"Just One More Clip":
Short Videos, Big Self-Control Problems*

RENJIE BAO†
Princeton University

December 8, 2025

Abstract

I study self-control problems in media consumption and their amplification by short-form content. Using microdata from a U.S. short drama series, I show viewers watch 23 episodes (82%) more than intended and overspend by $5.51 (23%). The structural estimation indicates that temptation takes effect over a short horizon, averaging 6.6 minutes per decision, making minute-long videos repeatedly trigger self-control problems. Policy interventions such as default time limits and mandatory breaks can meaningfully improve consumer welfare. Extrapolating the analysis to TikTok underscores the broader relevance of these findings.

JEL codes: D12, D91, I31, L82

*I am grateful to my advisors, Sylvain Chassang, Kate Ho, and Adam Kapor, for their guidance and support, as well as Nick Buchholz, Jan Eeckhout, Faruk Gul, Jakub Kastl, Alessandro Lizzeri, Pietro Ortoleva, Wolfgang Pesendorfer, and Leeat Yariv for many insightful discussions. The paper benefits from valuable feedback from Jason Abaluck, Guy Aridor, Roland Bénabou, Doug Bernhein, Francesco Bilotta, Pierre Bodéré, Eduard Boehm, Federica Carannante, Juan Camilo Castillo, Alex Chan, Allan Collard-Wexler, Jacques Crémer, Janet Currie, Jan De Loecker, Stefano DellaVigna, Mira Frick, Jeff Gortmaker, Hao Huang, Ryota Iijima, Yuyang Jiang, Botond Köszegi, David Laibson, Quan Le, Sanxi Li, Ranie Lin, Ben Lockwood, Rong Luo, Monica Morlacco, Kirby Nielsen, Andrew Olsen, Linda Ouyang, Giacomo Ponzetto, Alejandro Sabal, Yinshan Shang, Charles Sprenger, Philipp Strack, Dmitry Taubinsky, Martin Vaeth, Jo Van Biesebroeck, Frank Verboven, Duncan Webb, Thomas Wollmann, Wei Xiong, Xinmei Yang, Sohye Yoon, Yihao Yuan, Jidong Zhou, Yuci Zhou, and seminar participants at BSE Summer Forum, China Agricultural University, KU Leuven, Princeton, Renmin University of China, SITE, TSE Economics of Platform Seminar, Tsinghua, and Wuhan University. A special thanks to Luyin Zhang, who, with good humor and enthusiasm, has given permission to be acknowledged as an exemplar of the behavior explored in this paper. I am grateful to SD, SG, and J for their assistance with data access, and to Bobbie Goettler and ChatGPT for copy editing.

†renjie.bao@princeton.edu

1 Introduction

Many of us are familiar with the regret that follows giving in to temptation. We promise to stop after just one more episode, beer, or book chapter, only to find hours slipping away. Despite recognizing the harm, resisting proves difficult. This phenomenon, known as the *self-control problem*, has been identified in many contexts such as gym membership (e.g., DellaVigna and Malmendier, 2006) and retirement saving decision (e.g., Laibson, Repetto, Tobacman, Hall, Gale, and Akerlof, 1998). Sadly, social media platforms such as TikTok, YouTube and Instagram exploit this vulnerability by delivering endless streams of short, engaging content that intensifies users' temptation to consume more than intended. The American Psychiatric Association classifies this issue as "internet use disorder," linking it to negative effects on mental health and well-being.[1]

The welfare loss resulting from self-control problems in short-form content consumption could be enormous. For example, in 2023, over 170 million people in the U.S. used TikTok for an average of 58.4 minutes per day—an amount of time equivalent to 7.8% of U.S. GDP annually.[2] Given this scale, quantifying how much time is unintentionally wasted on social media, i.e., the magnitude of efficiency loss, and how much of this inefficiency is driven by short-form content design is crucial. More importantly, understanding the magnitude of this mechanism is essential for effective policy design. Current policy discussions range from time-limit tools to outright social media bans, requiring systematic evaluation that accounts for inefficiencies caused by self-control problems.[3] My paper contributes by analyzing self-control problems through a structural approach, providing insights into both the magnitude of welfare loss and potential policy solutions.

I begin by analyzing a stylized model of temptation in the context of short video platforms, demonstrating how self-control problems emerge and how short-form content exacerbates them. Users derive intrinsic utility from each clip of short videos but also experience a short-term, irrational temptation toward these clips for a limited period, the duration of which varies. The central assumption is that temptation continually shifts toward upcoming clips while users keep watching.[4] As a result, self-control problems arise among tempted users with naïve perception, leading them to consume more content than they would have preferred ex ante. Short-form formats shrink each decision window below the duration of temptation, thereby worsening self-control problems. By contrast, longer content is less addictive, because users with temptation durations shorter than the video length perceive true utility in the later segments, thereby mitigating the effects of irra-

[1] The APA defines "internet use disorders" as a pattern of behavior characterized by an inability to control use, difficulties with personal and professional responsibilities, and continued use despite negative consequences. See Allcott, Braghieri, Eichmeyer, and Gentzkow (2020); Rosenquist, Morton, and Weinstein (2021); Braghieri, Levy, and Makarin (2022); Allcott, Gentzkow, and Song (2022) for more evidence on these negative consequences.

[2] The number of U.S. users is based on TikTok's public statement during the Senate Judiciary Committee Hearing on January 31, 2024. The average daily time spent on TikTok is from a 2024 eMarketer survey. The average hourly earnings in January 2025 were $35.87, according to the Bureau of Labor Statistics (BLS).

[3] For instance, in 2023, TikTok introduced a feature that sets a default 60-minute daily screen time limit for all accounts held by users under the age of 18 (Source: TikTok NewsRoom). In a more extreme move, on November 21, 2024, Australia introduced a bill in parliament to ban social media use—including TikTok—for children under 16 (Reuters, November 20, 2024, "Australia launches 'landmark' bill to ban social media for children under 16").

[4] For example, when a user watches the first video, they feel tempted to watch the second; upon reaching the third, the temptation moves to the fourth, creating an ongoing cycle of impulsive consumption.

tional temptation.[5] I further show that a default time limit can effectively mitigate self-control problems for users whose temptation durations are shorter than the limit, i.e., those who do not ex ante plan to watch beyond it. The model's key innovation is to examine how content length interacts with temptation duration, shaping the magnitude of self-control problems. The underlying economic insight is general: breaking a tempting good into smaller units (e.g., a small pack of cigarettes) can increase demand by exacerbating consumers' self-control problems.

Building on these insights, I analyze a popular short drama series on a mobile platform leading in the U.S. market where self-control problems are salient and broadly representative of behavior on other media. The series consists of 80 one-minute episodes, closely resembling short videos in both length and vertical mobile optimization. Users must purchase non-refundable platform tokens through a non-linear top-up menu and use them to unlock episodes one by one. This platform design enables the identification of ex-ante preferences and self-control problems, which manifest as users making suboptimal top-up decisions and repeatedly purchasing lower-priced token packages.[6] On average, users who purchase at least one package intend to watch for only 28 minutes but ultimately spend 51 minutes on the series. This behavior also results in overpayment, because users could have saved money by purchasing a larger package upfront.[7] Indeed, 39.1% of these users spend more than the rational benchmark, with an average overpayment of $5.51 for the drama series. Moreover, by swapping the smallest token packages with larger ones, I estimate a lower bound on overpayment that abstracts from learning and uncertainty, and find that users buying the smallest packages between episodes 40 and 70 systematically overspend. These findings indicate self-control problems in this setting are widespread, persistent, and economically significant.

I further develop a structural demand model for this short drama by extending the stylized framework to incorporate empirically relevant features. In addition to intrinsic utility and temptation, I introduce rational addiction through habit formation, where consuming the current episode temporarily increases habit stock, thereby enhancing the marginal utility of watching subsequent episodes (Becker and Murphy, 1988). I model drama-watching behavior as a single-agent dynamic discrete-choice problem. Each episode represents a time period in which the user makes two decisions based on their habit stock and token balance. First, if their token balance is insufficient to unlock the current episode, they choose from the menu of token packages. Second, given their updated token balance, they decide whether to unlock the episode. If they choose to unlock, they pay the token price, consume the episode, and proceed to the next period with updated state variables; if they choose not to, they stop watching.

[5] For example, individuals deliberate more carefully before committing to a two-hour movie versus a one-minute clip. This economic mechanism also extends beyond short videos: it explains why people frequently purchase small beer packages despite the cost savings of larger packs.

[6] My setting parallels the literature on health clubs (e.g., DellaVigna and Malmendier, 2006), where researchers infer exercisers' ex-ante preferences from observed membership choices. It differs by offering a richer menu of top-up options, the market shares of which allow for identification of the duration distribution of temptation (or present bias).

[7] Users in my sample spend between $29.99 and $69.96 to watch the full 80-minute drama, suggesting a significant level of addiction. For comparison, an adult ticket for a two-hour movie at the Princeton Garden Theatre cost $13.50 in 2024. The much higher willingness to pay for this short drama highlights the strong addictive nature of the content.

Leveraging the platform's non-linear top-up menu enables identification of not only the magnitude but also the *duration* distribution of temptation. As discussed earlier, users who repeatedly purchase small packages are inferred to suffer from self-control problems, which primarily inform the magnitude of temptation (DellaVigna and Malmendier, 2006). Among these users, those with shorter temptation horizons expect to watch only a few episodes and thus choose cheaper, smaller packages, whereas those with longer horizons anticipate more consumption and are more likely to buy relatively larger ones. The market shares of these differently priced packages further reveal the distribution of temptation durations, which is central to the analysis, as both the severity of self-control problems and the welfare impact of policy interventions depend on how this distribution interacts with content length.

I then take the model to the data using the Simulated Method of Moments, targeting moments related to unlocking activity, top-up package purchases, and habit stock dynamics. On average, users derive a negative intrinsic utility of $-1.2¢$ per one-minute episode relative to their outside option, with heterogeneous tastes exhibiting a standard deviation of $11.4¢$. Temptation effects are substantial, increasing perceived flow utility by $18.5¢$ per minute and lasting on average 6.6 minutes. The duration distribution displays considerable heterogeneity, with a standard deviation of 10.2 minutes and a long right tail from log-normal distribution. Habit-formation utility ranges from 0 to \$1.1 per minute, with each additional episode of habit stock increasing per-episode utility by an average of $1.4¢$ in subsequent episodes.

The estimated model reveals substantial user surplus losses from self-control problems, partly driven by the short episode length. On average, users earn a surplus of \$0.21 from the drama series, reflecting the combined utility from intrinsic value and habit formation, with 49.2% experiencing negative surplus. Eliminating temptation raises average surplus to \$0.74, implying a per-user loss of \$0.53 attributable to temptation. At the individual level, the severity of self-control problems follows a U-shaped pattern with respect to drama valuation, with medium-valuation users most likely to be drawn into videos they do not actually enjoy. Counterfactual simulations show extending episode length from one to seven minutes raises average user surplus to \$0.30, reducing temptation-related inefficiency by 17.0%. I evaluate three policy interventions: removing the lowest-priced top-up package, introducing a default time limit with a small opt-out cost, and imposing a mandatory break at the top-up stage. All three increase user surplus (from \$0.21 to between \$0.27 and \$0.79) while reducing platform profits (from \$5.38 to between \$3.88 and \$5.13).

Finally, I extend my structural analysis to a broader context of short video platforms, such as TikTok, by reverting to the stylized model without payment or habit formation.[8] When each video lasts one minute, the average user derives an hourly surplus of \$0.21, whereas increasing video length to the average YouTube video length (12 minutes) raises hourly surplus to \$1.49. Short formats nearly triple the efficiency loss, as hourly surplus loss due to temptation rises from \$0.70

[8] I abstract from habit formation for two reasons. First, habit formation is not central to the mechanisms analyzed in this paper and appears to be of second-order importance in the context of the application on self-control problems. Second, TikTok videos are less temporally and thematically connected than drama episodes, rendering habit formation less salient. Incorporating habit formation would amplify the platform's addictiveness, suggesting that the consumer surplus loss estimated here represents a lower bound.

to $1.98. I show the optimal length for a default time limit is 12 minutes, raising hourly surplus by $1.31 and reducing temptation-induced surplus loss by 86.2%. A back-of-the-envelope calculation estimates TikTok's monthly user surplus at $1.1 billion, whereas self-control problems result in a $10.2 billion monthly welfare loss. Imposing a 12-minute default time limit could recover $6.8 billion (66.7%) per month. This extension underscores the large-scale welfare implications of my findings, highlighting their policy relevance beyond the specific short drama series analyzed.

This project makes several contributions. First, it introduces a new perspective on self-control problems, emphasizing the *duration* of temptation, and apply this framework to explain why short-form content is particularly addictive for users who lack sophistication. Second, by estimating the distribution of temptation durations, the paper provides novel field-based insights into the understudied question of how long the temptation lasts. Third, it develops an approach to separately identify self-control problems and habit formation, demonstrating that both mechanisms influence user decision-making. This finding contributes to the debate over whether addiction is rational or not. Finally, the paper quantifies the amplified self-control problems and substantial welfare loss caused by short-form content in monetary terms, highlighting the potential for significant policy gains when interventions like default time limit are well designed.

Related literature. Building on the behavioral framework of present bias and self-control problems (Laibson, 1997; O'Donoghue and Rabin, 1999; Gul and Pesendorfer, 2001; DellaVigna and Malmendier, 2004), my model extends the application of these theories to explain addictive behaviors on short-form contents (e.g., Gruber and Köszegi, 2001; Gul and Pesendorfer, 2007). Theoretically, this paper contributes by examining how content length interacts with self-control problems, providing new insights into the role of short format design in exacerbating impulsive consumption.[9] Second, my work provides further evidence of self-control problems on short-form video platforms, which have been documented in other contexts, such as health club memberships (e.g., DellaVigna and Malmendier, 2006) and retirement savings decisions (e.g., Laibson, Repetto, Tobacman, Hall, Gale, and Akerlof, 1998; Laibson, Chanwook Lee, Maxted, Repetto, and Tobacman, 2024).[10] Additionally, interpreting temptation as present bias, this paper contributes to the understudied question of how to define the duration of the present by estimating its distribution in the field, where existing evidence comes largely from laboratory experiments (Augenblick, 2018; Balakrishnan, Haushofer, and Jakiela, 2020).

My work also complements the discussion on digital addiction. It builds on the work of Allcott, Gentzkow, and Song (2022), who use a field experiment to estimate a model incorporating habit formation, self-control problems, and naïveté. Other studies on digital addiction primarily focus on its consequences (e.g., Vanman, Baker, and Tobin, 2018; Allcott, Braghieri, Eichmeyer, and Gentzkow, 2020; Mosquera, Odunowo, McNamara, Guo, and Petrie, 2020; Braghieri, Levy, and Makarin, 2022; Collis and Eggers, 2022) and the adoption of self-control tools (Hoong, 2021).

[9] This insight connects to the concept of "skewness preference in the small," as proposed by Ebert and Strack (2015) in the context of prospect theory and naïveté.

[10] The existence of money payment makes the identification of the time-inconsistent preference possible in my setting. See Strack and Taubinsky (2021) for a complete discussion.

This paper contributes to the literature by offering the first revealed-preference evaluation of digital addiction using field data, which addresses sample-selection issues, abstracts from strategic behaviors by experiment subjects, captures users' willingness to pay for different channels of addiction based on real monetary choices, and enables evaluation of counterfactual policies.

More broadly, the paper contributes to the literature on addiction across various contexts. By distinguishing between habit formation and self-control problems, I contribute to the debate on whether addiction is rational (e.g., Spinnewyn, 1981; Becker and Murphy, 1988; Orphanides and Zervos, 1995) or irrational (e.g., Gruber and Köszegi, 2001). This paper also complements studies of addiction in other markets, such as cigarettes (Chaloupka, 1991; Becker, Grossman, and Murphy, 1994; Giné, Karlan, and Zinman, 2010), alcohol (Cook and Moore, 2002; Baltagi and Geishecker, 2006), drugs (Gul and Pesendorfer, 2007; Maclean, Mallatt, Ruhm, and Simon, 2020), and sugar-sweetened beverages (Zhen, Wohlgenant, Karns, and Kaufman, 2011).

Finally, this paper complements the literature on entertainment goods. Experimental evidence from recent studies has documented key features of the short-video industry, including its addictiveness (Gao, Gao, Meng, and Yu, 2024) and content-supply dynamics (Xia, 2025). The context of short videos also relates to studies on other types of entertainment goods, such as movies (Michalopoulos and Rauh, 2024), cable TV (Crawford and Yurukoglu, 2012; Crawford, Lee, Whinston, and Yurukoglu, 2018), and video games (Lee, 2012). In addition, my work contributes to the literature on social media and digital platforms (e.g., Allcott, Braghieri, Eichmeyer, and Gentzkow, 2020; Liu, Sockin, and Xiong, 2020; Aridor, Jiménez-Durán, Levy, and Song, 2024; Beknazar-Yuzbashev, Jiménez-Durán, and Stalinski, 2024).

The remainder of the paper is organized as follows. Section 2 presents the stylized model to illustrate the economic point of how short-form contents amplify self-control problems. In section 3, I provide background on the short drama industry and document evidence of self-control problems. Section 4 introduces the demand model of short drama series. I estimate the model in section 5 and discuss the results in section 6. Section 7 extends my analysis of short dramas into a general short video platform like TikTok. Section 8 concludes.

2 Stylized Model of Temptation

I present a stylized model of temptation to demonstrate how could short-form contents amplify self-control problems. I later extend this model into an empirically relevant framework for short drama series in section 4.

2.1 Setup

Consider a platform offering a large set of videos $\{1, \ldots, T\}$, each lasting one minute. Users make decisions at the video level, watching sequentially and deriving a constant intrinsic utility $x \in \mathbb{R}$ from each. The outside option is normalized to 0. Additionally, users experience a *temptation* to

watch such videos, valued at $\kappa > 0$ per minute for a duration of $\chi \in \mathbb{N}$ minutes, which could be microfounded by present bias (e.g., Laibson, 1997), love for novelty (e.g., Modirshanechi, Lin, Xu, Herzog, and Gerstner, 2025), inertia, or other behavioral mechanisms. Let the superscript t be the *current video* when the user is making a decision, and let subscript $\tau \geq t$ denote the *evaluated video*, i.e., the video whose utility is being assessed from the perspective of current video t. The perceived utility from video τ at current time t is thus:

$$\widetilde{u}_\tau^t(x, \chi) = \underbrace{x}_{\text{intrinsic utility}} + \underbrace{\mathbb{1}\{\tau < t + \chi\}\,\kappa}_{\text{temptation } \kappa \text{ with duration } \chi}. \tag{1}$$

For example, starting at current video t, a user with $\chi = 0$ perceives true utility x from each subsequent videos. In contrast, a user with $\chi = 2$ perceives a biased utility of $x + \kappa$ from videos t and $t + 1$, and the true utility x from all remaining videos. For this section, I focus on users with $x \in (-\kappa, 0)$, who inherently dislike the short videos relative to outside options but may still choose to watch them due to temptation.

This temptation framework leads to a systematic overweighting of immediate entertainment relative to long-run utility. Depending on the temporal perspective t, users may perceive the same video τ differently, often overvaluing those in the near future (i.e., when $\tau < t + \chi$). For instance, Table 1a illustrates the preferences of users with $\chi = 2$. Before watching the first video ($t = 1$), they perceive the true utility $\widetilde{u}_3^1 = x$ for the third one ($\tau = 3$), because it lies outside their temptation horizon. However, at $t = 2$, the temptation effect emerges, biasing their perception to $\widetilde{u}_3^2 = x + \kappa$. This model therefore generates similar predictions to the quasi-hyperbolic model of present bias from Laibson (1997) and Gruber and Köszegi (2001), but it fits naturally with the discrete consumption of short videos and yields simpler estimation with the separably additive temptation (Banerjee and Mullainathan, 2010; Allcott, Gentzkow, and Song, 2022).[11]

The contribution of this paper is to examine the *duration* of temptation. When interpreting temptation as present bias, this duration captures the length of the "present" as perceived by a present-biased agent.[12] In section 2.3, I show the short format shrinks agents' decision windows below the temptation duration, thereby amplifying the resulting self-control problem. In my structural analysis, I estimate the distribution of temptation durations using variation induced by a non-linear pricing scheme, providing novel empirical insights.

Finally, my analysis focuses on naïve users who are unaware of the time-inconsistency in their preferences. At any given current video t, they mistakenly believe their future preferences will mirror their current ones, failing to recognize their tastes will shift as they continue watching. In section 3.4, I show empirical evidence on self-control problems, which aligns well with the model prediction based on this naïveté assumption. I further discuss this assumption in section 4.4 and characterize the behavior of sophisticated users in Appendix A.2.

[11] In Appendix A.1, I show how to represent the preference (1) from the quasi-hyperbolic discounting model, which micro founds this irrational temptation.

[12] See section 5.1 of DellaVigna (2018) for a detailed discussion regarding the duration of present bias.

Table 1: Example: Self-control problems with $\chi = 2$ and $x \in (-\kappa, -\kappa/2)$

(a) Baseline

Evaluated video τ	Current video t = 1	2	3	4
1	$\boxed{x+\kappa}$			
2	$\boxed{x+\kappa}$	$\boxed{x+\kappa}$		
3	x	$\boxed{x+\kappa}$	$\boxed{x+\kappa}$	
4	x	x	$\boxed{x+\kappa}$	$\boxed{x+\kappa}$
5	x	x	x	$\boxed{x+\kappa}$
6	x	x	x	x

(b) Comparative static: Video length $n = 4$

Evaluated video τ	Current video t = 1	2	3	4
1	$4x + 2\kappa$			
2	$4x$	$4x + 2\kappa$		
3	$4x$	$4x$	$4x + 2\kappa$	
4	$4x$	$4x$	$4x$	$4x + 2\kappa$
5	$4x$	$4x$	$4x$	$4x$
6	$4x$	$4x$	$4x$	$4x$

Notes: The tables provide examples of a naïve user with intrinsic value $x \in (-\kappa, -\kappa/2)$ and temptation duration $\chi = 2$. The perceived utility $\widetilde{u}_\tau^t$ from watching video τ at perspective t is displayed in the table, with each column representing a current video and each row an evaluated video. Positive utilities are marked with a box, and the highlighted box indicates the videos the user actually watches. Panel (a) illustrates the baseline scenario where the user, facing self-control problems, plans to watch two videos at every decision point but ends up watching all videos. Panel (b) depicts the case with four-minute videos, where the user perceives a flow utility of $4x + 2\kappa < 0$ for each video and thus chooses not to watch.

