Dissecting Netflix's Self-Preferencing: Evidence from Viewer-Level Data*

Tin Cheuk Leung† Wake Forest University

Shi Qi‡ William & Mary

Koleman Strumpf§ Wake Forest University

November 2025

Abstract

Self-preferencing by dominant digital platforms has become a focal point for antitrust scrutiny, yet little empirical work has examined this behavior in the context of video streaming. This paper provides the first systematic analysis of self-preferencing on a subscription-based streaming platform, focusing on Netflix. We assemble a novel dataset that combines a weekly panel of Netflix's U.S. catalog from 2016 to 2025, official Top 10 rankings since 2021, Wikipedia page views as an external proxy for popularity, and device-level streaming data from tens of millions of U.S. smart TVs. We begin by showing that the exit of licensed series significantly increases the likelihood of subscriber churn, whereas the effect of movie exits is small and even slightly negative. This underscores the risks of dependence on non-original serialized content and motivates Netflix's incentives to promote Originals. We then document that Netflix Originals are substantially more likely to appear in the Top 10 rankings than non-originals, conditional on popularity and availability. The magnitude of this self-preferencing effect is comparable to the influence of popularity itself, especially for serialized content. Finally, using a difference-in-differences design with matched titles, we show that Top 10 inclusion has a significant causal impact on subsequent viewer engagement, with stronger effects for Originals. Taken together, our findings suggest that Netflix leverages interface prominence to steer attention toward its proprietary content while insulating itself from the risks associated with expiring licenses, raising important implications for content competition and platform governance in the streaming era.

JEL Classification: D22, K21, L40, L82, M21

Keywords: Self-Preferencing, Netflix, Digital Platforms, Platform Bias

*We are grateful to seminar participants at Bates College, CREED and the Mallen Conference for helpful comments. We thank Andy Shen for his research assistance. We thank the NET Institute `www.NETinst.org` for financial support. The usual disclaimer applies.

†Contact: Department of Economics, Kirby Hall, Wake Forest University, Box 7505, Winston-Salem, NC 27109. Email: leungtc@wfu.edu

‡Email: sqi01@wm.edu

§Email: strumpks@wfu.edu

1 Introduction

In recent years, self-preferencing by dominant digital platforms has become a central concern for antitrust authorities on both sides of the Atlantic. In the United States, the proposed American Innovation and Choice Online Act (AICOA) would make it unlawful for large online platforms to favor their own products or services in ways that materially harm competition (U.S. Congress, 2021). Similarly, the European Union's Digital Markets Act, implemented in 2022, prohibits gatekeepers from engaging in "any form of differentiated or preferential treatment in ranking on the core platform service... in favor of products or services it offers itself" (Council of the European Union, 2022). While a growing body of academic research and media scrutiny has examined self-preferencing in digital marketplaces such as Amazon, Google's web search, Apple's App Store, and Spotify, little is known about such behavior in the context of video streaming. This paper addresses that gap by providing the first systematic analysis of self-preferencing on Netflix, the world's largest subscription-based streaming platform.

To investigate whether (and why) Netflix systematically promotes its own original content, we compile a new dataset that combines three key components. First, we construct a weekly panel of the Netflix U.S. catalog from 2016 to 2025, tracking when titles enter and exit the platform and classifying each as a Netflix Original or not. Second, we collect weekly Top 10 rankings from Netflix's official Tudum website starting mid-2021, which provide direct evidence of interface prominence. Third, we incorporate large-scale device-level streaming data collected by a third-party firm, which monitors tens of millions of smart TVs in the U.S. using digital fingerprinting technology. This dataset records the exact titles, timestamps, and durations of viewing sessions across Netflix and other platforms, enabling a detailed analysis of consumer behavior at the device level. It allows us not only to validate proxies for popularity, but also to estimate engagement and reach directly, and will play a larger role in future empirical extensions of this work.

We complement this with Wikipedia page view data for each title as an externally observable proxy for latent viewer interest. Prior literature has shown that Wikipedia traffic reflects public engagement and often precedes actual viewership. We match approximately 18,000 titles in our Netflix catalog to their corresponding Wikipedia pages

and aggregate daily page views to the weekly level to align with our catalog and Top 10 data.

Our empirical strategy proceeds in three parts. First, we examine whether the exit of licensed titles affects the likelihood of subscriber churn, using device-level data combined with Wikipedia pageviews as a proxy for the popularity of departing content. This provides direct evidence on the costs to Netflix of relying on non-original shows that may be withdrawn. Second, we estimate whether Netflix Originals are more likely to appear in the weekly Top 10 rankings than non-originals, conditional on Wikipedia popularity and other observables. This "conditioning on observables" approach detects potential self-preferencing by identifying asymmetries in promotion. Finally, we assess the causal effect of Top 10 inclusion on subsequent viewer interest using a difference-in-differences design with matched title pairs, where matching is based on predicted Top 10 propensity scores and semantic similarity of content.

Our analysis begins by examining the consequences of content exits for subscriber churn. We find that the departure of licensed series significantly increases the unsubscribe rate, while the effect of movie exits is small and, if anything, negative. Quantitatively, a 10% increase in Wikipedia views of exiting series raises the likelihood of unsubscribing in the subsequent three weeks by about 0.4%, compared to an average churn rate of 3–4%. By contrast, the effect of movie exits is an order of magnitude smaller and in the opposite direction. These results highlight the importance of serialized content in retaining subscribers and underscore why Netflix may wish to reduce its dependence on non-original series.

Our findings also reveal consistent evidence of self-preferencing. Conditional on external popularity, Netflix Originals are significantly more likely to be featured in the Top 10—particularly for series—but the magnitude of this bias has declined over time. We also document a positive and sizable causal impact of Top 10 inclusion on subsequent viewership, using both Wikipedia traffic and title-level views data from device-level streaming records. The effect is notably larger for Originals. These results suggest that Netflix's promotional interface plays an active role in shaping viewership, and that preferential visibility may amplify the reach and durability of its proprietary content. Our paper offers

the first empirical documentation of self-preferencing on a subscription-based streaming platform, contributing to ongoing debates about digital gatekeeping and platform neutrality. While our current analysis uses aggregated title-level streaming data, in future work we plan to exploit the rich individual-level streaming data to more directly estimate the effect of Top 10 inclusion on viewership and to model viewer demand based on historical viewing patterns and user-level preference heterogeneity.

A substantial literature has emerged to examine the economic consequences of self-preferencing in hybrid marketplaces like Amazon, where the platform competes with third-party sellers while also monetizing both through commissions and direct product sales. This dual role creates well-defined channels through which self-preferencing might harm competition—for example, by raising rivals' costs, deterring entry, or exploiting behavioral biases of myopic consumers (Hagiu and Wright, 2015; Zhu and Liu, 2018; Etro, 2023). These concerns are amplified by the transaction-based nature of e-commerce platforms, where the platform's profitability is directly linked to each individual sale. The resulting theoretical and empirical studies provide nuanced insights into when self-preferencing might reduce consumer welfare or merely reflect efficient vertical integration.

However, many of these arguments do not transfer cleanly to the streaming context, where users pay a flat subscription fee and face no marginal cost for consuming individual titles. Still, important analogies remain. Netflix may self-preference its Originals to reduce reliance on licensed content and enhance subscriber retention, as these titles are exclusive to its platform. Its control over interface elements—such as the "Top 10" list or homepage placement—can serve as a form of algorithmic steering, similar in function to Amazon's Buy Box. And while there are no per-view fees or seller commissions, third-party content providers often depend on Netflix for distribution and visibility; preferential treatment of Originals could subtly weaken their incentives to invest in high-quality content, raising concerns about long-run impacts on market competition and innovation.

The remainder of the paper is as follows. Section 2 discuss how our work connects to and advances several active research areas, and Section 3 provides the industry background in the video streaming industry and the business model of Netflix. We describe the data in Section 4. In Section 5, we analyze whether exiting titles affect subscriptions. We then

discuss our test on whether Netflix engages in self-preferencing in Section 6, and quantify the causal impact of Top 10 inclusion in Section 7. The last section concludes.

2 Literature

This paper relates to several strands of literature on digital platform bias. A growing number of empirical studies test for self-preferencing, particularly on Amazon. These include evidence of bias in "frequently bought together" recommendations and the buy box (Chen and Tsai, 2024; Raval, 2023), as well as direct measurements of bias in search rankings on Amazon's platform (Farronato, Fradkin, and MacKay, 2023; Jürgensmeier and Skiera, Forthcoming; Reimers and Waldfogel, 2023). Other studies quantify the welfare implications of self-preferencing on Amazon (Gutiérrez, 2022; Lam, 2021; Lee and Musolff, 2023). Beyond Amazon, several papers examine search bias in other digital ecosystems, such as Apple's App Store (Teng, 2023), online travel agencies (Hunold, Kesler, and Laitenberger, 2020), meta-search platforms (Cure et al., 2022), and Spotify (Aguiar, Waldfogel, and Waldfogel, 2021). While most existing work focuses on retail and app platforms, we are the first to study self-preferencing in a video streaming context.

Second, our paper is related to the debate on vertical integration and self-preferencing. Before becoming the chair of the FTC, Lina Khan argues that Amazon has exploited its market power and its preference towards Fulfilled by Amazon or "FBA" is anti-competitive (Khan, 2017), and discusses how structural separation can remedy competitive harms (Khan, 2019). Etro (2024) summarizes the literature. Several theoretical papers explore the reasons behind why platforms might bias the rankings of products (Armstrong and Zhou, 2011; Hagiu and Jullien, 2014; Parker, Petropoulos, and Van Alstyne, 2020; de Corniére and Taylor, 2019; Bourreau and Gaudin, 2022). Anderson and Bedre-Defolie (2022) analyze the condition under which a platform would sell its own products. Jürgensmeier, Bischoff, and Skiera (2024) argues that self-preferencing in online marketplaces Amazon primarily self-preference through retailing and fulfillment, revealing vertical integration as a key driver of platform bias. Zou and Zhou (2023) analyze the effects of self-preferencing on outcomes.

Our paper contributes to the growing literature on the economics of digital platforms.

A substantial body of research has examined e-commerce marketplaces such as Amazon (Farronato, Fradkin, and MacKay, 2023; Gutiérrez, 2022; Lam, 2021; Lee and Musolff, 2023) and eBay (Einav et al., 2018), sharing economy platforms like Uber (Lam, Liu, and Hui, 2021) and Airbnb (Zervas, Proserpio, and Byers, 2017), and music streaming services such as Spotify (Aguiar and Martens, 2016; Aguiar, Waldfogel, and Waldfogel, 2021). In contrast, empirical research on video streaming platforms—especially Netflix—remains limited, largely due to a lack of publicly accessible viewing data. Among the few existing studies, Aguiar and Waldfogel (2018) show that while Netflix distributes titles from a wide range of countries, U.S.-origin content continues to dominate in terms of geographic reach. McKenzie et al. (2019) use stated-preference data to estimate demand for Netflix, while Danaher and Doosti (2025) find that the launch of Netflix in Australia led to a 52% drop in sales for the average film licensed to the platform and Frick, Fletcher, and Smith (2023) examine the impact of Netflix availability on piracy intent using Google search data, leveraging the removal of Epix movies from Netflix as a natural experiment. Yuan (2025) analyzes exclusive contracts in the video streaming market, including Netflix, using Nielsen demand data. However, that study relies on aggregated viewership data rather than device-level records, and does not include a full catalog of Netflix content as we do. Our paper is the first to construct a weekly panel of the Netflix catalog across multiple countries and develop a scalable approach to measure the popularity of Netflix titles using Wikipedia page views and device-level streaming data, providing a new foundation for empirical research on streaming platforms.

3 Industry Background

Netflix is the dominant player in the subscription-based video streaming market. With over 260 million global subscribers as of early 2025, the company has maintained a leading position through its vast content library, international expansion, and continued investment in proprietary technology. Unlike ad-supported platforms or pay-per-view services, Netflix operates under a flat-rate subscription model, where revenue is largely decoupled from the consumption of any specific title. This business model provides strong incentives to retain subscribers over the long term by curating a highly engaging and differentiated

content offering.

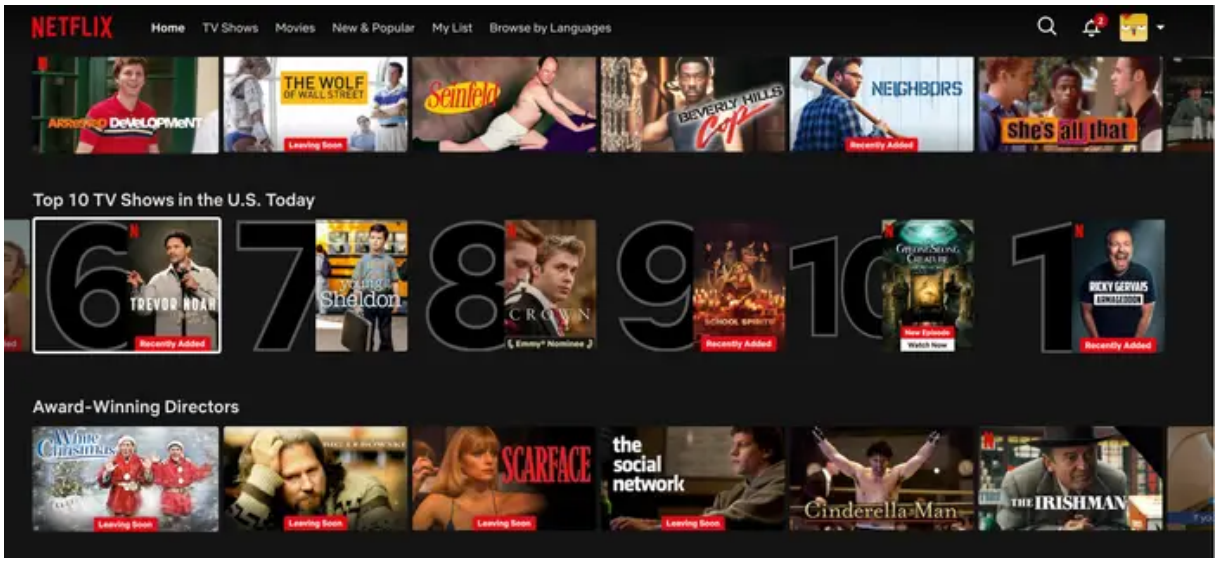
A key feature of Netflix's user experience is its algorithmic content recommendation system, which relies heavily on individual viewing data. Recommendations appear in various forms: a "Continue Watching" row that prompts users to resume unfinished titles; personalized rows such as "Because You Watched The Crown"; and curated selections that highlight genres, trending themes, or titles that are popular among viewers with similar tastes. These recommendations are dynamically ranked and displayed based on each subscriber's viewing history, preferences, and engagement patterns, and Netflix itself has emphasized that the accuracy of these predictions improves with the volume of user data it can draw upon.

