The Price of Engagement: Measuring Platform Market Power Through Recommendation Algorithm Audits*

Calvin Jahnke and Vartan Shadarevian

November 2025

[Preliminary and incomplete; please do not circulate.]

Abstract

We study platform market power in attention markets, where recommendation algorithms mediate exchange through user time rather than prices. We develop a novel methodology combining randomized audits with structural estimation to measure market power without requiring user data or platform cooperation. Our approach models how platforms can manipulate search costs to induce time expenditures above competitive levels—a "time markup" analogous to price markups in traditional markets. In over 2,000 systematic audits of a video-sharing platform, we find that recommendations systematically respond to multiple dimensions of user behavior consistent with engagement maximization. Simulations from our full model demonstrate that engagement-maximizing strategies favor longer videos than would maximize user welfare, with the gap shrinking as users' outside options improve. A simplified model fits observed platform strategies and suggests a small degree of market power that is slightly higher among older viewers. Our structural audit methodology offers regulators a template for monitoring algorithmic behavior.

*We extend our thanks to Jakub Kastl for his past and future advising on this project. We also thank the students and faculty in the Department of Economics at Princeton University for their comments on this project.

1 Introduction

Platforms in two-sided digital markets increasingly employ machine learning systems for matching. In particular, the use of recommendation algorithms in social media is ubiquitous and accounts for significant increases in usage of these products (Guess et al., 2023); however, it is unclear whether this increased usage is driven by increased match value or if these systems allow platforms to manipulate user incentives to extract consumer surplus. Concern about this potential manipulation has led to laws in at least two states targeted at banning the use of algorithmic recommendations on social media for kids and teenagers.[1][2] Measuring this extraction through market power presents a challenge in this setting lacking traditional prices and where markets are difficult to define, and effective regulation is hampered by both algorithmic opacity that impedes external monitoring and difficulty in accessing proprietary user data.

In this paper, we overcome these challenges to measure platform market power through audits of recommendation algorithms. Just as firms exercise market power to increase prices above competitive levels, we model how a video sharing platform induces users to spend time on their platform above competitive levels. To do this, we create a model of user choices over search and watch decisions and platform recommendation strategy conditional on user choices. We develop a novel demand estimation procedure in which we match predicted recommendation strategies to actual recommendations collected through random audits of a video sharing platform. The process allows us to back out the platform's prior over the distribution of demand parameters, run counterfactuals comparing time-maximizing to value-maximizing recommendations under the competitive benchmark. Our revealed preference approach allows us to test for market power without the need to define the market and without the need for user data. In results presented here, we show that the platform responds to several dimensions of user viewing and search choices including the length of videos the user has watched, how much they have searched, the fraction of videos they finish, and whether a user has left the platform. We use a toy model to illustrate the platform's strategic decision, and we present a simplified model that matches platform behavior dynamics. Comparing to utility-maximizing counterfactual, we estimate a small mark-up over user time of around 2%, slightly higher for our sample of older users. Our next steps are to estimate the full model and run these counterfactuals along with estimating bootstrapped standard errors.

Our paper relates to literature on matching in two-sided markets, audit studies, and consumer search. Two sided markets have long been a focus of study (Rochet and Tirole, 2006), but two-sided markets in attention have only recently come into focus (Chen, 2022). Smaller is the literature on the role that recommendation algorithms play in shaping incentives, (Jain and Qian, 2021; Bhargava, 2022; Qian and Jain, 2024) and ours is the first study to bring empirics to this setting. Audits have long played a role in economics research (Ayres and Siegelman, 1995; Bertrand and Mullainathan, 2004), and algorithmic

[1] Stop Addictive Feeds Exploitation (SAFE) Act
[2] Protecting Our Kids from Social Media Addiction Act

audits are increasingly common (Metaxa et al., 2021). While most such audits vary a single dimension to test for a specific behavior, we expand the audit to multiple dimensions. This allows us to map out algorithmic responses, providing interpretability to these otherwise opaque systems. We also contribute to the literature on digital overuse and well-being (Allcott et al., 2020; Braghieri et al., 2022; Allcott et al., 2022), providing a mechanism through which social media platforms can induce overuse in a model of rational user behavior. Relating to the literature on personalized pricing, price discrimination, and consumer welfare (Lambrecht et al., 2012; Dubé and Misra, 2023; Buchholz et al., 2020), we show how social media recommendation algorithms act as implicit price discrimination mechanisms. Finally, we contribute to the literature on digital market power, developing a novel conduct testing procedure that does not require user data (Calvano and Polo, 2020; Allcott et al., 2025).

2 Setting and Data

Our setting is YouTube, the largest video sharing platform, containing billions of hours of content and employing a recommendation algorithm with an explicitly defined architecture and objective: to maximize the expected watch time of a user's next video (Covington et al., 2016). Since publishing their paper detailing the algorithm's objectives, some updates to the algorithm have been made, in particular, using surveys to analyze quality as well as demoting content that the platform deems sensationalist and making tweaks to promote fairness across protected groups[3]. For our purposes, we assume that these tweaks are independent of video length, reasonable for these content-based changes. Finally, we have found no evidence that the platform's objective extends beyond the watchtime of the single next video, allowing us to rule out behavior like experimentation to learn users' preferences.

On each video page, the platform offers a set of recommended 'next videos.' In our audits, we select from this list at random, collecting subsequent recommended videos. We collect our data by first visiting a video and watching for a randomly determined amount of time, we then collect the recommended 'next videos' and select a video at random from that list, watching that video for a randomly determined amount of time. At a randomly determined point, we leave the platform and return, collecting the next set of recommended videos. These starting videos were selected from the platform's list of 'Trending' videos at the start of each audit. This ensured that the largest number of users would be watching these videos and that the platform would have the least informative priors on their demand parameters and thus react more to user choices. Because the platform maximizes the expected watch time of the next video, we can rule out complex strategies like experimentation in our model of the platform. Still, our model will allow us to test the robustness of our results in varying the platform objectives over multiple videos and watch sessions. Table 1 contains our video sample, with title along with length, views, and

[3] YouTube

an indicator for whether we use this video in our 'Teen Sample.' In estimating our model and testing for conduct in this sample, we will be able to test for whether the platform has more or less market power among younger viewers. Here we define markets as markets for the attention of the viewers of particular videos. While it is true that these samples do not cleanly separate age, we need only that our Teen Sample contains viewers that are younger than our non-Teen Sample on average to test for differences in market power by age.