2.2 Self-control problems

The stylized model generates self-control problems that are summarized in Proposition 1. At current time t, agents with $x \in (-\kappa, 0)$ predict that they will only watch the next χ videos ($\widetilde{u}_\tau^t = x + \kappa > 0$) and then stop ($\widetilde{u}_\tau^t = x < 0$). However, when they reach video $t + \chi$, they are once again tempted to continue watching for another χ videos, revealing a misprediction and self-control problem: although they intend to stop ex ante, they fail to follow through ex post. Table 1a illustrates an example where a user always plans to watch two videos ($\chi = 2$) at each point in time, but ultimately ends up consuming all the videos provided.

Proposition 1 (Self-control problems) *Users with intrinsic value $x \in (-\kappa, 0)$ and temptation duration $\chi \geq 1$ mispredict future consumption and exhibit self-control problems. At any time t, they predict that they will consume only χ additional videos, but ultimately consume all available videos.*

This self-control problem induces surplus loss on the user side. In the benchmark case without temptation ($\kappa = 0$), users watch the videos if and only if they derive positive intrinsic utility $x \geq 0$. Let (x, χ) represent a user, and let $S \subset \mathbb{R} \times \mathbb{N}$ be the set of users who have self-control problems. The *average surplus loss due to temptation* from short videos can be defined as

$$\Delta(S) = -\sum_{\chi \in \mathbb{N}} \left[T \int_{\{x:(x,\chi)\in S\}} x f_{x|\chi}(x)\, \mathrm{d}x \right] P(\chi), \tag{2}$$

where $P(\chi)$ is the probability mass of χ and $f_{x|\chi}(x)$ is the conditional probability density of x. Proposition 1 yields $S^* = (-\kappa, 0) \times \mathbb{N}_+$ for this baseline case, resulting in a positive user surplus loss $\Delta^* := \Delta(S^*) > 0$.

2.3 Short length exacerbates self-control problems

To understand how short video length exacerbates self-control problems, I consider a hypothetical scenario where the platform extends each video to n minutes.[13] Assume users still make decisions at the video level. The perceived flow utility is now defined as:

$$\widetilde{u}_t^t(x, \chi; n) = nx + \min\{n, \chi\}\kappa, \tag{3}$$

where intrinsic value scales with video length n, and temptation is adjusted by the lesser of video length n or temptation duration χ. If the video length is shorter than the temptation duration ($n \leq \chi$), the perceived utility increases proportionally. However, when the video length exceeds the temptation duration ($n > \chi$), users anticipate losing interest before the video ends, which effectively dampens the influence of temptation.

Given length n, users will watch the videos iff they perceive a positive flow utility, that is, $\widetilde{u}_t^t(x, \kappa, \chi; n) \geq 0$. Self-control problems can thus be characterized by the following condition:

$$\max\left\{-\kappa, -\frac{\chi\kappa}{n}\right\} < x < 0. \tag{4}$$

The self-control problem affects all users with $x \in (-\kappa, 0)$ when the video length n is shorter than the temptation duration χ. However, when $n > \chi$, users with low valuations for the videos ($x < -\chi\kappa/n$) will avoid watching. As n increases, the cutoff $-\chi\kappa/n$ approaches zero, indicating the self-control problem gradually diminishes with video length. Table 1b provides an example with $\chi = 2$ and $n = 4$, where users with $x < -\kappa/2$ escape from the self-control problem.

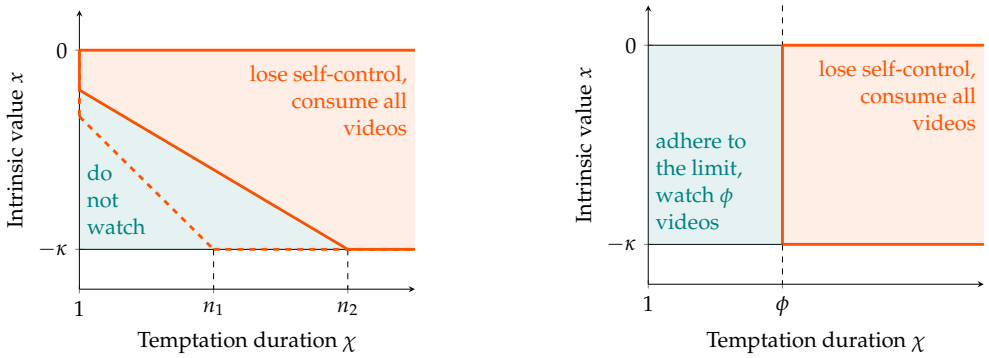
As a result, increasing video length mitigates welfare loss due to temptation. Let S_n^N be the set of users with self-control problems for a given length n, which is characterized by condition (4) and graphically represented by the orange area of Figure 1a. Let $\Delta_n^N := \Delta(S_n^N)$ be the associated average surplus loss due to temptation. For any n_1 and n_2 such that $1 \leq n_1 \leq n_2$, we have:

$$S^* = S_1^N \supset S_{n_1}^N \supset S_{n_2}^N \quad \Rightarrow \quad \underbrace{\Delta(S^*)}_{\Delta^*} \geq \underbrace{\Delta(S_{n_1}^N)}_{\Delta_{n_1}^N} \geq \underbrace{\Delta(S_{n_2}^N)}_{\Delta_{n_2}^N}.$$

When n increases, more users are exempt from the self-control problem, reducing the surplus loss due to temptation defined in equation (2). Proposition 2 summarizes this comparative static.

Proposition 2 (Shorter format exacerbates self-control problems) *Shorter videos engage more users in the self-control problem, leading to greater surplus loss due to temptation. Formally, for any two video lengths n_1 and n_2 such that $1 \leq n_1 \leq n_2$, it holds that $S_{n_1}^N \supset S_{n_2}^N$ and $\Delta^* \geq \Delta_{n_1}^N \geq \Delta_{n_2}^N$.*

[13] This exercise analyzes the impact of video length on self-control problems by merging n one-minute clips into a single, longer video. A real-world analogue is the U.S. law on "loosies," which prohibits the sale of individual cigarettes and mandates a minimum pack size of 20. Such a policy mitigates the effect of momentary temptation to smoke "just one more" cigarette and compels buyers to more fully evaluate the overall value of consumption.



(a) Longer video length n

(b) Default time limit of ϕ minutes

Figure 1: Illustration: User behavior under counterfactual scenarios

Notes: This figure illustrates user behavior under various counterfactual scenarios, where users are characterized by intrinsic value (x) and temptation duration (χ). Panel (a) presents the comparative static for video length (n) adjustments. Users with self-control problems, depicted in the orange region, watch all videos despite having a negative intrinsic value, whereas users in the teal triangular area choose not to watch any videos. Panel (b) shows user activity under a default screen time limit of ϕ minutes. The orange region represents users who continue watching all videos, whereas the teal area highlights users who adhere to the limit and watch less videos.

2.4 Counterfactual policy

To address the inefficiency caused by temptation, I consider a counterfactual policy that imposes a default screen time limit of ϕ minutes, allowing users to opt out ex ante but not ex post.[14] To break ties, I assume users must pay a small cost ε to opt out. The user response is illustrated in Figure 1b. For users with $\delta \in (-\kappa, 0)$, those with temptation duration $\chi \leq \phi$ will adhere to the limit because they only plan to watch for χ minutes, whereas those with $\chi > \phi$ expect to watch more and thus choose to opt out. The policy improves consumer surplus by mitigating self-control problems among agents with relatively short temptation durations.

Given user behavior under this policy, I define the average surplus loss due to temptation as:

$$\Delta_{\phi}^{TL} := -\sum_{\chi \leq \phi} \left[\phi \int_{-\kappa}^{0} x f_{x|\chi}(x), \mathrm{d}x \right] P(\chi) - \sum_{\chi > \phi} \left[T \int_{-\kappa}^{0} x f_{x|\chi}(x), \mathrm{d}x \right] P(\chi), \tag{5}$$

where users with $\chi \leq \phi$ consume only ϕ minutes of short videos. The baseline corresponds to both $\phi = 0$ (no compliers) or $\phi = T$ (non-binding limit), in which case $\Delta_0^{TL} = \Delta_T^{TL} = \Delta^*$. The improvement in surplus is given by:

$$\Delta^* - \Delta_{\phi}^{TL} = \sum_{\chi \leq \phi} \left[(T - \phi) \int_{-\kappa}^{0} x f_{x|\chi}(x), \mathrm{d}x \right] P(\chi) > 0, \quad \forall \phi \in (0, T) \tag{6}$$

highlighting a trade-off in choosing ϕ: increasing ϕ expands policy coverage (i.e., the set $\chi \in (0, \phi]$), but reduces the surplus gain per complier, since each watches more content. Proposition 3

[14] A similar policy was implemented by TikTok in 2023, setting a default 60-minute daily limit for users under age 18, with the option to opt out.

formalizes these insights.

Proposition 3 (Default time limit reduces self-control problem) *For users with self-control problems as defined in Proposition 1, introducing a default time limit of $\phi \in (0, T)$ minutes reduces video consumption to ϕ minutes for those with $\chi \leq \phi$, thereby increasing consumer surplus. The optimal value of ϕ balances a trade-off between broader policy coverage and the surplus gain per complier.*

3 Data and Institutional Background

I take the short-drama-series industry as the empirical context for studying the self-control problem and its interaction with content length. I introduce the short drama industry in section 3.1 and provide details about the platform of my study in section 3.2. The individual-level data for my analysis are described in section 3.3, which I use to establish empirical evidence on misprediction and self-control problems in section 3.4.

3.1 Short drama series

Short drama series are a mixture of drama series and short-form videos. They typically have scripts adapted from web novels and are optimized for vertical viewing on smart phones.[15] The most prominent feature of those mini-dramas is their *shortness*, with each episode of a 40- to 100-episode series lasting about one minute. With a genre similar to soap operas, these short dramas, described as "high on drama, low on glam, and full of plot twists," are designed to get viewers hooked fast, which makes this industry an ideal context for studying temptation and the resulting self-control problem.

Demand for short drama series is high and growing rapidly. In China—the origin of this industry in early 2022—the short drama market generated $7.0 billion in revenue in 2024, surpassing box office sales of $5.9 billion.[16] The trend is quickly expanding globally. For example, global short drama platforms such as DramaBox, ReelShort, and ShortMax respectively ranked 6th, 18th, and 33rd in the U.S. Apple App Store Entertainment Top Charts for free apps on April 26, 2024.[17] Users are also willing to pay for these dramas. For instance, the superstar drama series on the platform I study reached $3.5 million revenue in one month after its first release.

Yet, the production of these short dramas is generally low in quality. The scripts are usually bought and adapted from popular web novels, which typically emphasize immediate engagement over depth and complexity. The plots often rely on sensational or emotionally charged themes such as romance, conflict, or dramatic twists to maintain viewer interest. Character development

[15] Because of its vertical display mode, this type of series is sometimes also called "vertical drama."

[16] Source: *2024 China Short Drama Industry Research Report*, by Shenzhen Media Group, the Audiovisual Art Research Center of the Communication University of China, and the China Television Drama Production Industry Association.

[17] For comparison, the Apple App Store Entertainment Top Charts on the same day had other related apps such as TikTok (2nd), Netflix (5th), YouTube TV (16th), and AMC Theatres (46th).

and nuanced storytelling are minimal, because the focus is on creating visually appealing and attention-grabbing content that can be easily consumed on mobile devices. Short dramas feature relatively unknown actors, avoiding A-list or B-list celebrities, and are directed and edited by mid-level or inexperienced professionals. As a result, production costs are low, mostly below $80,000 per drama, with the entire production cycle completed in seven to 10 days.

The short-drama-series industry thus shares key characteristics with platforms such as TikTok and other short video formats, where low-cost content is designed to be rapidly consumed and highly engaging, often at the expense of quality or depth. This industry offers a unique opportunity to study consumer behavior in the context of short-form video consumption.

3.2 The platform

The data for my study are provided by one of the leading short drama platforms in the industry. I outline the basic structure of this platform in Figure 2, where activities occurring on the platform are highlighted in the shaded area. Dramas are the platform's primary assets. On the supply side, scripts are purchased and adapted from web novels and then produced into short dramas. The platform can acquire existing dramas from external producers through a negotiated two-part tariff or act as an integrated producer to create dramas itself.[18] To attract users, the platform constantly advertises its dramas on major social media platforms such as TikTok and Facebook. Users need to unlock each drama episode by purchasing tokens, subscribing, or watching external ads, which form the platform's three main revenue streams. Below, I provide more details on the components relevant to my study. A more comprehensive description of the platform is available in Appendix B.

User demographics. This platform operates in global markets. The daily active users (DAU) averaged 120,000 in 2024, with the US contributing 23% of users and 63% of revenue. I therefore focus on the US market. Additionally, 86% of the user base is female, and the majority are from younger generations: 36% of users are under 30, 63% are under 40, and 85% are under 50.[19] User overlap with Facebook and TikTok, the two primary social media platforms on which this platform advertises, is substantial, suggesting this sample is representative of a wider audience susceptible to social media addiction.

The (superstar) drama. By March 2024, the platform offered over 400 dramas in all major languages. Adapted from web novels, these short dramas feature dramatic genres, with the most popular being "love after marriage" and "toxic love."[20] On average, 1.7 million episodes are con-

[18]The two-part tariff includes an upfront payment and a revenue-sharing rule from the platform.

[19]The platform calculated the age for the period between March 1, 2024, and May 14, 2024. The app is labeled "18+," so only users aged over 18 are included. The age distribution may skew even younger because teenagers may access the platform using their parents' phones.

[20]The "love after marriage" genre explores romantic relationships that develop and evolve post-marriage. "Toxic love" refers to a dysfunctional relationship characterized by unhealthy behaviors, such as manipulation, control, emo-

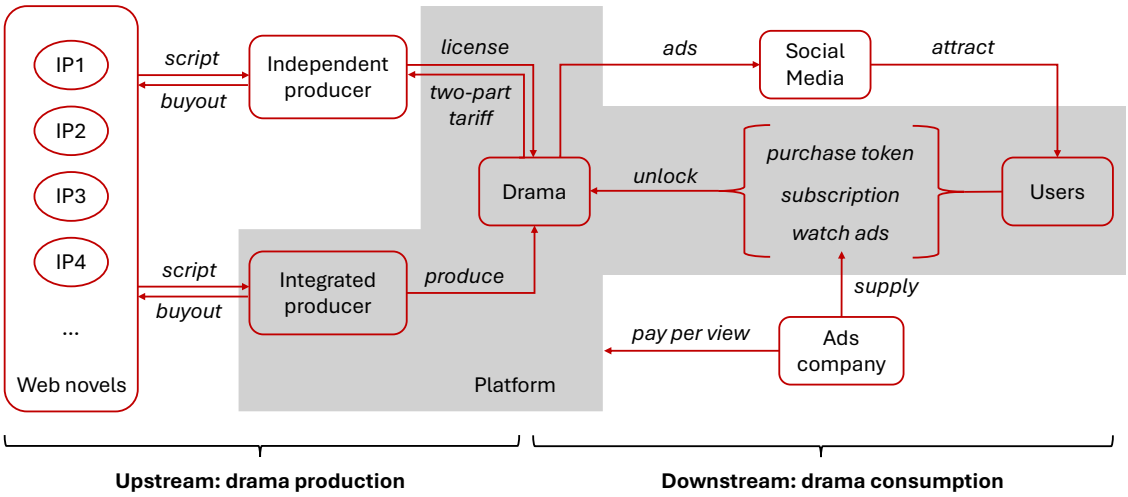


Figure 2: Institutional background of the platform

Notes: The figure summarizes the core business model of the platform. Activities occurring on the platform are highlighted in the shaded area, with short dramas serving as the platform's primary assets. On the supply side, scripts are acquired and adapted from web novels, then produced into short dramas—either by the platform itself or by independent producers. In the latter case, the platform secures licensing rights through an upfront payment and a revenue-sharing agreement. On the demand side, the platform actively promotes its dramas on major social media platforms such as TikTok and Facebook to acquire users. Users unlock each episode by purchasing tokens, subscribing, or watching external ads—these constitute the platform's three main revenue streams. Ads are provided by advertising firms, which pay the platform based on viewership.

sumed daily, with a high level of concentration where a single "superstar" drama accounts for 58.8% of viewership. Spillover across dramas is minimal, with only 3.6% (or 1.5%) of users who have unlocked the superstar drama also unlocking (or topping up for) another drama. Therefore, my main analysis focuses on user behavior related to this superstar drama, which was introduced in mid-January 2024. Consisting of 80 episodes, each lasting one minute, the first nine episodes are free to watch, whereas each subsequent episode costs 65 platform tokens. Produced in-house by the platform, this drama generated $3.5 million in revenue within a month, despite having a production cost of just $70,000. Its success can also be attributed to significant advertising efforts, with the platform investing over $600,000 in promotions within a month.

Drama demand. The payoff-relevant user activity for the platform is unlocking episodes, which can be done in three ways, as summarized in Figure 2: paying tokens, subscribing, or watching ads. My analysis focuses on tokens, the primary revenue source, generating $31,028 in daily income and accounting for 78.0% of total revenue between May 2023 and May 2024.[21] Users must unlock each episode sequentially by spending a specified amount of platform tokens, providing

tional abuse, and dependency.

[21]The second-largest revenue source is subscriptions, a model commonly used by drama platforms such as HBO. The platform offers three subscription options varying by duration: a week ($29.99), a month ($59.99), and a year ($199.99), contributing an average daily income of $6,631 and 16.7% of total revenue. The final revenue stream comes from ads, where users watch 30-second ads to unlock episodes, earning the platform $4,329 in daily income and 5.4% of total revenue. The time series of these income sources is reported in Appendix B.

ideal variation to identify willingness to pay on an episode-by-episode basis.

Additionally, users must purchase top-up packages with real money when they run out of tokens.[22] The platform offers four packages: $4.99 for tokens equivalent to 7.50 episodes, $9.99 for 16.62 episodes, $19.99 for 36.74 episodes, and $29.99 for 81.49 episodes. Random promotions introduce some variation in the number of tokens for each package across users and over time. As shown in Table 3, the platform employs a non-linear pricing structure, where the average price per episode decreases with larger package sizes. This structure creates a trade-off between lower costs and greater commitment, providing a source of identification for users' ex-ante preferences. The top-up page appears whenever users deplete tokens, and payment can be made seamlessly through the App Store or Google Play within seconds, making the process nearly frictionless.

User behavior. After unlocking an episode, users may rewatch it an unlimited number of times. However, I focus on the initial viewing of each episode and do not account for repeated viewing behavior. Users in the sample are required to unlock episodes sequentially, and they typically watch an episode in its entirety at least once before unlocking the next. This observation supports the assumption that users make viewership decisions at the video level.

3.3 Data and summary statistics

For a representative, randomly selected 10% sample of global users, I have access to individual-level log data that encompass users' full activities, including top-ups, unlocking, subscriptions, and viewership, from November 1, 2023, to March 31, 2024.[23] I focus on two key dimensions: top-up package purchase and unlocking. Each time users purchase platform tokens, I observe the transaction time, the price paid, and the number of tokens acquired. Additionally, I record the time, method, and number of tokens spent to unlock each drama episode.[24] This information fully characterizes users' drama-watching activities via platform tokens.

Data cleaning procedure. I outline the data cleaning procedure below and provide full details in Appendix C.1, including the rationale and resulting sample size for each step, and a discussion of potential selection introduced throughout the process. Based on institutional knowledge, I focus on top-up and unlocking activities related to the superstar drama by users located in the U.S. I exclude users who unlocked another drama before or during their engagement with the superstar drama, as well as those who accessed episodes via subscriptions or ad viewing. To focus on regular usage patterns, I drop users who did not watch the episodes in their natural sequential order. To ensure users had sufficient time to complete the drama, I retain only those who began

[22] As part of the platform's advertising strategy, users can also earn a small amount of gift tokens through daily log-ins, sharing content on social media, and following the platform's TikTok account, but these amounts are negligible compared with top-up packages.

[23] Each user is randomly assigned a unique user ID upon first opening the platform app. Users in my sample are those whose user ID ends with the number "8," effectively creating a 10% random sample of all users.

[24] For free episodes, I consider the first-time viewership as the equivalent unlocking activity.

Table 2: Summary statistics on the user level

Variable	mean	std	min	p25	p50	p75	p90	p95	p99	max	count
Episode	19.03	23.79	1	9	9	12	80	80	80	80	11,512
Top-up ($)	5.80	13.35	0	0	0	0	29.99	39.98	49.97	69.96	11,512
Round	1.37	0.73	1	1	1	2	2	3	4	5	11,512
Gift	0.29	0.96	0	0	0	0	1	2	5	12	11,512
Endowment	0.43	0.91	0	0	0	0	2	3	4	4	11,512

Table 3: Top-up packages: Tokens and market share

Price ($)	tokens	[p5,p95]	price/episode	market share
4.99	7.50	[7,9]	0.67	0.34
9.99	16.62	[16,20]	0.60	0.36
19.99	36.74	[36,46]	0.54	0.20
29.99	81.49	[62,92]	0.37	0.11

watching before January 31, 2024, and exclude users who had not finished the series but remained active during the final week of the sample period (March 24–31, 2024). To ensure that tokens were primarily acquired through top-up purchases, I exclude users whose initial or gift token endowment exceeded the equivalent of five episodes.[25]

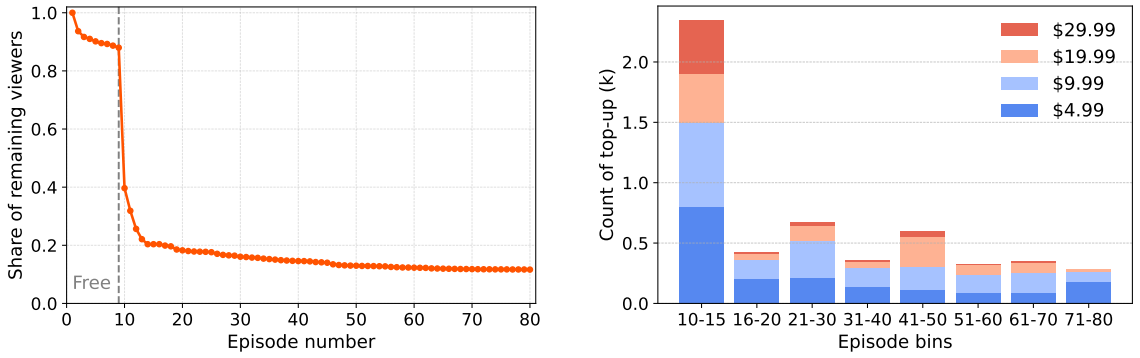
For this sample, I define a "round" as a sequence in which any two consecutive episodes are unlocked within two hours, capturing the behavior of users who pause and later resume viewing.[26] I further refine the sample by excluding outliers in the top percentile of the round distribution, users who purchased top-up packages not listed in Table 3, and those whose packages contained an unusually rare token quantity (frequency < 5%, conditional on prices).

This data-cleaning process results in a sample of 219,064 unlocking events and 5,360 top-up activities by 11,512 users. For ease of interpretation, I normalize all token quantities by the per-episode cost in the analysis that follows.

