

November 2025

“Mechanism Design for Ad-Supported Platforms”

Shota Ichihashi, Doh-Shin Jeon and Byung-Cheol Kim

Mechanism Design for Ad-Supported Platforms*

Shota Ichihashi Doh-Shin Jeon Byung-Cheol Kim

November 10, 2025

Abstract

Many digital platforms earn revenues by selling access to content and displaying ads. We study monopoly screening for a platform allocating heterogeneous content and ads to heterogeneous consumers. The platform faces a trade-off between advertising and rent extraction: Broadening content access increases advertising revenue but reduces access revenue by raising information rents. Unlike standard screening, the platform serves consumers with negative virtual types but allocates content and ads to limit the reduction in access revenue. The optimal mechanism rationalizes contracts used by diverse ad-supported platforms and explains why higher ad profitability might curtail their incentives to invest in content quality.

*Ichihashi: Queen's University, Department of Economics, shotaichihashi@gmail.com; Jeon: University of Toulouse Capitole, Toulouse School of Economics, dohshin.jeon@tse-fr.eu; Kim: University of Alabama, Culverhouse College of Business, bkim34@ua.edu. We are grateful to Gaurab Aryal, Jacques Crémer, Sungha Hwang, Bruno Jullien, Michihiro Kandori, Heiko Karle, Jeong-Yoo Kim, Jinwoo Kim, Adrien Raizonville, Markus Reisinger, Patrick Rey, Andrew Rhodes, Robert Somogyi, Tat-How Teh, Yiqing Xing, and Yuichi Yamamoto, and audiences at the ACE 2024 (Rome), IIOC 2025, Postal Economics Conference on E-commerce, Digital Economics and Delivery Services 2024 at TSE, the Workshop on Search and Platforms (Kyoto), and the KER International Conference (2024), as well as seminar participants at KAIST, Keio University, Kyung Hee University, Nanyang Technological University, Peking University, Renmin University, Seoul National University, the University of Tokyo, and Yonsei University. Jeon acknowledges the funding from ANR under grant ANR-17-EUR-0010 (Investissements d'Avenir program). This paper supersedes our previous working paper, [Ichihashi, Jeon, and Kim \(2024\)](#). An earlier version was presented at the Paris Conference on Digital Economics 2022 and the MaCCI Annual Conference 2022 at Mannheim.

Keywords: ad-supported platform, mechanism design, innovation, quality

JEL Codes: D42, D82, L15, O31

1 Introduction

Digital platforms adopt diverse business models, and evaluating consumer harm from their market power requires accounting for these differences (Scott Morton, Etro, Latham, and Caffarra, 2020). Existing reports on digital platforms¹ focus especially on ad-supported platforms and agree that their market power can harm consumers through lower quality and reduced innovation. However, these reports do not distinguish purely ad-funded platforms from ad-supported hybrid ones, which earn revenue from both selling access to content and advertising. This paper addresses the question of how hybrid platforms optimally balance their multiple revenue sources and whether we need to be concerned about their incentives to invest in quality.

There are many ad-supported hybrid platforms—such as Netflix, YouTube, Facebook, and The New York Times. They typically offer consumers a menu of plans, which vary in how heavily they rely on each source of revenue. For example, the New York Times offers free and paid plans, both of which are ad-supported and grant differential access to articles. In contrast, X (formerly Twitter) grants free access to all content and charges users to remove ads. Streaming services—such as Netflix and Prime Video—offer plans that differ in both content access and advertising intensity.

In this paper, we study a mechanism design problem for an ad-supported hybrid platform. The optimal mechanism we derive rationalizes various monetization strategies based on a unified economic force: the *trade-off between advertising and rent extraction*. To monetize attention through advertising, the platform must broaden content access. However, broadening access increases information rents of consumers and limits the platform’s ability to extract their surplus through content pricing (Mussa and Rosen, 1978; Myerson, 1981). This is a classical rent extraction-efficiency trade-off, adapted to digital platforms that control content access and advertising (cf. Laffont and Martimort 2009). The trade-off also yields a general intuition for why better ad-

¹Examples include Competition & Markets Authority (2020), Crémer et al. (2019), and Stigler Committee on Digital Platforms (2019).

vertising technologies or lax regulations surrounding advertising may curtail platforms' incentives to invest in content quality. This insight is highly relevant to policy. For example, when applied to the media industry, where many news organizations use the hybrid business model, our insight implies that better advertising technology may lead to lower-quality journalism.

Our model is a version of a monopoly screening problem. The platform hosts items and ads. Items represent content, such as news articles, videos, or social media posts. Items are vertically differentiated. Ads are differentiated along two dimensions: (i) the revenue they generate for the platform and (ii) their disutility level. Consumers have a one-dimensional private type that represents their taste for item quality and distaste for ads. Their utility depends on monetary transfers, allocated items, and ads.

The platform designs a menu of contracts (or a mechanism) that specifies, for each consumer type, a set of items to be allocated, an advertising policy, and a monetary transfer to the platform. An advertising policy determines whether to match each item with an ad, and if so, which one. An ad reduces the net quality of the matched item according to its disutility level. The platform maximizes total revenue—i.e., the sum of (i) advertising revenue generated by ads allocated to consumers, and (ii) monetary transfers from consumers, which we call sales revenue.

The optimal mechanism balances the trade-off between advertising and rent extraction in two ways. First, unlike in standard screening models, the platform may serve consumers with negative virtual types to earn ad revenue. At the same time, the platform designs content allocation and advertising policies for negative virtual types to reduce information rents. Specifically, the platform matches every item with some ad, because ads not only generate revenue but also reduce the quality of associated items and thereby information rents. Also, the increase in information rents is smaller when the allocated items are of low quality. As a result, the platform only allocates items whose quality levels are below a type-dependent threshold.

Second, the optimal mechanism bundles items with ads based on type-dependent

virtual advertising profits. In general, ads affect the platform’s total revenue in three ways: They generate advertising revenue, reduce consumers’ willingness to pay for the associated items, and reduce information rents of higher types. The magnitudes of these effects depend on the consumer’s type and the characteristics of ads. The virtual advertising profit captures their net effect as a single value for each ad, and the platform uses it to determine how to bundle items and ads for each consumer type.²

The optimal mechanism rationalizes various menus of contracts (Section 5). In practice, ad-supported platforms engage in widely different forms of discrimination in terms of content access and ad exposure. Some platforms differentiate only in terms of content access, e.g., the Wall Street Journal and the Financial Times restrict article access for non-subscribers while showing ads to everyone. Other platforms differentiate only in ad exposure; for example, YouTube and X remove ads for paying users while keeping content access essentially the same. Still other platforms, such as Netflix, Spotify, and Peacock, discriminate along both dimensions—content access and ad exposure. The optimal mechanism takes any of the three forms under certain conditions, and we argue that these conditions align with the characteristics of typical platforms that use each type of menu.

The optimal mechanism also shapes the platform’s incentives to invest in content quality. In the baseline model, the platform takes the quality of each item as given. In Section 6, we allow the platform to choose item quality at a cost. We first show that if virtual advertising profits uniformly increase across all ads and consumer types, the platform invests less in content quality. We then show that such a change occurs if the platform hosts a larger set of ads or faces a higher ad revenue associated with each

²More precisely, for positive virtual types, advertising reduces information rents, which induces the platform to serve too many ads. For negative virtual types, by contrast, the bundling of ads with content items—combined with the under-provision of items—tends to make the platform serve too few ads. In addition, there is another distortion for negative virtual types. Suppose all ads generate the same revenue. Given a fixed number of ads, the platform allocates ads with the highest nuisance levels, whereas social surplus is maximized by allocating ads with the lowest nuisance levels. This occurs because the virtual advertising profit increases with nuisance for negative virtual types but decreases for positive ones. For more details, see the discussion after Theorem 1.

ad. The intuition is linked to the rent extraction-advertising trade-off. If advertising becomes more profitable, the platform allocates more items to types with negative virtual valuation, and the resulting increase in information rents is larger when the allocated items have higher quality. Thus, greater advertising profitability makes it less attractive for the platform to invest in content quality.

The results on the platform’s investment incentives have policy implications. First, the results allow us to view various policy and technological changes that surround advertising in terms of their impact on platforms’ content quality. In [Section 7](#), we illustrate this point by examining a recent bill proposed in California that aims to limit the loudness of ads on streaming platforms. Second, our results provide a general rationale for the existing concerns that ad-supported business models may curtail platforms’ incentives to invest in content quality ([Competition & Markets Authority, 2020](#); [Stigler Committee on Digital Platforms, 2019](#)). More precisely, we show that the concerns are relevant for both purely ad-funded platforms and hybrid platforms. Our result is particularly relevant for the media industry because most news organizations use hybrid business models. In fact, there have been considerable concerns about the impact of advertising on the quality of journalism ([OECD, 2021](#); [Latham et al., 2022](#)).

Our contribution is twofold. First, we provide a general framework that rationalizes a spectrum of business models adopted by digital platforms, from social media and streaming services to news organizations. Second, we uncover a determinant of the platforms’ investment incentives and link it to policy or technological changes surrounding advertising. The paper contributes to policy and academic discussions that recognize business models and monetization strategies as fundamental for understanding digital platforms and formulating relevant policies ([Caffarra, 2019](#); [Scott Morton, Etro, Latham, and Caffarra, 2020](#)).

2 Related Literature

We propose a mechanism design problem for ad-supported platforms with three screening instruments—content access, advertising, and price. We uncover how the platform balances the trade-off between advertising and rent extraction by tailoring screening instruments for consumers, including those with negative virtual types. The optimal mechanism rationalizes observed platform strategies and explains why the increased profitability of advertising might reduce the platform’s content quality.

Monopoly Screening. Our work is related to monopoly screening and, more broadly, to mechanism design for selling goods (Mussa and Rosen, 1978; Myerson, 1981). The key departure from standard screening models is that the optimal mechanism may serve consumers with negative virtual values. This property also arises in two-sided screening in which serving negative virtual values generates positive externalities for other agents (e.g., Damiano and Li (2007), Johnson (2013b), Choi, Jeon, and Kim (2015), Gomes and Pavan (2016, 2024), Jeon, Kim, and Menicucci (2022), Corrao, Flynn, and Sastry (2023), Bar-Isaac, Deb, and Mitchell (2025)).

Compared with these studies, our model is simpler in that we do not model advertisers as strategic agents.³ However, we allow richer heterogeneity in content and ads and give the platform multiple screening instruments. This richness enables us to understand digital platforms from an angle that is relatively underexplored. First, we rationalize various ad-supported contracts observed in practice, including those that discriminate among consumers in both content access and advertising exposure—a common business model for streaming services. Second, we examine how advertising affects the platform’s investment in content quality, without imposing a priori restrictions on which business model the platform adopts.

Our work also differs from the literature on damaged goods (Deneckere and McAfee,

³The way we capture an advertising market is closer to the approach taken by Corrao, Flynn, and Sastry (2023), who, as one of their various applications, capture advertising in a reduced-form way with an external revenue function for the seller.

1996). In our model, the platform can “damage” its offering in several ways—e.g., granting access to fewer or lower-quality content items and displaying annoying ads—and we characterize how the platform combines these instruments.

Advertising-funded Platform. In the platform-economics literature, several papers study how ad-funded platforms discriminate among users by advertising exposure (Sato, 2019; Lin, 2020; Zenny, 2020; Cai and Spulber, 2023). This body of work typically abstracts from the platform’s content-allocation problem and costly investment in quality. Our model accommodates these features, which allow broader applications.

Our result on content quality (Proposition 1) is related to the platform-design literature, which studies platforms’ incentives to design policies such as service quality, the quality of hosted sellers, and the degree of competition within the platform (e.g., Casner (2020), Liu, Yildirim, and Zhang (2022), Teh (2022), Johnen and Somogyi (2024), Madio and Quinn (2024)). In particular, Etro (2021) shows that a purely ad-funded platform underinvests in quality compared with a device-funded platform. Choi and Jeon (2023) compare a purely ad-funded business model with a hybrid one in terms of design bias and find that the former is biased toward the advertising side whereas the latter is biased toward the consumer side. Although their first result is aligned with that of Etro (2021) and ours, their second result is opposite to ours.⁴ Their major difference with respect to our paper is that they consider homogeneous consumers and hence do not consider menus. Therefore, they cannot capture our main trade-off between advertising and rent extraction and this is why our prediction about the effect of the hybrid business model on innovation incentive is opposite to theirs.

Finally, the literature on media markets studies how a platform’s advertising decision can deviate from the social optimum and affect markets for content and products (Gabszewicz et al., 2004; Anderson and Coate, 2005; Peitz and Valletti, 2008; Berge-

⁴If we think that the opportunity cost of investment in content quality is investment in advertising technology, then their result implies that a hybrid platform invests too much in quality while a purely ad-funded one invests too little in quality.

mann and Bonatti, 2011; Prat and Valletti, 2022). We complement this literature by providing an intuition—that consumers’ information rents can incentivize the platform to provide a socially excessive level of “bads,” i.e., advertising.

3 Model

A monopoly platform allocates items and ads to consumers. Items vary in quality, while ads differ in both the revenue they generate for the platform and the disutility they impose on consumers. Each consumer has a one-dimensional private type that determines their preferences for item quality and aversion to ads. The platform offers a menu of contracts, each of which specifies the set of items, matching between items and ads, and monetary transfer from consumers to the platform. The platform maximizes the total revenue from consumer payments and advertising.

Items and Ads. The platform hosts a fixed set of items, denoted by $\mathcal{I} := \{1, \dots, I\}$, and a fixed set of ads, denoted by $\mathcal{J} := \{1, \dots, J\}$, where $I, J \in \mathbb{N}$. The quality of each item $i \in \mathcal{I}$ is denoted by $q(i) \geq 0$. Without loss, assume that item i has the i -th lowest quality, i.e., $q(1) \leq q(2) \leq \dots \leq q(I)$. The platform’s marginal cost of providing access to items is equal to zero as we consider digital content (see Section 3.1). Each ad $j \in \mathcal{J}$ is characterized by its revenue $r(j) \geq 0$ and disutility level $d(j) \geq 0$.

