This article was downloaded by: [2a00:8a60:5:111::1002] On: 26 October 2026, At: 08:16
Publisher: Institute for Operations Research and the Management Sciences (INFORMS)
INFORMS is located in Maryland, USA



Information Systems Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Hit the GAS: Designing Optimal Generalized Ad-Supported Subscription Mechanisms

Dominik Gutt, Sameer Mehta, Martin Quinn

To cite this article:
Dominik Gutt, Sameer Mehta, Martin Quinn (2026) Hit the GAS: Designing Optimal Generalized Ad-Supported Subscription Mechanisms. Information Systems Research

Published online in Articles in Advance 13 Jul 2026

. https://doi.org/10.1287/isre.2025.2329

This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as "*Information Systems Research*. Copyright © 2026 The Author(s). https://doi.org/10.1287/isre.2025.2329, used under a Creative Commons Attribution License: https://creativecommons.org/licenses/by/4.0/."

Copyright © 2026 The Author(s)

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.
For more information on INFORMS, its publications, membership, or meetings visit http://www.informs.org

https://pubsonline.informs.org/journal/isre

INFORMATION SYSTEMS RESEARCH
Articles in Advance, pp. 1–23
ISSN 1047-7047 (print), ISSN 1526-5536 (online)

Hit the GAS: Designing Optimal Generalized Ad-Supported Subscription Mechanisms

Dominik Gutt,[a] **Sameer Mehta,**[b,*] **Martin Quinn**[c]

[a] School of Business and Economics, RWTH Aachen University, Chair of Business Information Systems and Digital Transformation, 52062 Aachen, Germany; [b] Department of Technology and Operations Management, Rotterdam School of Management, Erasmus University, 3062 PA Rotterdam, Netherlands; [c] Marketing and Management Strategy, Institut Mines-Télécom Business School, 91000 Évry-Courcouronnes, France

*Corresponding author

dominik.gutt@rwth-aachen.de (D. Gutt), mehta@rsm.nl (S. Mehta), martin.quinn@imt-bs.eu (M. Quinn)

https://orcid.org/0000-0002-2633-9518 (D. Gutt), https://orcid.org/0000-0002-0410-3248 (S. Mehta), https://orcid.org/0000-0003-0800-2214 (M. Quinn)

Received: July 14, 2025
Revised: March 12, 2026; June 2, 2026
Accepted: June 19, 2026
Published Online in Just Accepted: July 13, 2026
Published Online in Articles in Advance: August 19, 2026

https://doi.org/10.1287/isre.2025.2329

Copyright: © 2026 The Author(s)

Open Access Statement: This work is licensed under a Creative Commons Attribution 4.0 International License. You are free to copy, distribute, transmit and adapt this work, but you must attribute this work as "*Information Systems Research*. Copyright © 2026 The Author(s). https://doi.org/10.1287/isre.2025.2329, used under a Creative Commons Attribution License: https://creativecommons.org/licenses/by/4.0/."

Abstract. Digital content platforms (DCPs), such as Netflix and Spotify, rely on subscriptions and advertising as their primary revenue sources. Beyond pure subscription-only and ad-only revenue models, DCPs increasingly blend these models by offering a two-tier menu: a free, ad-supported tier for price-sensitive users and a paid, ad-free tier for ad-sensitive users. In this paper, we introduce the generalized ad-supported subscription (GAS) mechanism—a broad class of subscription fee and ad intensity combinations that spans the continuum from ad only to subscription only—and nests the traditional mechanisms as special cases. Using a mechanism design framework, we characterize the revenue-maximizing GAS mechanism and compare its performance with the optimal ad-only, subscription-only, and two-tier mechanisms. Although GAS is optimal within this broad class, the simpler two-tier mechanism can achieve near-optimal revenue. We then characterize conditions under which the GAS mechanism delivers a material revenue advantage over the two-tier mechanism. Finally, we estimate our model parameters and empirically validate the theoretical results in the context of video-on-demand platforms.

History: Martin Bichler, Senior Editor; Pallab Sanyal, Associate Editor.
Funding: This work was supported by the Nederlandse Organisatie voor Wetenschappelijk Onderzoek [VI.Veni.231E.063].
Supplemental Material: The online appendix is available at https://doi.org/10.1287/isre.2025.2329.

Keywords: advertising • subscription • digital content platforms • revenue model • mechanism design

1. Introduction

Digital content platforms (DCPs), such as Netflix, Spotify, and the *New York Times*, have become foundational to the digital-media industry. In 2023, revenue in the digital content market reached $498.59 billion, with projections indicating a strong compound annual growth rate of 9.9% from 2024 through 2027 (Statista 2023). As these platforms continue to grow and shape the digital content market, their revenue models play a crucial role in defining both their business sustainability and consumer accessibility. Traditionally, there are two primary sources of revenue for the DCPs: advertising revenue from advertisers and subscription revenue from consumers. DCPs leverage these two levers strategically and typically adopt one of the following two revenue models.

The first is the subscription-only model, where customers pay a recurring fee (e.g., $9.99 per month for Apple TV+) to access content. This model provides DCPs with a stable revenue stream, but confines demand to consumers whose willingness-to-pay (WTP) for the service is at or above the subscription fee. On the upside, the consumers who subscribe to the service benefit from an ad-free experience.

Second, there are ad-supported or ad-only models, where consumers can access the content for free, and the DCP (e.g., Roku) generates revenues through ad display. The revenue collected in this model depends on the price the DCP can charge the advertiser and the number of consumers who view the ads. Compared with the subscription-only model, the ad-only model enables a larger consumer base to access the DCP's services because the consumers do not have to pay an access fee. However, under this model, consumers face an ad nuisance cost (Goldstein et al. 2014, Todri et al. 2020), which reduces their utility.

The tradeoff between these two widely used revenue models is straightforward: Subscription-only models provide a stable revenue stream by drawing from a

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

smaller consumer base, whereas ad-only models can potentially reach a broader audience (when consumers are not too ad sensitive), but rely on fluctuating advertising revenue and expose consumers to ad nuisance.

To address this tradeoff, DCPs have recently adopted a hybrid model: an ad-supported subscription model, which offers users the choice between an ad-free experience at a higher subscription fee and a lower or even zero fee with ads. A well-known example is YouTube, which is free with ads and $13.99 per month without ads. This model combines the advantages of both subscription-only and ad-only models by allowing customers to select their preferred option. In doing so, DCPs can attract high-WTP consumers who favor ad-free subscriptions and low-WTP consumers willing to view ads for free or reduced-cost access.

Given that these widely implemented revenue models form the foundation of most DCPs delivering information goods, a compelling question arises: With both advertising and subscription as revenue levers, what is a revenue-maximizing mechanism for the DCPs? Although the ad-supported subscription model provides a step toward flexibility by offering consumers a binary choice between ad-free and ad-supported options, it remains limited in scope. This limitation stems from the fact that offering a binary choice to the consumers may not fully capture the diverse range of their preferences and willingness to pay. In particular, the two choices may not completely align with the consumers' tradeoff between paying a subscription fee and incurring ad nuisance cost.

To address this question, our goal in this paper is to examine a broader class of mechanisms that leverage the two revenue sources to better meet consumer preferences and maximize revenue. Although a growing body of literature has emerged on revenue models for DCPs (Wang et al. 2023; Goli et al. 2024, 2025), the question of designing a revenue-maximizing mechanism for DCPs remains unexplored.

Identifying new revenue models is important not only for the growth of DCPs but also for understanding the broader implications for consumers. A natural concern is as follows: How do consumers fare under different revenue models? Simpler models like subscription only and ad supported offer limited flexibility in capturing consumer preferences. As a result, they may leave significant consumer welfare and revenue untapped. In contrast, although a hybrid model such as ad-supported subscription, enables better price discrimination than the two simpler models and thus helps generate higher revenue, it raises the question whether this additional revenue comes at the expense of consumer welfare. Although this is possible, hybrid models could also expand the overall economic pie, benefiting both the platform and its consumers. In our study, although analyzing a broader class of mechanisms, we also study how these mechanisms impact consumer welfare.

1.1. Key Results and Contributions

1.1.1. Comprehensive Model. As discussed above, our goal in this paper is to analyze a broader class of mechanisms that leverage the two sources of revenue—advertising and subscription—for DCPs. We refer to this broader class as the generalized ad-supported subscription (GAS) mechanism. In Section 3, we introduce the key elements of our model. We incorporate context-specific features such as subscription fees, ad rate, ad intensity, ad nuisance cost to the users, and endogenous usage time of the service. Additionally, we show that the broader class of GAS mechanisms provides a unified framework that encompasses subscription-only (hereafter, SUB), ad-supported (hereafter, ADS), and binary ad-supported subscription (hereafter, BAS) revenue models as special cases. This helps us conduct a head-to-head analysis of all the revenue models. Figure 1 illustrates a schematic representation of the four mechanisms that we study, with an illustration of subscription fees and ad intensities.

1.1.2. Optimal Mechanisms. Our model consists of a unit mass of users who are heterogeneous in the utility they derive from the content offered by the DCP. The type of the users is private information and the DCP has only distributional information about the user types. The goal of the DCP is to design a GAS mechanism—a set of rules that specify the ad intensity for each user type and the corresponding subscription payment from the consumers—that maximizes the total expected revenue from provisioning the content. In Section 4, we formulate and obtain the revenue-maximizing ADS, SUB, BAS, and GAS mechanisms and denote them as ADS*, SUB*, BAS*, and GAS*, respectively. By construction, the broader class of GAS mechanism encompasses the ADS, SUB, and BAS mechanisms; therefore, GAS* yields weakly higher revenue than the other three optimal mechanisms. More importantly, we show that the GAS* is operationally feasible and can be implemented as a posted-price mechanism (Proposition 1; Remark 1).

Figure 1. (Color online) Schematic Representation of Four Pricing Mechanisms

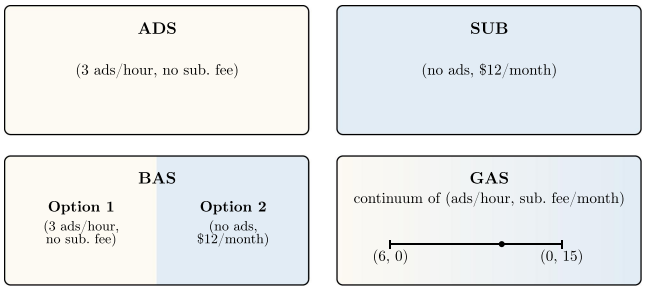

Notes. Numeric values are illustrative. Yellow denotes the ad-supported tier and blue denotes the subscription tier. ADS offers an ad-supported tier. SUB offers a subscription-only tier. BAS offers two discrete tiers (ad supported versus subscription). GAS offers a menu spanning a continuum of ad intensity and corresponding subscription fee combinations.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

1.1.3. Comparative Analysis of Mechanisms. In Section 5, we comprehensively compare the revenues under the four optimal mechanisms by analyzing their sensitivity to key model parameters. We find that BAS* and GAS* mechanisms better adjust to variation in model parameters, because they can partially offset losses in one revenue lever by shifting users toward the other. We also show that the GAS* mechanism generates weakly larger total welfare across all parameter combinations (Remark 2). Given that BAS* is a natural coarse approximation to GAS*, we numerically study the revenue gap between the two mechanisms and identify the parameter regions in which GAS* delivers a quantitatively meaningful improvement over BAS* (Section 5.2). In Section 8, we synthesize these comparative insights into a practical implementation guide for the DCP.

1.1.4. Model Extensions. In addition to our baseline framework, we explore several extensions in Section 6 that enhance both the practical applicability and policy relevance of our model.

- Capped Ad Intensity: Present-day digital environments are increasingly shaped by regulatory and ethical considerations that discourage excessive or intrusive advertising. For instance, frameworks such as the European Union's Digital Services Act emphasize the need for digital platforms to ensure a safe, transparent, and user-friendly experience. To align with these considerations, we introduce an upper bound on ad intensity and formulate a constrained mechanism design problem for the DCP. We show that when the cap binds, the expected revenue to the DCP under the optimal capped-GAS mechanism falls as the cap shrinks.
- Endogenous Ad Revenue Rate: Our base model assumes that the ad revenue rate is constant. Although this assumption resonates with practice, excessive ad intensity may reduce user engagement, leading to diminishing returns for advertisers and lower revenue per ad (Yang et al. 2025). To account for this effect, we endogenize the ad revenue rate as a decreasing function of ad intensity. Interestingly, we show that the optimal GAS mechanism for this setting scales back ad intensity for intermediate- and high-type users and assigns higher ad load on the low-type users.
- Two-Dimensional Heterogeneity: We also extend our base model to allow users to be privately heterogeneous in both content valuation and ad nuisance. Because optimal mechanism design under multidimensional private information is generally intractable, we use the discretize-and-discount idea from Madarász and Prat (2017) to construct a posted menu and obtain an approximation bound on its performance. Importantly, the approximation bound improves with the menu length. We further specialize this result to a menu of a given finite length and provide an explicit bound on its performance.
- Unary Ad-Supported (UAS) Mechanism: Beyond multitier mechanisms, the DCP may commit to a single service tier that combines a subscription fee with an ad load. In Section 6.4, we characterize this mechanism and identify the parameter regions in the (r, α) space where the UAS* mechanism dominates the ADS* and SUB* mechanisms.
- Competition: Finally, we study how the DCP operates under competition. We model competing platforms with an outside option to users and characterize the optimal GAS mechanism for the DCP. This helps us show how competitive pressure shapes the ad-intensity/subscription fees mix relative to the monopoly benchmark.

1.1.5. Empirical Validation. To validate our theoretical framework, we conduct an empirical study in the context of Video-on-Demand (VoD) streaming platforms using data collected on the Prolific platform (Section 7). We estimate the distribution of user types and other model parameters and test key assumptions in our model. Our empirical results corroborate the model primitives and contextualize the theoretical insights:

- Revenue Comparison: In the VoD setting, the GAS* mechanism yields 148%, 20%, and 1.8% higher expected revenue per consumer than ADS*, SUB*, and BAS* mechanisms, respectively. The near-optimality of BAS* has practical appeal because VoD platforms can obtain near-optimal revenue with a simple two-tier menu. The incremental advantage of GAS* stems primarily from intermediate-type users; because this segment is narrow, the revenue gain over BAS* is quantitatively small.
- Consumer Welfare: The SUB*, BAS*, and GAS* mechanisms generate comparable levels of expected consumer surplus—approximately 45% higher than ADS*. Although richer mechanisms sharpen price discrimination, the better matching of user preferences under BAS* and GAS* offsets the information rent extracted, so that the additional revenue need not come at the expense of consumer welfare.

2. Literature Review

The seminal work on information goods pricing has primarily focused on applications such as software sales and examined strategies including bundling (Bakos and Brynjolfsson 1999), product differentiation (Bhargava and Choudhary 2001), and versioning (Bhargava and Choudhary 2008), among others. With the rise of DCPs and streaming technology in recent years, more attention has been paid to developing pricing models tailored to this specific context, typically featuring subscription fees and ads. The early study by Fan et al. (2007) shows that the revenue of the DCP is higher when it offers an

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

ad-supported model and a subscription-supported model than either of the two in isolation. In contrast, we consider a continuum of ad-intensity and subscription fee combinations for users, which leads to a mechanism design problem for the DCP. Lin (2020) develops a two-sided monopoly versioning model with advertising, showing that price discrimination on the consumer side can strengthen the incentives to price-discriminate on the advertising side. Riggins (2002) finds that when advertising fees on DCPs decrease, the quality of ad-supported content declines, but at the same time allows platforms to raise prices for subscription-supported content. This aligns with the widely accepted notion that ads create nuisance that represents a cost to consumer's utility (Fan et al. 2007, Goldstein et al. 2014). DeValve and Pekeč (2022) show that when consumers differ only in a binary tolerance for advertising, the revenue-maximizing menu collapses to exactly two corner bundles: a free, ad-heavy tier and a paid, ad-free tier. Their model considers only binary user types and does not take into account the usage of the service (which is critical in determining the ad exposure to the users and the ad revenue to the DCP). In contrast, we model a continuum of user types and explicitly endogenize each user's usage of content. The ad revenue to the DCP is then determined by the usage of the users, the ad revenue rate, and the ad intensity chosen by the users. Ichihashi et al. (2024) hold ad exposure fixed at one ad per item and instead optimize which quality level of content is served and how items are matched with advertisers.

