



**HAL**  
open science

## Enhancing Transparency: Internet Video Quality Inference from Network Traffic

Paul Schmitt, Francesco Bronzino, Renata Teixeira, Tithi Chattopadhyay,  
Nick Feamster

► **To cite this version:**

Paul Schmitt, Francesco Bronzino, Renata Teixeira, Tithi Chattopadhyay, Nick Feamster. Enhancing Transparency: Internet Video Quality Inference from Network Traffic. TPRC 46: The 46th Research Conference on Communication, Information and Internet Policy, Sep 2018, Washinton D.C., United States. hal-02074828

**HAL Id: hal-02074828**

**<https://inria.hal.science/hal-02074828v1>**

Submitted on 21 Mar 2019

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Enhancing Transparency: Internet Video Quality Inference from Network Traffic

PAUL SCHMITT, Princeton University  
FRANCESCO BRONZINO, Inria  
RENATA TEIXEIRA, Inria  
TITHI CHATTOPADHYAY, Princeton University  
NICK FEAMSTER, Princeton University

## 1 INTRODUCTION

The 2017 FCC Restoring Internet Freedom Order [10] removes the “enhanced” transparency obligations introduced by the 2015 Open Internet Order and aims to return net neutrality policy to transparency rules based on the 2010 Open Internet Order. The ruling states that the burden of additional network performance disclosures exceed the benefits, and that the most salient metrics to report are those that involve consumer quality of experience (QoE) for the applications that they commonly use. Unfortunately, however, internet service providers (ISPs) will typically have difficulty reporting on application performance and QoE metrics, both of which are notoriously difficult to estimate from network traffic.

To address this shortcomings, we present the initial development of Network Microscope, a tool that estimates QoE for Internet video streaming from passively collected network traffic. Our system sits inline, on path, and analyzes traffic in real time as it traverses the network to (1) identify which traffic flows belong to a specific video streaming service; (2) estimate critical quality-of-experience metrics for streaming video such as: bitrate, changes in bitrate, startup delay, and rebuffering. When deployed on a commodity embedded device, the tools is suitable for deployment in consumer home networks, as well as near various network endpoints.

Because Internet video traffic accounts for majority of the global internet traffic, this approach of passively observing traffic has two significant policy implications:

- It reduces the administrative and operational burden on ISPs, because traffic collection and analysis is passive, in-line, and in homes, and does not introduce additional test traffic.
- The approach offers application QoE metrics that are complementary to the lower-level network performance metrics that ISPs already collect.

We discuss the capabilities of Network Microscope, its initial deployment to over 50 consumer homes, our initial findings concerning the reporting of application QoE metrics, and broader implications for policy surrounding ISP transparency reporting requirements. Network Microscope can provide a deeper understanding of several concepts often discussed both in the context of net neutrality and encouraging competitive forces in the market.

First, such a tool can shed more light on the nature of streaming traffic from consumers, including which video streaming services are most popular with consumers, and how those popular services perform on different networks. Measurements based on service-specific usage is particularly meaningful to consumers because consumers often understand their network needs better when it is tied to the applications they use often. More specific information about the performance of popular video services can help consumers make more informed choices about the network services that they purchase. Second, the tool facilitates the analysis of application performance for network traffic at multiple locations along a single end-to-end path, enabling both consumers and regulators to independently

---

Authors' addresses: Paul Schmitt, Princeton University, pschmitt@cs.princeton.edu; Francesco Bronzino, Inria, francesco.bronzino@inria.fr; Renata Teixeira, Inria, renata.teixeira@inria.fr; Tithi Chattopadhyay, Princeton University, tithic@princeton.edu; Nick Feamster, Princeton University, feamster@cs.princeton.edu.

, Vol. 1, No. 1, Article . Publication date: August 2018.

verify ISP reports about application performance. Ultimately, the type of information about application performance that our tool exposes can affect consumer decision-making; we explore and discuss these effects, and how they may ultimately interact with switching costs, market competition, and other commercial considerations.