We assume that $d(j) \geq 0$ for every $j \in \mathcal{J}$, i.e., ads reduce consumer utility. The assumption—that ads are utility-decreasing bads—is also adopted by various papers and consistent with the idea that advertising is an implicit price consumers pay for using a digital service (e.g., Stigler Committee on Digital Platforms 2019, p. 63).⁵ However, many of our results—including those on the characterization of the optimal

⁵For papers that adopt a similar assumption, see, e.g., Anderson and Coate (2005); Anderson and Gans (2011); Johnson (2013a); Gomes and Pavan (2016); Sato (2019). Our assumption does not exclude cases in which some ads bring (unmodeled) benefits to consumers, such as informing them about new products. The assumption of nonnegative disutility only implies that the nuisance from ads exceeds other possible benefits of ads, the net effect of which is reflected in each $d(j) \geq 0$.

mechanism (Lemma 1) and the platform’s quality choice (Proposition 1)—hold even if some ads have negative disutility levels. That said, the key intuitions become clearer and our model’s predictions have stronger applied relevance when consumers dislike ads. For this reason, we assume upfront that all ads entail nonnegative disutility.

We also assume that the net quality of any pair of an item and an ad is nonnegative: For every $(i, j) \in \mathcal{I} \times \mathcal{J}$, $q(i) - d(j) \geq 0$. The condition ensures that a consumer’s utility is monotone in their type. The assumption mirrors the one in standard one-dimensional screening models, in which quality or quantity is nonnegative.

Contracts. The platform offers consumers a menu of contracts. A contract is a tuple (T, B, π) , where $T \in \mathbb{R}$ denotes the monetary transfer from a consumer to the platform, and $B \subseteq \mathcal{I}$ denotes the set of allocated items.

The last component of the contract, π , denotes an *advertising policy*, which specifies a one-to-one matching between items and ads with the possibility of matching some items with no ads. Formally, an advertising policy π is a map

$$\pi : \mathcal{I} \rightarrow \mathcal{J} \cup \{\emptyset\} \tag{1}$$

such that there are no distinct items $i, i' \in \mathcal{I}$ with $\pi(i) = \pi(i') \in \mathcal{J}$. Given an advertising policy π , each item $i \in \mathcal{I}$ is bundled with ad $\pi(i)$, where $\pi(i) = \emptyset$ means that item i is not bundled with any ad. We call any element of $\mathcal{I} \times (\mathcal{J} \cup \{\emptyset\})$, including (i, \emptyset) , an item-ad pair. Let \mathcal{M} denote the set of all advertising policies.

One-to-one matching implies that (i) each item can be tied with at most one ad, and (ii) each ad can be shown to each consumer at most once. In Section 8.3, we discuss how we may relax the first restriction (i). The second restriction (ii) is without loss of generality, because ads in \mathcal{J} can represent copies of the same ad. For example, if the platform hosts a single ad but can show it J times to a given consumer, we can

interpret each ad $j \in \mathcal{J}$ as the j -th opportunity to display the ad.⁶

According to (1), the domain of an advertising policy π is always equal to the set of all items, \mathcal{I} , regardless of the underlying set of items B specified in the contract. We might think this is redundant because, as will become clear later, how π varies outside of B does not affect a consumer’s utility from a contract. However, we take the domain of π to be \mathcal{I} instead of B to streamline the exposition and simplify proofs.

Consumers. A unit mass of consumers interacts with the platform. Each consumer is privately informed of her type θ , which is distributed according to $F \in \Delta[0, 1]$ that has a positive density f on $[0, 1]$. The platform knows the type distribution. Assume that the virtual value, $v(\theta) := \theta - \frac{1-F(\theta)}{f(\theta)}$, is strictly increasing and continuous in θ . Define $\theta^0 \in (0, 1)$ as the unique type that has zero virtual value, $v(\theta^0) = 0$.⁷

A consumer’s utility from a contract (T, B, π) is defined as follows:

$$\theta \sum_{i \in B} [q(i) - d(\pi(i))] - T \quad (2)$$

with the notational convention that $d(\emptyset) = 0$.⁸ Thus, consumers who place greater value on content quality (i.e., those with higher types) also experience greater disutility from ads. In Section 3.1, we motivate our specification and discuss alternatives. Also, we assume that consumers cannot ignore ads while consuming items.⁹

Mechanism Design. The platform chooses a direct mechanism, represented by

$$\{(T(\theta), B(\theta), \pi(\cdot|\theta))\}_{\theta \in [0, 1]},$$

⁶By taking $r(j)$ to be decreasing in j , we can capture a situation in which the effectiveness of advertising declines over time as the platform shows a given ad repeatedly to the same consumer.

⁷The assumption—that $v(\theta)$ is strictly increasing and continuous—ensures that θ^0 uniquely exists. If $v(\cdot)$ is nonmonotone, we use the ironed virtual valuation (Myerson, 1981).

⁸Recall that $d(\pi(i))$ is the disutility level of ad $\pi(i)$, which is matched with item $i \in B$. If item i is not matched with any ad (i.e., if $\pi(i) = \emptyset$), we have $d(\pi(i)) = 0$.

⁹See, e.g., Anderson and Gans (2011) and Johnson (2013a), for implications of costly ad-avoidance technologies on market outcomes.

which for each type θ specifies a contract, i.e., monetary transfer $T(\theta)$ to the platform, item bundle $B(\theta)$, and advertising policy $\pi(\cdot|\theta)$. Aside from incentive compatibility and individual rationality constraints, we impose no restrictions on feasible mechanisms.¹⁰ Thus, in general, consumers may receive different sets of items, see different ads even for the same item, and pay different prices.

Given any direct mechanism, if a consumer with type θ reports to be θ' , she obtains a payoff of

$$\theta \sum_{i \in B(\theta')} [q(i) - d(\pi(i|\theta'))] - T(\theta'). \quad (3)$$

If all consumers report their type truthfully, the platform's revenue is

$$\int_0^1 T(\theta) + \sum_{i \in B(\theta)} r(\pi(i|\theta)) dF(\theta).$$

The first term $T(\theta)$ of the integrand is the monetary transfer from type θ . This is the revenue from selling access to content and we call it the sales revenue. The second term $\sum_{i \in B(\theta)} r(\pi(i|\theta))$ is the total advertising revenue generated by type θ .

The platform chooses a mechanism to maximize total revenue subject to incentive compatibility (IC) and individual rationality (IR) constraints. The resulting problem is as follows.

$$\max_{\{(T(\theta), B(\theta), \pi(\cdot|\theta))\}_{\theta \in [0,1]}} \int_0^1 T(\theta) + \sum_{i \in B(\theta)} r(\pi(i|\theta)) dF(\theta) \quad (M)$$

subject to
$$\theta \sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))] - T(\theta) \geq \theta \sum_{i \in B(\theta')} [q(i) - d(\pi(i|\theta'))] - T(\theta'), \forall \theta, \theta' \in [0, 1] \quad (IC)$$

$$\theta \sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))] - T(\theta) \geq 0, \forall \theta \in [0, 1]. \quad (IR)$$

¹⁰Nonetheless, a more realistic model might impose some restrictions that prevent platforms from differentiating content access, such as contracts between the platform and content producers.

3.1 Discussion of Assumptions

Consumer Utility Specification. Our utility specification (2) follows the standard formulation in one-dimensional screening such as [Mussa and Rosen \(1978\)](#), in which a buyer’s gross utility is given by the product of her type and the quality of the good. In our model, the net quality of an item is given by its quality minus disutility from the associated ad. As a result, consumers with a higher θ have greater willingness to pay for item quality and stronger aversion to ads.

Type-dependent utility is crucial to our results. Our model predicts that the platform may screen consumers in content access and advertising (in addition to price), consistent with practices observed in some real-world ad-supported platforms. Such a screening pattern may not arise if some components of utility are type-independent.

To illustrate this point, consider alternative specifications. For example, if a consumer’s utility is given by $\theta \sum_{i \in B} q(i) - \sum_{i \in B} d(\pi(i)) - T$, so that advertising disutility is independent of θ , the optimal advertising policy depends only on $r(j) - d(j)$ and is thus independent of θ . Such a model is not only subsumed by our model (see [Section 8.2](#)), but also fails to rationalize a strategy whereby a platform allows consumers to pay for ad removal (e.g., YouTube and Netflix). Similarly, a model with type-independent utility from items, such as $\sum_{i \in B} q(i) - \theta \sum_{i \in B} d(\pi(i)) - T$, fails to explain the phenomenon that some platforms offer a cheaper plan with only partial access to available content (Peacock, Netflix, and The New York Times).

We could also consider a model in which a consumer has a multidimensional type and her utility is given by $\theta_1 \sum_{i \in B} q(i) - \theta_2 \sum_{i \in B} d(\pi(i)) - T$, where (θ_1, θ_2) is drawn from some joint distribution.¹¹ While this paper does not pursue such a multidimensional screening problem, it could be a fruitful avenue for future research.

¹¹This specification brings our model closer to that of [Yang \(2021\)](#). However, his Theorem 1 does not apply to our baseline model or the specification discussed here, for at least two reasons: (i) [Yang \(2021\)](#) assumes that the “productive component” of screening instruments is one-dimensional, but in our model, the space of productive allocations, which are advertising policies, is multidimensional; and (ii) his Theorem 1 requires θ_1 and θ_2 to be negatively correlated, whereas in our model, θ_1 and θ_2 are perfectly positively correlated (as $\theta_1 = \theta_2$).

Cost of Producing Content. In the model, the platform does not incur any marginal or fixed cost of providing consumers with access to items. The assumption of zero marginal cost is natural if (i) items are freely replicable digital goods and (ii) there is no expense tied to content dissemination on the platform’s side, such as payments to content producers. The assumption of zero fixed cost is relaxed in [Section 6](#).

4 Optimal Mechanism

We first define a few concepts that help us describe the platform’s optimal mechanism. We then characterize the optimal mechanism and present its key properties.

4.1 Virtual Advertising Profits

We begin by defining *virtual advertising profits*, which measure an ad’s contribution to the platform’s total revenue when we take into account the impact of advertising on the sales revenue in any incentive-compatible mechanism.

Formally, for any type $\theta \in [0, 1]$ and ad $j \in \mathcal{J}$, the θ -*virtual advertising profit* of ad j , denoted by $\rho_\theta(j)$, is defined as

$$\rho_\theta(j) := r(j) - v(\theta)d(j). \quad (4)$$

Unpacking the virtual value as $v(\theta) = \theta - \frac{1-F(\theta)}{f(\theta)}$, we can express $\rho_\theta(j)$ as

$$\rho_\theta(j) = \underbrace{r(j)}_{\text{ad revenue}} - \underbrace{\theta d(j)}_{\text{lower price for } \theta} + \underbrace{\frac{1-F(\theta)}{f(\theta)}d(j)}_{\text{reduction in information rents}}. \quad (5)$$

Virtual advertising profit $\rho_\theta(j)$ captures the net effect of allocating ad j to type θ on the platform’s total revenue. The first term $r(j)$ is the direct contribution of ad j to the platform revenue, i.e., advertising revenue that ad j generates for the platform.

The rest of the expression (5) captures the impact of advertising on the platform’s

sales revenue. Specifically, the second term, $-\theta d(j)$, captures the impact of allocating ad j on the sales revenue $T(\theta)$ from type θ , i.e., the nuisance from ad j reduces type θ 's willingness to pay for an item by $\theta d(j)$. The last term, $\frac{1-F(\theta)}{f(\theta)}d(j)$, captures the impact of advertising on information rents for types higher than θ . Displaying ad j to type θ reduces the incentives of higher-type consumers to misreport and take the contract intended for type θ . As a result, the platform can increase prices for higher types while keeping the mechanism incentive compatible.

Virtual advertising profits illustrate the idea that the true contribution of an advertisement to the platform's revenue typically differs from either (i) its contribution to social welfare (i.e., ad revenue minus nuisance, $r(j) - \theta d(j)$) or (ii) its ad revenue (i.e., $r(j)$). Below we provide an intuition for each comparison.

First, an ad's contribution $\rho_\theta(j)$ to total revenue exceeds the contribution to social welfare (i.e., $\rho_\theta(j) > r(j) - \theta d(j)$ for every $\theta < 1$), because nuisance enables the platform to reduce information rents, which is captured by the last term in (5). Thus, for any fixed set of items, the platform has a socially excessive incentive to display ads (except for the highest type). For example, suppose that ad j satisfies $r(j) = \hat{\theta}d(j)$ for some $\hat{\theta} \in (0, 1)$. Then, for any type $\theta > \hat{\theta}$, the ad revenue $r(j)$ falls below the disutility $\theta d(j)$; thus, in the first best, the platform never allocates ad j to type θ . However, in the second best, the platform may strictly prefer allocating ad j to allocating no ad in order to reduce information rents. In other words, the same trade-off between efficiency and rent extraction, which generates downward distortions in “goods,” may generate upward distortions in “bads.”

Second, the virtual ad profit $\rho_\theta(j)$ is larger (smaller) than the ad revenue $r(j)$ for consumers with negative (positive) virtual types. For example, for negative virtual types, advertising is more profitable than suggested by their ad revenue—that is, $\rho_\theta(j) > r(j)$ when $v(\theta) < 0$ —because the reduction in information rents from ad nuisance is larger than the nuisance itself. In fact, the virtual ad profit increases (decreases) with the disutility level for negative (positive) virtual types. Therefore, if

all ads generate the same revenue, the platform prefers showing the ads with highest (lowest) disutility levels for negative (positive) virtual types. Generally, the platform’s ranking of ads by virtual ad profits depends on both ad characteristics, $(r(j), d(j))$, and a consumer’s type, θ .