Prior literature in information systems has focused on *why* and *when* users gravitate toward ad-supported versus subscription-based access and how they transition from one to the other. For instance, Oestreicher-Singer and Zalmanson (2013) show that users who heavily use social network features of a platform are more inclined to use subscription-based access rather than ad-supported access. In turn, Bapna et al. (2018) and Oh et al. (2016) further characterize the behavior of consumers who switch to subscription-based access, finding that they further increase their social engagement with the platform and provide more word-of-mouth on the platform about its content. Wang et al. (2023) show that the specific design of the platform content can either promote or inhibit subscription model adoption. In the context of gaming platforms, they show that tailoring the game elements to consumer preferences increases their willingness to adopt the subscription model. Finally, Goli et al. (2024) study the effect of varying ad load on content consumption and subscriptions. They conduct a field experiment on the Pandora platform and find that increasing ad load to users of the ad-supported model decreases content consumption and increases the number of consumers adopting the subscription-supported model. Collectively, these findings point to a spectrum of content usage and ad tolerance rather than a strict free-versus-paid divide. Our GAS framework operationalizes that spectrum by posting a revenue-maximizing continuum of subscription fee and ad intensity bundles, allowing each user to self-select the option that best matches their preferences.

A complementary stream of literature has examined how DCPs adjust ad exposure, ad inventory allocation, subscription pricing, and other levers to improve their revenue. On the demand side, Lambrecht and Misra (2017) show that loosening paywalls in peak-interest periods can raise total earnings by stimulating both ad impressions and future subscriptions. Chiou and Tucker (2013) find that a hard paywall cuts news site visits by roughly half; this highlights the perils of aggressive access restrictions. Yoganarasimhan et al. (2023) demonstrate that shortening and appropriately personalizing free-trial length may improve long-run subscriptions. On the supply side, DCPs manage the pool of ads and its delivery by dynamically allocating impressions between long-term contracts and real-time markets (Balseiro et al. 2014, Shen et al. 2021). Other levers include pacing ad delivery to smoothen ad spend (Conitzer et al. 2022), partitioning ad impressions (Mehta et al. 2020), optimizing real-time bidding and reserve prices for ad auctions (Ostrovsky and Schwarz 2023), and targeting (Goldfarb and Tucker 2011), among others. Our GAS framework combines demand-side features such as ad exposure, the associated nuisance, and subscription fees with supply-side controls such as ad intensity and ad revenue rate to yield a revenue-maximizing mechanism for the DCP.

The mechanism design methodology used in our paper is based on a few classical sources. The seminal works of Maskin and Riley (1984), Myerson (1981), and Mussa and Rosen (1978) analyze how a monopolist should design an optimal menu of offerings when serving a market in which consumer preferences are private information. In these papers, the monopolist maximizes over a single source of revenue (i.e., posted prices) and the consumption of the users is exogenously given. In our GAS framework, the monopolist DCP *jointly* optimizes the revenue from subscription fees and advertising, with the latter determined endogenously by each user's viewing time. This dual-sided setting transforms the mechanism design problem into one of simultaneously screening users' preferences and shaping their usage incentives.

3. Model Preliminaries

We begin by describing the key elements of our model.

- **User Types:** Our model consists of a monopolist DCP and a continuum (unit mass) of users indexed by their type $\theta \in [\theta_L, \theta_H]$, which represents the users' level of interest in the content offered by the DCP. The DCP does not observe the type of any user but knows $F(\theta)$,

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

the probability distribution of types across the user base, and the corresponding density function $f(\theta)$. We assume that the distribution of user types belongs to the family of monotone-hazard-rate (MHR) distributions; that is, the hazard rate $\frac{f(\theta)}{(1-F(\theta))}$ is nondecreasing for all $\theta \in [\theta_L, \theta_H]$. Under MHR, the virtual value function, defined as $\psi(\theta) := \theta - \frac{1-F(\theta)}{f(\theta)}$, is nondecreasing for all $\theta \in [\theta_L, \theta_H]$. This is a common regularity assumption in the mechanism design literature (Krishna 2009) and is satisfied by commonly used distributions such as the uniform, exponential, truncated normal, among others.

- **User Participation:** A user's interaction with the DCP proceeds in two steps. First, the user makes a binary participation decision: whether to access the content of the DCP or not. We denote this by the variable $d \in \{0,1\}$. If the user does not participate ($d = 0$), then there is no interaction between the DCP and that user. However, if the user decides to access the content ($d = 1$), then in the second step, the user decides to consume the content offered by the DCP.
- **User Consumption:** For a user who participates, let $t \geq 0$ denote the per-day average usage time of the service (i.e., the amount of time spent on the DCP's platform) by a user. Further, let $n \geq 0$ be the ad intensity (i.e., the number of ads displayed to the user per unit usage time) chosen by the DCP. Because users have a negative perception of the ads (Aseri et al. 2020), we let $\alpha \geq 0$ denote the ad nuisance cost per unit ad intensity per unit usage time. Finally, let $T > 0$ denote the subscription period in days (e.g., a monthly subscription service corresponds to $T = 30$).
- **User Utility:** For a user who does not participate ($d = 0$), we normalize her utility to zero. For a participating user ($d = 1$) of type θ, the utility from consuming t units of content per day for a period of T days with ad intensity n is given by

$$T \cdot \theta \cdot \overbrace{\left(\underbrace{t - \frac{t^2}{2k}}_{\text{usage utility}} - \underbrace{\alpha \cdot n \cdot t}_{\text{ad-nuisance disutility}} \right)}^{\text{net usage benefit}} - p, \tag{1}$$

where $k > 0$ is a scaling parameter. Equation (1) captures the key factors affecting a user's utility while consuming content. The usage utility $(t - \frac{t^2}{2k})$ quantifies the benefit a user obtains from consuming t units of content. The concave form captures diminishing marginal returns to the users from consuming content (Xu et al. 2019). Next, the user incurs ad nuisance disutility $(\alpha n t)$ from viewing ads, which is proportional to ad-intensity n and the usage time t. Collectively, the net usage benefit from viewing t units of content is $(t - \frac{t^2}{2k} - \alpha n t)$. The users are heterogeneous in their valuation of this net usage benefit. Specifically, each user's type θ scales the valuation of this benefit. Finally, p denotes the subscription fee paid by that user for a period of T days.

- **DCP Revenue:** The DCP has two potential sources of revenue: (i) subscription fees (p) from users and (ii) ad revenue from the advertisers. For the latter component, let $r \geq 0$ denote the average revenue per ad impression. Thus, if a participating user consumes t units of content per day at ad-intensity n for a period of T days, then the ad revenue to the DCP from that user is $T n r t$. With this setup in place, our goal in this paper is to study four distinct classes of mechanisms that the DCP can employ to leverage these two revenue levers: (i) ADS, a single ad-supported tier with fixed ad intensity and no subscription fee; (ii) SUB, a single subscription-only tier with no ads; (iii) BAS, a binary menu that lets the user select between (a) a free, ad-supported tier and (b) a paid, ad-free tier; and finally (iv) GAS, a flexible menu of ad intensities and subscription fees from which users can select an option of their choice. For each of these four classes, we will obtain a revenue-maximizing mechanism for the DCP.

We now introduce the mechanism design framework for analyzing these four mechanisms.

3.1. Mechanism Design Setup for the DCP

Using the revelation principle (Myerson 1981), we restrict our attention, without loss of generality, to *direct* mechanisms that are incentive-compatible and individually rational for the users. That is, mechanisms in which the users truthfully reveal their type (θ) to the DCP constitute a Bayesian Nash equilibrium and the users obtain a nonnegative payoff upon participating. Let $\mu \in \{\text{ADS}, \text{SUB}, \text{BAS}, \text{GAS}\}$ denote the four classes of direct mechanisms described above.

A direct mechanism μ is characterized by a pair of functions $(n_\mu(\theta), p_\mu(\theta))$, where $n_\mu : [\theta_L, \theta_H] \to [0, \infty)$ is the ad-intensity function and $p_\mu : [\theta_L, \theta_H] \to [0, \infty)$ is the subscription fee function. Thus, in a direct mechanism μ, if a user of type θ participates and reveals her type as $\hat{\theta}$, then that user receives and ad-intensity of $n_\mu(\hat{\theta})$ while consuming the content and pays $p_\mu(\hat{\theta})$ as the subscription fee to the DCP. Let π_μ denote the total expected revenue to the DCP under mechanism μ.

The sequence of events proceeds as follows:

1. The DCP publicly announces mechanism μ, that is, the functions (n_μ, p_μ).
2. A user of type θ may reveal[1] her type as $\hat{\theta}$.
3. The user either participates and accesses the content by receiving an ad intensity of $n_\mu(\hat{\theta})$ and paying a subscription fee of $p_\mu(\hat{\theta})$ or decides not to participate and walks away.

Note that under mechanism μ, the utility (described in Equation (1)) to a user of type θ who reveals her type as $\hat{\theta}$ and consumes t units of content per day for a

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

period of T days is

$$T\theta\left(t-\frac{t^2}{2k}-\alpha n_\mu(\hat{\theta})t\right)-p_\mu(\hat{\theta}). \tag{2}$$

The optimal usage time for this user under mechanism μ is

$$\begin{aligned} t^*_\mu(\hat{\theta};\theta) &= \arg\max_{t\geq 0}\left(T\theta\left(t-\frac{t^2}{2k}-\alpha n_\mu(\hat{\theta})t\right)-p_\mu(\hat{\theta})\right)\\ &= k\left(1-\alpha n_\mu(\hat{\theta})\right)_+, \end{aligned} \tag{3}$$

where the shorthand $(x)_+$ is defined as $\max\{x,0\}$. As expected, users consume less content as the ad intensity and the ad nuisance cost increase.[2] If the ad intensity $n_\mu(\hat{\theta})$ exceeds $1/\alpha$, then the user does not consume any content. Let $V_\mu(\hat{\theta};\theta)$ denote the net utility of a user of type θ who reveals her type as $\hat{\theta}$ upon consuming the optimal level of content. Then, substituting the optimal usage time $t^*_\mu(\hat{\theta};\theta)$ in the user's utility function gives us

$$V_\mu(\hat{\theta};\theta)=\frac{1}{2}T\theta k\left(\left(1-\alpha n_\mu(\hat{\theta})\right)_+\right)^2-p_\mu(\hat{\theta}). \tag{4}$$

A user of true type θ first decides whether to participate in mechanism μ and, conditional on participating, chooses a report $\hat{\theta}$ to maximize her utility. Because the outside option is normalized to zero, user θ participates if and only if that user can obtain nonnegative utility from some feasible report, that is, $\max_{\hat{\theta}} V_\mu(\hat{\theta};\theta)\geq 0$. Conditional on submitting a report $\hat{\theta}$, the participation decision is therefore given by

$$d_\mu(\hat{\theta};\theta)=\begin{cases}1 & \text{if } V_\mu(\hat{\theta};\theta)\geq 0,\\ 0 & \text{otherwise.}\end{cases} \tag{5}$$

Before proceeding with the analysis, we summarize the notation introduced in our model in the Online Appendix.

4. Optimal Mechanisms

In this section, we formulate and derive the optimal mechanisms for each of the four classes of mechanisms. Each mechanism μ is characterized by a distinct set of decision variables: n_{ADS} under ADS, p_{SUB} under SUB, the pair $(n_{\text{BAS}}, p_{\text{BAS}})$ under BAS, and the type-dependent functions $(n_{\text{GAS}}(\cdot), p_{\text{GAS}}(\cdot), d_{\text{GAS}}(\cdot))$ under GAS. For ease of reference, the notation introduced in the model is summarized in the Online Appendix.

4.1. ADS Mechanism

Under the ad-supported mechanism (μ = ADS), the DCP monetizes users solely through ads without charging any subscription fee. Formally, for all participating users, an ADS mechanism is characterized as $(n_{\text{ADS}}(\hat{\theta}), p_{\text{ADS}}(\hat{\theta})) = (n_{\text{ADS}}, 0)$ for all $\hat{\theta}\in[\theta_{\text{L}},\theta_{\text{H}}]$, where $n_{\text{ADS}}\geq 0$ is the fixed ad intensity for all participating users. Because all participating users face the same ad intensity and pay no subscription fee, irrespective of their revealed type, this mechanism trivially satisfies incentive compatibility.

Using (3) and (4), the optimal usage time and net utility of a participating user under the ADS mechanism are

$$\begin{aligned} t^*_{\text{ADS}}(\theta) &\equiv t^*_{\text{ADS}}(\hat{\theta};\theta) = k(1-\alpha\, n_{\text{ADS}})_+,\\ V_{\text{ADS}}(\theta) &\equiv V_{\text{ADS}}(\hat{\theta};\theta) = \frac{1}{2}T\theta k\left((1-\alpha\, n_{\text{ADS}})_+\right)^2. \end{aligned}$$

We note that the usage time decreases linearly with ad intensity n_{ADS}, whereas net utility declines quadratically: This highlights users' aversion to ads. Next, using (5), it is straightforward to see that

$$d_{\text{ADS}}(\theta)\equiv d_{\text{ADS}}(\hat{\theta};\theta)=\begin{cases}1 & n_{\text{ADS}}\leq 1/\alpha,\\ 0 & n_{\text{ADS}}>1/\alpha.\end{cases} \tag{6}$$

Thus, either all the users participate when the ad intensity is low ($n_{\text{ADS}}\leq 1/\alpha$) or nobody participates when it is high ($n_{\text{ADS}}>1/\alpha$). The expected revenue to the DCP by setting an ad intensity of n_{ADS} is

$$\begin{aligned} \pi_{\text{ADS}} &= \int_{\theta_{\text{L}}}^{\theta_{\text{H}}}(T\, n_{\text{ADS}}\, r\, t^*_{\text{ADS}}(\theta))d_{\text{ADS}}(\theta)f(\theta)d\theta\\ &= T\, n_{\text{ADS}}\, r\, k(1-\alpha n_{\text{ADS}})_+. \end{aligned} \tag{7}$$

Thus, the revenue-maximizing ad intensity, n^*_{ADS}, is

$$n^*_{\text{ADS}}=\arg\max_{n_{\text{ADS}}\geq 0}\ T\, n_{\text{ADS}}\, r\, k(1-\alpha n_{\text{ADS}})_+=\frac{1}{2\alpha}. \tag{8}$$

Because $n^*_{\text{ADS}}\leq\frac{1}{\alpha}$, we get $d^*_{\text{ADS}}(\theta)=1$ for all θ. Thus, all the users participate under the optimal ADS mechanism. The optimal ADS* mechanism prioritizes market coverage over premium pricing by monetizing users through ads rather than subscriptions. The optimal ad intensity n^*_{ADS} balances marginal ad revenue (rTk) against the marginal loss from reduced usage time ($-2\alpha rTkn$). This reflects a classic monopoly tradeoff: Higher ad intensity increases revenue per user but reduces usage time, limiting total revenue. The mechanism is effective when ad nuisance costs (α) are low or ad revenue rates (r) are high: This enables broad user participation without excessive attrition.

