Congestion vs. Content Provision in a Live-Streaming Video Platform: Trade-offs between Prioritization and Neutrality

José F. Tudón M.*

May 6, 2019

Abstract

This paper studies the trade-off between entry and congestion in Amazon's Twitch.tv, which prioritizes popular content using a fast lane. I specify and estimate supply and demand models for live video, and a congestion model. Using technological shocks, I identify congestion costs for content providers and their consumers. Using shocks in prioritization, I identify its benefits. With estimated preferences and technological parameters, I construct counterfactuals. Without congestion, demand potentially doubles. A supply-side Pigouvian tax is preferred to a demand-side one. Without prioritization, consumer welfare drops up to 10%. I consider a rent-extractive platform, and discuss parallels with net-neutrality policy.

1 Introduction

In many telecommunication networks, content provision generates congestion externalities. During periods of congestion, network managers improve the perfor-

*ITAM, Río Hondo 1, Ciudad de México 01080, Mexico (e-mail: jtudon@itam.mx). Thanks to Ali Hortaçsu, Michael Dinerstein, Dennis Carlton, Brent Hickman, and Pietro Tebaldi for their guidance and support. Financial support from the NET Institute (www.netinst.org) and from the Becker Friedman Institute is gratefully acknowledged. I am also grateful to the members of the IO and

mance of private networks by prioritizing time-sensitive data. However, like internet bandwidth, priority is a scarce resource. This paper studies the impact of prioritization on congestion, content provision, and consumer welfare. I focus on the trade-off between entry and congestion, and use data on a popular live-streaming video platform to identify short-run supply and demand elasticities with respect to congestion.

Quantifying congestion externalities and their effects on supply and demand remains an open question not only in economics, but also in the engineering sciences (Malone, Nevo and Williams, 2017; Sundaresan et al., 2017). Moreover, in the presence of a two-sided market structure, a Pigouvian tax could have different outcomes depending on which side of the market is levied (Rochet and Tirole, 2006). Furthermore, platforms might be interested in extracting rents from prioritized service, or in prioritizing their own products and services. By decreasing content provision, such discriminatory policies might increase consumer welfare if congestion externalities are large, even as these policies break net-neutrality rules.

To address these questions, this paper studies the world's leader in live-streaming content of any kind, Twitch.tv. Bought by Amazon.com in 2014 for nearly $1 billion, Twitch is an internet platform where people broadcast live video and where people watch the broadcasts of other people. In February 2014, Twitch was the fourth-largest source of internet traffic during peak times in the U.S., at 1.8% of total traffic, behind Netflix, Google, and Apple but ahead of Hulu and Facebook.[1] Broadcasters streamed from 170 countries, spoke 43 different languages, and spanned 37 time zones. Currently, Twitch has over 1.7 million unique broadcasters per month, and over 100 million unique viewers per month, consuming, on average, 106 minutes of video a day.[2]

The data feature the global state of Twitch each 10 minutes during 90 days of 2014. Importantly, the data, described in section 2, include viewership, video bitrates, and whether the stream has access to a transcoder, which is a form of a fast lane that allows viewers with low-speed internet connections to watch high-bitrate videos. The high-frequency of the data allows me to construct a volatility measure

applied micro working groups at the University of Chicago. Thanks to the Twitch users and its community who provided me with valuable insights into the platform.

[1]See *The Wall Street Journal*, `www.wsj.com`. In this paper, URLs are abbreviated in print, but their hyperlinks are complete.

[2]See `www.twitch.tv/p/about`.

of the bitrate. The demand for live videos is sensitive to such volatility, which increases with congestion and manifests as buffering, stuttering, and other anomalies.

In section 3.1, I estimate a technological equation in which the volatility of individual streams is a function of real-time aggregate traffic and idiosyncratic shocks. I identify the elasticity of volatility with respect to traffic by exploiting a software upgrade that exogenously increased aggregate traffic. In section 3.2, I estimate a static, discrete-choice model of demand in which viewers choose their preferred channel, given the individual channels' volatilities and their access to the fast lane. Upgrades in the platform's infrastructure provide exogenous variation to identify the demand elasticity for volatility. To identify the benefits of the fast lane, I exploit the high frequency of the data to recover unexpected prioritization allocations. In section 3.3, I estimate a binary-choice model of supply in which content providers decide to be online as a function of their individual audience and of other suppliers' decisions. The interaction between entry and congestion can be effectively studied, because I observe content providers turning on and off with high frequency.

The estimates, presented in section 4, suggest that if traffic increases by 1%, the volatility of every channel increases by 0.8%. In addition, they suggest the elasticity of demand with respect to volatility is -2.5; and that cross-side network externalities exist, but aggregate supply is more elastic to aggregate demand than aggregate demand is to aggregate supply.

The structural estimation allows me to quantify the equilibrium effects of congestion and to consider counterfactual prioritization in section 5. Congestion externalities are substantial; due to cross-side network effects, the platform could be twice as large if not for congestion. A small Pigouvian tax on the demand side decreases consumer surplus, whereas one on the supply side does not. If the platform allocates prioritization randomly, as opposed to prioritizing the most popular broadcasters, consumer welfare drops up to 10% in the worst-case scenario.

Finally, in section 5.1, I discuss the parallels with the net-neutrality debate (Greenstein, Peitz and Valletti, 2016). I consider a rent-extractive monopolistic platform with all the bargaining power, which charges content providers just enough to leave them indifferent between having priority and not. Net neutrality prohibits these charges as well as prioritization. Content provision increases with net neutrality. However, even in this worst-case scenario with a monopolistic platform, consumer surplus decreases with net neutrality when prioritization is scarce. Moreover, neutrality erodes the platform's incentives to expand its fast-lane infrastructure, whereas the monopolistic platform retains those incentives when prioritization is scarce. Intuitively, in the presence of congestion externalities, oversupply is possible, especially because prioritization incentivizes individual supply. Thus, when fast lanes are scarce, efficiency requires them to be allocated to the most popular content providers. The results depend on the empirical finding that consumers value quality over variety.

This paper is related to a literature that studies congestion externalities (Duranton and Turner, 2011), especially in telecommunications. Malone, Nevo and Williams (2017) study congestion in broadband networks, and find that peak-use pricing, combined with local-cache technology, effectively reduces congestion. However, they assume congestion follows a day-to-day first-order Markov process. My contribution is to relax the assumption that congestion is independent of aggregate traffic. I rely on high-frequency data to estimate a 10-minute-window congestion model that depends on the equilibrium level of traffic, which in turn responds to congestion.

This paper is also related to a literature studying discriminatory platforms and two-sided markets. Zhu and Liu (2018) analyze Amazon's entry into third-party sellers' product spaces; Weeds (2016) studies the pay-TV case; and Genakos, Kühn and Van Reenen (2018) explore Microsoft's incentives to reduce interoperability. Also related is Lee (2013), who studies the US video-gaming industry and estimates cross-side network effects.

Finally, this paper also contributes to the net-neutrality literature. Nurski (2012) uses UK data on consumers and internet service providers (ISPs) to estimate the demand for internet connections. Her preliminary results suggest a fast lane to a dominant content provider would increase the welfare of consumers and the profits of both the ISP offering the fast lane and the dominant content provider. This paper further models congestion, as well as content provision as a function of the fast lane.

2 Empirical context and data

Twitch specializes in broadcasting live video-game sessions and content, but other categories include entertainment, social, news, sports, animals, creativity, and poker. In 2014, Twitch already dominated the video-gaming industry, broadly defined, with 43% of the market share. In 2019, Twitch remains the largely uncontested leader in live-streaming content of any kind and, as *The Economist* reports, "e-sports, in which computer gamers compete before thousands of fans in person and millions more online, is on the rise," attracting a global audience of almost 400 million in an industry worth $700 million annually.[3]

The data used in this research come from two main sources: Pires and Simon (2015) and my own calls to the Twitch REST API.[4] The data consist of the global state of Twitch channels roughly every 10 minutes from January 6, 2014, to April 6, 2014. The dataset cointains more than 1 million unique broadcasters. Each panel follows a broadcaster and contains channel ids, session ids, time stamps, real-time number of viewers, video bitrates, creation date, country of origin, language, time zone, and other information. The data do not contain the actual content of the stream. The following platform's description applies to January-April 2014.

Twitch broadcasters manage one channel, registered to them when they join the platform. The platform does not charge for signing up or requesting a channel. A channel can either be online at a given time, which means the uploader is broadcasting a live video, or offline. No videos, live or otherwise, are available when the channel is offline. A channel being online corresponds to a session. The number of viewers changes during the session, and the broadcaster observes the changes in real time. Watching a channel is also free, except for the time spent watching ads, which amounts to the main source of revenue for the platform.

Viewers decide which channel to watch by searching and clicking a channel. Some channels are featured on the homepage of the website, which makes them salient. Viewers can change channels, or turn off, at any time.

Besides deciding what to broadcast, the broadcaster also decides when to stream, how long, and at what average speed, measured in bits per second. Because the majority of channels stream video-game content, a chess match is a useful analogy. Rather than streaming for a fixed amount of time or a set schedule as in a soccer match, broadcasters stream, say, one or two chess matches, which take an unknown amount of time. If broadcasters expect a high viewership, they start their streams.

[3]See `superdataresearch.com`; and see *The Economist*, (a) `economist.com`, (b) `economist.com`.

[4]Pires and Simon (2015) data can be found in `dash.ipv6.enstb.fr/dataset/live-sessions/`. The do files that I used to parse it, complement it, and analyze it are available upon request. REST API stands for REpresentational State Transfer Application Programming Interface. It is a communication method between the programmer and the platform and allows data requests.

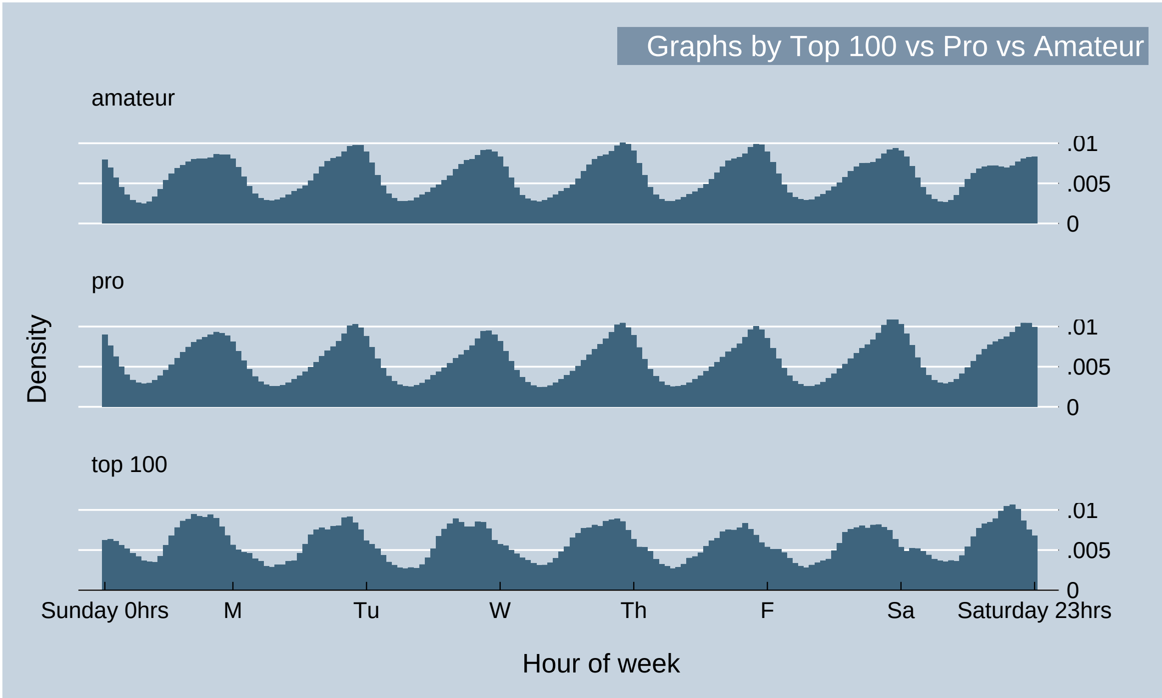

Note: Histograms of chosen hours-of-week to broadcast. Local times, 60,571,499 observations.

FIGURE 1: Pro vs. Amateur Schedules

If viewership is actually high, they play an extra match.

In practice, major broadcasters start their streams at a set schedule. They are not required to commit to any schedule, but consistency helps build viewership. However, scheduling is largely not strategic; broadcasters are motivated more by their live viewership and personal reasons than by competition from similar content providers. If we define professional streamers as being in the top 5% of popularity, Figure 1 shows the schedule distribution of pros versus amateurs. Schedules are similar except perhaps around 6pm, before primetime. No evidence of, say, American streamers targeting European consumers exists. If one worries about aggregation effects, Figure 1 also shows the schedule distribution for the top 100 channels.

The second decision—how long—depends on live viewership. All else equal, streamers are motivated to broadcast if they have a higher viewership. In turn, live viewership depends on their current competition from other streamers. A minority of streamers, called partners, have an extra financial motivation, because the platform shares ad revenues with them. In sum, all broadcasters respond to viewership.

The third decision—the bitrate—is not strategic. All else equal, a higher average bitrate implies a better quality, because it allows for a higher resolution. However, a viewer requires high-speed internet to be able to watch a live high-bitrate video. Otherwise, the viewer would experience video buffering and delays. Therefore, the broadcaster faces a trade-off between "viewability" and quality. Twitch guidelines and community suggest a target bitrate of 1.5 Mbps (SD video) to optimize viewership, which the data confirm.[5]

The platform gives preference to a minority of streams by transcoding their signal. A transcoder allows viewers to select the video quality that best suits their internet speeds. If a video is not transcoded, viewers must watch it at its original speed. Essentially, a transcoder works as a fast lane that allows content to reach viewers with low connection speeds. At any given moment, less than 10% of live channels are transcoded, in contrast to YouTube, where all non-live videos are transcoded by default. A live transcoding server is a scarce resource, because it uses specialized hardware to encode a single stream into several in real time. Reportedly, a transcoding server that costs about $1,000 can provide transcoding to four live channels at once. A popular channel, broadcasting HD video to 250 viewers, would cost $2,000 per hour using a third-party cloud-based transcoding service at current market rates.[6] The efficient allocation of transcoders is a first-order problem for the platform. In this sense, the platform is not neutral, because it gives preference to a minority of channels.