4.2 Assortative Advertising Policies

The optimal mechanism matches items and ads in a negatively assortative way, based on item quality and virtual advertising profits. To formalize this notion, we first define a class of advertising policies. Take any map $\rho : \mathcal{J} \rightarrow \mathbb{R}$, which assigns each ad $j \in \mathcal{J}$ a score, $\rho(j)$. Call a permutation $\mu : \mathcal{J} \rightarrow \mathcal{J}$ a ρ -permutation if μ labels ads in such a way that an ad with a larger index has a lower score—i.e.,

$$\rho(\mu(1)) \geq \rho(\mu(2)) \geq \dots \geq \rho(\mu(J)).$$

If multiple ads have the same score under $\rho(\cdot)$, the ρ -permutation may not be unique.

Definition 1. An advertising policy is a ρ -negative assortative policy if the following holds for some ρ -permutation, μ : Item k , which has the k -th lowest quality, is matched with ad $\mu(k)$ if $\rho(\mu(k)) \geq 0$ and not matched with any ad if $\rho(\mu(k)) < 0$.

As shown below, the optimal mechanism uses a ρ_θ -negative assortative policy for every type θ , where ρ_θ is the virtual ad profit function for type θ (see (4)). According to [Definition 1](#), we can construct the optimal advertising policy for each type θ as follows. First, we assign each ad j its θ -virtual advertising profit, $\rho_\theta(j)$. Then, starting from the lowest-quality item, we match items and ads in a negatively assortative manner, according to the quality $q(i)$ of each item and $\rho_\theta(j)$ of each ad. Once we exhaust ads that have nonnegative $\rho_\theta(j)$, we stop this process and leave the remaining higher-quality items unmatched with any ads.

4.3 Characterization of the Optimal Mechanism

We now solve the platform's problem and describe the optimal mechanism. First, we rewrite the platform's problem (M) as maximization of virtual surplus subject to the monotonicity constraint, i.e., $\sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))]$ is weakly increasing in θ (e.g., Myerson 1981). We then define the relaxed problem as the one in which we maximize the virtual surplus without the monotonicity constraint:

$$\max_{\{(B(\theta), \pi(\cdot|\theta))\}_{\theta \in [0,1]}} \int_0^1 v(\theta) \sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))] + \sum_{i \in B(\theta)} r(\pi(i|\theta)) dF(\theta). \quad (\text{P})$$

Solving Problem (P) and applying the standard technique to recover monetary transfers from local IC constraints, we obtain the optimal mechanism.

Lemma 1. *The platform's problem (M) has the following solution: For each type $\theta \in [0, 1]$, the platform adopts a ρ_θ -negative assortative matching, $\pi^*(\cdot|\theta)$, and allocates all items that generate non-negative virtual surplus, i.e., $B(\theta) := \{i \in \mathcal{I} : v(\theta)q(i) + \rho_\theta(\pi^*(i|\theta)) \geq 0\}$. The monetary transfer from type θ equals*

$$T(\theta) = \theta \sum_{i \in B(\theta)} [q(i) - d(\pi^*(i|\theta))] - \int_0^\theta \sum_{i \in B(t)} [q(i) - d(\pi^*(i|t))] dt.$$

We sketch the proof of Lemma 1 by describing how to solve the relaxed problem, Problem (P) (see Appendix A for omitted details). First, the virtual surplus is separable across types, so we can reduce Problem (P) to pointwise maximization of each type's virtual surplus, $v(\theta) \sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))] + \sum_{i \in B(\theta)} r(\pi(i|\theta))$.

Second, we calculate the contribution of each item-ad pair to type θ 's virtual surplus. Suppose that the platform matches item $i \in \mathcal{I}$ with ad $j \in \mathcal{J} \cup \{\emptyset\}$. If the platform allocates an item-ad pair (i, j) , the virtual surplus from type θ increases by

$$v(\theta)[q(i) - d(j)] + r(j), \quad (6)$$

where $d(\emptyset) = r(\emptyset) = 0$.

The platform allocates an item-ad pair (i, j) if and only if its contribution to the virtual surplus is nonnegative. As a result, the contribution of the pair (i, j) to the virtual surplus, given the platform's optimal decision for whether to allocate the pair, is equal to $\max(0, v(\theta)[q(i) - d(j)] + r(j))$, or equivalently,

$$\max(0, v(\theta)q(i) + \rho_\theta(j)). \quad (7)$$

with $\rho_\theta(\emptyset) = 0$. Thus, the platform's problem of maximizing type θ 's virtual surplus is written as

$$\max_{\pi \in \mathcal{M}} \sum_{i \in \mathcal{I}} \max(0, v(\theta)q(i) + \rho_\theta(\pi(i))). \quad (\mathbf{R}-\theta)$$

Here, the outside max operator is the problem of choosing an advertising policy for type θ , and the inside max operator reflects the problem of whether to allocate each item-ad pair to type θ .

The problem $(\mathbf{R}-\theta)$ admits the following solution. First, consider the outside max operator for a negative virtual type, $v(\theta) < 0$. The platform never displays ad j when $\rho_\theta(j) < 0$. Also, the virtual surplus (7) from an item-ad pair (i, j) is submodular in $(q(i), \rho_\theta(j))$. Consequently, the platform matches items in \mathcal{I} and ads in $\{j \in \mathcal{J} : \rho_\theta(j) > 0\}$ in a negatively assortative way, which reduces to a ρ_θ -negative assortative matching. The case of positive virtual types follows the same logic, except that the submodularity of the objective in $(\mathbf{R}-\theta)$ is trivial: If $v(\theta) > 0$, we have $v(\theta)q(i) + \rho_\theta(j) = v(\theta)[q(i) - d(j)] + r(j) \geq 0$. Thus, the ρ_θ -negative assortative matching is merely one of multiple solutions. Finally, the inside max operator of $(\mathbf{R}-\theta)$ determines whether to allocate each matched item-ad pair to type θ , the solution to which reduces to the item bundle $B(\theta)$ described in [Lemma 1](#).

We build on [Lemma 1](#) and establish key properties of the optimal mechanism. To streamline exposition, assume that no two items have the same quality level.

Theorem 1. *The optimal mechanism has the following properties:*

1. *Consumers with negative virtual types (i.e., $\theta < \theta^0$) receive all items whose quality levels are at or below a type-dependent threshold, $\bar{q}(\theta)$, where $\bar{q}(\theta)$ is weakly increasing in θ . For each type, the platform matches every allocated item with some ad or, otherwise, displays all available ads.*
2. *Consumers with positive virtual types (i.e., $\theta > \theta^0$) receive all items, and higher types see fewer items matched with ads.*

In standard screening problems, the seller excludes negative virtual values. In contrast, the platform in our model may serve consumers with negative virtual values to earn ad revenue.¹² However, whenever the platform tries to do so, it faces a trade-off between rent extraction and advertising: Allocating items bundled with ads increases ad revenue but increases information rents and decreases sales revenue.

Part 1 of [Theorem 1](#) illustrates how this trade-off shapes the allocation of content and ads for negative virtual types. First, the platform allocates items whose quality levels are below some threshold in order to earn advertising revenue while reducing information rents. The type-dependent threshold $\bar{q}(\theta)$ is increasing in θ ; thus, higher types gain access to higher-quality items (in addition to lower-quality items) and correspondingly pay higher prices.¹³ In practice, such a contract might arise as a plan that has a low price but excludes access to some premium content. Second, for negative virtual types, the platform matches every item with some ad, because without ads the platform will exclude these types from accessing any content.

¹²To be precise, Part 1 of [Theorem 1](#) includes two extreme possibilities: One is the case in which $\bar{q}(\theta)$ is so low for every $\theta < \theta^0$ that all consumers with negative virtual types are excluded; this is possible, for example, if all ads generate negligible revenue and disutility. The other extreme is when $\bar{q}(\theta)$ is so high that consumers receive all items regardless of their type. This case could arise if all ads generate sufficiently high ad revenues.

¹³The result that higher types gain access to a larger set of items is not a direct consequence of incentive compatibility. IC constraints imply the monotonicity of allocation, which only requires $\sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))]$ to be increasing.

Part 2 states that the platform allocates all items to consumers with positive virtual types. Indeed, the platform in our model does not incur marginal costs of production, and thus the expansion of total surplus due to item allocation exceeds the associated increase in information rents. Also, within positive virtual types, higher types are exposed to fewer ads, because they incur greater disutility from ads.

Our results also predict how the platform distorts advertising allocation relative to the first best. For consumers with positive virtual types, Part 2 of [Theorem 1](#) shows that they receive all items and higher types see fewer ads. Because the virtual ad profit of each ad is greater than its contribution to total surplus (i.e., $\rho_\theta(j) > r(j) - \theta d(j)$), the platform shows more ads than the efficient level to each type, except for the highest type, which receives the first-best set of ads.

For consumers with negative virtual types, Part 1 shows that they receive too few items. This implies that, contrary to the case of positive virtual types, the platform may serve fewer ads than the efficient level, because ads are bundled with items. Moreover, the platform is biased toward high-annoyance ads. To see this, suppose that all ads generate the same revenue. Then, given the number of ads to be displayed, the platform prefers displaying ads with highest annoyance, even though total surplus is maximized by allocating ads with lowest annoyance. This is because for negative virtual types, virtual ad profits increase with disutility levels. Such a distortion is absent for positive virtual types, for which virtual ad profits are decreasing in disutility levels.

The following example illustrates the typical structure and the distortion in the allocation of items and ads induced by the platform's optimal mechanism.

Example 1. Suppose that consumer types are uniformly distributed on $[0, 1]$, so the virtual type is given by $v(\theta) = 2\theta - 1$; the platform has the same number of items and ads, i.e., $I = J \geq 2$; each item $i = 1, \dots, I$ has quality $q(i) = 1 + \frac{i-1}{I-1}$; and each ad $j = 1, \dots, I$ has revenue $r = 0.1$ and disutility level $d(j) = \frac{j-1}{I-1}$. The example is a discrete version of a situation in which item quality and ad disutility are uniformly distributed on $[1, 2]$ and $[0, 1]$, respectively. Hereafter, we identify each item and ad by

its quality level and disutility level, respectively.

We use [Lemma 1](#) to derive the optimal allocation policy. For consumers with negative virtual types, the platform matches every item q with ad $d = 2 - q$, and allocates an item-ad pair (q, d) if and only if $v(\theta)(q - d) + r \geq 0$, i.e., $q \leq \bar{q}^P(\theta) := 1 + \frac{1}{20(1-2\theta)}$. Correspondingly, the platform allocates to each type θ all ads whose disutility levels exceed $\underline{d}^P(\theta) = 2 - \bar{q}^P(\theta)$. For consumers with positive virtual types, the platform allocates all items and displays only ads that yield a positive virtual advertising profit, i.e., $r - dv(\theta) \geq 0$, or equivalently $d \leq \bar{d}^P(\theta) := \frac{1}{10(2\theta-1)}$.

[Figure 1](#) depicts the resulting policy. The left panel shows that in the optimal mechanism, each consumer receives all items whose quality levels lie in $[1, \bar{q}^P(\theta)]$. Because the first best is to allocate all items to all types, the second-best allocation to negative virtual types reflects the standard underprovision of goods. However, unlike standard screening models, the platform allocates a nonempty set of items to every type.

The right panel presents the set of ads allocated to each type θ under the first best and the second best. Under the first best (i.e., the orange dashed area), the platform allocates all ads whose disutility levels are below $\bar{d}^{FB}(\theta) = \frac{r}{\theta}$. Hence, the amount of ads decreases with θ . As all ads generate the same revenue in this example, the first-best policy allocates less annoying ads to minimize welfare loss.

The right panel also presents the second-best advertising policy, where each type θ receives all ads whose disutility levels lie in $[\underline{d}^P(\theta), \bar{d}^P(\theta)]$ (i.e., the light blue area). For relatively low types, the platform distorts the advertising policy in terms of both the disutility levels and volume of advertising. Given a volume of advertising, the platform allocates ads with the highest disutility levels, even though total surplus would be maximized if it allocated ads with the lowest disutility levels. At the same time, these consumers are exposed to fewer ads (i.e., $\bar{d}^{FB}(\theta) < 1 - \underline{d}^P(\theta)$), because they view fewer items. For relatively high types (which include all positive-virtual types), the distortion is in advertising volume: The platform prioritizes allocating ads with low disutility levels, but it now allocates more ads than under the first best. Overall, the

example illustrates that our model can capture rich ways in which the platform distorts its advertising policy away from the first best.




Figure 1: The left panel depicts the set of allocated items in terms of their quality for each type under the second best. The right panel depicts the set of allocated ads in terms of their disutility levels for each type under the first best (dashed orange area) and the second best (light blue area).

5 Application: Rationalizing Contracts in Practice

We now leverage our results to rationalize real-world contracts used by ad-supported platforms. Specifically, we provide a sufficient condition under which the optimal menu of contracts discriminates among consumers in terms of (i) content access only, (ii) advertising exposure only, or (iii) both content access and advertising exposure.¹⁴ We then argue that each sufficient condition captures typical platforms that offer the corresponding type of contracts in practice.

Toward this goal, we impose two assumptions. First, assume that there is a constant $c > 0$ such that $d(j) = cr(j)$ for every $j \in \mathcal{J}$. The parameter c is called the disutility coefficient and represents the disutility level per unit of ad revenue. A higher c means that ads are overall more annoying relative to the revenue generated for the platform.

¹⁴To be precise, the platform also price discriminates consumers in all cases. Thus, for example, “discrimination in content access only” should be read as “discrimination in content access and price.”

The linearity assumption implies that ads that generate more revenue are also more annoying. For example, a platform may earn a high revenue from (unmodeled) advertisers by displaying long, loud, or unskippable ads, because they are effective at grabbing users' attention; however, such ads impose high nuisance on consumers.

Under the linearity assumption, the θ -virtual advertising profit (4) of each ad j is written as $\rho_\theta(j) = r(j)(1 - cv(\theta))$. As a result, there is some cutoff $\theta^* > \theta^0$ that solves $1 - cv(\theta^*) = 0$, such that for any ad j , we have $\rho_\theta(j) > 0$ if $\theta < \theta^*$ and $\rho_\theta(j) < 0$ if $\theta > \theta^*$ (recall θ^0 is the unique type with zero virtual value). Thus, the sign of the virtual advertising profit for ad j , which determines whether the platform prefers displaying ad j to displaying no ad, depends on consumer type θ but not j .