4.2. SUB Mechanism

Under the subscription (μ = SUB) mechanism, the DCP forgoes ad revenue and relies only on subscription fees and offers users an ad-free experience. As all participating users can access the content for a fixed subscription fee without ads, the SUB mechanism is characterized as $(n_{\text{SUB}}(\hat{\theta}), p_{\text{SUB}}(\hat{\theta})) = (0, p_{\text{SUB}})$ for all $\hat{\theta}\in[\theta_{\text{L}},\theta_{\text{H}}]$, where $p_{\text{SUB}}\geq 0$ is the fixed subscription fee for all participating users. Because these users receive identical ad-free service at the same subscription fee irrespective of their revealed type, it follows that the SUB mechanism is trivially incentive compatible.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

For a user of type θ, the optimal usage time and the net utility upon participation under the SUB mechanism are

$$t^*_{\mathrm{SUB}}(\theta) \equiv t^*_{\mathrm{SUB}}(\hat{\theta};\theta) = k$$
$$V_{\mathrm{SUB}}(\theta) \equiv V_{\mathrm{SUB}}(\hat{\theta};\theta) = \frac{T\theta k}{2} - p_{\mathrm{SUB}}.$$

Let $\theta_{\mathrm{S}} := \min\{\theta_{\mathrm{H}}, \frac{2p_{\mathrm{SUB}}}{Tk}\}$ denote the marginal user type, for whom $V_{\mathrm{SUB}}(\theta) = 0$. Then, the participation decision for a user of type θ is

$$d_{\mathrm{SUB}}(\theta) \equiv d_{\mathrm{SUB}}(\hat{\theta};\theta) = \begin{cases} 1 & \theta \geq \theta_{\mathrm{S}}, \\ 0 & \theta < \theta_{\mathrm{S}}. \end{cases} \tag{9}$$

Thus, the SUB mechanism excludes low-type users ($\theta < \theta_{\mathrm{S}}$) from participation. This helps the DCP charge a higher fee to the remaining high-type users. The expected revenue to the DCP under this mechanism is

$$\pi_{\mathrm{SUB}} = \int_{\theta_{\mathrm{S}}}^{\theta_{\mathrm{H}}} p_{\mathrm{SUB}} f(\theta) d\theta = p_{\mathrm{SUB}} \cdot \left(1 - F(\theta_{\mathrm{S}})\right)$$
$$= p_{\mathrm{SUB}} \cdot \left(1 - F\left(\frac{2p_{\mathrm{SUB}}}{Tk}\right)\right).$$

The DCP's problem of obtaining the optimal subscription fee is simply

$$p^*_{\mathrm{SUB}} = \underset{p_{\mathrm{SUB}} \geq 0}{\arg\max}\ p_{\mathrm{SUB}} \cdot \left(1 - F\left(\frac{2p_{\mathrm{SUB}}}{Tk}\right)\right). \tag{10}$$

The closed-form expression for p^*_{SUB} depends on the distribution $F(\cdot)$. Based on the optimal subscription fee p^*_{SUB}, we can obtain the marginal user type $\theta^*_{\mathrm{S}} = \min\{\theta_{\mathrm{H}}, \frac{2p^*_{\mathrm{SUB}}}{Tk}\}$.

The optimal SUB* mechanism prioritizes revenue from high-type users, who have a higher willingness-to-pay for ad-free experience while excluding low-type users with low willingness-to-pay. By sacrificing ad revenue and setting a subscription fee, the mechanism creates a tradeoff between premium pricing and market coverage. The optimal fee p^*_{SUB} balances the marginal gain from a higher price against the marginal loss from reduced demand. Naturally, the optimal SUB* mechanism is effective when user types are skewed toward high valuations, the ad nuisance parameter α is high, or the ad revenue rate r is low relative to the subscription fee.

4.3. BAS Mechanism

Although the ADS and SUB mechanisms are widely popular, DCPs have recently started to combine the two revenue models and let users select their preferred model. In particular, the DCP offers the following two-tiered menu to the users:

- **Ad-supported tier:** No subscription fee, but the user faces a fixed ad-intensity.
- **Subscription tier:** A fixed subscription fee, and no ads are displayed to the user.

We refer to this mechanism as the binary ad-supported subscription (μ = BAS) mechanism. Formally, a BAS mechanism is characterized as

$$\left(n_{\mathrm{BAS}}(\hat{\theta}), p_{\mathrm{BAS}}(\hat{\theta})\right) = \begin{cases} (n_{\mathrm{BAS}}, 0) & \text{if user selects ad-supported tier,} \\ (0, p_{\mathrm{BAS}}) & \text{if user selects subscription tier,} \end{cases} \tag{11}$$

where $n_{\mathrm{BAS}} \geq 0$ is the fixed ad intensity in the ad-supported tier and $p_{\mathrm{BAS}} \geq 0$ is the subscription fee in the subscription tier.

For a participating user of type θ who reveals her type as $\hat{\theta}$, the utility from selecting the ad-supported tier is $(\frac{T\theta k}{2}((1-\alpha n_{\mathrm{BAS}})_+)^2)$, whereas that from selecting the subscription tier is $\left(\frac{T\theta k}{2} - p_{\mathrm{BAS}}\right)$. Because the utility to the user in either case is independent of the revealed type $\hat{\theta}$, any BAS mechanism trivially satisfies incentive compatibility. Next, a type θ user who participates will select the subscription tier over ad-supported tier if

$$\left(\frac{T\theta k}{2} - p_{\mathrm{BAS}}\right) - \left(\frac{T\theta k}{2}((1-\alpha n_{\mathrm{BAS}})_+)^2\right) \geq 0$$
$$\iff \frac{T\theta k}{2}\left(1 - ((1-\alpha n_{\mathrm{BAS}})_+)^2\right) - p_{\mathrm{BAS}} \geq 0.$$

Define $\theta_{\mathrm{B}} := \min\left\{\theta_{\mathrm{H}}, \frac{2p_{\mathrm{BAS}}}{Tk(1-((1-\alpha n_{\mathrm{BAS}})_+)^2)}\right\}$ as the indifferent user type. Thus, a participating user of type θ opts for a subscription tier if $\theta > \theta_{\mathrm{B}}$ and for ad-supported tier if $\theta \leq \theta_{\mathrm{B}}$. A type θ user participates if they obtain a nonnegative utility from selecting either of the two options in the menu. Thus,

$$d_{\mathrm{BAS}}(\theta) \equiv d_{\mathrm{BAS}}(\hat{\theta}, \theta) = \begin{cases} 1 & \text{if } \max\left\{\frac{T\theta k}{2} - p_{\mathrm{BAS}}, \frac{T\theta k}{2}((1-\alpha n_{\mathrm{BAS}})_+)^2\right\} \geq 0, \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

Here, it is instructive to distinguish between low and high ad-intensity regimes, because user participation flips when n_{BAS} crosses the threshold $1/\alpha$:

- Low Ad Intensity ($n_{\mathrm{BAS}} \leq 1/\alpha$): In this range, the utility from the ad-supported option is $\frac{T\theta k}{2}((1-\alpha n_{\mathrm{BAS}})_+)^2 \geq 0$ for all user types. Consequently, every user finds at least one of the two tiers attractive and the DCP attains full market coverage.
- High Ad Intensity ($n_{\mathrm{BAS}} > 1/\alpha$): In this case, no user obtains a positive utility from opting for the ad-supported tier. Thus, a user will select the subscription tier only if it yields a nonnegative utility to her; that is, if $\frac{T\theta k}{2} - p_{\mathrm{BAS}} \geq 0$. Notice that in this case, the threshold θ_{B} takes the value $\min\left\{\theta_{\mathrm{H}}, \frac{2p_{\mathrm{BAS}}}{Tk}\right\}$ with the interpretation that when ($n_{\mathrm{BAS}} > 1/\alpha$), users with $\theta \leq \theta_{\mathrm{B}}$ do not participate, whereas all the users with $\theta > \theta_{\mathrm{B}}$ participate and opt for the subscription tier.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

The optimal usage time of a participating user is simply:

$$t^*_{\mathrm{BAS}}(\theta) \equiv t^*_{\mathrm{BAS}}(\hat{\theta};\theta) = \begin{cases} k(1-\alpha n_{\mathrm{BAS}})_+ & \text{if } \theta \le \theta_{\mathrm{B}} \quad \text{(ad-supported tier)},\\ k & \text{if } \theta > \theta_{\mathrm{B}} \quad \text{(subscription tier)}. \end{cases}$$

This implies that when the ad intensity is high ($n_{\mathrm{BAS}} > 1/\alpha$), the usage time of users who choose the ad-supported option drops to zero, and consequently, the DCP earns no ad revenue from this group. Thus, without loss of generality, the DCP can restrict the ad intensity to $0 \le n_{\mathrm{BAS}} \le 1/\alpha$.

Using these components, the revenue-maximization problem of the DCP can be written as

$$\pi^*_{\mathrm{BAS}} = \max_{0 \le n_{\mathrm{BAS}} \le 1/\alpha,\ p_{\mathrm{BAS}} \ge 0} \underbrace{\left(T k r n_{\mathrm{BAS}}(1-\alpha n_{\mathrm{BAS}})\right) \cdot \int_{\theta_{\mathrm{L}}}^{\theta_{\mathrm{B}}} f(\theta)d\theta}_{\text{revenue from ad-supported tier}} + \underbrace{p_{\mathrm{BAS}} \cdot \int_{\theta_{\mathrm{B}}}^{\theta_{\mathrm{H}}} f(\theta)d\theta}_{\text{revenue from subscription tier}} . \tag{13}$$

The revenue-maximization problem stated in (13) involves balancing subscription revenue and advertising revenue for the DCP. Note that the decision variables, ad intensity n_{BAS} and subscription fee p_{BAS}, are linked through the indifference threshold θ_{B}. Ceteris paribus, increasing subscription fee p_{BAS} raises the threshold θ_{B}, which leads to a reduction in the set of users selecting the subscription tier, albeit collecting higher subscription fee from each of them. Here, the DCP trades off a smaller paying user base against a higher per-subscriber margin. On the other hand, ceteris paribus, increasing the ad intensity n_{BAS} lowers the threshold θ_{B}, which makes the ad-supported tier less attractive to the users. Here, the DCP trades off higher per-user ad revenue against a smaller user base who select the ad-supported tier.

The objective function in (13) is analytically challenging because the threshold θ_{B} depends nonlinearly on both n_{BAS} and p_{BAS}. Consequently, we employ a numerical grid search approach—systematically exploring a range of feasible $(n_{\mathrm{BAS}}, p_{\mathrm{BAS}})$ pairs—to obtain the optimal subscription fee and ad intensity that maximizes total expected revenue.

4.4. GAS Mechanism

We now turn our attention to the generalized ad-supported subscription ($\mu = \mathrm{GAS}$) mechanism that allows the DCP to offer a more granular menu of ad intensities and corresponding subscription fees to the users (compared with the three mechanisms discussed above). By doing so, the GAS mechanism encompasses ADS, SUB, and BAS mechanisms as special cases. More importantly, the optimal GAS mechanism achieves the maximum possible revenue under private information: classically known as the second-best outcome when first-degree price discrimination is ruled out.

Formally, a GAS mechanism is characterized by the pair of functions $(n_{\mathrm{GAS}}(\cdot), p_{\mathrm{GAS}}(\cdot))$. Thus, if a participating user of type θ reveals her type as $\hat{\theta}$, then that user receives an intensity of $n_{\mathrm{GAS}}(\hat{\theta})$ and pays a subscription fee of $p_{\mathrm{GAS}}(\hat{\theta})$ to the platform. Using (3), (4), and (5) from Section 3.1, the optimal usage time, net utility of a participating user, and the participation decision of a type θ user who reveals her type as $\hat{\theta}$ under the GAS mechanism are

$$t^*_{\mathrm{GAS}}(\hat{\theta};\theta) = k(1-\alpha n_{\mathrm{GAS}}(\hat{\theta}))_+,$$

$$V_{\mathrm{GAS}}(\hat{\theta};\theta) = \frac{1}{2}T\theta k((1-\alpha n_{\mathrm{GAS}}(\hat{\theta}))_+)^2 - p_{\mathrm{GAS}}(\hat{\theta}),$$

$$d_{\mathrm{GAS}}(\hat{\theta};\theta) = \begin{cases} 1 & \text{if } V_{\mathrm{GAS}}(\hat{\theta};\theta) \ge 0,\\ 0 & \text{otherwise.} \end{cases}$$

Notice that the participation decision of a user is completely determined by the functions $(n_{\mathrm{GAS}}(\cdot), p_{\mathrm{GAS}}(\cdot))$ announced by the DCP. Hence, without loss of generality, we can describe any incentive-compatible GAS mechanism by the triplet $(n_{\mathrm{GAS}}(\cdot), p_{\mathrm{GAS}}(\cdot), d_{\mathrm{GAS}}(\cdot))$, subject to the condition that

$$d_{\mathrm{GAS}}(\theta) := d_{\mathrm{GAS}}(\theta;\theta) = \begin{cases} 1 & \text{if } V_{\mathrm{GAS}}(\theta;\theta) \ge 0,\\ 0 & \text{otherwise} \end{cases} \qquad \forall \theta \in [\theta_{\mathrm{L}}, \theta_{\mathrm{H}}],$$

where $V_{\mathrm{GAS}}(\theta;\theta) = \frac{1}{2}T\theta k((1-\alpha n_{\mathrm{GAS}}(\theta))_+)^2 - p_{\mathrm{GAS}}(\theta)$ is the utility of type θ user under an incentive-compatible mechanism. This is a standard simplification used for analytical convenience in the mechanism design literature (Krishna 2009). Using these elements, the revenue-maximization mechanism design problem of the DCP can be formulated as follows:

$$\pi^*_{\mathrm{GAS}} = \max_{\substack{n_{\mathrm{GAS}}(\cdot) \ge 0\\ p_{\mathrm{GAS}}(\cdot) \ge 0\\ d_{\mathrm{GAS}}(\cdot) \in \{0,1\}}} \int_{\theta_{\mathrm{L}}}^{\theta_{\mathrm{H}}} \left(\underbrace{p_{\mathrm{GAS}}(\theta)}_{\text{subscription revenue}} + \underbrace{T r n_{\mathrm{GAS}}(\theta) t^*_{\mathrm{GAS}}(\theta)}_{\text{ad revenue}} \right) d_{\mathrm{GAS}}(\theta) f(\theta) d\theta \qquad (\mathcal{P}\text{–}GAS),$$

subject to :

$$V_{\mathrm{GAS}}(\theta;\theta) \ge V_{\mathrm{GAS}}(\hat{\theta};\theta) \quad \forall \theta \in [\theta_{\mathrm{L}}, \theta_{\mathrm{H}}] \qquad \text{(incentive compatibility)}$$

$$d_{\mathrm{GAS}}(\theta) = \begin{cases} 1 & \text{if } V_{\mathrm{GAS}}(\theta;\theta) \ge 0,\\ 0 & \text{otherwise} \end{cases} \qquad \forall \theta \in [\theta_{\mathrm{L}}, \theta_{\mathrm{H}}]. \qquad \text{(individual rationality)}$$

In the mechanism design problem $\mathcal{P}\text{–}GAS$, the DCP simultaneously chooses an ad-intensity function $n_{\mathrm{GAS}}(\cdot)$,

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

a subscription fee function $p_{\mathrm{GAS}}(\cdot)$, and a participation indicator function $d_{\mathrm{GAS}}(\cdot)$ to maximize the total expected revenue from the two sources across all the user types. The incentive compatibility constraints state that it is optimal for the users to reveal their type truthfully to the DCP. The individual rationality constraints state that the participating users receive a nonnegative payoff upon truthfully revealing their type. Because participation is voluntary and the outside option is normalized to zero, incentive compatibility only needs to bind against deviations that yield nonnegative utility. Accordingly, writing $V_{\mathrm{GAS}}(\theta;\theta) \geq V_{\mathrm{GAS}}(\hat{\theta};\theta)$ for all $\hat{\theta}$ is without loss of generality. The following result states an optimal solution to ($\mathcal{P}-GAS$).