If a channel is transcoded, the broadcaster does not face the viewability-quality trade-off, because all her viewers can dial down the source speed to match an appropriate bitrate to their connections. Thus, a transcoded channel will choose the maximum speed its connection can handle, which, on average, is 2.5 Mbps (HD video). For these reasons, I treat the choice of video bitrate as exogenous.

The video bitrate is one of two key dimensions of the objective quality of a stream. The second dimension is the volatility of the bitrate. Video encoding and decoding is CPU-intensive. To decode a stream, a computer expects consistency from the source, especially if the encoding protocol is "Constant BitRate," which Twitch requires. Network congestion at any level creates variation in the streamed

[5] A standard-definition YouTube 480p video requires around a 1.5Mbps upload speed, whereas a high-definition 720p video requires around 2.5Mbps (`support.google.com/youtube`). In the first quarter of 2014, global average internet download speed was 3.9Mbps; a download speed of 3Mbps is recommended to watch SD video (`webanalyticsworld.net`, and `lifewire.com`).

[6] See `reddit.com/r/Twitch`. For market prices, see `amazon.com` and `zencoder.com`.

data due to lost data packets and bottlenecks. This bitrate volatility manifests as video buffering or stuttering. Evidence suggests delays and loss of data packets, such as those created by volatility, significantly decrease the end-users' quality of experience while watching a video stream (Pankert, Faggiano and Taga, 2014). All else equal, viewers prefer less volatility in their streams.

This paper measures volatility with the noise-to-signal ratio of the stream's video bitrate. The ratio is a form of coefficient of variation. Let b_{jt} stand for the (upload) video bitrate of channel j at date t, and let $\bar{b}_j$ be its mean over t. Then, the noise-to-signal ratio σ_{jt} is defined as

$$\sigma_{jt} \equiv \frac{\left| b_{jt} - \bar{b}_j \right|}{\bar{b}_j}.$$

Because Twitch enforces a constant bitrate protocol, in principle, b_{jt} should be constant. Any deviations from $\bar{b}_j$ correspond to noise caused by congestion. Note that $\bar{b}_j$ is a consistent estimate of the target bitrate and that σ_{jt} is dimensionless. All else equal, a low noise-to-signal ratio implies better quality and is largely beyond the control of the broadcasters.[7]

When content providers broadcast, the uploaded signal reaches a platform's server, from which it is mirrored through the platform's network and then distributed to edge servers, from which viewers download the stream. See Figure 2 for reference. At any leg of the journey, noise may be added to the signal due to congestion. Noise can be generated while uploading or downloading and has both idiosyncratic and aggregate components. Also note that noise gets compounded downstream; that is, noise generated on the upload side is transmitted to all the viewers watching the broadcast. However, noise generated on the download side is idiosyncratic to each consumer and does not affect the upload side of the network or other downloads.

On the upload side, the "production" of noise is composed of two factors: an individual factor and an aggregate factor. The individual factor is created by the broadcaster's private internet connection, the quality of the connection, and shocks beyond the platform's influence. The aggregate factor is common to all broadcasters and is rooted in the platform's physical capacities. This aggregate factor is a function

[7]For applications of the noise-to-signal ratio in the engineering sciences, see Huynh-Thu and Ghanbari (2008) or Shivaldova, Winkelbauer and Mecklenbrauker (2014). Huynh-Thu and Ghanbari (2008) discuss the scope of the noise-to-signal ratio to assess video quality. See also `help.twitch.tv`.

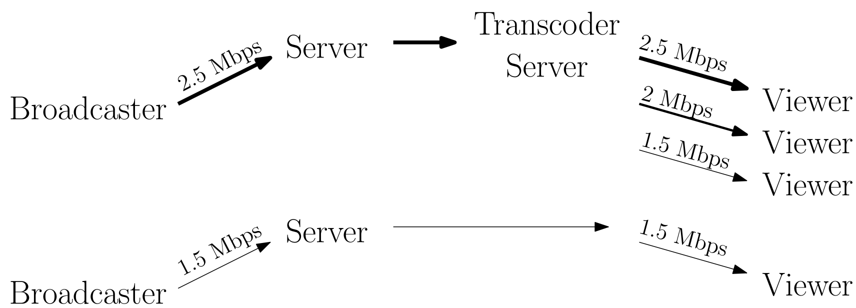

Note: An (abstracted) example of a broadcast's path from upload (↗) to download (↘). If the broadcast is transcoded, it reaches a transcode server that distributes the data at lower bitrates. If it is not, the broadcast is mirrored (→) at source speed. 1.5 Mbps $\sim$ SD video, 2.5 Mbps $\sim$ HD video.

FIGURE 2: Data flow with transcoding vs. without

of aggregate participation and the manifestation of congestion externalities.

Because the platform's physical network has a fixed throughput capacity, one would expect the noise-to-signal ratio to increase as more viewers and broadcasters participate in the platform. Indeed, increasing the uploads and downloads worsens queuing delays, packet loss, and connection blocking. However, its effect is confounded by selection, because viewers and broadcasters choose not to participate in the platform when it is congested. The noise ratio is thus an endogenous variable: Less noise attracts more people, and more people creates more noise. The identification strategy is composed of a combination of instruments to disentangle the effect of noise over participation and vice versa.

Twitch users can broadcast their streams from their personal computers or videogame consoles (e.g., PS4, Xbox One), and people can watch from any device. However, before March 11, 2014, viewers could watch broadcasts from the Xbox One, but broadcasters could not stream from it. On March 11, 2014, Twitch launched the complete app for the Xbox One, which allowed broadcasting.[8] This introduction of a new technology is the source of an exogenous variation that affected the broadcasters, but not the viewers. Figure 3a shows the time series of new accounts for the list of channels that are found in the data, which peaks on the date when the Xbox One broadcasting became available.

According to the platform's website, on March 11, 30% of broadcasters were streaming from the Xbox One, and within the first week, a total of 108,000 unique

[8] See `blog.twitch.tv`, published February 25, 2014.

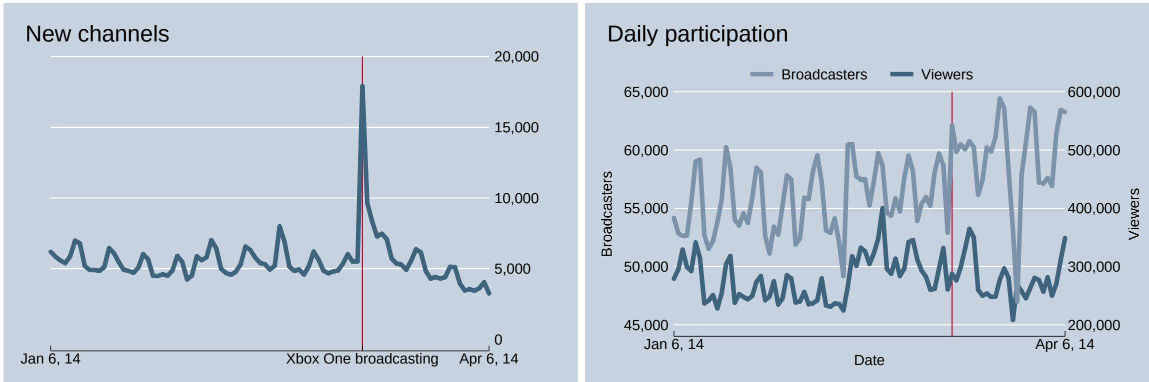

(A) Broadcaster account creation

(B) Daily broadcasting and viewership

Note: (A) Daily number of new broadcasting accounts created. (B) Daily number of unique channels in the data, and daily maximums of concurrent viewers in the data. Peaks correspond to weekends. Vertical lines mark Tuesday, March 11, 2014.

FIGURE 3: New channels and daily participation

broadcasters were streaming from the Xbox One, accounting for 22% of Twitch's broadcaster base.[9] Figure 3b shows the time series of daily broadcasting and viewership. On Tuesday, March 11, broadcasting numbers jumped, confirming the press release from the platform.

Additionally, the platform expanded its infrastructure four times during the sample period by upgrading their points of presence (PoP).[10] The expansion of PoPs increases the ingest capacity of the network and improves user experience in general. Effects are expected to be short-lived due to the "Iron Law of Congestion": improved capacity attracts more users, increasing noise (Duranton and Turner, 2011). The data show decreases in noise following each installment; see Figure 4, where the Xbox effect also is apparent.

Importantly, the Xbox One and PoP upgrades took months of planning and were released as soon as they were available.[11] That is, Twitch follows a long-run business plan instead of reacting to the short term. For instance, the Xbox One upgrade was unveiled on a Tuesday, which is the second-least popular day to broadcast. Instead of being correlated with current traffic, the update was planned to coincide with the release of Titanfall, a video game.[12]

[9]See `blog.twitch.tv`, published March 31, 2014.
[10]See `blog.twitch.tv`, published March 28, 2014.
[11]See `reddit.com`, `blog.twitch.tv`, or `blog.twitch.tv`.
[12]See `polygon.com`.

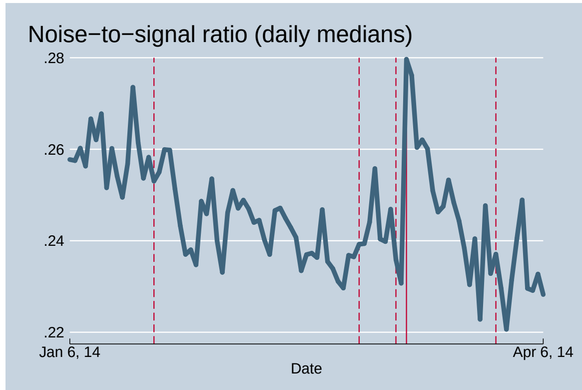

Note: Dashed lines signal the platform's point-of-presence upgrades. The solid line signals the Xbox One broadcasting shock.

FIGURE 4: Quality of service and point-of-presence upgrades

The platform incentivizes popular broadcasters through a partnership program. To join the program, the broadcaster needs to meet certain technical requirements, to be popular enough, and to be willing to join. Joining the partnership program is free, and its benefits include receiving a share of ad revenue, the ability to accept donations from viewers in the form of a subscription, and, more importantly, guaranteed video transcoding. From the viewer's perspective, beyond transcoding, the remaining perks of partnered channels, such as special chat interactions and "virtual badges," do not offer substantial benefits.

To be a partner, broadcasters must apply through the platform's website and be approved. Once they apply, broadcasters do not know if and when they will be approved. The platform approves partnerships on a rolling basis through an undisclosed method, which takes an unknown amount of time. From the broadcaster's point of view, partnership comes unexpectedly. Non-partners do not have incentives to behave differently until they have been upgraded to partners. Once gained, partnerships are never lost. In the data, I observe more than 1,600 partners, of whom more than 600 acquired partnerships during the sample period.

I do not observe partners' financial incentives. However, these incentives are proportional to their viewership numbers. In practice, all else equal, partners stream more than non-partners but are less sensitive to viewership, partly because they have more viewers.

Because transcoding is guaranteed for partners, the platform has roughly enough transcoding servers for them. When not all the servers are being used, the excess capacity is allocated to non-partnered channels according to their current viewership. This excess allocation is never announced and comes as a surprise for viewers and streamers. Thus, viewers and non-partnered streamers do not have incentives to behave differently before they are allocated a transcoder.

Table 1 shows the main statistics of the streaming sessions. The average broadcasting session has 15 viewers, lasts for 74 minutes, and streams at a video bitrate of 1.5Mbps. Partnership plays a role as expected: Partners have much longer sessions, better viewership numbers, and higher video quality. The difference between a partnered broadcast and a non-partnered one is roughly that between HD and SD video.

After the Xbox shock, the number of gaming channels and viewers increased while the number of non-gaming channels and viewers decreased. The Xbox shock did not affect the non-gaming segment directly, which suggests the presence of a negative externality.

Notably, the data contain no prices, because joining the platform is free. In practice, viewers pay with "eyeballs," because they have to watch ads to view the broadcasts. Moreover, this pricing and ad policy remained constant during the sampled time frame. When considering welfare implications, I focus on quantities.

The following section presents the structural models and identification arguments in detail.

3 The model

The model has three components. In the congestion model, individual noise-to-signal ratios are determined by aggregate participation in the platform. In the demand model, individual consumers decide which channel to watch, given a set of online channels, their noise-to-signal ratios, and whether they have access to transcoding. In the supply model, individual streamers decide to be online, depending on their individual demands. Section 3.1 introduces the congestion model, section 3.2 the demand model, and section 3.3 the content-provision model. Each component is estimated separately in section 4.