The other assumption is that the number of ads is weakly greater than the number of items, i.e., $J \geq I$. Let $r(I)$ denote the I -th highest ad revenue in \mathcal{J} .

We partition the space of parameters into three regions depending on the value of c and $r(I)$, and describe the optimal mechanism in each region. Concretely, Case 1 covers parameters such that $c < 1$; when $c > 1$, Case 2 or Case 3 applies depending on whether $r(I)$ is larger or smaller than a threshold, respectively.

Case 1 (Differentiated Content Access). Suppose that $c < 1$, i.e., the disutility coefficient is sufficiently low. In this case, every consumer will see some ad for every allocated item. This is because when $c < 1$, the virtual advertising profit $\rho_\theta(j) = r(j)(1 - cv(\theta))$ is positive for all ads and consumer types.

In this case, the optimal menu always contains a contract that allocates all items, each of which is bundled with some ad. However, whether, and if so what kind of, other contracts are offered depends on further details of parameters. If the platform offers more than one contract, the menu must also include a contract designed for sufficiently low types (including $\theta = 0$) and priced at zero. This contract excludes either all items or some items whose quality levels exceed a certain threshold. The platform may further discriminate among negative virtual types by assigning them different quality

thresholds (see [Theorem 1](#)).

Consumers with higher types are mechanically exposed to more ads, because their contract grants access to a greater number of items, which equals the number of ads. In other words, the variation in ad exposure arises solely from the variation in item allocation across types. Although higher types suffer more from advertising disutility, no one has the option to pay to remove ads.

This case roughly fits media platforms, such as the New York Times, the Wall Street Journal, and the Financial Times. They offer multiple subscription plans that differ in the number or the set of articles subscribers can read, where a higher-priced plan allows access to more articles. However, none of them offer an ad-free plan.¹⁵

Whether the condition $c < 1$ fits media platforms is an empirical question. However, we believe that—relative to social media and streaming services discussed below—people typically find ads displayed alongside news articles less annoying, because they can easily scroll down to continue reading the article.¹⁶

Case 2 (Differentiated Ad Exposure). We now consider the case in which ads are annoying but generate high revenue, i.e.,

$$c > 1 \quad \text{and} \quad r(I) > \frac{-v(0)q(I)}{1 - cv(0)}, \quad (8)$$

where $v(0) < 0$ is the virtual type of the lowest type ($\theta = 0$) and $q(I)$ is the highest item quality. The second inequality is equivalently written as $v(0)q(I) + r(I)[1 - cv(0)] > 0$.

The inequality implies that the platform allocates the highest-quality item to the lowest type, $\theta = 0$, which, in turn, implies that the platform allocates all items to

¹⁵See, e.g., <https://help.nytimes.com/hc/en-us/articles/360001712553-Ads-on-The-New-York-Times> (accessed on November 10, 2025).

¹⁶For printed newspapers, empirical evidence suggests that advertising disutility is low, if not negative; see, e.g., Footnote 1 in [Filistrucchi, Klein, and Michielsen 2012](#).

all consumers for free.¹⁷ At the same time, because $c > 1$, the virtual ad profit $r(j)(1 - cv(\theta))$ is negative for types above $\theta^* > \theta^0$. Thus, any type above θ^* pays a positive price to remove ads.

This case captures digital platforms such as YouTube, Instagram, Facebook, and X. These platforms allow users to access all content for free, but also offer paid plans that remove ads (Instagram and Facebook offer such a plan in the EU).

As in **Case 1**, whether the parameter restriction captures these platforms is an empirical question. However, a combination of high disutility and highly profitable ads seems to fit these large digital platforms, especially when we view $r(j)$ as the value of attention to a platform based on not only advertising but also on data collection, and $d(j)$ as a consumer’s disutility from targeted advertising and privacy concerns.

Case 3 (Differentiation in Both Dimensions). The remaining case is when advertising disutility is not too low and not all ads are sufficiently profitable, i.e.,

$$c > 1 \quad \text{and} \quad r(I) < \frac{-v(0)q(I)}{1 - cv(0)}. \quad (9)$$

Then, the platform offers contracts that differ both in content access and ad exposure (see [Figure 2](#); [Appendix B](#) provides details). First, consumers with a sufficiently high type choose a plan that gives them access to all items without ads. Second, consumers with negative virtual types have limited access to items or are excluded, and every allocated item comes with some ad. Thus, the platform discriminates consumers in terms of both content access and ad exposure.

The optimal mechanism in this case fits streaming platforms such as Netflix, Peacock, Hulu, Spotify, Pandora, and Amazon Music. Their lower-priced plan—which may be free or have a positive price—limits access to content and carries ads, whereas their higher-priced plan features unlimited access to content and no ads.

¹⁷Recall that type $\theta = 0$ receives the set of items whose quality levels are below some threshold ([Theorem 1](#)). Thus, if type $\theta = 0$ receives the highest-quality item (i.e., item I), it also receives all items. As higher types obtain a larger set of items, all types must be receiving all items.

In our model, the platform also offers a plan that gives full access to content but still displays ads for all items. Such a plan is offered to consumers whose types are between θ^0 and θ^* , i.e., their types are not too high but still have positive virtual types. Such a plan is used by streaming services such as Peacock and Amazon Prime Video.

Additionally, similar to **Case 1**, consumers with negative virtual types may select different contracts. They all receive items whose quality levels are below a threshold, and higher types choose contracts that correspond to higher quality thresholds.




Figure 2: Optimal mechanism when $c > 1$ and $r(I) < \frac{-v(0)q(I)}{1-cv(0)}$.

As in the previous two cases, we argue that condition (9) fits streaming platforms better than news organizations or social media platforms. The condition $c > 1$ is more suitable for streaming platforms than news organizations, for reasons discussed in **Case 1**. Moreover, streaming platforms—such as Netflix, Peacock, Hulu, and Spotify—are not (in)famous for intensive data collection and hyper-targeting, at least not to the same extent as Google or Meta. If we think of data collection and targeting as factors that make advertising more profitable, the condition $r(I) < \frac{-v(0)q(I)}{1-cv(0)}$ may be more appropriate for streaming platforms. Our model then predicts the pattern of discrimination adopted by these streaming services.

Our model is not the only way to rationalize different margins of discrimination discussed above. For example, **Case 2**—in which all consumers can access all content—may stem from network externalities between consumers. The point of our exercise is to show that one way to rationalize different contracts is to view them as the optimal resolution of the trade-off between advertising and rent extraction.

6 Platform's Innovation Incentives

A key feature of our model is that the platform may serve consumers with negative virtual types. This feature has an implication on the platform's incentive to improve content quality. To formally investigate this, we extend the model by incorporating the platform's choice of quality profile.

Throughout this section, we fix the number of items, I . We also take as given an arbitrary set of ads, \mathcal{J}^* , along with their characteristics, $(r(j), d(j))_{j \in \mathcal{J}^*}$. The set of ads hosted by the platform, which we call the *advertising set*, is a subset of \mathcal{J}^* and denoted by \mathcal{J} . Given I and \mathcal{J} , the platform chooses the quality of each item, then adopts the optimal mechanism characterized in [Lemma 1](#). Our goal is to examine how the platform's optimal quality choice depends on the advertising set.

We introduce several pieces of notation and define the platform's problem. First, let \mathcal{Q} denote the set of feasible quality profiles:

$$\mathcal{Q} := \{(q(1), \dots, q(I)) \in \mathbb{R}_+^I : q(I) \geq \dots \geq q(1) \geq \max_{j \in \mathcal{J}^*} d(j)\}.$$

The inequality $q(1) \geq \max_{j \in \mathcal{J}^*} d(j)$ implies that for any feasible quality profile, the net quality of any item-ad pair is nonnegative.

For any quality profile $\mathbf{q} \in \mathcal{Q}$ and advertising set $\mathcal{J} \subseteq \mathcal{J}^*$, let $\Pi(\mathbf{q}, \mathcal{J})$ denote the platform's revenue from the optimal mechanism. Let $C(\mathbf{q}) \in \mathbb{R}$ denote the cost of choosing a quality profile \mathbf{q} . Assume that $C(\cdot)$ is submodular, which holds, e.g., if $C(\cdot)$ is additively separable, i.e., $C(\mathbf{q}) = \sum_{i=1}^I C_i(q(i))$. Finally, for any $x, y \in \mathbb{R}_+^K$, we write $x \geq y$ to mean $x_i \geq y_i$ for all $i = 1, \dots, K$.

Our specification implies that the cost of increasing the quality of a given item does not depend on the number of consumers who receive the item. This assumption is natural for digital goods, which feature free replicability.

Define $\mathcal{Q}^*(\mathcal{J})$ as the set of solutions to the platform's quality choice problem:

$$\mathcal{Q}^*(\mathcal{J}) := \arg \max_{\mathbf{q} \in \mathcal{Q}} \Pi(\mathbf{q}, \mathcal{J}) - C(\mathbf{q}). \quad (\mathcal{Q})$$

Assume that the problem has a solution, i.e., $\mathcal{Q}^*(\mathcal{J}) \neq \emptyset$ for any $\mathcal{J} \subseteq \mathcal{J}^*$.

The following lemma describes how a change in the advertising set affects the platform's quality choice through a shift in virtual advertising profits. We say that advertising set \mathcal{J}_H *attains uniformly higher virtual ad profits than* \mathcal{J}_L if for every type θ and $i = 1, \dots, I$, $\max(0, \rho_\theta(j_{\theta,i}^H)) \geq \max(0, \rho_\theta(j_{\theta,i}^L))$, where for each $x \in \{L, H\}$, $j_{\theta,i}^x$ is the ad that generates the i -th highest θ -virtual ad profit in \mathcal{J}_x if it exists; otherwise (i.e., if $i > |\mathcal{J}_x|$), $j_{\theta,i}^x = \emptyset$.

Lemma 2. *If advertising set \mathcal{J}_H attains uniformly higher virtual ad profits than \mathcal{J}_L , the platform has a lower incentive to improve content quality under \mathcal{J}_H than \mathcal{J}_L , i.e., for any $\bar{\mathbf{q}}, \underline{\mathbf{q}} \in \mathcal{Q}$ such that $\bar{\mathbf{q}} \geq \underline{\mathbf{q}}$, we have $\Pi(\bar{\mathbf{q}}, \mathcal{J}_H) - \Pi(\underline{\mathbf{q}}, \mathcal{J}_H) \leq \Pi(\bar{\mathbf{q}}, \mathcal{J}_L) - \Pi(\underline{\mathbf{q}}, \mathcal{J}_L)$. Consequently, $\mathcal{Q}^*(\mathcal{J}_H)$ is smaller than $\mathcal{Q}^*(\mathcal{J}_L)$ in the strong set order.*

The intuition is as follows. According to [Lemma 1](#), the platform adopts a negative assortative advertising policy, i.e., it matches the item with the i -th lowest quality with the ad with the i -th highest virtual ad profit, provided that the latter is nonnegative. Thus, if \mathcal{J}_H attains uniformly higher virtual ad profits than \mathcal{J}_L , the platform facing \mathcal{J}_H finds it more profitable to allocate each matched item-ad pair to consumers. As a result, the platform allocates more items to consumers with negative virtual types. However, the expanded allocation toward negative virtual types makes it more costly for the platform to improve item quality, because the loss due to information rents is proportional to item quality. Therefore, the uniform increase in virtual ad profits discourages the platform from investing in content quality.

[Figure 3](#) illustrates [Lemma 2](#) for a platform that hosts one item and one ad. The ad imposes zero disutility, and thus the virtual ad profit is equal to its ad revenue, r . The platform's cost of investment is given by $0.1q^2$, and consumer types are uniformly

distributed on $[0, 1]$. As ad revenue r increases, the platform invests less in item quality (the left panel) and allocates the item to more consumers (the right panel). Once the ad revenue exceeds a threshold (slightly below 0.5), the platform sets the item quality to 0 and allocates the item to all consumers for free.




Figure 3: The left panel depicts the optimal quality as a function of $r \in [0, 1]$ when $C(q) = 0.1q^2$ and $\theta \sim U[0, 1]$. At a unique cutoff $r \approx 0.47$, where the optimal quality drops to 0, there are two optimal quality levels. The right panel depicts the set of types that receive the item at each r , where the blue line captures the lowest type that receives the item.

[Lemma 2](#) implies that the platform invests less in content quality if it hosts a larger set of ads or faces a higher ad revenue for each ad.

Proposition 1. *Take two advertising sets, \mathcal{J}_H and \mathcal{J}_L . Suppose we have (i) $\mathcal{J}_H \supset \mathcal{J}_L$ or (ii) we obtain \mathcal{J}_H by increasing the ad revenue $r(j)$ of each ad $j \in \mathcal{J}_L$ while keeping disutility levels the same. Then, \mathcal{J}_H attains uniformly higher virtual ad profits than \mathcal{J}_L , and thus the platform chooses lower quality profiles under \mathcal{J}_H than \mathcal{J}_L in the strong set order.*

Proof. If the advertising set expands, the i -th highest virtual ad profit increases for every i . Increasing the ad revenue for each ad has the same effect because it increases virtual ad profits (4) across all consumer types. Thus, either change leads to uniformly higher virtual ad profits. By [Lemma 2](#), the platform chooses lower quality profiles under \mathcal{J}_H than under \mathcal{J}_L . \square

As we discuss in the next section, [Proposition 1](#) allows us to examine how policy or technological changes surrounding advertising may affect the platform’s quality choice. In addition, [Proposition 1](#) allows us to examine the relation between the platform’s business model and its quality choice. To formalize this idea, we categorize mechanisms into three business models: The platform adopts a *subscription* model if it displays no ads (which occurs if and only if $\mathcal{J} = \emptyset$); a *purely ad-funded* model if the optimal mechanism sets a price of zero for all types; and a *hybrid* model if the platform adopts neither a subscription nor purely ad-funded model.