Summary statistics. Summary statistics for these users are presented in Table 2. On average, a user unlocks 19.03 episodes and spends $5.80 on this drama. Considerable heterogeneity exists, with over three-quarters of users not spending anything and only unlocking free episodes. More than half users consume the videos in a single round. The average value of gift tokens and initial

[25]Users could obtain initial or gift tokens through platform promotions such as daily log-ins or sharing the drama on social media. I exclude users whose non–top-up token value exceeds five episodes, which accounts for 8.7% of the sample. This small share suggests that most users did not rely heavily on such tokens, and that the exclusion assumption is unlikely to affect the quantitative results. A full discussion, including the distribution of initial and gift tokens, is provided in Appendix C.1.

[26]I model round formation as a stochastic and exogenous process, introduced in the structural model. This simplifying assumption abstracts from pausing behavior, which is not central to the analysis. Section 4.4 discusses round dynamics in greater detail, including the underlying rationale, robustness checks, and potential biases. All results remain robust to defining rounds using different (e.g., 12-hour) thresholds.



(a) Share of remaining viewers by episode

(b) Package purchases by episode bins

Figure 3: The dynamic patterns in unlocking and top-up package purchasing decisions

Notes: Panel (a) reports the share of users who remain viewing each episode, which is monotonically decreasing because users have to unlock sequentially. The total number of users is 11,512. The first nine episodes are free to watch, accounting for the sharp drop at the 10th episode. Panel (b) summarizes the dynamic pattern in top-up package sales by bins of five or 10 episodes.

endowments per user corresponds to 0.29 and 0.43 episodes, respectively—both small relative to top-up purchases.[27] For simplicity, I model these gifts as extra tokens included with all the top-up packages in my quantitative analysis.

Table 3 presents summary statistics related to top-up choices. The platform employs non-linear pricing, where the average price per episode decreases with higher-price packages. This structure highlights the trade-off between a lower average cost and a larger ex-ante commitment in the top-up decision. All four packages have a positive market share, allowing for the identification of users' ex-ante preferences and temptation durations based on observed top-up choices through a revealed-preference argument. Additionally, I observe variation in the number of tokens across top-up packages due to promotional offers randomly presented to users upon accessing the top-up page, which facilitates the identification of price sensitivities.

Figure 3a presents the share of unlocking activity across episodes. By construction, all 11,512 users in my sample watch the first episode. After each episode, some users discontinue watching, and the share of unlocks per episode stabilizes relatively quickly. The most significant drop in viewership occurs at the 10th episode, coinciding with the end of the nine-episode free trial, when users must purchase tokens to continue.[28] Three-quarters of users exit the series without purchasing any tokens. Ultimately, 1,340 users (11.64%) complete the entire drama. Figure 3b illustrates the dynamic pattern of top-up package choices, grouped in bins of five or 10 episodes. The largest package, priced at $29.99, is predominantly purchased between the 10th and 15th episodes by highly engaged users who are willing to commit to watching many episodes. The

[27] Users can receive free tokens as rewards for engaging in platform activities such as daily logins or sharing content on social media. The quantity of these gift tokens is small, rendering such promotions negligible for my analysis.

[28] The probability of discontinuing peaks at the 10th episode but remains elevated in subsequent episodes, as some users possess a small token endowment (sufficient for one to four episodes) and delay their top-up decision until exhausting their token stock.

other three packages exhibit positive sales across all episode bins. The persistent market share of lower-priced packages across episodes suggests a self-control problem: users initially plan to watch only a few episodes ex ante but fail to stop ex post. I further investigate this pattern in section 3.4.

3.4 Evidence of self-control problems

In this section, I examine patterns in user behavior related to self-control problems, which provide essential variation for identifying temptation in section 4. The identification arises from the non-linear pricing scheme in top-up package purchases. Figure 4a illustrates this idea with an actual user who repeatedly purchases the smallest top-up package and completes the entire drama in a single night.[29] The choice of the $4.99 package suggests the user initially intended to watch only the next seven episodes at each purchase; otherwise, opting for a larger package with a lower per-episode cost would have been more rational. However, their ex-post behavior—continuing to the end—reveals an inability to commit to stopping, indicating a severe self-control problem.[30] This pattern is widespread in my sample, as reflected in the persistently high market share of small top-up packages throughout the drama in Figure 3b.

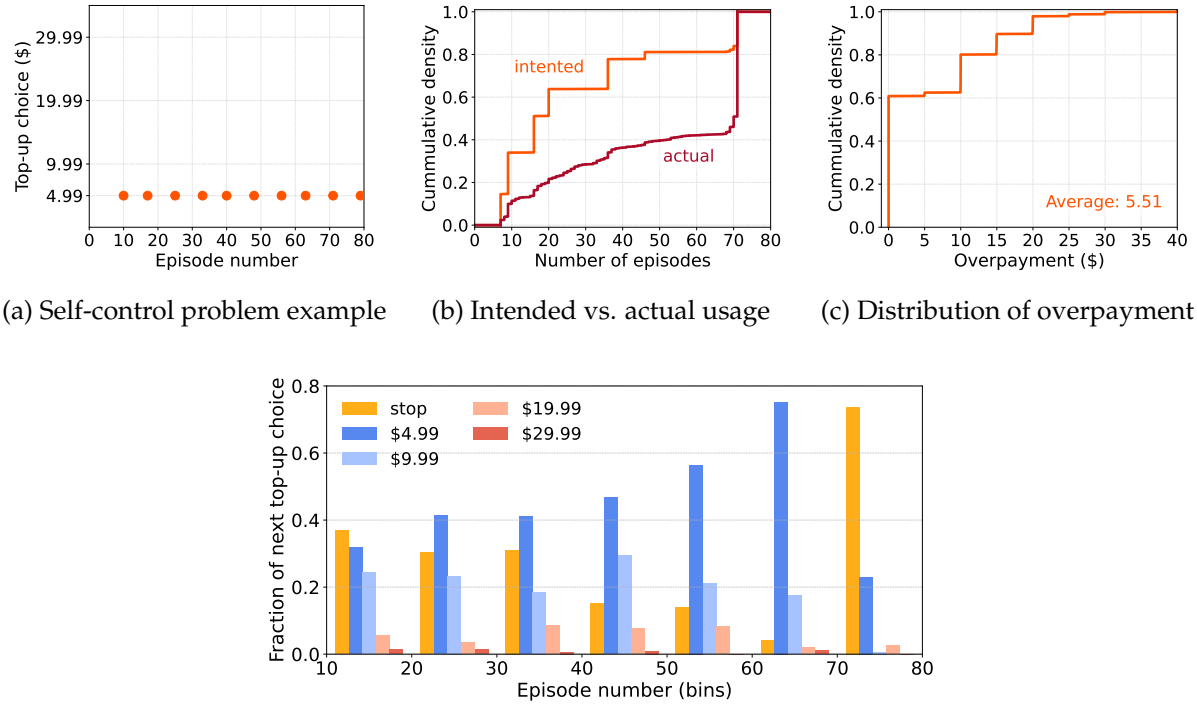
To further illustrate time inconsistency, in Figure 4b I compare the distribution of intended versus actual number of episodes watched for users who have purchased tokens. Intended consumption, represented by the orange cumulative density curve, is inferred from users' first token package choice. On average, users initially intend to watch 28.0 episodes when making their first top-up decision. By contrast, the actual number of episodes watched, shown by the red curve, exhibits first-order stochastic dominance over the intended distribution. The average number of episodes actually viewed is 51.0, which is 82.1% higher than users' initial intention. This significant discrepancy—where actual consumption far exceeds intended consumption—provides strong evidence of self-control problems.

Moreover, due to the platform's non-linear pricing scheme, the fact that the example user in Figure 4a could have spent less motivates the use of *overpayment* as an indicator of self-control problems. Figure 4c presents the distribution of overpayment, measured relative to optimal spending, among users who made at least one package purchase.[31] The results indicate 39.1% of these users spend more than the rational benchmark to unlock episodes, suggesting self-control problems are prevalent in my sample. Specifically, 17.6% overpay by approximately $10, 9.5% by $15, and 8.2% by $20. On average, these users overpay by $5.51, which is 22.7% more than the optimal expenditure, underscoring the substantial impact of self-control problems in this setting.

[29] The fact that this user finishes the series in one night rules out alternative explanations, such as "budgeting," which might apply if they had spaced out their viewing over several weeks.

[30] An alternative explanation is that users anticipate a higher likelihood of stopping as they progress further into the series, making them more cautious in token purchases. For example, if a user is waiting for a call, the longer they wait, the more likely they are to receive it in the next few minutes. However, the aggregate viewing pattern does not support this explanation. As shown in Figure 3a and Figure C.5 in Appendix C.4, the probability of continuing to the next episode increases quickly and then remains consistently high throughout the rest of the series.

[31] Optimal spending is imputed using the average token quantity in each package from Table 3.



(a) Self-control problem example (b) Intended vs. actual usage (c) Distribution of overpayment

(d) Next top-up package choice following $4.99 purchase

Figure 4: Evidence of self-control problems

Notes: Panel (a) illustrates an example of self-control problems, where the user repeatedly purchases the smallest top-up package. The y-axis lists the available top-up options, and the x-axis indicates the episode before which each top-up occurs. Panel (b) compares the cumulative distribution of intended vs. actual number of episodes watched among users who have purchased token packages. Intended consumption is inferred from the first package they purchase. Panel (c) presents the cumulative distribution of overpayment among users who made at least one package purchase. Overpayment is measured relative to the optimal spending, which is imputed using the average token quantity in each package from Table 3. The natural lower bound for overpayment is $0. Panel (d) reports the distribution of users' next top-up decisions, conditional on having previously purchased the $4.99 package, across episode groups. A time-consistent user would choose to stop. The sharp increase in stopping behavior between episodes 70 and 80 is mechanically driven, because users must stop after completing the series.

Figure 4d provides additional evidence of self-control problems through conditional choice probabilities. After purchasing the $4.99 package, a time-consistent user would stop watching when faced with the next top-up decision. However, as shown in Figure 4d, most users deviate from their ex-ante preferences and continue topping up. Among those who purchased the $4.99 package between episodes 10 and 70, only 30.7% manage to stop thereafter.[32] The majority instead purchase another $4.99 or $9.99 package, aligning with my stylized model's prediction that self-control problems arise from short-term temptation. Moreover, time inconsistency becomes more pronounced as the series progresses, with the conditional probability of stopping (the yellow bar) declining, suggesting users do not become more sophisticated in managing their self-control within this context.[33]

[32]The spike in stopping behavior between episodes 70 and 80 is mechanically high, because users must stop after completing the series.

[33]The drama lasts only 80 minutes in total, limiting users' ability to recognize and adjust for their self-control prob-

Table 4: Lower bound for overpayment: Savings per user by swapping $4.99 to lager packages

Episode (bin)		10-20	20-30	30-40	40-50	50-60	60-70	70-80
Switch $4.99 to:	$9.99	–1.72	–0.72	–0.84	0.20	0.30	2.62	–3.28
	$19.99	–6.71	–4.75	–4.04	1.75	–2.05	–6.16	–13.28
	$29.99	–10.05	–8.52	–8.70	–7.65	–12.05	–16.16	–23.28
Number of observations		797	203	140	97	74	83	154

Notes: In this exercise, I hold fixed each user's total number of episodes and replace their first $4.99 purchase within a given episode bin with a larger token package ($9.99, $19.99, or $29.99). All other top-ups are kept at their original choices (before the switch) or approximated by the closest original decision (after the switch). To account for the discreteness of token purchases, I adjust unused tokens from additional packages by imputing their value at the user's average token price. The table reports average net savings per user, where negative values indicate higher costs. These estimates provide a lower bound on overpayment from self-control problems, as they abstract from uncertainty, learning, and other rational motives for choosing the $4.99 package. I highlight positive values in orange, indicating evidence of self-control problems even when alternative explanations may exist.

Alternative explanation: uncertainty and learning. Learning offers a competing explanation for the patterns in Figure 4, where Bayesian users gradually update their beliefs about drama quality. Under this view, small top-up purchases reflect information acquisition under uncertainty. To consider these alternatives, I follow DellaVigna and Malmendier (2006) and estimate a lower bound on overpayment from self-control problems by replacing the first $4.99 package in a given episode bin with a larger package for all users. This exercise yields a lower bound because the estimates abstract from uncertainty, learning, and other rational motives for choosing the $4.99 package, while also preventing agents from reoptimizing.[34] Table 4 reports the average net savings. Despite this underestimation, the results indicate that users purchasing the $4.99 package between episodes 40 and 70 systematically overpay, suggesting that self-control problems persist even when alternative explanations are possible.[35]

In Appendix C.4, I provide extra evidence showing that learning is unlikely to be the main driver. Figure C.5 indicates that continuation probabilities stabilize after just two episodes, suggesting rapid learning on drama quality. Figure C.6a shows that even after watching 40 episodes, 33.6% of users continue to overpay, with an average overpayment of $2.79 for the remaining 40 episodes—comparable to the unconditional average of $5.51 across all 80 episodes reported in Figure 4c. Moreover, Figure C.6b suggests average overpayment remains seven to ten cents per episode for users who have already watched 10–40 episodes. That users persist in such behavior despite having acquired ample information suggests that self-control problems, not learning, are the primary force at play.

lems over time.

[34]Some users may purchase a $4.99 package to acquire information or because of uncertainty about outside options, and then stop watching. For these users, the exercise produces negative savings. The average net savings across all users therefore provide only a lower bound on overspending due to self-control problems.

[35]Negative values in Table 4 should not be interpreted as evidence against self-control problems, since the exercise only provides lower bounds. For example, high uncertainty about outside options early in viewership may amplify the underestimation. In the structural model, this mechanism is explicitly captured as an exogenous probability of stopping, which disrupts habit formation.

Self-control problems: potential mechanisms. In this section, I present evidence of self-control problems while remaining agnostic about the precise mechanisms underlying the observed time inconsistency. Within the stylized model, short-term behavioral temptation may be interpreted as present bias that overweights immediate entertainment, unperceived habituation that hinders stopping, or simple inertia. Although the data lack the variation needed to distinguish among these mechanisms, they are isomorphic in terms of behavioral predictions and welfare implications. I therefore adopt the general framework of temptation to analyze the problem.

4 Structural model

I adopt a structural approach to systematically examine self-control problems in the short-drama-series industry. Building on the stylized model from section 2, I develop a single-agent dynamic discrete-choice model for this setting in section 4.1. The model solution is characterized in section 4.2, followed by a discussion of modeling assumptions in section 4.4.

4.1 Setup

The superstar drama consists of T episodes indexed by t, with each one-minute episode naturally defining a time period in which the user must decide whether to purchase a top-up package and unlock the next episode. The drama-watching activity can thus be framed as a finite-horizon dynamic discrete-choice problem. The platform sets token price as $c_t = \mathbb{1}\{t > 9\}$, meaning users begin paying a normalized amount of tokens after the ninth episode.

Perceived flow utility. Following the stylized model in Section 2, I capture self-control problems as short-term temptation with naïve perception, which can be microfounded by behavioral mechanisms such as present-biased preferences or inertia.[36] The perceived flow utility from episode t is defined as:

$$u_t(h_t, \delta, \psi_t, \chi, \epsilon_t) := \underbrace{\overbrace{\delta}^{\text{drama}} - \overbrace{\psi_t}^{\text{outside}}}_{\text{intrinsic utility}} + \underbrace{\mathbb{1}\{\chi \geq 1\}\kappa}_{\text{temptation}} + \underbrace{\alpha(h_t)}_{\text{habit formation}} + \underbrace{\epsilon_t}_{\text{taste shifter}}, \tag{7}$$

which comprises intrinsic utility, temptation, habit formation, and a random taste shifter. The intrinsic utility reflects the drama's value relative to the user's outside option. The drama value δ is constant over episodes but varies across users.[37] Given the rapid convergence in viewership (Figure 3a) and consistent top-up behaviors across episodes (Figure 4d), I abstract from learning by assuming users have perfect information on δ.[38] The outside value ψ_t is time-varying to capture

[36] I remain agnostic about the precise mechanism. These mechanisms are modeling-wise isomorphic, because different interpretations of the temptation term yield the same welfare implications.

[37] For simplicity, I omit the user subscript i throughout this section because this is a single-agent problem.

[38] I could incorporate learning into the framework by modeling Bayesian users who gradually infer their value of δ from episodes they watch. Estimating this extended model reveals that the learning effect is minimal, because users

potential auto-correlation in the user's decision-making process from external factors. As defined in the stylized model, the per-minute temptation κ exists when its duration χ exceeds one-minute episode length.[39] The taste shifters ϵ_t captures random shocks on users' perceived utility from each episode and is unknown to them until the episode is watched.

I allow for rational addiction modeled as habit formation (Becker and Murphy, 1988).[40] Encountering each episode t, the user has habit stock h_t, defined as the number of previous episodes watched within the current round r_t, which depreciates across rounds. This habit stock contributes to utility through a non-linear functional form $\alpha(h_t) = \alpha_1 h_t(h_t - \alpha_2)$.[41] Each round evolves exogenously, with a probability ρ that, after watching an episode, the user must start a new round ($r_{t+1} = r_t + 1$).[42] By design, habit formation is temporary, building up within each round and capturing the transient nature of digital addiction. I assume users have rational expectations regarding habit formation and round dynamics, fully internalizing these dynamics when making their drama-watching decisions.

In practice, I assume the outside value ψ_t is round-specific, depending on the time of day when the user begins watching. I normalize $\psi_t = 0$ for rounds that start during the day (11am-12am) and parametrize the outside value for nighttime (1am-10am) as $\overline{\psi}$.[43] The probability that a new round begins at night ($\varphi = 0.27$) is directly calibrated from the data. I also assume the random taste shifter ϵ_t follows a Gumbel distribution, $Gumbel(-\beta_\epsilon\gamma, \beta_\epsilon)$.[44]

Dynamics. I formalize the dynamic components of my model. To align with the data, I assume all users choose to watch the first episode. Additionally, I assume each user experiences a temptation period lasting $\chi \in \{0, 1, ..., 60\}$ minutes and holds a naïve perception of temptation beyond this period.[45] The 60-minute limit reflects the short-term nature of temptation, is not quantitatively

can reasonably infer δ after one or two episodes. For simplicity, I exclude this mechanism in my baseline analysis. See section 4.4 and Appendix C.4 for more detailed discussions for this learning mechanism.

[39] For simplicity, I assume temptation remains constant every minute within its duration, though a more flexible model could allow it to decay over time. Although I cannot estimate the rate at which temptation decays, incorporating this feature would only strengthen the present bias, further amplifying the self-control problem and making the short video length an even greater factor in driving addiction.

[40] In drama series, episodes are linked by a continuous storyline, which helps keep users engaged and make habit formation relevant. Appendix C.2 provides evidence on habit formation by examining the relationship among habit stock, the probability of continuing watching, and the probability of purchasing higher-priced packages.

[41] I adopt this quadratic functional form for the following reasons. First, it satisfies $\alpha(0) = 0$, implying no additional utility in the absence of habit stock. Second, when $\alpha_1 < 0$, it generates an inverted-U shape, capturing diminishing returns due to fatigue. Third, the parameters are straightforward to interpret. This functional form is not critical to my quantitative results.

[42] The assumption of exogenous round dynamics is crucial for identifying habit formation. It is reasonable because if users stop watching because they dislike the series, they should not resume in the next round after losing all accumulated habit stock. Thus, any observed temporary disruptions must stem from exogenous factors—for example, a user might receive an unexpected email about a paper decision, interrupting their drama-watching session. See section 4.4 for a complete discussion on this round dynamic.

[43] This choice of time is motivated by viewing patterns in the data, which show below-average viewership between 1am and 10am. A detailed description can be found in Appendix C.3.

[44] The parameter β_ϵ is the scale parameter that controls the dispersion of a Gumbel distribution. The location parameter is set to $-\beta_\epsilon\gamma$, where γ is the Euler constant, such that the expectation of ϵ_t is 0 (e.g., Rust, 1987).

[45] As shown in Appendix A.2, sophisticated users who dislike the videos would never engage with the platform. Therefore, I assume all users in my sample are naïve. In Section 4.4, I discuss potential bias from this assumption and

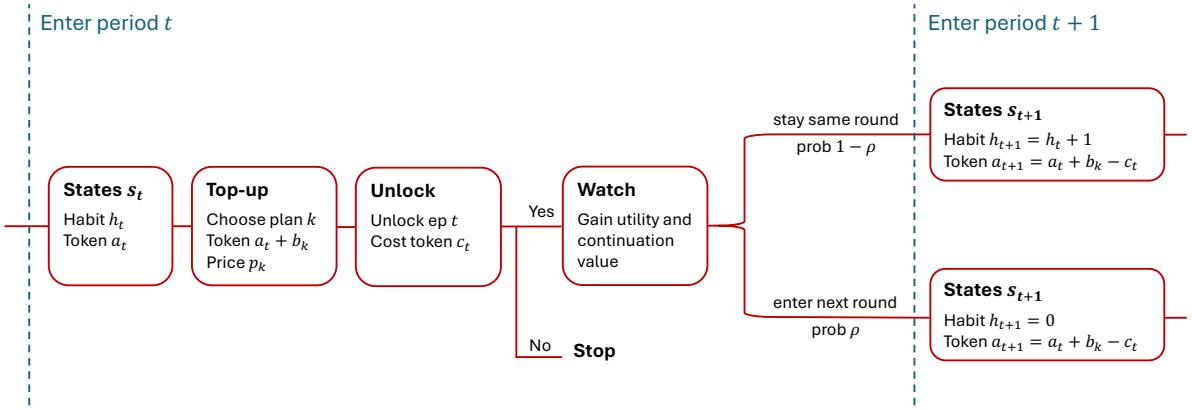


Figure 5: Timing: Decisions on top-up and unlocking within a time period

critical, and improves the tractability of the solution algorithm.[46] At the start of each episode $t \geq 2$, the user has unlocked and watched the previous $t-1$ episodes, with a state $\boldsymbol{s}_t$ that includes their habit stock h_t and token stock a_t. The decision-making process is summarized in Figure 5, where two key decisions are made within each period: top-up package purchasing and unlocking.

First, users decide whether to top up when their current token stock a_t is lower than the cost of unlocking the next episode, c_t.[47] I assume this top-up process is frictionless, because users can complete purchases within seconds.[48] They choose from a menu of $K = 4$ packages, indexed by $k \in \{1, ..., K\}$, with price p_k and token quantity b_{kt}, along with an outside option denoted as $p_0 = b_0 = 0$. The packages are ordered such that price p_k increases with k. The token amounts are randomly drawn for each user-period pair from a distribution $F_{\boldsymbol{b}}(\cdot)$ set by the platform, reflecting random promotions. In practice, the platform employs a non-linear pricing strategy, where the average price per token, $\mathbb{E}[p_k/b_{kt}]$, decreases as the package price increases, highlighting a trade-off between lower token prices and greater ex-ante commitment to watching more episodes. Thus, top-up choices reveal users' ex-ante preferences for how many episodes they intend to watch. Additionally, I assume an idiosyncratic taste shifter ν_{kt} follows a standard Gumbel distribution for each top-up package. I denote the value and expected value functions associated with the top-up decision by $V_t(h, a; \delta, \psi, \chi, \boldsymbol{b}, \boldsymbol{\nu})$ and $EV_t(h, a; \delta, \psi, \chi)$, respectively, with state variables listed before the semicolon. The full solution is detailed in Section 4.2.

