

The Supply Side of Social Media*

Shakked Noy[†]

MIT

Job Market Paper

Aakaash Rao[‡]

Northwestern Kellogg

October 29, 2026

Updated frequently; please click [here](#) for the latest version

Abstract

We study social media’s distinctive model of content production. Unlike legacy media outlets, social media platforms do not professionally produce content optimized to appeal to viewers. Instead, they harvest the output of large numbers of amateur users creating content in pursuit of nonpecuniary rewards. The resulting scale and diversity of platforms’ content libraries, combined with their ability to experimentally test new pieces on subsets of viewers, allows them to ex-post discover the pieces most appealing to each viewer. We develop a model of this process and show that it leads the content consumed in equilibrium to depend on viewer preferences (demand) but also on creator preferences (supply), because the scale powering platforms’ discovery concentrates where creators are eager to produce. Collecting new data on Instagram, TikTok, and X/Twitter, we show that viewers’ feeds are dominated by unpaid amateurs who produce envy-provoking “lifestyle” and partisan political content in pursuit of social status and expressive satisfaction. We experimentally generate supply- and demand-side shifters by manipulating viewers’ feeds and eliciting creators’ incentivized preferences while paying them to produce content. On the demand side, pieces of lifestyle and partisan content are less likely to highly engage viewers than more benign entertainment. But creators enjoy creating lifestyle and partisan content and therefore produce it at enormous scale. This supply-driven scale enables platforms to cherry-pick the minority of such pieces viewers find more engaging than the marginal piece of entertainment. Estimating our model using the experimental variation, we calculate that supply-side preferences account for one- to two-thirds of the prevalence of lifestyle and partisan content in viewers’ feeds.

*We are indebted to Daron Acemoglu, Frank Schilbach, Sendhil Mullainathan, and Tobias Salz for their guidance and support. For helpful comments, we thank Charles Angelucci, David Autor, Lucas Barros, Doug Bernheim, Valerie Chuang, John Conlon, Ben Enke, Nathan Lazarus, Alexey Makarin, James Okun, Nicola Persico, Charlie Rafkin, Chris Roth, Jesse Shapiro, Andrei Shleifer, David Yang, and Whitney Zhang. We gratefully acknowledge financial support from the George and Obie Shultz Fund, J-PAL North America, and the Kenneth C. Griffin Fund. Data collections in this paper were approved by the MIT Committee on the Use of Humans as Experimental Subjects under protocol numbers 2506001673 and 2606002035 and pre-registered on the AEA RCT Registry at AEARCTR-0019565 and on the Open Science Framework at <https://osf.io/8ctq7> and <https://osf.io/24tcz>.

[†]snoy@mit.edu

[‡]aakaash.rao@kellogg.northwestern.edu

1 Introduction

Over the past twenty years, social media platforms have gradually displaced legacy media outlets to become many users’ primary source of news and entertainment. Critics allege that by serving users content provoking political grievances and personal insecurities, and fragmenting their attention spans with short-form feeds, these platforms drive recent declines in political cohesion, mental wellbeing, and academic achievement.¹

The common view in research and popular commentary is that these problems originate on the demand side.² As the story goes, platforms’ algorithms excel at identifying what will engage viewers, and outrage and envy are highly engaging; moreover, platforms keep users compulsively scrolling by designing short-form feeds that, like slot machines, offer a stream of variable and unpredictable rewards. On this view, what differentiates social from legacy media is superior technology for identifying and exploiting latent demand.

This paper develops a new and complementary way of thinking about what makes social media different from legacy media. We focus on an understudied feature of social media: its distinctive production side. Unlike legacy outlets, social media platforms do not professionally produce material that viewers like to *watch*; instead, they harvest material that amateur users like to *create* in pursuit of social status, expressive satisfaction, or creative fulfillment. This lets them cheaply accumulate enormous libraries of content. They then sift through these libraries to engage in granular ex-post discovery of pieces that resonate with viewers.

This business model has two implications. First, because the production scale fueling platforms’ discovery concentrates selectively in regions of the content space where creators are most eager to produce, creators’ supply-side preferences can end up shaping what viewers consume. For example, the expressive satisfaction that partisan users derive from posting extreme political content, or the social status that attractive users gain from posting self-centered content, can carry polarization- or envy-provoking content into viewers’ feeds. Second, due to its emphasis on ex-post discovery, the business model favors media formats disaggregated enough to facilitate maximally granular cherry-picking, such as short posts and videos.

In this paper, we study the business model of social media in three steps. First, we develop a framework nesting legacy and social media. We characterize the technological advances enabling social media’s business model, explain why the model magnifies the influence of supply-side preferences and favors short-form content, and describe how it leads supply- and demand-side preferences to jointly determine consumption.

Second, we document the importance of nonpecuniary supply-side preferences by collecting data covering viewers and creators on Instagram, TikTok, and X/Twitter. We show that the majority of content consumed on these platforms is produced for nonpecuniary reasons by a minority of users selected on attractiveness and ideological extremity. These characteristics predict production of status-oriented “lifestyle” and partisan political content. Lifestyle and partisan content receive *lower* viewership on average and their production drops (lifestyle) or stays flat (partisan) when facing stronger incentives to attract views. But they are produced at such high volume that enough successful pieces of these types emerge to account for a third of consumption.

Finally, we conduct supply- and demand-side experiments to estimate the parameters required to quantitatively decompose equilibrium consumption of these content types into supply- and demand-driven components. We find that one- to two-thirds of consumption of these types is driven by the supply side.

¹ See Allcott et al. (2020); Levy (2021); Braghieri et al. (2022); Haidt (2024); Bursztyn et al. (2025); *People v. TikTok Inc.* (2024).

² See Eyal (2014); Allcott et al. (2022); Beknazar-Yuzbashev et al. (2024); Bao (2026); Solheim and Wang (2026); Bolte et al. (2026).

Conceptual Framework

Our framework features three groups of agents: creators who produce content in pursuit of pecuniary and nonpecuniary rewards; viewers with heterogeneous tastes; and a profit-maximizing outlet that pays creators to produce, distributes produced content to viewers, and extracts viewer surplus.

We assume content production is a noisy process (Goldman, 1983). A content “type” k is distinguished by a few features such as style, slant, or topic. Pieces of type k have a baseline appeal m_k reflecting viewers’ taste for these distinguishing features. When created, a piece draws a quality shock q and vector of taste shocks e_i for each viewer i , reflecting other features that emerge serendipitously during production, such as comic timing or visual or narrative resonance. Viewer i ’s utility from a produced piece is $m_k + q + e_i$.

The q and e_i shocks are realized only once a piece is produced, and even after production are not automatically visible to the outlet. Following canonical characterizations of legacy media (Waldfogel, 2019), we assume that legacy outlets face technological constraints on testing and distribution that prevent them from measuring and acting on the q and e_i shocks when disseminating content. They therefore optimally procure a few pieces of content with high ex-ante viewer appeal m_k and distribute them widely.

The Internet relaxes these constraints and births a new business model premised on procuring large numbers of pieces and testing them on subsets of viewers to discover their q and e_i realizations. This ex-post discovery step facilitates cherry-picking of pieces with high q shocks and personalized distribution of pieces to viewers based on e_i matches. By increasing the returns to production scale, it induces a shift towards content types that are cheap to procure in large quantities by virtue of being easy or satisfying to produce. It also rewards disaggregation of content into the smallest units over which quality and taste shocks are realized. The result is the emergence of new outlets harvesting and distributing short “social” posts.

Under this new business model, equilibrium consumption of content type k depends on *demand*—its distribution of appeal to viewers—and *supply*—creators’ nonpecuniary taste for producing it. Even if a content type has lower baseline appeal m , if creators like producing it, many pieces of that type will be produced, yielding more opportunities for high q or e draws to carry pieces of this type into viewers’ feeds.

For example, consider a viewer who seeks out political news or commentary and *ceteris paribus* prefers ideologically moderate content. If moderate creators dislike arguing about politics online, they will produce only a thin pool of political content, limiting a platform’s ability to cherry-pick moderate posts that this viewer will find unusually interesting or amusing. If, meanwhile, political extremists enjoy posting constantly about politics, the sheer scale of their output will throw up more unusually appealing outliers on dimensions other than slant. Our viewer may hence find herself with a feed full of extreme posts so much more amusing or interesting than the best available moderate posts that this overcomes her distaste for extremity.

Or suppose attractive users pursue social status by posting self-centered “lifestyle” content, so the supply of such content dramatically outpaces the supply of more effortful and less ego-gratifying comedy sketches. Then even if our viewer finds lifestyle content typically less interesting than comedy and more prone to triggering unfavorable social comparisons, she may end up watching cherry-picked lifestyle content so entertaining she prefers it to the best available comedy despite the envy it provokes.

The influence of creators’ tastes on equilibrium consumption is mediated by two factors. First, supply-

side tastes matter more when the variances of the q and e shocks are high. High-variance shocks represent viewer preferences that are difficult to target and are high-dimensional in the sense of depending on many characteristics other than the few captured by variation in m . If viewer preferences are predictable and low-dimensional, supply-side constraints never bind: professional creators enter, supply viewers with abundant content of their preferred types, and reap stable pecuniary rewards. By contrast, unpredictable and high-dimensional viewer preferences make it difficult to consistently attract viewers, eroding the scope for professionalization. Production becomes driven by nonpecuniary motives, and appealing content becomes locally scarce in the upper tail of $q + e$ draws. Consumption becomes driven not by a content type’s baseline appeal m , but by how much creators like producing it and therefore how many q and e draws it generates.

Second, the lower the pecuniary and nonpecuniary rewards creators derive from attracting viewership on the margin, the weaker the feedback from viewer taste for a content type to realized supply of that type.

Observational Evidence

If content production is dominated by professional creators seeking to maximize their audiences, or if super-abundance of appealing content allows platforms to construct feeds exactly matching viewers’ preferences, equilibrium consumption will be unaffected by nonpecuniary supply-side preferences. The comparative static predictions just outlined indicate both of these conditions will fail if viewer preferences are sufficiently unpredictable and dispersed. We therefore document a sequence of empirical facts confirming these predictions and providing smoking-gun evidence of the influence of specific supply-side preferences.

We collect four new datasets. First, we pay 383 young Americans to download their personal data on Instagram, TikTok, and YouTube and install a Chrome extension recording their X/Twitter feeds, yielding second-by-second histories of the posts they see on these platforms. We scrape the posts in participants’ histories, use LLMs to classify their contents, and scrape information about their creators. Second, we use TikTok’s Research API to construct a representative sample of the videos *created* on TikTok. Third, we contact 301 of the creators appearing in our participants’ histories and survey them about their motivations. Fourth, we measure selection into content creation by surveying 3,490 users of these platforms. For data quality reasons, our main results focus on TikTok; we discuss results for other platforms below.

We document three facts. First, most of the content viewers consume is produced for nonpecuniary reasons. Platforms benefit from a broad base of production driven by small creators who post in pursuit of status, expressive satisfaction, or creative fulfilment, and do not make money. Viewers’ feeds are dominated by one-off extreme hits that emerge in a sparse and unpredictable way from across this base of production. While large professional creators exist, they fill less than a third of the typical feed. These patterns are suggestive of high-variance ex-post q and e shocks. Our model explains that such shocks deprofessionalize production by making it hard to consistently optimize for viewership and predicts they create local scarcity of the best content. Indeed, although viewers engage predictably less with “stale” content posted longer ago, after 30 minutes of scrolling, half their feed is stale, indicating a paucity of appealing fresh material. We show that this scarcity lets unusually successful posts from disliked content types make it into viewers’ feeds.

Second, the small creators dominating feeds are positively selected relative to viewers on self-reported ideological extremity, attractiveness, and lifestyle desirability. These characteristics respectively predict

posting of partisan political content and of self-centered material we refer to as “lifestyle” content. Lifestyle and partisan content make up a third of the typical TikTok feed and about half the content created on TikTok.

Third, the production of lifestyle and partisan content does not appear to reflect attempted catering to viewer demand. We exploit the fact that TikTok’s Creator Rewards program, which pays eligible creators for each view they receive, applies only to videos at least a minute long. Financially motivated creators chasing viewership should therefore bunch their videos above 60 seconds. We show that crossing from 59 to 60 seconds substantially and discontinuously *reduces* the prevalence of lifestyle content and does not change the partisanship of political content. This suggests these content types are not deployed strategically to stir up engagement. Bunching is specific to accounts eligible for Creator Rewards and emerges after its rollout.

Overall, conditions appear ripe for supply-side preferences to influence consumption. The unpredictability of viewership deprofessionalizes production and makes appealing content locally scarce, while creators’ pursuit of status and expressive satisfaction seems to drive production of status-oriented and partisan content.

Experimentally Disentangling Supply and Demand

While directionally indicative of the importance of supply-side preferences, the evidence described above does not provide the variation necessary to quantitatively estimate the parameters of our model and thereby decompose equilibrium consumption of lifestyle and partisan content into supply- and demand-driven components. To generate this variation, we conduct two experiments.

In the canonical supply-demand setting, identification requires a supply shifter that traces out the local density of consumer surplus among consumers who do not purchase the product in equilibrium, and a demand shifter that traces out the elasticity of supply to the price for which produced units can be sold.

Unlike the canonical setting, our setting lacks product-specific prices that clear markets by transferring utility from consumers to producers. But identification can be described analogously. Identifying demand-side preferences requires a shifter that traces out viewers’ engagement with videos that in equilibrium they do not see in their feeds (and which may not even be created). To this end, we recruit viewers to scroll through a TikTok clone whose feed we manipulate. Identifying supply-side preferences requires a shifter that traces out the elasticity of supply to expected viewership. We generate this in an experiment where we elicit creators’ willingness to supply different content types while manipulating expectations about viewership.

For our demand-side experiment, we build a phone-based TikTok clone and populate it with videos randomly sampled from all of TikTok, plus videos whose creation was induced by our supply-side experiment. We recruit 586 TikTok users to scroll through it, measure their engagement behavior, and use it to estimate a parametric model of viewer utility from different content types. Our experimentally-derived demand estimates pass a range of overidentification tests in the observational data.

We have three results. First, lifestyle content faces the lowest demand of the types we study. For example, lifestyle videos are about 5 percentage points, or 25%, less likely than entertainment videos to be watched for at least ten seconds. Second, partisan political content enjoys a small demand premium over nonpartisan content.³ Third, consistent with across-type variation in baseline appeal m being small relative to the within-

³Both our experimental and discontinuity-based observational evidence suggest a demand penalty for lifestyle content. Observationally we found no partisanship premium, in contrast to a slight partisanship premium in our experiment. This is perhaps because, as we show, the general population of TikTok users is more moderate than the participants in our experiment. We therefore view our experimental estimates as an upper bound on the demand premium enjoyed by partisan content.

type variance of q and e shocks, engagement differences across content types are small and viewers like a substantial minority of videos of each type.⁴ This creates room for supply to matter. For example, if lifestyle videos are 25% less likely to highly engage viewer i than entertainment videos, a pool of lifestyle videos only needs to be 33% larger than a pool of entertainment videos in order to contain more videos that will highly engage i . In fact, the lifestyle pool produced on TikTok is 71% larger than the entertainment pool.

Overall, the results of our viewership experiment suggest that neither lifestyle nor partisan content enjoy a substantial demand premium, suggesting their considerable supply advantage may reflect nonpecuniary creator preferences rather than attempts to chase demand. We confirm this with our supply-side experiment.

We recruit 979 TikTok creators and offer to pay them to produce and post different types of content while also offering interventions that will increase their viewership. We thereby elicit their incentivized willingness to make posts of each type; their incentivized expectations about viewership; and their incentivized marginal valuation of viewership. We use these objects to back out the creator utility parameters in our model. The offers we extend induce the creation of 106 TikTok posts.

We find substantial variation in creators’ willingness to post different content types, but almost none of this variation is explained by expectations about demand. Creators expect only slightly lower viewership for content types they are unwilling to post, and in fact receive slightly *higher* viewership when paid to post these types. Moreover, although creators strongly value attention, their marginal valuation of viewership is low, reflecting sharp concavity of utility from attention.⁵ Instead, 82% of creators say they create certain types of content primarily because they find them easy and satisfying to create, and are unwilling to supply others because doing so would be difficult or uncomfortable.

Overall, our experimental results suggest that heavy production of lifestyle and partisan content by enviable and ideologically extreme users is driven by supply-side preferences rather than attempted catering to demand. These results derive from complex willingness-to-pay objects, but are robust to a much simpler design that varies whether we ask about creators’ binary willingness to *post* a piece or shut down expectations about demand by asking their willingness to *privately create* a piece to be shown to a few other participants.

We combine the estimates from our viewer and creator experiments to estimate the parameters of our model and perform decompositions. We have two sources of demand-side data—our viewership experiment, and viewership expectations from our creator experiment—and results are robust to using either source. We find that equalizing creators’ intrinsic preferences for producing different content types would reduce equilibrium consumption of both lifestyle and partisan political content by one- to two-thirds.

Cross-Platform Comparisons

For data quality reasons, our experiments and descriptive results focus on TikTok. However, we show that the descriptive patterns we document replicate on Instagram and X/Twitter. YouTube looks closer to television than to the other social media platforms: it pays creators more generously, creators are much more professionalized, and viewership is more consistent and predictable. Perhaps commensurately, participants in our surveys express much more positive views of YouTube than of the other platforms.

⁴To distinguish preference dispersion from noisy behavior, we ask participants at the end to choose between pairs of videos they saw, and use concordance of their choices with their prior engagement to estimate and purge behavioral noise; this barely affects results.

⁵This finding is consistent with prior work on social media (Srinivasan, 2024), and with a range of facts about how creators derive utility from attention, sourced from 105 AI-conducted open-ended 15-minute interviews we collect from creators in our sample.

Welfare Implications

We close by providing tentative evidence on welfare. Our framework predicts social media viewers may be matched to content they find aversive on certain creator-driven dimensions (lower m), but this should be compensated by greater appeal on other dimensions (higher $q + e$). To test this, we pay the participants providing their viewership histories to first spend a week using an iPhone app we built that detects when they are on social media and pings to ask how they are feeling. We link their responses to their contemporaneous viewing and exploit high-frequency variation in the content appearing in their feeds for identification. Political and enviable lifestyle content drive the bulk of negative emotions participants experience while scrolling. However, participants rarely report regretting having watched such videos and find them entertaining and interesting. In a survey where we show participants creators from their feeds and ask if they would like to block them, participants do not want to block enviable or political creators at a higher rate than other creators.

These results are consistent with our framework’s predictions, but the outcomes we collect are too narrow and short-run to undergird confident welfare conclusions. Our model allows the shift from legacy to social media to *reduce* welfare if particular internalities or externalities load on dimensions of content driven by amateur supply—for example, if viewers underestimate the long-run psychological effects of repeated exposure to unfavorable social comparisons, or if toxic news diets have negative externalities for democracy.

Related Literature

We contribute to two literatures. The first studies the political economy of media. The canonical question in this literature, reaching back to Mullainathan and Shleifer (2005), Gentzkow and Shapiro (2006, 2010), and Besley and Prat (2006), is whether features of equilibrium media content (such as slant) reflect supply or demand. Papers on this topic include Durante and Knight (2012); Martin and Yurukoglu (2017); Qin et al. (2018); Martin and McCrain (2019); Beattie et al. (2021); Szeidl and Szucs (2021); Angelucci et al. (2024); Boxell et al. (2026); Cagé et al. (2025); Noy and Rao (2025).

Because most papers in this literature study legacy media, the supply-side distortions they look for—including editorial pressure from wealthy owners or repressive governments—are examples of *top-down* capture of media by elites (Herman and Chomsky, 1988). We are the first to study a different kind of supply-side distortion: *bottom-up* capture of media by the masses of people willing to spend much of their time posting online, who may be *negatively* selected on factors like institutional enfranchisement and satisfaction with the status quo. The Internet-precipitated transition from media susceptible to top-down pressures to those susceptible to bottom-up pressures offers one lens through which to interpret political trends over the last 20 years in liberal democracies, including the rise of populist parties and fringe ideologies and declining popularity of mainstream parties and institutions.

A growing literature both within and outside of economics examines the relationship between social media and political outcomes (reviewed by Zhuravskaya et al., 2020; Aridor et al., 2024); a similarly large and interdisciplinary literature examines the relationship between social media and mental health and wellbeing (Allcott et al., 2020; Braghieri et al., 2022; Bursztyn et al., 2025; Beknazar-Yuzbashev et al., 2026). These two literatures have mostly proceeded independently and emphasized different and almost exclusively demand-side mechanisms. These include political homophily (Allcott and Gentzkow, 2017;

Levy, 2021; Müller and Schwarz, 2021; Acemoglu et al., 2024b; Braghieri et al., 2024), self-control problems (Allcott et al., 2022; Beknazar-Yuzbashev et al., 2024; Bao, 2026; Solheim and Wang, 2026; Bolte et al., 2026), other consumption frictions (Acemoglu et al., 2024a; Braghieri et al., 2026; Moehring, 2026), or collective traps (Bursztyn et al., 2025, 2026).

Our primary contribution is to provide a new theory of social media centering its distinctive *supply* side. Authors in the literature have pointed out social media’s unique amateur production model (Zhuravskaya et al., 2020; Aridor et al., 2024) and noted its upsides as a potential “liberation technology” empowering repressed populations (Enikolopov et al., 2020; Fergusson and Molina, 2019) and downsides as a source of toxicity or misinformation potentially necessitating content moderation (Guriev et al., 2026; Beknazar-Yuzbashev et al., 2025a). Other authors have studied the effects of attention on users’ production (Srinivasan, 2024; Filippas et al., 2026; Balbuzanov et al., 2025), political selection into production (Bail, 2021; Schulz et al., 2026), and the effects of exposure to user comments (Donati and Song, 2026).

But the literature lacks a unified account of the social media production model’s origins, purpose, and implications. We are the first to provide a comprehensive account of social media’s business model as arising from a *complementarity* between production scale enabled by reliance on user-generated content and technological advantages that facilitate cherry-picking from large content pools. This lets us do three things. First, we clarify and microfound the key difference between social and legacy media. Second, we trace out one implication of this difference: an increase in the influence of creators’ supply-side preferences on equilibrium content consumption. Third, we show how this implication could help tie together the political and mental health effects of social media via a common supply-side mechanism: self-selection into production by creators prone to causing negative reactions in viewers.

Finally, as we discuss later, our comparison of social and legacy media extends the insights of a small literature studying the effects of the Internet on legacy entertainment industries (Brynjolfsson et al., 2003; Waldfogel, 2016, 2017, 2019; Aguiar and Waldfogel, 2018).

Roadmap

Section 2 develops our model of social media and compares social to legacy media. Section 3 compiles observational evidence on social media corroborating our model’s predictions and motivating the importance of nonpecuniary supply-side preferences. Section 4 uses demand- and supply-side experiments to estimate the parameters of our model and decompose the equilibrium prevalence of envy-provoking “lifestyle” content and partisan political content into supply- versus-demand driven components.

2 Conceptual Framework

In this section, we write a simple model of social media platforms and use it to analyze how supply- and demand-side preferences interact to determine equilibrium content consumption on these platforms. We then compare social media platforms to legacy media outlets, analyzing how they differ technologically and what this implies for the types of content they favor.

Section 2.1 presents our model of social media. Section 2.1.1 sets up our assumptions; Section 2.1.2

discusses the empirical facts underlying them; Section 2.1.3 characterizes equilibria in terms of supply and demand; and Section 2.1.4 performs comparative statics. Section 2.2 compares social to legacy media.

2.1 A model of social media platforms

2.1.1 Setup

Agents. The model involves a finite set of viewers $i \in I$, a continuum of potential creators $j \in J$, and a platform intermediating between creators and viewers.

Content types. Pieces of content belong to “types” $k = 1, \dots, K$. As we discuss in Section 2.1.2, we interpret these types as representing cells in a coarse partition of an underlying high-dimensional characteristic space. Across different types, pieces differ systematically from each other on certain dimensions such as topic or style, but pieces *within* a given type also differ from each other on many residual dimensions.

Viewer utility. Viewer i ’s utility from post p of type k is:

$$u_{ip} = m_k + \tilde{m}_{ik} + q_p + e_{ip} \quad (1)$$

Here, m_k is viewers’ average taste for type k ; $\tilde{m}_{ik}$ is individual i ’s personal taste for type k relative to other viewers, and has mean zero across viewers; q_p is piece p ’s vertical quality; and e_{ip} is its idiosyncratic appeal to viewer i . q_p and $\{e_{ip}\}_{i \in I}$ are drawn upon piece p ’s creation from a creator-specific joint distribution and average to zero across potential pieces of each type.

Variation in viewers’ average tastes for content types is captured by m_k . Residual variation in utility from a piece conditional on its type—in $\tilde{m}_{ik}, q_p, e_{ip}$ —captures viewer preference heterogeneity and residual variation, both predictable and unpredictable, in a piece’s characteristics. For example, if viewers generally dislike partisan content, its m_k may be low; but some pieces of partisan content may have high q_p, e_{ip} , representing other characteristics viewers find appealing, such as well-placed jokes or a resonant style.

A viewer’s utility from post p in a sequence is $u_{ip} - \gamma h_{ik(p)}$, where $h_{ik(p)}$ is the share of posts in the sequence of p ’s type and $\gamma \geq 0$ is a parameter capturing diminishing marginal utility from a content type.

Creator utility. If creator j produces a unit of content of type k which receives v views, their payoff is:

$$\pi_{jk} = \mu_k + \tilde{\mu}_{jk} + \beta v^\rho + wv \quad (2)$$

Here, μ_k represents creators’ average taste for producing pieces of type k , and $\tilde{\mu}_{jk}$ is a creator-specific deviation. βv^ρ , where $\beta \geq 0$ and $\rho > 0$, represents rewards for attracting views, and wv , where $w \geq 0$ is a wage-per-view, represents payments from the platform.

Supply-side preferences are captured by μ_k , which allows the fixed cost, expressive satisfaction, or status utility associated with producing a piece of content to vary with its type. Variation in μ_k drives a wedge between viewer and creator preferences if certain content types have high μ_k but low m_k , or vice versa.

βv^ρ represents both nonpecuniary satisfaction from receiving attention and the ability to monetize attention through third-party sponsorships. Higher values of β and ρ help align viewer and creator preferences by encouraging creators to produce content types that viewers like and will therefore receive higher viewership.⁶

⁶We could allow β and ρ to vary by content type, thereby introducing another place where wedges between viewer and creator preferences could emerge. We impose homogeneous β and ρ for two reasons. First, it is simpler and, as we show later, empirically inconsequential: allowing for type-specific β and ρ in our estimation does not affect our results. Second, if content type k had higher

Per-view wages w similarly align viewer and creator preferences. Imposing a common (rather than type-specific) wage is inconsequential for our theoretical results. Empirically, there is little evidence platforms wage-discriminate based on the type of content a creator produces.

Creator decisions. Each creator j draws a separate cost shock associated with creating each content type, and a shock for not creating, from a common type-I extreme value distribution with scale $1/\psi$. They then create up to one unit of content of one type, maximizing their expected utility net of the cost shocks.

Creator heterogeneity. Creators vary in their skill at producing different types of content. We capture this by assuming that within type k , each creator belongs to one of a finite number of internally homogeneous groups indexed by g . The quality q_p and shocks $\{e_{ip}\}_{i \in I}$ of a piece of type k created by a member of group g are drawn from a type- and group-specific joint distribution F_{kg} . A creator's group can differ across types; when we discuss creator j producing type k , g denotes j 's group for type k .

Platform. So far, the model could describe either social or legacy media. Our assumptions now about the platform's *distribution technology* single out social media.

We assume the platform sets a wage-per-view $w \geq 0$, and, given this wage, potential creators decide what to produce, resulting in a pool of produced content. The platform then allocates pieces to viewers.

Our model is intended to describe platforms such as Instagram, TikTok, and X/Twitter, where viewers scroll one by one through a vertical feed of algorithmically curated posts. We therefore assume that viewers have a reservation value of time $\underline{u}_i(T)$ increasing in the number of posts T they have consumed, and, upon entering the platform, consume the posts they are served until their utility from the marginal post falls below $\underline{u}_i(T)$. Formally, the platform observes the characteristics of each produced piece p and serves each viewer i a set of posts P_i . Viewer i consumes the subset of P_i that maximizes her utility accounting for her opportunity cost of time and diminishing marginal utility within content types.

We are implicitly assuming that the platform can discover the characteristics of each produced piece via an ex-post *experimentation* step, and can act on this information via *personalized distribution*. These are the technological capabilities that we argue in Section 2.2 distinguish post-Internet from pre-Internet media and gave rise to the social media business model.

Platform's objective. We assume the platform extracts viewer surplus through price discrimination.⁷ It therefore sets the wage w and assigns posts to viewers to maximize viewer surplus net of wage costs.

For simplicity we assume the platform ignores the wage when assigning posts, meaning the marginal assigned post satisfies $u_{ip} \geq \underline{u}_i(T)$ rather than $u_{ip} \geq \underline{u}_i(T) + w$. This is inconsequential for our results and adds descriptive realism.⁸ We also assume the platform does not assign the viewer posts she will not

β_k because it created more sponsorship opportunities, it would not be obvious whether to interpret β_k as a supply-side preference. On the one hand, it would ultimately derive from appropriation of *viewer* surplus arising from demand for the sponsored product. On the other hand, creators could be biased towards content types where sponsorship availability makes viewer surplus more appropriable, particularly if wages hit a zero lower bound. Imposing homogeneous β and ρ allows us to sidestep these interpretive issues.

⁷This need not occur via a literal price and could, for example, occur via personalized ad loads (Beknazar-Yuzbashev et al., 2025b).

⁸This matches two empirical facts: (i) platforms do not algorithmically penalize posts eligible for per-view payments and (ii) platforms serve viewers an infinite feed of posts rather than cutting them off before they wish to stop scrolling.

consume. The platform therefore assigns viewer i the viewer-surplus-maximizing set P_i^* solving:

$$P_i^* = \arg \max_{P_i} \int_{P_i} (u_{ip} - \gamma h_{ik(p)}) dp - \int_0^{|P_i|} \underline{u}_i(t) dt \quad (3)$$

We will see below that P_i^* is unique and simple to characterize.

Intuitively, the platform serves the viewer her favorite posts of each type until their appeal drops so low, or her opportunity cost of time rises so high, that she switches off. Because she has an increasing opportunity cost of time, pieces of different types compete for her attention: if the platform gets access to a better pool of comedy sketches, the additional sketches inserted into her feed will crowd out pieces of other types.

Timing and equilibrium concept. The game proceeds as follows. (1) The platform sets a wage $w \geq 0$. (2) Each creator draws a cost shock for each content type and decides whether and what to produce. (3) Created pieces draw q_p and $\{e_{ip}\}_{i \in I}$. (4) Viewers i are assigned feeds P_i^* satisfying Equation (3).

An equilibrium consists of a wage w^* , and, for each $w \geq 0$, creator strategies mapping cost shock realizations $\{\kappa_{jk}\}_k$ into actions, such that creators best-respond to each other given each wage, and w^* maximizes expected viewer surplus minus wage payments given creator strategies.

Assumptions. We assume that for each i , $\underline{u}_i(\cdot)$ is continuous and weakly increasing and $\underline{u}_i(0)$ is low enough that i consumes a strictly positive mass of posts; and that the distribution of $q + e_i$ under each F_{kg} is atomless.

Characterization. Our first result characterizes the equilibria of this game. Proofs are in Appendix G.

Proposition 1. An equilibrium exists. In any equilibrium, working backwards:

1. Viewer i sees the unique (up to measure zero) subset P_i^* of produced pieces maximizing her utility.

Specifically, let T_i be the mass of posts consumed by viewer i and let h_{ik} be the share of viewer i 's consumption going to type k . Define viewer i 's "attention price" τ_i as $\tau_i := \underline{u}_i(T_i) - \gamma \sum_{\ell} h_{i\ell}^2$ and her type-specific cutoff t_{ik} as $t_{ik} := \tau_i + 2\gamma h_{ik}$. Then:

$$P_i^* = \{p : u_{ip} \geq t_{ik(p)}\}$$

2. Production results from creators' utility maximization given expectations about viewership.

Specifically, given others' behavior, creator j produces type k with probability σ_{jk} satisfying

$$\log \frac{\sigma_{jk}}{\sigma_{j0}} = \psi \left(\mu_k + \tilde{\mu}_{jk} + \mathbb{E}_{F_{kg}} \left[\beta \left(\sum_i \mathbb{1}\{u_{ip} \geq t_{ik}\} \right)^\rho + w \sum_i \mathbb{1}\{u_{ip} \geq t_{ik}\} \right] \right)$$

where σ_{j0} is the probability that j does not produce and $1/\psi$ is the scale of the cost shocks.

3. The platform's wage w^* trades off the increase in expected viewer surplus from the additional production induced by higher wages against the increase in the wage bill.

Specifically, let $V := \sum_i T_i$ be total views. If $w^* > 0$, it satisfies the first-order condition

$$\int_J \sum_k \frac{d\sigma_{jk}}{dw} \mathbb{E}_{F_{kg}} \left[\sum_i (u_{ip} - t_{ik})^+ \right] dj = V + w^* \frac{dV}{dw}$$

2.1.2 Discussion of assumptions

Before moving on to our substantive results, we discuss the interpretation of, and factual basis for, several of the assumptions we have made.

Content characteristics. We have written the model in terms of coarse content “types” that capture some, but not all, variation in viewer utility from pieces of content. Underlying this is the assumption that pieces of content are drawn from a high-dimensional characteristic space $\mathcal{X}$ over which viewers have preferences $u_i(x)$. Content “types” k refer to a coarse partition of $\mathcal{X}$ into subsets $X_1, \dots, X_K$. For example, given a binary characteristic $x^1 \in \{1, 2\}$, one such partition is $X_1 = \{x \in \mathcal{X} : x^1 = 1\}$, $X_2 = \{x \in \mathcal{X} : x^1 = 2\}$. In this example, the content types $k = 1, 2$ differ along the focal dimension x^1 ; *within* each type, different pieces still vary on all other dimensions, generating residual variation in their appeal to viewers. If $x^1 = 1$ represents partisan political content and $x^1 = 2$ represents moderate content, residual variation in characteristics could include variation in topic, style, phrasing, visuals, and so on, all of which may matter to viewers.

Preference dimensionality. When we refer to “high-dimensional preferences,” we mean there are many residual dimensions viewers care about beyond the focal dimensions captured by our content types. Such preferences increase the residual variance of viewer utility *within* content types, which in our model is captured by the variance of $q_p + e_{ip}$ relative to m_k .

Unpredictability and heterogeneity. We have assumed that content production is an unpredictable process: a creator producing a piece may target a general region of $\mathcal{X}$, but the final characteristics of the piece—and hence its appeal to potential viewers—emerge with some randomness, captured by the stochastic $q_p + e_{ip}$.