TABLE 1: SUMMARY STATISTICS

	MEAN	STD	P10	P25	P50	P75	P90	BEFORE	AFTER
All Sessions:									
DURATION (MIN)	73	242	0	0	22	80	177	74	71
MEAN VIEWERS	14.9	436.3	0.0	1.0	1.5	3.0	7.7	15	14
MAX VIEWERS	22	686	0	1	2	4	11	22	21
BITRATE (KBPS)	1486	1132	379	708	1148	2092	3026	1460	1545
NOISE-TO-SIGNAL (%)	17	22	3	5	8	21	47	17	18
AVG ONLINE CHANNELS								5,266	5,562
AVG ONLINE VIEWERS								308,524	312,364
UNIQUE BROADCASTERS	1,520,406								
OBSERVATIONS	11,614,966								
Partner:									
DURATION (MIN)	223	305	20	80	170	299	450	223	224
MEAN VIEWERS	798.4	2450.4	30.0	79.7	212.2	582.9	1549.5	850	696
MAX VIEWERS	1164	3577	40	107	294	824	2245	1239	1016
BITRATE (KBPS)	2330	1061	1022	1585	2261	3066	3559	2317	2357
NOISE-TO-SIGNAL (%)	12	14	4	5	7	11	26	12	12
AVG ONLINE CHANNELS								136	164
AVG ONLINE VIEWERS								169,231	162,123
UNIQUE BROADCASTERS	1,736								
OBSERVATIONS	97,775								
Non-Partner:									
DURATION (MIN)	72	241	0	0	22	79	172	73	69
MEAN VIEWERS	8.2	368.4	0.0	1.0	1.5	2.9	7.0	8	8
MAX VIEWERS	12	595	0	1	2	4	10	12	12
BITRATE (KBPS)	1479	1130	377	704	1143	2086	3016	1453	1538
NOISE-TO-SIGNAL (%)	18	22	3	5	8	21	47	17	18
AVG ONLINE CHANNELS								5,130	5,398
AVG ONLINE VIEWERS								139,294	150,241
UNIQUE BROADCASTERS	1,519,351								
OBSERVATIONS	11,517,191								
Gaming:									
DURATION (MIN)	63	128	0	0	22	78	169	63	62
MEAN VIEWERS	14.8	448.0	0.0	1.0	1.5	2.9	7.0	15	14
MAX VIEWERS	22	704	0	1	2	4	10	22	21
BITRATE (KBPS)	1535	1132	432	762	1192	2118	3053	1510	1588
NOISE-TO-SIGNAL (%)	17	22	3	5	8	20	46	17	18
AVG ONLINE CHANNELS								4,592	4,970
AVG ONLINE VIEWERS								284,550	290,583
UNIQUE BROADCASTERS	1,467,585								
OBSERVATIONS	10,986,717								
Non-Gaming:									
DURATION (MIN)	252	870	0	0	35	139	533	254	248
MEAN VIEWERS	16.5	99.5	0.0	0.1	1.5	6.5	37.2	17	16
MAX VIEWERS	24	157	0	1	2	9	48	24	23
BITRATE (KBPS)	641	738	70	197	439	857	1358	626	679
NOISE-TO-SIGNAL (%)	23	25	4	7	14	30	51	23	23
AVG ONLINE CHANNELS								674	593
AVG ONLINE VIEWERS								23,975	21,781
UNIQUE BROADCASTERS	53,353								
OBSERVATIONS	628,249								

Notes: Data collected at 5-minute intervals. Duration in minutes with a 5-minute margin of error. Video bitrate in kbps. Before and after with respect to the Xbox One shock.

3.1 Congestion externalities

I assume the following model for channel j's noise-to-signal ratio, σ_{jt},

$$\log \sigma_{jt} = \underbrace{\mu_0^\sigma t + \alpha_1^\sigma \log B_t^{up} + \boldsymbol{\beta}^{\sigma\prime} \mathbf{PoP}_t}_{\text{aggregate component}} \quad \underbrace{+\, \mu_j^\sigma + \varepsilon_{jt}^\sigma}_{\text{idiosyncratic component}}, \tag{1}$$

where μ_0^σ is a trend, B_t^{up} is the total video bitrate uploaded to the platform's servers at time t, $\mathbf{PoP}_t$ is a vector of four indicators for the point of presence upgrades, and μ_j^σ is a fixed effect. The parameter of interest is α_1^σ, which measures the elasticity of noise to traffic. In particular, $\mu_0^\sigma t$ captures technological progress and updates that naturally reduce noise at the aggregate level. The time-step for t is 10 minutes.

The data constrain the congestion model to focus on the upload side. I implicitly assume total download video bitrate affects the noise-to-signal ratio in a way proportional to that of B_t^{up}. That is, I assume B_t^{up} is a proxy for total traffic. Section 3.2 discusses the implied assumptions and their relevance for the demand-side model.

Identification. Although individual channels are too small to have an effect on aggregate variables, individual noise and traffic are simultaneously determined. Traffic increases noise. However, if noise is too high, channels turn off, which reduces total traffic. I use the Xbox One shock as an instrument for $\log B_t^{up}$ because the shock increased the number of broadcasters, which increased the upload traffic.

A threat to identification comes from the possibility that the Xbox One substantially changes demand or the characteristics of supply. However, the channels attracted by the Xbox One shock mostly have low or no viewership; for the next few days after the Xbox shock, the number of viewers per channel is below trend due to an increase in the number of channels. Moreover, I can perform the estimation only for the professional channels (those in the top 5% of popularity), which are arguably not changing their behavior because of the Xbox One broadcasting availability. Results are similar in this subsample.

Finally, I include fixed effects in this estimation because no cross-sectional variation exists over B_t^{up}. The fixed effects also control for unobserved channel characteristics such as location, internet connection, or hardware.

Formally,

Assumption 1. $\log B_t^{up} \perp \varepsilon_{jt}^\sigma | \text{Xbox}_t, \mathbf{PoP}_t, \mu_j^\sigma, t$.

All estimation results are presented in section 4.

3.2 Viewers (demand)

Each period t, a representative consumer i decides which channel to watch according to a discrete-choice model with a nested structure (McFadden, 1974; Berry, 1994). Let $\mathcal{J}$ be the set of all channels that have an account on the platform. The set $\mathcal{J}$ exogenously grows over time as new channels sign up; no new notation is introduced to account for this fact. Let $\mathcal{J}_t \subset \mathcal{J}$ be the set of online channels at time t, and it always includes the outside option, which is watching nothing. The consumer's outside option has a normalized utility of ε_{i0t}, whereas the payoff for tuning into channel j is $u_{ijt} \equiv \delta_{jt} + \varepsilon_{ijt}$. To allow for more flexible substitution patterns, four nests are defined: (1) the outside option, (2) non-gaming channels, (3) amateur gaming channels, (4) pro gaming channels. The subscript $g = 1, \ldots, 4$ labels nests.

Let σ_{ijt} be the final noise-to-signal ratio experienced by consumer i watching channel j at time t. I assume

$$\log \sigma_{ijt} = \log \sigma_{jt} + \eta_{ijt},$$

for some unobservable η_{ijt}, with $\eta_{ijt} \perp \log \sigma_{jt}$. In other words, I assume the final noise-to-signal ratio can be decomposed into the "upload" side noise, $\log \sigma_{jt}$, plus an idiosyncratic shock for consumer i, η_{ijt}, which is absorbed by ε_{ijt}. Because noise is compounded downstream, this assumption is close to reality.

The consumer solves $\max_{j \in \mathcal{J}_t} u_{ijt}$ and, assuming type-1 extreme-value errors, the market-share equation becomes

$$\log\left(\frac{s_{jt}}{s_{0t}}\right) = \delta_{jt} \equiv \mu_j + \alpha_1 \log \sigma_{jt} + \alpha_2 \tau_{jt} + \alpha_3 \log s_{jt|\mathrm{g}} + \boldsymbol{\beta}' \mathbf{x}_{jt} + \xi_{jt}, \tag{2}$$

where μ_j is a fixed effect, σ_{jt} is the noise-to-signal ratio of channel j, τ_{jt} indicates whether channel j's broadcast is transcoded, $s_{jt|\mathrm{g}}$ is the market share of channel j as a fraction of group g (Berry, 1994; Cardell, 1997), and ξ_{jt} is unobservable. Channel j characteristics, $\mathbf{x}_{jt}$, include the following: indicators for partnership, for a channel being featured on the platform's homepage, for weekends, and for peak hours (between 6pm and 1am local time); the log of tenure in the platform (as a proxy for experience, quality, and fan-base size); and the log of uptime.

The quality of a stream can be divided into subjective and objective qualities. The unobserved subjective quality of a video (how good the content is) is absorbed by

μ_j and ξ_{jt} and may be regarded as a horizontal attribute. As a vertical attribute, the objective quality is composed of three parts: the video bitrate, the noise-to-signal ratio, and whether the stream is transcoded. The video bitrate is absorbed by the fixed effect, because, in principle, the bitrate should be constant. All else equal, I expect a high noise-to-signal ratio to decrease the demand for a channel, and a transcoded stream to have more viewers.

Let M_t be the market size at time t; then, $s_{0t}M_t$ viewers tune out of the platform at t. Therefore, quasi-demand is $(1 - s_{0t})M_t$ and quasi-supply is $\#\mathcal{J}_t$ (Rochet and Tirole, 2006). Streamers compete for viewers with a Twitch account. The relevant market size is then the number of Twitch viewers. In February 2014, Twitch.tv had about 45 million unique viewers each month. By the end of 2014, the number had reached 100 million. Assuming linear growth, by the end of April 2014, the platform would have grown to around 50 million viewers. The platform also reports that each viewer watches an average of 106 minutes per day. I assume viewers have two hours available per day to watch videos. The data span 90 days. Assuming a uniform distribution over the day, I define the market size at date $d = 1, \ldots, 90$ as $M_d = 45(2/24) + 5(2/24)(d/90)$ million viewers available at date d. Thus, $M_t = M_d$ if time t is part of day d. The results are not sensitive to this assumption. Demand equation (2) does not include a time trend, because platform growth is already embedded in M_t.

Identification. First, as an equilibrium outcome, noise is simultaneously determined with demand: Noise decreases demand, but demand increases supply, which increases noise. Also, noise could be correlated with the error term through selection, because noisier channels go offline. Yet another concern is measurement error, because σ_{jt} is constructed using the average bitrate as a proxy for the target bitrate.

To mitigate these concerns, the regression includes a fixed effect for channel j; moreover, I use two sets of instruments for $\log \sigma_{jt}$; and, finally, I correct for selection. The first set of instruments are the PoP upgrades, which improved the network infrastructure and decreased noise all over the platform. These upgrades should not increase viewership directly but only through noise reduction. A second instrument comes from the introduction of the Xbox One broadcasting ability, which increased upload video bitrate. The spike in traffic increased the noise-to-signal ratio across the platform. Importantly, the Xbox shock attracted fringe channels with low viewership: Channels created after March 11, 2014, have 2.1 viewers on average, whereas the 90th percentile is 2.5 viewers, compared with 15 and 7.7, respectively, for the whole sample. Thus, consideration sets are virtually unchanged after the shock.

A more serious threat to identification comes from the possibility of the platform reacting strategically to viewership when it deployed the PoP and Xbox upgrades. However, as argued in section 2, these upgrades require months of planning, rather than days or weeks.

The instruments do not contain cross-sectional variation. I address selection separately using a control function based on Heckman (1979) and Olsen (1980). In a first step, the decision of channel j to be online at t is modeled as a function of a channel fixed effect, whether the stream was interrupted, viewership, and other observables at $t-1$. In a second step, a control function is included in the regression to deal with the fact that noisier channels choose to turn off.

Equation (2) includes no random coefficients that allow for more flexible substitution patterns. However, the logic that applies to common price elasticities does not translate directly into noise elasticities. A classic problem in a BLP-type model is that own-price elasticities are roughly constant (when the model includes log prices), which implies similar markups even for products with different prices (Berry, Levinsohn and Pakes, 1995). An additional problem is that cross-price elasticities are proportional to market shares. However, back to this setting, channels do not choose their noise-to-signal ratio. Therefore, we have no a priori reason to discard constant own-noise elasticities. More importantly, the relevant substitution direction is toward the outside option. The reason the nested logit is favored over a random-coefficients approach is a practical one. Because viewers can watch any channel around the world, at any given time, 5,000 "products" are available in about 13,000 "markets" t.

Second, transcoder allocation is correlated with quality. Identification of the transcoder effect comes from comparing non-partnered channels before and after receiving a transcoder. For non-partners, transcoders come as a surprise. Formally, I assume transcoders are orthogonal to the error conditional on fixed effects. That is, conditional on quality, the allocation of transcoders is as good as random. A possible threat to identification comes from the fact that receiving a transcoder means fewer partnered channels are online, which means less competition. The estimation could be picking up the effect of less competition, instead of the causal effect of a transcoder, and could bias the target coefficient upward. To mitigate this concern, I include controls such as weekend and peak-hour dummies. The results are also robust to the inclusion of day-of-week and hour-of-day fixed effects.

Third, by construction, $\log s_{jt|g}$ is an endogenous regressor. Following Berry (1994), I instrument it with the characteristics of other channels in the group. Namely, I instrument $\log s_{jt|g}$ with the number of channels: in English, from the US, with transcoders, featured, and partnered, as percentages of group g, excluding j, at time t. The characteristics of other channels in the group should not affect directly the market share of channel j but should shift the within-group shares.

Although the data set is poor on channel characteristics, it is rich in the time dimension. Channel fixed effects control for important but unobserved attributes such as content, subjective quality, and other horizontal characteristics. The following summarizes the identification assumptions.

Assumption 2.
(a) $\log \sigma_{jt} \perp \xi_{jt} | \text{Xbox}_t, \mathbf{PoP}_t, \mu_j, \mathbf{x}_{jt}$;
(b) $\tau_{jt} \perp \xi_{jt} | \mu_j, \mathbf{x}_{jt}$;
(c) $\log s_{jt|g} \perp \xi_{jt} | \mu_j, \mathbf{x}_{jt}$, English$\%_{-jt|g}$, US$\%_{-jt|g}$, Transcoder$\%_{-jt|g}$, Featured$\%_{-jt|g}$, Partner$\%_{-jt|g}$.

3.3 Content providers (supply)

Broadcasters are strategic. Each period t, they play a simultaneous game in which each broadcaster $j \in \mathcal{J}$ decides to be online or not. Let y_{jt} indicate if channel j is online at t. Let u^C_{jt} be the utility of broadcasting at t. That is, the payoffs are $\{u^C_{jt}(y_{jt}, y_{-jt})\}_{j \in \mathcal{J}}$, which depend on own strategies (y_{jt}) and on others' strategies (y_{-jt}), and are defined below. The utility of not broadcasting is normalized to zero. Thus, broadcasters solve a static problem: $\max \left\{0, u^C_{jt}\right\}$. Lemma 1 in the Appendix ensures a Nash equilibrium exists for this game.

Payoffs u^C_{jt} depend on the information available at t. In particular, the decision depends on the number of viewers watching j. However, this information is not available when channel j is offline. I thus separate broadcasters' decisions by their information sets.

I assume the broadcaster's static problem is divided into an extensive and an intensive margin: (1) When the channel is offline, the broadcaster decides whether to **Turn On**. (2) If the broadcaster has decided to be online, this decision will be followed by the decision to **Keep On** or not, in the following periods.

First, when the channel is offline, the channel does not know exactly how many viewers it could have if it were online, but it can make an educated guess. I model this decision as a function of a fixed effect, and observables such as the aggregate number of viewers and channels, which proxy as potential market and competition. Importantly, a deciding factor is the promise of a transcoder, which is only guaranteed for partners. Time effects are included to account for schedules.