3 Model

In the following sections, we build from a toy model describing platform beliefs, updating, and strategy to a simplified model over fewer dimensions and to our full model of search.

3.1 Toy Model

To illustrate the platform's recommendation strategy decision, we consider a simplified model over two demand dimensions—preference for length (short or long) and outside option value (high or low)—and two history dimensions—average lengths of past watched videos (short or long) and platform exit (left or stayed). We consider four user types:

$$\theta_1 = (d_{short}, V^o_{low})$$
$$\theta_2 = (d_{long}, V^o_{low})$$
$$\theta_3 = (d_{short}, V^o_{high})$$
$$\theta_4 = (d_{long}, V^o_{high})$$

Who are distributed among the population according to $P_\theta(\theta)$, who make different choices that lead to different likelihoods of watch histories, $P_{h|\theta}(\theta)$, and who will respond to recommended video lengths with different expected watchtimes, $\mathbb{E}[t(r)|\theta]$. As an example, the user that prefers short videos and has a low outside option, θ_1, may have a history distribution of

$$P_{h|\theta_1}(short, left) = 0$$
$$P_{h|\theta_1}(long, left) = 0$$
$$P_{h|\theta_1}(short, stayed) = 0.5$$
$$P_{h|\theta_1}(long, stayed) = 0.5$$

and expected watchtimes of

$$\mathbb{E}[t(short, short)|\theta_1] = short$$
$$\mathbb{E}[t(short, long)|\theta_1] = short$$
$$\mathbb{E}[t(long, long)|\theta_1] = long$$

Whereas the user that prefers short videos and has a high outside option may have the following history distribution and expected watchtimes:

$$P_{h|\theta_3}(short, left) = 0.5$$
$$P_{h|\theta_3}(long, left) = 0$$
$$P_{h|\theta_3}(short, stayed) = 0.5$$
$$P_{h|\theta_3}(long, stayed) = 0$$
$$\mathbb{E}[t(short, short)|\theta_3] = short$$
$$\mathbb{E}[t(short, long)|\theta_3] = short$$
$$\mathbb{E}[t(long, long)|\theta_3] = 0$$

The second user has left the platform at a higher rate and will leave if presented with two long videos, owing to their higher outside option value. Because users who leave the platform have no watchtime, we construct the history distribution through multiple simulations.

The platform recommends two videos, each either short or long. In the case where all demand parameter combinations are equally likely, the platform responds to user histories as follows

$$r(short, stayed) = (long, long)$$
$$r(short, left) = (short, long)$$
$$r(long, stayed) = (long, long)$$
$$r(long, left) = (long, long)$$

maximizing the time of low outside option and users with preferences for longer videos. The platform offers a shorter video to the user who has left to avoid the user leaving the platform and having a watchtime of zero. In the case where the user with preferences for short videos and high outside option is exceedingly rare in the population, the platform may respond with

$$r(short, stayed) = (long, long)$$
$$r(short, left) = (long, long)$$
$$r(long, stayed) = (long, long)$$
$$r(long, left) = (long, long)$$

given the lower base rate of users with preferences for shorter videos and higher outside option. In this way, the platform's prior over user demand parameters affects its response to user choices.

User history is a function of recommendations, but we simplify our state space and avoid the need to define equilibria over strategies by taking the initial strategy the platform employs when the user has no watch history as given and then calculating conditional histories using that initial recommendation strategy.

The value of the outside option governs platform market power. Marginal costs are negligible in digital settings like this one such that under perfect competition with free entry, users have higher outside option values. In particular, the outside option value is equal to the maximum expected value the user gets from the platform. Our model allows us to estimate this value and compare to the counterfactual where the platform maximizes expected value instead of watchtime. That difference is our measure of market power, and provides market power quantification in terms of both utility and user time. Performing this estimation for different audiences (i.e. markets for attention) allows us to compare conduct between these markets as well as make inferences about the source of market power (e.g. barriers to entry, switching costs, network effects).

3.2 Simplified Continuous Model

As a proof of concept, we present and solve a simplified continuous model over two user demand parameters: preferences over video length, γ, and their outside option, V^O. The user, presented with a choice over 20 videos will fully watch the video that maximizes their utility or select the outside option and leave the platform:

$$v^* = \max\{u(d_1, \epsilon_1), u(d_2, \epsilon_2), ..., u(d_{20}, \epsilon_{20}), V^O\}$$
$$u(d, \epsilon) = \frac{\gamma}{2}d^2 + \epsilon d$$

Where d is the length of the video, $\gamma < 0$ reflects decreasing marginal utility in video length, and ϵ represents idiosyncratic match value. This term represents the additional matching that the platform does on all other video characteristics other than length, including content. We normalize $\epsilon \sim N(0, 1)$ and hold this distribution fixed in our counterfactual, and video lengths d are drawn from a random normal distribution with mean μ and variance σ_d^2. The platform's decision is the value $r = \mu$, and we estimate $\hat{\sigma}_d^2$ from the data. This represents the constraint on the recommendation strategy from the platform performing matching on other content dimensions. We assume that this constraint is independent of these other content dimensions.

The result of user choice is a history vector, $\mathbf{h} = (d, o)$, where d is the length of the video chosen by the user and o is an indicator for whether the user has left the platform. Simulating this choice many times given the platforms initial recommendation strategy, we construct the conditional probability of any history given demand parameters, $P_{h|\theta}(h)$. We again simplify our estimation by taking the initial platform as given and computing conditional histories given that strategy. Simulating over recommendation strategies, we construct the expected watchtime of a user over demand parameters and recommendation strategies, $\mathbb{E}[t(r)|\theta]$.