Theorem 1 (Optimal GAS Mechanism). *Let the virtual-value function be defined by* $\psi(\theta) = \theta - \frac{1-F(\theta)}{f(\theta)}$. *Then, the mechanism GAS* defined by*

$$n^*_{\mathrm{GAS}}(\theta) = \begin{cases} \dfrac{r-\alpha\psi(\theta)}{\alpha\left(2r-\alpha\psi(\theta)\right)} & \text{if } \theta < \psi^{-1}\left(\dfrac{r}{\alpha}\right) \\ 0 & \text{if } \theta \geq \psi^{-1}\left(\dfrac{r}{\alpha}\right), \end{cases} \tag{14}$$

$$p^*_{\mathrm{GAS}}(\theta) = \frac{\theta Tk\left(1-\alpha n^*_{\mathrm{GAS}}(\theta)\right)^2}{2} - \int_{\theta_{\mathrm{L}}}^{\theta} \frac{Tk\left(1-\alpha n^*_{\mathrm{GAS}}(x)\right)^2}{2} dx$$

$$d^*_{\mathrm{GAS}}(\theta) = 1 \quad \forall \theta \in [\theta_{\mathrm{L}}, \theta_{\mathrm{H}}]. \tag{15}$$

is an optimal solution to Problem ($\mathcal{P}-GAS$).

The proofs and analysis of all technical results are provided in the Online Appendix. The GAS* mechanism in Theorem 1 characterizes a revenue-maximizing mechanism that balances subscription fees and ad revenue across a heterogeneous user base. First, we note that the GAS* mechanism covers the entire market ($d^*_{\mathrm{GAS}}(\theta) = 1$ $\forall\theta$); that is, all the users participate and obtain nonnegative payoff. The optimal ad intensity function $n^*_{\mathrm{GAS}}(\theta)$ is nonincreasing in user type θ, whereas the optimal subscription fee function $p^*_{\mathrm{GAS}}(\theta)$ is nondecreasing in θ (see the left panel of Figure 2 for illustration). Specifically, high-type users ($\theta \geq \psi^{-1}(r/\alpha)$) receive no ads ($n^*_{\mathrm{GAS}}(\theta) = 0$) but pay a higher subscription fee to the DCP, whereas low-type users ($\theta < \psi^{-1}(r/\alpha)$) are exposed to a positive ad intensity ($n^*_{\mathrm{GAS}}(\theta) > 0$) and pay a lower subscription fee. By aligning ad intensity and subscription fees with user preferences, the GAS* mechanism effectively screens users by extracting higher rents from high-type, ad-sensitive users and monetizing low-type, ad-tolerant users through ads.

Note that the revenue-maximizing GAS* mechanism derived above is a *direct* mechanism in which the users reveal their private information (θ) to the DCP. The outcomes—ad intensity, subscription fee, and participation decision ($n^*_{\mathrm{GAS}}(\theta), p^*_{\mathrm{GAS}}(\theta), d^*_{\mathrm{GAS}}(\theta)$)—are then determined solely based on the revealed user types. From a practical implementation viewpoint, it is often challenging to obtain truthful revelation of private information. Instead, a *posted-price* mechanism, in which the DCP publicly posts a menu of options and the users self-select the option that best matches their preferences, offers a simpler, more transparent alternative. The result below shows that the GAS* mechanism characterized in Theorem 1 can be implemented as a posted-price mechanism.

Proposition 1 (GAS* as a Posted-Price Mechanism). *The GAS* mechanism can be implemented as a posted-price mechanism as follows. Define the inverse mapping from ad-intensity to type* $\theta(n)$ *as*

$$\theta(n) = \psi^{-1}\left(\frac{r(1-2\alpha n)}{\alpha(1-\alpha n)}\right),$$

where $\psi(\theta) = \theta - \frac{1-F(\theta)}{f(\theta)}$ *is the virtual-value function. Then, the menu of ad-intensity and subscription price*

Figure 2. (Color online) Optimal GAS* Mechanism

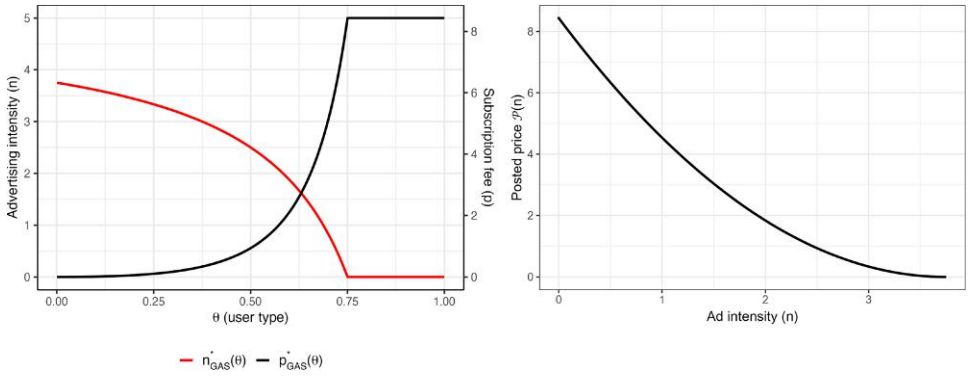

Notes. (Left) Illustration of GAS* mechanism: ad intensity $n^*_{\mathrm{GAS}}(\theta)$ (red) and subscription fee $p^*_{\mathrm{GAS}}(\theta)$ (black) as a function of user type θ (here, $\theta \sim Beta(1,1)$ on support [0,1]). (Right) A posted-price implementation of the GAS* mechanism. The parameter values underlying both panels are as specified in Table 1.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

$\{(n, \mathcal{P}(n)) : n \in [0, \frac{1}{\alpha}]\}$, *where*

$$\mathcal{P}(n) = \frac{Tk\theta(n)(1-\alpha n)^2}{2} - \int_{\theta_L}^{\theta(n)} \frac{Tk\left(1-\alpha n^*_{\mathrm{GAS}}(x)\right)^2}{2} dx,$$

is an optimal posted-price mechanism in that it yields the same outcomes as the GAS* *mechanism.*

Proposition 1 (illustrated in the right panel of Figure 2) bridges a critical gap between the theoretical design of an optimal mechanism and its practical implementation. In many settings, especially in digital markets, it is challenging for firms to elicit truthful private information from users. The direct GAS* mechanism characterized in Theorem 1, whereas optimal in theory, requires users to reveal their private types: a process that can be both complex and prone to misreporting. The posted-price mechanism, by contrast, offers a transparent menu from which users self-select based on their preferences, without ever having to disclose private information explicitly. This result is nontrivial: In multidimensional screening problems, implementability is generally complex, and optimal mechanisms need not admit a simple posted-price representation without careful construction of the menu (Armstrong 1996, Rochet and Choné 1998). By showing that the GAS* mechanism can be equivalently implemented as a posted-price mechanism, the result provides a powerful insight: Even in the presence of information asymmetry, the DCP can obtain optimal revenue using a simple and practical posted-price mechanism.

Remark 1 (Implementation). In practice, the posted-price menu in Proposition 1 can be operationalized through an interactive interface (e.g., sliders): The user selects a desired ad intensity level n along the slider, and the corresponding subscription fee $\mathcal{P}(n)$ is displayed dynamically (as illustrated schematically in the GAS panel of Figure 1). Several digital platforms already employ such continuous posted-price menus. For instance, Amplitude[3] (a digital analytics provider) lets customers set their subscription price via a slider that trades off monthly price against the number of users served. Similarly, Zapier[4] (an artificial intelligence (AI) workflow platform) uses a slider to let customers choose the number of tasks in their plan, with the price adjusting accordingly. We note, however, that a slider-based menu still presents users with infinitely many options; this may increase decision effort and lead to user fatigue or suboptimal choices. The finite-menu approximation in Corollary 1 (Section 6.3) alleviates this concern: the DCP can offer a finite-sized menu of length at most L, and the resulting revenue loss decreases proportionally to $L^{-1/4}$ as L increases.

It is instructive to note that the GAS mechanism provides a unified framework for studying revenue models that leverage both advertising and subscription fees. By appropriately choosing its ad-intensity and subscription price functions, the GAS mechanism can be specialized to yield the ADS, SUB, and BAS mechanisms. As a consequence, the optimal GAS* mechanism better price discriminates and yields weakly higher revenue compared with the optimal ADS*, SUB*, and BAS* mechanisms.

5. Comparative Analysis of Mechanisms

In this section, we compare the revenue of the DCP under the four optimal mechanisms studied in the previous section, ADS*, SUB*, BAS*, and GAS*, by analyzing how it varies with key parameters: the ad nuisance rate (α), the ad revenue rate (r), the variance in the distribution of user types (θ), and the baseline usage (k). Table 1 summarizes the nominal values and ranges considered for each parameter.

We choose nominal values as a representative baseline and vary parameters over an economically meaningful range for sensitivity analysis. The range of the ad revenue rate r is based on empirical magnitudes typical in digital content settings,[5] and the range of ad nuisance reflects a variety of potential situations, from low ad nuisance to very strong ad nuisance. We fix $T = 30$ to reflect a standard monthly plan. We model user types as $\theta \sim \mathrm{Beta}(a, b)$ on support [0, 1], with density proportional to $\theta^{a-1}(1-\theta)^{b-1}$. The shape parameters a and b control the mean and dispersion of θ (and skewness when $a \neq b$). This makes the Beta family a flexible way to represent different shapes and degrees of heterogeneity. In particular, the symmetric distribution $\mathrm{Beta}(m, m)$ is uniform at $m = 1$ and becomes increasingly concentrated around 1/2 as m grows. We vary $a, b \in [1, 6]$ and note that all specifications considered satisfy the MHR condition required by our analysis.

5.1. Sensitivity Analysis

We briefly summarize the key sensitivity insights here.

- **Sensitivity to ad nuisance (α) and ad revenue rate (r):** As expected, SUB* is invariant to both α and r, whereas the revenue under ADS* declines sharply with α and increases linearly with r. BAS* and GAS*

Table 1. Parameter Values and Ranges Used in Sensitivity Analysis

Parameter	Nominal value	Range
α (ad nuisance rate)	0.25	(0, 1]
r (ad revenue rate)	0.1	(0, 0.2]
$\theta \sim \mathrm{Beta}(a, b)$ (distribution of user types)	$a = 1$	a:[1, 6]
	$b = 1$	b:[1, 6]
T (subscription period in days)	30	—
k (average content usage without ads)	1	(0, 5]

Note. All $\mathrm{Beta}(a, b)$ specifications considered in the table satisfy the MHR condition.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

better adjust to variation in α and r as they can partially offset losses in ad revenue by shifting high-type users toward subscription.

- **Sensitivity to user-type distribution (Beta(a,b)):** In the symmetric case $\theta \sim \text{Beta}(m,m)$, ADS* remains essentially flat because every user faces the same ad intensity, whereas SUB*, BAS*, and GAS* increase modestly as higher m concentrates mass toward moderate-to-high valuations (and reduces uncertainty about how many users will subscribe). In the asymmetric case $\theta \sim \text{Beta}(a,6)$, increasing a shifts the distribution from right-skewed to near-normal, which increases revenues under subscription-based mechanisms, whereas the incremental gain in ADS* is relatively small.
- **Sensitivity to baseline usage (k):** Expected revenues under all mechanisms increase with k because higher baseline usage raises total consumption and enables the DCP to extract greater surplus via subscriptions and/or ads. The level of k amplifies revenue uniformly and does not alter the qualitative ranking among mechanisms. The relative ranking of SUB* and ADS* depends on the ad nuisance environment: SUB* dominates ADS* when ad nuisance is high, whereas ADS* dominates when ad nuisance is low.
- **Sensitivity of optimal access options under GAS*:** Although the revenue under GAS* is relatively stable across parameter values, the underlying access options—subscription fees $p^*_{\text{GAS}}(\theta)$ and ad intensities $n^*_{\text{GAS}}(\theta)$—shift substantially. As the user-type distribution becomes more concentrated, the maximum ad intensity decreases and the cutoff type shifts to a lower θ. Further, as ad nuisance increases, both the level and range of ad exposure shrink.

Before examining the revenue gap between GAS* and BAS* mechanisms, we briefly comment on the total welfare.

Remark 2 (Total Welfare). The total welfare analysis compares the sum of DCP's revenue and consumer surplus under all four mechanisms. We observe that the GAS* mechanism, beyond offering a revenue advantage over the other three mechanisms, also generates a weakly larger total welfare across all parameter combinations.

5.2. Revenue Gap Between GAS* and BAS* Mechanisms

Although GAS* weakly dominates BAS* by construction, BAS* is a natural coarse approximation to GAS*: The two discrete options in BAS* closely align with the preferences of extreme types of users. The sensitivity analysis indicates that the revenue advantage of GAS* over BAS* may be modest in certain parametric regions. This raises a natural question: Under which conditions does GAS* deliver a quantitatively meaningful revenue improvement over BAS*? Because BAS* does not admit a closed-form solution, an analytical comparison is not feasible. Therefore, we study the revenue gap numerically over the parameter ranges in Table 1.

To identify the parameter regions associated with large revenue gaps, we use a classification tree–based approach. We find that a higher revenue gap between GAS* and BAS* is driven primarily by two features.

- **Ad attractiveness** (high r/α)**:** The ratio r/α captures the attractiveness of ads: A higher value of the ratio indicates that ads generate more revenue while causing lower nuisance to the users. When r/α is high, advertising becomes a stronger monetization lever compared with charging subscription fees. Accordingly, it becomes important for the DCP to determine how ad exposure should be allocated across user types. Because GAS* offers a continuum of ad intensities, it screens users more effectively than BAS* (which is limited to a coarse two-tier menu) and can therefore tailor ad intensity more closely to users' types. As a result, the revenue advantage of GAS* over BAS* increases with r/α.
- **Proportion of high-type users** (high a/b when $\theta \sim \text{Beta}(a,b)$)**:** A higher value of a and/or a lower value of b shifts the distribution of user types toward higher valuations. As the ratio a/b increases, revenue becomes more sensitive to how effectively the mechanism screens high-type users. BAS* relies on a single subscription fee, so when more users have high valuations it cannot effectively price discriminate among higher-type users. In contrast, GAS* offers a finer mix of ad intensities and subscription fees, which allows the DCP to better differentiate among the higher-type users. Consequently, the revenue advantage of GAS* over BAS* increases with a/b.

Importantly, these two drivers are complementary: Neither alone is sufficient to generate a large gap. The largest revenue gaps arise when both ad attractiveness and the proportion of high-type users are high. Figure 3 summarizes how the revenue difference between the two mechanisms varies jointly with ad attractiveness and the proportion of high-type users.

6. Model Extensions

In this section, we extend our baseline model in several directions. In particular, we (i) introduce an exogenous cap on the ad-intensity, (ii) endogenize the ad revenue rate based on the ad intensity, (iii) obtain an approximation bound when users are endowed with two-dimensional heterogeneity, (iv) consider a single-tier menu consisting of ad intensity and subscription fees, and (v) incorporate competition.

