case in which the platform does not exist. Let us start with the naïfs. Naïve agents have lower *ex post* welfare than the case without the platform because they

receive a high number of digital ads and given their misperception about the signal generating process, they are *de facto* manipulated and spend more on the product of the monopolist firm than they would have done with the correct beliefs. This in particular pushes them into consuming more of the product when it is low-quality for them, lowering their *ex post* welfare. Their inflated demand for the product also leads to a higher monopoly price $\hat{p}^* > \bar{p}^*$ than in regime (a) (as well as in the hypothetical case without the platform), and this reduces consumer surplus for both naïve and sophisticated agents. Consequently, the welfare of sophisticated users is lower than the base case, purely because of the negative spillover the naïfs create on them, working through the higher prices of the monopoly firm. In fact, recall that in this case sophisticated users do not even participate in the platform (and thus do not enjoy the entertaining content or receive ads), so the only impact on their welfare relative to the environment without the platform is via the price effect.

Another noteworthy feature is that the negative welfare effects occur when there is a larger gap between ϕ_1 and ϕ_0, which roughly corresponds to the informativeness of digital ads. The comparative static of increasing ϕ_0 (holding ϕ_1 fixed) then gives the following paradoxical result: the ad-based business model emerges when digital ads are less useful.

Remark — We conclude this section by highlighting the role of two assumptions in the sharp welfare comparative result in Proposition 4. The first is the linear-quadratic utility function, which we discussed already. Recall that this structure implies that the firm and thus the platform cannot extract any surplus from the sophisticated agents. If we adopted a different utility function, expected consumption may increase or decrease with informative digital ads. In this case, the platform may be able obtain additional revenues from sophisticates with the ad-based business model. Nevertheless, the source of our main result—the fact that digital ads generate more revenues from naïve agents—continues to hold in this case, highlighting that the linear-quadratic utility function is just a simplifying assumption. Second, it is important that naïve users become over-optimistic about product quality after seeing digital ads. If they naïvely became over-pessimistic, then the mechanism we emphasize here would not apply. We do not see this as a shortcoming, however, since naïve agents being *de facto* manipulated by ads into believing that products featured in advertisements are higher quality than they are in reality is the plausible case.

4 General Platform Business Models

In this section, we relax the assumption that the platform cannot offer both subscription fees and digital ads. We will see that all of the main insights from the previous section generalize to this case. We start with equilibrium characterization in Section 3 when the platform is allowed to offer a menu consisting of multiple plans, each of which is monetized either with advertisements or subscription fee. In Section 4.2, we allow the most general case where the platform can offer multiple plans some of which intermix ads and subscription.

4.1 Profit-Maximizing Menus of Business Models

We now allow the platform to offer a menu of multiple plans that specify either an ad intensity α_ℓ or a subscription price P_ℓ, and each user is allowed to self select into one of these plans. We first observe that, with fully-rational users, Proposition 3 still applies and the unique profit-maximizing strategy is to offer a single subscription-based business model with $P^* = T - v$.

We next consider the case where there are both sophisticated and naïve agents. Our main result is provided in the following proposition.

Proposition 5. *When the platform is allowed to offer a menu of plans, there exists* $\phi_0^*(\lambda, \phi_1, \phi_{0,N}) < \hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$ *such that*

(a) If $\phi_0 < \phi_0^*(\lambda, \phi_1, \phi_{0,N})$*, the profit-maximizing business model is subscription-based, with* $P^* = T - v$*, and product price* $p^* = \bar{p}^*$.

(b) If $\phi_0 > \phi_0^*(\lambda, \phi_1, \phi_{0,N})$*, the profit-maximizing business model involves a subscription plan with* $P^* = T - v$*, an ad-based plan with intensity* $\hat{\alpha}^*$*, and product price* $p^* = \hat{p}^*$.

Moreover, $\phi_0^*(\lambda, \phi_1, \phi_{0,N})$ *is increasing in* λ *and* $\phi_{0,N}$*, and decreasing in* ϕ_1.

Recall that in Proposition 2, the separating equilibrium took the form of sophisticated agents being completely excluded from the platform. This was a consequence of the fact that the platform could offer either a subscription fee or digital ads, but not both. Nor could it offer a menu of different plans. Now that such a menu is feasible, the platform can always attract the sophisticates with a subscription-based service. This does not impact the profit-maximizing plan offered to naïve users, which remains the same as in Proposition 2(b). The key spillover from naïfs to sophisticates identified in Proposition 2 is present here as well. When naïve agents receive the ad-based plan (with high intensity), this increases their demand for the product and thus the monopoly price of the firm, which then hurts sophisticates.

Interestingly, $\phi_0^*(\lambda, \phi_1, \phi_{0,N})$ is increasing in λ, despite the platform being able to perfectly segment naïfs and sophisticates. This is because the amount of advertising revenue the platform can extract from the firm is (convex) quadratic in the fraction of naïve agents, whereas the subscription revenue is linear. Thus, as more naïve agents participate on an ad-based platform, the more profit the firm generates from *other* naïve agents, because with more naïve agents, it can charge higher prices. This implies that, for the platform, there is more than a one-for-one loss in advertising revenue as we increase the fraction of sophisticates, λ. It is easiest to see this when thinking about the extreme cases. When there are very few naïve agents, the value of advertising is also very little because the firm has to simultaneously pick its product price to appeal to sophisticated agents, who do not see digital ads. The presence of sophisticates will deflate the surplus the firm can extract from naïfs because it charges a single price. On the other hand, when most of the population is naïve, digital advertising is most profitable because it can price almost entirely to the crowd that viewed digital ads. A different interpretation of this result is that a higher fraction of sophisticates in the population provides some protection for naïve agents, even though the two types of agents participate in different plans.

Corollary 1.

(a) When $\phi_0 < \phi_0^*(\lambda, \phi_1, \phi_{0,N})$*, user welfare is* $\hat{W}^*(\tau) = W_{\textit{fully-rational}}(\tau) = W_{\textit{base}}(\tau)$ *for both* $\tau \in \{S, N\}$.

(b) When $\phi_0 > \phi_0^*(\lambda, \phi_1, \phi_{0,N})$*, user welfare is* $\hat{W}^*(\tau) < W_{\textit{fully-rational}}(\tau) = W_{\textit{base}}(\tau)$ *for both* $\tau \in \{S, N\}$.

Therefore, the welfare results from Section 3.3 extend immediately to the setting where the platform can offer menus to their users. Even though sophisticates are now included on the platform with the subscription fee, they are still held down to their outside option, so when $\phi_0 < \phi_0^*(\lambda, \phi_1, \phi_{0,N})$, welfare effects are analogous to those in Proposition 4. Even more importantly, when $\phi_0 > \phi_0^*(\lambda, \phi_1, \phi_{0,N})$, both types of agents are pushed to levels of welfare worse than the benchmark without the platform at all—just as in Proposition 4. As shown in Corollary 1, more informative digital ads generally lead to lower welfare. Specifically, a higher ϕ_1, while holding ϕ_0 constant, indicates that the digital ads are more informative. The comparative static with respect to ϕ_1 implies that as digital ads become more informative, the platform is more likely to adopt an ad-based plan, resulting in reduced welfare for both sophisticated and naïve agents. This outcome aligns with the findings from Section 3.

A final noteworthy observation is that $\phi_0^*(\lambda, \phi_1, \phi_{0,N})$ is lower than the threshold discussed in the previous section, $\hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$. This implies that for values of ϕ_0 between these two thresholds, the equilibrium in the previous section was subscription-based, whereas in the current section, it involves digital ads targeted at naïve agents. Intuitively, since the platform could not previously segment the market between sophisticated and naïve agents without excluding the former, it was less inclined to adopt an ad-based model. However, with the ability to offer a menu of multiple plans, the platform is now more willing to target naïve agents with ads. This means that sophisticated agents, who previously provided a form of protection for naïve agents, no longer do so when segmentation is possible through multiple plans. This shift has significant implications for welfare, as will be explored in the next proposition.

Proposition 6. *Consumer welfare* $\hat{W}^*(\tau)$ *is monotonically decreasing in both* ϕ_0 *and* ϕ_1*, for both user types* $\tau \in \{S, N\}$.

Unsurprisingly, increases in ϕ_0 lead to reductions in user welfare because it allows the platform more leverage to *de facto* manipulate the agents' consumption. A higher ϕ_0 means more false-positives and thus lower welfare for naïve agents. It also implies more inflated demand from these agents and thus a higher monopoly price for the firm, which indirectly reduces the welfare of sophisticated agents as well. Perhaps more surprising is that a higher ϕ_1 also leads to an unambiguous reduction in welfare. The reasons are twofold. First, greater ϕ_1 can induce a switch from a subscription-based business model to a mixed one where naïve users receive digital ads and thus have lower *ex post* welfare, for the same reasons as we saw in Proposition 4. Second, when the platform advertises, it always extracts maximal surplus from naïfs, and this implies in particular that any surplus from the greater informativeness of ads is captured by the platform.

4.2 Mixed Business Models

The findings of Section 4.1 readily generalize to the case where the platform can offer a richer set of plans that mix subscription and advertising. In other words, we can allow the platform to either offer a single plan (α^*, P^*) to all users or to offer two plans (α_1^*, P_1^*) and (α_2^*, P_2^*). We consider these more general mixed business models throughout the remainder of the paper. The next proposition characterizes the equilibrium in this more general case.

Proposition 7. *There exists* $\tilde{\phi}_0(\lambda, \phi_1, \phi_{0,N}) < \hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$ *such that*

(a) If $\phi_0 < \tilde{\phi}_0(\lambda, \phi_1, \phi_{0,N})$*, the profit-maximizing business model offers two (not necessarily distinct) plans* (α_1^*, P_1^*) *and* (α_2^*, P_2^*)*, where* $\alpha_S^{FB} \leq \alpha_2^* < \hat{\alpha}^*$ *with the implied welfare level for users being* $\tilde{W}_{(a)}^*(\tau) \leq W_{\textit{fully-rational}}^*(\tau) = W_{\textit{base}}^*(\tau)$ *for both* $\tau \in \{S, N\}$*.*

(b) If $\phi_0 > \tilde{\phi}_0(\lambda, \phi_1, \phi_{0,N})$*, the profit-maximizing business model involves a subscription plan with* $P^* = T - v$*, and ad-based plan with intensity* $\hat{\alpha}^*$*, and product price* $p^* = \hat{p}^*$*, with the implied welfare level for users being* $\tilde{W}_{(b)}^*(\tau) < W_{\textit{fully-rational}}^*(\tau) = W_{\textit{base}}^*(\tau)$ *for both* $\tau \in \{S, N\}$*.*

Moreover, $\phi_0^*(\lambda, \phi_1, \phi_{0,N})$ *is increasing in* λ *and* $\phi_{0,N}$*, and decreasing in* ϕ_1*.*

The structure of the business model in Proposition 7 is analogous to Proposition 5, but also richer. Once again, like in Proposition 5, the profit-maximizing business model turns on the rate of false positives, as regulated by the parameter ϕ_0. When ϕ_0 is smaller than the key threshold (as in regime (a)), the platform offers plans that have less advertising than $\hat{\alpha}^*$, which is the amount that makes naïve agents indifferent between platform participation and not. This is because the ad technology's power to manipulate naïve users' consumption behavior is weaker, and the platform prefers to collect more subscription fee from both types of users and consequently chooses a lower ad intensity. Once ϕ_0 reaches a critical threshold (as in regime (b)), then the platform can generate more revenue from maximally advertising to naïfs, leaving them with no platform surplus, and correspondingly charging no subscription fee to them (while still collecting subscription fees from sophisticates). This situation leads to the lowest user welfare for both naïve and sophisticated agents, with the impact on the latter again being driven by the spillovers through the product price.

There are two interesting observations about $\tilde{\phi}_0(\lambda, \phi_1, \phi_{0,N})$ to note. The first is that $\tilde{\phi}_0(\lambda, \phi_1, \phi_{0,N})$ is increasing in λ. The reason again relates to the fact that the firm charges a single price, which caters both to sophisticates and naïfs. This implies that when more of the population is sophisticated, who view fewer digital ads, the firm must price more to this audience and can extract less from naïfs who view a higher quantity of ads. This makes the high-intensity ad-based model of Proposition 7(b), which specifically targets naïfs, less attractive to the platform. In this sense, sophisticates protect their naïve counterparts from business models that specifically target them, because they are less effective when only a small fraction of the population is naïve. Second, $\tilde{\phi}_0(\lambda, \phi_1, \phi_{0,N})$ is also decreasing in ϕ_1, meaning that once again high-ad intensity plans are chosen for naïfs when ads are, paradoxically, more informative. This corroborates the findings of Proposition 4, further reinforcing the idea that more informative advertising does not necessarily lead to better outcomes for the users.

5 Firm-Level and Platform-Level Competition

In this section, we show that the insights emphasized so far generalize to an environment in which there are multiple platforms and firms and that the fundamental reason for this is related to the presence of digital ads. We first establish that in a generalized version of our model with multiple firms and platforms, there exists a unique (robust) equilibrium. In Section 5.2, we then study the case with multiple firms in greater detail, and subsequently in Section 5.3, we study competition between multiple platforms.

5.1 Existence and Uniqueness

We extend our model in Section 2 to allow for $N \geq 1$ firms and $M \geq 1$ platforms. At $t = 1$, each of the platforms $\rho \in \{1, \ldots, M\}$ makes contract proposals to each one of the N firms (and not offering a contract to a subset of these firms is a special case).

Because there can be uninteresting multiple equilibria based on coordination between firms, we adopt two less standard features. First, we will assume that firms set their prices sequentially rather than simultaneously. As in standard voting models, this will eliminate equilibria supported by weakly-dominated strategies and the exact sequence in which firms make their offers will turn out to be completely irrelevant (e.g., see Moldovanu and Winter (1995)). Second, we will impose a robustness refinement, which we define and explain below.

The exact timing is as follows:

- At $t = 1.1$, each platform ρ simultaneously offers menus $\{(\alpha_{1,\rho}^{(j)}, P_{1,\rho}^*), (\alpha_{2,\rho}^{(j)}, P_{2,\rho}^*), m_\rho^{(j)}\}_{j=1}^N$ of entertainment plans to every firm j to advertise at intensity $\alpha_{\ell,\rho}^{(j)}$ for product j while charging subscription price $P_{\ell,\rho}^*$ if the user selects plan ℓ on its platform. The variable $m_\rho^{(j)}$ specifies the total amount transferred from firm j to the platform ρ if accepted. Notice also the restriction that the subscription fees associated with each plan offered by platform ρ is the same across all different firms.

- At $t = 1.2$, each of the firms $j \in \{1, \ldots, N\}$ accepts or rejects these proposals (a firm can accept multiple proposals simultaneously). If the proposal from platform ρ is accepted by firm j, advertisement rates $(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})$ for firm j's product is implemented and firm j transfers $m_\rho^{(j)}$ to the platform ρ. The platform also charges subscription fees $P_{1,\rho}^*$ and $P_{2,\rho}^*$ as promised. If the proposal is rejected, the platform collects no transfer but can advertise at whatever rate it likes.

- At $t = 2$, all firms set their price p_j^* for the product sequentially.

- At $t = 3$, users decide which platform and plan to participate in, if any. Formally, user i chooses $x_{i,\ell,\rho} \in \{0, 1\}$ for all ρ and ℓ with $\sum_{\ell=1}^{2} \sum_{\rho=1}^{M} x_{i,\ell,\rho} \leq 1$.

- At $t = 4$, users enjoy the platform content and watch the ads on the platform in which they participate. We assume that the probability of each ad appearing is independent across users and across multiple ads seen by the same user.

- At $t = 5$, user i makes purchasing decisions to maximize her utility

$$\mathbb{E}^i \left[\beta \sum_{j=1}^{N} \theta_i^{(j)} z_i^{(j)} - \frac{\left(\sum_{j=1}^{N} z_i^{(j)} \right)^2}{2} - \sum_{j=1}^{N} p_j z_i^{(j)} \right],$$

  where $\mathbb{E}^i$ denotes this user's expectation given her type and $\theta_i^{(j)}$ is drawn i.i.d. according to the distribution in Section 2.

In this section, we focus on *robust* (perfect) Berk-Nash equilibria. Berk-Nash are defined exactly analogously to before, so here we only explain our robustness notion, which is used to eliminate various equilibria that emerge on the basis of miscoordination.[11] For this purpose, let us consider an ε-variant of our game where prices for all firms have to belong to the discrete grid, $p_j \in \{0, \varepsilon, 2\varepsilon, \ldots\}$ for some $\varepsilon > 0$. Define $\mathcal{E}(\varepsilon)$ as the set of Berk-Nash equilibria of this discretized game. We say that equilibrium $\{p_j^*\}_{j=1}^N$ is a *limit equilibrium* if there exists a sequence $\{\varepsilon_n\}_{n=1}^\infty$ with $\lim_{n\to\infty} \varepsilon_n = 0$ and a sequence of (Berk-Nash) equilibria $\{p_j^*(\varepsilon_n)\}_{j=1}^N \in \mathcal{E}(\varepsilon_n)$ where $\lim_{n\to\infty} p_j^*(\varepsilon_n) = p_j^*$ for all firms j. We say that an equilibrium is *robust* if it is an equilibrium of our original game (with continuous pricing decisions) and also a limit of the equilibria of the discretized game. We also say that a (robust perfect Berk-Nash) equilibrium is *essentially unique* if the resulting allocation in terms of advertisements, consumption levels and expected payoffs are uniquely pinned down (though the strategies may not be unique and, if there are any asymmetries in the equilibrium, which is possible, the identity of firms or platforms taking one type of action versus another may be non-uniquely determined).

Proposition 8. *Generically, there exists an essentially unique robust (perfect) Berk-Nash equilibrium for any number of firms N and platforms M, which is independent of the sequence of pricing decisions at $t = 2$.*

Proposition 8 establishes the existence of and essentially unique robust equilibrium, which is useful because it pins down the equilibrium allocation uniquely and thus simplifies our description of comparative statics and welfare with multiple firms and platforms. We use the result from this proposition in the next two subsections.

5.2 Advertising with Firm-Level Competition

To isolate the effects of firm-level competition on welfare, we focus on the case of $N = 2$ and $M = 1$. Each of the firms has a single product (product 1 and product 2). The two products are *ex ante* identical, because they have an identical distribution of being high-quality or low-quality for each user. If there was any competition between the two firms at this *ex ante* stage, Bertrand competition would

[11] In particular, the intuition for multiple equilibria in this case is similar to those that arise in voting models, whereby everybody else voting for a less preferred outcome renders it a weak best response for each voter to also do so. Here, too, despite product differentiation, there can be multiple equilibria whereby each firm sets a very low price expecting the other firms to set a very low price. In voting models, sequential actions are sufficient to restore uniqueness (and do so in an order-independent manner). Here, due to the more complicated nature of the game, we need impose one more robustness refinement in order to achieve the same objective.

drive prices down to marginal cost. However, once some users start receiving additional information from advertisements, the products become horizontally differentiated—some users will have weak or strong signals about the quality of either or both of the products' quality for them.

