Algorithmic Attention and Content Creation on Social Media Platforms*

Yi Chen Fei Li Marcel Preuss

June 3, 2026

Abstract

We study the revenue-maximizing allocation of attention on an ad-funded social media platform governed by content recommendation algorithms. Attention is costly and serves a dual role: it can be directly monetized through advertising, or it can be allocated to creators to increase exposure, strengthen effort incentives, and thereby generate more attention thanks to better content. This tension creates a trade-off between monetization and production incentives. In a two-sided model with heterogeneous viewers and creators under private information, the optimal recommendation mix includes content that is ex post suboptimal for some viewers in order to leverage cross-side network externalities. Two-sided complementarities amplify informational distortions and reshape the platform's content landscape, influencing both the variety of content and its average quality. Our analysis yields policy implications for personalization regulation and advertising markets.

Keywords: Social Media; Platform; Attention; Mechanism; Network Effect
JEL Codes: D81, D82, D83

---

[*] Chen: SC Johnson Graduate School of Management, Cornell University (yc2535@cornell.edu). Li: Department of Economics, University of North Carolina (lifei@email.unc.edu). Preuss: SC Johnson Graduate School of Management, Cornell University (preuss@cornell.edu). The order of authors is alphabetical. We are grateful to Nageeb Ali, Simon Anderson, Ian Ball, Arjada Bardhi, Heski Bar-Isaac, Gary Biglaiser, Özlem Bedre-Defolie, Gabriel Carroll, Roberto Corrao, Yeon-Koo Che, Daniel Chen, Dongkyu Chung, Alexandre de Corniere, Tom Conningham, Rahul Deb, Piotr Dworczak, Mira Frick, Yingni Guo, Kevin He, Wei He, Shota Ichihashi, Justin Johnson, Bruno Jullien, Tony Ke, Kyungmin Kim, Chiu-Yu Ko, Jiangtao Li, Yingkai Li, Yunan Li, Jeanine Miklos-Thal, Ellen Muir, Ron Siegel, Yangbo Song, Bruno Strulovici, Tat-How Teh, Guoqiang Tian, Udayan Vaidya, Jiamingmei Wang, Jiaming Wei, Kai Hao Yang, Leeat Yariv, Huseyin Yildrim, Nathan Yoder, Hanzhe Zhang, Jidong Zhou and seminar and conference participants for their helpful comments and valuable discussions. All remaining errors are our own.

1 Introduction

Social media platforms hold a prominent position in the digital economy. More than five billion users worldwide spend over two hours per day consuming user-generated content, including short-form videos, live streams, commentary, tutorials, and entertainment. This massive content consumption is fueled by an unmatched supply of content by digital creators: on TikTok alone, approximately 34 million videos are uploaded each day. Social media platforms capitalize on this popularity by selling their users' attention to advertisers, leading to vast amounts of ads embedded in the users' organic feeds. In 2025, social media advertising spending exceeded USD 200 billion.

To manage user interactions at such scale, major platforms such as TikTok, Instagram, YouTube, or X rely on centralized recommendation algorithms to determine which content and sponsored messages appear in each user's feed.[1] By directing exposure across heterogeneous content and advertisements, these systems govern how attention is generated, allocated, and monetized.

Designing such systems raises important economic questions about user attention allocation. Attention serves two distinct roles: it generates advertising revenue and shapes creators' exposure and incentives to produce content. At the same time, attention is scarce and costly, and attracting it requires engaging content. The platform's recommendation system must therefore balance monetization incentives against the need to provide creators with sufficient exposure to sustain the production of engaging content.

This paper develops a framework to study the optimal allocation of attention on a social media platform. The platform operates as a two-sided marketplace that hosts content across multiple horizontally differentiated categories and serves viewers and creators. Viewers differ in their interests across categories and in their privately known valuation of content quality, and they incur opportunity cost of attention when consuming both content and sponsored advertisements. Creators differ in category and privately known ability, benefit from exposure, and choose costly effort to determine quality. The platform commits to a *recommendation mechanism* that allocates attention across content and advertisements subject to participation and incentive constraints. The platform faces a *trade-off* between harvesting attention through advertising and cultivating it by strengthening production incentives.

We characterize the advertisement-revenue-maximizing allocation of attention. Allocating attention to creators activates network effects by strengthening effort incentives and improving content quality, thereby increasing the monetizable attention base of all interested viewers. Leveraging

---

[1] On Instagram, the Reels and Explore tabs consist exclusively of recommended content. Even the main feed, which before 2022 was reserved for followed accounts, is now algorithmically ranked (Instagram Announcement, 2022). On TikTok, recommended content is even more prominent.

network effects requires reallocating some viewers' attention toward content that they value less since this raises creators' effort and quality, benefiting other viewers who value that content more. Informational frictions on both sides therefore interact, and the cost of using attention to stimulate production depends jointly on viewer preferences and creator abilities.

The optimal design diverts substantial attention away from direct monetization to subsidize creator effort. For low-ability creators, however, generating high-quality content would require too much subsidized attention, so they are excluded from the recommendation system. For included creators, these *attention subsidies* induce some viewers to watch content that is not worth their individual attention cost, with the extent of the subsidy depending on creator ability. That is, the optimal algorithm creates matches with negative virtual surplus. Medium-ability creators receive exposure primarily within their category, whereas sufficiently high-ability creators receive additional cross-category attention when the marginal return to effort is large. Through this careful allocation of attention, the platform endogenously determines the strength and scope of network effects.

In this two-sided market framework, private information on one side amplifies distortions on the other. To limit information rents, the platform allocates attention more selectively toward higher-value viewers and higher-ability creators than under complete information. As a result, participating creators exert less effort and overall viewer engagement declines. Because attention and effort are complementary across the two sides, these distortions propagate through the network: changes in incentives on one side feed back into the other, and distortions extend even to the top of the user-type space on both sides. At the extensive margin, private information raises the bar both for inclusion and for receiving cross-category reach. Hence, whether a category survives or thrives depends not only on the size and value composition of its viewer base, but also on the severity of information asymmetry. Categories with large, high-valuation audiences and limited informational distortions may sustain creators with broad algorithmic reach, whereas categories with small audiences or severe informational frictions receive too little exposure to operate at scale.

Despite the richness of the allocation, the mechanism admits a simple implementation. On the creator side, exposure depends on realized quality, sustaining effort incentives under private information. On the viewer side, the nested structure allows the platform to present a common ordered feed of content and advertisements. Viewers consume sequentially and choose a stopping point that replicates the optimal allocation. We consider this implementation natural as it is consistent with the "infinite scroll" interface dominating almost all social media platforms.

The comparison between complete and private information offers a new perspective on the role of user data in social-media platforms' ability to generate engagement. It highlights a regulatory trade-off between sustaining category diversity and protecting user privacy through policies such as GDPR, CCPA or the recent EU's Digital Markets Act (DMA). When platforms can use rich user data to allocate attention, they can target recommendations more precisely, exploit attention-effort complementarities more effectively, and sustain a broader set of content categories. By contrast, restrictions on the use of user data limit personalization and thus leave participating viewer privacy benefit in the form of information rent. Since it is more costly to provide incentives and leverage network effects, the platform reduces participation on the extensive margin and shrink the scale and diversity of content provision: contents from niche categories are less likely to be recommended, some viewers lose access to content they value, and talented creators in those categories may be unable to find a viable audience.

We then extend the model to allow direct monetary payments to creators. In general, the optimal mechanism may use exposure and payment jointly, and the relative reliance on the two instruments can vary across creator types. Under a mild condition, however, the pattern becomes especially simple. For each category, there is a cutoff ability: below the cutoff, creators are incentivized only through exposure; above the cutoff, exposure is held fixed and additional incentives are provided through positive monetary transfers. Thus, monetization programs are targeted toward high-ability creators precisely when additional exposure is a relatively inefficient way to reward them. This is consistent with the widespread use of multiple creator incentive instruments on platforms such as TikTok and YouTube, which combine algorithmic exposure with direct monetization programs to reward creators.

As the platform earns more from each unit of advertising, for example because of improved ad-targeting technology, richer profiling-based targeting, or greater market power in the advertising market, viewer attention becomes more valuable in direct monetization. The platform therefore has a stronger incentive to preserve attention for advertising and substitutes toward monetary payments when incentivizing creators, which reduces the use of cross-category exposure-based incentives and improves recommendation accuracy.

This comparative static provides a potential rationale for why platforms choose different mixes of creator payments and exposure-based rewards. Platforms with stronger advertising monetization rely more on direct payments and less on distorted cross-category exposure, whereas platforms with weaker advertising monetization rely more heavily on reach-based rewards. It yields a testable implication: improvements in ad monetization should be associated with greater use of creator payment programs and more accurate recommendations.

Finally, it raises a policy trade-off between protecting users from profiling-based advertising and preserving recommendation accuracy. For example, the EU DMA restricts profiling-based advertising based on special categories of personal data, and chokes off cross-app tracking and targeting at major tech platforms, including Meta and Alphabet. These rules are motivated by privacy, safety, and user autonomy, but they may also affect the platform's attention-allocation trade-off: if they reduce the profitability of advertising, exposure becomes a relatively cheaper incentive instrument, potentially increasing the platform's reliance on algorithmic amplification and other recommendation distortions.

Contribution to Literature This paper belongs to the economics of social media platforms, which host user-generated content and monetize attention by selling advertising (see, e.g., Allcott et al. (2020), or Aridor et al. (2024), who provide a comprehensive survey). Much of this literature takes the algorithm as given and studies its welfare, political, and advertising implications. Our contribution is to provide a unified framework where algorithmic attention allocation is the platform's primary incentive instrument for shaping content production and viewer participation.

We are not the first to study how recommendation algorithms influence users incentives. For example, Berman and Katona (2020) demonstrate that improving filtering technology can improve content quality by encouraging users to manually expand their social networks and intensify creator competition. Qian and Jain (2024) argue that recommendation algorithms can mitigate creators' moral hazard problem by threatening to withdraw attention from low-quality content, even at the expense of viewers' ex post surplus. While these studies examine recommendation design and incorporate some related elements in our study, the key trade off of attention allocation in our setting and the corresponding two-sided strategic interaction and endogenous determination of network effects are not their focus. Acemoglu et al. (2024) consider a model with exogenous content and shows how engagement-maximizing network design shapes sharing incentives and can generate echo chambers. Bar-Isaac et al. (2025) adopt a mechanism design approach to study visibility allocation and certification, focusing on how exposure and certification can screen content providers' willingness to pay for visibility. By contrast, we analyze algorithm design in a heterogeneous two-sided environment, where screening on both viewer and creator sides and the scale of network effects are jointly determined.[2]

More broadly, our paper adds to the literature on platform design, which studies platforms as multi-sided marketplaces and analyzes how to set fees and access rules to maximize profits (Jullien et al., 2021). One related strand examines media platforms in which content is produced by the platform itself and revenue is generated through advertising, beginning with Anderson and Coate (2005); see Anderson and Jullien (2015) for a survey and Choi and Jeon (2023); Ichihashi et al. (2025) for recent developments. In these models, platforms either host or directly produce content by themselves to attract viewers and monetize attention through advertising. The strategically active sides are advertisers and consumers, and content provision is typically modeled as centralized rather

---

[2] Beyond algorithmic recommendation design, platforms may also shape content supply through pricing and access rules imposed on content providers. See, for example, Choi et al. (2015) and Casner and Teh (2025), who study how platform pricing structures toward content providers affect entry, participation, and network size.

than decentralized, so creator incentives do not arise as a strategic margin.

A second strand studies product-market platforms that design ranking rules, consideration sets, or search order to influence seller's pricing decision (e.g., Johnson et al. (2023) and Ichihashi and Smolin (2025)). In these models, platform design shapes seller strategies by controlling which sellers consumers observe, effectively using allocation as a screening and incentive device. As a result, platforms may optimally distort ex post efficient allocation to induce truthful revelation or to influence seller action, similar to our use of attention allocation as an incentive instrument for creators. In contrast to both strands, we analyze attention allocation under decentralized content production and focus on the implications of creator–viewer incentive complementarities on platform design.

Also related is the literature on matching design in two-sided markets, where platforms distort matching patterns to extract surplus (e.g., Damiano and Li (2007); Gomes and Pavan (2016, 2024); Fershtman and Pavan (2022)). See Section 6 of Jullien et al. (2021) for a survey. In these models, platforms typically monetize participation through entry fees or transaction-based payments and use matching rules as instruments for price discrimination. We also study how a platform matches content to consumers through attention allocation. However, unlike these matching settings, digital content is non-rival and non-exclusive, and viewer attention rather than content itself is the scarce resource. As a result, allocation in our environment shapes incentive complementarities and the scope of network effects, rather than serving primarily as a rent-extraction device. Also, we study the endogenous content production, which is typically absent from these models.

Finally, our optimal algorithm shares some common features with the literature on mechanism design with externalities. The algorithm addresses consumer's incentive problems arising from private information and free-riding by restricting their access to different types of content, akin to mechanisms that use exclusion to screen willingness to pay in the provision of excludable public goods (e.g., Hellwig (2003), Norman (2004), and Meisner and Pillath (2024)). Also, the interaction of network externalities and private information leads to allocation distortions even at the top, which mirrors a broader insight of the literature studying screening with interdependent payoffs or externalities (e.g., Oren et al. (1982) and Lockwood (2000)).

Organization The rest of the paper is organized as follows. Section 2 sets up the model. Section 3 studies the optimal algorithm in the complete information benchmark. Section 4 characterizes the optimal design under incomplete information and discusses its implementation. Section 5 extends the model by allowing for monetary incentives for creators. Section 6 concludes.

2 Model

We model a social media platform as a monopolistic two-sided online marketplace where users produce and consume digital content, such as articles, music, and videos. The platform employs a personalized recommendation algorithm to distribute user-generated content and sponsored advertisements (ads) in order to maximize advertising revenue.

2.1 Environment

The platform intermediates between two groups of users: content *creators* (he), who incur costs to produce content and benefit from attracting viewer attention, and content *viewers* (she), who spend attention and derive utility from content but not from advertisements. Content is organized into $N > 1$ horizontal *categories*, representing distinct topical domains (e.g., news, entertainment, sports). Let $\mathcal{N} \equiv \{1, \ldots, N\}$.

Content Production A creator is characterized by a vertical type called *ability* $\theta \in \Theta \equiv [0, \overline{\theta}]$, and a horizontal *category* $i \in \mathcal{N}$ which is the only category in which he can produce content. There is a continuum of creators in category i with mass $\alpha_i > 0$. Let $\alpha \equiv \sum_i \alpha_i$ denote the total mass of creators. Within category i, abilities follow a CDF F_i with a continuous PDF f_i. We assume strictly increasing hazard rate $\frac{f_i(\theta)}{1-F_i(\theta)}$. We refer to a creator of ability θ in category i as a (θ, i)-creator, sometimes simply θ-creator when no confusion arises.

A creator exerts effort $e \in \mathbb{R}_+$ to produce one unit of content. At ability level θ, effort increases *quality* of this piece of content according to:

$$q = \theta e.$$

Ability and effort are thus complementary: higher ability enables a higher marginal return to effort. The cost of effort is linear and normalized to e.

A creator derives utility from the total mass of viewer attention received. Receiving $X \in \mathbb{R}_+$ mass of attention while exerting effort e, a creator's net payoff is:

$$u(X) - e,$$

where $u(0) = 0$ and $u'(\cdot) > 0$. The attention-driven utility $u(X)$ captures creators' desire for reach and influence, which may reflect psychological satisfaction (Kranton and McAdams, 2024; Filippas et al., 2025), career concerns (Petrova et al., 2021) or external monetization opportunities, i.e., payments outside the platform.[3] In Sections 3 and 4, we focus on the main model where any utility creators receive stems directly or indirectly from attention. Section 5 allows creators to receive payment from the platform.

Sponsored Advertising The platform profits from viewer exposure to *ads* that are displayed alongside content. We assume the platform earns a constant marginal revenue $z > 0$ per unit of attention allocated to advertisements. This reduced-form revenue can be interpreted as the equilibrium per-impression price generated by an underlying auction among advertisers.

Content Consumption A viewer is characterized by a vertical type called *valuation* $v \in V \equiv [\underline{v}, \overline{v}]$ with $\underline{v} > 0$, and a horizontal category $j \in \mathcal{N}$ which is the only category *relevant* to her (generating positive utility). There is a continuum of viewers in category j with mass $\beta_j > 0$. Let $\beta \equiv \sum_j \beta_j$ denote the total mass of viewers. Within category j, valuations follow a CDF G_j with a continuous PDF g_j. We assume strictly increasing hazard rate $\frac{g_j(v)}{1-G_j(v)}$. We refer to a viewer of valuation v in category j as a (v, j)-viewer, sometimes simply v-viewer when no confusion arises.

Paying attention is costly to viewers: each unit of attention spent on content or ads costs a constant $c > 0$, regardless of its use. This *attention cost* captures the opportunity cost of time and potential "nuisance cost" of advertising that is often mentioned in the literature (Anderson and Coate, 2005).

A (v, j)-viewer, when spending a unit of attention on some content of quality q in category i, obtains net payoff from this unit:[4]

$$vq\mathbb{1}\{i = j\} - c.$$

When the viewer instead spends a unit of attention on ads, we normalize the intrinsic utility to zero, so that the net payoff is simply $-c$. We assume that both content and ads are *experience goods*, in the sense that viewers derive payoffs only after already spending the attention cost and therefore cannot cherry-pick. We will discuss this assumption in Section 2.3.

