Video Advertising by Twitch Influencers*

Yufeng Huang
University of Rochester
Simon Business School

Ilya Morozov
Northwestern University
Kellogg School of Management

March 23, 2022

Abstract

We study the effectiveness of influencer marketing in the video game industry. To this end, we construct a novel dataset on video game streaming on Twitch.tv, the largest video game streaming platform in the world, by monitoring live streams every 10 minutes for eight months. Leveraging these high-frequency data, we isolate plausibly exogenous variation in streamers' daily schedules and use it to estimate the extent to which live streaming brings additional players into broadcasted games. We find that organic live streams only marginally increase the number of concurrent players in these games. We also find that sponsored streams solicited by game publishers are even less effective than organic streams, implying that sponsored streams generate, on average, negative return on investment (ROI). This result suggests that influencer promotions are less effective than previously thought. We then examine heterogeneous returns to streaming by estimating generalized random forests, and we find that sponsored streaming can significantly benefit games released by small publishers, inexpensive games, and "niche" games that strongly appeal to small groups of consumers. Therefore, despite the negative average ROI, influencer promotions may generate high returns when they promote games by lesser-known publishers or inform consumers about appealing game attributes.

*Contacts: Yufeng Huang, email: yufeng.huang@simon.rochester.edu; Ilya Morozov, email: ilya.morozov@kellogg.northwestern.edu. We thank seminar and conference participants at University of Delaware, HKUST, Marketing Science Conference, and University of Rochester for their helpful comments and suggestions. We also thank James Ryan for excellent research assisstance. All errors are our own.

1 Introduction

Over the last decade, influencer marketing has grown into a $14 billion industry.[1] This rapid growth is partly driven by the increasing popularity of major video content platforms, such as Twitch, YouTube, and TikTok, where influencers distribute their content and build their own fan communities. By 2022, YouTube has become the second most visited website in the world that attracts 30 million visitors per day, and Twitch has grown into a video game streaming giant that hosts over 100,000 live channels and attracts over 2.5 million viewers at any point in time.[2] These large audiences create a unique opportunity for companies to promote products and services by having video influencers showcase them during live streams. In fact, many practitioners believe that because influencers can better engage with their audiences and appear more trustworthy, influencer promotions have greater return on investment (ROI) than traditional advertising.[3]

This general excitement about influencer marketing, as it turns out, has little empirical support. When explaining why influencer promotions are effective, many practitioners cite anecdotes in which influencer marketing supposedly generated extraordinarily high ROI.[4] Nevertheless, industry observers caution against putting too much weight on these anecdotes, emphasizing that the industry has yet to devise reliable ways of measuring the effectiveness of influencer promotions.[5] Absent experimental variation, measuring the effects of influencer marketing poses a significant empirical challenge. Influencers can often choose which products to promote and when to promote them. They may choose to promote high-quality products at a time when product sales are already trending up, which generates a simultaneity bias familiar from the literature on the effects of word-of-mouth (Seiler et al., 2018) and advertising (Shapiro et al., 2020). In fact, the prior work on influencer marketing emphasizes this simultaneity bias as the central empirical challenge (Li et al., 2021; Yang et al., 2021). Addressing this challenge would help researchers and practitioners understand whether influencer promotions are indeed a highly effective marketing channel.

We study the effectiveness of influencer marketing in the video game industry, and we collect unique high-frequency data that enable us to address the main identification challenge. Specifically, we ask how video game streaming on Twitch affects the popularity of broadcasted games, defined as the number of their concurrent players. Twitch is the largest video game streaming platform in the world, hosting over 90% of all streaming content.[6] Streamers broadcast their gaming sessions

[1] Influencer marketing industry has grown to $13.8 billion by 2021 (Influencer Marketing Hub, 2021).

[2] YouTube statistics come from the official platform's blog (www.blog.youtube), and Twitch statistics are taken from a third-party Twitch monitoring website (twitchtracker.com/statistics).

[3] See the 2019 industry report by MediaKix (Bailis, 2019).

[4] Nielsen Catalina Solutions conducted a correlational study, concluding that influencer marketing generates 11 times higher ROI than online banner ads.

[5] See https://influencermarketinghub.com/influencer-marketing-roi/.

[6] See the Streamlabs and StreamHatchet Quarterly Report for Q3 (2020).

live while commenting on the game and reacting to viewers' chat messages. We construct an original high-frequency dataset on Twitch streaming and video game usage to estimate the causal effect of streaming. For eight months, from May to December 2021, we continuously monitor the broadcasting activities of 60,000 pre-selected Twitch streamers and gather information about their live streams every 10 minutes. In each 10-minute window, we record which streamers are live, what games they are streaming, and how many viewers are watching each stream. We also collect data on the number of concurrent players in each game with the same 10-minute frequency. When combined, these high-frequency data allow us to examine how the number of players in a game responds to broadcasting activities in real time, which proves crucial to our empirical strategy.

Video games gain and lose popularity due to general market trends as well as in-game events, new version releases, and esports competitions. Responding to these temporal changes, streamers tend to broadcast trending games, thus generating a simultaneity bias. In an ideal experiment, we would make streamers broadcast random games at random times of the day and measure the corresponding lift in the number of concurrent players in the games they broadcast. Our empirical strategy mimics this ideal experiment. We assume that, although streamers strategically decide which days to work on and which games to broadcast, their exact streaming schedules within a day do not respond to real-time changes of game popularity. Streamers often shift their work hours to accommodate other demands on their time, such as university classes and part-time jobs. They may finish broadcasting earlier than planned due to fatigue or later than planned if events that occur in the game lead to a longer gaming session. Because of these idiosyncratic decisions, many Twitch streamers follow irregular streaming schedules when broadcasting a given game. Our strategy uses this within-day variation in streaming schedules as an exogenous shifter of the game's viewership on Twitch. Using this variation, we estimate how increasing the Twitch viewership of a given game affects the number of this game's concurrent players.

We find that *organic* (non-sponsored) streams bring additional players into the broadcasted games. We estimate an elasticity of 0.033, implying that doubling the number of people watching a game on Twitch in a given hour increases the number of this game's concurrent players by 3.3%. This effect is both small and short-lived: the initial effect becomes 17% weaker in every subsequent hour and dissipates to 15% of its initial magnitude within about ten hours. These estimates inform the ongoing debate in the video game industry on whether live streaming benefits game publishers. We find that organic live streams have positive effects on the short-term popularity of video games. This finding contradicts a common belief that live streams divert consumers from playing games by providing an alternative source of entertainment and allowing them to experience the game-play without paying (Johnson and Woodcock, 2019).[7] In other words, we find that watching live

[7]Concerned about this substitution, Nintendo restricted streaming of their games, and a creative director of Google Stadia suggested that publishers should charge streamers royalty fees (PC Gamer, 2020; SVG News 2020).

streams and gaming are complementary leisure activities rather than substitutes. In this sense, our results reinforce those of Li et al. (2021), who empirically show that YouTube videos and e-sports tournaments serve as complements to gaming.

We also find that *sponsored* streams are much less effective than *organic* streams at bringing additional players to the broadcasted games. Specifically, organic streams are four times as effective as sponsored streams, with the estimated elasticities of 0.030 and 0.007. In fact, the estimated elasticity of sponsored streams is half the size of the median effect of TV advertising on sales estimated by Shapiro et al. (2020). We therefore do not find any evidence that influencer promotions are more effective than traditional advertising. To assess how profitable it is for publishers to sponsor live streams, we collect daily data on the subscription revenues of individual streamers and use it to proxy their hourly income. Combining this hourly income with the estimated streaming elasticities, we find that sponsored streams generate negative ROI for most games in our sample. Put another way, a typical game publisher in this industry would find that sponsoring influencer promotions is not worth the investment.

Although Twitch influencer promotions generate negative average returns, they might still be effective at promoting relatively unknown games and at informing consumers about the quality and gameplay of these games. We enrich our analysis by allowing streaming elasticities to depend on game attributes that might be revealed during live broadcasts and on variables proxying which games are relatively unknown. We then estimate heterogeneous streaming elasticities using generalized random forests (Athey et al., 2019). This heterogeneity analysis helps us understand when and why influencer promotions work, distinguishing our paper from the prior literature that focuses on estimating the average effect of such promotions on product usage and sales (Li et al., 2021; Yang et al., 2021).

Our results show substantial heterogeneity in streaming elasticities across games. We find that live streaming has large positive effects on games released by small publishers. Traditionally, these publishers have significantly lower marketing budgets than major video game conglomerates, and consumers might often be unaware of the games they release. By contrast, live streaming may increase this awareness by informing consumers about "indie" games released by these publishers. Such an effect is consistent with the prior work that documents the awareness effect of advertising (Honka et al., 2017; Tsai and Honka, 2021). Thus, one implication of our result is that the growth of live streaming might reduce the entry barriers in this industry by providing small publishers a new promotion channel.

We also find that Twitch streaming has large positive effects on inexpensive games as well as "niche" games that strongly appeal to some consumers. This result suggests that, by watching several hours of gameplay on Twitch, consumers might learn about both vertical and horizontal game attributes. In this sense, our results suggest that one can view live streams as informative advertising for video games (Grossman and Shapiro, 1984; Ackerberg, 2003). As a whole, our results show that despite the negative average ROI, Twitch influencer promotions may generate high positive returns when they are used to promote games by lesser-known publishers or to inform consumers about appealing game attributes.

Our paper contributes to the literature on the effectiveness of marketing media. Much of the existing literature focuses on estimating the effects of TV advertising (Lodish et al., 1995; Liaukonyte et al., 2015; Shapiro et al., 2020), online advertising (Johnson et al., 2017; Gordon et al., 2019), and word-of-mouth (Lovett and Staelin, 2016; Seiler et al., 2017). We contribute by studying the effects of influencer marketing, an emerging marketing channel that has gained momentum over the past few years. The closest related papers are Li et al. (2021) on the effect of YouTube influencer content about gaming and Yang et al. (2021) on the effect of TikTok influencer videos. Both papers use within-product variation in video uploads across days and examine the impact of these uploads on the usage and sales of promoted products. We attempt to improve upon their analysis by constructing a novel dataset and developing an identification strategy that leverages high-frequency data. By leveraging within-day variation in streamers' schedules, our approach brings the analysis to the level of granularity at which the simultaneity bias is less plausible.[8] Given the increasing availability of high-frequency data, this empirical strategy might prove helpful in future research on influencer marketing. Using our empirical strategy, we find that sponsored streams only marginally increase the popularity of broadcasted games. Our findings therefore question the conventional wisdom that influencer marketing generates much greater returns on investment (ROI) than that of traditional advertising.[9]

2 Measuring Video Streaming on Twitch

2.1 What is Twitch?

Twitch.tv is an Amazon-owned video live streaming platform mostly dedicated to streaming video games and broadcasting esports competitions. The platform has experienced tremendous growth over the last decade. In 2012, it had only several thousand registered channels and 70,000 concurrent viewers. By 2022, it grew into a video streaming giant that hosts over 100,000 live channels and attracts over 2.5 million viewers at any point in time.[10] This growth was mainly driven by

[8] In this sense, our paper resembles the prior work that measures the effectiveness of TV ads using high-frequency data and discontinuity-in-time research designs (Liaukonyte et al., 2015; He and Klein, 2019).

[9] Outside of studying the promotional effect, Simonov et al. (2020); Rajaram and Manchanda (2020); Hwang et al. (2021); Ershov and Mitchell (2020) examine characteristics of influencer media and the extent to which these characteristics explain how viewers choose what to watch and whom to follow.

[10] These aggregate statistics are from *twitchtracker.com*, a third-party website that monitors the streaming activity on Twitch.tv and reports historical data going back to 2012.

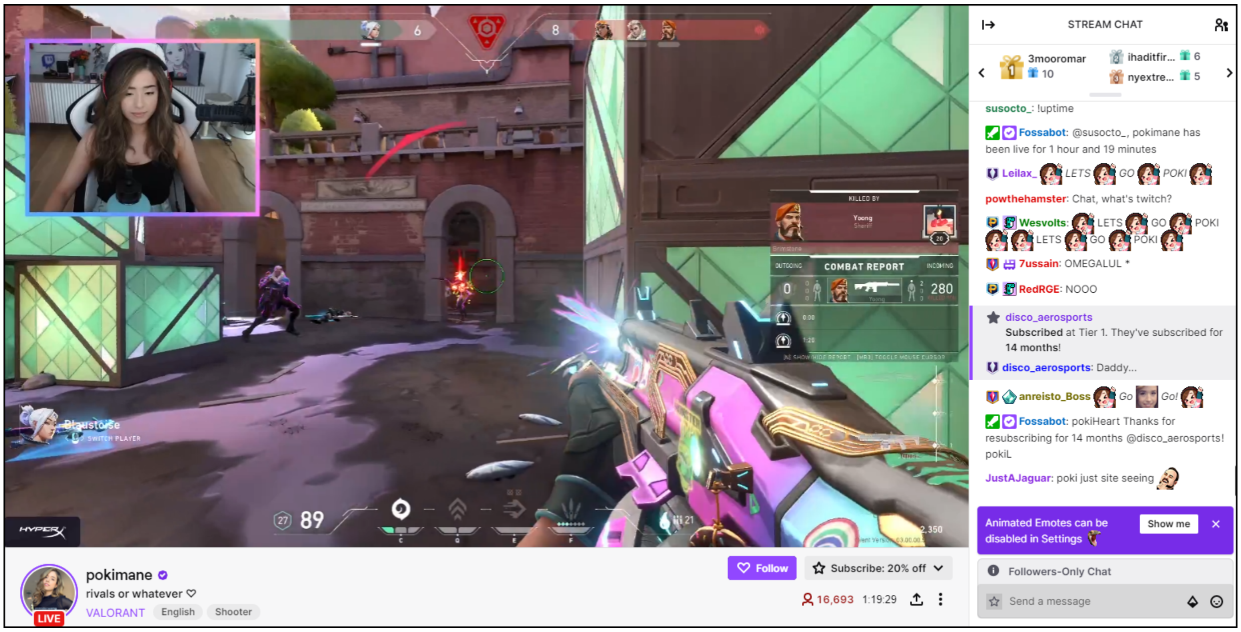

Figure 1: **A famous streamer *Pokimane* broadcasting her gameplay of the shooter game *Valorant* on Twitch.tv.** The main window shows the gameplay that Pokimane is broadcasting live. The window in the top left corner shows the streamer's web camera video. The vertical window on the right is the chat where viewers can send the streamer text messages in real time.

the increasing worldwide popularity of esports and was further enhanced by stay-at-home orders of 2020-2022, which made millions of people around the globe look for new entertainment options. At the time of this writing, Twitch remains the largest video game streaming platform in the world, accounting for 63.6% of total hours of content watched and hosting 91.1% of all video game streaming content.[11] During peak hours, Twitch often attracts 5-6 million viewers, more than major TV networks such as Fox News, CNN, or MSNBC.

Twitch streamers broadcast their gameplay live by sharing their screen and web camera video on the platform. Figure 1 shows what these streams look like from the viewers' perspective. A typical streamer plays a video game live while commenting on the gameplay and engaging with viewers by replying to their chat messages. Popular Twitch streamers follow different styles of live streaming. Some are professional esports players who focus on streaming one game and impress viewers with their gaming skills. One example is a professional "Fortnite" player Ninja, who once had 650,000 people watching him play Fortnite live (The Verge, 2018). Other streamers, such as Auronplay and Summit1G, stream different games every week to introduce variety to their content and expose viewers to video games of different genres. To make their broadcasts more engaging, streamers often attempt to make their content funny and entertaining (see Figure F.1 for some

[11] See the Streamlabs and StreamHatchet Quarterly Report for Q3 (2020). https://blog.streamlabs.com/streamlabs-stream-hatchet-q3-live-streaming-industry-report-a49adba105ba

notable examples). One such example is a streamer DrDisrespect whose vibrant personality and unique look – mullet hairstyle, 80s-style mustache, and polarized sunglasses – made him one of the most popular streamers on the platform.

2.2 Data on Twitch Streaming and Game Usage

To study how Twitch streaming affects the popularity of broadcasted games, we require a dataset that contains information about both live streaming and game usage. We construct a novel dataset by combining several data sources. First, we collect high-frequency data from Twitch by continuously monitoring live streaming and game viewership on the platform. These data describe when individual streamers go live, what games they stream, and how many viewers they attract at each point in time. Second, we also collect high-frequency data on the number of people currently playing each game. To obtain these data, we continuously monitor Steam, the largest online video game platform in the world. These data help us track changes in the popularity of several hundred games during periods when streamers broadcast these games on Twitch. Finally, we complement these two datasets with information on daily (self-reported) subscription counts of individual streamers, which helps us estimate their hourly income.

Video streaming data from Twitch. We monitored video game streaming on Twitch for almost eight months between May 11, 2021 and December 31, 2021. We first pre-selected a list of 60,000 streamers during a three-week pilot period. Each streamer was selected with the probability proportional to the total number of viewers they attracted on Twitch during this pilot period (see Appendix A for details). Then, during the following eight months, we sent high-frequency requests to Twitch API to retrieve information about each streamer. Every 10 minutes, we requested the status of each streamer (online or offline), the concurrent number of viewers, the game they were streaming, and the title of each stream. Most streamers went live on Twitch at least once during the sample period, so our sample covers 96.8% of the streamers we attempted to track (58,060 out of 60,000). Tracking streamers at high-frequency enables us to record the exact times at which a streamer starts and ends broadcasting any game, which serves as the main input to our empirical strategy (see Section 3).