The platform can advertise for both products, and we denote the advertising rates in plan ℓ by $\alpha_\ell^{(1)}$ and $\alpha_\ell^{(2)}$ for the two products, respectively. This implies that if a user subscribes to plan ℓ, then with probability $1 - e^{-\alpha_\ell^{(1)} T}$ she views an ad for product 1 and with probability $1 - e^{-\alpha_\ell^{(2)} T}$ she views an ad for product 2; since these events are independent, it is possible that she may view both ads or neither. At the same time, the platform charges subscription fees $\{P_\ell^*\}_{\ell=1}^2$. The platform offers simultaneous contracts $\{(\alpha_\ell^{(1)}, m_\ell^{(1)})\}_{\ell=1}^2$ and $\{(\alpha_\ell^{(2)}, m_\ell^{(2)})\}_{\ell=1}^2$ to both firms who then either accept or reject the contracts.

To build some intuition, we first discuss the fully-rational benchmark and the equilibrium prices and quantities when there is no advertising.

Fully-Rational Benchmark. Under the fully-rational benchmark ($\phi_{0,N} = \phi_0$), each agent i subscribes to the same plan and hold Bayesian beliefs that firm j's quality for her is $\theta_i^{(j)} = 1$ denoted by $\pi_S^{(j)}$. Each user will then choose $j^* \in \arg\max_j \beta\pi_i^{(j)} - p^{(j)*}$ and consume a quantity $z_i^{(j^*)^*} = \beta\pi_i^{(j^*)} - p^{(j^*)*}$ of product j^* and none of the other product, i.e., $z_i^{(-j^*)^*} = 0$. We let $\bar{p}_1^*$ and $\bar{p}_2^*$ denote the prices offered under the fully-rational benchmark in equilibrium. Note that this fully-rational benchmark may or may not involve some amount of advertising in this single plan. If there is no advertising in the single plan, then additionally, we have that prices are equal to marginal cost—$\bar{p}_1^* = \bar{p}_2^* = c$—because both products have the same *ex ante* appeal to all users.

Benchmark with No Advertising. When there is no advertising plan offered (or selected) in equilibrium, then $\pi_S^{(1)} = \pi_S^{(2)} = q$, which substantiates our above claim that *ex ante* products are identical. Therefore, without advertising, the two firms will compete à la Bertrand and in equilibrium will charge prices equal to marginal cost, i.e., $p_1^* = p_2^* = c$. In this benchmark, therefore, competition from multiple firms reduces prices and increases consumer surplus relative to the monopoly scenario. We will next see that the situation is different in the presence of digital ads.

Equilibrium in the Full Model. We now consider our full model with both sophisticated and naïve agents, in the presence of digital ads. The first consequence of digital ads is that, in general, $\pi_S^{(1)} \neq \pi_S^{(2)}$ because (i) the user may view an ad for product 1 but not product 2 or vice-versa, and (ii) the user may receive different signals from product 1's ad versus product 2's ad. Hence, there will be (partial) differentiation between the products for different segments of the population. As a consequence of this differentiation, firms regain market power and equilibrium prices will typically have $p_1^* > c$ and $p_2^* > c$, as we show next.

Proposition 9. *There exist* $\phi_1^{\mathcal{F}}, \phi_0^{\mathcal{F}}(\phi_1) \in [0,1]$ *such that*

(i) If $\phi_1 \leq \phi_1^{\mathcal{F}}$*, the platform offers one subscription-based plan with* $P^* = T - v$ *and no advertising* $\alpha^{(1)*} = \alpha^{(2)*} = 0$ *and equilibrium prices are* $p_1^* = p_2^* = c$*;*

(ii) If $\phi_1 > \phi_1^{\mathcal{F}}$ and $\phi_0 \geq \phi_0^{\mathcal{F}}(\phi_1)$, the platform offers a subscription-based plan with $P_1^* = T - v$ and no advertising $\alpha_1^{(1)*} = \alpha_1^{(2)*} = 0$, and an ad-based plan ($\alpha_2^{(1)*} + \alpha_2^{(2)*} > 0$) with product prices $\hat{p}_1^* > \bar{p}_1^*$ and $\hat{p}_2^* > \bar{p}_2^*$.

Proposition 9 extends our main characterization result, Proposition 5, to the case with multiple firms. It establishes that, in contrast to the situation without advertising, market power in the product market and prices above marginal cost can persist, despite the presence of multiple firms. In fact, it is possible that advertising is higher when there are multiple competing firms than in the monopoly case, as we demonstrate next using an example:

Example 1. Let us set ϕ_0 and ϕ_1 so that $I_S'(0) = T$, but naturally, $\alpha_S^{FB} = 0$. Advertising at rate ε reduces the maximal subscription price the platform could extract from sophisticates by a rate of only $o(\varepsilon)$. In other words, the marginal informational value that sophisticates gain from advertising is equal to the loss in entertaining content, making them nearly indifferent between no advertising and a very small amount of advertising $\varepsilon > 0$.

For a sufficiently small naive population ($\lambda < \underline{\lambda}$ for some $\underline{\lambda} < 1$), the platform with a single advertiser charges the subscription fee $P^* = T - v$, offers no advertising, and the firm charges $\bar{p}^* = (\beta q + c)/2$ (as in regime (a) of Proposition 7). This is because it cannot extract any surplus from the sophisticates via advertising anyway (Lemma 1), and with the small naive population, does not collect sizable digital ad revenues by targeting them with an ad-based plan, preferring to collect the subscription fee from them instead. The result is a fully subscription-based business model for the platform, with no advertising for a monopolist firm.

With firm-level competition, if there is no advertising, then both firms compete to offer prices $p_1^* = p_2^* = c$, leading to zero economic profits. However, each firm would be willing to pay $A_1\varepsilon$ to advertise at rate ε in both plans (for some constant $A_1 > 0$), because then it can make economic profits on the order of ε — it can charge a markup on p_1^* proportional to $\beta\pi_i^{(j)}(G)$, where $\pi_i^{(j)}(G)$ is the belief of a sophisticated agent i conditional on a good signal, at still capture market share $(q\phi_1 + (1-q)\phi_0)(1 - e^{-\varepsilon T}) \sim \varepsilon$. At the same time, the loss in subscription revenue from the sophisticates on the platform from implementing such an advertising scheme is $o(\varepsilon)$, which is necessarily less than the digital ad revenue it can generate from charging the firm. Thus, in equilibrium, there will be positive advertising in all plans offered with competition (in order to relax competitive pressures), even though there was no advertising with a monopolist. ∎

In terms of the general characterization, a major difference between Proposition 5 and Proposition 9 is that in Proposition 9, not just the rate of false positives, ϕ_0, but also the informativeness of positive signals, ϕ_1, is critical for the equilibrium business model, advertising, and prices. In essence, ϕ_0 still matters for the same reasons (digital ads are most useful to firms and the platform because they *de facto* manipulate naïve agents), but in addition, ϕ_1 needs to be sufficiently large, since otherwise digital ads are not sufficiently powerful in creating endogenous differentiation between products and do not allow firms to charge sufficiently high markups.

When both ϕ_0 and ϕ_1 are above their respective thresholds, the equilibrium resembles part (b) of Proposition 5: sophisticates, who understand that both parameters are large, find most ads to give

positive signals regardless of the true underlying $\theta_i^{(j)}$, which means their demand barely responds to signals from ads, and this makes it profitable for the platform to segregate the market and monetize its services to sophisticates based on subscription fees (and no advertisement). In contrast, naïve agents continue to be responsive to positive signals from ads and are the main targets for these ads. Because firms know that naïve agents will view their ads and have differentiated willingness to pay for these products, they charge prices above marginal cost, which then creates a negative spillovers on sophisticates, as before. The platform extracts the value of entertaining content from sophisticates using a subscription fee and extracts the value of digital ads targeting the nave agents from both firms. Consequently, as in Proposition 5, digital ads harm not just naïve agents but also sophisticates, who now have to pay higher prices for products.

We next discuss welfare in the presence of firm competition. Let $\bar{W}^*_{2,1}(\tau)$ denote the welfare of users of type $\tau \in \{S,N\}$ under the fully-rational benchmark with two firms and a single platform. Define $\hat{W}^*_{2,1}(\tau)$ analogously as as the welfare of the two types of users in the equilibrium characterized in Proposition 9. Our main welfare result in this section is:

Proposition 10.

(i) If $\phi_1 \leq \phi_1^{\mathcal{F}}$*, then* $\hat{W}^*_{2,1}(\tau) = \bar{W}^*_{2,1}(\tau)$ *for both* $\tau \in \{S,N\}$.

(ii) If $\phi_1 > \phi_1^{\mathcal{F}}$ *and* $\phi_0 \geq \phi_0^{\mathcal{F}}(\phi_1)$*, then* $\hat{W}^*_{2,1}(\tau) < \bar{W}^*_{2,1}(\tau)$ *for both* $\tau \in \{S,N\}$.

The first part of this proposition simply reiterates that when the platform opts for a subscription-based plan and there are no digital ads, the two firms will compete fiercely and drive prices down to marginal cost. This gives us the same welfare to both types of agents as in the fully-rational benchmark—which also features no digital ads and positive subscription fees—because, without digital ads, naïve and sophisticated agents behave identically and all products are identical in the eyes of all consumers, driving their prices down to marginal cost.

The second part of the proposition is more interesting. It shows that, analogously with Proposition 4, when both ϕ_0 and ϕ_1 are sufficiently large, the equilibrium involves digital ads targeted to naïve users, and these digital ads lead to higher prices—in this case restoring positive markups relative to the equilibrium without digital ads where prices are equal to marginal cost.

The intuition for higher prices is more nuanced in this case, however. In Proposition 4, digital ads increased equilibrium prices because they raised naïve users' willingness to pay for the product. Here, digital ads do not just over-inflate naïve users' valuations, but also relax competition between the two firms by generating (endogenous) differentiation. This is the reason why the standard Bertrand logic does not apply and competition does not protect naïve agents from *de facto* manipulation from firms and the platform (and this is also the reason why competition can increase the demand for digital ads from firms, as shown in Example 1). As a result of these forces, welfare is still substantially lower than in the fully-rational benchmark.

The comparative statics of Propositions 9 and 10 are also interesting. As before, a higher rate of false positives, ϕ_0, makes digital ads more likely, which reiterates the same result that digital ads are more likely to emerge when they are less informative. But now we also have digital ads being more

likely when informativeness of positive signals from ads, ϕ_1, is higher, because with low ϕ_1, digital ads are not impactful on user valuations and the subscription-based business plan becomes more likely.

5.3 Platform-Level Competition

We now discuss the implications of between-platform competition and to simplify the analysis, this time we focus on the case where there are two platforms and a single firm, i.e., $M = 2$ and $N = 1$. The two platforms simultaneously offer plans $\{\alpha^*_{1,\ell}, P_{1,\ell}\}_{\ell=1}^2$ with associated transfer m_1 to the firm and $\{\alpha^*_{2,\ell}, P_{2,\ell}\}_{\ell=1}^2$ with associated transfer m_2 to the firm. Following this stage, the firm decides which plan(s), if any, to accept. The next proposition characterizes the equilibrium in this case. We again denote the price for the unique final good in the corresponding fully-rational benchmark by $\bar{p}^*$.

Proposition 11. *There exists $\phi_1^{\mathcal{P}}$ such that:*

(i) If $\phi_1 < \phi_1^{\mathcal{P}}$, the platforms offer competing plans with no advertising and no subscription fee, and the product is priced at $p^* = \bar{p}^*$;

(ii) If $\phi_1 > \phi_1^{\mathcal{P}}$, the platforms offer two ad-based plans $\alpha_1^* = \alpha_S^{FB}$ and $\alpha_2^* \in (\alpha_S^{FB}, \hat{\alpha}^*)$ with no subscription fees, and the product is priced at $p^* = \hat{p}^*_{\mathcal{P}} > \bar{p}^*$.

The results in this proposition are even more striking than those in Proposition 9. Our results turn on the informativeness of of positive signals from ads, ϕ_1. If this parameter is below the critical threshold $\phi_1^{\mathcal{P}}$, platforms are able to extract no surplus from either the firm or the users, and completely forgo digital ads but also do not charge a subscription fee, because competition between them drives their prices down to marginal cost—which is equal to zero for the entertaining content that they offer.

In contrast, when ϕ_1 is greater than $\phi_1^{\mathcal{P}}$, digital ads on the platforms are again present. Digital ads once more target naïve users, who now interpret positive signals from ads as evidence of high quality, which over-inflates their valuation for the product. The reason why there are no subscription-fee-based plans is that each platform can steal sophisticates from the other platform by offering a subscription fee together with the first-best level of advertisements for these agents (α_S^{FB}). Equally importantly, both platforms compete for naïve users by segmenting the market between them and the sophisticates, and the naïfs again receive more frequent digital ads ($\alpha_2^* > \alpha_S^{FB}$), which is the reason why their valuation for the products becomes over-inflated and the firm becomes encouraged to charge even higher prices (analogously to the situation in Proposition 5). Platforms are able to extract surplus from naïfs, because these users have a greater willingness to pay for digital ads and once we are in an ad-based equilibrium, the platforms cannot compete against each other by cutting prices (subscription fees are equal to zero already). Although it might first appear that they can compete against each other by reducing ad intensity, this would not work because naïve agents prefer the high intensity of digital ads (because they think ads are more informative than they are in actuality), and hence no platform has an incentive to reduce the intensity of digital ads in the plan targeted at naïve agents below α_2^*.

We next study the welfare properties of this equilibrium. For this purpose, let $\bar{W}^*_{1,2}(\tau)$ once again denote the welfare of the fully-rational benchmark (now with a single firm and two platforms) for $\tau \in \{S, N\}$, and $\hat{W}^*_{1,2}(\tau)$ denote user welfare in the equilibrium of Proposition 11 for $\tau \in \{S, N\}$.

Proposition 12.

(i) If $\phi_1 < \phi_1^{\mathcal{P}}$*, then* $\hat{W}^*_{1,2}(\tau) = \bar{W}^*_{1,2}(\tau)$ *for both* $\tau \in \{S, N\}$.

(ii) If $\phi_1 > \phi_1^{\mathcal{P}}$*,* $\hat{W}^*_{1,2}(\tau) < \bar{W}^*_{1,2}(\tau)$ *for both* $\tau \in \{S, N\}$.

As in the case of firm-level competition analyzed in Propositions 9 and 10, if there are no digital ads, welfare is restored back to the fully-rational benchmark because platforms compete with each other and this limits their ability to extract surplus from users. In contrast, as soon as ϕ_1 is above the threshold $\phi_1^{\mathcal{P}}$, digital ads targeted at naïve agents reappear and this now softens the competition between platforms. In particular, as explained above, once ad-based business models are being used, the two platforms no longer undercut each other.

It is again interesting to note that the second regime, where welfare is low, is more likely when ϕ_1 is high. This is for the same reasons as in Proposition 11: when ϕ_1 is very low, digital ads are not sufficiently appealing for the firm, because they do not expand the demand for its product sufficiently, and only a subscription-based business model can be sustained in equilibrium.

6 Digital Ad Taxation

Our analysis so far has identified a fundamental market failure in platform economies with digital ads and naïve agents who may misinterpret the meaning of these ads. A natural question therefore is whether there are obvious regulatory or tax-based solutions to this market failure. This is a question we investigate in this section. While the first best allocation is not implementable, we show that the second best, where user type is private information and users will choose which plain to join themselves, can be decentralized through nonlinear digital ad taxes and product subsidies. At the same time, simple policies that just linearly tax digital advertising revenue can improve welfare in the decentralized equilibrium. For simplicity, we present all results this section for a single platform and single firm. These results generalize in a straightforward way to multiple platforms and multiple firms, as we demonstrate in Appendix C.

6.1 Second-Best User Welfare

Recall that our first-best user welfare, W_{FB}, was obtained by allowing the social planner to fully control the allocation, while observing user types. This meant that there were no "self-selection" constraints, and the planner could choose the level of digital ads different types of users would observe. This is of course unrealistic and another natural benchmark may be one where the social planner cannot observe user types.

Specifically, we define *second-best* user welfare are the welfare level when the planner chooses the menu of advertising, subscription prices, and product prices offered, (α_1, P_1, p_1) and (α_2, P_2, p_2). The

planner is the subject to the incentive compatibility constraints of every agent i. That is, agent i selects plan $\ell^*(i)$ if and only if

$$\begin{aligned}
&(1-\alpha_{\ell^*(i)})T - P_{\ell^*(i)} + q\mathbb{E}^{\tau_i}_{\pi_i}[U(z^*(\pi_i, p_{\ell^*(i)}); \theta_i = 1) - pz^*_i(\pi_i)|\alpha_{\ell^*(i)}] \\
&\quad + (1-q)\mathbb{E}^{\tau_i}_{\pi_i}[U(z^*(\pi_i, p_{\ell^*(i)}); \theta_i = 0) - pz^*_i(\pi_i)|\alpha_{\ell^*(i)}] \\
\geq &(1-\alpha_{\ell})T - P_{\ell} + q\mathbb{E}^{\tau_i}_{\pi_i}[U(z^*(\pi_i, p_{\ell}); \theta_i = 1) - pz^*_i(\pi_i)|\alpha_{\ell}] \\
&\quad + (1-q)\mathbb{E}^{\tau_i}_{\pi_i}[U(z^*(\pi_i, p_{\ell}); \theta_i = 0) - pz^*_i(\pi_i)|\alpha_{\ell}]
\end{aligned}$$

for all other plans ℓ.

As already anticipated in the discussion following Proposition 1, assigning naïve users to the advertising intensity $\alpha^{FB}_N < \alpha^{FB}_S$, as the planner wishes to do, is not incentive compatible because naïfs prefer an even higher advertising intensity than α^{FB}_S. In the second best, the planner can select different plans than the platform will offer, say with advertising intensities $(\alpha^{SB}_1, \alpha^{SB}_2)$, but has to allow users themselves decide among these plans. Our next result characterizes the second best.