Not all promotional placements, however, are personalized. In 2021, Netflix introduced a globally visible "Top 10" ranking feature, which highlights the most-watched movies and TV shows in each country on a weekly basis. Unlike the tailored rows described above, this list is identical across users within the same region and is prominently featured on the home screen. Figure 1 illustrates an example of the "Top 10 TV Shows in the U.S. Today" as displayed on the Netflix interface. Notably, many of the titles featured—such as The Crown and Ricky Gervais: Armageddon—are Netflix Originals, while others are third-party licensed content. The introduction of this uniform ranking raises the possibility that its influence may be strongest among new or less active subscribers, while for heavy users—whose personal data allows for finely tuned recommendations—the incremental effect of the Top 10 may be less pronounced.

Over the past decade, Netflix has strategically shifted toward producing its own content, reducing its reliance on licensed titles. Beginning with *House of Cards* in 2013, Netflix has steadily increased its investment in Originals across both movies and series. High-profile examples include *Stranger Things*, *The Witcher*, *Bridgerton*, *Extraction*, and *The Irishman*. This strategy offers multiple advantages: Originals are globally licensed, can be promoted without external negotiation, and serve as unique draws to retain and attract subscribers. Simultaneously, Netflix has reduced its catalog of non-original titles, many of which are now tied up in competing platforms' exclusive agreements.

Crucially, Netflix retains full control over the algorithms that govern both personalized

Figure 1: Top 10 TV Shows on Netflix Interface



recommendations and the visibility of the Top 10 rankings. This allows the platform to steer user attention toward its original content through a variety of means—ranging from homepage placement and category prominence to editorial decisions on what appears in the Top 10. While such steering can reflect underlying consumer demand or efficient vertical integration, it also raises concerns about biased visibility and reduced opportunities for third-party content, particularly in a market where discoverability is tightly mediated by platform-controlled interfaces.

4 Data

4.1 Netflix's Weekly Catalog

To analyze Netflix's content promotion strategy, we construct a comprehensive weekly panel dataset documenting the U.S. catalog of Netflix titles from 2016 through early 2025. This dataset tracks when titles enter and exit the platform, as well as their key characteristics, and serves as the backbone for our empirical analysis.

Our construction process begins with the New on Netflix website (`https://usa.newonnetflix.info/`), which offers daily records of new arrivals to Netflix's U.S. catalog. These records are used to identify when titles are first made available on the

platform. To capture when titles are removed, we rely on data from the uNoGS API (`https://unogs.com/`), which provides accurate exit dates along with detailed title-level metadata. This metadata includes genre, country of origin, and primary language, all of which are useful for understanding content patterns. We also determine whether a title is designated as a "Netflix Original" based on its listing status on the New on Netflix site.

Starting from a full catalog snapshot in March 2025, we reconstruct historical availability by iteratively moving backward week-by-week. In each step, titles that have not yet entered are excluded, while those with expiration dates beyond the current week remain in the catalog. This recursive process allows us to trace the evolution of the Netflix U.S. catalog on a weekly basis going back to 2016.

For each title-week observation, we collect data on whether the title is a Netflix Original, whether it is scheduled to leave the platform within the next four weeks, and its geographic availability in Canada and the U.K. These variables allow us to control for international licensing patterns and content lifecycle effects. We also flag titles produced in India to capture a segment of regionally targeted or niche content that may be treated differently in Netflix's promotion algorithm.

In addition, we incorporate data from the Movie of the Night API (`https://www.movieofthenight.com/`) beginning in 2024 to determine whether a title is concurrently available on other major subscription streaming platforms, including Amazon Prime Video, Hulu, Disney+, HBO Max, and Apple TV+. While this availability data is not available prior to 2024, it enables us to assess the competitive positioning of titles in the most recent period.

Through this process, we generate a clean, weekly dataset of the Netflix U.S. catalog spanning nearly a decade. This catalog construction enables us to trace changes in platform offerings over time, observe patterns in content entry and exit, and prepare the groundwork for analyzing whether Netflix differentially promotes its own original titles.

4.2 Netflix Official Weekly Top Ten Lists

To evaluate whether Netflix systematically promotes its own original content, we collect data on Netflix's official weekly Top Ten rankings, which serve as a direct measure of the

platform's promotional choices. These lists identify the most-watched titles on Netflix in the United States and are published separately for movies and television series.

Our primary source for this data is Netflix's Tudum website (`https://www.netflix.com/tudum/top10`), which began publishing weekly Top Ten rankings starting from the week of June 28 to July 4, 2021. Each weekly update includes two ranked lists: one for the Top 10 movies and one for the Top 10 series in the U.S. market. The rankings are based on views during the corresponding week and reflect viewing behavior from Monday through Sunday.[1] Importantly, these rankings represent Netflix's own internal measure of content success and are prominently displayed on the platform's interface, often influencing viewer choices.

We download these rankings for each week since their inception. For each entry, we record the title, its rank (1 through 10), and whether it appeared in the movie or series list. We also collect the official synopsis provided by Netflix, which will later be used to construct semantic similarity metrics for matching analyses. To align the Top 10 data with our broader weekly catalog, we convert the weekly Top 10 lists into binary indicators at the title-week level, flagging whether a title appeared in the Top 10 in a given week. Titles that appear in the Top 10 list are matched to our weekly Netflix catalog using unique Netflix IDs.

These Top 10 indicators serve two purposes in our analysis. First, they act as the dependent variable in our regression models testing for self-preferencing—specifically, whether Netflix Originals are more likely to appear on these lists after controlling for external measures of popularity. Second, they define treatment status in our difference-in-differences analysis estimating the causal impact of Top 10 placement on subsequent viewer interest. Because Netflix selects and publishes these rankings itself, they provide a transparent window into the platform's algorithmic or editorial promotion of content, and hence offer a meaningful way to detect potential biases in visibility.

[1] In June 2023, Netflix revised its definition of "views." Previously, titles were ranked based on total viewing time. Under the new metric, views are calculated by dividing the total time spent watching a series or movie in a given week by its runtime, yielding an estimate of the number of Netflix accounts that watched the title (Porter, 2023).

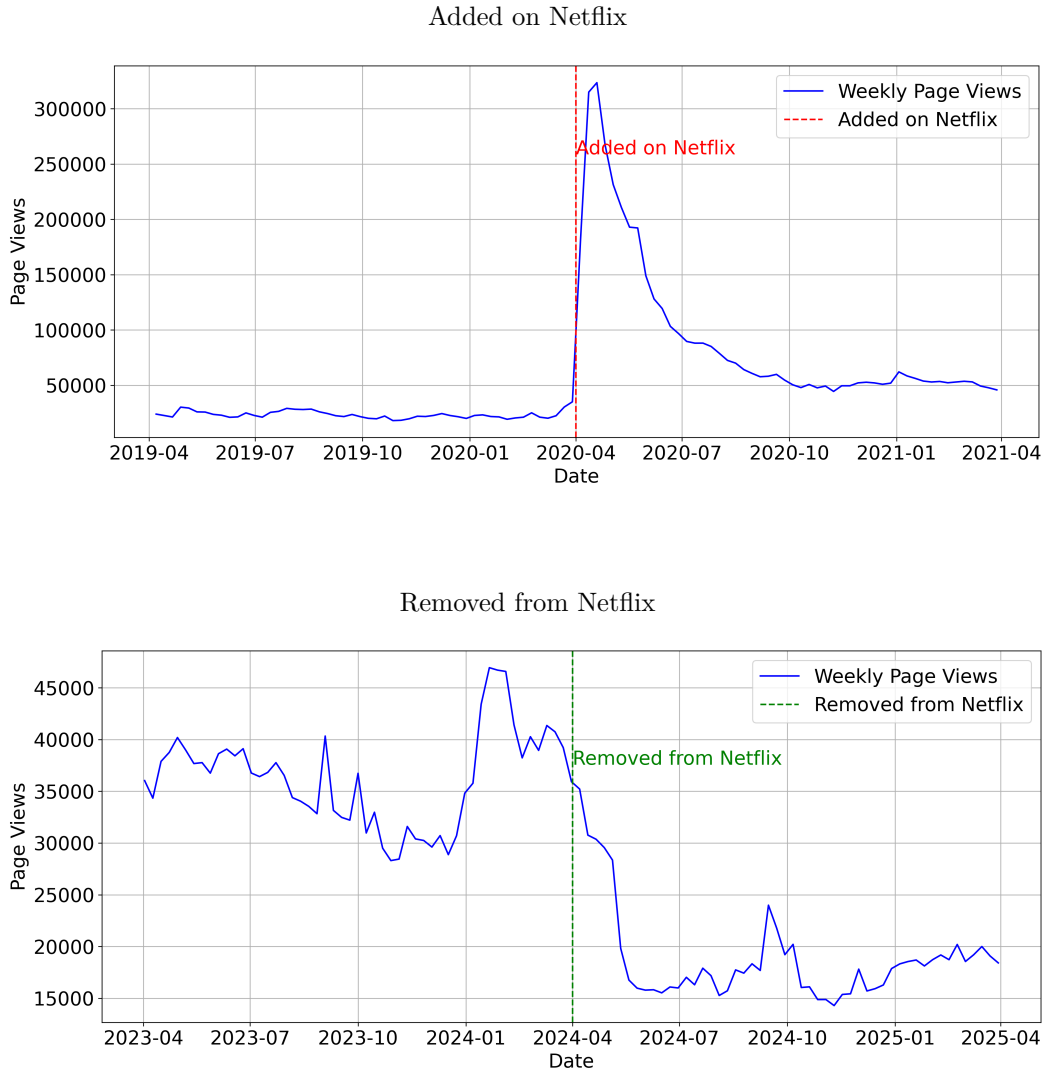
4.3 Wikipedia Page Views

To proxy for latent viewer interest independent of Netflix's own promotional choices, we collect Wikipedia page view data for each title in our Netflix catalog. Wikipedia is a widely used source of information for movies and TV series, and its traffic patterns often mirror public curiosity or engagement triggered by trailers, reviews, word of mouth, or general online discussion. Mestyán, Yasseri, and Kertész (2013) demonstrates that Wikipedia activity, including page edits and views, can serve as an early predictor of box office performance. Leung, Qi, and Strumpf (2025) shows that releases of movie trailers are correlated with a surge in views and edits in the Wikipedia pages of the movies. Importantly, unlike viewership figures controlled by Netflix, Wikipedia traffic is generated externally and is publicly observable, making it a valuable tool for gauging popularity outside the platform.

The responsiveness of Wikipedia traffic to streaming availability can be illustrated using the example of Community (TV series), which was added to Netflix on April 1, 2020 and removed exactly four years later on April 1, 2024 (the show originally aired on NBC between 2009 and 2014). Figure 2 plots the weekly Wikipedia page views of the show one year before and after the date the show was added on and removed from Netflix. When Community was added to Netflix in the week of April 1, 2020, its Wikipedia views spiked from fewer than 50,000 to over 300,000—an immediate sixfold increase in public interest. Conversely, after the show was removed from Netflix in April 2024, weekly Wikipedia views dropped from nearly 40,000 to roughly half that level. These patterns support the use of Wikipedia page views as a timely and externally valid proxy for consumer engagement and visibility, especially around key moments of platform availability.

We identify each title's corresponding English-language Wikipedia page by conducting automated Google searches using the title name in conjunction with keywords such as "film," "TV series," or "Wikipedia." This matching process is validated manually to ensure accuracy, particularly for titles with common names, foreign-language entries, or remakes. Out of roughly 22,000 unique titles in our weekly catalog, we successfully match approximately 18,000 titles to a distinct Wikipedia page. We continue to work on completing the remaining matches.

Figure 2: Wikipedia Weekly Page Views: *Community* (TV series)



For each matched title, we retrieve daily Wikipedia page view counts using the Wikimedia REST API. The API returns the number of views for each day, which we aggregate to the weekly level to correspond with the timing of our Netflix Top 10 and catalog data. The result is a panel dataset of weekly Wikipedia traffic for each title during the weeks it is available on Netflix U.S.

4.4 Device-Level Streaming and Linear TV Viewership Data

To complement our analysis of externally observable popularity measures, we incorporate device-level streaming and linear TV viewership data. This dataset enables us to observe

actual viewer behavior at the device level and to validate the effectiveness of Netflix's Top 10 lists in driving engagement. It also provides an opportunity to estimate causal effects of platform visibility on consumption outcomes with greater precision than is possible using aggregated or proxy measures alone.

The data come from a third-party which collects viewing data from dozens of smart TV brands, and includes tens of millions such devices in the US between 2021 and 2024. The third-party then uses a digital fingerprinting technology to identify each specific content that is viewed, and each data record also includes a timestamp as well as various metadata. For each device, the data record the exact titles watched, the duration of each viewing session, and the timing of those sessions. The dataset includes Netflix as well as other streaming platforms, though only Netflix Original titles can be reliably identified by name due to metadata restrictions imposed by Netflix on third-party access.

Our initial validation exercise focuses on a representative week in March 2024, during which the device-level panel captures viewing activity from over 400,000 U.S. devices that accessed any streaming service. For each matched Netflix Original title, we compute two key measures of viewer engagement: the total number of seconds watched across all devices and the total number of distinct viewing events. We then merge these measures with Wikipedia traffic data and find a strong positive correlation between Wikipedia page views and both viewing intensity and reach, particularly for series. This supports the use of Wikipedia views as a reasonable proxy for underlying popularity in weeks when direct viewership data are unavailable.