Table 1 summarizes live streaming activity on Twitch, both overall and for the most popular streamers. Throughout the paper, we measure streamer popularity using the average number of concurrent viewers. In the last two rows, the table reports averages (1) across top 5% most popular streamers, and (2) across all 58,060 streamers. The average streamer attracts only around 150 viewers and streams about 5.4 hours per day conditional on working on that day (14.6 hours per week). By contrast, the top 15 streamers attract over 30,000 viewers at any given time and often

Table 1: **Streaming activity on Twitch.**

#	Twitch Streamer	Primary Game Streamed on Twitch.tv	Average Concur. Viewers	Maximum Concur. Viewers	Daily Hours Stream	Weekly Hours Stream	Stream Start Time Avg (S.D.)	Stream End Time Avg (S.D.)
1	AuronPlay	GTA V	104,930	318,181	4.0	5.3	15:20 (0:50)	19:20 (0:50)
2	RanbooLive	N/A	75,695	234,626	2.5	21.0	21:00 (2:50)	23:40 (3:10)
3	Ibai	GTA V	75,513	1,538,645	4.9	10.0	16:30 (2:50)	21:20 (2:30)
4	Sapnap	UNO	72,992	186,592	2.0	3.0	22:20 (3:20)	0:20 (2:40)
5	xQcOW	GTA V	72,481	175,453	15.7	52.3	14:00 (7:20)	5:40 (4:30)
6	loud_coringa	GTA V	65,142	307,450	7.6	10.8	21:10 (5:30)	4:50 (3:10)
7	RocketLeague	Rocket League	58,233	208,124	5.0	30.8	18:00 (3:30)	23:10 (3:40)
8	Flashpoint	CS: GO	52,655	128,800	7.3	22.3	15:10 (1:50)	22:30 (3:30)
9	Asmongold	Final Fantasy	52,135	135,042	7.4	18.3	15:10 (1:20)	22:40 (2:20)
10	thisisnotgeorge	Among Us	50,243	113,707	3.5	3.2	21:40 (7:00)	1:10 (6:20)
11	MontanaBlack88	GTA V	48,447	159,731	6.3	5.5	15:20 (2:50)	21:40 (2:20)
12	Rubius	GTA V	44,377	207,592	5.3	16.5	17:30 (1:00)	22:50 (1:50)
13	Mizkif	Jump King	35,246	189,851	5.9	13.5	20:20 (2:30)	2:20 (3:10)
14	karlnetwork	Golf w/Friends	33,127	93,412	5.5	3.3	20:40 (8:50)	2:10 (7:50)
15	shroud	N/A	30,979	334,836	11.3	30.0	16:30 (5:20)	3:50 (5:10)
	Average (top 5% streamers)		2,375	17,988	8.2	22.0	15:20 (4:00)	23:40 (4:20)
	Average (all streamers)		148	1,277	5.4	14.6	17:40 (4:20)	23:00 (4:20)

This table summarizes live streaming activity on Twitch.tv, both overall and for the most popular streamers, across all games and non-game activities. We measure streamer popularity using the average number of concurrent viewers in our dataset. The primary game of each streamer is defined as a game that this streamer broadcasted for the largest number of hours and that attracted at least 25% of total viewers of this streamer. The last two columns show the average start and end times of live streams and report the standard deviations of start and end times in the parentheses. The time is in the UTC time zone (Coordinated Universal Time). The last two rows report the averages (1) across top 5% most popular streamers, and (2) across all 58,060 streamers.

attract hundreds of thousands of viewers in peak times. A Spanish streamer Ibai, for example, peaked at 1,538,645 active viewers by streaming a series of boxing matches that featured other Spanish streamers, establishing a viewership record in our sample. On average, top streamers work 8.2 hours a day and 22 hours a week (see the second-to-last row of Table 1). Column 3 reports the primary game of each streamer, which we define as a game that this streamer broadcasted for the largest number of hours and that attracted at least 25% of total viewers of this streamer. Although many top streamers broadcast hit games with international following, such as "GTA V" and "CS:GO," several of them have risen to the top by broadcasting relatively unknown games like "Jump King" and "Golf with Friends."

Table 2 shows summary statistics for this sample of 599 games, describing how often these games are streamed and watched on Twitch. Panel A reports the total streaming and viewership numbers, whereas Panel B reports the same statistics in terms of daily averages. As daily numbers show, the average game was streamed 102 hours a day in 33 different streams, which generated 24,700 hours of the total watch time. This distribution is largely skewed toward popular games, as the median game in the sample was only streamed three-four times a day for eight hours in total.

Game usage data from Steam. We also collected data on the number of people playing video games at any given point in time. We collected these data from Steam, the largest online video game platform in the world that attracts 62.6 million active players every day (steampowered.com).[12] Because consumers who buy games on Steam have to be logged into their accounts to play, the platform collects accurate data on the number of concurrent players in each game. We first pre-selected 599 games that were most frequently streamed and watched on Twitch during the three-week pilot period (see Appendix A for details). We then sent high-frequency requests to Steam API, retrieving the number of concurrent players of the pre-selected games every 10 minutes, synchronized with Twitch data collection. Tracking the player counts at a high-frequency allows us to study how Twitch broadcasts affect the number of players in a given game in real time, which is critical for our empirical strategy.

Video game attributes. We additionally collected data on video game attributes. Specifically, we gathered each game's publisher name, release date, and price history from Steam's official website, as well as customer ratings and professional critic ratings from the website metacritic.com. We use these attribute data in Section 4 when studying the heterogeneous effects of live streaming across games. Panel C in Table 2 summarizes the distribution of game attributes, and Table F.1 in the Appendix provides specific game examples. The average game is around 3.9 years old, have a Metacritic rating of 78.8 out of 100, and has been released by a small publisher that produced five other games among the 599 titles. Most games in the sample are relatively new and have been released within the last two to three years. Some notable examples are games "Apex Legends" and "Among Us" that experienced a massive surge of popularity often attributed to Twitch streaming (PC Gamer, 2020). Almost half of all games have been released by indie publishers that released only one other game. Many of these indie games have become remarkably popular and established a significant presence on Twitch. For instance, the game "Rocket League," which can be best described as "soccer, but with rocket-powered cars," has become an internationally recognized esports game with prize pools reaching $6 million.

Subscription counts of individual streamers. To estimate the ROI of sponsoring live streams in Section 4.4, we need to estimate how much publishers would need to pay streamers per hour

[12] See "Steam 2020 Year in Review" (Steamcommunity, 2021).

Table 2: **Streaming activity, usage, and attributes of games.**

	Mean	Std.dev	P5	P25	P50	P75	P95
Panel A. Streaming activity, viewership, and usage (total during the observation period)							
No. Streams (All Streamers)	9,079	42,206	14	178	959	4,481	31,464
No. Streams (Top Streamers)	666	3794	2	15	61	239	2093
Hours Streamed (thousands)	28.2	135.8	0	0.4	2.2	12.5	100.8
Hours Viewed (thousands)	6,842	52,482	5	72	394	1,689	17,109
Hours Played (thousands)	30,165	167,518	5	252	1,710	12,267	101,523
Panel B. Streaming activity, viewership, and usage (per day)							
No. Streams (All Streamers)	32.8	152.4	0.1	0.6	3.5	16.2	113.6
No. Streams (Top Streamers)	2.4	13.7	0.0	0.1	0.2	0.9	7.6
Hours Streamed	101.7	490.3	0.1	1.3	8.0	45.0	363.8
Hours Viewed	24,700	189,467	19	262	1,423	6,097	61,767
Hours Played	130,585	725,187	23	1,092	7,406	53,105	439,493
Panel C. Game attributes							
Publisher Size (no. games)	4.9	5.6	1.0	1.0	2.0	9.0	17.0
Years Since Release	3.9	4.2	0.1	0.8	2.7	5.5	12.5
Rating Metacritic Score	78.8	9.7	62.0	75.0	80.0	85.0	91.0
Regular Price (dollars)	21.1	16.3	0.0	10.0	20.0	30.0	60.0
Customer Rating St. Dev.	2.3	0.7	1.0	1.9	2.4	2.8	3.1

This table shows streaming, viewership, and usage statistics for 599 Steam games in our sample. To estimate the number of people watching a stream at any given point in time, we multiply the number of current viewers obtained from Twitch API by 10 minutes (the frequency of data collection). We construct the streaming time statistics in a similar fashion. The variable "years since release" captures the number of years passed between the official game release and the first day of our data collection. The publisher size is the number of games a publisher released among 599 titles in our sample. The regular price is the 95-th percentile of the distribution of daily prices for a given game, which usually corresponds to the non-discounted price.

of broadcasting on Twitch. To this end, we gathered additional data on the number of active subscribers of each streamer. Because viewer subscriptions represent a significant share of streamers' regular income, the estimate we obtain gives us an informative lower bound on the daily and hourly income of top streamers (see Appendix A.2 for details). Twitch does not publish any official data on subscriptions, so we instead obtained subscription data from a third-party website twitchtracker.com, which tracks the number of active subscriptions for Twitch streamers who chose to publicly disclose this information. By tracking 10,000 most-subscribed streamers on a daily basis, we collected the current number of active subscriptions and its breakdown by subscription type (i.e., Tier 1, Tier 2, Tier 3, or Amazon Prime), each of which has a fixed dollar value per subscription. We then converted these subscription counts into daily and monthly income estimates. The resulting estimates capture pre-tax income after the streamer has paid the commission to Twitch.

The main limitation of these data is that streamers self-select into disclosing subscription counts, so we cannot assume our dataset on subscription revenues includes a random subsample of streamers.

We find that the average streamer in our sample has 983 active subscriptions and earns around \$3,800 per month in subscription revenues. In contrast, the average top 5% streamer earns about \$20,000-30,000 in subscription revenues each month. Dividing their monthly subscription income on streaming hours, we obtain that the average top streamer earns about \$144 per hour of live streaming. In Section 4.4, we use this number as an estimate of the hourly income of top streamers.

3 The Effect of Stream Viewership on Game Usage

3.1 Empirical strategy

We aim to estimate the causal effect of Twitch stream viewership on the broadcasted games' popularity, defined as the number of concurrent players. In an ideal experiment, we would make streamers broadcast random games at random times of the day and measure the corresponding lift in the number of concurrent players. Such an experiment is impossible to implement because neither Twitch nor we can control when streamers go live and which games they broadcast. As Seiler et al. (2017) point out, this lack of experimental variation makes it difficult to measure the effect of organic content on product popularity.

We adopt an instrumental variable strategy that mimics the ideal experiment and control for confounding factors using fixed effects. We leverage our high-frequency data and focus on variation in the exact broadcast schedules of top streamers within a given day. Although streamers can strategically decide which days to work on and which games to broadcast, our main identifying assumption is that within a given day, their streaming schedules do not respond to real-time changes of game popularity. This assumption is reasonable in the context of Twitch streaming. Instead of working regular hours, many streamers start working whenever it is convenient for them given other demands on their time, such as university classes and part-time jobs. They may finish broadcasting earlier than planned due to fatigue or later than planned if events that occur in the game lead to a longer gaming session. Because of these idiosyncratic decisions, many Twitch streamers follow irregular streaming schedules when broadcasting a given game. By leveraging this unpredictable variation in streaming schedules, we estimate how Twitch viewership affects game popularity. We now present model-free evidence that supports this empirical strategy.

We first demonstrate that individual streamers follow irregular broadcast schedules. Figure 2 visualizes the variation in broadcast schedules of three streamers, randomly drawn from the pool of the 5% most popular streamers ("top streamers"). In each graph in Figure 2, a row corresponds to 24 hours of a given day, and square markers indicate whether the streamer was live on Twitch in

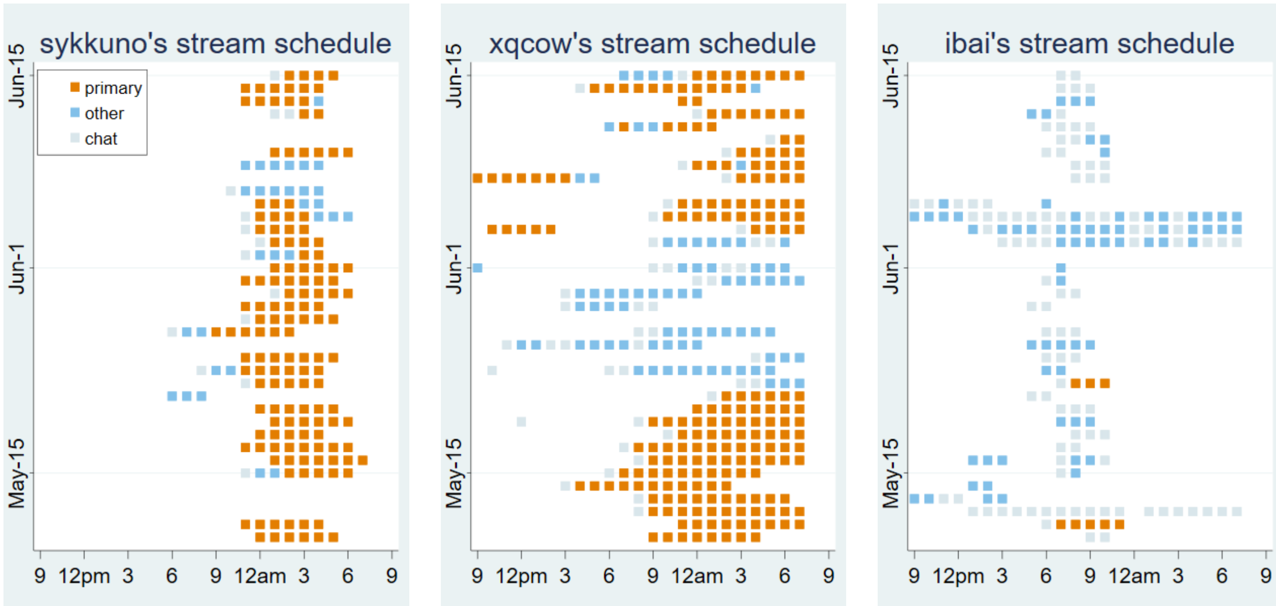

Figure 2: **Daily work schedules of top Twitch streamers.** This figure visualizes the daily work schedules of three streamers, which we randomly selected from the pool of 5% most popular Twitch streamers. Each of the three graphs visualizes the time slots in which each person was streaming live on Twitch (colored squares), with the horizontal axis showing the time of the day and the vertical axis showing different dates during the first month of our sample. We use color to depict whether a streamer was broadcasting their primary Steam game (dark orange), or some other game (light blue), or was "just chatting" (light gray).

each hour of that day. Marker colors represent primary Steam games (dark orange), other games (light blue), and "just chatting" with the audience (light gray). The figure shows that streamers xQcOW and Ibai have haphazard schedules that change every day. They often shift their work schedules by several hours and switch between games. Streamer Sykkuno follows a more stable schedule, but even his schedule exhibits substantial variation across days. In the Appendix, we visualize schedules of several other top streamers and show that most of them have highly variable streaming schedules, similar to the three examples here (see Figure F.4).

Table 3 further summarizes this variation in broadcast schedules by reporting the standard deviations of start times, end times, and stream durations. We decompose this variation into three components: (a) variation across games, (b) variation across streamers within a game, and (c) variation across dates within a streamer-game combination. Within-streamer-game variation explains 50-60% of the total variation of broadcast timing. In contrast, about 40% of the total variation is across streamers for a given game and less than 10% is across games. This decomposition confirms that individual streamers follow variable schedules and frequently change the timing of their live broadcasts within a day, thus further supporting our empirical strategy.

Next, we show that, when a top streamer starts to broadcast a game, both Twitch viewership

Table 3: **Variance decomposition for start times, end times, and duration of streams.**

		Variance decomposition (% of total variance)		
	Std.dev	Across games	Across streamers (for a given game)	Across dates (for a given game-streamer)
All streamers:				
Stream start time	6.08	2.0%	46.7%	51.3%
Stream end time	6.16	1.9%	42.8%	55.3%
Stream duration (hrs)	3.31	4.4%	35.2%	60.4%
Top 5% streamers:				
Stream start time	6.06	3.3%	47.4%	49.3%
Stream end time	5.92	4.7%	42.9%	52.3%
Stream duration (hrs)	5.67	6.1%	34.7%	59.2%

The table shows the standard deviations of start times, end times, and duration of streams (column 1) for the main 599 Steam games in our sample. We decompose this variance into three components: (a) variation across games (column 2); (b) variation across streamers within a game (column 3); and (c) variation across dates within a given streamer (column 4). The table reports statistics for all streamers in the upper panel and the same statistics for top 5% streamers in the lower panel.

and the number of players sharply increase. To isolate the effect of a single broadcast, we identify all days on which a given game was broadcasted by at most one top 5% streamer. We call the top streamer's broadcast a *focal stream*. We select focal streams that do not overlap with any other stream of the same game—that is, no other top streamer broadcasts the same game within 10 hours before the start and 20 hours after the start of the focal stream. We then examine the change in the number of viewers and players of the broadcasted game during this 30-hour window. Although these selection criteria help us simplify visualization, we only use them for illustration purposes and relax them in the formal regression analysis.[13]

Panel A in Figure 3 shows the number of viewers before and after the start of the focal stream. The figure plots the regression coefficients from a linear model that regresses the log number of viewers on a set of dummies, each capturing one-hour time periods for the 10 hours before and 20 hours after each stream. To account for systematic variation in game popularity that might correlate with broadcast activities, both across days and within a given day, the regression also controls for game-date, game-hour of the day, and time fixed effects—the same set of fixed effects as in our formal model, which we explain in detail below. As Panel A shows, the number of viewers remains roughly constant before the start of the focal stream, increases sharply after its start, remains high for two hours, and then gradually declines to its initial level before the stream. Panel B shows that

[13]This strategy of isolating discontinuous responses using high-frequency data is similar to that of He and Klein (2019) who estimate the effects of radio ads on the sales of national lottery tickets and Liaukonyte et al. (2015) who leverage exogenous variation in the timing of TV ads to estimate how advertising affects online purchases.