Our model of the platform allows us to back out the implicit distribution, $P_\theta(\theta)$, over the demand parameters $\theta = (\gamma, V^O)$. The objective of the platform is to use a recommendation strategy, r, that maximizes the expected watch time of the next video given their history, $\mathbb{E}[t(r)|\mathbf{h}]$. We model the platform as having a prior over demand parameters $P_\theta(\theta)$ that they update with user history $\mathbf{h}$ to form a posterior. That posterior is then used to create a recommendation strategy to maximize expected watchtime:

$$
\begin{aligned}
r(\mathbf{h}, P_\theta(\theta)) &= \arg\max_r \mathbb{E}[t(r)|\mathbf{h}] \\
&= \arg\max_r \int \mathbb{E}[t(r)|\theta] \frac{P_{\mathbf{h}|\theta}(\mathbf{h}) P_\theta(\theta)}{\int P_{\mathbf{h}|\theta}(\mathbf{h}) P_\theta(\theta) d\theta} d\theta
\end{aligned}
$$

We use our demand model to estimate $\mathbb{E}[t(r)|\theta]$ and $P_{\mathbf{h}|\theta}(\mathbf{h})$ for all values of r, $\mathbf{h}$, and θ. Then given $P_\theta(\theta)$, we calculate $r(\mathbf{h}, P_\theta(\theta)$. To estimate $P_\theta(\theta)$, we match our model predictions to observed recommendations from algorithm audits, $r_{obs}(\mathbf{h})$:

To solve this model, we minimize the mean squared error between our modeled recommendations and the recommendations we observe in our audits.

3.3 Full Model

Our fuller model of user demand and platform recommendations incorporating user search as a central element and separates the video choice from watch amount decision. The user first performs a sequential search along the list of recommended videos, comparing their outside option, V^O, the highest-value video they have found thus far, V^{max}, and the expected value of searching:

$$
\begin{aligned}
V^s(s, V^{max}; \mathbf{h}) = \max\{&V^O, \\
&V^{max}, \\
&-\alpha + \mathbb{E}[V^s(s+1, \max\left\{V^{max}, V^w(\epsilon, d; \mathbf{h})|s+1, \mathbf{h}\right\}])\}
\end{aligned}
$$

Where s is the number of videos searched, V^w is the watch-value of a video, ϵ is an idiosyncratic match value, d is the length of a video, α is the disutility of search and $\mathbf{h}$ represents the user's past watch history. ϵ represents the match value along all dimensions other than video length, including matching based on user content preferences. The user chooses how much of a video to watch based on the following value function:

$$
V^w(\epsilon, d; \mathbf{h}) = \max_t \{\epsilon t + \gamma t^2 + \phi[t = d]d + e^{-\rho t}\left(V^s(1, 0; \mathbf{h})\right)\}
$$

Where γ reflects diminishing marginal utility of a video a la Allcott et al. (2022), ϕ governs the value to the user of finishing a particular video, and the final term giving the value of the next search-watch decision following the conclusion of the video.

We consider a recommendation strategy to consist of the length of the first recommended video, μ, and the way in which recommended video length changes over the list of recommended videos, ν, such that $r = (\mu, \nu)$. The added change along the list of recommended videos allows the platform to engage in screening, providing shorter videos first if demand for shorter videos is correlated with higher search cost, for example.

3.4 Identification and Estimation

To ensure that we can identify the platform's prior, we match demand parameters to history dimensions. In particular, we use user search choices to identify α, choices over video lengths to identify γ, choices of how much of a given video to watch to identify ϕ, and the choice to leave the platform to identify V^O. The intermediate step in our model, the vector of recommendations, is of fewer dimensions than our parameters. As such, there are several values of these parameters that will result in the same maximizing strategy. Our moments are not simply the recommendations we observe but the map from user histories to recommendations. As such, matching these history dimensions to demand parameters ensures we observe sufficient variation in this map. Our formal identification argument is presented in a brief appendix.

With our priors and model, we measure platform market power as the change in welfare, watchtime, and strategy when the platform maximizes user utility rather than expected watchtime. By holding ϵ constant, we hold constant the many ways in which the recommendation algorithm matches users to content other than video length. Our main counterfactual compares user time on the platform and welfare under a value-maximizing recommendation algorithm that still updates its posterior over demand parameters given user choices. In repeating this process for different groups of videos in our sample, we identify how this market power differs among these audiences. In particular, audiences that contain younger viewers. We can also compare counterfactuals where the platform has sufficiently rich content to choose maximizing strategies ($\sigma_d^2 = 0$) as well as when the platform has sufficient data to directly identify demand parameters.

In the results presented today, we solve our models on grids for parameters and histories. To simplify the prior in our simple model we parametrize the prior using 5 knots for each γ and V^O and add a regularization term to our loss function that penalizes priors that result in strategies that do not vary with user histories.

4 Results

4.1 Audit

First, we show that randomized user choices over these history dimensions predict subsequent recommendations from the platform. To do this, we regress the length, d, of a particular video v on the list of recommended videos on several dimensions of user history interacted with the position of that video on the list of recommended videos, s. Our history dimensions are the average length of videos the user has watched, $\overline{d}_h$, the average position on the list of recommended videos the user has selected $\overline{s}_h$, the fraction of videos the user has finished $\overline{f}_h$, and whether the user has left the platform and returned, o_h:

$$\begin{aligned} d = &\alpha_1 \overline{d}_h + \alpha_2 \overline{s}_h + \alpha_3 \overline{f}_h + \alpha_4 o_h \\ &+ \beta_1 s \times \overline{d}_h + \beta_2 s \times \overline{s}_h + \beta_3 s \times \overline{f}_h + \beta_4 s \times o_h \\ &+ \gamma_0 + \gamma_v + \gamma_p + \epsilon \end{aligned}$$

Where γ_0 is a constant and γ_r and γ_p are fixed effects for the initial video and number of videos watched, respectively. All continuous variables are log-transformed except for the fraction of videos the user has finished, $\overline{f}_h$.