Second, given that the channel is already online at time t, the number of viewers is realized and the channel decides whether to Keep On given this information. Thus, being online at $t+1$ is decided at t with information at t.

Specifically, channel j observes the state of the world and decides whether to Turn On at the beginning of period t, according to a utility associated with each decision. Let u_{jt}^{IO} be the utility of Turning On, and let 0 be the utility of staying offline. Thus, channel j Turns On iff $u_{jt}^{IO} > 0$. That is, I assume

$$y_{j,t+1} = 1 \,|\, y_{jt} = 0 \iff u_{jt}^{IO} > 0,$$

$$u_{jt}^{IO} \equiv \mu_j^{IO} + \alpha_1^{IO} \mathbb{1}\{\text{Partner}_{jt}\} + \boldsymbol{\beta}^{IO\prime} \mathbf{x}_{jt}^{IO} + \varepsilon_{jt}^{IO}, \tag{3}$$

where μ_j^{IO} is a fixed effect, and $\mathbb{1}\{\text{Partner}_{jt}\}$ is an indicator for partnership. The controls, $\mathbf{x}_{jt}^{IO}$, include information available before the transmission starts: an indicator of the Xbox One broadcasting availability; its interaction with a gaming dummy; indicators for weekend and for peak hours; and the log of tenure in the platform, as a proxy for experience or for developing a fan base. Channel j perfectly observes u_{jt}^{IO}. The fixed effects μ_j^{IO} and the error ε_{jt}^{IO} are unobservable to the researcher.

All else equal, one would expect partners to be more willing to broadcast, because they know they will have a transcoder, which will increase their viewership.

Given that the channel decided to Turn On at t, the viewers of channel j are realized and observed by j at t. With this information, at the end of period t, the channel decides whether to Keep On or to Turn Off at $t+1$. Let u_{jt}^{B} be the utility of Keeping On, and let 0 be the utility of Turning Off. Thus, channel j Keeps On iff $u_{jt}^{B} > 0$:

$$y_{j,t+1} = 1 \,|\, y_{jt} = 1 \iff u_{jt}^{B} > 0,$$

$$u_{jt}^{B} \equiv \mu_j^{B} + \alpha_1^{B} \log n_{jt}^{V,\text{Partner}} + \alpha_2^{B} \log n_{jt}^{V,\text{NonPartner}} + \alpha_3^{B} \mathbb{1}\{\text{Partner}_{jt}\} + \boldsymbol{\beta}^{B\prime} \mathbf{x}_{jt}^{B} + \varepsilon_{jt}^{B}, \tag{4}$$

where μ_j^B is a fixed effect, n_{jt}^V is the number of viewers watching j at t, and it is interacted with the partnership dummy. The controls, $\mathbf{x}_{jt}^B$, include the following: indicators for Xbox broadcasting availability, for weekends, and for peak hours; the log of tenure; and y_{jt}. Channel j perfectly observes u_{jt}^B. The fixed effects μ_j^B and the error ε_{jt}^B are unobservable to the researcher.

Because only information sets change between equations (3) and (4), I impose equal fixed effects in both equations. That is, I discipline the model by assuming

$$\mu_j^{IO} = \mu_j^B.$$

These fixed effects reflect channel j's unconditional expectation of viewership. The effect of own viewership, n_{jt}^V, is separated between partners and non-partners to allow different marginal valuations.

At this point, the likelihood of observing a given history of a channel can be written. Let T_j^s define the set of periods when channel j had its s-th broadcasting session. That is, $T_j^s \equiv \left\{t_{1j}^s, \ldots, t_{end,j}^s\right\}$, where the channel was offline at time $t_{1j}^s - 1$, Turned On at t_{1j}^s and was online until $t_{end,j}^s$, when it decided to Turn Off. That is, the typical s-th session is $\{y_{jt}\}_{t \in T_j^s} = \{1, 1, \ldots, 1, 1, 0\}$. Any channel j could have S_j of such broadcasting sessions. The likelihood of session s is

$$L_j^s\left(\boldsymbol{\theta}_0^{IO}, \boldsymbol{\theta}_0^B\right) = P\left[y_{jt_{1j}^s} = 1 \,\middle|\, y_{j,t_{1j}^s-1} = 0, \boldsymbol{\theta}_0^{IO}\right] \left[\prod_{t=t_{1j}^s}^{t_{end,j}^s-2} P\left[y_{j,t+1} = 1 \,\middle|\, y_{jt} = 1, \boldsymbol{\theta}_0^B\right]\right]$$
$$\times P\left[y_{j,t_{end,j}^s} = 0 \,\middle|\, y_{j,t_{end,j}^s-1} = 1, \boldsymbol{\theta}_0^B\right],$$

where $\boldsymbol{\theta}_0^{IO}$ and $\boldsymbol{\theta}_0^B$ are the vectors of parameters from the decisions of Turning On and Keeping On.

This likelihood is conditional on the fact that (by definition) the channel was offline before the session starts. To wrap up the likelihood of the data, define $T_j^{\text{off}} \equiv T \setminus \bigcup_{s=1}^{S_j} T_j^s$ as the inter-session periods when the channel was offline, and where T is the set of periods in the data. Then,

$$L_j\left(\boldsymbol{\theta}_0^{IO} | y_{j0}\right) = \prod_{t \in T_j^{\text{off}}} P\left[y_{jt} = 0 \,\middle|\, y_{j,t-1} = 0, \boldsymbol{\theta}_0^{IO}\right]$$

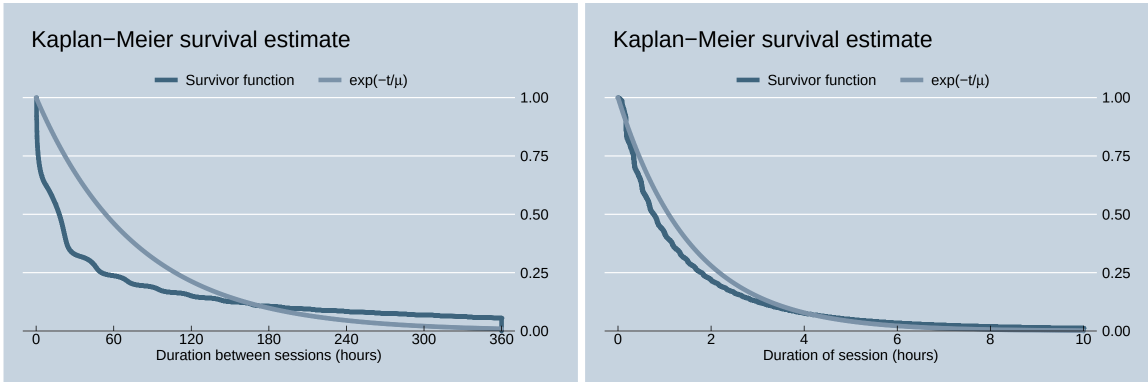

Note: Downtime and uptime comparison. Kaplan-Meier survival estimates and a fitted exponential curve.

FIGURE 5: Comparison of downtime and uptime

is the likelihood of observing channel j being offline in the inter-session periods, conditional on the starting value y_{j0}.

Finally, the likelihood of the data can then be written by considering all sessions for all channels:

$$L\left(\boldsymbol{\theta}_0^{IO}, \boldsymbol{\theta}_0^{B}\right) = \prod_{j \in \mathcal{J}} L_j\left(\boldsymbol{\theta}_0^{IO} | y_{j0}\right) \prod_{s=1}^{S_j} L_j^s\left(\boldsymbol{\theta}_0^{IO}, \boldsymbol{\theta}_0^{B}\right). \tag{5}$$

In particular, this model is consistent with a survival model with exponential decay, in which each period the broadcaster can change state. Figure 5 presents non-parametric Kaplan-Meier survival estimates for downtime, where Turning On is the change of state, and for uptime, where Turning Off is the change of state. The figure also shows a simple exponential curve, with no covariates, approximating the Kaplan-Meier curve.

Identification. A simplifying assumption is implicit in (5): I assume equations (3) and (4) are independent of each other. That is, the working assumption is that the errors from equations (3) and (4) are independent.

The potential problem with this assumption is that a shock that pushes a channel to Turn On could be correlated with the shocks a channel experiences once it is online. That is, the errors could be correlated across time. I do not consider this issue to be a major setback, because I include individual effects on both equations and I allow for panels in the likelihood. I also include tenure data, which allows me to control for unobserved dynamics. Moreover, the likelihood conditions on the outcome of the previous period, and the presence of autocorrelation does not bias the estimates. Dynamic considerations are even less important if one assumes channels can predict their viewers in the next 10 minutes by observing their current number of viewers.

Because the likelihoods are separable, the parameters can be estimated with a stacked logit. The advantage is that I can include channel fixed effects in both equations, which allows me to control for unobserved but important attributes such as content, quality, or ability that may be correlated with partnership or viewership. Because the time series is long for each panel, the incidental parameters problem is a minor issue. Alternatively, one could estimate a conditional logit in which the channel fixed effects drop out, so they are never actually estimated. However, the conditional logit is impractical.

Now, consider the problem of estimating the effect of partnership in the decision to be online. Being a partner is endogenous, because partners are expected to have a higher quality and to be the most popular channels. Thus, to identify the causal effect of partnership, I exploit within-panel differences of channels before and after they become partners. Identification comes from non-partnered channels that suddenly are upgraded to partners. The tenure of the channel controls for growth in quality or in the fan base over time.

Finally, consider the problem of estimating the effect of the number of viewers, n^V_{jt}. Viewership could be correlated with the error term for a number of reasons. First, endogenous scheduling could dampen the elasticity of supply to viewership. Second, demand and supply are simultaneously determined. Third, unobserved quality could have a positive effect on the utility of the broadcaster. Fourth, dynamic considerations, such as inertia of viewers, could imply lagged variables or differences are omitted variables.

To identify the effects of n^V_{jt}, first, I include peak-hour and weekend dummies. If schedule choice is endogenous, supply's elasticity to viewership would be underestimated, especially for larger channels, because larger channels would not change their broadcasting decisions as much, because they are predetermined by their schedules. Appendix section A.1 defines a two-stage game, where streamers choose schedules in a first stage and choose to keep online in a second stage. The upshot is that a Nash equilibrium exists and that the two-stage game can be approximated by the benchmark if we control for scheduling.

Second, I use the demand-side model to derive exogenous variation in viewership. Specifically, in a first stage, I instrument $\log n_{jt}^V$ with the noise-to-signal ratio, a dummy for transcoding, and a dummy for being featured, and I construct a control function (Petrin and Train, 2010). That is, I estimate

$$\log n_{jt}^V = \mu_j^F + \alpha_1^F \log \sigma_{jt} + \alpha_2^F \tau_{jt} + \alpha_2^F \mathbb{1}\{\text{Featured}_{jt}\} + \boldsymbol{\beta}^{F\prime} \mathbf{x}_{jt}^F + \eta_{jt}, \tag{2'}$$

where η_{jt} and ε_{jt}^B are correlated with each other. I then construct $\widehat{\eta}_{jt}$ and include it as an additional regressor in (4).

The following summarizes the assumptions.

Assumption 3.
(a) $\varepsilon_{jt}^{IO}, \varepsilon_{jt}^B$ are iid;
(b) $\mathbb{1}\{\text{Partner}_{jt}\} \perp \varepsilon_{jt}^{IO} | \mu_j^{IO}, \mathbf{x}_{jt}^{IO}$;
(c) $\mathbb{1}\{\text{Partner}_{jt}\} \perp \varepsilon_{jt}^{B} | \mu_j^{B}, \mathbf{x}_{jt}^{B}$;
(d) the error ε_{jt}^B can be decomposed into the part that can be explained by a general function of an error η_{jt} and a residual: $\varepsilon_{jt}^B \equiv \text{CF}(\eta) + \widetilde{\varepsilon}_{jt}^B$;
(e) $\log n_{jt}^V \perp \varepsilon_{jt}^B | \mu_j^B, \mathbf{x}_{jt}^B, \widehat{\eta}_{jt}$.

In practice, for computational concerns, I separate the estimation of the supply decision by partners and non-partners. Partners provide identification of the effect of a partnership change. Moreover, a separate estimation allows partners and non-partners to have different marginal valuations of viewership.

4 Estimation results

Congestion externalities. Table 3 shows the estimation results of the noise equation (1). Column 1 shows an OLS. Columns 2 to 5 instrument B_t^{up} with an Xbox One dummy; first stages can be found in Table 2. Column 2 restricts the sample to pro channels (defined by being in the top 5% of popularity), and column 3 restricts it for non-gaming channels. These subsamples should not be affected by the Xbox shock directly under any circumstance. They are only affected by congestion externalities. Finally, column 5 corrects for selection. The coefficient of interest is not statistically different between columns 4 and 5.

The coefficients of the PoP upgrades loosely align with the relative importance of these upgrades, even if we only establish correlations. Although precise numbers were not released, the first upgrade was the largest, followed by the third, followed by the second. No information is available on the relative size of the fourth upgrade.[13]

The results imply increasing the aggregate, upload video bitrate by 1% will increase the noise-to-signal ratio of *every* channel by about 0.8%. An increase of 10% in B_t^{up} would translate into an increase of 8% in the noise-to-signal ratio of every channel. In a single day, B_t^{up} can fluctuate by up to 50%.

Demand-side results. Table 4 shows the estimation results of the demand equation (2). Column 1 shows an OLS. Columns 2 and 3 instrument $\log \sigma_{jt}$ and $\log s_{jt|g}$. Column 3 controls for selection. First stages are reported in Table 6 in the appendix.

As expected, an increase in the noise-to-signal ratio decreases demand. Because market shares are small, $(1 - \alpha_3 s_{jt|g} - (1 - \alpha_3)s_{jt}) \approx 1$, then $\alpha_1/(1 - \alpha_3)$ can be interpreted as the elasticity of own noise-to-signal ratio, where α_1 and α_3 are the coefficients of $\log \sigma_{jt}$ and $\log s_{jt|g}$. In other words, if σ_{jt} increases by 1%, demand decreases approximately 2.5%. Demand is elastic to noise. The partial-equilibrium effect of a 10% increase in B_t^{up} translates into an increase of 8% in σ_{jt}, which in turn implies a 20% reduction in demand.