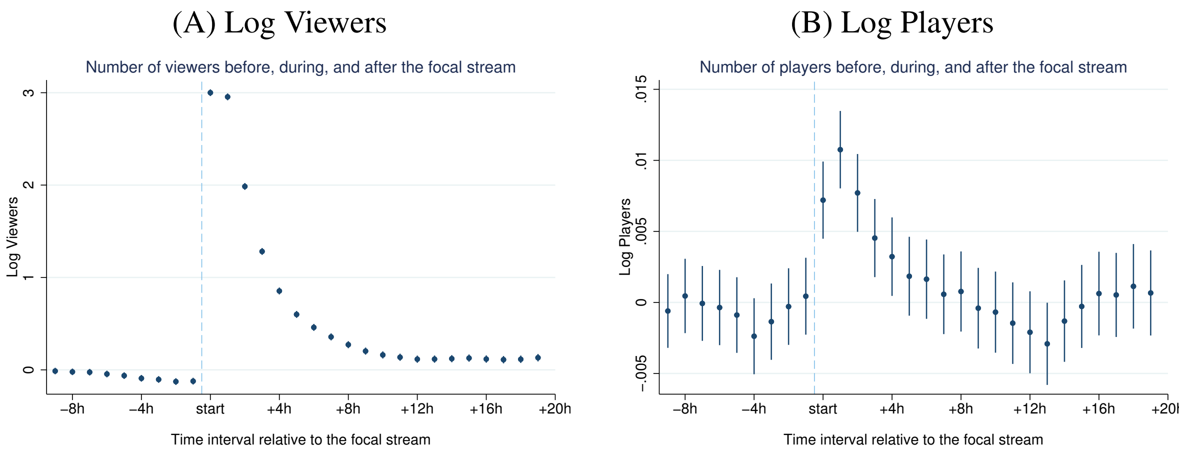

Figure 3: **The number of viewers and players before, during, and after a focal stream.** Both graphs plot the estimated regression coefficients from a linear model that regresses the log number of viewers (left panel) or the log number of players (right panel) on a set of dummy variables for one-hour time intervals, which capture 10 hours before and 20 hours after each stream. The regression controls for game-date, time interval, and hour of the day fixed effects in order to account for predictable changes in game popularity both across days and within a given day. The average focal stream lasts for about four hours.

the number of players of the broadcasted game changes in a similar fashion, suggesting that streams indeed bring additional players into the broadcasted game. Notably, the number of players does not immediately peak after the start of the focal stream when the lift in viewership is the highest, and it slowly returns to its initial level in about eight hours. We interpret this apparent lagged response as preliminary evidence that the increase in stream viewership generates persistent effects on game usage. Motivated by this observation, we now specify a more formal model that allows for (but does not assume) the persistent effects of Twitch viewership.

3.2 Empirical model and estimation

We specify a model that captures both the immediate effect of Twitch viewership on game usage as well as its persistent effects in subsequent time periods. Let j index games, and let t index hour-long time periods. We model the number of players in game j in time t as follows:

$$\log\left(1+players_{jt}\right)=\beta\log\left(1+V_{jt}\right)+\lambda_{j,d(t)}+\mu_{j,h(t)}+\eta_t+\varepsilon_{jt} \tag{1}$$

where V_{jt} is the *stock of total viewership* explained below; $\lambda_{j,d(t)}$ are game-date fixed effects, where $d(t)$ is the date to which the time period t belongs; $\mu_{j,h(t)}$ are game-hour of the day fixed effects; η_t are time fixed effects; and ε_{jt} are idiosyncratic shocks. The game-date fixed effects $\lambda_{j,d(t)}$ capture that games often become more or less popular over time, e.g., due to in-game events or new version releases, which may affect both viewership and game usage. The game-hour of the day fixed effects $\mu_{j,h(t)}$ account for predictable within-day shifts in game popularity, which might occur if different games are played from different time zones (see Figures F.2-F.3 in the Appendix). Finally, time fixed effects η_t, common across games, capture unobserved events that affect the opportunity cost of time of both players and streamers, such as the effects of holidays or major sports events. In this model and the subsequent analysis, we aggregate data to one-hour time intervals to reduce computational complexity. We obtain almost identical estimates when using the original 10-minute time intervals (see Appendix C.1).

The variable V_{jt} in (1) represents the total Twitch viewership stock for game j in time t. One can think of V_{jt} as measuring the cumulative amount of time that viewers have recently spent watching the streams of game j on Twitch, which might cause them to play the game. We define V_{jt} as a weighted sum of the recent number of viewers with geometrically decaying weights:

$$V_{jt} = \sum_{\tau=0}^{T} \delta^{\tau} viewers_{j,t-\tau}. \tag{2}$$

where $viewers_{jt}$ is the total number of people watching Twitch streams of game j in time t, and δ is between zero and one. In estimation, we assume $T = 72$, which allows the effects of viewership to persist for up to three days. This geometric weight specification is similar to the models used to capture persistent advertising effects in the prior work (Shapiro et al., 2020). When put together, equations (1) and (2) yield a model in which an increase in Twitch viewership can influence game usage both concurrently and in future periods.

The parameters of interest are β and δ. We interpret β as the elasticity of game usage with respect to the cumulative viewership stock V_{jt}. We refer to β as the *streaming elasticity*. By contrast, we interpret δ as the decay parameter that captures the persistent effects of viewership. Some viewers might immediately download the game once they see it on Twitch (an immediate effect), while others might research the game later, or they can immediately download it but keep playing it after the stream (a carry-over effect). The decay parameter δ captures the magnitude of the carry-over effect relative to the immediate effect.

The primary source of identification is within-day variation in the broadcast schedules of top streamers. Having controlled for game-date fixed effects $\lambda_{j,d(t)}$ and for other fixed effects in (1), we ask how these live broadcasts drive the game's viewership on Twitch, and how a lift in viewership translates into changes in game usage, both during and after the broadcast. To this end, we construct a vector of instruments $Z_{jt} = \left\{z_{j,t}, z_{j.t-1}, \ldots, z_{j,t-12}\right\}$ that capture how many top streamers broadcast game j in time t and in the 12 hours preceding t (Appendix C.1 shows that our results are robust to including a different number of lagged variables z_{jt}). As we show in Appendix B.1, most variation in these instruments comes from days on which a game was either broadcasted by one top streamer or not broadcasted by top streamers at all. The main identifying assumption is that, conditional on fixed effects, the broadcast decisions of top streamers in the past 12 hours are orthogonal to idiosyncratic shocks in game popularity, ε_{jt}, so that

$$\mathbb{E}\left[\varepsilon_{jt}|Z_{jt}, \lambda_{j,d(t)}, \mu_{j,h(t)}, \eta_t\right] = 0. \tag{3}$$

One might worry that streamers strategically schedule their broadcasts to coincide with significant in-game events (e.g., new version releases or tournaments), which may create a correlation between the instruments Z_{jt} and shocks ε_{jt} even within a day. However, if this was the case, the number of players would increase even before a top streamer goes live, reflecting that the game is already trending up on Twitch by the time the live stream starts. By contrast, Figure 3 shows that the number of players remains relatively constant before the stream and increases sharply right after the start of the stream. We also demonstrate in Appendix C.1 that controlling for game-week instead of game-date fixed effects produces similar estimates of β and δ, suggesting that streamers do not strategically choose their broadcast days within a week based on game-specific demand shocks. By implication, it might be even less likely that they respond to these demand shocks when choosing the exact broadcasting time within a day. We therefore do not find any indication that within-day endogeneity poses a substantial threat to our empirical strategy.

The identifying assumption in (3) implies the condition $\mathbb{E}\left[\varepsilon_{jt}Z_{jt}|\lambda_{j,d(t)}, \mu_{j,h(t)}, \eta_t\right] = 0$, which we use to estimate parameters β and δ. Specifically, we minimize the sum of squared interactions between residuals ε_{jt} and instruments Z_{jt}:

$$(\hat{\beta}, \hat{\delta}) = \arg\min_{(\beta,\delta)} \sum_j \sum_t Z'_{jt} Z_{jt} \left(\log\left(1 + players_{jt}\right) - \beta \log\left(1 + V_{jt}(\delta)\right) - \lambda_{j,d(t)} - \mu_{j,h(t)} - \eta_t\right)^2. \tag{4}$$

This objective function corresponds to GMM estimation with the identity weighting matrix, which under the assumption (3) yields consistent estimates of β and δ. To solve the minimization problem in (4), we perform a golden-section search for the parameter δ, and for a given guess of δ, we estimate parameter β using a closed form 2SLS formula (see Appendix B.2 for details). We obtain clustered standard errors via bootstrap by sampling game-date combinations with replacement.

3.3 The average effect of Twitch streaming

Table 4 presents parameter estimates from the model in (1). The first column shows the estimates from an OLS regression that assumes away persistent effects (i.e., setting $\delta = 0$) and does not include any fixed effects. Because viewership and game usage are highly correlated, the OLS estimate without controls returns a high estimated elasticity of 0.606. The next two columns show

Table 4: **The effect of Twitch viewership on video game usage**

Variable	Parameter	OLS	IV	IV
Log Viewership Stock V_{jt}	β	0.606***	0.015***	0.033***
		(0.002)	(0.002)	(0.009)
Effect Persistence	δ			0.828***
				(0.060)
Game-Date FE		No	Yes	Yes
Game-Hour of day FE		No	Yes	Yes
Time FE		No	Yes	Yes
Observations		3,277,728	3,277,728	3,277,728

Column 1 shows results from an OLS regression that fixes the persistence parameter to zero ($\delta = 0$) and does not control for any fixed effects. Columns 2-3 shows results from our main specification in (1), without persistence (column 2) and with persistence (column 3). Bootstrap standard errors are clustered at the game-date level. *, **, and *** represent significance at the 10%, 5%, and 1% level.

IV estimates obtained using the GMM estimator in equation (4). Column 2 reports IV estimates from a specification that sets persistence to zero ($\delta = 0$), which yields an elasticity estimate of 0.015. This estimate is much lower than that from the OLS regression, consistent with the idea that including instruments Z_{jt} and fixed effects helps remove the simultaneity bias. But because this specification ignores persistent effects, it might still generate a biased estimate of the streaming elasticity β. For example, if the true effect persists several hours after a stream, this model will fail to attribute the elevated post-stream game usage to the effect of the broadcast and might bias the estimated β toward zero. Consistent with this observation, we find that allowing for persistent effects raises the estimated streaming elasticity from $\hat{\beta} = 0.015$ to $\hat{\beta} = 0.033$ (column 3 in Table 4). We also estimate the persistence parameter to be $\hat{\delta} = 0.828$, suggesting that the initial effect becomes 17% weaker in every subsequent hour and dissipates to 15% of its initial magnitude within about ten hours.

The estimated streaming elasticity of 0.033 is broadly in line with the previous work on the effects of word-of-mouth and advertising. For example, Seiler et al. (2017) estimate the elasticity of 0.016 when studying how the number of organic comments on a microblogging platform increases the viewership of discussed TV shows. Their estimates can be compared to ours because they study the impact of word-of-mouth activity, which is analogous to Twitch viewership in our analysis, and they also adopt a measure of consumption (i.e., TV show audience) as the main outcome variable. Similarly, Shapiro et al. (2020) report the mean elasticity of 0.023 and the median elasticity of 0.014 when estimating the effect of TV advertising viewership on the sales of packaged goods.[14]

[14] Seiler et al. (2017) find no strong evidence of persistent effects. Since their estimation leverages one specific shock, i.e., the blocking of microblogging platform Sina Weibo in China, these results might reflect that the persistent effects are difficult to identify from their data. Shapiro et al. (2020) use a similar construction of viewership stock

Table 5: **Streaming elasticities of sponsored and partnered streams.**

Log Viewership Stock V_{jt}	Parameter	All sponsored streams (IV)	Partner program streams (IV)
Organic stream viewership	β^{ns}	0.030***	0.032***
		(0.005)	(0.005)
Sponsored stream viewership	β^{s}	0.007***	0.006***
		(0.002)	(0.001)
Game-Date FE		Yes	Yes
Game-Hour of day FE		Yes	Yes
Time FE		Yes	Yes
Observations		3,277,728	3,277,728

Standard errors are clustered at the game-date level. *, **, and *** represent significance at the 10%, 5%, and 1% level. All specifications include game-date, game-hour of day, and time fixed effects.

In both papers, the authors argue that addressing the simultaneity bias substantially reduces the estimated elasticities. Our results point in the same direction (e.g., see columns 1 and 3 in Table 4), suggesting that one needs to address the simultaneity bias in order to estimate the causal effect of influencer marketing.

Robustness analyses. Our geometric decay model imposes that the concurrent effect of streaming has the same sign and is stronger than the carry-over effect. In practice, however, streaming might divert viewers from playing the broadcasted game, thus creating a negative concurrent effect and a positive carry-over effect on game usage. To explore this possibility, in Appendix C.2, we estimate a distributed lag specification that nonparametrically regresses game usage on the lagged values of $viewers_{jt}$. We find positive concurrent and carry-over effects. Additionally, the estimated effects roughly follow a geometric decay pattern, confirming that our geometric decay model specification does not impose overly rigid assumptions regarding the lagged streaming effects.

Another concern is that game-date fixed effects absorb the carry-over effects of Twitch streams, especially if these effects materialize several days after the stream. If this was the case, our estimated streaming effect would be too small and too short-lived. To examine this possibility, we estimate an alternative specification in which, instead of controlling for game-date and game-hour of the day fixed effects, we control for game-week, game-day of the week, and game-hour of the day fixed effects (see Appendix C.1). We obtain the estimates $\hat{\beta} = 0.034$ and $\hat{\delta} = 0.794$, remarkably similar to our main specification.

V_{jt} and find a persistence parameter around 0.9 at the weekly level. Our estimated $\delta = 0.828$ is at the hourly level, suggesting that the effect of Twitch viewership has a much shorter time span compared to TV advertising of consumer packaged goods.

3.4 Sponsored versus organic streams

Most of the streams in our data are *organic* in the sense that streamers broadcast their gaming sessions without getting paid by publishers. If a publisher were to *sponsor* streams, compensating a top streamer for the time spent broadcasting their game, the effect of such a sponsored stream might be substantially lower than that of organic streams. This might happen because sponsored streams tend to be less engaging or because viewers might negatively react to sponsored content (Ershov and Mitchell, 2020). To empirically assess this conjecture, we identify sponsored streams in our data and estimate the effect of these sponsored streams separately from the effect of organic streams.

We identify sponsored streams in our data by searching for specific keywords in stream titles. Streamers are required to disclose the sponsored status of their streams due to several official requirements. First, the Federal Trade Commission requires all influencers to publicly disclose the sponsorship status (FTC, 2019).[15] Second, Twitch facilitates the match between streamers and game publishers via the internal platform "Bounty Board," which also requires streamers to disclose the sponsored status to viewers. We can therefore identify sponsored streams in our data by searching for the keywords in stream titles that directly reveal the sponsorship status (e.g., *#sponsored*, *#ad*). Additionally, several publishers offer partnership programs to streamers who are willing to regularly broadcast their games. Because partnered streamers are also required to disclose their partner agreement on Twitch, we can identify partnered streams by locating stream titles that mention an official partnership program (e.g., *#ApexLegendsPartner*, *#PubgPartner*). Using these criteria, we find that around 3% of broadcasts by top streamers are sponsored, and 1% are partnered. Although we do not consider how prominent these disclosures are, the recent work of Li (2022) suggests that some Twitch streamers use long stream titles to make the sponsorship appear less prominent on the platform.

We then separately estimate the streaming elasticities for sponsored and organic streams using the following model:

$$\log\left(1+players_{jt}\right)=\beta^{ns}\log\left(1+V_{jt}^{ns}\right)+\beta^{s}\log\left(1+V_{jt}^{s}\right)+\lambda_{j,d(t)}+\mu_{j,h(t)}+\eta_t+\varepsilon_{jt}. \qquad (5)$$

where V_{jt}^{ns} is the cumulative viewership stock of non-sponsored streams, V_{jt}^{s} is the viewership stock of sponsored streams, and both viewership stocks are constructed based on the previously estimated decay parameter $\delta=0.828$. We interpret β^{ns} and β^{s} as streaming elasticities of organic and sponsored streams. We estimate equation (5) via two-stage least squares, using as instruments the current and lagged number of sponsored and organic streams by top 5% streamers. As Table

[15] Source: https://www.ftc.gov/news-events/press-releases/2019/11/ftc-releases-advertising-disclosures-guidance-online-influencers (accessed in February, 2022).

5 shows, organic streams are four times as effective as sponsored streams, with the estimated elasticities of 0.030 and 0.007. We find similar results when we consider sponsored streams funded by official partnership programs (estimated elasticities 0.032 vs 0.006).