Our assumption that pieces of media emerge stochastically from a high-dimensional characteristic space is motivated by two stylized facts from the literature on media economics (Caves, 2003). First, it is extremely difficult to predict ex-ante which pieces of media will resonate with audiences and which will not. This is sometimes called “Goldman’s Law,” after the screenwriter William Goldman’s insistence that “the single most important fact, perhaps, of the entire movie industry [is that] *nobody knows anything*” about which projects will appeal to consumers (Goldman, 1983). Quantitative studies of legacy entertainment industries bear out this assertion (De Vany and Walls, 1999; Waldfogel, 2019). Goldman’s Law also seems to be in force on social media, where creators routinely express profound puzzlement about what draws audiences to their content and why seemingly-identical videos end up differing wildly in performance (Andrews, 2020; Sweet, 2023; Tait, 2023; Cruz, 2024; Ronson, 2026; O’Meara, 2026).

Second, viewer preferences exhibit extreme heterogeneity: for example, different evaluators’ assessments of media correlate only very weakly for movies (Wallisch and Whritner, 2017), images (Vessel and Rubin, 2010), jokes (Rosenbusch et al., 2022), and poems (Belfi et al., 2018).

These facts both arguably stem from the impossibility of tracing the experiences evoked by media to a small set of identifiable dimensions of content (Sibley, 1959). Pieces of media succeed when viewers find them funny, moving, suspenseful, or transporting, reactions produced by collisions between subtle characteristics of the media and consumers’ personal traits.⁹

⁹An alternative interpretation of the empirical regularities underlying Goldman’s Law is that they reflect behavioral cascades created by herding or strategic complementarities in consumption (Salganik et al., 2006; Gilchrist and Sands, 2016). As long as the property of “being good enough to generate such cascades” is distributed in an ex-ante unpredictable way, our results, which

The role of unpredictability and heterogeneity. The key function of unpredictability and preference heterogeneity is to increase the returns to production *scale*. The greater the ex-ante unpredictability of the vertical quality of pieces of media, the higher the returns to producing many pieces and ex-post cherry-picking the best ones. Similarly, the greater the heterogeneity of viewer preferences, the higher the returns to procuring many pieces so as to find each viewer pieces close to their personal ideal point.

Note that preference heterogeneity does not need to be ex-ante unpredictable in order to increase the returns to production scale. Instead of the stochastic e_{ip} , we could represent preference heterogeneity in our model with a deterministic viewer-creator match value e_{ij} and derive similar results. These two modelling choices yield essentially the same results and will be empirically indistinguishable given the variation we later generate, so we adopt the stochastic e_{ip} formulation for simplicity.

Measuring ex-post characteristics. We have assumed that the platform can observe q_p and e_{ip} when assigning pieces to viewers. We assume this is accomplished through a preliminary experimentation step where it serves a new piece p to a small subset of viewers and infers its quality q_p through their average engagement and the distribution of its e_{ip} by combining their engagement with an estimate of the covariance matrix of viewer preferences based on past behavior (“viewers like you liked...”). We do not model this experimentation step, but we trace it out empirically later.

Algorithmic markups. We have assumed that the platform profits equally from consumption of all content types conditional on their effect on immediate viewer surplus. In principle, the platform may wish to apply an algorithmic “markup” δ_k to certain types of content, serving them to viewers whenever $u_{ip} \geq t_{ik(p)} - \delta_k$, if consumption of those types raises its profits beyond immediate viewer surplus. For example, this would be the case if consuming those types has positive spillover effects on future consumption or engagement with advertisements (Becker and Murphy, 1988, 1993), or if creators of that type particularly value viewership and increasing the reach of their content encourages them to produce. We show in Appendix D that such algorithmic markups can be detected observationally using a method following Kleinberg et al. (2018) and Rambachan (2024). The idea is that if a content type enjoys a markup, viewers’ engagement with marginal posts of that type in their feeds should be predictably lower than for marginal posts of other types. We find that advertisements and sponsored videos are algorithmically marked up according to this measure, as expected, but find no evidence of algorithmic markups favoring any of our content types of interest.

Competition between platforms. For simplicity we consider a monopolist platform. Because the mechanisms we are interested in operate through Coasean surplus maximization rather than some kind of rent-shifting that could be undermined by competition, our results are robust to allowing perfect competition.

are fundamentally about the returns to production scale, will go through. Moreover, several facts support the interpretation that Goldman’s Law reflects ex-ante unpredictability of the intrinsic appeal of pieces of media (Waldfogel, 2019). First, the Law holds in settings where herding and strategic complementarities are likely strong, such as movie sales and repost-based platforms like X/Twitter. But it holds with equal strength in settings where they are likely much weaker, such as TikTok, where reposting plays almost no role and viewers’ feeds differ dramatically from their friends’. Indeed, Appendix Figure A3 shows that in our viewership experiment, hiding the “like” counts of TikTok videos does not change the slope of engagement with respect to video popularity. Second, the Law holds for final consumption outcomes but also at earlier steps in the production process. For example, it is very hard to predict whether a script will turn into a movie liked by executives or test audiences (Goldman, 1983). Third, if new pieces of media did not have unpredictable intrinsic appeal, methods such as the TikTok algorithm of discovering underlying appeal would have little commercial value.

2.1.3 Equilibrium consumption in terms of supply and demand

We now move on to our first substantive result: equilibrium consumption of each content type can be described in terms of separate components depending only on demand- versus supply-side preferences.

Proposition 2. At any equilibrium wage w , equilibrium consumption shares are a fixed point that can be characterized as follows. Begin by stipulating a vector of cutoffs $\tau = \{\tau_i\}_i$ for appearing in each viewer's feed and a vector $\mathbf{h} = \{h_{ik}\}_{i,k}$ of shares of each content type in each viewer's feed.

1. **(Demand)** A creator from group g producing type k faces a cumulative distribution D_{kg} of potential views that depends only on demand-side parameters and equilibrium cutoffs and shares:

$$D_{kg}(v, \tau, \mathbf{h}) = \mathbb{P}_{(\{e_i\}_{i \in I}, q) \sim F_{kg}} \left[\sum_{i \in I} \mathbb{1}\{m_k + \tilde{m}_{ik} + q + e_i - 2\gamma h_{ik} \geq \tau_i\} \leq v \right] \quad (4)$$

2. **(Supply)** The supply of pieces of type k from group g can be written as a function $S_{kg}(\mathbf{D})$ of the viewership distribution $\mathbf{D} = \{D_{kg}\}_{k,g}$, where $S_{kg}(\cdot)$ depends only on supply-side parameters:

$$S_{kg}(\mathbf{D}) = \int_{\{j: g_k(j)=g\}} \frac{\exp(\psi(\mu_k + \tilde{\mu}_{jk} + \mathbb{E}_{D_{kg_k(j)}}[\beta V^\rho + wV]))}{1 + \sum_{\ell} \exp(\psi(\mu_\ell + \tilde{\mu}_{j\ell} + \mathbb{E}_{D_{\ell g_\ell(j)}}[\beta V^\rho + wV]))} dj \quad (5)$$

3. **(Consumption shares)** Consumption of type k can be written as a function of supply and demand:

$$C_k = \sum_g S_{kg}(\mathbf{D}) \mathbb{E}_{D_{kg}}(v) \quad (6)$$


Aggregate consumption results from a vector of individual consumption $\{C_{ik} = \sum_g S_{kg}(\mathbf{D}) \mathbb{P}_{F_{kg}}[m_k + \tilde{m}_{ik} + q + e_i - 2\gamma h_{ik} \geq \tau_i]\}_{i,k}$ generating the stipulated vectors of cutoffs and feed shares: $\tau_i = \underline{u}_i(\sum_{\ell} C_{i\ell}) - \gamma \sum_{\ell} h_{i\ell}^2$ and $h_{ik} = C_{ik} / \sum_{\ell} C_{i\ell}$ for all i, k .

Proposition 2 says that three sets of parameters determine equilibrium consumption of content type k . The interaction of these parameters is visualized in Figure 1 Panel (a), which plots the quantity of type- k posts produced on the x -axis and the average viewership of type- k posts on the y -axis. Aggregate consumption is the area of the shaded rectangle: the product of quantity of posts and average viewership per post.


First are the demand-side parameters capturing viewers' taste for type k : the more viewers like type k 's distinctive characteristics, the more views a piece of type k will receive. In Figure 1 Panel (a), these parameters determine the height of the blue demand curve—the average viewership a post of type k receives.

Figure 1: Supply, demand, and equilibrium content consumption

(a) Equilibrium consumption C_k of content type k



(b) Comparative statics: parameters mediating the effect of an increase in average appeal m_k



Notes: This figure plots the “demand” and “supply” curves determining equilibrium consumption of content type k under the simplifying assumptions that there is no diminishing marginal utility ($\gamma = 0$), type k is a small enough share of feeds that it does not affect viewers’ τ_i , and creators make entry decisions according to the mean of expected viewership rather than its full distribution.

Second are the supply-side parameters μ_k affecting creators’ taste for producing type k . The more creators enjoy producing type k , the more pieces they will create. Even if m_k is low, meaning viewers dislike type k ’s distinctive characteristics, pieces of type k can make it into viewers’ feeds through lucky draws of q and e . More production entails more shots at that. In Figure 1, higher values of μ_k shift the orange supply curve outwards by making creators more willing to supply content at any given level of expected viewership.

Third are the supply parameters β, ρ, w governing creators’ utility from attracting viewership. These parameters mediate the feedback from demand to supply: the more that viewers like watching type k , the more viewership a creator will expect for producing type k , and therefore the larger the number of creators who will choose to produce type k in pursuit of wages and attention. In Figure 1, higher values of these parameters flatten the supply curve, amplifying the effects of a given vertical demand shift on consumption.

2.1.4 Comparative statics

Two factors govern the influence of supply- relative to demand-side preferences on equilibrium consumption. It is easiest to see this in a simplified environment, nested by our general model, in which we can derive closed-form expressions for equilibrium objects.

Suppose we have a fixed wage $w > 0$, a unit mass of homogeneous creators, and $n \geq 2$ homogeneous viewers who each consume their T favorite posts. There is no diminishing marginal utility within type ($\gamma = 0$) and $u_{ip} = m_{k(p)} + q_p + e_{ip}$ where $q_p \stackrel{\text{iid}}{\sim} U[-s, s]$ and $e_{ip} \stackrel{\text{iid}}{\sim} U[-\lambda s, \lambda s]$ for $0 < \lambda < 1$. Here, $s > 0$ is a parameter controlling the *dispersion* of the unpredictable component of viewer preferences.

There are two types of content, L and H , where $m_H = m_L + \Delta$ (viewers prefer consuming H content) and $\mu_L > \mu_H > 0$ (creators prefer producing L content). We assume that $\lambda + \frac{\Delta}{s} < T \leq \frac{1}{2}$.

Proposition 3. Under the assumptions above, the expected share of produced pieces that are type H is:

$$\alpha = \frac{\exp\left(\psi\left[\mu_H - \mu_L + (\beta n^\rho + wn)\frac{\Delta}{2s}\right]\right)}{1 + \exp\left(\psi\left[\mu_H - \mu_L + (\beta n^\rho + wn)\frac{\Delta}{2s}\right]\right)} \quad (7)$$

And the expected share of consumed pieces that are type H is $\alpha + \frac{Q}{T} \frac{\Delta}{2s} \alpha (1 - \alpha)$, where Q is the number of pieces produced in equilibrium. In particular we have that:

1. **(Viewer preference dispersion attenuates the effect of demand.)** The production and consumption shares of H are decreasing in viewer preference dispersion s .
2. **(Creator valuation of viewership amplifies the effect of demand.)** The production and consumption shares of H are increasing in β , ρ , and w .
3. **(Preference dispersion deprofessionalizes production.)** Let $R(s) := \frac{wnT}{Q[(1-\alpha)\mu_L + \alpha\mu_H]}$ denote the equilibrium ratio of creator utility from wages to intrinsic utility from μ . Then $R(s)$ is decreasing in s .
4. **(Preference dispersion creates local scarcity of the most appealing content.)** For any $s' > s$ and any $\varepsilon > 0$ small enough that $\varepsilon < s(T - \lambda) - \Delta$, the expected equilibrium mass of produced posts within distance ε of the cutoff required to enter a viewer's feed is smaller under s' than s .

Proposition 3 says that as in Figure 1, increasing the average appeal of content type k has an effect on its consumption that is *attenuated* by the degree of residual viewer preference dispersion and *amplified* by creators' valuation of viewership, for reasons described below.


Preference dispersion. Consider first the residual dispersion of viewer preferences conditional on a piece's type, meaning the variance of the vertical quality q and horizontal preference e shocks.

As above, consider two types of content, L and H . In Figure 2 we plot the distributions of utility that pieces of each offer a viewer. Their average appeals m_L, m_H determine the distributions' means, while the dispersion of the q and e shocks affects the distributions' variance. We fix $m_H - m_L = \Delta$.

Suppose the variance of the q and e shocks is low, as in Panel (a). One interpretation is that viewer i has low-dimensional preferences: there are few residual characteristics of content they care about conditional on a piece being L or H . Then no matter how many L pieces are produced, none stand a shot at appealing to i enough on residual dimensions to overcome her average distaste for L and end up in her feed. Even a large supply-side force favoring L will not pass through to consumption.

By contrast, in Panel (b), suppose viewer i cares about many dimensions of content other than those captured by the L vs. H distinction, so the variance of q and e is high. Then the first-order determinant of type k 's representation in consumption becomes not its ex-ante average appeal to viewer m_k but *how many pieces of type k are produced* and therefore how many q and e draws it generates.

Figure 2: Preference dispersion and pass-through of production to consumption



Notes: This figure shows the effects of increasing viewer preference dispersion—the variance of the q and e shocks—holding fixed $m_H - m_L$ and the expected amount of content falling above the focal viewer’s opportunity cost of time.

Scarcity. A more intuitive way to think about the importance of viewer preference dispersion is through the lens of scarcity. Suppose viewer i , all else equal, prefers moderate to partisan political content. Our hypothesized mechanism is that if partisan content enjoys a supply-side advantage, the platform can cherry-pick pieces of partisan content that viewer i prefers to the marginal available moderate pieces because they better match her preferences on *other* dimensions (such as style, topic, or personality).

This presupposes local *scarcity* of content closely matching viewer i ’s preferences. If the platform’s content library were infinite, for any given partisan post p , it could find a moderate post p' perfectly matching p on all residual dimensions. It would hence serve viewer i partisan content only if i had a taste for partisanship. Given the enormity of production on modern social media platforms (Steel et al., 2026), we might expect their content libraries to be effectively infinite.

Preference dispersion endogenously creates local scarcity of the content most appealing to a viewer. In a high-dispersion world, the content in a viewer’s feed comes from the thin upper tail of q and e draws, where pieces are sparsely distributed and have no close substitutes. The curse of dimensionality offers one way to make this kind of scarcity quantitatively intuitive. If content is characterized by fifty binary dimensions over which viewers have preferences, the probability that a random piece is in a viewer’s top preference bucket is $\frac{1}{2^{50}}$, or about one in a thousand trillion. If a platform has a billion pieces, a viewer’s favorite hundred will on average differ from her ideal on seven out of fifty dimensions, and from each other on twelve.

Deprofessionalization. Greater viewer preference dispersion also “deprofessionalizes” production in the sense that viewership becomes ex-ante unpredictable and it becomes difficult for a creator to consistently earn high wages by producing a content type with large m_k and attracting high and stable viewership. Production instead becomes driven by the intrinsic satisfaction parameters μ_k .

Creator valuation of viewership. Proposition 3 also notes that the parameters β , ρ , and w mediate the indirect feedback effect from demand to supply. If creators face a high wage-per-view or derive substantial nonpecuniary satisfaction from attracting more views, they will be pulled towards producing the content types demanded by viewers, amplifying the importance of demand-side tastes m_k and diminishing the importance of supply-side tastes μ_k . By contrast, if wages are low and creators derive little satisfaction from viewership *on the margin*—either because β is low or because ρ is low, reflecting strongly concave or scope-insensitive nonpecuniary satisfaction from viewership—this feedback effect weakens.

Equilibrium wages on platforms such as Instagram, TikTok, and X/Twitter are very low. This can occur for multiple reasons. It can occur if nonpecuniary motivations suffice to induce substantial production, meaning higher wages would largely transfer money from the platform to inframarginal creators. It can also occur due to a market failure *exacerbated* by inter-platform competition (Jeon and Rey, 2024). Suppose platforms cannot enforce exclusivity contracts on creators. (They in fact do not, likely due to practical difficulties enforcing such contracts.) Then a platform could raise its wages to induce new content creation, but could not appropriate the resulting viewer surplus because this content would also be cross-posted on rival platforms—a practice our survey of creators shows is commonplace. This could explain why YouTube, which has no close competitor for longer-form video content, pays much higher wages, a point to which we return below. We remain agnostic between these explanations for low prevailing wages.

2.1.5 Summary

The model just outlined shapes our subsequent empirical investigation in two ways.

First, it helps define our empirical objectives. We have shown how to write equilibrium content consumption in terms of separate objects depending only on demand- and on supply-side preferences. In Section 4 we estimate these objects and analyze counterfactuals in which we equalize content types’ supply- or demand-side preferences to decompose observed consumption into supply- vs. demand-driven components.

Second, it generates two predictions about the conditions under which supply-side preferences will influence equilibrium consumption. In Sections 3.2 and 4.2, we show that viewer preferences exhibit high dispersion that creates local scarcity of content matching viewer preferences and attenuates the effects of average taste differences on consumption. In Section 4.3, we show that creators’ marginal valuation of viewership is low due to low wages and strongly concave nonpecuniary rewards to attention.

2.2 Social versus legacy media

The previous section described a social media business model premised on harvesting a large number of posts and engaging in ex-post personalized distribution to viewers based on the realization of q and e shocks. We now explain how this business model differs from that of legacy media outlets and why it magnifies the influence of supply-side preferences and encourages short-form content.

Institutional background. Pre-Internet legacy media outlets faced the same unpredictable and heterogeneous consumer preferences as social media platforms (Caves, 2003; Waldfogel, 2017). But they were constrained by two factors (Waldfogel, 2019).

First were technological bottlenecks on distribution. The broadcast spectrum only supported three TV channels; cable supported only up to a few hundred, due to cables’ limited bandwidth. Movie theaters had only a few screens. Bookstores could only stock so many books. This meant legacy outlets could distribute only a limited amount of content, preventing them from fully catering to preference heterogeneity.

Second were high costs of audience testing. The expense and inconvenience of disseminating content prevented outlets from efficiently testing new pieces of content on large samples of consumers to learn their vertical quality before making a decision about wider distribution (Goldman, 1983). For example, test

screenings for movies were expensive and tended to yield small and noisy samples (Lumet, 1995).

Faced with these difficulties, outlets professionally produced small portfolios of content ex-ante optimized, to the extent possible, to appeal to viewers. They then distributed them widely without extensive testing, gambling on opening weekend or the first post-release weeks of a book or album (Waldfoegel, 2019).

The Internet eliminated physical constraints on both personalized distribution and cheap at-scale collection of feedback. Unlike a TV network, a website with a large content library could serve a distinct feed to each visitor at zero marginal cost, and quickly collect feedback in the form of views and shares. This enabled sites with large libraries to rapidly test new pieces of content on large subsets of consumers and use the resulting feedback to inform selective further distribution to other consumers.¹⁰

Platforms such as Netflix and Amazon deployed these capabilities to streamline distribution of movies, TV shows, and books, increasing the productivity of legacy media and leading it to expand in quantity and quality (Waldfoegel, 2019; Aridor et al., 2026). But as with many new technologies (David, 1990), fully harnessing the Internet’s potential for media required a more fundamental redesign of the underlying product. By giving platforms the tools to sift through large libraries of content and locate the best pieces to distribute to each viewer, the Internet raised the returns to content library *scale* and thereby rewarded content that is *cheap to procure at scale*—for example, low-effort content that is satisfying to produce—and *disaggregated enough to facilitate granular cherry-picking and personalization*—for example, short-form content.

Supply-side preferences. A simplified version of our model helps illustrate. Suppose there are two types of content, L and H , with mean appeal to viewers $m_L < m_H$ and per-unit procurement costs $c_L < c_H$. That is, L content is less appealing but is also cheaper to procure. When produced, a piece j of L or H content draws vertical q_j and horizontal $\{e_{ij}\}_i$ appeal shocks from a distribution common to both types. An outlet procures some pieces, distributes one to each viewer, and maximizes viewer surplus net of procurement costs.

Suppose we represent the technological constraints of legacy media as an inability to observe or act on the q, e shocks before distributing content to viewers. In this case, an outlet optimally procures a single unit of H and distributes it to everyone. This represents legacy media: absent the ability to ex-post cherry-pick or personalize, an outlet should produce a small amount of the content ex-ante most appealing to consumers.

When we introduce the Internet—allowing the outlet to experimentally uncover the q and e shocks and act on them via personalized distribution—the outlet is incentivized to procure several pieces of H content rather than just one. This strategy represents Netflix and Amazon: expand production of legacy media types.

But the outlet may be able to do even *better* by spending the same budget procuring a larger number of L pieces, compensating for their lower average appeal with a larger number of draws enabling superior cream-skimming of the best q and e draws. This represents social media.

The shift from content types with high average appeal m to those with low procurement costs c —which ultimately derive from the supply-side preferences determining creators’ reservation wages—is the sense in which social media’s business model magnifies the influence of supply- relative to demand-side preferences.

¹⁰This strategy was pioneered by Facebook and Twitter, neither of which initially used recommendation algorithms. Instead, they relied on users to surface and spread high-quality content by deciding what to share or repost, and to personalize their own feeds by deciding who to follow or connect with. They thereby reproduced the decentralized “testing” and dissemination of products through repeated marketplace interactions and social learning, but at vastly greater speed than physical markets. Fueled in part by data from these networks, advances in machine learning then led to algorithms that greased the wheels of this process.

Welfare. This shift leads viewers to consume content with lower m but higher q, e . In principle, this increases viewer surplus on net. But if particular internalities or externalities load on the dimensions distinguishing L from H , welfare could decrease. For example, if $m_L < m_H$ because L content is more likely to trigger unfavorable social comparisons, and viewers underestimate the long-run psychological effects of repeated exposure to such comparisons, they could end up less happy. Or if L content represents partisan vitriol and H represents factual reporting, this shift could reduce welfare via negative externalities for politics.

Short-form content. Finally, suppose the q and e draws are separately realized across different subunits of longer pieces of content: for example, some of the jokes in a comedian’s stand-up routine are better than others, or some scenes in a movie appeal more to i and others to i' . Then an outlet able to observe, and condition distribution on, q and e benefits from splitting content into the smallest subunits over which these shocks are realized, to enable maximally granular ex-post discrimination. Appendix G.2 provides evidence that such discrimination is extensive on social media and summarizes two other mechanisms: (i) short-form content allows more rapid learning about viewer preferences by linking higher-frequency user interaction data to more disaggregated variation in exposure to content characteristics;¹¹ and (ii) reducing the minimum required length of content lowers barriers to content production.¹²

3 Observational Evidence

The previous section outlined a model in which social media platforms harvest large quantities of content produced for a mix of pecuniary and nonpecuniary reasons and engage in ex-post cherry-picking of the most appealing produced pieces. It predicted that creators’ supply-side preferences can influence viewer consumption because content types creators enjoy producing will be supplied at higher rates, generating more draws of the unpredictable components of appeal q and e and hence facilitating superior cherry-picking.


The unpredictability and dispersion of viewer preferences emerged as a key mediator: the more unpredictable and dispersed viewer preferences are, the harder it is for viewership to consolidate around a few professional creators, the more that supply becomes driven by nonpecuniary motivations, the greater the local scarcity of highly appealing content, and the greater the importance of generating more q and e draws.

In this section, we provide observational evidence corroborating each step of this story. In Section 3.2, we show that the majority of content consumed on social media is produced by unpaid amateurs. We show that this is because appealing content emerges in a sparse and unpredictable way from across a broad base of production; that highly appealing content is locally scarce; and that this scarcity lets unusually successful

¹¹The nascent “microdrama” industry operates according to this logic: produce large volumes of short, cheap entertainment, rely on algorithms to engage in granular experimentation and ex-post matching to viewers, and use variation in performance to learn which specific story beats and set pieces do well with viewers (Kolitz, 2026); see also Bao (2026).

¹²Conventional explanations often claim social media content is short-form because the format hijacks our brain’s susceptibility to being tempted by uncertain, variable, and immediate rewards, like those offered by slot machines. But this fails to explain what makes social media different from TV, where content could also be structured to exploit these vulnerabilities, and plausibly *is*. Indeed, 20th-century critics accused TV channels of fragmenting viewers’ attention spans by serving them loosely connected sequences of short clips: Postman (1985) writes that “what [TV viewers] watch, and like to watch, are moving pictures: millions of them, of short duration and dynamic variety,” and Pynchon (1990) describes a satirical TV-detoxification facility as serving “a retinal diet of scientifically calculated short video clips of what in full dosage would, according to theory, have destroyed [the subject’s] sanity.” One way to rescue conventional theories is to append an additional claim about how the tactile sensation of swiping or scrolling somehow activates or exacerbates these psychological vulnerabilities.

Figure 3: Overview of data collections



videos from content types viewers appear to dislike make it into viewers’ feeds. In Section 3.3, we introduce our content types of interest: lifestyle and partisan political content. We show that the amateur creators dominating feeds are selected on enviability and ideological extremity, which predict production of these content types. And we show that these types account for half of all production on TikTok and about a third of consumption. In Section 3.4, we exploit discontinuities in financial incentives to attract viewership to show that, by creators’ revealed preference, lifestyle content faces a demand *penalty* relative to other content types and partisan political content does not appear to enjoy a premium compared to nonpartisan content.

3.1 Data

We provide an overview of the data underlying results in this section, with further details in Appendix C. Figure 3 summarizes our data collections, including the experiments introduced in Section 4.

Recruitment. We recruit social media *viewers* from the survey platforms Prolific and MobLab, as often done in the literature (e.g., Allcott et al., 2025). We recruit *creators* via Prolific and through an email campaign targeting creators who appear in the viewership histories of our respondents (see below).

Our recruitment procedures involve one innovation relative to preceding literature: across most data collections, we require respondents to begin by logging in with their social media accounts using an authentication flow we developed drawing on the TikTok Login Kit, the X/Twitter `users.read` API, and Google OAuth. The tool collects and authenticates respondents’ social media profiles and checks their numbers of posts and followers before letting them into our surveys. This allows us to verify that they genuinely use these platforms, to screen out respondents whose accounts fall below certain size thresholds for our creator experiments, and to scrape respondents’ accounts to recover the contents of their recent posts.

Appendix Tables A1–A4 compare our samples to nationally representative surveys from Pew, YouGov, and the American National Election Studies that ask about usage of specific social media platforms. Relative to the general youth population, users of our platforms, especially TikTok, skew more female, left-wing, and ideologically extreme. Our samples skew even further in these directions as well as being more educated.

Viewership histories and emotion-sampling. Our data on the contents of users’ social media feeds and their contemporaneous emotions come from a panel of 383 Americans aged 18-30 recruited from Prolific.¹³

¹³Our hypotheses and empirical strategy for estimating the effects of social media content on emotions were pre-registered on the Open Science Framework at <https://osf.io/8ctq7>.

We pay participants to download our iPhone app, which pings them randomly throughout the day—with particularly high frequency when it detects they are using social media—and asks about their emotions. After a week of using it, participants download and send us their viewership histories on Instagram, TikTok, and YouTube. Their final payment depends on submission of these histories and their response rate to the app’s notifications. Appendix Table A5 traces out attrition from our screening survey through to completion. Persistence selects on gender (being female) and politics (being liberal) but not education, income, or mental health. Because we have no randomized treatment, standard concerns about selective attrition do not apply.

Instagram, TikTok, and YouTube comply with data ownership regulations in the EU and California by giving users the option to download their data, which is where our viewership histories come from. These histories contain second-by-second records of each post the user has viewed and each search they have made, reaching back one week, six months, and all time, respectively, for the three platforms, totalling 24 million posts across all respondents and platforms. We link these histories by timestamp to the emotion responses collected via our app, allowing us to observe exactly which post a respondent was viewing when they were pinged by our app. Appendix Table A6 describes the characteristics of these histories.

X/Twitter’s “request my data” feature does not yield comparable data, so we follow up with participants and pay them to briefly use a Chrome extension that records the top 300 posts on their X/Twitter feed; we obtain these data, which are not linked to emotions, for 57 respondents.

Random sampling of produced posts. In addition to viewership histories for a panel of users, TikTok offers a Research API that allows us to randomly sample from the pool of posts produced on the platform in the US in any specified date range. This allows us to compare a representative sample of the content *produced* on TikTok to the TikToks appearing in viewers’ feeds. For this reason, as well as other data-quality reasons described in Appendix C, we focus most of our results on TikTok, showing results for other platforms in Appendix B (to the extent our data allow).

Post metadata and classification. We take the post URLs appearing in our viewership histories and API-based random sampling and scrape characteristics of those posts, including thumbnail images, captions, transcripts, and metadata (including view counts and flags for commercial sponsorships or ad spending). We give each post’s imagery and text metadata to GPT-5-mini and ask it to assign the post a content category based on a classification schema described in Appendix C. The schema was developed on the basis of open-ended data gathered from social media users. The LLM classifications are validated against 100 human hand-coders recruited on Prolific, as shown in Appendix Table A9; human codings correlate strongly with the LLM codings, and humans agree slightly more with the LLM than with each other. Appendix C applies a sparse auto-encoder (Movva et al., 2025; Carlson, 2026) to our experimental and observational data stripped of classifications and shows that it autonomously discovers categories similar to our classifications.

Creator survey. We conduct a descriptive survey of creators who produce content on the platforms we study. We contact creators in two ways. First, we scrape the profiles of creators who appear in our respondents’ viewership histories and extract email addresses where available. We email creators and offer Amazon gift cards in exchange for completing our survey, collecting 301 responses at a 1.7% response rate. Second, we recruit 607 creators on Prolific, using our social media log-in flow to automatically screen out small accounts.

Our viewership history data, which allow us to see the characteristics of creators appearing in viewers’ feeds, allow us to re-weight our survey responses to create a sample representative (in terms of observables such as follower counts and types of content posted) of the creators appearing in viewers’ feeds. In our emails, we randomly offer either \$5 or \$10 gift cards to test whether responses are selected on unobservables (Dutz et al., 2026). Appendix Table A12 shows the \$10 offers significantly increase response rates and the resulting test shows no evidence of selection on unobservables.

User survey. To measure selection into content creation, we survey 3,490 users of our platforms recruited through Prolific. We ask respondents to log in with their accounts, allowing us to automatically grab their numbers of followers and posts and to scrape the contents of recent posts from their profiles.¹⁴ We then elicit survey outcomes and correlate those outcomes with features of their posting behavior, as described later.

Qualitative interviews. We follow up with a sample of 105 creators who participated in our surveys or experiments and ask them to complete a 15-minute AI-conducted qualitative interview (Chopra and Haaland, 2023; Geiecke and Jaravel, 2024). The interviewer asks a series of open-ended questions about what types of content the interviewee posts and why they post. It then asks probing followups seeking concrete details about the satisfaction interviewees derive from creating posts or attracting likes or views and about why they might be reluctant to post certain types of content. Quotes from these interviews are collected in Appendix F and we refer to them throughout the paper to enrich our quantitative results.

3.2 Amateur production, experimentation, and scarcity

In this section, we show that most of the content social media viewers consume is produced for nonpecuniary reasons and link this fact to our model’s comparative-static predictions about viewer preference dispersion.

We show (i) platforms benefit from a broad base of production mostly driven by nonpecuniary motivations; (ii) appealing pieces of content emerge sparsely and unpredictably from across this base of production; and (iii) as a consequence, most of the content in viewers’ feeds comes from small, non-professional creators.

Our model predicts that successful videos will emerge in this kind of sparse and unpredictable way when the variance of q and e shocks is high. It predicts that, in these circumstances, highly appealing videos will be locally scarce, and videos from low- m content types will consequently make it into viewers’ feeds thanks to high q and e draws. Consistent with these predictions, we show that (iv) platforms quickly run out of content that will highly engage a given viewer and (v) as this scarcity kicks in, viewers see larger numbers of unusually successful videos from content types they appear to dislike on average.

Because our richest linked data on production and consumption come from our TikTok viewership histories plus TikTok’s Research API, results in this section focus on TikTok. At the end of the section, we briefly discuss the analogous results for other platforms, which are presented in Appendix B.

Distribution of production. The grey bars in Figure 4 Panel (a) use a representative sample of all videos produced on TikTok between November 2025 and February 2026. Panel (a) plots the distribution of the follower counts of these videos’ creators, a measure of those creators’ size and prominence. The panel shows that nearly 90% of production on TikTok comes from very small creators with fewer than 10,000 followers.

¹⁴Since Instagram does not expose any API or toolkit allowing us to build a login flow, we ask users directly for their handles.

If TikTok viewers mainly saw recent posts from accounts they follow, these small creators would account for a large share of production but only a small share of consumption—their share of followership is 7.5%. However, all social media platforms have mechanisms for a creator’s post to propagate past their followers. This could be a simple “share” button, as with pre-algorithmic versions of Facebook and Twitter. In TikTok’s case, this mechanism is the TikTok algorithm, which fills viewers’ feeds mainly with videos from accounts they do *not* follow but that the algorithm predicts they will engage with (Zannettou et al., 2024).

The algorithm’s recommendations are informed by an experimentation stage. When a video is posted, TikTok “featurizes” it, observing its creator and hashtags, generating embeddings from its transcript and imagery, and using these to predict its potential appeal (Boeker and Urman, 2022). However, this step still leaves TikTok with incomplete information. It therefore serves the video to a pool of “test” viewers and uses their engagement to resolve its residual uncertainty about the video’s quality and distribution of appeal.

Identifying test videos. How deep into the creator size distribution does TikTok reach to draw its test videos? Does it assume that users will never be interested in a video from a small account they do not follow, or does it believe that broadly-appealing gems regularly emerge from such accounts? Answering this requires us to identify instances where TikTok serves a video to a viewer as a test. Panel (b) of Figure 4 does this.

We use our participants’ TikTok histories and plot, on the x -axis, the temporal gap between a video’s appearance in a viewer’s feed and its creation. Low x -values represent videos that appear in a viewer’s feed shortly after being posted. The blue coefficients show that for videos appearing in viewers’ feeds between 6 and 72 hours after their creation, average engagement is roughly flat—as we would expect from an engagement-maximizing algorithm that equates viewers’ marginal propensity to engage with different types of videos. But viewers predictably dislike videos that appear in their feeds less than 6 hours after being posted. This suggests these videos are being served not because the algorithm predicts they will engage the viewer but for another purpose. Indeed, the orange coefficients show that a viewer’s engagement with such a video is highly predictive of the total viewership later accumulated by the video. This predictiveness declines rapidly for videos appearing in a viewer’s feed a longer interval after being posted, consistent with TikTok having already learned these videos’ quality and gradually serving them to increasingly marginal viewers.


“Appearing in a feed less than 6 hours after being posted” therefore seems to be a reasonable proxy for being a test video. These videos make up 7% of the typical feed. Half of them are from creators the viewer has never seen before; in these cases they tend to come from small creators, with a median of 1,697 followers.

Distribution of testing and consumption. The orange bars in Panel (a) of Figure 4 show that although TikTok does statistically discriminate in favor of large creators when deciding which videos to test, it draws test videos from across the distribution of creators. Even though creators in our smallest bin make up less than a tenth of followership, they produce half of all tested videos.