## 2 MOTIVATION

Under the Restoring Internet Freedom Order, the older conduct rules of the 2015 Open Internet Order are eliminated. The FCC argument is that cost of complying with open internet rules outweigh the benefits. It was stated that the transparency requirements adopted by the commission are sufficient to protect consumer interests and promote innovation and infrastructure investments. While the order allows for some flexibility, ISPs are required to disclose information about blocking, throttling, affiliated prioritization, paid prioritization, congestion management, application specific behavior, device attachment rules and practices used to ensure end user security. ISPs are also required to disclose information about performance characteristics such as service description and the impact of specialized services on performance and disclose commercial terms such as price, privacy policy and redress options. Current FCC rules require that ISPs disclose this information either on their website, or file with the FCC. However, if these disclosures are meant to protect consumers, help them make informed choices, and boost provider investment, understanding the consumer and producer decision making processes while presenting these disclosures is crucial.

In this part of the study, we focus on a few salient consumer and provider decision making features that have been extensively studied in other markets where disclosures are often the most popular method of promoting informed consumer choice.

### 2.1 Designing an effective tool

#### **Important considerations in the consumer decision making process:**

- **Limited attention of consumers:** There is growing literature in behavioral economics to suggest that consumer attention is a scarce resource, and the format and method of presentation of information has a direct impact on their elasticity of demand. There are several key factors that have already been found to be important when it comes to disclosures. The importance of simplifying the information along with personalization or the presentation of the most valuable information to customer is crucial to making these transparency rules effective. This is particularly important when it comes to buying a product like internet service that is extensively used, and is most often paid for in the form of monthly bills. If the purchase of a product or service is considered small and repetitive like that of internet service, consumers are often less likely to pay attention to all the available information and important details of the commercial terms [7]. Additionally, Kling et al. found that even when information was readily available, personalization and providing a menu of choice was a critical to improving consumer choice particularly in markets where the goal is to use disclosures to reduce government interventions or regulations [17].
- **Cognitive considerations:** Literature in several other industries like the financial industry where disclosures are typically required by providers indicate that disclosures are more effective when cognitive biases and other limitations of consumers decision making processes are taken into consideration. The importance of how data is presented (ranking, comparative dashboards etc.) has a direct impact on consumer demand [19].
- **Comparative information:** In the absence of regulations, comparative information is particularly useful in helping consumers understand their choices [5, 12]. In their study of a student's school choice, Hastings & Weinstein 2008 find that comparative performance information of schools had a significant impact on school choice [13]. In the absence of performance information, other factors that are more directly observable had a bigger impact on choice.

- **Aggregated information:** Two factors that are most important to consumers while making a choice is the price/commercial terms of the product and the quality of the product. Even when information is readily available to consumers about both price and the quality of product, consumers often struggle to aggregate the available information. The importance of presenting aggregated information is also said to have an impact on a consumer's demand elasticity [11].
- **Peer information:** While we have discussed the importance of providing comparative information of the providers above, the value of gathering information from peers is equally critical to the consumer decision making process [6]. For example, in markets like the internet service market, where the consumer's ability to process any available information is highly variable, peer information can go a long way in helping consumers understand their choices.
- **Information asymmetry:** When consumers are presented with information about their internet service, there are additional factors to consider given the complex nature of some of the information that is presented. As discussed earlier consumers have varying degrees of technical expertise, and this leads to various issues like misinterpretation of observed performance due to a lack of understanding of the impact of various non ISP related causes. Similarly some consumers also do not pay attention to missing information or the presence of motivated information.

If transparency is to be an effective consumer protection tool, the above consumer decision making features point towards the importance of developing tools for measuring ISP performance. Furthermore, it is important to make these performance metrics available to consumers along with the already publicly available information about consumer terms. The availability of aggregated information provides consumers a clearer picture of what quality of service is available at the different price points. While thinking about transparency rules and enhanced informational tools, the following ways of presenting information should be considered to improve overall effectiveness of the tool.

- Provide simple and standardized information
- Provide a comparative dashboard for consumers so that they can compare performance metrics of competing providers
- Pay attention to personalization of information and provide consumers with tools to visualize their own preferred rankings
- Provide aggregated information that helps consumer access financial impact alongside product quality.