6.1. Capped Ad Intensity

Recall from Theorem 1 that in our baseline model, the ad intensity under the GAS* mechanism, namely $n^*_{\text{GAS}}(\cdot)$, is endogenously determined and can reach up

Figure 3. (Color online) Revenue Gap Between GAS* and BAS* Mechanisms

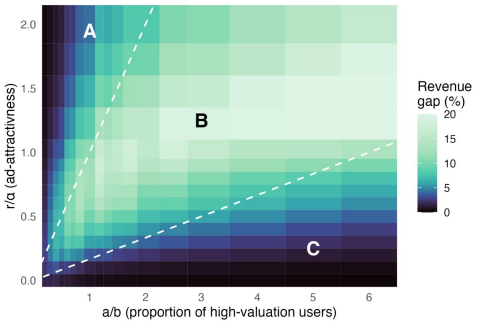

Region A ($\frac{r}{\alpha} > \frac{a}{b}$; gap ≈ 0–13%): Ad attractiveness is high enough for both mechanisms to push most users towards high ad intensities; this limits the scope for GAS* to screen better over BAS*.

Region B ($\frac{a}{6b} < \frac{r}{\alpha} \leq \frac{a}{b}$; gap ≈ 2–25%): Ads are effective and sufficient high-type users allow GAS* to better differentiate the users.

Region C ($\frac{r}{\alpha} \leq \frac{a}{6b}$; gap ≈ 0–7%): Ads are low-paying or impose substantial nuisance; neither mechanism generates substantial ad revenue.

Notes. (Left) Revenue gap (in %) as a joint function of ad attractiveness (r/α) and the user-type distribution (a/b). Each cell reports the average revenue gap. Lighter colors indicate a larger revenue gap. (Right) Three regions delineated by the dashed boundaries in the figure.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

to $1/\alpha$. When the ad nuisance parameter α is low (implying that users are not very sensitive to advertising), the GAS* mechanism incentivizes the DCP to allocate high ad intensity to the low-type users. Such high levels of advertising may run counter to emerging regulatory standards. For instance, both the European Union's *Digital Services Act*[6] and *Digital Fairness Act*[7] emphasize the responsibility of digital platforms to provide a safe, transparent, and user-centric online environment. This discourages exploitative or excessive advertising practices by DCPs.

To reflect these emerging regulatory constraints, we impose an upper bound, namely, $n_{\max}$, on the ad intensity that the DCP can allocate to any user type. In particular, we add the constraint $0 \leq n_{\text{GAS}}(\theta) \leq n_{\max}$ for all $\theta \in [\theta_{\text{L}}, \theta_{\text{H}}]$ to Problem ($\mathcal{P}$–*GAS*) and characterize the optimal mechanism under this additional feasibility condition. This constraint limits the extent to which the platform can rely on advertising. Theorem 2 presents the structure of the optimal GAS mechanism in this constrained setting.

Theorem 2 (Capped GAS Mechanism). *Define* $\theta_{\text{C}} := \max\{\theta_{\text{L}}, \psi^{-1}(\frac{-2rn_{\max}}{1-\alpha n_{\max}})\}$, *where* $\psi(\theta) = \theta - \frac{1-F(\theta)}{f(\theta)}$ *denotes the virtual-value function. Then, if the ad intensity is capped at* $n_{\max}$, *the mechanism GAS-C* defined as*

$$n^*_{\text{GAS–C}}(\theta) = \begin{cases} \min\left\{n_{\max}, \dfrac{r-\alpha\psi(\theta)}{\alpha\left(2r-\alpha\psi(\theta)\right)}\right\} & \text{if } \theta < \psi^{-1}\left(\dfrac{r}{\alpha}\right) \\ 0 & \text{if } \theta \geq \psi^{-1}\left(\dfrac{r}{\alpha}\right), \end{cases} \tag{16}$$

$$p^*_{\text{GAS–C}}(\theta) = \frac{\theta Tk\left(1-\alpha n^*_{\text{GAS–C}}(\theta)\right)^2}{2} - \int_{\theta_{\text{C}}}^{\theta} \frac{Tk\left(1-\alpha n^*_{\text{GAS–C}}(x)\right)^2}{2} dx, \tag{17}$$

$$d^*_{\text{GAS–C}}(\theta) = \begin{cases} 0 & \text{if } \theta < \theta_{\text{C}} \\ 1 & \text{if } \theta \geq \theta_{\text{C}}, \end{cases} \tag{18}$$

is an optimal GAS mechanism with capped ad intensity.

Theorem 2 formalizes how an exogenous cap $n_{\max}$ on ad intensity reshapes the GAS mechanism. The left panel in Figure 4 plots expected revenue to the DCP under GAS* and GAS – C* mechanisms as a function of the ad cap parameter $n_{\max}$. As expected, the revenue under the constrained GAS – C* mechanism (dashed curve) is weakly lower than under the unconstrained GAS* mechanism (solid curve). The vertical line marks the maximum unconstrained ad intensity, $n^*_{\text{GAS}}(\theta_{\text{L}})$, under the GAS* mechanism. When the cap is above this threshold ($n_{\max} > n^*_{\text{GAS}}(\theta_{\text{L}})$), the constraint is slack and the GAS – C* mechanism coincides with the GAS* mechanism. When the cap value is below the threshold, the constraint becomes binding, and the DCP experiences revenue loss due to the reduced flexibility in screening users through ads.

Interestingly, in the limiting case where $n_{\max} = 0$ (i.e., the case where DCP is entirely prohibited from showing ads and thus relies only on the subscription fee), the GAS – C* mechanism reduces to the SUB* mechanism. In contrast, the unconstrained GAS* mechanism never excludes any user type and therefore never converges to the SUB* mechanism, even when the users are highly ad sensitive. This shows that imposing a cap on ad intensity can drive the GAS – C* mechanism to converge to the SUB* mechanism, whereas no level of user ad nuisance by itself can compel the unconstrained GAS* mechanism to converge to the same outcome.

The right panel of Figure 4 contrasts the GAS* and GAS – C* mechanisms when the cap is fixed at $n_{\max} = 1$. The dashed vertical line marks the critical type θ_{C} for the GAS – C* mechanism: Only users with type $\theta \geq \theta_{\text{C}}$ participate. The cap on ad intensity has differential effect depending on user type, and we see three regions emerging.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

Figure 4. (Color online) Comparison of the GAS* and GAS − C* Mechanisms

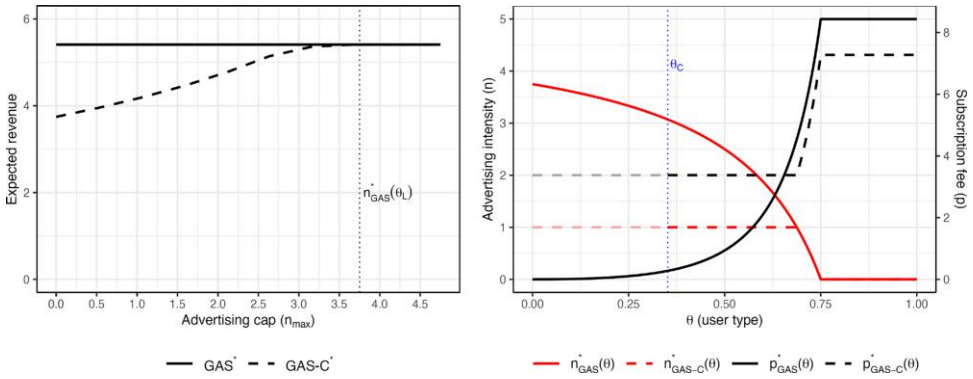

Notes. (Left) Expected revenue as a function of the advertiser cap parameter $n_{\max}$, for both GAS* (solid) and GAS − C* (dashed) mechanisms. The dotted vertical line gives the optimal $n^*_{\mathrm{GAS}}(\theta_{\mathrm{L}})$. (Right) Ad intensity (red) and subscription fee (black) across user types θ, for both GAS* (solid) and GAS − C* (dashed) mechanisms when $n_{\max} = 1$. The blue dashed line indicates the threshold user type at which users are indifferent between participating and opting out under the GAS − C* mechanism.

- **Low-Type Users:** For users with low valuation ($\theta < \theta_{\mathrm{C}}$), the cap prevents the DCP from displaying sufficiently high ad load. As a result, the DCP cannot offer those types a contract that is both individually rational and profitable; therefore, all types $\theta < \theta_{\mathrm{C}}$ rationally decline to participate.
- **Intermediate-Type Users:** For user types slightly higher than θ_{C}, the cap is binding, and the participating users receive the maximum ad intensity of $n_{\max}$. Because types with $\theta < \theta_{\mathrm{C}}$ do not participate, the incentive compatibility constraints for these types are lifted. This allows the DCP to raise the subscription fee for type θ_{C} up to the point where individual rationality binds. As a result, the subscription fee under the GAS − C* mechanism at this threshold exceeds that under the GAS* mechanism.
- **High-Type Users:** For sufficiently high user types, the cap on ad intensity no longer binds; therefore, GAS − C* and GAS* mechanisms yield the same ad intensity for these user types. However, the subscription fee under the GAS − C* mechanism now becomes lower than that under the GAS* mechanism. This is because the ad intensity cap removes the DCP's ability to threaten high ad intensity on lower types, due to which it loses leverage over high types. Consequently, the DCP compensates these high-type users with larger information rents, which translates to lower subscription fees compared with the GAS* mechanism.

6.2. Endogenous Ad Revenue Rate

In our base model, we assume that the ad revenue rate (r) is constant. However, a higher value of ad intensity (n) can lead to ad cluttering. As a consequence, excessive ads may erode user engagement, which in turn may diminish the value to the advertisers and lead to lower ad revenue rates. To capture this phenomenon, we endogenize the ad revenue rate as a decreasing function of ad intensity. Specifically, we let $r(n) = r_0 - \beta n$, where $r_0, \beta \geq 0$ are constants. The term r_0 is the base ad revenue rate and β reflects ad-cluttering parameter, that is, the marginal decline in the ad revenue rate due to clutter. To ensure that the ad revenue rate is always nonnegative $r(n) \geq 0$ for all $n \in [0, 1/\alpha]$, we impose the condition that $\beta \leq r_0\alpha$. We note that $\beta = 0$ represents our base model.

Broadly, endogenizing the ad revenue rate fundamentally alters the tradeoff between DCP's two revenue levers: ads and subscription fees. In particular, it reduces the attractiveness of displaying high ad intensity to users and, in turn, compels the DCP to rely on subscription fees. The mechanism design problem for this setting can be formulated as follows:

$$\max_{\substack{n_{\mathrm{GAS}}(\cdot)\geq 0\\ p_{\mathrm{GAS}}(\cdot)\geq 0\\ d_{\mathrm{GAS}}(\cdot)\in\{0,1\}}} \int_{\theta_{\mathrm{L}}}^{\theta_{\mathrm{H}}} \Bigg(\underbrace{p_{\mathrm{GAS}}(\theta)}_{\text{subscription revenue}} + \underbrace{T\left(r_0 - \beta n_{\mathrm{GAS}}(\theta)\right) n_{\mathrm{GAS}}(\theta)\, t^*_{\mathrm{GAS}}(\theta)}_{\text{ad revenue}} \Bigg) d_{\mathrm{GAS}}(\theta) f(\theta) d\theta \quad (\mathcal{P}\text{–GAS – E})$$

subject to :

$$V_{\mathrm{GAS}}(\theta;\theta) \geq V_{\mathrm{GAS}}(\hat{\theta};\theta) \quad \forall \theta \in [\theta_{\mathrm{L}}, \theta_{\mathrm{H}}]\text{'} \qquad \text{(incentive compatibility)}$$

$$d_{\mathrm{GAS}}(\theta) = \begin{cases} 1 & \text{if } V_{\mathrm{GAS}}(\theta;\theta) \geq 0, \\ 0 & \text{otherwise} \end{cases} \qquad \forall \theta \in [\theta_{\mathrm{L}}, \theta_{\mathrm{H}}]. \qquad \text{(individual rationality)}$$

Theorem 3 characterizes an optimal solution to Problem ($\mathcal{P}$–GAS – E).

Theorem 3 (Endogenous Ad Revenue Rate). *Define* $\Delta := (\alpha^2\psi(\theta) - 2\alpha r_0 - 2\beta)^2 - 12\alpha\beta(r_0 - \alpha\psi(\theta))$, *where* $\psi(\theta) = \theta - \frac{1-F(\theta)}{f(\theta)}$ *denotes the virtual-value function. Then,*

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

the mechanism GAS − E* *defined by*

$$n^*_{\mathrm{GAS-E}}(\theta) = \begin{cases} \dfrac{2\alpha r_0 + 2\beta - \alpha^2\psi(\theta) - \sqrt{\Delta}}{6\alpha\beta}, & \text{if } \theta < \psi^{-1}\left(\dfrac{r_0}{\alpha}\right), \\ 0, & \text{if } \theta \geq \psi^{-1}\left(\dfrac{r_0}{\alpha}\right), \end{cases}$$

$$p^*_{\mathrm{GAS-E}}(\theta) = \frac{Tk\theta}{2}\left(1 - \alpha n^*_{\mathrm{GAS-E}}(\theta)\right)^2 - \int_{\theta_L}^{\theta} \frac{Tk}{2}\left(1 - \alpha n^*_{\mathrm{GAS-E}}(x)\right)^2 dx,$$

$$d^*_{\mathrm{GAS-E}}(\theta) = 1,$$

is an optimal solution to Problem ($\mathcal{P}$−*GAS* − *E*).

The left panel in Figure 5 shows that expected revenue under the GAS − E* mechanism declines monotonically as β increases. There are two distinct effects that occur as β increases: (i) each displayed ad earns less revenue to the DCP and (ii) the optimal ad intensity falls for most user types. The former is a direct effect that lowers the marginal revenue on each impression, and the latter effect reflects an endogenous adjustment in the optimal ad intensity. Collectively, both these effects reduce the platform's total revenue.

The right panel in Figure 5 contrasts the ad intensities and the subscription fees under the GAS − E* and GAS* mechanisms. Because of the presence of ad cluttering, the ad revenue lever for the DCP becomes less attractive compared with the subscription fee lever. Here, one might intuitively expect that due to ad cluttering, all the users would receive a lower ad intensity under the GAS − E* mechanism compared with that under the GAS* mechanism. Contrary to that intuition, the ad intensities under the GAS − E* mechanism (dashed, red curve) do not lie uniformly below that under the GAS* mechanism (solid, red curve). Instead, three regions emerge.