This passes through to final consumption. The blue bars plot the distribution of follower counts in participants’ full viewership histories. The blue distribution looks similar to the orange, suggesting that, conditional on being tested, videos from large and small accounts are similarly likely to be broadly disseminated. Overall, the bars show that *nearly half* of the videos in viewers’ feeds are from very small accounts.

Figure 4: Production, experimentation, and consumption on TikTok

(a) Distribution of produced, tested, and consumed videos across account sizes



(b) Identifying tested videos



Notes: Panel (a) plots the percentage of videos falling into each creator-size bucket in three different samples: (1) a random sample of the videos produced on TikTok, from its Research API; (3) all videos appearing in the histories of participants who sent us their personal TikTok data; and (2) the subset of history videos that appeared in a participant's feed less than 5 hours after being posted. Panel (b) uses videos in participants' feeds and plots the number of hours elapsing between the video being posted and appearing in the focal viewer's feed against (i) the average number of seconds watched by the participant (controlling for viewer by video duration fixed effects and time since the beginning of the session) and (ii) the slope between the focal viewer's engagement and the video's total log viewership scraped several months afterwards.

Amateurs vs. professionals. This suggests that, contrary to what we might expect, it is not true that viewers mostly see content produced by professional “influencers” who cater to viewers in pursuit of financial rewards.

Indeed, Panel (a) of Figure 5 uses scraped platform data to estimate the percentage of videos in viewers' feeds that make money. Green bars represent income from Creator Rewards, TikTok's program paying creators based on how many views they receive. Eligibility for the program requires an account with at least 10,000 followers, at least 100,000 views in the preceding 30 days, and a video that is at least a minute long. Only 9% of the videos in people's feeds satisfy these criteria and are hence eligible to earn Creator Rewards.

Purple bars represent income from third-party sponsors who pay a creator in exchange for featuring their product. These sponsorships must legally be tagged in a way our scraper can detect. The figure shows that only 8% of the videos seen by viewers are tagged as sponsored. (Our sample includes ads, meaning videos someone has paid to boost into viewers' feeds; 7 out of 8 of these percentage points are ads.)

Appendix Figure A29 instead asks: how many of the videos people see come from creators who *ever* earn money, as suggested by their other videos? The figure also incorporates data from our survey of creators, allowing us to get at income sources we cannot directly observe in the administrative platform data, such as unreported sponsorships, videos produced for sponsors' accounts, or donations from viewers. Such income sources are commonly reported by creators in our survey, but are almost exclusively reported by creators who *also* earn money in administratively observable ways, suggesting that our extensive-margin estimates based on administrative tags of the share of creators earning money are close to the truth. The figure shows that, depending on the conservativeness of our assumptions, up to around 45% of the videos viewers see on TikTok come from accounts earning money on the platform. This is higher than the 15% of videos viewers see that directly earn money but still a minority. And even creators who make money likely create content for a mix of pecuniary and nonpecuniary reasons, as confirmed by our creator survey.

Creator survey. More than half of the videos TikTok viewers see are therefore produced by non-professionals producing for nonpecuniary reasons. What are these reasons? Panel (b) of Figure 5 plots responses from our creator survey, weighted to be representative of the creators in viewers’ feeds.

We first ask how respondents choose what to create: do they produce the content that will get the most views? That will attract the highest financial payouts? Or that they personally find satisfying to produce?

Responses to such a question may be influenced by social desirability. We therefore randomize the wordings of the response options between two sets intended to flip their perceived social desirability, and confirm in a supplementary survey detailed in Appendix Figure A31 that our manipulations succeed. The bars in the figure plot average responses, while the + and – symbols show average responses in each condition.

The figure shows that respondents report the most important consideration determining their production choices is what they personally find satisfying to produce.

We also ask creators to indicate which kinds of nonpecuniary satisfaction they derive from content production. Panel (b) shows that social connection, creative fulfillment, and expressive satisfaction are among the most popular answers, while a third of creators report gaining social status.¹⁵ Appendix Tables A13 and A14 collect responses from our qualitative interviews to open-ended questions about why creators enjoy producing content or having people watch it. Creators report a range of rewards, including social connection, pride in creation, admiration from viewers, self-expression, and the ability to relive experiences. We return to the nature of these rewards in Section 4.3 when we discuss why the utility creators derive from posting appears to be so strongly concave in the number of views they receive.

Viewer preference dispersion. The videos that receive positive engagement signals when tested, and end up widely disseminated to viewers, are unpredictable outliers relative to their creator’s typical output.

Panel (c) of Figure 5 shows that the median video in a viewer’s feed receives 37 times more views than its creator’s median viewership on their other videos. For videos posted by creators with fewer than 10,000 followers, this ratio is 212. Panel (d) shows that nearly half of the videos in viewers’ feeds are from creators they see once and never again (in the six-month window covered by their TikTok history). These creators have a median of 6,716 followers; half of viewers’ feeds are hence filled by one-hit wonders from small creators whose content does not typically perform well, as evidenced by their low follower count.

The ex-ante unpredictable origins of the videos in viewers’ feeds suggest that viewer preferences are hard to target in advance. This explains the failure of professional creators to dominate viewership and suggests viewer preferences cannot be reduced to a few optimizable dimensions such as video style, topic, or slant.


Fat-tailed and ex-ante unpredictable viewership is in turn suggestive of high-variance q and e shocks in the language of our model.¹⁶ But fat-tailed viewership could also reflect small differences in appeal being amplified into enormous differences in dissemination à la Rosen (1981). This would manifest as a large density of videos whose qualities place them close to viewers’ opportunity cost of time $\underline{u}$. In this case, videos in viewers’ feeds could be outliers in terms of viewership but have many close substitutes in terms of quality.

¹⁵Appendix Figure A10 shows that creators who report “expressive satisfaction” or “persuasion” are more likely to post political content on their accounts; those who report “social status” are more likely to post lifestyle content.


¹⁶The connection to preference dimensionality is that if viewers’ preferences are high-dimensional and viewers have a non-lexicographic utility function in which they gain utility from each additional dimension of a piece matching their tastes (Tversky, 1969; Fishburn, 1976), this will result in what looks like higher-variance q and e shocks.

Figure 5: Amateur production, viewer preference dispersion, and local content scarcity on TikTok


(a) Amateur production: What fraction of posts in viewers' feeds earn money for their creators? (Admin data)




(b) Amateur production: Which considerations influence the content creators choose to make? (Creator survey)




(c) Dispersion: The videos viewers see are in the far upper tail of their creators' views




(d) Dispersion: Nearly half the videos viewers see are from creators they see once and never again



(e) Scarcity: Stale videos receive less engagement, but...



(f) Scarcity: ... nearly half a typical viewer's feed is stale



Notes: Panel (a) plots videos from the histories of participants who sent us their TikTok data. An observation is a video occurring in a participant's feed. Blue bars are the percent of observations in each creator size bucket. The green and purple bars plot how many observations in that bucket are (i) eligible for Creator Rewards (creator has at least 10k followers and 100k views in past 30 days and video is at least 60 seconds long and not sponsored) or (ii) tagged as sponsored. Bars are in percentage points, meaning, for example, that 1ppt out of 43ppt in the first bucket is sponsored. Panel (b) plots responses from our creator survey, with responses weighted by follower count and posting type to be representative of the creators viewers see in their feeds. The first sub-panel asks respondents to allocate 100 points between three possible considerations determining which types of content they choose to post. Answer options are 50/50 randomized between two sets intended to flip rankings of perceived social desirability as detailed in Appendix Figure A31. The second asks them to "tick all that apply" out of a list of nonpecuniary rewards they derive from producing content. "Social connection" was added to the list partway through fielding, so its percentage is among the respondents who were shown it. Panels (c) and (d) use TikTok histories and the same units of observation as Panel (a). Panel (c) plots the ratio of a video's view count to the median view count of other videos posted by the same creator over the same period spanned by our histories. Panel (d) plots the number of times this observation's creator appears in the viewer's history. Panel (e) uses TikTok histories and plots the share of a video a viewer watches against the number of days since the video was posted (controlling for viewer by video duration fixed effects and time since the beginning of the session). Panel (f) plots the share of videos in a viewer's feed that were posted at least 7 days ago as a function of time since the viewer started this scrolling session, controlling for viewer fixed effects.

Scarcity. To distinguish these possibilities, we test directly whether the videos in viewers' feeds lack close substitutes in terms of utility. Panel (e) of Figure 5 shows that viewers engage predictably and substantially

less with “stale” videos—videos posted more than a week before they are served to a viewer. If, each day, TikTok receives an abundant supply of fresh content closely matching a viewer’s preferences, it should never serve her stale videos. But Panel (f) shows that from the moment the typical viewer begins scrolling, more than a third of the videos in her feed are stale; by the time she has scrolled for 30 minutes, this has risen to more than half. This indicates substantial scarcity of highly appealing content, and the magnitude of this relationship—stale videos receive at least a 30% engagement penalty yet occupy up to half of the typical feed—suggests that the marginal appeal of available content falls very rapidly.¹⁷

Scarcity and supply. Our model predicted that such local scarcity means that even if viewer i dislikes content type k , pieces of type k with sufficiently high q or e draws can make it into her feed because of a paucity of comparably appealing videos from types the viewer likes more.

Appendix Figure A5 shows evidence consistent with this prediction. The figure draws on our classification of videos into seventeen genres described in Appendix C, such as “sports,” “pets,” or “food and cooking.” We proxy for whether viewer i likes or dislikes a genre—her $\tilde{m}_{ik}$ —by whether it appears in her feed more or less commonly than the average among other viewers. Panels (a) and (b) show that videos from a genre viewer i dislikes that nevertheless appear in her feed score higher on the dispersion measures from Figure 5—consistent with higher q or e realizations being needed to carry those videos into her feed. Panels (c) and (d) of Appendix Figure A5 show that as scarcity bites more keenly—as a viewer scrolls further into her feed, or as we compare lighter to heavier TikTok users—TikTok relies more and more on these unusually appealing outliers from disliked content types.¹⁸

Summary. The descriptive patterns documented in this section validate the predictions of our model: the unpredictability and inscrutability of viewer preferences prevent professional creators from consistently optimizing for those preferences and instead lead platforms to engage in ex-post filtering from a broad base of amateur production. Moreover, preference dispersion creates local content scarcity that allows exceptional outliers from unappealing content types to sneak into viewers’ feeds.

Other platforms. Appendix B reproduces these analyses, to the extent possible given our data, for X/Twitter, Instagram, and YouTube. X/Twitter looks exactly like TikTok: viewership is highly volatile, feeds are dominated by one-off hits from small accounts, and (to an even greater extent) very few creators make money. Instagram is similar to TikTok and X/Twitter in the sense that most consumed content is produced by amateur accounts that do not make money. However, Instagram consumption is more skewed towards large accounts and viewership is less volatile. Unlike TikTok, Instagram has historically engaged in less algorithmic experimentation on small creators, which has been blamed for Instagram Reels’ underperformance relative to TikTok (Instagram, 2024; TikTok, 2020). And unlike X/Twitter, Instagram is less reliant on sharing and

¹⁷ The final alternative interpretation is that videos in viewers’ feeds do lack close substitutes in terms of utility, but this is endogenous to the fact that they have already gone viral—in other words, there is a utility premium to engaging with popular content. As previously mentioned, this interpretation is compatible with our story if “potential virality” is a property distributed in the same sparse and unpredictable way as intrinsic appeal. Our results in Section 4 do not use observational view dispersion and therefore do not require assuming that it represents dispersion in true appeal. We also do not find evidence of a utility premium for engaging with popular content in our viewership experiment, where we randomize across participants whether videos’ “like” counts are visible; Appendix Figure A3 shows that making like counts visible does not increase engagement with highly-liked videos.

¹⁸ The formal prediction underlying this test is that if we pick a viewer i and rank produced pieces according to u_{ip} , as we move down the ranking, the prevalence of content types with lower $\tilde{m}_{ik}$ should increase relative to other viewers as long as the $q + e$ shocks have an increasing hazard.

reposting, the other mechanism for surfacing hidden gems from low-follower accounts.

YouTube looks very different: it is much more professionalized, with most consumed content making money for its creators. This is both because YouTube pays creators for viewership much more generously and because YouTube is more dominated by large creators with more stable viewership.¹⁹ In this sense, YouTube is more similar to a TV channel that professionally subcontracts content production than to a social media platform. Consistent with reliance on amateur content helping drive the pathologies of social media platforms, Appendix Figure A12 shows that only 10% of 18–30-year-olds we surveyed wish that YouTube did not exist, compared to more than 40% for each of Instagram, TikTok, and X/Twitter.

3.3 Selection into content production and the composition of social media content

The previous section established the presence of the basic patterns—amateur-dominated production, unpredictable viewer preferences, local scarcity of appealing content—required for supply-side preferences to potentially influence social media consumption. It did not zoom in on any specific type of content.

Production and consumption of content types. We begin by describing our content types of interest. We take samples of content *produced* and *consumed* on TikTok from the Research API and our viewership histories, respectively.²⁰ We scrape the posts’ text and imagery and use LLMs to classify them according to a schema of content types described in Appendix C. We aggregate the detailed classifications into four broad categories depicted in Appendix Figure A18: **politics** (posts about political events or issues); **lifestyle** (posts visually emphasizing features of the creator’s life, such as their appearance or routine); **stories/advice** (posts where a creator talks to the camera, relating a story, gossip, or advice); and **entertainment** (a broad category spanning comedy sketches, hobby demonstrations, edited montages, and so on).

Lifestyle. Panel (a) of Figure 6 plots the prevalence of these content types on TikTok. Observe first that lifestyle content makes up just under 30% of videos consumed on TikTok and nearly 50% of the videos produced; conversely, entertainment accounts for nearly 40% of consumption but less than 30% of production. This is suggestive of a mismatch between viewer and creator preferences on these axes, and the first object we will focus on empirically is relative consumption of lifestyle content compared to entertainment. Appendix Figure A17 shows that the lifestyle-to-entertainment consumption ratio is 0.74 on TikTok and 0.81 on Instagram, but only 0.21 on the more professionalized YouTube.

Politics. Political content makes up just under 5% of viewers’ feeds. This is slightly lower than the 8% of TV consumption going to political news according to Nielsen ratings and our own calculations (Appendix H). This is perhaps unsurprising given that social media viewers are much younger than TV viewers. Appendix Figure A17 shows that politics is more common on Instagram (6%) and X/Twitter (10%) than on TikTok.

Panel (b) of Figure 6 plots the partisanship of political content: an indicator for whether a video expresses a clear partisan perspective. The intensive-margin partisanship of political content is the second object we will focus on. About 53% of consumed and 60% of produced political content is partisan. This compares to only 22% of political TV news minutes that are partisan according to the same measure (Appendix H).

¹⁹ YouTube’s equivalent of Creator Rewards pays out to channels above 1,000 subscribers (as opposed to TikTok’s 10,000 follower requirement) and pays out for all videos (as opposed to TikTok’s ≥ 1 minute restriction).

²⁰ An alternative way to measure consumption patterns is to weight the random sample of produced videos by their view counts. Appendix Figure A1 shows this produces nearly identical results.

Political content posted by users vs. news organizations. Posts by news organizations make up 27% of the political videos in viewers’ TikTok feeds. Only 14% of news organizations’ political videos are partisan according to our measure, compared to 67% of user-generated political content in feeds.

This is not mechanical: our data contain partisan news outlets, such as Fox News, which has a 69% partisanship rate for its TikTok videos. Instead, Appendix Figure A2 Panel (a) shows it is because TikTok viewers prefer nonpartisan outlets. The top outlet for American political news among both Democrats and Republicans in our sample is the Daily Mail, a British tabloid; the next slots are occupied by CBS, NBC, and CNN. Panel (b) compares the partisanship of the clips news organizations post on their TikTok accounts to random political news clips selected from their broadcasts. The clips news organizations post on TikTok are *less* partisan than their typical broadcasting, and their partisan posts on TikTok get fewer views.

Appendix Figure A6 shows that early in a TikTok session, a higher proportion of the political videos viewers see are from news organizations; as viewers scroll and the scarcity from Figure 5 Panel (f) kicks in, TikTok increasingly relies on user-generated political content. Overall, this pattern of results suggests the partisanship of consumed political content may be partly driven by the supply side of amateur production.


Selection into content creation. Section 3.2 showed that most of what viewers see on social media is produced by small, non-professional creators. Figure 6 Panels (c)-(e) show that these creators are heavily selected, relative to the broader user population, on *enviability* and *ideological extremity*.

The figure draws on our survey of 3,490 users of Instagram, TikTok, X/Twitter, and YouTube. Respondents log in with their accounts, allowing us to scrape their account information and recent posts, which we link to two survey-elicited characteristics. “Enviability” averages responses to questions about how attractive the respondent is and how desirable other people would find their lifestyle. “Extremity” is the respondent’s distance from the midpoint of a seven-point scale from “very liberal” to “very conservative.”

Panel (c) shows that, conditional on using Instagram, users who post more heavily are more enviable and ideologically extreme. On TikTok, heavier posters are more enviable but not more extreme. On X/Twitter, a text- rather than image-based platform, heavy posters are substantially more extreme but *less* enviable. On YouTube—a platform that is, as previously described, highly professionalized—creators are not selected on either characteristic. Panel (d) shows that across platforms, enviable users are more likely to produce *lifestyle* content specifically, while ideologically extreme users are more likely to produce *political* content.

The selection of enviable and extreme users into posting lifestyle and political content admits a supply-side explanation according to which enviable or partisan users derive more status or satisfaction from posting lifestyle or political content holding demand fixed. It also admits a demand-side interpretation according to which all creators pursue viewership, and enviable or extreme users expect their lifestyle or political content will attract more views. Distinguishing between these possibilities is the central objective of our creator experiment in Section 4.3, but our user survey provides a preliminary piece of evidence. For each of our video types, we ask 50% of respondents if they would be willing to *post* such a video on their account for \$3. We ask the other 50% if for \$3 they would be willing to privately *create* such a video and send it to us to be shown to a few other participants, but not post it. Willingness to perform the first task should be affected by expectations about the viewership such a post would receive, but willingness to perform the second task

Figure 6: Selection into content production, and content production patterns, on social media



Notes: Panels (c)-(e) use data from our survey of users of our four platforms of interest. Enviability is measured as the average response to two questions about self-rated attractiveness and lifestyle desirability. Ideological extremity is the distance from the midpoint of a 7-point self-reported ideology scale. Posting volume and the contents of recent posts are measured by requiring users to log in with their accounts, collecting their usernames, and scraping their accounts. Panels (c) and (d) plot the predictiveness of enviability and ideological extremity for total posting volume on a platform (conditional on being a user) and indicators for having recently posted specific content types, controlling for the number of posts we scrape. Appendix Figure A9 reproduces Panel (d) using the share of scraped posts of each type. Panel (e) plots responses to questions offering respondents \$3 to either post (for 50% of respondents) or privately create (other 50%) different types of content. It correlates willingness to post/create lifestyle content with enviability and political content with extremity, controlling for average willingness to produce other types. Panel (a) uses a random sample of videos created on TikTok from the Research API (production) and TikTok viewer histories (consumption) and plots the percent of videos falling into each of our four aggregated content categories, described in Appendix C. Appendix Figure A1 reproduces Panel (a), measuring consumption by weighting produced videos by their view counts instead of by using participants' histories. Panel (b) plots the percent of videos classified as being about politics that are classified as partisan, separately for created videos, history videos, and videos in histories that are versus are not posted by news organizations.

should only be influenced by the cost or enjoyment of producing such a video. Panel (e) of Figure 6 shows that enviability and ideological extremity predict willingness to perform these two tasks with equal strengths.

Summary. We have shown three facts suggesting that supply-side preferences might drive production of lifestyle and partisan political content. First, both content types are produced at a higher rate than they are consumed. Second, user-generated political content is much more partisan than the political content viewers prefer from news organizations. Third, content creators are selected on personal enviability and ideological extremity in a way that predicts production of these content types and appears to reflect lower costs or higher satisfaction of supplying them, rather than higher expectations about demand.

3.4 Discontinuities in incentives to attract viewership

We close this section by providing one more piece of observational evidence that production of lifestyle and partisan political content is not driven by attempted catering to viewer preferences.

This evidence is presented in Figure 7, which shows systematic bunching in the production of different types of content at different video lengths. It exploits two facts.

First, TikTok’s interface for filming videos offers default duration options of 15 seconds and 1 minute, as well as a normal “record indefinitely” button. After selection of one of the default options, the camera will start filming and stop automatically after 15 or 59 seconds unless manually stopped earlier. If creators using these default options tend to be *unsophisticated* producers whose content is not planned out, we should observe bunching of unsophisticated creators at and below 15 and 59 seconds.

Second, TikTok’s Creator Rewards program—which pays eligible accounts with over 10,000 followers for views—pays out only for videos at least 60 seconds long.²¹ So creators who are trying to maximize viewership in order to earn Creator Rewards payments should bunch at 60 seconds and above.

Summing up, if we sort videos by length and study their underlying creators, we should see a discontinuous increase in creator sophistication as we cross from 15 to 16 seconds and a discontinuous increase in sophistication *and* financially-motivated attempts to cater to demand as we cross from 59 to 60 seconds.

Figure 7 Panels (a) and (b) show that crossing these thresholds brings large, sharp decreases in the prevalence of lifestyle content and corresponding increases in the prevalence of entertainment. Panel (c) shows that the prevalence of partisan compared to non-partisan political content is smooth across these thresholds, suggesting partisanship does not enjoy a demand premium. Neither lifestyle nor partisan content therefore appear to be produced strategically to maximize engagement.

Two further results corroborate identification. First, only accounts with at least 10,000 followers are eligible for Creator Rewards. Figure 7 shows that the sharp discontinuity at 60 seconds appears only for such accounts. Second, Panel (d) shows that both bunching of video durations and of content types just above 60 seconds appear rapidly in the months following rollout of Creator Rewards in May of 2023.

3.5 Summary of observational evidence

This section has documented a series of observational facts about content production and consumption on social media, each supporting a different step in our narrative.


First, conditions are ripe for supply-side preferences to influence consumption: unpredictable viewer preferences prevent consolidation of viewership around professional creators, create local scarcity of appealing content, and allow exceptional videos from disliked content types to make it into viewers’ feeds.

Second, users who opt into content creation are selected on enviability and ideological extremity, which predict production of lifestyle and partisan political content in a way that appears driven by supply-side preferences rather than expectations about demand.

Third, viewer preferences do not seem to favor lifestyle or partisan political content. Incentives to cater to viewers reduce production of lifestyle content and do not increase the partisanship of political content.

²¹This is presumably why the 60-second default recording option automatically stops at 59, rather than 60, seconds.

Figure 7: Lifestyle videos are posted by unsophisticated accounts with weak incentives to attract viewership; the opposite is true for entertainment



Notes: This figure plots data from TikTok videos. To maximize precision we pool both videos randomly sampled from the Research API and videos in our viewership histories and use, for lifestyle and entertainment, predicted classifications from a classifier trained on our LLM codings. Appendix Figure A11 shows results separately by data source and by LLM versus predicted codings—all show the same patterns. The x-axis for the first three panels is the duration of the video in seconds. The y-axis shows the share of videos of this duration classified as lifestyle (Panel (a)) or entertainment (Panel (b)), or the share of political videos classified as partisan. Solid lines plot videos produced by accounts with at least 10,000 followers and shaded lines videos from accounts with fewer than 10,000 followers. Panel (d) uses data on videos produced during 2023 scraped from the Research API. It plots, over time, the share of videos between 60 and 69 seconds long (green line, left axis) and the entertainment-minus-lifestyle share gap when comparing videos 60-69 seconds long to videos 50-59 seconds long (black line, right axis). We use a 7-day rolling average on either side of the rollout date. Shaded regions show 95% confidence intervals.

And among political videos posted by news organizations, viewers appear to prefer nonpartisan content.

4 Experimentally Disentangling Supply and Demand on Social Media

Section 3 described observational evidence showing that conditions are ripe for supply-side creator preferences to influence consumption and directionally suggesting that the prevalence of two specific types of content—lifestyle and partisan political content—is driven by supply-side preferences.

However, our observational data do not contain the variation necessary to estimate the parameters of our model and thereby quantitatively decompose the equilibrium prevalence of lifestyle and partisan political content into supply- versus demand-driven components. For example, the analysis of bunching in Section 3.4

suggested directionally that viewers prefer entertainment to lifestyle content, but converting these directional results into parameter estimates would require strong quantitative assumptions about how costly it is for creators to manipulate the length of their videos.

In this section, we conduct two experiments, one on the demand side and one on the supply side, to generate the variation necessary to estimate our model. Section 4.1 discusses the identification challenges these experiments are designed to overcome. Section 4.2 describes the demand experiment. Section 4.3 describes the supply experiment. Section 4.4 combines the experimental results to estimate our model and perform counterfactual decompositions.

4.1 Identification

Proposition 2 established that equilibrium consumption of a given content type can be written as a function of two objects: the distribution of expected viewership of potential pieces of this type, which depends only on demand-side preference parameters; and the distribution of creators’ willingness to produce pieces of this type as a function of their expected viewership, which depends only on supply-side preferences.

Our goal in this section is to separately estimate the demand- and supply-side parameters underlying each of these objects. Doing so will enable us to conduct Oaxaca decompositions where we break equilibrium gaps in consumption between content types into components driven by differences in their demand- versus supply-side parameters. We now discuss the identification challenges we face in estimating these parameters.

Canonical supply-demand identification. It helps to begin by recapitulating the canonical supply-demand identification problem. In this problem, the econometrician observes an equilibrium price and quantity in each of several markets. Each market is cleared by a product-specific price that transfers utility from consumers to producers. As Panel (a) of Figure 8 illustrates, a given equilibrium is consistent with many possible demand and supply curves. Identifying the true curves requires supply and demand shifters.

The supply shifter traces out the demand curve. The demand curve itself derives from the distribution of utility from the product across all potential consumers. In principle, it can be identified by approaching each consumer and directly measuring their utility from the product. This will essentially be our approach.


The demand shifter traces out the elasticity of supply to the reward producers receive for supplying, which is the price they can sell produced units for. Again, this will essentially be our approach.

Our setting. In our setting, a “product”—a collection of units that are ex-ante interchangeable from the perspective of consumers—is a content type k crossed with a creator group g . Recall that groups capture variation in creators’ ability to produce type k , so group g may produce type k at a different average quality than g' . From now on, we will parametrize this by assuming that the quality q_p of piece p produced by group g is $\bar{q}_{kg} + \xi_p$, where $\bar{q}_{kg}$ is a scalar and ξ_p is drawn from some group-specific mean-zero distribution.


Identifying supply. Unlike the canonical setting, our setting lacks a product-specific price that transfers utility directly from consumers to producers and therefore operates identically on both sides of a market. Nevertheless, identification of supply-side parameters can be described in an analogous way. The product-specific “price” that creators face is the (probability distribution of) viewership that creators of group g expect for their units of type k . The higher the viewership creators expect, the more likely they will be to

Figure 8: Identifying supply and demand


(a) Canonical identification problem



(b) Our setting: identifying supply



(c) Our setting: barriers to identifying demand



produce, and the strength of this response will depend on their marginal valuation of viewership. Tracing out their elasticity of supply requires manipulating their expected viewership as in Panel (b) of Figure 8. Intuitively, this lets us determine whether certain content types are supplied at high rates because (i) creators enjoy producing them or (ii) creators value views and expect these types to receive high viewership.

Identifying demand. Identification on the demand side can also be described analogously to the canonical problem. For each type k and group g , we can trace out demand for the “product” kg by taking pieces of type k produced by group g , showing them to each potential viewer, and measuring the viewer’s utility.

Where the analogy breaks down is in articulating the identification *problem* this overcomes. There is no product-specific price on the demand side: the “price” viewers pay for watching a video is their marginal opportunity cost of time, which is the *same* across all the “products” we consider. Moreover, in any given market, we directly observe the “demand curve” in the Figure 8 Panel (b) sense, meaning the viewership distribution of pieces in this market. The problem is that demand for a given content *type* aggregates across multiple markets corresponding to different creator groups, and this is difficult to capture empirically.

More intuitively, the problem we face is a *selection* problem operating at two levels displayed in Figure 8 Panel (c). Recall that a viewer’s utility from a piece p of type k is $m_k + \tilde{m}_{ik} + q_p + e_{ip}$, or, using the parametrization of q introduced just above, $m_k + \tilde{m}_{ik} + \bar{q}_{kg} + \xi_p + e_{ip}$.

The first problem arises if we look at the videos *in viewers’ feeds*. Since the platform fills viewer i ’s feed with their favorites among produced pieces, pieces from lower- m content types will be more positively selected on $\bar{q}_{kg} + \xi_p + e_{ip}$ conditional on appearing in viewer i ’s feed. This means cross-type variation in a viewer’s engagement *within their own feed* is not informative.

Showing viewer i random videos sampled from the pool of all produced content ameliorates this first selection problem. But the videos that are *produced* are still selected in a way that depends on supply-side parameters. If creators’ taste μ_k for producing content type k is high, then even low- $\bar{q}_{kg}$ groups will opt into producing type k because it will still be worthwhile even at lower viewership. So we also need to sample posts that are *not produced* in equilibrium.

4.2 Demand-side experiment

Parametrization. Recall from Section 2 that viewer i ’s utility from piece p of type k is given by $u_{ip} = m_k + \tilde{m}_{ik} + q_p + e_{ip}$. We are interested in five content types: lifestyle, stories/advice, entertainment, partisan

politics, and nonpartisan politics. We parametrize the components of u_{ip} as:

$$u_{ip} = m_k + \tilde{m}_{ik} + \underbrace{\bar{q}_{kg} + \xi_p}_{q_p} + \underbrace{e_{G(i)p}^G + e_{P(i)p}^P + e_{ip}^o}_{e_{ip}} \quad (8)$$

e^G and e^P allow for commonality in the preferences of same-gender $G(i)$ and same-political-leaning $P(i)$ viewers. We assume that $\tilde{m}$, e^G , and e^P are normally distributed with mean zero and type-specific variances, and e^o is drawn from a type I extreme value distribution shifted to have mean zero with a type-specific scale. The stochastic component of quality ξ_p follows a lognormal distribution shifted to have mean zero and has a type-specific variance and common shape.²² We assume that across types, the shares of utility variance attributable to each of $\tilde{m}$, q , e^G , e^P , and e^o are the same. All stochastic components are independent.

Within each type, we assume there are ten creator groups g that can be empirically approximated as follows. In the creator experiment described in Section 4.3, we ask creators whether they would be willing to produce a post of content type k in the next 24 hours for \$0.25. We assume a creator’s group for type k can be approximated by their follower count (in five bins) crossed with whether they answer “yes” to this question. This is an approximation because a creator’s willingness to produce type k on a given day is derived from the creator’s primitive tastes and abilities plus today’s realization of their k -specific cost shock. These shocks should wash out on average, though this will slightly attenuate our estimates of inter-group differences.

In our creator experiment, we randomly pay some creators \$10 to post content types they do not normally produce. We use these videos to estimate the latent quality of content from groups unlikely to produce in equilibrium. Because we have only 106 experimentally-induced posts, we cannot allow for arbitrary heterogeneity in the quality distributions of different groups. We therefore assume each unwilling-to-produce group faces a constant penalty δ_k to their quality relative to the willing-to-produce group in the same follower bin, and pool this parameter for lifestyle, entertainment, and politics.²³ This amounts to a moderate restriction on the shape of inter-group quality heterogeneity.²⁴ In our alternative demand estimates sourced from creators’ expectations in the creator experiment, we do not impose this restriction, and Figure 11 shows that the two produce almost identical results.

Viewership experiment. The design and reduced-form analysis of this experiment were pre-registered on the AEA RCT Registry at AEARCTR-0019565 and are described in more detail in Appendix I.

Building a TikTok clone. We build a phone-browser-based clone of TikTok. We download real TikTok videos, host them on a server, and pipe them into users’ feeds. Appendix Figure A36 shows a screenshot of the clone. It reproduces all of TikTok’s main features: viewers swipe at their leisure through an infinite

²²Appendix Figure A7 shows the four content types have very similar viewership distribution shapes.

²³Because of variation in creators’ willingness to post content types, our creator experiment generates 23 stories posts that creators were willing to produce absent payment versus 19 unwilling; but for lifestyle, entertainment, and politics, these are 16/4, 21/3, and 6/14, respectively, limiting the power of willing/unwilling comparisons.

²⁴Consider a simple deterministic version of the model where creators’ tastes for content type k generate a threshold $\underline{q}_k$ such that only groups g with $\bar{q}_{kg} > \underline{q}_k$ are willing to produce. (“Groups” here mean the actual primitive groups we approximate empirically.) The restriction we impose is twofold. First, within each type k the distribution of $\bar{q}_{kg}$ is such that for varying values of q near $\underline{q}_k$, $\mathbb{E}(\bar{q}_{kg} | \bar{q}_{kg} > q) - \mathbb{E}(\bar{q}_{kg} | \bar{q}_{kg} \leq q)$ is constant. For example, this will be satisfied if the local distribution of $\bar{q}_{kg}$ across creators is uniform. Second, the variance of $\bar{q}_{kg}$ across creators is the same for each k . Importantly, we still allow the supply-side parameters determining $\underline{q}_k$ to vary across types. Under these restrictions, $\underline{q}_k$ affects the share of creators falling above the threshold but not the average quality gap between above- and below-threshold creators.

vertical feed, and they can “like” a video, tap to open the top comments or expand the caption, or copy the video’s URL to share it with a friend. We record every aspect of users’ behavior, including how long they watch a video, whether they “like” it, whether they open its comments or expand its caption, and so on.

Sampling videos. We randomly sample 260,000 videos posted in the days preceding the experiment using TikTok’s Research API, classify them, and select a stratified subset covering our four content types, yielding 1,900 politics videos and 5,000 videos each for lifestyle, entertainment, and stories/advice. We also scrape 400 videos recently posted by the 100 news organizations appearing most commonly in the feeds of participants who donated their TikTok viewership histories. Finally, we add 106 videos, including at least 20 of each type, whose creation we induced by paying a participant in our creator experiment \$10 (described next section).

Each viewer receives a feed of videos randomly drawn from this pool, stratified by content type. The shares of lifestyle and political content in a viewer’s feed are drawn independently from $\{0, 0.1, 0.2, 0.3\}$ with probabilities $\{0.2, 0.2, 0.4, 0.2\}$. Residual videos are split equally between entertainment and stories, with 5% reserved for non-political videos posted by news organizations. Political videos are stratified to be 80% user-generated and 20% posted by news organizations.

Recruiting viewers. We recruit Americans aged 18-30 from Prolific, verifying their status as TikTok users by requiring them to log in with their account to enter the experiment. After login, 624 participants start the experiment and 586 complete it. Appendix Figure A37 shows that attrition is not predicted by the content early in a person’s feed and that including incomplete sessions does not affect our results.

Our sample includes 94 participants who had earlier donated their viewership histories; for these participants, we include in their experimental feeds several dozen videos recently posted by creators who appeared regularly in their personal history, and use their engagement with these videos for an overidentification test.