Proposition 13. *The second best involves a single plan with advertising intensity $\alpha^{SB} \in [\alpha^{FB}_N, \alpha^{FB}_S]$. Whenever $\alpha^{FB}_S > 0$, the second-best level of welfare is less than the first best, but at least that of the fully-rational benchmark and the base case: $W_{FB}(\tau) > W_{SB}(\tau) \geq W_{\text{fully-rational}}(\tau) = W_{\text{base}}(\tau)$ for both $\tau \in \{S, N\}$.*

The intuition for Proposition 13 is closely related to our discussion of the first best in Proposition 1. Ideally, the planner would like to offer a menu with lower ad intensity for naïve agents than what will be offered to sophisticates, α^{FB}_S. However, naïfs actually prefer an even higher ad intensity than α^{FB}_S, because in their assessment ads are more informative than sophisticates consider them to be and this makes it more attractive for naïve agents to trade-off a little less entertainment and a little more ads starting in the neighborhood of α^{FB}_S. But this reasoning also suggests that whenever the planner offers a menu with different options, naïve users will have a preference for the one with greater ad intensity—which is the exact opposite of what the planner would like. Hence the planner is forced to choose a single plan. Because this plan will cater to both naïve and sophisticated agents, its ad intensity is intermediate between α^{FB}_N and α^{FB}_S, trading off the utility of naïve and sophisticated agents.

6.2 Decentralizing the Second Best

To decentralize the second best, we consider a nonlinear taxation policy on digital advertising quantities. Let $\zeta(\alpha_1, \alpha_2)$ denote a tax on digital ad quantity α_1 in plan 1 and quantity α_2 in plan 2. We consider a separable policy where $\zeta(\alpha_1, \alpha_2) = \tilde{\zeta}(\alpha_1) + \tilde{\zeta}(\alpha_2)$, which does not tax advertising at or below α^{SB}, but taxes advertising at intensities higher than α^{SB} at rate $\mu > 0$. In other words,

$$\tilde{\zeta}(\alpha) = \begin{cases} 0, & \text{if } \alpha \leq \alpha^{SB} \\ \mu(\alpha - \alpha^{SB}), & \text{if } \alpha > \alpha^{SB} \end{cases}$$

At the same time, we offer a per-unit product subsidy δ to the firm and a subscription-fee subsidy η to platform. For the firm, we subsidize the cost of the good by δ, whereas for the platform we provide a per-user subsidy of η for every user who subscribes, conditional on no subscription fee being charged (and zero subsidy otherwise). Formally, we provide a $\delta \int_0^1 z_i^* \, di$ subsidy of the product to the firm and for the platform give a subsidy

$$\begin{cases} \eta \left(\int_0^1 x_{i,1}^* \, di + \int_0^1 x_{i,2}^* \, di \right), & \text{if } P_1^* = 0 \text{ and } P_2^* = 0 \\ \eta \int_0^1 x_{i,1}^* \, di, & \text{if } P_1^* = 0 \text{ and } P_2^* > 0 \\ \eta \int_0^1 x_{i,2}^* \, di, & \text{if } P_1^* > 0 \text{ and } P_2^* = 0 \\ 0, & \text{if } P_1^* > 0 \text{ and } P_2^* > 0. \end{cases}$$

Our next result characterizes the corresponding decentralized equilibrium under such a policy.

Proposition 14. *There exists $\bar{\mu} > 0$, $\bar{\eta} > 0$, and $\delta^* > 0$ such that if the platform's digital ad tax policy satisfies $\mu > \bar{\mu}$, the firm subsidy is given by δ^*, and the platform subsidy satisfies $\eta > \bar{\eta}$, then the decentralized equilibrium implements the second-best.*

The intuition for Proposition 14 is straightforward. In a decentralized equilibrium before the policy, the platform prefers to advertise to naifs at a rate higher than α^{SB} because naifs prefer a higher advertising intensity than the one the planner would choose for them. The higher ad rate thus allows the platform to collect more subscription revenue from these agents. However, if there is a sufficiently high tax on ad quantities larger than α^{SB}, then the cost of the tax will more than offset the additional subscription fees collected. The subsidy to the firm then ensures that the firm does not use the platform's advertising as a way to extract more surplus from consumers using a higher price. In particular, the subsidy guarantees that the firm will charge marginal cost in equilibrium, and all of the informational rents from viewing ads stays with the user.

6.3 Flat Digital Ad Tax

In the previous subsection we focused on nonlinear taxes and subsidies that use information on digital ads quantities and prices, which enabled us to implement the second-best allocation. In this subsection we show that the simpler intervention of a proportional tax on digital ad revenues improves welfare (though does not restore to the second best).

We define a flat digital ad tax as a linear tax on digital ad revenues. In particular, letting m denote the digital ad revenue (transfer from the firm) and $\mathcal{S}$ denote the total subscription revenue collected from the platform, $\mathcal{S} = \sum_{\ell=1}^{2} \int_0^1 P_i x_{i,\ell}^* \, di$, we can write the total income of the platform as $m + \mathcal{S}$. A digital ad tax is one that results in post-tax income $\varphi(m) + \mathcal{S}$ for some function $\varphi(m) < m$. In this case, a flat digital add tax corresponds to the special case where $\varphi(m) = (1 - \gamma)m$ for some linear tax rate $\gamma \in (0, 1)$. Alternatively, a non-linear $\varphi(m)$ would correspond to a non-linear tax scheme.

Proposition 15. *There exists $\bar{\gamma} < 1$ such that a flat digital ad tax with $\gamma > \bar{\gamma}$ improves welfare starting from the decentralized equilibrium.*

Proposition 15 establishes that a sufficiently large digital ad tax improves welfare (without any other policy instrument used). Clearly, this does not achieve the second best we characterized earlier in this section, but it still provides a useful result because it is a simple and easy-to-implement policy.

7 Conclusion

Digital advertising has become the dominant business model for online platforms, reaching revenue of nearly half a trillion dollars in 2022.[12] Platforms have also enriched their offerings by combining subscription-based and advertisement-based plans. Despite the growing importance of the ecosystem defined by digital ads and growing concerns that the "attention economy" created by the desire of platforms to increase the profitability of digital ads has led to mental health problems, digital addiction and polarization (references), little is known about the welfare consequences of digital advertisement and how it affects user beliefs and demand, and via these channels, the prices that users face for other goods and services.

Our paper is a first attempt to explore these issues. We developed a parsimonious two-sided platform model, where an online platform brings together users and a firm wishing to sell a product to users. The platform offers both entertaining content and potentially informative ads about the match-specific quality of the product marketed by the firm. More ads mean less time for enjoying the content on the platform, and hence, all else equal, users would prefer less ads. Nevertheless, users also value the information that they get from the ads.

The main non-standard feature we introduce is that while some users are sophisticated and understand the exact data generating process for signals from ads, some users are naïve and under-estimate the probability with which a low-quality product still generates a positive signal. We interpret this under-estimation to be due to both the naïveté of some users (which can be affected by the salience of the ads) and to a lack of understanding that the ads are being specifically targeted and tailored for them on the basis of their personal data, which can make the advertised products appear more appealing or more favorable-looking than there truly are.

This pattern has a number of important and to the best of our knowledge novel implications. First, naïve agents will have a greater demand for digital ads than sophisticated users because they think that the ads are more informative than they truly are. Second, naïve users will be *de facto* manipulated by digital ads, because a higher digital ad intensity means a greater likelihood that naïve users will over-estimate the quality of the product. Third, as a result of these forces, targeting digital ads to naïve users is more profitable than targeting sophisticates. In fact, in our baseline model with linear-quadratic utility, expected purchases from sophisticated users do not change when they are exposed to informative ads, whereas expected purchases from naïve users increases because of their over-estimation of the quality of the product after they view digital ads.

These observations are at the root of the systemic inefficiency of the decentralized equilibrium in this model. Unless digital ads are considered to be very uninformative by both types of agents, the

[12] For exact figures, please see https://www.globenewswire.com/en/news-release/2022/09/28/2524217/28124/en/Global-Digital-Advertising-and-Marketing-Market-to-Reach-786-2-Billion-by-2026-at-a-CAGR-of-13-9.html.

equilibrium involves market segmentation between naïve and sophisticated users. Sophisticates are either left out of the platform (when we do not allow the platform to offer a menu of plans) or they sign up for a plan that involves a subscription fee, while naïfs are assigned to an ad-based plan without a subscription fee. This latter plan has very high ad intensity, targeting specifically these naïve users. When naïve agents sign up for such a plan, this increases the sales of the firm and also enables it to charge a higher price. The resulting greater profits are clawed back by the platform from the firm.

We evaluate the welfare of users by looking at their *ex post* utility, which depends on the actual quality of the product. While, at the interim stage where they see the ads, naïve users have an inflated assessment of the informativeness of digital ads and, consequently, the quality of the product, their *ex post* utility is simply a function of the quality of the product they consume. This implies that digital ads have a first-order welfare cost for naïve users. Notably, the misspecified model of the naïve users only applies to how they interpret signals from digital ads, so with a subscription-based model, these welfare costs are not present.

Even though sophisticates are not misled by digital ads and do not sign up for the high ad intensity plans, ads targeted at naïve users also have welfare costs for sophisticates. This is because, when the firm knows that it will be able to target naïfs, it prefers to charge a higher price, and this price is also paid by sophisticated users. Consequently, digital ads make both naïve and sophisticated users worse off.

We show that all of these results generalize to environments in which there are multiple platforms and multiple firms. Beyond this generalization, our analysis reveals that digital ads have an important role in softening the competition between firms and platforms. Without digital ads, firms would compete a la Bertrand, driving their prices to marginal cost. Digital ads enable endogenous differentiation of their products—users that see positive signals about the quality of these products have different valuations than those who do not, and naïve users will have a particularly distorted evaluation. This endogenous differentiation breaks Bertrand competition leads to equilibrium markups. In fact, competition increases the desire to target naïve users, because this is a way of escaping the competitive pressure from other firms. Consequently, the demand for digital ads could be even higher under competition. Similarly, digital ads also relax competition between platforms. Without naïve agents, the platforms would compete by making their offerings cheaper and more attractive to users. However, because digital ads appear more informative to naïve users, platforms generally have no incentive to reduce their digital ads, and the same type of market segmentation we saw with the monopoly platform occurs even when there are multiple competing platforms.

We also explored various policy options to counteract these systemic inefficiencies. First-best allocations where a planner can directly control the amount of digital ads served tonight even sophisticated agents cannot generally be implemented because the planner or policy authorities do not observe who is naïve and who is sophisticated, and naïve agents have a greater willingness to consume ads than sophisticates, because they think that such ads are more informative than they truly are. Nevertheless, we show that a second-best allocation, where the social planner chooses entertainment and advertisement diets and resulting product demands, is easy to characterize and can be implemented using nonlinear taxation and subsidy schemes. Even more simply, a flat digital ad

tax (on digital ad revenue) can always improve welfare starting from an equilibrium in which there is an ad-based plan. Both of these results leverage the fact that the equilibrium often features excessive ad intensity, over-inflating the valuation of naïve users and inducing higher product market prices, reducing the welfare of both types of users. By taxing revenues from digital ads, the planner makes it more attractive for platforms to monetize through subscription fees and thus reduce the excessive digital ad intensity.

We view our paper as a first step in the exploration of the positive and normative implications of new business models and information interactions that have become important over the last two decades. In this context, there are several interesting topics we have not touched on and many promising avenues for future research. Here we briefly list a few of these research directions.

1. In a first attempt to explore these issues, we abstracted from other social consequences of digital ads, including those related to mental health problems and digital addiction (Lukianoff and Haidt (2019) and Wu (2017)). An interesting direction for future research would be to model and incorporate some of these issues and see how competition between firms and platforms and informational exchange between platforms and users influences these social consequences.

2. Relatedly, we took the content offered by platforms as given. How the platform monetizes itself may have first-order implications for the kind of content that it offers. In fact, some of the major concerns mentioned above are rooted in the fact that monetizing data via digital ads becomes more profitable when people spend more time on the platform, which can encourage the platform to offer content that is more addictive or emotionally triggering in order to increase user engagement (this is in fact the argument for digital ad taxes in Johnson and Acemoglu (2023)). One attempt to study these questions is Acemoglu et al. (2023c), where platform algorithms modify the degree of homophily by political beliefs in order to affect engagement, which in turn has first-order implications for the spread of misinformation. Similar issues may become even more important when one considers a broader menu of content that can be offered, such as low-quality clickbait (Immorlica et al. (2024)), envy-generating content from friends (Beknazar-Yuzbashev et al. (2022)), or politically-provocative content (Mostagir and Siderius (2023)).

3. Once agents are naïve, there may also be additional strategies platforms can utilize for extracting surplus from them. One such possibility is explored in Acemoglu et al. (2023b), where the platform can engage in behavioral manipulation by steering users towards products where they are more likely to overestimate quality (either statically or dynamically as in Acemoglu et al. (2023b)). A more general treatment of these issues in the context of two-sided platforms would be an interesting area for future research.

4. We simplified the analysis by ignoring how digital ads are constructed and targeted. A more in-depth analysis of this question requires us to study how user data is leveraged to tailor and target ads, and this opens the door to a broader discussion of how data access should be regulated, who owns the data generated in the process of social media interactions unrelated to one's own social network, and whether individuals can and should control their own data—especially taking

into account both data externalities and other aspects of their naïvety or lack of information (Acemoglu et al. (2022), Acemoglu et al. (2023a), and Mostagir and Siderius (2022)).

5. Another major simplification was achieved by abstracting from social networks. Individuals often like to join platforms where their friends and acquaintances are active. Introducing this element in the competition between platforms and the business model choices of platforms would be another interesting direction for future research.

6. Last but not least, our exploration raises a number of new empirical questions about how different platform plans/offerings influence product market competition and prices. An important direction for future research is to explore both some of the foundational assumptions we have imposed (such as how naïve individuals process information from ads) and these new implications.

References

Acemoglu, Daron, Alireza Fallah, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar (2023a), "How good are privacy guarantees? platform architecture and violation of user privacy." Technical report, National Bureau of Economic Research.

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asu Ozdaglar (2022), "Too much data: Prices and inefficiencies in data markets." *American Economic Journal: Microeconomics*, 14, 218–256.

Acemoglu, Daron, Ali Makhdoumi, Azarakhsh Malekian, and Asuman Ozdaglar (2023b), "A model of behavioral manipulation." Technical report, National Bureau of Economic Research.

Acemoglu, Daron, Asuman Ozdaglar, and James Siderius (2023c), "A model of online misinformation." *Review of Economic Studies*, rdad111.

Allcott, Hunt, Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow (2020), "The Welfare Effects of Social Media." *American Economic Review*, 110, 629–676.

Allcott, Hunt, Matthew Gentzkow, and Lena Song (2022), "Digital Addiction." *American Economic Review*, 112, 2424–2463.

Beknazar-Yuzbashev, George, Rafael Jiménez Durán, Jesse McCrosky, and Mateusz Stalinski (2022), "Toxic content and user engagement on social media: Evidence from a field experiment." *Available at SSRN 4307346*.

Bennett, Colin and Jesse Gordon (2020), "Understanding the "Micro" in Micro-Targeting: An Analysis of Facebook Digital Advertising in the 2019 Federal Canadian Election." *Available at SSRN 3589687*.

Berger, Benedikt, Christian Matt, Dennis M. Steininger, and Thomas Hess (2015), "It Is Not Just About Competition with "Free": Differences Between Content Formats in Consumer Preferences and Willingness to Pay." *Journal of Management Information Systems*, 32, 105–128.

Blackwell, David (1953), "Equivalent comparisons of experiments." *The annals of mathematical statistics*, 265–272.

Budzinski, Oliver, Sophia Gaenssle, and Nadine Lindstädt-Dreusicke (2021), "The battle of YouTube, TV and Netflix: an empirical analysis of competition in audiovisual media markets." *SN Business & Economics*, 1, 116.

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

Chyi, Hsiang Iris (2005), "Willingness to Pay for Online News: An Empirical Study on the Viability of the Subscription Model." *Journal of Media Economics*, 18, 131–142.

Chyi, Hsiang Iris and Yee Man Margaret Ng (2020), "Still Unwilling to Pay: An Empirical Analysis of 50 U.S. Newspapers' Digital Subscription Results." *Digital Journalism*, 8, 526–547.

Cinelli, Matteo, Walter Quattrociocchi, Alessandro Galeazzi, Carlo Michele Valensise, Emanuele Brugnoli, Ana Lucia Schmidt, Paola Zola, Fabiana Zollo, and Antonio Scala (2020), "The COVID-19 social media infodemic." *Scientific Reports*, 10, 16598.

Danciu, Victor et al. (2014), "Manipulative marketing: persuasion and manipulation of the consumer through advertising." *Theoretical and Applied Economics*, 21, 591.

De Jans, Steffi, Dieneke Van de Sompel, Liselot Hudders, and Veroline Cauberghe (2019), "Advertising targeting young children: an overview of 10 years of research (2006–2016)." *International Journal of Advertising*, 38, 173–206.

De Reuver, Mark, Carsten Sørensen, and Rahul C Basole (2018), "The digital platform: a research agenda." *Journal of information technology*, 33, 124–135.

Deng, Yiting and Carl F. Mela (2018), "TV viewing and advertising targeting." *Journal of Marketing Research*, 55, 99–118.

Dixit, Avinash and Victor Norman (1978), "Advertising and welfare." *The Bell Journal of Economics*, 1–17.

Esponda, Ignacio and Demian Pouzo (2016), "Berk–Nash Equilibrium: A Framework for Modeling Agents With Misspecified Models." *Econometrica*, 84, 1093–1130.

Eyal, Nir (2014), *Hooked: How to build habit-forming products*. Penguin.

Flew, Terry (2021), "Willingness to Pay: News Media." *International Institute of Communication*.

Galán, Aritz Arrate, José González Cabañas, Ángel Cuevas, María Calderón, and Rubén Cuevas Rumin (2019), "Large-scale analysis of user exposure to online advertising on facebook." *IEEE access*, 7, 11959–11971.

Golab-Andrzejak, Edyta (2023), "The impact of generative ai and chatgpt on creating digital advertising campaigns." *Cybernetics and Systems*, 1–15.

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

Gupta, Aastha (2023), "Deceptive advertising, regulation and naive consumers." *International Journal of Industrial Organization*, 91, 103026.

Hattori, Keisuke and Keisaku Higashida (2012), "Misleading advertising in duopoly." *Canadian Journal of Economics/Revue canadienne d'économique*, 45, 1154–1187.

Immorlica, Nicole, Meena Jagadeesan, and Brendan Lucier (2024), "Clickbait vs. quality: How engagement-based optimization shapes the content landscape in online platforms." In *Proceedings of the ACM on Web Conference 2024*, 36–45.

Johnson, Simon and Daron Acemoglu (2023), *Power and progress: Our thousand-year struggle over technology and prosperity*. Hachette UK.

Kruikemeier, Sanne, Minem Sezgin, and Sophie C Boerman (2016), "Political microtargeting: relationship between personalized advertising on facebook and voters' responses." *Cyberpsychology, Behavior, and Social Networking*, 19, 367–372.