Information Structure The platform observes users' horizontal categories i, j but not their vertical types θ, v. We make these assumptions because a creator's horizontal category is the topic on which they post content, which machine learning techniques can easily identify. The platform

---

[3] For simplicity, we assume creators value total attention regardless of its source. What matters is that both within- and cross-category attention benefit creators, even at different rates; see Section 2.3.

[4] Allowing heterogeneous attention costs does not change the analysis, since participation depends only on the ratio v/c. We therefore normalize c across viewers and attribute heterogeneity to v.

can also easily learn a viewer's horizontal category (favorite topic) after only a few visits by simply tracking the user's engagement with content from different topics. On the other hand, a creator's ability and a viewer's valuation for quality are considerably more difficult to infer. Although social media platforms collect extensive behavioral data, these data cannot fully capture the residual heterogeneity in viewers' intrinsic preferences or the transient shocks—such as mood, attention, or available leisure time—that affect their willingness to engage with content. Similarly, a creator's effective productivity may fluctuate over time due to changing ideas, effort, or external circumstances, preserving an element of private information even in the long run. For example, many creators post content from their daily lives. When they are traveling, it may be easier for them to post exciting (high quality) content. Finally, the private-information model allows us to study how legal and regulatory constraints on personalization affect equilibrium outcomes. Even when the platform observes characteristics that are informative about users' types, such constraints may restrict its ability to condition recommendations on them.

We assume that the platform observes the realized quality $q = \theta e$ of each piece of content but cannot identify the components θ and e. This mirrors the real-world practice: quickly conducting experiments of various scales on randomly selected viewers and inferring content performance from engagement data.

2.2 Algorithm Design Problem

Viewers and creators do not interact directly. Instead, the platform intermediates both sides through a recommendation algorithm that allocates viewers' attention across content of different categories and quality levels, as well as advertisements. By controlling how attention is allocated, the algorithm essentially transforms digital content into a *club good* — excludable (since access can be selectively restricted) yet non-rival (since one viewer's engagement does not diminish others'). Throughout, we use the term "algorithm" to refer to the induced allocation rule rather than the computational procedure that implements it.

Examples Before moving further, we discuss users' strategic interaction under a few stylized recommendation algorithms. The purpose is to show how attention allocation governs and links users' incentives on both sides.

First, consider an algorithm that recommends to each viewer only content from the category she is interested in, regardless of the quality. At the same time, to generate advertising revenue, the platform mixes some ads in the recommended feed. While such an algorithm maximizes relevance from the viewer's perspective, it provides no incentive for creators to exert effort, and the equilibrium quality collapses. Anticipating the lack of quality, viewers anticipate no utility from either content or ads, and thus choose not to engage.

Next, suppose the algorithm still recommends only relevant content but imposes a minimum quality threshold for exposure. This policy partially restores incentives on the supply side as high-ability creators exert just enough effort to meet the threshold, while low-ability creators opt out. Now with meaningful qualities on the platform, high-valuation viewers choose to consume the content-ads mix. In turn, viewers' engagement also encourages more creators to participate. Network effects arise in this case, albeit confined to individual categories. By adjusting the quality threshold, the algorithm endogenously determines the magnitude of these within-category network effects.

Finally, consider an algorithm that deliberately introduces biased recommendations. It recommends all relevant content above a minimum quality, while also promoting irrelevant content to viewers if its quality exceeds a higher threshold. Creators now expect an even larger audience at high quality levels, and those with high ability will exert more effort. While such irrelevant content lowers the utility of viewers, the heightened effort from creators indirectly benefits those viewers who *are* interested in those categories.

The last case features cross-category spillovers generated by the algorithm: a viewer's attention affects the welfare of viewers in other categories. The key is that attention from both within- and cross-category viewers is valued by creators and can therefore be used to provide incentives. This example illustrates how, by precisely controlling who is exposed to whose content, the algorithm endogenously determines the scope and magnitude of network effects on a social media platform. In other words, both the direction and the strength of network effects are artifacts of design, which can be leveraged to spur content production and ultimately increase advertising revenue.

Recommendation Algorithm The preceding examples are useful for building intuition, but they represent only special cases of feasible recommendation rules. In practice, the platform implements personalized recommendations that condition on users' characteristics, selectively leveraging network effects. To characterize the optimal algorithm without loss of generality, we invoke the Revelation Principle (Myerson, 1986) and study the platform's problem over a broad class of direct mechanisms. An algorithm is formulated as a (direct) mechanism that elicits users' information and distributes content and ads, while respecting users' incentive compatibility and individual rationality. Formally, an algorithm $\mathscr{A} \equiv \left(\{e_i\}_{i\in\mathcal{N}}, \{x_{ij}\}_{i,j\in\mathcal{N}}, \{a_j\}_{j\in\mathcal{N}}\right)$ consists of three sets of functions:

$$
\begin{aligned}
&e_i : \Theta \to \mathbb{R}_+,\\
&x_{ij} : \Theta \times V \to [0,1],\\
&a_j : V \to \mathbb{R}_+,
\end{aligned}
$$

where, given reported types, $e_i(\theta)$ denotes the recommended effort for a (θ, i)-creator, $x_{ij}(\theta, v)$ the probability that a (v, j)-viewer receives recommendation on content produced by a (θ, i)-creator, and $a_j(v)$ the total mass of ads mixed in the recommendation bundle for a (v, j)-viewer. For notational simplicity, define:

$$
X_i(\theta) \equiv \sum_j \beta_j \int_V x_{ij}(\theta, v)\mathrm{d}G_j(v) \tag{1}
$$

as the total amount of attention a (θ, i)-creator receives, determined by $\{x_{ij}\}_{i,j}$. Then *incentive compatibility* (IC) requires, for any θ, θ', i:

$$
u(X_i(\theta)) - e_i(\theta) \geqslant u(X_i(\theta')) - \frac{\theta' e_i(\theta')}{\theta}, \tag{2}
$$

and for any v, v', j:

$$
\begin{aligned}
&\sum_i \alpha_i \int_\Theta x_{ij}(\theta, v)\Big(v\theta e_i(\theta)\mathbb{1}\{i = j\} - c\Big)\mathrm{d}F_i(\theta) - c a_j(v)\\
\geqslant\ &\sum_i \alpha_i \int_\Theta x_{ij}(\theta, v')\Big(v\theta e_i(\theta)\mathbb{1}\{i = j\} - c\Big)\mathrm{d}F_i(\theta) - c a_j(v').
\end{aligned} \tag{3}
$$

Intuitively, constraint (2) says that a (θ, i)-creator, who can produce content with quality $\theta' e_i(\theta')$ and receive attention according to $\{x_{ij}(\theta', v)\}_{v,j}$, finds it optimal to truthfully report $\theta' = \theta$ and follow the recommended effort $e_i(\theta)$. Similarly, constraint (3) states that a (v, j)-viewer must prefer recommendation $\{x_{ij}(\theta, v)\}_{\theta,i}$ mixed with $a_j(v)$ amount of ads to the recommendation $\{x_{ij}(\theta, v')\}_{\theta,i}$ mixed with $a_j(v')$ amount of ads, intended for any other v'. *Individual rationality* (IR) requires that any creator and viewer must receive payoff greater than their outside option, normalized to zero:

$$
u(X_i(\theta)) - e_i(\theta) \geqslant 0,\ \forall\, \theta, i, \tag{4}
$$

$$
\sum_i \alpha_i \int_\Theta x_{ij}(\theta, v)\Big(v\theta e_i(\theta)\mathbb{1}\{i = j\} - c\Big)\mathrm{d}F_i(\theta) - c a_j(v) \geqslant 0,\ \forall\, v, j. \tag{5}
$$

The platform maximizes the income from the ads market, that is:

$$
\max_{\mathscr{A}}\ z \sum_j \beta_j \int_V a_j(v)\mathrm{d}G_j(v) \qquad \text{s.t.} \quad (2), (3), (4), (5).
$$

2.3 Model Discussion

We discuss some simplifying assumptions made in the model setup that allow us to focus on the core insight.

First, in a viewer's feed (recommendation), ads and various content are seamlessly blended like coffee and milk, so that viewers cannot selectively pick only high-quality relevant content without wasting attention cost on others. This assumption is common in the literature (see, e.g., de Corniere and Sarvary (2023)) and is particularly appropriate for static content such as short text or photos, where by the time viewers determine whether they like the content, the attention has already been spent. In this sense, evaluation and consumption are nearly indistinguishable. Longer videos or texts differ somewhat because viewers may attempt to screen content after a few seconds. However, such screening is imperfect as "clickbait" is pervasive. Moreover, some platforms require viewers to watch content or advertisements for a minimum duration before skipping: skippable ads on YouTube and in-feed ads on TikTok. As long as viewers cannot screen content without attention cost, the no-cherry-picking assumption is a reasonable approximation.[5]

Second, we assume that creators care only about total exposure $X_i(\theta)$ and do not distinguish across audiences. This assumption is in the same spirit as our no-cherry-picking simplification. Because content is an experience good, a viewer must spend *some* attention on the content before learning its worth. Such engagement already generates *some* exposure, contributing to the creator's psychological gratification and monetizable traffic. More generally, we can allow a category-i creator to benefit differently on within-category versus cross-category exposure:

$$u\left(\beta_i \int_V x_{ii}(\theta, v)\,\mathrm{d}G_i(v) + \delta \sum_{j\neq i} \beta_j \int_V x_{ij}(\theta, v)\,\mathrm{d}G_j(v)\right),$$

where $\delta > 0$ captures the relative value of cross-category attention.

Third, we assume that viewers sharply distinguish between relevant and irrelevant content: any positive-quality content in a viewer's preferred category generates positive utility (i.e., $\underline{v} > 0$), whereas cross-category content and advertisements generate zero payoff regardless of quality. This parsimoniously captures heterogeneity in preferences over digital content on the platform, including both organic and sponsored material. Under this formulation, a viewer's private information is one-dimensional—the value she assigns to content in her preferred category—which keeps the mechanism design problem tractable. This provides a conservative benchmark for our results on

---

[5] Suppose the platform shows a viewer a mass $a > 0$ of ads if the viewer cannot cherry pick at all. Now suppose the user can detect and skip ads with half the attention cost, then the platform can just raise the mass of ads to $2a$ to achieve the same outcome.

algorithmic exposure: if viewers obtained positive, quality-dependent utility from cross-category content, or simply preferred such content to advertisements, generating cross-category reach would be less costly for the platform, strengthening the attention distortions we characterize. Thus, our framework isolates the incentive motive by treating irrelevant content as a "tax" on viewer attention equivalent to advertising.

Fourth, we leave viewers' attention capacity unbounded. This choice reflects our focus on environments in which time spent on social media is well below any physical capacity constraint, so the relevant margin is the opportunity cost of time rather than a hard cap. Leaving attention uncapped also allows the platform's extractable attention per viewer to be determined endogenously; imposing an upper bound would mechanically fix this margin and obscure the trade-off between monetization and the expansion of monetizable attention. Technically, introducing attention capacity complicates screening in a way analogous to mechanism design with financially constrained agents (see, e.g., Pai and Vohra (2014); Li (2021)).

Lastly, in our model, viewers can only read within their tailored algorithm recommendation. This does not rule out viewers following creators of their choice; rather, the constraint is that viewers cannot read "followed content" exclusively. In reality, recommended content represents the lion's share of what people consume on social media platforms such as TikTok and Instagram, while "followed content" is in decline. However, even where the "follow" function still exists, viewers intending to consume "followed content" are nevertheless exposed to ads and potentially even other recommended content. After all, the platform has complete control over what viewers can and cannot see.[6] In sum, the platforms realize the profitability of strengthening the algorithm and taking control of the attention flow. As such, we omit the "follow" function to simplify analysis.

3 Complete Information

As a benchmark, suppose the platform has complete information about both creators' and viewers' types. This exercise allows us to set aside the screening issues and focus on how attention allocation endogenously shapes the network effect and thereby maximizes advertising revenue.

In this case, the platform's problem can be written as follows:

$$\max_{\mathscr{A}} \; z \sum_j \beta_j \int_V a_j(v) \mathrm{d}G_j(v) \qquad \text{s.t.} \qquad (4), (5).$$

The algorithm serves as a recommendation mechanism in which obedience constraints ensure that

---

[6] On Quora and Reddit, ads and recommended threads are interspersed with pertinent content.

users voluntarily follow the prescribed attention and effort choices.

3.1 Simplifying the Design Problem

We begin with two observations of constraints (4) and (5). First, an optimal mechanism does not leave a positive measure of viewers' participation constraints slack; otherwise, it could profitably increase their ads exposure. Second, for any (θ, i)-creator, it is without loss of generality to bind constraint (4); otherwise, the algorithm can extract higher $e_i(\theta)$ to further entertain the viewers who watch and value this creator's content, and ultimately squeeze in more ads. These insights are summarized in the following lemma.

Lemma 1 (Binding Constraints)
Under complete information, there exists an optimal algorithm such that:

(i) the creator participation constraint (4) *binds for every* (θ, i)*;*

(ii) the viewer participation constraint (5) *binds for every* (v, j).

Now we substitute out $a_j(v)$ using the viewers' binding IR constraints, omit the constant multiplier $\frac{z}{c}$, and rewrite the problem as:

$$\max_{\substack{\{e_i(\cdot)\}_i,\\ \{x_{ij}(\cdot,\cdot)\}_{i,j}}} \sum_{i,j} \alpha_i \beta_j \int_\Theta \int_V x_{ij}(\theta, v) \Big(\underbrace{v\theta e_i(\theta) \mathbb{1}\{i = j\} - c}_{\text{Surplus}} \Big) \mathrm{d}G_j(v) \mathrm{d}F_i(\theta) \tag{6}$$

$$\text{s.t. } e_i(\theta) = u(X_i(\theta)), \ \forall\, \theta, i, \tag{IR-C}$$

The objective in problem (6) is the total viewer surplus from watching content only, which is then extracted by the platform using ads. For each combination of (θ, i, v, j), $v\theta e_i(\theta)\mathbb{1}\{i = j\} - c$ represents viewer (v, j)'s net payoff from watching creator (θ, i)'s content. A positive net payoff leaves room for the platform to show more ads. Constraint (IR-C) pins down a creator's effort. Rigorously speaking, this is a relaxed problem of the platform's original optimization. To ensure that the solution to the relaxed problem satisfies the non-negativity constraint $a_j(v) \geqslant 0$ for advertising in the original problem, we require:

$$\sum_i \alpha_i \int_\Theta x_{ij}(\theta, v) \Big(v\theta e_i(\theta) \mathbb{1}\{i = j\} - c \Big) \mathrm{d}F_i(\theta) \geqslant 0, \ \forall\, v, j. \tag{NN-C}$$

This is necessary because, after the replacement, $a_j(v)$ is no longer an explicit variable in the optimization. In the remainder of the paper, we assume c is sufficiently small so that this constraint never binds.

To understand the platform's trade off in attention allocation, we examine the marginal impact of $x_{ij}(\theta, v)$ on the platform's profit. Substituting constraint (IR-C) into the objective of problem (6), we decompose the marginal effect of $x_{ij}(\theta, v)$ into two terms below.

Definition 1 (Network-Effect-Adjusted Surplus)
In the optimal algorithm, the value of $x_{ij}(\theta, v)$ *is determined by the sign of*

$$\underbrace{v\theta e_i(\theta)\mathbb{1}\{i=j\} - c}_{\text{Local Surplus}} + \underbrace{u'(X_i(\theta))\theta\beta_i \int_V x_{ii}(\theta, v')v'\mathrm{d}G_i(v')}_{\text{Network Effect}} . \tag{7}$$

In particular, $x_{ij}(\theta, v) = 1$ *if expression* (7) *is positive,* $x_{ij}(\theta, v) = 0$ *if it is negative, and any value in* $[0, 1]$ *is optimal if it is zero.*

The derivative in expression (7) captures two effects when recommending type-(θ, i) creator's content to viewer (v, j). The first term captures the direct pointwise surplus of (v, j)-viewer from watching content produced by (θ, i)-creator. This effect is *local*: it depends only on the creator-viewer pair in question and is independent of other allocations. When positive, this surplus relaxes the participation constraint of viewer (v, j) and thus allows for more inserted ads. The second term captures the indirect network effect. By assigning attention $x_{ij}(\theta, v) = 1$, the creator (θ, i) receives $u'(X_i(\theta))$ marginal utility, which in turn allows the platform to extract more effort. The increased effort then boosts the quality of content by a factor of θ and benefits all viewers in the same category who spend attention on it, eventually enabling more ads to be blended in for these viewers. Due to this non-negative network effect, attention can improve profit even if the local surplus alone is negative, provided the production-side externality is strong enough. In sum, pointwise *local* surplus maximization is no longer the right principle to follow.

The optimal algorithm is challenging to solve because problem (6) is inherently high-dimensional: with n categories, the algorithm consists of $n^2 + 2n$ functions. Moreover, due to network effects, categories are interlinked, and one cannot simply maximize viewer surplus within each category in isolation. Fortunately, we can reduce the dimensionality of the platform's problem by exploiting the symmetry of viewer preferences and the separability of the objective function.