Experimental flow. Once a participant enters the experiment, we explain its procedures. We tell them they will scroll through videos for 20 minutes, and should behave as they ordinarily would on TikTok, swiping past videos they dislike and watching ones they like. We incentivize genuine engagement by telling them they may be invited back to a future scrolling session, and we will use their engagement behavior in the first session to design a personalized feed for the session just like the TikTok algorithm would operate. To ensure respondents actively watch and scroll, we randomly flash several attention checks throughout the 20-minute window that respondents must tap within a few seconds of their appearance in order to receive their payment.

After 20 minutes of scrolling, we present the respondent with 5 pairs of videos they saw and ask them to indicate which they liked more. Comparing their choices to their prior engagement behavior lets us measure behavioral noise. Finally, participants answer a few demographic questions.

Identification of parameters. Identifying utility parameters from viewers’ engagement behavior requires a model linking engagement to utility and accounting for the fact that decisions to swipe past a video are endogenous to the viewer’s outside option of watching the next video. We assume that when viewer i swipes into video p , her flow utility of watching it starts at $u_{ip} - \gamma h_{ik}$, where h_{ik} is the share of already watched videos that are of the same type. We assume her utility declines linearly at rate a as she continues watching. Opportunities to consider swiping to the next video arrive as a Poisson process with rate r_i . When an opportunity arrives after t seconds, she compares $u_{ip} - \gamma h_{ik} - at + v_{ip}$, where v_{ip} is a behavioral noise term,

to her value of starting the next video V_i minus a swiping cost c_i , and decides whether to swipe.

A more parsimonious engagement model could be written but would struggle to fit observed engagement behavior, which exhibits several patterns: (i) viewers vary systematically in how rapidly they drop out of watching a video; (ii) about 40% of videos are swiped past within the first two seconds; (iii) conditional on watching two seconds, viewers exit gradually rather than always finishing the full video.

Figure 9 Panel (a) summarizes how each parameter in Equation (8) is identified from experimental engagement behavior. Intuitively, viewers' average taste for a content type is identified from their average watch time when encountering a video of this type. Variances of vertical quality and horizontal fit are identified through the covariance matrix across viewers of taste for videos within each type. And diminishing returns within type are identified from our extensive-margin randomization of the share of a viewer's feed drawn from each type. The engagement model parameters are identified from the shapes of watch time distributions: the reconsideration arrival rate r_i is identified from the earliest times a participant is observed swiping past videos, the swiping cost c_i is identified from a participant's average watch duration, and a is identified from the slope with which participants gradually exit videos.

Reduced-form results. Figure 9 Panel (c) shows results for one of our pre-registered engagement measures: the probability that a viewer watches at least 10 seconds of a video. Our other pre-registered measures, which produce consistent patterns, are plotted in Appendix Figure A38.

Three patterns are evident. First, lifestyle content suffers a substantial demand penalty: it is the least engaging type. Partisan political content, particularly content concordant with a viewer's ideology, is more engaging than non-partisan content, but this difference is quantitatively modest.²⁵

Second, differences in engagement rates across types are small, with a substantial minority of videos of each type engaging viewers. For example, 21% of entertainment videos, the most engaging type, are watched for at least 10 seconds, while the figure is 15% for the least engaging type, lifestyle. This is consistent with high-dimensional viewer preferences: if our coarse content categorizations capture variation along only a few of the dimensions viewers care about, all types will have some videos that appeal to viewers on enough dimensions to appear in their feeds.²⁶ This amplifies the importance of the supply side: if lifestyle videos are $\approx 25\%$ less likely to highly engage a viewer than entertainment videos, a 33% supply advantage would suffice to make lifestyle videos appear in viewers' feeds more commonly than entertainment.

Modest variation in engagement rates could result from genuine preference dispersion or from noisy behavior. We separate the two with the end-of-experiment module where we ask participants to retrospectively indicate which video they preferred out of five pairs they saw. Respondents affirm their earlier revealed preference in 84% of cases. Assuming the 16% disagreement rate reflects noise in engagement behavior

²⁵Our discontinuity-based evidence suggested no demand premium for partisan content. This discrepancy is partly explained by the fact that our experimental participants are more politically partisan than the general population of TikTok users. Appendix Figure A40 shows that re-weighting accordingly shrinks the experimental partisanship premium, but only very slightly. This suggests our experiment if anything *overestimates* the demand advantage of partisan content.

²⁶Since viewers watch 19% of experimental videos for at least ten seconds, we are capturing attenuation in the probability that a video enters their top 20%, rather than their top 0.0001%, which is what is relevant for feed assignment. But this means our results should understate the degree of relevant attenuation as long as the taste shock distribution has a decreasing hazard rate. Appendix Table A21 compares distributions fitted to the upper tail of video viewership. Distributions with a decreasing hazard (lognormal, Pareto, and Weibull) fit well, while those with a constant hazard (exponential) do not.

(rather than gaps between revealed and stated preference), we conduct an empirical Bayes procedure to disattenuate engagement differences across types. Appendix Figure A39 shows this barely changes results.

Third, note that if content production primarily caters to demand, videos produced in equilibrium should engage viewers substantially more than videos not produced in equilibrium. By contrast, if there is important variation in supply-side preferences, some videos may not be created in equilibrium despite facing high latent demand because they would be difficult or unsatisfying to create. The green bars in Figure 9 Panel (c) show that when we pay a participant in our experiment \$10 to produce a video of a type they were *unwilling* to produce for \$0.25, these videos receive just as much engagement on average as videos participants *were* willing to create for \$0.25, and just as much as actual scraped videos. This suggests variation in production is driven by supply- rather than demand-side preferences, a conclusion our creator experiment affirms.²⁷

Parameter estimates. Panel (b) of Figure 9 shows estimates of our key utility parameters. We normalize the mean utility of entertainment to zero and the variance of the stochastic component of utility across all pieces to 1. We estimate the model in two steps. First, we fit the viewer-specific engagement parameters r_i, c_i and average utilities by content type to observed viewership behavior by maximum likelihood. Second, we split the residual variance of utility into its components by maximizing the joint likelihood of the engagement behaviors of different participants seeing the same videos.

Results line up qualitatively with Panel (c), though stories/advice looks worse because the parameter estimates exploit richer variation in engagement than whether videos are watched for at least 10 seconds.

Two conclusions are again evident. First, entertainment is the highest-utility type, lifestyle is the lowest, and partisan political content enjoys a modest advantage over nonpartisan content. Second, residual preference dispersion is large enough to substantially attenuate the effects of variation in average appeal. For example, the normalized utility gap between entertainment and lifestyle is 0.56, but a single standard deviation of vertical quality q or half a standard deviation of idiosyncratic appeal e each almost suffice to overcome this mean disadvantage. Our results imply that the posts appearing in a viewer’s feed are about 6.7 standard deviations of appeal above the average entertainment post, and lifestyle posts are about half as likely as entertainment posts to clear this threshold. This means that, all else equal, lifestyle posts would only need to be produced twice as much as entertainment to appear in feeds at higher rates.

Simulating consumption. Using our parameter estimates to simulate consumption requires one more step: estimating the utility threshold above which videos appear in viewers’ feeds. We calibrate this to the probability that a video we randomly sample from TikTok’s Research API appears in a given participant’s viewership history (about 1 in 44,000).

This allows us to take the observed distribution of production across content types (and account sizes within content types) from Figures 4 and 6 and use our demand estimates to simulate the resulting distribution of consumption. Figure 9 Panel (d) shows results. Despite using only experimental variation, our demand estimates closely reproduce actual consumption patterns across content types. This is our first and most important overidentification test.

²⁷ Marginal created and uncreated videos within a type should experience equal viewership, but the blue bars in the figure capture average inframarginal engagement with created videos, and the first green bar captures average engagement with uncreated videos down to \$10 below the creation threshold.

Figure 9: Viewership experiment: identification, viewer utility parameters, engagement, and simulated consumption

(a) Identification of viewer utility parameters

Parameter	Variation that identifies it
Mean appeal m_k	Average engagement with videos of this type.
Variance of $\tilde{m}_{ik}$	Cross-viewer variation in average engagement with this type.
Group-specific part $\bar{q}_{kg}$ of q_p	Average viewer engagement with videos of this type created by the relevant group, using videos induced by our creator experiment from creators who do and do not regularly produce this type.
Variance and shape of stochastic part ξ_p of q_p	Cross-video engagement variation that is correlated across viewers.
Variance of e_{ip}	Shared components: within-video within-demographic-group correlation in engagement. Residual component: remainder after partialling out other components of utility and noise.
Diminishing marginal utility γ	How engagement with a type declines with its randomized share of the viewer's feed.


Engagement model (experiment only)

Swiping costs c_i	Variation in viewers' average viewing times.
Within-video boredom a	Slope with which viewers gradually exit videos conditional on staying past the likely arrival of a reconsideration opportunity.
Reconsideration rate r	The earliest times that different viewers swipe past videos.
Variance of noise ν_{ip}	How often viewers re-pick the video they engaged with more in the re-pick module.


(b) Parameter estimates

Parameter	Estimate	SE
<i>Mean appeal m_k, relative to entertainment</i>		
Lifestyle	-0.56	(0.06)
Stories	-0.26	(0.15)
Co-partisan politics	0.14	(0.09)
Counter-partisan politics	-0.12	(0.08)
Nonpartisan politics	-0.17	(0.09)
<i>Standard deviations of shocks</i>		
Quality q_p	0.42	(0.02)
Idiosyncratic appeal e_{ip}	0.91	(0.01)

(c) Probability of watching ≥ 10 seconds by content type



(d) Production, consumption, and simulated consumption of content



Notes: Panel (a) explains how each viewer utility parameter is identified from our viewership experiment. Panel (b) displays estimates of the parameters of the viewer utility function from our experiment, with bootstrapped standard errors. Panel (c) plots the average probability that a viewer watches a video of each type for at least 10 seconds, calculated from regressions controlling for viewer, sampling-source, and video-duration-decile fixed effects. Panel (d) plots, for each of our content types, the type's observed share of *production* (measured by randomly sampling from produced videos using TikTok's Research API), its observed share of *consumption* (measured from the viewership histories of our participants), and its *simulated* share of consumption, obtained by taking the observed pool of produced content and applying our model-derived demand estimates to it. In Panel (d), we restrict to user-generated political content, excluding news organizations.

Other fit and overidentification tests. Appendix I describes and shows results of our other fit and overidentification tests. Our model closely fits experimental watch-time distributions and variation in watch times across content types. Moreover, its predictions pass a series of overidentification tests in our observational data. The most salient concern about our experimentally-derived estimates is that they will not extrapolate to the thin upper tail of utility where the videos assigned to feeds come from. We therefore focus on whether our model quantitatively reproduces relationships between engagement and feed assignment in the observational data. This includes the predictiveness of experimental engagement for actual viewership of our experiment’s videos on TikTok; participants’ experimental engagement with videos from creators who appear regularly in their actual TikTok histories; and the relationship between videos’ age and their prevalence in viewer feeds.

Demand estimates from creators’ expectations. An alternative approach to estimating the latent distribution of demand for content is to ask each creator j about their subjective probability distribution over the views they would receive if they created a piece of content of type k . We approximate this in our creator experiment, described below, by asking creators incentivized questions about their average expected viewership if they created content of different types. These viewership expectations do not let us cleanly separate demand into components stemming from the different primitives m_k, q, e^G, e^P, e^O , and so on, but they are sufficient statistics for the decompositions and counterfactuals we conduct in Section 4.4. Appendix Figure A8 repeats the overidentification test from Figure 9 Panel (d) using these expectations and shows that these expectations-sourced demand estimates reproduce observed consumption from observed production about as well as demand estimates sourced from our viewership experiment.

4.3 Supply-side experiment

Objective. Recall from Proposition 3 in Section 2 that a key factor mediating the influence of supply-side preferences on equilibrium consumption is creators’ marginal valuation of viewership. If creators primarily care about attracting attention, production will chase demand: certain content types will be produced at a higher rate because creators expect those types to receive higher viewership. By contrast, if creators care little about viewership on the margin, supply-side tastes will be a more important driver of production.

Our objective in this section is to assess whether creators’ production decisions are driven by expectations about demand or by supply-side considerations like the difficulty of, or satisfaction derived from, creating different content types. We conduct an experiment that elicits creators’ willingness to accept to post different types of content and decomposes it into a component driven by supply-side preferences and a component driven by expectations about, and valuation of, demand.

Recruitment. We recruit 979 Prolific users aged 18-30 who have posted at least once before on TikTok, as verified by logging in with their account. The median participant has 160 followers and 16 historical posts, while the average participant has 1,755 followers and 88 posts. Appendix J shows that our lifestyle results are robust to restricting to the 19% of participants with at least 1,000 followers, while our politics results become very noisy in the smaller sample because political posting is rare. We also show that our results are qualitatively consistent with descriptive results from our creator survey (Figure 5 Panel (b)), which has a sample representative of the distribution of creators that viewers see on TikTok. As pre-registered, we

re-weight responses to match the observed distribution of TikTok production.

Experimental design. The experiment, weighting scheme, and reduced-form analyses were pre-registered on the Open Science Framework at <https://osf.io/24tcz>. See Appendix J for further details. The survey can be clicked through at <https://demo.mitsocialmediastudy.com/>.

The flow of the experiment is depicted in Appendix Figure A47. Participants complete a survey about content creation on TikTok. We begin by explaining that it is incentive-compatible to answer all questions truthfully, link participants to the full implementation procedure, and tell them they are guaranteed to receive a paid posting offer at the end of the survey. We then define the four content types we will ask about, providing definitions, examples, and links to TikToks in each category, randomly sampled from our observational data.

The first section elicits participants' baseline willingness to accept to post different content types and their expectations about viewership. We start with a binary question: for each content type, would you be willing to make a post of this type on your account in exchange for \$0.25? We then elicit incentivized expectations about how many views the post would receive. To enable supplementary analyses, we also elicit participants' willingness for \$0.25 to privately *create*, but not post, a video of this type and send it to us to be shown to other participants. And we ask two questions probing the reasons behind participants' answers.

Next, we elicit respondents' scalar willingness to accept to post each type. For respondents *unwilling* to post for \$0.25, this is the minimum payment they would need to be offered to agree to make a post of this type. For respondents *willing* to post who indicate they were planning to make a post of this type anyway, we elicit the minimum payment they would require to *abstain* from making posts of this type on their account for the next week. (Since we have their account details, we can enforce this.) We divide their willingness to accept by their average number of posts of this type in weeks where they make at least one, to calculate a per-post willingness to pay to continue posting.²⁸ Our offer generation procedure ensures responding truthfully at each step in this sequence is incentive-compatible; essentially, the procedure is a BDM mechanism and we use a participant's responses to construct a willingness to accept that can either be positive or negative.

The second section measures valuation of viewership. For each content type, we tell participants that if we offer them a payment high enough to induce a post of this type, we may offer them an intervention intended to boost their viewership for that post. We randomly offer one of two interventions: a payment from us to boost their post algorithmically (an option TikTok offers), or tips from a professional content creator. Insofar as the latter also boosts their viewership on future posts, we should overestimate participants' marginal valuation of viewership. We elicit participants' willingness to pay for this intervention and their incentivized expectation about its effects on viewership.²⁹ Their willingness to pay for the intervention divided by their expectation about its effects captures their marginal valuation of viewership. To measure their full viewership valuation function βv^ρ , we estimate ρ by asking participants to indicate indifference points between various increments of views starting at different viewership levels. For example, we ask:

²⁸This may underestimate their true valuation if they would simply postpone these posts until after the week ended. Appendix Table A22 shows that results are robust to restricting to creators who report that they would not make these posts at all rather than postponing, and to inflating valuations in these arms to account for potential underestimation.

²⁹For 50% of respondents, this is an unincentivized question about its expected effects on their own viewership. (This cannot be incentivized because we would not observe both potential outcomes.) For the other 50%, this is an incentivized question about its expected effects on the viewership of other participants, where aggregating across multiple participants lets us construct this estimate. Appendix Table A23 shows results are robust to restricting to either group.

“starting at 1,000 views, how many additional views would a creator need to get so that their satisfaction from this increase in viewership matches the satisfaction of going from 0 to 100 views?” This last object is the only one that is not incentivized.

At the end of the survey, each creator is randomly offered a payment to produce a certain type of content. With 1% probability this offer is drawn based on the BDM mechanism that makes truthful responding incentive-compatible. With 99% probability it is a \$10 offer to produce one of our four content types, drawn at random. This is because the goal of these offers is to induce creation of videos for use in our viewership experiment. We began by setting equal draw probabilities for each of the four content types and, as data collection progressed, adjusted these probabilities to try and balance created videos across the four types. 33% of respondents indicate an intention to accept the offer and one-third of them end up submitting a video. We end up with 20 lifestyle, 24 entertainment, 42 stories, and 20 politics videos.

Parametrization and identification. Recall that creator j 's payoff from creating a piece of type k which ends up receiving v views is $\pi_{jk} = \mu_k + \tilde{\mu}_{jk} + \beta v^\rho + wv - \kappa_{jk}$. We assume that κ_{jk} is drawn from a type I extreme value distribution with the same scale for each content type. To map the experiment into our static model, we assume that a period in the model corresponds to a 24-hour period, so the willingness-to-accept answers covering the next 24 hours represent one-period utilities. We also assume that when creating political content, self-identified ideologically moderate creators produce nonpartisan content and creators with ideological leanings produce partisan content.

As pre-registered, we re-weight creators so that the distribution of production implied by their responses matches the observed distribution of production on TikTok. Since our objective is to *decompose* observed production into components driven by different rewards rather than to reproduce observed production, this re-weighting is innocuous. It implicitly imposes a conditional independence assumption: that conditional on follower count and observed posting, creators in our experiment are representative in terms of supply parameters of actual creators on TikTok.

Identification of parameters is simple and summarized in Figure 10 Panel (a). A creator's overall willingness to accept to create a post captures their overall utility from that post π_{jk} . Their viewership-valuation parameters β and ρ are estimated from their marginal valuation of viewership and indifference responses; combining these with their expected viewership and Creator Rewards eligibility yields their reward from viewership $\beta v^\rho + wv$. Subtracting this out of their overall willingness to accept yields $\pi_{jk} - \beta v^\rho - wv = \mu_k + \tilde{\mu}_{jk} - \kappa_{jk}$. Calibrating to the creator's observed historical posting behavior yields the scale of κ_{jk} .

The experimental elicitation procedure is complex. And we estimate creator preferences via a residualization: we ask how much of observed variation in willingness to post remains after subtracting out the product of expected viewership with implied valuations of viewership. Classical measurement error in the latter will not bias our results because it enters additively. But a natural worry is that some form of non-classical response noise on the right-hand side may bias our results. Our approach is to corroborate our parameter estimates with simple, transparent supplementary quantitative and qualitative results.

Results. Parameter estimates are displayed in Figure 10 Panel (b). We estimate them in two independent steps. First, we estimate an average ρ across creators by fitting creators' viewership indifference responses

by least squares, and β using creators' willingness to pay for, and expectations about the effects of, the viewership interventions. Second, we estimate each creator's total expected utility from posting each content type, plus the scale of the cost shocks κ , by fitting creators' willingness to accept answers from the survey and their historical posting frequencies by maximum likelihood. Subtracting results from the two steps, and using information about wages imputed from creators' follower counts, separates $\beta v^\rho + wv$ from $\mu + \tilde{\mu}$.

There is substantial variation in willingness to post across content types: the median gap in willingness to post between types a creator is willing to post for \$0.25 versus types they are not is \$23.33. But the median gap in implied utility from viewership is only \$0.64, suggesting that almost all of the variation in willingness to post is driven by variation in supply-side parameters rather than expectations about viewership.


Two facts drive this. First, the gap in average expected viewership between content types respondents are versus are not willing to post is small: 813 versus 666 views, or 22%. Appendix Figure A48 shows that when paid to post different content types, creators actually get slightly *higher* (albeit imprecisely estimated) viewership for types they were initially unwilling to post.

Second, participants' implied marginal valuation of viewership—median of \$0.23 per 100 views—is small. This is not because creators do not care about viewership. Panel (b) shows that for the median post creators are willing to make for \$0.25, their implied payoff from viewership is \$9.47 while their overall utility is \$5.00, indicating a \$4.47 cost of producing and posting the video. Rather, it is because returns to viewership are strongly concave, with an estimated $\rho \approx 0.3$, meaning viewership is not that valuable *on the margin*.

Such concavity is consistent with scope-insensitive satisfaction from attention: the first few dozen views or likes are exciting, but after that, a creator's ability to psychologically differentiate between the attention from 200 views and the attention from 500 views is limited. Our qualitative interviews (Appendix Table A15), which ask creators open-ended questions about why they enjoy posting, suggest four additional reasons for concavity. First, many creators report deriving most of their satisfaction from small-scale, two-way interactions with small communities of commenters and viewers who regularly return to their posts. Second and similarly, early views tend to come from friends or consistent followers whose attention the creator values, while later views come from strangers who matter less. This suggests many creators post in pursuit of small-scale bilateral social connection, and the price they pay for accessing a platform that facilitates those connections is that the platform occasionally distributes their posts more widely to entertain other viewers. Third, many creators only pay attention to a post for a short while after it is posted and therefore only attend to the first few views and likes. Fourth, several creators report deriving most of their satisfaction from simply knowing that at least some people like what they posted.

Concave rewards from attention are consistent with a large body of work in psychology on diminishing returns to social rewards (Eisenberger, 1970; Ren et al., 2022). They are also consistent with prior work on social media. Srinivasan (2024) studies TikTok creators whose posts go unexpectedly viral and finds a subsequent increase in production that is almost completely insensitive to the size of the virality shock—the effects of receiving 100 likes and receiving 3,500 are indistinguishable. He interprets this as evidence of extremely concave returns to viewership. Outside of economics, Lindström et al. (2021) and da Silva Pinho

Figure 10: Creator experiment: identification, creator payoff parameters, and drivers of willingness to post



Notes: Panel (a) explains how creator utility parameters are identified and Panel (b) shows parameter estimates and bootstrapped standard errors. Panels (c) and (d) show the shares of creators willing to post or privately create a lifestyle or political video for \$0.25, separately by their self-reported enviability in the case of lifestyle or ideological extremity in the case of politics. Panel (e) shows the percent of creators selecting each of three multi-choice options in response to a question about why they were willing to post some content types for \$0.25 and not others (asked to all creators for whom that was true). Panel (f) plots responses to a question we ask separately by content type: hypothetically, if the viewership a creator expected for posts of this content type doubled, how would this change their likelihood of posting this type. We split responses by whether this was a type the creator was versus was not willing to post for \$0.25.

et al. (2024) find similar results using the same strategy and Scissors et al. (2016) find similar results when surveying Facebook users about satisfaction from receiving “likes.”

Three additional facts corroborate the result that expectations about viewership play little role in explaining variation in willingness to supply different types of content. First, as Figure 10 Panels (c) and (d) show, willingness to *post* a content type is strongly correlated with willingness to privately *create* it, replicating Figure 6 Panel (e)’s implementation of the same exercise in our user survey. The panels show that enviable creators are substantially more willing to post lifestyle content, but also substantially more willing to privately create such content. The same is true for ideologically extreme creators and political content. Respondents’ binary willingness to post and willingness to create in exchange for \$0.25 agree in 73% of cases.³⁰

Second, we ask respondents a qualitative question about why they were willing to post certain types for \$0.25 but not others. Figure 10 Panel (e) shows that 82% of respondents say this is because “creating posts of [certain types] would be easier or more enjoyable, or [I] would be more comfortable posting them,” rather than because they expect those posts would attract more likes and views. We randomize the wording of the latter option to manipulate its social desirability as detailed in Appendix Figure A50, which changes answer rates by only one percentage point despite successfully flipping perceived social desirability.

Third, for each content type, we ask respondents to imagine that their expected viewership for this content type doubled. We ask whether this would make them more likely to post content of this type. Figure 10 Panel (f) shows that for content types respondents were already willing to post, 71% say yes. For content types respondents were *not* willing to post, only 40% say yes, and 10% say this would make them *less* likely—concentrated among ideological moderates asked about posting political content, who are evidently uncomfortable having others see their political content.

The main result of the experiment—that variation in willingness to supply different types of content is mostly driven by variation in the costs or expressive benefits of production, rather than expectations about demand—is therefore corroborated by several supplementary pieces of evidence. It is also corroborated by our representative survey of TikTok creators, which showed (Figure 5 Panel (b)) that supply-side considerations are rated as being much more important than expectations about viewership in determining production choices, and that this is robust to a social-desirability wording manipulation (Appendix Figure A31).

4.4 Quantification

We now combine our demand- and supply-side parameter estimates and conduct decompositions and counterfactuals. We have two sources of demand-side information. First is our viewership experiment, which yields estimates of utility parameters for each content type. Second is our creator experiment, which collects creators’ expectations about viewership for each type of content they could potentially produce. We can therefore show robustness of our results to using either source of demand-side information.

Panel (a) of Figure 11 conducts Oaxaca decompositions on two objects of interest: the gap in equilibrium consumption between entertainment and lifestyle content, and the gap between partisan and nonpartisan

³⁰Participants are in general more willing to privately create these content types for \$0.25 than to post these content types. One way to make sense of this is that creating a video yields expressive satisfaction that outweighs creation costs, while publicly posting a video brings utility from viewership *and* a fixed cost of exposure to embarrassment.

political content. The decompositions operate by asking how much consumption would change if we re-assigned one content type to have the other’s supply (or demand) parameters and mechanically recalculated consumption without accounting for factors like diminishing marginal utility from a given content type. As expected given the directions of our demand estimates, the gap in consumption between entertainment and lifestyle content reflects a strong supply-side force favoring lifestyle content colliding with a strong demand-side force favoring entertainment. Meanwhile, the gap between partisan and nonpartisan political content is between half and two-thirds explained by a supply-side force favoring partisan content, with the rest explained by demand favoring partisan content.

Panel (b) performs counterfactual simulations that do account for diminishing marginal utility from content types. We perform two separate counterfactuals.

In our lifestyle-focused counterfactual, we assign entertainment’s supply-side parameters to lifestyle and stories, hence eliminating lifestyle’s supply-side advantage over the other two types. In this counterfactual, the equilibrium prevalence of lifestyle content in viewers’ feeds drops by about a third. The magnitude of the drop is attenuated by diminishing marginal utility within content types, which pushes towards equalization of consumption across types.

In our politics-focused counterfactual, we assign partisan political content’s supply-side parameters to nonpartisan content, eliminating partisan content’s supply-side advantage. In this counterfactual, the partisanship of consumed political content drops by more than half. The supply side hence appears more important for partisan than lifestyle content. This is because—even though lifestyle faces a demand disadvantage and partisan content a demand advantage—the supply advantage of partisan over nonpartisan content is quantitatively much larger than the advantage of lifestyle over entertainment.

Overall, results suggest the prevalence of both lifestyle and partisan political content is one- to two-thirds supply-driven.

5 Implications: Emotion-Sampling


Section 4 decomposed equilibrium consumption of lifestyle and partisan political content on TikTok into a component driven by supply-side preferences and one driven by demand-side preferences.

A separate question is whether the demand-side preferences we estimate reflect true viewer preferences or successful exploitation of viewers’ self-control problems. This is a secondary question: our model and results are meaningful whether one interprets viewer “utility” as experienced or decision utility. But it may be interesting to collect some evidence on this question.

In particular, our model in Section 2 suggested that even if viewers dislike some features of a content type, pieces of this type may end up in their feeds thanks to having high residual appeal on *other* dimensions. Appendix K probes this possibility by using emotion data collected via the iPhone app we asked participants to use for a week before sending us their social media histories.

The app detects when a participant is using TikTok, Instagram, or YouTube and pings her at random intervals to ask how she is feeling. This is “experience-sampling,” a measurement technique with a history in economics (Kahneman and Krueger, 2006). We link participants’ responses to the posts they were viewing

Figure 11: Oaxaca decomposition of observed consumption gaps, and counterfactual consumption rates after equalizing supply parameters



Notes: Panel (a) displays a Oaxaca decomposition of observed gaps in consumption rates between lifestyle content and entertainment content, and between partisan and nonpartisan political content, into a supply contribution (how different would consumption rates look if we equalized both types’ demand parameters?) and a demand contribution (how different would they look if we equalized supply parameters?). We plot two versions of the decomposition, one using demand estimates calculated from our viewership experiment and the other using creators’ expected views from our creator experiment. Panel (b) displays, in blue, the “status quo” distribution of consumption simulated from our model using creator-experiment supply estimates and viewership-experiment demand estimates. The red bars show two separate counterfactuals. Left of the axis break, they show the shares of lifestyle, stories, and entertainment in consumption if we give lifestyle and stories the supply parameters of entertainment. Right of the axis break, they show the partisan share of political consumption if we give nonpartisan political content the supply parameters of partisan political content, holding all other supply parameters fixed. We consider only user-generated content when analyzing political content, excluding content created by news organizations.

at the time the pings were randomly delivered. For identification, we restrict to sessions where a participant is scrolling through an infinite vertical feed and therefore not actively choosing which video will appear next. We implement event study specifications that exploit high-frequency video-to-video variation in the content viewers are served. Such local exogenous variation should exist as long as viewers have a taste for unpredictable moment-to-moment variation conditional on the overall composition of their feeds (Eyal, 2014). We evaluate the exogeneity of the variation by examining pre-trends in our event study plots and implementing balance tests checking whether a viewer’s emotions predict the content they are served *next*.

Appendix K reports results. The two content types with the largest negative effects on contemporaneous emotional affect are political content and the subset of lifestyle content that is “envious,” meaning coded as featuring an unusually attractive or wealthy person. This is driven by anger and anxiety in response to political videos and envy- or insecurity-related emotions in response to lifestyle videos.

However, participants seem to like the political and envious lifestyle videos that appear in their feeds even though they trigger certain negative emotions. Our app occasionally asks participants if they regret watching the video they just watched. Panel (c) of Appendix Figure A53 shows that participants only very rarely report that they regret having watched an envious lifestyle or partisan political video. After participants submit their viewership histories, we conduct a retrospective survey where we show them a sequence of randomly selected creators from their feeds and ask if they would like the opportunity to block those creators from appearing in their feeds in future. Panel (d) shows that participants say “yes” in a substantial minority of cases but are *less* likely to say yes for envious lifestyle or political creators compared to random other creators. This is consistent with the possibility that the lifestyle and political creators that do appear in viewers’ feeds are sufficiently positively selected on other dimensions to compensate for the negative emotions they trigger.

These results are limited in several respects. They represent only short-run measures of emotions and could not, for example, capture the possibility that repeated exposure to unfavorable social comparisons over a longer period of time could be psychologically deleterious. They capture only emotions, not other outcomes that might matter, like affective polarization or partisan attitudes. And participants may be reluctant to give answers that admit self-control problems. We therefore do not draw firm conclusions and simply note that this pattern of results is consistent with our theory.

6 Conclusion

Two contrasting perspectives marked the early era of social media. An optimistic view saw social media platforms as “liberation technologies” democratizing communication and freeing society from the tyranny of the elites gatekeeping newspapers, radio, and television (Diamond, 2010). The pessimistic view, which has come to dominate both popular discourse and academic research, sees these platforms as algorithmically tapping into users’ personal insecurities and political grievances, driving polarization and social discontent.

This latter view sees the prevalence of politically divisive and “enviable” lifestyle content on social media as a consequence of viewer demand, and attributes the negative effects of social media to platforms’ sophistication at identifying and exploiting latent demand. From this perspective, the fact that social media enables content creation and distribution by amateurs, whereas legacy media serves content produced by professionals and selected by elite gatekeepers, is largely incidental.

This paper shows that the business model of social media fundamentally derives from a *complementarity* between the production-side differences emphasized by the optimists and the algorithmic sophistication emphasized by the pessimists. By relying on unpaid amateurs rather than paid professionals, platforms amass enormous pools of the types of content that creators are intrinsically motivated to produce. As we have shown, what creators enjoy creating can differ from what viewers enjoy watching. But experimentation and algorithmic personalization enable platforms to convert the enormous scale of amateur content production into viewer surplus by cherry-picking the videos that will most engage each viewer.

The result has been a fundamental transformation of media consumption in liberal democracies. In the 20th century, citizens got their political news from professional journalists and entertained themselves by reading or watching books, television shows, or movies produced by writers, directors, and actors paid to cater to their tastes. In the 2020s, they are likely to get their political news from the tiny slice of the population enthusiastic about diving into political arguments online, and their entertainment from the minority of their peers attractive enough to enjoy showing off in front of a camera. They also spend more time on media consumption than ever before (U.S. Bureau of Labor Statistics, 2026). This is because the turn to amateur production was precipitated by the Internet’s facilitation of an enormously successful business model. This paper has characterized this distinctive business model and taken the first step towards tracing out its consequences.

References

- Acemoglu, D., Huttenlocher, D., Ozdaglar, A., and Siderius, J. (2024a). Online Business Models, Digital Ads, and User Welfare. Working Paper.
- Acemoglu, D., Ozdaglar, A., and Siderius, J. (2024b). A Model of Online Misinformation. *Review of Economic Studies*, 91(6):3117–3150.
- Aguiar, L. and Waldfogel, J. (2018). Quality Predictability and the Welfare Benefits from New Products: Evidence from the Digitization of Recorded Music. *Journal of Political Economy*, 126(2):492–524.
- Allcott, H., Braghieri, L., Eichmeyer, S., and Gentzkow, M. (2020). The Welfare Effects of Social Media. *American Economic Review*, 110(3):629–676.
- Allcott, H., Castillo, J. C., Gentzkow, M. A., Musolff, L., and Salz, T. (2025). Sources of Market Power in Web Search: Evidence from a Field Experiment. *American Economic Review*. Forthcoming.
- Allcott, H. and Gentzkow, M. (2017). Social Media and Fake News in the 2016 Election. *Journal of Economic Perspectives*, 31(2):211–236.
- Allcott, H., Gentzkow, M., and Song, L. (2022). Digital Addiction. *American Economic Review*, 112(7):2424–2463.
- Andrews, T. M. (2020). Charli D’Amelio Is TikTok’s Biggest Star. She Has No Idea Why. *The Washington Post*.
- Angelucci, C., Cagé, J., and Sinkinson, M. (2024). Media competition and news diets. *American Economic Journal: Microeconomics*, 16(2):62–102.
- Aridor, G., Chou, W., Kallus, N., Scheid, A., Tren, A., and Zielinski, K. (2026). Recommendation Quality and the Concentration of Consumption: Experimental Evidence from Netflix. *arXiv preprint arXiv:2608.21274*.
- Aridor, G., Jiménez-Durán, R., Levy, R., and Song, L. (2024). The Economics of Social Media. *Journal of Economic Literature*, 62(4):1422–1474.
- Bail, C. (2021). *Breaking the Social Media Prism: How to Make Our Platforms Less Polarizing*. Princeton University Press, Princeton.
- Balbuzanov, I., Gars, J., Stalinski, M., and Tjernström, E. (2025). Incentivizing Engagement: Experimental Evidence on Journalist Performance Pay. Working Paper.
- Bao, R. (2026). “Just One More Clip”: Short Videos, Big Self-Control Problems. Working Paper.
- Beattie, G., Durante, R., Knight, B., and Sen, A. (2021). Advertising Spending and Media Bias: Evidence from News Coverage of Car Safety Recalls. *Management Science*, 67(2):698–719.
- Becker, G. S. and Murphy, K. M. (1988). A Theory of Rational Addiction. *Journal of Political Economy*, 96(4):675–700.
- Becker, G. S. and Murphy, K. M. (1993). A Simple Theory of Advertising as a Good or Bad. *Quarterly Journal of Economics*, 108(4):941–964.
- Beknazar-Yuzbashev, G., Capozza, F., Rutherford, K., and Stalinski, M. (2026). Social Media Toxicity and Mental Health. *SSRN Working Paper*.
- Beknazar-Yuzbashev, G., Jiménez-Durán, R., McCrosky, J., and Stalinski, M. (2025a). Toxic Content and User Engagement on Social Media: Evidence from a Field Experiment. Working Paper.