Figure 1a) presents the estimates of the α coefficients, and Figure 1b) presents the estimates for the β coefficients. There are 14 total sets of estimates, 13 for each individual starting video and one for the pooled sample. The pooled estimates include the full video sample, but we omit individual estimates from starting videos with insufficient variation for estimation; in particular, the fraction of videos a user has finished is often non-zero, resulting in errors of multicollinearity. We see this reflected in the noisiness of the estimates on video finishes, but we see that the platform increases the length of recommended videos in response to a user finishing more videos in the pooled estimates. The platform responds consistently to user choices over video lengths, increasing the length of recommended videos in response longer videos in a user's history. We also see that the platform responds to the user leaving by offering shorter videos upon their return. These two points are consistent with the platform recommending slightly longer videos than the user would prefer, with the platform positively updating its posterior over the value the user places on finishing videos when they choose to finish a video, subsequently recommending longer videos. On the other hand, the platform updates its belief over the user's outside option when they leave, providing shorter videos upon their return.

As an initial test of varying conduct by video audience (attention market), we test whether the platform responds differently to the user leaving the platform for the teen and non-teen sample. When controlling for other history characteristics, we see no statistically significant difference in the platform's

updating on the user's outside option. While this provides initial evidence that the platform does not have more market power among younger viewers, our model is necessary to fully capture the correlations between demand parameters as well as to quantify the extent of market power in welfare and watchtime terms.

4.2 Simplified Model Results

Figure 2 shows how user choices vary with our demand parameters. As expected, users with lower values of γ (steeper diminishing utility in video length) prefer shorter videos. Users with higher outside option values are more likely to leave the platform; however, they watch longer videos when they stay. Figure 3a shows the distribution of our optimal prior, and Figure 3b shows how our model recommendations align with observed recommendations. In this figure, each dot represents an observed user history.

Table 2 presents our conduct test. We repeat the estimation for the Teen Sample and non-Teen Sample and compare the optimized parameters as well as strategies, watchtimes, and welfare under the status quo and competitive regimes. Here we see that users from our teen sample have a slightly higher expected value of γ, preferring slightly longer videos. They also have a higher outside option value. Both watchtime and utility differences across both samples are relatively small: on the order of 2-3%, indicating a small degree of platform market power. As expected, the non-Teen Sample with lower expected value of the outside option is faced with a slightly higher mark-up on their time (mark-down on welfare). These preliminary results are consistent with younger users having access to a greater array of outside options in consuming online content, and their estimated preference for slightly longer videos is consistent with them choosing other platforms for short-form content and coming to YouTube for longer-form content. Additional conduct comparisons based on differences in audience engagement can shed light on the role of network effects in this market power, and counterfactuals that vary the ability of the platform to target users can shed light on the role of data to train recommendation algorithms.

4.3 Full Model Simulations

Figure 4 presents results from our simulation exercise, where we show how the watchtime-maximizing strategy varies over the user's value of the outside option. We can see that as the value of the user's outside option increases, the recommendation strategy that maximizes their watch time tends toward shorter videos and aligns more closely with the recommendation strategy that would maximize their value, reflecting the inability of the platform to push users with higher outside option values toward longer videos. We also document in our simulations that there is a unique watchtime-maximizing strategy that varies across our parameter values, necessary for our estimation procedure.

5 Conclusion

We develop a novel demand estimation strategy to estimate the welfare loss from recommendation algorithms that maximize user engagement. The procedure, combining structural modeling and multidimensional audits, allows us to use responses from a recommendation algorithm to synthetic user choices to estimate the distribution of demand parameters and measure implied market power. We show that the platform's recommendations respond to user choices in ways consistent with maximizing user watchtime and that the watchtime-maximizing recommendation strategy deviates from the value-maximizing strategy. In preliminary estimations using a simplified model we show that the platform exhibits a small degree of market power, that is slightly lower for our sample of younger viewers. Estimation of our full model and comparison between more video audiences will allow us to measure market power among these audiences more robustly and allow us to make inferences about the sources of platform market power.

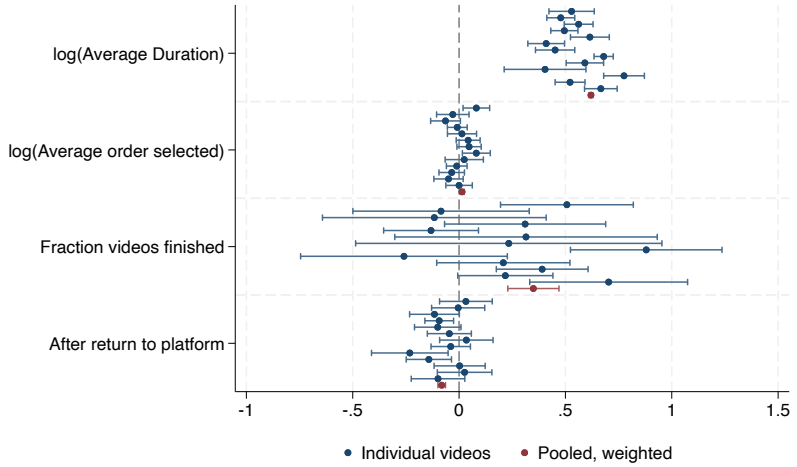
In our audits, we find also that the platform is responsive to user choices along dimension of content that we construct through word-vectorized video titles, for instance. By focusing solely on video length, our approach provides a lower bound on the engagement-value gap, as additional content-specific factors such as controversial content would likely further increase the divergence between engagement maximization and user welfare. We leave the measurement and modeling of these additional factors for future work. Our methodology can be used to monitor algorithmic systems as well as provide a framework for antitrust policies in settings without formal prices.