Second, users decide whether to unlock episode t, given their habit stock h_t and token stock $a_t + b_{k^*t}$. By unlocking, users expect to receive their perceived flow utility u_t as defined in equation

conclude that including (partially) sophisticated users would further strengthen the case for self-control problems.

[46] This 60-minute upper bound on temptation duration is not restrictive: under the estimated distribution of χ, 98.8% of users have $\chi < 60$. See Section 5.2 for details.

[47] In practice, the top-up page only appears when users have insufficient token balance. Users could, in principle, stop watching, return to the main page, and manually access the top-up page at any time, but fewer than 1% of users choose to do so when they have sufficient tokens.

[48] This assumption is quantitatively conservative because introducing friction would reduce users' willingness to top up, implying that the structural model would estimate an even stronger present bias to account for the time-inconsistent behaviors observed in the data.

(7), along with a perceived continuation value, denoted $\mathcal{C}_t(h, a; \delta, \psi, \chi)$ for notational simplicity. This continuation value stems from the ability to unlock the next episode, $t+1$, where users perceive themselves to be one minute less tempted under the naïveté assumption. Notation-wise, I use $W_t(h, a; \delta, \psi, \chi, \boldsymbol{\epsilon})$ and $EW_t(h, a; \delta, \psi, \chi)$ to represent the perceived value and expected value functions associated with the unlocking decision.

Finally, I assume no continuation value exists for habits and tokens outside of this superstar drama, which yields the terminal condition:

$$EV_{T+1}(\cdot) \equiv EW_{T+1}(\cdot) \equiv 0. \tag{8}$$

This assumption is based on the empirical observation that spillover across dramas is minimal, and most users do not watch anything else after completing this one.[49] Thus, condition (8) converts the user's dynamic problem into a finite-horizon framework.

4.2 Solution: Recursive definition of value functions

Following O'Donoghue and Rabin (1999), I use the concept of *perception-perfect equilibrium strategies* to account for the behavioral forces in my model. Under this framework, users maximize their *perceived* utility, which can lead to irrational self-control problems. I define the value functions iteratively by using backward induction, which, combined with the terminal condition (8), fully characterizes the solution.

For any episode $t \leq T$, the perceived value function associated with the top-up decision for a user with habit stock h, token stock a, drama value δ, outside value ψ, token amounts b_k, perceived temptation duration χ, and taste shifter ν_k is given by:

$$\begin{aligned} V_t(h, a; \delta, \psi, \chi, \boldsymbol{b}, \boldsymbol{\nu}) = \quad & \mathbb{1}\{a < c_t\} \max_{k \in \{0,1,\ldots,K\}} \{EW_t(h, a + b_k; \delta, \psi, \chi) - \omega p_k + \nu_k\} \\ & + \mathbb{1}\{a \geq c_t\} EW_t(h, a; \delta, \psi, \chi), \end{aligned} \tag{9}$$

where $EW_t(\cdot)$ is the perceived value from unlocking and ω represents price sensitivity. If the token balance is insufficient ($a < c_t$), the user must choose one of the top-up options k; otherwise, they skip topping up and proceed to the unlocking stage. Because the taste shifter ν_{kt} follows a standard Gumbell distribution, the expected value can be written as:

$$\begin{aligned} EV_t(h, a; \delta, \psi, \chi) = \quad & \mathbb{1}\{a < c_t\} \mathbb{E}_{\boldsymbol{b}} \left[\log \left(\sum_{k=0}^{K} \exp \left[EW_t(h, a + b_k; \delta, \psi, \chi) - \omega p_k \right] \right) \right] \\ & + \mathbb{1}\{a \geq c_t\} EW_t(h, a; \delta, \psi, \chi), \end{aligned} \tag{10}$$

where the expectation $\mathbb{E}_{\boldsymbol{b}}(\cdot)$ is taken over the token amount $\boldsymbol{b}$.

In the unlocking stage, the user decides whether to unlock episode t. By unlocking, they derive

[49] See more details in Appendix C.5.

the flow utility $\widetilde{u}_t(h, \delta, \psi, \chi, \epsilon)$ and gain the perceived continuation value defined as:

$$\begin{aligned} \mathcal{C}_t\left(h, a; \delta, \psi, \chi\right) = \quad & (1-\rho)\, EV_{t+1}\left(h+1, a-c_t; \delta, \psi, \max\left\{\chi-1, 0\right\}\right) \\ & + \rho \mathbb{E}_{\psi'}\left[EV_{t+1}\left(0, a-c_t; \delta, \psi', \max\left\{\chi-1, 0\right\}\right)\right], \end{aligned} \tag{11}$$

where the user forms rational expectation over the future habit stock and outside value according to the round dynamics. Because users are naïve about temptation, they perceive themselves as one minute less tempted when approaching the next episode $t+1$, resulting in a perceived next-period temptation duration of $\max\{\chi-1, 0\}$ in the continuation value (11). The value function associated with this unlocking stage is therefore:

$$W_t\left(h, a; \delta, \psi, \chi, \boldsymbol{\epsilon}\right) = \mathbb{1}\{a \geq c_t\} \max\left\{ \underbrace{u_t(h, \delta, \psi, \chi, \epsilon)}_{\text{flow}} + \underbrace{\mathcal{C}_t(h, a; \delta, \psi, \chi)}_{\text{continuation}}, \epsilon_0 \right\} + \mathbb{1}\{a < c_t\}\epsilon_0, \tag{12}$$

which, with $\epsilon \sim Gumbel(-\beta_\epsilon \gamma, \beta_\epsilon)$, yields the expected value function:

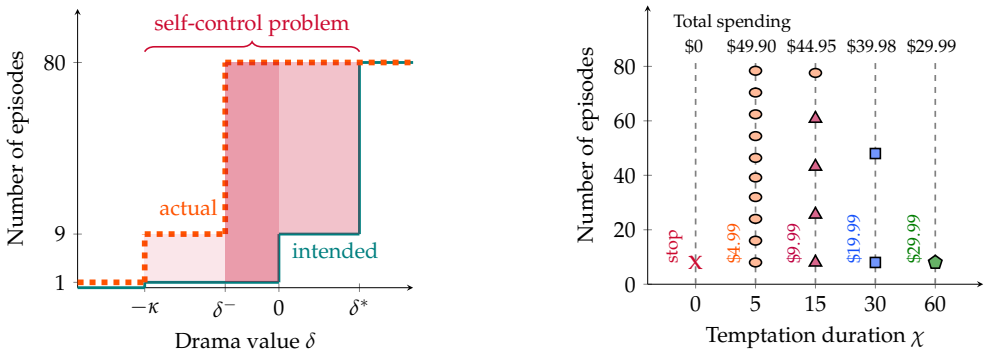
$$EW_t\left(h, a; \delta, \psi, \chi\right) = \mathbb{1}\left\{a \geq c_t\right\} \beta_\epsilon \log\left(1 + \exp\left(\frac{u_t(h, \delta, \psi, \chi, \epsilon) + \mathcal{C}_t\left(h, a; \delta, \psi, \chi\right)}{\beta_\epsilon}\right)\right). \tag{13}$$

The solution of this single-agent dynamic model is fully characterized by the recursive system of equations (9)–(13) and the terminal condition (8). Numerically, one can solve these value functions and the corresponding policy rules by iterating over perceived temptation duration χ and episode t. I further characterize model solutions in Appendix D.1. There, I analyze users' top-up and unlocking choices, establish the cutoff values of intrinsic quality that govern package selection, and show how variation in temptation duration shapes demand for different package sizes. I also demonstrate that the unlocking decision follows a threshold rule in intrinsic value, habit stock, and temptation.

4.3 Predictions of self-control problems

In this section, I examine patterns of self-control problems as functions of drama value δ and temptation duration χ, which manifest as overconsumption and variation in top-up package choices, respectively. This analysis provides the foundation for model identification. To isolate these effects, I assume no taste shocks (ϵ and ν), no initial endowment, and that each top-up package contains the average number of tokens reported in Table 3.

Figure 6a shows an inverted-U relationship between the magnitude of self-control problems, measured by the shaded gap between intended and actual episodes watched, and drama value δ for users with temptation. Users with low valuation ($\delta < -\kappa$) do not continue watching regardless of temptation, resulting in no self-control problem. Those with high valuation ($\delta > \delta^*$) rationally choose to watch all episodes, meaning temptation does not distort their behavior. Self-control problems arise primarily among users with intermediate valuations ($-\kappa < \delta < \delta^*$), with darker shading in Figure 6a indicating a larger discrepancy between ex ante preferences and ex post



(a) Medium-valued users overconsume most

(b) Longer-tempted users overpay less

Figure 6: Illustrative examples: Self-control problems as functions of δ and χ

Notes: This figure illustrates examples of self-control problems under varying drama value δ and temptation duration χ. To isolate the effects, all taste shocks (ϵ and ν) are set to zero, and no initial endowment is assumed. Panel (a) shows the intended (teal, solid line) and actual (orange, dashed line) number of episodes watched by a tempted user with one-minute temptation ($\chi = 1$) as a function of δ. The gap between the two lines, shaded in red, represents the magnitude of self-control problems—the darker the shading, the more severe the problem. Panel (b) illustrates the sequence of top-up choices for a user who, absent temptation, would watch only the free episodes but instead completes the entire series when subject to temptation of varying duration χ. The figure also reports the corresponding total spending, showing that longer-lasting temptation is associated with lower overpayment.

behavior. In the example with temptation duration $\chi = 1$, users with $-\kappa < \delta < 0$ intend to watch only the first episode, while those with $0 < \delta < \delta^*$ would plan to watch the free episodes (teal line). However, temptation causes these users to lose self-control and overconsume, as shown by the upward shift in the orange line. Users with $-\kappa < \delta < \delta^-$ watch 9 episodes (compared to an intended 1), while those with $\delta^- < \delta < \delta^*$ watch all episodes—despite intending to watch only 1 if $\delta^- < \delta < 0$, or 9 if $0 < \delta < \delta^*$. This pattern has important welfare implications: users with mid-range valuations suffer the most from temptation-driven self-control problems—a result confirmed by the structural estimates in Section 6.1.

I examine the effect of temptation duration χ on top-up choices in Figure 6b, highlighting how package selection and overpayment vary. Without temptation ($\chi = 0$), the user would only watch the free episodes. With temptation ($\chi > 0$), however, the user invariably pays to finish the entire series, falling into the self-control trap. The anticipated number of episodes at the time of purchase depends on χ, so longer-tempted users are more likely to buy larger token packages. Consequently, although all tempted users overconsume by completing the series, the magnitude of overpayment declines with χ, since higher-χ types choose cheaper per-episode contracts. This variation, generated by the rich menu of nonlinear contracts, identifies the distribution of temptation duration, providing novel contribution to the literature (e.g., DellaVigna and Malmendier, 2006).

This discussion illustrates how variation in drama value δ and temptation duration χ shapes self-control problems, as reflected in user behavior. It also provides an identification argument within the model: unlocking and top-up decisions inform the distribution of δ and χ. I build my identification strategy around this insight, as detailed in Section 5.1.

4.4 Discussion

I discuss my model assumptions in this section, with the goals of being transparent in my modeling choices, reasonings, and potential biases.

Sophistication. I assume users are naïve about their temptation beyond its duration, due to insufficient variation in the data to identify the degree of sophistication. The observed overpayment behavior suggests that users are not fully sophisticated. The assumption of naïveté aligns with most findings in the behavioral literature (DellaVigna, 2018; Augenblick and Rabin, 2019) and is commonly imposed in structural analysis (DellaVigna, Lindner, Reizer, and Schmieder, 2017; Laibson, Chanwook Lee, Maxted, Repetto, and Tobacman, 2024). In Appendix D.5, I characterize the solutions for sophisticated users and estimate the structural model assuming that 20% of users are sophisticated. The results indicate that including sophisticated users implies stronger (11.4% higher) and longer-lasting (54.5% higher) temptation, because naïve users must exhibit greater temptation to generate the same level of self-control problems observed in the data. Incorporating sophistication therefore only reinforces my quantitative findings.

Naïveté about habit formation. I model irrationality in my framework through temptation and the associated naïveté. As summarized in Allcott, Gentzkow, and Song (2022), irrational digital addiction can also stem from naïveté about habit formation. Although this type of irrationality could also lead to repeated top-up purchases, we would expect users to gradually shift to higher-priced top-up plans as their habit stock grows over time. However, this pattern is inconsistent with what is shown in Figure 4d, where users frequently purchase another $4.99 package, regardless of how many episodes they have watched. Thus, I conclude naïve temptation is the more relevant story in the context of this short drama.

Learning. For simplicity and tractability, I abstract from the model of learning, where Bayesian users infer the unknown drama quality δ based on the utility from previously watched episodes, and instead assume that users learn δ immediately after watching the first episode. Estimating an earlier version of this extended model reveals the learning effect is minimal, because users can reasonably infer δ after one or two episodes. Indeed, empirical evidence in section 3.4 and Appendix C.4 suggests learning cannot account for the observed irrational spending behavior. Figures 3a and Figure C.5 show that viewership and continuation probabilities stabilize after the first two episodes, indicating that users assess drama quality quickly and that learning effects are minimal relative to price effects. Moreover, the time-inconsistent pattern of package purchases persists even after many episodes, and the distribution of overpayment beyond the 40th episode mirrors the overall distribution.

Round dynamics. I model round dynamics as exogenous, capturing users who pause and later resume viewing—potentially due to a temporary increase in the outside option. This assumption

provides a tractable way to account for pausing behavior, which is not the focus of the analysis. It is not critical for the quantitative results, because all findings remain robust when rounds are defined using alternative thresholds, such as 12 hours. Alternatively, one could model pausing endogenously by assuming decreasing returns to viewing due to fatigue. My non-linear habit formation specification $\alpha(h)$ allows for this mechanism, whereas the estimation results in Section 5.2 reject decreasing returns.

Constant episode value. Instead of allowing the drama value to vary by episode, I assume users have a constant value δ to all episodes. This simplification is supported by the empirical observation that viewership evolves smoothly across episodes, as shown in Figure 3a. If episode values differed systematically, we would expect greater volatility in viewership patterns. Additionally, the random taste shifter ϵ_t captures idiosyncratic shocks for each episode, and its estimated magnitude (in section 5.2) is quite small, suggesting episode-specific value is not a key factor in explaining user behavior.

Network externality. Social media platforms often benefit from network externalities, which can contribute to user addiction (Barwick, Chen, Fu, and Li, 2024) or heighten the fear of "missing out" (Bursztyn, Handel, Jimenez, and Roth, 2023). I exclude this mechanism from my structural model for two main reasons. First, the platform analyzed here is small relative to mainstream social media, limiting potential network effects. Second, user utility on this platform is primarily driven by content–the short dramas–rather than social interaction, because users cannot message each other or leave comments on episodes.

Pacing. An alternative explanation for the repeated purchase of small top-up packages is that users are pacing themselves for reasons such as budgeting—that is, they lack sufficient cash on hand and therefore opt for lower-priced packages. However, this explanation is unlikely to be dominant in my context for two reasons. First, even the most expensive package ($29.99) is affordable for most adults. Second, as shown in Table 2, over half of users complete their viewership within a single round (a two-hour window), and more than 90% do so within two rounds. It is unlikely that users would engage in budgeting over such short time frames.

Disutility from monetary payments. The behavioral literature emphasizes that when modeling present bias, the timing of payoffs should be linked to consumption rather than money receipt (DellaVigna, 2018). In my model, the price-sensitivity parameter, ω, can be interpreted as the discounted present value of spending each dollar as perceived by the user. Because present bias arises from the immediate entertainment of video consumption rather than the timing of monetary payments, this assumption is not critical for my analysis. Accounting for delayed consumption yields qualitatively similar results, reinforcing the robustness of my findings. DellaVigna and Malmendier (2006) present a similar argument.

5 Estimation

I estimate my demand model using individual-level data from the short drama platform. The estimation approach is detailed in section 5.1, and the results are reported in section 5.2.

5.1 Estimation approach

To map my model to the data, I assume users, indexed by i, differ in their values δ_i, temptation durations χ_i, and initial endowments a_{i1}. For each individual, the data includes observations on their unlocking history $\boldsymbol{d}_i$, round $\boldsymbol{r}_i$, habit stock $\boldsymbol{h}_i$, token stock $\boldsymbol{a}_i$, and selected top-up packages $\boldsymbol{k}_i^*$ and token amounts $\boldsymbol{b}_{ik^*}$ for every episode t they have watched. I define my estimator and then discuss the identification strategy.

Estimator. The parameters to be estimated are:

$$\boldsymbol{\theta} := \{ \underbrace{\mu_\delta, \sigma_\delta, \overline{\psi}, \beta_\epsilon}_{\text{intrinsic value}} , \underbrace{\kappa, \mu_\chi, \sigma_\chi}_{\text{temptation}} , \underbrace{\alpha_1, \alpha_2, \rho}_{\text{habit}} , \underbrace{\omega}_{\text{price sensitivity}} \}.$$

For intrinsic value, the parameters μ_δ and σ_δ represent the mean and standard deviation of the drama value δ, which follows a Gaussian distribution $\mathcal{N}(\mu_\delta, \sigma_\delta^2)$. The outside value during nighttime is captured by $\overline{\psi}$, and the dispersion of the unlocking taste shifter ϵ_t is captured by β_ϵ. Regarding temptation, κ represents its per-minute magnitude, whereas its duration χ follows a log-normal distribution $\log \chi \sim \mathcal{N}(\mu_\chi, \sigma_\chi^2)$. The habit-formation utility is modeled with a quadratic function parameterized by (α_1, α_2), and ρ controls for the likelihood that a round will end after watching an episode. Finally, ω represents price sensitivity. To normalize utility, I set the scale parameter of the top-up taste shifter, β_ν, to 1, so that it follows a standard Gumbel distribution.

Because the user-specific value δ_i is unobserved, the standard Maximum Likelihood Estimator (MLE) requires computationally intensive integration and is not well-suited to this setting.Therefore, I estimate my dynamic demand model using the Simulated Method of Moments (SMM) by minimizing a simulated criterion function:

$$\widehat{\boldsymbol{\theta}}_{SMM} = \arg\min_{\boldsymbol{\theta}} \left(\widehat{\boldsymbol{m}} - \widehat{\boldsymbol{m}}^S(\boldsymbol{\theta}) \right)' \boldsymbol{W} \left(\widehat{\boldsymbol{m}} - \widehat{\boldsymbol{m}}^S(\boldsymbol{\theta}) \right), \tag{14}$$

where $\widehat{\boldsymbol{m}}$ is a set of moments constructed from the observed data, and $\widehat{\boldsymbol{m}}^S(\boldsymbol{\theta})$ represents the average moments from S simulations of user behavior in the model (McFadden, 1989; Pakes and Pollard, 1989). The weighting matrix $\boldsymbol{W}$ is constructed via bootstrap, following Agarwal (2015). In each simulation, I simulate the sequence of top-up and unlocking decisions made by each user i, taking their initial endowment a_{i1} and the observed part of round sequence from the data. To ensure my estimator reaches a global solution to problem (14), I first conduct a series of grid searches to find an initial value close to the global optimizer, and then apply a local gradient-free Nelder-Mead

algorithm to refine the optimization. Detailed estimation routines and intermediate results are documented in Appendix D.2.

I construct three categories of moments for estimation: habit stock, top-up package purchase, and unlocking. For habit stock, I target the probability of stopping after episode 9 among users with zero initial endowment, conditional on the user's last habit stock being $h_9 \leq 4$ or $h_9 \geq 5$. In addition, I include the average habit stock conditional on purchasing each top-up package. Second, for top-ups, I target the sales of each top-up package across different episode bins (10–20, 21–40, 41–60, and 61–80), which reveal the level and evolution of top-up choices. Regarding unlocking, I track the share of users unlocking the second and third episodes, as well as episodes 4–9, 10–14, and 15–80. These episode ranges are intended to capture the most informative changes in unlocking behavior, as shown in Figure 3a. I also target the fraction of viewership that occurs during nighttime. A complete list of moments is provided in Appendix D.2. All parameters are jointly identified from the selected moments, though certain moments provide more information for specific parameters, for which I discuss the intuition below.

Intuition for identification. I first discuss the identification of temptation, which is the only source of irrationality in the model that generates self-control problems. Therefore, the positive market share of small-sized top-up packages ($4.99, $9.99, $19.99) across all episodes, which reflect time inconsistency, provides insights into the magnitude of per-minute temptation κ. Furthermore, as discussed in section 4.3, users with longer-lasting temptation are more likely to choose higher-priced package, making the shares of different packages informative about the distribution of temptation duration (μ_χ, σ_χ).

Habit formation (α_1, α_2), representing the channel of rational addiction, is identified through variation in habit stock. Empirical evidence suggests users with higher habit stock are more likely to continue watching, which makes the conditional probability of stopping for users at different habit stock levels informative.[50] Moreover, habit formation influences both the flow and continuation value, affecting users' top-up package purchasing decisions. As a result, the habit stock conditional on top-up choices provides information about (α_1, α_2) and the perceived probability of losing habit stock in the next period, ρ.

The intrinsic value is identified from the "residual" variation in unlocking and top-up choices after accounting for the temptation and habit formation. As discussed in section 4.2, users with higher δ are more likely to purchase the $29.99 top-up package and unlock the entire series, whereas those with lower δ are more inclined to stop watching. Thus, the top-up shares of the $29.99 package, the outside option, and the fraction of users who complete the series provide variation to identify this distribution $(\mu_\delta, \sigma_\delta)$. The nighttime outside value $\overline{\psi}$ is informed by the fraction of users unlocking episodes during this period. Lastly, β_ϵ captures the impact of the unlocking taste shifter on each episode, with more negative shocks increasing the likelihood that users will stop watching. As a result, the overall viewership dynamics offer information on the magnitude of this static utility shock.

[50] See Appendix C.2 for a complete discussion of related empirical patterns.

Finally, the price sensitivity, ω, is reflected in users' top-up choices when they pay with real money, making the market shares of top-up plans at different prices informative. Because these packages are virtual products with no distinguishing characteristics beyond their observed prices and token amounts, the identification of ω is free from the usual concerns about endogeneity.[51]

5.2 Results

I report the estimation results for $\widehat{\boldsymbol{\theta}}_{SMM}$ in Table 5 and graphically in Figure 7. The targeted moments are displayed in Appendix D.3. The estimated price sensitivity is $\omega = 0.81$, which translates utility into real monetary terms (USD). I use this numéraire to report all quantitative results.