- Beknazar-Yuzbashev, G., Jiménez-Durán, R., Simonov, A., and Stalinski, M. (2025b). Social Media Advertising Loads as Prices. Working Paper.
- Beknazar-Yuzbashev, G., Jiménez-Durán, R., and Stalinski, M. (2024). A Model of Harmful yet Engaging Content on Social Media. In *AEA Papers and Proceedings*, volume 114, pages 678–683.
- Belfi, A. M., Vessel, E. A., and Starr, G. G. (2018). Individual Ratings of Vividness Predict Aesthetic Appeal in Poetry. *Psychology of Aesthetics, Creativity, and the Arts*, 12(3):341.
- Besley, T. and Prat, A. (2006). Handcuffs for the Grabbing Hand? Media Capture and Government Accountability. *American Economic Review*, 96(3):720–736.
- Boeker, M. and Urman, A. (2022). An Empirical Investigation of Personalization Factors on TikTok. In *Proceedings of the ACM Web Conference 2022*, pages 2298–2309.
- Bolte, L., Conlon, J. J., and Hue, J. (2026). Social Media and the Limits to Selective Attention. Working Paper.
- Boxell, L., Conway, J., and Longuet-Marx, N. (2026). Journalist Ideology and the Production of News: Evidence from Movers. Working Paper.
- Braghieri, L., Eichmeyer, S., Levy, R., Mobius, M. M., Steinhardt, J., and Zhong, R. (2024). Article-Level Slant and Polarization of News Consumption on Social Media. Working Paper.
- Braghieri, L., Levy, R., and Makarin, A. (2022). Social Media and Mental Health. *American Economic Review*, 112(11):3660–3693.
- Braghieri, L., Levy, R., and Trachtman, H. (2026). Frictions in News Consumption: Evidence from Social Media. Working Paper.
- Brynjolfsson, E., Hu, Y., and Smith, M. D. (2003). Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers. *Management Science*, 49(11):1580–1596.
- Bursztyn, L., Gentzkow, M., Jiménez-Durán, R., Leonard, A., Milojević, F., and Roth, C. (2026). Measuring Markets for Network Goods. *Review of Economic Studies*.
- Bursztyn, L., Handel, B., Jiménez-Durán, R., and Roth, C. (2025). When Product Markets Become Collective Traps: The Case of Social Media. *American Economic Review*, 115(12):4105–4136.
- Cagé, J., Hengel, M., Hervé, N., and Urvoy, C. (2025). Political Bias in the Media: Evidence from the Universe of French Broadcasts, 2002–2020. Working Paper.
- Carlson, J. (2026). Interpretable Discovery from Unstructured Data: A High-Dimensional Approach. Working Paper.
- Caves, R. E. (2003). Contracts Between Art and Commerce. *Journal of Economic Perspectives*, 17(2):73–83.
- Chopra, F. and Haaland, I. (2023). Conducting Qualitative Interviews with AI. Working Paper.
- Cruz, C. (2024). Ironmouse Is Twitch’s Top Streamer, but She Doesn’t Know Why You’re Watching. *Rolling Stone*.
- da Silva Pinho, A., Céspedes Izquierdo, V., Lindström, B., and van den Bos, W. (2024). Youths’ Sensitivity to Social Media Feedback: A Computational Account. *Science Advances*, 10(43):eadp8775.
- David, P. A. (1990). The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox. *American Economic Review*, 80(2):355–361.

- De Vany, A. and Walls, W. D. (1999). Uncertainty in the Movie Industry: Does Star Power Reduce the Terror of the Box Office? *Journal of Cultural Economics*, 23(4):285–318.
- Diamond, L. (2010). Liberation Technology. *Journal of Democracy*, 21(3):69–83.
- Donati, D. and Song, L. (2026). The Influence of the Vocal Few: Evidence from Social Media Comments. Working Paper.
- Durante, R. and Knight, B. (2012). Partisan Control, Media Bias, and Viewer Responses: Evidence from Berlusconi's Italy. *Journal of the European Economic Association*, 10(3):451–481.
- Dutz, D., Huitfeldt, I., Lacouture, S., Mogstad, M., Torgovitsky, A., and Van Dijk, W. (2026). Selection in Surveys: Using Randomized Incentives to Detect and Account for Nonresponse Bias. *Review of Economic Studies*, 93(5):3086–3132.
- Eisenberger, R. (1970). Is There a Deprivation-Satiation Function for Social Approval? *Psychological Bulletin*, 74(4):255–275.
- Enikolopov, R., Makarin, A., and Petrova, M. (2020). Social Media and Protest Participation: Evidence from Russia. *Econometrica*, 88(4):1479–1514.
- Eyal, N. (2014). *Hooked: How to Build Habit-Forming Products*. Portfolio/Penguin, New York.
- Fergusson, L. and Molina, C. (2019). Facebook causes protests. *Documento Cede*, (41).
- Filippas, A., Horton, J. J., Lipnowski, E., and Parasurama, P. (2026). The Production and Consumption of Social Media. *Management Science*, 72(7):6368–6388.
- Fishburn, P. C. (1976). Noncompensatory Preferences. *Synthese*, 33(1):393–403.
- Freyaldenhoven, S., Hansen, C., Pérez Pérez, J., and Shapiro, J. M. (2026). Visualization, Identification, and Estimation in the Linear Panel Event-Study Design. In Chernozhukov, V., Hörner, J., La Ferrara, E., and Werning, I., editors, *Advances in Economics and Econometrics: Twelfth World Congress*, pages 225–268. Cambridge University Press, Cambridge.
- Geiecke, F. and Jaravel, X. (2024). Conversations at Scale: Robust AI-Led Interviews with a Simple Open-Source Platform. *Review of Economic Studies*. Forthcoming.
- Gentzkow, M. and Shapiro, J. M. (2006). Media Bias and Reputation. *Journal of Political Economy*, 114(2):280–316.
- Gentzkow, M. and Shapiro, J. M. (2010). What Drives Media Slant? Evidence from U.S. Daily Newspapers. *Econometrica*, 78(1):35–71.
- Gilchrist, D. S. and Sands, E. G. (2016). Something to Talk About: Social Spillovers in Movie Consumption. *Journal of Political Economy*, 124(5):1339–1382.
- Goldman, W. (1983). *Adventures in the Screen Trade: A Personal View of Hollywood and Screenwriting*. Warner Books, New York.
- Guriev, S., Henry, E., Marquis, T., and Zhuravskaya, E. (2026). Curtailing False News, Amplifying Truth. *Econometrica*, 94(5):1563–1601.
- Haidt, J. (2024). *The Anxious Generation: How the Great Rewiring of Childhood Is Causing an Epidemic of Mental Illness*. Penguin Press, New York.
- Herman, E. S. and Chomsky, N. (1988). *Manufacturing Consent: The Political Economy of the Mass Media*. Pantheon Books, New York.

- Instagram (2024). Helping Creators Find New Audiences. Instagram for Creators.
- Jeon, D.-S. and Rey, P. (2024). Platform Competition and Innovation. Working Paper.
- Kahneman, D. and Krueger, A. B. (2006). Developments in the Measurement of Subjective Well-Being. *Journal of Economic Perspectives*, 20(1):3–24.
- Kleinberg, J., Lakkaraju, H., Leskovec, J., Ludwig, J., and Mullainathan, S. (2018). Human Decisions and Machine Predictions. *Quarterly Journal of Economics*, 133(1):237–293.
- Kolitz, D. (2026). Microdramas, TV Shows, and Hollywood Entertainment. *The New York Times Magazine*.
- Levy, R. (2021). Social Media, News Consumption, and Polarization: Evidence from a Field Experiment. *American Economic Review*, 111(3):831–870.
- Lindström, B., Bellander, M., Schultner, D. T., Chang, A., Tobler, P. N., and Amodio, D. M. (2021). A Computational Reward Learning Account of Social Media Engagement. *Nature Communications*, 12(1):1311.
- Lumet, S. (1995). *Making Movies*. Alfred A. Knopf, New York.
- Martin, G. J. and McCrain, J. (2019). Local News and National Politics. *American Political Science Review*, 113(2):372–384.
- Martin, G. J. and Yurukoglu, A. (2017). Bias in Cable News: Persuasion and Polarization. *American Economic Review*, 107(9):2565–2599.
- Moehring, A. (2026). Personalization, engagement, and content quality on social media: An evaluation of reddit’s news feed. *Management Science*.
- Movva, R., Peng, K., Garg, N., Kleinberg, J., and Pierson, E. (2025). Sparse Autoencoders for Hypothesis Generation. In *Proceedings of the 42nd International Conference on Machine Learning*, volume 267 of *Proceedings of Machine Learning Research*, pages 44997–45023.
- Mullainathan, S. and Shleifer, A. (2005). The Market for News. *American Economic Review*, 95(4):1031–1053.
- Müller, K. and Schwarz, C. (2021). Fanning the Flames of Hate: Social Media and Hate Crime. *Journal of the European Economic Association*, 19(4):2131–2167.
- Noy, S. and Rao, A. (2025). The Business of the Culture War. Working Paper.
- O’Meara, V. (2026). The Many Meanings of Measure: Social Media Metrics and Their Implications for Collective Organizing in the Influencer Economy. *New Media & Society*.
- People v. TikTok Inc.* (2024). Complaint for Injunctive and Other Relief. Superior Court of California, County of Santa Clara, No. 24CV449203. Filed October 8, 2024. Public-redacted version.
- Postman, N. (1985). *Amusing Ourselves to Death: Public Discourse in the Age of Show Business*. Viking Penguin, New York.
- Pynchon, T. (1990). *Vineland*. Little, Brown and Company, Boston.
- Qin, B., Strömberg, D., and Wu, Y. (2018). Media Bias in China. *American Economic Review*, 108(9):2442–2476.
- Rambachan, A. (2024). Identifying Prediction Mistakes in Observational Data. *Quarterly Journal of Economics*, 139(3):1665–1711.

- Ren, D., Stavrova, O., and Loh, W. W. (2022). Nonlinear Effect of Social Interaction Quantity on Psychological Well-Being: Diminishing Returns or Inverted U? *Journal of Personality and Social Psychology*, 122(6):1056–1074.
- Ronson, J. (2026). He Was a Top Lawn Care Influencer on YouTube. Then Everything Went Haywire. *The New York Times Magazine*.
- Rosen, S. (1981). The Economics of Superstars. *American Economic Review*, 71(5):845–858.
- Rosenbusch, H., Evans, A. M., and Zeelenberg, M. (2022). The Relative Importance of Joke and Audience Characteristics in Eliciting Amusement. *Psychological Science*, 33(9):1386–1394.
- Salganik, M. J., Dodds, P. S., and Watts, D. J. (2006). Experimental Study of Inequality and Unpredictability in an Artificial Cultural Market. *Science*, 311(5762):854–856.
- Schulz, W., Robertson, C., Sjøstad, H., Ashokkumar, A., Oswald, L., Yang, J., Munzert, S., Van Bavel, J., and Guess, A. (2026). Online Discourse Is Falsely Polarized by Self-Selection. Working Paper.
- Scissors, L., Burke, M., and Wengrovitz, S. (2016). What’s in a Like? Attitudes and Behaviors Around Receiving Likes on Facebook. In *Proceedings of the 19th ACM Conference on Computer-Supported Cooperative Work & Social Computing*, pages 1501–1510. Association for Computing Machinery.
- Sibley, F. (1959). Aesthetic Concepts. *The Philosophical Review*, 68(4):421–450.
- Solheim, H. and Wang, N. (2026). Algorithms and Parasitic Content. Working Paper.
- Srinivasan, K. (2024). Paying Attention. Working Paper.
- Steel, B., Schirmer, M., Ruths, D., and Pfeffer, J. (2026). Just Another Hour on TikTok: ID Sampling to Obtain a Complete Slice of TikTok. *Journal of Quantitative Description: Digital Media*, 6.
- Sweet, J. (2023). The Instagram Reels Gold Rush. *The New Yorker*. Rabbit Holes column.
- Szeidl, A. and Szucs, F. (2021). Media Capture Through Favor Exchange. *Econometrica*, 89(1):281–310.
- Tait, A. (2023). Unboxing Is So 2012. The Internet Wants Packing Videos. *Wired*.
- TikTok (2020). How TikTok Recommends Videos #ForYou. TikTok Newsroom.
- Tversky, A. (1969). Intransitivity of Preferences. *Psychological Review*, 76(1):31–48.
- U.S. Bureau of Labor Statistics (2026). American Time Use Survey, Table A-1: Time Spent in Detailed Primary Activities and Percent of the Civilian Population Engaging in Each Activity, Averages per Day by Sex, 2025 Annual Averages.
- van der Maaten, L. and Hinton, G. (2008). Visualizing Data using t-SNE. *Journal of Machine Learning Research*, 9:2579–2605.
- Vessel, E. A. and Rubin, N. (2010). Beauty and the Beholder: Highly Individual Taste for Abstract, but Not Real-World Images. *Journal of Vision*, 10(2):18–18.
- Waldfogel, J. (2016). Cinematic Explosion: New Products, Unpredictability and Realized Quality in the Digital Era. *Journal of Industrial Economics*, 64(4):755–772.
- Waldfogel, J. (2017). How Digitization Has Created a Golden Age of Music, Movies, Books, and Television. *Journal of Economic Perspectives*, 31(3):195–214.
- Waldfogel, J. (2019). *Digital Renaissance: What Data and Economics Tell Us About the Future of Popular Culture*. Princeton University Press.

- Wallisch, P. and Whritner, J. A. (2017). Strikingly Low Agreement in the Appraisal of Motion Pictures. *Projections*, 11(1):102–120.
- Zannettou, S., Nemes-Nemeth, O., Ayalon, O., Goetzen, A., Gummadi, K. P., Redmiles, E. M., and Roesner, F. (2024). Analyzing User Engagement with TikTok’s Short Format Video Recommendations Using Data Donations. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, pages 1–16.
- Zhuravskaya, E., Petrova, M., and Enikolopov, R. (2020). Political Effects of the Internet and Social Media. *Annual Review of Economics*, 12:415–438.

Online Appendix

“The Supply Side of Social Media”


Shakke Noy and Aakaash Rao

Contents

A	Additional Figures	56
B	Results for Instagram, X/Twitter, and YouTube	65
C	Observational Data and Classification	69
C.1	Descriptive statistics on samples	71
C.2	Content classification	78
D	Identifying Algorithmic Markups	86
E	Creator Survey, Monetization, and Motivations	89
F	Quotes from Qualitative Interviews	95
G	Model: Proofs and Extensions	108
G.1	Proofs	108
G.2	Short-form content	114
H	TV Consumption	118
I	Viewership Experiment	120
I.1	Design details	120
I.2	Additional exhibits	121
J	Creator Experiment Details	127
J.1	Design details	127
J.2	Additional exhibits	129
J.3	Social desirability	130
K	Emotion-Sampling Details and Results	132

A Additional Figures

Appendix Figure A1: Production and consumption of content on TikTok, weighting videos by their view counts




Notes: This figure reproduces Panel (a) of Figure 6, but instead of measuring consumption using our participants’ viewership histories, it measures consumption by taking the Research API produced-video sample and weighting each video by its view count.

Appendix Figure A2: Viewers prefer nonpartisan news outlets on TikTok


(a) Top political news outlets in Democrats’ and Republicans’ feeds (b) News organizations’ more partisan videos receive fewer views on TikTok

Democrats	Republicans
1. Daily Mail	1. Daily Mail
2. CBS News	2. CBS News
3. CNN	3. NBC News
4. NBC News	4. Fox News
5. ABC News	5. ABC News
6. MSNOW	6. CNN
7. The Daily Show	7. New York Post
8. CBS Mornings	8. CBS Mornings
9. NBC New York	9. Metro UK
10. New York Post	10. Reuters




Notes: These panels are constructed from our data on participants’ viewership histories plus TV data from the TV News Archive (Noy and Rao, 2025). Panel (a) lists the ten news outlets responsible for the largest numbers of videos about American politics in participants’ feeds, separately for participants who self-identify as Democrats versus Republicans. The blue dots in Panel (b) use political news videos posted on TikTok by ABC, CBS, CNN, NBC, and Fox News, weighting each outlet equally. The panel shows a binned scatterplot of whether the video was partisan against its view count on TikTok. The orange diamond instead plots the percent of 60-second political news segments broadcast on these channels in 2026 that we code as partisan.

Appendix Figure A3: *Revealing like counts does not increase engagement with highly-liked videos*




Notes: This figure plots data from our viewership experiment. We randomize at the participant level whether the “like” count on TikTok of the videos they watch is visible or not. The panels show coefficients from regressions of our pre-registered engagement measures on the video’s like count separately by whether the count was hidden or not, using our main specification that includes fixed effects for participants, video length decile, and sampling source, the latter interacted with the hidden condition.

Appendix Figure A4: Identifying tested videos: additional evidence



Notes: This figure plots data from our viewership histories, focusing on “test videos” (Figure 4), which are videos that appear in a respondent’s feed within 5 hours after being posted. Panel (a) shows that test viewers’ average engagement with a video (on the y-axis) is predictive of the video’s ultimate views (x-axis) even, for example, above 1 million views. Engagement is measured as seconds watched, controlling for viewer by video duration fixed effects. Panel (b) shows that the engagement of the first test viewer in our data to get a given test video is predictive of the engagement of other viewers in our data who see the same video over the next seven days. Panels (c) and (d) replicate the orange dots from Figure 4 Panel (b) but splitting test videos by whether they come from small or large creators (Panel (c)) and whether they come from a creator who has appeared in the viewer’s history before or not (Panel (d)).


Appendix Figure A6: News organizations’ share of political videos falls as a session goes on




Notes: This figure replicates Panel (f) of Figure 5 but plots on the y-axis the share of political videos posted by news organizations.

Appendix Figure A5: Preference dispersion $\rightarrow$ pass-through of all content types to consumption


(a) Pass-through: Videos from categories a viewer doesn't usually see are larger outliers




(b) Pass-through: ... and come from creators the viewer rarely sees again



(c) Scarcity: Heavier users get more of their feed from styles they see less than the average viewer...




(d) Scarcity: ... and so do viewers further into a session




Notes: This figure uses data from our TikTok viewership histories and uses our schema categorizing videos into 17 categories described in Appendix C. The x-axes in Panels (a) and (b) plot the percent of a participant's feed in a given category, residualized on participant and category fixed effects, so values above zero indicate categories the participant sees more frequently than other participants. The y-axes plot the dispersion measures from Panels (c) and (d) of Figure 5: the ratio of the view count of a video appearing in the participant's feed to the creator's median view count on their videos, and whether this creator is observed in this participant's feed more than once. The y-axes on Panels (c) and (d) plot the percent of videos a participant sees from categories they see less often than average. Both x-axes plot proxies for content scarcity. The x-axis in Panel (c) plots the participant's average daily hours spent on TikTok based on their viewership history. The x-axis in Panel (d) is the same as in Panel (f) of Figure 5: how far into a given TikTok session a participant is.

Appendix Figure A7: *The distribution of views has a similar shape for every content type*




Notes: This figure plots the view count distributions of videos randomly sampled from the pool of all videos produced on TikTok, via the Research API. We plot distributions separately for our four broad content types and by bins of creator follower counts. We center distributions at the type-by-follower bin's median viewership.

Appendix Figure A8: *Production, consumption, and simulated consumption of content types, using demand estimates from creators' expectations*




Notes: This figure replicates Figure 9 Panel (d) but simulates consumption using demand estimates sourced from the expectations of creators in our creator experiment, rather than from our viewership experiment.

Appendix Figure A9: Enviability, ideological extremity, and the share of posts of each content type




Notes: This figure reproduces Figure 6 Panel (d), but the outcome is the share of a user’s scraped posts falling into each content type rather than an indicator for having posted that type. We scrape every TikTok post users made from April to June 2026 through the Research API, users’ recent posts on X/Twitter, and users’ 12 most recent posts plus up to 30 Reels on Instagram. Each coefficient is from a regression of the share on enviability and ideological extremity (both standardized) with platform fixed effects, with one observation per user and platform.

Appendix Figure A10: *Creators who say they get different rewards from creating post different types of content*




Notes: This figure uses data from our creator survey. It plots coefficients from regressions of the shares of a creator's posts (which we scrape from their account and classify) falling into each of our content quadrants on indicators for whether the creator reported receiving a specific nonpecuniary reward from content production.

Appendix Figure A11: Discontinuities in content by video length, separately for each sample and content measure




Appendix Figure A11: Discontinuities in content by video length, separately for each sample and content measure (continued)



Notes: To maximize sample size, Panels (a)-(c) of Figure 7 pool data on produced videos randomly sampled from TikTok’s Research API and data on consumed videos from viewership histories. Panels (a) and (b) also use predicted classifications based on a classifier trained on our LLM codings, to extend the classifications to a larger sample without incurring monetary cost. This figure shows results separately for each data source and each coding method.


B Results for Instagram, X/Twitter, and YouTube

Appendix Figure A12: Attitudes towards Instagram, TikTok, X/Twitter, and YouTube




Notes: This figure shows responses from a survey of 301 Americans aged 18–30 recruited on Prolific. Panel (a) plots the share of respondents who say “yes” to “Would you prefer to live in a world where [platform] did not exist,” separately for each of Instagram, TikTok, X/Twitter, and YouTube, and separately by respondents who report using versus not using the relevant platform. Panels (b) and (c) plot the share of respondents who say they associate this platform with “having a lot of content about politics and political opinions” and “having a lot of lifestyle content, meaning content focused on appearance, fashion, fitness, wealth, and so on,” respectively.

Appendix Figure A13: *Distribution of posts in viewers' feeds across account sizes, by platform*



Notes: This figure replicates Figure 4 Panel (a) for each of our four platforms of interest. Private Instagram accounts whose follower counts cannot be scraped are allocated to the < 10k bucket.


Appendix Figure A14: What fraction of posts in viewers’ feeds earn money for their creators? Instagram, X/Twitter, and YouTube




Notes: This figure reproduces Appendix Figure A29, asking what percentage of posts viewers see on each platform come from creators who appear to ever earn money on those platforms. Instagram has no program analogous to Creator Rewards. We therefore plot the percentage of creators offering paid subscriptions via their profile (“Subscriptions”) and the percentage of creators with any Reel tagged as a paid partnership (“Paid Partnership”) over the 7 months covered by our data. X/Twitter has a revenue sharing program, so the green bars plot the percentage of creators estimated to make at least \$5/hour assuming the commonly reported payout rate of \$8.50/million views and assuming a text-based post takes 6 minutes, a photo-based post takes 15, and a video 30. The purple bars plot the share of creators who have any posts with a #ad, #sponsored, or other similar tag suggesting third-party money. YouTube also has a revenue sharing program for channels with at least 1,000 subscribers. The green bars assume \$2 per 1,000 views for main videos and \$0.05 for Shorts based on common estimates, and assume conservatively that a Short takes 30 minutes and a long video takes 2 hours to produce. The purple bars use YouTube’s paid-promotion tags.

Appendix Figure A15: Dispersion on Instagram, X/Twitter, and YouTube


(a) Instagram: Reels' views relative to their creator's median views




(b) X/Twitter: posts' views relative to their creator's median views




(c) YouTube: videos' views relative to their channel's median views



(d) Instagram: how many times viewers see each creator




(e) YouTube: how many times viewers see each channel




Notes: This figure reproduces Panels (c) and (d) of Figure 5 for our other platforms. Panels (a)-(c) plot the ratio of a consumed video's view count to its creator's median view count across their other videos. For YouTube and X/Twitter, this is straightforward. For Instagram, we do not see which specific post from a creator appeared in a person's feed, so instead we sample recent Reels by this creator, weight them in proportion to their views, and divide by the creator's median Reels views. Panels (d) and (e) plot the number of times a creator appears. This is not possible on X/Twitter because our scraped data for a participant's feed only contains 300 recent posts, not enough to observe repeated appearances by creators over any meaningful time period.

Appendix Figure A16: Scarcity on YouTube: stale videos within sessions



Notes: This figure reproduces Figure 5 Panel (f) for YouTube. Panel (a) includes all YouTube videos and Panel (b) includes YouTube Shorts. This is not possible for Instagram because we observe only the creators appearing in a viewer’s feed and not which specific posts appear, and not possible on X/Twitter because we only observe an automated scroll of the first 300 posts in a feed.

Appendix Figure A17: Content consumed on Instagram, TikTok, X/Twitter, and YouTube



Notes: This figure reproduces Figure 6 Panel (a) for our other three platforms. All four platforms are classified using GPT-5-mini and the same classification prompt. For Instagram, we see which creators appear in a participant’s feed but not which specific posts, so we scrape creators’ recent posts, classify them, and weight by their view counts. For YouTube, we exclude videos classified as music videos (which are presumably listened to but not watched).

C Observational Data and Classification

Recruitment for viewership histories, user survey, and experiments. Participants in our surveys and experiments are 18-30 year old Americans predominantly recruited from Prolific (with a few recruited from the survey platform MobLab). Appendix Tables A1–A4 compare them to nationally representative samples of users of the same platforms. Relative to the general youth population, users of our platforms, especially TikTok, skew more female, left-wing, and ideologically extreme. Our samples skew even further in these directions as well as being more educated.

Data for our viewership history collection and emotion sampling were collected between October 2025 and January 2026 and our hypotheses and empirical strategy for estimating the effects of social media content on emotions were pre-registered on the Open Science Framework at <https://osf.io/8ctq7>. We screened 6,529 Prolific users for eligibility (asking about their social media usage) and invited 2,513 to an onboarding

survey. 673 register for our emotion-sampling iPhone app and 383 complete a week of using the app and submit their viewership histories. Appendix Table A5 traces out sample characteristics at each step of this process.

Data for our user survey were collected between June and July 2026. Users were required to log in with either a TikTok or X/Twitter account to access the survey and were paid bonuses for also logging in with the other account and/or giving us their Instagram handle (since we could not build a login tool for Instagram). Our YouTube login tool was not approved at the time so we retrospectively re-survey participants in our user survey to get them to log in with YouTube in September 2026.

Social media histories. We collect social media histories for Instagram, TikTok, and YouTube.

The common features of these histories are as follows: each contains a chronological list of the posts a user has seen on the platform, with the post’s URL (TikTok and YouTube) or just the username of the account that made the post (Instagram), and with a timestamp indicating when the user swiped into or clicked onto this post. These records include ads the user sees as they swipe or watch. For Instagram and YouTube these ads are in a separately labelled file, whereas for TikTok they appear as URLs and we scrape metadata indicating whether a video is an ad. We use the URLs to scrape video metadata, including duration, view and like counts, transcripts, captions, thumbnails, top comments, commercial metadata (ad and sponsorship tags), and details about the creator such as their follower count, posting volume, bio, and so on.

We use the timestamps to infer how long a user spends on a particular post. For example, if a user swipes into a 65-second video and then swipes into another video 18 seconds later, we infer that they watched the first video for 18 seconds. We cap this at the length of the video (65 seconds). We also use the timestamps to separate out contiguous “sessions” of usage; a session ends at a video where the next video begins at least 5 minutes later (TikTok, Instagram, YouTube Shorts) or at least 30 minutes later (YouTube). The last video in a session is not assigned a watch time. Appendix Table A6 shows descriptive statistics about videos and sessions.

Where these histories differ is on the following points. The Instagram histories reach back 7 days, the TikTok histories 6 months, and the YouTube histories back to the creation of an account. The Instagram histories also only allow us to observe which account made the post seen by the user, not the URL of the specific post itself. For this reason, our descriptive results using Instagram scrape a sample of recent posts for these accounts and take a view-weighted average across them.

C.1 Descriptive statistics on samples

Appendix Table A1: Our samples of TikTok users compared to nationally representative surveys

	Our samples			Pew ATP 2023-25		ANES 2024		YouGov 2025	
	Viewership Histories	Viewership Experiment	User Survey	All	TikTok Users	All	TikTok Users	All	TikTok Users
<i>Demographics</i>									
Male	0.21	0.28	0.28	0.47	0.36	0.50	0.42	0.51	0.42
Age	25.14	25.18	24.26	–	–	24.53	24.62	24.84	24.64
Bachelor's Degree	0.60	0.32	0.37	0.32	0.28	0.29	0.31	0.25	0.25
<i>Party (Including Leaners)</i>									
Democrat	0.64	–	–	0.56	0.63	0.47	0.55	0.55	0.59
Republican	0.18	–	–	0.38	0.33	0.43	0.38	0.34	0.31
Strong Democrat or Strong Republican	0.32	–	–	0.26	0.30	0.37	0.34	0.41	0.44
<i>Ideology (7-Point Scale)</i>									
Liberal	0.69	0.56	0.62	–	–	0.37	0.46	–	–
Conservative	0.18	0.19	0.16	–	–	0.35	0.29	–	–
Ideological Extremity (0-3)	1.76	1.50	1.56	–	–	1.35	1.41	–	–
<i>Ideology (5-Point Scale)</i>									
Liberal	–	–	–	0.35	0.39	–	–	0.41	0.44
Conservative	–	–	–	0.24	0.21	–	–	0.24	0.21
<i>Ideology (Either Scale)</i>									
Most Liberal or Most Conservative Point	0.24	0.23	0.23	0.21	0.21	0.19	0.18	0.27	0.27
Observations	383	586	2,657	2,468	1,372	514	292	4,573	2,498

Note: This table compares the samples of 383 Prolific users donating their viewership histories, 586 completing our viewership experiment, and 2,657 (TikTok users) completing our user survey, all of whom are Americans aged 18-30, to samples of 18-30 year olds from nationally representative surveys that ask about TikTok usage. The surveys are Pew's American Trends Panel from August 2023, March 2024, and March 2025; the American National Election Studies 2024 Time Series sample; and a 2025 YouGov survey.

Appendix Table A2: Our samples of Instagram users compared to nationally representative surveys

	Our samples		Pew ATP 2024		ANES 2024		YouGov 2025	
	Viewership Histories	User Survey	All	Instagram Users	All	Instagram Users	All	Instagram Users
<i>Demographics</i>								
Male	0.20	0.33	0.46	0.44	0.50	0.45	0.51	0.48
Age	25.04	23.97	–	–	24.53	24.68	24.84	24.68
Bachelor's Degree	0.62	0.36	0.34	0.36	0.29	0.34	0.25	0.27
<i>Party (Including Leaners)</i>								
Democrat	0.65	–	0.59	0.62	0.47	0.53	0.55	0.59
Republican	0.18	–	0.35	0.34	0.43	0.37	0.34	0.31
Strong Democrat or Strong Republican	0.32	–	–	–	0.37	0.30	0.41	0.42
<i>Ideology (7-Point Scale)</i>								
Liberal	0.71	0.60	–	–	0.37	0.44	–	–
Conservative	0.17	0.17	–	–	0.35	0.29	–	–
Ideological Extremity (0-3)	1.78	1.56	–	–	1.35	1.41	–	–
<i>Ideology (5-Point Scale)</i>								
Liberal	–	–	0.37	0.39	–	–	0.41	0.45
Conservative	–	–	0.24	0.23	–	–	0.24	0.22
<i>Ideology (Either Scale)</i>								
Most Liberal or Most Conservative Point	0.24	0.23	0.22	0.21	0.19	0.20	0.27	0.27
Observations	335	1,725	771	614	514	370	4,573	3,238

Note: This table replicates Appendix Table A1 for Instagram. It omits the viewership experiment, which used TikTok viewers only, and it keeps respondents from the user survey who gave their Instagram handle and respondents from the viewership history panel who reported using Instagram regularly in our screening survey.

Appendix Table A3: Our samples of YouTube users compared to nationally representative surveys

	Our samples		Pew ATP 2025		ANES 2024		YouGov 2025	
	Viewership Histories	User Survey	All	YouTube Users	All	YouTube Users	All	YouTube Users
<i>Demographics</i>								
Male	0.23	0.42	0.49	0.50	0.50	0.54	0.51	0.53
Age	25.14	24.46	–	–	24.53	24.59	24.84	24.70
Bachelor's Degree	0.58	0.36	0.31	0.31	0.29	0.34	0.25	0.24
<i>Party (Including Leaners)</i>								
Democrat	0.65	–	0.48	0.49	0.47	0.52	0.55	0.57
Republican	0.17	–	0.48	0.48	0.43	0.39	0.34	0.33
Strong Democrat or Strong Republican	0.32	–	0.26	0.26	0.37	0.32	0.41	0.41
<i>Ideology (7-Point Scale)</i>								
Liberal	0.69	0.60	–	–	0.37	0.44	–	–
Conservative	0.16	0.15	–	–	0.35	0.31	–	–
Ideological Extremity (0-3)	1.74	1.54	–	–	1.35	1.40	–	–
<i>Ideology (5-Point Scale)</i>								
Liberal	–	–	0.29	0.30	–	–	0.41	0.43
Conservative	–	–	0.26	0.26	–	–	0.24	0.23
<i>Ideology (Either Scale)</i>								
Most Liberal or Most Conservative Point	0.24	0.26	0.18	0.19	0.19	0.19	0.27	0.30
Observations	332	685	760	725	514	378	4,573	3,294

Note: This table replicates Appendix Table A1 for YouTube. It omits the viewership experiment, which used TikTok viewers only, and it keeps respondents from the user survey who logged in with their YouTube account and respondents from the viewership history panel who reported using YouTube regularly in our screening survey.

Appendix Table A4: Our samples of X/Twitter users compared to nationally representative surveys

	Our samples		Pew ATP 2024-25		ANES 2024		YouGov 2025	
	Viewership Histories	User Survey	All	X/Twitter Users	All	X/Twitter Users	All	X/Twitter Users
<i>Demographics</i>								
Male	0.28	0.46	0.47	0.59	0.50	0.61	0.51	0.61
Age	25.32	24.25	–	–	24.53	24.68	24.84	24.89
Bachelor's Degree	0.60	0.39	0.32	0.33	0.29	0.40	0.25	0.28
<i>Party (Including Leaners)</i>								
Democrat	0.61	–	0.54	0.58	0.47	0.58	0.55	0.56
Republican	0.21	–	0.42	0.40	0.43	0.34	0.34	0.33
Strong Democrat or Strong Republican	0.30	–	0.26	0.30	0.37	0.30	0.41	0.41
<i>Ideology (7-Point Scale)</i>								
Liberal	0.64	0.64	–	–	0.37	0.47	–	–
Conservative	0.20	0.14	–	–	0.35	0.29	–	–
Ideological Extremity (0-3)	1.70	1.63	–	–	1.35	1.43	–	–
<i>Ideology (5-Point Scale)</i>								
Liberal	–	–	0.33	0.36	–	–	0.41	0.42
Conservative	–	–	0.25	0.27	–	–	0.24	0.23
<i>Ideology (Either Scale)</i>								
Most Liberal or Most Conservative Point	0.25	0.27	0.20	0.23	0.19	0.19	0.27	0.27
Observations	156	1,667	1,531	628	514	217	4,573	1,744

Note: This table replicates Appendix Table A1 for X/Twitter. It omits the viewership experiment, which used TikTok viewers only, and it keeps respondents from the user survey who logged in with their X/Twitter account and respondents from the viewership history panel who reported using X/Twitter regularly in a followup survey run after their participation in the emotion sampling and history donation.