First, consider cross-category attention x_{ij} where $j \neq i$. Because these viewers derive a constant net utility of $-c$ from such irrelevant content, and because digital content is non-rival, the platform's profit and the creator's incentives are independent from the specific composition of irrelevant audience. We can therefore aggregate the attention from all those “mismatched” viewers into a single variable called creator (θ, i)’s *extended reach*, the mass of attention from outside category i:

$$R_i(\theta) \equiv \sum_{j\neq i} \beta_j \int_V x_{ij}(\theta, v)\mathrm{d}G_j(v) \in [0, \beta - \beta_i], \tag{8}$$

so that total attention received can be recast in a dichotomous way as:

$$X_i(\theta) = \beta_i \int_V x_{ii}(\theta, v)\mathrm{d}G_j(v) + R_i(\theta).$$

Second, the constant marginal cost of attention implies that the cost of incentivizing one creator is separable from the cost of incentivizing another. Since the platform extracts the full viewer surplus via binding participation constraints, the total ads profit is proportional to the sum of the viewer surplus generated by creators in each category. Thus, the aggregate maximization problem decomposes into n additive sub-problems, one for each creator category. We can therefore characterize the optimal schedule $\{e_i(\cdot), x_{ii}(\cdot), R_i(\cdot)\}$ for each category i independently, as stated in Lemma 2.

Lemma 2 (Separability)
Consider the sub-problem for each category $i \in \mathcal{N}$,

$$\begin{aligned} \max_{e_i, x_{ii}, R_i} \quad & \beta_i \int_\Theta \int_V x_{ii}(\theta, v)\Big(v\theta e_i(\theta) - c\Big)\mathrm{d}G_i(v)\mathrm{d}F_i(\theta) - c\int_\Theta R_i(\theta)\mathrm{d}F_i(\theta) \\ \text{s.t.} \quad & e_i(\theta) = u\left(\beta_i \int_V x_{ii}(\theta, v)\mathrm{d}G_i(v) + R_i(\theta)\right), \ \forall\, \theta, i \end{aligned} \tag{9}$$

and denote the solution $\{\hat{e}_i, \hat{x}_{ii}, \hat{R}_i\}_i$. There exists $\{e_i^*, x_{ij}^*\}_{i,j}$ such that

$$e_i^*(\theta) = \hat{e}_i(\theta), \quad x_{ii}^*(\theta, v) = \hat{x}_{ii}(\theta, v), \quad \sum_{j\neq i} \beta_j \int_V x_{ij}^*(\theta, v)\mathrm{d}G_j(v) = \hat{R}_i(\theta), \forall\, \theta, i, v,$$

and solves problem (6).

Although summing the objective functions in (9) across categories recovers the platform’s total viewer surplus, the objective for a given category i does not correspond to viewer surplus in category i. Rather, it equals all viewers’ total surplus from consuming content generated by creator category i. The first term captures the surplus of within-category viewers, while the second term accounts for that of cross-category viewers (incurring cost c without deriving utility). This representation disentangles the roles of e_i, x_{ii}, and R_i in the objective function and the constraints, rendering the original optimization problem additively separable across categories.

Notice that the extended reach does not uniquely determine the underlying attention allocation $\{x^*_{ij}\}_{j\neq i}$, as what matters for optimality is the induced aggregate exposure $R^*_i(\theta)$ of each creator. Any $\{x^*_{ij}\}_{j\neq i}$ that generates the same $R^*_i(\theta)$ yields the same incentives and the same objective value, and is therefore optimal. As a result, the lemma characterizes a class of optimal policies rather than a unique attention assignment across categories.

3.2 Optimal Design under Complete Information

We now characterize the optimal allocation when recommendations condition on users' categories and vertical types. A key insight from problem (9) is that within-category attention x_{ii} and creator effort e_i are complementary. If a creator reaches only a subset of viewers in his category, the participation constraint limits his effort and thus content quality. Expanding within-category attention relaxes this constraint, inducing higher effort and improving quality. Importantly, the resulting quality gains benefit not only newly reached viewers but also those already paying attention. This positive feedback between attention and effort generates increasing returns to within-category attention.

Proposition 1 (Algorithm: Complete Information)
For each category i, the optimal algorithm is characterized by an increasing total attention function $X_i(\cdot)$ such that:

$$x_{ii}(\theta, v) = \mathbb{1}\left\{G_i(v) \geqslant 1 - \tfrac{X_i(\theta)}{\beta_i}\right\}, \quad R_i(\theta) = (X_i(\theta) - \beta_i)^+, \quad e_i(\theta) = u(X_i(\theta)).$$

Moreover, there exist cutoff abilities $0 < \theta^l_i \leqslant \theta^m_i \leqslant \theta^h_i$ such that:

$$X_i(\theta)\begin{cases} = 0 & \text{if } \theta < \theta^l_i, \\ \in (0, \beta_i] & \text{if } \theta^l_i \leqslant \theta \leqslant \theta^m_i, \\ \in (\beta_i, \beta) & \text{if } \theta^m_i < \theta < \theta^h_i, \\ = \beta & \text{if } \theta \geqslant \theta^h_i. \end{cases}$$

Proposition 1 shows that the optimal algorithm can be summarized by a single increasing total-attention function, which specifies how much attention each creator type receives on the path. The function is illustrated in Figure 1(a). Given this required amount of attention, the platform sources viewers' attention according to a *pecking order*: creators with ability just above the inclusion

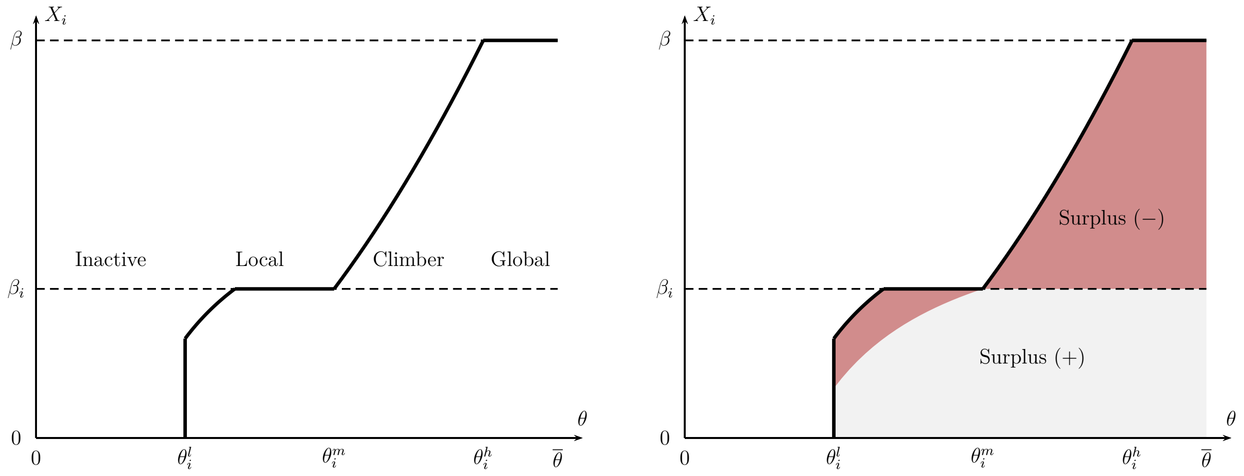

(a) Attention (b) Surplus

Figure 1. Optimal Allocation and Viewer Surplus in Category i.

(a) The total attention X_i as a function of θ. When below β_i, it consists of only with-category attention. The part above β_i is the extended reach R_i. (b) The gray area denotes total attention allocation if only positive-surplus attention is recommended, and the red area represents negative-surplus attention that should be recommended after accounting for network effects.

threshold receive attention only from high-value viewers in their own categories. As creator ability rises, the algorithm expands their reach to lower-value within-category viewers. Only sufficiently high-ability creators receive cross-category reach. Regardless of the creator's ability, the effort is designed to deplete creator surplus due to the binding participation constraint. This characterization naturally segments creators by ability:

- *Inactive*: types $\theta < \theta_i^l$, with zero effort and zero attention;
- *Local Entertainers*: types $\theta_i^l \leqslant \theta \leqslant \theta_i^m$, with positive increasing effort and attention only from own category;
- *Ladder Climbers*: types $\theta_i^m < \theta < \theta_i^h$, with positive and increasing effort, full attention from own category, and increasing but partial reach outside own category;
- *Global Influencers*: types $\theta \geqslant \theta_i^h$, with constant effort and full attention from all viewers.

We now unpack the four creator segments induced by the optimal algorithm. First, the optimal algorithm gives up creators with abilities below θ_i^l. In particular, a creator of ability θ is worth including only if there exists some attention level X_i such that

$$\theta u(X_i) X_i \mathbb{E}\left[v \middle| G_i(v) \geqslant 1 - \frac{X_i}{\beta_i} \right] \geq c X_i. \tag{10}$$

The left-hand side is the total benefit from allocating X_i units of within-category attention to this creator: attention induces effort, effort raises content quality, and the resulting content generates viewer utility and hence additional attention. The right-hand side is the corresponding opportunity cost, namely, the monetization value of the same attention elsewhere on the platform. The cutoff θ_i^l is therefore the lowest ability for which this condition holds for some $X_i \in [0, \beta]$. The inequality also makes clear that the cutoff θ_i^l is lower when category i has a larger viewer base β_i, or when the valuation distribution G_i becomes more favorable. Intuitively, when a category contains a larger mass of high-value viewers, the creator's effort is more effective generating values for the viewers and thus profits for the platform.

Second, total attention $X_i(\theta)$ exhibits *discontinuity* at θ_i^l: if a creator receives attention at all, it must be bounded away from zero. This jump reflects the complementarity between creator effort and viewer attention within the same category: recommending content to infinitesimal mass of viewers would generate vanishing utility for the creator and hence induces vanishing effort, and the resulting low quality in turn cannot justify viewers' attention. Formally, in (10), the benefit from including a θ-creator is $O(X_i^2)$, whereas the opportunity cost is $O(X_i)$. Hence, for sufficiently small X_i, the benefit is dominated by the cost, and the total attention, if allocated at all, must be bounded away from zero.

Third, the pecking-order principle determines how attention is allocated among active creators. Viewer attention differs in its opportunity cost as an incentive instrument. For a given creator, high-valuation viewers in his own category gain the most from his content, lower-valuation viewers in the same category gain less, and cross-category viewers receive no direct consumption value. The optimal algorithm therefore grants attention according to this ranking.

Specifically, a category-i creator with ability $\theta \leqslant \theta_i^m$ receives attention $X_i(\theta) \leqslant \beta_i$; thus, it suffices to allocate attention only from high-valuation viewers within his own category, i.e., $G_i(v) \geqslant 1 - \frac{X_i(\theta)}{\beta_i}$. As ability rises, the creator's audience expands to include progressively lower-valuation viewers in that category. This generates a *nested* within-category allocation: higher-ability creators reach larger within-category audiences, and higher-valuation viewers watch larger set of same-category contents.

For sufficiently high-ability creators, within-category attention alone is no longer sufficient to provide the desired incentives. The platform therefore provides additional incentives through extended reach, $R_i(\theta) = X_i(\theta) - \beta_i$, represented in Figure 1(a) by the vertical distance between $X_i(\theta)$ and β_i. Additional exposure raises effort, and higher ability makes the induced effort increase especially valuable. Thus $R_i(\theta)$ increases with θ and reaches its maximum at $\theta \geqslant \theta_i^h$.

Importantly, the optimal algorithm directs viewers to pay *excessive* attention to content with *negative pointwise surplus*, i.e., $v\theta e_i(\theta)\mathbb{1}\{i = j\} - c < 0$. In Figure 1(b), the red area shows the excessive attention that seemingly generates negative surplus. From the platform's perspective, such distortions on the viewer side function as *subsidies* to creators. The optimal mechanism subsidizes almost every participating creator to leverage both within- and cross-category network effects, and those subsidies are discriminatory across creators. When a creator barely qualifies for production, positive surplus is achieved only for very high-valued within-category viewers. Yet the algorithm subsidizes him by recommending the content to some lower-valued viewers within the category to take advantage of within-category network effects. As ability increases, higher content quality reduces the need for within-category subsidies, which eventually become unnecessary. Once ability becomes sufficiently high ($\theta > \theta_i^m$), the platform further subsidizes production by utilizing cross-category network effects, and the subsidy grows with ability. Overall, the magnitude of the subsidy is non-monotonic in creator ability, and the source of the subsidy differs across creator types.

Proposition 1 implies the increasing pattern of effort and quality in ability. Under complete information, the creator's participation constraint binds, so effort fully dissipates the utility from attention. Since total attention is increasing in ability, so is effort. As a consequence, higher-ability creators produce strictly higher-quality content.

Under stronger assumptions, the optimal algorithm admits a sharper characterization, making the role of network effects especially transparent.

Corollary 1 *If u is strictly concave and $u(X)\ \mathbb{E}_i\left[v \middle| G_i(v) \geqslant 1 - \frac{X}{\beta_i}\right]$ is strictly increasing for all $X \in (0, \beta_i]$, then:*

$$X_i(\theta) = \begin{cases} \beta_i & \text{if } \theta_i^l \leqslant \theta \leqslant \theta_i^m, \\ u'^{-1}\left(\frac{c}{\theta\beta_i\mathbb{E}_i[v]}\right) & \text{if } \theta_i^m < \theta < \theta_i^h, \end{cases}$$

where $\theta_i^l = \frac{c}{u(\beta_i)\mathbb{E}_i[v]}$, $\theta_i^m = \frac{c}{\beta_i u'(\beta_i)\mathbb{E}_i[v]}$, and $\theta_i^h = \frac{c}{\beta_i u'(\beta)\mathbb{E}_i[v]}$.

The monotonicity on $u(X)\ \mathbb{E}_i\left[v \middle| G_i(v) \geqslant 1 - \frac{X}{\beta_i}\right]$ requires that as a piece of content is recommended to more viewers, the increase of creator's utility must compensate the declining *average* valuation among viewers. Then, the marginal incentive effect of additional within-category attention always dominates its marginal cost. This generates global increasing returns with respect to within-category attention, so the optimal allocation is bang-bang: if a creator's ability exceeds the inclusion cutoff, the optimal allocation expands his exposure until all viewers in the category are included and the maximum feasible within-category network effects are activated. This *all-or-nothing* structure yields a closed-form expression for θ_i^l, obtained by imposing equality in condition (10) at $X_i = \beta_i$.

The closed-form cutoffs also give simple conditions under which a category can receive some attention: some creator in the category can be included only if the inclusion cutoff is no larger than the highest creator ability, or

$$\mathbb{E}_i[v]u(\beta_i) \geqslant \tfrac{c}{\theta},$$

which requires the category to have a sufficiently large viewer base and a sufficiently high average viewer value. For each category, the formula gives the *critical mass* of viewers required for the category to remain active, taking as given its value distribution. This observation also clarifies the role of viewer-base expansion. Beyond increasing traffic in already active categories, a larger viewer base can bring marginal categories above their survival thresholds and expand the variety of categories on the platform. The same logic determines whether a category can generate influencers with cross-category exposure.

Since u is concave, when $\theta_i^m < \theta < \theta_i^h$, the optimal $X_i(\theta)$ is solved from the first-order condition:

$$-c + u'\big(X_i(\theta)\big)\theta\beta_i\mathbb{E}_i[v] = 0. \tag{11}$$

This condition equates the platform's marginal cost of cross-category attention to its marginal incentive benefit. On the cost side, directing a cross-category viewer to the creator displaces attention that could otherwise be monetized. On the benefit side, the additional attention raises the creator's effort, and hence content quality, which benefits the creator's within-category audience. In this case, the pecking order property yields the cross-category reach:

$$R_i(\theta) = u'^{-1}\left(\frac{c}{\theta\beta_i\mathbb{E}_i[v]}\right) - \beta_i.$$

The cutoffs θ_i^m and θ_i^h can be solved from the two boundary cases of (11): θ_i^m is the type at which cross-category reach first becomes positive, $X_i(\theta) = \beta_i$, and θ_i^h is the type at which cross-category reach reaches its maximum, $X_i(\theta) = \beta$. These cutoffs are lower when category i has a larger or more valuable viewer base, because cross-category attention generates stronger incentive benefits when the resulting quality improvement is enjoyed by more, or higher-value, within-category viewers. For such categories, the platform therefore extends the use of cross-category network effects to a broader range of creator types.

4 Incomplete Information

We now return to the design under incomplete information. We study how screening on both sides of the market interacts with attention allocation in the presence of network effects. We first simplify the incentive compatibility constraints using a Myersonian approach and then characterize the resulting optimization problem. We next examine how private information on both sides constrains the platform's ability to manage network effects for advertising monetization, and the resulting implications for user participation. Finally, we discuss how to implement the optimal direct mechanism.

4.1 Simplifying the Design Problem

We begin with characterizing the set of feasible direct mechanisms. Fix a direct mechanism $\mathscr{A}$, write:

$$U_i(\theta) \equiv u(X_i(\theta)) - e_i(\theta)$$

as (θ, i)-creator's indirect utility under truthful report, and write

$$W_j(v) \equiv \sum_i \alpha_i \int_\Theta x_{ij}(\theta, v)\Big(v\theta e_i(\theta)\mathbb{1}\{i = j\} - c\Big)\mathrm{d}F_i(\theta) - ca_j(v)$$

as (v, j)-viewer's indirect utility under truthful report. The standard Myersonian approach provides the necessary and sufficient conditions for a direct mechanism to be incentive compatible:

Lemma 3 (User IC)
Creators' IC constraint, (2)*, is satisfied if and only if for all* θ, i*:*

- $q_i(\theta) = \theta e_i(\theta)$ *is increasing in* θ*, and*
- $U_i(\theta) = U_i(0) + \int_0^\theta \frac{q_i(\tilde{\theta})}{\tilde{\theta}^2}\mathrm{d}\tilde{\theta}$.