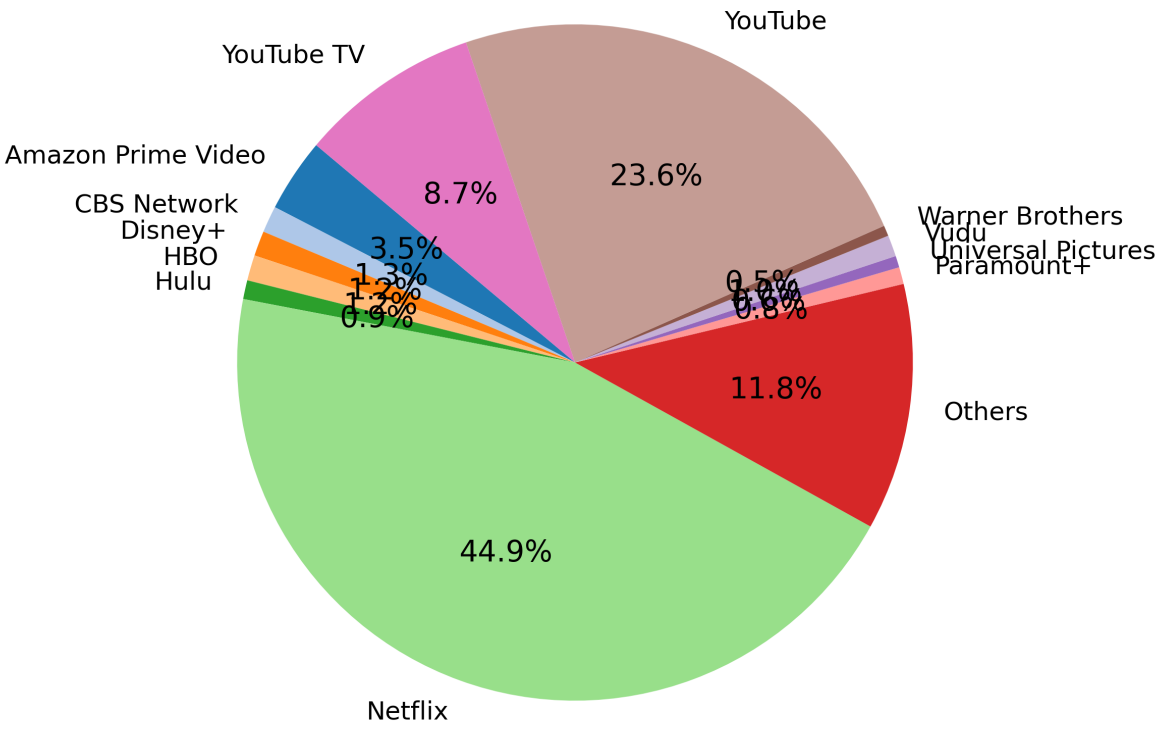
4.5 Summary Statistics

4.5.1 Market Shares of OTT Streaming Services

Figure 3 presents the market shares of major streaming platforms between September 2021 and October 2024, using device-level viewership data from Samba TV. We calculate market shares based on the total duration watched (in seconds) for each platform, aggregating viewership at the network level. This approach offers a more behaviorally grounded measure of market activity than subscriber counts or app installations.

According to the data, Netflix dominates the streaming landscape with a time-watched

Figure 3: Streaming Market Share by Network Sep 2021–Oct 2024 (Total Duration)



market share of 44.9%, followed by YouTube (23.6%) and YouTube TV (8.7%). Other notable platforms include Amazon Prime Video (3.5%), CBS Network (1.3%), Disney+ (1.3%), HBO (1.2%), Hulu (0.9%) and Warner Brothers (0.5%).

While Netflix's position as the market leader is not surprising, the extent of its dominance—especially in terms of total duration watched—is noteworthy. Its viewership share is an order of magnitude larger than that of other major platforms such as Amazon Prime Video, HBO, and Hulu. This is particularly striking given that subscriber-based estimates, such as those from FlixPatrol, show a much narrower gap:[2] Netflix reportedly has about 81 million U.S. subscribers, compared to 75 million for Amazon Prime, 54 million for Hulu, and 50 million for HBO. The disconnect between subscriber counts and actual viewing time suggests that relying on subscription figures to infer market power or user engagement may be misleading. Our results underscore the value of behavioral data in

[2]See `https://flixpatrol.com/streaming-services/subscribers/united-states/`.

accurately assessing platform performance and competitive dynamics in the streaming ecosystem.

Figure 4 presents the average duration watched per device per week on Netflix and across all streaming platforms combined, offering a more detailed view of user engagement over time. Rather than showing market shares alone, the figure highlights absolute viewing levels at the device level. On average, each device in the sample spent around 3.5 to 4 hours per week watching content on any platform. For Netflix specifically, weekly per-device viewing time declined from about 2 hours in late 2021 and early 2022 to roughly 1 hour by the end of 2024.

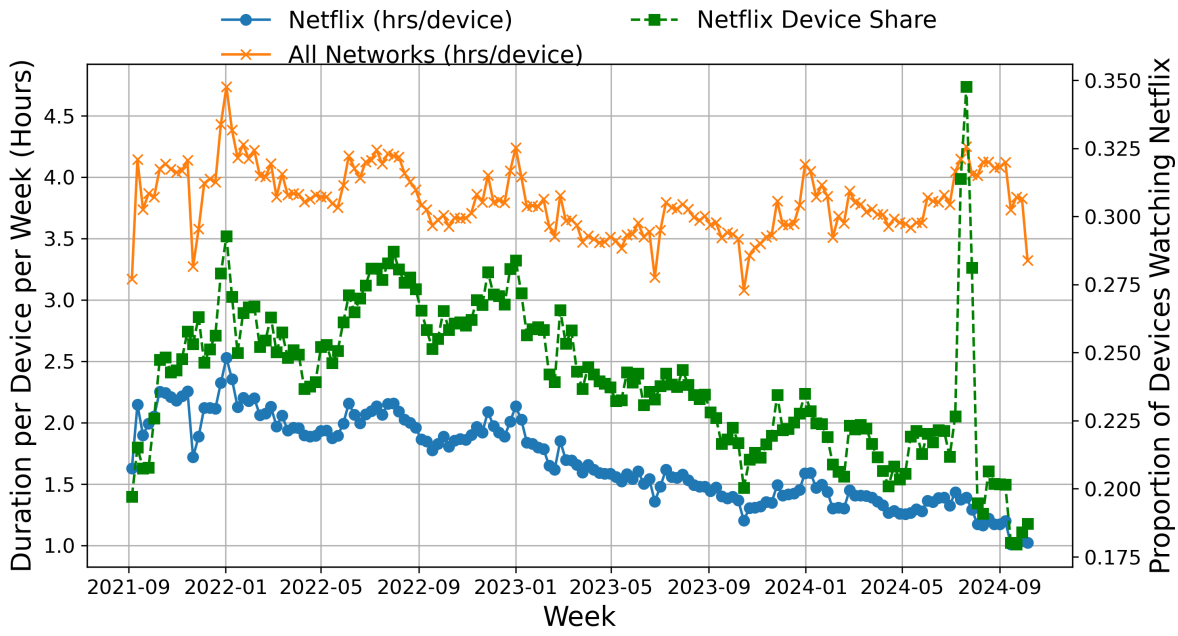
In other words, Netflix's share of total viewing time was around 50%—peaking at approximately 55%—throughout 2021 and much of 2022, before declining steadily to about 30% by late 2024. Importantly, the decline in Netflix engagement appears to be driven more by the extensive margin—the proportion of devices watching Netflix—rather than the intensive margin, or the amount of time each user spends once subscribed. As shown in the figure, the proportion of devices engaging with Netflix content hovered between 25% and 35% prior to 2023, but began to fall in 2023, dropping to around 15% by the end of 2024. This pattern suggests that Netflix's erosion in market share is largely attributable to a shrinking user base, rather than a drop in engagement among existing users.

This trend coincides with the rise of competing platforms such as Amazon Prime Video, Disney+, and HBO's Peacock, which have gained traction through aggressive content investments and platform expansion. The decline in both relative market share and absolute per-device engagement suggests that Netflix's dominance in the streaming ecosystem has waned in the face of increasing competition.

4.5.2 Netflix Original Content Over Time

Figure 5 illustrates the evolution of Netflix's content composition and its promotional behavior by plotting the share of Netflix Original titles in both the overall catalog and the weekly Top 10 lists for movies and series from 2016 through 2025. For movies, the share of originals in the catalog rose steadily from near zero in 2016 to around 50% by

Figure 4: Netflix's Market Share Over Time (Total Duration)



2025. For series, the growth has been even more pronounced, with the share of originals exceeding 70% by the end of the sample period.

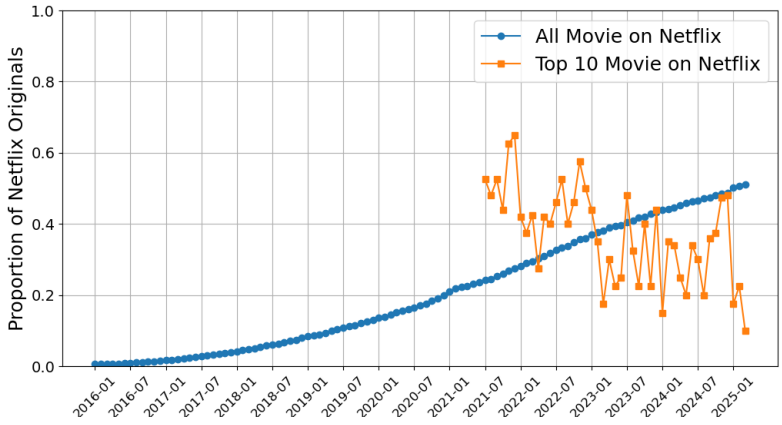
When we turn to the Top 10 lists—available starting in mid-2021—the share of originals is consistently higher than in the catalog, especially for series. More than 80% of series featured in the Top 10 each week are Netflix Originals, compared to roughly 70% in the catalog. For movies, the gap between catalog share and Top 10 share was substantial until 2023, when the share of originals in the Top 10 dropped below their catalog share. These trends are suggestive of potential preferential treatment of Netflix Originals, particularly for serialized content. However, they are not conclusive on their own, as the elevated presence of originals in the Top 10 may also reflect differences in investment, quality, or viewer preferences. Subsequent regression analyses will help determine whether these patterns persist after controlling for popularity and other observables.

4.5.3 Validation of Wikipedia Page Views

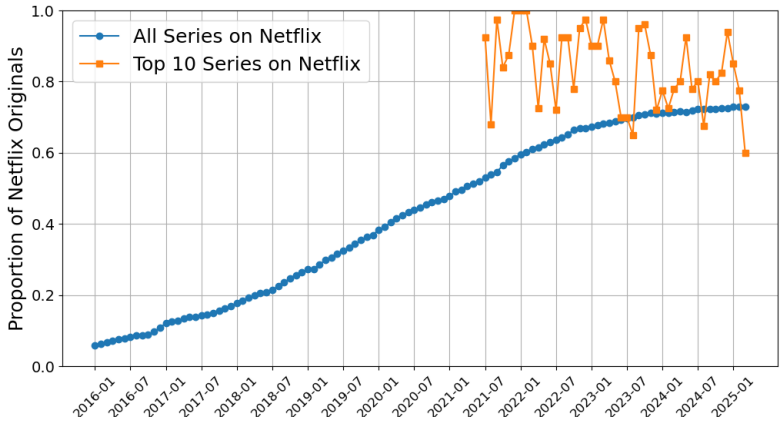
To validate our use of Wikipedia page views as a proxy for title popularity and viewer interest, we leverage the device-level streaming and linear TV viewership data for the weeks

Figure 5: Trends in the Share of Netflix Original Titles in the Overall Catalog and Top 10 Lists, 2016–2025

Movies



Series



between September 2021 and October 2024. During this period, the device-level dataset captures viewing activity from approximately 440,000 unique U.S. devices watching any Netflix program in any given week.

We aggregate the device-level data at the title level, focusing on Netflix original content. For each title, we compute three measures of viewer engagement: (1) the total duration watched (in seconds), (2) the total number of distinct view events, and (3) the number of Netflix views, defined as "total time spent watching a movie or season of a TV series, divided by the running time" (Porter, 2023).[3] Among the 559,990 Netflix original

[3]Note that the number of TV series with this measure drops by nearly half, primarily because runtime

title-week observations available on the platform between September 2021 and October 2024, we successfully match 286,586 observations in the device-level dataset.

Figure 6 illustrates the relationship between our Wikipedia proxy and the three device-level measures. The top left panel plots the log of Wikipedia page views against the log of total watch duration; the top right panel plots against the log of watch count; the low panel plots against the log of Netflix views. All panels reveal strong positive correlations, supporting the validity of Wikipedia traffic as a measure of actual viewer interest.