We draw several conclusions from these results. First, organic live streams have positive effects on the short-term popularity of video games. This finding contradicts the belief of some practitioners in the video game industry that live streams divert consumers from playing games by providing an alternative source of entertainment and allowing them to experience the gameplay without paying (Johnson and Woodcock, 2019). In other words, we find that live streaming and gaming should be viewed as complementary leisure activities rather than substitutes. Second, our results show that sponsored streams are relatively ineffective at bringing additional players to the broadcasted games. Later in Section 4.4, we use our elasticity estimates to compute the return on investment (ROI) and find that paying for sponsored streams is not profitable for the average game. Our findings therefore question the conventional wisdom that influencer marketing generates much greater ROI than that of traditional advertising.

Although Twitch influencer promotions generate negative returns on average, they might still be effective at promoting relatively unknown games. Given the long format of Twitch broadcasts, such promotions might also reveal rich information about the game's price, quality, and gameplay. In the next section, we empirically explore these conjectures and study what kind of games might substantially benefit from being promoted in sponsored streams on Twitch.

4 Which Games Benefit the Most from Twitch Streaming?

4.1 Overview of potential mechanisms

To understand whether sponsored streams might be effective at promoting certain games, we need to first understand what kind of games are most likely to benefit from live broadcasts. To answer this question, we need to first understand what mechanisms drive the effect of streaming on game popularity. We hypothesize that Twitch streams either inform consumers about the existence of the broadcasted games or reveal information about their price, quality, or gameplay. First, consumers face an enormous choice set and might not be aware of all offered titles. On Steam alone, they encounter an assortment of more than 60,000 games, which grows with thousands of new titles introduced every year. Twitch streams can draw consumers' attention to specific games, thus generating an awareness effect (Honka et al., 2017; Tsai and Honka, 2021). An example of such an awareness effect is the game "Among Us" by a small indie studio, which stayed dormant on Steam for almost two years and only became popular when consumers learned about it from Twitch broadcasts. Additionally, even if consumers are aware of a game, they might learn something about its price, quality, or gameplay by watching streamers play it live on Twitch. For example, while many consumers know that "Dark Souls" is a monster fighting game, they might realize after watching Twitch streams that this game stands out with its intricate level design and combat depth. In other words, live streams might serve as informative advertisements that reveal the horizontal and vertical attributes of broadcasted games (Grossman and Shapiro, 1984; Ackerberg, 2003).

If streams provide any information at all, they should have stronger effects on relatively unknown games. To proxy consumer knowledge, we assume that consumers are less informed about new games and games released by small publishers, who lack the reputation and marketing budgets of large publishers. We start by testing whether the streaming elasticity β negatively correlates with game age and publisher size. To understand what information consumers might learn after watching Twitch streams, we also study whether β varies across games with different vertical game attributes, such as price and quality. We measure quality using Metacritic ratings, a widely recognized quality metric analogous to Rotten Tomatoes ratings for movies. Streams may either directly reveal information about quality, or they can encourage consumers to research the game and learn its quality. A similar effect may arise with prices: although our informal observation suggests that streamers rarely talk about game prices during live streams, consumers might still learn the game's price while doing their own research after the stream. Both mechanisms, direct and indirect, would make consumers more likely to adopt inexpensive or high-quality games after seeing them on Twitch. Finally, we analyze whether streaming disproportionately benefits "niche" games that strongly appeal to some consumers despite the mediocre quality. We proxy "niche" games by using the standard deviation of customer ratings on Metacritic.

4.2 Estimates from median sample splits

We start by subsampling games and comparing the estimated streaming elasticities across subsamples. Table 6 presents the estimated streaming elasticities $\hat{\beta}$ from the main specification in (1)-(2) for different subsamples. Throughout this section, we simplify estimation by fixing the persistence parameter at the level estimated in Section 3 ($\hat{\delta} = 0.828$).

In rows 1-4 of Table 6, we estimate stream elasticities by game age and publisher size. To this end, we define new games as those released within 2.7 years prior to our sample period. We additionally define small publishers as those who only sell one game from our sample of 599 titles. Both definitions roughly correspond to the median splits of variables "game age" and "publisher size." We find that new games have a somewhat higher estimated elasticity, although the difference is small and not statistically significant (elasticities of 0.034 vs 0.032). Nevertheless, we find that the games produced by small publishers benefit from streaming almost three times more than the games of large publishers (elasticities of 0.052 vs 0.019). This result is consistent with the

Table 6: **Streaming elasticities by game characteristics**

		Log viewership stock V_{jt}		Number of observations	
		$\hat{\beta}$ Estimate	$\hat{\beta}$ S.E.	No. games	No. obs.
By game age:	New games (<2.7 years old)	0.034***	(0.004)	299	1,636,128
	Old games (≥2.7 years old)	0.032***	(0.008)	300	1,641,600
By publisher size:	Small publisher (1 game)	0.052***	(0.011)	283	1,548,576
	Large publisher (2+ games)	0.019***	(0.003)	316	1,729,152
By price:	Inexpensive (<$20)	0.044***	(0.010)	281	1,537,632
	Expensive (≥$20)	0.023***	(0.004)	318	1,740,096
By quality:	High quality (Metacritic score >80)	0.041***	(0.010)	219	1,198,368
	Low quality (Metacritic score ≤80)	0.028***	(0.004)	247	1,351,584
By rating variance:	Niche games (rating std. >2.4)	0.050***	(0.010)	212	1,160,064
	Mainstream games (rating std. ≤2.4)	0.008***	(0.002)	213	1,165,536

This table presents the estimates of streaming elasticity $\hat{\beta}$ for games with different characteristics, holding the persistence parameter δ at the level estimated in Table 4. All specifications include game-date, game-hour of day, and time fixed effects. Standard errors are clustered at the game-date level. *, **, and *** represent significance at the 10%, 5%, and 1% level.

information effect, considering that small publishers have modest advertising budgets and do not get the same media coverage as big conglomerates like EA Games and Ubisoft. Consumers might therefore be unaware of these publishers' games, and Twitch broadcasts might break this awareness barrier.

We further study what kind of information consumers acquire from Twitch broadcasts. Rows 5-8 of Table 6 show that Twitch streams are more effective for inexpensive games and for high-quality games. We define inexpensive games as having the regular price below the median level of $20. We estimate an elasticity of 0.044 for inexpensive games, twice larger than the elasticity for expensive games. Similarly, we obtain a higher estimated elasticity for high-quality games, defined as games whose Metacritic ratings are above the median. The estimated elasticity is about 50% larger for high-quality than low-quality games (elasticities of 0.041 vs 0.028). These results suggest that by watching live streams, consumers acquire information – directly or indirectly – about the vertical attributes of broadcasted games.

Twitch streams might also help consumers understand whether the game matches their idiosyncratic preferences. This mechanism might be especially relevant for low or medium-quality games, whose mediocre ratings reflect that not all consumers enjoy them. These might be "niche" games that appeal to some consumers but leave others indifferent. Using the standard deviation of user ratings to proxy "niche" games, we compare the estimated elasticity for games above and below the median value of this proxy. Rows 9-10 of Table 6 show that niche games have a much higher streaming elasticity of 0.050. On the other hand, Twitch broadcasts barely affect mainstream games

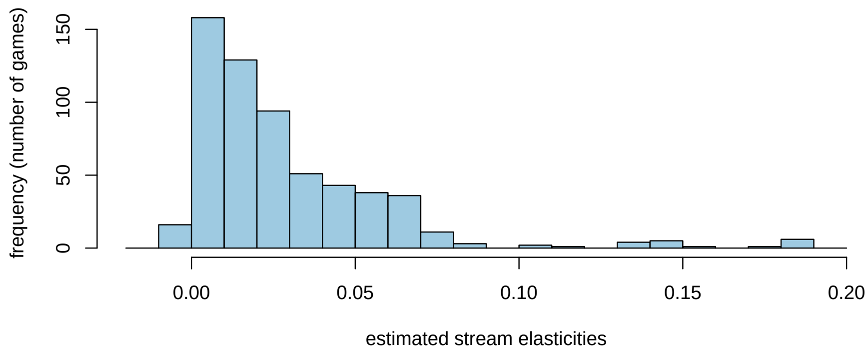

Figure 4: **Distribution of estimated streaming elasticities.** The graph shows the distribution of estimated elasticities $\hat{\beta}$ across games obtained from the generalized random forests.

with less dispersed user ratings (elasticity 0.008). This stark contrast supports our hypothesis that consumers learn about horizontal game attributes by watching Twitch broadcasts.

4.3 Estimates from generalized random forests

The median splits in Section 4.2 divide games based on somewhat arbitrary criteria. Additionally, splitting on each variable does not rule out the possibility that the effect is driven by other game attributes omitted from each pair of regressions. To address these concerns, we estimate heterogeneous streaming effects using generalized random forests, a nonparametric estimation method proposed by Athey et al. (2019). The method estimates heterogeneous treatment effects as a flexible function of observables using a set of local moment conditions. The local moment conditions are weighted across nearby observations, and the weights are adaptively computed using the random forest algorithm (Breiman, 2001). By using these local moment conditions, this approach effectively generalizes our instrumental variable strategy. The method also generates estimates with known asymptotic distributions, enabling us to construct confidence intervals for the estimates.

We generalize our main model in (1) by making the streaming elasticity β a function of game characteristics X_j:

$$\log\left(1+players_{jt}\right) = \beta(X_j)\cdot\log\left(1+V_{jt}\right)+\lambda_{j,d(t)}+\mu_{j,h(t)}+\eta_t+\varepsilon_{jt} \tag{6}$$

Because we fix the persistence parameter δ, the viewership stock V_{jt} is observed. The vector of game attributes, X_j, includes five variables: game age, publisher size, regular price, Metacritic rating, and the standard deviation of consumer ratings. These are the same five variables we used for the univariate median splits. Following Athey et al. (2019), we estimate the function $\beta(X_j)$

using local moment conditions stated in Appendix D.

Figure 4 visualizes the distribution of the estimated elasticities $\hat{\beta}(X_j)$. The average estimated elasticity is 0.029, somewhat lower than the average elasticity of 0.033 from our main specification in Table 4. We also find substantial heterogeneity: the estimated elasticities vary between -0.003 and 0.183 with the interquartile range from 0.008 to 0.040. Out of 599 games in our sample, 403 games (67%) have statistically significant estimates $\hat{\beta}$, suggesting that the game-specific streaming elasticities are precisely estimated. The estimated elasticities are positive for 97% of games, implying strong complementarity between Twitch viewership and game usage. Table F.2 in the Appendix compares the average attributes X_j between the games with highest and lowest estimated elasticities $\hat{\beta}$. In line with our results below, we find that games with the highest elasticities are on average newer, produced by smaller publishers, less expensive, and have more dispersed consumer ratings.

Figure 5 visualizes the estimated streaming elasticities. The three graphs on the left visualize the estimated function $\hat{\beta}(X_j)$ by focusing on two dimensions at a time while holding all other attributes X_j fixed at their average levels (see details in Appendix D.4). One might worry that some attribute combinations in these graphs are unrealistic because they never appear in the actual data. To address this concern, in the right panel of Figure 5 we plot the empirical density of the same game attributes, thus showing which estimated elasticities $\hat{\beta}(X_j)$ rely on extrapolation.

Panel A visualizes the estimates by game age and publisher size. We confirm the previous finding that games by small publishers benefit more from Twitch broadcasts. In contrast, we find more nuanced patterns for game age than those discussed above. Streaming elasticities $\hat{\beta}$ are highest for the newest games (less than one year old) and the oldest games (more than eight years old). Whereas the high elasticities among new games are aligned with our median-split results, one plausible explanation for the high elasticities among old games is that Twitch streams make consumers aware of the games forgotten by the current generation of players. Indeed, because Twitch viewers are between 16 and 24 years old, in 2021 many of them had not yet reached the legal age when these older games were launched.[16]

[16]This statistic is based on a 2019 survey. Source: https://www.statista.com/statistics/634057/twitch-user-age-worldwide/.

(A) By publisher size and game age

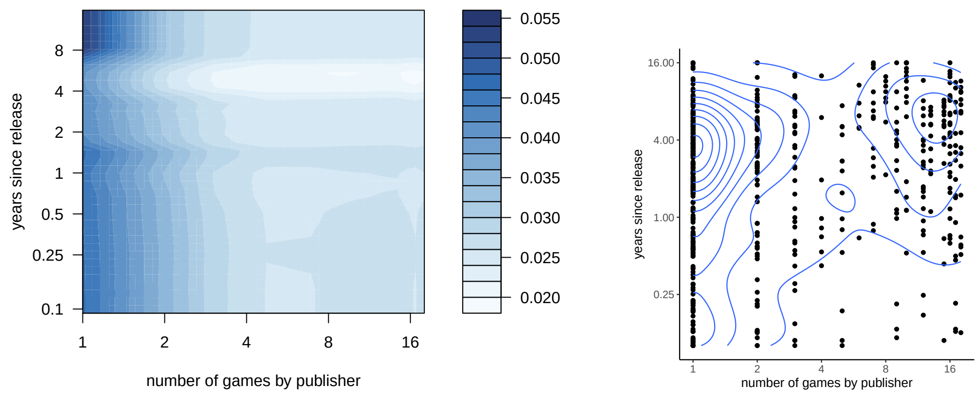

(B) By Metacritic rating and price

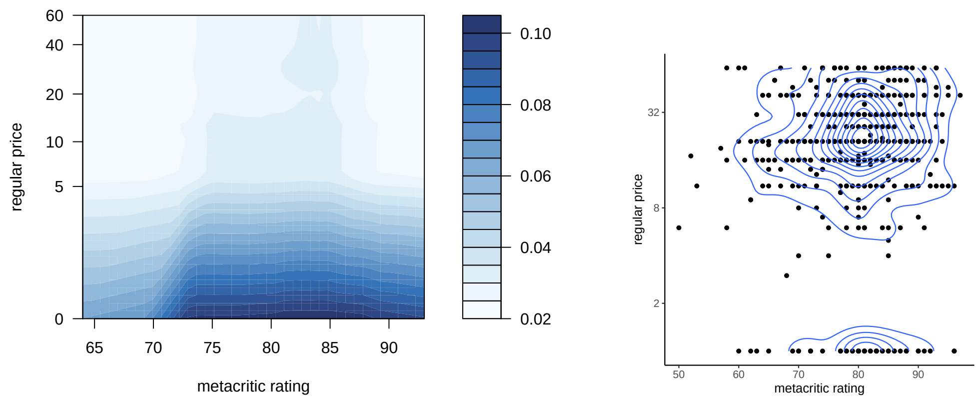

(C) By Metacritic rating and std.dev. of the user rating

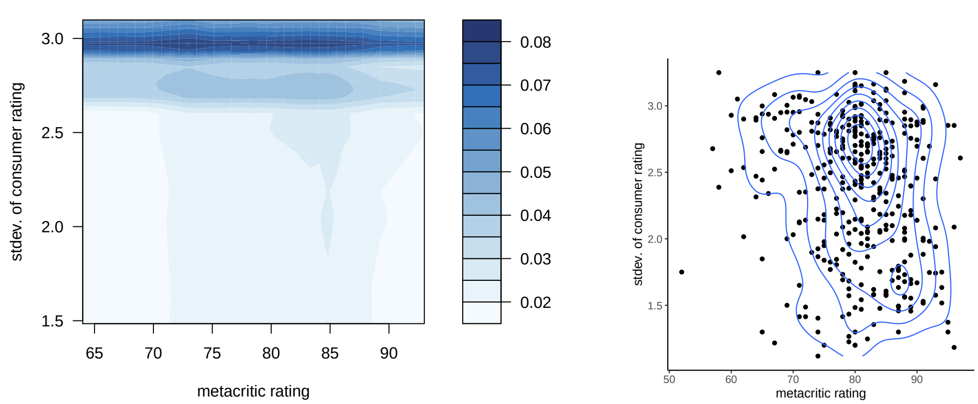

Figure 5: **Estimated streaming elasticities from Generalized Random Forests.** Each graph on the left visualizes the estimated function $\hat{\beta}(X_j)$ by focusing on two dimensions at a time, while holding all other attributes X_j fixed at their average levels (see details in Appendix D.4). Each figure on the right shows the empirical distribution of the same two game attributes and presents the contour lines of the estimated density function.

Panel B shows the estimated elasticities $\hat{\beta}$ by game price and rating. We find higher elasticities for inexpensive games, especially those priced under \$5, which is consistent with our earlier analysis. The density plot on the right suggests that this pattern is mostly driven by the high streaming elasticities of free games. As for the ratings, we find a nonlinear pattern whereby games with ratings around 80-85, slightly above the median, have the highest estimated elasticities. Somewhat counterintuitively, estimated elasticities drop as the rating approaches 90-95. We observe a similar pattern in Panel C. It is possible that games with high ratings already get extensive media coverage and therefore do not benefit from additional exposure on Twitch, although we cannot directly test this hypothesis with our data.

Panel C visualizes the estimated elasticities by Metacritic expert rating and the standard deviation of consumer ratings. Consistent with our previous results, we find that games with highly dispersed consumer ratings benefit more from Twitch streaming than games with more uniform consumer ratings. This result suggests Twitch streams reveal the gameplay of the broadcasted game, thus helping consumers understand whether the game matches their preferences.

When put together, these results suggest that despite the modest average streaming elasticity, a small fraction of games benefit considerably from streaming. For example, Twitch streams might be effective at promoting games by lesser-known publishers, informing consumers about appealing game attributes (e.g., low price), or promoting niche games that appeal to a small group of consumers. In fact, the promotional effects for some games might be sufficiently high for sponsored streams to generate positive ROI. We now demonstrate this point by calculating ROI separately for each game.