#### **Important considerations of provider motivation:**

- **Social impact of direct consumer access:** While the evidence of impact of transparency tools and disclosures on consumer choice is debatable, the impact that these tools have on provider behavior is far larger. Provider behavior is impacted when consumers have direct access to information about their performance and behavior. The fear of regulation or the fear of losing consumers to other providers promotes better behavior and performance. In the internet service market, this is likely to have an impact on provider behavior as it relates to blocking and throttling of services. If consumers have access to information about these practices, providers are unlikely to engage in behavior that could potentially result in them losing their market share.
- **Importance of third party validation:** While consumers are often reluctant to trust information disclosed by providers, the use of a third party validation tool helps improve trust between providers and consumers. This can often be beneficial to providers if they find government intervention or regulation to be administratively burdensome.
- **Importance of comparisons and standardization:** In order for transparency rules and disclosures to help boost investment, the information disclosed needs to be standardized for ease of comparison. As mentioned above, the provision of comparative information to consumers encourages better behavior and performance. The availability of simple and standardized information is helpful in promoting healthy competition and

thereby boosting investments in the technology. Consumers are also more likely to understand certain provider practices. For example: clearly available standardized information about performance could help providers justify paid prioritization on certain critical applications.

## 2.2 The need for a system

These requirements motivate the design of a system that can provide consumers and providers information necessary for understanding internet performance in a meaningful way. Likewise, the metrics collected by such a system should be aggregated and presented in a manner that allows for comparison between services. In the following section we discuss the design of a lightweight system intended for placement in the consumer network. Our tool estimates application performance metrics for internet video content, which could be used for comparison purposes.

## 3 TECHNICAL BACKGROUND

In this section, we discuss passive measurement of internet traffic and briefly introduce background on Internet video streaming. Then, we review the video quality metrics introduced in the literature and discuss the challenges of inferring these metrics from within the network.

### 3.1 Active versus passive internet measurement

Traditionally, active measurement, whereby some network endpoint tests the performance of an end-to-end network path, an application, or a segment of the network by introducing additional “probe” traffic towards some network destination and measuring properties of the measurements that are introduced, is the *de facto* method for estimating network performance. Although common, such metrics often do not correlate to the actual performance of application traffic, which is ultimately the goal according to the 2017 FCC order.

Intuitively, passive measurement offers more realistic insight into network activity as it is based on real user traffic. However, inferring performance metrics using passively observed traffic can be more challenging. For example, each service or application might exhibit different traffic characteristics as they perform poorly or well. Passive inference requires an in-depth understanding of application behavior in order to translate network traffic to application performance. We detail the challenges specific to passive estimation of internet video in Section 3.4

### 3.2 Internet video streaming

Video streaming services over the Internet have largely converged to using Dynamic Adaptive Streaming over HTTP (DASH) [26]. DASH-based systems (and its proprietary counterparts [21, 23]) rely on web protocols to retrieve “segments” of video from web servers distributed at different locations of the network. Segments are slices of time in a video that have been encoded and stored at multiple bitrates. The client player adapts the video stream during playback by requesting the segments with the appropriate bitrate / encoding combinations as determined by the application and network performance.

### 3.3 Internet video quality metrics

Video quality of experience depends on many factors that range from user expectations and context to network and device performance [16]. Given the prescribed goal of estimating QoE that users experience, our goal is to estimate common QoE metrics from network traffic. Resultingly, we focus on technical video quality metrics that depend on the network performance and have shown to affect the video quality of experience [1–3, 8, 9, 15, 18].

- **Startup delay.** This metric represents the time elapsed from the moment the player initiates a connection to a video server until the time it starts rendering video frames. Prior work has shown that higher startup delays lead to more users abandoning video sessions [18] as well as shorter total duration that users spend on a video service [8].

Service	Collected Sessions
<i>Netflix</i>	5,310
<i>YouTube</i>	10,086
<i>Amazon Prime Video</i>	6,416
<i>Twitch</i>	3,336

Table 1. Collected video sessions over seven months of a pilot deployment.