The first set of parameters describes the distribution of intrinsic value, which is also represented in Figure 7a. The estimated per-episode drama value δ has a mean of -0.0097 ($-1.2¢$) and a standard deviation of 0.092 (11.4¢). The outside value during nighttime is 0.022 (2.7¢) per minute, whereas the daytime outside value is normalized to 0. The idiosyncratic taste shifter ϵ has a scale parameter of 0.055, resulting in a standard deviation of 0.071 (8.8¢). These results indicate that, on average, users experience a net loss of 1.9¢ per one-minute episode compared with their outside options, highlighting the potential loss in user surplus.[52]

Temptation is estimated at 0.15 (18.5¢) per minute, placing it at the 96th percentile of the drama value distribution. As shown in Figure 7b, temptation lasts an average of 6.6 minutes, with a standard deviation of 10.2 minutes and a long right tail from the log-normal distribution. Approximately 25.0% of users experience temptation lasting more than 8 minutes (the duration covered by the \$4.99 package), 10.0% exceed 17 minutes (covered by the \$9.99 package), and 2.9% exceed 37 minutes (covered by the \$19.99 package). Only 1.2% of users exhibit temptation durations over 60 minutes, for whom I set $\chi = 60$, which thereby supports my parametric assumption that temptation lasts less than 60 minutes. The large magnitude and short duration of temptation make self-control problems an important factor in drama-watching behaviors.

I identify a substantial degree of habit formation, as shown in Figure 7c. The estimated habit-formation function (depicted by the red curve) is convex and nearly linear in relation to habit stock, ranging from 0 to 0.87 (\$1.08) per episode. Watching an additional episode in the round contributes, on average, 1.4¢ in per-episode utility via habit formation. This strong level of habit formation aligns with the high willingness to pay observed among users in the data. The estimated probability of a round change, ρ, is 3.9%, influencing users' decisions through their expectations of future habit formation.

Estimation results reveal that both temptation and habit formation are significant in this short-drama-series industry. The structural model leads to further decomposition of the perceived flow

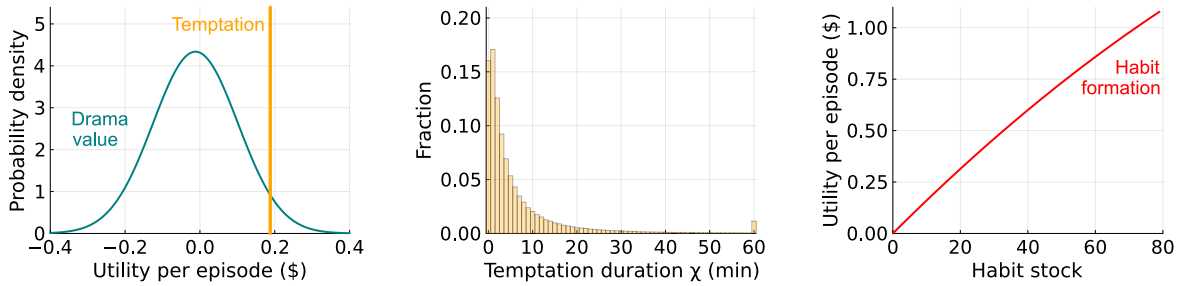
[51] In traditional markets (e.g., used cars), the identification of price sensitivity is subject to omitted variable bias: some attributes (e.g., car condition) are observed by buyers but unobserved by econometricians, and these attributes affect both choice and price. In contrast, top-up options in my setting differ only in two observable dimensions—price and number of tokens—eliminating concerns about such bias in identification.

[52] The average loss from watching one episode equals its average value of $-1.2¢$ minus the average value of the outside option, $0.27 \times 2.7¢$.

Table 5: Estimation: SMM estimators and standard error

Parameter	Estimator	Std	Meaning
1. Intrinsic value			
μ_δ	-0.0097	0.0021	Mean of drama value distribution
σ_δ	0.092	0.0023	Standard deviation of drama value distribution
β_ϵ	0.055	0.0077	Magnitude of unlocking taste shifter
$\overline{\psi}$	0.022	0.0045	Outside value during nighttime
2. Temptation			
κ	0.15	0.0030	Magnitude of naive temptation
μ_χ	1.24	0.022	Mean of temptation duration (log)
σ_χ	1.25	0.033	Standard deviation of temptation duration (log)
3. Habit formation			
α_1	-2.8e-5	4.3e-7	Habit formation utility $\alpha(h) = \alpha_1 h(h - \alpha_2)$
α_2	471.94	2.9e-5	Habit formation utility $\alpha(h) = \alpha_1 h(h - \alpha_2)$
ρ	0.039	0.00049	Probability of entering next round after watching
4. Price sensitivity			
ω	0.81	0.0071	Price sensitivity

Notes: The GMM robust standard errors are reported in the table.



(a) Relative magnitude (b) Temptation duration (c) Habit formation

Figure 7: Estimation: Intrinsic value, temptation, and habit formation

Notes: Panel (a) shows the estimated distribution of drama value δ (teal curve) alongside the temptation level (orange bar), all normalized in dollars using the estimated price sensitivity. Panel (b) displays the estimated distribution of temptation duration χ, which follows a lognormal distribution with an average of 6.57 minutes. For tractability, I cap χ at 60 for all values exceeding this threshold, affecting 1.15% of users. Panel (c) illustrates the estimated function for habit formation $\alpha(h)$, which is nearly linear with respect to users' habit stock.

utility and comparison between the relative magnitudes of these two components, which I report in Appendix D.4.

6 Main results

In this section, I present my quantitative findings on how short video length intensifies self-control problems. I begin by defining in section 6.1 the user surplus loss due to self-control problems. In section 6.2, I assess the comparative static over video length, demonstrating the magnitude of surplus loss due to the short form. Finally, in Section 6.3, I examine counterfactual policies aimed at mitigating self-control problems on the platform.

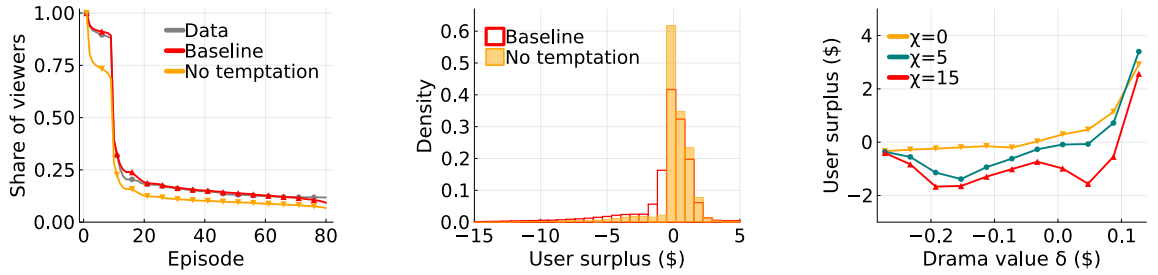
6.1 Surplus loss due to self-control problems

To evaluate welfare implications and efficiency loss due to self-control problems, I define user surplus as the utility from both intrinsic value and habit formation. This metric represents the surplus that a rational user would seek to maximize, treating temptation as an irrational component. My analysis is qualitatively robust with other measures of welfare, such as the pure surplus using only the intrinsic value.

Temptation can induce users to continue watching multiple episodes despite deriving negative intrinsic utility, leading to a potential loss in surplus. Figure 8a illustrates the share of remaining viewers per episode. The baseline model prediction (red line) closely replicates the observed data pattern (gray line), validating the effectiveness of my structural estimation. Removing temptation by setting $\kappa = 0$ results in a substantial decline in viewer retention across episodes. For instance, without temptation, 19.7% of users stop watching after the first episode, versus just 5.7% in the baseline. By the final episode, only 6.7% of users would complete the series, which is a 28.0% reduction from the baseline completion rate of 9.3%. This gap represents users who succumb to temptation and continue watching despite wishing they had stopped, indicating a significant surplus loss attributable to self-control problems.

I compute user surplus and report its distribution in the histogram of Figure 8b. In the baseline scenario (red bars), the average user derives a surplus of \$0.21 from the drama series. However, 49.2% of users experience a negative surplus, with the most extreme one losing \$41.1. Removing temptation (orange bars) significantly reduces the share of negative-surplus users, while those with positive surplus remain unaffected, indicating temptation primarily impacts users who dislike the drama. Without temptation, the average surplus rises to \$0.74. The resulting surplus gap of \$0.53 per user quantifies the *average surplus loss due to temptation*, serving as a key welfare metric for evaluating its broader implications.

Figure 8c further reports the baseline user surplus with respect to drama value δ, which follows a U-shaped relationship as is predicted in section 4.3. Intuitively, users with the highest drama values δ achieve the greatest surplus regardless of temptation duration, whereas those



(a) Share of viewers by episode (b) Surplus distribution (c) Surplus by user type

Figure 8: User surplus and the loss due to temptation

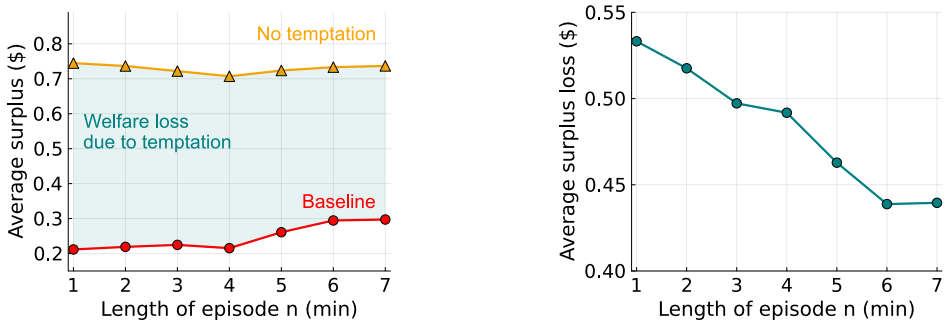
Notes: I present the model outcomes with and without temptation. The counterfactual scenario without temptation is obtained by setting $\kappa = 0$. Panel (a) displays the predicted share of remaining viewers per episode with temptation (red line) and without temptation (orange line), alongside the observed data (gray line). Panel (b) shows the histogram of user surplus under both conditions, with red bars representing the baseline model and orange bars depicting the no-temptation scenario. Panel (c) shows baseline surplus for users with different temptation durations $\chi \in \{0, 5, 15\}$, revealing an U-shape relationship with drama value δ. Users with no temptation ($\chi = 0$) and negative value ($\delta < 0$) incur slightly negative surplus due to the assumption that they must watch the first episode.

with the lowest δ watch only one episode by assumption and incur small negative surplus close to zero. Users with mid-range δ, however, suffer the most from self-control problems, consuming many episodes they do not actually value, leading to the largest surplus losses. Additionally, user surplus generally declines with longer temptation duration χ. Users with no temptation ($\chi = 0$, orange line) generally exhibit positive surplus, providing a rational benchmark, whereas those with longer χ and lower δ mostly experience surplus losses. The greatest surplus loss occurs for users with long temptation duration ($\chi = 15$, red line) and slightly positive δ, who watch many episodes that do not justify their cost. This result aligns with the stylized model that users who are more strongly tempted remain engaged with content they do not truly value.

6.2 Short length amplifies self-control problems

I now explore the mechanism by which short episode length amplifies self-control problems, using a comparative static exercise that varies episode duration. To avoid confounding effects from the platform's dynamic pricing rule, I hold the length of the first nine free episodes fixed. For subsequent episodes, let $n \in \mathbb{N}$ denote the counterfactual episode length, effectively merging n consecutive one-minute episodes into a single, longer episode. Let τ index these counterfactual episodes, with each τ corresponding to a set $Q_\tau \subset \{1, \ldots, T\}$ of original episodes. As in the stylized model in Section 2, I assume that once an episode is unlocked, the user watches it in full. The perceived flow utility is given by:

$$\widetilde{u}_\tau(h, \delta, \psi, \chi, \epsilon; n) = \mathbb{E}_\tau \left[n\delta - \sum_{t \in Q_\tau} \psi_t + \sum_{t \in Q_\tau} \alpha(h_t) + \max\{\chi, n\}\kappa + \sum_{t \in Q_\tau} \epsilon_t \right], \tag{15}$$



(a) Average user surplus

(b) Surplus loss due to temptation

Figure 9: Comparative static: Episode length and user surplus

Notes: These figures illustrate the counterfactual effects of setting a minimum episode length on user surplus. Panel (a) shows the average user surplus with temptation (red line) and without temptation (orange line). The gap in between (teal area) represents the surplus loss due to temptation. Panel (b) plots the evolution of this surplus loss with regard to the episode length n.

where the expectation is taken with respect to the information available in the first original episode of Q_τ, including habit stock, outside option, and taste shifter.[53] As in the stylized model, the key mechanism is a reduction in temptation due to the $\max\{\chi, n\}$ term, which dampens the effect of short-term bias. The counterfactual model is solved analogously to the baseline, using the same top-up framework.

Figure 9 presents the welfare outcomes of this comparative static exercise as episode length (n) increases from one to seven minutes.[54] As n increases, baseline user surplus gradually rises from \$0.21 to \$0.30, reflecting a 42.9% increase. By contrast, the no-temptation surplus remains stable at approximately \$0.73, as episode length only affects outcomes through its interaction with temptation. Consequently, the average surplus loss due to temptation (teal line) declines from \$0.53 to \$0.44, representing an efficiency gain of 17.0% by increasing the length to seven minutes. This result quantitatively confirms the mechanism proposed in Proposition 2, demonstrating shorter episode lengths exacerbate self-control problems.

6.3 Counterfactual policy

I now turn to policy interventions designed to mitigate self-control problems. From the platform's perspective, I examine a minimum token package size requirement, which reduces self-control problems by increasing the episode bundle sold at each top-up. At the individual level, I consider a default time-limit policy, discussed in section 2.4, and a mandatory break following each top-up purchase, both intended to correct individual mistakes.

[53] I assume users observe the habit stock, outside value, and taste shifter associated with the first original episode in Q_τ and form beliefs about the rest, ensuring comparability between the counterfactual and baseline scenarios.

[54] We exclude cases where $n > 8$ to avoid interactions with the top-up package menu, because the \$4.99 package covers an average of seven episodes.

Table 6: Policy evaluations: Minimum package size, default time limit, and mandatory break

Outcomes	Baseline	I. Min. Pack. Size		II. Time Limit		III. Break	
		$> \$4.99$	$> \$9.99$	$\Phi = 0.1$	$\Phi = 0.5$	1 min	5 min
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
I. *Per-user surplus*							
User surplus ($)	0.21	0.79	0.92	0.27	0.53	0.28	0.42
Profit ($)	5.38	3.88	3.49	5.13	3.87	5.11	4.52
II. *User activity*							
Top-up per token buyer ($)	22.42	28.12	29.72	22.12	19.92	22.39	22.31
Frac. token buyer	0.24	0.14	0.12	0.23	0.19	0.23	0.22
Frac. completion	0.09	0.10	0.10	0.09	0.09	0.09	0.08
Frac. time limit complier				0.64	0.95		

Notes: This table presents counterfactual outcomes under three classes of policy interventions. Column (1) reports the baseline results without any policy. Columns (2) and (3) assess minimum package size restrictions by sequentially removing lower-priced top-up options, columns (4) and (5) evaluate default time limits with increasing opt-out costs Φ (in dollars), while columns (6) and (7) analyze mandatory break at top-up with different length. For each scenario, I report average user surplus, average platform profit, average top-up spending per token purchaser, the share of users who purchase tokens, the share who complete the entire series, and, where applicable, the share who comply with the time limit.

Minimum package size requirement. I first consider a class of policy that addresses the shortness of content by restricting the minimum size of token packages.[55] Because episodes are bundled and sold via these packages, raising the minimum package size effectively increases the duration of content per purchase, thereby mitigating self-control problems. I evaluate this policy by sequentially removing the lower-valued top-up packages ($4.99 and $9.99) in counterfactual simulations. Results are reported in section I of Table 6.

Column (2) reports the effect of removing the $4.99 package, which increases average user surplus by $0.58 (from $0.21 to $0.79) and reduces average profit by $1.50 (from $5.38 to $3.88). As predicted by the model, eliminating the smallest package mitigates self-control problems, as reflected by a decline in the share of token purchasers (from 24.0% to 13.8%). Those who do purchase tokens generally value the drama more, with average spending rising to $28.12 and 71% completing the entire series. Columns (3) further excludes the $9.99 package. User surplus increases to $0.92, driven by further reductions in self-control problems, as a growing share of token buyers (83.3%) now complete the entire drama.

Default time limit. Building on the screen time limit concept in Section 2.4, I consider a default time limit policy tied to the duration implied by the user's chosen top-up package. By default, users may watch only the free episodes and the number of episodes covered by their initial top-up. Before viewing begins, they may opt out of the time limit by paying a cost valued at Φ dollars.[56]

[55] Regulating product size is a practical policy tool. For instance, U.S. "loosies" law prohibits the sale of individual cigarettes and mandates a minimum pack size of 20 to discourage youth smoking and facilitate tax enforcement.

[56] This opt-out cost can be interpreted as a friction—such as a time delay—valued at Φ, or alternatively as a monetary fee charged by the platform to override the limit. Due to random taste shifters, adherence to this time limit requires a

The decision problem for users who opt out is identical to the one specified in Section 4. I therefore focus on the solution for users who comply with the time limit, followed by their decision of whether to opt out.

In the model, this time limit is equivalent to restricting users to a single top-up. I introduce an additional state variable $\ell_t \in \{0, 1\}$ to denote the number of remaining top-up opportunities. The system of value functions under this policy is defined below, with the superscript "TL" denoting the time limit policy:

$$\begin{aligned} V_t^{TL}(h, a, \ell; \delta, \psi, \chi, \boldsymbol{b}, \boldsymbol{\nu}) = \quad & \mathbb{1}\{a < c_t\}\mathbb{1}\{\ell > 0\} \max_k \left\{ EW_t^{TL}(h, a + b_k, \ell - 1; \delta, \psi, \chi) - \omega p_k + \nu_k \right\} \\ & + \mathbb{1}\{a \geq c_t\} EW_t^{TL}(h, a, \ell; \delta, \psi, \chi), \end{aligned} \tag{16}$$

$$\begin{aligned} EV_t^{TL}(h, a, \ell; \delta, \psi, \chi) = \quad & \mathbb{1}\{a < c_t\}\mathbb{1}\{\ell > 0\} \mathbb{E}_{\boldsymbol{b}} \left[\log \left(\sum_{k=0}^{K} e^{EW_t^{TL}(h, a + b_k, \ell - 1; \delta, \psi, \chi) - \omega p_k} \right) \right] \\ & + \mathbb{1}\{a \geq c_t\} EW_t^{TL}(h, a, \ell; \delta, \psi, \chi), \end{aligned} \tag{17}$$

where $W_t^{TL}(\cdot)$ and $EW_t^{TL}(\cdot)$ are defined following equations (12) and (13), with $EV_t(\cdot)$ replaced by $EV_t^{TL}(\cdot)$ throughout. The key change is that users may top up only once, i.e., $\ell_1 = 1$. A user with initial endowment a_1, valuation δ, temptation duration χ, and outside option ψ will choose to opt out of the time limit if and only if:

$$\underbrace{EV_1^{TL}(0, a_1, 1; \delta, \psi, \chi)}_{\text{Expected value with time limit}} < \underbrace{EV_1(0, a_1; \delta, \psi, \chi)}_{\text{Expected value without limit (baseline)}} - \underbrace{\omega\Phi}_{\text{Opt-out cost in util}}, \tag{18}$$

which typically applies to users who expect to top up multiple times. Users most affected by self-control problems—those who intend to watch only a few episodes but repeatedly purchase small packages—are most likely to comply with and benefit from the default time limit. High-valuation users remain unaffected, as they consistently buy the largest package and complete the series.

I report the counterfactual policy outcomes under varying opt-out costs (Φ) in columns (4) and (5) of Table 6. Comparing column (4) with the baseline in column (1), even a small opt-out cost substantially improves user surplus: a \$0.10 cost increases surplus by 28.6% (from \$0.21 to \$0.27). However, the policy reduces per-user profit by 4.6% (from \$5.38 to \$5.13), as it lowers top-up purchases on both the intensive and extensive margins. On the intensive margin, paying users reduce their average top-up amount from \$22.42 to \$22.12 due to mitigation of self-control problems. On the extensive margin, some compliers, who would have previously purchased the \$4.99 package anticipating no continuation value, now choose not to top up, further alleviating self-control problems. As the opt-out cost Φ increases to 0.5 in columns (5), the fraction of compliers rises from 64.3% to 94.8%, further increasing user surplus to \$0.47, while reducing platform profits more substantially to \$3.87.

positive cost Φ.

Mandatory break during top-ups. Finally, I consider a counterfactual policy that mandates a break following a top-up purchase by requiring users to watch an ad of ϕ minutes before resuming content. During this period, users experience their outside value but derive no intrinsic utility from the ad itself and do not feel tempted to continue watching it.[57] In this way, breaks serve as a mechanism to shorten the duration of temptation at the moment of the top-up decision.

To incorporate such a feature in my model, I adjust the (expected) value function of top-up equations (9) and (10) with a mandatory break accordingly:

$$\begin{aligned} V_t^B(h,a;\delta,\psi,\chi,\boldsymbol{b},\boldsymbol{\nu}) = \quad & \mathbb{1}\{a < c_t\} \max_k \left\{ EW_t(h, a+b_k;\delta,\psi,\max\{\chi-\phi,0\}) - \omega p_k + \nu_k \right\} \\ & + \mathbb{1}\{a \geq c_t\} EW_t(h,a;\delta,\psi,\chi), \end{aligned} \tag{19}$$

$$\begin{aligned} EV_t^B(h,a;\delta,\psi,\chi) = \quad & \mathbb{1}\{a < c_t\} \mathbb{E}_{\boldsymbol{b}} \left[\log \left(\sum_{k=0}^{K} e^{EW_t(h,a+b_k;\delta,\psi,\max\{\chi-\phi,0\})-\omega p_k} \right) \right] \\ & + \mathbb{1}\{a \geq c_t\} EW_t(h,a;\delta,\psi,\chi), \end{aligned} \tag{20}$$

which, combined with the unlocking value functions (12) and (13) and terminal condition (8), define the new system of value functions. This mandatory break effectively reduces the user's perceived temptation duration into $\max\{\chi-\phi,0\}$ when making top-up decisions. The solutions to this system yields the decision rules for each individual user under this counterfactual policy.