Both the platform’s business model and quality choice are endogenous and depend on the advertising set. Thus, by varying the advertising set, we obtain different pairs of optimal business models and quality profiles. The following result illustrates the resulting relationship between the platform’s business model and quality choice.

Proposition 2. *A platform with a subscription model chooses a higher quality profile than a platform with a hybrid model (in the strong set order), which, in turn, chooses a higher quality profile than a purely ad-funded platform.*

Intuitively, a platform whose business model relies more heavily on advertising also allocates more items to negative virtual types. Such an allocation rule discourages the platform from investing in content quality. We also note that a purely ad-funded model could be interpreted as a platform that cannot use monetary transfers for an exogenous reason, rather than one that optimally sets a zero price for all types. With this interpretation, [Proposition 2](#) holds verbatim.

7 Policy Implications

[Propositions 1](#) and [2](#) provide two sets of policy implications. First, the results clarify the possible impacts of policy and technological changes on platforms’ incentives to invest in content quality. We illustrate this first implication with two examples.

The first example concerns a recent bill introduced in California that aims to apply the Commercial Advertisement Loudness Mitigation (CALM) Act to streaming services.¹⁸ The CALM Act is a federal law that requires television commercials to be no louder than the average volume of the program. The proposed bill seeks to apply the same restriction to platforms such as Netflix and Hulu.

To examine the potential impact of the CALM Act in our model, suppose the platform faces a baseline set of ads, \mathcal{J} . For each ad $j \in \mathcal{J}$, $(r(j), d(j))$ represents the ad revenue and disutility level when the platform plays the ad at the same loudness as the main content. The platform can also play each ad more loudly. The louder version of ad j is defined as ad \hat{j} such that $r(\hat{j}) > r(j)$ and $d(\hat{j}) > d(j)$, i.e., louder ads are more annoying but also more effective at grabbing users' attention. Thus, without any regulation, the advertising set is given by $\hat{\mathcal{J}} = \mathcal{J} \cup \{\hat{j}\}_{j \in \mathcal{J}}$.

The CALM Act shrinks the advertising set from $\hat{\mathcal{J}}$ to \mathcal{J} . As formalized by [Proposition 1](#), such a regulation limits the platform's ability to degrade the quality of offerings to negative virtual types, which, in turn, incentivizes the platform to invest more in content quality and allocate it to consumers who have higher valuations.

The second example concerns improvements in advertising technology. While we do not explicitly model the relevance of each ad to consumers, we may capture better targeting—such as the platform's ability to display relevant ads when users are likely to respond—as an increase in the ad revenue associated with each $j \in \mathcal{J}$. [Proposition 1](#) then implies that such improvements, which expand the allocation of items to negative virtual types, may reduce the platform's incentives to invest in content quality. An increase in each $r(j)$ might also arise if the platform can extract more surplus from (unmodeled) advertisers, either through greater market power or through improved

¹⁸For details of the bill, see https://leginfo.legislature.ca.gov/faces/billNavClient.xhtml?bill_id=202520260SB576 (accessed on November 10, 2025).

information about advertisers’ willingness to pay.¹⁹ In summary, while the policy and technological changes described above may seem unrelated, our model provides a unified way to understand their potential impacts as a shift of virtual ad profits.

The second policy implication is that existing concerns about the negative impact of the ad-supported business model on quality and innovation ([Competition & Markets Authority, 2020](#); [Stigler Committee on Digital Platforms, 2019](#)) apply not only to purely ad-funded platforms but also to hybrid platforms, which combine both advertising and direct pricing. We provide a general rationale for why this concern is relevant for hybrid platforms from the mechanism-design perspective: Reliance on advertising makes it costly for the platform to improve content quality because of the associated increase in information rents.

The underinvestment result for hybrid platforms in [Proposition 2](#) is particularly relevant to media platforms. [Latham et al. \(2022\)](#) note that “an ad-funded business model (which) might have knock-on effects for media plurality by, for example, reducing incentives for investment in content generation.” Moreover, if artificial intelligence disproportionately lowers the cost of producing low-quality news articles relative to high-quality ones, this strengthens the negative impact of increased ad profitability on the quality of journalism.²⁰

Finally, while [Propositions 1](#) and [2](#) identify changes in the advertising environment or business model that lead to lower content quality, the resulting impact on consumers is ambiguous. Suppose, for example, that advertising becomes more profitable as in [Proposition 1](#). On the one hand, the platform reduces its quality profile, which lowers

¹⁹In an earlier version of this paper, we explicitly modeled advertisers as strategic players privately informed of their valuation of impressions. We showed that when the platform learns advertisers’ types, which eliminates their private information, this can be captured as an increase in each $r(j)$. See [Ichihashi, Jeon, and Kim \(2024\)](#).

²⁰For instance, AI can generate content automatically by converting data into informative and narrative texts, leading to the production of thousands of stories with little or no human intervention ([Noain-Sánchez, 2022](#)).

every consumer’s utility for a fixed allocation rule.²¹ On the other hand, a platform that hosts many profitable ads allocates more items to low-valuation consumers, which raises their utilities—and can also increase the utilities of higher-valuation consumers. Thus, although greater reliance on advertising may shift the platform’s focus from quality toward the quantity of allocation, the resulting consumer impacts are ambiguous.

8 Extensions

We discuss a few extensions that relax some of the key assumptions.

8.1 Attention Cost and Free Disposal

In our model, any item-ad pair yields nonnegative utility. In [Appendix D](#), we assume that a consumer’s utility from an item-ad pair is given by $\theta[q(i) - d(j)] - a$, where $a > 0$ is an attention cost per unit of item. The attention cost limits the space of implementable allocation rules, because consumers may now ignore item-ad pairs whose net quality is lower than a . Nonetheless, under certain conditions, the key structure of the optimal mechanism remains the same. The main technical challenge is that IC constraints involve double deviations, whereby consumers misreport their type and then consume a strict subset of allocated items. To solve the problem, we adopt an approach inspired by [Corrao et al. \(2023\)](#), where we first solve the problem without double deviations and then verify that the solution to this problem satisfies all IC constraints.

8.2 Type-independent Advertising Disutility

In our model, advertising disutilities are type dependent. However, our model subsumes a general setup in which ads impose both type-dependent and type-independent

²¹In any incentive-compatible mechanism such that the IR constraint for the lowest type binds, the interim payoff of type θ is given by $\int_0^\theta \sum_{i \in B(t)} [q(i) - d(\pi(i|t))] dt$. For a fixed allocation rule, this expression decreases as the quality of each item decreases.

nuisance. Each ad j is characterized by $(r(j), d_1(j), d_2(j)) \in \mathbb{R}_+^3$, and a consumer’s utility from an item-ad pair (i, j) is $\theta[q(i) - d_1(j)] - d_2(j)$. The new component is $d_2(j)$, which captures a type-independent but ad-dependent disutility from ad j . This model is equivalent to ours, because we can redefine advertising revenue as $r(j) - d_2(j)$ and then normalize each $d_2(j)$ to 0. This modification does not affect the optimal mechanism, consumers’ interim payoffs, or the platform’s revenue.²²

8.3 Many-to-One Matching

We have assumed one-to-one matching between items and ads. However, our results hold even when each item can be matched with at most $K \in \mathbb{N}$ ads. In such a case, the platform matches, for each θ , the lowest-quality item with the K ads that have the highest virtual ad profits, provided they are nonnegative. The platform then matches the second-lowest quality item with the next K highest nonnegative virtual ad profits, and so on. In this model, [Theorem 1](#) and the results in [Section 6](#) hold verbatim.

9 Conclusion

Digital platforms, such as social media, streaming services, and news organizations, all serve content and ads, but they adopt widely different business models. Meanwhile, ad-supported platforms have become a focus of policy debates, amid concerns that ad-funded business models may weaken platforms’ incentives to invest in quality ([Scott Morton et al., 2020](#); [Competition & Markets Authority, 2020](#); [Crémer et al., 2019](#); [Stigler Committee on Digital Platforms, 2019](#)).

In this paper, we study mechanism design for ad-supported platforms. Our results

²²The modification might affect the analysis if the platform is constrained to using only nonnegative monetary transfers. Indeed, the nonnegative price constraint binds whenever the platform tries to induce the participation of the lowest type $\theta = 0$ and exposes type 0 to ads with positive $d_2(j) > 0$. In contrast, the constraint never binds if $d_2(j) = 0$. However, it seems natural to assume a constant type-independent consumption utility if we introduce a constant type-independent disutility. Then, the constraint may not bind.

rationalize various monetization strategies observed in practice and clarify how advertising affects a platform’s incentives to invest in content quality. The results are derived from unified economic forces that are captured by two key concepts: virtual advertising profits and the rent extraction–advertising trade-off. These concepts will be relevant in richer models, such as those with platform competition or strategic content providers, which we leave for future research.

References

- Simon P Anderson and Stephen Coate. Market provision of broadcasting: A welfare analysis. *Review of Economic Studies*, 72(4):947–972, 2005.
- Simon P Anderson and Joshua S Gans. Platform siphoning: Ad-avoidance and media content. *American Economic Journal: Microeconomics*, 3(4):1–34, 2011.
- Heski Bar-Isaac, Rahul Deb, and Matthew Mitchell. Selling certification, content moderation, and attention. *arXiv preprint arXiv:2506.12604*, 2025.
- Dirk Bergemann and Alessandro Bonatti. Targeting in advertising markets: implications for offline versus online media. *RAND Journal of Economics*, 42(3):417–443, 2011.
- Cristina Caffarra. ‘follow the money’—mapping issues with digital platforms into actionable theories of harm. *Art. N° 91579 e-Competitions*, 2019.
- Gaoyang Cai and Daniel F Spulber. The freemium pricing strategy and the opportunity cost of time. *Available at SSRN 4516760*, 2023.
- Ben Casner. Seller curation in platforms. *International Journal of Industrial Organization*, 72:102659, 2020.
- Jay Pil Choi and Doh-Shin Jeon. Platform design biases in ad-funded two-sided markets. *RAND Journal of Economics*, 54(2):240–267, 2023.

- Jay Pil Choi, Doh-Shin Jeon, and Byung-Cheol Kim. Net neutrality, business models, and internet interconnection. *American Economic Journal: Microeconomics*, 7(3): 104–141, 2015.
- Competition & Markets Authority. Online platforms and digital advertising: Market study final report. *London: Competition and Markets Authority*, 2020.
- Roberto Corrao, Joel P Flynn, and Karthik A Sastry. Nonlinear pricing with underutilization: A theory of multi-part tariffs. *American Economic Review*, 113(3):836–860, 2023.
- Jacques Crémer, Yves-Alexandre de Montjoye, and Heike Schweitzer. Competition policy for the digital era. *Report for the European Commission*, 2019.
- Ettore Damiano and Hao Li. Price discrimination and efficient matching. *Economic Theory*, 30(2):243–263, 2007.
- Raymond J Deneckere and Preston R McAfee. Damaged goods. *Journal of Economics & Management Strategy*, 5(2):149–174, 1996.
- Federico Etro. Device-funded vs ad-funded platforms. *International Journal of Industrial Organization*, 75:102711, 2021.
- Lapo Filistrucchi, Tobias J Klein, and Thomas O Michielsen. Assessing unilateral merger effects in a two-sided market: an application to the dutch daily newspaper market. *Journal of Competition Law and Economics*, 8(2):297–329, 2012.
- Jean J Gabszewicz, Didier Laussel, and Nathalie Sonnac. Programming and advertising competition in the broadcasting industry. *Journal of Economics & Management Strategy*, 13(4):657–669, 2004.
- Alfred Galichon. *Optimal transport methods in economics*. Princeton University Press, 2018.

- Renato Gomes and Alessandro Pavan. Many-to-many matching and price discrimination. *Theoretical Economics*, 11(3):1005–1052, 2016.
- Renato Gomes and Alessandro Pavan. Price customization and targeting in matching markets. *RAND Journal of Economics*, 55(2):230–265, 2024.
- Shota Ichihashi, Doh-Shin Jeon, and Byung-Cheol Kim. Mechanism design and innovation incentive for an ad-funded platform. *Available at SSRN 5016687*, 2024.
- Doh-Shin Jeon, Byung-Cheol Kim, and Domenico Menicucci. Second-degree price discrimination by a two-sided monopoly platform. *American Economic Journal: Microeconomics*, 14(2):322–369, 2022.
- Johannes Johnen and Robert Somogyi. Deceptive features on platforms. *The Economic Journal*, 134(662):2470–2493, 2024.
- Justin P Johnson. Targeted advertising and advertising avoidance. *RAND Journal of Economics*, 44(1):128–144, 2013a.
- Terence R Johnson. Matching through position auctions. *Journal of Economic Theory*, 148(4):1700–1713, 2013b.
- Jean-Jacques Laffont and David Martimort. The theory of incentives: the principal-agent model. In *The theory of incentives*. Princeton university press, 2009.
- Oliver Latham, Gaber Burnik, Rafael Jiménez Durán, and Hasnain Khaki. Media plurality and online intermediation of news consumption, 2022.
- Song Lin. Two-sided price discrimination by media platforms. *Marketing Science*, 39(2):317–338, 2020.
- Yi Liu, Pinar Yildirim, and Z John Zhang. Implications of revenue models and technology for content moderation strategies. *Marketing Science*, 41(4):831–847, 2022.

- Leonardo Madio and Martin Quinn. Content moderation and advertising in social media platforms. *Journal of Economics & Management Strategy*, 2024.
- Paul Milgrom and Ilya Segal. Envelope theorems for arbitrary choice sets. *Econometrica*, 70(2):583–601, 2002.
- Michael Mussa and Sherwin Rosen. Monopoly and product quality. *Journal of Economic Theory*, 18(2):301–317, 1978.
- Roger B Myerson. Optimal auction design. *Mathematics of Operations Research*, 6(1):58–73, 1981.
- Amaya Noain-Sánchez. Addressing the impact of artificial intelligence on journalism: the perception of experts, journalists and academics. *Communication Society*, 35(3):105–121, 2022.
- OECD. Competition issues in news media and digital platforms. *OECD Roundtables on Competition Policy Papers*, 2021. URL <https://doi.org/10.1787/a877a656-en>.
- Martin Peitz and Tommaso M Valletti. Content and advertising in the media: Pay-tv versus free-to-air. *International Journal of Industrial organization*, 26(4):949–965, 2008.
- Andrea Prat and Tommaso Valletti. Attention oligopoly. *American Economic Journal: Microeconomics*, 14(3):530–557, 2022.
- Susumu Sato. Freemium as optimal menu pricing. *International Journal of Industrial Organization*, 63:480–510, 2019.
- Fiona Scott Morton, Federico Etro, Oliver Latham, and Cristina Caffarra. Designing regulation for digital platforms: why economists need to work on business models. *Vox-CEPR Policy Portal*, 2020.