4.4 When is it profitable to sponsor live streams?

To calculate the implied returns on investment, we consider a counterfactual in which a publisher pays top streamer a fixed fee for broadcasting a game for one hour. Using the estimated streaming effects $\hat{\beta}(X_j)$ from equation (6), we predict how many new players will be brought into the game in this hour of streaming and how this increase will affect the expected profits of the sponsoring publisher.

We predict the increase in the expected profit as

$$\text{Profit lift}_j = \text{Conversion Rate} \times \Delta\text{Players}_j \times \text{Profit Margin}_j - \text{Streaming Costs} \tag{7}$$

where the first term captures the expected revenue lift and the second term is the fixed fee that the publisher pays for one hour of streaming. We calculate the lift in the number of players, $\Delta\text{players}_j$, by combining the game-specific streaming elasticity $\hat{\beta}(X_j)$ with the game-specific lift in the number of viewers generated by the broadcast of a top streamer. Since sponsored streams

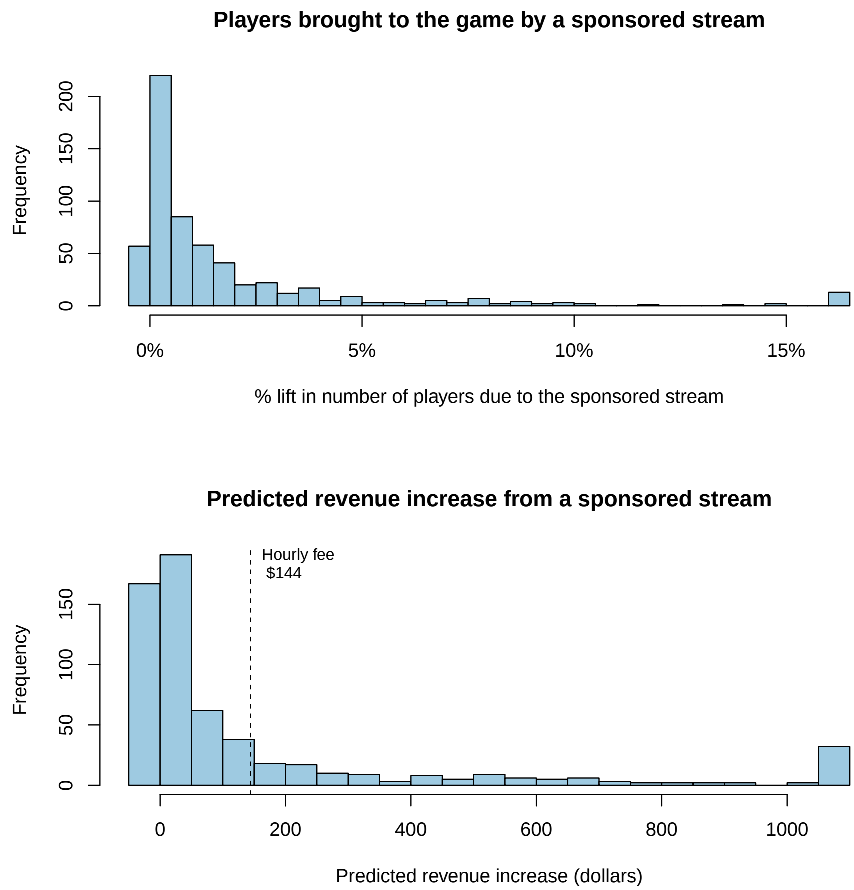

Figure 6: **Is It Profitable to Sponsor Live Streams?** The figure shows the predicted increase in the number of active players (top panel) and in the expected dollar revenues (bottom panel) from a one-hour live stream on Twitch. The predicted increase the number of active players in the top panel is computed using the formula (11), whereas the predicted revenue increase in the bottom panel corresponds to the left-hand side expression in the profit lift formula (7).

are less effective than organic ones, we discount the streaming elasticity by a factor of $0.233 = 0.007/0.030$, the ratio between the average sponsored and organic stream elasticities in Table 5. See Appendix E for details on how we compute $\Delta\text{players}_j$. To obtain the profit margin of paid games, we assume that a per-unit profit margin equals 70% of the game's price because Steam charges a 30% fee for publishing games on its platform. Although some games in our sample are free, they often bring comparable margins to their publishers via in-game transactions. To compute the profit margin for such games, we assume that their sales revenues are equal to the median sales revenue among the paid games.

The conversion rate in (7) reflects the share of players brought into the game by the stream who will buy a game copy. Because we do not have sales data to estimate the conversion rate, we need to make an assumption about its value. We consider an optimistic scenario by fixing the conversion rate at 100%, noting that this number is not far from the 79% conversion rate estimated by Li et al. (2021).[17] As we show below, even this optimistic scenario implies negative ROI for most games in the sample. Finally, we assume that to incentivize one additional hour of streaming, the publisher has to pay the equivalent of this streamer's average hourly wage on Twitch. We proxy hourly wage using data on individual revenues from subscriptions, which yields the estimated hourly wage of $144 (see Section 2.2 for details). This estimated hourly fee of $144 is close to the standard fee that publishers pay streamers on the internal Twitch's platform "Bounty Board" (see Figure F.5).

Figure 6 summarizes our results by plotting the distribution of the lift in the number of players (top panel) and the lift in the expected revenue (bottom panel) across games. In line with our estimates in Table 5, we find that for the median game, a one-hour sponsored stream increases the number of players by around 0.7% relative to the baseline. However, the predicted lift in the number of players differs dramatically across games, reaching 1.8-2% for at least a quarter of games. The bottom panel of Figure 6 additionally shows the implied increase in the expected revenues. Three quarters of games have a negative expected profit from sponsored streams because their predicted revenue lift falls short of the $144 hourly wage. Notably, most games obtain negative ROI despite the optimistic assumption of 100% conversion rate, which is unsurprising given that sponsored streams only marginally increase the short-term popularity of games.

Nevertheless, we predict substantial returns for a small set of games in the right tail of Figure 6 (bottom panel). At the 80th percentile of the profit distribution, we predict that sponsored streams increase the expected profit by $60, implying ROI of 42%. Around the 90th percentile, this predicted profit increase reaches $418, thus predicting extremely high ROI of 290%. Most games in the right tail of this distribution are games for which we estimated high estimated streaming elasticities in Section 4.3. Our estimates suggest that publishers of these games could extract substantial benefits from sponsoring additional live streams on Twitch.

4.5 Discussion and implications

The ROI calculations above have two main managerial implications. First, despite the negative average ROI, some game publishers may find it highly profitable to sponsor live streams on Twitch. Sponsored streams might be especially profitable for games that have high predicted streaming elasticities (see Section 4), including games released by small publishers, inexpensive games, and niche games that strongly appeal to some consumers. For these games, the ROI of sponsored streams might be as high as several hundred percent, which makes sponsored Twitch streams profitable for the publisher.

[17] Li et al. (2021) find that a release of a YouTube video about a game increases this game's purchase rate by 6.2% and usage by 7.8%. The ratio between the two effects implies a conversion rate of 79.4%.

Second, we find that most publishers benefit little from directly sponsoring Twitch streams. One practical implication is that, instead of directly sponsoring streams, publishers may want to encourage streaming of their games in some other way. For example, they can make their games more streaming-friendly, offer Twitch viewers in-game perks (e.g., unique items or characters), or create Twitch apps that allow viewers to participate in live streamed game sessions.[18] Interestingly, such alternative strategies have become more common in the last few years. A recent example was the developer Amazon Games that released its new online multiplayer game "Lost Ark" in February 2022. Using its corporate affiliation with Twitch, Amazon Games launched a special event letting Twitch streamers compete with each other in teams in exchange for in-game perks and special items (see Figure F.6). Over the new few weeks, the game became a huge hit and reached over one million active players, becoming the first most popular game on Steam. The game developer reported having to add additional 15 servers to deal with the influx of players (NME News, 2022). This remarkable success of "Lost Ark" may encourage other game developers to make their games more Twitch-friendly, which might, in turn, simulate streamers to generate additional organic content.

5 Concluding remarks

We use within-day variation in top streamers' broadcasting activities to estimate the effect of live streaming on game usage. Our results reveal that organic streams bring additional players into the broadcasted games, but the estimated effect is small and short-lived. We also find that organic streams are four times as effective at sponsored streams at bringing additional players to the broadcasted games. Exploring plausible mechanisms behind the streaming effects, we find evidence that live streams make consumers aware of lesser-known games from small publishers and help consumers learn about the game attributes. Although sponsored streams generate negative average ROI, they might generate much higher returns for a small set of games.

[18]Consistent with this direction, Twitch has recently released a so-called "Game Developer Playbook," which gives developers a long list of suggestions for how to make their games more Twitch-friendly (Twitch Official Blog, 2019).

References

ACKERBERG, D. A. (2003): "Advertising, learning, and consumer choice in experience good markets: an empirical examination," *International Economic Review*, 44, 1007–1040.

ALDRIDGE, T. (2019): "Twitch for Game Developers," *Twitch Official Blog*.

ALLEN, J. (2022): "Lost Ark helps Asmongold beat Twitch viewer records," *NME News*.

ATHEY, S., J. TIBSHIRANI, AND S. WAGER (2019): "Generalized random forests," *The Annals of Statistics*, 47, 1148–1178.

BAILIS, R. (2019): "The state of influencer marketing: 10 influencer marketing statistics to inform where you invest," *Big Commerce*.

BREIMAN, L. (2001): "Random forests," *Machine learning*, 45, 5–32.

CORREIA, S. (2016): "A feasible estimator for linear models with multi-way fixed effects," *Preprint at http://scorreia. com/research/hdfe. pdf*.

EFRON, B. AND R. J. TIBSHIRANI (1994): *An introduction to the bootstrap*, CRC press.

ERSHOV, D. AND M. MITCHELL (2020): "The effects of influencer advertising disclosure regulations: Evidence from instagram," in *Proceedings of the 21st ACM Conference on Economics and Computation*, 73–74.

FENLON, W. (2020): "How Among Us became so wildly popular," *PC Gamer*.

GAURE, S. (2013): "OLS with multiple high dimensional category variables," *Computational Statistics & Data Analysis*, 66, 8–18.

GEYSER, W. (2021): "The state of influencer marketing 2021: Benchmark report," *Influencer Marketing Hub*.

GOODMAN, L. (2021): "How much do Twitch streamers make?" *Stream Scheme*.

GORDON, B. R., F. ZETTELMEYER, N. BHARGAVA, AND D. CHAPSKY (2019): "A comparison of approaches to advertising measurement: Evidence from big field experiments at Facebook," *Marketing Science*, 38, 193–225.

GREENBAUM, A. (2020): "Nintendo continues to target streamers," *SVG News*.

GROSSMAN, G. M. AND C. SHAPIRO (1984): "Informative advertising with differentiated products," *The Review of Economic Studies*, 51, 63–81.

GUIMARAES, P. AND P. PORTUGAL (2010): "A simple feasible procedure to fit models with high-dimensional fixed effects," *The Stata Journal*, 10, 628–649.

HE, C. AND T. J. KLEIN (2019): "Advertising as a reminder: Evidence from the Dutch State Lottery," .

HONKA, E., A. HORTAÇSU, AND M. A. VITORINO (2017): “Advertising, consumer awareness, and choice: Evidence from the US banking industry,” *The RAND Journal of Economics*, 48, 611–646.

HWANG, S., X. LIU, AND K. SRINIVASAN (2021): “Voice Analytics of Online Influencers – Soft Selling in Branded Videos,” *Available at SSRN 3773825*.

JOHNSON, G. A., R. A. LEWIS, AND E. I. NUBBEMEYER (2017): “Ghost ads: Improving the economics of measuring online ad effectiveness,” *Journal of Marketing Research*, 54, 867–884.

JOHNSON, M. R. AND J. WOODCOCK (2019): “The impacts of live streaming and Twitch. tv on the video game industry,” *Media, Culture & Society*, 41, 670–688.

LI, I. (2022): “Sponsorship Disclosure in Livestreaming: The Behavior of Streamers on Twitch,” *Working Paper*.

LI, N., A. HAVIV, AND M. J. LOVETT (2021): “Digital Marketing and Intellectual Property Rights: Leveraging Events and Influencers,” *Available at SSRN 3884038*.

LIAUKONYTE, J., T. TEIXEIRA, AND K. C. WILBUR (2015): “Television advertising and online shopping,” *Marketing Science*, 34, 311–330.

LODISH, L. M., M. ABRAHAM, S. KALMENSON, J. LIVELSBERGER, B. LUBETKIN, B. RICHARDSON, AND M. E. STEVENS (1995): “How TV advertising works: A meta-analysis of 389 real world split cable TV advertising experiments,” *Journal of Marketing Research*, 32, 125–139.

LOVETT, M. J. AND R. STAELIN (2016): “The role of paid, earned, and owned media in building entertainment brands: Reminding, informing, and enhancing enjoyment,” *Marketing Science*, 35, 142–157.

RAJARAM, P. AND P. MANCHANDA (2020): “Video Influencers: Unboxing the Mystique,” *arXiv preprint arXiv:2012.12311*.

SCULLION, C. (2021): “The entirety of Twitch has reportedly been leaked,” *Video Games Chronicle*.

SEILER, S., S. YAO, AND W. WANG (2017): “Does online word of mouth increase demand?(and how?) evidence from a natural experiment,” *Marketing Science*, 36, 838–861.

SEILER, S., S. YAO, AND G. ZERVAS (2018): “Causal inference in word-of-mouth research: Methods and results,” *Working Paper*.

SHAPIRO, B., G. J. HITSCH, AND A. TUCHMAN (2020): “Generalizable and robust tv advertising effects,” .

SIMONOV, A., R. URSU, AND C. ZHENG (2020): “Do suspense and surprise drive entertainment demand? evidence from twitch. tv,” *Evidence from Twitch. tv (October 19, 2020)*.

TSAI, Y.-L. AND E. HONKA (2021): “Informational and Noninformational Advertising Content,” *Marketing Science*, 40, 1030–1058.

VINCENT, J. (2018): “Drake Drops in to Play Fortnite on Twitch and Breaks the Record for Most-Viewed Stream,” *The Verge*.

WILDE, T. (2020): “Should streamers pay game developers to stream their games?” *PC Gamer*.

YANG, J., J. ZHANG, AND Y. ZHANG (2021): “First Law of Motion: Influencer Video Advertising on TikTok,” *Available at SSRN 3815124*.

Online Appendix

A Industry and Data: Additional Details

A.1 Selection of streamers and games

During a preliminary three-week period on 6-26 April 2021, we sent repeated high-frequency requests to Twitch API to record which video games are streamed on Twitch most often and attract the largest audiences. To this end, each 10 minutes we requested from Twitch API the list of 200 games that had the largest total number of viewers on Twitch at the time of the request. Each of these games had a corresponding game ID assigned to it by Twitch. We selected all game IDs that appeared in our top-200 list at least once, which generated the list of 1010 game IDs. We then manually removed game IDs of non-Steam games (e.g., "Fortnite" and "League of Legends"), for which we have no way of tracking player counts; game IDs for non-gaming streams (e.g., "Just Chatting," "Music," "Sports"); and game IDs for which we could not access data on player counts from Steam API (23 games for which Steam does not report player counts). Only eight games from our sample had zero active players throughout the whole period, so these games are omitted from our analysis. This selection process yields a list 599 games which constitutes our main sample throughout the paper.

In the same preliminary period on 6-26 April 2021, we sent hourly requests to Twitch API collecting the lists of most popular active streamers. Each hour, we requested the list of 100,000 most-viewed streamers that were live on Twitch at the time of the request, and we recorded which game they streamed and how many viewers they had at that time. We then made a list of all streamers in these data, and we randomly sampled 60,000 streamers using weighted sampling in which weights were proportional to the total number of viewers the streamer attracted in three weeks. This sampling procedure allowed us to construct a comprehensive sample that covers both popular streamers and lesser-known streamers who often broadcast small "indie" games. We could not simultaneously track more than 60,000 streamers without exceeding the daily quota of Twitch API requests. We kept this list of pre-selected streamers constant throughout the main data collection period in May-December 2021.

A.2 Monetary incentives of Twitch streamers

One might wonder why streamers spend hours broadcasting video games on Twitch, considering that anyone can watch their broadcasts free of charge. As it turns out, however, live streaming video games can be a lucrative business for the top streamers and a steady source of extra income for many others. Streamers begin their careers as *Twitch Affiliates* who mostly earn income from the donations and subscriptions of viewers.[19] Most donations are direct PayPal transfers from viewers (streamer gets 100% of income) or "virtual bits" that viewers donate through the chat window (streamer gets 75% of income). In addition, viewers can subscribe to a channel for a $4.99, $9.99, or $24.99 monthly fee or use one free subscription that comes with Amazon Prime. Subscribing to a streamer gives viewers access to exclusive content, custom emotes, and sometimes to ad-free viewing of the streamer's content. The most popular streamers can make *Twitch Partner*, which requires them to stream at least three times a week and consistently average at least 75 concurrent viewers. While Twitch Affiliates retain only 50% of subscription income, Twitch Partners retain 60-70% and get access to exclusive contracts for promoting games and non-gaming products on their channels. According to survey data, a typical Twitch Affiliate earns $100-$1,000 per month, and a full-time streaming Partner earns $3,000-$6,500 per month (Goodman, 2021). Although most streamers do not publicly disclose their incomes, the payout data recently leaked from Twitch suggests that the most popular streamers earn as much as $115,000-370,000 per month (Scullion, 2021).