Table 1: Video Sample

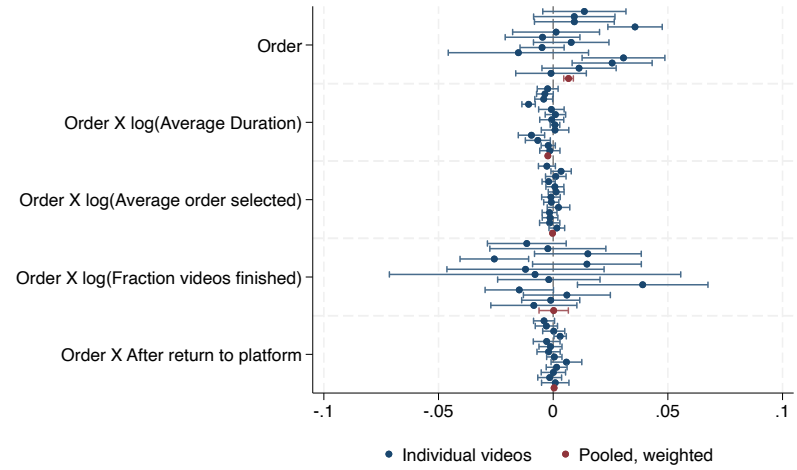
Title	Length (min)	Views (mil)	Teen Sample
EXPOSING MY PRIVATE TIKTOK LIKES.	23	1.00	X
The Results of our IVF	12	0.67	
Being A Professor At Streamer University	64	1.98	
The new Hay Day letters are here!	2.7	1.41	
4-Year-Old Miss Shirley Stuns JHud & Niecy Nash-Betts with Her Line Dancing!	5.9	0.85	
Schedule 1	50	4.08	
CboysTV Cheap Car Challenge (Part 2)	73	2.13	
#BlackChicagoBeLikePart119 Part 34	27	0.67	
Beat Neymar, Win $500,000	28	56.4	X
We Snuck Into a Waterpark OVERNIGHT	31	8.25	X
10 NEW Costco Deals You NEED To Buy in April 2025	17	0.90	
Athletic Club vs. Man. United: Extended Highlights...	11	0.73	
Superman — Official Trailer — DC	3.0	15.6	
How One Company Secretly Poisoned The Planet	54	3.94	
DDG - what i prefer (Official Music Video)	2.1	1.25	
Imperial Takeover — Fortnite Galactic Battle Gameplay Trailer	0.57	0.82	X
Loving Mother & Herd Do All They Can to Save Baby Buffalo from Lions	2.58	1.72	
The April 28, 2025 Severe Weather Coverage, As It Happened...	559	1.48	
WAR 2 — Official Teaser...	1.57	18.6	
Joey Bada$$ - MY TOWN feat. Loaded Lux (Official Audio)	5.4	0.82	
Indiana Fever Put on a SHOW ... FULL Game Highlights	7.9	1.15	
I Spent $5,000,000 So You Can Go To Space For FREE	21.2	5.72	X
I Tested 1-Star Hotels (finale)	42	3.90	
Wednesday: Season 2 — Official Teaser Trailer — Netflix	2.3	8.61	
Judge weighing throwing out Sean "Diddy" Combs trial testimony	4.2	2.92	
Lego Car Vs Jet Car Soapbox Race	13.4	1.76	X

Figure 1: Audit Results

(a) Baseline Coefficients $\alpha_1, \alpha_2, \alpha_3, \alpha_4$



(b) Interaction Coefficients $\beta_1, \beta_2, \beta_3, \beta_4$



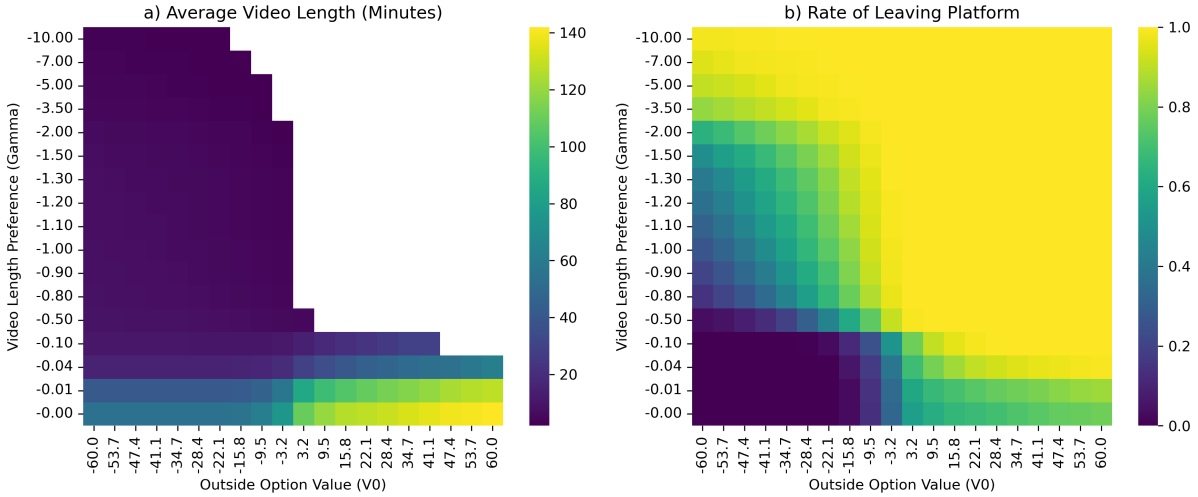
Notes: Coefficient estimates from regressions on output from audits for each of eight starting videos (each in blue) and a single pooled regression (red). The individual audits contain 250, 227, 259, 99, 91, 102, 204, and 105 audits with 18,618; 16,846; 19,028; 8,195; 7,604; 8,259; 15,319; and 8,604 total observations, respectively. Bars represent 95% confidence intervals constructed using standard errors clustered at the level of recommendation set.