Appendix Table A5: Characteristics of emotion-sampling and viewership history sample

	Screened Mean (SD)	Invited Mean (SD)	Registered for App Mean (SD)	Completed Mean (SD)
<i>Demographics</i>				
Male	0.30 (0.46)	0.20 (0.40)	0.20 (0.40)	0.21 (0.41)
Age	25.20 (3.28)	24.95 (3.27)	25.13 (3.36)	25.14 (3.29)
Bachelor's Degree	0.57 (0.50)	0.55 (0.50)	0.58 (0.49)	0.60 (0.49)
Employed	0.80 (0.40)	0.82 (0.38)	0.79 (0.40)	0.77 (0.42)
White	0.67 (0.47)	0.66 (0.47)	0.67 (0.47)	0.64 (0.48)
Black	0.13 (0.34)	0.17 (0.37)	0.16 (0.36)	0.17 (0.38)
Income (\$)	46,735 (40,141)	45,935 (39,222)	44,896 (38,750)	45,992 (41,177)
Financially Distressed (Fed Measure)	0.45 (0.50)	0.52 (0.50)	0.51 (0.50)	0.49 (0.50)
Democrat	0.53 (0.50)	0.57 (0.50)	0.61 (0.49)	0.64 (0.48)
Republican	0.20 (0.40)	0.20 (0.40)	0.19 (0.40)	0.18 (0.39)
<i>Social Media Usage</i>				
Uses TikTok Regularly	0.72 (0.45)	1.00 (0.06)	0.99 (0.11)	0.99 (0.10)
Uses YouTube Regularly	0.83 (0.38)	0.81 (0.39)	0.85 (0.35)	0.87 (0.34)
Uses Instagram Regularly	0.81 (0.39)	0.86 (0.35)	0.87 (0.33)	0.87 (0.33)
<i>Social Media Attitudes</i>				
Prefer TikTok Didn't Exist	0.53 (0.50)	0.39 (0.49)	0.39 (0.49)	0.37 (0.48)
Prefer YouTube Didn't Exist	0.12 (0.33)	0.13 (0.34)	0.14 (0.34)	0.14 (0.35)
Prefer Instagram Didn't Exist	0.53 (0.50)	0.52 (0.50)	0.52 (0.50)	0.50 (0.50)
<i>Wellbeing</i>				
Subjective Wellbeing (1-3)	1.81 (0.60)	1.81 (0.57)	1.82 (0.56)	1.78 (0.59)
Anxiety (0-6, PHQ-4)	2.42 (1.87)	2.61 (1.87)	2.53 (1.87)	2.51 (1.82)
Depression (0-6, PHQ-4)	1.91 (1.84)	2.03 (1.87)	2.02 (1.89)	2.02 (1.89)
Subjective Percentile Rank Among Peers	53.27 (20.71)	52.12 (20.45)	51.53 (20.65)	50.04 (20.75)
Observations	6,529	2,513	673	383

Note: This table displays descriptive statistics of our screened and recruited sample throughout the attrition process. All outcomes are collected in the screening survey. The first column shows all screening survey respondents; the second shows screened respondents invited to the onboarding survey; the third shows respondents who completed the onboarding survey and registered on the app; the fourth shows respondents for whom we have cleaned social media viewership histories.

Appendix Table A6: Characteristics of viewership histories

	TikTok			YouTube			Instagram		
	Mean	p10	p90	Mean	p10	p90	Mean	p10	p90
<i>Sessions</i>									
Sessions per day	9.6	3.1	16.8	2.4	0.2	5.5	9.8	1.7	19.1
Minutes per day	100.4	26.2	189.2	72.4	1.3	178.5	18.3	2.1	39.9
Minutes per session	10.6	0.2	28.2	24.1	0.8	63.1	1.8	0.2	5.2
Videos or posts per session	37.2	1.0	96.0	7.9	1.0	17.0	8.4	1.0	19.0
Searches per session	0.75	0.00	2.00	3.32	0.00	9.00	0.03	0.00	0.00
% of videos or posts right after a search	2.4			8.5			0.3		
<i>Videos and posts</i>									
Video length (seconds)	52	7	121	1,617	29	3,722	–		
Watch time (seconds)	22	1	59	446	8	1,401	–		
Share of video watched	0.51	0.02	1.00	0.61	0.02	1.00	–		
% watched to the end	31.2			41.4			–		
% Shorts	–			20.4			–		
% ads	19.2			32.3			14.2		
% of watch time on ads	8.9			4.0			21.9		
% paid partnerships (non-ads)	0.3			6.1			–		
People	369			271			324		

Note: This table uses the viewership histories from participants who participated in our emotion-sampling and history-donation data collection. It shows descriptive statistics for the “usage sessions” in these histories — meaning contiguous lengths of time where a participant was watching videos on the platform, as inferred from the sequences of timestamps of video viewing — and for the videos appearing in these histories.

Appendix Table A7: Characteristics of user survey and creator survey samples

	User Survey	Creator Survey	
		Unweighted	Weighted
<i>Demographics</i>			
Male	0.35 (0.48)	0.25 (0.44)	0.33 (0.47)
Age	24.39 (3.83)	25.66 (5.06)	27.21 (6.95)
Bachelor’s Degree	0.39 (0.49)	0.48 (0.50)	0.52 (0.50)
Employed Full-Time	0.37 (0.48)	0.38 (0.49)	0.36 (0.48)
Employed Part-Time	0.15 (0.36)	0.11 (0.32)	0.10 (0.31)
Student	0.18 (0.38)	0.14 (0.35)	0.13 (0.33)
Not Employed	0.13 (0.33)	0.07 (0.26)	0.06 (0.23)
Content Creation Is Main Work	0.02 (0.13)	0.09 (0.29)	0.15 (0.36)
<i>Politics</i>			
Liberal	0.62 (0.49)	0.62 (0.49)	0.64 (0.48)
Conservative	0.16 (0.36)	0.15 (0.36)	0.14 (0.34)
Ideological Extremity (0-3)	1.56 (1.08)	1.59 (1.10)	1.61 (1.11)
Favorability Toward Democrats (0-100)	54.63 (25.41)	55.82 (25.93)	54.87 (26.91)
Favorability Toward Republicans (0-100)	25.10 (26.14)	24.64 (27.31)	24.44 (27.38)
Often Angry About Politics (1-5)	3.58 (1.07)	3.64 (1.12)	3.61 (1.11)
Often Anxious or Sad About Politics (1-5)	3.59 (1.08)	3.63 (1.15)	3.55 (1.15)
People Like Me Lack a Political Voice (1-5)	2.62 (1.09)	2.62 (1.12)	2.55 (1.13)
<i>Self-Image</i>			
Rank Among Peers (1-10 Ladder)	5.35 (1.77)	5.50 (1.81)	5.38 (1.84)
Others Consider Me Attractive (1-5)	3.32 (0.98)	3.54 (0.99)	3.47 (0.99)
Others Consider My Lifestyle Desirable (1-5)	3.01 (1.08)	3.25 (1.12)	3.38 (1.10)
<i>Social Media Use</i>			
Logged In with TikTok	0.77 (0.42)	–	–
Logged In with X/Twitter	0.49 (0.50)	–	–
Regularly Uses Instagram	0.84 (0.37)	–	–
1+ Hours/Day on TikTok	0.72 (0.45)	–	–
1+ Hours/Day on X/Twitter	0.30 (0.46)	–	–
1+ Hours/Day on Instagram	0.58 (0.49)	–	–
1+ Hours/Day on YouTube	0.68 (0.47)	–	–
<i>Content Creation</i>			
Has Posted on TikTok	0.62 (0.48)	0.99 (0.09)	1.00 (0.03)
TikTok Followers, Median [p10, p90]	82 [10, 1,308]	4,119 [288, 217,700]	20,200 [638, 1,300,000]
TikTok Posts, Median [p10, p90]	3 [0, 103]	224 [20, 1,429]	445 [33, 2,525]
Observations	3,490	923	752

Note: This table shows descriptive statistics for participants in our user survey and creator survey. For the latter, we show statistics separately when weighting all respondents equally versus when using TikTok respondents and weighting them to match the distribution of follower counts in the TikTok histories of our participants. Follower and post counts are scraped directly from their accounts; all other outcomes are survey outcomes.

C.2 Content classification

We classify posts and videos using GPT-5-mini. For TikTok videos, we give the model the video’s transcript, caption, creator’s name and bio, text that appears on-screen, and the video’s thumbnail image. For Instagram posts or Reels, we give the model the post’s caption, hashtags, creator username and bio, location and music tags, and an image (either the video’s thumbnail or the first image in an image post). For YouTube videos, we give the model the video’s title, description, transcript, and category, the channel’s name, and a thumbnail image.

The model assigns each video one category out of the seventeen mutually exclusive lower-level categories displayed in Table A8. The bold lines are then our aggregation of those classifications into higher-level categories. In addition to assigning a category, the model classifies whether a video is “enviable” (meaning prominently featuring an unusually attractive or wealthy non-celebrity) and tags any news topics the video discusses. A follow-up classification step applied to political videos classifies whether the video expresses a clear partisan perspective.

Appendix Table A8: Content types and the style categories they contain

Content type and style category	Share of videos (%)
Lifestyle	28.7
Appearance	15.5
Wealth	1.2
Day in the life	3.6
Music, dance, or performance	6.2
Travel or outdoors	2.3
Stories, gossip, and advice	23.7
Advice or product recommendation	10.3
Storytime or gossip	13.4
Entertainment	38.6
Meme, joke, or comedy sketch	14.0
Video gaming	1.8
Pop culture, movies, or film	6.7
Sports	2.3
Food or cooking	5.5
Pets or animals	4.3
DIY, crafts, or art	2.2
Tech, science, or education	2.0
Politics	4.1
Political news or discussion	3.4
Other styles flagged as political	0.7
Not assigned to a content type	4.8
Other	4.8
Videos	498,890


Notes: This table lists our four broad content types and the underlying specific content types from our schema that are aggregated into these broad types. Our LLM classification assigns each video the one specific content type that best fits it. The percentages show the shares of each content type in participants' viewership histories.

Appendix Table A9: Agreement between human coders and the LLM classifier

	Videos	Human vs. LLM	Human vs. human
<i>Panel A: All test videos</i>			
Broad content type agreement	992	53%	52%
Broad content type κ	992	0.39	0.37
Specific content type agreement	992	42%	40%
<i>Panel B: Agreement by the LLM's broad and specific content type</i>			
Lifestyle	240	56%	53%
Appearance	79	41%	34%
Wealth	34	43%	34%
Day in the life	41	48%	34%
Music, dance, or performance	49	11%	31%
Travel or outdoors	37	56%	44%
Stories, gossip, and advice	236	34%	43%
Advice or product recommendation	101	31%	31%
Storytime or gossip	135	27%	33%
Entertainment	240	67%	56%
Meme, joke, or comedy sketch	30	41%	38%
Video gaming	30	73%	67%
Pop culture, movies, or film	30	48%	29%
Sports	30	64%	58%
Food or cooking	30	56%	45%
Pets or animals	30	35%	28%
DIY, crafts, or art	30	51%	40%
Tech, science, or education	30	31%	32%
Politics	237	60%	54%
Political news or discussion	196	64%	56%
Other styles flagged as political	41	41%	47%
Not assigned to a content type	39	15%	59%
Other	39	15%	59%

Notes: This table shows agreement rates between our LLM and Prolific-recruited human handcoders. The “human vs LLM” column shows the percentage of the time (across videos) that a human’s pick matches the LLM’s pick, while “human vs human” is the equivalent for two humans coding the same video. κ is Cohen’s kappa.

Appendix Figure A18: Examples of TikTok content types



These next figures report results of a sparse auto-encoder applied to (i) our viewership experiment data, (ii) our production-vs-consumption observational data (Figure 6), and (iii) our discontinuities-at-60-seconds data (Figure 7).


We prepare the underlying data by taking each video's caption, on-screen text, and transcript and adding a description of its imagery obtained by having GPT-5-mini describe four frames from the video (experimental videos, for which we have MP4s) or the thumbnail of the video (observational videos). We use the Gemma Scope sparse auto-encoder on the Gemma-2-2B language model to break videos into concepts.

For the experimental data, we regress engagement (residualized on the fixed effects from our main specifications) on each concept and keep those that predict engagement after the multiple-testing correction

of Carlson (2026). We use Claude Opus 5.5 to label concepts based on the videos they appear in and group concepts into themes. We then check which concepts are predictive of which content types.


For the observational data, we predict whether a video is consumed versus produced or whether a video is just over versus just under 60 seconds long, and then follow a similar procedure.

Appendix Figure A19: Concepts most predictive of each content type




Notes: This figure shows, for each of our four content types, the concepts discovered by the sparse auto-encoder whose presence in a video most strongly predicts this content type. The sample consists of videos from our viewership experiment and we retain concepts that predict engagement at the 5% level after the multiple testing correction of Carlson (2026) and that appear in at least 2% and under 50% of videos. Claude groups concepts into themes and the coefficients here show themes, with the label of a theme being the concept in the theme that is most predictive of the relevant content type.

Appendix Figure A20: Concepts most predictive of engagement




Notes: This figure shows the concepts discovered by the sparse auto-encoder that are most predictive of engagement (specifically, those with the largest positive and largest negative t -statistics). The x-axis shows coefficients from regressions of experimental engagement (the share of viewers seeing a video who watched it for at least 10 seconds) on an indicator for the relevant concept being present in the video, using our main engagement specification.

Appendix Figure A21: Clustering of discovered concepts




Notes: This figure shows clustering of the concepts from Appendix Figure A19 in a t-SNE plot (van der Maaten and Hinton, 2008), which positions concepts so that concepts that tend to appear in the same videos are placed near each other. The cardinal distance between clusters is not meaningful and the axes have no units or interpretation.

Appendix Figure A22: Clustering of discovered concepts: principal components




Notes: This figure replicates Appendix Figure A21 using principal components as the axes. They are the first two principal components of a matrix where a row is a concept, a column is a video, and the entries record whether the concept appears in the video. That is, they capture the two principal components of the pattern of concepts' appearances in videos.

Appendix Figure A23: Concepts most predictive of a video being consumed rather than produced



Notes: This figure is analogous to Appendix Figure A20, but uses results from a sparse auto-encoder applied to our observational consumption vs. production data and shows which concepts are most predictive of whether a video appears in the “consumed” set as opposed to the “produced” set.

Appendix Figure A24: Clustering of concepts predicting consumption rather than production




Notes: This figure replicates Appendix Figure A21, but uses results from our sparse auto-encoder applied to our observational “produced versus consumed videos” data.

Appendix Figure A25: Concepts most predictive of a video lasting just over rather than just under 60 seconds


(a) More common at 60-65 seconds

(b) More common at 55-59 seconds



Notes: This figure is analogous to Appendix Figure A20, but uses results from a sparse auto-encoder applied to our observational video-duration data and shows which concepts are most predictive of whether a video is just under or just over 60 seconds long.

Appendix Figure A26: Clustering of concepts predicting video length around 60 seconds



Notes: This figure replicates Appendix Figure A24, but uses concepts from Appendix Figure A25.

D Identifying Algorithmic Markups

This section tests for algorithmic markups following the ideas of Kleinberg et al. (2018) and Rambachan (2024). We take the pool of videos in our participants’ TikTok viewership histories, encompassing 20.6 million total video appearances, and train a prediction model. The inputs are video metadata (length, age when shown, time of day, position in scrolling session, whether it is an ad, caption, transcript, on-screen text, creator username) and viewer information (aggregate engagement rate, earlier engagement with same creator, similarity of video’s text embeddings to earlier videos in the session, viewer-by-content interactions). The model predicts for how many seconds a viewer watches a video. We generate embeddings using `gte-multilingual-base` fine-tuned using contrastive learning on 20,000 TikToks. The model achieves an out-of-sample R^2 of 0.187 for “number of seconds watched” and 0.406 for “share of a video watched.”

We train the model on an 80% sample and hold out 20% of the sample for evaluation. In the figure below, we split videos from the evaluation set falling into different categories by ventiles of predicted engagement according to the model and plot their actual engagement on the y-axis.


The intuition is that if the platform does not apply algorithmic markups—it equates engagement with the marginal video of each type—the lines should intersect at the bottom ventile. If one line remains persistently elevated above the other, this indicates the line below it is algorithmically subsidized—the platform serves

videos of that type even when the marginal video it serves will predictably receive lower engagement than the marginal video of another type.

Panel (a) shows that advertisements are algorithmically subsidized relative to organic videos according to this measure, which makes sense as advertisements definitionally are videos a creator or sponsor has paid to subsidize algorithmically. Panel (b) shows that non-advertisement videos with third-party sponsorship tags are also subsidized algorithmically. 80% of these are videos with TikTok Shop affiliate links; TikTok takes a direct cut on these sales, so it makes sense for those videos to be algorithmically subsidized. Figure A28 shows that it is indeed TikTok Shop videos that drive this subsidization, and videos sponsored by third parties that do *not* link to the TikTok Shop are not subsidized at all.


Panels (c) and (d) show that, by contrast, none of the content types we are interested in appear to be algorithmically subsidized relative to the others: the bottom ventiles' confidence intervals all cross.

Appendix Figure A27: Testing for algorithmic markups: actual versus predicted engagement



Notes: This figure plots data from our participants' viewership histories. In each panel, the x-axis plots ventiles of predicted engagement. The y-axis plots actual engagement by viewers with videos in this ventile, residualized on participant-by-video-length fixed effects. Shaded bands are 95% confidence intervals clustered by participant.


Appendix Figure A28: Testing for algorithmic markups: TikTok Shop versus third-party sponsored videos



Notes: This figure replicates Panel (b) of Figure A27 but splits videos with commercial disclosures by whether the disclosure is a link to a TikTok Shop product or a different kind of third-party sponsorship.


E Creator Survey, Monetization, and Motivations

Appendix Figure A29: What fraction of the videos viewers see come from creators who ever earn money?



Note: This figure replicates Figure 5 Panel (a), except the green and purple bars—instead of plotting indicators for whether the focal *video* earns money for its creator—plot indicators for whether the video’s creator appears to *ever* earn money on TikTok. We draw a representative sample of creators who appear in viewers’ feeds and scrape all videos they posted from July to September 2025. We estimate their Creator Rewards wage rate in each month by calculating their earnings from videos at least a minute long (imputed from its view count and public information about Creator Rewards’ rates) and dividing this by the time spent on all their videos. The latter is calculated using responses to a question in our creator survey about how long it takes to create videos, from creators in the same follower count bin. Reported third-party sponsorship numbers are based on sponsorship, paid-partnership, or advertising tags scraped from the creator’s videos. Unreported third-party sponsorship numbers are based on responses to a question in our creator survey about third-party income sources that do not show up in official tags, such as unreported sponsorships or videos made and posted on a sponsor’s account rather than the creator’s own. We use average responses from creator-survey responses in the same follower bin. The orange bars aggregate over the preceding three bars.

Appendix Figure A30: Creators who earn money: scraped posting histories versus the creator survey



Note: The orange bars in this figure replicate those in Appendix Figure A29. The red bars plot the share of TikTok creators in our creator survey with the same follower count who say they made money or received free products from posting on TikTok in the last three months. We reweight survey respondents so that within each follower count bin, the share with an email address in their bio matches the share of feed occurrences coming from creators with an email address in their bio.

Appendix Table A10: Creator Survey respondents by follower count

	All	Recruited Through			Platform	
		Email	Prolific	TikTok	Instagram	X/Twitter
Under 1,000	332	10	308	235	45	52
1,000–9,999	329	61	264	231	56	42
10,000–99,999	176	143	32	167	7	2
100,000–999,999	65	62	3	63	2	0
1,000,000 or More	25	25	0	25	0	0
Respondents	927	301	607	721	110	96

Note: This table shows the distribution of follower counts of respondents to our creator survey, separately by whether they were recruited via Prolific or email and separately by the primary platform they post on.

Appendix Table A11: Characteristics of Creator Survey respondents

	All	Recruited Through			Platform	
		Email	Prolific	TikTok	Instagram	X/Twitter
<i>Demographics</i>						
Male	0.25 (0.43)	0.29 (0.45)	0.22 (0.41)	0.24 (0.43)	0.19 (0.39)	0.33 (0.47)
Age	25.61 (4.96)	27.27 (6.99)	24.89 (3.41)	25.67 (5.31)	25.18 (3.87)	25.64 (3.14)
Bachelor’s Degree	0.49 (0.50)	0.65 (0.48)	0.42 (0.49)	0.49 (0.50)	0.59 (0.49)	0.43 (0.50)
Employed Full-Time	0.39 (0.49)	0.36 (0.48)	0.40 (0.49)	0.38 (0.49)	0.43 (0.50)	0.42 (0.50)
Employed Part-Time	0.12 (0.33)	0.10 (0.30)	0.14 (0.34)	0.11 (0.32)	0.18 (0.39)	0.12 (0.33)
Student	0.15 (0.36)	0.16 (0.37)	0.14 (0.35)	0.15 (0.35)	0.21 (0.41)	0.12 (0.33)
Not Employed	0.07 (0.25)	0.03 (0.16)	0.09 (0.29)	0.07 (0.25)	0.03 (0.16)	0.15 (0.35)
Content Creation Is Main Work	0.08 (0.27)	0.19 (0.39)	0.02 (0.16)	0.10 (0.29)	0.03 (0.16)	0.01 (0.10)
<i>Politics</i>						
Liberal	0.62 (0.49)	0.64 (0.48)	0.60 (0.49)	0.61 (0.49)	0.56 (0.50)	0.75 (0.44)
Conservative	0.17 (0.37)	0.13 (0.34)	0.18 (0.38)	0.16 (0.36)	0.26 (0.44)	0.12 (0.32)
Ideological Extremity (0-3)	1.60 (1.10)	1.56 (1.09)	1.62 (1.11)	1.53 (1.09)	1.70 (1.06)	2.05 (1.06)
Often Angry About Politics (1-5)	3.63 (1.10)	3.64 (1.05)	3.65 (1.11)	3.62 (1.10)	3.60 (1.00)	3.75 (1.17)
Often Anxious or Sad About Politics (1-5)	3.64 (1.13)	3.64 (1.09)	3.65 (1.14)	3.63 (1.13)	3.65 (1.03)	3.71 (1.24)
People Like Me Lack a Political Voice (1-5)	2.58 (1.11)	2.50 (1.12)	2.63 (1.10)	2.61 (1.11)	2.31 (1.00)	2.68 (1.18)
<i>Self-Image</i>						
Rank Among Peers (1-10 Ladder)	5.47 (1.78)	5.85 (1.75)	5.36 (1.78)	5.58 (1.78)	5.31 (1.57)	4.93 (1.91)
Others Consider Me Attractive (1-5)	3.53 (0.98)	3.74 (0.92)	3.44 (0.99)	3.59 (0.96)	3.51 (0.90)	3.16 (1.14)
Others Consider My Lifestyle Desirable (1-5)	3.27 (1.10)	3.66 (0.98)	3.08 (1.10)	3.33 (1.08)	3.39 (0.93)	2.68 (1.21)
Respondents	927	301	607	721	110	96

Note: This table reports descriptive statistics for our creator survey respondents, weighting all responses equally.

Appendix Table A12: Creator Survey answers by randomized gift-card offer

	\$5 Offer	\$10 Offer	Difference	(SE)	p-Value
<i>Response to the email (% of emails sent)</i>					
Answered the survey	1.40	2.17	0.77	(0.20)	<0.01
<i>How I choose what to create (points out of 100)</i>					
What I enjoy creating	51.6	50.8	1.8	(4.1)	0.66
What will maximize views	29.8	26.8	-3.2	(3.0)	0.29
What is most monetizable	18.6	22.3	1.5	(3.0)	0.63
<i>Rewards I get from creating (% selecting)</i>					
Entertaining others	88.9	90.1	-0.8	(4.9)	0.88
Social connection	82.9	84.9	4.1	(10.6)	0.70
Creative fulfillment	79.5	80.7	-0.4	(6.3)	0.95
Expressive satisfaction	60.7	61.9	6.5	(7.9)	0.41
Social status	45.3	43.1	-11.7	(7.7)	0.13
Persuasion	45.3	40.9	-5.3	(8.0)	0.51
Respondents	117	181			
Emails sent	8,328	8,329			
Joint test, all answers			$\chi^2(8) = 8.0$		0.43

Note: This table uses randomized incentives to check for selection on unobservables into responding to our emailed creator survey invitations (Dutz et al., 2026). We randomized whether each email offered a \$5 or \$10 Amazon gift card for completion. The first row of the table shows a first-stage effect of this randomization on response rates. The subsequent rows show no difference in average responses to the survey questions from creators in the two conditions, suggesting no selection into responding on unobservables correlated with the characteristics elicited in the survey. The first two columns report unweighted average answers among respondents offered each amount. The third column reports the coefficient on the \$10 offer from a regression of the answer on an indicator for being offered \$10 and email-wave-by-follower-decile fixed effects.

Social desirability. Panel (b) of Figure 5 plots responses to a question in our creator survey about which considerations influence what content creators choose to make. Creators are given 3 options and asked to allocate 100 points between them in proportion to the importance of those considerations. To gauge the extent to which answers reflect social desirability, we randomize the wording of the three options between two versions intended to flip their perceived social desirability.

In one version, the question and answer options are:

Consider the following 3 extreme statements:

1. I only care about the enjoyment I get from creating the kinds of content I love and people seeing it. I wouldn't switch to making other types of content that I loved creating less, even if this meant more views, followers, or monetization.
2. I only care about getting views, followers, or view-based creator rewards payments, regardless of what types of content I'm creating. I wouldn't switch to content that would result in fewer views or followers, even if I enjoyed creating it more or could get more third-party monetization.
3. I only care about being able to monetize my content through sponsorships, affiliate links, or other types of sales, and I wouldn't switch to making types of content that would make

this harder to do, even if I enjoyed creating it more or could get more views or followers.


In the flipped version, the three statements are instead:

1. I make content mainly to please myself, and I don't pay much attention to whether they're what people want to read or watch. I wouldn't switch to making other types of content that I'd enjoy less, even if this meant more views, followers, or monetization.
2. What matters most to me is reaching and entertaining as large an audience as I can—views, followers, and view-based creator rewards. I wouldn't switch to content that would result in fewer views or followers, even if I enjoyed creating it more or could get more third-party monetization.
3. I treat this as a real job and focus on building third-party income I can support myself and my family on—through sponsorships, affiliate links, or other sales. I wouldn't switch to content that would make this harder, even if I enjoyed creating it more or could get more views or followers.

In both versions, creators then read: “Most people fall somewhere between these extremes. If you had to allocate 100 points indicating how much you agree with each statement, how would you do it? For example, allocating 100 points to one statement means you fully agree with that statement and do not agree at all with the others. Allocating 50 points to each of two statements means you agree with those two statements equally, and don't agree at all with the third.”

Figure A31 plots results from a supplementary survey of 140 Prolific users. Each respondent was shown one of the two question versions and asked to rank the statements according to “which one makes a creator agreeing with it look best in other people's eyes.” The figure shows that our manipulation successfully flips the perceived social desirability of the nonpecuniary “enjoyment” reward compared to the pecuniary views/monetization rewards.


Appendix Figure A31: Perceived social desirability of the creator survey's content-choice statements



F Quotes from Qualitative Interviews

The tables below display quotes from qualitative interviews, conducted retrospectively by an AI interviewer, with participants in our user survey and creator experiment who have posted at least 3 videos on TikTok. Appendix Figure A32 shows the sequence of questions asked by the AI.

Appendix Figure A32: Flow of the qualitative interviews



Appendix Figure A33: Excerpts of qualitative interviews as they appear to participants

(a) Answering an open-ended question



Northwestern

Why are you always aiming for more views and likes?

I like the possibility of potentially making some extra income off my Tiktok

Can you tell me more about how more views and likes connect to making extra income for you?

Well if you get more views and likes, eventually your follower count also increases along with it, albeit at a much slower rate. This also makes you eligible for monetization and the creator fund.

Let's move on to something different. Is this kind of satisfaction specific to the type of content you make? Or could you get similar satisfaction from making very different types of content? Why?

I think I could get a similar level of satisfaction doing other types of content, but this is what I know I'm best at so I would rather stick with it.

Send

(b) Answering a multiple-choice question



Northwestern

What is your main strategy for trying to grow it? Posting a lot, or posting specific types of content you think are likely to do well or go viral?

A bit of both. Ive had four "viral" videos and all four are completely different

Can you walk me through what you actually do when you're trying to grow your account?

I will see what is performing well on others accounts, what does well on my account and what is trending overall

Here's a different question. Is hoping to make it big someday (becoming a well-known creator, or making a living from posting) also a major reason you post? Or do you just enjoy posting, even if you never become really big?

CHOOSE ONE

☐ I just enjoy posting, even if I never become really big

☐ Making it big is also a major reason I post

Notes: This figure shows screenshots of the interview interface from the perspective of a participant.

Appendix Table A13: Satisfaction from producing a video

Quote	Followers
<i>Making something and watching it come together</i>	
"I enjoy the creating process of the videos and the comedy of adding specific sounds to funny videos. [...] I added a sound of "get it get it get it" over a video of my siblings dancing and it was funny [...] I thought the way the sound lined up with the movement was funny"	30
"i enjoy being able to create something for entertainment and being able to see the outcome [...] i liked being able to plan out my day and figure out what parts i wanted to film and put together in my video to be able to show the variation of events happening and then adding aethetic edits and music over lay [...] it was fun being able to create something enjoyable to watch"	120
"I like having an idea and impliementing it. It is the creative process I enjoy [...] I had an idea for a tiktok I wanted to do so I spent time filming all the parts and then edited them all together [...] I enjoyed filming the most as I got to experience the idea as I filmed. Editing can become choresome"	150

(continued)

Quote	Followers
“I like editing content in a funny way and making a video look the best it can, it feels fulfilling. [...] I was editing a concert video and wanted to include some captions showing my inner feelings during a performance, and I really liked how the final product came out. [...] It felt like I was making something that my community could connect with and find great entertainment from.”	320
“The post-production side [...] Synchronizing the music to combat or emotional scenes [...] Being able to fully convey the experience I curated through footage repurposing and synchronizing”	410
“I just like the creative process and the creative flow that I have when I’m creating post. It just allows me to really show my personality and in my interest in the things that I like to see other people getting engaged in their similar interest as well. [...] I enjoy putting it together the music lyrics and then also the video footage you know I tied in you know as I was watching the game and seeing the play develop you know I thought about that song and I kind of put the two together and posted it [...] It just feel natural and exciting like I’m The novel person I am responsible for the creation”	500
“I enjoy the process of filming and editing and seeing a good video come together that I enjoy watching, like a baking video with the finished product at the end [...] I enjoy setting up the camera to get different angles of the ingredients and the process, then taking the final shots of the baked good. I also enjoy editing and feeling like my video will be entertaining or helpful to others [...] It feels satisfying to cut the clips down and see the raw footage transform into a finished video”	650
“What I enjoy most is being able to edit the video such as the thumbnails and audio [...] I ended up changing the design and adding special effects in the thumbnail through Canva and other editing apps. What I enjoyed the most about it is being able to test out many templates.”	5,200
“The editing, creation, finding the inspiration and sounds etc [...] I enjoyed the process of sourcing the footage, finding the trend and audio and the process of perfecting the hashtags and the caption [...] It was a professional post, so getting the caption and hashtags right was necessary”	11,000
“The creation process and seeing the finished product. So doing the transition and editing it and seeing it come together. [...] I did a makeup transition and filming the transition was fun. I enjoyed watching the video as it developed [...] Doing the makeup itself was fun, and watching myself go from myself to made up”	16,000
<i>Expressing themselves</i>	
“I enjoy being able to express myself freely and display my personality [...] I enjoyed filming the angles of traveling and all the get ready with men’s for the day. I really enjoy just sharing glimpses of my lifestyle [...] It feels like producing something I am proud of”	2,500
<i>Absorption and escape</i>	
“it engages the creative part of my mind. for example thinking of new ideas, recording [...] I found the famous cat, and spent ages making the correct look/making it look funnier (on my face) [...] basically just focusing on anything but real life - escapism”	150

(continued)

Quote	Followers
<i>Reliving and capturing experiences</i>	
“It is cool to put something up that I did and can look back on [...] I enjoyed trying out some new ideas about how the video could go together with the music [...] It felt satisfying”	65
“I don’t really care who views my posts. I mainly use these apps as an archive or just to put it somewhere [...] I might end up losing it or breaking the device containing the original footage [...] It was ok, I always like to appreciate others but sometimes I feel like I like my solitude [...] Felt like someone viewed something that I did.”	65
“It is fun making videos to look back on as memories”	210
“I enjoy making travel related content, which I share about my travel journey. when I making it I can relive with it [...] I post some pictures of trip I took in [a region]. I add notes about my experiences with each photos. My view about places I visit, my opinion about food I had [...] I felt the excitement that I had when I am experiencing those moments. Then , after post it many people had similar experiences”	260
“I do it for myself anyway. I maintain it as my memory book. I can log in to it in future and can enjoy it. I don’t have large follower base anyway. [...] It make me happy little bit and feel like small accomplishment [...] I put some effort to create it. So, when someone enjoy it or use it, its give me joy”	260
“I love trying the trends that go around on booktok , for example “books that define my reading taste” [...] Putting all the pictures of the book covers on the post as it reminds me of how many books I really have read and absolutely loved! [...] That I’m just immersed in a world of books and fictional stories and it really relaxes me!”	1,500
“I post to document my life and also interact with viewers”	2,500
<i>Looking and feeling good</i>	
“Creating something that I can share with others [...] Feeling sassy and hot gave me empowerment [...] Just looked cute and lip synced to a song I like”	39
“I just like looking good and I like when people agree [...] Cause I looked so good that day just did my hair and also it was my birthday [...] Just getting dressed up”	99

Notes: This table shows answers to “What specifically do you enjoy about making a post? Please be specific and give some examples.” Quotes are verbatim and “[...]” indicates places where the AI interviewer asked a follow-up question; answers to successive follow-ups are joined into one quote.