Lukianoff, Greg and Jonathan Haidt (2019), *The coddling of the American mind: How good intentions and bad ideas are setting up a generation for failure*. Penguin.

Marwick, Alice E. and Rebecca Lewis (2017), "Media manipulation and disinformation online." *Data & Society Research Institute*.

Meurer, Michael and Dale O. Stahl (1994), "Informative advertising and product match." *International Journal of Industrial Organization*, 12, 1–19.

Moldovanu, Benny and Eyal Winter (1995), "Order independent equilibria." *Games and Economic Behavior*, 9, 21–34.

Mostagir, Mohamed and James Siderius (2022), "Naive and bayesian learning with misinformation policies." Technical report, Massachusetts Institute of Technology.

Mostagir, Mohamed and James Siderius (2023), "When should platforms break echo chambers?" Technical report, Massachusetts Institute of Technology.

Piccolo, Salvatore, Piero Tedeschi, and Giovanni Ursino (2018), "Deceptive advertising with rational buyers." *Management Science*, 64, 1291–1310.

Richter, Felix (2019), "The generation gap in tv consumption." *Statista*.

Rochet, Jean-Charles and Jean Tirole (2006), "Two-sided markets: a progress report." *The RAND journal of economics*, 37, 645–667.

Sato, Susumu (2019), "Freemium as optimal menu pricing." *International Journal of Industrial Organization*, 63, 480–510.

Sherman, Ryland and David Waterman (2016), "The economics of online video entertainment." *Handbook on the Economics of the Internet*, 458–474.

Simchon, Almog, Matthew Edwards, and Stephan Lewandowsky (2024), "The persuasive effects of political microtargeting in the age of generative artificial intelligence." *PNAS nexus*, 3, pgae035.

Susser, Daniel and Vincent Grimaldi (2021), "Measuring automated influence: between empirical evidence and ethical values." In *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 242–253.

Tirole, Jean (1988), *The Theory of Industrial Organization*, 1st edition edition. The MIT Press, Cambridge, Mass.

Tran, Trang P (2017), "Personalized ads on facebook: An effective marketing tool for online marketers." *Journal of retailing and consumer services*, 39, 230–242.

Twenge, Jean M., Gabrielle N. Martin, and Brian H. Spitzberg (2019), "Trends in U.S. Adolescents' media use, 1976–2016: The rise of digital media, the decline of TV, and the (near) demise of print." *Psychology of Popular Media Culture*, 8, 329–345.

Vock, Marlene, Willemijn van Dolen, and Ko de Ruyter (2013), "Understanding Willingness to Pay for Social Network Sites." *Journal of Service Research*, 16, 311–325.

Wu, Tim (2017), *The attention merchants: The epic scramble to get inside our heads*. Vintage.

Zennyo, Yusuke (2020), "Freemium competition among ad-sponsored platforms." *Information Economics and Policy*, 50, 100848.

A Proofs

Lemma A.1. *For type τ_i, the (interim) informational value is given by $I_{\tau_i}(\alpha) = q\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=1}[\pi_i(\beta^2 - \beta^2\pi_i/2)\,|\,\alpha] - (1-q)\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=0}[(\beta\pi_i)^2/2\,|\,\alpha] - (\beta q)^2/2$. In particular, $I_{\tau_i}(\alpha)$ does not depend on price p.*

Proof. Note that by definition,

$$\begin{aligned}
I_{\tau_i}(\alpha) =& q\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=1}[U(C_i(\pi_i,p^*);\,\theta_i=1) - p^*C_i(\pi_i,p^*)\,|\,\alpha] \\
&+ (1-q)\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=0}[U(C_i(\pi_i,p^*);\,\theta_i=0) - p^*C_i(\pi_i,p^*)\,|\,\alpha] - (U(C_i(q,p^*);\,\theta_i=q) - p^*C_i(q,p^*))\,.
\end{aligned}$$

Notice that $C_i(\pi,p^*) = \max_{z_i}\pi_i\beta z_i - z_i^2/2 - pz_i = \pi_i\beta - p^*$. Thus,

$$\begin{aligned}
I_{\tau_i}(\alpha) =& q\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=1}[(\pi_i\beta-p^*)\beta - (\pi_i\beta-p^*)^2/2 - p^*(\pi_i\beta-p^*)] + (1-q)\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=0}[-(\pi_i\beta-p^*)^2/2 - p^*(\pi_i\beta-p^*)] \\
&- q\beta(q\beta-p^*) + (q\beta-p^*)^2/2 + p^*(q\beta-p^*) \\
=& q\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=1}[\pi_i(\beta^2-\beta^2\pi_i/2)] - q\beta p^* + q(p^*)^2/2 - (1-q)\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=0}[(\beta\pi_i)^2/2] + (1-q)(p^*)^2/2 \\
&+ q\beta p^* - (\beta q)^2/2 - (p^*)^2/2 \\
=& q\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=1}[\pi_i(\beta^2-\beta^2\pi_i/2)] - (1-q)\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=0}[(\beta\pi_i)^2/2] - (\beta q)^2/2,
\end{aligned}$$

where for simplicity of notation we have dropped the dependence on α. ∎

A.1 Proofs from Section 2

Proof of Proposition 1. The first-best user welfare for sophisticates occurs where $I_S(\alpha) + (1-\alpha)T$ is maximized, which is attained at some value we denote by α_S^{FB}. Let us denote by $\Delta(\alpha,c)$ the surplus loss associated with consumption under naïve beliefs:

$$\begin{aligned}
\Delta(\alpha,c) =& \Bigg[q\underbrace{\mathbb{E}_{\pi\sim F_{S1}(\alpha)}\left[U(z^*(\pi,c);\,\theta_i=1) - cz^*(\pi,p^*)\right]}_{\text{Sophisticate Consumption Utility Conditional on }\theta_i=1} + (1-q)\underbrace{\mathbb{E}_{\pi\sim F_{S0}(\alpha)}\left[U(z^*(\pi,c);\,\theta_i=0) - cz^*(\pi,c)\right]}_{\text{Sophisticate Consumption Utility Conditional on }\theta_i=0}\Bigg] \\
&- \Bigg[q\underbrace{\mathbb{E}_{\pi\sim F_{N1}(\alpha)}\left[U(z^*(\pi,c);\,\theta_i=1) - cz^*(\pi,p^*)\right]}_{\text{naïve Consumption Utility Conditional on }\theta_i=1} + (1-q)\underbrace{\mathbb{E}_{\pi\sim F_{N0}(\alpha)}\left[U(z^*(\pi,c);\,\theta_i=0) - cz^*(\pi,c)\right]}_{\text{naïve Consumption Utility Conditional on }\theta_i=0}\Bigg] \\
=& e^{-\alpha T}\cdot 0 + (1-e^{-\alpha T})\Big[q\phi_1\Big(U(z^*(\pi^S(G),c);\,\theta_i=1) - cz^*(\pi^S(G),p^*) \\
&- U(z^*(\pi^N(G),c);\,\theta_i=1) + cz^*(\pi^S(G),p^*)\Big) \\
&+ q(1-\phi_1)\left(U(z^*(\pi^S(B),c);\,\theta_i=1) - cz^*(\pi^S(B),p^*) - U(z^*(\pi^N(B),c);\,\theta_i=1) + cz^*(\pi^N(B),p^*)\right) \\
&+ (1-q)\phi_0\left(U(z^*(\pi^S(G),c);\,\theta_i=0) - cz^*(\pi^S(G),p^*) - U(z^*(\pi^N(G),c);\,\theta_i=0) + cz^*(\pi^N(G),p^*)\right) \\
&+ (1-q)(1-\phi_0)\left(U(z^*(\pi^S(B),c);\,\theta_i=0) - cz^*(\pi^S(G),p^*) - U(z^*(\pi^N(B),c);\,\theta_i=0) + cz^*(\pi^N(B),p^*)\right)\Big]
\end{aligned}$$

It remains just to show that $\Delta(\alpha,c)$ is increasing in α; then, the maximizer of $I_S(\alpha) + (1-\alpha)T - \Delta(\alpha,c)$ must occur at some $\alpha_N^{FB} \le \alpha_S^{FB}$. Of course, the difference in the previous expression is an $\exp(-\alpha T)$-

linear combination of 0 and a strictly positive value, given by the welfare loss when naïve agents see an ad relative to their sophisticated counterparts, with the latter maximizing ex-post welfare conditional on ad viewership by definition. Moreover, the $\exp(-\alpha T)$-linear combination puts weight $e^{-\alpha T}$ on 0 and $1-e^{-\alpha T}$ weight on a strictly positive value, which is therefore increasing in α.

Finally, note that without the platform the agent gets v and faces product price $p^*_{\text{base}} \geq c$. By assumption that $T > v$, we know that there exists some value of α such that $I_S(\alpha) + (1-\alpha)T > v$, and moreover that $I_S(\alpha^{FB}_S) + (1-\alpha^{FB}_S)T > v$, so $W_{FB}(S) > W_{\text{base}}(S)$. For naïve agents, we need only consider $I_S(\alpha) + (1-\alpha)T - \Delta(\alpha, c)$ as before. Because naïve agents make the same consumption decision as sophisticates when $\alpha = 0$, we know that $I_S(0) + T - \Delta(0,c) = T > v$. Thus, once again because there exists a value of α where $I_S(\alpha) + (1-\alpha)T - \Delta(\alpha, c) > v$, we know that $I_S(\alpha^{FB}_N) + (1-\alpha^{FB}_N)T - \Delta(\alpha^{FB}_N, c) > v$, and thus $W_{FB}(N) > W_{\text{base}}(N)$. ∎

A.2 Proofs from Section 3

Proof of Lemma 1. This follows immediately from the Martingale property of Bayesian beliefs. Let us denote by $\bar{\pi}_F$ the expected belief under distribution F, i.e., $\bar{\pi}_F = \mathbb{E}_{\pi\sim F}[\pi]$. Then,

$$
\begin{aligned}
&q\bar{\pi}_{F_{S1}(\alpha)} + (1-q)\bar{\pi}_{F_{S0}(\alpha)} \\
&= q\left(e^{-\alpha T}q + \phi_1(1-e^{-\alpha T})\frac{\phi_1 q}{\phi_1 q + \phi_0(1-q)} + (1-\phi_1)(1-e^{-\alpha T})\frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_0)(1-q)}\right) \\
&+ (1-q)\left(e^{-\alpha T}q + \phi_0(1-e^{-\alpha T})\frac{\phi_1 q}{\phi_1 q + \phi_0(1-q)} + (1-\phi_0)(1-e^{-\alpha T})\frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_0)(1-q)}\right) \\
&= q.
\end{aligned}
$$

Thus, $\int_{\tau_i=S}(p^*-c)z_i^*\,di = \int_{\tau_i=S}(p^*-c)(\beta q - p^*)\,di$, which is independent of α. ∎

Proof of Lemma 2. By the fact that $\pi_i^S|\theta_i \preceq_{FOSD} \pi_i^N|\theta_i$, we know that $q\bar{\pi}_{F_{N1}(\alpha)} + (1-q)\bar{\pi}_{F_{N0}(\alpha)} > q$ for $\alpha > 0$. To show that Π^*_N is increasing in α, we note that

$$
\begin{aligned}
&q\bar{\pi}_{F_{N1}(\alpha)} + (1-q)\bar{\pi}_{F_{N0}(\alpha)} \\
&= q\left(e^{-\alpha T}q + \phi_1(1-e^{-\alpha T})\frac{\phi_1 q}{\phi_1 q + \phi_{0,N}(1-q)} + (1-\phi_1)(1-e^{-\alpha T})\frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_{0,N})(1-q)}\right) \\
&+ (1-q)\left(e^{-\alpha T}q + \phi_0(1-e^{-\alpha T})\frac{\phi_1 q}{\phi_1 q + \phi_{0,N}(1-q)} + (1-\phi_0)(1-e^{-\alpha T})\frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_{0,N})(1-q)}\right)
\end{aligned}
$$

Of course, by virtue of $\phi_{0,N} < \phi_0$, we know that

$$
\begin{aligned}
&q\phi_1\frac{\phi_1 q}{\phi_1 q + \phi_{0,N}(1-q)} + q(1-\phi_1)\frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_{0,N})(1-q)} \\
&+ (1-q)\phi_0\frac{\phi_1 q}{\phi_1 q + \phi_{0,N}(1-q)} + (1-q)(1-\phi_0)\frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_{0,N})(1-q)}
\end{aligned}
$$

$$>q\phi_1\frac{\phi_1 q}{\phi_1 q+\phi_0(1-q)}+q(1-\phi_1)\frac{(1-\phi_1)q}{(1-\phi_1)q+(1-\phi_0)(1-q)}$$
$$+(1-q)\phi_0\frac{\phi_1 q}{\phi_1 q+\phi_0(1-q)}+(1-q)(1-\phi_0)\frac{(1-\phi_1)q}{(1-\phi_1)q+(1-\phi_0)(1-q)}=q$$

which implies that $q\bar{\pi}_{F_{S1}(\alpha)}+(1-q)\bar{\pi}_{F_{S0}(\alpha)}=e^{-\alpha T}q+(1-e^{-\alpha T})\chi$ for some $\chi>q$. Thus, we see this expression is increasing in α and that $\int_{\tau_i=N}(p^*-c)z_i^*\,di=\int_{\tau_i=N}(p^*-c)(e^{-\alpha T}q+(1-e^{-\alpha T})\chi-p^*)$, which is increasing in α for the profit-maximizing $p^*=\arg\max_p(p-c)(\lambda\beta q+(1-\lambda)\beta(e^{-\alpha T}q+(1-e^{-\alpha T})\chi)-p)$, provided that $I_N(\alpha)+(1-\alpha)T\geq v$ (the naïve agent's participation constraint is met). ∎

Proof of Lemma 3. Since both agents are Bayesian under their subjective models, we know by Blackwell's theorem we can rank $I_N(\alpha)$ and $I_S(\alpha)$ if we can show the naïve agent's subjective model is more informative than the sophisticate's subjective model in the Blackwell order (Blackwell (1953)). Consider a signal s_i generated according to the naïve agent's subjective model and consider an alternative signal generation process that is strictly less informative, constructed as follows. If $s_i=G$ given $\theta_i=0$, retain the signal $s_i'=G$. If $s_i=B$ and $\theta_i=0$, then with probability $\frac{1-\phi_0}{1-\phi_{0,N}}$, retain the signal as $s_i'=B$, otherwise switch the signal to $s_i'=G$. Clearly this construction is a garbling process that makes the sophisticates' signal structure Blackwell dominated by the naïve agents' signal structure. Consequently, $I_N(\alpha)>I_S(\alpha)$ for all $\alpha>0$.

Using Lemma A.1, we note that expanding $I_{\tau_i}(\alpha)=q\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=1}[\pi_i(\beta^2-\beta^2\pi_i/2)\,|\,\alpha]-(1-q)\mathbb{E}^{\tau_i}_{\pi_i|\theta_i=0}[(\beta\pi_i)^2/2\,|\,\alpha]-(\beta q)^2/2$ leaves us with $I_N(\alpha)$ and $I_S(\alpha)$ of the form:

$$I_N(\alpha)=(1-e^{-\alpha T})\chi_N,$$
$$I_S(\alpha)=(1-e^{-\alpha T})\chi_S,$$

for some constants $\chi_N>\chi_S$ that depend on model primitives (e.g., ϕ_0, ϕ_1) but not on α. Concavity and monotone increasing in α follow immediately. To observe $\arg\max I_N(\alpha)+(1-\alpha)T>\arg\max I_S(\alpha)+(1-\alpha)T$, we note that $I_N'^{-\alpha T}=\frac{1}{\chi_N}$, which has an intersection point that occurs later than $I_S'^{-\alpha T}=\frac{1}{\chi_S}$, because $1/\chi_N<1/\chi_S$ and $e^{-\alpha T}$ is a decreasing function in α. ∎

Proof of Lemma 4. For $t=1$ and $t=2$, we can collapse the decision problem of the platform and firm into one joint decision problem, because the platform gives a take-it-or-leave-it offer. In particular, the platform will solve a set of maximization problems with various participation constraints, and then compare the outcomes. Then the platform solves:

- The optimal product price p^* when no agent participates on the platform, and learns nothing about her preferences.
- The optimal product price p^* and advertising intensity α^* when only the sophisticated agents participate on the platform. Sophisticates learn about their preferences from ads, whereas naïves operate under their prior.
- The optimal product price p^* and advertising intensity α^* when only the naïve agents participate

on the platform. naïves learn about their preferences from ads, whereas sophisticates operate under their prior.