- **Video bitrate.** The bitrate of video segments is an indicator of the quality of the displayed content and is the key quality metric for live video content [8].
- **Bitrate switches.** Video players can adapt to changes in network conditions by adjusting the requested video bitrates with the goal to select the best possible bitrate for any given condition. Nevertheless, excessive bitrate switching has a negative effect on the quality of experience [2, 14].
- **Rebuffering events.** This metric captures the periods of time the video stalls because the buffer is empty.

### 3.4 Challenges facing video quality inference.

Inferring video quality metrics from encrypted network traffic presents many difficulties. First, we do not have access to the application-layer details such as the segment and session information as previously discussed in Section 3.2. For example the content being downloaded is uncertain—the length of Internet videos varies widely and users can interrupt playback, seek to a different part of a video, or change the stream to a different video altogether. The silent periods, when we observe no traffic, could correspond to the end of a short video, poor network performance that triggers a rebuffering event, or simply that the video is paused. Finally, modern DASH-based streaming services employ variable bitrate algorithms. The sizes of video segments can vary based not only on the resolution, but also on the type of displayed video. For example, a scene containing little or no movement, such as the end-credits for a movie, will result in relatively small bitrates compared with scenes the contain more motion.

## 4 DATA

We have developed and deployed a prototype of Network Microscope. This system is built for the Raspberry Pi [24] and Odroid [22] platforms, and we currently manage a deployment in the homes of volunteer users. As of July 2018, approximately 45 devices are installed in homes in the United States, and nine devices are deployed in homes France. The deployment places our measurement devices at the home network gateway. This allows us to view individual devices within the home as well as performance within the home network itself.

During the initial seven months of this pilot deployment, we have collected more than 20,000 video sessions across different major streaming services. Table 1 summarizes the data that we have collected to date.

We have collected ground truth information by recording packet traces as well as information collected from a web browser extension we developed that allows us to record the application-level video quality information we are seeking to estimate (*e.g.*, startup delay, bitrate, etc.). The ground truth dataset was collected from in-lab experiments using multiple devices while varying the network conditions in a controlled environment.

## 5 PASSIVE VIDEO QUALITY INFERENCE

In this section, we describe and evaluate methods for inferring video quality metrics described in Section 3.3 from network traffic. For brevity, we limit our focus to startup delay and video bitrate for Netflix and YouTube videos. We rely on the datasets from Section 4 to identify features from encrypted traffic that capture each of the quality metrics we infer.




Fig. 1. Relationship between segment download rate and startup delay.

Function	Netflix $R^2$	YouTube $R^2$
Exponential	0.79	0.91
Power Law	0.31	0.54
Logarithmic	0.31	0.54
Quadratic	0.21	0.47
Linear	0.07	0.23

Table 2. Goodness of fit for startup delay curve fitting functions.

## 5.1 Startup delay

A video’s startup delay is the time the video player needs to receive enough segments to begin playback with low probability of stalling or quality degradations. It is hence intuitive that there is a correlation between the download of segments and the startup delay. The blue dots in Figure 1 represent the segment download rate (expressed in segments per second) against the ground truth startup delay captured with our browser extension for Netflix and YouTube video experiments. This figure offers two insights. First, there appears to be a strong relationship between segment rate and the corresponding startup delay: as expected, a video session experiencing a poor (*i.e.*, low) segment download rate is likely to have a longer startup delay. Second, the trends for the two services seem to differ. Thus, we believe we must treat each video service separately as they may employ different algorithms for buffering and playback. Such algorithmic differences between services can be logically explained by different goals for the services: Netflix may value higher quality playback and the avoidance of video stalls at the expense of a longer startup delay whereas YouTube may prefer the opposite.

Given our initial findings, we can generate a regression line for each service based on the segment rate in order to predict the startup delay for a video. We generate optimal least-squares fit lines using the python scipy library and several fit line functions for each video service using experiments conducted in a lab environment as our training dataset. Table 2 shows the  $R^2$  values for each function we tested. We find that an exponential function maximizes the  $R^2$  values for both services, as shown in Figure 1. We posit that the exponential relationship is likely attributable to TCP’s exponential backoff mechanism due to timeouts.