Table 2: Leave-Return Effect by Teen Sample

	(1)	(2)
After Leave/Return=1	-0.078***	-0.080***
	(0.014)	(0.009)
Teen Sample=1	0.045***	-0.060***
	(0.016)	(0.012)
After TikTok Page × Teen Sample	-0.043***	0.012
	(0.013)	(0.010)
Observations	462800	462800
Start Video, # Watched FE	Yes	Yes
History Controls	No	Yes

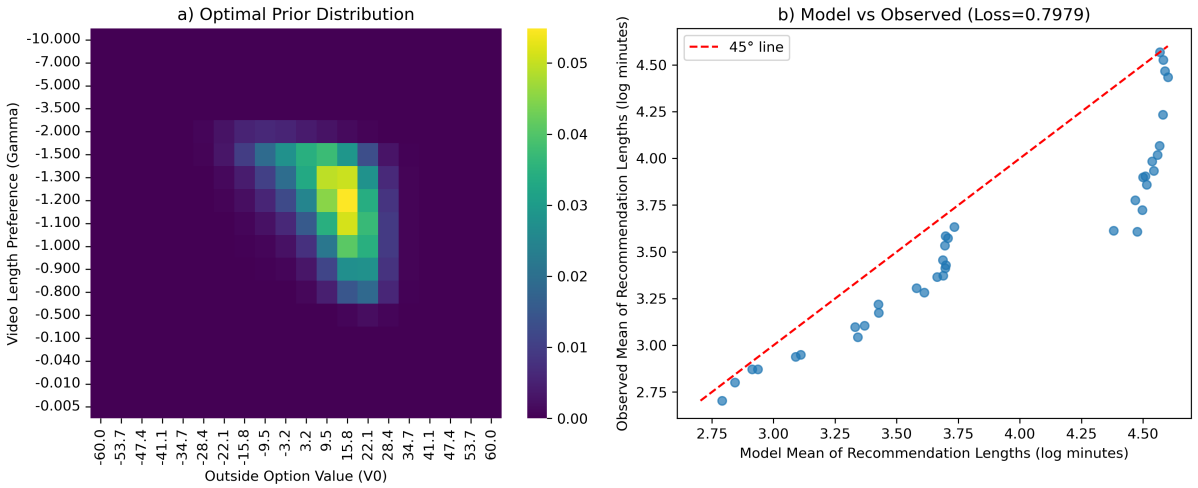
Notes: Coefficient estimates from regression on recommended video duration. History controls include length of videos, search amount, video finish rate, and whether user has left. Standard errors clustered at the run-pull level robust to heteroskedasticity are presented in parentheses.

Figure 2: User History



Notes: Simulations of the simplified continuous model showing average length (duration) of watched videos and the rate of leaving the platform over the demand parameters γ and V^O.

Figure 3: Model Fit

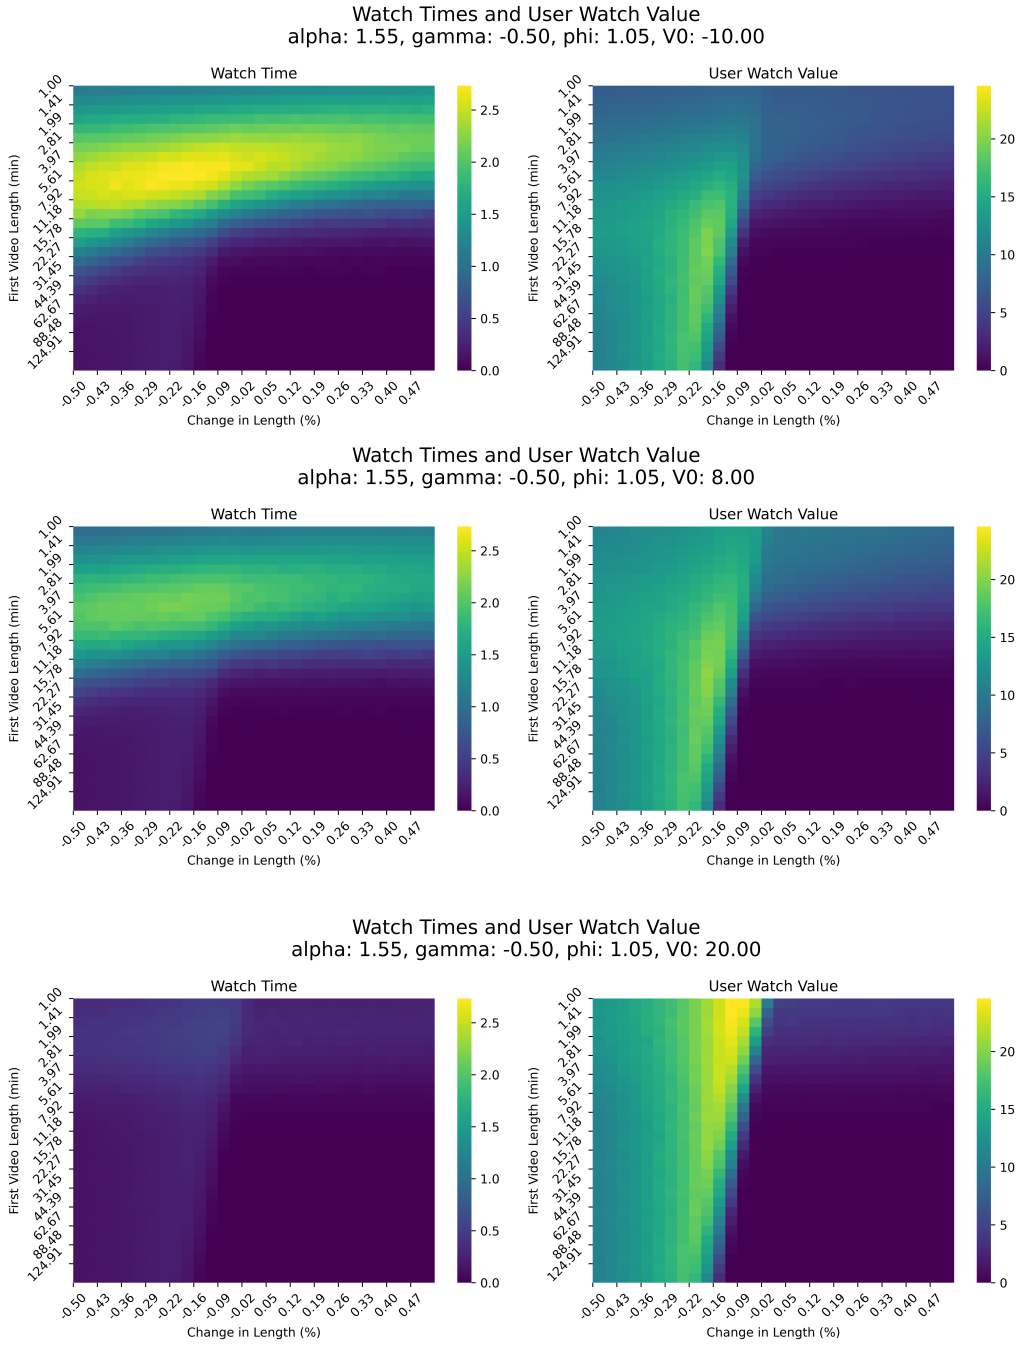


Notes: Estimated distribution over demand parameters (left) and fit of observed and modeled recommendations (right).