Figure 6: Relationship Between Wikipedia Page Views and Device-Level Viewership Metrics for Netflix Originals (Sep 2021–Oct 2024)

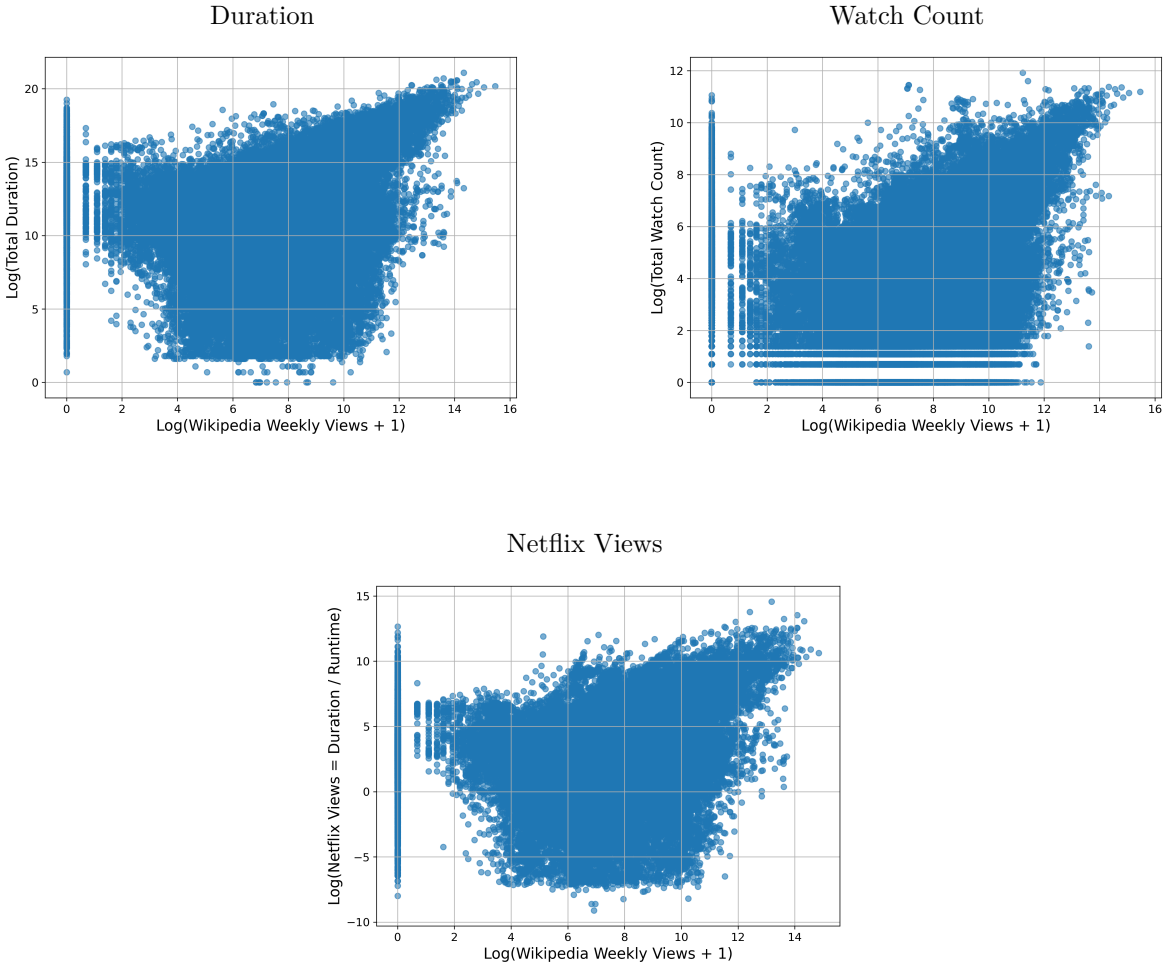


Table 1 presents regression results where we regress log device-level viewership measures (log total duration watched, log watch count, and log Netflix views) on the log of

information is missing for about 50% of series in the uNoGS API. We plan to address this limitation by supplementing the missing data.

Wikipedia page views for the matched titles. We include controls for whether the titles are available in the U.S. and/or Canada and whether they were produced in India to account for cross-regional availability. The first three columns present estimates for movies, and the last three for series.

Across all specifications, we find that Wikipedia page views are positively and statistically significantly correlated with all device-level measures. The strength of this relationship is notably stronger for series, with coefficients between 0.34 and 0.57 (i.e. a one percent increase in a title's Wikipedia page views is correlated with a 0.34-0.57 percent increase in the device-level measures), compared to 0.29 to 0.34 for movies. These results support the validity of Wikipedia traffic as a proxy for actual viewer engagement and suggest that Wikipedia views may better capture interest in serialized or episodic content.

Table 1: OLS Regression of Device-Level Measures on Wikipedia Views

	Movie			Series		
	log(Duration)	log(Watch Count)	log(Netflix View)	log(Duration)	log(Watch Count)	log(Netflix View)
log(Wikipedia Pageviews)	0.335***	0.290***	0.330***	0.455***	0.340***	0.565***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.004)
Constant	9.131	2.033	0.440	8.741	1.609	-0.648
	(0.049)	(0.035)	(0.047)	(0.052)	(0.039)	(0.059)
Observations	115014	115014	114855	139252	139252	74030
R^2	0.085	0.113	0.093	0.145	0.145	0.215

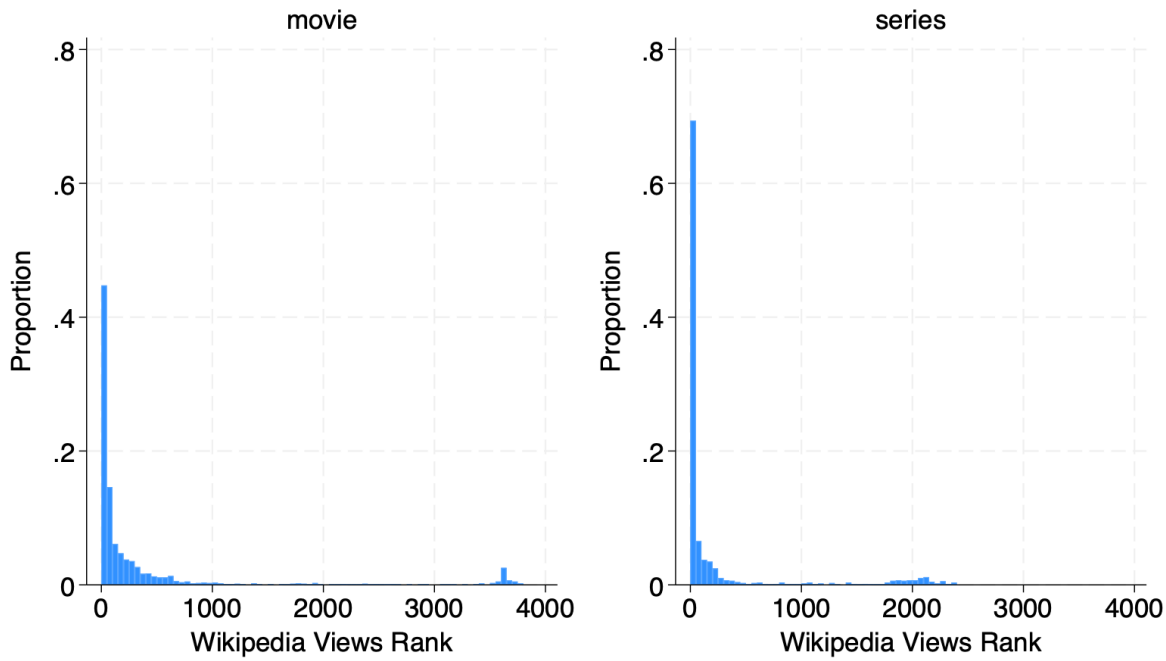
Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

To validate Wikipedia page views as a proxy for title popularity, we examine how titles that make the Top 10 rank in terms of Wikipedia traffic. Figure 7 presents the histogram of weekly Wikipedia rank for titles that appear in Netflix's official Top 10 list, for both movies and series. Titles are ranked each week based on their total Wikipedia page views within the movie or series category, and grouped into bins of 50 (e.g., ranks 1–50, 51–100, etc.).

The majority of top 10 titles are highly ranked in Wikipedia page views. Nearly 50% of top 10 movies fall within the top 50 by Wikipedia traffic, and close to 70% of top 10 series do. This strong alignment supports the use of Wikipedia page views as a reasonable and informative proxy for actual viewing behavior on the platform.

Figure 7: Distribution of Wikipedia Page View Ranks Among Top 10 Titles



4.5.4 Top 10 vs Non-Top 10

We summarize key differences between Top 10 and non-Top 10 titles using two complementary sources of data. Table 2 presents summary statistics based on Wikipedia page views, while Table 3 reports analogous statistics derived from device-level streaming data. Both tables provide separate breakdowns for movies and series.

In Table 2, we observe that Top 10 titles receive substantially more Wikipedia page views than non-Top 10 titles. On average, Top 10 movies garner around 110,000 views during the relevant week, compared to just 6,500 for non-Top 10 movies. The contrast is even more pronounced for series, with Top 10 entries averaging 210,000 views versus 5,800 for non-Top 10 titles. This pattern holds when we examine Wikipedia views as a share of total weekly traffic.

Other differences include the composition of content: while around 40% of movies in both groups are Netflix Originals, the share of Originals is much higher among Top 10 series (86%) than non-Top 10 series (71%). Geographic availability also varies, with Top 10 movies being slightly less available in the U.K. and Canada than their non-Top 10 counterparts. Finally, Indian titles—particularly movies—are underrepresented in the

Table 2: Summary Statistics of Netflix Catalog (Wikipedia)

	Movie		Series	
	Non-Top 10	Top 10	Non-Top 10	Top 10
Weekly Wikipedia Pageviews (10,000)	0.6589 (2.6843)	11.0642 (18.6053)	0.5885 (1.7187)	21.4052 (37.6050)
Weekly Wikipedia Pageviews Share (%)	0.0167 (0.0628)	0.280 (0.454)	0.0151 (0.0441)	0.547 (0.926)
Netflix Original	0.410 (0.492)	0.381 (0.486)	0.706 (0.455)	0.863 (0.344)
Available (UK)	0.753 (0.431)	0.537 (0.499)	0.892 (0.311)	0.898 (0.303)
Available (Canada)	0.814 (0.389)	0.655 (0.475)	0.908 (0.289)	0.921 (0.269)
Indian Movie	0.128 (0.334)	0.0222 (0.147)	0.0219 (0.146)	0.00411 (0.0640)
N	857801	1864	475362	1568

mean coefficients; sd in parentheses

Top 10: they constitute 12% of all movies but just 2

Table 3 provides a parallel summary using device-level streaming data. We report two engagement metrics: (1) average duration watched per device (in seconds), and (2) average number of Netflix Views per device, where one Netflix View is defined as total duration divided by the runtime of the title.[4] The patterns mirror those seen in the Wikipedia data—Top 10 titles consistently outperform non-Top 10 titles—but the magnitudes are even more striking. Both duration watched and Netflix Views for Top 10 movies and series are approximately 100 times greater than for non-Top 10 titles, compared to a roughly 50-fold difference in the Wikipedia-based metrics. This highlights the stronger contrast captured in behavioral data and reinforces the importance of using direct viewership measures in evaluating platform impact.

While we do not present these figures in a formal table, data obtained from the *Movie of the Night* API reveal that in 2024, only a small proportion of non-Netflix original titles were also available on the subscription plans of other major streaming platforms.[5] For instance, fewer than 1% of these titles were simultaneously available on Apple TV+ or Disney+, and only 1–2% appeared on Hulu, Paramount+, Peacock, or Amazon Prime

[4]Note that the number of observations in Table 3 is smaller because the device-level data covers a shorter time window and includes only Netflix Originals, due to metadata limitations.

[5]Almost none of the Netflix original titles were available in other platform's subscription plans.

Video. The proportion is slightly higher for series than for movies. Overall, the overlap in availability suggests that only a tiny fraction of titles on Netflix were concurrently accessible via other platforms' subscription services during 2024.

Table 3: Summary Statistics of Netflix Catalog (Device-Level Streaming Data)

	Movie		Series	
	Non-Top 10	Top 10	Non-Top 10	Top 10
Duration (sec per device)	0.2490	34.3908	0.6029	59.1587
	(1.1217)	(30.1420)	(1.9467)	(65.5033)
Netflix Views (per device)	0.0000	0.0052	0.0004	0.0285
	(0.0002)	(0.0042)	(0.0024)	(0.0619)
Available (UK)	0.9600	0.9413	0.9805	0.9877
	(0.1959)	(0.2353)	(0.1381)	(0.1104)
Available (Canada)	0.9589	0.9545	0.9764	0.9910
	(0.1986)	(0.2085)	(0.1517)	(0.0943)
Indian Title	0.0368	0.0208	0.0195	0.0011
	(0.1882)	(0.1430)	(0.1382)	(0.0335)
N	129546	528	155619	893

mean coefficients; sd in parentheses

5 Do Exiting Titles Affect Subscriptions?

A central motivation for Netflix to engage in self-preferencing—prominently recommending its own original content—stems from the risks associated with reliance on licensed, non-original shows. Unlike originals, which Netflix controls, licensed content is vulnerable to contract expiration and the growing tendency of rival streaming platforms to reclaim their own titles. When such shows leave, Netflix faces the possibility of heightened subscriber churn. For example, analysts estimated that the departure of The Office in 2021 alone could cost Netflix nearly $935 million in annual subscriptions (Shuman, 2021). This illustrates the strategic importance of retaining viewers through content that cannot be withdrawn by competitors.

By investing in and preferentially promoting original programming, Netflix can reduce its exposure to these risks. Originals ensure long-term catalog stability, and promoting them through the recommendation algorithm further entrenches user habits around content that Netflix fully controls. In this sense, self-preferencing is not only about favoring

in-house productions but also about safeguarding the subscription base against the volatility of licensing arrangements.

To examine the impact of such dynamics on subscriber behavior, we leverage device-level streaming data in conjunction with Netflix's catalog and Wikipedia pageview statistics. Specifically, we estimate the following regression model:

$$\begin{aligned} \text{Unsubscribe}_{it} \;=\; & \sum_{k=0}^{K} \beta_{1k}^{movies} \ln(\text{Views}_{t-k}^{movies}) + \sum_{k=0}^{K} \beta_{1k}^{series} \ln(\text{Views}_{t-k}^{series}) \\ & + \beta_2 \ln(X_{it}^{duration}) + \beta_3 \ln(X_{it}^{week}) + \phi_i + \epsilon_{it}, \end{aligned} \tag{1}$$

where the dependent variable Unsubscribe_{it} is an indicator for whether device i unsubscribes in week t, defined as a cessation of Netflix viewing for at least ten consecutive weeks. The key regressors, $\text{Views}_{t-k}^{movies}$ and $\text{Views}_{t-k}^{series}$, capture cumulative Wikipedia views of movies and series exiting Netflix in week $t-k$, which serve as proxies for the popularity of departing non-original titles. Since the device-level data do not report the specific non-original titles viewed, Wikipedia pageviews offer an external measure of audience interest in departing content. Lagged values are included to capture delayed responses, as cancellations may occur weeks after the exit event. Control variables include cumulative watch duration ($X_{it}^{duration}$) and the number of active weeks (X_{it}^{week}) for each device, with device fixed effects (ϕ_i) absorbing time-invariant heterogeneity.

This framework allows us to assess whether the loss of popular non-original titles, as proxied by Wikipedia activity, significantly increases the likelihood of unsubscribing. A strong positive relationship would provide evidence that promoting original content—through algorithmic self-preferencing—serves as a strategic hedge against churn driven by catalog instability.