Viewers' IC constraint, (3)*, is satisfied if and only if for all* v, j*:*

- $Q_j(v) \equiv \alpha_j \int_\Theta x_{jj}(\theta, v)q_j(\theta)\mathrm{d}F_j(\theta)$ *is increasing in* v*, and*
- $W_j(v) = W_j(\underline{v}) + \int_{\underline{v}}^v Q_j(\tilde{v})\mathrm{d}\tilde{v}$.

In order for creators and viewers to truthfully report their types θ and v, exact amount of information rent must be granted. This limits the platform's ability to extract effort from creators and insert ads to viewers. Also, incentive compatibility requires the equilibrium quality $q_i(\theta)$ to be increasing in θ, and the quality-weighted total attention $Q_j(v)$ to be increasing in v.

Define:

$$\Psi_j(v) \equiv v - \frac{1 - G_j(v)}{g_j(v)}, \quad \Phi_i(\theta) \equiv \theta - \frac{1 - F_i(\theta)}{f_i(\theta)}$$

as the *virtual types* of a (v, j)-viewer and a (θ, i)-creator, respectively. Following the standard procedure, we set the lowest type user's payoff to be zero for optimality. By Lemma 3, all users' payoffs are monotone in their vertical types, so all user types' IR are automatically satisfied. Using the users' IC conditions, we can replace $a_j(v)$ in the objective and simplify the platform's problem to the following:

$$\max_{\substack{\{e_i(\cdot)\}_i,\\ \{x_{ij}(\cdot,\cdot)\}_{i,j}}} \sum_{i,j} \alpha_i\beta_j \int_{\Theta}\int_V x_{ij}(\theta,v)\Big(\underbrace{\Psi_j(v)\theta e_i(\theta)\mathbb{1}\{i=j\}-c}_{\text{Virtual surplus}}\Big)\mathrm{d}G_j(v)\mathrm{d}F_i(\theta), \tag{12}$$

$$\text{s.t. } U_i(\theta) = \int_0^{\theta} \frac{e_i(\tilde{\theta})}{\tilde{\theta}}\mathrm{d}\tilde{\theta},\ \forall\,\theta,i, \quad \theta e_i(\theta) \text{ is increasing,} \tag{IC-I}$$

Comparing the objective of problem (12) to its counterpart (6) with the complete information, the viewer type v in the parentheses of the objective in problem (12) is now replaced by its virtual counterpart $\Psi_j(v)$. Indeed, from the platform's perspective, the ads profit created by each match between (θ,i) and (v,j) is exactly the *virtual viewer surplus* in the parentheses, that is, the viewer's surplus minus the associated information rent the platform has to leave to higher-type viewers. Using Lemma 3, viewers' IC can be replaced by constraint (IC-I).

As before, to understand the trade off in determining attention allocation, we plug constraint (IC-I) into the objective and take pointwise first derivative with respect to $x_{ij}(\theta,v)$. After some algebra, we decompose the marginal effect of incremental attention $x_{ij}(\theta,v)$ into three terms.

Definition 2 (Network-Effect-Adjusted Virtual Surplus)
In the optimal algorithm, the value of $x_{ij}(\theta,v)$ is determined by the sign of

$$\underbrace{\Psi_j(v)\theta e_i(\theta)\mathbb{1}\{i=j\}-c}_{\text{Virtual Surplus}} + \underbrace{u'(X_i(\theta))\theta\beta_i \int_V x_{ii}(\theta,\tilde{v})\Psi_i(\tilde{v})\mathrm{d}G_i(\tilde{v})}_{\text{Network Effect: Current Type}}$$

$$\underbrace{-\,u'(X_i(\theta))\frac{\beta_i}{f_i(\theta)}\int_{\theta}^{\overline{\theta}}\int_V x_{ii}(\tilde{\theta},\tilde{v})\Psi_i(\tilde{v})\mathrm{d}G_i(\tilde{v})\mathrm{d}F_i(\tilde{\theta})}_{\text{Network Effect: Higher Types}}\,. \tag{13}$$

In particular, $x_{ij}(\theta,v)=1$ if expression (13) is positive, $x_{ij}(\theta,v)=0$ if it is negative, and any value in $[0,1]$ is optimal if it is zero.

The first term represents the direct pointwise effect of assigning attention, mirroring its counterpart in (7). Instead of extracting all viewer surplus, the platform now only captures the *virtual surplus* due to information rent. The second term echoes the indirect network effect in (7). Holding constant the information rent of the current type, the incremental attention allows the platform to extract more effort, globally benefiting all relevant viewers who watch this content. The third term is new to the incomplete information setting, as it represents the "doubly global" effect. An increase in effort by the current type globally raises information rent for all higher types, causing them to reduce effort; this reduction, in turn, globally affects all viewers who watch their content. The resulting aggregate impact is summarized by the double integral term. The second term, as in the complete-information benchmark, favors greater attention allocation; whereas the third term, arising under incomplete information, works against this force, discouraging excessive attention assignment.

4.2 Optimal Design under Incomplete Information

We now characterize the optimal allocation when users' vertical types are private information. Throughout the rest of the paper, we assume $\Psi_j(\underline{v}) > 0$. As in section 3, we still assume the attention cost c to be small so that the non-negativity condition (NN-C) never binds.

Proposition 2 (Algorithm: Incomplete Information)
Suppose $\frac{u(X)}{X}$ is decreasing and $u(X)\ \mathbb{E}_i\left[\Psi_i(v)\middle|G_i(v) \geqslant 1 - \frac{X}{\beta_i}\right]$ is strictly increasing for all $X \in [0, \frac{\beta_i}{2}]$. Then for each category i, the optimal algorithm is characterized by an increasing total attention $X_i(\cdot)$ such that:

$$x_{ii}(\theta, v) = \mathbb{1}\left\{G_i(v) \geqslant 1 - \tfrac{X_i(\theta)}{\beta_i}\right\}, \quad R_i(\theta) = (X_i(\theta) - \beta_i)^+, \quad e_i(\theta) = \frac{1}{\theta}\int_0^\theta \tilde{\theta}\, \mathrm{d}u\big(X_i(\tilde{\theta})\big).$$

Moreover, there exist cutoff abilities $0 < \theta_i^L \leqslant \theta_i^M \leqslant \theta_i^H$, with $\theta_i^L > \theta_i^l$, $\theta_i^M > \theta_i^m$, $\theta_i^H > \theta_i^h$, such that:

$$X_i(\theta)\begin{cases} = 0 & \text{if } \theta < \theta_i^L, \\ \in (0, \beta_i] & \text{if } \theta_i^L < \theta \leqslant \theta_i^M, \\ \in (\beta_i, \beta) & \text{if } \theta_i^M < \theta < \theta_i^H, \\ = \beta & \text{if } \theta \geqslant \theta_i^H. \end{cases}$$

The monotonicity on $\frac{u(X)}{X}$ and $u(X)\ \mathbb{E}_i\left[\Psi_i(v)\middle|G_i(v) \geqslant 1 - \frac{X}{\beta_i}\right]$ is a regularity assumption to guarantee that total attention is increasing without ironing. The assumption on $\frac{u(X)}{X}$ requires that the creator's utility cannot grow too fast when his exposure is relatively small. The monotonicity on $u(X)\ \mathbb{E}_i\left[\Psi_i(v)\middle|G_i(v) \geqslant 1 - \frac{X}{\beta_i}\right]$ is an incomplete information analogy of the assumption in Corollary 1. It requires that as a piece of content is recommended to more and more viewers, the increase of creator's utility must compensate the declining *average virtual* valuation among viewers.

Similar to the complete information case, the creators in each category are partitioned into four intervals based on ability. Within-category attention jumps from zero to a strictly positive level at θ_i^L and increases thereafter. Within a given category, attention allocation is monotone in viewer value: higher-value viewers consume everything that lower-value viewers do, and possibly more. Cross-category attention will be assigned only if the creator ability is significantly high ($\theta > \theta_i^M$), and is maxed out at θ_i^H. Except for creators at the entry cutoff, all participating creators earn information rents; accordingly, recommended effort is distorted away from the level that would fully extract their surplus. Again, the optimal algorithm partitions creators of each category into four segments as before:

- *Inactive*: types $\theta < \theta_i^L$, with zero effort and zero attention;
- *Local Entertainers*: types $\theta_i^L \leqslant \theta \leqslant \theta_i^M$, with positive and increasing quality and attention only from own category;
- *Ladder Climbers*: types $\theta_i^M < \theta < \theta_i^H$, with positive and increasing quality, full attention from own category, and increasingly but partial reach outside own category;
- *Global Influencers*: types $\theta \geqslant \theta_i^H$, with constant quality and full attention from all viewers.

Relative to complete information, private information on both sides makes it more costly for the platform to expand and monetize network effects. As a result, the optimal algorithm features both lower viewer attention and lower creator effort. Given the complementarity between attention and effort, this joint contraction is natural, but it is driven by two layers of forces. First, standard informational distortions require leaving information rents to both viewers and creators, limiting ads extraction and weakening effort incentives. Second, private information interacts with network effects: it becomes more costly to distort viewer attention and leverage network effects, while the incentive gains from such distortions are reduced. Lower effort and content quality further depress viewers' willingness to allocate attention, generating a feedback loop that dampens activity on both sides of the market.

The interaction between network effects and private information is most clearly reflected in two distortions of the optimal allocation: one on the intensive margin and one on the extensive margin. First, the distortion caused by asymmetric information extends to top types of both sides. For a highest-valued viewer of $v = \overline{v}$, we have $\Psi_j(\overline{v}) = \overline{v}$ so that normally there is no concern about leaving information rent to even higher types. However, with two-sided incomplete information her allocation is nevertheless distorted: she watches less and worse content due to information friction on the creator side. A similar logic applies to the most able creator of $\theta = \overline{\theta}$: effort is distorted downward and even attention is reduced if $\overline{\theta}$ falls below θ_i^H. Although the most able creator does not have an even higher type above, he still faces viewers who allocate less attention due to their own information rent, and therefore reduces effort.

Second, private information raises the creator-inclusion cutoff θ_i^L above the complete information counterpart θ_i^l. Under private information, the benefit of including a creator is now evaluated using viewers' virtual types that reflect the information rents left to the viewers. Meanwhile, the cost of including a creator includes the information rent paid to higher types of creators. Both forces require the algorithm to be more selective, shutting down more low-ability creators. Global Influencers are more difficult to sustain in niche categories for the same reason.

Category Survival and Reach Information asymmetry contracts the platform along both intensive and extensive margins, and these distortions are especially pronounced for niche categories with thinner and lower-valued viewer bases: two-sided private information threatens its existence from both directions by further raising an already high inclusion threshold. The following corollary imposes additional assumptions and makes this point more transparent.

Corollary 2 *If u is strictly concave and $u(X)\,\mathbb{E}_i\left[\Psi_i(v)\middle|G_i(v) \geqslant 1-\frac{X}{\beta_i}\right]$ is strictly increasing for all $X \in (0,\beta_i]$, then:*

$$X_i(\theta) = \begin{cases} \beta_i & \text{if } \theta_i^L < \theta \leqslant \theta_i^M, \\ u'^{-1}\left(\frac{c}{\Phi_i(\theta)\beta_i\mathbb{E}_i[\Psi_i(v)]}\right) & \text{if } \theta_i^M < \theta < \theta_i^H, \end{cases}$$

where $\Phi_i(\theta_i^L) = \frac{c}{u(\beta_i)\mathbb{E}_i[\Psi_i(v)]}$, $\Phi_i(\theta_i^M) = \frac{c}{\beta_i u'(\beta_i)\mathbb{E}_i[\Psi_i(v)]}$, *and* $\Phi_i(\theta_i^H) = \frac{c}{\beta_i u'(\beta)\mathbb{E}_i[\Psi_i(v)]}$.

The extra strictly concavity on u grants us a closed-form characterization of the optimal allocation. For the inclusion cutoff, information asymmetry implies the following upward shift:

$$\theta_i^L \geqslant \Phi_i(\theta_i^L) = \frac{c}{u(\beta_i)\mathbb{E}_i[\Psi_i(v)]} > \frac{c}{u(\beta_i)\mathbb{E}_i[v]} = \theta_i^l.$$

This first inequality holds by definition of virtual type, capturing the increased incentive costs of creators due to private information. The second inequality arises because viewers' private information reduces the platform's ability to profit from attention on ads, reflected in $\Psi_i(v) < v$ for almost all v. Information asymmetry on both sides jointly raises the participation cutoff. The same logic applies to the remaining cutoffs on the production side. Notice that these cutoffs are determined by group characteristics of the same category and are shaped by information asymmetry on both sides.

The closed-form expression above yields a simple survival condition for a category: some creator in the category can be included only if the inclusion cutoff is no larger than the highest creator ability, or

$$\mathbb{E}_i[\Psi_i(v)]u(\beta_i) \geqslant \frac{c}{\overline{\theta}}.$$

For each category, this formula implicitly gives the *critical mass* of viewers required for the category to remain active, taking as given its value distribution and informational frictions. Importantly, the critical mass is affected by viewer-side private information through $\mathbb{E}_i[\Psi_i(v)]$, but not by creator-side private information, since the top creator's virtual type equals her true type. For categories with viewer base β_i such that $u(\beta_i) \in \left(\frac{c}{\overline{\theta}\mathbb{E}_i[v]}, \frac{c}{\overline{\theta}\mathbb{E}_i[\Psi_i(v)]}\right)$, viewer-side private information can make these categories extinct, even though they would remain viable under complete information.

The discussion above points to a stark policy implication. In practice, social media platforms may observe or infer user characteristics that predict preferences, yet legal and regulatory constraints can limit how such information is used in recommender systems. For example, by granting users the right to opt out of data sharing (CCPA) and requiring explicit consent for data collection (GDPR), these regulations severely shrink the volume of first-party and third-party data platforms can legally collect. Additionally, the EU's Digital Markets Act's (DMA) prohibition on cross-app data combination prevents platforms such as Meta from merging the data fragments about their users, leaving them with incomplete user profiles. Under this interpretation, our results suggest a regulatory trade-off: limiting the use of user characteristics in personalized recommendations can raise participating users' welfare by leaving them information rents, but may also reduce diversity in the platform's content ecosystem by making it harder for niche content to reach its relevant audience.

Also, the above discussion further clarifies the role of viewer-base expansion: by bringing marginal categories above their survival thresholds, a larger viewer base can mitigate the loss of network-effect leverage caused by restrictions on personalized recommendations and sustain greater category variety.

Non-Monotone Effort Information asymmetry also reshapes the relationship between creator ability, quality, and effort. In the optimal allocation, some creators are sufficiently able to be shown to all viewers in their own category, but not sufficiently able to justify costly cross-category reach. Since the attention reward is flat for these creators, incentive compatibility can only induce a constant quality level among them. When the within-category allocation is all-or-nothing, all Local Entertainers fall into this fixed-reward region. A similar logic applies to sufficiently able creators ($\theta \geqslant \theta_i^H$). Since they already reach all viewers, the platform runs out of incentive instruments and the resulting quality bunches at the maximum level. By contrast, Ladder Climbers face active trade-offs: they can marginally increase quality in exchange for marginally further reach among cross-category viewers. A typical quality function is depicted in Figure 2(a), where bunching appears as flat segments for Local Entertainers and Global Influencers.

Non-decreasing qualities does not necessarily translate into non-decreasing efforts. In the flat-quality segments in Figure 2(a), more able creators exert less effort to sustain the same quality

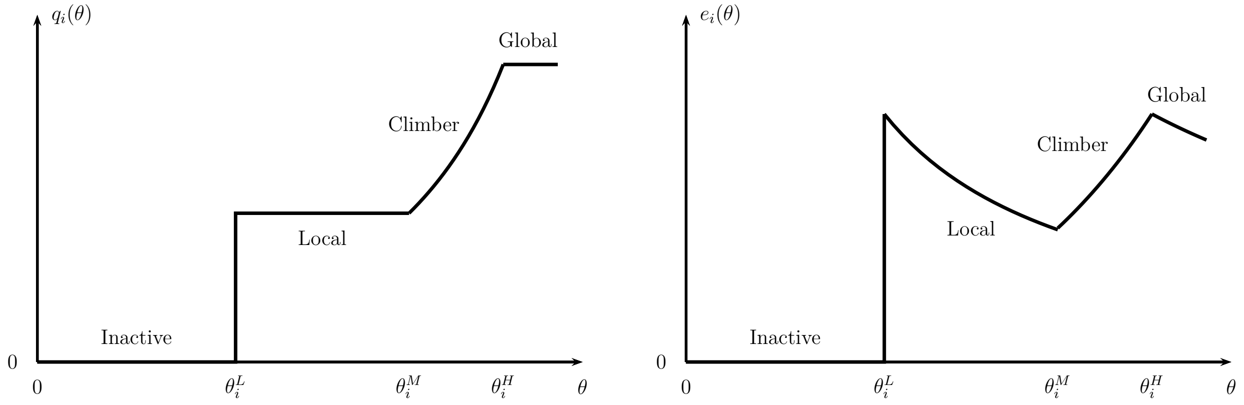

(a) Quality

(b) Effort

Figure 2

Quality $\theta e_i(\theta)$ and Effort $e_i(\theta)$ as functions of ability θ for category-i creators. (a) Quality is increasing. (b) Effort is not monotone.

level, whereas for Ladder Climbers, effort typically increases with ability. Intuitively, Local Entertainers and Global Influencers face insufficient marginal return from quality improvement and thus slack off, while Ladder Climbers aspire to increase their extended reach by working harder. Overall, effort is non-monotone in ability and typically double peaked.