To estimate the cost of sponsoring live streams, we collected daily data on the number of active subscriptions from twitchtracker.com, a third-party website that reports the number of active subscriptions of Twitch streamers who chose to publicly disclose this information. We tracked on the daily basis the 10,000 most-subscribed streamers reported on Twitchtracker and collected the current number of active subscriptions and their breakdown by subscription type (i.e., Tier 1, Tier 2, Tier 3, or Amazon Prime). Using these subscription counts, we then constructed estimates of monthly income by computing the pre-tax income after the streamer has paid the commission to Twitch. One might worry that subscription revenues are not the primary source of income for streamers. Whereas we do not have systematic data on all income sources, anecdotal evidence suggests that subscription revenue is the most important income source for top streamers. Mediakix estimates that streamer Ninja makes three quarters of his income from subscription.[20] Additionally, streamers xQcOW, Auronplay, and Asmongold received total monthly payouts from Twitch of $325,170, $117,436, and $98,139 in subscription and advertising revenues, about two times as high as the subscription income we observe in our data.[21]

[19] Streamers can also run ads during their live streams, which earn them about $250 monthly per 100 subscribers.

[20] Source: https://influencermarketinghub.com/twitch-money-calculator/.

[21] These monthly figures are from the article "The entirety of Twitch has reportedly been leaked" which uses dataset of payouts leaked from Twitch in September 2021 (Video Games Chronicle, 2021).

B Additional Estimation Details

	Distribution of daily averages across games						
	Mean	S.E.	Q 5%	Q 25%	Q 50%	Q 75%	Q 95%
No. unique streamers	2.29	12.49	0.01	0.05	0.22	0.85	8.12
Max no. streamers live	1.80	5.33	1.00	1.00	1.01	1.14	4.30
Stream duration (hrs)	2.89	1.86	0.81	1.67	2.58	3.75	5.78

Appendix Table B.1: **Broadcasting activity of top streamers across different games.** These summary statistics illustrate the variation we isolate with instrumental variables z_{jt}, which measure how many top 5% Twitch streamers are broadcasting game j in time period t. We first compute the daily averages across all days in our sample and then report the distribution of these averages across 599 games.

B.1 Variation captured by instruments z_{jt}

In this section, we describe variation isolated by instruments $z_{j,t}, z_{j.t-1}, \ldots, z_{j,t-12}$ in Section 3.2. Table B.1 describes the variation in the number of active top 5% streamers captured by these instruments. In this table, we compute the daily numbers of unique streamers broadcasting a specific game, maximum number of streamers broadcasting simultaneously at any point in time, and average stream duration. We then report the distribution of these game-specific averages in Table B.1. As these statistics show, the average game is broadcasted by only 2-3 unique streamers on a given day and is broadcasted by no more than 1-2 streamers at the same time. Additionally, when top streamers do broadcast a game, those broadcasts are on average almost three hours long. This long average stream time suggests that most top streamers spend several hours exploring a given game and showcasing its gameplay rather than picking it up for a few minutes during breaks in their main activity. Lastly, we note a median game in this sample is never broadcasted by more than one streamer at any given time. Therefore, for a typical game, our instrument z_{ij} captures the variation between time periods when nobody broadcasts game j and other periods when one of the top streamers picks up a game and plays it for several hours.

B.2 Grid search algorithm for δ and β

To solve the minimization problem in (4), we use the following algorithm. We use a golden-section algorithm for the persistence parameter δ starting from a wide interval $(0.001, 0.999)$ and terminating search when the length of the search interval falls below the tolerance level 0.01. Given a candidate value of δ, we compute a point estimate for the parameter β using the standard closed form 2SLS formula. The main challenge is how to deal with three high-dimensional fixed effects (game-date, time interval, and game-hour of the day). Given that our sample includes approximately 600 games and 5,300 one-hour time intervals, we need to include approximately 150,000 million fixed effects. To handle the problem of this scale, we use packages *reghdfe* and *ivreghdfe* that rely on an iterative algorithm that was proposed by Guimaraes and Portugal (2010) and was further optimized by Correia (2016). This algorithm relies on a simple fixed-point iteration principle whereby all regression coefficients are partitioned into groups (e.g., by the class of fixed effects), and the algorithm iterates group-specific first order conditions while fixing the values of coefficients in all other groups. Guimaraes and Portugal (2010) show that this algorithm converges to the correct least squares estimates but manages memory more efficiently than the standard estimators.

By combining this algorithm with the golden-section search described above, we obtain point estimates of parameters β and δ. To compute standard errors clustered at the game-date level, we perform a block bootstrap that samples game-date pairs from the original sample with replacement and draws in total 50 bootstrap samples (Efron and Tibshirani, 1994, p. 86). For specifications in Sections 3-4 in which we do not estimate δ, we do not use bootstrap and instead obtain standard errors using the standard asymptotic theory of the 2SLS estimators.

C Robustness Analyses

C.1 Alternative specifications

We explore the robustness of our main results from Section 3 with respect to the (a) included fixed effects, (b) definition of the instrument Z_{jt}, (c) definition of the time period t, and (d) sample definition. Table C.1 presents the estimates $\hat{\beta}$ and $\hat{\delta}$ for different specifications. Rows 1-4 show that our results are robust to controlling for game-week fixed effects and changing the definition of the instrument Z_{jt}. Our initial motivation for including game-date fixed effects was that they allow us to better control for unobserved game-specific events. Nevertheless, the estimate $\hat{\beta}$ barely changes when we instead use game-week fixed effects (an estimate of 0.034 vs 0.033). This result suggests that, even if game popularity changes within a given week, streamers do not systematically schedule their live streams on days when a game is trending. Similarly, constructing the instrument Z_{jt} using a different number of lagged values z_{jt} does not seem to affect our estimates. Row 5 shows how the estimates change when we switch from 1-hour time to 10-minute time intervals t. While this switch makes the estimation computationally burdensome, it changes our results only marginally. We obtain the estimated elasticity of 0.032 and an implied hourly persistence parameter of $(0.956)^6 = 0.763$, which are reasonably similar to our main specification. Row 6 shows that dropping games that are never broadcasted by the top 5% streamers (i.e., zero variation in the instrument Z_{jt}) returns estimates 0.034 and 0.849, which, once again, are close to those in the main specification.

Finally, in row 7 we explore whether our functional form assumptions in (1), especially the log expressions $\log(1 + players_{jt})$ and $\log(1 + V_{jt})$, generate bias. When removing all games that have on average less than 10 concurrent viewers or less than 10 concurrent players, we obtain results similar to our main specification. This finding suggests that our estimates are unlikely to be mainly driven by the functional form assumptions. Note, however, that we find a somewhat higher estimate of the elasticity β (0.042 vs 0.033 in the main specification), likely because we are focusing on more popular games that are likely to be of higher quality.

Appendix Table C.1: **Robustness analyses for the main specification in Section 3.**

Specification	Elasticity $\hat{\beta}$	Persistence $\hat{\delta}$
(1) Main specification	0.033	0.828
(2) Main + game-week FEs	0.034	0.794
(3) Main + 6 lags in the instrument Z_{jt}	0.035	0.831
(4) Main + 18 lags in the instrument Z_{jt}	0.033	0.831
(5) Main + 10-minute intervals t	0.032	0.956
(6) Main + drop games without Z_{jt} variation	0.034	0.849
(7) Main + only non-zero $players_{jt}$ and $viewers_{jt}$	0.042	0.897

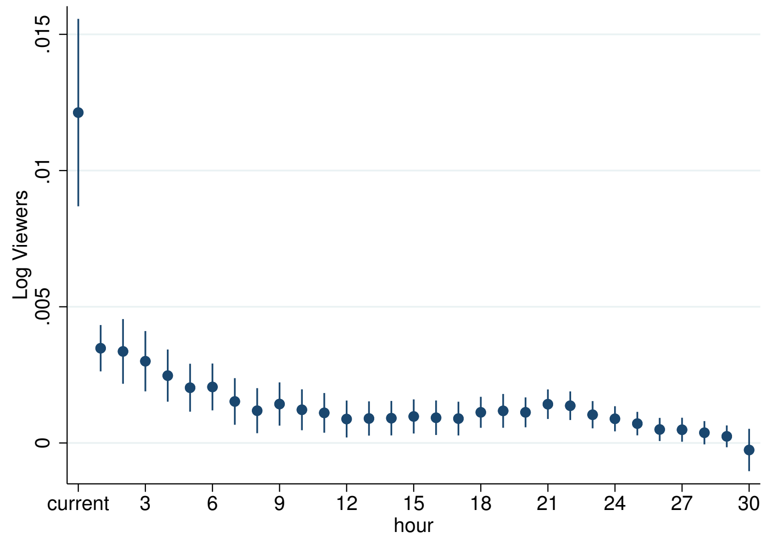

Appendix Figure C.1: **Distributed lag regression results.** This figure shows the coefficient estimates of model (8). Vertical bars around point estimates are the 95% confidence intervals.

C.2 Nonparametric estimation of lagged effects

Our geometric decay model in (1) imposes that the concurrent effect of streaming has the same sign and is stronger than the carry-over effect. However, live streams might distract viewers from playing the game, thus creating a negative concurrent effect and a positive carry-over effect. To explore this possibility, we now estimate a flexible distributed lag model that does not impose this assumption. Specifically, we estimate the following model:

$$\log\left(1 + players_{jt}\right) = \sum_{\tau=0}^{30} \beta_\tau \log\left(1 + viewers_{jt-\tau}\right) + \lambda_{j,d(t)} + \mu_{j,h(t)} + \eta_t + \varepsilon_{jt} \tag{8}$$

We include the same fixed effects as our main specification but model the number of concurrent players as a flexible function of lagged viewership counts. Although the model (8) does not nest our main model in (1) as a special case, it gives us an opportunity to examine the shape of the response function estimated in a more flexible specification. Figure C.1 visualizes the estimates. Consistent with our main specification, we find that both the concurrent effect and the carry-over effects are positive. In addition, the estimated effect quickly declines and dissipates almost to zero within about 10 hours. Therefore, we do not find evidence that our geometric decay in Section 3 imposes overly rigid assumptions on streaming effects.

D Implementation and Results from Generalized Random Forests (GRF)

D.1 GRF implementation details

To estimate a causal forest, we use the `grf` R package developed by Athey et al. (2019). The model we estimate using GRF is specified in equation (6) of Section 4.3. We estimate the heterogeneous streaming effects $\beta(X_j)$ using a set of local moment conditions:

$$\mathbb{E}\left[Z_{jt} \cdot \varepsilon(x) | \lambda_{j,d(t)}, \mu_{j,h(t)}, \eta_t, X_j = x\right] = 0 \tag{9}$$

where the conditional error term $\varepsilon_{jt}(x)$ is defined as $\varepsilon_{jt}(x) = \log(1 + \text{players}_{jt}) - \beta(x) \cdot \log(1 + V_{jt}) - \lambda_{j,d(t)} - \mu_{j,h(t)} - \eta_t$. Note that the model is additive and separable in game-date, game-hour of the day, and time fixed effects, Using this fact, we first demean the log player counts $Y_{jt} \equiv \log(1 + \text{players}_{jt})$, weighted viewership counts $W_{jt} \equiv \log(1 + V_{jt})$, and instruments Z_{jt} from these three sets of fixed effects. To demean these variables from three sets of high-dimensional fixed effects, we use the alternating projection method proposed by Gaure (2013) and implemented in the `lfe` R package. Applying the Frisch-Waugh-Lovell Theorem, we then estimate the function $\beta(x)$ by using the following local moment conditions:

$$\mathbb{E}\left[\tilde{Z}_{jt}\left(\tilde{Y}_{jt} - \beta(x) \times \tilde{W}_{jt}\right) | X_j = x\right] = 0 \tag{10}$$

where $\tilde{Y}_{jt}$, $\tilde{W}_{jt}$ and $\tilde{Z}_{jt}$ represent the demeaned values of variables Y_{jt}, W_{jt}, and Z_{jt}. Equation (10) then represents a standard set of local moment conditions in Athey et al. (2019), which allows us to estimate $\beta(x)$ using their `grf` R package. The `grf` package only permits one endogenous variable and one instrument. For this reason, we take the sum $Z_{jt} = \sum_{\tau=t-12}^{t} \delta^{t-\tau} z_{j\tau}$, where $z_{j\tau}$ is the number of top streamers broadcasting game j at time interval τ. With these new instruments, we get similar results to Table 6, where we use lags $Z_{jt} = \left\{z_{jt}, z_{jt-1}, \ldots, z_{jt-12}\right\}$ as separate instruments.

D.2 Tuning

The primary tuning parameter we focus on is the "leaf size" of the forest, defined as the minimum number of observations in each branch of each tree. Choosing a large leaf size stops the tree-splitting process prematurely and leads to heavily regularized estimates. On the other hand, allowing for smaller leaves leads to more flexible but potentially noisy estimates.

In our estimation, we split trees using the time-invariant attributes of each game, X_j. One should ideally include all observations of each game into the same leaf, but this is not guaranteed due to the random splitting of trees in random forest algorithms. We estimate the heterogeneous streaming elasticities using a range of tuning parameters that take values $n \times 5,471$, with $n = 1, 2, ..., 8$. We find that choosing a high leaf size regularizes the tail end of the $\hat{\beta}$ distribution, whereas choosing a low leaf size leads to noisier results in our validation exercise (see section D.3). We pick $n = 2$ to balance between the two considerations. We note, however, that our main qualitative findings change little when we move to other values of the tuning parameter.

D.3 Validation of GRF estimates

We validate the generalized random forest estimates in two ways. First, we compute the average elasticities that GRFs predict for games above and below the median of each game attribute X_j, and we compare these elasticities to our median split results from Section 4.2. To this end, we compute elasticities $\hat{\beta}_j$ predicted by the generalized random forests for the exact same subsamples. If the generalized random forests recover the true heterogeneity in streaming elasticities, the average predicted elasticities should be roughly the same as our median split 2SLS estimates in Section 4.2. Instead, if the forest does not capture meaningful heterogeneity, the average predicted elasticities need not be aligned with the 2SLS estimates. Table D.1 compares these two sets of estimates. We find that the conditional average elasticities $\hat{\beta}$ generated by the two methods have the same order of magnitude and follow analogous patterns. Based on this comparison, we do not find that GRFs overfit the data and produce an implausible large dispersion of predicted streaming effects.

We also compare the distributions of estimated elasticities obtained in two different ways. We first use GRF estimates to split games into four groups that correspond to the quartiles of the estimated elasticities $\hat{\beta}$, and we re-estimate the streaming effect by using a 2SLS regression with all observations in each quartile group. To obtain the second set of analogous estimates, we estimate a GRF on 90% of randomly selected games and then predict streaming elasticities $\hat{\beta}$ for the remaining 10% games. We then visualize the averages of these out-of-sample predictions for the same quartile groups as before, averaging across 10 iterations, each of which holds out a 10% subsample of games and estimates a GRF on the remaining data. Figure D.1 compares the two sets of estimates by reporting quartile averages and confidence intervals. We find that the out-of-sample elasticities predicted by GRF monotonically increase across the four groups, which are defined by in-sample GRF predicted elasticities. We also find that average out-of-sample elasticities in the four groups align with in-sample 2SLS elasticity estimates. These findings suggest that generalized random forests indeed recover meaningful heterogeneity in streaming effects across games.

Appendix Table D.1: **Comparison of estimated elasticities $\hat{\beta}_j$ from median splits and the GRF.**

		Median Splits		Generalized Random Forests	
		Estimate $\hat{\beta}$	S.E.	Average $\hat{\beta}_j$	Average S.E.
Game age:	New games ($\leq$2.5 years old)	0.034	(0.004)	0.030	(0.015)
	Old games (>2.5 years old)	0.032	(0.008)	0.028	(0.008)
Publisher size:	Small publisher (1 game)	0.053	(0.011)	0.038	(0.016)
	Large publisher (2+ games)	0.018	(0.003)	0.021	(0.008)
Price:	Inexpensive (<$20)	0.046	(0.010)	0.037	(0.017)
	Expensive ($\geq$$20)	0.023	(0.004)	0.021	(0.007)
Quality:	High quality (meta score >80)	0.041	(0.010)	0.024	(0.005)
	Low quality (meta score $\leq$80)	0.028	(0.004)	0.031	(0.014)
Variance:	Niche games (rating std. >2.4)	0.050	(0.010)	0.038	(0.009)
	Mainstream games (rating std. $\leq$2.4)	0.008	(0.002)	0.012	(0.006)

D.4 Visualizing the GRF results

In Figure 5 of Section 4.3, we visualize the estimated elasticities $\hat{\beta}(X_j)$ for two game attributes at a time, while holding all other attributes fixed at their average levels. To construct these graphs, we first generate a grid of values for each dimension in X_j with a step equal to a 5% quartile of the observed unconditional distribution. We take the Cartesian product of these uni-dimensional grids to create a five-dimensional grid in the space of five attributes in X_j. When visualizing estimated elasticities in Figure 5, we show the estimated elasticities $\hat{\beta}(X)$ for each point on a two-dimensional grid, for the two selected game attributes, and we average the predicted $\hat{\beta}$'s across all other dimensions of X_j. One can interpret these graphs as visualizing the conditional average treatment effects (CATE) for selected pairs of game attributes. One concern is that some combinations of attributes in these graphs are unrealistic and are never observed in the actual data. To this end, in the right panel of Figure 5 we additionally visualize the empirical density of game attributes, which helps to understand which areas of the estimated elasticities $\hat{\beta}(X)$ rely on the extrapolation of GRF estimates.