Table 4 reports the estimates of the relationship between content exits and subscriber churn. The first specification includes only the cumulative Wikipedia views of movies and series exiting Netflix in the current week, while the second adds lagged exits, and the final specification further incorporates device-level characteristics such as cumulative viewing duration and the number of active weeks.

The results yield two key findings. First, the effect of movie exits on the unsubscribe rate is negative, which runs counter to the expectation that the loss of popular films

Table 4: Exiting Titles on Unsubscribe Rates

	(1) Unsubscribe	(2) Unsubscribe	(3) Unsubscribe
$\ln(\text{Views}_t^{movie})$	-0.000167***	0.000186***	-0.000154***
	(0.0000162)	(0.0000179)	(0.0000179)
$\ln(\text{Views}_t^{series})$	0.000812***	0.000957***	0.00111***
	(0.0000172)	(0.0000174)	(0.0000175)
$\ln(\text{Views}_{t-1}^{movie})$		0.000280***	-0.0000431**
		(0.0000172)	(0.0000172)
$\ln(\text{Views}_{t-2}^{movie})$		0.000159***	-0.000119***
		(0.0000175)	(0.0000175)
$\ln(\text{Views}_{t-3}^{movie})$		-0.000256***	-0.000566***
		(0.0000172)	(0.0000172)
$\ln(\text{Views}_{t-1}^{series})$		0.000892***	0.000878***
		(0.0000183)	(0.0000183)
$\ln(\text{Views}_{t-2}^{series})$		0.000913***	0.000869***
		(0.0000178)	(0.0000179)
$\ln(\text{Views}_{t-3}^{series})$		0.000987***	0.000952***
		(0.0000178)	(0.0000179)
ln(Duration)			-0.0186***
			(0.0000631)
# of Weeks Active			0.000665***
			(0.00000170)
Constant	0.0368***	0.0287***	0.245***
	(0.0000622)	(0.000189)	(0.000789)
Observations	61611833	61611833	61611833
Device FE	Yes	Yes	Yes
# Device	1757660.00	1757660.00	1757660.00
R^2	0.12	0.12	0.13

Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

would increase cancellations. The magnitude of this effect is modest, however. In the full specification, a 10% increase in the cumulative Wikipedia views of exiting movies reduces the unsubscribe rate over the next three weeks by less than 0.1%, compared to an average churn rate of 4% in the sample. Second, the effect of series exits is positive and quantitatively larger. In the full specification, a 10% increase in the cumulative Wikipedia views of exiting series raises the unsubscribe rate over the next three weeks by roughly 0.4%. These results highlight the greater importance of serialized content in retaining subscribers, consistent with the idea that audiences are more attached to ongoing series than to films.

The controls behave as expected. Conditional on the number of active weeks, devices that accumulate more viewing hours are less likely to unsubscribe, while, conditional

on cumulative viewing duration, devices with longer tenure on Netflix are more likely to churn. Together, the findings suggest that promoting original series through self-preferencing may be particularly effective in mitigating subscriber loss when non-original content exits the platform.

6 Testing for Netflix Self-Preferencing

6.1 OLS for Conditioning on Observables (COO) Approach

In this part of the paper, we will examine whether Netflix original titles are more likely to be featured in the weekly Top 10 list than non-original titles, conditional on their popularity and other observables using the conditioning on observables (COO) approach defined in Reimers and Waldfogel (2023). To do this, we estimate a regression of the following form:

$$\begin{aligned} TopTen_{it} &= \beta_0 + \beta_1 X_{it}^{wiki} + \beta_2 X_{it}^{wiki} \times X_i^{original} + \beta_3 X_{it}^{wiki} \times X_{it}^{Expiring} \\ &\quad + \sum_{c \in C} \beta_{4c} X_{it}^{wiki} \times X_{ict}^{Available} + \beta_5 X_{it}^{wiki} \times X_i^{India} + \phi_t + \epsilon_{it}, \end{aligned} \tag{2}$$

where $TopTen_{it}$ is a binary indicator for whether title i makes the Top 10 in week t, X_{it}^{wiki} is the natural log of Wikipedia page views the title receives that week, and $X_i^{original}$ indicates whether the title is a Netflix original. The interaction term $X_{it}^{wiki} \times X_i^{original}$ captures whether equally popular Netflix originals are more likely to enter the Top 10, providing a direct test of self-preferencing. We also interact page views with whether the title is expiring in the next four weeks, availability indicators for Canada and the UK, and with an indicator for Indian titles, in order to account for streaming availability, regional availability and cultural targeting. Week fixed effects ϕ_t will be included to absorb weekly variation in the overall visibility or competitiveness of the platform.

6.2 Results

Table 5 presents the regression results. The first three columns report estimates for movies, while the last three columns correspond to series. Columns 2 and 5 incorporate

an interaction term between Wikipedia page views and whether the title is set to expire within four weeks. Columns 3 and 6 further include controls for geographic availability and whether the title is an Indian production.

The results indicate that Wikipedia page views are a strong predictor of Top 10 appearance. Titles with higher Wikipedia traffic are significantly more likely to be featured in the Top 10. On average, a 10% increase in Wikipedia page views is associated with a 1.5–1.7% higher likelihood for movies and a 2.3–2.4% higher likelihood for series to enter the Top 10. This relationship remains robust after controlling for whether a title is set to expire within four weeks and becomes even stronger when accounting for international availability and whether the title is an Indian production.

Most importantly, we find evidence consistent with self-preferencing: the interaction between Wikipedia page views and Netflix Original status is positive and statistically significant in both the movie and series regressions. This indicates that, conditional on popularity (as proxied by Wikipedia page views), Netflix Originals are more likely to appear in the Top 10 than comparable non-original titles. Under the full specification (columns 3 and 6), a 10% increase in Wikipedia page views is associated with an additional 0.15% and 0.75% higher likelihood of appearing in the Top 10 for Netflix Original movies and series, respectively—representing roughly a 10% and 30% increase over their baseline effects. This suggests that self-preferencing exists for both movies and series, but Netflix may be pushing its own series content more aggressively, consistent with its broader strategic shift toward expanding original series production.[6]

We also find that the interaction between Wikipedia page views and whether a title is set to expire within four weeks is negative and statistically significant. This suggests that Netflix may be systematically deprioritizing or "punishing" content that is about to be removed from the platform, even if those titles are receiving substantial viewer interest. Such a pattern could reflect strategic choices in how Netflix promotes content with limited

[6]We conduct four robustness checks to validate this result: (1) using actual views rather than log of views, (2) restricting the sample to the top 300 titles by weekly Wikipedia page views in each week, (3) using the share of Wikipedia page views among all titles in that week as an alternative measures, and (4) adding controls for whether titles are available on other streaming platforms. The last check is limited to weeks in 2024, the period for which we can trace platform availability. Appendix A shows that all robustness checks yield qualitatively similar results.

remaining licensing windows.[7]

Table 5: Self-Preferencing Regression Results

	(1) Movie	(2) Movie	(3) Movie	(4) Series	(5) Series	(6) Series
ln(Weekly Wikipedia Pageviews)	0.00154***	0.00152***	0.00174***	0.00242***	0.00237***	0.00244
	(0.0000268)	(0.0000272)	(0.0000296)	(0.0000476)	(0.0000484)	(0.00005
ln(Wikipedia Pageviews)XNetflix Original	0.0000536***	0.0000746***	0.000153***	0.000691***	0.000692***	0.000750
	(0.0000163)	(0.0000167)	(0.0000198)	(0.0000273)	(0.0000282)	(0.00003
ln(Wikipedia Pageviews)XExpiring in 4 Weeks		-0.000210***	-0.000312***		-0.000327**	-0.00032
		(0.0000506)	(0.0000508)		(0.000152)	(0.00015
ln(Wikipedia Pageviews)XAvailable (UK)			-0.000302***			-0.00022
			(0.0000233)			(0.00005
ln(Wikipedia Pageviews)XAvailable (Canada)			-0.0000385			0.00008
			(0.0000245)			(0.00005
ln(Wikipedia Pageviews)XIndian Movie			-0.000352***			-0.00062
			(0.0000252)			(0.00008
Constant	-0.00875***	-0.00866***	-0.00830***	-0.0175***	-0.0170***	-0.0170
	(0.000201)	(0.000203)	(0.000204)	(0.000344)	(0.000348)	(0.00034
Observations	711329	683316	683316	399718	380803	380803
R^2	0.00	0.00	0.01	0.01	0.01	0.01

Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

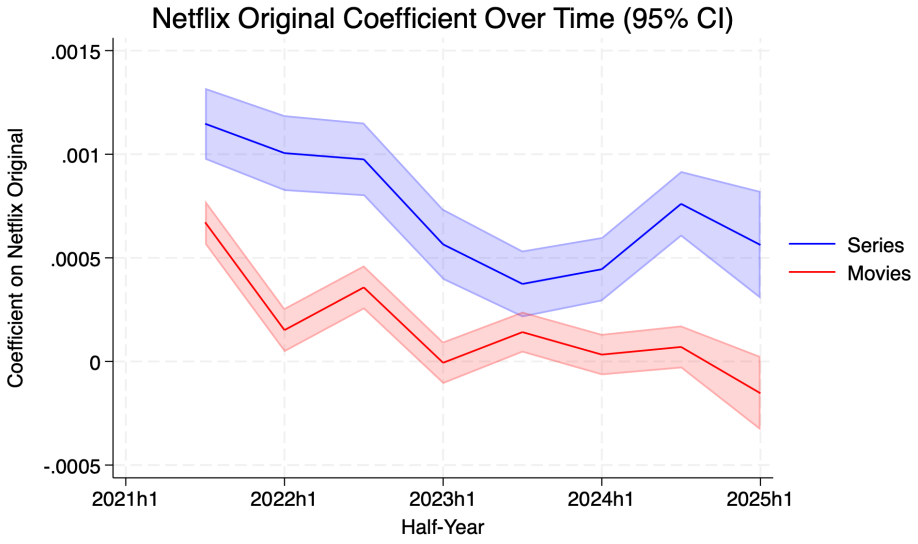
We also investigate how the self-preferencing behavior has evolved over time by estimating the interaction effect separately for each half-year. Figure 8 plots the estimated coefficients of the interaction between Wikipedia page views and the Netflix original indicator, along with 95% confidence intervals, for movies (in red) and series (in blue).

The results show that the self-preferencing effect has been consistently positive and statistically significant across all years for both movies and series, with one important exception. Starting in 2023, the self-preferencing effect for movies disappears, while the effect for series remains positive and statistically significant. Moreover, there appears to be a structural break in 2023, during which the magnitude of the self-preferencing effect declines. This timing coincides with Netflix's introduction of a new metric for measuring viewership. In June 2023, Netflix shifted from reporting total hours viewed to a normalized measure: total time spent watching a title divided by its runtime. Netflix explained that this new metric is "anchored in engagement — our best measure of member satisfaction

[7]We also run alternative specifications where the interaction term is defined based on titles expiring within 1 to 8 weeks instead of just 4 weeks. The results, shown in Appendix B, remain qualitatively similar.

and a key driver of retention (which in turn drives our business)," and that it "levels the playing field" for titles with shorter runtimes (Porter, 2023). This change may have altered Netflix's internal promotion strategies, leading to a more balanced representation of titles in the Top 10—especially reducing the relative advantage previously enjoyed by Netflix Original movies. Notably, this period also coincides with the beginning of a decline in Netflix's overall market share, as shown in Figure 4, suggesting broader competitive pressures that may have prompted further adjustments in its promotional and content strategies.

Figure 8: Estimated Self-Preferencing Effects Over Time: Interaction Between Wikipedia Views and Netflix Original Status (with 95% Confidence Intervals)



7 Causal Impact of Top 10 Inclusion on Viewership

We begin by assessing the relationship between Netflix Top 10 placement and subsequent viewer engagement using ordinary least squares (OLS) regressions. The following regression specification is estimated:

$$\text{Views}_{i,t+1} = \beta_1 \text{Views}_{i,t} + \beta_2 \text{Top10}_{i,t} + \mathbf{X}'_{i,t}\gamma + \alpha_i + \lambda_t + \epsilon_{i,t}, \tag{3}$$

where $\text{Views}_{i,t+1}$ is the number of views for title i in week $t+1$, $\text{Views}_{i,t}$ is the number of views in the current week, and $\text{Top10}_{i,t}$ is an indicator variable equals to 1 if the title appeared in Netflix's US Top 10 list during week t. The vector $\mathbf{X}_{i,t}$ includes controls for whether the title was also available in the UK or Canada in that week. We include both title fixed effects (α_i) and week fixed effects (λ_t) to control for persistent differences across titles and common shocks over time.

Table 6 presents results using Wikipedia page views as the dependent variable, while Table 7 uses device-level streaming data to construct behavioral outcomes, including total duration watched and a normalized measure of Netflix views (duration divided by runtime). In both tables, we find that current views are strongly predictive of next-week views, with coefficients on $\text{Views}_{i,t}$ around 0.7 across all specifications.

Table 6: OLS of Top 10 Inclusion (Wikipedia)

	(1) Movie Non-Original	(2) Movie Original	(3) Movie Overall	(4) Series Non-Original	(5) Series Original	(6) Series Overall
Wikipedia Pageviews$_t$	0.650***	0.609***	0.636***	0.780***	0.714***	0.721***
	(0.00114)	(0.00152)	(0.000904)	(0.00186)	(0.00140)	(0.00115)
Top 10	0.238***	1.883***	0.685***	0.781***	1.955***	1.685***
	(0.0610)	(0.0801)	(0.0481)	(0.0652)	(0.0682)	(0.0541)
Observations	414378	289429	705319	115992	280631	398810
Title FE	Yes	Yes	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
R^2	0.57	0.55	0.56	0.80	0.64	0.65

Standard errors in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

The estimates on the Top 10 variable are also significant and economically meaningful. In Table 6, Top 10 inclusion is associated with an increase of about 7,000 Wikipedia page views for movies—roughly a 6.4% increase relative to the average of 110,000 among Top 10 movies—and 17,000 page views for series, or about 7.9% compared to the 214,000 average.

In Table 7, Top 10 inclusion increases total viewing duration for movies by about 1,600 hours (relative to 1,900 hours for Top 10 movies) and Netflix views by 900 (relative to 10,000). For series, Top 10 inclusion leads to 1,800 additional viewing hours (relative to 3,300) and an increase of 14,000 Netflix views (relative to 57,000), corresponding to an effect size of about 25%.