Appendix Table A14: Satisfaction from having people see one’s posts

Quote	Followers
<i>I don’t care about this</i>	
“nothing [...] i didn’t really consider others when i posted [...] a reality tv stars husband responded to my comment with a video [...] it was pretty cool. i started thinking about the six degrees of separation”	24

(continued)

Quote	Followers
“I don’t really care who views my posts. I mainly use these apps as an archive or just to put it somewhere [...] I might end up losing it or breaking the device containing the original footage [...] It was ok, I always like to appreciate others but sometimes I feel like I like my solitude [...] Felt like someone viewed something that I did.”	65
“i don’t really pay much attention to who sees my post, i post for fun not views [...] it’s nice to see someone like or interact with the fun videos [...] nice that someone enjoyed the video as much as me”	120
“I don’t really get anything from it. I do enjoy the occasional compliment. However, I mainly enjoy going back and watching them myself. [...] I enjoyed how it made me feel [...] It made me feel happy and appreciated.”	150
“nothing [...] same as it was before i posted it. makes no difference to me [...] it doesn’t effect me”	540
<i>Connection and small-scale, two-way interaction</i>	
“I want to be a gamer content creator online so it’s nice to bring the gaming community together [...] I don’t have a lot of gamer friends so I’m trying to be more in the community [...] Enjoyment. I was happy that people had fun with my content [...] Their responses were positive and engaging”	25
“not much just a small bit of connection [...] a like or a comment made me feel connected to the viewer [...] I got one just hyping up the video it made me feel happy”	30
“I get a community that likes the same things I like [...] People who finds what I post funny and are ready for more post from me [...] It was fun because I got to see so many people happy about it”	43
“I don’t get much out of anything. Just more followers and likes. I don’t care about those things but just putting something out someone could relate to, and maybe become mutuals. [...] Friendship [...] Having someone that has something in common with you. Having someone you could talk to about things only you understand is cool. [...] Honestly, the ice video. A girl commented and said she loves chewing ice too. She followed me, I followed her back, and we instantly connected and started chatting.”	54
“make new friends [...] they found my post funny or relatable and we would meet based on those interests [...] it makes me feel happy because it’s like people like me”	150
“I am able to interact with others in the comments section, either feeling happy when they express that I helped them or sometimes they offer me advice too that helps me for the future [...] I posted a video learning how to use a sewing machine and someone commented with a lot of good tips and advice for how to better use it to it’s fullest possibilities [...] It felt good and I was happy they chose to take the time to share that with me”	650
“I meet new people! [...] Really nice, I like talking to people about what books that they like and therefore we give eachother some book recommendations! [...] That I can have a conversation with someone about something I love and that I am so passionate about as a lot of my family and friends don’t understand why people love to read!”	1,500

(continued)

Quote	Followers
“I get to interact with individuals that see my post which builds a little community. [...] Someone made a funny comment reply to my post and it was hilarious [...] It was just a little laugh to brighten my day and I gained a potential consistent viewer”	2,500
<i>Pride or general validation for making a video people liked</i>	
“i do not get much out of it, but it is nice to watch my likes rise [...] it was really exciting honestly [...] seeing the numbers go up and watching my phone blow up with likes”	150
“I get validation that people enjoy my videos [...] I posted a horror tiktok showing clips and a montage from a film i like it got a lot of views and nice comment [...] It made me feel comfortable [...] It made me proud that other people like my videos”	170
“I do it for myself anyway. I maintain it as my memory book. I can log in to it in future and can enjoy it. I don't have large follower base anyway. [...] It makes me happy little bit and feel like small accomplishment [...] I put some effort to create it. So, when someone enjoys it or uses it, it gives me joy”	260
“I get recognition from a wide audience and it feels good seeing people like my content. [...] Recently I had a video go viral with over half a million views and over 150,000 likes! That was a huge milestone for me because it was my most viewed video by far. I felt super proud that people liked it and found me funny.”	320
“Just self satisfaction Observing the increase of view and likes would also offer an validation to myself [...] Very satisfying to watch [...] Observing the increase of view and likes would also offer an validation to myself [...] It shows that this video is a piece that the public like to watch and want more of it Which is created by myself”	340
“it is nice seeing when you get a larger amount of views than normal, but i just enjoy sharing certain fun things or whatever that i'm doing [...] it was a year ago during [a hurricane], i posted a crazy video and it got a lot of views (for me) and it was pretty exciting! [...] i just don't get many views so it was cool to see that my post reached such a large amount of people [...] kinda because it was about a current event and because the intensity of the video”	450
“Pride and satisfaction [...] A lot of people hearted it and commented how cute my rats were and I was so happy to share my joy with others”	1,100
“I mean I obviously like getting big on tiktok like getting a bunch of likes is the best part for sure [...] it makes me pretty happy because i'm like wow they really think i'm funny”	1,500
<i>Pride or validation for abilities shown in a video</i>	
“No, there is no need for that [...] I do need to get out to have people see my content [...] Sorry I misunderstood the question! I don't get anything out of my content. Mostly are joy to share my car [...] That people would comment and like my posts and appreciating what I have or what I have modeled on my car”	34

(continued)

Quote	Followers
“It boosts my self confidence and confidence in what I offer, and motivates me to keep going and move forward. [...] Yes, when I posted a clip of my gameplay, there was a comment saying that I am a talented person. [...] I felt proud.”	540
“I feel good knowing my creativity is seen [...] Its a feeling of success and pride, because ive spent a long time developing my skills and my creative side”	16,000
<i>Appearance-related social validation</i>	
“Nothing just a confirmation that I look good [...] Yess by commenting [...] Hearts eyes [...] Nothing really I’m use to comments like that”	99
“liking, commenting and giving compliments or asking what products, clothing I might be wearing and being able to be positive and friendly to others [...] I feel like I am giving them assistance and acknowledging their question [...] I felt good because they seemed to like what I was wearing and they could go out and purchase it themself”	210
<i>Making a difference to viewers</i>	
“I get some sort of satisfaction from the fact that i could have possibly inspired them and entertained them, it makes mee happy and feel like ive done a good job [...] Yes, once a viewer said they loved my reaction to getting a kill on call of duty and it really made me feel happy that im enteratining my viewers”	530
“I want them to be able to have good feelings or have the song take them back to a specific time where their life was something that the reminisce about [...] It makes me feel good [...] They dont even have to comment i appreciate a like or even a share or save [...] It means my post made that viewer feel something good”	1,300
“i get a sense of pride and satisfaction from my peers relating and consuming my content [...] yes, i have multiple post that i made during the early stages of my pregnancy that included product reviews about items that i received from medicaid for example and gave insight on how to obtain the different services and items and that helped one young lady specifically outgain a carseat as she couldnt afford it on her own and that made me feel so proud”	1,400

Notes: This table shows answers to “What, if anything, do you get out of having people see your posts? Please be specific.” Quotes are verbatim and “[...]” indicates places where the AI interviewer asked a follow-up question; answers to successive follow-ups are joined into one quote.

Appendix Table A15: Reasons for concave returns to attention

Quote	Followers
<i>Small-scale, two-way interaction</i>	
“I don’t get much out of anything. Just more followers and likes. I don’t care about those things but just putting something out someone could relate to, and maybe become mutuals. [...] Friendship [...] Having someone that has something in common with you. Having someone you could talk to about things only you understand is cool. [...] Honestly, the ice video. A girl commented and said she loves chewing ice too. She followed me, I followed her back, and we instantly connected and started chatting.”	54
“make new friends [...] they found my post funny or relatable and we would meet based on those interests [...] it makes me feel happy because it’s like people like me”	150
“finding likeminded people with similar interests [...] I found other people who like crochet and we talked about our projects and shared tips with one another [...] I find it enjoyable and helpful [...] I learnt quite a lot about different crochet styles and how to do certain techniques”	1,400
“I meet new people! [...] Really nice, I like talking to people about what books that they like and therefore we give eachother some book recommendations! [...] That I can have a conversation with someone about something I love and that I am so passionate about as a lot of my family and friends don’t understand why people love to read!”	1,500
“its exciting seeing your friends interacting with your content [...] it was so much fun we were just bantering and laughing in the comments section [...] its just how we used to interact back in school also its very hilarious [...] My friend was telling me to remember a time when we swore we would be body built and maintain abs even as adults”	1,800
“I get to interact with individuals that see my post which builds a little community. [...] Someone made a funny comment reply to my post and it was hilarious [...] It was just a little laugh to brighten my day and I gained a potential consistent viewer”	2,500
“I have the ability to be able to interact with my followers and join their live streams [...] A recent time I interacted with a follower was when they mentioned me in their story and after that happened. Me and them been playing games together and streaming anime [...] It was a peaceful experience because the person was friendly and caring [...] What made it rewarding was the friendship I gained from that experience”	5,200
“The positivity and good vibes, a wholesome community in which we can discuss what we share interests in. [...] They commented and helped each other relate to a difficult part of a game I was playing, it was very close connected. [...] A sense of joy that people could relate to one another.”	5,600
“I have people who regularly watch my content so I get joy in seeing how they react and comment to my videos [...] Me and my wife recently posted an image trend showing all the places we’ve traveled together and it got a lot of comments from people who wish to have similar travel experiences one day. [...] It made me feel happy that I was able to share my travel experiences with my fans”	70,000
<i>Early viewers matter more than later ones</i>	
“because the first comments and likes shows that the person cared to do it [...] boring”	57

(continued)

Quote	Followers
“posting often with consistency so I can build a following that actually enjoys my videos instead of quick viral hits”	650
“i usually just pay attention to those the most cause thats like the real reach, after that its just people going on my profile to spam me and not liking or commenting because they first saw it and liked it [...] commenting and liking every post [...] it does not feel genuine”	1,400
“I think the first views, likes, and comments are usually from my friends and the people I’m closer to, so I value their comments a lot more than comments from people randomly stumbling across my content. [...] Sometimes I make videos that I’m scared will anger people I’m close to, like where I’m discussing a hot take or an unpopular opinion. I feel satisfied when these people actually agree with what I’m saying, and can see my perspective and where I’m coming from.”	17,000
“Just seeing that people saw it at all and wanted to leave behind a comment right at the beginning shows me that people are genuinely interested. Of course, the more views it gets the more satisfying as well, but those first few people matter the most because they were the ones who saw it before the video might have gone viral [...] The first few comments always seem to be more interested and put more effort into showing how they feel about the video and leaving behind comments that are more thoughtful. Later comments tend to just be random things that don’t matter as much or not much thought was put into them, more of a casual reaction than an actual thoughtful comment”	35,000
<i>Paying attention only shortly after posting</i>	
“because im usually only focused on the views or comments right after I post [...] I was paying attention to who viewed and liked the post [...] I typically lose interest and forget to go back and check because many of my posts dont get many views”	30
“The first few views means the people who initially saw the video liked it and it appealed to them, later on I would have moved past the fact that I posted and no longer care”	39
“Because I often forget that about what I posted much later after the fact and If I see the post isn’t doing well in the beginning, I assume it won’t do well [...] If I am not getting many views [...] The last time I posted I saw the view count continuously rising at a steady pace so I was happy to see that it was reaching people”	150
“It’s the instant gratification, you dont have to wait [...] it was cool to see how fast other people see the stuff i post [...] it gets old fast, i normally am focused on something else as time goes on”	1,100
“Because I’m usually on the app in that moment, when I first post. So, seeing those first views and likes and comments make me happy to engage with others. I usually won’t look at the post again for a couple days [...] I usually notice the views or likes first. [...] It usually will feel either worse or good, depending on how well the video does.”	1,200
“Because it feels somewhat engaging. I don’t normally care how well a post does. [...] I just turned my attention to other things. [...] It feels somewhat exciting but nothing else really.”	1,400

(continued)

Quote	Followers
<i>Knowing that at least some people saw and liked it</i>	
“That at least a handful of people connected with my content. The rest are like extra credit [...] It looks like views. Comments and likes are nice too [...] Gratifying”	51
“It means people is actually seeing the post and is engaging [...] That it was a very popular song that had everyone feeling good”	1,300
“I feel like the first few views, likes, or comments are the most satisfying because it often dictates how the post is going to do in general, whereas the later views and likes are expected because it just follows a general trend/curve. I tend to actually put my phone away and not care about the first few reactions, because it’s scary and I could feel dejected/rejected if it doesn’t do as well as I am hoping [...] It depends, but usually I try to come back with no expectations, so it is always nice seeing some reactions rather than no reactions at all.”	5,200
“Because I already know people have seen my video [...] The excitement of posting it was gone [...] I get a rush from hitting post”	6,200
“I have a big following so I know I’ll get some view regardless, but it’s nice to know that there are people who still care [...] They see it quickly and like and comment [...] I’m still appreciative of any engagement I get on my posts”	41,000

Notes: This table shows answers to “What, if anything, do you get out of having people see your posts? Please be specific” (first group) and “Why are the first few views, likes, or comments the most satisfying? What makes later views and likes less important?” (other groups, except for one answer to “What is your main strategy for trying to grow it?”). Quotes are verbatim and “[...]” indicates places where the AI interviewer asked a follow-up question; answers to successive follow-ups are joined into one quote.

Appendix Table A16: *Is the satisfaction specific to the content one makes?*

Quote	Followers
<i>Yes: specific to the content I make</i>	
“No, it’s different when I post about experiences it’s more a shared satisfaction from people who have the same interests or went to the same event appreciate the post”	39
“It’s typically what I post because it’s my own interest I wouldn’t post anything I don’t find interesting [...] Yes [...] Just trying new things out”	43
“It’s more satisfying cause I actually enjoy gaming [...] It’s something I like so it seems less like an effort”	91
“I could get this with any content as long as it is regarding an topic i like [...] Movies”	170
“I feel like if it was different content but still something I was interested in I would still feel good because im there helping people”	210
“It dependes. I already enjoy travelling eventhough I make it content or not. I don’t like to make joke video, I may not feel the same with making those kinds of contents. However, there can be other interest I prefer to make it as a conten”	260

(continued)

Quote	Followers
“Good question but I believe it’s only on the specific type I made [...] Because it is something I genuinely like to explore and want to share”	340
“it’s specific to the content because that’s what i enjoy doing”	540
“I am not sure, I think it wouldn’t feel the same because the content I make now feels like me and different content that doesn’t feel like me probably wouldn’t feel satisfying to get interactions from”	650
“I think its more related to the specific content that i make because i mainly make content that is aimed at helping people or informing people and thats where i get that satisfaction in return.”	1,400
“I think all types of content could get similar satisfaction, but mine might be differnet because its comedy [...] because its validation on something that is innate like you cant train humor i think”	1,500
“I get satisfaction from all my content because it is something I made and created from my own brain [...] If it was still something relevant to me or my life i would always feel satisfaction. [...] Anything political, mental health and neurodivergence and other stigmatised topics”	16,000
<i>No: similar satisfaction from other content</i>	
“I get satisfaction making different types of content too because content making and sharing is fun [...] I would make GIFs to post on blogs and the internet because I think people enjoy looking at them”	25
“Yes I can get a similar satisfaction by posting something funny and getting reactions”	99
“i would get it for any content. its just fun to see people relating to it”	150
“I think I could get a similar level of satisfaction doing other types of content, but this is what I know I’m best at so I would rather stick with it.”	320
“Any form of content, as long as I enjoy the creative process [...] It must be both challenging and exciting [...] A project that puts your creative abilities to the test, attempting techniques and styles you wouldn’t usually do”	410
“I could probably get similar satisfaction from making other content as I love watching TikTok videos and I spend a lot of time on TikTok and I get inspired from other peoples posts!”	1,500
“I like the process of creating content in general therefore, no matter what the content I think I would enjoy it”	11,000

Notes: This table shows answers to “Is this kind of satisfaction specific to the type of content you make? Or could you get similar satisfaction from making very different types of content? Why?” Quotes are verbatim and “[...]” indicates places where the AI interviewer asked a follow-up question; answers to successive follow-ups are joined into one quote.

Appendix Table A17: Strategies for growing one's account

Quote	Followers
<i>I don't deliberately grow my account</i>	
"wait no i dont deliberately try to [...] ill sometiems try to boost it or something or share it or repost so i get more engagement [...] i repost it to other platforms"	1,500
<i>Through posting volume or general production quality</i>	
"Posting frequently at the right time with the right hashtags [...] I look online for analytics [...] I search how to post at the right time and listen to other popular content creators' advice"	25
"Posting a lot and being consistent [...] Posting everyday at the same time"	43
"posting everyday [...] i just make drafts [...] i recorded videos"	150
"Posting semi-consistent quality media [...] Polished the project as well as my technical skills [...] Instead of intense effects, I refined the experience through more intentional choices"	410
"posting often with consistency so I can build a following that actually enjoys my videos instead of quick viral hits"	650
<i>Through posting specific content types I think will do well</i>	
"Posting specific types of content that I think will do well! [...] Because I have seen it all over booktok and I seen that it went viral so I wanted to try it too! [...] The ones that I think fit well with my account or ones that I enjoy watching so therefore I hope the people who view my posts will also enjoy watching!"	1,500
"My main strategy is to post new videos once every week. I mainly post anime and gaming content because it is more likely to go viral."	5,200
"posting consistently and frequently [...] usually a few posts per day occasionally I get slack and post less, but generally I try to stay consistent [...] Depending on the account, it depends on what viral trends I have seen or what content seems to be performing well"	11,000
"A bit of both. Ive had four "viral" videos and all four are completely different [...] I will see what is performing well on others accounts, what does well on my account and what is trending overall"	16,000

Notes: This table shows answers to "What is your main strategy for trying to grow it? Posting a lot, or posting specific types of content you think are likely to do well or go viral?" Quotes are verbatim and "[...]" indicates places where the AI interviewer asked a follow-up question; answers to successive follow-ups are joined into one quote.

Appendix Table A18: Why participants don't post lifestyle content

Quote	Followers
<i>Privacy and anonymity</i>	
"I don't want to share my face or personal information on the internet"	25
"I prefer to keep anonymous on social media platform [...] like i don't like putting myself on camera"	56
"I like my privacy, feels to invasive to my life"	1,100

(continued)

Quote	Followers
<i>My life isn't interesting or aesthetic enough</i>	
"I dont find my life very interesting or aesthetic [...] It is just typically. I live in an older home in a small town. No fancy clothes or fancy activities. Just normal adult stuff going to work and coming home"	30
"i don't think my life is aesthetic enough [...] it's not "camera worthy" i don't do anything on the day to day that would be interesting to record"	150
"i dont have an exciting life [...] all i do is work [...] i would be giving away personal infromation"	150

Notes: This table shows answers to "Why not?" from participants who said they do not post Life Highlights content (aesthetic or interesting casual moments, activities, belongings, surroundings, or features from one's own life). Quotes are verbatim and "[...]" indicates places where the AI interviewer asked a follow-up question; answers to successive follow-ups are joined into one quote.

Appendix Table A19: Why participants don't post political content

Quote	Ideology	Followers
<i>Avoiding arguments, conflict and abuse</i>		
"I dont feel strongly politacally and do not want to fight with people online about it"	Slightly conservative	30
"don't care for the argument [...] politics makes people argue"	Moderate	120
"i don't want to have arguments about politics. i have my views and i don't really care about what other people think of them"	Liberal	150
"I dont like posting about politics as everyone has their own opinions and it can cause a lot of judgement and arguments online"	Moderate	210
"I don't like getting into politics as I feel like it can be very controversial [...] Everyone has a different option and if you aren't on their side then it can cause a bit of tension!"	Conservative	1,500
"Too much arguing around that topic [...] I'm not real political myself and I don't want the confrontation around it"	Moderate	9,800
"I find politics a very negative topic. Its very polarising in the current climate. I posted one video and was met with a barrage of abuse so it is not even worth the mental strain"	Liberal	16,000
<i>Not interested or not knowledgeable enough</i>		
"I don't care about politics"	Liberal	25
"I'm not really into politics."	Moderate	54
"I just don't want to share my political opinions as I'm not an expert"	Liberal	63
"I don't have very strong political opinions nor would I like to push them upon others"	–	65
"I dont really understand politics i dont get involved in it"	Moderate	170

(continued)

Quote	Ideology	Followers
<i>Keeping politics private or off my page</i>		
“I don’t want my page to be associated with politics”	Liberal	650
“In my lifestyle vlogs I may post myself going to vote but I don’t ever state my stance because I don’t care to open my platform up for debate regarding politics or important topics. I have my views and May state them but it isn’t a platform for me to want people to say their views. My content is my life so I only care to show that”	Slightly liberal	2,500
<i>Worried about future employers</i>		
“im trying to get a job iin the future [...] if you post something that employers don’t agree with youre cooked”	Moderate	1,500

Notes: This table shows answers to “Why not?” from participants who said they do not post Political Opinions and Discussion content (an opinion, observation, or useful information about politics or current events). Quotes are verbatim and “[...]” indicates places where the AI interviewer asked a follow-up question; answers to successive follow-ups are joined into one quote.

G Model: Proofs and Extensions

G.1 Proofs

Proof of Proposition 1. We work backwards from the last step of the game, proving each statement in the proposition and then proving equilibrium existence.

Statement 1. Fix any viewer i and any pool of produced pieces and suppress the subscript i for convenience. Let N_k be the mass of produced pieces of type k and define $U_k(n_k)$ as the focal viewer’s total utility from the best n_k pieces of type k . Note that $U_k(\cdot)$ is concave and that because the distribution of utility is atomless, this concavity is strict.

Let $c(\mathbf{n})$ denote the total cost to the viewer—combining her opportunity cost of time and boredom cost from diminishing marginal utility—of consuming a feed with $\mathbf{n} = \{n_k\}_k$ pieces of each type. Then:

$$c(\mathbf{n}) = \gamma \sum_k \frac{n_k^2}{\sum_k n_k} + \int_0^{\sum_k n_k} \underline{u}(t) dt \quad (9)$$

Notice that this depends only on the number of pieces of each type. The utility-maximizing feed given $\mathbf{n}$ therefore obviously picks the user’s favorite n_k pieces of each type and the user’s resulting utility is $\sum_k U_k(n_k) - c(\mathbf{n})$.

The problem of finding the user’s utility-maximizing feed therefore reduces to solving $\max_{\mathbf{n}} \sum_k U_k(n_k) - c(\mathbf{n})$ for $\mathbf{n} \in \Pi_k[0, N_k]$, where N_k is the produced mass of each type. The function being maximized is continuous on a compact set. $U_k(\cdot)$ is strictly concave as previously established, and the components of $c(\mathbf{n})$ are convex (the latter term because $\underline{u}(\cdot)$ is increasing), so the function being maximized is also strictly concave. It therefore has a unique maximizer $\mathbf{n}^*$, which yields a unique P^* .

To establish the inequality in the statement, notice that:

$$\frac{\partial c}{\partial n_k} = \gamma(2h_k - \sum_{\ell} h_{\ell}^2) + \underline{u}(\sum_k n_k) = \tau + 2\gamma h_k = t_k \quad (10)$$

The marginal cost of adding a type- k piece to the feed is t_k . This means the marginal piece of each type in the feed has $u_p = t_k$.

Statement 2. By the result above, the viewership of piece p is $v_p = \sum_i \mathbb{1}(u_{ip} \geq t_{ik})$. Since each creator j is infinitesimal and their production decisions hence cannot move viewers' cutoffs, creator j 's expected payoff from producing type k before cost shocks and holding fixed the behavior of other creators, is:

$$\bar{\pi}_{jk} = \mu_k + \tilde{\mu}_{jk} + \mathbb{E}_{F_{kg}} [\beta v_p^{\rho} + w v_p] \quad (11)$$

Her corresponding payoff from not producing is zero. The standard logit formula then immediately gives $\sigma_{jk} = \frac{\exp(\psi \bar{\pi}_{jk})}{1 + \sum_{\ell} \exp(\psi \bar{\pi}_{j\ell})}$.

Statement 3. Let $W(w)$ and $V(w)$ be total viewer surplus and total viewership at wage w , respectively. The platform's profit is then $W(w) - wV(w)$ and an interior profit-maximizing solution satisfies the first-order condition $W'(w^*) = V + w^*V'(w^*)$. The right-hand side is the right-hand side of the equation in the proposition. It remains to unpack the left-hand side.

Because we have assumed the platform ignores w when assigning feeds to viewers, w affects $W(w)$ only through the pool of produced content. Imagine we add to this pool a mass ε of type- k pieces that each give viewer i utility u . By the envelope theorem, viewer i 's surplus rises to first order by $\varepsilon(u - t_{ik})^+$. The same is true for removing masses of pieces. Since raising w affects creator j 's production of type k at rate $\frac{d\sigma_{jk}}{dw}$ and creator j 's pieces are drawn from F_{kg} , we get:

$$W'(w^*) = \int_J \sum_k \frac{d\sigma_{jk}}{dw} \mathbb{E}_{F_{kg}} \left[\sum_i (u_{ip} - t_{ik})^+ \right] dj \quad (12)$$

Existence of equilibrium. To begin with, fix some $w \geq 0$. Let $\mathbf{Q} := \{Q_{kg}\}_{k,g}$ be any set of masses of type- k posts produced by groups g . Note that $\mathbf{Q}$ lies in a compact, convex set. The proofs of Statements 1 and 2 demonstrated mappings from $\mathbf{Q}$ to viewer feeds and from viewer feeds to a new $\mathbf{Q}$, respectively.

Moreover, these mappings are continuous. Each viewer's utility-maximizing $\mathbf{n}_i$ is the unique maximizer of a function over $\Pi_k[0, N_k(\mathbf{Q})]$ that is continuous in $\mathbf{n}_i$ and $\mathbf{Q}$, so it and the cutoffs t_{ik} are continuous in $\mathbf{Q}$. Since utilities are atomless, the probability that a piece lands in any given set of feeds is continuous in the cutoffs t_{ik} . Since views take finitely many values, $\bar{\pi}_{jk}$ and the logit probabilities are continuous too.

Brouwer's fixed point theorem then gives an equilibrium $\mathbf{Q}$ at each w .

So it remains to be shown that the platform has an optimal w . The same arguments as above show that W and V are continuous in $\mathbf{Q}$ and that the set of $(w, \mathbf{Q})$ pairs where $\mathbf{Q}$ is an equilibrium under w is closed. Note there is some positive constant $\underline{V}$ that bounds total viewership from below at every wage. Also, for some $\bar{W}$, for any $\mathbf{Q}$, $W(\mathbf{Q}) \leq \bar{W}$. Thus any $w > \bar{W}/\underline{V}$ yields negative profit and performs strictly worse than $w = 0$. The platform's profit $W(\mathbf{Q}) - wV(\mathbf{Q})$ is therefore continuous on a compact set of $(w, \mathbf{Q})$ pairs with $w \leq \bar{W}/\underline{V}$, so it has a maximizer $(w^*, \mathbf{Q}^*)$, which gives an equilibrium.

□

Proof of Proposition 2. Fix any wage w . We start by showing that any equilibrium at w generates a fixed point of the kind described in the proposition. We then show that any such fixed point describes an equilibrium at w .

Notationally, we will let Q_{kg} denote the mass of type- k pieces produced by group g and $Q_{kg}\mathbb{P}_{F_{kg}}[A]$ the mass of produced pieces satisfying an arbitrary condition A on the q and e_i draws.

We begin with demand. Fix an equilibrium at w and let τ_i and h_{ik} be the attention prices and feed shares in this equilibrium. By Proposition 1, viewer i sees piece p of type k if and only if:

$$m_k + \tilde{m}_{ik} + q_p + e_{ip} - 2\gamma h_{ik} \geq \tau_i \quad (13)$$

A piece's views v_p are the number of viewers i for whom this inequality holds. So the views of a piece of type k created by a creator from group g are distributed according to $D_{kg}(\cdot, \tau, h)$ as described in the proposition.

Next, consider supply. By Proposition 1, a creator j producing type k as a member of group g expects utility $\mathbb{E}_{F_{kg}}[\beta v_p^\rho + w v_p]$ from views, which we can rewrite as $\mathbb{E}_{D_{kg}}[\beta v_p^\rho + w v_p]$. Substituting into the logit formula from Proposition 1 and integrating over creators yields $Q_{kg} = S_{kg}(\mathbf{D})$ as defined in the proposition.

Finally, consider consumption shares. Viewer i 's feed contains all pieces of type k satisfying Inequality (13). A total mass $S_{kg}(\mathbf{D})\mathbb{P}_{F_{kg}}[m_k + \tilde{m}_{ik} + q_p + e_{ip} - 2\gamma h_{ik} \geq \tau_i]$ satisfy this inequality. Summing over g gives viewer i 's consumption C_{ik} and summing over viewers gives $C_k = \sum_g S_{kg}(\mathbf{D})\mathbb{E}_{D_{kg}}(v)$.

The total consumption of viewer i is $T_i = \sum_\ell C_{i\ell}$, so it follows by definition that $h_{ik} = \frac{C_{ik}}{\sum_\ell C_{i\ell}}$ and by Proposition 1 that $\tau_i = \underline{u}_i(\sum_\ell C_{i\ell}) - \gamma \sum_\ell h_{i\ell}^2$.

So this equilibrium is a fixed point of the kind described in the proposition. It remains to show that any such fixed point is an equilibrium.

Take any fixed point (τ, h) generating $\mathbf{D}$ and suppose creators' production decisions are described by $S_{kg}(\mathbf{D})$ and each viewer i receives the feed P_i of pieces satisfying Inequality (13). To demonstrate this is an equilibrium, we first show that P_i maximizes viewer i 's utility and then that creator production decisions are optimal in expectation.

Note that P_i contains C_{ik} type k pieces and its total mass is $\sum_\ell C_{i\ell}$ and shares of each type of h_{ik} . Hence as in the proof of Proposition 1, the marginal cost of adding a piece of type k is $\tau_i + 2\gamma h_{ik}$. By Inequality (13), each piece of type k in P_i offers utility exceeding this cost and each piece of type k outside P_i falls below this threshold. This means no small change in any C_{ik} can raise i 's utility. Since $\sum_k U_k(n_k)$ is additively separable across content types, we are at a local optimum. Since, as we established earlier, i 's utility is concave in the masses composing $\mathbf{n}_i$, P_i must be globally optimal.

Given this, each type k piece made by a creator of group g has an expected distribution of views given by D_{kg} , so such a creator's expected utility from producing this type is exactly as in $S_{kg}(\mathbf{D})$ and the S_{kg} logit probabilities represent her best response.

This means both production and consumption form an equilibrium at w given that they satisfy the fixed-point property in the proposition.

□

Proof of Proposition 3. We begin by solving for the equilibrium and, once we have done so, prove each statement in the proposition.

Since creators are homogeneous, each creator produces type k with probability σ_k , and $Q = \sigma_L + \sigma_H$ and the H production share is $\alpha = \frac{\sigma_H}{Q}$. Let $R_k = \mathbb{E}(\beta v^\rho + wv)$ be a creator's expected reward from views when producing type k given the equilibrium behavior of other creators.

Step 1: Viewer cutoffs. We begin by observing that enough is produced for viewers to fill their feeds. By Proposition 1, a creator chooses to produce nothing with probability $\sigma_0 = 1/(1 + \exp(\psi(\mu_L + R_L)) + \exp(\psi(\mu_H + R_H))) < \frac{1}{3}$, so $Q > \frac{2}{3}$ and hence $Q > T$ by our parameter assumptions.

Since viewers are homogeneous and there are no diminishing returns, each viewer watches a piece of type k exactly when $q_p + e_{ip} \geq \tau - m_k$ for some common cutoff τ . Let $M(\tau)$ denote the mass of posts satisfying this inequality given a hypothesized τ . Since viewers consume T posts, solving for τ means solving $M(\tau) = T$. We do this via guess-and-verify.

To guess, note that the density of $q + e$ is $\frac{1}{2s}$ on the interval $[-(1-\lambda)s, (1-\lambda)s]$. If both $\tau - m_L$ and $\tau - m_H$ lie in this interval, the mass of pieces a viewer sees is therefore given by $\frac{Q}{2s} [\alpha(s - \tau + m_H) + (1-\alpha)(s - \tau + m_L)]$. Setting this equal to T and rearranging to isolate τ yields our guess:

$$\tau^* = m_L + s + \alpha\Delta - \frac{2sT}{Q} \quad (14)$$

We now verify that both $\tau^* - m_L$ and $\tau^* - m_H$ lie in the aforementioned interval. Since $m_H > m_L$ it suffices to show that:

$$m_H - s < \tau^* - \lambda s \quad \text{and} \quad \tau^* + \lambda s < m_L + s \quad (15)$$

Substituting in Equation (14) and rearranging, the first inequality is equivalent to $\frac{2T}{Q} < 2 - \lambda - \frac{(1-\alpha)\Delta}{s}$. This is true: since $Q > \frac{2}{3}$ and we have assumed $T \leq \frac{1}{2}$, the left-hand side is $< 3T \leq \frac{3}{2}$; and on the right-hand side, $\lambda + \frac{(1-\alpha)\Delta}{s} \leq \lambda + \frac{\Delta}{s} < T$, where the latter inequality is assumed in the proposition; so the right-hand side is larger than $2 - T \geq \frac{3}{2}$.

The second inequality, meanwhile, is equivalent to $\frac{2T}{Q} > \lambda + \frac{\alpha\Delta}{s}$. This is true because $Q \leq 1$ and we have assumed that $T > \lambda + \frac{\Delta}{s}$.

We can conclude that $M(\tau^*) = T$ and since M is nonincreasing and strictly decreasing near τ^* , it is the only cutoff satisfying this equation.

Step 2: Creator rewards. Consider a piece of type k with quality q . If $m_k + q \geq \tau + \lambda s$, all n viewers see this piece, even those with the least favorable e draws, and its creator therefore receives $\beta n^\rho + wn$ viewership utility. Similarly, if $m_k + q \leq \tau - \lambda s$ no viewer sees the piece. In between, each viewer sees the piece independently with probability $\frac{m_k + q - \tau + \lambda s}{2\lambda s}$, so its viewership is binomially distributed with this success probability. Since $m_k + q$ is uniformly distributed over this intermediate range, viewership is uniform on $\{0, \dots, n\}$. The creator's expected viewership reward over this range is hence $\frac{1}{n+1} \sum_{v=0}^n (\beta v^\rho + wv)$. By Inequalities (15), the distribution of $m_k + q$ covers this whole intermediate range for both H and L , meaning

that:

$$R_k = \frac{1}{2s}((\beta n^\rho + wn)(m_k + s - \tau - \lambda s) + 2\lambda s \frac{1}{n+1} \sum_{v=0}^n (\beta v^\rho + wv)) \quad (16)$$

This equation in turn shows that $R_H - R_L = (\beta n^\rho + wn) \frac{\Delta}{2s}$ which does not depend on τ . The creator logit choice formula then gives $\frac{\sigma_H}{\sigma_L} = \exp(\psi(\mu_H - \mu_L + R_H - R_L))$ and, as in the proposition,

$$\alpha = \frac{\sigma_H}{\sigma_H + \sigma_L} = \frac{\exp(\psi[\mu_H - \mu_L + (\beta n^\rho + wn) \frac{\Delta}{2s}])}{1 + \exp(\psi[\mu_H - \mu_L + (\beta n^\rho + wn) \frac{\Delta}{2s}])} \quad (17)$$

We also have that, substituting Equation (14) into the equation for R_k and using $\frac{1}{n+1} \sum_{v=0}^n (\beta v^\rho + wv) \geq 0$ and $Q \leq 1$, $R_L \geq (\beta n^\rho + wn)(T - \frac{\Delta}{2s} - \frac{\lambda}{2})$, so since $T > \lambda + \frac{\Delta}{s}$, $R_L \geq R_H - R_L$.