4.3 Implementation

Creator Side The optimal direct mechanism specifies effort and total attention $(e_i(\theta), X_i(\theta))$ for each creator ability θ in category i. As an outcome, a (θ, i)-creator produces observable content quality at $q_i(\theta) = \theta e_i(\theta)$ and earns total attention $X_i(\theta)$. To implement this, the platform can use an attention allocation rule χ_i that maps observed quality q into the promised total attention $\chi_i(q)$ for each creator in category i. Specifically, we can write $q_i(\theta) = \theta e_i(\theta) = \int_0^\theta \tilde{\theta}\, \mathrm{d}u\left(X_i(\tilde{\theta})\right)$ from Proposition 2. Eliminating the dependence on θ, we have the following result.

Proposition 3 (Implementation: Creators)
For creators in category i, the platform promises total attention $\chi_i(q)$ implicitly defined by:

$$q = \int_0^{\chi_i(q)} X_i^{-1}(X)\mathrm{d}u(X),$$

where X_i^{-1} is the generalized inverse. Then effort $e_i(\theta)$ in Proposition 2 is optimal for θ-creators.

Proposition 3 proposes one implementation directly from the optimal algorithm. Each creator maximizes $u(\chi_i(q)) - \frac{q}{\theta}$ and self-sort into the intended quality level for their ability. The implementation can be further simplified for practice, by truncating quality levels that are not chosen by any creator. For example, each category can set a minimum quality threshold $\underline{q}_i \equiv q(\theta_i^L)$, below which no attention is granted. It can also set a quality cap $\overline{q} \equiv q(\theta_i^H)$ at which all viewer attention is deployed to reward the creator.

Viewer Side The optimal direct mechanism assigns a bundle content-ads mix for each (v, j)-viewer. Define:

$$\vartheta_j(v) \equiv \{\theta : x_{jj}(\theta, v) = 1\} = \left\{\theta : G_j(v) \geqslant 1 - \tfrac{X_j(\theta)}{\beta_j}\right\}$$

as the set of creator types in category j whose content is recommended to (v, j)-viewers. Then a (v, j)-viewer receives all content from $\vartheta_j(v)$ plus a total amount $\frac{\alpha_j}{c}\int_{\vartheta_j(v)}(v\theta e_j(\theta) - c)\mathrm{d}F_j(\theta) - \frac{W_j(v)}{c}$ of irrelevant content and ads.

To implement this naturally, we use a single *ordered feed* for all creators in each category. Ads are blended into the content, and their frequency varies along the feed. All items are arranged in a fixed order, and a viewer consumes them along this order and chooses where to stop. To gain intuition, although the model is static, the feed can be interpreted as unfolding over time. Each moment on the platform consumes one unit of attention. Viewers move through an ordered mix of content and ads and choose where to stop, with the stopping point determining total attention. By varying the composition and intensity of content along the feed, the platform induces viewers with different willingness to stay to choose different stopping points, thereby replicating the optimal allocation.

We already defined $Q_j(v) = \alpha_j \int_{\vartheta_j(v)} \theta e_j(\theta)\mathrm{d}F_j(\theta)$ as the within-category quality-weighted content. Further, define:

$$A_j(v) \equiv \frac{\alpha_j v}{c}\int_{\vartheta_j(v)} \theta e_j(\theta)\mathrm{d}F_j(\theta) - \frac{W_j(v)}{c} = \frac{vQ_j(v) - W_j(v)}{c}$$

as the total attention a (v, j)-viewer spends. The next result proposes the exact ordered feed for all viewers in category j, described by the cumulative within-category quality-weighted content Q as a function of cumulative attention A spent in the feed.

Proposition 4 (Implementation: Viewers)
For viewers in category j, the platform recommends the same ordered feed characterized by function K_j, such that cumulative attention A is rewarded with cumulative within-category quality-weighted

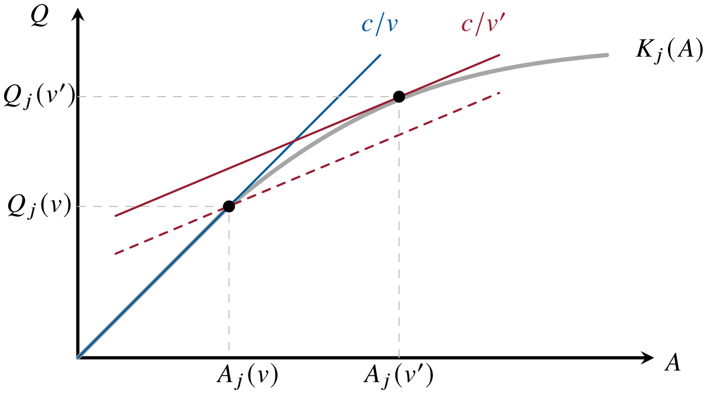

Figure 3. Implementation: Viewer Side

Notes: $v' > v$. The gray thick curve represents $K_j(A)$: category j viewer's feasible choices of (Q, A). The solid red and blue curves represent the indifference curves of the two viewers, with upward shifts indicating higher utility levels. The red dashed curve depicts the utility level that a v'-viewer would obtain if she were to stop at $A_j(v)$.

content $K_j(A)$*, where:*

$$K_j(A) \equiv \sup\left\{\sum_k \rho_k Q_j(v_k) : \sum_{k=1}^{n} \rho_k A_j(v_k) = A, \rho_k \geqslant 0, \sum_{k=1}^{n} \rho_k \leqslant 1\right\}$$

is the concave envelope of the graph $\{(0,0)\} \cup \{(A_j(v), Q_j(v))\}_{v\in V}$.

Proposition 4 uses variables in the optimal algorithm to construct the implementation. Each (v, j)-viewer maximizes $vK_j(A) - cA$ and optimally stops at the intended "stopping time" $A_j(v)$. By construction, incentive compatibility already ensures that a (v, j)-viewer weakly prefers stopping at $A_j(v)$ attention to stopping at $A_j(v')$ for any v'. We illustrate this construction in Figure 3.

Notice, our implementation can allow viewers to choose a level of total attention A that does not correspond to any type's allocation specified by the direct mechanism. First, the set of equilibrium total attention $\{A_j(v)\}_{v\in V}$ may be disconnected. The concave envelope $K_j(A)$ construction extends the viewers' feasible choice to convex hull of $\{A_j(v)\}_{v\in V} \cup \{0\}$ but ensures that the viewer does not stop before consuming her intended expanded set. Second, $A_j(\cdot)$ is maximized at $\bar{v}$. One can extend viewer's feasible set by allowing her to browse beyond the amount intended for the highest type $\overline{v}$ in her category without violating IC: for $A > A_j(\overline{v})$, display only cross-category content or ads so that $K_j(A)$ plateaus at $Q_j(\overline{v})$ for all $A > A_j(\overline{v})$ and the viewer will automatically choose not to do so.

5 Platform Payments

In this section, we extend the baseline model by introducing creator-side monetary transfers. In addition to allocating attention through recommendations, the platform can now make monetary transfer to creators. One possible implementation of the transfer instrument is direct monetary payments from the platform to creators, as in monetization programs such as TikTok Creator Rewards. More broadly, however, the platform may share monetization revenue with creators through a variety of arrangements, including advertising revenue-sharing programs, affiliate commissions, sponsored content opportunities, or embedded commerce features such as live shopping. For example, creators may be allowed to insert advertisements or product promotions into their own content and retain part of the resulting revenue. We show that monetary payments often complement algorithmic attention: even when monetary transfers are feasible, attention remains a valuable incentive instrument, and the platform may optimally combine exposure with transfers to motivate content production. Moreover, the analysis helps explain why different creator incentive models coexist across platforms, and yields policy implications for regulations that restrict profiling-based advertising.

To formally incorporate the monetary channel, we assume that a creator values both attention and money. His payoff from total attention $X \geqslant 0$ and transfer $T \geqslant 0$ from the platform is $u(Y)$ where $Y \equiv \pi(X) + T$ is the *total money-metric reward*, decomposed into $\pi(X)$, the money-metric payoff from attention X, and direct payment T. We assume u is strictly concave to capture overall diminishing returns while allowing π to be any increasing and continuously differentiable function.

With monetary transfers, the direct mechanism receives report θ and recommends effort $e_i(\theta)$. If the observable quality q matches $\theta e_i(\theta)$, then attention profile $x_{i\cdot}(\theta, \cdot)$ is assigned and a direct transfer $T_i(\theta) \geqslant 0$ is made.[7] As in the main model, we substitute out the platform-inserted ads $a_j(v)$ and simplify the platform's problem to:

$$\max_{\substack{\{e_i(\cdot)\}_i, \{T_i(\cdot)\}_i, \\ \{x_{ij}(\cdot,\cdot)\}_{i,j}}} \sum_i \alpha_i \int_\Theta \left(\sum_j \beta_j \int_V x_{ij}(\theta, v) \Big(\Psi_j(v) \theta e_i(\theta) \mathbb{1}\{i = j\} - c \Big) \mathrm{d}G_j(v) - \frac{c}{z} T_i(\theta) \right) \mathrm{d}F_i(\theta),$$

$$\text{s.t.} \quad e_i(\theta) + U_i(\theta) = u\left(\pi(X_i(\theta)) + T_i(\theta)\right), \ \forall\, \theta, i,$$

$$U_i(\theta) = \int_0^\theta \frac{e_i(\tilde{\theta})}{\tilde{\theta}} \mathrm{d}\tilde{\theta}, \ \forall\, \theta, i, \quad \theta e_i(\theta) \text{ is increasing},$$

$$\sum_i \alpha_i \int_\Theta x_{ij}(\theta, v) \Big(\Psi_j(v) \theta e_i(\theta) \mathbb{1}\{i = j\} - c \Big) \mathrm{d}F_i(\theta) \geqslant 0, \ \forall\, v, j.$$

The incentives enter the problem in two places. On one hand, $T_i(\theta)$ is directly deducted from the

---

[7] We restrict attention to nonnegative transfers. Allowing $T_i(\theta) < 0$ would amount to allowing the platform to sell recognition or visibility to creators, as in models of paid certification or status provision; see, e.g., Bar-Isaac et al. (2025).

platform's objective. On the other hand, it enters the utility of the creators so that higher effort can be extracted.

Proposition 5 (Optimal Algorithm with Money)
Suppose $\frac{u(\pi(X)+T)-u(T)}{X}$ *is decreasing in* X *and* $u(\pi(X)+T)\ \mathbb{E}_i\left[\Psi_i(v)\middle|G_i(v)\geqslant 1-\frac{X}{\beta_i}\right]$ *is strictly increasing for all* $X\in[0,\frac{\beta_i}{2}]$ *and* $T\geqslant 0$*. Then for each category* i*, the optimal algorithm is characterized by an increasing total attention* $X_i(\cdot)$ *such that* $x_{ii}(\theta,v)$ *and* $R_i(\theta)$ *are similarly defined as in Proposition 2, and a potentially non-monotone transfer* $T_i(\cdot)$ *such that* $e_i(\theta)=\frac{1}{\theta}\int_0^\theta \tilde{\theta}\ \mathrm{d}u\left(\pi(X_i(\tilde{\theta}))+T_i(\tilde{\theta})\right)$*. For all* θ *such that* $X_i(\theta)\geqslant\beta_i$*, the ability space is partitioned into alternating intervals of two patterns:*

- $X_i(\theta)$ *strictly increases while* $T_i(\theta)=0$*, and*
- $X_i(\theta)$ *stays constant while* $T_i(\theta)$ *strictly increases.*

If π *is concave and* $\pi'(\beta_i)>z$*, then there exists* $\theta_i^\$>\theta_i^M$ *such that* $T_i(\theta)=0$ *for* $\theta\leqslant\theta_i^\$$*, and* $T_i(\theta)>0$ *and* $X_i(\theta)=X_i(\theta_i^\$)$ *for* $\theta>\theta_i^\$$*.*

Proposition 5 predicts a similar segmentation of creator abilities by total attention received. Its main novelty is that monetary transfers interact with attention allocation. In general, transfers follow an *increase-and-reset* pattern, and when the money-metric payoff function π exhibits diminishing marginal returns, only sufficiently high-ability creators receive monetary payments.

To better understand the efficient use of X_i and T_i as alternative resources, we examine what total money-metric reward $Y_i=\pi(X_i)+T_i$ can be achieved from any arbitrary level $\hat{X}_i\equiv X_i+\frac{T_i}{z}$ of total attention-metric resources, whenever $X_i>\beta_i$. For creator incentives, only the money-metric reward Y_i matters; for platform cost, only the attention-metric resources $\hat{X}_i$ matters. To proceed, define $\hat{\pi}$ as the pointwise smallest function satisfying:

$$\begin{aligned}&\hat{\pi}(X)\geqslant\pi(X)\quad\forall\ X\geqslant\beta_i,\\&\hat{\pi}(X')-\hat{\pi}(X)\geqslant z(X'-X)\quad\forall X'>X\geqslant\beta_i.\end{aligned}$$

Geometrically, $\hat{\pi}$ is the z-ironing of π, achieved by the upper envelope of all z-sloped rays starting from π and pointing to the right (Figure 4). The ironing guarantees that whenever the return from attention is locally lower than z, the platform should withhold this portion of attention and pay the creators instead.

The following lemma characterizes the resource allocation in the optimal algorithm when $X_i\geqslant\beta_i$, and explains why money T_i displays resetting behavior.

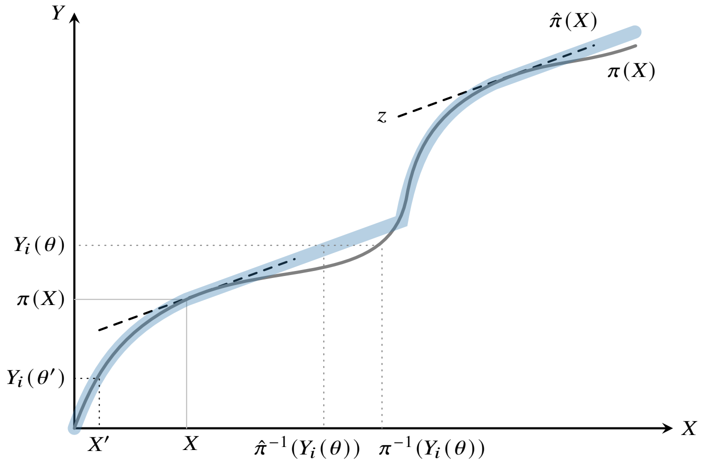

Figure 4. Illustration of Optimal Allocation

Notes: The solid thin black curve represents $\pi(X)$, and the black dashed lines have slope z and are tangent to $\pi(X)$ from above. The thick blue curve represents $\hat{\pi}(X)$. If the optimal allocation specifies $Y_i(\theta)$ such that $\hat{\pi}^{-1}(Y_i(\theta)) < \pi^{-1}(Y_i(\theta))$, then $X_i(\theta) = X$ and $T_i(\theta) = Y_i(\theta) - \pi(X)$. If instead $Y_i(\theta') = \pi(X') = \hat{\pi}(X')$ for some X', then $X_i(\theta') = \pi^{-1}(Y_i(\theta'))$ and $T_i(\theta') = 0$.

Lemma 4 (Resource Optimization)
When $X_i(\theta) \geqslant \beta_i$, we must have $Y_i(\theta) = \hat{\pi}(\hat{X}_i(\theta))$. Given $Y_i(\theta)$, we have $X_i(\theta) = \max\{X : \pi(X) + z(\hat{\pi}^{-1}(Y_i(\theta)) - X) = Y_i(\theta)\}$.

The lemma claims that any optimal algorithm must feature efficient resource allocation such that the attention-metric resource $\hat{X}_i$ and the money-metric reward Y_i must lie on the graph of $\hat{\pi}$. It also gives the efficient resource allocation once the money-metric reward Y_i: if $Y_i(\theta) = \pi(\hat{\pi}^{-1}(Y_i(\theta)))$, then $(\hat{\pi}^{-1}(Y_i(\theta)), Y_i(\theta))$ lies on the graph of π where the z-ironing has no bite, and pure attention is the most efficient use of resources. Then, as ability increases, attention strictly increases but money is fixed at zero. Instead, if $Y_i(\theta) > \pi(\hat{\pi}^{-1}(Y_i(\theta)))$, then the platform strictly benefits from some money. As ability increases, the ironed part of $\hat{\pi}$ continues to allocate more money while holding total attention at the same level.

The following proposition can then focus on the $Y_i(\theta), X_i(\theta), T_i(\theta)$ and how it is affected by z.

Proposition 6 (Comparative Statics on Ads Income)
For any (θ, i)-creator, $X_i(\theta)$ decreases in z while $T_i(\theta)$ increases in z whenever $X_i(\theta) \geqslant \beta_i$.

This comparative static follows naturally from the paper's core trade-off in attention allocation. A higher z means that each unit of viewer attention is more valuable in direct advertising monetization. As a result, diverting attention away from viewers' preferred content in order to incentivize creators becomes more costly. The platform therefore relies more on monetary transfers and less on distorted attention allocation to provide creator incentives.