Stigler Committee on Digital Platforms. *Final Report*. Stigler Center for the Study of Economy and the State, 2019.

Tat-How Teh. Platform governance. *American Economic Journal: Microeconomics*, 14(3):213–254, 2022.

Frank Yang. Costly multidimensional screening. *arXiv preprint arXiv:2109.00487*, 2021.

Yusuke Zennyō. Freemium competition among ad-sponsored platforms. *Information Economics and Policy*, 50:100848, 2020.

Appendix

A Proofs of [Lemma 1](#) and [Theorem 1](#)

A.1 Proof of [Lemma 1](#)

We solve problem (R- θ) for any given type θ . Hereafter, we refer to the θ -virtual ad profit simply as the virtual ad profit. We begin with the case of negative virtual types.

First, the platform prioritizes displaying ads with higher virtual ad profits and never displays ads with negative ones. Also, the platform is restricted to displaying at most one ad for each item. Thus, instead of the full set \mathcal{J} of ads, we can focus on ads that could potentially be displayed. Specifically, let J^+ be the number of ads that have a strictly positive virtual ad profit. If $J^+ < I$ —i.e., if the number of such ads is strictly smaller than the number of items—then let \mathcal{J}^* be the set that consists of (i) all ads with strictly positive virtual ad profits and (ii) $I - J^+$ fictitious ads, which have zero virtual ad profits. If $J^+ \geq I$, then let \mathcal{J}^* be the set that consists of ads that have the first to the I -th highest virtual ad profit. The set \mathcal{J}^* depends on type θ .

The platform can now focus on a subclass of advertising policies that specify a one-to-one exhaustive matching between the items in \mathcal{I} and the ads in \mathcal{J}^* , where matching with a fictitious ad corresponds to matching with no ad in our original setup.²³ Let \mathcal{M}^* be the set of such advertising policies. Then, we can write the platform's relaxed problem for type θ as

$$\max_{\pi \in \mathcal{M}^*} \sum_{i \in \mathcal{I}} \max(0, v(\theta)q(i) + \rho_\theta(\pi(i))). \quad (\text{A.1})$$

We show that the function $\hat{V}_\theta(q, \rho) = \max(0, v(\theta)q + \rho)$ is submodular in (q, ρ) . Take any q_L and q_H such that $q_H > q_L$. We have

$$\hat{V}_\theta(q_H, \rho) - \hat{V}_\theta(q_L, \rho) = \begin{cases} 0 & \text{if } \rho < -v(\theta)q_L \\ -v(\theta)q_L - \rho & \text{if } -v(\theta)q_L \leq \rho \leq -v(\theta)q_H \\ v(\theta)(q_H - q_L) & \text{if } \rho > -v(\theta)q_H, \end{cases}$$

which is overall decreasing in ρ . Thus $\hat{V}_\theta(q, \rho)$ is submodular in (q, ρ) . As a result, a negative assortative matching between items in \mathcal{I} and ads in \mathcal{J}^* solves (A.1) (e.g., Galichon, 2018, Theorem 3.4 and Theorem 4.3).

We now translate the solution of (A.1) to that of (R- θ). First, the negative assortative matching between items in \mathcal{I} and ads in \mathcal{J}^* is equivalent to the ρ_θ -negative assortative policy, i.e., they match lower-quality items with ads that have higher values of $\rho_\theta(j)$, until we exhaust ads that have positive values of $\rho_\theta(j)$. Let $\pi^*(\cdot|\theta) : \mathcal{I} \rightarrow \mathcal{J} \cup \{\emptyset\}$ denote the optimal advertising policy (for type θ) in our original formulation. The platform then allocates each item-ad pair $(i, \pi^*(i|\theta))$ if and only if $v(\theta)q(i) + \rho_\theta(\pi^*(i|\theta)) \geq 0$, which reduces to the allocation policy described in the lemma.

The case of positive virtual types follows the same logic. In this case, $\hat{V}_\theta(q_H, \rho) - \hat{V}_\theta(q_L, \rho) = v(\theta)(q_H - q_L)$ is independent of ρ and thus trivially submodular.

²³By one-to-one exhaustive matching, we mean that every item is matched with some ad in \mathcal{J}^* .

Having solved the relaxed problem, we now verify the monotonicity of allocation. Given the allocation that solves the relaxed problem, we adopt the following notation: $Q(\theta) \triangleq \sum_{i \in B(\theta)} q(i)$, $D(\theta) \triangleq \sum_{i \in B(\theta)} d(\pi^*(i|\theta))$, and $R(\theta) \triangleq \sum_{i \in B(\theta)} r(\pi^*(i|\theta))$. The monotonicity constraint requires that $Q(\theta) - D(\theta)$ is non-decreasing in θ .

Suppose to the contrary that there are some θ_L, θ_H such that $\theta_H > \theta_L$ but $Q(\theta_L) - D(\theta_L) > Q(\theta_H) - D(\theta_H)$. At the solution to the relaxed problem, we have

$$v(\theta_L)[Q(\theta_L) - D(\theta_L)] + R(\theta_L) \geq v(\theta_L)[Q(\theta_H) - D(\theta_H)] + R(\theta_H).$$

If we replace $v(\theta_L)$ by $v(\theta_H)$ in both sides, the LHS increases more than the RHS, because $Q(\theta_L) - D(\theta_L) > Q(\theta_H) - D(\theta_H) \geq 0$. The resulting inequality is

$$v(\theta_H)[Q(\theta_L) - D(\theta_L)] + R(\theta_L) > v(\theta_H)[Q(\theta_H) - D(\theta_H)] + R(\theta_H).$$

This is a contradiction, because the platform could have increased its virtual surplus by replacing the outcome for type θ_H with that for θ_L . Because the item and ad allocation satisfies monotonicity, the solution to the relaxed problem constitutes the optimal mechanism. Finally, the binding IR constraint for type $\theta = 0$ and local IC constraints pin down the monetary transfer as shown in the lemma (see, e.g., Lemma 2 of [Myerson 1981](#)). \square

A.2 Proof of [Theorem 1](#)

We adopt the following notation: For any item $i \in \mathcal{I}$, ad $j \in \mathcal{J} \cup \{\emptyset\}$, and type $\theta \in [0, 1]$, let $\hat{V}_\theta(i, j) = v(\theta)q(i) + \rho_\theta(j)$ be the contribution to the virtual surplus when the platform allocates an item-ad pair (i, j) to type θ .

First, we prove Part 1, the case of negative virtual types. The platform allocates item i to type θ iff $v(\theta)q(i) + \rho_\theta(\pi^*(i|\theta)) \geq 0$. The LHS is strictly decreasing in $q(i)$, because $v(\theta)q(i)$ is decreasing in $q(i)$ and a higher $q(i)$ is matched with a lower virtual ad profit. Thus, the platform allocates items whose quality levels are below some type-

dependent threshold, $\bar{q}(\theta)$. Let us define $\bar{q}(\theta)$ to be 0 if the platform allocates no items to type θ ; and $\bar{q}(\theta) = q(i)$ if the highest quality level allocated to type θ equals $q(i)$.

We show that $\bar{q}(\theta)$ is non-decreasing in types on $[0, \theta^0]$. Suppose to the contrary that there are some θ_L and θ_H with $\theta_L < \theta_H \leq \theta^0$ such that the lower type θ_L faces a strictly higher quality threshold than type θ_H . Let i_L denote the highest-quality item that type θ_L receives, which implies that type θ_H does not receive item i_L . Type θ_L receives item i_L along with some ad $j_L \neq \emptyset$. Indeed, if type θ_L received item i_L with no ad, the contribution to the platform's virtual surplus would be $v(\theta_L)q(i_L) < 0$.

There are now two cases to consider. First, if ad j_L is not allocated to type θ_H , then we obtain a contradiction, because $\hat{V}_{\theta_L}(i_L, j_L) \geq 0$ implies $\hat{V}_{\theta_H}(i_L, j_L) \geq 0$, so the platform could increase its profit by allocating (i_L, j_L) to type θ_H .

Second, suppose that the platform allocates ad j_L to type θ_H along with some item other than item i_L . The fact that type θ_L receives pair (i_L, j_L) implies that any item allocated to type θ_L , whose quality is lower than i_L , comes with some ad, because the advertising policy for θ_L is the ρ_{θ_L} -negative assortative matching. Then, it holds that type θ_L receives a strictly greater number of ads than type θ_H , so there must be some ad, say \hat{j}_L , which is allocated to θ_L but not to θ_H . We have $\rho_{\theta_L}(\hat{j}_L) \geq \rho_{\theta_L}(j_L)$, because ad \hat{j}_L is matched with a lower-quality item than ad j_L . Thus, we have $\hat{V}_{\theta_L}(i_L, \hat{j}_L) \geq \hat{V}_{\theta_L}(i_L, j_L) \geq 0$, which implies $\hat{V}_{\theta_H}(i_L, \hat{j}_L) \geq \hat{V}_{\theta_L}(i_L, \hat{j}_L) \geq 0$, i.e., the platform can increase its profit by allocating (i_L, \hat{j}_L) to type θ_H . This is a contradiction.

The second part of the theorem holds because for any θ with $v(\theta) < 0$, we have $\rho_\theta(j) = r(j) - v(\theta)d(j) > 0$ for every $j \in \mathcal{J}$.

We now prove Part 2, the case of positive virtual types. The first part holds because $v(\theta), q(i), \rho_\theta(\pi^*(i|\theta)) \geq 0$ implies $B(\theta) = \{i \in \mathcal{I} : v(\theta)q(i) + \rho_\theta(\pi^*(i|\theta)) \geq 0\} = \mathcal{I}$. For the second part, note that the number of ads allocated to type θ is equal to $\min(I, J_\theta^+)$, where J_θ^+ is the number of ads with positive virtual ad profits, i.e., $J_\theta^+ = |\{j \in \mathcal{J} : \rho_\theta(j) > 0\}|$. As θ increases, $\rho_\theta(j) = r(j) - v(\theta)d(j)$ decreases for any j , so J_θ^+ is non-increasing in θ , and so is $\min(I, J_\theta^+)$. Thus, higher types will see fewer ads. \square

B Details for Case 3 in Section 5

We provide details of the analysis for Case 3 in Section 5, which assumes $c > 1$ and $r(I) < \frac{-v(0)q(I)}{1-cv(0)}$. Because $c > 1$, there is a unique type $\theta^* \in (0, 1)$ such that $\rho_{\theta^*}(j) = r(j)(1 - cv(\theta^*)) = 0$ for every ad j . We have $\theta^* > \theta^0$, because $\rho_{\theta^0}(j) = r(j)(1 - cv(\theta^0)) = r(j) > 0$. Thus, any type $\theta > \theta^*$ receives all items without ads, and any type $\theta \in (\theta^0, \theta^*)$ receives all items and every item is matched with some ad.

Type $\theta = 0$ —and generically, a positive measure of types that includes $\theta = 0$ —receives a contract with zero price, because type 0’s value of any item-ad pair is 0. Also, Theorem 1 implies that type $\theta = 0$ receives some set of items whose quality levels are below some threshold. The set indeed excludes at least item I , which has the highest quality. The reason is as follows. For any $\theta < \theta^0$, the virtual advertising profit $r(j)(1 - cv(\theta))$ is proportional to $r(j)$. Thus, the negative assortative matching between items and ads implies that item I is matched with ad I , which (by notation) has the I -th highest ad revenue. Then, $r(I) < \frac{-v(0)q(I)}{1-cv(0)}$ implies that the contribution of item I under the optimal item-ad matching to type 0’s virtual surplus is negative.

C Proof of Lemma 2 and Proposition 2

To prove Lemma 2 and Proposition 2, we first establish a lemma. Recall that in the optimal mechanism, the platform’s virtual surplus from each type θ is written as the optimal value of the one-to-one matching problem, (A.1), where (i) the platform matches I items and I ads, (ii) each ad is expressed by its virtual advertising profit, and (iii) some ads could be fictitious ads with zero virtual advertising profit. Therefore, we first establish comparative statics for such a problem.

Abusing notation, let \mathcal{M}^* denote the set of all one-to-one exhaustive matching policies between I items and I ads, where items are represented by the profile of their quality levels, \mathbf{q} , and ads are represented by the profile of their virtual ad profits, $\boldsymbol{\rho}$. For any $\pi \in \mathcal{M}^*$ and $(i, j) \in \{1, \dots, I\}^2$, $\pi(i, j) = 1$ if item i and ad j are matched and

0 otherwise.

Lemma 3. Take any submodular function $\hat{V} : \mathbb{R}^2 \rightarrow \mathbb{R}$, and define

$$V(\mathbf{q}, \boldsymbol{\rho}) := \max_{\pi \in \mathcal{M}^*} \sum_{(i,j) \in \{1, \dots, I\}^2} \hat{V}(q(i), \rho(j)) \pi(i, j).$$

Then, V has decreasing differences in $(\mathbf{q}, \boldsymbol{\rho})$, i.e., for any $\bar{\boldsymbol{\rho}}, \underline{\boldsymbol{\rho}}, \bar{\mathbf{q}},$ and $\underline{\mathbf{q}}$ such that $\bar{\boldsymbol{\rho}} \geq \underline{\boldsymbol{\rho}}$ and $\bar{\mathbf{q}} \geq \underline{\mathbf{q}}$, we have $V(\bar{\mathbf{q}}, \bar{\boldsymbol{\rho}}) - V(\bar{\mathbf{q}}, \underline{\boldsymbol{\rho}}) \leq V(\underline{\mathbf{q}}, \bar{\boldsymbol{\rho}}) - V(\underline{\mathbf{q}}, \underline{\boldsymbol{\rho}})$.