- **Low-Type Users:** Because the low-type users have low valuation for the service, the DCP cannot levy a high enough subscription fee on them, and consequently, relies on the advertising channel for revenue from these users. Here, the only significant cost of displaying an additional ad to a low-type user is the corresponding reduction in the revenue from that ad, because there is hardly any revenue from the subscription fee to forego. Consequently, these low-type users receive a higher ad intensity in the GAS − E* mechanism compared with the GAS* mechanism. The higher ad intensity also reduces the gap between the utilities of these low-type users, which in turn loosens the incentive compatibility constraints and enables the DCP to charge a higher subscription fee compared with the GAS* mechanism. In sum, the low-type users experience a higher ad intensity and a modestly higher subscription fee under the GAS − E* mechanism compared with the GAS* mechanism.
- **Intermediate-Type Users:** Compared with low-type users, the DCP can better leverage both its advertising and subscription fee levers for the intermediate-type of users. Under the GAS − E* mechanism, displaying an additional ad to these users reduces user engagement (due to ad nuisance) and the per-ad revenue (due to ad cluttering): This makes displaying high ad intensities to the users progressively less lucrative for the DCP. Consequently, the DCP scales down the ad intensities for the users in this segment more sharply under the GAS − E* mechanism compared with the GAS* mechanism. The GAS − E* mechanism partially compensates for the foregone ad revenue by raising the subscription fees for the users in these segments.
- **High-Type Users:** For these users, the ad intensities under both the mechanisms drop to zero. This is because the cluttered and uncluttered ad revenue under the GAS − E* and GAS* mechanisms, respectively, cannot cover the high ad nuisance cost faced by these users. Although both mechanisms serve no ads to these users, the subscription fee under the GAS − E* mechanism is modestly lower than that under the GAS* mechanism. Because the GAS − E* mechanism

Figure 5. (Color online) Comparison of the GAS* and GAS − E* Mechanisms

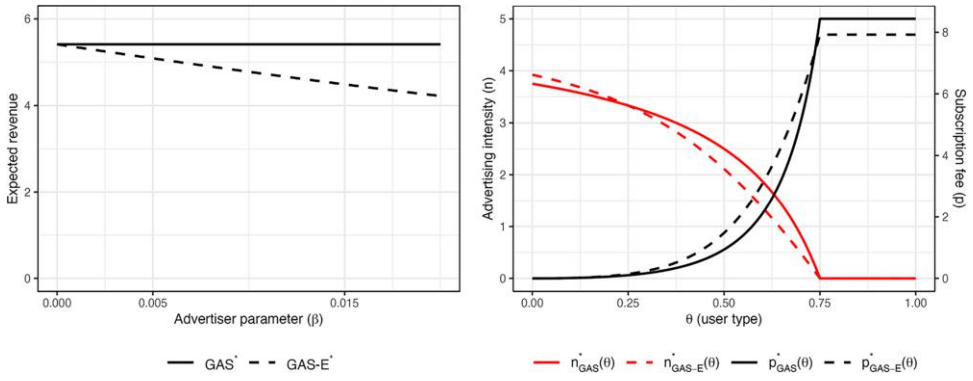

Notes. (Left) Expected revenue as a function of the advertiser sensitivity parameter β, for both GAS* (solid) and GAS − E* (dashed) mechanisms, with $r_0 = r$. (Right) Ad intensity (red) and subscription fee (black) across user types θ, for both GAS* (solid) and GAS − E* (dashed) mechanisms when $\beta = 0.01$.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

extracts a higher surplus from low- and intermediate-type users—via high ad intensities and steeper subscription fees—it shrinks the surplus available from high-type users by offering a lower subscription fee compared with the GAS* mechanism.

6.3. Two-Dimensional Heterogeneity: Approximation Bound

Thus far, we have allowed users to be heterogeneous only in their valuation for content (θ) and assumed that the ad nuisance α is common across users. We now consider the setting in which users are heterogeneous in both dimensions. In particular, each user is endowed with two-dimensional type $(\theta,\alpha)\in[\theta_{\mathrm{L}},\theta_{\mathrm{H}}]\times[\alpha_{\mathrm{L}},\alpha_{\mathrm{H}}]$, drawn from a joint distribution with density $h(\theta,\alpha)$. The DCP observes the distribution but not individual realizations.

With two-dimensional private information, the exact characterization of a revenue-maximizing mechanism is generally intractable (Daskalakis et al. 2014). To address this difficulty, we develop an approximation bound, inspired by a discretize-and-discount technique in Madarász and Prat (2017) and its application in Mehta et al. (2022), for a simple posted-menu mechanism. In particular, for this two-dimensional setting, we construct a posted-price menu consisting of a finite collection of ad intensity–subscription fee pairs (n,p) from which users self-select, and we provide an explicit additive bound on the resulting revenue loss relative to the optimal two-dimensional mechanism. We summarize the primitives, the construction of the menu, and the resulting approximation guarantee below.

Formally, consider a direct mechanism μ that maps a type $(\hat\theta,\hat\alpha)$ to an ad intensity $n_\mu(\hat\theta,\hat\alpha)$ and a subscription fee $p_\mu(\hat\theta,\hat\alpha)$. Users may opt out (outside option, represented by $\varnothing$), which yields zero utility to the users and zero revenue to the DCP. Thus, conditional on participation, a user of true type (θ,α) who selects the menu corresponding to type $(\hat\theta,\hat\alpha)$ obtains net utility:

$$V_\mu(\hat\theta,\hat\alpha;\theta,\alpha)=\frac{1}{2}T\theta k(1-\alpha\, n_\mu(\hat\theta,\hat\alpha))_+^2-p_\mu(\hat\theta,\hat\alpha). \quad (19)$$

For the approximation bound, we impose a uniform upper bound $\overline{n}$ on ad intensity, so $n\in[0,\overline{n}]$. This delivers uniform Lipschitz constants for V_μ in each type dimension and for the ad revenue term in α: $\kappa_\theta:=\frac{1}{2}Tk$, $\kappa_\alpha:=Tk\,\theta_{\mathrm{H}}\,\overline{n}$, and $\gamma_\alpha:=Trk\,\overline{n}^2$. We also use the uniform per-user revenue bound

$$\overline{R}:=\frac{1}{2}Tk\theta_{\mathrm{H}}+Trk\,\overline{n}, \quad (20)$$

which will appear directly in the approximation error. The posted menu is obtained via a three-step discretize-and-discount construction.

- **Step 1 (Discretization).** Partition the type space $[\theta_{\mathrm{L}},\theta_{\mathrm{H}}]\times[\alpha_{\mathrm{L}},\alpha_{\mathrm{H}}]$ into an $M_\theta\times M_\alpha$ grid of cells, indexed by (i,j), where $i\in\{1,\dots,M_\theta\}$ and $j\in\{1,\dots,M_\alpha\}$ denote the grid indices along the θ and α dimensions, respectively. For each cell (i,j), fix a representative type (θ^i,α^j). By Lipschitz continuity, any type's utility differs from its cell representative's by at most

$$\Delta_V:=\kappa_\theta\cdot\frac{\theta_{\mathrm{H}}-\theta_{\mathrm{L}}}{M_\theta}+\kappa_\alpha\cdot\frac{\alpha_{\mathrm{H}}-\alpha_{\mathrm{L}}}{M_\alpha} \quad (21)$$

within each cell. That is, for any menu item (n,p) and any (θ,α) in cell (i,j), $|V_\mu(\hat\theta,\hat\alpha;\theta,\alpha)-V_\mu(\hat\theta,\hat\alpha;\theta^i,\alpha^j)|\le\Delta_V$.

- **Step 2 (Representative Problem).** In this step, we assign each cell its probability mass ρ^{ij} under $h(\theta,\alpha)$ and solve the mechanism design problem restricted to the finite set of representative types. In this problem, we impose incentive compatibility and individual rationality only on this grid. Solving this problem yields a pair (n_*^{ij},p_*^{ij}) and a participation indicator $d_*^{ij}\in\{0,1\}$ for each cell (i,j). Here, $d_*^{ij}=0$ corresponds to the outside option $\varnothing$. Let $\mathcal{M}:=\{(n_*^{ij},p_*^{ij}):d_*^{ij}=1\}$ denote the menu.
- **Step 3 (Uniform Discounting).** The menu $\mathcal{M}$ is incentive compatible only on the grid of representative types. In this final step, we modify the subscription fees so that deviations are controlled for all types in the continuous space. Let $\eta\in(0,1]$ denote the discount parameter for the subscription fees. For each menu item (i,j), define its representative revenue

$$\pi_*^{ij}:=p_*^{ij}+Trk\,n_*^{ij}(1-\alpha^j n_*^{ij})_+. \quad (22)$$

We then discount the corresponding fee by the fraction η of π_*^{ij}, that is, $\widetilde{p}^{ij}:=p_*^{ij}-\eta\pi_*^{ij}$, and define the discounted posted menu $\widetilde{\mathcal{M}}(\eta):=\{(n_*^{ij},\widetilde{p}^{ij})\}_{i,j}\cup\{\varnothing\}$. When indifferent across menu items, users break ties in favor of the item with higher representative revenue π_*^{ij}.

The rationale behind discounting the subscription fees is to uniformly control off-grid deviations. Discounting adds the same constant $\eta\pi_*^{ij}$ to every type's utility from item (i,j). Hence, a type in cell (i,j) can profitably switch to (i',j') only if the revenue-based discount gap is large enough to overcome the within-cell utility error Δ_V. This implies $\pi_*^{i'j'}$ must be within order Δ_V/η of π_*^{ij}. This observation gives us the approximation guarantee in Theorem 4.

Theorem 4 (Two-Dimensional Approximation Bound). *Let $\pi_{2\mathrm{D}}^*$ denote the optimal expected revenue in the two-dimensional problem and let $\pi(\widetilde{\mathcal{M}}(\eta))$ denote the expected revenue when types best-respond to the discounted posted menu $\widetilde{\mathcal{M}}(\eta)$. Choosing*

$$\eta^*:=\min\left\{1,\sqrt{\frac{2\Delta_V}{\overline{R}}}\right\} \quad (23)$$

the discounted menu $\widetilde{\mathcal{M}}(\eta^*)$ satisfies

$$\pi_{2\mathrm{D}}^*-\pi(\widetilde{\mathcal{M}}(\eta^*))\le 2\left(\eta^*\overline{R}+\frac{2\Delta_V}{\eta^*}+\gamma_\alpha(\alpha_{\mathrm{H}}-\alpha_{\mathrm{L}})\right). \quad (24)$$

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

Theorem 4 provides an explicit posted menu and an additive bound on its revenue loss relative to the optimal two-dimensional mechanism. The discount parameter η captures a key tradeoff: On one hand, a larger η makes profitable off-grid deviations harder (reflected in the term $2\Delta_V/\eta$). On the other hand, it also lowers revenue by uniformly reducing subscription fees (reflected in $\eta\overline{R}$). The choice η^* balances these two forces. Further, as the discretization is refined ($M_\theta, M_\alpha \to \infty$), $\Delta_V \to 0$ and hence $\eta^* \to 0$. Thus, the first two components of the bound vanish. The last term $\gamma_\alpha(\alpha_H - \alpha_L)$ arises from the dependence of the ad revenue on α and remains even as the grid becomes arbitrarily fine. We note here that for the last term to converge to zero, we need additional structure on the model primitives.

Theorem 4 shows that a finer discretization improves the performance guarantee of the posted menu. However, refining the grid comes at the expense of an increase in the size of the menu (ad intensity–subscription fee pairs), which may be undesirable in practice. We therefore consider the setting where the DCP limits the size of the menu to at most L items and obtain a bound on the revenue loss.

Notice that a menu of size at most L can be obtained by running Steps 1–3 above on an $M_\theta \times M_\alpha$ grid with $M_\theta M_\alpha = L$. Denote this menu by $\widetilde{\mathcal{M}}_L(\eta)$. Allocating the L grid points across the two dimensions to minimize Δ_V yields

$$M_\theta^*(L) = \sqrt{L}\sqrt{\frac{\kappa_\theta(\theta_H - \theta_L)}{\kappa_\alpha(\alpha_H - \alpha_L)}},$$
$$M_\alpha^*(L) = \sqrt{L}\sqrt{\frac{\kappa_\alpha(\alpha_H - \alpha_L)}{\kappa_\theta(\theta_H - \theta_L)}}, \tag{25}$$

with discretization error

$$\Delta_V^*(L) := 2\sqrt{\frac{\kappa_\theta\kappa_\alpha(\theta_H - \theta_L)(\alpha_H - \alpha_L)}{L}}. \tag{26}$$

Substituting $\Delta_V^*(L)$ into Theorem 4 gives the following guarantee for $\widetilde{\mathcal{M}}_L(\eta)$.

Corollary 1 (Limited-Size Menu). *For a menu of size at most $L \in \mathbb{N}$ obtained as above, let*

$$\eta_L^* := \min\left\{1, \sqrt{2\Delta_V^*(L)\overline{R}}\right\}. \tag{27}$$

Then the menu $\widetilde{\mathcal{M}}_L(\eta_L^*)$ satisfies

$$\pi_{2D}^* - \pi(\widetilde{\mathcal{M}}_L(\eta_L^*)) \le 2\left(\eta_L^*\overline{R} + \frac{2\Delta_V^*(L)}{\eta_L^*} + \gamma_\alpha(\alpha_H - \alpha_L)\right). \tag{28}$$

Corollary 1 implies that if the DCP offers a menu of length at most L, then the discretization-discounting component of the revenue loss bound decreases proportionally to $L^{-1/4}$ as the menu size L increases. This shows that longer menus deliver progressively tighter approximation guarantees.

6.4. UAS Mechanism

In addition to the ADS and SUB mechanisms, the DCP may commit to a single service tier that combines a subscription fee with an ad load. We refer to this as the unary ad-supported mechanism (μ = UAS). Under this mechanism, all participating users face a common ad intensity and a common subscription fee. Formally, a UAS mechanism is characterized by the tuple (n_{UAS}, p_{UAS}), where $n_{UAS} \ge 0$ denotes the ad intensity and $p_{UAS} \ge 0$ denotes the subscription fee. The UAS mechanism can be viewed as an extreme simplification of the GAS mechanism: Whereas GAS offers a continuum of ad intensity–subscription fee pairs, UAS collapses this continuum to a single pair (n_{UAS}, p_{UAS}) offered uniformly to all users. It therefore eliminates the continuous-menu concern that arises while implementing GAS mechanism, at the cost of the ability to price discriminate across user types.

Under the UAS mechanism, a type θ user chooses usage time $t_{UAS}^*(\theta) = k(1 - \alpha n_{UAS})_+$ and obtains net utility $V_{UAS}(\theta) = \frac{1}{2}T\theta k((1 - \alpha n_{UAS})_+)^2 - p_{UAS}$. Let θ_U denote the marginal user type for which $V_{UAS}(\theta_U) = 0$. Then,

$$\theta_U := \min\left\{\theta_H, \frac{2p_{UAS}}{Tk\left((1 - \alpha n_{UAS})_+\right)^2}\right\}. \tag{29}$$

Note that $V_{UAS}(\theta)$ is an increasing function of θ. Thus, a user participates if and only if $\theta \ge \theta_U$. The DCP's revenue-maximization problem can then be written as

$$\max_{n_{UAS}\ge 0,\, p_{UAS}\ge 0} \left(p_{UAS} + Tr\, n_{UAS} k(1 - \alpha n_{UAS})_+\right)\left(1 - F(\theta_U)\right). \tag{30}$$

The UAS* mechanism involves a tradeoff between monetization and participation. Raising the subscription fee p_{UAS^*} increases revenue per participating user but also raises the participation cutoff θ_U. Raising ad intensity n_{UAS^*} increases ad revenue per user, but it also reduces usage and lowers users' net utility, thereby raising θ_U. Therefore, both instruments increase monetization conditional on participation, but they do so at the cost of reduced participation. We also note that the UAS* mechanism nests ADS* and SUB* as special cases: Setting $p_{UAS^*} = 0$ recovers the ADS mechanism, whereas setting $n_{UAS^*} = 0$ recovers the SUB* mechanism. Comparing UAS* against these special cases identifies the regions in the (r, α) parameter space where $\pi_{UAS^*} = \pi_{ADS^*}$, where $\pi_{UAS^*} = \pi_{SUB^*}$, and where $\pi_{UAS^*} > \max\{\pi_{ADS^*}, \pi_{SUB^*}\}$.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

Although UAS* weakly revenue dominates ADS* and SUB*, it remains more constrained than GAS*: It offers a single ad intensity–subscription fee pair $(n_{\mathrm{UAS}}, p_{\mathrm{UAS}})$ uniformly to all users, whereas GAS* offers a type-dependent menu and can therefore better price discriminate across users. Comparing the two mechanisms, we find that GAS* retains a revenue advantage that is largest at low-to-intermediate ad attractiveness (r/α) and a low-to-moderate share of high-type users (a/b). This is the region where tailoring ad intensities and subscription fees across types is most valuable to GAS*.

We also note that this pattern differs from the GAS*-BAS* comparison, which naturally motivates a comparison of UAS* and BAS* mechanisms. This comparison shows that neither the UAS* nor the BAS* mechanism uniformly dominates: UAS* is preferred when ads are attractive and high-type users predominate, whereas BAS* is preferred when ad attractiveness is low and low-type users predominate.