Table 3: Conduct Test by Sample

Item	Counterfactual	Teen Sample	non-Teen Sample
loss		0.994	0.798
$\mathbb{E}[\gamma]$		-0.60	-1.21
$\mathbb{E}[V^o]$		21.91	11.25
$corr(\gamma, V^O)$		0.74	0.48
μ	Status Quo	3.85	3.88
(log minutes)	Competitive	3.63	3.63
Watchtime	Status Quo	16.76	18.35
(minutes)	Competitive	16.41	17.93
	% Difference	-2.10	-2.27
Welfare	Status Quo	27.21	28.22
	Competitive	27.83	29.19
	% Difference	2.27	3.40

Figure 4: Simulations



Notes: Simulations vary the value of the outside option faced by the simulated user. The demand model is simulated under static recommendations and search and watch decisions for a single video.

References

Allcott, Hunt, Juan Camilo Castillo, Matthew Gentzkow, Leon Musolff, and Tobias Salz, "Sources of market power in web search: Evidence from a field experiment," Technical Report, National Bureau of Economic Research 2025.

_ , Luca Braghieri, Sarah Eichmeyer, and Matthew Gentzkow, "The Welfare Effects of Social Media," *American Economic Review*, March 2020, *110* (3), 629–676.

_ , Matthew Gentzkow, and Lena Song, "Digital Addiction," *American Economic Review*, 2022, *112* (7), 2424–2463. Publisher: American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Ayres, Ian and Peter Siegelman, "Race and Gender Discrimination in Bargaining for a New Car," *The American Economic Review*, 1995, *85* (3), 304–321. Publisher: American Economic Association.

Bertrand, Marianne and Sendhil Mullainathan, "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination," *American Economic Review*, September 2004, *94* (4), 991–1013.

Bhargava, Hemant K., "The Creator Economy: Managing Ecosystem Supply, Revenue Sharing, and Platform Design," *Management Science*, July 2022, *68* (7), 5233–5251.

Braghieri, Luca, Ro'ee Levy, and Alexey Makarin, "Social media and mental health," *American Economic Review*, 2022, *112* (11), 3660–3693. Publisher: American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

Buchholz, Nicholas, Laura Doval, Jakub Kastl, Filip Matějka, and Tobias Salz, "Personalized Pricing and the Value of Time: Evidence from Auctioned Cab Rides," May 2020.

Calvano, Emilio and Michele Polo, "Market power, competition and innovation in digital markets: A survey," 2020.

Chen, Daniel, "The market for attention," *Available at SSRN 4024597*, 2022.

Covington, Paul, Jay Adams, and Emre Sargin, "Deep Neural Networks for YouTube Recommendations," in "Proceedings of the 10th ACM Conference on Recommender Systems" ACM Boston Massachusetts USA September 2016, pp. 191–198.

Dubé, Jean-Pierre and Sanjog Misra, "Personalized Pricing and Consumer Welfare," *Journal of Political Economy*, January 2023, *131* (1), 131–189.

Guess, Andrew M., Neil Malhotra, Jennifer Pan, Pablo Barberá, Hunt Allcott, Taylor Brown, Adriana Crespo-Tenorio, Drew Dimmery, Deen Freelon, Matthew Gentzkow, Sandra González-Bailón, Edward Kennedy, Young Mie Kim, David Lazer, Devra Moehler, Brendan Nyhan, Carlos Velasco Rivera, Jaime Settle, Daniel Robert Thomas, Emily Thorson, Rebekah Tromble, Arjun Wilkins, Magdalena Wojcieszak, Beixian Xiong, Chad Kiewiet de Jonge, Annie Franco, Winter Mason, Natalie Jomini Stroud, and Joshua A. Tucker, "How do social media feed algorithms affect attitudes and behavior in an election campaign?," *Science*, July 2023, *381* (6656), 398–404. Publisher: American Association for the Advancement of Science.

Jain, Sanjay and Kun Qian, "Compensating Online Content Producers: A Theoretical Analysis," *Management Science*, November 2021, *67* (11), 7075–7090.

Lambrecht, Anja, Katja Seim, Naufel Vilcassim, Amar Cheema, Yuxin Chen, Gregory S. Crawford, Kartik Hosanagar, Raghuram Iyengar, Oded Koenigsberg, Robin Lee, Eugenio J. Miravete, and Ozge Sahin, "Price discrimination in service industries," *Marketing Letters*, June 2012, *23* (2), 423–438.

Metaxa, Danaë, Joon Sung Park, Ronald E. Robertson, Karrie Karahalios, Christo Wilson, Jeff Hancock, and Christian Sandvig, "Auditing Algorithms: Understanding Algorithmic Systems from the Outside In," *Foundations and Trends® in Human–Computer Interaction*, November 2021, *14* (4), 272–344. Publisher: Now Publishers, Inc.

Qian, Kun and Sanjay Jain, "Digital Content Creation: An Analysis of the Impact of Recommendation Systems," *Management Science*, March 2024. Publisher: INFORMS.

Rochet, Jean-Charles and Jean Tirole, "Two-sided markets: a progress report," *The RAND Journal of Economics*, September 2006, *37* (3), 645–667.