Proof. The submodularity of \hat{V} implies that the platform adopts negative assortative matching, i.e., the k -th lowest q is matched with the k -th highest ρ . We write $q(k)$ for the k -th lowest quality in \mathbf{q} and $\rho(k)$ for the k -th highest virtual ad profit in $\boldsymbol{\rho}$. Then, we have

$$V(\mathbf{q}, \boldsymbol{\rho}) = \sum_{k=1}^I \hat{V}(q(k), \rho(k)). \quad (\text{C.2})$$

Because $\bar{\mathbf{q}} \geq \underline{\mathbf{q}}$, we have $\bar{\mathbf{q}}_k \geq \underline{\mathbf{q}}_k$ for each $k = 1, \dots, I$. Similarly, $\bar{\boldsymbol{\rho}} \geq \underline{\boldsymbol{\rho}}$ implies $\bar{\rho}_k \geq \underline{\rho}_k$. The submodularity of \hat{V} implies

$$\hat{V}(\bar{\mathbf{q}}_k, \bar{\rho}_k) - \hat{V}(\bar{\mathbf{q}}_k, \underline{\rho}_k) \leq \hat{V}(\underline{\mathbf{q}}_k, \bar{\rho}_k) - \hat{V}(\underline{\mathbf{q}}_k, \underline{\rho}_k).$$

Adding up both sides with respect to $k = 1, \dots, I$, we conclude that V has decreasing differences in $(\mathbf{q}, \boldsymbol{\rho})$. \square

We are now ready to prove [Lemma 2](#).

Proof of Lemma 2. Suppose that advertising set \mathcal{J}_H attains uniformly higher virtual ad profits than \mathcal{J}_L . As we have shown in the proof of [Proposition 1](#), for each $\theta \in [0, 1]$, type θ 's virtual surplus $\hat{V}_\theta(q, \rho)$ is submodular in (q, ρ) (trivially so if $\theta > \theta^0$). Thus for any $\theta \in [0, 1]$, [Lemma 3](#) implies that

$$\hat{V}_\theta(\mathbf{q}, \boldsymbol{\rho}_\theta) := \max_{\pi \in \mathcal{M}^*} \sum_{(i,j) \in \{1, \dots, I\}^2} \hat{V}_\theta(q(i), \rho_\theta(j)) \pi(i, j) \quad (\text{C.3})$$

has decreasing differences in $(\mathbf{q}, \boldsymbol{\rho}_\theta)$. Here, $\boldsymbol{\rho}_\theta \in \mathbb{R}_+^I$ is the profile of virtual ad profits (including those of fictitious ads) constructed in the proof of [Theorem 1](#) for each θ . Let $\hat{V}_\theta(\mathbf{q}, \mathcal{J})$ denote the maximized virtual surplus from type θ when written as a function of the quality vector \mathbf{q} and advertising set \mathcal{J} .

If the advertising set changes from \mathcal{J}_L to \mathcal{J}_H , then for every θ , the corresponding vector of the top I virtual ad profits (ranked in the descending order) increases from $\boldsymbol{\rho}_\theta$ to $\bar{\boldsymbol{\rho}}_\theta$. By the decreasing-difference property of $\hat{V}_\theta(\mathbf{q}, \boldsymbol{\rho}_\theta)$, $\hat{V}_\theta(\mathbf{q}, \mathcal{J}_H) - \hat{V}_\theta(\mathbf{q}, \mathcal{J}_J)$ decreases in \mathbf{q} . Therefore, $\Pi(\mathbf{q}, \mathcal{J}_H) - \Pi(\mathbf{q}, \mathcal{J}_J)$ also decreases in \mathbf{q} .

Finally, for any fixed \mathcal{J} , $\Pi(\mathbf{q}, \mathcal{J})$ is trivially supermodular in $\mathbf{q} = (q(1), \dots, q(I))$. Also, by assumption, $-C(\mathbf{q})$ is supermodular. Thus, $\Pi(\mathbf{q}, \mathcal{J}) - C(\mathbf{q})$ is supermodular in \mathbf{q} . The standard argument of monotone comparative statics establishes the result. \square

Proof of Proposition 2. To show the first part, note that a subscription model arises if and only if $\mathcal{J} = \emptyset$, whereas a hybrid model arises only if $\mathcal{J} \neq \emptyset$. By Part (i) of [Proposition 1](#), a platform with a subscription model chooses a higher quality profile.

To show the second part, suppose the platform adopts a hybrid model with an advertising set \mathcal{J} . We then expand \mathcal{J} to $\hat{\mathcal{J}}$ by adding I ads, each of which has a sufficiently high ad revenue and zero disutility level so that (a) for every type θ , its θ -virtual ad profit exceeds any of the θ -virtual ad profits that could arise from \mathcal{J} , and (b) the platform is willing to allocate every item to all types by matching it with one of the newly added ads. Because the lowest type is 0, the platform can implement such an allocation rule only by setting a price of 0 to all types. In such a case, the platform's total revenue equals advertising revenue, i.e., the platform adopts a purely ad-funded model. By Part (i) of [Proposition 1](#), such a platform chooses a lower quality profile than a hybrid model.

Finally, any purely ad-funded platform must have the same set of optimal quality profiles as the purely ad-funded platform described in the previous paragraph. Indeed, if two platforms are purely ad-funded, their total revenues are equal to respective ad-

vertising revenues, which depend on their advertising sets and the number of allocated items (which equals I), but not on their quality levels. Thus, the set of optimal quality profiles for these platforms is equal to $\arg \min_{\mathbf{q}} C(\mathbf{q})$; otherwise, we obtain a contradiction, because a platform could strictly increase total revenue by choosing a quality profile in $\arg \min_{\mathbf{q}} C(\mathbf{q})$ while maintaining the same optimal mechanism. \square

D Omitted Materials for Section 8.1

In this appendix, we assume that a consumer's utility from each item-ad pair (i, j) equals

$$\theta[q(i) - d(j)] - a,$$

where $a > 0$ is the attention cost of consuming any item-ad pair. The attention cost a is exogenous and common across all consumers. Therefore, if a consumer consumes the set of items B under an advertising policy π and pays T , she obtains a payoff of

$$\theta \sum_{i \in B} [q(i) - d(\pi(i))] - a|B| - T, \tag{D.4}$$

where $|B|$ is the cardinality of the set B .

The specification in (D.4), combined with the assumption that consumers can freely dispose of allocated item-ad pairs, poses a new challenge for the platform. To see this, suppose that the platform wishes to assign to a consumer a pair consisting of a low-quality item and an annoying ad, such that $q(i) - d(j) - a < 0$. If consumption were contractible, the platform could offer a negative price (i.e., $T < 0$) in exchange for consumption. However, in this new setup, the consumer facing such an offer would receive the monetary transfer and then ignore the item-ad pair.

In this extended setup, we need to modify the IC constraint to reflect the possibility that a consumer consumes some strict subset of the allocated item-ad pairs after reporting her type. Thus, we replace the incentive compatibility constraint (IC) with

the following stronger constraint:

$$\theta \sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))] - a |B(\theta)| - T(\theta) \geq \max_{\theta', B' \subseteq B(\theta')} \left\{ \theta \sum_{i \in B'} [q(i) - d(\pi(i|\theta'))] - a |B'| - T(\theta') \right\}, \quad (\text{ICO})$$

where ICO stands for “incentive compatibility and obedience.” Here, the LHS is the consumer’s utility from truthfully reporting her type θ and consuming all allocated item-ad pairs $\{(i, \pi(i|\theta))\}_{i \in B(\theta)}$. The RHS subsumes two possibilities: (i) the consumer should not strictly benefit from truthfully reporting her type θ and then ignoring some items in $B(\theta)$; and (ii) the consumer should not strictly benefit from any *double deviation*, whereby she misreports her type as $\theta' \neq \theta$ and consumes a subset $B' \subseteq B(\theta')$ of the allocated item-ad pairs. Our baseline model assumes $a = 0$, so constraints (IC) and (ICO) coincide.

General characterization of the optimal mechanism is beyond the scope of the paper, but we characterize the optimal mechanism under two different assumptions. We begin with the following result.

Proposition 3. *Suppose that every ad has zero disutility level, i.e., $d(j) = 0$ for all $j \in \mathcal{J}$. Then, in the optimal mechanism, the platform allocates to each type θ all items whose quality levels $q(i)$ belong to some interval $[\underline{q}(\theta), \bar{q}(\theta)]$, where $\underline{q}(\theta)$ is decreasing, $\bar{q}(\theta)$ is increasing, and $\bar{q}(\theta)$ equals the highest possible quality $q(I)$ whenever θ has a positive virtual value. Moreover, every item is matched with some ad.*

The main difference between the mechanism in Proposition 3 and that in Theorem 1 is that under non-contractible consumption, each type’s consumption bundle has a lower bound, $\underline{q}(\theta)$, in terms of quality level, because consumers would ignore items if their quality is too low. At the same time, the optimal mechanism preserves two key features from Theorem 1: (i) the platform may serve negative virtual types; and (ii) whenever the platform accommodates negative virtual types, it prioritizes allocating low-quality items in order to save information rents, which leads to an upper bound, $\bar{q}(\theta)$. As one’s type θ increases, a consumer progressively gains access to, and thus

consumes, lower-quality items as well as higher-quality items.

First, we describe the part of the proofs that is common between [Proposition 3](#) and [Proposition 4](#), which we will show later. Under noncontractible consumption, the platform's problem becomes as follows:

$$\max_{\{(\pi(\cdot|\theta), B(\theta), T(\theta))\}_{\theta \in [0,1]}} \int_0^1 T(\theta) + \sum_{i \in B(\theta)} r(\pi(i|\theta)) dF(\theta)$$

$$\begin{aligned} \text{subject to} \quad & \theta \sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))] - a|B(\theta)| - T(\theta) \\ & \geq \max_{\theta', B' \subseteq B(\theta')} \theta \sum_{i \in B'} [q(i) - d(\pi(i|\theta'))] - a|B'| - T(\theta'), \end{aligned} \quad (\text{ICO})$$

$$\theta \sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))] - a|B(\theta)| - T(\theta) \geq 0, \forall \theta \in [0, 1]. \quad (\text{IRO})$$

Both for [Proposition 3](#) and [Proposition 4](#), we solve this problem in two steps. First, we solve a problem in which we do not allow double deviations. We call such a problem the weak problem. Second, we verify that the weak problem's solution satisfies the full constraint, [\(ICO\)](#).

The weak problem is written as follows:

$$\max_{\{(\pi(\cdot|\theta), B(\theta), T(\theta))\}_{\theta \in [0,1]}} \int_0^1 T(\theta) + \sum_{i \in B(\theta)} r(\pi(i|\theta)) dF(\theta) \quad (\text{P1})$$

$$\begin{aligned} \text{subject to} \quad & \theta \sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))] - a|B(\theta)| - T(\theta) \\ & \geq \theta \sum_{i \in B(\theta')} [q(i) - d(\pi(i|\theta'))] - a|B(\theta')| - T(\theta'), \forall \theta, \theta' \in [0, 1] \end{aligned} \quad (\text{IC})$$

$$\theta [q(i) - d(\pi(i|\theta))] - a \geq 0, \forall \theta, \forall q(i) \in B(\theta) \quad (\text{O})$$

$$\theta \sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))] - a|B(\theta)| - T(\theta) \geq 0, \forall \theta \in [0, 1] \quad (\text{IRO})$$

The constraint (IC) is the same as the original IC constraint. The obedience constraint (O) requires that users optimally consume all allocated items after truthful reporting.

The weak problem is equivalent to the following problem of virtual surplus maximization:

$$\max_{\{(\pi(\cdot|\theta), B(\theta))\}_{\theta \in [0,1]}} \int_0^1 \sum_{i \in B(\theta)} v(\theta)[q(i) - d(\pi(i|\theta))] + r(\pi(i|\theta)) - a \, dF(\theta), \quad (\text{P2})$$

subject to the obedience constraint, (O), and the monotonicity of allocation. Hereafter, the relaxed problem refers to problem (P2) without the monotonicity constraint.

D.1 Proof of Proposition 3

We solve (P2), verify the monotonicity of allocation, then verify the constraint (ICO).

Allocation for positive virtual types. Take any $\theta \geq \theta^0$. By assumption, the disutility level of every ad equals 0. Thus, regardless of whether item i is matched with an ad, a user will consume item i if and only if $q(i) \geq \frac{a}{\theta}$. Thus, once we restrict our attention to the set of items defined by $\mathcal{I}_+(\theta) \triangleq \{i \in \mathcal{I} : q(i) \geq \frac{a}{\theta}\}$, problem (P2) is equivalent to our baseline setup with contractible consumption. Because $\max(v(\theta)q + r - a, 0)$ is supermodular in (q, r) , the platform will adopt a positive assortative matching between the items in $\mathcal{I}_+(\theta)$ and ads in \mathcal{J} (in terms of $q(i)$ and $r(j)$). That is, the platform allocates the i -th highest-quality item, $I - i + 1$, with the i -th highest-revenue ad, denoted by $\pi(i)$. The platform then allocates any item-ad pair whose contribution to virtual surplus is nonnegative, i.e., $v(\theta)q(I - i + 1) + r(\pi(i)) - a \geq 0$. This inequality is more likely to hold for a higher θ and a higher-quality item. Thus, the set of allocated items is written as $[\underline{q}(\theta), \bar{q}(\theta)]$ with $\bar{q}(\theta) = q(I)$, and $\underline{q}(\theta)$ decreases in θ .