The novelty is to link the platform's advertising profitability to the accuracy of its recommendation algorithm. When advertising monetization is weak, attention is relatively cheap as an incentive instrument, so the platform has a stronger reason to use extended reach to reward creators. When advertising monetization is strong, attention is too valuable to waste on incentive provision, and the platform substitutes toward direct payments. Hence, platforms with greater advertising profitability should, all else equal, make more accurate recommendations. The parameter z can be interpreted broadly. It may reflect the platform's market power in advertising, the nature of the content environment, such as long videos versus short posts, or technological and legal constraints on profiling-based advertising. Restrictions that limit the platform's ability to monetize attention through targeted ads reduce z, making exposure-based incentives relatively more attractive and potentially increasing recommendation distortions.

This comparative static provides a potential rationale for why different platforms rely on different creator compensation models. Platforms with weak advertising monetization may rely primarily on exposure-based rewards, such as recommendation amplification and external reach, because viewer attention is relatively cheap to use as an incentive instrument. By contrast, platforms with strong advertising profitability have a greater incentive to preserve attention for direct monetization and therefore rely more heavily on explicit monetary compensation programs for creators. Hence, differences in creator payment models across platforms may arise endogenously from differences in advertising environments rather than purely from technological or managerial choices.

The result also has policy implications. Regulations that restrict targeted or profiling-based advertising may affect not only advertising revenue directly, but also the equilibrium design of recommendation algorithms. For example, the DMA's restrictions on cross-app tracking or the EU's Digital Services Act (DSA) restrictions on certain profiling-based ads using special category data heavily degrade the algorithmic precision Meta relies on to command high ad prices. Similarly, privacy regulation (including GDPR and CCPA) limits any platform's ability to sell targeted ads. Because advertisers see significantly lower click-through and conversion rates on less-personalized ads, the bidding value (yield) for each individual ad slot drops (see, e.g., Johnson et al., 2020).

By lowering the profitability of monetizing viewer attention, such regulations make exposure-based incentives more attractive and induce platforms to distort recommendations more aggressively to motivate content creation. Importantly, this effect arises in equilibrium through the platform's endogenous choice of incentive instruments. The model therefore provides a new perspective for policy discussions on advertising regulation and recommendation systems, highlighting a trade-off between limiting data-driven monetization and preserving recommendation accuracy.

6 Conclusion

This paper builds a framework to study a monopolistic social media platform's optimal content recommendation algorithm. The key trade-off is between immediate monetization of attention and allocating attention to amplify endogenous network effects that increase content quality and, in turn, expand the revenue base.

Our framework can be extended in many ways. One direction is to add monetary transfers between the platform and viewers. For example, YouTube offers an ad-free premium membership, allowing viewers to pay to reduce ads exposure (see, e.g., Johnson (2013), Bisceglia et al. (2025)). One could allow viewers to pay to avoid ads and even cross-category content. While such options would limit the platform's ability to leverage network effects, they would generate direct revenue, raising the question of how to optimally price the network effects a viewer generates.

A related extension is to add personalized advertising. In our baseline model, each unit of attention generates a constant advertising return. In practice, however, the value of exposure depends on the viewer's characteristics and the category of content consumed. Allowing advertising revenue to vary with viewer type and category would introduce an additional design margin: the algorithm would simultaneously manage production incentives and allocate attention to maximize heterogeneous advertising surplus. This would connect our framework to the literature on targeted advertising and data-driven ad markets (see, e.g., Bergemann and Bonatti (2011, 2024), and Ichihashi et al. (2025)), while embedding targeting decisions within a unified two-sided screening environment.

One can also introduce limited attention capacity to generate competition among content items for viewer time, creating crowding-out effects even within a monopoly platform. This becomes especially relevant in environments where technological progress—such as advances in AI—raises content productivity and amplifies differences across creators, making attention scarcity central to the determination of algorithmic thresholds and market structure. Attention capacity is also essential for studying platform competition: with multi-homing, viewers must allocate limited attention across platforms (as in, e.g., Ambrus et al. (2016), Prat and Valletti (2022), and Chen (2025)), and algorithm design would interact with cross-platform competition for viewer time.

Finally, our analysis abstracts from dynamic considerations. In practice, user characteristics may evolve over time, and platforms learn about users' private information from their historical behavior. When platforms lack commitment, dynamic interaction may generate ratchet effects that dampen incentives for truthful revelation. Alternatively, users themselves may initially be uncertain about their types and learn through experience. Extending the framework to incorporate dynamic mechanism design (e.g., Bergemann and Välimäki (2019); Fershtman and Pavan (2022)) would deepen our understanding of how algorithm design interacts with long-run incentives.

Appendix: Proofs

Proof of Lemma 1. Suppose (5) is slack for a positive G_j-measure of v's in some category j. Then the platform should increase $a_j(v)$ for these v's to increase profit.

Now suppose (4) is slack for a positive F_i-measure of θ's in some category i. Without loss of generality, the platform can increase $e_i(\theta)$ for these θ's to increase the quality of their contents. This in turn weakly relaxes (5) for all viewers (v, j) and thus increases profit by inserting more ads for them. ■

Proof of Lemma 2. Let $x^*_{ij}(\theta, v) = \frac{\hat{R}_i(\theta)}{\beta-\beta_i}$ for all θ, i, v, j so that $\sum_{j\neq i}\beta_j \int_V x^*_{ij}(\theta, v)\mathrm{d}G_j(v) = \hat{R}_i(\theta)$. Notice that (6) can be rewritten as:

$$\sum_i \alpha_i \left(\beta_i \int_\Theta \int_V x_{ii}(\theta, v)\Big(v\theta e_i(\theta) - c\Big)\mathrm{d}G_i(v)\mathrm{d}F_i(\theta) - c\int_\Theta R_i(\theta)\mathrm{d}F_i(\theta)\right).$$

If $\{e^*_i, x^*_{ij}\}_{i,j}$ does not solve problem (6), then there exists $i \in \mathcal{N}$, e_i and $\{x_{ij}\}_j$ such that $R_i(\theta) \equiv \sum_{j\neq i}\beta_j \int_V x_{ij}(\theta, v)\mathrm{d}G_j(v)$, $e_i(\theta) = u\Big(\beta_i \int_V x_{ii}(\theta, v)\mathrm{d}G_j(v) + R_i(\theta)\Big)$, and:

$$\begin{aligned}
&\beta_i \int_\Theta \int_V x_{ii}(\theta, v)\Big(v\theta e_i(\theta) - c\Big)\mathrm{d}G_i(v)\mathrm{d}F_i(\theta) - c\int_\Theta R_i(\theta)\mathrm{d}F_i(\theta)\\
\geqslant\ &\beta_i \int_\Theta \int_V \hat{x}_{ii}(\theta, v)\Big(v\theta \hat{e}_i(\theta) - c\Big)\mathrm{d}G_i(v)\mathrm{d}F_i(\theta) - c\int_\Theta \hat{B}_i(\theta)\mathrm{d}F_i(\theta).
\end{aligned}$$

This implies that $\{\hat{e}_i, \hat{x}_{ii}, \hat{R}_i\}_i$ is not a solution to problem (9), a contradiction. ■

Proof of Proposition 1. The effort choice is derived from Lemma 1: $e_i(\theta) = u\left(\beta_i \int_V x_{ii}(\theta, v)\mathrm{d}G_i(v) + R_i(\theta)\right)$. Plugging (IR-C) into (6) and taking derivative w.r.t. $x_{ij}(\theta, v)$, we have (7). For $i = j$ and $X_i(\theta) > 0$, then $e_i(\theta) > 0$ and (7) is strictly increasing in v for all $\theta > 0$ and therefore $x_{ii}(\theta, \cdot)$ must be a step function. If $X_i(\theta) = 0$ or $i \neq j$, (7) does not depend on v.

Notice that if $x_{ii}(\theta, \underline{v}) = 0$, then (7) $\leqslant 0$ for $x_{ii}(\theta, \underline{v})$ implies (7) < 0 for $x_{ij}(\theta, v)$ for any v and $j \neq i$, and thus $R_i(\theta) = 0$. Therefore, we use only one variable $X_i(\theta)$ to characterize optimal algorithm without ambiguity, where $X_i(\theta) \leqslant \beta_i$ implies $R_i(\theta) = 0$ and $x_{ii}(\theta, v) = \mathbb{1}\left\{G_i(v) \geqslant 1 - \frac{X_i(\theta)}{\beta_i}\right\}$, and $X_i(\theta) > \beta_i$ implies $R_i(\theta) > 0$ and $x_{ii}(\theta, v) = 1$ for all v.

Category i's sub-problem simplifies to:

$$\begin{aligned}
&\max_{X_i} \int_\Theta H^C_i(\theta, X_i)\mathrm{d}\theta,\\
\text{where}\qquad &\frac{H^C_i(\theta, X_i)}{f_i(\theta)} \equiv \beta_i\theta u(X_i)\int_{G_i^{-1}\left(1-\frac{\min\{X_i,\beta_i\}}{\beta_i}\right)}^{\overline{v}} v\mathrm{d}G_j(v) - cX_i(\theta).
\end{aligned}$$

Note that

$$\frac{\partial^2}{\partial\theta\partial X}\frac{H_i^C(\theta,X)}{f_i(\theta)} = \begin{cases} u(X)G_i^{-1}\left(1-\frac{X}{\beta_i}\right)+\beta_i u'(X)\int_{G_i^{-1}(1-\frac{X}{\beta_i})}^{\overline{v}} v\mathrm{d}G_j(v) > 0 & \text{if } X \leqslant \beta_i,\\ \beta_i u'(X)\mathbb{E}_i[v]>0 & \text{if } X>\beta_i.\end{cases}$$

According to Topkis' theorem, $X_i(\theta)$ is increasing in θ. Define $\theta_i^l \equiv \inf\{\theta : X_i(\theta) > 0\}$, $\theta_i^m \equiv \inf\{\theta : X_i(\theta) > \beta_i\}$, and $\theta_i^h \equiv \inf\{\theta : X_i(\theta) = \beta\}$. By definition. $\theta_i^h \geqslant \theta_i^m \geqslant \theta_i^l$. Moreover, $\frac{\partial}{\partial X}\frac{H_i^C(0,X)}{f_i(0)} = -c$, so $X_i(\theta) = 0$ for θ close enough to 0, implying that $\theta_i^l > 0$. Finally, $\frac{\partial}{\partial X}\frac{H_i^C(\theta_i^l,X)}{f_i(\theta_i^l)} = -c$ at $X = 0$, which means $X_i(\theta_i^l) > 0$. ■

Proof of Corollary 1. If the inequality holds, then for any $\theta > 0$ and any $X \in (0,\beta_i)$, we have:

$$\frac{H_i^C(\theta,X)}{f_i(\theta)X} = \theta u(X)\mathbb{E}_i\left[v \middle| G_i(v) \geqslant 1-\tfrac{X}{\beta_i}\right] - c.$$

By assumption, this is strictly increasing, and therefore we have $X_i(\theta_i^l) = \beta_i$. Define $\theta_i^l = \frac{c}{u(\beta_i)\mathbb{E}_i[v]}$, then $H_i^C(\theta_i^l,\beta_i) = 0$.

Moreover, when u is strictly concave, $\frac{\partial}{\partial X_i}\frac{H_i^C(\theta,X_i)}{f_i(\theta)} = \beta_i\mathbb{E}_i[v]\theta u'(X_i) = c$ for $X_i \geqslant \beta_i$. Setting $X_i(\theta_i^m) = \beta_i$ and $X_i(\theta_i^h) = \beta$ yields the thresholds. ■

Proof of Lemma 3. (i) Necessity. If (2) holds, then for any $\theta > \theta'$:

$$\begin{aligned} U_i(\theta) &\geqslant U_i(\theta') + \frac{q_i(\theta')}{\theta'} - \frac{q_i(\theta')}{\theta},\\ U_i(\theta') &\geqslant U_i(\theta) + \frac{q_i(\theta)}{\theta} - \frac{q_i(\theta)}{\theta'}. \end{aligned}$$

Adding up, we have $(\theta-\theta')(q_i(\theta)-q_i(\theta')) \geqslant 0$. Therefore, monotonicity holds. Taking the limit as $\theta' \to \theta$, we use Sandwich theorem to have $U_i'(\theta) = \frac{q_i(\theta)}{\theta^2}$, which implies the integral condition.

Sufficiency. Suppose monotonicity and integral conditions hold. Then for any θ, θ':

$$\begin{aligned} U_i(\theta) &= U_i(\theta') + \int_{\theta'}^{\theta}\frac{q_i(\tilde\theta)}{\tilde\theta^2}\mathrm{d}\tilde\theta \geqslant U_i(\theta') + \int_{\theta'}^{\theta}\frac{q_i(\theta')}{\tilde\theta^2}\mathrm{d}\tilde\theta\\ &= U_i(\theta') + q_i(\theta')\left(\frac{1}{\theta'}-\frac{1}{\theta}\right) = u\Big(\sum_j \beta_j\int_V x_{ij}(\theta',v)\mathrm{d}G_j(v)\Big) - \frac{\theta' e_i(\theta')}{\theta}. \end{aligned}$$

(ii) Necessity. If (3) holds, then for any v, v':

$$\begin{aligned} W_j(v) &\geqslant W_j(v') - (v'-v)Q_j(v'),\\ W_j(v') &\geqslant W_j(v) - (v-v')Q_j(v). \end{aligned}$$

Adding up, we have $(v-v')(Q_j(v)-Q_j(v')) \geqslant 0$. Therefore, monotonicity holds. Taking the limit as $v' \to v$, we use Sandwich theorem to have $W_j'(v) = Q_j(v)$, which implies the integral condition.

Sufficiency. Suppose monotonicity and integral conditions hold. Then for any v, v':

$$\begin{aligned} W_j(v) &= W_j(v') + \int_{v'}^{v} Q_j(\tilde{v})\mathrm{d}\tilde{v} \geqslant W_j(v') + \int_{v'}^{v} Q_j(v')\mathrm{d}\tilde{v} \\ &= W(v') + (v - v')Q_j(v') \\ &= \sum_i \alpha_i \int_{\Theta} x_{ij}(\theta, v')\Big(v\theta e_i(\theta)\mathbb{1}\{i = j\} - c\Big)\mathrm{d}F_i(\theta) - ca_j(v'). \end{aligned}$$

■

Proof of Proposition 2. We first ignore the constraints that $\theta e_i(\theta)$ is increasing and $e_i(\theta) \geqslant 0$, solve the problem, and then verify later. The problem is additively separable in creator categories. Plugging in $W_i(v) = W_i(\underline{v}) + \int_{\underline{v}}^{v} Q_i(\tilde{v})\mathrm{d}\tilde{v}$ and integrating by parts, we can rewrite category i's Hamiltonian as:

$$f_i(\theta)\sum_j \left(\beta_j \int_V x_{ij}(\theta, v)(\Psi_j(v)\theta e_i(\theta)\mathbb{1}\{i = j\} - c)\mathrm{d}G_j(v)\right) + \gamma_i(\theta)\frac{e_i(\theta)}{\theta},$$

$$\text{where} \qquad e_i(\theta) = u\left(\sum_j \beta_j \int_V x_{ij}(\theta, v)\mathrm{d}G_j(v)\right) - U_i(\theta).$$

Taking derivative w.r.t. $x_{ij}(\theta, v)$, we have (13). For $i = j$ and $X_i(\theta) > 0$, then $e_i(\theta) > 0$ and (13) is strictly increasing in v for all $\theta > 0$ and therefore $x_{ii}(\theta, \cdot)$ must be a step function. If $X_i(\theta) = 0$ or $i \neq j$, (13) does not depend on v.

Notice that if $x_{ii}(\theta, \underline{v}) = 0$, then (13) $\leqslant 0$ for $x_{ii}(\theta, \underline{v})$ implies (13) < 0 for $x_{ij}(\theta, v)$ for any v and $j \neq i$, and thus $R_i(\theta) = 0$. Therefore, we use only one variable $X_i(\theta)$ to characterize optimal algorithm without ambiguity, where $X_i(\theta) \leqslant \beta_i$ implies $R_i(\theta) = 0$ and $x_{ii}(\theta, v) = \mathbb{1}\left\{G_i(v) \geqslant 1 - \frac{X_i(\theta)}{\beta_i}\right\}$, and $X_i(\theta) > \beta_i$ implies $R_i(\theta) > 0$ and $x_{ii}(\theta, v) = 1$ for all v.

Define:

$$J_i(X) \equiv \min\{X, \beta_i\}G_i^{-1}\left(1 - \tfrac{\min\{X, \beta_i\}}{\beta_i}\right),$$

which satisfies $J_i(0) = 0$, $J_i'(X) > 0$ for $X < \beta_i$ and $J_i(X) = J_i(\beta_i)$ for $X \geqslant \beta_i$. The Hamiltonian can be simplified to:

$$\frac{H_i^I(\theta, X_i, U_i, \gamma_i)}{f_i(\theta)} \equiv \theta e_i(\theta)J_i(X_i) - cX_i + \frac{\gamma_i(\theta)}{f_i(\theta)}\frac{e_i(\theta)}{\theta}.$$

Optimality requires $\frac{\partial H_i^I}{\partial U_i} = -\gamma_i'(\theta)$, which means:

$$\gamma_i'(\theta) = \frac{\gamma_i(\theta)}{\theta} + \theta f_i(\theta)J_i(X_i).$$

Together with the transversality condition $\gamma_i(\overline{\theta}) = 0$, we know $\gamma_i(\theta) \in [-\beta_i\underline{v}\theta(1 - F_i(\theta)), 0]$.