We consider a recommendation strategy $r(h)$ for a given prior $P_\theta(\theta)$ as the map from any history $h = \overline{s}, \overline{d}, \overline{f}, o$ to a strategy $r = \mu, \nu$. Identifying the prior distribution using the platform's response to variation in history relies on invertibility from priors, $P_\theta(\theta)$ to maximizing recommendation strategies, $r(h)$. We start with uniqueness of recommendation strategy:

Assumption 5.1. $\mathbb{E}[t(r)|\theta]$ *is continuous and strictly concave in* r *for all* θ

We satisfy this assumption in our case, as changing the recommendation strategy will mechanically change watch times for users. We need to also consider the case in which changing the recommendation strategy causes users to not watch a given video, breaking continuity. For this, we rely on the noise in recommendations and idiosyncratic match value ϵ along with the fact that we are in expectation to get continuity and strict concavity.

Assumption 5.2. *The prior distribution* $P_\theta(\theta)$ *is a probability measure over* θ

Theorem 5.3. *Each prior* $P_\theta(\theta)$ *has a unique maximizing strategy,* $r(h)$

Proof. Consider the prior over the empty history, $h = \emptyset$. $P_{\theta|h=\emptyset}(theta) = P_\theta(\theta)$. Then the platform objective is

$$\int \mathbb{E}[t(r)|\theta] P_{\theta|h=\emptyset}(theta) d\theta = \int \mathbb{E}[t(r)|\theta] P_\theta(\theta) d\theta$$

By assumptions 4.1 and 4.2, this objective function must also be continuous and strictly concave in r and must have a unique maximum over r. □

For the proof in the other direction, we add two additional assumptions:

Assumption 5.4. *For any two different priors* $P_{\theta_1}(\theta)$ *and* $P_{\theta_2}(\theta)$*, there exists at least one history* h^* *such that:*

$$\frac{P_{h|\theta_1}(h^*) P_{\theta_1}(\theta)}{P_{\theta_1}(h^*)} \neq \frac{P_{h|\theta_2}(h^*) P_{\theta_2}(\theta)}{P_{\theta_2}(h^*)}$$

for a set of θ *with positive measure,* Θ

Again in the trivial case, of $h = \emptyset$, we have

$$\frac{P_{h|\theta_1}(h^*) P_{\theta_1}(\theta)}{P_{\theta_1}(h^*)} = P_{\theta_1}(\theta) \neq P_{\theta_2}(\theta) = \frac{P_{h|\theta_2}(h^*) P_{\theta_2}(\theta)}{P_{\theta_2}(h^*)}$$

More generally under normal priors, posteriors are defined as

$$P_{\theta|h}(\theta) \propto L(\theta; h) \cdot \exp\left(-\frac{1}{2} \begin{pmatrix} \alpha - \mu_\alpha \\ \gamma - \mu_\gamma \\ \phi - \mu_\phi \\ V^O - \mu_{V^O} \end{pmatrix}^T \begin{pmatrix} \sigma_\alpha^2 & \sigma_{\alpha\gamma} & \sigma_{\alpha\phi} & \sigma_{\alpha V^O} \\ & \sigma_\gamma^2 & \sigma_{\gamma\phi} & \sigma_{\gamma V^O} \\ & & \sigma_\phi^2 & \sigma_{\phi V^O} \\ & & & \sigma_{V^O}^2 \end{pmatrix}^{-1} \begin{pmatrix} \alpha - \mu_\alpha \\ \gamma - \mu_\gamma \\ \phi - \mu_\phi \\ V^O - \mu_{V^O} \end{pmatrix} \right)$$

Provided two demand parameters have distinguishable effects on histories as reflected in the $L(\theta; h)$ term, the difference in priors will be reflected in the posteriors for at least some values of h.

Theorem 5.5. *Each maximizing recommendation strategy, $r(h)$ has a unique associated prior, $P_\theta(\theta)$*

Proof. 1. By Assumption 1.4, there exists a history h^* such that:

$$\frac{P_{h|\theta_1}(h^*)P_{\theta_1}(\theta)}{P_{\theta_1}(h^*)} \neq \frac{P_{h|\theta_2}(h^*)P_{\theta_2}(\theta)}{P_{\theta_2}(h^*)}$$

for a set of θ with positive measure.

2. We define:

$$g(\theta) = \frac{P_{h|\theta_1}(h^*)P_{\theta_1}(\theta)}{P_{\theta_1}(h^*)} - \frac{P_{h|\theta_2}(h^*)P_{\theta_2}(\theta)}{P_{\theta_2}(h^*)}$$

where $g(\theta) \neq 0$ for a set of θ with positive measure.

3. Consider the first-order conditions for the maximizing recommendation strategy for history h^* and one dimension of r, say μ without loss of generality:

$$\int \frac{\partial \mathbb{E}[t(r)|\theta]}{\partial \mu} \frac{P_{h|\theta_i}(h^*)P_{\theta_i}(\theta)}{P_{\theta_i}(h^*)} d\theta = 0, \quad i = 1, 2$$

Where we'll call this expectation-maximizing value of r μ^{**},

4. Subtracting these equations for $i = 1$ and $i = 2$:

$$\int \frac{\partial \mathbb{E}[t(r)|\theta]}{\partial \mu} g(\theta) d\theta = 0$$

5. Now we consider the set Θ of nonzero measure, where we have that $g(\theta) \neq 0$ and $\frac{\partial \mathbb{E}[t(r)|\theta]}{\partial \mu} \neq 0$ for values of γ such that $\mu^*(\gamma) \neq \mu^{**}$ such that

$$\int \frac{\partial \mathbb{E}[t(r)|\theta]}{\partial \mu} g(\theta) d\theta \neq 0$$

Implying that $g(\theta) \neq 0$ Essentially, with the maximizing value of μ varying over values of γ, the strategy that maximizes the posterior expectation will differ when the posterior differs.

6. This contradicts the equation in step 5, which states that this integral should be zero for all r and for both μ and ν.

Therefore, our initial assumption that there exist two different priors resulting in the same maximizing recommendation strategy must be false.

Thus, under Assumptions 1 and 2, we have proven that each maximizing recommendation strategy has a unique prior associated with it.

□

Having proved uniqueness in both directions, we can conclude that with sufficient data on histories, we will be able to identify the platform's prior distribution over demand parameters for viewers of a particular video.