Transcoding increases demand in a significant way, approximately by 25%. As expected, conditional on transcoding, features, and a fixed effect, a channel's partnership status does not increase viewership, because the remaining perks of partnered channels are trivial for viewers.

Finally, it can be shown that increasing the number of online channels by 1 increases viewership by $(s_0 - s_0^+)/(1 - s_0)$ percent in expectation, where s_0 is the share of the outside option and s_0^+ is the share of the outside option when the extra channel is online. Translated into elasticities, increasing the number of channels by 1% increases the number of viewers by 0.12%.

Supply-side results. Results are in Table 5. Both columns show logit models. Translated into probabilities, a 1% increase in viewership increases the probability of keeping on by about 0.9% for partners, and by about 1% for non-partners. That is, partners' supply is less elastic to viewership than non-partners', even after controlling for scheduling. Because partners have many more viewers than non-partners, a lower marginal valuation is expected. Gaining partnership increases the probabil-

[13] See `blog.twitch.tv`.

TABLE 2: NOISE ESTIMATIONS, FIRST STAGES

	Models				
DEP VAR:	(1) Online$_{j,t+1} = 1$	(2) $\log B_t^{Upload}$	(3) $\log B_t^{Upload}$	(4) $\log B_t^{Upload}$	(5) $\log B_t^{Upload}$
XBOX AVAILABLE	-0.003^{***}	0.041^{***}	0.043^{***}	0.044^{***}	0.043^{***}
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
POP1	-0.003^{***}	-0.011^{***}	-0.006^{***}	-0.012^{***}	-0.013^{***}
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
POP2	0.002^{***}	0.007^{***}	0.007^{***}	0.004^{***}	0.004^{***}
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
POP3	0.002^{***}	0.011^{***}	0.006^{***}	0.009^{***}	0.010^{***}
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
POP4	0.013^{***}	0.053^{***}	0.057^{***}	0.051^{***}	0.055^{***}
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
TREND	0.000^{***}	0.001^{***}	0.001^{***}	0.001^{***}	0.001^{***}
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
$\log n_{jt}^V$ (PARTNER)	0.010^{***}				
	(0.000)				
$\log n_{jt}^V$ (NON-PARTNER)	0.031^{***}				
	(0.000)				
PARTNER	0.113^{***}				
	(0.002)				
TRANSCODER	-0.014^{***}				
	(0.000)				
FEATURED	-0.052^{***}				
	(0.010)				
INTERRUPTED	-0.028^{***}				
	(0.000)				
WEEKEND	-0.008^{***}				
	(0.000)				
AFTER 6PM	-0.016^{***}				
	(0.000)				
log TENURE	0.004^{***}				
	(0.000)				
log UPTIME	-0.004^{***}				
	(0.000)				
HECKMAN					-0.375^{***}
					(0.005)
CONS	0.854^{***}	15.942^{***}	15.924^{***}	15.943^{***}	15.907^{***}
	(0.001)	(0.000)	(0.001)	(0.000)	(0.001)
CHANNEL FE	YES	YES	YES	YES	YES
SUBSAMPLE OF	.	PROS	NONGAMING	.	.
R_a^2	0.12	0.18	0.16	0.16	0.17
F-STAT	14,934	38,098	10,323	89,968	74,229
OBS	48,570,885	15,861,549	2,649,865	38,283,191	38,283,191

Notes: Standard errors in parentheses, clustered at channel id. Time series of 10-minute windows. First column predicts channel j keeping online at time $t+1$; predictions used to contruct a control function to correct for selection bias, based on Heckman (1979) and Olsen (1980). B_t^{up} is the aggregate, upload video bitrate. "Pros only" refers to a restricted sample of channels in the top 5% of popularity. "Heckman" indicates the inclusion of the control function derived from column 1, which controls for selection. Bootstrapped standard errors in last column with 500 replications. Stars: *** significant at the 1% level; ** at 5%; * at 10%.

ity of starting a stream at any given t by about 0.2% from a baseline of about 0.7%. Partners are more likely to start a broadcast because the promise of a transcoder increases expected viewership, and because viewers bring about monetary compensation.

In sum, for both partners and non-partners, the elasticity of supply with respect

Table 3: Noise estimations

Dep var: $\log \sigma_{jt}$	Models (1)	(2)	(3)	(4)	(5)
$\log B_t^{Upload}$	0.003	0.958^{***}	0.594	0.780^{***}	0.760^{***}
	(0.003)	(0.142)	(0.382)	(0.081)	(0.076)
Pop1	-0.062^{***}	-0.060^{***}	-0.042^{**}	-0.050^{***}	-0.051^{***}
	(0.004)	(0.006)	(0.018)	(0.004)	(0.004)
Pop2	0.022^{***}	0.015^{**}	0.022	0.021^{***}	0.022^{***}
	(0.004)	(0.006)	(0.018)	(0.004)	(0.003)
Pop3	0.028^{***}	-0.012^{**}	0.000	-0.007^{*}	-0.006^{*}
	(0.003)	(0.006)	(0.015)	(0.004)	(0.004)
Pop4	0.050^{***}	0.016^{*}	0.018	0.008	0.013^{**}
	(0.003)	(0.009)	(0.027)	(0.005)	(0.005)
Trend	-0.001^{***}	-0.002^{***}	-0.002^{***}	-0.002^{***}	-0.002^{***}
	(0.000)	(0.000)	(0.001)	(0.000)	(0.000)
Heckman					-0.309^{***}
					(0.042)
cons	-2.432^{***}	-17.664^{***}	-11.631^{*}	-14.818^{***}	-14.518^{***}
	(0.050)	(2.270)	(6.081)	(1.294)	(1.214)
Channel FE	yes	yes	yes	yes	yes
2SLS	.	yes	yes	yes	yes
Subsample of	.	Pros	NonGaming	.	.
Obs	38,283,191	15,861,549	2,649,865	38,283,191	38,283,191

Notes: Standard errors in parentheses, clustered at channel id. Time series of 10-minute windows. B_t^{up} is the aggregate, upload video bitrate, which is being instrumented with the Xbox One shock in the 2SLS columns. "Pros only" refers to a restricted sample of channels in the top 5% of popularity. "Heckman" indicates the inclusion of a control function, based on Heckman (1979) and Olsen (1980), which controls for selection. First stages, for both selection-correction and 2SLS, are shown in Table 2. Bootstrapped standard errors in last column with 500 replications. Stars: *** significant at the 1% level; ** at 5%; * at 10%.

to viewers is about 1; a 1% increase in viewers increases the number of online channels by about 1%. Together with the demand-side results, the estimation implies the broadcaster side is more elastic to viewers than the viewer side is to broadcasters. Therefore, for platform growth, efforts to attract viewers have more impact than efforts to attract broadcasters.

5 Counterfactuals

Transcoding is valuable to consumers and, by extension, to producers. However, providing transcoders increases supply, which increases congestion. The platform might be interested in assessing the importance of congestion externalities and how they interact with the level of transcoding infrastructure. To the extent that the plat-

TABLE 4: DEMAND ESTIMATIONS

	MODELS		
DEP VAR: $\log\left(s_{jt}/s_{0t}\right)$	(1)	(2)	(3)
$\log \sigma_{jt}$	-0.001^{***}	0.166^{***}	-1.677^{***}
	(0.000)	(0.041)	(0.125)
TRANSCODER	-0.051^{***}	0.328^{***}	0.291^{***}
	(0.004)	(0.011)	(0.014)
PARTNER	0.052^{***}	-0.188^{***}	-0.266^{***}
	(0.009)	(0.034)	(0.046)
FEATURED	0.346^{***}	1.837^{***}	1.085^{***}
	(0.048)	(0.312)	(0.235)
log TENURE	-0.018^{***}	0.029^{***}	-0.184^{***}
	(0.002)	(0.006)	(0.012)
WEEKEND	0.132^{***}	0.054^{***}	0.174^{***}
	(0.001)	(0.001)	(0.002)
AFTER 6PM	0.082^{***}	0.050^{***}	0.308^{***}
	(0.002)	(0.002)	(0.005)
$\log s_{jt\|g}$	0.807^{***}	-0.048^{***}	0.320^{***}
	(0.001)	(0.007)	(0.010)
HECKMAN			16.661^{***}
			(0.253)
CONS	-5.284^{***}	-13.718^{***}	-11.874^{***}
	(0.015)	(0.116)	(0.332)
CHANNEL FE	YES	YES	YES
2SLS	NO	YES	YES
OBS	38,283,191	38,283,191	38,283,191

Notes: Standard errors in parentheses, clustered at channel id. Time series of 10-minute windows. "Heckman" indicates the inclusion of a control function, based on Heckman (1979) and Olsen (1980), which controls for selection. First stages, for both selection-correction and 2SLS, are shown in Table 6. Instruments for $\log \sigma_{jt}$: Xbox One broadcasting availability, PoP server upgrades. Instruments for $\log s_{jt|g}$: channels in English, from the US, with transcoders, featured, and partnered, as percentages of group g, excluding j. Excluded variables for Heckman correction: if a channel was interrupted, and $\log n^V_{jt}$ for partners and non-partners. Bootstrapped standard errors in last column with 500 replications. Stars: *** significant at the 1% level; ** at 5%; * at 10%.

form can price both sides of the market, the effects of a Pigouvian tax on each side might also be of interest. Moreover, if the platform can charge broadcasters differentially, it might also be interested in extracting rents from content providers and in its implied equilibrium effects.

The counterfactuals will consist of three regimes: neutral, preferential, and preferential with a rent-extractive platform. In the neutral regime, transcoders will be allocated randomly through the population.

In the preferential regime, which represents the platform's status quo, channels will be ranked by their quality and transcoders will be given to the highest-ranked

TABLE 5: SUPPLY ESTIMATIONS

DEP VAR: $\text{Online}_{j,t+1} = 1$	MODELS (1)	(2)
$\log n^V_{jt}$ (NON-PARTNER)	0.311^{***}	0.084^{**}
	(0.046)	(0.033)
$\log n^V_{jt}$ (PARTNER)	0.230^{***}	
	(0.019)	
PARTNER^1	0.517^{**}	
	(0.212)	
PARTNER^0	0.308^{***}	
	(0.035)	
$\log \text{TENURE}^1$	-0.548^{***}	-0.442^{***}
	(0.044)	(0.025)
$\log \text{TENURE}^0$	-0.605^{***}	-0.623^{***}
	(0.028)	(0.023)
WEEKEND^1	0.035^{***}	-0.015
	(0.009)	(0.029)
WEEKEND^0	0.011	0.150^{***}
	(0.014)	(0.022)
PEAK HOURS^1	-0.122^{***}	-0.069^{*}
	(0.025)	(0.038)
PEAK HOURS^0	0.708^{***}	0.638^{***}
	(0.035)	(0.037)
ONLINE_{jt}	6.219^{***}	6.667^{***}
	(0.333)	(0.124)
XBOX^1	0.039^{***}	0.415^{*}
	(0.013)	(0.230)
XBOX^0	-0.030^{*}	-0.880^{***}
	(0.017)	(0.230)
XBOX, GAMING^1		-0.218
		(0.231)
XBOX, GAMING^0		0.904^{***}
		(0.227)
CONTROL F., NON PARTNER	-0.067^{***}	0.183^{***}
	(0.017)	(0.025)
CONTROL F., PARTNER	-0.103^{***}	
	(0.015)	
CHANNEL FE	YES	YES
MODEL	LOGIT	LOGIT
SAMPLE OF	PARTNERS	NONPARTNERS
OBS	20,134,391	20,650,762

Notes: Bootstrapped standard errors in parentheses with 200 replications, panel defined at channel id. Time series of 10-minute windows. Logit models for j's decision to be online at t. Variables "X^1" refer to X_{jt} when j is online; "X^0" refer to X_{jt} when j is offline. The control function is the residual of a regression of $\log n^V_{jt}$ on $\log \sigma_{jt}$, indicators for transcoding and features, fixed effects, and the second-stage covariates; first stages are shown in Table 7. Sample of Partners refers to all broadcasters that eventually were partnered; NonPartners refers to a random sample of the rest. Coefficients already rescaled to achieve a standard logit's variance (Guevara and Ben-Akiva, 2012). Stars: *** significant at the 1% level; ** at 5%; * at 10%.

channels. As a proxy for quality, I use the residual of a regression of the estimated fixed effects from the demand-side model on the observable time-invariant characteristics of the channel: average bitrate, average log noise-to-signal ratio, average viewers, log tenure, cost of internet in home country, and dummies for the gaming category, English language, US country, eventually becoming a partner, and being in the pro category.

Finally, in the preferential regime with a rent-extractive platform, transcoders will also be given to the highest-ranked channels, but the platform will charge them just enough to offset the benefits accrued from transcoding. This last regime is featured in the next section because it has important parallels with the net-neutrality debate.

In all counterfactuals, I focus on transcoders and turn off the partnerships. For a given regime, label as c the counterfactual of guaranteeing enough transcoders for c% of the population. For example, the platform could have resources to guarantee transcoders to $c = 1\%$ of the population. I consider $c \in \{0, 1, \ldots, 100\}$.

Appendix section A defines the equilibrium and ensures, in a practical sense, existence. I do not impose an equilibrium-selection mechanism. Instead, I randomly select starting points, and then find a local fixed point by iterating on the equilibrium conditions. I draw 10,000 different starting points for each counterfactual c.

For each regime and for each counterfactual c, the algorithm to find the equilibrium outcome is as follows. Take a random sample $J \subset \mathcal{J}$ of channels. For each simulation $s = 1, \ldots, S$:

1. Allocate partnerships to c% of J according to the given regime.
2. Given observed probabilities of being online, simulate a starting point of initial online/offline status, $\{y_{j0}^s\}$. Simulate relevant covariates taken from observed distributions. Simulate structural, random shocks $\varepsilon_j^{IO,s}$, $\varepsilon_j^{B,s}$ and $\varepsilon_j^{\sigma,s}$, $\forall j \in J$.
   (a) Given two of {Demand, Supply, Congestion}, predict the third one and update. That is, given $\{y_{j0}^s, \sigma_j^s\}$, shocks, and covariates, predict $\{n_j^{V,s}\}$. Then, given $\{n_j^{V,s}, \sigma_j^s\}$, predict $\{y_j^s\}$. Then, given $\{n_j^{V,s}, y_j^s\}$, predict $\{\sigma_j^s\}$.
   (b) Repeat 2a until convergence.