Fig. 2. Startup delay estimation errors.

*Validation.* Using the fit-line formula, we estimate startup delays for in-home video experiments and compare the estimate with the ground truth collected by the browser extension. Figure 2 shows the absolute error for our startup delay estimates with video sessions binned by startup delay ranges. Overall, our estimate is within three seconds of the ground truth in 86.4% of the sessions and within five seconds in 94.6% of sessions. We observe that the error range tends to increase as the estimated startup delay increases, whereas our technique performs well (*i.e.*, median errors of 0.32, 0.81, and 1.91 seconds for startup delay ranges 1–2, 2–4, and 4–8 seconds, respectively) in cases where the startup delay is estimated to be low. We also see that error ranges for YouTube tend to be larger than those for Netflix sessions. We investigated the cause of this and attribute it to the fact that YouTube uses variable-length segments while Netflix uses fixed-length segments. For example, a session that uses variable-length segments may rapidly buffer many small segments leading to a high segment rate, yet the buffer contains a relatively short amount of playable content. However, we consider that our estimates are useful in practice as we are able to estimate with high confidence when startup delays are less than or greater than two seconds, which is an acceptable threshold under normal circumstances [18].

## 5.2 Video bitrate

The estimation of video bitrate from network traffic is challenging because video players attempt to fill a playback buffer at the beginning of a streaming session, which results in high network utilization during moments where the player is attempting to buffer content faster than it is playing the content back. Figure 3 shows this download behavior. We plot the bitrate for a video session in 5-second windows. At the beginning we see that the bitrate is high as the player fills the buffer before settling to a steady state bitrate. In order to accurately infer video bitrate throughout a session, and account for the bursty network behavior of adaptable video protocols, we seek to track video bitrate throughout the course of a video session while identifying changes in player state such as those related to buffering.

We must identify changes in player state over the lifetime of a session as instantaneous network throughput will result in large fluctuations in a bitrate estimate. Intuitively, such changes should be reflected in the segment download behavior. For instance, during steady state playback segments will be downloaded at the same rate in which the buffer empties during normal playback; whereas a buffering event will result in segments that are downloaded at a higher rate. Following this logic, we use an online Bayesian change detection algorithm [25] with segment interarrival times as the input data. The algorithm assumes our observations of the session may be divided



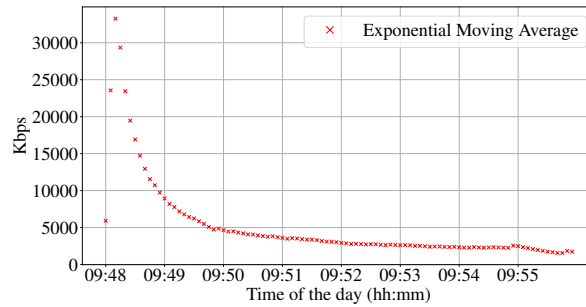


Fig. 3. Bitrate over time for a video session.




Fig. 4. Bitrate change detection and bitrate tracking.

into non-overlapping partitions [4], where the edges of partitions are defined as changepoints in the data stream and that the datapoints within each partition are independent and identically distributed from some probability distribution.

The algorithm examines a sliding window consisting of segment interarrival values to output a probability indicating the likelihood that a change has occurred. We select a window size of 15 segments as we seek to balance responsiveness to change events with sensitivity to natural fluctuations in interarrival times. Netflix commonly uses 4 second segments while YouTube uses variable-length segments with lengths varying from 1 second to




Fig. 5. Bitrate error comparison.

nearly 60 seconds [20] depending on network conditions. This means that a 15 second window allows us to view performance over the relatively recent past. The top plot in Figure 4 shows the segment interarrival times for a Netflix session along with dashed vertical lines indicating the detected changepoints using the Bayesian detector using a probability threshold of 0.4. The plot illustrates the significant differences for segment interarrival times between partitions. We see that during the initial moments of the video the interarrival times are quite small as the player attempts to buffer content faster than playback. Around 20:43 the player enters steady state where we see segment interarrivals fall into a periodic pattern before changing to a less-uniform distribution around 20:46 which coincides with a shift to a lower video bitrate before returning to steady state at 20:50.