- The optimal product price p^* and advertising intensity α^* when both sophisticates and naïves participate on the platform. All agents learn about their preferences from ads.

$$\mathcal{A}_1 \equiv \max_{\alpha,p} \quad \underbrace{(p-c)}_{\text{Marginal Profit}} \cdot \underbrace{(\beta q - p)}_{\text{Consumer Demand without Ads}} \qquad \text{(No user participation)}$$

$$\mathcal{A}_2 \equiv \max_{\alpha,p} \quad \lambda(p-c)\underbrace{(\beta q \bar{\pi}_{F_{S1}(\alpha)} + \beta(1-q)\bar{\pi}_{F_{S0}(\alpha)} - p)}_{\text{Sophisticate Demand with Ads}} + (1-\lambda)(p-c)(\beta q - p) \qquad \text{(Sophisticates participate)}$$

$$\text{subject to} \quad \underbrace{I_S(\alpha) + (1-\alpha)T - v}_{\text{Sophisticate Surplus from Platform Participation}} \geq 0$$

$$\mathcal{A}_3 \equiv \max_{\alpha,p} \quad (1-\lambda)(p-c)\underbrace{(\beta q \bar{\pi}_{F_{N1}(\alpha)} + \beta(1-q)\bar{\pi}_{F_{N0}(\alpha)} - p)}_{\text{naïve Demand with Ads}} + \lambda(p-c)(\beta q - p) \qquad \text{(naïves participate)}$$

$$\text{subject to} \quad \underbrace{I_N(\alpha) + (1-\alpha)T - v}_{\text{naïve Surplus from Platform Participation}} \geq 0$$

$$\begin{aligned} \mathcal{A}_4 \equiv \max_{\alpha,p} \quad & (p-c)(\lambda\beta q \bar{\pi}_{F_{S1}(\alpha)} + \lambda\beta(1-q)\bar{\pi}_{F_{S0}(\alpha)} \\ & + (1-\lambda)\beta q \bar{\pi}_{F_{N1}(\alpha)} + (1-\lambda)\beta(1-q)\bar{\pi}_{F_{N0}(\alpha)} - p) \qquad \text{(All users participate)} \\ \text{subject to} \quad & I_S(\alpha) + (1-\alpha)T - v \geq 0 \\ & I_N(\alpha) + (1-\alpha)T - v \geq 0 \end{aligned}$$

We can further simplify the platform's problem by noting that it is without loss to restrict attention to just $\mathcal{A}_3$. First, one can observe that $q\bar{\pi}_{F_{S1}(\alpha)} + (1-q)\bar{\pi}_{F_{S0}(\alpha)} = q$ because sophisticated agents have a properly specified Bayesian model (by Lemma 1), and thus, $\mathcal{A}_1 \geq \mathcal{A}_2$. At the same time, $\bar{\pi}_{F_{N1}(\alpha)} > \bar{\pi}_{F_{S1}(\alpha)}$ and $\bar{\pi}_{F_{N0}(\alpha)} > \bar{\pi}_{F_{S0}(\alpha)}$ because $\pi^N \,|\, \theta_i \succeq_{FOSD} \pi^S \,|\, \theta_i$, so $q\bar{\pi}_{F_{N1}(\alpha)} + (1-q)\bar{\pi}_{F_{N0}(\alpha)} > q$ (via Lemma 2). We also note that $I_N(\alpha) + (1-\alpha)T - v \geq 0$ can be feasibly satisfied at $\alpha = 0$ (by assumption that $T > v$), so it must be that $\mathcal{A}_3 \geq \mathcal{A}_1 \geq \mathcal{A}_2$. Finally, notice that because $q\bar{\pi}_{F_{S1}(\alpha)} + (1-q)\bar{\pi}_{F_{S0}(\alpha)} = q$, the objective of $\mathcal{A}_3$ and $\mathcal{A}_4$ are identical, but $\mathcal{A}_4$ is subject to an additional constraint, implying that $\mathcal{A}_3 \geq \mathcal{A}_4$. Putting these pieces together, we observe that conditional on adopting an advertising model, the platform will advertise to attract only naïves ($\mathcal{A}_3$). Moreover, because the objective of $\mathcal{A}_3$ is increasing in α (by Lemma 2), the platform will choose an intensity α^* that satisfies $I_N(\alpha^*) + (1-\alpha^*)T - v = 0$. Consequently, because $I_S(\alpha) < I_N(\alpha)$ by Lemma 3, sophisticates will not participate on account of $I_S(\alpha^*) + (1-\alpha^*)T - v < 0$. ∎

Proof of Proposition 2. The only two business models possible via Lemma 4 are those in Proposition 2(a) and (b), so it sufficient to prove it has the cutoff structure described. We let

$m^*(\alpha, \phi_0, \lambda, \phi_1, \phi_{0,N})$ denote the advertising revenue the platform can extract from the firm as a function of the advertising intensity α, false positive rate ϕ_0, sophisticated fraction λ, true positive rate ϕ_1, and naïve agent's false positive rate $\phi_{0,N}$. With advertising, we know the firm will solve the pricing problem post-advertising:

$$\max_p (p-c)(\lambda\beta q+(1-\lambda)\beta(q\bar{\pi}_{F_{N1}(\alpha)}+(1-q)\bar{\pi}_{F_{N0}(\alpha)}-p)) = \left(\frac{\lambda\beta q + (1-\lambda)\beta(q\bar{\pi}_{F_{N1}(\alpha)} + (1-q)\bar{\pi}_{F_{N0}(\alpha)}) - c}{2}\right)^2,$$

by charging $\hat{p}^* = \frac{1}{2}\left(\lambda\beta q + (1-\lambda)\beta(q\bar{\pi}_{F_{N1}(\hat{\alpha}^*)} + (1-q)\bar{\pi}_{F_{N0}(\hat{\alpha}^*)}) + c\right)$. Without advertising, the platform will charge a subscription fee and the firm will make a pre-advertising profit of $\left(\frac{\beta q - c}{2}\right)^2$ by charging $\bar{p}^*$. The difference in the firm's two profit expressions pre- and post-advertising corresponds to the maximum transfer m^* the platform can extract from the firm.

By Lemma 1 and Lemma 2, we know that m^* is increasing in α. For comparative statics on other primitives, it useful to study

$$\begin{aligned}
&q\bar{\pi}_{F_{N1}(\alpha)} + (1-q)\bar{\pi}_{F_{N0}(\alpha)} \\
&= q\left(e^{-\alpha T}q + \phi_1(1-e^{-\alpha T})\frac{\phi_1 q}{\phi_1 q + \phi_{0,N}(1-q)} + (1-\phi_1)(1-e^{-\alpha T})\frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_{0,N})(1-q)}\right) \\
&+ (1-q)\left(e^{-\alpha T}q + \phi_0(1-e^{-\alpha T})\frac{\phi_1 q}{\phi_1 q + \phi_{0,N}(1-q)} + (1-\phi_0)(1-e^{-\alpha T})\frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_{0,N})(1-q)}\right)
\end{aligned}$$

We see this is linear in ϕ_0, and ϕ_0 does not affect $\hat{\alpha}^*$. Moreover, $\frac{\phi_1 q}{\phi_1 q + \phi_{0,N}(1-q)} > \frac{(1-\phi_1)q}{(1-\phi_1)q + (1-\phi_{0,N})(1-q)}$ by assumption that $\phi_1 > \phi_0 > \phi_{0,N}$. Thus, m^* is increasing in ϕ_0 and the platform trades off the advertising revenue m^* with $T - v$, which will lead to a cutoff strategy in ϕ_0 given by $\hat{\phi}_0$. To determine the dependence of $\hat{\phi}_0$ on other model primitives, we perform comparative statics. For λ, because we know $q\bar{\pi}_{F_{N1}(\alpha)} + (1-q)\bar{\pi}_{F_{N0}(\alpha)} > q$, we see that m^* is decreasing in λ. Thus, the advertising-based business model becomes less attractive as λ increases, so $\hat{\phi}_0$ is increasing in λ. For ϕ_1, we notice that holding α fixed,

$$\frac{\partial(q\bar{\pi}_{F_{N1}(\alpha)} + (1-q)\bar{\pi}_{F_{N0}(\alpha)})}{\partial\phi_1} = (1-e^{-\alpha T})\frac{(1-q)^2 q(\phi_0 - \phi_{0,N})(q^2\phi_{0,N}^2 - 2q^2\phi_{0,N}\phi_1 + q^2\phi_1^2 - \phi_{0,N}^2 + \phi_{0,N})}{(q\phi_{0,N} - q\phi_1 - \phi_{0,N} + 1)^2(q\phi_{0,N} - q\phi_1 - \phi_{0,N})^2},$$

which is positive because $\phi_0 > \phi_{0,N}$, $q^2\phi_1^2 + q^2\phi_{0,N}^2 > 2q^2\phi_{0,N}\phi_1$, and $\phi_{0,N} > \phi_{0,N}^2$. At the same time, higher ϕ_1 increases $I_N(\alpha)$, so a greater advertising intensity can also be adopted while satisfying the naifs' participation constraint, leading to even higher m^*. Thus, the ad-based business model becomes even more likely as ϕ_1 increases, or in other words, the cutoff $\hat{\phi}_0$ decreases in ϕ_1.

Finally, we can see that $q\bar{\pi}_{F_{N1}(\alpha)} + (1-q)\bar{\pi}_{F_{N0}(\alpha)}$ is decreasing in $\phi_{0,N}$:

$$\frac{\partial(q\bar{\pi}_{F_{N1}(\alpha)} + (1-q)\bar{\pi}_{F_{N0}(\alpha)})}{\partial\phi_{0,N}} = (1-e^{-\alpha T})(1-q)q\left(-\frac{\phi_1(\phi_0(1-q) + q\phi_1)}{(\phi_{0,N} + q\phi_1 - q\phi_{0,N})^2} - \frac{(1-\phi_1)(1-\phi_0(1-q) - q\phi_1)}{(1-\phi_{0,N} - q(\phi_1 - \phi_{0,N}))^2}\right),$$

which is strictly negative. Increasing $\phi_{0,N}$ also makes the advertising intensity tolerated by naïves, $\hat{\alpha}^*$, decrease (because $I_N(\alpha)$ decreases), depressing m^* further. So increasing $\phi_{0,N}$ makes the subscription-

based model more attractive, increasing $\hat{\phi}_0$. ∎

Proof of Proposition 3. From Lemma 1, we know that there is no equilibrium transfer from the firm due to advertising, $m^* = 0$. Thus, the subscription model is more attractive to the platform, which yields $T - v$ profit, which is positive by assumption. Because there is no advertising, the firm once again solves $\max_p (p-c)(\beta q - p)$, which occurs when the price is set to $\bar{p}^* = (\beta q + c)/2$. This is the same price of the product as in the base case with no platform. Moreover, the consumer surplus is determined by $I_S(\alpha) + (1-\alpha)T - (T-v)$ for $\alpha = 0$, which is exactly equal to the agent's outside option v. Thus, welfare is given by $v + qU(z_i^*(q,\bar{p}^*);\, \theta_i = 1) + (1-q)U(z_i^*(q,\bar{p}^*);\, \theta_i = 0) - \bar{p}^* z_i^*(q,\bar{p}^*)$ for both types of users in both the base case and the fully-rational benchmark. This allows us to establish $W_{\text{fully-rational}}(\tau) = W_{\text{base}}(\tau)$, which is independent of the agent's type τ because there is no advertising in equilibrium. ∎

Proof of Proposition 4. For regime (a), we have the platform adopts the same subscription model it does in Proposition 3, leading to identical welfare as in the fully-rational benchmark, which as we showed also has the same welfare as the base case without a platform.

For regime (b), advertising intensity is at $\hat{\alpha}^*$ and prices are at $\hat{p}^*$. We have the user welfare for the sophisticated agents is given by

$$\begin{aligned} W(S, x_i = 0) &= v + qU(z_i^*(q,\hat{p}^*);\, \theta_i = 1) + (1-q)U(z_i^*(q,\hat{p}^*);\, \theta_i = 0) - \hat{p}^* z_i^*(q,\hat{p}^*) \\ &< v + qU(z_i^*(q,\bar{p}^*);\, \theta_i = 1) + (1-q)U(z_i^*(q,\bar{p}^*);\, \theta_i = 0) - \bar{p}^* z_i^*(q,\bar{p}^*) \\ &= W_{\text{base}}(S) \end{aligned}$$

Recall from Proposition 1 that $\Delta(\alpha, p)$ is the naïve agent's ex-post consumer surplus lost relative to a sophisticated agent under the same advertising intensity α and product price p. Consumer welfare for the naïve agents is given by

$$\begin{aligned} W(N, x_i = 1) =& (1-\hat{\alpha}^*)T + q\mathbb{E}_{\pi \sim F_{N1}}\left[U(z^*(\pi,\hat{p}^*);\, \theta_i = 1) - \hat{p}^* z^*(\pi,\hat{p}^*)\right] \\ &+ (1-q)\mathbb{E}_{\pi \sim F_{N0}}\left[U(z^*(\pi,\hat{p}^*);\, \theta_i = 0) - \hat{p}^* z^*(\pi,\hat{p}^*)\right] \\ =& (1-\hat{\alpha}^*)T + q\mathbb{E}^S_{\pi_i|\theta_i=1}[U(C_i(\pi_i,\hat{p}^*);\, \theta_i = 1) - \hat{p}^* C_i(\pi_i,\hat{p}^*) \,|\, \hat{\alpha}^*] \\ &+ (1-q)\mathbb{E}^S_{\pi_i|\theta_i=0}[U(C_i(\pi_i,\hat{p}^*);\, \theta_i = 0) - \hat{p}^* C_i(\pi_i,\hat{p}^*) \,|\, \hat{\alpha}^*] - \Delta(\hat{\alpha}^*,\hat{p}^*) \\ =& (1-\hat{\alpha}^*)T + I_S(\hat{\alpha}^*) - \Delta(\hat{\alpha}^*,\hat{p}^*) + U(C_i(q,\hat{p}^*)) - \hat{p}^* C_i(q,\hat{p}^*) \\ =& v - \Delta(\hat{\alpha}^*,\hat{p}^*) + (I_S(\hat{\alpha}^*) - I_N(\hat{\alpha}^*)) + \frac{1}{2}(\beta q - \hat{p}^*)^2 \\ <& v + \frac{1}{2}(\beta q - \bar{p}^*)^2 = W_{\text{fully-rational}}(N) \end{aligned}$$

Finally, we note from Proposition 3 that $W_{\text{fully-rational}}(\tau) = W_{\text{base}}(\tau)$ for both $\tau \in \{S, N\}$. ∎

A.3 Proofs from Section 4

Proof of Proposition 5. We now allow the platform to offer both subscription fees and advertising intensities (but not in the same plan). By Lemma 1, we know that the platform can maximize revenue

from sophisticated agents by offering a subscription service and charging $T - v$. Such a subscription service leaves naïves and sophisticates both indifferent between platform participation and not; thus, an ad-based plan for naïves is still chosen if and only if $I_N(\alpha) + (1-\alpha)T - v \geq 0$. The firm-platform transfer $m^*(\alpha, \phi_0, \lambda, \phi_1, \phi_{0,N})$ conditional on advertising is identical to that of Proposition 2, and in particular, is maximized as α is larger (by Lemma 2). Consequently, an ad-based plan will extract all the surplus from the naïf's participation constraint, letting the only two candidates for business models be regimes (a) and (b) of Proposition 5.

The revenue generated from business model (a) is given by $T - v$ whereas the revenue generated from business model (b) is given by $\lambda(T-v) + m^*(\alpha, \phi_0, \lambda, \phi_1, \phi_{0,N})$. The difference between the latter and the former is thus $m^*(\alpha, \phi_0, \lambda, \phi_1, \phi_{0,N}) - (1-\lambda)(T-v)$. The comparative statics on ϕ_0, ϕ_1, and $\phi_{0,N}$ then follow immediately from Proposition 2, giving the cutoff characterization and that $\phi_1^*(\lambda, \phi_1, \phi_{0,N})$ that the cutoff is increasing in $\phi_{0,N}$ but decreasing in ϕ_1. To see that it is increasing in λ, we note that $m^*(\alpha, \phi_0, \lambda, \phi_1, \phi_{0,N}) - (1-\lambda)(T-v)$ is equal to 0 when $\lambda = 1$, and either $-\partial m/\partial \lambda^*\Big|_{\lambda=1} > T - v$ or $-\partial m/\partial \lambda^*\Big|_{\lambda=1} < T - v$. In the former case, we know that $-\partial m/\partial \lambda^* > T - v$ for all λ, which implies that the ad-based plan offered to naïfs generates more revenue than the subscription-based plan for all $\lambda \in (0,1)$. On the other hand, if $-\partial m/\partial \lambda^*\Big|_{\lambda=1} < T - v$, then because m^* is quadratic and $\lambda(T-v)$ is linear with a shared intersection at $\lambda = 1$, there exists a unique single crossing at $\lambda^* < 1$ where $m^*(\alpha, \phi_0, \lambda^*, \phi_1, \phi_{0,N}) = (1-\lambda^*)T - v$, and the ad-based plan for naïfs is more profitable when $\lambda < \lambda^*$ but the subscription-based plan is more profitable when $\lambda > \lambda^*$. This implies the subscription-based model is more likely as λ increases, which means the corresponding cutoff $\phi_0^*(\lambda, \phi_1, \phi_{0,N})$ is also increasing in λ.

Finally, we note that $m^*(\alpha, \hat{\phi}_0, \lambda, \phi, \phi_{0,N}) = T - v$ defines the cutoff for $\hat{\phi}_0$ whereas $m^*(\alpha, \phi_0^*, \lambda, \phi_1, \phi_{0,N}) = (1-\lambda)(T-v)$ defines the cutoff for ϕ_0^*. Because $m^*(\alpha, \cdot, \lambda, \phi_1, \phi_{0,N})$ is increasing and $(1-\lambda)(T-v) < T - v$, it is necessarily the case that $\phi_0^*(\lambda, \phi_1, \phi_{0,N}) < \hat{\phi}_0(\lambda, \phi_1, \phi_{0,N})$. ∎

Proof of Corollary 1. Regime (a) is exactly the same as it was in Proposition 4. For regime (b), the user welfare of the sophisticates is given by

$$\begin{aligned}
W(S, x_i = 0) &= T - P^* + qU(z_i^*(q, \hat{p}^*);\, \theta_i = 1) + (1-q)U(z_i^*(q, \hat{p}^*);\, \theta_i = 0) - \hat{p}^* z_i^*(q, \hat{p}^*) \\
&= v + qU(z_i^*(q, \hat{p}^*);\, \theta_i = 1) + (1-q)U(z_i^*(q, \hat{p}^*);\, \theta_i = 0) - \hat{p}^* z_i^*(q, \hat{p}^*) \\
&< v + qU(z_i^*(q, \bar{p}^*);\, \theta_i = 1) + (1-q)U(z_i^*(q, \bar{p}^*);\, \theta_i = 0) - \bar{p}^* z_i^*(q, \bar{p}^*) \\
&= W_{\text{base}}(S)
\end{aligned}$$

The user welfare of the naïve agents in regime (b) is exactly as in regime (b) of Proposition 4. ∎

Proof of Proposition 6. We know that if $\phi_0 < \phi_0^*(\lambda, \phi_1, \phi_{0,N})$ the platform business model is fully subscription-based and $\hat{W}^*(\tau) = W_{\text{base}}$; in particular, it is constant over this entire range for ϕ_0. By Corollary 1, there is a discontinuous jump in welfare down to $\hat{W}^*(\tau) < W_{\text{base}}$ at $\phi_0 = \phi_0^*(\lambda, \phi_1, \phi_{0,N})$. Thus, it just remains to show that $\hat{W}^*(\tau)$ is decreasing in ϕ_0 when $\phi_0 > \phi_0^*(\lambda, \phi_1, \phi_{0,N})$. We know that $\hat{p}^* = \frac{1}{2}\left(\lambda\beta q + (1-\lambda)\beta(q\bar{\pi}_{F_{N1}(\hat{\alpha}^*)} + (1-q)\bar{\pi}_{F_{N0}(\hat{\alpha}^*)}) + c\right)$ in equilibrium, and moreover we observed in

Proposition 2 that $q\bar{\pi}_{F_{N1}(\hat{\alpha}^*)} + (1-q)\bar{\pi}_{F_{N0}(\hat{\alpha}^*)}$ is increasing in ϕ_0. Sophisticate welfare is given by

$$\hat{W}^*(S) = v + qU(z_i^*(q,\hat{p}^*);\, \theta_i = 1) + (1-q)U(z_i^*(q,\hat{p}^*);\, \theta_i = 0) - \hat{p}^* z_i^*(q,\hat{p}^*)$$

which is monotonically decreasing in $\hat{p}^*$ (and thus ϕ_0). On the other hand, from Proposition 2, the welfare of naïve agents is given by

$$\hat{W}^*(N) = v - \Delta(\hat{\alpha}^*,\hat{p}^*) + (I_S(\hat{\alpha}^*) - I_N(\hat{\alpha}^*)) + \frac{1}{2}(\beta q - \hat{p}^*)^2$$

Since $\hat{\alpha}^*$ is constant in ϕ_0, changing ϕ_0 has no bearing on $I_S(\hat{\alpha}^*)$ or $I_N(\hat{\alpha}^*)$. As with Lemma A.1, it can also be shown that $\Delta(\hat{\alpha}^*,\hat{p}^*)$ has no dependence on $\hat{p}^*$, thus ϕ_0 has no effect on $\Delta(\hat{\alpha}^*,\hat{p}^*)$. Consequently, $\hat{W}^*(N)$ is monotonically decreasing in $\hat{p}^*$ (and thus ϕ_0).