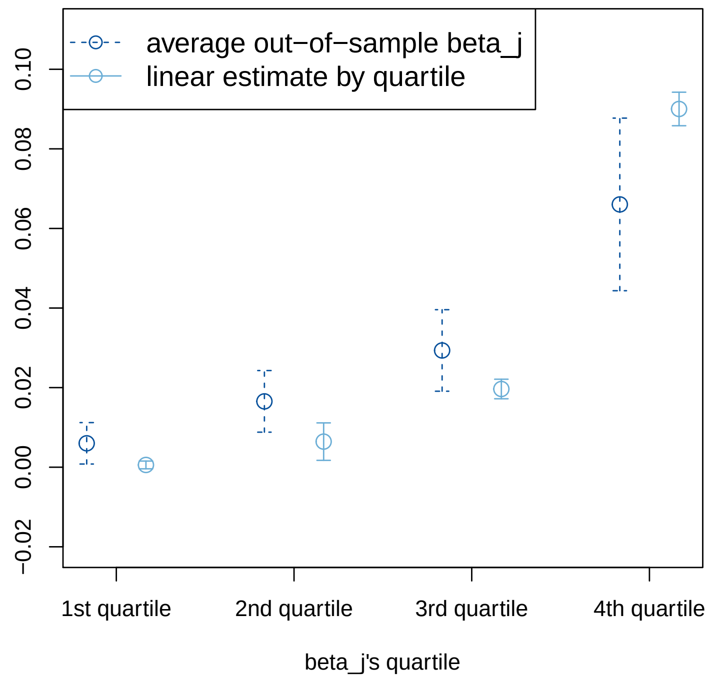

Appendix Figure D.1: **Out-of-sample validation of GRF results.** The graph splits games into four quartiles based on the estimated distribution of streaming elasticities $\hat{\beta}$. We compare two sets of estimates in each quartile. First, we estimate a GRF on 90% of randomly selected games and then predict streaming elasticities $\hat{\beta}$ for the remaining 10% games. We then use the GRF estimates to predict elasticities out-of-sample and present the mean predicted elasticities by group. Second, we estimate linear models using 2SLS, separate by each quartile, and present the in-sample average elasticities. We fix the leaf size at 2×5471 (see Section D.2).

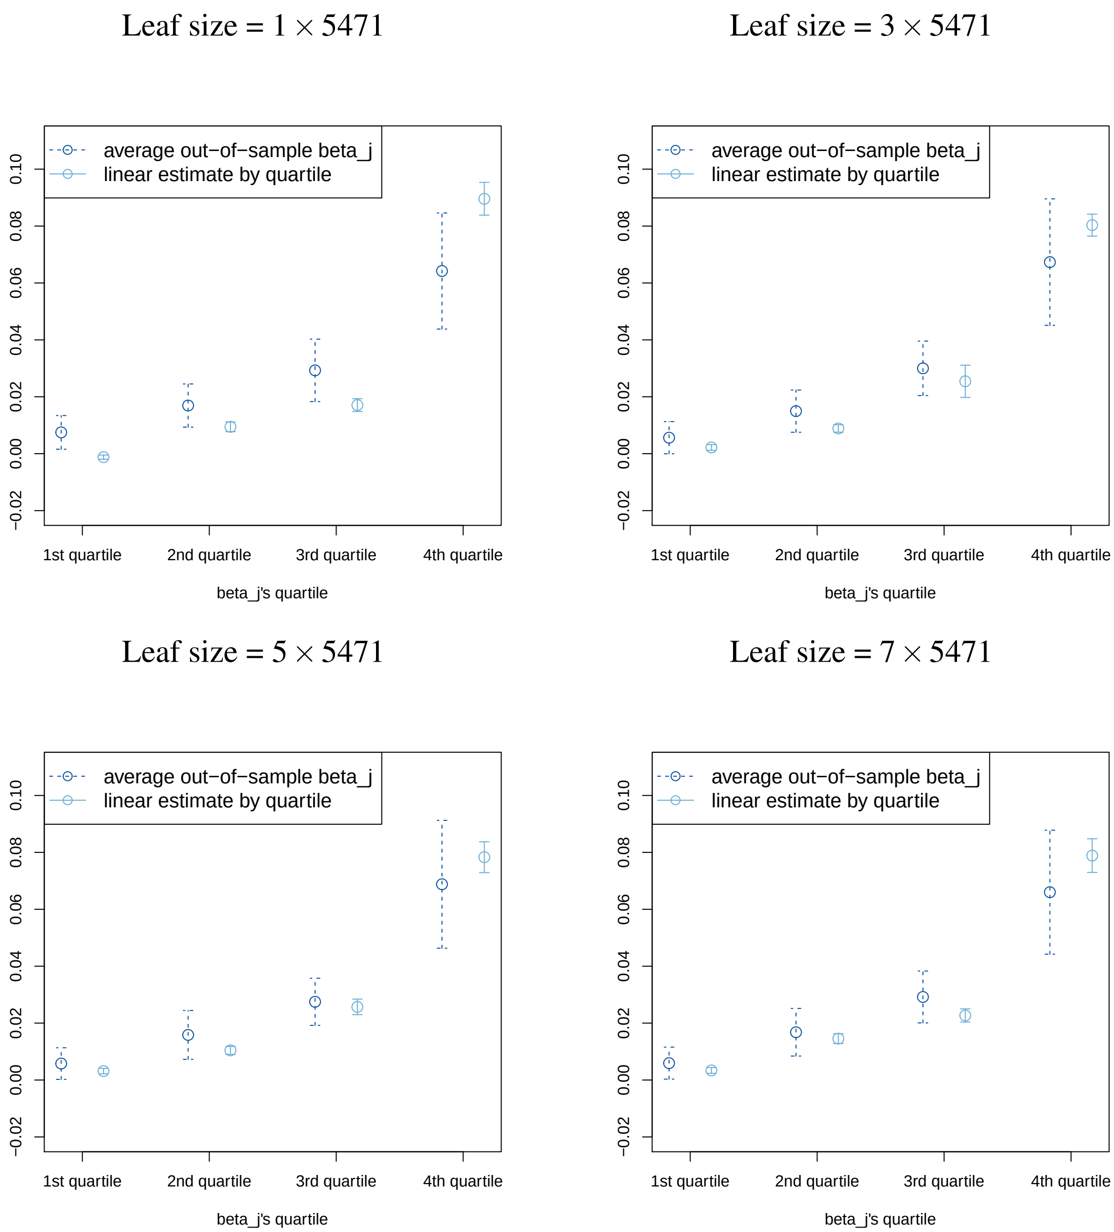

Appendix Figure D.2: **GRF validation results: alternative leaf-size parameters.** These four panels report the validation exercise in Figure D.1 under alternative leaf-size parameters $n \times 5{,}471$, where $n = 1, 3, 5, 7$.

E ROI Calculation Details

We compute the predicted lift in the number of players ($\Delta\text{players}_j$) in (7) as follows. We assume that prior to the sponsored broadcast, the number of viewers of game j, $viewers_{jt}$, is in a "steady-state" and is equal to the average number of viewers we observe in our sample. Therefore, the viewership stock V_{jt} is also constant over time prior to the sponsored broadcast. We fix the persistence parameter at $\delta = 0.828$ to simplify the analysis. When the broadcast starts, a top streamer goes live and increases the number of viewers who watch game j on Twitch. To predict the size of this increase, we regress the log number of viewers $\log(1 + viewers_{jt})$ on the number of active top streamers z_{jt} observed in our data, estimating the regression separately for each game to obtain game-specific viewership lifts. We then compute the predicted lift in the number of players, $\Delta\text{players}_j$, as

$$\Delta\text{players}_j = \beta_j^{spons} \times \left(\log\left(1 + V_{j1}\right) - \log\left(1 + V_{j0}\right)\right) \times \text{players}_j^{baseline}. \tag{11}$$

where V_{j0} is the observed "steady-state" viewership stock absent the sponsored stream, V_{j1} is the counterfactual viewership stock with the sponsored stream, and β_j^{spons} is the streaming elasticity of game usage with respect to the number of viewers brought into the game by sponsored streams.

F Additional Figures and Tables

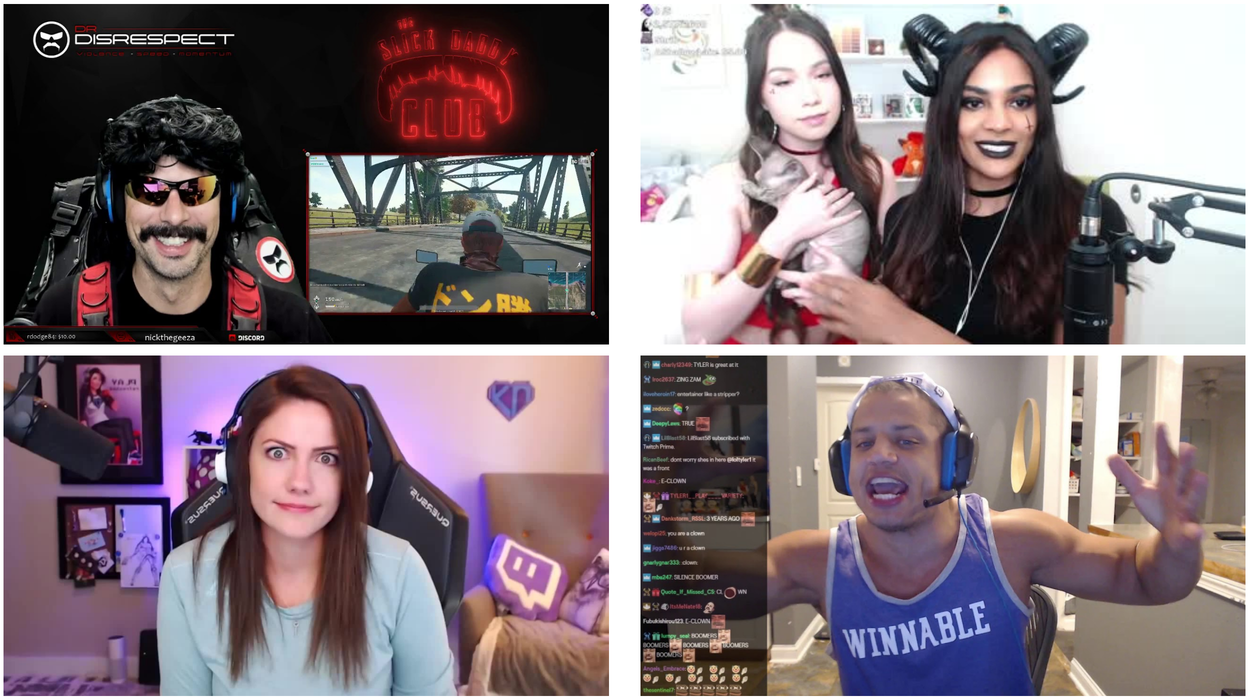

Appendix Figure F.1: **Many popular Twitch streamers have exuberant and memorable personalities, which makes their content funny and entertaining.** The figure shows screenshots from broadcasts of Twitch streamers Dr. Disrespekt (top left), Sydeon (top right), KayPea (bottom left), and Tyler1 (bottom right).