Table 7: OLS of Top 10 Inclusion (Device-Level Streaming Data)

	(1) Movie Duration (100 Hrs)	(2) Movie Netflix Views (1000)	(3) Series Duration (100 Hrs)	(4) Series Netflix Views (1000)
Duration/Views$_t$	0.605***	0.598***	0.740***	0.576***
	(0.00186)	(0.00189)	(0.00187)	(0.00292)
Top 10	16.13***	0.924***	17.99***	14.26***
	(0.495)	(0.0267)	(0.949)	(0.433)
Observations	124718	124561	151970	75231
Title FE	Yes	Yes	Yes	Yes
Week FE	Yes	Yes	Yes	Yes
R^2	0.70	0.70	0.69	0.47

Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

While informative, these OLS results remain correlational. Titles selected into the Top 10 may differ systematically from others in unobservable ways. To address this, we next employ a propensity score matching approach to construct better counterfactual comparisons between treated (Top 10) and untreated titles.

To estimate the causal impact of appearing in Netflix's Top 10 list on subsequent viewer interest, we employ a propensity score matching approach to compare viewer behavior across two groups of titles: those that appeared in the Top 10 during the previous week (treated group) and those that did not (control group). We construct matched pairs using a combination of propensity score matching (Rosenbaum and Rubin, 1983) and semantic similarity in content. A title is classified as "treated" if it appeared in Netflix's Top 10 list in a given week.

7.1 Propensity Score Matching and Semantic Similarity

The matching is performed separately within each stratum defined by content type (movie or series), Netflix Original status (original vs. non-original), and calendar week. This ensures that treated and control titles are compared only to others that share the same structural and temporal context.

Within each stratum, we first estimate propensity scores using a logistic regression model, where the dependent variable indicates whether a title appears in the Top 10 list during a given week. The covariates include the title's weekly Wikipedia page views or weekly duration watched/Netflix views from the device-level streaming data (proxies for

underlying popularity), Netflix Original status, a set of indicators on the availability of the title in Canada and the UK, and a set of one-hot encoded indicators for genre and country of production. The fitted probability from this model represents the estimated likelihood that a title appears in the Top 10 based on its observable characteristics. Matching is then performed using nearest-neighbor matching on the propensity scores, restricted to titles within the same content type and calendar week.

In order to assess the quality of the matching, we show in Appendix C that the propensity score distributions of treated and matched control titles align closely in both datasets.

To further refine the match, we impose a semantic similarity constraint on content. We obtain dense vector representations of each title's synopsis using Sentence-BERT (`all-MiniLM-L6-v2`), a pre-trained language model that produces semantically meaningful embeddings (Reimers and Gurevych, 2019). We compute the cosine similarity between embeddings of treated and control titles, and restrict the control pool to only those titles whose synopsis exceeds a fixed similarity threshold (e.g., 0.1).

7.1.1 Difference-in-Difference Regression

From this restricted pool, each treated title is matched to the control title with the closest propensity score, yielding one-to-one matched pairs. For each pair, we compute the difference in Wikipedia page views or weekly duration watched/Netflix views during the treatment week and the following week. These are used in the following difference-in-differences (DiD) specification:

$$\Delta\text{Views}_{i,t+1} = \alpha + \beta\Delta\text{Views}_{i,t} + \varepsilon_i, \tag{4}$$

where $\Delta\text{Views}_{i,t}$ is the difference in views between the treated and control titles in the treatment week, and $\Delta\text{Views}_{i,t+1}$ is the corresponding difference in the following week. The intercept α captures the average treatment effect of Top 10 placement on subsequent Wikipedia interest.

7.1.2 Results

The results, summarized in Table 8, show that Top 10 placement has a positive average treatment effect on subsequent viewer interest. For movies (column 3), inclusion in the Top 10 list increases Wikipedia page views in the following week by approximately 4,800 (marginally insignificant statistically), which represents a 4.3% increase relative to an average of about 110,000 for Top 10 movies overall. The effect is larger for series (column 6), with an average increase of roughly 13,000 Wikipedia views—a 6% increase compared to a Top 10 series average of 214,000.

Table 8: Difference-in-Difference Results of Top 10 Impact (Wikipedia Data)

	(1) Movie Non-Original	(2) Movie Original	(3) Movie Overall	(4) Series Non-Original	(5) Series Original	(6) Series Overall
$\triangle$Wikipedia Pageviews$_{t-1}$	0.661***	0.614***	0.618***	0.971***	0.697***	0.698***
	(0.018)	(0.029)	(0.016)	(0.039)	(0.021)	(0.019)
Constant (α)	0.377***	0.484	0.476	-0.004	1.272	1.251*
	(0.131)	(0.779)	(0.29)	(0.278)	(0.822)	(0.711)
Observations	1046	656	1702	181	1186	1367
R^2	0.56	0.41	0.46	0.77	0.48	0.48

Standard errors in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

Moreover, for both movies and series, the treatment effect is stronger for Netflix original titles than for non-originals. Although the difference in effect size is not statistically significant for movies, this pattern suggests that Top 10 placement may be more effective at boosting viewer interest in original content. Alternatively, it may reflect additional promotional effort by Netflix for its original titles—an influence not directly observable in our data.

Additionally, the estimates on $\Delta\text{Views}_{i,t}$ are significantly positive at around 0.7, indicating that the difference in views at the time of Top 10 placement is highly predictive of the difference in views in the following week. This correlation highlights the persistence of relative viewer interest across time.

In total, we are able to match 1,702 movies (1,046 non-original and 656 original) and 1,251 series (181 non-original and 1,186 original) for this analysis.

To complement these results using a more direct measure of viewer behavior, Table 9 presents analogous estimates based on device-level streaming data. Unlike Wikipedia

page views, these data capture actual consumption, measured as either total duration watched (in 100 hours) or Netflix views (in 1,000s), where Netflix views are defined as total duration divided by runtime, at the title level. Due to data limitations, this analysis is restricted to Netflix Originals only. We successfully match 506 original movies and 834 original series for the duration-based outcome. For the Netflix views outcome, the sample for series drops to 505 and 443, for movies and series respectively, as runtime data is missing for some titles—a limitation we plan to address in the next draft.

Table 9: Difference-in-Difference Results of Top 10 Impact (Device-Level Streaming Data)

	(1) Movie Duration (100 Hrs)	(2) Movie Netflix Views (1000)	(3) Series Duration (100 Hrs)	(4) Series Netflix Views (1000)
$\triangle$Duration/Views$_{t-1}$	0.589***	0.576***	0.746***	0.437***
	(0.031)	(0.031)	(0.021)	(0.049)
Constant (α)	2.067***	1.2***	1.402	18.199***
	(0.671)	(0.358)	(0.942)	(5.349)
Observations	506	505	834	443
R^2	0.42	0.40	0.60	0.15

Standard errors in parentheses
* $p < .1$, ** $p < .05$, *** $p < .01$

The smaller sample size in Table 9 relative to Table 8 is primarily due to the shorter time span covered by the device-level data. Nonetheless, the estimates of $\Delta\text{Views}_{i,t}$ are reassuringly similar across the two tables—both around 0.6-0.7—which supports the robustness of our findings. More importantly, we find that Top 10 inclusion also significantly increases actual viewing behavior. For movies, the Top 10 placement increases total duration watched by 200 hours, representing a 10.5% increase over the treated-group mean of 1,900 hours, and Netflix views by 1,200, or 12% over the mean of 10,000. For series, we observe an increase of 140 hours (marginally insignificant statistically) in duration watched (a 4.2% increase relative to the 3,300-hour mean) and 18,000 in Netflix views (a 31.6% increase relative to the 57,000 mean). These results provide direct evidence that algorithmic prominence not only signals popularity but also meaningfully shifts consumption on the platform.

7.2 The Effect of Top 10 Rankings on Viewing at the Device Level (Very Preliminary)

To further examine whether Netflix's Top 10 rankings causally influence audience behavior, we turn to the device-level streaming data. This dataset allows us to track the week-by-week viewing decisions of individual devices, which provides a granular perspective on how the promotional visibility of the Top 10 feature affects title-level demand.

We estimate the following panel regression model:

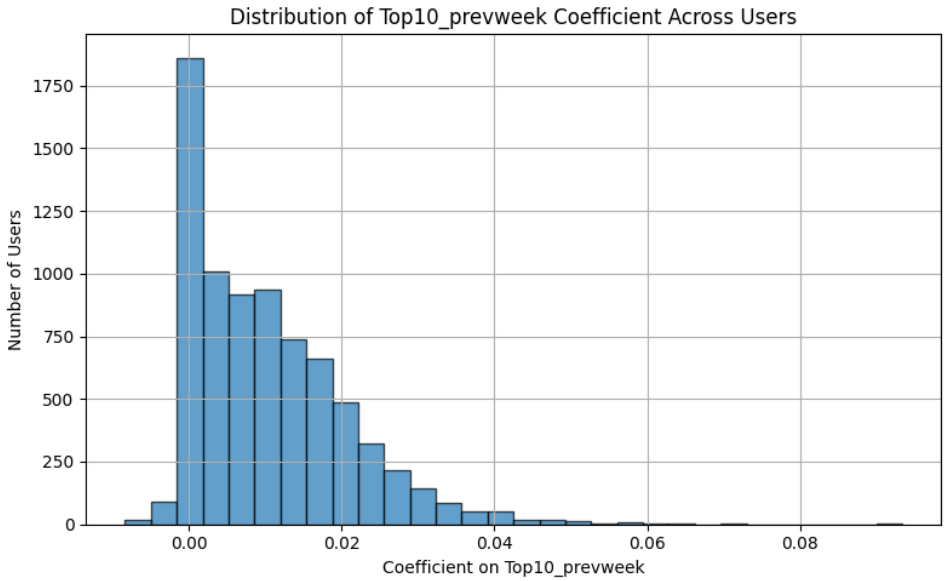
$$\text{Watch}_{i,j,t} = \beta_{i,1}\,\text{Watch}_{i,j,t-1} + \beta_{i,2}\,\text{Top10}_{j,t-1}\beta_{i,2}\,\text{Top10}_{j,t-1} + \mathbf{X}'_{i,j,t}\gamma_i + \alpha_i + \epsilon_{i,j,t}, \quad (5)$$

where the dependent variable $\text{Watch}_{i,j,t}$ is an indicator equal to one if device i viewed title j in week t. The main variable of interest is $\text{Top10}_{j,t-1}$, an indicator for whether title j appeared in the Top 10 list in the previous week. The model also controls for persistence in viewing, captured by the lagged outcome $\text{Watch}_{i,j,t-1}$, as well as a vector of device–title–week controls $\mathbf{X}_{i,j,t}$, which include proxies for prior exposure to related titles or genres. Device fixed effects α_i absorb time-invariant heterogeneity in viewing preferences.

Because each device may exhibit highly idiosyncratic viewing patterns, we estimate the regression separately for each device i. This approach yields device-specific estimates of both the persistence coefficient $\beta_{i,1}$ and the Top 10 effect $\beta_{i,2}$. For computational feasibility, we currently implement this estimation procedure on a 0.5% random sample of the approximately 1.8 million devices that were active on Netflix for at least eight consecutive weeks during the observation window.

The distribution of estimated $\beta_{i,2}$ across devices provides a direct measure of the extent to which Top 10 visibility increases the likelihood of viewing a given title. Preliminary results in Figure 9 show that $\beta_{i,2}$ is positive and statistically significant for a substantial share of devices, indicating that inclusion in the Top 10 list in week $t-1$ meaningfully raises the probability that a device will watch the title in week t. This effect is robust to the inclusion of lagged outcomes and rich controls, suggesting that the Top 10 feature functions as an attention-shifting mechanism rather than merely reflecting underlying

Figure 9: Distribution of Top 10 Impact Across Devices



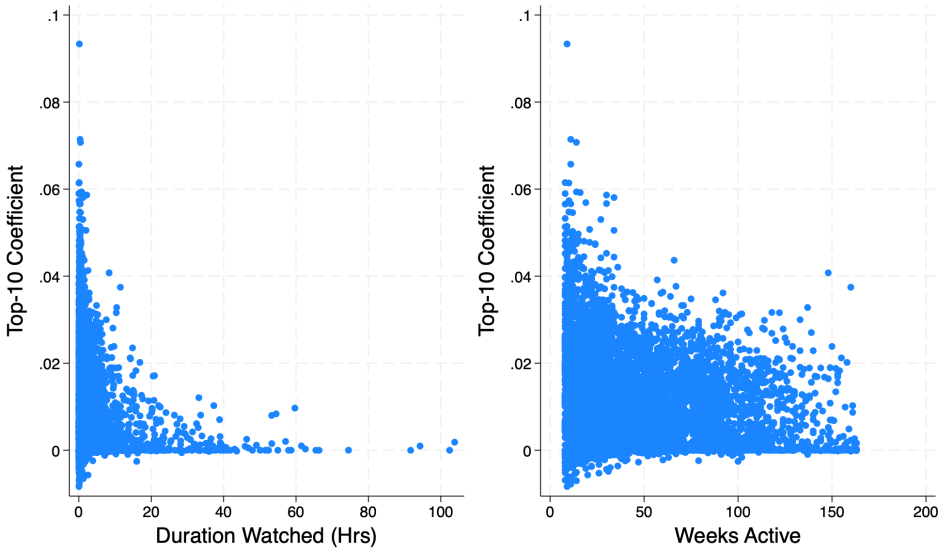
popularity trends.

While the average effect of Top 10 rankings on viewing is positive, there is notable heterogeneity across devices. Table 9 shows the distribution of estimated Top 10 coefficients $\beta_{i,2}$ across devices. Although a significant share of devices exhibit positive and statistically meaningful coefficients, a sizable portion of the distribution—approximately 22% of devices—clusters near zero, indicating little measurable effect of Top 10 rankings on subsequent viewing.

To better understand this heterogeneity, Figure 10 plots the device-level Top 10 coefficients against proxies for each device's overall engagement with Netflix. The left panel uses total viewing duration (in hours) as the proxy, while the right panel uses the number of weeks a device has been active on the platform. Both plots reveal a consistent pattern: devices with shorter watching histories tend to exhibit both higher variance and higher mean values of $\beta_{i,2}$. By contrast, as Netflix accumulates more data on a device's viewing behavior, the estimated Top 10 coefficients converge toward zero.