For each s, I calculate the consumer surplus, total channels online, total viewers online, the average noise-to-signal ratio, and the average quality of online channels. I perform $S = 10,000$ simulations per counterfactual c, and I draw a random sample of $\#J = 50,000$ channels from the population for the simulations. I assume each channel has a weight of 20 for the total population size to be 1,000,000.

Benchmark comparison: Neutrality vs. Status Quo. Figure 6 shows the results for the baseline comparison between the neutral and the preferential regime. The X-axis shows the estimated cost of purchasing enough transcoders to guarantee one for c% of the population, up to $250 million, the face-value cost of acquiring 1,000,000 transcoders. For clarity, all figures show a fractional polynomial fit that runs through 101 points; each point represents a different counterfactual $c \in \{0, 1, \ldots, 100\}$. The platform's actual infrastructure is approximately $c = 1$.

In this benchmark case, neutrality yields a lower consumer surplus by construction. Because transcoders are scarce, their random allocation must yield a less efficient outcome than an allocation that favors the most proficient channels. However, this benchmark is useful to quantify a drop of up to 10% in consumer surplus, measured in utils, caused by an immediate shift to neutrality. Assuming an opportunity cost for viewers of 7.25USD/hour, 30 seconds of ads cost viewers about 6 cents. On average, 10 minutes of streaming imply 30 seconds of ads. We can thus assume viewers are willing to pay at least 6 cents for 10 minutes. Up to 40,000 viewers may drop due to an immediate shift to neutrality. Thus, consumer welfare decreases by about 2,400USD each 10 minutes in the worst-case scenario.

Under the preferential regime, the first transcoders are allocated to the channels with higher quality. The largest gains are here, because those channels make the best use of transcoders by reaching more viewers. These initial gains offset the congestion externalities from increased traffic. However, as the platform expands its transcoding infrastructure, and transcoders become less scarce, the relative efficiency gains shrink, because the marginal channel's quality starts to decrease while the noise-to-signal ratio keeps increasing. In effect, the congestion externality increases the opportunity cost of transcoders.

The counterfactual exercises have two main caveats. First, the nested logit model might overestimate gains from variety. Therefore, viewership and consumer-welfare estimations should be interpreted as an upper bound, especially for the neutral regime, where the number of broadcasters is higher. Second, for the preferential regime, my allocation of transcoders is based on a quality proxy. In reality, the platform has better information and can thus allocate transcoders optimally. Therefore, the preferential regime should be interpreted as a lower bound on consumer welfare and viewership.

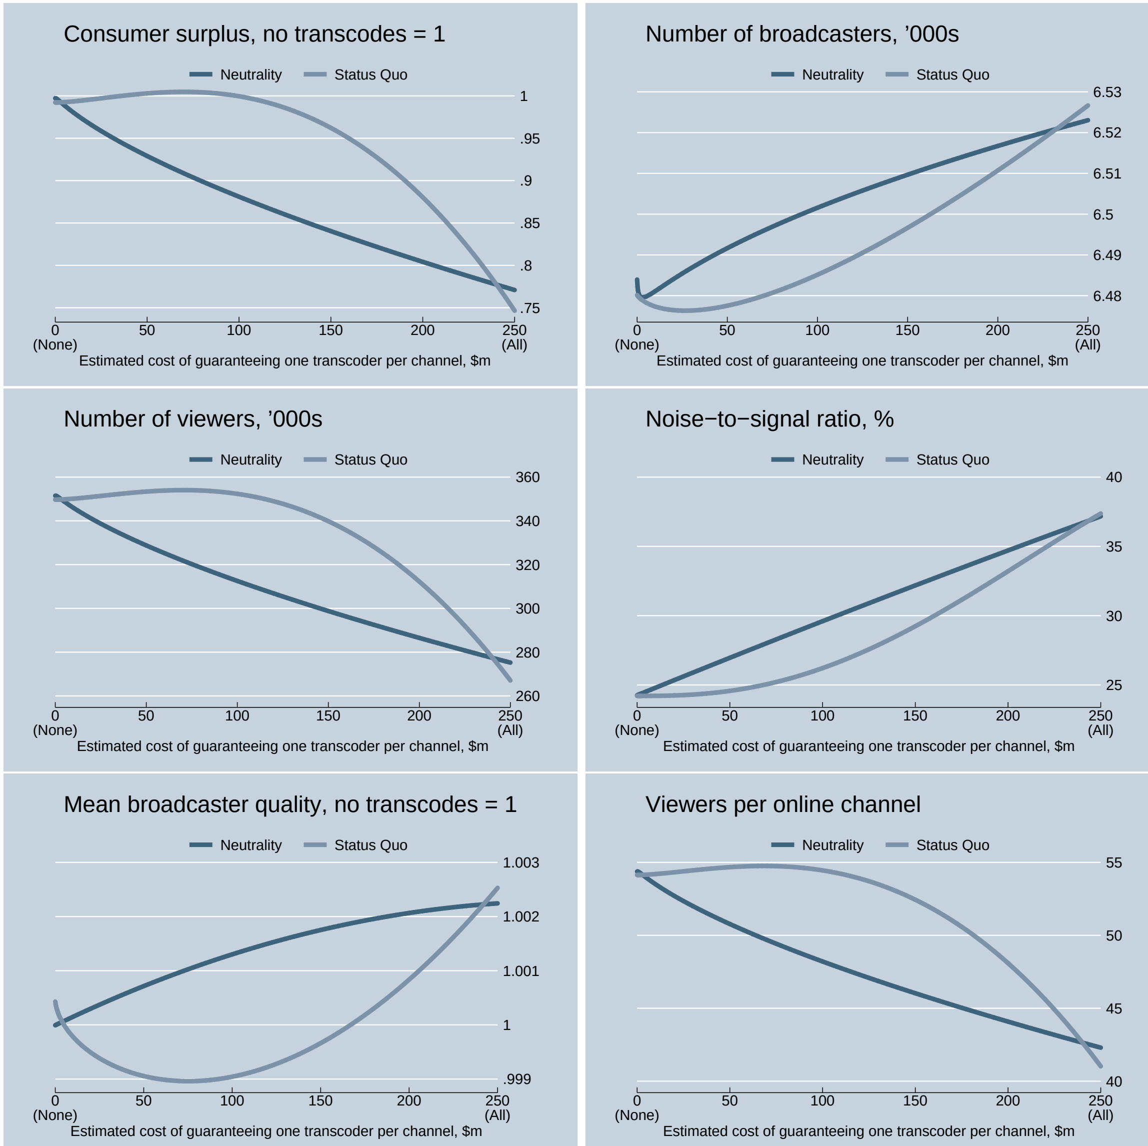

Note: Fractional polynomial curves fitted over averages of counterfactual simulations. Averages computed from 10,000 equilibrium simulations.

FIGURE 6: Neutral vs. preferential regimes

Pigouvian taxation. Congestion externalities play a major role in the growth of the platform. Figure 7 compares the preferential regime with and without congestion externalities. When the externality is turned off, the noise-to-signal ratio is assumed to be constant and equal to its 10th percentile in the data, regardless of aggregate traffic. The exercise shows that if not for congestion, participation in the platform would be significantly higher.

Interestingly, viewership, which is directly affected by the congestion externality, increases substantially without congestion, up to 100%. However, supply slightly decreases, by about 3%. Thus, with congestion, demand is inefficiently drawn toward those channels that are less valuable but that have a lower noise-to-signal ratio, prompting an over-supply of said channels. In contrast to Berry, Eizenberg and Waldfogel (2016), where excessive entry occurs in radio markets (with no congestion externalities), here, excessive entry has a high opportunity cost due to congestion.

Assuming perfect information, the platform can tax viewers to offset the congestion externality. As an example, Figure 8 compares the equilibria with and without a consumer's Pigouvian tax, equivalent to a (partial-equilibrium) reduction of 5% in demand. In equilibrium, viewership decreases, supply increases, but congestion is at similar levels. Content providers are better off, but viewers are not. Note, however, that uploaded traffic that is never consumed is a significant source of congestion. Taxing consumers would not alleviate this channel. Also, the two-sided structure matters: A dollar taxed on the consumer side does not have the same effect as a dollar taxed on the producer side. In particular, the estimations imply cross-side network externalities are larger for broadcasters than for viewers. That is, the platform's size decreases by less if broadcasters rather than viewers are taxed. I revisit this idea in the next section, where I discuss net neutrality.

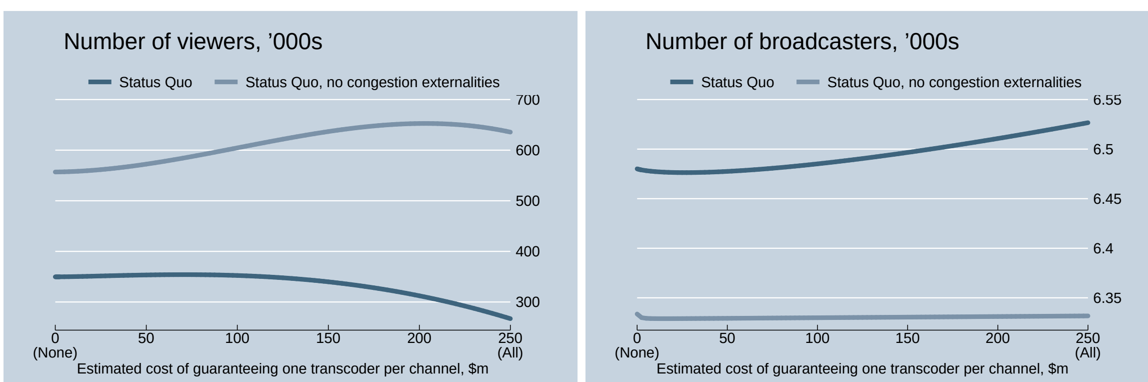

FIGURE 7: Preferential regime with and without congestion externalities

[14]See Lee and Wu (2009) and the Federal Communications Commission's website, `fcc.gov`.

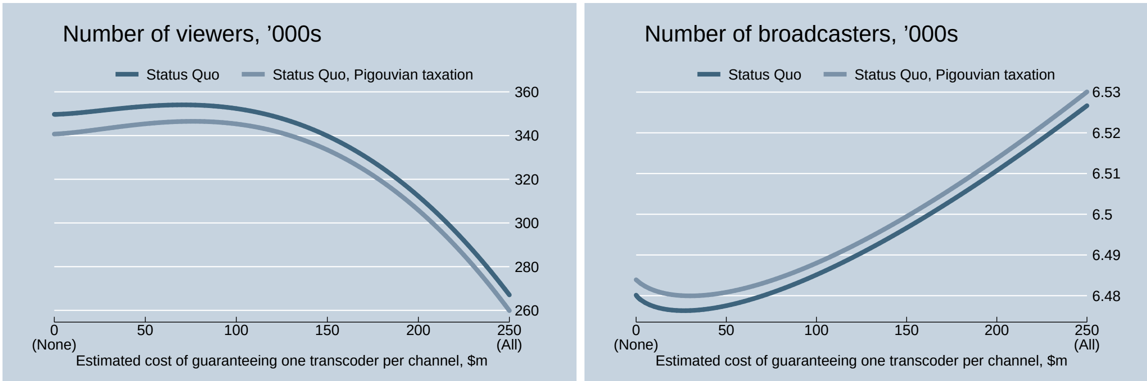

FIGURE 8: Preferential regime with and without Pigouvian taxation on viewers

5.1 Net neutrality

Net neutrality is the principle that internet service providers (ISPs) and governments should treat all data on the internet the same, not prioritizing traffic, charging differentially by priority status, or imposing congestion charges.[14] Proponents of net neutrality argue prioritization risks being concentrated over larger firms, which will decrease the variety of content provision and thus decrease consumer welfare. Proponents of net neutrality also worry about ISPs using second-degree price discrimination as a device to extract rents from innovators and content providers. For instance, ISPs might offer content providers the option to purchase priority, and could intentionally slow down delivery of data from content providers without priority. On the margin, a small fee could deter entry of small firms and start-ups. By prohibiting prioritization, net neutrality "levels the playing field," and entry becomes easier for content providers at large.

However, although net neutrality encourages entry, it also creates congestion externalities from the increase in data traffic. Critics of net neutrality argue antitrust enforcement or more limited regulatory mechanisms provide a better framework for addressing competitive concerns.[15] Moreover, by differentiating between types of traffic, network managers can avoid congestion and improve utilization and quality of high-bandwidth services. Thus, even without a formal pricing system, prioritization can boost the quality of service, which will encourage content provision and investment in innovations that take advantage of improved network performance. Finally, if ISPs can extract rents from content providers in exchange for a "fast lane,"

[15]See Becker, Carlton and Sider (2010).

they would have a greater incentive to invest in higher-quality infrastructure, thus providing faster and more reliable internet service.

In recent years, policymakers and the general public have escalated the debate about net neutrality. The FCC's website has received more than 23 million comments since April 2017.[16] Europe's telecommunications regulator held a public consultation in 2016 and gathered more than 500,000 comments in six weeks—the previous record was less than 100.[17] Published on August 30, 2016, its final guidelines strongly protect net neutrality in the European Union.[18]

Lack of data is the major challenge in empirically analyzing net neutrality. With the exception of Nurski (2012), virtually no empirical evidence in favor or against net neutrality is available (Greenstein, Peitz and Valletti, 2016). Twitch presents itself as a useful laboratory to study how net neutrality affects consumer welfare. Twitch users face trade-offs similar to those of internet users. Consumers decide what to consume at an opportunity cost of their time. Content providers decide what and when to produce, and their willingness to do so depends on their audience penetration.

Admittedly, Twitch is not likely to be representative in terms of the elasticity of its content supply. On the internet at large, potential entrants face high sunk costs of entry and high fixed costs. On Twitch, entry is free, fixed costs are low, and supply is generally unconcentrated and of low quality.