We use the boundaries defined by the changepoint detection algorithm to slice a video session into different partitions. We then track video bitrate for each partition independently as segment downloads complete. At the end of a segment, we calculate the bitrate using the average of all of the segments since the last changepoint. The bottom plot in Figure 4 shows our output. We see that after the initial buffering, when the network throughput is high, we are able to closely track even small shifts in bitrate using this technique. Our method is successful even as ground truth bitrate varies across the video session.

*Validation.* We validate our bitrate estimation method against ground truth data gathered by the browser extension. Figure 5 depicts a distribution of the percent error obtained by subtracting the ground truth bitrate from the estimated bitrate calculated for each segment in the video sessions. Our online estimation technique estimates video bitrate with an error percentage of less than 10% for 83.5% of the datapoints, though we have a long tail representing overestimates. We attribute this overestimation to the high network utilization during the initial buffering of a video or from buffering events. For characterization purposes, we anticipate that such errors can be removed by accurately identifying buffering events.

## 6 PROPOSED CONSUMER SURVEY

While the current transparency rules makes it possible for sophisticated third parties like consumer interest group or another ISP to evaluate an ISP's network management practices and performance, it does very little to help the individual consumer. We suggest that the development of a system such as Network Microscope simplifies the process. In the next stage of this study we will use the deployment process to understand what information is valuable, how we can standardize for easier comparison and how to present financial information alongside performance information. This deployment process will help us understand what important attributes to focus on

<b>Recommended Features: Consumer Decision Making</b>	<b>Recommended Features: Provider Motivation</b>
Simplification & Personalization <ul style="list-style-type: none"> <li>• Limited Attention</li> </ul>	Third party validation <ul style="list-style-type: none"> <li>• Better implementation of transparency rules</li> <li>• More consumer trust</li> </ul>
Aggregated information including peer information <ul style="list-style-type: none"> <li>• Missing data, misinterpretations and motivated information</li> <li>• Lack of technical knowhow including impact of non ISP related causes</li> </ul>	Comparative information <ul style="list-style-type: none"> <li>• Standardized data can be compared easily</li> <li>• Improves competition</li> <li>• Improves network investment</li> </ul>
Comparative Information <ul style="list-style-type: none"> <li>• Direct impact on consumer choice</li> <li>• Simpler for consumers to understand</li> </ul>	Direct consumer access <ul style="list-style-type: none"> <li>• Better behavior on blocking and throttling if consumers can compare</li> <li>• Healthy provider competition</li> </ul>
Aggregated information with the option of personalizing the order <ul style="list-style-type: none"> <li>• Impact of presentation</li> <li>• Importance of presenting information about quality alongside price.</li> </ul>	Comparative and Standardized Information <ul style="list-style-type: none"> <li>• Promotes competition</li> <li>• Paid prioritization more justifiable on certain critical applications</li> </ul>

Table 3. Recommended disclosure features.

while constructing a comparative dashboard. We propose to survey users to understand how they process these disclosures in the context of their internet service, and whether these disclosures in their current form have an impact on their decision making processes. Based on the discussion of the consumer decision making process and producer behavior, Table 3 summarizes recommended disclosure features to incorporate in the consumer survey process.

Based on the above information, the following survey questionnaire is designed to gather more information about next steps regarding deployment of the above measurement tool. In the coming months we intend to distribute this survey.

Survey Questions:

- (1) Are you aware that ISPs are required to provide information about network management practices, performance and commercial terms?
- (2) Do you care about an ISPs network management practices?
- (3) Do you care about their performance?
- (4) Do you care about their performance only as it pertains to video?
- (5) What performance metrics are most meaningful to you?
- (6) Do you care about their commercial terms?
- (7) Would you go to an ISPs website to review this information?
- (8) Would you go to the FCC's website to review this information?
- (9) Is it useful to see comparative performance information? As in would you be interested in comparing across providers that serve in your area?
- (10) Would you be interesting in comparing across video service providers?
- (11) Would you like to see their commercial terms (pricing and contract) alongside performance metrics?