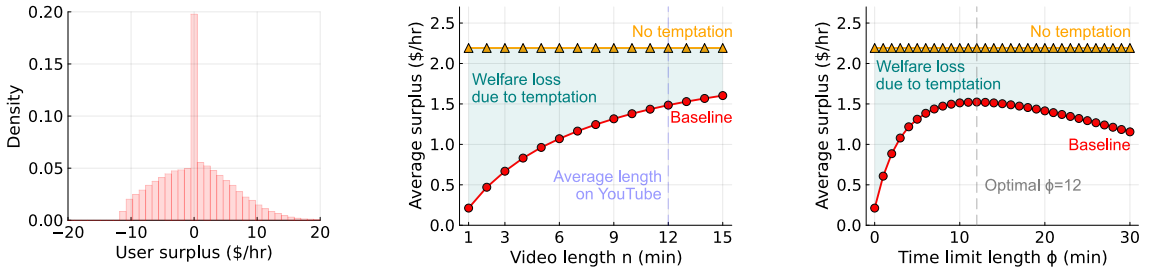
The counterfactual outcomes are reported in columns (6) and (7) for varying break lengths. Relative to the no-break benchmark, column (6) shows a one-minute break increases user surplus from \$0.21 to \$0.28, a 33.3% gain. The policy alleviates self-control problems by reducing the share of users who purchase tokens irrationally, as reflected in the decline in the fraction of token buyers. Furthermore, column (7) suggests extending the break into five minutes raises user surplus to \$0.42, doubling the baseline level.

7 Extension: From short dramas to short videos

I extend my structural analysis on the short drama platform to a broader short video platform such as TikTok. For this purpose, I made following essential assumptions. First, I revert to the stylized model in section 2, omitting both payment and habit formation from the structural framework.[58] Second, for the quantitative analysis, I use the estimated temptation parameter, κ, along with the distributions for intrinsic value ($x := \delta - \psi$) and temptation duration (χ) from my structural estimation reported in Table 5 and Figure 7. Because the stylized model does not predict stopping

[57] The assumption that users derive outside value during ads reflects the reality that they often engage in other activities while ads play, ensuring the policy does not mechanically alter welfare. This assumption likely underestimates the policy's effectiveness—if users actively dislike ads and experience negative utility, the policy's impact on reducing addiction would be even greater.

[58] The key distinction is that TikTok does not charge users, and TikTok videos are not interconnected like episodes in a drama series. Incorporating habit formation would amplify the platform's addictiveness, suggesting that the consumer surplus loss estimated here represents a lower bound.



(a) Distribution of user surplus (b) Increasing video length n (c) Policy: Default time limit ϕ

Figure 10: Extension: Users' hourly surplus on short video platforms

Notes: Panel (a) shows the distribution of users' hourly surplus from short video consumption, with a spike at \$0 reflecting the 14.9% of users who choose not to watch. Panel (b) and (c) report the counterfactual effects of increasing video length n and introducing a default time limit of length ϕ in the short-video-platform setup, respectively. Both panels show the average user surplus with temptation (red line) and without temptation (orange line). The gap in between (teal area) represents the surplus loss due to temptation.

behavior, I focus on welfare outcomes under the scenario in which users are given one hour of short videos—the average daily time spent on TikTok.[59]

User surplus. When given an hour of short videos, the average sample user derives a surplus of \$0.21. Approximately 14.9% of users choose not to watch any videos. However, as shown in Figure 10a, 41.9% of users experience welfare losses due to self-control problems, with an average loss of \$4.73 and the most affected losing up to \$11.3 in surplus per hour. These losses contribute to an overall average hourly surplus reduction of \$1.98 due to temptation, highlighting opportunities for policy interventions.

Short length amplifies self-control problem. I begin by quantifying the effects of increasing video length, as outlined in Proposition 2. Figure 10b presents the average hourly surplus for different video lengths (n). The baseline hourly surplus (red line) rises significantly with n, increasing from \$0.21 at $n = 1$ to \$1.49 at $n = 12$, the average length of YouTube videos.[60] In the efficient benchmark without temptation (orange line), the average user surplus is \$2.19 per hour for all n. Consequently, the hourly surplus loss due to temptation declines from \$1.98 to \$0.70, representing a 64.6% efficiency gain when extending video length to the average YouTube length.

Counterfactual policy: Default time limit. I now consider a practical policy: implementing a default time limit on the short video platform. As described in Section 2.4, the default ϕ-minute time limit is modeled as a default screen time cap with the option to opt out. Proposition 3 shows that users with $x \in (-\kappa, 0)$ begin watching because temptation outweighs their disutility. However, those with temptation duration $\chi \leq \phi$ adhere to the limit and stop after ϕ minutes, while

[59] Source: https://www.emarketer.com/content/5-charts-on-video-marketing-momentum

[60] As of December 2018, the average video length on YouTube was 11.7 minutes (Source: Statista).

users with $\chi > \phi$ opt out and continue watching for the full hour. The policy reduces usage among users with short temptation durations, thereby mitigating self-control problems and increasing user surplus. As discussed in Proposition 3, the optimal time limit balances policy coverage and surplus gain per complier, depending on the distributions of video value δ, temptation duration χ, and temptation magnitude κ.

Figure 10c evaluates the welfare effects of this policy using the estimated model. The no-policy benchmark corresponds to $\phi = 0$, where all users with $\chi > 0$ opt out. Average user surplus is shown by the red curve, which follows an inverted-U relationship with the time limit ϕ: surplus initially rises with ϕ due to broader policy coverage, but eventually declines as each complier consumes more short videos. The optimal time limit is 12 minutes, raising hourly surplus from \$0.21 to \$1.52—a \$1.31 improvement that reduces temptation-induced surplus loss by 86.2%.

Aggregate welfare implication. Finally, I examine the welfare implications for TikTok through a back-of-the-envelope calculation. According to a 2024 eMarketer survey, the average U.S. TikTok user spends 58.4 minutes per day, or 29.2 hours per month, on the platform.[61] This usage translates to an average monthly surplus of \$7.3 per active user relative to their outside options. However, self-control problems lead to a substantial reduction in surplus, with an average monthly loss of \$67.9 due to temptation. Comparative static analysis suggests short video length is a key driver of this efficiency loss. Extending video length to the average YouTube length (12 minutes) could increase the average monthly surplus to \$50.9, recovering \$43.6 (64.2%) of the surplus lost due to temptation. Regarding policy, implementing a default 12-minute time limit could raise the average monthly surplus to \$52.3, recovering \$45.0 (66.2%) of the loss.

Given TikTok's large user base and the addictive nature of short videos, the aggregate welfare effects are substantial. With 150 million monthly active US users reported in March 2023, TikTok's total monthly consumer surplus is estimated at \$1.1 billion. However, 73.8 million of these users (49.2%) suffer from self-control problems, contributing to an estimated monthly welfare loss of \$10.2 billion. Imposing a 12-minute default time limit could recover \$6.8 billion per month, demonstrating significant potential to enhance user welfare. Although this calculation relies on strong assumptions, the magnitude of the potential welfare effect offers useful guidance for policy discussions.[62]

8 Conclusion

This paper explores how small quantities amplify addiction, applying a theoretical framework to the short video industry. I find social media platforms offering short-form videos exploit users' self-control problems, leading to substantial losses in consumer surplus. Comparative statics reveal that shorter video lengths intensify user addiction, whereas policies mandating breaks can

[61] Source: https://www.emarketer.com/content/5-charts-on-video-marketing-momentum

[62] The magnitude of welfare calculation on TikTok is comparable to the estimates for Facebook in Allcott, Braghieri, Eichmeyer, and Gentzkow (2020), providing a sanity check for this exercise.

effectively reduce these self-control problems and alleviate the economic impact of social media addiction. My quantitative works provide an evaluation of potential policies, such as a default screen time limit, in reducing self-control problems. A follow-up experiment is desirable to verify the economic mechanisms and the magnitude of policy gain.

The implications of this study extend beyond social media. Evidence shows consumers frequently purchase small packages of goods (e.g., beer) at a premium, even when buying larger quantities would save money—a behavior likely driven by self-control problems similar to those observed in short video consumption. By connecting these diverse contexts, this paper highlights a broader economic principle: restrictions on quantity or access may serve as effective tools in curbing addictive behaviors across various industries.

References

AGARWAL, N. (2015): "An Empirical Model of the Medical Match," *American Economic Review*, 105, 1939–1978.

ALLCOTT, H., L. BRAGHIERI, S. EICHMEYER, AND M. GENTZKOW (2020): "The Welfare Effects of Social Media," *American economic review*, 110, 629–676.

ALLCOTT, H., M. GENTZKOW, AND L. SONG (2022): "Digital Addiction," *American Economic Review*, 112, 2424–2463.

ARIDOR, G., R. JIMÉNEZ-DURÁN, R. LEVY, AND L. SONG (2024): "The Economics of Social Media," .

AUGENBLICK, N. (2018): "Short-Term Time Discounting of Unpleasant Tasks," *Unpublished manuscript*, 7.

AUGENBLICK, N. AND M. RABIN (2019): "An experiment on time preference and misprediction in unpleasant tasks," *Review of Economic Studies*, 86, 941–975.

BALAKRISHNAN, U., J. HAUSHOFER, AND P. JAKIELA (2020): "How Soon is Now? Evidence of Present Bias from Convex Time Budget Experiments," *Experimental Economics*, 23, 294–321.

BALTAGI, B. H. AND I. GEISHECKER (2006): "Rational Alcohol Addiction: Evidence from the Russian Longitudinal Monitoring Survey," *Health economics*, 15, 893–914.

BANERJEE, A. AND S. MULLAINATHAN (2010): "The Shape of Temptation: Implications for the Economic Lives of the Poor," Tech. rep., National Bureau of Economic Research.

BARWICK, P. J., S. CHEN, C. FU, AND T. LI (2024): "Digital Distractions with Peer Influence: The Impact of Mobile App Usage on Academic and Labor Market Outcomes," *NBER Working Paper*.

BECKER, G., M. GROSSMAN, AND K. MURPHY (1994): "An Empirical Analysis of Cigarette Addiction," *American Economic Review*, 84, 396–418.

BECKER, G. S. AND K. M. MURPHY (1988): "A Theory of Rational Addiction," *Journal of Political Economy*, 96, 675–700.

BEKNAZAR-YUZBASHEV, G., R. JIMÉNEZ-DURÁN, AND M. STALINSKI (2024): "A Model of Harmful Yet Engaging Content on Social Media," in *AEA Papers and Proceedings*, American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, vol. 114, 678–683.

BRAGHIERI, L., R. LEVY, AND A. MAKARIN (2022): "Social Media and Mental Health," *American Economic Review*, 112, 3660–3693.

BURSZTYN, L., B. R. HANDEL, R. JIMENEZ, AND C. ROTH (2023): "When Product Markets Become Collective Traps: The Case of Social Media," Tech. rep., National Bureau of Economic Research.

CHALOUPKA, F. (1991): "Rational Addictive Behavior and Cigarette Smoking," *Journal of political Economy*, 99, 722–742.

COLLIS, A. AND F. EGGERS (2022): "Effects of Restricting Social Media Usage on Wellbeing and Performance: A Randomized Control Trial among Students," *PloS one*, 17, e0272416.

COOK, P. J. AND M. J. MOORE (2002): "The Economics of Alcohol Abuse and Alcohol-Control Policies," *Health affairs*, 21, 120–133.

CRAWFORD, G. S., R. S. LEE, M. D. WHINSTON, AND A. YURUKOGLU (2018): "The Welfare Effects of Vertical Integration in Multichannel Television Markets," *Econometrica*, 86, 891–954.

CRAWFORD, G. S. AND A. YURUKOGLU (2012): "The Welfare Effects of Bundling in Multichannel Television Markets," *American Economic Review*, 102, 643–685.

DELLAVIGNA, S. (2018): "Structural Behavioral Economics," in *Handbook of Behavioral Economics: Applications and Foundations 1*, Elsevier, vol. 1, 613–723.

DELLAVIGNA, S., A. LINDNER, B. REIZER, AND J. F. SCHMIEDER (2017): "Reference-dependent job search: Evidence from Hungary," *The Quarterly Journal of Economics*, 132, 1969–2018.

DELLAVIGNA, S. AND U. MALMENDIER (2004): "Contract design and self-control: Theory and evidence," *The Quarterly Journal of Economics*, 119, 353–402.

——— (2006): "Paying not to go to the gym," *american economic Review*, 96, 694–719.

EBERT, S. AND P. STRACK (2015): "Until the bitter end: On prospect theory in a dynamic context," *American Economic Review*, 105, 1618–1633.

GAO, Q. C., Y. GAO, J. MENG, AND S. YU (2024): "Algorithmic Personalization and Digital Addiction: A Field Experiment on Douyin (TikTok)," *Available at SSRN 4885043*.

GINÉ, X., D. KARLAN, AND J. ZINMAN (2010): "Put Your Money Where Your Butt Is: A Commitment Contract for Smoking Cessation," *American Economic Journal: Applied Economics*, 2, 213–235.

GRUBER, J. AND B. KÖSZEGI (2001): "Is Addiction "Rational"? Theory and Evidence," *The Quarterly Journal of Economics*, 116, 1261–1303.

GUL, F. AND W. PESENDORFER (2001): "Temptation and Self-Control," *Econometrica*, 69, 1403–1435.

——— (2007): "Harmful Addiction," *The Review of Economic Studies*, 74, 147–172.

HOONG, R. (2021): "Self Control and Smartphone Use: An Experimental Study of Soft Commitment Devices," *European Economic Review*, 140, 103924.

LAIBSON, D. (1997): "Golden Eggs and Hyperbolic Discounting," *The Quarterly Journal of Economics*, 112, 443–478.

LAIBSON, D., S. CHANWOOK LEE, P. MAXTED, A. REPETTO, AND J. TOBACMAN (2024): "Estimating Discount Functions with Consumption Choices over the Lifecycle," *The Review of Financial Studies*, hhae035.

LAIBSON, D. I., A. REPETTO, J. TOBACMAN, R. E. HALL, W. G. GALE, AND G. A. AKERLOF (1998): "Self-control and saving for retirement," *Brookings papers on economic activity*, 1998, 91–196.

LEE, R. S. (2012): "Home Video Game Platforms," in *The oxford handbook of the digital economy*, Oxford University Press Oxford, 83–107.

LIU, Z., M. SOCKIN, AND W. XIONG (2020): "Data Privacy and Temptation," Tech. rep., National Bureau of Economic Research.

MACLEAN, J. C., J. MALLATT, C. J. RUHM, AND K. SIMON (2020): "Economic Studies on the Opioid Crisis: A Review," .

MCFADDEN, D. (1989): "A method of simulated moments for estimation of discrete response models without numerical integration," *Econometrica: Journal of the Econometric Society*, 995–1026.

MICHALOPOULOS, S. AND C. RAUH (2024): "Movies," Tech. rep., National Bureau of Economic Research.

MODIRSHANECHI, A., W.-H. LIN, H. A. XU, M. H. HERZOG, AND W. GERSTNER (2025): "Novelty as a drive of human exploration in complex stochastic environments," *Proceedings of the National Academy of Sciences*, 122, e2502193122.

MOSQUERA, R., M. ODUNOWO, T. MCNAMARA, X. GUO, AND R. PETRIE (2020): "The Economic Effects of Facebook," *Experimental Economics*, 23, 575–602.

O'DONOGHUE, T. AND M. RABIN (1999): "Doing It Now or Later," *American Economic Review*, 89, 103–124.

ORPHANIDES, A. AND D. ZERVOS (1995): "Rational Addiction with Learning and Regret," *Journal of Political Economy*, 103, 739–758.

PAKES, A. AND D. POLLARD (1989): "Simulation and the Asymptotics of Optimization Estimators," *Econometrica: Journal of the Econometric Society*, 1027–1057.

ROSENQUIST, J. N., F. M. S. MORTON, AND S. N. WEINSTEIN (2021): "Addictive technology and its implications for antitrust enforcement," *NCL Rev.*, 100, 431.

RUST, J. (1987): "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," *Econometrica: Journal of the Econometric Society*, 999–1033.

SPINNEWYN, F. (1981): "Rational habit formation," *European Economic Review*, 15, 91–109.

STRACK, P. AND D. TAUBINSKY (2021): "Dynamic preference "reversals" and time inconsistency," Tech. rep., National Bureau of Economic Research.

VANMAN, E. J., R. BAKER, AND S. J. TOBIN (2018): "The Burden of Online Friends: The Effects of Giving Up Facebook on Stress and Well-being," *The Journal of Social Psychology*, 158, 496–508.

XIA, T. (2025): "Supporting Content Creators on Two-Sided Markets: Experimental Evidence from a Short-Form Video Platform," *Available at SSRN 5118719*.

ZHEN, C., M. K. WOHLGENANT, S. KARNS, AND P. KAUFMAN (2011): "Habit Formation and Demand for Sugar-Sweetened Beverages," *American Journal of Agricultural Economics*, 93, 175–193.

Appendix

Appendix A Model

Appendix A.1 Represent temptation with present-biased preferences

In this section, I demonstrate my model of temptation can be micro-founded by present bias within a class of standard quasi-hyperbolic discounting models. This formulation shows that the economic insights derived from my baseline model can be generalized within the classic present-biased preference framework.

Consider a large set of one-minute videos $\{1, ..., T\}$, where each minute represents a unit of time. At any given moment t, the agent chooses between watching videos and working (the outside option). Per-minute work provides delayed utility b after period $t + \chi + 1$, whereas watching a video yields immediate utility $b + x$. Consistent with the baseline model, x captures the intrinsic utility of watching a video relative to the agent's outside option.

Regarding time preferences, I assume the agent is present-biased, with the present defined as a window of χ minutes:

$$U_t = \sum_{\tau=t}^{t+\chi} u_\tau + \beta \sum_{\tau=t+\chi+1}^{T} u_\tau, \tag{A.1}$$

where the present-bias parameter $\beta < 1$ applies to all utility realized after χ minutes. This utility specification represents a quasi-hyperbolic discounting framework where the regular discount factor is set to 1 for clarity.

Given the preference structure in (A.1), the set of users experiencing self-control problems can be characterized as:

$$-(1-\beta)b \leq x \leq 0. \tag{A.2}$$

Agents with $x \leq 0$ should not watch any videos. However, due to present bias, those with $x \geq -(1-\beta)b$ find video-watching more appealing because $b + x \geq \beta b$, making them willing to watch videos for the next χ minutes. These users fall into a self-control trap, as their bias continually shifts toward the next set of future videos, leading them to consume more than intention.

This quasi-hyperbolic discounting specification serves as the micro foundation for my baseline model of temptation. The temptation duration χ can be equivalently interpreted as the span of present bias, whereas the magnitude of temptation κ is micro-founded in present bias, taking the value $(1-\beta)b$. This equivalence demonstrates that the economic insights derived from my baseline model are directly applicable to the standard present-biased preference framework.

Appendix A.2 Solution for sophisticated users

Agents might not be fully naïve about their temptation. Let $\widetilde{\kappa} \in [0, \kappa]$ be the perceived temptation, where $\widetilde{\kappa} = \kappa$ corresponds to sophistication and $\widetilde{\kappa} \in (0, \kappa)$ corresponds to partial sophistication. The perceived flow utility (1) can be generalized into:

$$\widetilde{u}_\tau^t(x, \chi) = \underbrace{x}_{\text{intrinsic utility}} + \underbrace{\mathbb{1}\{\tau < t + \chi\}\,\widetilde{\kappa}}_{\text{perceived temptation}} + \underbrace{\mathbb{1}\{\tau < t + \chi\}\,(\kappa - \widetilde{\kappa})}_{\text{unperceived temptation}}. \tag{A.3}$$

For users who are susceptible to self-control problems ($x \in (-\kappa, 0)$) in Proposition 1, I consider the following two situations:

1. **Users with $x \in (-\widetilde{\kappa}, 0)$ decide not to watch.** Because temptation $\widetilde{\kappa}$ is perceived, these users anticipate that once they begin, the irrational temptation will persist for the next χ videos, preventing them from stopping. They thus evaluate their total utility as $\chi(x + \kappa) + (T - \chi)x < 0$ and choose not to watch.

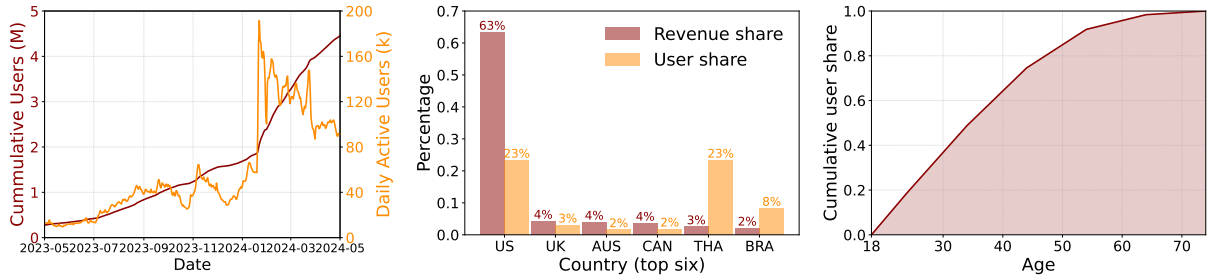
2. **Users with $x \in (-\kappa, -\widetilde{\kappa})$ fall into self-control problems.** Although these users recognize the temptation $\widetilde{\kappa}$, they do not expect to keep watching beyond the temptation duration χ, since the perceived temptation is insufficient to justify continuation. Yet the unperceived component is strong enough to trigger impulsive consumption, resulting in the self-control problem described in Proposition 1.

The general takeaway is that (partial) sophistication diminishes self-control problems. Under full sophistication ($\kappa = \widetilde{\kappa}$), users behave rationally and watch videos if and only if $x \geq 0$.

Appendix B Background

I provide further details on the platform in this section, including the descriptive facts on users, dramas, revenue, pricing, and the supply-side details.

Users. The platform size and user characteristics are summarized in Figure B.1. According to panel (a), the total number of users has reached 4 million worldwide and the Daily Active Users (DAU) remains around 100,000 by the May of 2024. Figure B.1b shows a large fraction of users are from US (23%), Thailand (23%), and Brazil (8%), whereas most of the platform revenue is contributed by US users (63%). Country boundaries create a natural form of market segmentation, and I therefore focus on the biggest market, the US market, in our main analysis. Panel B.1c indicates users of this platform are mainly young people, with 36% aged below 30, 63% below 40, and 85% below 50.