For any θ and any $X \in (0, \frac{\beta_i}{2}]$, we have:

$$\frac{H_i^I(\theta, X, U_i, \gamma_i) - H_i^I(\theta, 0, U_i, \gamma_i)}{f_i(\theta) X} = \frac{\theta J_i(X)}{X}(u(X) - U_i(\theta)) + \frac{\gamma_i}{\theta f_i(\theta) X} u(X) - c.$$

The above has derivative w.r.t. X of the same sign as:

$$\frac{\gamma_i}{\theta^2 f_i(\theta) J_i(X)} \left(\frac{X u'(X)}{u(X)} - 1 \right) + \frac{\Psi_i(G_i^{-1}(1 - \frac{X}{\beta_i}))}{G_i^{-1}(1 - \frac{X}{\beta_i})} + \frac{X u'(X)}{u(X)} - 1 + \frac{U_i(\theta)}{u(X)} \left(1 - \frac{\Psi_i(G_i^{-1}(1 - \frac{X}{\beta_i}))}{G_i^{-1}(1 - \frac{X}{\beta_i})} \right).$$

Note that the first term is positive, the sum of the next three is positive due to the assumption, and the last term is non-negative. Then, whenever $H_i^I(\theta, \frac{\beta_i}{2}, U_i, \gamma_i) > 0$, we have $H_i^I(\theta, \frac{\beta_i}{2}, U_i, \gamma_i) > H_i^I(\theta, X, U_i, \gamma_i)$ for all $X \in [0, \frac{\beta_i}{2})$, and therefore $X \in (0, \frac{\beta_i}{2})$ is never optimal for any θ.

Fix any θ such that $X_i(\theta) = X_i \geqslant \frac{\beta_i}{2}$. Then for any $\tilde{X}_i \in [\frac{\beta_i}{2}, X_i)$:

$$\frac{\mathrm{d}}{\mathrm{d}\theta} \frac{H_i^I(\theta, X_i, U_i, \gamma_i)}{f_i(\theta)} - \frac{\mathrm{d}}{\mathrm{d}\theta} \frac{H_i^I(\theta, \tilde{X}_i, U_i, \gamma_i)}{f_i(\theta)} = (u(X_i) - u(\tilde{X}_i)) \left(J_i(X_i) + J_i(\tilde{X}_i) - \frac{\gamma_i(\theta) f_i'(\theta)}{\theta f_i(\theta)^2} \right).$$

If $f_i'(\theta) \geqslant 0$, then the above is positive. If $f_i'(\theta) < 0$, then because $\gamma_i \geqslant -\beta_i \underline{v} \theta (1 - F_i(\theta))$ and $f_i'(\theta) \geqslant -\frac{f_i(\theta)^2}{1 - F_i(\theta)}$, we know

$$J_i(X_i) + J_i(\tilde{X}_i) - \frac{\gamma_i(\theta) f_i'(\theta)}{\theta f_i(\theta)^2} > 2 \frac{\beta_i}{2} \underline{v} - \beta_i \underline{v} = 0.$$

By Topkis' theorem, $X_i(\theta)$ is increasing in θ. Define $\theta_i^M \equiv \inf\{\theta : X_i(\theta) > \beta_i\}$ and $\theta_i^H \equiv \inf\{\theta : X_i(\theta) = \beta\}$.

Note that $e_i(\theta) = u(X_i(\theta)) - U_i(\theta)$ and $e_i(0) = 0$. Differentiating w.r.t. θ and using $U_i'(\theta) = \frac{e_i(\theta)}{\theta}$, we have an ODE:

$$\mathrm{d}(\theta e_i(\theta)) = \theta \, \mathrm{d}u(X_i(\theta)),$$

which yields the expression in Proposition 2. Because X_i is increasing, we verify that $\theta e_i(\theta) \geqslant 0$ and is increasing in θ. Because $e_i(0) = 0$ and $\theta e_i(\theta) \geqslant 0$ for $\theta > 0$, we verify that $e_i(\theta) \geqslant 0$.

If $X_i(\hat{\theta}) > 0$ for all $\hat{\theta} \in (\theta, \theta + \varepsilon)$, then $\hat{\theta} U_i(\hat{\theta}) = \int_0^{\hat{\theta}} u(X_i(\theta)) \mathrm{d}\theta > 0$, and $X_i(\theta) \geqslant u^{-1}(U_i(\theta)) \geqslant u^{-1}(U_i(\hat{\theta})) > 0$ for all $\theta > \hat{\theta}$. Define $\theta_i^L \equiv \inf\{\theta : X_i(\theta) > 0\}$. By definition, $\theta_i^H \geqslant \theta_i^M \geqslant \theta_i^L$. Moreover, $\frac{\partial}{\partial \tilde{X}_i} \frac{H_i^I(0, \tilde{X}_i, U_i, \gamma_i)}{f_i(0)} = -c - \frac{u'(\tilde{X}_i)}{\theta f_i(0)} \int_0^{\overline{\theta}} J_i(X_i(\theta)) \mathrm{d}F_i(\theta) < 0$, so there exists a neighborhood of 0 where $X_i(\theta) = 0$, i.e., $\theta_i^L > 0$. Finally, $\frac{\partial}{\partial X_i} \frac{H_i^I(\theta_i^L, 0, U_i, \gamma_i)}{f_i(\theta_i^L)} = -c + \frac{\gamma_i(\theta_i^L)}{\theta_i^L f_i(\theta_i^L)} u'(0) < 0$, which means $X_i(\theta_i^L) > 0$.

Now we compare the thresholds between complete and incomplete information cases. For any θ

such that $H_i^I(\theta, X_i, U_i, \gamma_i) \geqslant H_i^I(\theta, 0, U_i, \gamma_i)$, we have

$$\frac{\gamma_i u(X_i)}{\theta f_i(\theta)} + \beta_i \theta (u(X_i) - U_i(\theta)) \int_{G_i^{-1}(1-\frac{\min\{X_i,\beta_i\}}{\beta_i})}^{\overline{v}} \Psi_i(v) \mathrm{d}G_i(v) \geqslant cX_i,$$

implying $\beta_i \theta u(X_i) \int_{G_i^{-1}(1-\frac{\min\{X_i,\beta_i\}}{\beta_i})}^{\overline{v}} v \mathrm{d}G_i(v) > cX_i$, i.e., $H_i^C(\theta, X_i) > H_i^C(\theta, 0)$ with complete information. By continuity of H_i^C, we have $\theta_i^L > \theta_i^l$. For any θ such that $H_i^I(\theta, X^H, U_i, \gamma_i) - H_i^I(\theta, X^L, U_i, \gamma_i) \geqslant 0$ whenever $X^H > \beta_i \geqslant X^L$, we know $\gamma_i = -\beta_i \underline{v} \theta (1 - F(\theta))$, and $H_i^I(\theta, X^H, 0, \gamma_i) - H_i^I(\theta, X^L, 0, \gamma_i) \geqslant 0$. Notice that

$$\frac{H_i^C(\theta, X^H) - H_i^C(\theta, \beta_i)}{u(X^H) - u(\beta_i)} - \frac{H_i^I(\theta, X^H, 0, \gamma_i) - H_i^I(\theta, \beta_i, 0, \gamma_i)}{u(X^H) - u(\beta_i)} = \beta_i (\mathbb{E}_i[v]\theta - \underline{v}\Phi_i(\theta)) > 0.$$

Moreover,

$$\frac{\partial}{\partial X^L}(H_i^C(\theta, X^H) - H_i^C(\theta, X^L)) - \frac{\partial}{\partial X^L}(H_i^I(\theta, X^H, 0, \gamma_i) - H_i^I(\theta, X^L, 0, \gamma_i)) < 0$$

implies $H_i^C(\theta, X^H) - H_i^C(\theta, X^L) > H_i^I(\theta, X^H, 0, \gamma_i) - H_i^I(\theta, X^L, 0, \gamma_i) \geqslant 0$ for all $X^L < \beta_i$. By continuity of H_i^C and $\frac{\partial H_i^C}{\partial X}$, we have $\theta_i^M > \theta_i^m$. Finally, for any θ such that $H_i^I(\theta, \beta, U_i, \gamma_i) - H_i^I(\theta, X_i, U_i, \gamma_i) \geqslant 0$ for all X_i, the previous inequalities imply $H_i^C(\theta, \beta) - H_i^C(\theta, X_i) > 0$ for all $X_i \leqslant \beta_i$, and furthermore, $H_i^C(\theta, \beta) - H_i^C(\theta, X_i) > 0$ for all $X_i > \beta_i$ because $U_i > 0$, $\mathbb{E}_i[v] > \underline{v}$ and $\theta \geqslant \Phi_i(\theta)$. By continuity of H_i^C, we have $\theta_i^H > \theta_i^h$. ■

Proof of Corollary 2. Define $J_i(X)$ as in the proof of Proposition 2. If the inequality holds, then for any $\theta < \theta_i^L$ and any X, we have:

$$\frac{H_i^I(\theta, X, U_i, \gamma_i)}{f_i(\theta) X} = \left(\theta \frac{J_i(X)}{X} + \frac{\gamma_i}{\theta f_i(\theta) X}\right)(u(X) - U_i(\theta)) - c.$$

Note $U_i(\theta) = 0$. The above is negative at $\theta = 0$, and its derivative w.r.t. $X < \beta_i$ has the same sign as:

$$\frac{\gamma_i}{\theta^2 f_i(\theta) J_i(X)} \left(\frac{X u'(X)}{u(X)} - 1\right) + \frac{\Psi_i(G_i^{-1}(1 - \frac{X}{\beta_i}))}{G_i^{-1}(1 - \frac{X}{\beta_i})} + \frac{X u'(X)}{u(X)} - 1.$$

Note that the first term is positive, and the sum of the last three is positive due to the assumption. Then, we must have $X_i(\theta_i^L) \geqslant \beta_i$. According to Proposition 2, $X_i(\theta) \geqslant \beta_i$ for all $\theta \geqslant \theta_i^L$, and $\gamma_i(\theta) = -\beta_i \underline{v} \theta (1 - F_i(\theta))$. Define $\theta_i^L = \Phi_i^{-1}\left(\frac{c}{u(\beta_i)\mathbb{E}_i[\Psi_i(v)]}\right)$ so that $H_i^I(\theta_i^L, \beta_i, U_i, \gamma_i) = 0$. The above derivative w.r.t. $X > \beta_i$ is negative at θ_i^L, so indeed $X_i(\theta) = \beta_i$.

Moreover, when u is concave, $\frac{\partial}{\partial X} \frac{H_i^I(\theta, X, U_i, \gamma_i)}{f_i(\theta)} = \beta_i \mathbb{E}_i[\Psi_i(v)] \Phi_i(\theta) u'(X_i) = c$ for $X_i \geqslant \beta_i$. Setting $X_i(\theta_i^M) = \beta_i$ and $X_i(\theta_i^H) = \beta$ yields the thresholds. ■

Proof of Proposition 3. Facing function $\chi_i(q)$, a (θ, i)-creator maximizes:

$$u(\chi_i(q)) - \frac{q}{\theta}.$$

Implicit function theorem gives $\chi_i'(q) = \frac{1}{X_i^{-1}(\chi_i(q))u'(\chi_i(q))}$. Derivative w.r.t. q reads $\frac{1}{X_i^{-1}(\chi_i(q))} - \frac{1}{\theta}$. If $q = \theta e_i(\theta)$, then by definition $\chi_i(q) = X_i(\theta)$ and $\frac{1}{X_i^{-1}(\chi_i(q))} - \frac{1}{\theta} = 0$. If $q > \theta e_i(\theta)$, then $\chi_i(q) > X_i(\theta)$ and $X_i^{-1}(\chi_i(q)) \geqslant \theta$, and the derivative is non-positive. If $q < \theta e_i(\theta)$, then $\chi_i(q) < X_i(\theta)$ and $X_i^{-1}(\chi_i(q)) \leqslant \theta$, and the derivative is non-negative. ■

Proof of Proposition 4. Facing function $K_j(A)$, a (v, j)-viewer maximizes:

$$vK_j(A) - cA.$$

By construction of the direct mechanism, the (v, j)-viewer is willing to participate and does not mimic any (v', j)-type. Then:

$$\begin{aligned} vK_j(A_j(v)) - cA_j(v) \geqslant vQ_j(v) - cA_j(v) &\geqslant \sum_k \rho_k(vQ_j(\hat{v}_k) - cA_j(\hat{v}_k)) \\ &= v\sum_k \rho_k Q_j(\hat{v}_k) - c\sum_k \rho_k A_j(\hat{v}_k) \end{aligned}$$

for any $\{\rho_k\}_k \subset \mathbb{R}_+$ with $\sum_k \rho_k \leqslant 1$ and $\{\hat{v}_k\}_k \subset [\underline{v}, \overline{v}]$. By definition, this means:

$$vK_j(A_j(v)) - cA_j(v) \geqslant vK_j(A) - cA$$

for all $A \in [0, A_j(\overline{v})]$, and $Q_j(v) = K_j(A_j(v))$. Therefore, total attention $A_j(v)$ is implemented for viewer (v, j) even though $A \notin \{A_j(v)\}_v$ is allowed. ■

Proof of Proposition 5. We first ignore the constraints that $\theta e_i(\theta)$ is increasing and $e_i(\theta) \geqslant 0$, solve the problem, and then verify later. The problem is additively separable in creator categories. We can rewrite category i's objective as:

$$f_i(\theta)\sum_j \beta_j \int_V x_{ij}(\theta, v)(\Psi_j(v)\theta e_i(\theta)\mathbb{1}\{i = j\} - c)\mathrm{d}G_i(v) - f_i(\theta)\frac{c}{z}T_i(\theta) + \gamma_i(\theta)\frac{e_i(\theta)}{\theta},$$

$$\text{where} \quad e_i(\theta) = u\left(\pi\left(\sum_j \beta_j \int_V x_{ij}(\theta, v)\mathrm{d}G_i(v)\right) + T_i(\theta)\right) - U_i(\theta).$$

With the same argument as in the proof of Proposition 2, $x_{ii}(\theta, \cdot)$ is a step function. Moreover, $X_i(\theta) \leqslant \beta_i$ implies $R_i(\theta) = 0$, and $X_i(\theta) > \beta_i$ implies $x_{ii}(\theta) = 1$ for all v. Define $J_i(X)$ as in the proof of Proposition 2. The Hamiltonian can be simplified to:

$$\frac{H_i^z(\theta, X_i, Y_i, U_i, \gamma_i)}{f_i(\theta)} \equiv \theta e_i(\theta)J_i(X_i) - cX_i - \frac{c}{z}(Y_i - \pi(X_i)) + \frac{\gamma_i(\theta)}{f_i(\theta)}\frac{e_i(\theta)}{\theta},$$

where $e_i(\theta) = u(Y_i(\theta)) - U_i(\theta)$ and $Y_i = \pi(X_i) + T_i$.

Optimality requires $\frac{\partial H_i^z}{\partial U_i} = -\gamma_i'(\theta)$, which means:

$$\gamma_i'(\theta) = \frac{\gamma_i(\theta)}{\theta} + \theta f_i(\theta) J_i(X_i).$$

Together with the transversality condition $\gamma_i(\overline{\theta}) = 0$, we know $\gamma_i(\theta) \in [-\beta_i \underline{v}\theta(1 - F(\theta)), 0]$.

For any θ and any $X \in (0, \frac{\beta_i}{2}]$, we have:

$$\begin{aligned} & \frac{H_i^z(\theta, X, \pi(X) + T, U_i, \gamma_i) - H_i^z(\theta, 0, T, U_i, \gamma_i)}{f_i(\theta) X} \\ = \ & \theta \frac{J_i(X)}{X}(u(\pi(X) + T) - U_i(\theta)) + \frac{\gamma_i}{\theta f_i(\theta) X}(u(\pi(X) + T) - u(T)) - c. \end{aligned}$$

Its total derivative w.r.t. X has the same sign as:

$$\begin{aligned} & \frac{\gamma_i}{\theta^2 f_i(\theta) J_i(X)} \left(\frac{u(T) + X\pi'(X) u'(\pi(X) + T)}{u(\pi(X) + T)} - 1 \right) \\ + \ & \frac{\Psi_i(G_i^{-1}(1 - \frac{X}{\beta_i}))}{G_i^{-1}(1 - \frac{X}{\beta_i})} + \frac{X\pi'(X) u'(\pi(X) + T)}{u(\pi(X) + T)} - 1 + \frac{U_i(\theta)}{u(\pi(X) + T)} \left(1 - \frac{\Psi_i(G_i^{-1}(1 - \frac{X}{\beta_i}))}{G_i^{-1}(1 - \frac{X}{\beta_i})} \right). \end{aligned}$$

Note that the first term is positive, the sum of the next three is positive due to the assumption, and the last term is non-negative. Then, whenever $H_i^z(\theta, \frac{\beta_i}{2}, \pi(\frac{\beta_i}{2}) + T, U_i, \gamma_i) > H_i^z(\theta, 0, T, U_i, \gamma_i)$, we have $H_i^z(\theta, \frac{\beta_i}{2}, \pi(\frac{\beta_i}{2}) + T, U_i, \gamma_i) > H_i^z(\theta, X, \pi(X) + T, U_i, \gamma_i)$ for all $X \in [0, \frac{\beta_i}{2})$, and therefore $X \in (0, \frac{\beta_i}{2})$ is never optimal for any θ.