One possible interpretation is that Netflix leverages extensive viewing histories to fine-tune its personalized recommendation system. For devices with longer and richer viewing records, the platform can rely more heavily on personalized predictions, which may diminish the relative influence of the Top 10 list—a more generic, non-personalized

Figure 10: Top 10 Coefs and Watching History



form of recommendation. Conversely, for devices with limited viewing histories, Netflix has less information to customize recommendations effectively. In these cases, the Top 10 list plays a more prominent role in guiding viewing decisions, resulting in larger and more variable Top 10 coefficients.

This interpretation underscores an important interaction between platform-level promotional visibility and user-level personalization: the same interface element may exert substantially different effects depending on the richness of Netflix's information about a given user. The Top 10 list thus acts as a more influential signal for new or infrequent viewers, while its impact diminishes among highly engaged devices with established consumption patterns.

8 Conclusion

This paper presents the first empirical analysis of self-preferencing behavior on a subscription-based streaming platform. Leveraging a unique dataset that combines Netflix's weekly catalog, Top 10 rankings, Wikipedia page views, and large-scale device-level streaming data, we examine whether Netflix systematically promotes its own original content and whether such promotion shapes viewer behavior.

We begin by showing that the exit of licensed titles has asymmetric effects on sub-

scriber churn: while movie exits have a negligible and even slightly negative impact, the departure of popular series significantly raises the unsubscribe rate. This pattern highlights the risks Netflix faces when relying on non-original serialized content that may be withdrawn due to expiring licenses or competitive pressures.

Against this backdrop, our findings indicate that Netflix Originals are significantly more likely to appear in the platform's Top 10 lists, even after conditioning on external popularity and other observables. This self-preferencing effect is particularly pronounced for series and is comparable in magnitude to the effect of popularity itself. Moreover, using a matched difference-in-differences approach, we find that Top 10 inclusion causally increases subsequent viewer interest, with larger effects for Originals—suggesting that platform-driven visibility not only amplifies the reach of proprietary content but also serves as a strategic hedge against subscriber loss when non-original series exit the platform.

Our contribution is twofold. First, we extend the self-preferencing literature to the streaming domain, where incentives and market structure differ markedly from e-commerce or app platforms. Second, we introduce new measurement tools—such as Wikipedia traffic and device-level viewership logs—that provide a scalable and externally valid approach to studying digital content consumption.

In this draft, we incorporate device-level streaming data to estimate the causal impact of Top 10 inclusion on actual viewership, using title-level measures such as total duration watched and Netflix Views. While these results provide a more direct behavioral validation of platform influence, our current analysis does not yet exploit the full richness of the individual-level data. In the next draft, we plan to leverage the panel structure of the dataset to control for users' viewing histories and preferences, allowing us to estimate more accurate models of demand and viewer responsiveness to algorithmic promotion. These extensions will help us better quantify the effects of self-preferencing on content consumption and competitive dynamics in the streaming market.

References

Aguiar, Luis and Bertin Martens. 2016. "Digital Music Consumption on the Internet: Evidence from Clickstream Data." *Information Economics and Policy* 34:27–43.

Aguiar, Luis and Joel Waldfogel. 2018. "Netflix: Global Hegemon or Facilitator of Frictionless Digital Trade?" *Journal of Cultural Economics* 42 (3):419–445. URL https://doi.org/10.1007/s10824-017-9315-z.

Aguiar, Luis, Joel Waldfogel, and Sarah Waldfogel. 2021. "Playlisting favorites: Measuring platform bias in the music industry." *International Journal of Industrial Organization* 78:102765.

Anderson, Simon and Özlem Bedre-Defolie. 2022. "Hybrid Platform Model: Monopolistic Competition and a Dominant Firm." URL https://ssrn.com/abstract=3867851. SSRN working paper.

Armstrong, Mark and Jidong Zhou. 2011. "PAYING FOR PROMINENCE." *The Economic Journal* 121 (556):F368–F395.

Bourreau, Marc and Germain Gaudin. 2022. "Streaming platform and strategic recommendation bias." *Journal of Economics & Management Strategy* 31 (1):25–47.

Chen, Nan and Hsin-Tien Tsai. 2024. "Steering via algorithmic recommendations." *The RAND Journal of Economics* 55 (4):501–518. URL https://onlinelibrary.wiley.com/doi/abs/10.1111/1756-2171.12481.

Council of the European Union. 2022. "Regulation of the European Parliament and of the Council on Contestable and Fair Markets in the Digital Sector (Digital Markets Act)." https://www.consilium.europa.eu/media/56086/st08722-xx22.pdf. Document ST 8722/22. Accessed: 2025-05-09.

Cure, Morgane, Matthias Hunold, Reinhold Kesler, Ulrich Laitenberger, and Thomas Larrieu. 2022. "Vertical Integration of Platforms and Product Prominence." *Quantitative Marketing and Economics* 20 (4):353–395. URL https://doi.org/10.1007/s11129-022-09255-4.

Danaher, Brett and Shahryar Doosti. 2025. "Netflix and Chilling Effect? The Impact of Streaming Video Licenses on Paid DIgital Downloads." Working paper.

de Corniére, Alexandre and Greg Taylor. 2019. "A model of biased intermediation." *The RAND Journal of Economics* 50 (4):854–882.

Einav, Liran, Chiara Farronato, Jonathan D. Levin, and Neel Sundaresan. 2018. "Auctions Versus Posted Prices in Online Markets." *Journal of Political Economy* 126 (1):178–215.

Etro, Federico. 2023. "Hybrid Marketplaces with Free Entry of Sellers." *Review of Industrial Organization* 62 (2):119–148.

———. 2024. "e-Commerce platforms and self-preferencing." *Journal of Economic Surveys* 38 (4):1516–1543. URL `https://onlinelibrary.wiley.com/doi/abs/10.1111/joes.12594`.

Farronato, Chiara, Andrey Fradkin, and Alexander MacKay. 2023. "Self-Preferencing at Amazon: Evidence from Search Rankings." *AEA Papers and Proceedings* 113:239–43.

Frick, Sarah J., Deborah Fletcher, and Austin C. Smith. 2023. "Pirate and Chill: The Effect of Netflix on Illegal Streaming." *Journal of Economic Behavior & Organization* 209:334–347. URL `https://doi.org/10.1016/j.jebo.2023.03.013`.

Gutiérrez, Germán. 2022. "The Welfare Consequences of Regulating Amazon." URL `http://germangutierrezg.com/Gutierrez2021_AMZ_welfare.pdf`. Working paper, University of Washington, Foster School of Business.

Hagiu, Andrei and Bruno Jullien. 2014. "Search diversion and platform competition." *International Journal of Industrial Organization* 33:48–60.

Hagiu, Andrei and Julian Wright. 2015. "Marketplace or Reseller?" *Management Science* 61 (1):184–203.

Hunold, Matthias, Reinhold Kesler, and Ulrich Laitenberger. 2020. "Rankings of Online Travel Agents, Channel Pricing, and Consumer Protection." *Marketing Science* 39 (1):92–116. URL `https://doi.org/10.1287/mksc.2019.1167`.

Jürgensmeier, Lukas, Jan Bischoff, and Bernd Skiera. 2024. "Opportunities for Self-Preferencing in International Online Marketplaces." *International Marketing Review* 41 (5):1118–1132. URL https://doi.org/10.1108/IMR-03-2024-0094. Open Access under CC BY 4.0.

Jürgensmeier, Lukas and Bernd Skiera. Forthcoming. "EXPRESS: Measuring Self-Preferencing on Digital Platforms." *Journal of Marketing* 0 (ja):00222429251360772. URL https://doi.org/10.1177/00222429251360772.

Khan, Lina M. 2017. "Amazon's Antitrust Paradox." *Yale Law Journal* 126 (3):710–805. URL https://www.yalelawjournal.org/note/amazons-antitrust-paradox.

———. 2019. "The Separation of Platforms and Commerce." *Columbia Law Review* 119 (4):973–1098. URL https://www.columbialawreview.org/content/the-separation-of-platforms-and-commerce/.

Lam, Chungsang Tom, Meng Liu, and Xiang Hui. 2021. "The geography of ridesharing: A case study on New York City." *Information Economics and Policy* 57:100941. URL https://www.sciencedirect.com/science/article/pii/S0167624521000299.

Lam, H. Tai. 2021. "Platform Search Design and Market Power." URL https://www.anderson.ucla.edu/sites/default/files/document/2021-09/Lam_JMP_Platform_Search_Design.pdf. Working paper.

Lee, Kwok Hao and Leon Musolff. 2023. "Entry Into Two-Sided Markets Shaped By Platform-Guided Search." URL https://lmusolff.com/papers/Entry_and_Platform_Guided_Search.pdf. Working paper.

Leung, Tin Cheuk, Shi Qi, and Koleman Strumpf. 2025. "An Empirical Decomposition of Information vs Sentiment Shifts in Advertising." Working paper.

McKenzie, Jordi, Paul Crosby, Joe Cox, and Alan Collins. 2019. "Experimental evidence on demand for "on-demand" entertainment." *Journal of Economic Behavior & Organization* 161:98–113. URL https://www.sciencedirect.com/science/article/pii/S0167268119300873.

Mestyán, Márton, Taha Yasseri, and János Kertész. 2013. "Early Prediction of Movie Box Office Success Based on Wikipedia Activity Big Data." *PLOS ONE* 8 (8):1–8.

Parker, Geoffrey, Georgios Petropoulos, and Marshall W. Van Alstyne. 2020. "Digital Platforms and Antitrust." URL https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3608397. Working Paper.

Porter, Rick. 2023. "Netflix's New View Metric Is Easier to Grasp, but It Still Doesn't Say a Lot." *The Hollywood Reporter* URL https://www.hollywoodreporter.com/tv/tv-news/netflix-viewership-data-metric-explained-1235521282/. Accessed: 2025-05-09.

Raval, Devesh. 2023. "Steering in One Click: Platform Self-Preferencing in the Amazon Buy Box." URL https://deveshraval.github.io/buyBox.pdf. Working Paper.

Reimers, Imke and Joel Waldfogel. 2023. "A Framework for Detection, Measurement and Welfare Analysis of Platform Bias." Working Paper.

Reimers, Nils and Iryna Gurevych. 2019. "Sentence-BERT: Sentence embeddings using Siamese BERT-networks." In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

Rosenbaum, Paul R and Donald B Rubin. 1983. "The central role of the propensity score in observational studies for causal effects." *Biometrika* 70 (1):41–55.

Shuman, Taylor. 2021. "Loss of {The Office} Could Cost Netflix $935 Million in Annual Subscriptions." URL https://www.inmyarea.com/research/netflix-losing-subscribers-when-the-office-leaves. InMyArea Research.

Teng, Xuan. 2023. "Self-preferencing, Quality Provision, and Welfare in Mobile Application Markets." URL https://drive.google.com/file/d/1El3Staxp8-Z8XKgt1R1Su14ia4YfQWv7/view. Working Paper.

U.S. Congress. 2021. "S.2992 - American Innovation and Choice Online Act." https://www.congress.gov/bill/117th-congress/senate-bill/2992/text. 117th Congress (2021–2022), Introduced in Senate on 10/18/2021. Accessed: 2025-05-09.

Yuan, Yihao. 2025. "Exclusive Contracts in the Video Streaming Market." URL https://yihao-yuan.github.io/MyWebsite/streaming_exclusivity.pdf. Working paper, The Wharton School, University of Pennsylvania.

Zervas, Georgios, Davide Proserpio, and John W. Byers. 2017. "The Rise of the Sharing Economy: Estimating the Impact of Airbnb on the Hotel Industry." *Journal of Marketing Research* 54 (5):687–705.

Zhu, Feng and Qihong Liu. 2018. "Competing with Complementors: An Empirical Look at Amazon.com." *Strategic Management Journal* 39 (10):2618–2642.

Zou, Tianxin and Bobby Zhou. 2023. "Self-preferencing and Search Neutrality in Online Retail Platforms." URL https://papers.ssrn.com/abstract=3987361. SSRN working paper.

Appendix A Self-Preferencing Test: Robustness Checks

To assess the robustness of our main findings on Netflix self-preferencing reported in Section 6, particularly those in Table 5, we perform a series of sensitivity analyses using alternative sample restrictions and variable constructions. First, we replace log of Wikipedia page views with the raw page view counts. Second, we limit the sample to the top 300 titles by weekly Wikipedia page views, thereby focusing the analysis on content with the highest audience interest and reducing noise from low-view titles. Third, we replace the level of Wikipedia page views with a normalized measure—Wikipedia page view share—which captures the proportion of total Wikipedia traffic accounted for by each title in a given week. This controls for potential week-to-week variation in overall Wikipedia usage. Lastly, we augment our baseline specification by controlling for whether a title is available on other streaming platforms besides Netflix. This check is restricted to weeks in 2024, during which we are able to observe platform availability through third-party APIs. Across all these alternative specifications, we find that the positive association between title popularity and Top 10 inclusion remains robust, and that Netflix originals continue to be systematically favored, consistent with the main findings in the paper.

Appendix A.1 Actual Views

As a first robustness check, we replace the log of Wikipedia page views with the raw page view counts. The results, reported in Table A1, are qualitatively similar to those in Table 5, though the estimated self-preferencing effects are larger. In this specification, when two titles have the same number of Wikipedia page views, Netflix Originals—both movies and series—are roughly twice as likely to appear in the Top 10, compared with an increase of only 10–30% in Table 5.