Because $\phi_0^*(\lambda,\phi_1,\phi_{0,N})$ is decreasing in ϕ_1 by Corollary 1, we know there exists $\phi_1^*(\lambda,\phi_0,\phi_{0,N})$ such that if $\phi_1 < \phi_1^*(\lambda,\phi_0,\phi_{0,N})$ the platform business model is fully subscription-based and $\hat{W}^*(\tau) = W_{\text{base}}$; in particular, it is constant over this entire range of ϕ_1. There is a discontinuous jump in welfare down to $\hat{W}^*(\tau) < W_{\text{base}}$ at $\phi_1 = \phi^*(\lambda,\phi_1,\phi_{0,N})$. Therefore, it just remains to show that $\hat{W}^*(\tau)$ is decreasing in ϕ_1 when $\phi_1 > \phi_1^*(\lambda,\phi_1,\phi_{0,N})$. By Blackwell's theorem, we know that $I_N(\alpha)$ is increasing in ϕ_1, which means that the platform's choice of advertising intensity $\hat{\alpha}^*$ is increasing in ϕ_1. At the same time $\hat{p}^*$ is increasing in $q\bar{\pi}_{F_{N1}(\hat{\alpha}^*)} + (1-q)\bar{\pi}_{F_{N0}(\hat{\alpha}^*)}$, which is increasing in ϕ_1 and $\hat{\alpha}^*$ (which is, in turn, increasing in ϕ_1). Sophisticate welfare is given by

$$\hat{W}^*(S) = v + qU(z_i^*(q,\hat{p}^*);\, \theta_i = 1) + (1-q)U(z_i^*(q,\hat{p}^*);\, \theta_i = 0) - \hat{p}^* z_i^*(q,\hat{p}^*)$$

which is monotonically decreasing in $\hat{p}^*$ (and thus ϕ_1). The welfare of naïve agents is given by

$$\hat{W}^*(N) = v - \Delta(\hat{\alpha}^*,\hat{p}^*) + (I_S(\hat{\alpha}^*) - I_N(\hat{\alpha}^*)) + \frac{1}{2}(\beta q - \hat{p}^*)^2$$

where $-\Delta(\hat{\alpha}^*,\hat{p}^*)$ is decreasing in $\hat{\alpha}^*$ and constant in $\hat{p}^*$, and where $I_S(\hat{\alpha}^*) - I_N(\hat{\alpha}^*)$ is negative and proportional to $1 - e^{-\alpha T}$, which is decreasing in $\hat{\alpha}^*$. Thus, $\hat{W}^*(N)$ is monotonically decreasing in both $\hat{\alpha}^*$ and $\hat{p}^*$ (and thus ϕ_1). ∎

Proof of Proposition 7. The platform will always choose a plan that generates participation from both types of agents because it can always offer the fully subscription-based plan with $P^* = T - v$ to extract some surplus from any non-participating type. Thus, if offering a single plan is a best response for the platform, then the stricter of the two types' participation constraint will bind:

$$\begin{aligned} \max_{\alpha,p,P} \quad & (p-c)(\lambda\beta q + (1-\lambda)\beta q\bar{\pi}_{F_{N1}(\alpha)} + (1-\lambda)\beta(1-q)\bar{\pi}_{F_{N0}(\alpha)} - p) + P \\ \text{subject to} \quad & I_S(\alpha) + (1-\alpha)T - v - P \geq 0, \end{aligned}$$

for which there exists some α^* and corresponding $P^* = I_S(\alpha^*) + (1-\alpha^*)T - v$ that maximizes the above expression. Such (α^*, P^*) is the only candidate for a mixed business model where only one plan is offered. Note that $\alpha^* \geq \alpha_S^{FB}$ because the platform could increase P^* and $\beta q + (1-\lambda)\beta q\bar{\pi}_{F_{N1}(\alpha)} + (1-$

$\lambda)\beta(1-q)\bar{\pi}_{F_{N0}(\alpha)}$ by increasing α up to α_S^{FB}, which monotonically increases its objective (by Lemma 2 and 3). Note that $\alpha^* < \hat{\alpha}^*$ because $I_S(\hat{\alpha}^*) + (1-\hat{\alpha}^*)T - v < 0$ by construction in Proposition 2.

In the case of a mixed business model with two plans offered, (α_1^*, P_1^*) and (α_2^*, P_2^*), we maximize the firm's profit subject to the incentive compatibility constraint between the naïve and sophisticated users and their corresponding participation constraints:

$$\begin{aligned}
\max_{\alpha_1,\alpha_2,p,P_1,P_2} \quad & (p-c)(\lambda\beta q + (1-\lambda)\beta q\bar{\pi}_{F_{N1}(\alpha_2)} + (1-\lambda)\beta(1-q)\bar{\pi}_{F_{N0}(\alpha_2)} - p) + \lambda P_1 + (1-\lambda)P_2 \\
\text{subject to} \quad & P_2 - P_1 - I_S(\alpha_2) + I_S(\alpha_1) \geq 0 \\
& P_1 - P_2 + I_N(\alpha_2) - I_N(\alpha_1) \geq 0 \\
& I_S(\alpha_1) + (1-\alpha_1)T - v - P_1 \geq 0 \\
& I_N(\alpha_2) + (1-\alpha_2)T - v - P_2 \geq 0,
\end{aligned}$$

which yields some profit-maximizing business model (α_1^*, P_1^*) and (α_2^*, P_2^*) when setting the profit-maximizing price $p^* = \frac{1}{2}(\lambda\beta q + (1-\lambda)\beta q\bar{\pi}_{F_{N1}(\alpha_2)} + (1-\lambda)\beta q\bar{\pi}_{F_{N0}(\alpha_2)} + c)$. The platform then compares the profits under the single plan, (α^*, P^*), and under the two plans, (α_1^*, P_1^*) and (α_2^*, P_2^*), and chooses the better of them.

We show that α_2^* is increasing in ϕ_0, attains at least α_S^{FB}, and never exceeds $\hat{\alpha}^*$, in the profit-maximizing business model, thus establishing the cutoff form $\tilde{\phi}_0(\lambda, \phi_1, \phi_{0,N})$. The first one is observed because $\beta q\bar{\pi}_{F_{N1}(\alpha_2^*)} + \beta q\bar{\pi}_{F_{N0}(\alpha_2^*)}$ is increasing in ϕ_0 while leaving all of the constraints (both IC and IR) unaffected, which necessarily increases the firm's profit (which it can extract through an ad contract). For the second, if $\alpha_1^* \leq \alpha_2^* \leq \alpha_S^{FB}$, then the platform can increase the ad intensity of both α_1^* and α_2^* without needing to reduce P_1^* and P_2^*, which leads to higher profit, a contradiction. The third is a direct consequence of the IR constraint on naïve agents and that for any $\alpha > \hat{\alpha}^*$, we have $I_N(\alpha) + (1-\alpha)T - v < 0$. The same comparative statics with respect to $\phi_{0,N}$, ϕ_1, and λ also hold as in Proposition 5 for identical reasoning.

Finally, we establish that $\tilde{W}_{(a)}^*(\tau) \leq W_{\text{base}}^*(\tau)$ $\tilde{W}_{(b)}^*(\tau) < \tilde{W}_{\text{base}}^*(\tau)$. We leverage the reasoning from Proposition 4 and Corollary 1. When a single plan is offered, we know that the platform will choose α maximal so that $I_S(\alpha^*) + (1-\alpha^*)T - v - P^* = 0$. At the same time, $\beta q\bar{\pi}_{F_{N1}(\alpha^*)} + \beta q\bar{\pi}_{F_{N0}(\alpha^*)} \geq q$, making the product price $p^* \geq \bar{p}^*$. For the same reasons as in Corollary 1, this will lead to lower user welfare relative to the base case for both types of users. When multiple plans are offered, we know that either $I_S(\alpha_1^*) + (1-\alpha_1^*)T - v - P_1^* = 0$ or $I_N(\alpha_2^*) + (1-\alpha_2^*)T - v - P_2^* = 0$, because otherwise both P_1^* and P_2^* could be increased without affecting the incentive compatibility constraints and increasing the platform's profit. If $I_N(\alpha_2^*) + (1-\alpha_2^*)T - v = P_2^*$, then $I_S(\alpha_2^*) + (1-\alpha_2^*)T - v < P_2^*$ by Lemma 3, which means the sophisticates would choose their outside option over participating on the naïfs' plan, and thus the platform can also set $P_1^* = I_S(\alpha_1^*) + (1-\alpha_1^*)T - v$ without affecting incentive compatibility. Otherwise, $I_S(\alpha_1^*) + (1-\alpha_1^*)T - v = P_1^*$, but then the sophisticates' incentive compatibility constraint implies (by plugging in P_1^*) that $P_2^* - (I_S(\alpha_1^*) + (1-\alpha_1^*)T - v) - I_S(\alpha_2^*) + I_S(\alpha_1^*) \geq 0$ which in turn suggests $I_S(\alpha_2^*) + (1-\alpha_1^*)T - v - P_2^* < 0$, and since $\alpha_1^* \leq \alpha_2^*$, $I_S(\alpha_2^*) + (1-\alpha_2^*)T - v - P_2^* < 0$. In both cases, both sophisticated and naïve agents receive at most the welfare of $v + \frac{1}{2}(\beta q - p^*)^2$, where $p^* \geq \bar{p}^*$, which is no more than base case welfare. It is straightforward to see that $\tilde{W}_{(b)}^*(\tau) < \tilde{W}_{\text{base}}^*(\tau)$ because

$p^* > \bar{p}^*$ when there is a positive level of advertising $\hat{\alpha}^*$ for naïfs. ∎

A.4 Proofs from Section 5

Proof of Proposition 8. As in our baseline game, all decisions for $t = 3,4,5$ are generically unique by backward induction, so it suffices to consider just $t \leq 2$. At time $t = 2$, we take the platform advertising rates $\{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N$ as given. Because product prices are chosen sequentially, Zermelo's theorem guarantees there exists an equilibrium vector of prices $(p_1^*, \ldots, p_N^*)$ attained from backward induction. To show this is the unique choice of pricing vectors, we just need to show that no firm is indifferent between selecting two prices (generically). First, observe that no firm will choose a price below marginal cost, because all other firms would price above it and this firm would steal at least $\prod_{\rho=1}^M \prod_{j=1}^N \prod_{\ell=1}^2 e^{-\alpha_{\ell,\rho}^{(j)} T} > 0$ share of the market, making negative profit, which is a dominated strategy. This implies that in a robust equilibrium, every firm will price strictly above marginal cost and can make strictly positive profits by charging $p^* = \min\{p \,|\, p = k\varepsilon > c,\ k \in \mathbb{N}\}$. The firm j will choose p_j to maximize $\max_{p_j}(p_j - c)\mathbb{E}_{\boldsymbol{\pi}_i \sim \{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N}[(\pi_i^{(j)}\beta - p_j) \cdot \mathbb{1}_{\beta\pi_i^{(j)} - p_j \in \max_\omega\{\beta\pi_i^{(\omega)} - p_\omega\}}]$, where the prices $\{p_\omega\}_{\omega \neq j}$ are taken as given for all firms who price before firm j and taken as a best-response function of p_j for all firms that price after firm j, with $\mathbb{1}_{\beta\pi_i^{(j)} - p_j \in \max_\omega\{\beta\pi_i - p_\omega\}} = (|\{j \,:\, \beta\pi_i^{(j)} - p_j = \max_\omega \beta\pi_i^{(\omega)} - p_\omega\}|)^{-1}$. It is clear that p_j^* is upper bounded by the monopoly price, which we denote by $\tilde{p}^* = \arg\max_{p_j}(p_j - c)\mathbb{E}_{\boldsymbol{\pi}_i \sim \{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N}[\pi_i^{(j)}\beta - p_j]$, so $p_j^* \in \{p^*, p^* + \varepsilon, \ldots, (k^* + 1)\varepsilon\}$ where $k^* = \max\{k \in \mathbb{N} \,|\, k\varepsilon < \tilde{p}^*\}$.

We proceed by induction on the reverse order of sequential offers (and without loss of generality we suppose that firm N acts last). For firm N, all prices $(p_1^*, \ldots, p_{N-1}^*)$ are fixed. Note that unless $\mathbb{P}_{\boldsymbol{\pi}_i \sim \{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N}[\beta\pi_i^{(N)} - p_N \in \max_\omega\{\beta\pi_i^{(\omega)} - p_\omega\}] = 0$, generically none of the prices $p_N \in \{p^*, p^* + \varepsilon, \ldots, (k^*+1)\varepsilon\}$ yield the same profit. Moreover, if $\mathbb{P}_{\boldsymbol{\pi}_i \sim \{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N}[\beta\pi_i^{(N)} - p_N \in \max_\omega\{\beta\pi_i^{(\omega)} - p_\omega\}] = 0$ for some pricing strategy p_N, then the firm earns zero profits with certainty, and this strategy is dominated by pricing at p^*, which guarantees strictly positive profits. Thus, we can retain a subset, $\mathcal{P}_N$, of $\{p^*, p^* + \varepsilon, \ldots, (k^* + 1)\varepsilon\}$ where $\mathbb{P}_{\boldsymbol{\pi}_i \sim \{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N}[\beta\pi_i^{(N)} - p_N \in \max_\omega\{\beta\pi_i^{(\omega)} - p_\omega\}] > 0$, which generically yields positive and distinct payoffs for $(p_N - c)\mathbb{E}_{\pi_i \sim \{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N}[(\pi_i^{(N)}\beta - p_N) \cdot \mathbb{1}_{\beta\pi_i^{(N)} - p_N \in \max_\omega\{\beta\pi_i^{(\omega)} - p_\omega\}}]$ for all $p_N \in \mathcal{P}_N$. Thus, generically, there is a unique p_N^* that maximizes this expression. Finally, note that because the set $\{p^*, p^* + \varepsilon, \ldots, (k^*+1)\varepsilon\}$ is discrete, firm N's best response choice of p_N^* is insensitive to sufficiently small perturbations in model parameters.

We leverage this fact to argue the inductive step. For all $p_j \in \{p^*, p^* + \varepsilon, \ldots, (k^* + 1)\varepsilon\}$, we can once again rule out any p_j^* with $\mathbb{P}_{\boldsymbol{\pi}_i \sim \{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N}[\pi_i^{(j)} - p_j \in \max_\omega\{\pi_i^{(\omega)} - p_\omega\}] = 0$ as an equilibrium strategy, because pricing at p^* results in strictly positive profits. If we have two prices p_j', p_j'' with both $\mathbb{P}_{\boldsymbol{\pi}_i \sim \{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N}[\pi_i^{(j)} - p_j' \in \max_\omega\{\pi_i^{(\omega)} - p_\omega\}] > 0$ and $\mathbb{P}_{\boldsymbol{\pi}_i \sim \{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N}[\pi_i^{(j)} - p_j'' \in \max_\omega\{\pi_i^{(\omega)} - p_\omega\}] > 0$ that yield identical profits, then one can introduce a small perturbation to β which has no effect on the best response prices of firms $j+1, \ldots, N$ but breaks the exact inequality between $(p_j' - c)\mathbb{E}_{\pi_i \sim \{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N}[(\pi_i^{(j)}\beta - p_j') \cdot \mathbb{1}_{\beta\pi_i^{(\omega)} - p_j' \in \max_\omega\{\beta\pi_i^{(\omega)} - p_\omega\}}]$ and $(p_j'' - c)\mathbb{E}_{\pi_i \sim \{(\alpha_{1,\rho}^{(j)}, \alpha_{2,\rho}^{(j)})\}_{j=1}^N}[(\pi_i^{(j)}\beta - p_j'') \cdot \mathbb{1}_{\beta\pi_i^{(\omega)} - p_j'' \in \max_\omega\{\beta\pi_i^{(\omega)} - p_\omega\}}]$, just as in the base case step. Therefore, generically we once again get

a unique $p_j^* \in \mathcal{P}_j$, completing the inductive step and establishing the pricing vector $(p_1^*, \ldots, p_N^*)$ as unique. Moreover, this unique pricing vector will be order-independent (up to a relabeling of the firms) because all firms are ex-ante identical.

For $t = 1$, we consider the case of a single platform ($M = 1$) and multiple platforms ($M > 1$) separately. When there is a single platform, there generically exists a unique $(\alpha_*^{(1)}, \ldots, \alpha_*^{(N)})$ (up to a renumbering of the firms) which maximizes the transfers $\sum_{j=1}^N m^{(j)}$ the platform can collect from the firms, by extracting the full surplus the firms gain from the advertising. The firms always accept such a contract in equilibrium. When there are multiple platforms, they simultaneously compete to offer the plans most desired by the sophisticated and naïve users. As described in the equilibrium of Proposition 11, these plans have no subscription fees and cater exactly to the advertising levels desired by sophisticates and naïfs out of all feasible plans. These are given by $(\tilde{\alpha}_S^{FB}, \ldots, \tilde{\alpha}_S^{FB})$ and $(\tilde{\alpha}_N^*, \ldots, \tilde{\alpha}_N^*)$ for some $\tilde{\alpha}_S^{FB}$ and $\tilde{\alpha}_N^*$ that maximize $I_S(\alpha^{(1)}, \ldots, \alpha^{(N)}) + (1 - \sum_{j=1}^N \alpha^{(j)})T$ and $I_N(\alpha^{(1)}, \ldots, \alpha^{(N)}) + (1 - \sum_{j=1}^N \alpha^{(j)})T$, respectively (see Appendix C). Platforms simultaneously compete over transfers to firms, driving these transfers to zero in the unique equilibrium, which the firms always accept. The equilibrium is unique because we can fully characterize for the unique values of $\tilde{\alpha}_S^{FB}$ and $\tilde{\alpha}_N^*$. ∎

Proof of Proposition 9. Under firm competition, we have that $I_S(\alpha^{(1)}, \alpha^{(2)})$ takes the form of $I_S(\alpha^{(1)}, \alpha^{(2)}) = (1-e^{-\alpha^{(1)}T})e^{-\alpha^{(2)}T}\chi_S^{(1)} + (1-e^{-\alpha^{(2)}T})e^{-\alpha^{(1)}T}\chi_S^{(2)} + (1-e^{-\alpha^{(1)}T})(1-e^{-\alpha^{(2)}T})\chi_S^{(1,2)}$ for some $\chi_S^{(1)}, \chi_S^{(2)}, \chi_S^{(1,2)}$ that depend on model primitives (e.g., ϕ_0 and ϕ_1). Similarly, $I_N(\alpha^{(1)}, \alpha^{(2)})$ takes the form of $I_N(\alpha^{(1)}, \alpha^{(2)}) = (1 - e^{-\alpha^{(1)}T})e^{-\alpha^{(2)}T}\chi_N^{(1)} + (1 - e^{-\alpha^{(2)}T})e^{-\alpha^{(1)}T}\chi_N^{(2)} + (1 - e^{-\alpha^{(1)}T})(1 - e^{-\alpha^{(2)}T})\chi_N^{(1,2)}$ for some $\chi_N^{(1)} > \chi_S^{(1)}$, $\chi_N^{(2)} > \chi_S^{(2)}$, and $\chi_N^{(1,2)} > \chi_S^{(1,2)}$.