That being said, marginal costs are similar between Twitch and the internet at large. Consider Netflix, which has high sunk and fixed costs of production, but low marginal costs of distributing its data once they are produced. Moreover, Twitch features important parallels with the internet, such as a power law in the popularity of channels akin to the power law found on the internet (Newman, 2005; Gabaix, 2009; Pires and Simon, 2015). Figure 9 compares the internet with Twitch when popularity is plotted against ranking. Finally, as a subset of the internet, Twitch data flow through the same pipes on the ground that ISPs and the general public use. Moreover, video of any kind accounts for 70% of web traffic (Cisco, 2015). Thus, the estimates are revealing for a substantial proportion of internet content.

[16] Docket 17-108, "Restoring Internet Freedom," fcc.gov. Around 150,000 of those comments were received within 24 hours after HBO's John Oliver urged his viewers to intervene (usatoday.com). Shortly thereafter, the FCC's website was subject to a distributed denial-of-service attack (cnbc.com).

[17] See berec.europa.eu, medium.com.

[18] See savetheinternet.eu.

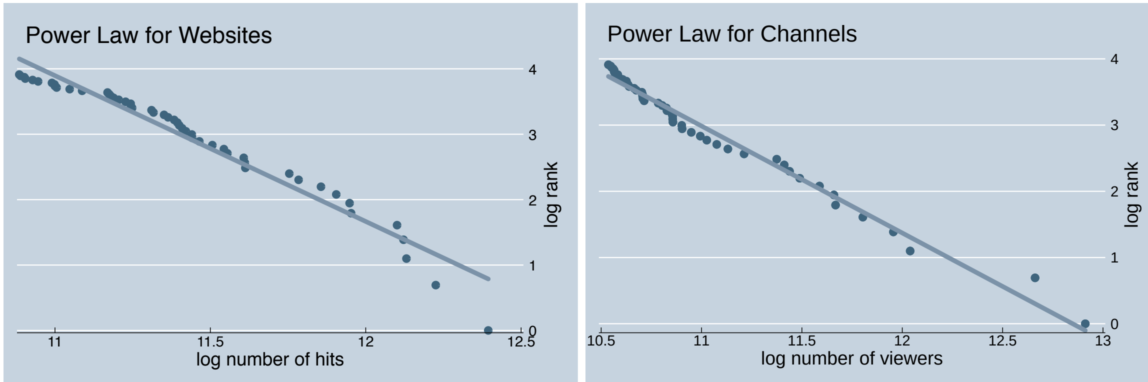

Note: Power law comparison for websites (left) and broadcasters (right). Size taken as unique visits for websites and peak viewers for broadcasters. Website data from comScore ranking of June 2017 (`comscore.com`).

FIGURE 9: Power laws on the internet and the platform

Using Twitch as an analogy for an ISP, the most popular channels correspond to the large content providers such as Netflix or YouTube, whereas the least popular channels correspond to small websites and amateur producers. Twitch is not a neutral platform, because transcoders are allocated based on popularity. Essentially, a transcoder works as a fast lane that allows content to reach viewers with low connection speeds. However, neutrality would require that the platform allocates transcoders without taking into account the identity of the broadcaster.

Rent-extractive platform vs. neutrality. In this counterfactual experiment, I explore a monopolistic platform with all the bargaining power, arguably the worst-case scenario for content providers. I assume the platform has perfect information and that it charges content providers just enough to leave them indifferent between being transcoded and not. In the absence of prices, the assumption implies the existence of quantities that correspond to the platform's chosen prices.

Figure 10 shows the main results. As expected, neutrality encourages content provision. However, in equilibrium, congestion is higher with neutrality. When transcoders are scarce, larger gains accrue from the efficient use of the infrastructure. But when transcoders become less scarce, they incentivize over-supply.

Investment. The theoretical literature emphasizes four lines of net-neutrality re-

[19] See Choi and Kim (2010), Economides and Hermalin (2012).

[20] See Economides and Tåg (2012), Caves (2012), Greenstein, Peitz and Valletti (2016).

[21] See Greenstein, Peitz and Valletti (2016).

[22] See Bourreau, Kourandi and Valletti (2015), Choi and Kim (2010).

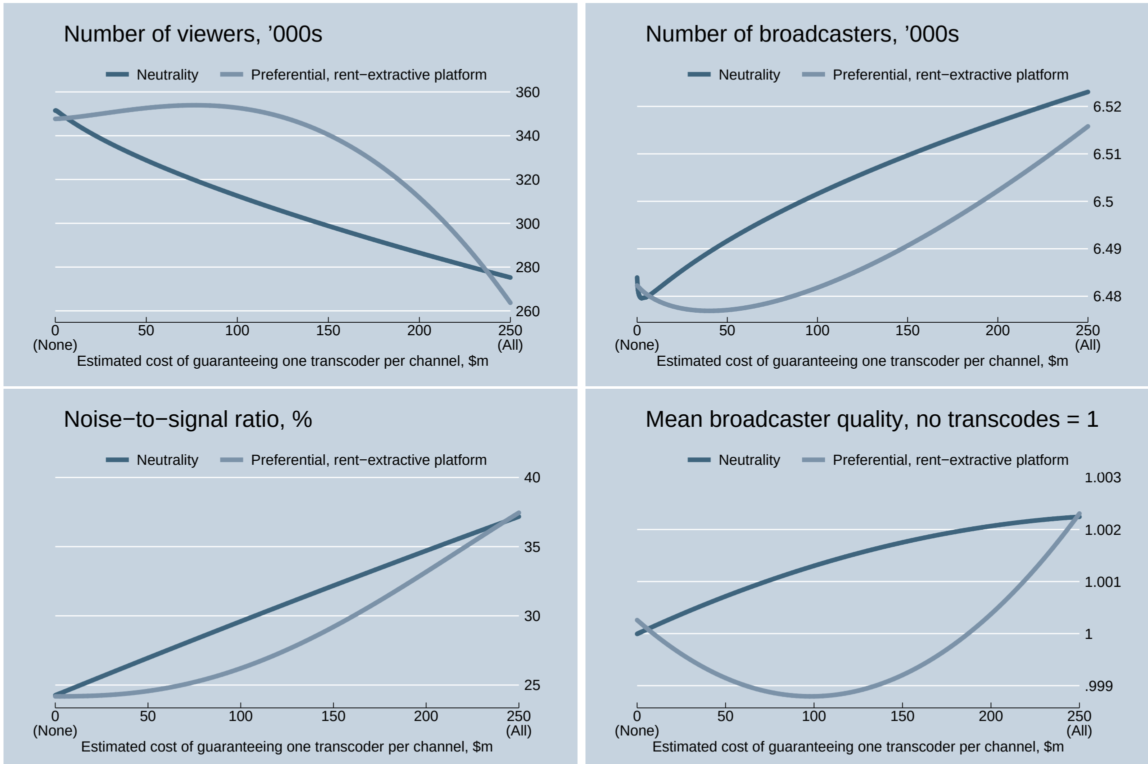

FIGURE 10: Neutral regime vs. preferential regime with rent-extractive platform

search: (1) congestion,[19] (2) the two-sided structure of the internet,[20] (3) the heterogeneity of consumers and content providers,[21] and (4) investment.[22] The structural model explicitly includes points 1–3, which gives us a framework to discuss point 4.

Consider the platform's incentives to invest. Twitch receives revenue mainly from ads and second from subscriptions. Both sources of revenue are proportional to the total number of viewers in the platform. Therefore, even if the actual optimization problem of the platform (or ISPs) is unknown, revenues as a function of infrastructure are proportional to the number of viewers in Figure 10.

Therefore, when transcoders are scarce, a rent-extractive platform has incentives to invest. When transcoders are abundant, the platform has incentives to disinvest. However, under net neutrality, the platform has incentives to disinvest at all levels of infrastructure.

Pigouvian taxation on the supply side. To the extent that the platform can charge broadcasters differentially, it might be interested in the effects of a Pigouvian tax to alleviate congestion. The rent-extractive regime can be reinterpreted as a platform that taxes popular channels and thereby the largest sources of traffic. Note, however, that this type of taxation is not allowed under net-neutrality rules.

Figure 11 shows that compared with the status quo, taxing the most popular channels (by an amount equivalent to the rents of transcoding) has virtually no effect on participation on either side of the market, at low levels of infrastructure. This result is driven in part by the tax falling on the market side with higher cross-side network elasticities. In this example, the platform is better off, because tax revenues are positive.

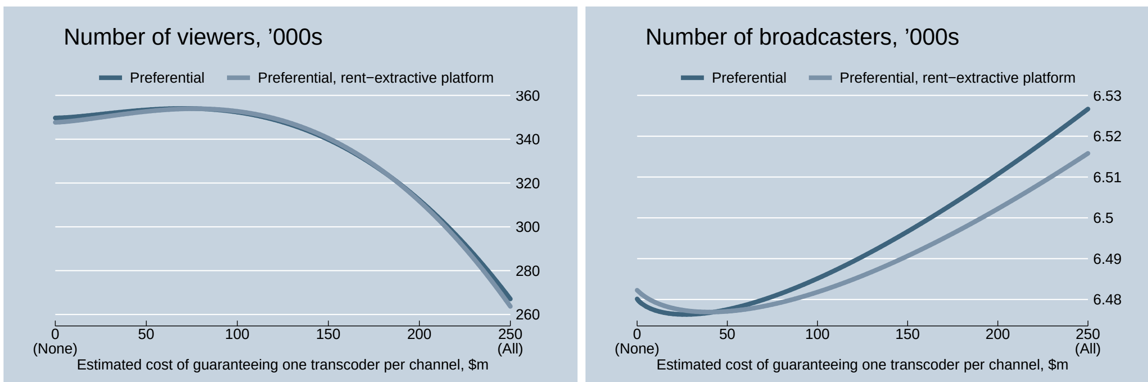

FIGURE 11: Preferential regime vs. preferential regime with rent-extractive platform

Heterogeneity. The internet at large features more heterogeneous content providers than Twitch. To explore the effects of heterogeneity, I consider a mean-preserving spread of quality. By adding a mean-zero shock, I double the variance of the content providers' quality. Then, I compute the counterfactual equilibrium outcomes, and I compare the results with the outcomes of the neutral regime and the rent-extractive platform regime.

Figure 12 shows the results. Each line represents the difference between (1) the equilibrium with a mean-preserving spread of quality and (2) the equilibrium with the baseline quality. Increasing heterogeneity increases viewership under both neutral and preferential regimes. However, the preferential regime achieves that outcome not because of more content providers, but because of the increase in the quality of online channels. Intuitively, when heterogeneity on the supply-side increases, the preferential regime allocates transcoders more efficiently, taking advantage of the fatter tails of the distribution. Therefore, if heterogeneity among content providers is pronounced, as in this extreme example, we would expect the benefits of prioritization to increase.

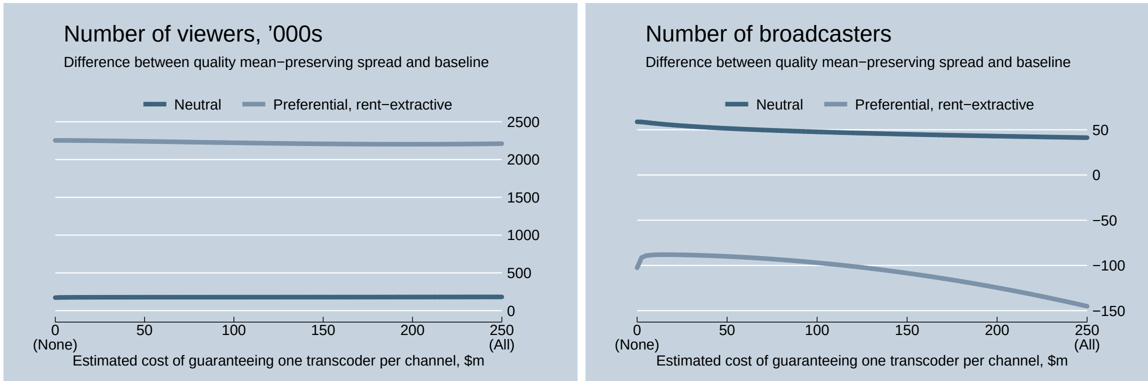

FIGURE 12: Effects of increasing quality heterogeneity

6 Concluding remarks

Internet platforms are interested in quantifying the effects of congestion. Moreover, because platforms can set prices on both sides of the market, they are also interested in the effects of pricing differentially on one side or the other. This paper finds Twitch could be twice as big if not for congestion externalities, and that a Pigouvian tax on the supply side is preferred to one on the demand side.

What can we learn about net neutrality and the internet at large? Intuitively, the results of this paper are supported by the empirical finding that consumers value quality over variety. For the results to flip, consumers in Twitch would need to have substantially different preferences than consumers on the internet, who would need to value variety much more. However, 82% of U.S. adults say they sometimes or always read online reviews for new purchases, suggesting consumers actively search for quality, at least for physical goods.[23]

On the technological side, video could reach 82% of total online traffic by 2020 (Cisco, 2015). Thus, one could think of Twitch as an approximation of online traffic or as a near-future state of the world. As the industry leader, Twitch boasts state-of-the-art infrastructure, built to minimize congestion problems. The internet at large, therefore, is likely to be even more susceptible to congestion.

Twitch offers a setting that features important trade-offs that parallel those faced by internet users. Moreover, net neutrality is easily translated into this setting, because Twitch has scarce resources to provide priority, and its network has a limited capacity. Although Twitch does not equate to the whole internet, potential lessons

[23]See www.pewinternet.org.

can learned from this first step.

This paper asks if net neutrality is beneficial to consumers. The results suggest an immediate shift to neutrality could cost up to 10% in consumer surplus. In the worst-case scenario, the platform extracts surplus from content providers and leaves them indifferent between having preferential treatment and not having it. In this case, net neutrality does increase content provision. However, consumer surplus is still lower under neutrality for plausible levels of infrastructure. Moreover, neutrality erodes the platform's incentives to invest in infrastructure.

On the internet at large, the policymaker could observe that fast lanes are, in fact, scarce. Only large platforms and internet service providers would have the infrastructure to support net neutrality at no cost to consumers. With scarce resources, however, the congestion generated by net neutrality will likely outweigh the gains in content provision.