Appendix A.2 Top 300 Titles

In our second robustness check, we restrict the analysis to the top 300 titles in each week based on Wikipedia page views. The results are reported in Table A2. This restriction substantially reduces the sample size—from approximately 700,000 observations for

Table A1: Self-Preferencing Regression Results (Actual Views)

	(1) Movie	(2) Movie	(3) Movie	(4) Series	(5) Series	(6) Series
Weekly Wikipedia Pageviews (10,000)	0.00150***	0.00172***	0.00312***	0.00342***	0.00344***	0.00670***
	(0.0000256)	(0.0000282)	(0.0000459)	(0.0000786)	(0.0000808)	(0.000166)
Wikipedia PageviewsXNetflix Original	0.00419***	0.00410***	0.00483***	0.00499***	0.00490***	0.00627***
	(0.0000401)	(0.0000420)	(0.0000479)	(0.0000831)	(0.0000854)	(0.000101)
Wikipedia PageviewsXExpiring in 4 Weeks		-0.00159***	-0.000976***		-0.00344***	-0.00276***
		(0.0000705)	(0.0000739)		(0.000521)	(0.000521)
Wikipedia PageviewsXAvailable (UK)			-0.00197***			-0.00330***
			(0.0000690)			(0.000159)
Wikipedia PageviewsXAvailable (Canada)			-0.000167**			-0.00135***
			(0.0000708)			(0.000189)
Wikipedia PageviewsXIndian Movie			-0.00132***			-0.000687
			(0.0000749)			(0.000484)
Constant	0.000421***	0.000324***	0.000263***	-0.000810***	-0.000757***	-0.000893***
	(0.0000598)	(0.0000609)	(0.0000609)	(0.0000897)	(0.0000922)	(0.0000924)
Observations	711329	683316	683316	399718	380803	380803
R^2	0.05	0.05	0.05	0.17	0.17	0.17

Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

movies and 400,000 for series in the full sample, to around 55,000 to 58,000 observations for each content type. Despite the reduced sample, the estimated coefficients remain essentially unchanged from those in Table A1. This result is not surprising, as the vast majority of titles appearing in the Netflix Top 10 lists also rank among the top 300 most-viewed titles on Wikipedia in their respective weeks. These findings confirm that our main results are not driven by noise in the lower tail of the popularity distribution.

Appendix A.3 View Shares

In a third robustness check, we replace the level of Wikipedia page views with a normalized measure—Wikipedia page view share—which captures the proportion of total Wikipedia traffic accounted for by each title in a given week. The results are presented in Table A3. While the interpretation of the coefficients changes slightly due to the rescaling of the dependent variable, the qualitative patterns remain very similar to those in Table A1.

We continue to find strong evidence of self-preferencing: Netflix original movies and series are significantly more likely to appear in the Top 10 list compared to non-originals with the same Wikipedia share. As before, the magnitude of self-preferencing is larger for

Table A2: Self-Preferencing Regression Results (Top 300)

	(1) Movie	(2) Movie	(3) Movie	(4) Series	(5) Series	(6) Series
Weekly Wikipedia Pageviews (10,000)	0.00102***	0.00121***	0.00258***	0.00340***	0.00340***	0.00693***
	(0.0000851)	(0.0000946)	(0.000151)	(0.000207)	(0.000212)	(0.000431)
Wikipedia PageviewsXNetflix Original	0.00435***	0.00431***	0.00499***	0.00504***	0.00497***	0.00630***
	(0.000125)	(0.000132)	(0.000150)	(0.000211)	(0.000216)	(0.000249)
Wikipedia PageviewsXExpiring in 4 Weeks		-0.00122***	-0.000660***		-0.00364***	-0.00294**
		(0.000221)	(0.000232)		(0.00127)	(0.00127)
Wikipedia PageviewsXAvailable (UK)			-0.00190***			-0.00340***
			(0.000221)			(0.000389)
Wikipedia PageviewsXAvailable (Canada)			-0.000127			-0.00145***
			(0.000226)			(0.000465)
Wikipedia PageviewsXIndian Movie			-0.00135***			-0.000250
			(0.000238)			(0.00121)
Constant	0.0114***	0.0109***	0.0103***	-0.000200	0.000122	-0.000964
	(0.000708)	(0.000730)	(0.000731)	(0.000633)	(0.000646)	(0.000657)
Observations	57380	54988	54988	57399	55004	55004
R^2	0.05	0.05	0.05	0.18	0.18	0.18

Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

Table A3: Self-Preferencing Regression Results (View Shares)

	(1) Movie	(2) Movie	(3) Movie	(4) Series	(5) Series	(6) Serie
Weekly Wikipedia Pageviews Share (%)	0.0718***	0.0778***	0.141***	0.137***	0.139***	0.281
	(0.00111)	(0.00119)	(0.00194)	(0.00309)	(0.00319)	(0.006
Wikipedia Pageviews ShareXNetflix Original	0.169***	0.169***	0.200***	0.204***	0.202***	0.262
	(0.00167)	(0.00174)	(0.00200)	(0.00327)	(0.00337)	(0.003
Wikipedia Pageviews ShareXExpiring in 4 Weeks		-0.0678***	-0.0461***		-0.138***	-0.114
		(0.00360)	(0.00371)		(0.0202)	(0.020
Wikipedia Pageviews ShareXAvailable (UK)			-0.0904***			-0.145
			(0.00278)			(0.006
Wikipedia Pageviews ShareXAvailable (Canada)			-0.00356			-0.058
			(0.00282)			(0.007
Wikipedia Pageviews ShareXIndian Movie			-0.0656***			-0.032
			(0.00300)			(0.019
Constant	0.000179***	0.000111*	0.0000351	-0.000975***	-0.000930***	-0.0010
	(0.0000599)	(0.0000610)	(0.0000609)	(0.0000895)	(0.0000919)	(0.0000
Observations	711329	683316	683316	399718	380803	38080
R^2	0.05	0.05	0.06	0.18	0.18	0.18

Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

movies than for series. Other key results are also consistent with our main findings—titles with higher Wikipedia shares are more likely to enter the Top 10, expiring titles are less likely to appear, titles available in both the UK and Canada are favored, and Indian

movies receive less visibility relative to their Wikipedia popularity. These results reinforce the robustness of our self-preferencing findings to alternative ways of measuring content popularity.

Appendix A.4 Availability on Other Streaming Platforms

The last robustness check augments our baseline specification by controlling for whether a title is available on other streaming platforms besides Netflix. This analysis is limited to weeks in 2024, the only period during which we can reliably observe platform availability using third-party APIs. Results are shown in Table A4, with the first two columns reporting results for movies and the last two for series.

For a fair comparison, the first column in each content type replicates the baseline regression from Table A1, but restricts the sample to 2024 only. The second column then adds controls for availability on other platforms. Across both movies and series, the inclusion of these additional controls has little impact on the estimates. The key findings remain unchanged: we continue to find evidence of Netflix self-preferencing, with original titles more likely to appear in the Top 10 than non-originals of similar popularity. The magnitudes of the coefficients and the significance levels are virtually identical before and after controlling for streaming competition, suggesting that our main results are not confounded by multi-platform availability.

Table A4: Self-Preferencing Regression Results

	(1) Movie	(2) Movie	(3) Series	(4) Series
Weekly Wikipedia Pageviews (10,000)	0.00790***	0.00808***	0.0132***	0.0125***
	(0.000115)	(0.000120)	(0.000269)	(0.000287)
Wikipedia PageviewsXNetflix Original	0.00274***	0.00271***	0.00607***	0.00713***
	(0.000107)	(0.000107)	(0.000179)	(0.000206)
Wikipedia PageviewsXAvailable (UK)	-0.00172***	-0.00197***	-0.00327***	-0.00412***
	(0.000186)	(0.000189)	(0.000378)	(0.000399)
Wikipedia PageviewsXAvailable (Canada)	-0.00411***	-0.00400***	-0.00814***	-0.00777***
	(0.000175)	(0.000175)	(0.000398)	(0.000408)
Wikipedia PageviewsXIndian Movie	0.000547***	0.000519***	-0.00726***	-0.00728***
	(0.000146)	(0.000147)	(0.000883)	(0.000882)
Wikipedia PageviewsXExpiring in 4 Weeks	-0.00491***	-0.00476***	-0.00443**	-0.00443**
	(0.000387)	(0.000387)	(0.00223)	(0.00223)
Wikipedia PageviewsXApple Subscription		0.000612		-0.00719
		(0.000489)		(0.00776)
Wikipedia PageviewsXDisney Subscription		-0.00765		-0.00450
		(0.0157)		(0.0236)
Wikipedia PageviewsXHulu Subscription		-0.00311***		-0.00220***
		(0.00110)		(0.000516)
Wikipedia PageviewsXParamount Subscription		-0.00505***		0.00255***
		(0.000545)		(0.000246)
Wikipedia PageviewsXPeacock Subscription		-0.00435***		0.00476***
		(0.000786)		(0.000583)
Wikipedia PageviewsXPrime Subscription		0.00196***		0.00703***
		(0.000521)		(0.000918)
Constant	-0.000469***	-0.000455***	-0.00121***	-0.00116***
	(0.000117)	(0.000117)	(0.000155)	(0.000155)
Observations	192265	192265	120834	120834
R^2	0.06	0.06	0.16	0.16

Standard errors in parentheses

* $p < .1$, ** $p < .05$, *** $p < .01$

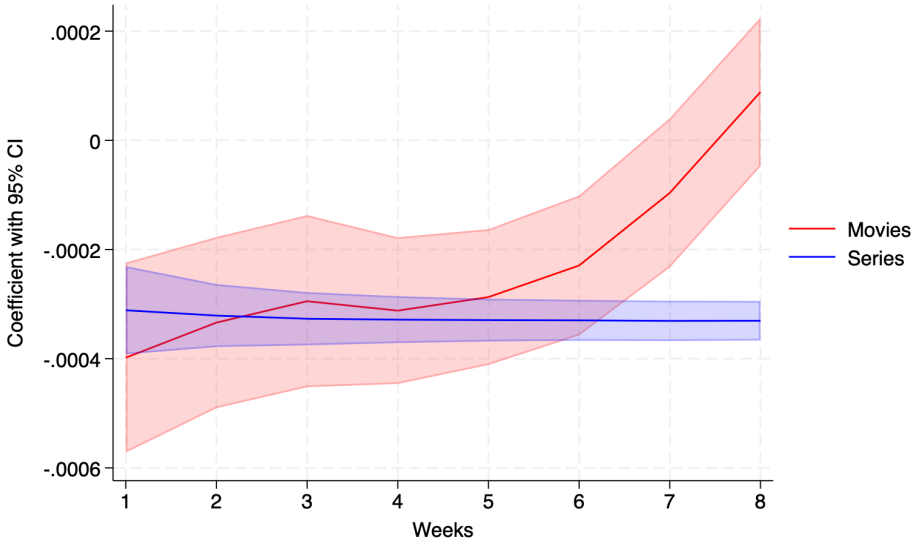
Appendix B Robustness Check: Varying Expiration Window Definition

In our main analysis (Table A1), we find that titles set to expire within four weeks are significantly less likely to appear in the Netflix Top 10 list, even after controlling for their Wikipedia page views. This suggests that Netflix may systematically deprioritize content nearing the end of its licensing window.

To assess the robustness of this finding, we re-estimate the same regression model while varying the expiration indicator across different window lengths, from 1 week to 8 weeks.

Specifically, we redefine the expiration dummy to equal one if the title is set to expire in k weeks, where k ranges from 1 to 8, and plot the resulting interaction coefficients in Figure A1.

Figure A1: Effect of Expiring by Weeks and Content Type



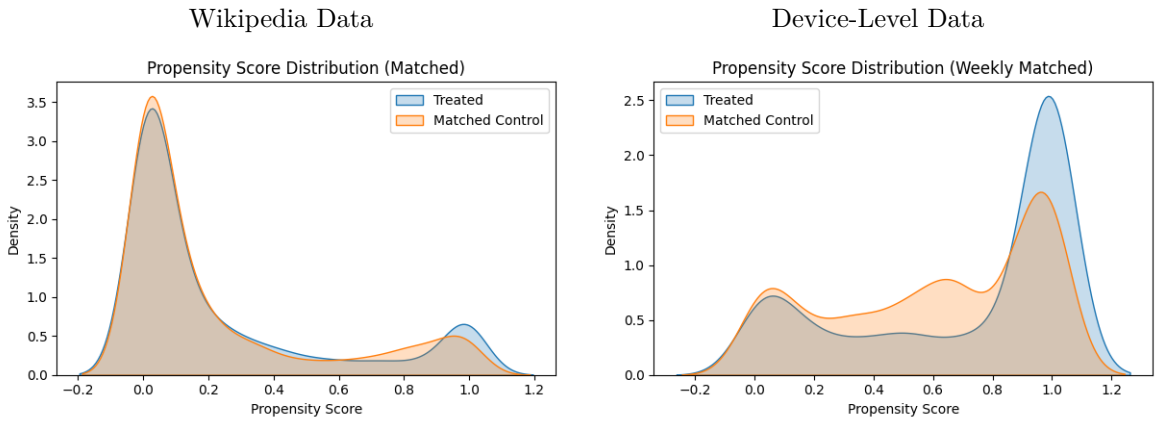
The results remain consistent across specifications: expiring titles are less likely to be featured in the Top 10, regardless of whether expiration is defined as occurring in 1, 2, or up to 8 weeks. The magnitude of the effect is stable and statistically significant for both movies and series across most expiration windows. The only exception is for movies when the expiration window is defined at 8 weeks, where the coefficient becomes statistically insignificant—likely reflecting that expiration effects weaken when the time horizon is relatively long.

Overall, these results confirm that our finding of negative visibility bias for expiring titles is robust to alternative definitions of the expiration window.

Appendix C Propensity Score Distribution

Figure A2 illustrates the distribution of estimated propensity scores for treated titles (in blue) and their matched control counterparts (in orange) across two datasets. The left panel shows results based on Wikipedia page view data, while the right panel presents the corresponding distribution using device-level streaming data.

Figure A2: Propensity Score Distribution



To construct these figures, we first estimate propensity scores within strata defined by content type (movie or series), Netflix Original status, and calendar week, using a logistic regression model. In the Wikipedia-based analysis, the model includes weekly page views as a proxy for popularity, along with availability indicators (UK and Canada), Netflix Original status, and one-hot encoded genre and country indicators. For the device-level analysis, we use weekly duration watched in place of page views, with the same additional covariates. Nearest-neighbor matching is then applied within each stratum to pair each treated title with a control title that has a similar propensity score.

The resulting density plots show that the distribution of propensity scores for treated and matched control titles align closely in both panels. This suggests that the matching procedure achieves good covariate balance and ensures that treated and control groups are comparable in terms of observable characteristics that predict Top 10 inclusion.