Allocation for negative virtual types. Take any $\theta < \theta^0$. By the same argument as the case of positive virtual types, the platform allocates items only if they belong to

$\mathcal{I}_+(\theta)$. Thus, once we restrict attention to the items in $\mathcal{I}_+(\theta)$, problem (P2) is equivalent to our baseline setup with contractible consumption. As a result, the platform will adopt a negative assortative matching between the items in $\mathcal{I}_+(\theta)$ and ads in \mathcal{J} . The platform then allocates any item-ad pair whose contribution to virtual surplus is nonnegative, i.e., $v(\theta)q(i) + r(\pi(i)) - a \geq 0$. The expression $v(\theta)q(i) + r(\pi(i)) - a$ is decreasing in i (i.e., it takes a lower value for a higher-quality item), so the platform allocates to type θ all items whose quality levels are below some threshold. Combined with the fact that \mathcal{I}_+ consists of items whose quality levels are above the type-dependent threshold $\frac{a}{\theta}$, we conclude that the platform allocates items whose quality levels belong to the interval $[\underline{q}(\theta), \bar{q}(\theta)]$, where $\underline{q}(\theta) = \frac{a}{\theta}$ is decreasing in θ , whereas $\bar{q}(\theta)$, the maximum quality threshold, is increasing in θ .

Monotonicity. Monotonicity requires that the term $\sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))]$, which equals $\sum_{i \in B(\theta)} q(i)$, is increasing in θ . This term is increasing in θ for $\theta < \theta^0$ and $\theta \geq \theta^0$ separately, because among the types whose virtual types have the same sign, a higher type receives a larger set of items. Thus, it suffices to show that the gross quality $\sum_{i \in B(\theta)} q(i)$ does not decrease when the type crosses the threshold θ^0 where $v(\theta^0) = 0$. Note that the number of items allocated to type θ^0 is $\min(|I_+(\theta^0)|, K)$, where $|I_+(\theta^0)|$ is the number of items that satisfy type θ^0 's obedience constraint, and K is the number of ads that satisfy $r(j) - a \geq 0$. In contrast, the number of items allocated to any type $\theta < \theta^0$ never exceeds $\min(|I_+(0)|, K)$, and they receive lower-quality items. Thus, the total quality $\sum_{i \in B(\theta)} q(i)$ allocated to type θ^0 is always weakly greater than that to $\theta < \theta^0$. Thus, the monotonicity of allocation holds.

Solution to the weak problem satisfies (ICO). We show that the solution to the weak problem (i.e., the allocation rule constructed above and the transfer rule induced by the local IC constraint) satisfies (ICO). The proof of this step consists of two parts. First, suppose that a consumer has type θ but misreports to be $\hat{\theta} < \theta$. Then, it is optimal for type θ to consume all items in $B(\hat{\theta})$. Indeed, any item $i \in B(\hat{\theta})$ satisfies

$\hat{\theta}q(i) - a \geq 0$, which also satisfies $\theta q(i) - a \geq 0$. Thus, there is no double deviation that involves downward misreporting.

Second, suppose to the contrary that there is some profitable double deviation for type θ that involves upward misreporting, i.e., type θ misreports to be type $\hat{\theta} > \theta$ and consumes some strict subset of allocated items. Let $B'(x) \triangleq \{i \in B(\hat{\theta}) : q(i) \geq a/x\}$ denote the optimal consumption that arises when type x misreports to be $\hat{\theta}$. Also, define $Q^x(\hat{\theta}) \triangleq \sum_{i \in B'(x)} q(i)$ and $N^x(\hat{\theta}) \triangleq |B'(x)|$.

The condition that type θ benefits from this double deviation is written as

$$\theta Q^\theta(\hat{\theta}) - aN^\theta(\hat{\theta}) - T(\hat{\theta}) > \theta Q(\theta) - aN(\theta) - T(\theta), \quad (\text{D.5})$$

which we can write as

$$\theta Q^\theta(\hat{\theta}) - aN^\theta(\hat{\theta}) - U(\theta) > T(\hat{\theta}), \quad (\text{D.6})$$

where

$$U(\theta) = \theta Q(\theta) - aN(\theta) - T(\theta). \quad (\text{D.7})$$

Also denote

$$U(\hat{\theta}) = \hat{\theta} Q(\hat{\theta}) - aN(\hat{\theta}) - T(\hat{\theta}). \quad (\text{D.8})$$

Solving the above equation with respect to $T(\hat{\theta})$ and plugging it into [Equation D.6](#), we obtain

$$U(\hat{\theta}) - U(\theta) > \hat{\theta} Q(\hat{\theta}) - aN(\hat{\theta}) - [\theta Q^\theta(\hat{\theta}) - aN^\theta(\hat{\theta})]. \quad (\text{D.9})$$

By the local IC, we get

$$U(\hat{\theta}) - U(\theta) = \int_{\theta}^{\hat{\theta}} Q(x) \, dx. \quad (\text{D.10})$$

To arrange the RHS of [Equation D.9](#), we apply the envelope theorem to

$$V(y) \triangleq \max_{x \in [0,1]} yQ^x(\hat{\theta}) - aN^x(\hat{\theta}). \quad (\text{D.11})$$

This is the problem of type y that faces the content bundle for type $\hat{\theta}$ and optimally chooses how to truncate the bundle. Because $y \in \arg \max_{x \in [0,1]} yQ^x(\hat{\theta}) - aN^x(\hat{\theta})$, the envelope theorem of [Milgrom and Segal \(2002\)](#) implies that

$$\hat{\theta}Q(\hat{\theta}) - aN(\hat{\theta}) - [\theta Q^\theta(\hat{\theta}) - aN^\theta(\hat{\theta})] = V(\hat{\theta}) - V(\theta) = \int_{\theta}^{\hat{\theta}} Q^y(\hat{\theta}) dy. \quad (\text{D.12})$$

Combining [\(D.9\)](#), [\(D.10\)](#), and [\(D.12\)](#), we obtain

$$\int_{\theta}^{\hat{\theta}} Q(x) dx > \int_{\theta}^{\hat{\theta}} Q^x(\hat{\theta}) dx. \quad (\text{D.13})$$

This is a contradiction for the following reason. For each $x \in [\theta, \hat{\theta}]$, the set of allocated items, which induces $Q(x)$, is written as an interval $[q_1, q_2]$, where $q_1 \geq \frac{a}{x}$. The set of qualities allocated to $\hat{\theta}$ (leading to $Q(\hat{\theta})$) is written as $[q_3, q_4]$, where $q_4 \geq q_2$. If $q_3 \geq \frac{a}{x}$, then $Q(x) \leq Q(\hat{\theta}) = Q^x(\hat{\theta})$. If $q_3 < \frac{a}{x}$, then $Q(x)$ contains items in interval $[q_1, q_2]$, whereas $Q^x(\hat{\theta})$ comes from all items in interval $[\frac{a}{x}, q_4]$. Because $\frac{a}{x} \leq q_1$ and $q_4 \geq q_2$, $Q^x(\hat{\theta}) \geq Q(x)$. To sum up, we get $Q^x(\hat{\theta}) \geq Q(x)$ for all $x \in [\theta, \hat{\theta}]$, so [Equation D.13](#) must be violated. \square

We conclude this subsection by presenting an example in which (i) the assumptions for [Proposition 3](#) (and those for [Proposition 4](#) shown below) do not hold and (ii) the approach based on the weak problem fails, because the solution to the relaxed problem (of the weak problem) violates the monotonicity constraint.

Example 2. Suppose there is only one item with quality q and one ad with disutility level d and ad revenue r . Types are binary but the low type has a positive virtual type. Suppose that

$$\frac{a}{\theta_L} + d > q > \max \left(\frac{a}{\theta_H} + d, \frac{a}{v(\theta_L)} \right). \quad (\text{D.14})$$

(To construct a more concrete example, assume that the fraction of θ_H is low enough so that $v(\theta_L) \approx \theta_L$. Then we can take $d \approx \frac{a}{\theta_L} - \frac{a}{\theta_H}$, and then $q = \frac{a}{\theta_H} + d + \epsilon$.) Suppose r is very high. Then, the platform allocates the item with the ad to θ_H and without

ad to θ_L . Indeed, $\frac{a}{\theta_L} + d > q > \frac{a}{v(\theta_L)}$ implies that it is profitable for the platform to allocate even without an ad, and it cannot allocate with an ad due to moral hazard. Then we have $Q(\theta_H) = q - d < q = Q(\theta_L)$. This violates the monotonicity condition.

D.2 The Statement and Proof of Proposition 4

Despite the challenge described in Example 2, we can accommodate ads with strictly positive disutility levels when types are binary.

Assumption 1. Disutility levels are constant across all ads, i.e., there is some $d \geq 0$ such that $d(j) = d \geq 0$ for all $j \in \mathcal{J}$. Types are binary (i.e., $\Theta = \{\theta_L, \theta_H\}$) and the low type has a negative virtual value, i.e., $\theta_L - \frac{1-\beta}{\beta}(\theta_H - \theta_L) < 0$ where $\beta = \Pr(\theta = \theta_L)$.

Proposition 4. *Under Assumption 1, the optimal mechanism is as follows:*

1. *For type θ_H , the platform allocates a set of items whose quality levels exceed some threshold.*
2. *For type θ_L , the platform allocates all items whose quality levels belong to some interval $[\underline{q}(\theta_L), \bar{q}(\theta_L)]$, and matches every item with some ad.*

Proof. First, we derive the allocation for positive virtual types. Take the high type, $\theta = \theta_H$. Note that under binary types, we have $v(\theta_H) = \theta_H$. Recall that under Assumption 1, all ads have the same disutility level, d . If $\theta_H[q(i) - d] - a < 0$, or equivalently, if $q(i) < \frac{a}{\theta_H} + d$, then a user will never consume item i when it is matched with an ad. Also, item i generates a nonnegative virtual surplus for type θ without ads if and only if $\theta_H q(i) - a \geq 0$. Combining these inequalities, we conclude that any item i whose quality is in $[\frac{a}{\theta_H}, \frac{a}{\theta_H} + d)$ will be allocated to type θ without ads.

A user will consume any item with $q(i) \geq \frac{a}{\theta} + d$ whether or not it is matched with an ad. Because $v(\theta_H) = \theta_H$, the same inequality also ensures that the allocation weakly increases virtual surplus. Thus, the platform will allocate all items such that $q(i) \geq \frac{a}{\theta} + d$ and match them with ads that have higher revenue. To sum up, the

allocation to type θ_H is given by some top-down allocation with cutoff $\frac{a}{\theta_H}$, and the items whose quality levels are above some cutoff will be matched with ads.

Second, we derive the allocation for negative virtual types. Consider the low type, θ_L with $v(\theta_L) < 0$. First, the platform never allocates items without ads. Thus, it allocates item i only if $q(i) \geq \underline{q}(\theta_L) \triangleq \frac{a}{\theta} + d$. Once we restrict our attention to the set of items defined by $\{i \in \mathcal{I} : q(i) \geq \underline{q}(\theta_L)\}$, the problem is equivalent to our baseline setup with contractible consumption. As a result, the platform will adopt a negative assortative matching between these items and the ads that have nonnegative virtual advertising profits. The platform then allocates items whose quality levels are below some threshold, $\bar{q}(\theta_L)$. To sum up, the platform allocates to type θ_L all items whose quality levels are in $[\underline{q}(\theta), \bar{q}(\theta)]$. By construction, every item is matched with some ad.

We now verify the monotonicity of allocation. Type θ_H receives all items such that $q(i) \geq \frac{a}{\theta_H}$. In contrast, type θ_L receives a subset of items such that $q(i) \geq \frac{a}{\theta_L} + d$ and all of them are matched with ads. Thus, the term $\sum_{i \in B(\theta)} [q(i) - d(\pi(i|\theta))]$ is higher for type θ_H than type θ_L .

Finally, we show that the solution to the weak problem satisfies (ICO). First, if type θ_H misreports to be θ_L , then θ_H optimally consumes all items allocated, because any item that satisfies θ_L 's obedience constraint also satisfies θ_H 's obedience constraint. Thus, type θ_H has no profitable double deviation.

Second, suppose that type θ_L misreports to be θ_H . Note that the IC constraint for θ_H binds at the optimum, i.e.,

$$\theta_H[Q(\theta_H) - D(\theta_H)] - aN(\theta_H) - T(\theta_H) = \theta_H[Q(\theta_L) - D(\theta_L)] - aN(\theta_L) - T(\theta_L),$$

where $D(\theta)$ is the total disutility level incurred by type θ . Note that the gross payoff (excluding monetary transfer) of type θ_H increases by $\theta_H[Q(\theta_H) - D(\theta_H)] - aN(\theta_H) - [\theta_H(Q(\theta_L) - D(\theta_L)) - aN(\theta_L)]$ by reporting θ_H instead of θ_L . If type θ_H reports truthfully but consumes as if type θ_L would, then type θ_H who adopts such behavior would ignore some item-ad pairs, so the gain decreases and becomes

equal to $\theta_H[Q^L(\theta_H) - D^L(\theta_H)] - aN^L(\theta_H) - [\theta_H(Q(\theta_L) - D(\theta_L)) - aN(\theta_L)]$. We have $Q^L(\theta_H) - D^L(\theta_H) \geq Q(\theta_L) - D(\theta_L)$ because the allocation for type θ_H contains more items, and some items do not come with ads.

If type θ_L reports to be θ_H and consumes item-ad pairs optimally, which is the optimal double deviation, the gross gain is

$$\begin{aligned} & \theta_L[Q^L(\theta_H) - D^L(\theta_H)] - aN^L(\theta_H) - [\theta_L(Q(\theta_L) - D(\theta_L)) - aN(\theta_L)] \\ & \leq \theta_H([Q^L(\theta_H) - D^L(\theta_H)] - (Q(\theta_L) - D(\theta_L))) - aN^L(\theta_H) + aN(\theta_L) \\ & \leq T(\theta_H) - T(\theta_L). \end{aligned}$$

The first inequality comes from $\theta_H > \theta_L$ and $Q^L(\theta_H) - D^L(\theta_H) \geq Q(\theta_L) - D(\theta_L)$. Combining the first and last lines, we conclude that type θ_L does not benefit from double deviation. \square