References

Becker, Gary S., Dennis W. Carlton, and Hal S. Sider. 2010. "Net neutrality and consumer welfare." *Journal of Competition Law and Economics*, 6(3): 497–519.

Berry, Steven, Alon Eizenberg, and Joel Waldfogel. 2016. "Optimal product variety in radio markets." *RAND Journal of Economics*, 47(3): 463–497.

Berry, Steven T. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *The RAND Journal of Economics*, 25(2): 242–262.

Berry, Steven T., James Levinsohn, and Ariel Pakes. 1995. "Automobile Prices in Market Equilibrium." *Econometrica*, 63(4): 841–890.

Bourreau, Marc, Frago Kourandi, and Tommaso Valletti. 2015. "Net neutrality with competing internet platforms." *Journal of Industrial Economics*, 63(1): 30–73.

Cardell, N. Scott. 1997. "Variance Components Structures for the Extreme-Value and Logistic Distributions with Application to Models of Heterogeneity." *Econometric Theory*, 13(02): 185.

Caves, Kevin W. 2012. "Modeling the welfare effects of net neutrality regulation: A Comment on Economides and Tåg." *Information Economics and Policy*, 24(3-4): 288–292.

Choi, Jay Pil, and Byung-Cheol Kim. 2010. "Net neutrality and investment incentives." *The RAND Journal of Economics*, 41(3): 446–471.

Cisco. 2015. "The Zettabyte Era: Trends and Analysis."

Duranton, Gilles, and Matthew A Turner. 2011. "The Fundamental Law of Road Congestion: Evidence from US Cities." *American Economic Review*, 101(6): 2616–2652.

Economides, Nicholas, and Benjamin E. Hermalin. 2012. "The Economics of Network Neutrality." *RAND Journal of Economics*, 43(4): 602–629.

Economides, Nicholas, and Joacim Tåg. 2012. "Network neutrality on the Internet: A two-sided market analysis." *Information Economics and Policy*, 24(2): 91–104.

Gabaix, Xavier. 2009. "Power Laws in Economics and Finance." *Annual Review of Economics*, 1(1): 255–294.

Genakos, Christos, Kai Uwe Kühn, and John Van Reenen. 2018. "Leveraging Monopoly Power by Degrading Interoperability: Theory and Evidence from Computer Markets." *Economica*, 85(340): 873–902.

Greenstein, Shane, Martin Peitz, and Tommaso Valletti. 2016. "Net Neutrality: A Fast Lane to Understanding the Trade-offs." *Journal of Economic Perspectives*, 30(2): 127–150.

Guevara, Cristian Angelo, and Moshe E. Ben-Akiva. 2012. "Change of Scale and Forecasting with the Control-Function Method in Logit Models." *Transportation Science*, 46(3): 425–437.

Heckman, James J. 1979. "Sample Selection Bias as a Specification Error." *Econometrica*, 47(1): 153–161.

Huynh-Thu, Q., and M. Ghanbari. 2008. "Scope of validity of PSNR in image/video quality assessment." *Electronics Letters*, 44(13): 800.

Jehle, Geoffrey, and Philip Reny. 2011. *Advanced Microeconomic Theory.* Prentice Hall.

Lee, Robin S. 2013. "Vertical Integration and Exclusivity in Two-Sided Markets." *The American Economic Review*, 103(7): 2960–3000.

Lee, Robin S., and Tim Wu. 2009. "Subsidizing Creativity through Network Design: Zero-Pricing and Net Neutrality." *Journal of Economic Perspectives*, 23(3): 61–76.

Malone, Jacob B., Aviv Nevo, and Jonathan W. Williams. 2017. "The Tragedy of the Last Mile: Congestion Externalities in Broadband Networks."

McFadden, Daniel. 1974. "Conditional logit analysis of qualitative choice behavior." In *Frontiers in Econometrics.* , ed. P. Zarembka, Chapter 4, 105–142. New York:Academic Press.

Newman, M. E. J. 2005. "Power laws, Pareto distributions and Zipf's law." *Power laws, Pareto distributions and Zipf's law. Contemporary physics*, 46(5): 323–351.

Nurski, Laura. 2012. "Net Neutrality, Foreclosure and the Fast Lane: An empirical study of the UK."

Olsen, Randall J. 1980. "A Least Squares Correction for Selectivity Bias." *Econometrica*, 48(7): 1815–1820.

Pankert, Gregory, Andrea Faggiano, and Karim Taga. 2014. "The Future of the Internet: Innovation and Investment in IP Interconnection." Arthur D. Little.

Petrin, Amil, and Kenneth Train. 2010. "A Control Function Approach to Endogeneity in Consumer Choice Models." *Journal of Marketing Research*, XLVI(February): 3–13.

Pires, Karine, and Gwendal Simon. 2015. "YouTube Live and Twitch: A Tour of User-Generated Live Streaming Systems." 1–6. Portland, Oregon:ACM.

Rochet, Jean-Charles, and Jean Tirole. 2006. "Two-Sided Markets: A Progress Report." *RAND Journal of Economics*, 37(3): 645–667.

Shivaldova, Veronika, Andreas Winkelbauer, and Christoph F. Mecklenbrauker. 2014. "Signal-to-noise ratio modeling for vehicle-to-infrastructure communications." *2014 IEEE 6th International Symposium on Wireless Vehicular Communications, WiVeC 2014 - Proceedings*.

Sundaresan, Srikanth, Xiaohong Deng, Yun Feng, Danny Lee, and Amogh Dhamdhere. 2017. "Challenges in inferring internet congestion using throughput measurements." 43–56. New York, New York, USA:ACM Press.

Weeds, Helen. 2016. "TV Wars: Exclusive Content and Platform Competition in Pay TV." *Economic Journal*, 126(594): 1600–1633.

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

Appendix For Online Publication

A Definition of equilibrium and existence

Definition 1 (Equilibrium.)**.** Let consumers have mass M. Given t and initial $y_{j,t-1}$, $\forall j \in J$, given an allocation of transcoders and partnerships, given shocks ξ_{jt}, ε_{jt}^{IO}, ε_{jt}^{B} and $\varepsilon_{jt}^{\sigma}$, $\forall j \in J$, and given parameters and covariates $\mathbf{x}_{jt}$, $\mathbf{x}_{jt}^{IO}$, $\mathbf{x}_{jt}^{B}$, $\forall j \in J$, an equilibrium is a set of demands $\{s_{jt}\}$, supplies $\{y_{jt}\}$, and noise-to-signal ratios $\{\sigma_{jt}\}$, such that $\forall j, t$,

1. Congestion levels satisfy (1);
2. Consumers maximize their utility, yielding demands according to (2);
3. Broadcasters maximize their utility, according to (3) and (4).

Lemma 1. If broadcasters are allowed to mix strategies, a Nash equilibrium exists.

Proof. First, congestion levels are well defined for any $B_t^{up} > 0$. Second, consumers are not strategic in their choices. Given any set of online channels $\mathcal{J}_t$, demands s_{jt} are well defined.

However, broadcasters are strategic. They must decide whether to be online or offline: $y_{jt} \in \{0, 1\}$. Note their payoffs, u_{jt}^{IO} and u_{jt}^{B}, are functions of n_{jt}^{V}, which are functions of s_{jt}, which are functions of $\log \sigma_{jt}$. Ultimately, $\log \sigma_{jt}$ is a function of B_t^{up}, which is a function of $\{y_{jt}\}_{j \in \mathcal{J}}$. In other words, one can write $u_{jt}^{IO}(y_{jt}, y_{-jt})$ and $u_{jt}^{B}(y_{jt}, y_{-jt})$. Because the set of strategies, $\{0, 1\}$, is finite, and $\mathcal{J}$ is finite, standard arguments (see Jehle and Reny (2011)) ensure the existence of a mixed-strategy Nash equilibrium. $\square$

A.1 Schedules in two-stage game for deciding when to Turn On

If schedule choice is endogenous and I ignore it, I expect I would be underestimating the sensitivity to viewership, especially for larger channels. "Pro" channels would not change their broadcasting decisions as much, because they are predetermined by their schedules. Here, I argue that controlling for scheduling is a reasonable solution to approximating a much more complex game.

Let $\mathcal{T}$ be the set of 10-minute periods in a week. Streamers play a simultaneous game at the beginning of the week. In a first stage, broadcasters choose a set $I_j \subseteq \mathcal{T}$

for turning on. Once they are online, in a second stage, viewership is realized and broadcasters decide to keep broadcasting or not: $y_{jt} = \mathbb{1}\{\text{Keep On}_{jt}\}$ for $\{t : t \notin I_j \text{ and } y_{j,t-1} = 1\}$. As above, one can write payoffs $\widetilde{u}(\mathbf{y}_j, \mathbf{y}_{-j}, I_j, I_{-j})$, where payoffs depend on the vector of j's decisions and the set I_j, taken as given the decisions and sets of other channels, $\mathbf{y}_j, I_{-j}$. Again, because the set of strategies is finite, and $\mathcal{J}$ is finite, a mixed-strategy Nash equilibrium exists.

This paper approximates $\widetilde{u}(\mathbf{y}_j, \mathbf{y}_{-j}, I_j, I_{-j})$ with

$$\{u^{IO}_{jt}(y_{jt}, y_{-jt}|\text{DoW}_{jt}, \text{PeakHours}_{jt})\}_{t\in\mathcal{T}} \text{ and } \{u^{B}_{jt}(y_{jt}, y_{-jt}|\text{DoW}_{jt}, \text{PeakHours}_{jt})\}_{t\in\mathcal{T}},$$

where DoW_{jt} and PeakHours_{jt} are day-of-week and peak-hour dummies, which approximate the role of scheduling.

Table 6: Demand estimations, First stages

	Models		
	(1)	(2)	(3)
Dep var:	$\text{Online}_{j,t+1} = 1$	$\log \sigma_{jt}$	$\log s_{jt\|g}$
$\log n^V_{jt}$ (Partner)	0.006***		
	(0.000)		
$\log n^V_{jt}$ (Non-Partner)	0.028***		
	(0.000)		
Interrupted	−0.023***		
	(0.000)		
Transcoder	−0.017***	0.009	0.546***
	(0.000)	(0.005)	(0.014)
Partner	0.117***	−0.037	−0.396***
	(0.002)	(0.023)	(0.036)
Featured	−0.041***	−0.015	1.544***
	(0.010)	(0.074)	(0.471)
log Tenure	0.003***	−0.050***	−0.108***
	(0.000)	(0.004)	(0.008)
Weekend	−0.009***	−0.009***	0.129***
	(0.000)	(0.001)	(0.003)
After 6pm	−0.012***	0.005***	0.259***
	(0.000)	(0.001)	(0.006)
Xbox Available	−0.002***	0.020***	0.128***
	(0.000)	(0.003)	(0.010)
Pop1	−0.008***	−0.092***	0.232***
	(0.000)	(0.003)	(0.006)
Pop2	−0.001***	0.000	−0.020***
	(0.000)	(0.003)	(0.007)
Pop3	0.000	−0.004	−0.022**
	(0.000)	(0.003)	(0.011)
Pop4	0.010***	0.045***	−0.272***
	(0.000)	(0.003)	(0.007)
$\text{English\%}_{-jt\|g}$	0.274***	0.143***	−6.047***
	(0.004)	(0.052)	(0.168)
$\text{USA\%}_{-jt\|g}$	−0.092***	−0.072***	2.562***
	(0.001)	(0.012)	(0.045)
$\text{Transcoder\%}_{-jt\|g}$	0.107***	0.005	12.006***
	(0.004)	(0.045)	(0.249)
$\text{Featured\%}_{-jt\|g}$	−0.524***	0.156	5.462***
	(0.009)	(0.147)	(0.421)
$\text{Partner\%}_{-jt\|g}$	−0.169***	−0.668***	−6.017***
	(0.007)	(0.144)	(0.347)
Heckman		−1.043***	28.075***
		(0.034)	(0.188)
cons	0.660***	−2.305***	−3.801***
	(0.003)	(0.044)	(0.138)
Channel FE	Yes	Yes	Yes
R^2_a	0.88	0.01	0.62
F-stat	57,514	261	29,391
Obs	48,570,885	38,283,191	38,283,191

Notes: Standard errors in parentheses, clustered at channel id. Time series of 10-minute windows. First column predicts channel j keeping online at time $t+1$; predictions used to contruct a control function to correct for selection bias, based on Heckman (1979) and Olsen (1980). Regressors $X\%_{-jt|g}$ refer to the number of channels with characteristic X, as a percentage of group g, excluding j. "Heckman" indicates the inclusion of the control function derived from column 1, which controls for selection. F-statistics of relevant instruments. Bootstrapped standard errors in last two columns with 500 replications. Stars: *** significant at the 1% level; ** at 5%; * at 10%.

TABLE 7: SUPPLY ESTIMATIONS, FIRST STAGES

	MODELS	
	(1)	(2)
DEP VAR:	$\log n_{jt}^V$	$\log n_{jt}^V$
$\log \sigma_{jt}$	-0.014^{***}	-0.011^{**}
	(0.003)	(0.006)
TRANSCODER	0.144^{***}	0.508^{***}
	(0.045)	(0.090)
FEATURED		3.870^{***}
		(0.088)
PARTNER	-0.040	
	(0.043)	
log TENURE	0.577^{***}	-0.011
	(0.126)	(0.033)
WEEKEND	0.030^{***}	0.079^{***}
	(0.007)	(0.017)
PEAK HOURS	0.034^{***}	0.004
	(0.013)	(0.019)
XBOX	-0.079^{***}	-0.146^{**}
	(0.020)	(0.067)
XBOX, GAMING		0.263^{**}
		(0.108)
CONS	1.624^{**}	1.165^{***}
	(0.812)	(0.154)
CHANNEL FE	YES	YES
SAMPLE OF	PARTNERS	NONPARTNERS
R_a^2	0.01	0.32
F-STAT	20	24,678
OBS	2,130,761	88,738

Notes: Standard errors in parentheses, clustered at channel id. Time series of 10-minute windows. Sample of Partners refers to all broadcasters that eventually were partnered; NonPartners refers to a random sample of the rest. F-statistics of instruments: $\log \sigma_{jt}$, Transcoder, and Featured. Stars: *** significant at the 1% level; ** at 5%; * at 10%.