For regime (i), notice that when $\phi_0 = \phi_{0,N} = \phi_1 = 0$, beliefs of the agents do not change as a result of advertising, and therefore the firms are not willing to transfer any amount to the platform from an ad contract. At the same time, we observe that there exists $\bar{\phi}_1$ such that for all $\phi_1 < \bar{\phi}_1$, the maximal subscription fee the platform can generate for both sophisticates and naïfs occurs where $\alpha^{(1)} = \alpha^{(2)} = 0$. This is a direct consequence of noting that $\partial I_S(\alpha^{(1)}, \alpha^{(2)})/\partial\alpha^{(1)} = Te^{-\alpha^{(1)}T}e^{-\alpha^{(2)}T}\chi_S^{(1)} + Te^{-\alpha^{(1)}T}(1-e^{-\alpha^{(2)}T})(\chi_S^{(1,2)} - \chi_S^{(2)}) < T$ for all $\alpha^{(1)}$ and $\alpha^{(2)}$, for ϕ_1 sufficiently close to 0, because $\chi_S^{(1)}$, $\chi_S^{(2)}$ and $\chi_S^{(1,2)}$ are sufficiently close to 0 (and the same for $\partial I_S(\alpha^{(1)}, \alpha^{(2)})/\partial\alpha^{(2)}$, $\partial I_N(\alpha^{(1)}, \alpha^{(2)})/\partial\alpha^{(1)}$, and $\partial I_N(\alpha^{(1)}, \alpha^{(2)})/\partial\alpha^{(2)}$). So, in particular, $\alpha^{(1)} = \alpha^{(2)} = 0$ maximizes both $I_S(\alpha^{(1)}, \alpha^{(2)}) + (1 - \alpha^{(1)} - \alpha^{(2)})T$ and $I_N(\alpha^{(1)}, \alpha^{(2)}) + (1 - \alpha^{(1)} - \alpha^{(2)})T$ for all $\phi_1 < \bar{\phi}_1$. Thus, at $\phi_1 = 0$, the profit-maximizing business model of the platform is to set $\alpha^{(1)} = \alpha^{(2)} = 0$ and extract full surplus from a subscription fee $P^* = T - v$. We also know the user will purchase all of the product which has the largest $\pi_i - p_j^*$, which is equivalent to largest $q - p_j^*$ when there is no advertising, reducing to standard Bertrand competition on price without horizontal differentiation. This leads to $p_1^* = p_2^* = c$ as the unique equilibrium under regime (i). Finally, by Blackwell's theorem, we know the transfers from the firms will be monotonically increasing in ϕ_1 for every pair $(\alpha^{(1)}, \alpha^{(2)})$, therefore, we can choose $\phi_1^{\mathcal{F}} \leq \bar{\phi}_1$ to be maximal such that for any $\phi_1 > \phi_1^{\mathcal{F}}$, there is an ad-based plan offered in equilibrium.

In regime (ii), let us consider $\phi_1 > \phi_1^{\mathcal{F}}$ and $\phi_0 = \phi_1$. We note that $\chi_S^{(1)} = \chi_S^{(2)} = \chi_S^{(1,2)} = 0$, and therefore the platform can extract the maximal amount of revenue from the sophisticates by once again charging the subscription fee $T - v$ and setting $\alpha^{(1)} = \alpha^{(2)} = 0$, as in regime (i). At the same time, under the fully-rational benchmark, for the exact same reasons as the previous paragraph, the

equilibrium involves the platform offering one subscription-based plan and the firms competing over price alone to offer $\bar{p}_1^* = \bar{p}_2^* = c$. However, by construction of $\phi_1^{\mathcal{F}}$, we know that $\alpha_2^{(1)*} + \alpha_2^{(2)*} > 0$ and the naïfs opt into the ad-based plan in the equilibrium. Without loss of generality, let $\alpha_2^{(1)*} > 0$. Then with probability $\phi_0(1 - e^{-\alpha_2^{(1)*}})e^{-\alpha_2^{(2)*}} > 0$, the user sees a positive ad from firm 1 and no ad from firm 2, leading to belief $\pi_N^{(1)}(G) > q$ about product 1 but belief $\pi_N^{(2)} = q$ about product 2. We know that firm 1 pricing below marginal cost is a dominated strategy, and pricing at marginal cost is dominated by a strategy that sets $p_1^* = c + (\pi_N(G) - q)$, which attains strictly positive profits; thus, $\hat{p}_1^* > \bar{p}_1^* = c$ in equilibrium. At the same time, with probability $(1-\phi_0)(1 - e^{-\alpha_2^{(1)*}})e^{-\alpha_2^{(2)*}} > 0$, the user sees a negative ad from firm 1 and no ad from firm 2, leading to a belief $\pi_N^{(1)}(B) < q$ about product 1 but belief $\pi_N^{(2)} = q$ about product 2. We know that firm 2 pricing below marginal cost is a dominated strategy, and pricing at marginal cost is dominated by a strategy that sets $p_2^* = c + (q - \pi_N(B))$, which attains strictly positive profits; thus, $\hat{p}_2^* > \bar{p}_2^* = c$ in equilibrium. The claim in (ii) follows by choosing the minimal $\phi_0^{\mathcal{F}}(\phi_1)$ where this property holds. ∎

Proof of Proposition 10. In the fully-rational benchmark of regimes (i) and (ii) we have that the platform offers a subscription-based plan with $P^* = T - v$ and firms price at marginal cost $\bar{p}_1^* = \bar{p}_2^* = c$. The equilibrium played in regime (i) corresponds exactly to the fully-rational benchmark, so we obtain that $\hat{W}_{2,1}(\tau)^* = \bar{W}_{2,1}^*(\tau)$ for both user types $\tau \in \{S, N\}$. For regime (ii), advertising intensity is at $\alpha_2^* = \alpha_2^{(1)*} + \alpha_2^{(2)*} > 0$ for naïfs and prices are at $\hat{p}_1^* > c$ and $\hat{p}_2^* > c$. Sophisticates subscribe to the same plan in equilibrium as they do under the fully-rational benchmark, but face strictly higher prices in the product market, so $\hat{W}_{2,1}^*(S) < \bar{W}_{2,1}^*(S)$. naïve agents, on the other hand, receive welfare at most $\hat{W}_{2,1}^*(S) + I_S(\alpha^{(1)*}, \alpha^{(2)*}) - (\alpha^{(1)*} + \alpha^{(2)*})T - \Delta(\alpha^{(1)*}, \alpha^{(2)*}, \hat{p}_1^*, \hat{p}_2^*) < \hat{W}_{2,1}^*(S) < \bar{W}_{2,1}^*(S) = \bar{W}_{2,1}^*(N)$, by construction of $\phi_1^{\mathcal{F}}$ and $\phi_0^{\mathcal{F}}(\phi_1)$ that $(\alpha^{(1)}, \alpha^{(2)}) = (0,0)$ maximizes $I_S(\alpha^{(1)}, \alpha^{(2)}) + (1 - \alpha^{(1)} - \alpha^{(2)})T$. ∎

Proof of Proposition 11. We claim that both platforms compete to offer two plans with advertising intensities $\alpha_N^{1,2}$ and $\alpha_S^{1,2}$ at prices $P_N^{1,2} = 0$ and $P_S^{1,2} = 0$ to naïfs and sophisticates, respectively, where $\alpha_N^{1,2} \in \arg\max_\alpha I_N(\alpha) + (1-\alpha)T$ and $\alpha_S^{1,2} \in \arg\max_\alpha I_S(\alpha) + (1-\alpha)T$. Note that if these two plans are offered then naïfs and sophisticates will choose their respective plan regardless of whether there is another plan $(\tilde{\alpha}, \tilde{P})$ offered. By way of contradiction, suppose that sophisticates participate on platform 1's plan $(\tilde{\alpha}_S^{1,2}, P_S)$ where $\tilde{\alpha}_S^{1,2} \neq \alpha_S^{1,2}$ or $P_S \neq 0$. If $P_S > 0$, then platform 2 has a profitable deviation to offer $(\tilde{\alpha}^{1,2}, P_S - \epsilon)$ for $\epsilon > 0$, whereas if $\tilde{\alpha}_S^{1,2} \neq \alpha_S^{1,2}$, then platform 2 has a profitable deviation to either offer $((\tilde{\alpha}^{1,2} + \alpha_S^{1,2})/2, P_S + \epsilon)$ for $\epsilon > 0$ sufficiently small. The same is true for naïve agents.

By Blackwell's theorem, $I_S(\alpha)$ and $I_N(\alpha)$ are both increasing in ϕ_1 and by Lemma A.1, $I_S(\alpha) = I_N(\alpha) = 0$ when $\phi_1 = 0$. Recall that by the same reasoning as in Lemma 3, we know that $I_N(\alpha) = (1 - e^{-\alpha T})\chi_N(\phi_1)$ and $I_S(\alpha) = (1 - e^{-\alpha T})\chi_S(\phi_1)$ with $\chi_N(0) = \chi_S(0) = 0$. We observe that $I_N'(\alpha) = Te^{-\alpha T}\chi_N(\phi_1) < I_N'(0) = T\chi_N(\phi_1)$. Moreover, there exists some $\phi_1^{\mathcal{P}} \in (0,1)$ such that for all $\phi_1 < \phi_1^{\mathcal{P}}$, $\chi_N(\phi_1) < 1$, implying that $I_N'(\alpha) < T$ for all $\alpha > 0$, thus $\alpha_N^{1,2} = 0$. This in turn suggests that $I_S'(\alpha) < T$ for all $\alpha > 0$, and $\alpha_S^{1,2} = 0$. Therefore, the equilibrium under regime (i) involves $\alpha_N^{1,2} = \alpha_S^{1,2} = 0$ and $P_N^{1,2} = P_S^{1,2} = 0$. In the fully-rational benchmark, all agents are sophisticated, so the firm solves $\max_p (p - c)(\beta q - p)$ by Lemma 1 regardless of the level of advertising in equilibrium. When there is no advertising, the firm solves exactly the same problem, so $p^* = \bar{p}^*$ in regime (i).

We choose $\phi_1^{\mathcal{P}}$ to be maximal so that for any $\phi_1 > \phi_1^{\mathcal{P}}$, $\chi_N(\phi_1) > 1$ and $I_N'(0) > T$, which implies the advertising level to naïve agents in equilibrium will be such that $\alpha_N^{1,2} > 0$. Note that $\alpha_N^{1,2}$ solves $I_N'(\alpha_N^{1,2}) = T$, which necessarily occurs before $\hat{\alpha}^*$ where $I_N(\hat{\alpha}^*) + (1-\hat{\alpha}^*)T - v = 0$. Similarly, because $I_S'(\alpha) < I_N'(\alpha)$, we know that $\alpha_N^{1,2} > \alpha_S^{1,2} = \alpha_S^{FB}$. In the fully-ration benchmark, we have that $\bar{p}^* = (\beta q + c)/2$ whereas the firm will charge $p^* = \frac{1}{2}\left(\lambda\beta q + (1-\lambda)\beta(q\bar{\pi}_{F_{N1}(\alpha_N^{1,2})} + (1-q)\bar{\pi}_{F_{N0}(\alpha_N^{1,2})}) + c\right) > (\beta q + c)/2$ for $\alpha_N^{1,2} > 0$. ∎

Proof of Proposition 12. In the fully-rational benchmark, all agents receive advertising at rate α_S^{FB}, zero subscription fees, and $p^* = \bar{p}^*$. Under regime (a), we have $\alpha_S^{FB} = 0$ so the equilibrium played is exactly that of the fully-rational benchmark; trivially, we have $\hat{W}_{1,2}^*(\tau) = \bar{W}_{1,2}^*(\tau)$ for both $\tau \in \{S, N\}$.

For regime (b), advertising intensity is at $\alpha_1^* = \alpha_S^{FB}$ and $\alpha_2^* \in (\alpha_S^{FB}, \hat{\alpha}^*)$ and prices are at $\hat{p}_{\mathcal{P}}^*$. We have the user welfare for the sophisticated agents is given by

$$\begin{aligned}
\hat{W}_{1,2}^*(S) =& (1-\alpha_S^{FB})T + q\mathbb{E}_{\pi\sim F_{S1}}\left[U(z^*(\pi,\hat{p}_{\mathcal{P}}^*);\,\theta_i=1) - \hat{p}_{\mathcal{P}}^* z^*(\pi,\hat{p}_{\mathcal{P}}^*)\right] \\
&+ (1-q)\mathbb{E}_{\pi\sim F_{S0}}\left[U(z^*(\pi,\hat{p}_{\mathcal{P}}^*);\,\theta_i=0) - \hat{p}_{\mathcal{P}}^* z^*(\pi,\hat{p}_{\mathcal{P}}^*)\right] \\
=& (1-\alpha_S^{FB})T + I_S(\alpha_S^{FB}) + U(C_i(q,\hat{p}_{\mathcal{P}}^*)) - \hat{p}_{\mathcal{P}}^* C_i(q,\hat{p}_{\mathcal{P}}^*) \\
=& \bar{W}_{1,2}^*(S) + \frac{1}{2}(\beta q - \hat{p}_{\mathcal{P}}^*)^2 - \frac{1}{2}(\beta q - \bar{p}^*)^2 < \bar{W}_{1,2}^*(S)
\end{aligned}$$

Consumer welfare for the naïve agents is given by

$$\begin{aligned}
\hat{W}_{1,2}^*(N) =& (1-\alpha_2^*)T + q\mathbb{E}_{\pi\sim F_{N1}}\left[U(z^*(\pi,\hat{p}_{\mathcal{P}}^*);\,\theta_i=1) - \hat{p}_{\mathcal{P}}^* z^*(\pi,\hat{p}_{\mathcal{P}}^*)\right] \\
&+ (1-q)\mathbb{E}_{\pi\sim F_{N0}}\left[U(z^*(\pi,\hat{p}_{\mathcal{P}}^*);\,\theta_i=0) - \hat{p}_{\mathcal{P}}^* z^*(\pi,\hat{p}_{\mathcal{P}}^*)\right] \\
=& (1-\alpha_2^*)T + q\mathbb{E}^S_{\pi_i|\theta_i=1}[U(C_i(\pi_i,\hat{p}_{\mathcal{P}}^*);\,\theta_i=1) - \hat{p}_{\mathcal{P}}^* C_i(\pi_i,\hat{p}_{\mathcal{P}}^*)\,|\,\alpha_2^*] \\
&+ (1-q)\mathbb{E}^S_{\pi_i|\theta_i=0}[U(C_i(\pi_i,\hat{p}_{\mathcal{P}}^*);\,\theta_i=0) - \hat{p}_{\mathcal{P}}^* C_i(\pi_i,\hat{p}_{\mathcal{P}}^*)\,|\,\alpha_2^*] - \Delta(\alpha_2^*,\hat{p}_{\mathcal{P}}^*) \\
=& (1-\alpha_2^*)T + I_S(\alpha_2^*) - \Delta(\alpha_2^*,\hat{p}_{\mathcal{P}}^*) + U(C_i(q,\hat{p}_{\mathcal{P}}^*)) - \hat{p}_{\mathcal{P}}^* C_i(q,\hat{p}_{\mathcal{P}}^*) \\
<& (1-\alpha_S^{FB})T + I_S(\alpha_S^{FB}) + (I_S(\alpha_2^*) - I_N(\alpha_2^*)) + \frac{1}{2}(\beta q - \hat{p}_{\mathcal{P}}^*)^2 \\
<& \bar{W}_{1,2}^*(N) + \frac{1}{2}(\beta q - \hat{p}_{\mathcal{P}}^*)^2 - \frac{1}{2}(\beta q - \bar{p}^*)^2 < \bar{W}_{1,2}^*(N),
\end{aligned}$$

which establishes the welfare claim. ∎

A.5 Proofs from Section 6

Proof of Proposition 13. First, we claim that if the menu offered is (α_1, P_1, p_1) and (α_2, P_2, p_2) with $\alpha_2 > \alpha_1$, and the sophisticates self-select into the plan with advertising, α_2, then so do the naïfs. By assumption, we know that $I_S(\alpha_1) + (1-\alpha_1)T - P_1 + \frac{1}{2}(\beta q - p_1)^2 < I_S(\alpha_2) + (1-\alpha_2)T - P_2 + \frac{1}{2}(\beta q - p_2)^2$, which implies that $(e^{-\alpha_1 T} - e^{-\alpha_2 T})\chi_S > (\alpha_2 - \alpha_1)T - P_1 + P_2 + \frac{1}{2}(\beta q - p_1)^2 - \frac{1}{2}(\beta q - p_2)^2$, where χ_S is defined as in Lemma 3. Of course, this suggests that $(e^{-\alpha_1 T} - e^{-\alpha_2 T})\chi_N > (\alpha_2 - \alpha_1)T - P_1 + P_2 + \frac{1}{2}(\beta q - p_1)^2 - \frac{1}{2}(\beta q - p_2)^2$, and rearranging gives us $I_N(\alpha_1) + (1-\alpha_1)T - P_1 + \frac{1}{2}(\beta q - p_1)^2 < I_N(\alpha_2) + (1-\alpha_2)T - P_2 + \frac{1}{2}(\beta q - p_2)^2$. Second, we claim that if the menu offered is again (α_1, P_1, p_1) and (α_2, P_2, p_2), but the

sophisticates opt for the plan with advertising α_1, then welfare is larger if naïve agents also choose α_1 instead of α_2. Recall that naïfs' welfare can be written as $I_S(\alpha)+(1-\alpha)T-\Delta(\alpha)-P+\frac{1}{2}(\beta q-p)^2$where $\Delta(\alpha)$ is an increasing function in α. However, $I_S(\alpha_1)+(1-\alpha_1)T-P_1+\frac{1}{2}(\beta q-p_1)^2 > I_S(\alpha_2)+(1-\alpha_2)T-P_2+\frac{1}{2}(\beta q-p_2)^2$ and $\Delta(\alpha_1)<\Delta(\alpha_2)$; thus, $I_S(\alpha_1)+(1-\alpha_1)T-\Delta(\alpha_1)-P_1+\frac{1}{2}(\beta q-p_1)^2 > I_S(\alpha_2)+(1-\alpha_2)T-\Delta(\alpha_2)-P_2+\frac{1}{2}(\beta q-p_2)^2$. Putting both of these pieces together implies that welfare-maximizing plan under self-selection involves just one advertising intensity we denote by α^{SB}. It is immediate to see there is just a single subscription price and single product price, with the former equal to zero and the latter equal to marginal cost, in the welfare-maximizing plan.