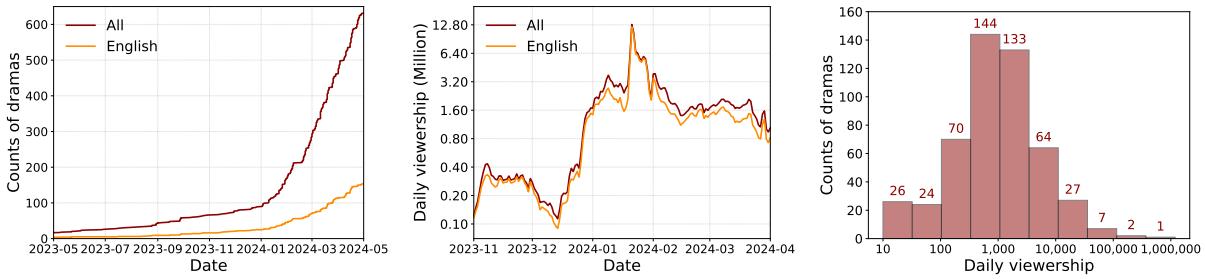


(a) Number of users (b) Geographical distribution (c) Cumulative distribution of age

Figure B.1: Users on the platform

Notes: Panel (a) shows the evolution of the accumulative number of users (left, red) and the daily active users (right, orange) between May 2023 and May 2024. My sample period is highlighted by shaded areas. Panel (b) reports the six countries with the highest revenues and most users up to May 8, 2024. Panel (c) displays the cumulative distribution function of age for new users who entered between March 1, 2024, and May 14, 2024. These aggregate data are directly provided by the platform.

Drama. I report the number and viewership of dramas in Figure B.2. Figure B.2a shows the evolution of the number of drama-language pairs over time. Over my sample period, the total number of drama-episode pairs increased from 67 in November 1, 2023, to 461 in March 31, 2024, where the number of English dramas increased enormously from 16 to 114. The corresponding viewership trends are displayed in Figure B.2b. Consistent with the geographical distribution in Figure B.1b, 83% of viewership during my sample period is in English. The daily viewership is 1.7 million episodes on average, which increased drastically in January 2024 due to the introduction of one "superstar" drama. I further report the distribution of daily viewership in Figure B.2c, which showcases a remarkable heterogeneity between dramas. The superstar drama, which was introduced on January 19, 2024, has about 1 million daily viewership, whereas most dramas have viewerships below 10,000 per day.



(a) Number of dramas (b) Aggregate viewership (c) Viewership distribution

Figure B.2: Drama entry and viewership

Notes: Panel (a) shows the evolution of the number of drama-language pairs on this platform over time. The shaded area indicates for the sample period for my study. Panel (b) reports the aggregate viewership of dramas over my sample period for all dramas and English dramas. Panel (c) plots the distribution of viewership on drama-language level. Note that viewership is displayed in log scale in both the y-axis of panel (b) and x-axis of panel (c).

Revenue. I present the time-series sequence of the three revenue sources in Figure B.3. The main source of all time is the token income, whereas the revenue from subscription has been increasing quickly since February 2024. The ad revenue remains at a relatively low level for the whole sample period.

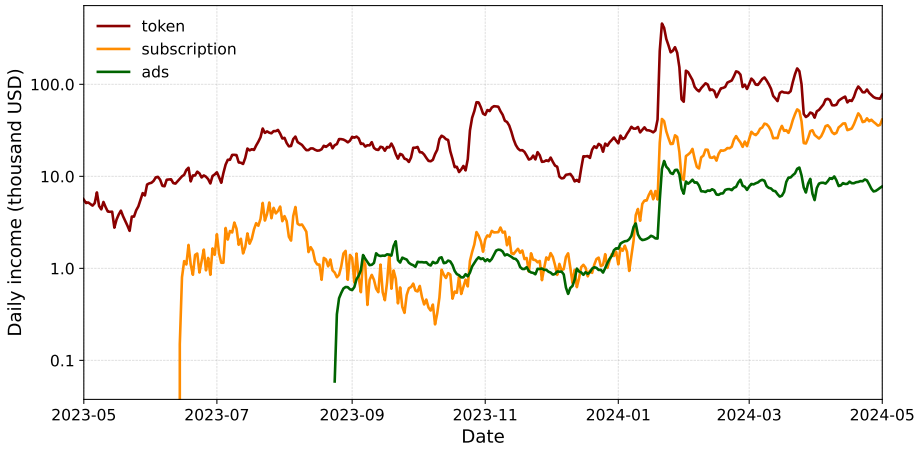
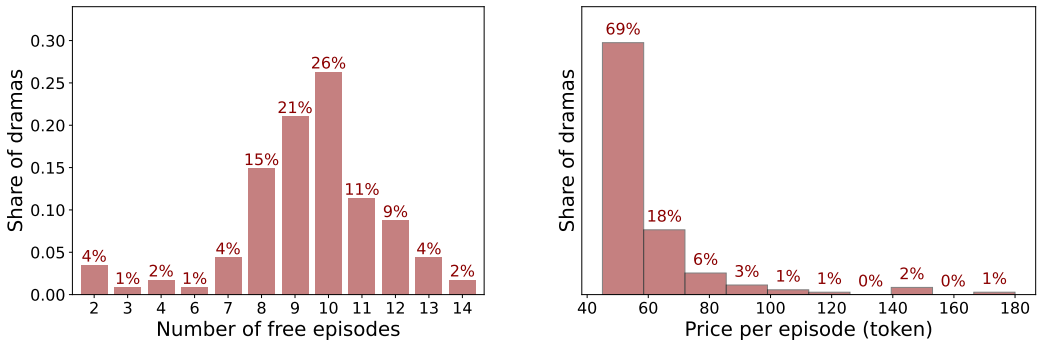


Figure B.3: Three sources of revenue: Token, subscription, and ads

Notes: This figure shows the evolution of different daily revenue sources. My data for subscription and ads started in June 2023 and August 2023, respectively. The level of revenue is displayed in log scale on the y-axis.

Drama pricing. In Figure B.4a, I report the length distribution of free episodes for each drama. Most dramas have eight to 10 free episodes to attract users, whereas the minimum and maximum numbers are two and 14. Figure B.4b displays the price distribution of episodes by dramas in terms of platform tokens. Within a drama, each episode costs the same amount of tokens to unlock, whereas price varies substantially across dramas. Most drama episodes are priced between 45 and 80 tokens.



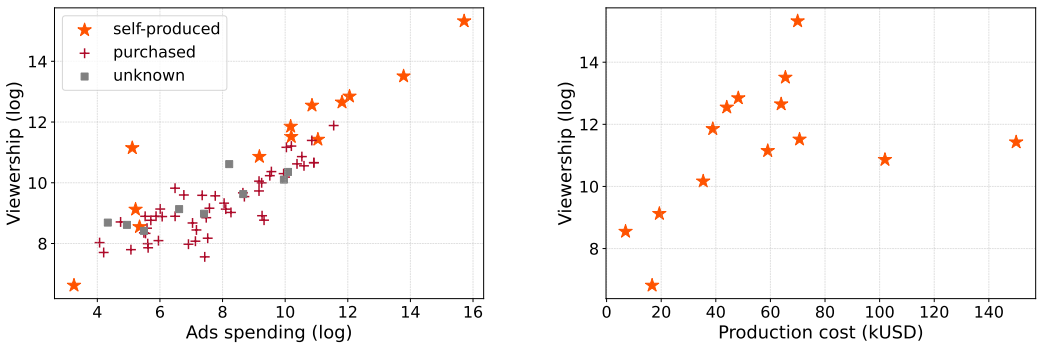
(a) Length of free episodes

(b) Price per episode

Figure B.4: Drama prices and token purchase

Notes: Panel (a) shows the distribution of the number of free episodes by dramas for all English dramas in my sample. Panel (b) plots the distribution of price per episode for those dramas.

Drama supply. In Figure B.5a, I show a positive relationship between ads spending and total viewership on drama level. Consistent with the aggregate description, self-produced dramas in general have higher ads spending and viewership. Conditional on ads spending, self-produced dramas also have higher viewerships than purchased ones, which indicates for some unobserved quality differences. I also observe the total production costs for platform-owned dramas, which are plotted in Figure B.5b against their total viewership. The production costs for dramas are fairly low, because most dramas are produced with costs below $80,000. In general, dramas that cost more are more popular among users, highlighting the need to consider heterogenous quality among different dramas.



(a) Ads spending

(b) Production cost

Figure B.5: Drama ownership, ads spending, and production cost

Notes: Panel (a) plots the total viewership and total ads spending between November 1, 2023, and March 31, 2024, for each drama. The plot distinguishes three types of ownership: self-produced (orange, star), purchased from independent producers (red, square), and dramas whose ownership is unknown to me (grey, circle). In panel (b), I plot the production costs for self-produced dramas against their total viewership.

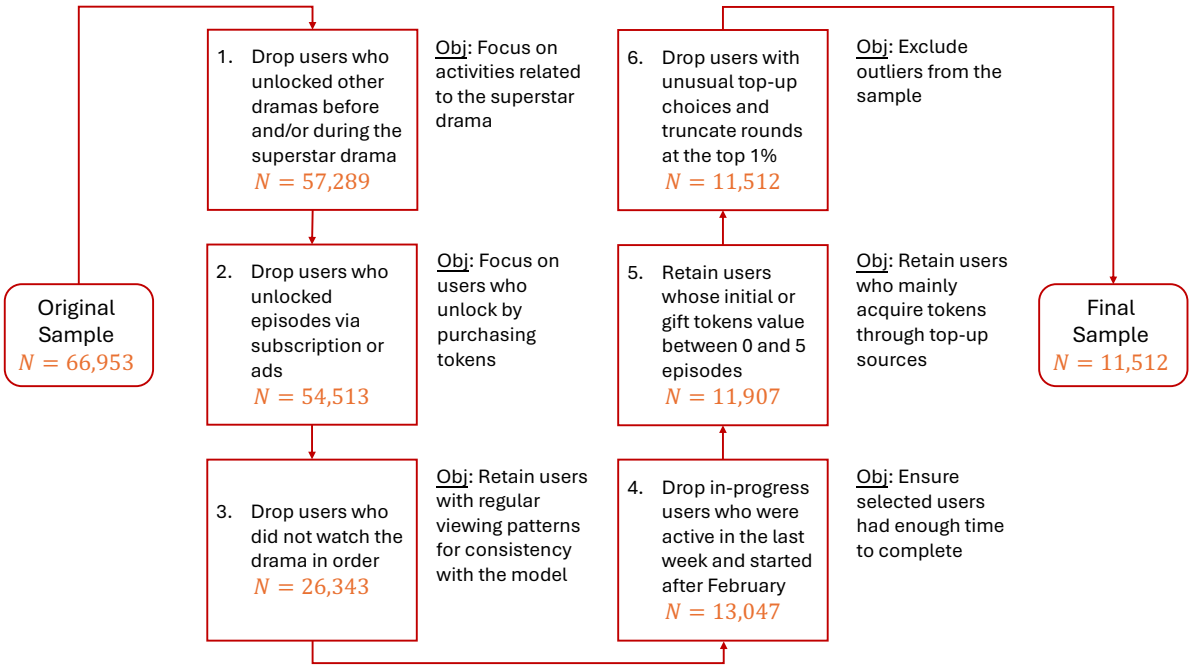


Figure C.1: Data cleaning procedure

Notes: This figure summarizes the data cleaning procedure step by step in the form of a diagram. Each box shows the main filtering step, the resulting sample size (in orange), and the objective of the step (in text alongside).

Appendix C Empirics

I present more empirical evidence from the platform to complements my main analysis.

Appendix C.1 Data cleaning procedure

This section provides a detailed discussion of the data cleaning procedure. The full process is summarized in six steps in Figure C.1, with each step accompanied by the resulting sample size and its underlying rationale. I elaborate on each step below.

1. *Drop users who unlocked other dramas before and/or during the superstar drama.* This step excludes users with prior experience watching other dramas on the platform, thereby removing potential confounds from pre-existing habit formation. It also ensures that all top-up activity is attributable solely to the superstar drama, aligning the sample with the structural model, which focuses exclusively on that title. The resulting sample remains broadly representative, retaining 85.6% of users, suggesting that this step is unlikely to introduce substantial selection bias.

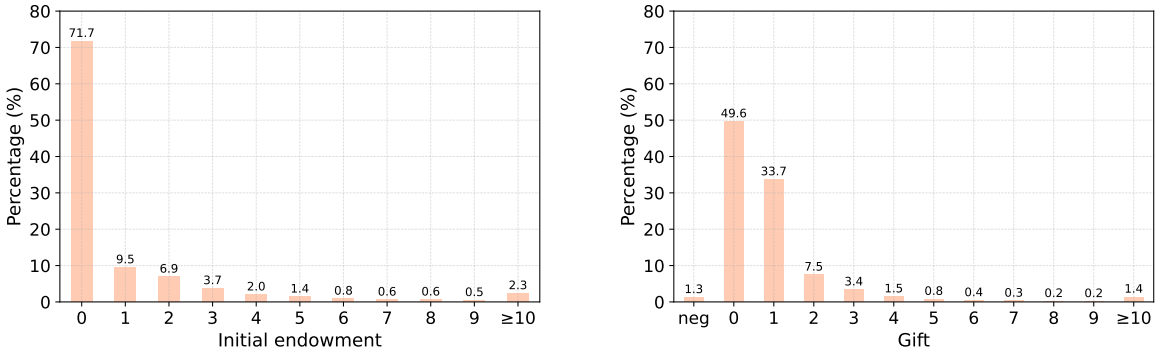
2. *Drop users who unlocked episodes via subscription or ads.* This step ensures that sample users unlock episodes exclusively through token purchases, which is the primary focus of the

analysis. As the vast majority of users (95.2%) use tokens only, this restriction is unlikely to introduce significant selection bias.

3. *Drop users who did not watch the drama in order.* This step retains users with regular, sequential viewing patterns, consistent with the structural model's assumption of episode-by-episode decision-making. It simplifies the modeling exercise by abstracting from the choice of viewing sequence. The cost is a substantial reduction in sample size, with 51.7% of users excluded. If users who watch out of order are less "rational" or more erratic, this selection likely results in a sample of more rational users—potentially leading to an underestimation of self-control problems, thereby reinforcing the paper's core argument.

4. *Drop in-progress users who were active in the last week (March 24–31, 2024) and those who started after February, 2024.* This step ensures that users in the sample had sufficient time to complete the drama, so that incomplete viewing reflects a choice not to continue. While this restriction reduces the sample size by 49.5%, it is unlikely to introduce systematic selection, because the timing of when users begin watching is plausibly random.

5. *Retain users whose initial or gift token value falls between 0 and 5 episodes.* This step excludes users who received a substantial number of tokens from non–top-up sources, such as daily log-ins, which appear as large initial endowments or gift token balances. Figure C.2 shows the distribution of these token values among users entering this step, indicating that most users receive minimal amounts from such sources. This restriction removes only 8.7% of users, suggesting that it is unlikely to introduce significant selection bias.



(a) Distribution of initial endowment

(b) Distribution of gift tokens

Figure C.2: Distribution of initial and gift tokens

Notes: Panels (a) and (b) report the distributions of initial endowments and gift tokens, respectively, among users entering Step 5 of the data cleaning procedure in Figure C.1. Token amounts are normalized by the number of tokens required to unlock one episode.

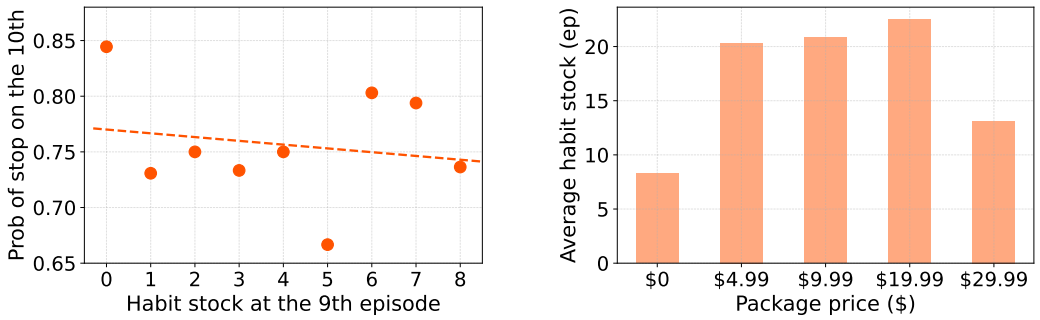
6. *Drop users with unusual top-up choices and truncate rounds at the top percentile.* This step removes outliers from the sample. Unusual top-up packages are defined as those with non-standard prices (i.e., outside the four main packages analyzed in this paper) or token amounts

with conditional frequencies below 5%. These packages typically originate from limited-time promotions or pricing experiments conducted by the platform. I also exclude users in the top percentile of the round distribution to eliminate extreme viewing patterns. This step introduces minimal selection, removing only 3.3% of users.

Appendix C.2 Rational addiction

In Figure C.3, I provide suggestive evidence of habit formation in users' drama-watching behaviors, which is commonly interpretted as rational addiction (Becker and Murphy, 1988). Panel C.3a shows the fits of regression described in section 3.4. Among users who have to top up before unlocking the 10^{th} episode of the superstar drama, a negative relationship exists between the likelihood of stopping and the corresponding habit stock. Those with one additional unit of habit stock are, on average, 0.3% more likely to continue watching.

Panel C.3b shows the average habit stock conditional on selecting each top-up package. A higher-priced package is in general selected by users with higher habit stock, supporting the habit-formation effect. The drop for the $29.99 package is mechanical, because users no longer need the largest package when they have already watched most episodes, and few remain to unlock.



(a) Conditional probability to stop

(b) Top-up package choice

Figure C.3: Suggestive evidence on rational addiction (habit formation)

Notes: Panel (a) shows that, among users who have to top up before unlocking the 10^{th} episode of the superstar drama, a negative relationship exists between the likelihood of stoping and the corresponding habit stock. Those with one additional unit of habit stock are, on average, 0.3% more likely to continue watching. Panel (b) presents the average habit stock based on the top-up package choice, where higher-priced packages are generally chosen by users with higher habit stock, indicating habit formation. The drop for the $29.99 package is mechanical, because users no longer need the largest package when they have already watched most episodes and few remain to unlock.

Appendix C.3 Viewership conditional on time of a day

Figure C.4 reports the number of unlockings conditional on hour of a day (Eastern time). It suggests the usage between 1am and 10am is below average, depending on which I define the "night-time" outside option for my structural analysis.

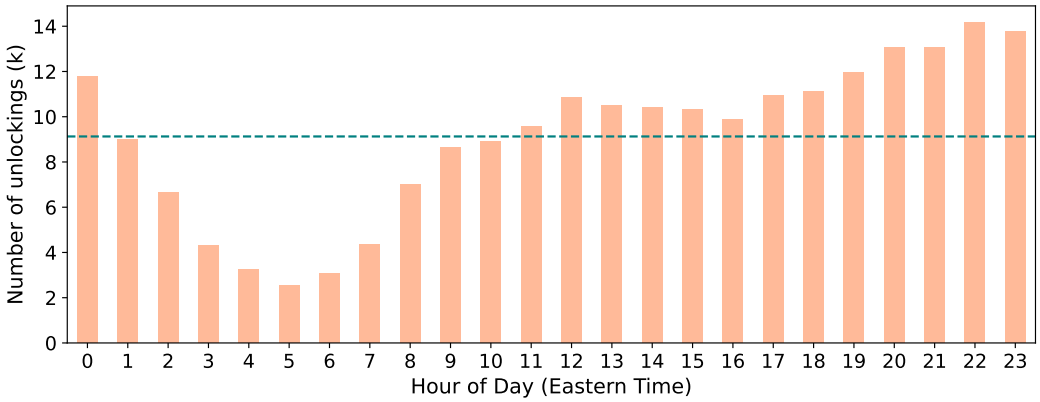


Figure C.4: Usage conditional on time of a day

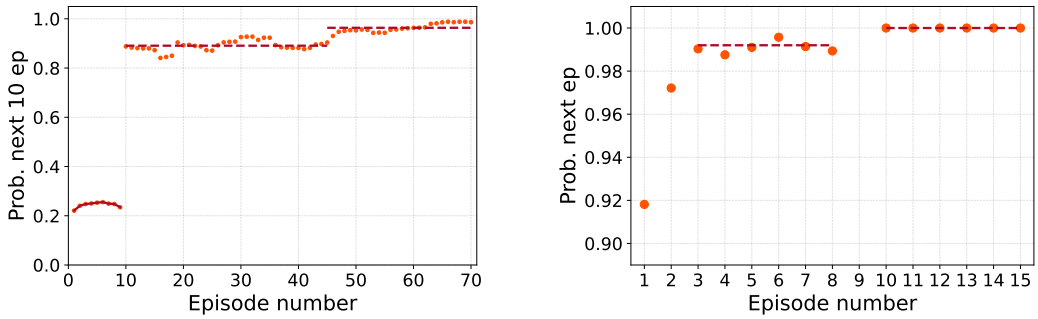
Notes: This figure represents the total number of unlocking activities that happened in each hour of a day within my sample. The hour is based on Eastern time. The teal, dashed line is the average number of unlocking per hour.

Appendix C.4 Minimal learning effect

In this section, I provide evidence that learning or information effects are not the primary drivers of users' drama-watching behavior. Two key observations support this conclusion. First, the conditional probability of continuing to watch stabilizes within one or two episodes, indicating learning happens quickly, making the assumption of perfect information reasonable. Second, drama-watching patterns remain consistent even after 40 episodes, suggesting the influence of learning is minimal.

Evidence 1: The probability of continuing watching stabilizes quickly. In Figure C.5a, I present the evolution of the probability of continuing to the next 10 episodes, conditional on each episode. This probability initially rises from 22.1% at the first episode to 24.8% by the third, and then stabilizes around 25% until the ninth episode. A sharp increase occurs at the 10th episode—when users must start paying for new episodes—pushing the probability to 89.1%, where it remains stable between episodes 10 and 45. As the show nears its conclusion (episodes 45–70), the probability of continuing rises further to 96.3%. This analysis yields two key insights. First, the conditional probability increases significantly only during the first two episodes, after which it stabilizes quickly, suggesting learning occurs rapidly. Second, the price effect at the 10th episode is substantial, causing a 64.1-percentage-point increase in the continuation probability, whereas the learning effect accounts for only a few percentage points overall. Thus, I conclude learning plays a minimal role in shaping users' drama-watching behavior.

These insights are further confirmed in Figure C.5b, which illustrates the evolution of the conditional probability of continuing to the next episode. Initially, this probability is relatively low, at 91.8% for the first episode and 97.2% for the second. From the third episode onward, it stabilizes around 99%, indicating any learning effect occurs primarily within the first two episodes.

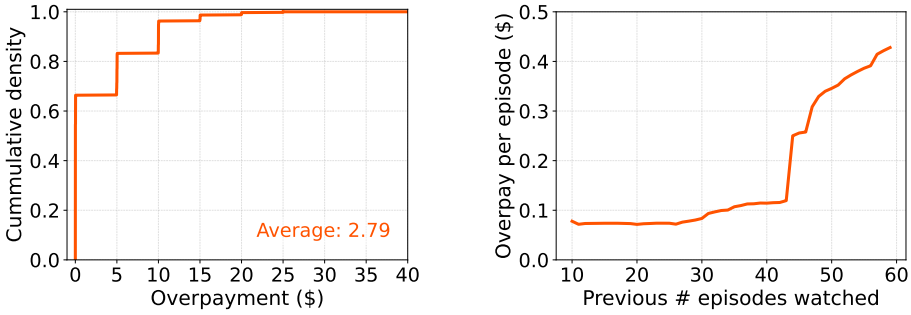


(a) Next 10 episodes (b) Next episode

Figure C.5: Conditional probability of watching the next few episodes

Notes: This set of figures presents the probability that a user will continue watching 10 (panel (a)) or one (panel (b)) conditional on watching each episode. The orange circles are the conditional probability, whereas the red, dashed lines provide fitted values. For a clean comparison, the conditional probability is calculated among users with zero initial token endowment.