6.5. GAS Mechanism Under Competition

Thus far, we analyzed mechanisms for a monopolist DCP. We now briefly discuss how the GAS mechanism is affected in the presence of competition. A tractable way to incorporate competition in our framework is to model the rival platform as an outside option that tightens the individual rationality constraint. In our baseline model, the outside option is normalized to zero. With a competitor, we can represent the rival platform as an outside option. Let V^{o} denote the utility a user can obtain from the best alternative platform. Then the participation condition in (5) becomes

$$d_{\mu}(\hat{\theta};\theta) = \begin{cases} 1 & \text{if } V_{\mu}(\hat{\theta};\theta) \geq V^{\mathrm{o}}, \\ 0 & \text{otherwise.} \end{cases} \tag{31}$$

The key insights from this analysis are as follows.

First, introducing an outside option primarily affects participation of the low-type users. Relative to the baseline case with $V^{\mathrm{o}} = 0$, low-type users may no longer find the DCP's ad-intensive offering attractive and may prefer to opt out. Consequently, full market coverage need not hold under competition. Second, with an outside option V^{o}, the DCP must ensure that participating users receive at least their reservation utility. This tighter participation requirement reduces the DCP's ability to extract surplus. Put differently, the competition forces DCP to lower subscription fees relative to the monopoly benchmark. Third, conditional on participation, the optimal ad intensity retains the same structure as in the baseline GAS* mechanism. Intuitively, the outside option changes the participation decision, but it does not change how the DCP trades off ad revenue against the reduction in user utility caused by ads. Finally, as the outside option improves, the served market shrinks. Thus, DCP primarily focuses on serving higher types. In particular, sufficiently strong competitive pressure can shift the optimal mechanism toward less ad-intensive (and potentially ad-free) offerings for the participating users.

7. Application to VoD Platforms

In this section, we validate the assumptions and results of our theoretical model in the context of VoD streaming platforms through two phases (Figure 6). In Phase 1, we collect data based on contextual features commonly associated with contemporary popular VoD streaming platforms in order to calibrate the model parameters. In Phase 2, we plug those estimates into our model and compute the expected performance of the GAS* mechanism vis-à-vis other mechanisms on the VoD market, both in terms of revenue to the DCP and consumer welfare. We briefly discuss the phases below.

7.1. Phase 1: Data Collection and Estimation of Model Parameters

To calibrate the optimal mechanisms, we estimate four key elements of our model: the ad nuisance parameter α, the distribution $F(\cdot)$ of user types θ, the usage parameter k, and the ad revenue rate r. For this purpose, we conduct a data collection exercise using the Prolific[8] platform. Here, participants were asked to report both their willingness-to-pay (WTP) for access to a fictitious VoD service (that draws on contextual features commonly associated with popular VoD streaming platforms) and their intended usage time for that service.

The study employed a Becker-DeGroot-Marschak (BDM) mechanism (Becker et al. 1964), a well-established tool in experimental economics designed to elicit truthful valuations. The BDM mechanism ensures incentive compatibility: Participants maximize their expected utility by revealing their true reservation price.

To estimate the ad nuisance parameter α, participants were presented with five distinct ad intensity levels, ranging from zero to four ads per hour.[9] For each level, users reported their usage time (without upper limit) under a fixed subscription fee, yielding a data set of $5 \times 316 = 1{,}580$ observations. This variation allows us to isolate and quantify the disutility from advertising.[10] We estimate the effect of ad intensity on reported usage time using Poisson regressions, as usage time is reported in minutes and thus represents count data. We consider different model specifications to test for potential nonlinearity. Our results indicate that usage time decreases linearly with ad intensity, with a nuisance parameter $\hat{\alpha}$ of approximately 0.188. As a robustness check, we also run a Tobit regression to account for censoring at zero and obtain similar results.

Next, we use the elicited WTP values to recover the distribution $F(\cdot)$ of user types θ. Our empirical findings suggest that user valuations follow a truncated

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

Figure 6. Schematic Representation of the Two-Phase Empirical Procedure

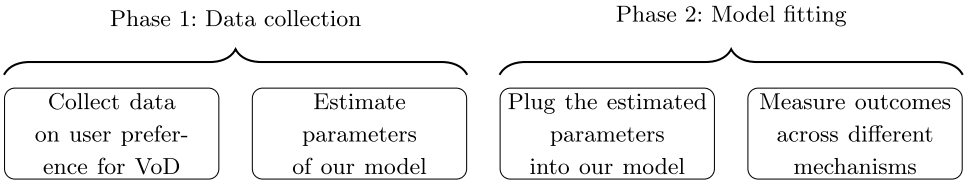

exponential distribution with rate parameter 2.13 on support [0, 1.68]. For ad revenue rate, we calibrate the cost per mille (CPM) at \$30, consistent with prevailing market rates.[11] This implies a per-impression revenue of $\hat{r} = \$0.03$ for the platform.

Based on these estimates, we parameterize the VoD market as follows: $\hat{\alpha} = 0.188$, $\hat{k} = 1.78$, $\hat{F}(\cdot)$ as a truncated exponential with rate 2.13 on support [0, 1.68], and $\hat{r} = 0.03$. These calibrated inputs enable us to compute the SUB*, ADS*, BAS*, and GAS* mechanisms for the VoD market, yielding the following outcomes:

- ADS*: \$0 with 2.64 ads per hour,
- SUB*: \$11.69 with 0 ads per hour,
- BAS*: ad-supported tier \$0 with 3.79 ads per hour; ad-free tier \$12.39 with 0 ads per hour, and
- GAS*: a continuum ranging from \$12.49 with 0 ads per hour to \$0 with 4 ads per hour.

7.2. Phase 2: Fitting the Model and Measuring Outcomes

Using the estimated parameters, we compute the expected revenue and its decomposition (i.e., whether the revenue to the DCP stems from ads or from subscription fee) under each of the four mechanisms. Figure 7 illustrates the per-user revenue decomposition. In line with our theoretical model, the calibrated GAS* mechanism outperforms the other three mechanisms in terms of expected revenue to the DCP. Specifically, GAS* generates 148%, 20%, and 1.8% more revenue per user compared with ADS*, SUB*, and BAS*, respectively.

We further examine the advantage of the GAS* mechanism by decomposing the per-user revenue gap against the other three optimal mechanisms (Figure 8). In the left panel of that figure, we observe that, except for low-type users, the GAS* mechanism yields a higher revenue than the ADS* mechanism. This is because the GAS* mechanism can use ad-intensity to screen users: Low-type users, who have low willingness-to-pay for reduced ad exposure, are assigned higher ad intensities to generate ad revenue, whereas high-type users are offered lower ad intensities in exchange for higher subscription fees. In contrast, the ADS* mechanism sets a single, uniform ad intensity for all users and cannot effectively screen users. However, the GAS* mechanism better exploits the intermediate- and high-type users to cumulatively yield a higher revenue than the

Figure 7. Expected Revenue per User for Each Mechanism

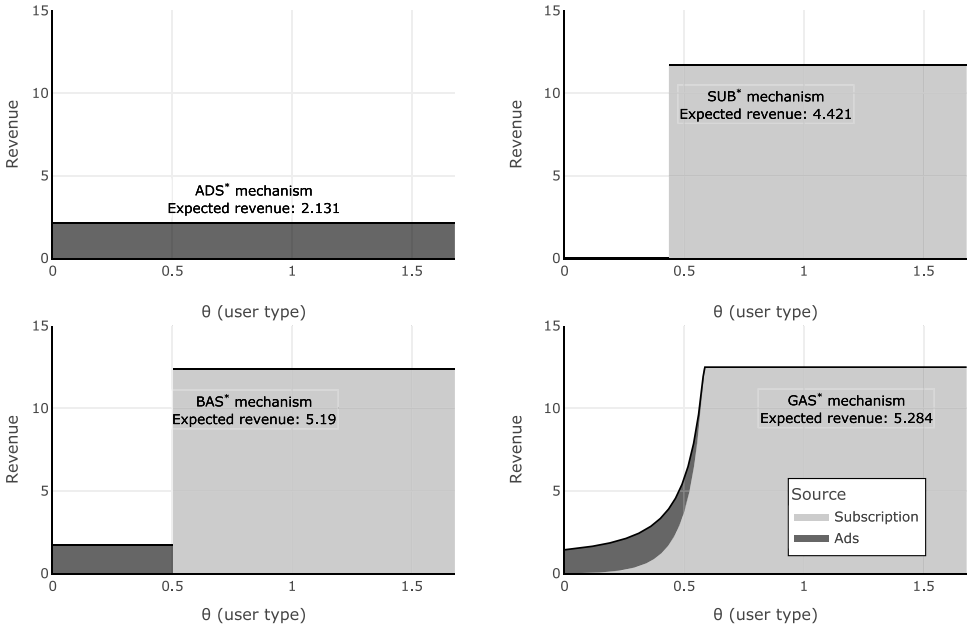

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

Figure 8. Difference in Revenue per User Between GAS* and Other Mechanisms

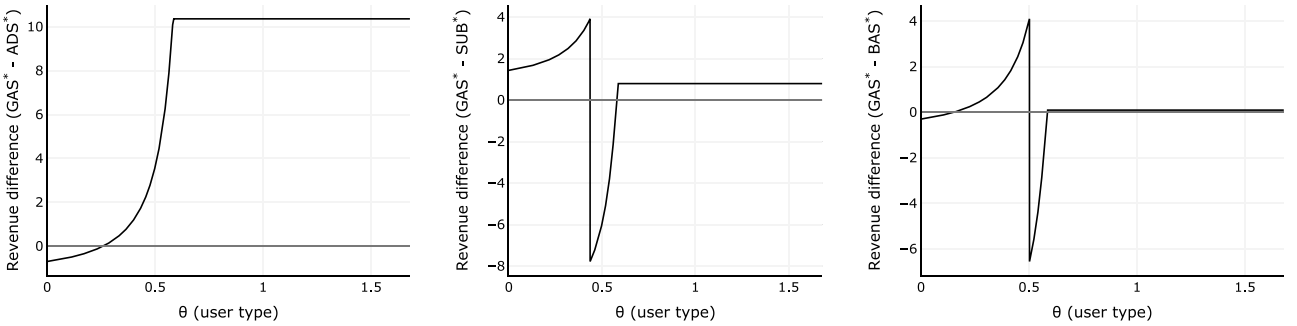

ADS* mechanism. The middle panel shows that the GAS* mechanism improves over the SUB* mechanism in two distinct ways: It provides ad-supported access to low-type users who did not participate under the SUB* mechanism, and it charges high-type users a steeper subscription fee relative to the SUB* mechanism. Finally, the right panel demonstrates the near-optimal performance of the BAS* mechanism. The two tiers of the BAS* mechanism align closely with the preferences of extreme user types, and the incremental improvement that GAS* provides comes predominantly from fine-tuning the mix of ad intensities and subscription fees for the intermediate types.

To understand how the users fare under each mechanism, we compute individual surplus (net usage utility less the subscription fee) for each type and aggregate it (weighted by the distribution) to obtain the consumer surplus (Figure 9). We observe that surplus increases with user type θ under all four mechanisms because (i) all the mechanisms satisfy incentive compatibility, which ensures that high-type users do not face heavier ad intensities or disproportionately higher subscription fees than the low-type users, and (ii) net usage utility increases with user type θ, which ensures that the individual surplus of higher-type users is at least as much as that of lower-type users. We also note that the SUB*, BAS*, and GAS* mechanisms deliver approximately 45% more consumer surplus than the ADS* mechanism. This shows that in the context of VoD, advertising is an expensive channel to raise revenue for the DCPs. The SUB*, BAS*, and GAS* mechanisms substitute a large share of their ad revenue with subscription fees, which helps reduce the ad nuisance cost for the users, which in turn improves welfare.

Figure 10 further explores the welfare implications by plotting, for each user type θ, the point-wise

Figure 9. Consumer Surplus per User for Each Mechanism

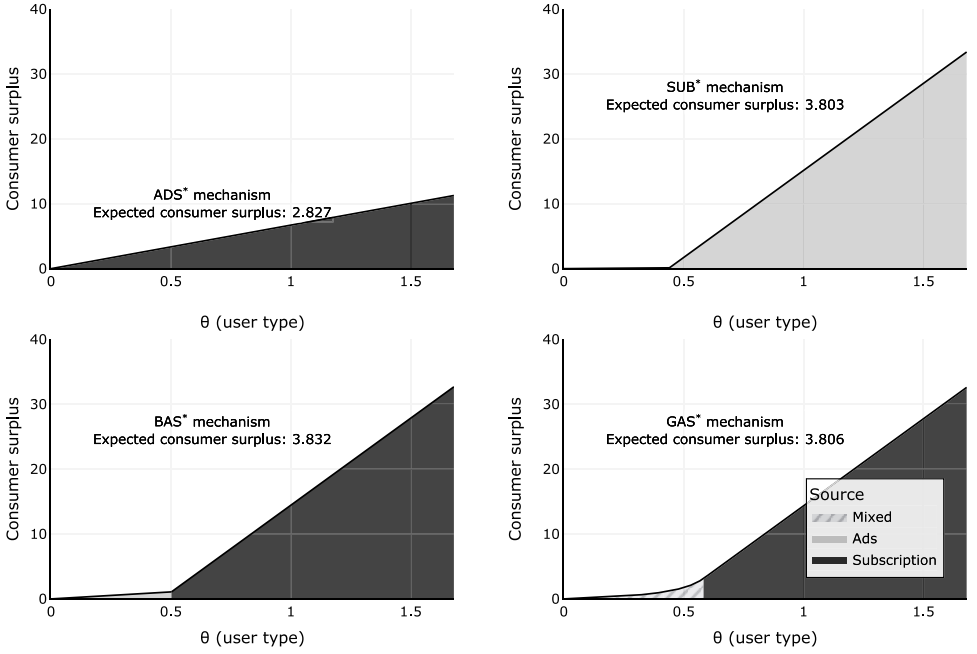

Figure 10. Consumer Surplus Difference Between GAS* and Other Mechanisms

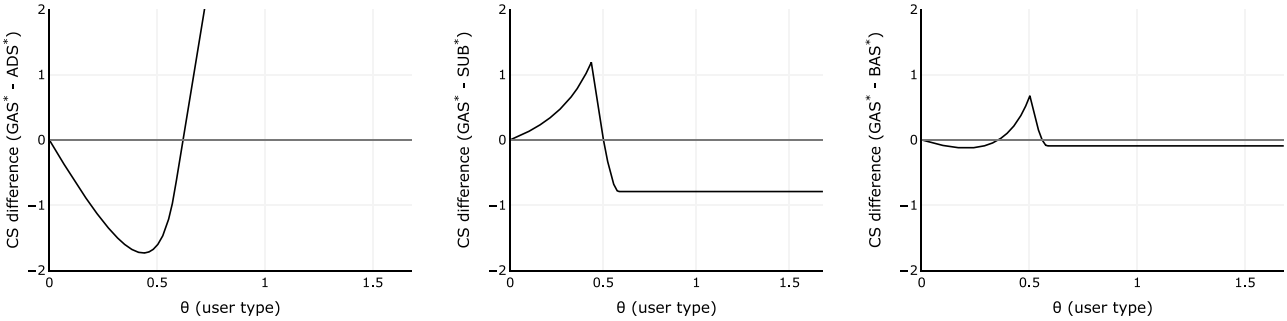

Notes. The vertical axis is fixed in the range [−2, 2] across all three plots for consistency. Therefore, the values in the left panel that exceed two are truncated.

difference in consumer surplus between the GAS* mechanism and each of the other three optimal mechanisms. The left panel shows that relative to the ADS* mechanism, the GAS* mechanism generates higher consumer surplus. Compared with the ADS* mechanism, GAS* yields lower surplus for low- and medium-type users but significantly higher surplus for high-type users: This is because the high-type users can completely avoid the ad nuisance cost under the GAS* mechanism. Overall, the increased consumer surplus enjoyed by high-type users offsets the lower surplus experienced by the other users. The middle panel compares the SUB* mechanism with the GAS* mechanism. The low-type users benefit under the GAS* mechanism because they can opt for an ad-supported tier whereas they are priced out under the SUB* mechanism. On the flip side, the high-type users benefit under the SUB* mechanism because they receive an ad-free experience. Finally, the right panel contrasts the GAS* mechanism with the BAS* mechanism and shows an almost flat differential except for a modest hump for intermediate-type users. This is because the BAS* mechanism already spans the two extreme options of ad-free tier and no-subscription fee tier. However, note that the ad-intensity on the free tier and the subscription fee on the ad-free tier need not be identical under BAS* and GAS*, as the intermediate options under GAS* change the binding incentive constraints. The GAS* mechanism better matches the preferences of the intermediate type users who value a reduced ad-intensity and are not willing to pay the full subscription fee. Therefore, the GAS* mechanism delivers a higher surplus to these users compared with the BAS* mechanism. However, the aggregate consumer surplus across user types is marginally higher under the BAS* mechanism relative to the GAS* mechanism.