Fix any θ such that $X_i(\theta) = X_i \geqslant \frac{\beta_i}{2}$. Then for any $\tilde{X}_i \geqslant \frac{\beta_i}{2}$ and any $\tilde{Y}_i$:

$$\frac{\mathrm{d}}{\mathrm{d}\theta} \frac{H_i^z(\theta, X_i, Y_i, U_i, \gamma_i)}{f_i(\theta)} - \frac{\mathrm{d}}{\mathrm{d}\theta} \frac{H_i^z(\theta, \tilde{X}_i, \tilde{Y}_i, U_i, \gamma_i)}{f_i(\theta)} = (u(Y_i) - u(\tilde{Y}_i)) \left(J_i(X_i) + J_i(\tilde{X}_i) - \frac{\gamma_i(\theta) f_i'(\theta)}{\theta f_i(\theta)^2} \right).$$

If $f_i'(\theta) \geqslant 0$, then the above is positive. If $f_i'(\theta) < 0$, then because $\gamma_i \geqslant -\beta_i \underline{v}\theta(1 - F_i(\theta))$ and $f_i'(\theta) \geqslant -\frac{f_i(\theta)^2}{1 - F_i(\theta)}$, we know

$$J_i(X_i) + J_i(\tilde{X}_i) - \frac{\gamma_i(\theta) f_i'(\theta)}{\theta f_i(\theta)^2} > 2\frac{\beta_i}{2}\underline{v} - \beta_i \underline{v} = 0.$$

Therefore, $H_i^z(\theta, X_i, Y_i, U_i, \gamma_i) \geqslant H_i^z(\theta, \tilde{X}_i, \tilde{Y}_i, U_i, \gamma_i)$ for some $\tilde{Y}_i \leqslant Y_i$ implies $H_i^z(\hat{\theta}, X_i, Y_i, U_i, \gamma_i) \geqslant H_i^z(\hat{\theta}, \tilde{X}_i, \tilde{Y}_i, U_i, \gamma_i)$ for all $\hat{\theta} > \theta$. Moreover, $H_i^z(\theta, X_i, Y_i, U_i, \gamma_i) \geqslant H_i^z(\theta, \tilde{X}_i, \tilde{Y}_i, U_i, \gamma_i)$ for some $\tilde{Y}_i > Y_i$ and $\tilde{X}_i < X_i$ implies $H_i^z(\hat{\theta}, X_i, \tilde{Y}_i, U_i, \gamma_i) \geqslant H_i^z(\hat{\theta}, \tilde{X}_i, \tilde{Y}_i, U_i, \gamma_i)$ for all $\hat{\theta} > \theta$. In sum, a higher θ can only lead to weakly higher $X_i(\theta)$ and $Y_i(\theta)$. Define $\theta_i^M \equiv \inf\{\theta : X_i(\theta) > \beta_i\}$ and $\theta_i^H \equiv \inf\{\theta : X_i(\theta) = \beta\}$.

Note that $e_i(\theta) = u(Y_i(\theta)) - U_i(\theta)$ and $e_i(0) = 0$. Differentiating w.r.t. θ and using $U_i'(\theta) = \frac{e_i(\theta)}{\theta}$, we have an ODE:

$$\mathrm{d}(\theta e_i(\theta)) = \theta \, \mathrm{d}u(Y_i(\theta)),$$

which yields the expression in Proposition 5. Because Y_i is increasing, we verify that $\theta e_i(\theta) \geqslant 0$ and is increasing in θ. Because $e_i(0) = 0$ and $\theta e_i(\theta) \geqslant 0$ for $\theta > 0$, we verify that $e_i(\theta) \geqslant 0$.

If $Y_i(\hat{\theta}) > 0$ for all $\hat{\theta} \in (\theta, \theta + \varepsilon)$, then $\hat{\theta} U_i(\hat{\theta}) = \int_0^{\hat{\theta}} u(Y_i(\theta))\mathrm{d}\theta > 0$, and $Y_i(\theta) \geqslant u^{-1}(U_i(\theta)) \geqslant u^{-1}(U_i(\hat{\theta})) > 0$ for all $\theta > \hat{\theta}$. Also note that $X_i(\theta) = 0$ implies $Y_i(\theta) = 0$ from the definition of H_i^z. Therefore, define $\theta_i^L \equiv \inf\{\theta : Y_i(\theta) > 0\} = \inf\{\theta : X_i(\theta) > 0\}$. By definition, $\theta_i^H \geqslant \theta_i^M \geqslant \theta_i^L$. Moreover, $\frac{\partial}{\partial \tilde{Y}_i} \frac{H_i^z(0, \tilde{X}_i, \tilde{Y}_i, U_i, \gamma_i)}{f_i(0)} = -\frac{c}{z} - \frac{u'(\tilde{Y}_i)}{\theta f_i(0)} \int_0^{\overline{\theta}} J_i(X_i(\theta))\mathrm{d}F_i(\theta) < 0$, so that $\tilde{Y}_i = \pi(\tilde{X}_i)$. Plugging back in H_i^z, we have $\frac{\mathrm{d}}{\mathrm{d}\tilde{X}_i} \frac{H_i^z(0, \tilde{X}_i, \pi(\tilde{X}_i), U_i, \gamma_i)}{f_i(0)} = -c - \frac{\pi'(\tilde{X}_i)u'(\pi(\tilde{X}_i))}{\theta f_i(0)} \int_0^{\overline{\theta}} J_i(X_i(\theta))\mathrm{d}F_i(\theta) < 0$. This means for θ close enough to 0 we have $X_i(\theta) = 0$, i.e., $\theta_i^L > 0$. Finally, $\frac{\partial}{\partial Y_i} \frac{H_i^z(\theta_i^L, 0, 0, U_i, \gamma_i)}{f_i(\theta_i^L)} = -\frac{c}{z} + \frac{\gamma_i(\theta_i^L)}{\theta_i^L f_i(\theta_i^L)} u'(0) < 0$, and plugging back into H_i^z, we have $\frac{\mathrm{d}}{\mathrm{d}X_i} \frac{H_i^z(\theta_i^L, 0, 0, U_i, \gamma_i)}{f_i(\theta_i^L)} = -c + \frac{\gamma_i(\theta_i^L)}{\theta_i^L f_i(\theta_i^L)} u'(0)\pi'(0) < 0$, which means $X_i(\theta_i^L) > 0$.

When $X_i(\theta) \geqslant \beta_i$, we know $\gamma_i = -\theta \beta_i \underline{v}(1 - F_i(\theta))$, and the relevant terms in the Hamiltonian reads:

$$\beta_i \underline{v} \Phi_i(\theta) u(Y_i) - \frac{c}{z} Y_i + \frac{c}{z} \pi(X_i) - c X_i.$$

If $T_i > 0$, then $Y_i > \pi(X_i)$. Unconstrained optimization over Y_i yields $Y_i = u'^{-1}(\frac{c}{z\beta_i \underline{v} \Phi_i(\theta)})$. Fixing Y_i, the relevant terms for X_i is always $\frac{c}{z}\pi(X_i) - cX_i$, independent of θ, except the constraint that $\pi(X_i) \leqslant u'^{-1}(\frac{c}{z\beta_i \underline{v} \Phi_i(\theta)})$. As θ continuously increases, so does $u'^{-1}(\frac{c}{z\beta_i \underline{v} \Phi_i(\theta)})$. Therefore, using the tie-breaking rule that favors higher attention, we know that $T_i > 0$ means the constraint is not binding, and at marginally lower θ, the optimal X_i does not change. Alternatively, if X_i is increasing in θ, it must be the case that X_i attained at the boundary $\pi(X_i) \leqslant u'^{-1}(\frac{c}{z\beta_i \underline{v} \Phi_i(\theta)})$, in which case $T_i(\theta) = 0$.

When π is concave and $\pi'(\beta_i) > z$, we know $T_i(\theta) = 0$ whenever $X_i(\theta) \leqslant \beta_i$ because the derivative of the Hamiltonian w.r.t. $X_i(\theta)$ reads $c(\frac{\pi'(X_i)}{z} - 1) + \theta u(Y_i) J'(X_i) > 0$ and $Y_i \geqslant \pi(X_i)$ must be binding. For $X_i(\theta) > \beta_i$, we need $\theta > \theta_i^M$. For the objective relevant for X_i, $\frac{c}{z}\pi(X_i) - cX_i$, first order condition requires $X_i = \pi'^{-1}(z)$. If $\pi'^{-1}(z) < \beta$, then $\theta_i^\$$ is determined by $u'^{-1}(\frac{c}{z\beta_i \underline{v} \Phi_i(\theta_i^\$)}) = \pi(\pi'^{-1}(z))$. If $\pi'^{-1}(z) \geqslant \beta$, then $\theta_i^\$$ is determined by $u'^{-1}(\frac{c}{z\beta_i \underline{v} \Phi_i(\theta_i^\$)}) = \pi(\beta)$. ■

Proof of Lemma 4. When $X_i(\theta) \geqslant \beta_i$, the Hamiltonian boils down to:

$$\frac{H_i^z(\theta, X_i, \pi(X_i) + T_i, U_i, \gamma_i)}{f_i(\theta)} \equiv \theta e_i(\theta) \beta_i \underline{v} - c(X_i + \tfrac{T_i}{z}) + \frac{\gamma_i(\theta)}{f_i(\theta)} \frac{u(\pi(X_i) + T_i) - U_i(\theta)}{\theta}.$$

Therefore, an optimal combination of X_i and T_i must maximize $Y_i = \pi(X_i) + T_i$ given any attention-metric total resource $\hat{X}_i = X_i + \frac{T_i}{z}$. In other words, $Y_i = \max\{\pi(X_i) + T_i : X_i + \frac{T_i}{z} = \hat{X}_i, T_i \geqslant 0\} = \hat{\pi}(\hat{X}_i)$ by definition of $\hat{\pi}$.

The dual problem minimizes $X_i + \frac{T_i}{z}$ subject to $\pi(X_i) + T_i = Y_i$. Therefore, $\hat{\pi}^{-1}(Y_i) = \hat{X}_i = X_i + \frac{Y_i - \pi(X_i)}{z}$. Using the tie-breaking rule that favors higher attention, we have $X_i(\theta) = \max\{X :$

$\pi(X) + z(\hat{\pi}^{-1}(Y_i(\theta)) - X) = Y_i(\theta)\}$. ■

Proof of Proposition 6. Notice that $\frac{\partial^2}{\partial z \partial X_i} \frac{H_i^z}{f_i(\theta)} = 0$, $\frac{\partial^2}{\partial z \partial T_i} \frac{H_i^z}{f_i(\theta)} = \frac{c}{z^2} > 0$, $\frac{\partial^2}{\partial X_i \partial U_i} \frac{H_i^z}{f_i(\theta)} = 0$, $\frac{\partial^2}{\partial T_i \partial U_i} \frac{H_i^z}{f_i(\theta)} = 0$, and finally $\frac{\partial^2}{\partial X_i \partial T_i} \frac{H_i^z}{f_i(\theta)} = \frac{\pi'(X_i) u''(\pi(X_i)+T_i)}{\theta f_i(\theta)} (\beta_i \underline{v} \theta^2 f_i(\theta) + \gamma_i) < 0$, where the last inequality follows from the fact that $H_i^z(\theta, X_i, \pi(X_i) + T_i, U_i, \gamma_i) \geqslant H_i^z(\theta, 0, 0, U_i, \gamma_i)$. Moreover, $T_i \geqslant 0$ and $X_i \geqslant \beta_i$ is a sub-lattice, and the continuous-time version of Topkis theorem implies $X_i(\theta)$ is decreasing in z and $T_i(\theta)$ is increasing in z for all θ. ■

References

Acemoglu, D., A. Ozdaglar, and J. Siderius (2024). A model of online misinformation. *Review of Economic Studies 91*(6), 3117–3150.

Allcott, H., L. Braghieri, S. Eichmeyer, and M. Gentzkow (2020). The welfare effects of social media. *American economic review 110*(3), 629–676.

Ambrus, A., E. Calvano, and M. Reisinger (2016). Either or both competition: A "two-sided" theory of advertising with overlapping viewerships. *American Economic Journal: Microeconomics 8*(3), 189–222.

Anderson, S. P. and S. Coate (2005). Market provision of broadcasting: A welfare analysis. *The review of Economic studies 72*(4), 947–972.

Anderson, S. P. and B. Jullien (2015). The advertising-financed business model in two-sided media markets. In *Handbook of media economics*, Volume 1, pp. 41–90. Elsevier.

Aridor, G., R. Jiménez-Durán, R. Levy, and L. Song (2024). The economics of social media.

Bar-Isaac, H., R. Deb, and M. Mitchell (2025). Selling certification, content moderation, and attention. *arXiv preprint arXiv:2506.12604*.

Bergemann, D. and A. Bonatti (2011). Targeting in advertising markets: implications for offline versus online media. *The RAND Journal of Economics 42*(3), 417–443.

Bergemann, D. and A. Bonatti (2024). Data, competition, and digital platforms. *American Economic Review 114*(8), 2553–2595.

Bergemann, D. and J. Välimäki (2019). Dynamic mechanism design: An introduction. *Journal of Economic Literature 57*(2), 235–274.

Berman, R. and Z. Katona (2020). Curation algorithms and filter bubbles in social networks. *Marketing Science 39*(2), 296–316.

Bisceglia, M., A. Bonatti, and F. S. Morton (2025). Regulating privacy policies on digital platforms. *Working Paper, Yale University*.

Casner, B. and T.-H. Teh (2025). Content-hosting platforms: Discovery, membership, or both? *The RAND Journal of Economics 56*(4), 668–688.

Chen, D. (2025). A model of the attention economy. *Princeton University, Working Paper*.

Choi, J. P. and D.-S. Jeon (2023). Platform design biases in ad-funded two-sided markets. *The RAND Journal of Economics 54*(2), 240–267.

Choi, J. P., D.-S. Jeon, and B.-C. Kim (2015). Net neutrality, business models, and internet interconnection. *American Economic Journal: Microeconomics 7*(3), 104–141.

Damiano, E. and H. Li (2007). Price discrimination and efficient matching. *Economic Theory 30*(2), 243–263.

de Corniere, A. and M. Sarvary (2023). Social media and news: Content bundling and news quality. *Management Science 69*(1), 162–178.

Fershtman, D. and A. Pavan (2022). Matching auctions. *The RAND Journal of Economics 53*(1), 32–62.

Filippas, A., J. J. Horton, E. Lipnowski, and P. Parasurama (2025). The production and consumption of social media. *Management Science*.

Gomes, R. and A. Pavan (2016). Many-to-many matching and price discrimination. *Theoretical Economics 11*(3), 1005–1052.

Gomes, R. and A. Pavan (2024). Price customization and targeting in matching markets. *The Rand journal of economics 55*(2), 230–265.

Hellwig, M. F. (2003). Public-good provision with many participants. *The Review of Economic Studies 70*(3), 589–614.

Ichihashi, S., D.-S. Jeon, and B.-C. Kim (2025). Mechanism design for ad-supported platforms. *Available at SSRN 5016687*.

Ichihashi, S. and A. Smolin (2025). Buyer-optimal algorithmic consumption. *Available at SSRN 4635866*.

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

Johnson, J. P. (2013). Targeted advertising and advertising avoidance. *The RAND journal of economics 44*(1), 128–144.

Johnson, J. P., A. Rhodes, and M. Wildenbeest (2023). Platform design when sellers use pricing algorithms. *Econometrica 91*(5), 1841–1879.

Jullien, B., A. Pavan, and M. Rysman (2021). Two-sided markets, pricing, and network effects. *4*(1), 485–592.

Kranton, R. and D. McAdams (2024). Social connectedness and information markets. *American Economic Journal: Microeconomics 16*(1), 33–62.

Li, Y. (2021). Mechanism design with financially constrained agents and costly verification. *Theoretical Economics 16*(3), 1139–1194.

Lockwood, B. (2000). Production externalities and two-way distortion in principal-multi-agent problems. *Journal of Economic Theory 92*(1), 142–166.

Meisner, V. and P. Pillath (2024). Monetizing digital content with network effects: A mechanism-design approach. *arXiv preprint arXiv:2408.15196*.

Myerson, R. B. (1986). Multistage games with communication. *Econometrica: Journal of the Econometric Society*, 323–358.

Norman, P. (2004). Efficient mechanisms for public goods with use exclusions. *The Review of Economic Studies 71*(4), 1163–1188.

Oren, S. S., S. A. Smith, and R. B. Wilson (1982). Nonlinear pricing in markets with interdependent demand. *Marketing Science 1*(3), 287–313.

Pai, M. M. and R. Vohra (2014). Optimal auctions with financially constrained buyers. *Journal of Economic Theory 150*, 383–425.

Petrova, M., A. Sen, and P. Yildirim (2021). Social media and political contributions: The impact of new technology on political competition. *Management Science 67*(5), 2997–3021.

Prat, A. and T. Valletti (2022). Attention oligopoly. *American Economic Journal: Microeconomics 14*(3), 530–557.

Qian, K. and S. Jain (2024). Digital content creation: An analysis of the impact of recommendation systems. *Management Science*.