- (12) Would this tool help you make a more informed decision while purchasing or switching your internet service connection?

## 7 DISCUSSION AND CONCLUSION

The above discussed tool, along with the current transparency rules have commercial considerations, and can affect consumer decision-making, and eventually promote provider investments. These tools are effective in promoting informed consumer choice only when consumers can easily and regularly access this information when making a choice about their subscription service. The presence of switching costs may make it cumbersome for consumers to switch even when they have a preference to switch. High switching costs are typical in the information industry, and the most common form of these costs come from long term contracts and investment in complementary products. The market for mobile service has successfully addressed these costs, but it is important to ensure that providers don't use monopoly leveraging tactics to raise these costs. Furthermore, as providers enter adjacent markets like the application service or content market, it is important for providers to make available information about any leveraging practices they use in the internet service market. Vertical integration and entering into adjacent markets is common in these industries, and transparency rules and other tools should strive towards aggregating information about practices in adjacent markets to provide consumers and other competing providers with a more complete picture of their practices.

In order to achieve the transparency goals laid out by the Restoring Internet Freedom Order, we must engineer new measurement solutions for characterizing ISP and application performance. This work presents our initial proposal for a tool that would provide consumers with metrics that correlate with internet application performance by collecting such information directly in consumer networks. The outputs of such a tool can be leveraged to continually update a comparative database, allowing consumers to compare ISPs and service tiers in order to make informed decisions.