8. Managerial Implications and Concluding Remarks

Our study offers a unified framework for understanding how digital content providers should leverage the two revenue levers (advertising and subscription fees) to design mechanisms that maximize the expected revenue in the presence of information asymmetry. Our GAS framework is a broader class that nests the commonly studied ADS, SUB, and BAS mechanisms as special cases. Therefore, this framework enables a head-to-head comparison of the mechanisms under a common set of primitives.

First, we show that the GAS* mechanism is implementable in practice: Proposition 1 establishes a posted-price self-selection menu, so users choose from a schedule of ad-intensity–subscription fee pairs. Such a continuous menu can be delivered through interactive interfaces such as sliders (Remark 1). Alternatively, we show that if DCP offers a finite-size menu with at most L options, then the revenue loss decreases proportionally to $L^{-1/4}$ as L increases (Corollary 1).

Moving from ADS*/SUB* to BAS* to GAS* represents increasing flexibility in how the DCP uses the two revenue levers to screen heterogeneous users. In Section 5, we comprehensively compare the four mechanisms. Table 2 synthesizes the insights from this analysis in the form of a practical implementation guide.

Beyond the comparative analysis of mechanisms, we leverage the tractability of our GAS framework to explore several policy- and market-motivated extensions and validate its foundation.

- In Sections 6.1 and 6.2, we study advertising caps and ad cluttering (endogenous ad revenue rates) and show how these forces can induce exclusion of low-type users and reshape the optimal allocation of ads across types.
- In Section 6.3, we relax the assumption of a common ad nuisance parameter and consider a setting where users are heterogeneous in both content valuation and ad sensitivity. Because exact optimization in this two-dimensional setting is generally intractable, we develop an approximation that yields an explicit finite-sized menu of ad-intensity–subscription fee pairs together with an additive bound on the revenue loss. Here, we also show that offering a menu with at most L

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

Table 2. Implementation Guide Based on the Comparative Analyses of the Mechanisms

	Revenue lever		
Mechanism	Advertising	Subscription fee	When to use
ADS*	✓	×	DCP prefers a simple mechanism (single revenue lever). Ad revenue rate (r) is high and users are not very sensitive to ads (i.e., low α).
SUB*	×	✓	DCP prefers a simple mechanism (single revenue lever). Advertising is unattractive (high α or low r) and the proportion of high-type users (who are willing to pay for content) is sufficiently high.
BAS*	✓	✓	Neither lever clearly dominates and DCP is comfortable combining the two levers in a simple way. The revenue gap with GAS* is small when ad attractiveness or the proportion of high-type users is low (see Figure 3).
GAS*	✓	✓	Both ad attractiveness (r/α) and the proportion of high-type users (a/b) are high (see Figure 3). Implementable via a slider interface (Remark 1) or a finite-menu approximation (Corollary 1).

options incurs a loss that decreases proportionally to $L^{-1/4}$. Thus, offering longer menus deliver progressively tighter guarantees.

- In Section 6.4, we analyze a single-tier mechanism that combines ads and subscription fees, denoted by UAS*. This mechanism naturally nests ADS* and SUB* as special cases and can be viewed as an extreme simplification of GAS*. The single-tier structure makes UAS* easier to implement than the continuous menu of GAS*. We characterize the parameter regions in which UAS* dominates the ADS* and SUB* mechanisms, examine the revenue gap between GAS* and UAS*, and identify when UAS* outperforms BAS*.
- Finally, we study competition by introducing an outside option that tightens the participation constraint (Section 6.5) and show that competitive pressure can shrink market coverage and shift the optimal GAS mechanism toward less ad-intensive offerings.
- We examine the robustness of GAS* revenue to estimation error in the ad nuisance parameter. Here, we show that if the ad nuisance parameter is estimated as $\hat{\alpha}$ while the true value is α, and the estimation error satisfies $|\hat{\alpha} - \alpha| \leq \epsilon$, for some $\epsilon > 0$, then the resulting revenue loss under the GAS* mechanism is bounded linearly in the error ϵ. Together with the numerical finding that GAS* revenue is relatively insensitive to α, this suggests that the GAS* mechanism is robust to moderate misspecification of the ad nuisance parameter.
- In Section 7, we ground our theoretical model and results with an empirical choice survey in a VoD setting. We estimate the distribution of user types and other model parameters such as ad nuisance cost and average usage time. The empirical study validates the primitives of our theoretical model and attests to the robustness of our framework for modeling user choices.

Although our framework offers robust insights into designing revenue-maximizing mechanisms for DCPs, we have abstracted away from a few real-world complexities that merit further exploration. On the modeling side, our framework captures ads via ad intensity but does not explore granular elements such as targeted advertising or content-level heterogeneity. We model the BAS mechanism as a simple two-tier menu—one ad-supported tier and one subscription fee tier—so that the ADS and SUB mechanisms occur as its extreme cases and can be compared in a head-to-head manner. In practice, DCPs could optimize the ad intensity and subscription fee levels of both tiers endogenously. Similarly, our theoretical model is static and assumes that the DCP has knowledge of the model parameters. In practice, a DCP can learn these parameters during an initial period before deploying a mechanism—Our two-phase empirical procedure in Section 7 illustrates one such approach. Jointly optimizing the learning and pricing phases through a dynamic mechanism design framework is a natural extension. On the operationalization front, although the implementation of the GAS* mechanism is feasible (Remark 1), user-facing considerations such as bounded rationality, choice overload, and anchoring biases may affect how users interact with such menus. More broadly, our empirical calibration validates the model primitives but does not assess deployment in operational environments. Testing the GAS* mechanism in a live setting, such as studying practical adoption barriers, user behavior, and long-run retention effects, is a promising direction for future work.

Acknowledgments

The authors thank the senior editor, associate editor, and three anonymous referees for constructive and thoughtful feedback. The authors also thank the participants at WISE and SCECR for the valuable comments and suggestions. Sameer Mehta gratefully acknowledges support from the Dutch Research Council (NWO) through a Veni Grant (VI.Veni.231E.063) during the course of this research. The authors are listed in alphabetical order.

Endnotes

[1] That is, the user self-selects an ad intensity and subscription fee pair from the menu offered by the DCP; this selection identifies the type of the user within the mechanism. The incentive-compatibility constraints (introduced below) ensure that each user type prefers the pair intended for that user over those intended for other types.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.

[2] The linear decay in the usage time with ad intensity is consistent with recent empirical literature on digital content platforms. For instance, Goli et al. (2025) conduct a field experiment with Pandora (an online streaming radio service) to study how ad intensity affects content consumption and find that listening time decreases linearly with ad intensity.

[3] See https://amplitude.com/pricing.

[4] See https://zapier.com/pricing.

[5] See https://www.emarketer.com/content/digital-video-forecast-trends-q4-2024.

[6] See https://commission.europa.eu/strategy-and-policy/priorities-2019-2024/europe-fit-digital-age/digital-services-act_en.

[7] See https://www.digital-fairness-act.com/.

[8] Prolific (https://www.prolific.com/) is a commonly used online research platform that connects researchers with a diverse pool of participants for academic studies.

[9] As a robustness check, we vary the length of the ads from 30 seconds to 1 minute, which does not change our estimates.

[10] We also collect demographic information for each respondent to ensure that the sample captures a diverse set of users in terms of age, gender, employment status, and viewing habits across different VoD services.

[11] See https://www.emarketer.com/content/digital-video-forecast-trends-q4-2024.

References

Armstrong M (1996) Multiproduct nonlinear pricing. *Econometrica* 64(1):51–75.

Aseri M, Dawande M, Janakiraman G, Mookerjee VS (2020) Ad-blockers: A blessing or a curse? *Inform. Systems Res.* 31(2): 627–646.

Bakos Y, Brynjolfsson E (1999) Bundling information goods: Pricing, profits, and efficiency. *Management Sci.* 45(12):1613–1630.

Balseiro S, Feldman J, Mirrokni V, Muthukrishnan S (2014) Yield optimization of display advertising with ad exchange. *Management Sci.* 60(12):2886–2907.

Bapna R, Ramaprasad J, Umyarov A (2018) Monetizing freemium communities. *MIS Quart.* 42(3):719–735.

Becker GM, DeGroot MH, Marschak J (1964) Measuring utility by a single-response sequential method. *Behav. Sci.* 9(3):226–232.

Bhargava HK, Choudhary V (2001) Information goods and vertical differentiation. *J. Management Inform. Systems* 18(2):89–106.

Bhargava HK, Choudhary V (2008) Research note—When is versioning optimal for information goods? *Management Sci.* 54(5): 1029–1035.

Chiou L, Tucker C (2013) Paywalls and the demand for news. *Inform. Econom. Policy* 25(2):61–69.

Conitzer V, Kroer C, Panigrahi D, Schrijvers O, Stier-Moses NE, Sodomka E, Wilkens CA (2022) Pacing equilibrium in first price auction markets. *Management Sci.* 68(12):8515–8535.

Daskalakis C, Deckelbaum A, Tzamos C (2014) The complexity of optimal mechanism design. *Proc. 25th Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM).

DeValve L, Pekeč S (2022) Optimal price/advertising menus for two-sided media platforms. *Oper. Res.* 70(3):1629–1645.

Fan M, Kumar S, Whinston AB (2007) Selling or advertising: Strategies for providing digital media online. *J. Management Inform. Systems* 24(3):143–166.

Goldfarb A, Tucker C (2011) Online display advertising: Targeting and obtrusiveness. *Marketing Sci.* 30(3):389–404.

Goldstein DG, Suri S, McAfee RP, Ekstrand-Abueg M, Diaz F (2014) The economic and cognitive costs of annoying display advertisements. *J. Marketing Res.* 51(6):742–752.

Goli A, Reiley DH, Zhang H (2024) Personalizing ad load to optimize subscription and ad revenues: Product strategies constructed from experiments on pandora. *Marketing Sci.* 44(2):327–352.

Goli A, Huang J, Reiley D, Riabov NM (2025) Measuring consumer sensitivity to audio advertising: A long-run field experiment on Pandora internet radio. *Quant. Marketing Econom.* 23(3):347–377.

Ichihashi S, Jeon D-S, Kim B-C (2024) Mechanism design and innovation incentive for an ad-funded platform. Preprint, submitted December 13, https://doi.org/10.2139/ssrn.5016687.

Krishna V (2009) *Auction Theory* (Academic Press, New York).

Lambrecht A, Misra K (2017) Fee or free: When should firms charge for online content? *Management Sci.* 63(4):1150–1165.

Lin S (2020) Two-sided price discrimination by media platforms. *Marketing Sci.* 39(2):317–338.

Madarász K, Prat A (2017) Sellers with misspecified models. *Rev. Econom. Stud.* 84(2):790–815.

Maskin E, Riley J (1984) Monopoly with incomplete information. *RAND J. Econom.* 15(2):171–196.

Mehta S, Dawande M, Janakiraman G, Mookerjee V (2020) Sustaining a good impression: Mechanisms for selling partitioned impressions at ad exchanges. *Inform. Systems Res.* 31(1):126–147.

Mehta S, Dawande M, Janakiraman G, Mookerjee V (2022) An approximation scheme for data monetization. *Production Oper. Management* 31(6):2412–2428.

Mussa M, Rosen S (1978) Monopoly and product quality. *J. Econom. Theory* 18(2):301–317.

Myerson RB (1981) Optimal auction design. *Math. Oper. Res.* 6(1):58–73.

Oestreicher-Singer G, Zalmanson L (2013) Content or community? A digital business strategy for content providers in the social age. *MIS Quart.* 37(2):591–616.

Oh H, Animesh A, Pinsonneault A (2016) Free versus for-a-fee. *MIS Quart.* 40(1):31–56.

Ostrovsky M, Schwarz M (2023) Reserve prices in internet advertising auctions: A field experiment. *J. Political Econom.* 131(12):3352–3376.

Riggins FJ (2002) Market segmentation and information development costs in a two-tiered fee-based and sponsorship-based web site. *J. Management Inform. Systems* 19(3):69–86.

Rochet J-C, Choné P (1998) Ironing, sweeping, and multidimensional screening. *Econometrica* 66(4):783–826.

Shen H, Li Y, Chen Y, Pan K (2021) Integrated ad delivery planning for targeted display advertising. *Oper. Res.* 69(5):1409–1429.

Statista (2023) Media–Worldwide. Accessed July 23, https://www.statista.com/outlook/amo/media/worldwide.

Todri V, Ghose A, Singh PV (2020) Trade-offs in online advertising: Advertising effectiveness and annoyance dynamics across the purchase funnel. *Inform. Systems Res.* 31(1):102–125.

Wang L, Lowry PB, Luo X, Li H (2023) Moving consumers from free to fee in platform-based markets: An empirical study of multiplayer online battle arena games. *Inform. Systems Res.* 34(1):275–296.

Xu L, Duan JA, Hu YJ, Cheng Y, Zhu Y (2019) Forward-looking behavior in mobile data consumption and targeted promotion design: A dynamic structural model. *Inform. Systems Res.* 30(2):616–635.

Yang J, Zhang J, Zhang Y (2025) Engagement that sells: Influencer video advertising on TikTok. *Marketing Sci.* 44(2):247–267.

Yoganarasimhan H, Barzegary E, Pani A (2023) Design and evaluation of optimal free trials. *Management Sci.* 69(6):3220–3240.

Dominik Gutt is a chaired professor at the School of Business and Economics, RWTH Aachen University. His research focuses on user-generated content, web3, and artificial intelligence. He received the AIS Early Career Award in 2024 and the Reviewer of the Year Award 2023 from *MIS Quarterly*, where he currently serves as an associate editor. His work is published in leading journals, such as *Information Systems Research* and *MIS Quarterly*.

Sameer Mehta is an Associate Professor at the Rotterdam School of Management, Erasmus University. He obtained his PhD from the University of Texas at Dallas. His research focuses on the design and analysis of technology markets, using tools from mechanism design, game theory, and optimization.

Martin Quinn is a professor at the Institut Mines-Télécom Business School. His research examines how platform design, regulation, and market mechanisms shape user behavior, firm strategy, and societal outcomes. He publishes in leading outlets including *The Economic Journal, Journal of Management Information Systems,* and *Journal of Economics and Management Strategy.* His work spans digital platforms, online advertising, content moderation, generative AI, privacy, and digital regulation worldwide.

Downloaded from informs.org by [2a00:8a60:5:111::1002] on 26 October 2026, at 08:16 . For personal use only, all rights reserved.