To maximize welfare, we want to choose α^{SB} such that $\lambda W(S;\alpha^{SB})+(1-\lambda)W(N;\alpha^{SB})$ is maximized. The lower bound on α^{SB} trivially holds when $\alpha_N^{FB}=0$, so suppose $\alpha_N^{FB}>0$. Then we know that $\lambda W'(S;\alpha_N^{FB})+(1-\lambda)W'(N;\alpha_N^{FB})>0$, so is dominated by choosing some $\alpha_N^{FB}+\epsilon$ for small $\epsilon>0$. Likewise, if $\alpha_S^{FB}=0$, then it straightforward to see $\alpha^{SB}=0$, so let us take α_S^{FB}. Then we know that $\lambda W'(S;\alpha_S^{FB})+(1-\lambda)W'(N;\alpha_S^{FB})<0$, so is dominated by choosing some $\alpha-\epsilon$ for small $\epsilon>0$. This implies that α^{SB} lies somewhere in the interval $[\alpha_N^{FB},\alpha_S^{FB}]$.

Finally, we claim that $W_{FB}(\tau) > W_{SB}(\tau) \geq W_{\text{base}}(\tau)$ for both $\tau\in\{S,N\}$. The latter inequality follows from noting that a feasible option for the planner in second-best is to offer a prohibitively expensive subscription plan that neither type opts into, which yields the same welfare as the base case. For the former inequality, the proof follows by construction of α^{SB}, α_S^{FB}, and α_N^{FB}. We know that when $\alpha_S^{FB}>0$, then $\alpha^{SB}\in(\alpha_N^{FB},\alpha_S^{FB})$, so for sophisticated agents we have $I_S(\alpha^{SB})+(1-\alpha^{SB})T < I_S(\alpha_S^{FB})+(1-\alpha_S^{FB})T$. For naïve agents we have that $\alpha^{SB}>\alpha_N^{FB}$, so for naïfs we have $I_S(\alpha^{SB})+(1-\alpha^{SB})T-\Delta(\alpha^{SB}) < I_S(\alpha_N^{FB})+(1-\alpha_N^{FB})T-\Delta(\alpha^{SB})$. ∎

Proof of Proposition 14. Consider $m^*(\alpha)$ to be the transfer to the platform as a function of ad intensity α, while holding all model primitives constant. We know that $\sup_{\alpha\in[0,1]}\partial m^*(\alpha)/\partial\alpha<L$ for some constant L because total demand $(\lambda\beta q+(1-\lambda)\beta(q\bar{\pi}_{F_{N1}(\alpha)}+(1-q)\bar{\pi}_{F_{N0}(\alpha)})-p)$ has a bounded derivative in α. Second, we see that the platform's potential gain in subscription revenue from the agents of type τ for all $\alpha>\alpha^{SB}$ is given by $(\lambda\mathbb{1}_{\tau=S}+(1-\lambda)\mathbb{1}_{\tau=N})\int_{\alpha^{SB}}^{\alpha}(I'_\tau(x)-T)\,dx=(\lambda\mathbb{1}_{\tau=S}+(1-\lambda)\mathbb{1}_{\tau=N})(I_\tau(\alpha)-I_\tau(\alpha^{SB})-T(\alpha-\alpha^{SB}))$, which has a derivative bounded above by $I'_\tau(\alpha)<I'_N(0)<\infty$ by Lemma 3. Thus, taking $\bar{\mu}=I'_N(0)+L$, we see that the platform's revenue from setting $\alpha>\alpha^{SB}$ is upper bounded by $R(\alpha)=m^*(\alpha)-m^*(\alpha^{SB})+\lambda(I_S(\alpha)-I_S(\alpha^{SB})-T(\alpha-\alpha^{SB}))+(1-\lambda)(I_N(\alpha)-I_N(\alpha^{SB})-T(\alpha-\alpha^{SB}))$ which is strictly less than the tax $\mu(\alpha-\alpha^{SB})$ for all $\alpha>\alpha^{SB}$. This implies that the platform will set an ad intensity no greater than α^{SB} in any of its plans. To see that it will never set an ad intensity strictly less than α^{SB}, note that because $\alpha^{SB}\leq\alpha_S^{FB}$, if $\alpha^*<\alpha^{SB}$ in one of the offered plans, the platform could instead offer $((\alpha^*+\alpha^{SB})/2,P^*+\epsilon)$ for sufficiently small ϵ and both the sophisticates and naifs would prefer $((\alpha^*+\alpha^{SB})/2,P^*+\epsilon)$ to (α^*,P^*). Moreover, because we know that $m^*(\alpha)$ is monotonically increasing in α, both its advertising and subscription revenue would increase if the platform instead offered $(\alpha^*+\alpha^{SB})/2$ instead of α^* (and it would not be subject to the tax). This represents a profitable deviation; thus, the platform implements exactly one plan and it is at α^{SB}.

Getting the product price down to marginal cost c in equilibrium can be accomplished with a per-unit subsidy δ as follows. Let $Z(\alpha^{SB})$ denote the demand for the good as a function of the advertising intensity to naïfs (recall by Lemma 1 that aggregate demand of sophisticates is unaffected by their

advertising intensity). Then the firm solves $\max_p (p-c)(Z(\alpha^{SB})-p)$. Instead let us provide a subsidy to the good in the amount $\delta = Z(\alpha^{SB})-c$. Then p solves $\arg\max_p (p-c+\delta)(Z(\alpha^{SB})-p) = \frac{1}{2}(Z(\alpha^{SB})+c-\delta) = c$, as desired.

Finally, we show that we can implement zero subscription fees with the appropriate per-user subsidy. Consider if the platform offers one plan at advertising intensity α^{SB}, and suppose it charges subscription fee $P^* = I_N(\alpha^{SB})+(1-\alpha^{SB})T-v$ to maximally extract consumer surplus from the naive agents (with sophisticates refraining from participation). Then setting $\bar{\eta} = P^*$ means that if $\eta > \bar{\eta}$, the platform can generate the most revenue from each user conditional on advertising at α^{SB} by charging no subscription fee, which implements the second best plan of $P^{SB} = 0$.

Proof of Proposition 15. We first argue that in equilibrium, there is a plan offered with $\alpha^* > \alpha_N^* = \arg\max_\alpha I_N(\alpha) + (1-\alpha)T$. Let us show the weak inequality. The claim holds vacuously if $\alpha_N^* = 0$, so suppose $\alpha_N^* > 0$ and the platform offers (α_2, P_2) which is selected by the naïfs, with $\alpha_S^{FB} < \alpha_2 < \alpha_N^*$ (the result is trivial when $\alpha_2 < \alpha_S^{FB}$, because the platform can extract more subscription and revenue from both sophisticates and naifs by offering a plan at least at α_S^{FB}). For the same reasons as in the proof of Proposition 14, because $I_N'(\alpha_2) > T$, the platform could instead offer the plan $((\alpha_2+\alpha_N^*)/2, P_2+\epsilon)$ for sufficiently small $\epsilon > 0$ and naïfs would prefer $((\alpha_2+\alpha_N^*)/2, P_2+\epsilon)$ to (α_2, P_2), which leads to strictly higher subscription revenue and advertising revenue $m^*(\alpha)$, which is a profitable deviation (as sophisticates would remain in their same plan). To show it is a strict inequality note that $\partial m^*(\alpha)/\partial\alpha > 0$ for all α, but the loss in subscription revenue from naifs is equal to $(1-\lambda)(T(\alpha-\alpha_N)-(I_N(\alpha)-I_N(\alpha_N^*))$, with $\frac{\partial}{\partial\alpha}(1-\lambda)(T(\alpha-\alpha_N)-(I_N(\alpha)-I_N(\alpha_N^*)) = T-I_N'(\alpha)$ which evaluated at $\alpha = \alpha_N^*$ is 0. Thus, there exists some small $\kappa > 0$ such that $m^*(\alpha_N^*+\kappa)-(1-\lambda)(T\kappa-(I_N(\alpha_N^*+\kappa)-I_N(\alpha_N^*)) > m^*(\alpha)$, thus the platform can generate more revenue by setting its ad intensity to at least $\alpha_N^*+\kappa$, establishing the strong inequality.

Next, we show that linear digital ad tax is possible of implementing advertising intensity $(\alpha_N^*+\alpha^*)/2$, which leads to lower consumer welfare because $\alpha^* > \alpha_N^* > \alpha_S^{FB}$ and product price p^* is increasing in the ad intensity of the naifs' plan. With a linear digital ad tax, the platform will solve $\max_{\alpha\geq\alpha_N^*}(1-\gamma)m^*(\alpha)-(1-\lambda)(T(\alpha-\alpha_N)-(I_N(\alpha)-I_N(\alpha_N^*)))$. Note that $\partial(\beta q\bar{\pi}_{F_{N1}(\alpha)}+\beta(1-q)\bar{\pi}_{F_{N0}(\alpha)})/\partial\alpha$ is bounded for all α, therefore there exists some $L > 0$ such that $\partial m^*/\partial\alpha < L$. But then notice that

$$\frac{\partial}{\partial\alpha}(1-\gamma)m^*(\alpha)-(1-\lambda)(T(\alpha-\alpha_N)-(I_N(\alpha)-I_N(\alpha_N^*))) < (1-\gamma)L-(1-\lambda)(T-I_N'((\alpha^*+\alpha_N^*)/2)$$

which is less than 0 for $\gamma > 1-\frac{(1-\lambda)(T-I_N'((\alpha^*+\alpha_N^*)/2))}{L}$ which is strictly between 0 and 1. Given this restriction on γ, the platform will always opt to select an advertising intensity less than or equal to $(\alpha^*+\alpha_N^*)/2$. Finally, the platform can maximize revenue from the sophisticates by offering the first-best plan α_S^{FB}, which yields the highest subscription fee the platform can extract from them (and note by Lemma 1 there are no transfers from advertising). To ensure sophisticates and naifs participate in their respective plans, the platform then just sets the subscription prices (P_1, P_2) to maximize $\lambda P_1 + (1-\lambda)P_2$ subject to the incentive compatibility constraints from Proposition 7. ∎

B Kinked Linear Demand

In the main text we had assumed β is large, allowing us to drop the $z_i \geq 0$ consumption restriction. Here, we relax this and suppose β is arbitrary and users are constrained to only non-negative consumption. In particular, agents now solve $\max_{z_i \geq 0} \mathbb{E}^{\tau_i}[U(z_i;\theta_i) - pz_i]$, which yields a kinked demand curve $z_i^* = (\beta\pi_i - p)_+$, instead of our previously linear demand curve $z_i^* = \beta\pi_i - p$. Our results are impacted as follows.

First, it is possible to extract some surplus from sophisticated agents by advertising to them when demand is kinked. This follows from Jensen's inequality, noting that $\mathbb{E}^S[(\beta\pi_i - p)_+] > (\beta\mathbb{E}^S[\pi_i] - p)_+ = (\beta q - p)_+$. Thus, in the baseline model, the platform is no longer completely indifferent between sophisticated agents' participation and not, as advertising to them leads to profits $\Pi_S(\alpha)$ which are increasing in α for the same reasons as Lemma 2. By Lemma 3, we still know that naïfs will always participate on the ad-based plan whenever sophisticates do, but not necessarily vice-versa. When an ad-based platform decides on the advertising intensity, it trades off a lower ad intensity $\hat{\alpha}' < \hat{\alpha}^*$ that keeps sophisticates on the platform (capturing informational surplus from a greater fraction of the population) with the higher ad intensity $\hat{\alpha}^*$ that sacrifices Π_S but can extract more surplus from naïfs at intensity $\hat{\alpha}^*$. This implies that our Proposition 2 now turns on both λ and ϕ_0. Fixing ϕ_0 (and other model primitives), there exists λ^* such that if $\lambda < \lambda^*$ the platform chooses ad intensity $\hat{\alpha}^*$ with $I_N(\hat{\alpha}^*) + (1 - \hat{\alpha}^*) - v = 0$ just as in the baseline model under Lemma 4. However, when $\lambda > \lambda^*$, there are sufficiently many sophisticates and the platform prefers to retain their participation, instead choosing $\hat{\alpha}'$ where $I_S(\hat{\alpha}') + (1 - \hat{\alpha}')T - v = 0$. Within each of these regimes, the dependence on ϕ_0 is different, however. For sufficiently small values of λ, we recover exactly the cutoff structure of Proposition 2, where a higher rate of false-positives results in the platform adopting an ad-based model with ad intensity $\hat{\alpha}^*$. On the other hand, when λ is sufficiently close to 1, we end up with a flipped cutoff structure: There exists $\hat{\hat{\phi}}_0$ such that the ad-based business model is adopted if and only if $\phi_0 < \hat{\hat{\phi}}_0$. The reason here is that Blackwell informativeness is one-to-one with Π_S, so the platform can extract more informational rent if the ads themselves are more informative, which happens when the gap between ϕ_0 and ϕ_1 is larger. Regardless of which regime we are in, we still have our central result from Proposition 4: The platform will leave the sophisticated agents with no excess surplus from platform usage, and product prices will be higher under an ad-based business model, so user welfare will fall below base case levels when there is advertising.

Most of these insights generalize immediately to mixed platform business models, competition, and the policy analysis. When the platform can offer a menu that segments sophisticates and naïfs, there will be a pair of cutoffs (ϕ_0^*, ϕ_0^{**}) where $\phi_0 > \phi_0^*$ leads to an advertising plan for naïve agents (as in Proposition 5) and where $\phi_0 < \phi_0^*$ leads to an advertising plan for sophisticated agents (as opposed to always offering a subscription-based plan to sophisticates). Because competition forces the firms and platforms to cater more to users (rather than extract full surplus), both Proposition 9 and 11 apply identically. Finally, our policy implications remain fully intact: The first, second, and third-best are allocations, so are unaffected by the surplus the platform can extract from sophisticates, and the flat digital ad tax implements the same exact third-best allocation as before.

C Digital Ad Taxation: Multiple Firms and Platforms

We consider how our digital ad taxation policy changes in the case of multiple firms (N) and platforms (M). When there are multiple firms, there will generally be a menu of advertising vectors $(\alpha^{(1)}, \ldots, \alpha^{(N)})$ instead of just a menu of advertising intensities α. The informational value from advertising will generally take the form of $I_S(\alpha^{(1)}, \ldots, \alpha^{(N)}) = \sum_{j \in \{0,1\}^N} \prod_{k=1}^{N} (e^{-\alpha^{(j)}T}(1 - j_k) + (1 - e^{-\alpha^{(j)}T}) j_k) \chi_S^{(j)}$ (for some $\chi_S^{(j)}$ that depend on model primitives) and analogously for $I_N(\alpha^{(1)}, \ldots, \alpha^{(N)})$. To solve for the first-best level of advertising, we maximize $I_S(\alpha^{(1)}, \ldots, \alpha^{(N)}) + (1 - \sum_{j=1}^{N} \alpha^{(j)} T)$ for sophisticates and maximize $I_S(\alpha^{(1)}, \ldots, \alpha^{(N)}) + (1 - \sum_{j=1}^{N} \alpha^{(j)} T) - \Delta(\alpha^{(1)}, \ldots, \alpha^{(N)})$ for naïfs, for an appropriately defined Δ under the multi-advertiser case, just as in the proof of Proposition 1. It is easy to see from there that there exists a unique α_S^{FB} such that $(\alpha_S^{FB}, \ldots, \alpha_S^{FB})$ is first-best for sophisticates, and a corresponding unique $\alpha_N^{FB} \leq \alpha_S^{FB}$ such that $(\alpha_N^{FB}, \ldots, \alpha_N^{FB})$ is first-best for naïve agents. Once again, self-selection constraints from the users will prevent us from implementing the first-best level of advertising, and there will be a tension that makes it impossible to separate sophisticates and naïfs in the second-best, leading to the offering of a single plan which takes the form of $(\alpha^{SB}, \ldots, \alpha^{SB})$ with $\alpha_N^{FB} \leq \alpha^{SB} \leq \alpha_S^{FB}$. Note that because our first-best and second-best are allocations, and users always participate on at most one plan, the existence of multiple platforms has no impact on our analysis.

To implement second-best, we consider the case of a single platform with multiple firms and multiple platforms and potentially multiple firms separately. For a single platform, the policymaker can similarly levy a sufficiently high advertising tax on ad quantities above α^{SB} for any individual advertiser. Note that in general it will not be sufficient to regulate total advertising, because the platform may not play a symmetric advertising strategy, so taxing the sum of advertising above $N\alpha^{SB}$ may not implement the second-best. However, a policy of the form of Proposition 14 with the tax applying to each individual advertiser will have the same desired effect as with the single advertiser case. With multiple platforms, our analysis from Proposition 11 extends here to show that the platforms will compete to offer the plans that are viewed as most desirable to sophisticates and naïve agents. In other words, the platforms will offer one plan for sophisticates with no subscription fee that has advertising at $(\alpha_S^{FB}, \ldots, \alpha_S^{FB})$ and will offer a second plan to naïfs with no subscription that has advertising at $(\alpha_N^*, \ldots, \alpha_N^*)$with $\alpha_N^* > \alpha^{SB}$, which is seen as maximizing naïfs' utility under their subjective model. In such a setting, a similar policy as in the single platform case will implement second-best, but the tax rate may need to be even higher. This is because competition drives platforms not to offer the most profitable business models, but the ones that are seen as utility-maximizing by the agents. Thus, the tax must be strict enough that platforms who advertise above $(\alpha^{SB}, \ldots, \alpha^{SB})$ in fact operate at loss, as opposed to just making less profit than when $(\alpha^{SB}, \ldots, \alpha^{SB})$ is played.