## REFERENCES

- [1] Adnan Ahmed, Zubair Shafiq, Harkeerat Bedi, and Amir Khakpour. 2017. Suffering from buffering? Detecting QoE impairments in live video streams. In *2017 IEEE 25th International Conference on Network Protocols (ICNP)*. IEEE, 1–10.
- [2] A. Balachandran, V. Sekar, A. Akella, S. Seshan, I. Stoica, and H. Zhang. 2012. A quest for an internet video quality-of-experience metric. In *Proceedings of the 11th ACM Workshop HotNets*. ACM, 97–102.
- [3] Athula Balachandran, Vyas Sekar, Aditya Akella, Srinivasan Seshan, Ion Stoica, and Hui Zhang. 2013. Developing a predictive model of quality of experience for internet video. In *ACM SIGCOMM Computer Communication Review*, Vol. 43. ACM, 339–350.
- [4] Daniel Barry and J. A. Hartigan. 1992. Product Partition Models for Change Point Problems. *Ann. Statist.* 20, 1 (03 1992), 260–279. <https://doi.org/10.1214/aos/1176348521>
- [5] Marianne Bertrand and Adair Morse. 2011. Information Disclosure, Cognitive Biases, and Payday Borrowing. *The Journal of Finance* 66, 6 (2011), 1865–1893. <http://www.jstor.org/stable/41305179>
- [6] John Beshears, James J. Choi, David Laibson, Brigitte C. Madrian, and Katherine L. Milkman. [n. d.]. The Effect of Providing Peer Information on Retirement Savings Decisions. *The Journal of Finance* 70, 3 ([n. d.]), 1161–1201. <https://doi.org/10.1111/jofi.12258> arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/jofi.12258>
- [7] Raj Chetty, Adam Looney, and Kory Kroft. 2009. Saliency and Taxation: Theory and Evidence. *American Economic Review* 99, 4 (September 2009), 1145–77. <https://doi.org/10.1257/aer.99.4.1145>
- [8] Florin Dobrian, Vyas Sekar, Asad Awan, Ion Stoica, Dilip Joseph, Aditya Ganjam, Jibin Zhan, and Hui Zhang. 2011. Understanding the impact of video quality on user engagement. In *ACM SIGCOMM Computer Communication Review*, Vol. 41. ACM, 362–373.
- [9] Driving 2017. Driving Engagement for Online Video. <http://events.digitallyspeaking.com/akamai/mddec10/post.html?hash=ZDIBSGhsMXBidnJ3RXNWSW5mSE1HZz09>.
- [10] FCC. 2017. Restoring Internet Freedom Order FCC-17-166.
- [11] Amy Finkelstein. 2009. E-ztax: Tax Saliency and Tax Rates\*. *The Quarterly Journal of Economics* 124, 3 (2009), 969–1010. <https://doi.org/10.1162/qjec.2009.124.3.969>
- [12] Justine S Hastings and Lydia Tejada-Ashton. 2008. *Financial Literacy, Information, and Demand Elasticity: Survey and Experimental Evidence from Mexico*. Working Paper 14538. National Bureau of Economic Research. <https://doi.org/10.3386/w14538>
- [13] Justine S. Hastings and Jeffrey M. Weinstein. 2008. Information, School Choice, and Academic Achievement: Evidence from Two Experiments\*. *The Quarterly Journal of Economics* 123, 4 (2008), 1373–1414. <https://doi.org/10.1162/qjec.2008.123.4.1373>
- [14] Tobias Hößfeld, Michael Seufert, Christian Sieber, and Thomas Zinner. 2014. Assessing effect sizes of influence factors towards a QoE model for HTTP adaptive streaming. In *Quality of Multimedia Experience (QoMEX), 2014 Sixth International Workshop on*. IEEE, 111–116.
- [15] ITU-method 2012. ITU Methodology for the subjective assessment of the quality of television pictures. Recommendation ITU-R BT.500-13.
- [16] Parikshit Juluri, Venkatesh Tamarapalli, and Deep Medhi. 2016. Measurement of Quality of Experience of Video-on-Demand Services: A Survey. *IEEE Communications Surveys & Tutorials* 18, 1 (2016).
- [17] Jeffrey R. Kling, Sendhil Mullainathan, Eldar Shafir, Lee C. Vermeulen, and Marian V. Wrobel. 2012. Comparison Friction: Experimental Evidence from Medicare Drug Plans \*. *The Quarterly Journal of Economics* 127, 1 (2012), 199–235. <https://doi.org/10.1093/qje/qjr055>
- [18] S Shunmuga Krishnan and Ramesh K Sitaraman. 2013. Video stream quality impacts viewer behavior: inferring causality using quasi-experimental designs. *IEEE/ACM Transactions on Networking* 21, 6 (2013), 2001–2014.
- [19] Michael Luca and Jonathan Smith. [n. d.]. Saliency in Quality Disclosure: Evidence from the U.S. News College Rankings. *Journal of Economics & Management Strategy* 22, 1 ([n. d.]), 58–77. <https://doi.org/10.1111/jems.12003>
- [20] Abhijit Mondal, Satadal Sengupta, Bachu Rikith Reddy, MJV Koundinya, Chandar Govindarajan, Pradipta De, Niloy Ganguly, and Sandip Chakraborty. 2017. Candid with YouTube: Adaptive Streaming Behavior and Implications on Data Consumption. In *Proceedings of the 27th Workshop on Network and Operating Systems Support for Digital Audio and Video*. ACM, 19–24.
- [21] MS-SSTR 2017. [MS-SSTR]: Smooth Streaming Protocol. <https://msdn.microsoft.com/en-us/library/ff469518.aspx>.
- [22] Odroid 2018. ODROID-C2 Wiki. <https://wiki.odroid.com/odroid-c2/odroid-c2>.
- [23] R. Pantos and W. May. 2017. *HTTP Live Streaming*. RFC 8216. RFC Editor.
- [24] Pi 2018. Raspberry Pi 3 Model B. <https://www.raspberrypi.org/products/raspberry-pi-3-model-b/>.
- [25] R. Prescott Adams and D. J. C. MacKay. 2007. Bayesian Online Change-point Detection. *ArXiv e-prints* (Oct. 2007). arXiv:stat.ML/0710.3742
- [26] Iraj Sodagar. 2011. The MPEG-DASH standard for multimedia streaming over the internet. *IEEE MultiMedia* 18, 4 (2011), 62–67.