Evidence 2: Persistent overpayment conditional on having previously watched many episodes. Figure C.6a provides further evidence of overpayment among users who have already watched 40 episodes. If learning played a central role, these users should have gained substantial information from their viewing history and, consequently, made more rational package purchasing decisions. However, 33.6% of users continue to overpay for tokens, with an average overpayment of $2.79 for the remaining 40 episodes. Both figures are substantial and closely align with the unconditional distribution reported in Figure 4c, where 39.1% of users overpay, with an average overpayment of $5.51 for 80 episodes. This consistency suggests that the observed irrational spending behavior cannot be attributed to learning.



(a) Distribution after 40th episode (b) Average overpayment

Figure C.6: Overpayment conditional on having previously watched many episodes

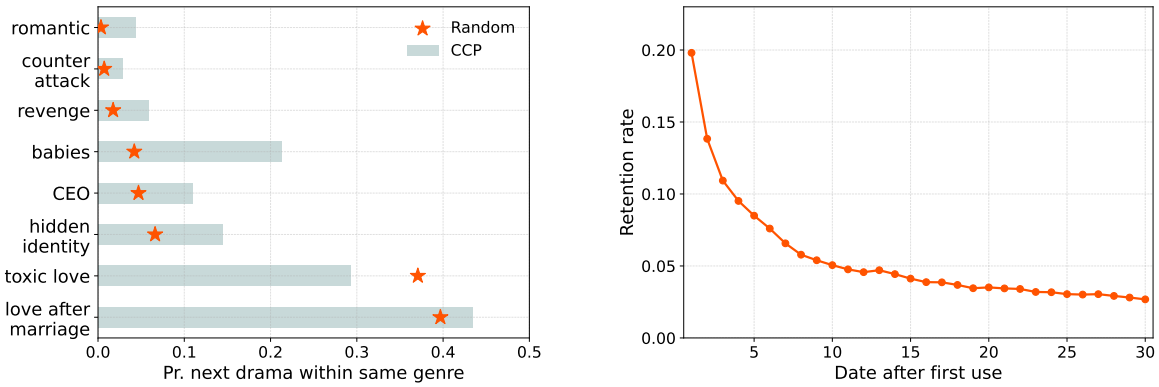
Notes: Panel (a) presents the cumulative distribution of overpayment among users who made at least one package purchase after the 40th episode, following the same idea as the unconditional distribution in Figure 4c. Overpayment is measured relative to the optimal spending, imputed using the average token quantity in each package from Table 3. The natural lower bound for overpayment is 0. Panel (b) plots the average overpayment per episode conditional on users who have previously watched 10-60 episodes.

Figure C.6b generalizes this idea, showing that average overpayment remains seven to ten

cents per episode for users who have already watched 10–40 episodes. This pattern suggests that learning is not the main driver of overpayment, since additional information does not reduce overspending. After 40 episodes, average overpayment rises sharply due to the finite-horizon structure and selection effects, as the remaining users are most vulnerable to self-control problems.

Appendix C.5 No spillover across dramas

I do not find evidence for addiction beyond a drama, which justifies my analysis on the single, superstar drama. I check the history dependence on the level of both genre choice and platform usage. In Figure C.7a, I report the conditional choice probability (CCP) of a user who has just finished one drama to watch another drama that shares the same genre (the teal bar), along with the probability if the user just randomly selects the next drama following the empirical distribution (the orange star). The results show no sizable difference between the two, let alone the fact that user selection may mechanically make CCP greater than random probability. I therefore conclude no significant addiction occurs on the genre level. On the platform level, I plot the retention rates in Figure C.7b, which is defined as the fraction of users who are still active on the X^{th} day after their first appearance. The retention rate is fairly low relative to other digital platforms. For example, the 30-day retention rate on this platform is 2.7%, which is lower than Netflix (30.0%), Netflix (14.4%), and GoodNovel (7.5%). I therefore conclude no strong evidence exists for addiction beyond the drama level.

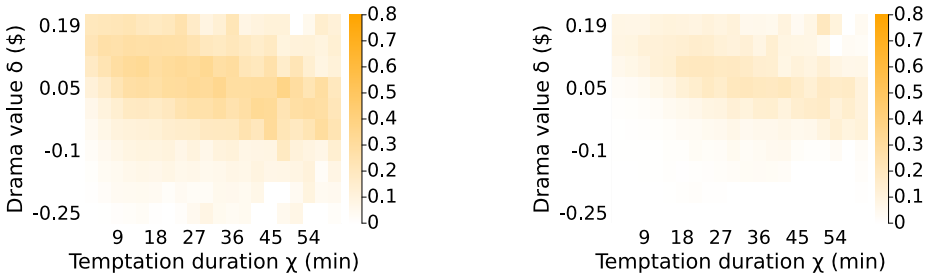


(a) No spillover within a genre

(b) No spillover within the platform

Figure C.7: No spillover across dramas

Notes: The teal bars in panel (a) show that, conditional on users just finishing a drama with the corresponding genre, the probability that the next drama they watch will share the same genre. The orange star indicates the probability if the users randomly select the next drama to watch based on the empirical distribution. Panel (b) plots the retention rate on the platform with varying number of days, which is the fraction of users who are still active X days after their first use. The 30-day retention rate is 2.7%, which is lower than other online platforms such as TikTok (30%), Netflix (14.4%), and GoodNovel (7.5%).



(a) Choice probability of $4.99

(b) Choice probability of $9.99

(c) Choice probability of $19.99

(d) Choice probability of $29.99

Figure D.8: Choice probability of each top-up package between episodes 10 and 15

Notes: The four heat maps report the estimated choice probability of each top-up package between episode 10 and 15 conditional on value δ and temptation duration χ. The warmer the color, the larger the choice probability. The probabilities from the four figures do not sum to 1, with the gap being the share that choose the outside option ($k = 0$).

Appendix D Structural Model

Appendix D.1 Solution characterization

Top-up package choices. With the full solution in hand, I can characterize the conditions under which users purchase each top-up package. For better intuition, consider the case without taste shifters ($\nu_k \equiv 0$). When $t > 9$ and the token balance is $a_t = 0 < c_t$, the perceived value of top-up package k is given by:

$$\pi_{kt}(h; \delta, \psi, \chi) = EW_t(h, b_k; \delta, \psi, \chi) - \omega p_k. \tag{D.1}$$

The top-up choice, therefore, depends on the shape of the value function $EW_t(h, b_k; \delta, \psi, \chi)$, which is monotonically increasing in all its arguments. For illustration, I plot the estimated choice probability for each package, conditional on user type, in Figure D.8.

The outside option ($k = 0$) yields a value of zero and will be selected if $\pi_k(h; \delta, \psi, \chi) < 0$ for all $k \in \{1, ..., K\}$, which typically occurs when users dislike the drama (low δ). When the intrinsic value δ is low enough, users will be unwilling to unlock the episode, regardless of their token stock, habit stock, or temptation duration. In this extreme scenario, $\lim_{\delta \to 0} \pi_{kt}(h; \delta, \psi, \chi) = -\omega p_k < 0$. Given the monotonicity and continuity, a cutoff value $\underline{\delta}(h, \psi, \chi)$ exists below which users will choose not to top up. As Figure D.8 shows, users with lower δ are more likely to select

the outside option, as indicated by the lower choice probability for all packages.

At the other extreme, the largest package becomes most desirable when users plan to watch many episodes, which occurs when they derive high intrinsic value δ. When δ is sufficiently large, users believe ex ante they will finish the entire series, turning the problem into cost minimization. The highest-priced package, which offers the lowest price per token, is thus chosen early on, for example, around period $t = 10$. By monotonicity and continuity, a cutoff value $\overline{\delta}(h, \psi, \chi)$ must exist above which users opt for the largest package K. This intuition is supported by Figure D.8d, where most high-value users purchase the $29.99 package between episode 10 and 15.

The medium-priced packages ($k \in \{1, \ldots, K-1\}$) are mostly purchased by users with medium intrinsic values, driven by varying durations of temptation. All else being equal, the longer this perceived temptation lasts, the more episodes the user anticipates watching. Mathematically, the value function $EW_t(h, a; \delta, \psi, \chi)$ increases more rapidly in a when $a \leq \chi$, and then slows down afterward. As a result, users with longer temptation durations are more likely to purchase higher-priced packages with more tokens. However, these users often continue topping up when running out of tokens again and ultimately finish the entire series, highlighting the self-control problem. This prediction is confirmed by the first three panels in Figure D.8, where the average temptation duration increases with the price of the top-up packages purchased by users. The choice probability of these medium-priced packages therefore informs the distribution of temptation duration.

Unlocking choices. With sufficient tokens, the user will unlock an episode if and only if their perceived value exceeds the outside option, namely, $u_t(h, \delta, \psi, \chi, \epsilon) + \mathcal{C}_t(h, a; \delta, \psi, \chi) > 0$. Because both the flow utility and continuation value are monotonically increasing in drama value δ, habit stock h, and temptation duration χ, the unlocking decision follows a threshold rule: users will continue unlocking if and only if the combination of intrinsic value, habit stock, and temptation is sufficiently high. For identification, the unlocking decision thus provides information on the relative magnitudes of habit formation, temptation, and the drama's intrinsic value.

Appendix D.2 Estimation algorithm

For estimation, I calculate model moments associated with each set of parameter $\boldsymbol{\theta}$ by simulating the economy 20 times and taking the average. To ensure my estimation routine converges to the global solution, I apply a two-step procedure that combines a global grid search with a subsequent local solver in approaching the SMM optimization problem (14).

Global grid search. I first conduct a global search over a large grid on my parameter space. I calculate the objective function on each grid point and locate the on-grid minimizer. The outcome-minimized objective function takes a value of 2361.51 and yields the initial estimators reported in Table D.1.

Local solver: Nelder-Mead algorithm. I then refine the first-step global minimizer $\widehat{\theta}^{(1)}_{SMM}$ with a local solver. I apply a gradient-free solver, Nelder-Mead algorithm packed in `Optim.jl`, using programming software Julia. This process yields a final estimator with an objective value 2353.94.

Table D.1: Two-stage estimation procedure

	Parameter											
	μ_δ	σ_δ	$\overline{\psi}$	β_ϵ	κ	μ_χ	σ_χ	α_1	α_2	ρ	ω	obj.
First stage	-0.0094	0.091	0.017	0.056	0.15	1.21	1.26	0.000028	470	0.039	0.81	2361.51
Second stage	-0.0096	0.092	0.022	0.055	0.15	1.24	1.25	0.000028	472	0.039	0.81	2353.94

Appendix D.3 Targeted moments

I report the targeted moments in Table D.2. I normalize the habit stock by the total number of episodes (80), top-up user share by the number of users in each episode bin, and unlocking share by the total number of users. My model predictions match reasonably well with data moments.

Appendix D.4 Decomposition of perceived utility

The estimated model suggests both rational and irrational addiction are significant in this short drama industry. I further decompose the perceived flow utility into three components: intrinsic value, temptation, and habit formation. Figure D.9 displays the average value of each component for users who have watched episode t. The teal area represents the average intrinsic value per episode, increasing from $-1.2¢$ for the first episode to $16¢$ for the last, reflecting user selection over episodes as low-valuation users gradually stop watching. The sharp rise between episodes 10 and 14 is due to an exogenous increase in episode prices from 0 to 1 token, leading over three-quarters of users to stop watching. On average, users derive \$1.43 intrinsic utility from this drama relative to their outside options.

Addiction utility is substantial relative to intrinsic value. The orange area represents the value of temptation, valued at $19¢$ per one-minute episode, providing an average perceived value of \$3.67 for each user from the drama series. However, this temptation is irrationally perceived and should be excluded from the subsequent welfare analysis. The red area at the top reflects the average level of habit formation for each episode, which gradually increases as users build habit stock over episodes. For most episodes, habit formation is the primary component of utility, explaining users' high willingness to pay for this drama, conditional on paying. However, the majority of users stop watching between episodes 10 and 15, which reduces the average impact of habit formation across users. On average, users derive \$4.30 from habit formation throughout the series, suggesting the realized value is similar between rational and irrational addiction.

Table D.2: Targeted moments

No	Moment	Note	Data	Model
1	E_ccp_j10_h_leq_4	Mean of stopping CCP for the 10th ep, conditional on last h <= 4 and no endowment	0.76	0.80
2	E_ccp_j10_h_5_8	Mean of stopping CCP for the 10th ep, conditional on 5 <= h <= 8 and no endowment	0.75	0.73
3	E_h_k_0	Mean of last habit stock when user chooses k = 0, normalized by T	0.09	0.11
4	E_h_k_1	Mean of last habit stock when user chooses k = 1	0.25	0.27
5	E_h_k_2	Mean of last habit stock when user chooses k = 2	0.27	0.29
6	E_h_k_3	Mean of last habit stock when user chooses k = 3	0.30	0.19
7	E_h_k_4	Mean of last habit stock when user chooses k = 4	0.17	0.10
8	q_k1_j_10_20	User share of package k=1 sold at t $\in$ [10,20]	0.34	0.55
9	q_k1_j_21_40	User share of package k=1 sold at t $\in$ [21,40]	0.18	0.32
10	q_k1_j_41_60	User share of package k=1 sold at t $\in$ [41,60]	0.12	0.30
11	q_k1_j_61_80	User share of package k=1 sold at t $\in$ [61,80]	0.19	0.29
12	q_k2_j_10_20	User share of package k=2 sold at t $\in$ [10,20]	0.29	0.17
13	q_k2_j_21_40	User share of package k=2 sold at t $\in$ [21,40]	0.23	0.16
14	q_k2_j_41_60	User share of package k=2 sold at t $\in$ [41,60]	0.21	0.15
15	q_k2_j_61_80	User share of package k=2 sold at t $\in$ [61,80]	0.17	0.07
16	q_k3_j_10_20	User share of package k=3 sold at t $\in$ [10,20]	0.15	0.06
17	q_k3_j_21_40	User share of package k=3 sold at t $\in$ [21,40]	0.09	0.04
18	q_k3_j_41_60	User share of package k=3 sold at t $\in$ [41,60]	0.21	0.01
19	q_k3_j_61_80	User share of package k=3 sold at t $\in$ [61,80]	0.07	0.00
20	q_k4_j_10_20	User share of package k=4 sold at t $\in$ [10,20]	0.15	0.28
21	q_k4_j_21_40	User share of package k=4 sold at t $\in$ [21,40]	0.02	0.00
22	q_k4_j_41_60	User share of package k=4 sold at t $\in$ [41,60]	0.04	0.00
23	q_k4_j_61_80	User share of package k=4 sold at t $\in$ [61,80]	0.01	0.00
24	unlock_2	Share of unlockings at ep 2	0.94	0.94
25	unlock_3	Share of unlockings at ep 3	0.92	0.93
26	unlock_4_9	Share of unlockings at ep 4-9	0.89	0.91
27	unlock_10_14	Share of unlockings at ep 10-14	0.28	0.30
28	unlock_15_80	Share of unlockings at ep 15-80	0.14	0.15
29	unlock_night	Fraction of unlockings in the night	0.26	0.26

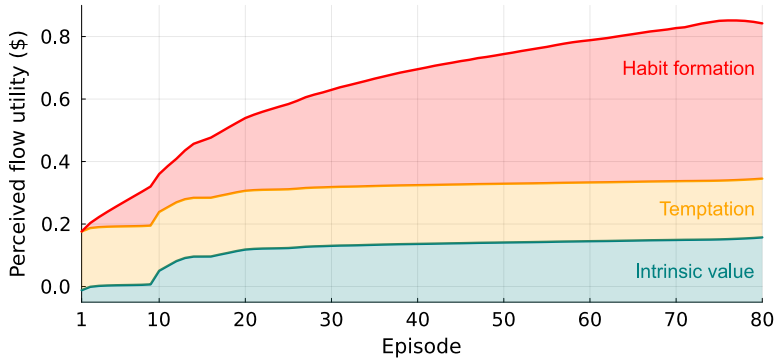


Figure D.9: Decomposition of perceived flow utility by episode

Notes: This figure decomposes the perceived flow utility into three components by episode: intrinsic value, temptation, and habit formation. For each episode, I report the average level of each utility component conditional on unlocking. All the components are normalized in dollars.

Appendix D.5 Robustness: Sophistication

In this section, I discuss the structural estimation under sophistication, where users recognize that their future selves will face the same temptations when making unlocking and top-up decisions as current selves. I first present the model solution for sophisticated users. Although the data do not allow me to identify the degree of sophistication, I estimate the model assuming that 20% of users are sophisticated, showing that the naïveté assumption is conservative and provides a lower bound on the impacts of temptation.

Model solution for sophisticated users. I define the system of value functions recursively for sophisticated users, denoted by a superscript "$*$," building on the naïve solutions presented in section 4.2. Let χ^* and χ denote the actual and perceived temptation durations, respectively: the actual duration reflects the user's true susceptibility and shapes expectations about future behavior, while the perceived duration governs the utility calculation at time t. This distinction is crucial, as sophisticated users correctly anticipate χ^* when reasoning about future decisions but perceive temptation lasting only χ when evaluating utilities for current or future selves, which will be clear when we write the recursive system.

Following the same idea for naïve users, the (expected) value associated with the top-up decision writes:

$$\begin{aligned} V_t^*(h,a;\delta,\psi,\chi,\chi^*,\boldsymbol{b},\boldsymbol{\nu}) = \quad & \mathbb{1}\{a < c_t\} \max_{k\in\{0,1,\ldots,K\}} \{EW_t^*(h,a+b_k;\delta,\psi,\chi,\chi^*) - \omega p_k + \nu_k\} \\ & + \mathbb{1}\{a \geq c_t\} EW_t^*(h,a;\delta,\psi,\chi,\chi^*), \end{aligned} \tag{D.2}$$

and

$$EV_t^*(h,a;\delta,\psi,\chi,\chi^*) = \mathbb{1}\{a<c_t\}\mathbb{E}_{\boldsymbol{b}}\left[\log\left(\sum_{k=0}^{K}\exp\left[EW_t^*(h,a+b_k;\delta,\psi,\chi,\chi^*)-\omega p_k\right]\right)\right] + \mathbb{1}\{a\geq c_t\}EW_t^*(h,a;\delta,\psi,\chi,\chi^*). \tag{D.3}$$

The continuation value is adapted as:

$$\begin{aligned}\mathcal{C}_t^*(h,a;\delta,\psi,\chi,\chi^*) = {} & (1-\rho)\,EV_{t+1}(h+1,a-c_t;\delta,\psi,\max\{\chi-1,0\},\chi^*) \\ & +\rho\mathbb{E}_{\psi'}\left[EV_{t+1}(0,a-c_t;\delta,\psi',\max\{\chi-1,0\},\chi^*)\right],\end{aligned} \tag{D.4}$$

where the user perceives themselves with the actual length of temptation duration χ^* when making decisions in $t+1$, but calculate the utility with a shorter duration $\max\{\chi-1,0\}$ as the naïve user. Using this continuation value, we can define the value functions regarding unlocking decisions:

$$W_t^*(h,a;\delta,\psi,\chi,\chi^*,\boldsymbol{\epsilon}) = \mathbb{1}\{a\geq c_t\}\max\left\{\underbrace{u_t(h,\delta,\psi,\chi,\epsilon)}_{\text{flow}}+\underbrace{\mathcal{C}_t^*(h,a;\delta,\psi,\chi,\chi^*)}_{\text{continuation}},\epsilon_0\right\}\mathbb{1}\{a<c_t\}\epsilon_0, \tag{D.5}$$

and

$$EW_t^*(h,a;\delta,\psi,\chi,\chi^*) = \mathbb{1}\{a\geq c_t\}\beta_\epsilon\log\left(1+\exp\left(\frac{u_t(h,\delta,\psi,\chi,\epsilon)+\mathcal{C}_t^*(h,a;\delta,\psi,\chi,\chi^*)}{\beta_\epsilon}\right)\right), \tag{D.6}$$

where the flow utility depends on the perceived temptation duration χ.

Model estimation: When 20% of users are sophisticated. To assess how the naïveté assumption might bias the analysis, I estimate the model assuming that 20% of users are sophisticated. Solving the full system of sophisticated value functions (D.3), (D.4), and (D.6) is computationally challenging because the additional state variable greatly expands the state space. I therefore approximate sophisticated behavior by that of rational users with $\chi = 0$. In Appendix A.2, I show that sophisticated users' behavior converges to rational behavior when T is large, as in early-stage decisions (e.g., the first top-up after free episodes). Moreover, sophisticated users purchase tokens only once, since they can correctly forecast future usage: they stop if they do not intend to continue, or purchase the largest package if they anticipate succumbing to temptation and finishing the series. This one-time decision makes subsequent top-ups less salient, supporting the suitability of the approximation.

Table D.3 reports the counterfactual estimation results. Assuming 20% of users are sophisticated, most parameter estimates remain close to the baseline. The main difference lies in the temptation distribution, with per-minute temptation valued at 20.6¢ (11.4% higher) and an average duration of 10.2 minutes (54.5% higher). These findings suggest that the naïveté assumption yields conservative estimates of temptation and its duration, providing a lower bound on the associated efficiency loss.

Table D.3: Counterfactual estimation: When 20% of users are sophisticated

Parameter	Baseline		Sophistication		Meaning
	Estimator	Std	Estimator	Std	
1. Intrinsic					
μ_δ	-0.0097	0.0021	-0.0045	0.00045	Mean of drama value
σ_δ	0.092	0.0023	0.082	0.0017	SD of drama value
β_ϵ	0.055	0.0077	0.058	0.0062	Magnitude of unlocking taste shifter
$\overline{\psi}$	0.022	0.0045	0.016	0.00063	Outside value during nighttime
2. Temptation					
κ	0.15	0.0030	0.18	0.0041	Magnitude of naive temptation
μ_χ	1.24	0.022	1.69	0.030	Mean of temptation duration (log)
σ_χ	1.25	0.033	1.66	0.045	SD of temptation duration (log)
3. Habit					
α_1	-2.8e-5	4.3e-7	-3.0e-5	3.6e-7	Habit utility $\alpha(h) = \alpha_1 h(h - \alpha_2)$
α_2	471.94	2.9e-5	451.97	9.3e-6	Habit utility $\alpha(h) = \alpha_1 h(h - \alpha_2)$
ρ	0.039	0.00049	0.039	0.0011	Prob of entering next round
4. Price					
ω	0.81	0.0071	0.86	0.0033	Price sensitivity

Notes: The GMM robust standard errors are reported in the table.