Step 3: Consumption shares. By Equation (14) and $m_H = m_L + \Delta$, each viewer's feed contains $\alpha T + \frac{Q\Delta}{2s}\alpha(1 - \alpha)$ type H pieces. Dividing by T , the type- H share is $h := \alpha + \frac{Q}{T} \frac{\Delta}{2s}\alpha(1 - \alpha)$ as stated in the proposition. We therefore have that:

$$\frac{h}{1-h} = \frac{\alpha}{1-\alpha} \left(1 + \frac{1}{\frac{2sT}{Q\Delta} - \alpha}\right) \quad (18)$$

where $\frac{2sT}{Q\Delta} > 2 > \alpha$.

Step 4: Aggregate production. By the logit choice formula, $\frac{Q}{1-Q} = \exp(\psi(\mu_L + R_L)) + \exp(\psi(\mu_H + R_H)) = \frac{\exp(\psi(\mu_L + R_L))}{1-\alpha}$. Substituting Equations (14) and (16) and taking logs yields:

$$\log \frac{Q}{1-Q} - \frac{\psi(\beta n^\rho + wn)T}{Q} = \psi\left(\mu_L - \alpha(R_H - R_L) + \lambda\left(\frac{1}{n+1} \sum_{v=0}^n (\beta v^\rho + wv) - \frac{\beta n^\rho + wn}{2}\right)\right) - \log(1-\alpha) \quad (19)$$

The right-hand side does not depend on Q and the left-hand side is strictly increasing in Q , so equilibrium production quantity Q is unique.

We note three facts about Q . First, let $G(Q)$ denote the left-hand side of Equation (19). Then $G'(Q) = \frac{1}{Q(1-Q)} + \frac{\psi(\beta n^\rho + wn)T}{Q^2} \geq \frac{1}{Q(1-Q)}$. Second, since $\mu_L > 0$ and $R_L \geq R_H - R_L$ by Step 2, $1 - Q = \sigma_0 < e^{-\psi(\mu_L + R_L)} < e^{-\psi(R_H - R_L)}$.

Finally, the right-hand side of Equation (19) equals $\psi\lambda\left(\frac{1}{n+1} \sum_{v=0}^n (\beta v^\rho + wv) - \frac{\beta n^\rho + wn}{2}\right) + \Phi(R_H - R_L)$, where $\Phi(x) := \psi\mu_L - \psi\alpha(x)x - \log(1 - \alpha(x))$ and $\alpha(x)$ is the formula for α in Step 2 with $R_H - R_L$ denoted by x . Differentiating gives $\alpha'(x) = \psi\alpha(x)(1 - \alpha(x))$, so:

$$\Phi'(x) = -\psi\alpha(x) - \psi x \alpha'(x) + \frac{\alpha'(x)}{1 - \alpha(x)} = -\psi^2 x \alpha(x)(1 - \alpha(x)) \quad (20)$$

which is negative for $x > 0$. So, all else equal, a wider gap in reward between H and L content lowers total production.

We can now move on to proving the statements.

Statement 1. Increasing s strictly lowers $R_H - R_L = (\beta n^\rho + wn) \frac{\Delta}{2s}$ since $w > 0$. This strictly decreases α since $\alpha'(x) > 0$.

Now consider Q . In Equation (19), increasing s leaves the left-hand side unchanged and increases the right-hand side, since $\Phi' < 0$. Since the left-hand side is increasing in Q , Q increases.

Moreover, we can see that Q rises less than proportionally with s : differentiating Equation (19),

$$G'(Q) \frac{dQ}{ds} = \frac{\psi^2(R_H - R_L)^2 \alpha(1 - \alpha)}{s} \quad (21)$$

Using the inequalities $G'(Q) \geq \frac{1}{Q(1-Q)}$, $\alpha(1 - \alpha) \leq \frac{1}{4}$, and $1 - Q < \exp(-\psi(R_H - R_L))$, we can therefore see that

$$\frac{s}{Q} \frac{dQ}{ds} \leq (1 - Q) \frac{\psi^2(R_H - R_L)^2}{4} < \frac{\psi^2(R_H - R_L)^2 e^{-\psi(R_H - R_L)}}{4} \leq e^{-2} < 1 \quad (22)$$

This therefore means that $\frac{Q}{s}$ strictly decreases in s as well as α . Consequently, by Equation (18), so does the H consumption share h .

Statement 2. Let θ denote the focal parameter out of β, ρ , or w . Raising θ increases $\beta n^\rho + wn$ (for ρ , if $\beta > 0$), so $R_H - R_L = (\beta n^\rho + wn) \frac{\Delta}{2s}$ increases. Since $\alpha'(x) > 0$, α rises too.

Next consider h . Differentiating h yields:

$$\frac{dh}{d\theta} = \left(1 + \frac{Q\Delta}{2sT}(1 - 2\alpha)\right) \frac{d\alpha}{d\theta} + \frac{\Delta}{2sT} \alpha(1 - \alpha) \frac{dQ}{d\theta} \quad (23)$$

The term in large brackets exceeds $1 - \frac{\Delta}{2sT} > \frac{1}{2}$, and $\frac{d\alpha}{d\theta} = \alpha'(R_H - R_L) \frac{\Delta}{2s} \frac{d(\beta n^\rho + wn)}{d\theta} = \psi \alpha(1 - \alpha) \frac{\Delta}{2s} \frac{d(\beta n^\rho + wn)}{d\theta} > 0$. So the first term is positive.

It remains to be shown that $\frac{dQ}{d\theta}$ cannot be negative enough to offset this. To bound $\frac{dQ}{d\theta}$, we differentiate both sides of Equation (19) with respect to θ . The left-hand side depends on θ through Q and $\beta n^\rho + wn$, and the right-hand side through the λ term and $R_H - R_L$. Using Equation (20):

$$G'(Q) \frac{dQ}{d\theta} = \left(\frac{\psi T}{Q} - \psi^2(R_H - R_L) \alpha(1 - \alpha) \frac{\Delta}{2s}\right) \frac{d(\beta n^\rho + wn)}{d\theta} + \psi \lambda \frac{d}{d\theta} \left(\frac{1}{n+1} \sum_{v=0}^n (\beta v^\rho + wv) - \frac{\beta n^\rho + wn}{2}\right) \quad (24)$$

Each term $\beta v^\rho + wv$ in the average rises by between zero and $\frac{d(\beta n^\rho + wn)}{d\theta}$, since $v \leq n$, so the last derivative lies between $-\frac{1}{2}$ and $\frac{1}{2}$ times $\frac{d(\beta n^\rho + wn)}{d\theta}$. Dropping the positive term $\frac{\psi T}{Q}$ and using $\alpha(1 - \alpha) \leq \frac{1}{4}$ and $G'(Q) \geq \frac{1}{Q(1-Q)}$ then yields:

$$\frac{dQ}{d\theta} \geq -(1 - Q) \psi \left(\frac{\lambda}{2} + \frac{\psi(R_H - R_L)\Delta}{8s}\right) \frac{d(\beta n^\rho + wn)}{d\theta} \quad (25)$$

Substituting both bounds into the expression for $\frac{dh}{d\theta}$:

$$\frac{dh}{d\theta} \geq \psi \alpha(1 - \alpha) \frac{\Delta}{2s} \left(\frac{1}{2} - \frac{(1 - Q)\lambda}{2T} - \frac{(1 - Q)\psi(R_H - R_L)\Delta}{8sT}\right) \frac{d(\beta n^\rho + wn)}{d\theta} \quad (26)$$

Since $1 - Q < \frac{1}{3}$ and $\lambda < T$, the second term in the large brackets is below $\frac{1}{6}$. Since $1 - Q < e^{-\psi(R_H - R_L)}$, $\frac{\Delta}{sT} < 1$, and $xe^{-x} \leq e^{-1}$, the third is below $\frac{1}{8e} < \frac{1}{20}$. So $\frac{dh}{d\theta} > 0$.

Statement 3. This is immediate. In the definition of $R(s)$, the numerator does not depend on s ; in the denominator, Q increases with s by Statement 1, and $(1 - \alpha)\mu_L + \alpha\mu_H$ increases because α falls and $\mu_L > \mu_H$. So $R(s)$ is decreasing in s .

Statement 4. A piece of type k lies within distance ε of a viewer's cutoff if and only if $q + e \in [\tau - m_k - \varepsilon, \tau - m_k + \varepsilon]$. If for both $k = L, H$, this interval lies in $[-(1 - \lambda)s, (1 - \lambda)s]$, which is the region where $q + e$ has density $\frac{1}{2s}$, the expected mass of produced pieces lying within ε of the cutoff is $Q \frac{2\varepsilon}{2s}$, which is strictly decreasing in s by Statement 1.

In fact, both of these intervals will lie within $[-(1 - \lambda)s, (1 - \lambda)s]$ at both s and s' . By Equation (14), $\tau - m_L + \varepsilon \leq (1 - \lambda)s$ is equivalent to $\varepsilon \leq \frac{2sT}{Q} - \lambda s - \alpha\Delta$, whose right-hand side is at least $2sT - \lambda s - \Delta > s(T - \lambda) - \Delta$. Similarly, $\tau - m_H - \varepsilon \geq -(1 - \lambda)s$ is equivalent to $\varepsilon \leq 2s - \lambda s - (1 - \alpha)\Delta - \frac{2sT}{Q}$, whose right-hand side exceeds $2s - \lambda s - \Delta - 3sT \geq s(T - \lambda) - \Delta$ because $\frac{2T}{Q} < 3T$ and $T \leq \frac{1}{2}$. Both conditions therefore hold whenever $\varepsilon < s(T - \lambda) - \Delta$. Since $T > \lambda$, $s'(T - \lambda) - \Delta > s(T - \lambda) - \Delta$, so they also hold at s' . $\square$

G.2 Short-form content

We write a short toy model formalizing the observations in Section 2.2 about why the social media business model complements short-form content.

Consumption of short and long content. For simplicity, we consider a single content type with a baseline appeal m and joint distribution F from which q and e shocks are drawn. *Subunits* of content are the smallest units over which the q and e shocks are realized. Viewer i 's utility from subunit p is $u_{ip} = m + q_p + e_{ip}$.

A piece of content can be *short*, consisting of a single subunit, or *long*, consisting of two subunits with independently drawn q and e shocks. Each viewer has a feed with two slots. A viewer's utility from a feed is her utility from the two subunits in the feed, plus χ if both subunits were part of the same long piece. χ can be positive, representing a preference for coherence, or negative, representing a preference for variety.

Production of short and long content. Pieces of content are produced by a finite number of homogeneous potential creators. Each creator chooses between producing nothing, producing a short piece, or producing a long piece, and draws independent logit shocks for each option from a common distribution with scale $1/\psi$.

Producing a piece of length $\ell \in \{1, 2\}$ incurs cost c_ℓ and yields nonpecuniary payoff μ_ℓ . Because creators are ex-ante homogeneous, we can roughly think of μ as subsuming expected utility from viewership.³¹

We allow both costs and rewards to be either concave or convex in length, depending on whether $c_2 < 2c_1$ and likewise whether $\mu_2 < 2\mu_1$. In general, convex costs and concave rewards will favor short-form content on both social and legacy media, and vice versa for concave costs and convex rewards. Because our results are about the *difference* in returns to short-form content *between* legacy and social media, we can remain agnostic between these possibilities.

Outlets' objectives and technological capabilities. As before, we assume outlets maximize expected viewer surplus net of procurement costs. Legacy outlets are unable to observe and act on the q and e shocks when distributing content to viewers, while social media platforms are able to.

Mechanism I: social media benefits more from breaking up produced content. Our first point is immediate. Fix some produced pool containing some number of long units of content and fix a viewer i . Consider whether the outlet wishes to serve viewer i a single long unit or two short subunits. (Either by serving two

³¹Technically it should then depend on the number of other creators being recruited, but we abstract away from this for simplicity.

produced short subunits, or “breaking up” long units into their subunits and serving viewer i subunits from two different long units.)

From the perspective of a legacy outlet unable to observe q and e realizations, all subunits are interchangeable. So the outlet will strictly prefer to serve long units when viewers prefer coherence ($\chi > 0$) and strictly prefer to serve pairs of short subunits when viewers prefer variety ($\chi < 0$).

From the perspective of a social media platform, the platform may want to serve viewer i two short subunits even if $\chi > 0$. This will happen if the two highest $q + e$ draws accrue to two subunits that do not belong to the same long unit, and the utility from serving these two subunits exceeds the utility from serving the best long unit by at least χ .

So conditional on a produced pool, social media platforms wish to distribute short-form content in a strictly larger set of cases than legacy outlets.

Endogenizing production. We now endogenize production. Each outlet sets a wage per view and creators make production decisions. Views accrue to subunits, so a long piece served to n viewers receives $2n$ views.

Mechanism II: social media benefits more from reducing barriers to entry. Consider a *minimum length requirement* where the platform refuses to accept short pieces.

Proposition 4. Suppose $q + e$ is nondegenerate and fix a wage.

1. If removing a minimum length requirement reduces total production of subunits, it is not profitable for a legacy or a social media outlet.
2. If doing so increases total production of subunits, the cases where it strictly raises a social media platform’s profit are a strict superset of the cases where it raises a legacy outlet’s profit.

Removing a minimum length requirement *lowers barriers to entry*, inducing a flow of creators from “produce nothing” to “produce a short subunit.” But it can also induce a flow of creators from “produce a long unit” to “produce a short subunit”.

If the latter flow outnumbers the former, such that removing a minimum length requirement reduces total production scale, it is never profitable for either type of media.

But suppose that removing a minimum length requirement increases total production scale. This could be because production costs are convex or rewards are concave, but it could also be because there is a large mass of creator taste shocks near the boundary between “produce nothing” and “produce a short unit.” For example, it could be the case that production costs are concave in length *but* there are also many people who never have an idea good enough for a long post but occasionally have an idea for a 30-second video or short post and will only post it if minimum requirements are low enough.

In this case, because social media platforms benefit more from production scale than legacy outlets do, social media platforms are more likely to favor removal of the minimum length requirement. (They will not always do so, because if $\chi > 0$ then removing a length requirement can be unprofitable even if it increases total production of subunits.)

Mechanism III: short content improves social media’s experimentation step. Throughout the paper so far we have abstracted away from the details of the experimentation step where the platform learns a piece’s

ex-post distribution of appeal to viewers. Our next mechanism requires us to unpack this step a bit more.

Learning about a piece’s potential appeal requires learning about the characteristics of the *piece* and learning about the preferences of different *viewers*. Short-form content makes the latter easier. If different subunits have different characteristics, long pieces mash together many different characteristics. This makes it hard to learn about a viewer’s tastes from her engagement with a long piece because the platform does not know which subunit’s characteristics drove the viewer’s engagement. Disaggregating pieces into the smallest units over which characteristics vary allows much more efficient learning about viewer preferences. To grasp the intuition, consider how much more TikTok knows about your preferences after you scroll for 10 hours, compared to how much Netflix knows about your preferences after you watch 10 hours of movies or TV. Netflix will not know whether you like a certain show because it has a certain famous actor, or because it is a thriller, or because it is critically acclaimed, or because it is set in London, or because. . .

Formally, suppose each subunit independently and uniformly draws a characteristic $x_p \in \{-1, 1\}$. Viewer i ’s utility is $u_{ip} = \theta_i x_p + m + q_p + e_{ip}$. Her taste for x_p is given by $\theta_i = \bar{\theta} + \tilde{\theta}_i$ where $\tilde{\theta}_i$ is independent and identically distributed across viewers. $q_p + e_{ip}$ has variance ω^2 .

In the experimentation step, the platform serves each viewer a random feed of T subunits and knows the characteristics of all served subunits. If it serves short pieces, it observes T draws of $\theta_i x_p + m + q_p + e_{ip}$ for an expected precision of $\frac{T}{\omega^2}$ about θ_i . If it serves long pieces, it observes $\frac{T}{2}$ draws of $\theta_i(x_a + x_b) + 2m + \chi + q_a + e_{ia} + q_b + e_{ib}$ and ends up with precision $\frac{T}{2\omega^2}$, or half as much precision.

We have focused on learning about viewer preferences given knowledge of piece characteristics. There is also the converse of learning about piece characteristics given knowledge of viewer preferences. But this is less interesting because if the platform is ultimately serving long pieces of content, there is actually no gain to doing the experimentation step with short pieces in the simple structure we have adopted with additive utility over subunits. However, such gains can exist with preferences that are nonadditive over the utility from subunits.

Other mechanisms. In asking why social media content is short-form, we have pitted two classes of explanations against each other. The conventional explanation in the literature is that short-form content hijacks features of our reward psychology just like slot machines do (e.g., Eyal, 2014; Bao, 2026). We have emphasized a class of alternative explanations focused on our characterization of social media’s business model as being primarily about ex-post experimentation and cherry-picking.

There are also other explanations that do not fit into either class. For example, smartphones enable people to scroll briefly on social media in their spare minutes. This complements short pieces of content that do not require much context to appreciate and does not fit well with longer pieces of content like hour-long TV episodes where there is some buildup, context-setting, and characterization before “payoffs” arrive.³² In our data, platforms vary in the extent to which viewers scroll on them briefly. The median minute spent on TikTok is part of a 22-minute session, while the median minute on Instagram is part of a 4.5-minute session.

Evidence on Mechanism I. Mechanism I relies on social media platforms discriminating between different subunits of the same long piece of content because of internal variation in a long piece’s vertical quality or horizontal characteristics. The extent to which this kind of variation exists is an empirical question. For

³²We thank Nikhil Agarwal for suggesting this explanation.

example, does somebody who likes one of a comedian’s jokes tend to like all the other jokes in the same set? Or does somebody who likes one of a movie’s scenes tend to like the others?


The ideal data for answering these questions consists of measurements of the same person’s utility from different subunits of a long piece of content, like their enjoyment of each joke from a comedy set. We can try to approximate this ideal dataset by looking at variation across the different pieces of short content posted by a creator on social media. Does a viewer who likes one of a creator’s videos tend to like the others? Do two viewers who both like some of the same creator’s videos like the *same* videos as each other?

Figures A34 and A35 show that the TikTok algorithm substantially discriminates in its assignment of videos *within* creator.

Even a viewer who sees ten of a creator’s posts in a month—indicating a fairly consistent taste for this creator—is still served only a third of the posts the creator makes over this time period. This suggests discrimination that may be partly about vertical quality: only a third of the creator’s posts are high-quality enough to make it into the viewer’s feed.


Meanwhile, compare two viewers who both see at least ten posts from the same creator over the same period. The average overlap between the posts each saw is only around 35%—indicating there is also substantial *horizontal* discrimination, with one viewer liking some of a creator’s posts and the other viewer liking others.

Appendix Figure A34: Viewers see only a small share of the videos posted by the creators they watch



Notes: This figure uses our TikTok viewership histories. For a random sample of creators in these histories we scrape every post they make between July and September 2025 using TikTok’s Research API. Conditional on a viewer seeing at least one video from a given creator during this period within three days of the video being posted, the figure plots the share of posts made by the creator in the same month that appear in the participant’s feed, as a function of the total number of posts from that creator appearing in the participant’s feeds this month. Results are weighted to be representative of the creators appearing in participants’ feeds.

Appendix Figure A35: Two viewers of the same creator see largely different videos from that creator



Notes: This figure uses our TikTok viewership histories, restricted to a period (August to September 2025) and set of viewership histories such that all histories span the full period. For each pair of participants i and i' and each creator j such that both participants see a video from j , we count n_i (the number of posts i sees from j), $n_{i'}$ (the number of posts i' sees from j), and n_{both} (the number of j 's videos seen by both i and i'). The y-axis of this plot shows $100 \times \max\{\frac{n_{\text{both}}}{n_i}, \frac{n_{\text{both}}}{n_{i'}}\}$ and the x-axis shows $\min\{n_i, n_{i'}\}$.

H TV Consumption

Table A20 estimates what percentage of TV viewing goes to (i) news, (ii) the subset of news that is about politics or political issues according to the definition our classification schema uses, and (iii) the subset of political news minutes that are partisan, again according to the definition our schema uses. Data are sourced as follows, mostly from Noy and Rao (2025).

1. Ratings are taken from Nielsen AdIntel C3 for 2019 and news programs are ABC, CBS, NBC, CNN, Fox News, and MSNBC.
2. We sample news minutes on those programs in 2026 from the TV archive and classify each minute using an LLM prompt analogous to the one used to classify TikTok videos. This allows us to detect whether a minute is political and, if so, whether it expresses a clear partisan perspective. As with our TikTok data, we drop ad minutes and compute measures as shares of non-ad minutes.
3. The “Linear TV” column shows our estimates as percentages of cable and broadcast viewing minutes. “All TV” adds streaming, which made up about 20% of TV time in 2019 according to Nielsen’s Total Audience Report. We assume streaming contains no political news.

Appendix Table A20: *News and political news as shares of TV viewing, and the partisan share of political news*

	Linear TV	All TV
News	14.4%	11.7%
Political news	10.3%	8.3%
Partisan (% of political news)	21.9%	21.9%

I Viewership Experiment


I.1 Design details

Appendix Figure A36: Viewership experiment interface




I.2 Additional exhibits

Appendix Figure A37: Attrition is unrelated to feed composition, and including incomplete sessions does not change results




Notes: This figure plots data from our viewership experiment. Panel (a) plots regressions predicting attrition by 0.5, 1, 2, 5, or 10 minutes using the share of a participant's feed up to that point assigned to a given content type. Panel (b) replicates Figure 9 Panel (c) keeping incomplete sessions instead of dropping them.

Appendix Figure A38: Engagement by content type: other pre-registered measures




Notes: This figure replicates Figure 9 Panel (c) for each of our other pre-registered outcome variables: (a) seconds watched, (b) whether the participant liked the video, (c) whether they watched the full video, and (d) whether they both liked it and watched all of it. As in that figure we control for participant, video length decile, and sampling source fixed effects.

Appendix Figure A39: Probability of watching ≥ 10 seconds by content type, correcting for behavioral noise




Notes: This figure replicates Figure 9 Panel (c) after disattenuating estimates by purging behavioral noise with an empirical Bayes correction using the end-of-experiment module where we show participants two videos they saw earlier and ask which they prefer.

Appendix Figure A40: *The partisanship premium, reweighting participants to the ideology of young TikTok users*




Notes: For each of our five pre-registered engagement outcomes, this figure plots the engagement difference between partisan and nonpartisan political videos using our main specification, either unweighted or weighting participants so they match the ideological distribution of 18-30 year old TikTok users from the American National Election Studies (Appendix Table A1).

Appendix Figure A41: *Probability of watching ≥ 10 seconds by content type, separately by the view-count floor used to sample random videos*




Notes: This figure replicates Figure 9 Panel (c) separately for two groups of videos. When we randomly sample videos from TikTok's Research API we use either a 10-view or a 10,000-view floor. Panels show our main estimates on videos sampled in each case.

Appendix Figure A42: Watch times in the data and in our model




Notes: The blue line in Panel (a) plots the share of views in which the viewer is still watching the video t seconds in (restricting to videos at least t seconds long) across all occurrences in our viewership experiment. The black dashed line shows the distribution simulated by our model. Panel (b) reproduces this in our observational viewership-history data, re-calibrating the reservation value of swiping to the next post to match the share of viewers watching videos for at least five seconds and assuming that each video in a person's actual feed is drawn from the model's utility distribution truncated above this reservation value.

Appendix Figure A43: Engagement by content type in the viewership experiment: data and model




Notes: These figures plot coefficients from pre-registered regressions of engagement outcomes on dummies for content types, with fixed effects for participant, video length decile, and sampling source. Next to these coefficients, we display the analogous coefficients using viewership simulated from our model estimates.

Appendix Figure A44: Predicted and actual TikTok views of the videos in our viewership experiment




Notes: This figure plots the realized view counts on TikTok of the videos used in our viewership experiment against the view counts predicted by our model based on utility estimates from the engagement behavior of experimental participants. It plots this slope within groups (follower count $\times$ content type $\times$ sampling source). The black dashed line is the 45 degree line, reflecting perfect concordance between our model estimates and actual viewership.

Appendix Figure A45: Stale videos' engagement and feed-appearance penalties: data and model



Notes: This figure buckets posts in our viewership history by their “age” (the time between their posting and their appearance in a viewer’s feed). On the x-axis it plots the average engagement penalty associated with a bucket relative to videos 1–2 days old, measured as the percent of viewers watching them for at least five seconds. On the y-axis it plots the feed appearance penalty of this bucket, meaning how much less common videos from this bucket are relative to videos 1–2 days old. The dashed line shows our model’s prediction about the effects of a given decrease in the engagement measure on a video’s probability of appearing in a viewer’s feed.

Appendix Figure A46: *Feed appearances of videos by creators viewers regularly see: data and model*



Notes: This figure studies the subset of participants in our viewership experiment for whom we pipe in “personalized” videos recently posted by a creator who appeared at least 3 times in their TikTok viewership history. It plots how much more likely a random video from such a creator is to appear in their feed than a random video from another creator from our experiment in the same number-of-followers bin, based on comparing the historical overlap between their viewership histories, the posting rates of the personalized creators, and the production and consumption shares of creators in this follower bin. It also plots our model’s prediction for this ratio.

Appendix Table A21: *Fit of different distributions to the upper tail of video views*

	Log-likelihood	Δ AIC	KS distance	Share of tail videos with $\geq v$ views (%)			
				$v = 10K$	100K	1M	10M
Data				37.3	9.00	1.19	0.054
Lognormal	-277,183	0	0.018	37.7	7.94	1.45	0.230
Pareto	-277,216	63	0.012	36.6	8.32	1.89	0.429
Weibull	-277,427	486	0.051	41.1	7.45	0.729	0.031
Exponential	-317,840	81,312	0.548	89.8	26.3	0.000	0.000

Notes: This table compares the fit of different parametric distributions to the upper tail of view counts in our random sample of produced TikToks drawn from the Research API. Each distribution is fit by maximum likelihood. The first column reports the log-likelihood, the second the difference in AIC from the best-fitting distribution (the log-normal), and the third the largest gap between the fitted distribution and empirical distribution.

J Creator Experiment Details


J.1 Design details

Our creator experiment was conducted in August 2026 and our reduced-form analyses were preregistered on the Open Science Framework at <https://osf.io/24tcz>. We invited participants who had completed our user survey and logged in with a public TikTok account with at least one post, and allowed new participants to enter the survey by logging in with a public TikTok account with at least one post. 1,115 creators entered the experiment after logging in and 979 completed it. As preregistered, we exclude participants who indicate at the end that they answered strategically rather than truthfully (29 participants) or reported any monetary value exceeding \$1,000 (36 further participants). We scrape participants' TikTok accounts and recover posting histories for 770 of them.

Appendix Figure A47 displays the survey flow and the full survey can be clicked through at <https://demo.mitsocialmediastudy.com/>. Completing the survey takes 14 minutes on average.


At the end of the survey each creator is extended a randomized offer. With 1% probability this offer is drawn from a BDM mechanism based on their willingness to post answers and with 99% probability it is a \$10 offer to post one of our four content types, chosen at random. We start the probabilities of each type at 25% each and then, as data collection continues, adjust them to achieve a balanced sample of created videos of each type given creators' offer acceptance rates across types.

Appendix Figure A47: Creator experiment design



J.2 Additional exhibits

Appendix Figure A48: TikTok views of paid creator-experiment posts, by creators' willingness to post for \$0.25



Appendix Figure A49: Creator payoff parameter estimates: creators with at least 1,000 followers

Parameter	Estimate	SE
<i>Mean taste μ_k, % relative to entertainment</i>		
Lifestyle	13.5	(3.2)
Stories	-6.5	(4.3)
<i>Mean taste μ_k, % relative to nonpartisan politics</i>		
Partisan politics	6.0	(25.8)
Value of views β	1.44	(0.25)
Curvature ρ	0.33	(0.005)
<i>Expected views, creators willing to post only some types</i>		
Types willing to post for \$0.25	1111	(117)
Types not willing to post for \$0.25	822	(115)
<i>Expected views, % relative to entertainment</i>		
Lifestyle	-37.1	(4.7)
Stories	-33.0	(5.2)
<i>Expected views, % relative to nonpartisan politics</i>		
Partisan politics	106.8	(55.7)
<i>For the median planned post</i>		
Total utility from post	\$7.14	(1.85)
Utility from viewership	\$12.65	(2.25)
Utility from first 100 views	\$6.63	(1.13)

Notes: This figure replicates Panel (b) of Figure 10, estimated only on creators in our experiment with at least 1,000 followers.

Appendix Table A22: Creator results are robust to the valuations of creators already planning to post

Specification	Creators	Mean taste, % cheaper		Oaxaca supply contribution (pp)		Supply-equalized counterfactual	
		Lifestyle vs. ent.	Partisan vs. nonpart.	Lifestyle – ent.	Partisan – nonpart.	Lifestyle share (%)	Partisan share (%)
Main estimates	770	8.8 (1.5)	35.3 (8.5)	22.8	14.8	19.4	32.2
Creators who would forgo rather than delay	486	9.4 (2.5)	37.7 (11.1)	23.6	17.4	18.2	42.5
Planning-branch payments ×1.1	770	8.9 (1.5)	34.9 (8.3)	22.5	14.8	19.4	32.6
Planning-branch payments ×1.2	770	8.9 (1.5)	34.4 (8.2)	22.2	14.9	19.5	32.8
Planning-branch payments ×1.5	770	8.9 (1.4)	33.1 (7.7)	21.5	15.0	19.5	33.3
Planning-branch payments ×2	770	9.0 (1.4)	31.3 (7.0)	20.6	15.0	19.6	33.5
Planning-branch payments ×3	770	9.0 (1.4)	28.6 (6.0)	19.5	14.9	19.7	33.6

Notes: This table plots robustness of our creator-experiment results to the treatment of the “willingness to pay to keep posting” question branch where creators who were already planning to post a given type of content tell us how much we would need to pay them not to post it for a week. Most creators who respond to this question answer a followup by saying this would lead them to postpone making the post rather than not make it at all, so treating their WTP divided by their expected number of posts as their per-post valuation may underestimate their true valuation. Here we simply test sensitivity of our results to this kind of underestimation by multiplying responses in this branch by increasing numbers.

Appendix Table A23: Creator results are robust to how we elicit the expected effects of the boost

Boost question	β	Mean taste, % cheaper		Oaxaca supply contribution (pp)		Supply-equalized counterfactual	
		Lifestyle vs. ent.	Partisan vs. nonpart.	Lifestyle – ent.	Partisan – nonpart.	Lifestyle share (%)	Partisan share (%)
Both groups (main)	1.20 (0.09)	8.8 (1.5)	35.3 (8.5)	22.8	14.8	19.4	32.2
Own views (unincentivized)	1.24 (0.14)	8.8 (1.5)	35.2 (8.5)	23.0	14.7	19.4	32.1
Others' views (incentivized)	1.16 (0.11)	8.9 (1.5)	35.4 (8.5)	22.5	14.9	19.4	32.5

Notes: This table shows robustness of our creator-experiment results to whether participants' expected effects of the intervention we offer on their viewership are given in response to an unincentivized question about its expected effects on their own viewership or an incentivized question about its expected average effects on other participants' viewership.

J.3 Social desirability


Panel (e) of Figure 10 asks creators who accepted the \$0.25 offer for some content types but not others why they did so. The question first reminds them of their choices and then asks: “What is the main reason you would accept the \$0.25 offer for [the first group of content types], but not for [the second group]?” It offers

four answer options and we randomize the wording of one of them to flip its perceived social desirability.

1. Creating posts of the former types would be easier or more enjoyable, or you would be more comfortable posting them.
2. *Original wording:* You care about how popular your post is (how many views or likes you would get), and you expect you would get more views and likes from posts of the former types.
Flipped wording: You want to create content that reaches and engages people. You expect posts of the former types would receive more views and likes.
3. You think posts of the former types would create more opportunities to earn money, attract sponsorships, or build an audience you could monetize.
4. You were already planning to post the former types of content, but were not planning to post the latter types.


To verify that we flip social desirability successfully we run a supplementary survey of 62 respondents recruited via Prolific. Each respondent sees one of the two wordings and ranks the first three options from the one where giving it as the main reason makes a creator look best in other people's eyes to the one where it makes a creator look worst. Figure A50 shows results.

Appendix Figure A50: Perceived social desirability of the creator experiment's stated reasons




K Emotion-Sampling Details and Results

Appendix Figure A51: Emotion-sampling app interface




Appendix Figure A52: Daily TikTok use before and during the emotion-sampling week

(a) Minutes per day




(b) Videos per day




Notes: Blue circles show average daily TikTok use among participants who donate their viewership histories, before versus after they enter our experiment by downloading our emotion-sampling app. It stacks users who have versus have not yet started the app.

Appendix Figure A53: Political and enviable lifestyle videos trigger negative emotions, but viewers do not regret consuming the ones appearing in their feeds

(a) Open-ended surveys and interviews: what content types drive negative emotions on social media?




(b) Emotion-sampling: what content types reduce contemporaneous emotional valence?



(c) App questions: what did you get out of the video you just watched? Did you regret watching it?


Reason	All	Enviably	Politics
Entertaining	55%	55%	34%
Funny	32%	22%	24%
Staying Informed	22%	18%	58%
Visually Appealing	17%	23%	7%
Self-Improvement Advice	8%	13%	4%
Product Recommendation	8%	11%	3%
Regret Watching (Negative Emotions)	4%	4%	12%
Regret Watching (Waste of Time)	4%	3%	2%

(d) Retrospective survey: which of these specific creators from your feed would you want to block?




Appendix Figure A54: Response delays to wellbeing notifications

(a) Notifications delivered while off social media




(b) Notifications delivered while on social media




Notes: This figure plots distributions of notification response delays, defined as the number of seconds between notification delivery and notification response for notifications that are responded to. Panel (a) plots the distribution for notifications delivered while the recipient is not on social media, while Panel (b) plots notifications delivered while the recipient is on social media. (Social media usage is measured through viewership histories.)

Appendix Figure A55: Balance tests: content on screen, response delays, and lagged wellbeing




Notes: This figure restricts to videos that were being viewed at the time that a user responded to a notification from the app containing a wellbeing question; the blue coefficients show the correlation between the amount of time it took the user to respond to the question (top-coded at 10 minutes) and the characteristics of the content being viewed at the time of the response; the orange coefficients show the correlation of the *most recent* wellbeing response submitted by the respondent (within the last 30 minutes) with the characteristics of the content being viewed at the time of the response.

Appendix Figure A56: Wellbeing effects of enviable and political content: linear panel models



Notes: Panels (a) and (b) plot estimates of the effects of enviable and political content exposure, respectively, on contemporaneous wellbeing using two linear panel model specifications. The purple dots use, as a panel unit, a social media usage session, and as a time unit the n th wellbeing response in that session; we estimate coefficients using the package developed by Freyaldenhoven et al. (2026) and de-cumulate the $t > 0$ coefficients. The orange dots regress contemporaneous wellbeing on the characteristics of the video being viewed while the wellbeing response was received ($t = 0$) and, one at a time, the characteristics of the videos being viewed at the next 3 wellbeing responses in the session ($t < 0$) or the previous 3 ($t > 0$), controlling for session fixed effects.

Appendix Figure A57: Robustness of the wellbeing effects of enviable and political content



Notes: This figure plots the coefficients on enviable and political content from regressions of contemporaneous wellbeing on indicators for enviable content and for our other content types (with entertainment as the omitted category). As in our standard specifications we control for session fixed effects at baseline, and then add a sequence of additional controls as described on the y-axis.