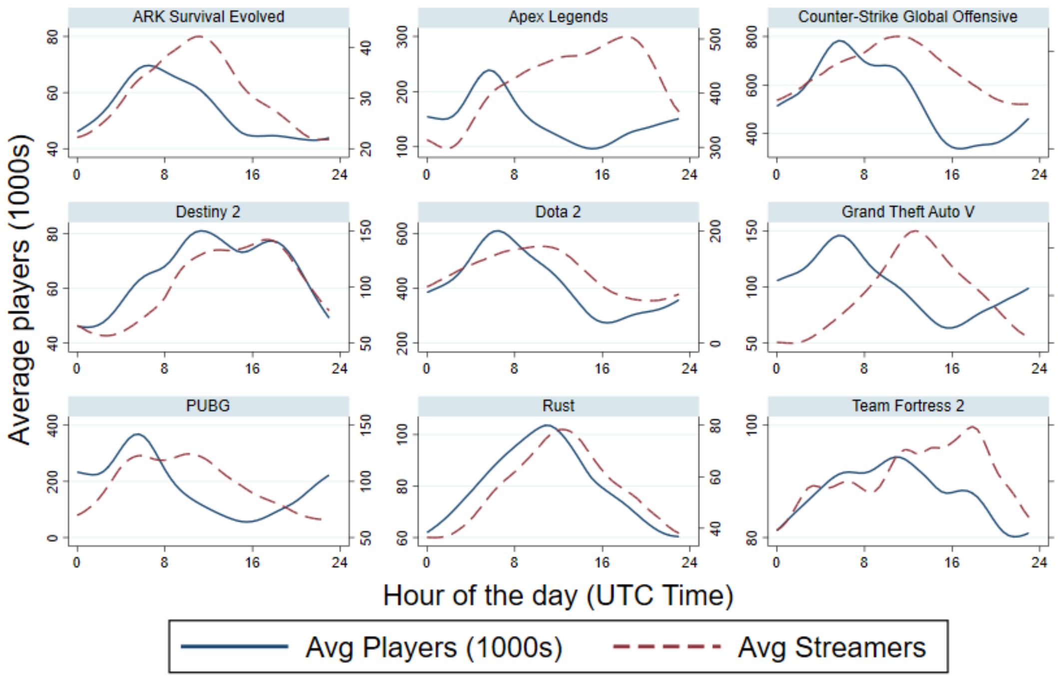

Appendix Figure F.2: **When are games played and streamed?**

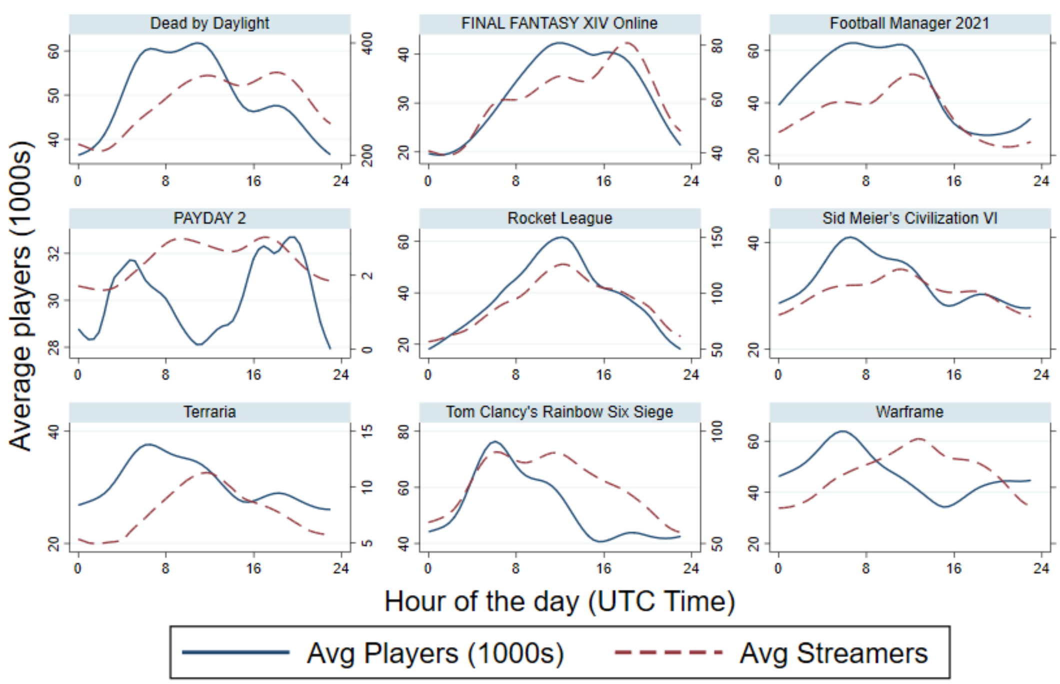

Appendix Figure F.3: **When are games played and streamed? (continued)**

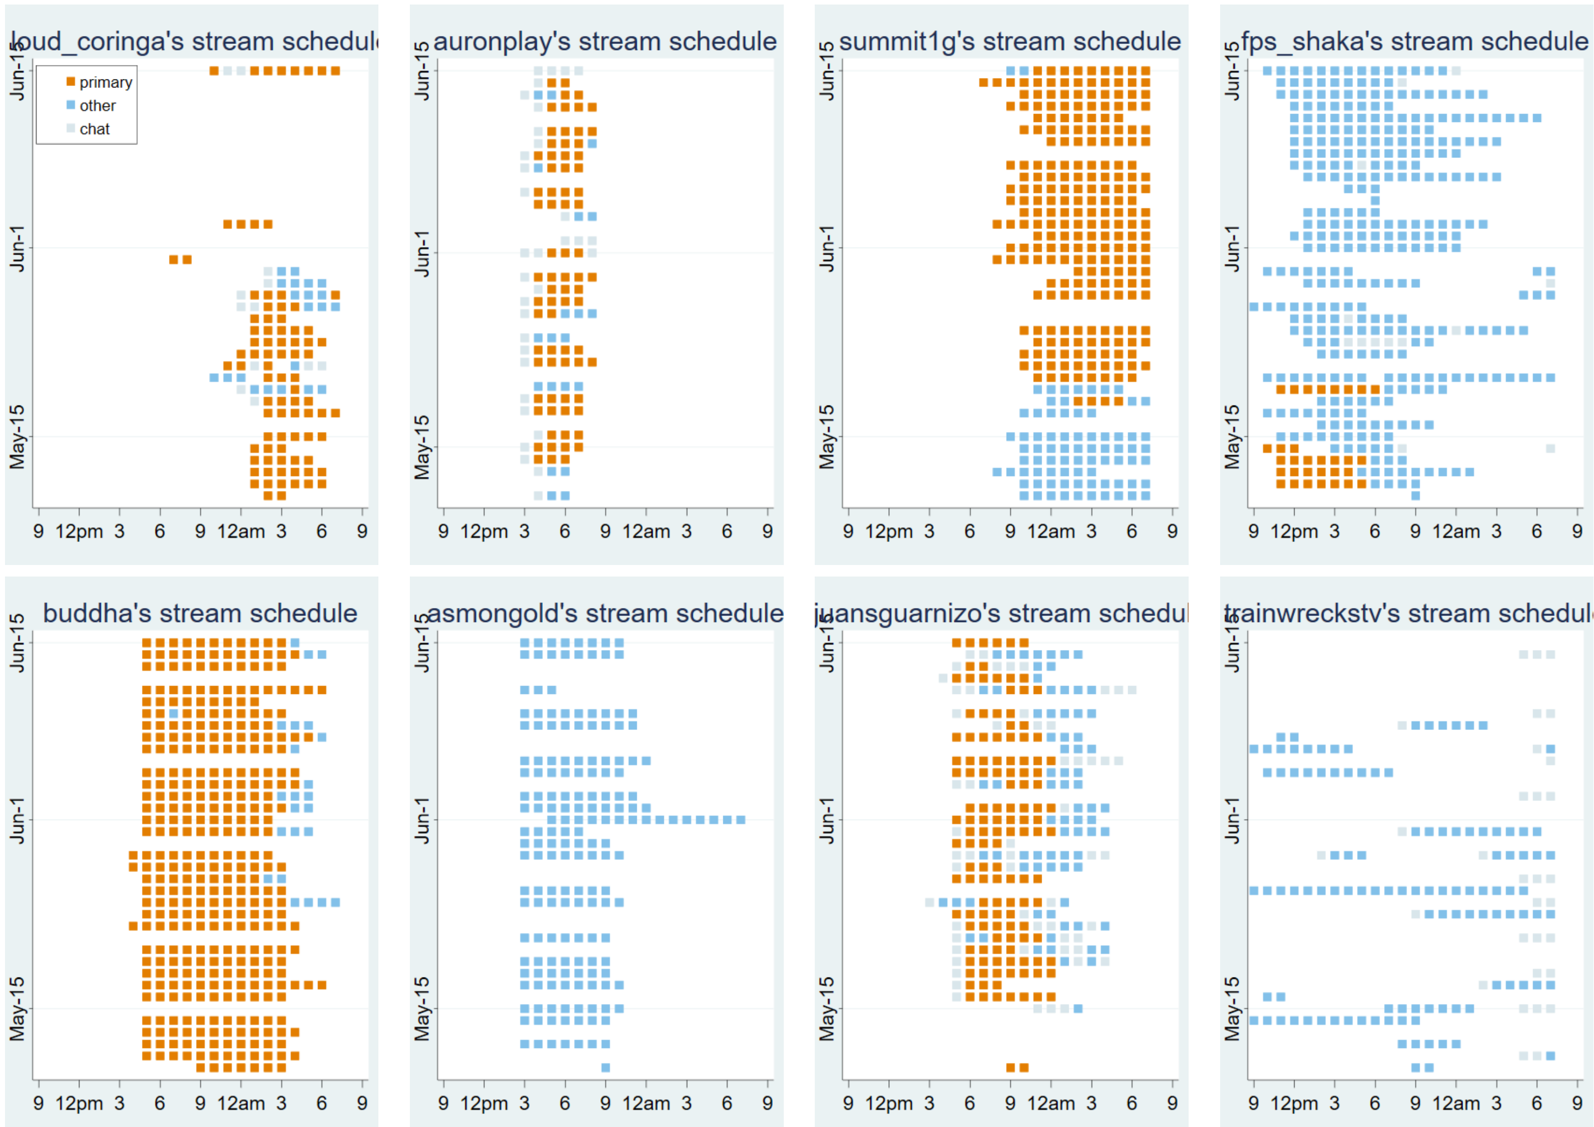

Appendix Figure F.4: **Exact stream schedules for a random set of top streamers.**

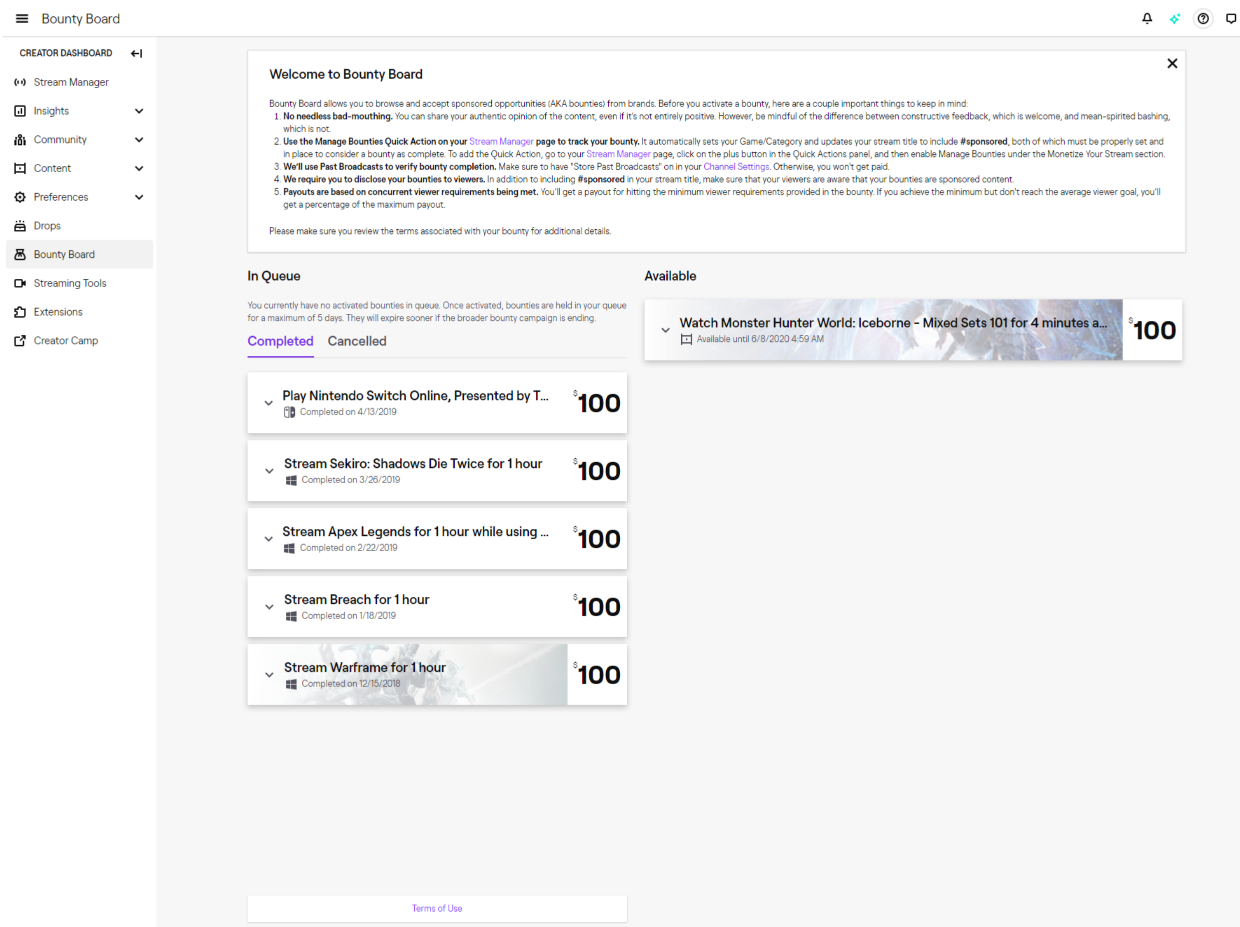

Appendix Figure F.5: **A screenshot of the streamer's view of the Twitch Bounty Board.** Source: https://help.twitch.tv/s/article/bounty-board-program-information-and-faq.

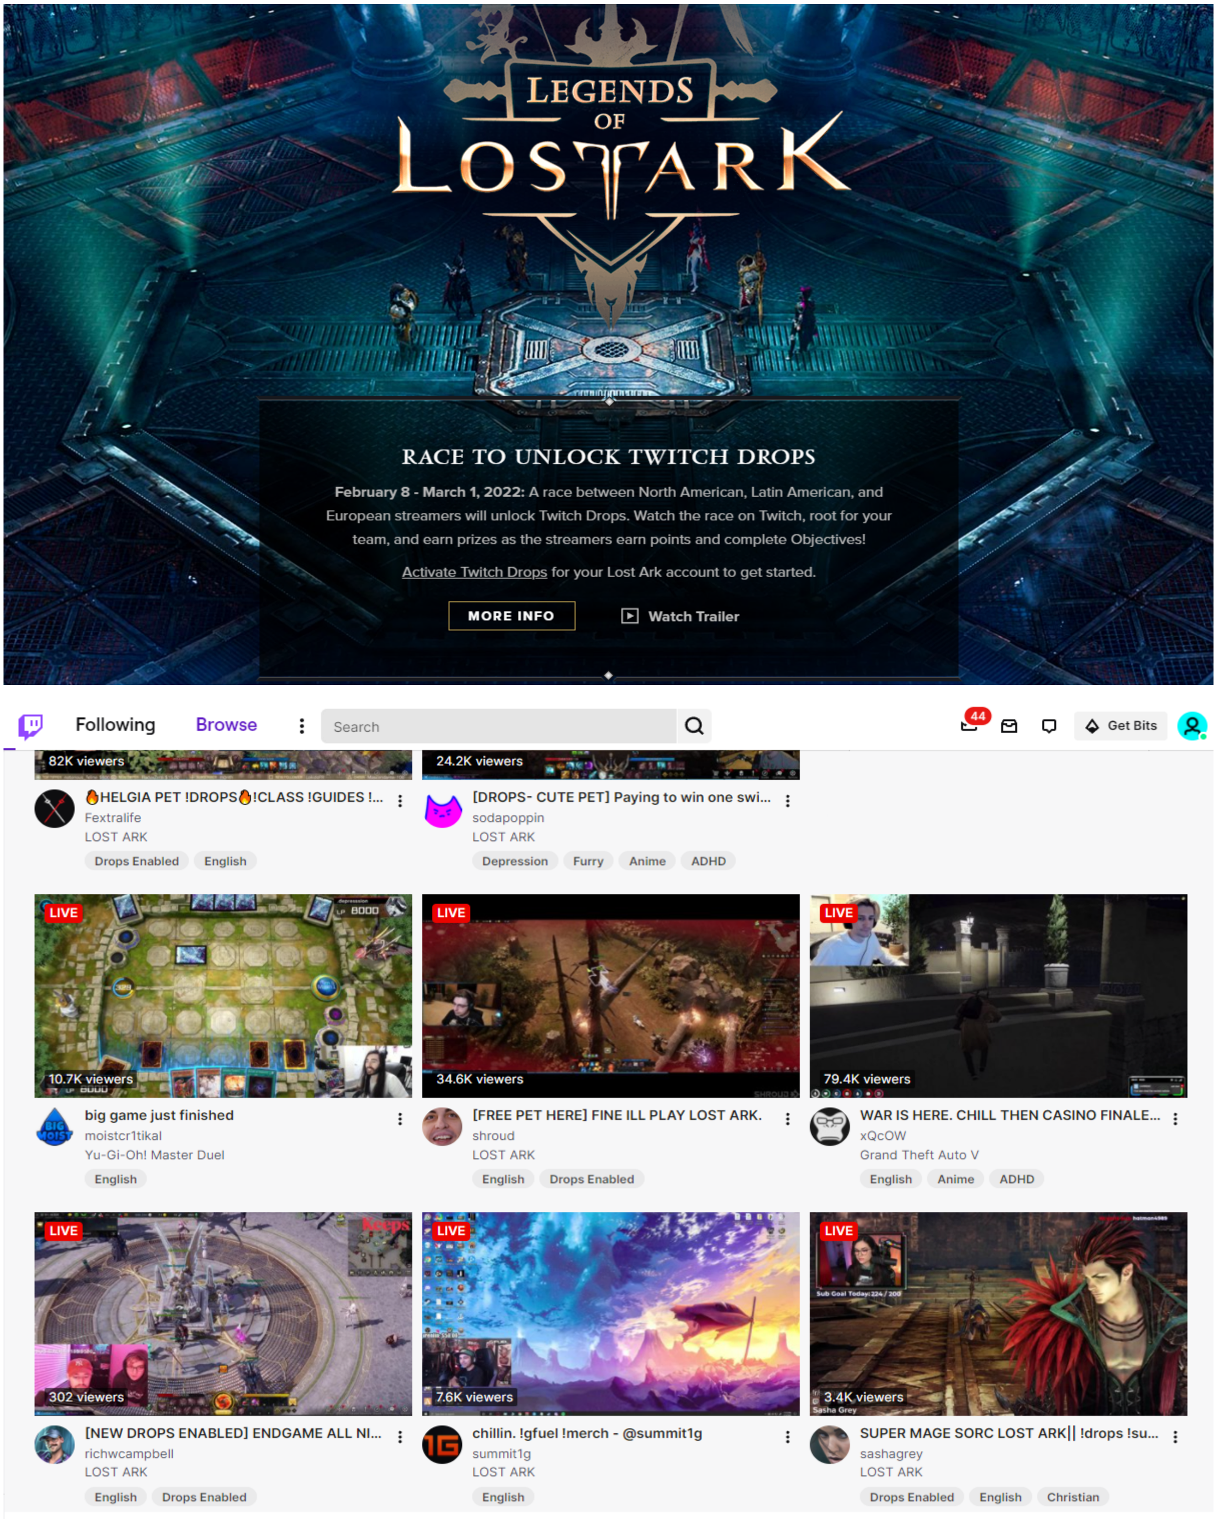

Appendix Figure F.6: **The promotion campaign of Amazon Games for the new game "Lost Ark".** The top panel shows a screenshot from the official website of the game (lostark.com), which describes in-game perks that streamers can earn by broadcasting the game on Twitch. The bottom panel is screenshot that shows the most popular live streams on Twitch on February 15, 2022, a few days after the release of the game. Most major streamers on that day were playing "Lost Ark" live to take advantage of the promotion campaign.

Appendix Table F.1: **Examples of games by age, rating, and developer size.**

	Value	No. games	Share games	Most streamed games (by total hours streamed)
By game age:	<1 years	173	28.8%	Apex Legends, FIFA 21, Sea of Thieves, Phasmophobia
	1-3 years	138	23.0%	Destiny 2, Day Z, Hunt Showdown, Red Dead Redemption 2
	3-5 years	107	17.9%	Dead by Daylight, PUBG, Rust, Black Desert
	5+ years	181	30.2%	Grand Theft Auto V, Counter-Strike Global Offensive, Dota 2, Rocket League
By developer size:	1 game	283	47.2%	Dead by Daylight, Destiny 2, PUBG, Rocket League
	2-5 games	115	19.2%	Apex Legends, FIFA 21, Rainbow Six Siege, Sea of Thieves
	6-10 games	135	22.5%	Old School RuneScape, World of Tanks, DayZ, SMITE
	11+ games	66	11.0%	Grand Theft Auto V, Counter-Strike Global Offensive, Dota 2, Red Dead Redemption 2
By Metacritic score:	0-75	130	21.7%	Dead by Daylight, FIFA 21, Sea of Thieves, Rust
	75-80	117	19.5%	Tom Clancy's Rainbow Six Siege, Phasmophobia, Old School RuneScape, World of Tanks
	80-85	116	19.4%	Counter-Strike Global Offensive, Destiny 2, Final Fantasy XIV Online, Hunt Showdown
	85-100	103	17.2%	Apex Legends, Grand Theft Auto V, Dota 2, PUBG
	no-rating	133	22.2%	VRChat, Marbles on Stream, Eternal Return
	Total:	599	100%	Apex Legends, Grand Theft Auto V, Dead by Daylight, Counter-Strike Global Offensive, Dota 2

The table is based on the sample of 599 Steam games we use for our main analysis. We define game age as the number of years since its official release date. The developer size is defined as the number of games a given developer released among the 599 titles in our sample. The variable game age captures the number of years passed between the official release of the game and the beginning of our data collection on May 11, 2021. The regular price corresponds to the 95-th percentile of the distribution of daily prices we observe for a given game, which usually captures the non-discounted price.

Appendix Table F.2: **Average attributes of games with the highest and lowest estimated streaming elasticities.**

	Estimated streaming elasticities $\hat{\beta}$		Welch t-statistic
	Highest 10%	Lowest 10%	(H_0: equal mean)
Game age (in years)	4.4	6.7	-3.60^{***}
Publisher size (# games)	3.1	10.4	-7.84^{***}
Regular price (dollars)	6.6	24.4	-5.00^{***}
Metacritic rating	77.2	80.9	-2.01^{**}
Std. consumer rating	2.8	1.9	7.42^{***}

This table describes the generalized random forest results by reporting the average characteristics X_j of the 60 games with the highest estimated elasticities $\hat{\beta}$ and the 60 games with the lowest estimated elasticities $\hat{\beta}$. The last column presents the Welch t-statistic (null hypothesis: the two samples have equal mean). ** and *** indicate that the mean of the two samples are statistically significantly different at the 5% and 1% level.