Working Paper #18-12

October 2018

Steering Incentives and Bundling Practices in the Telecommunications Industry

Brian McManus
University of North Carolina at Chapel Hill

Aviv Nevo
University of Pennsylvania

Zachary Nolan
Duke University

Jonathan W. Williams
University of North Carolina at Chapel Hill

Steering Incentives and Bundling Practices in the Telecommunications Industry*

Brian McManus† Aviv Nevo‡ Zachary Nolan§
Jonathan W. Williams¶

September 30, 2018

Abstract

We model mixed-bundle pricing by internet service providers (ISPs) to study their incentive to steer consumers across different subscription options and influence usage decisions. Using unique panel data from an ISP, we test predictions from the model. We find that the ISP's introduction of internet usage allowances and overage charges steered internet-only consumers into bundled TV and internet subscriptions; this effect was greatest for heavy users of streaming services most similar to conventional TV. Internet usage growth - especially in streaming video services - was curtailed for consumers who added TV subscriptions, and it also fell for consumers who did not upgrade their internet usage allowances. We discuss the implications of these findings for antitrust and regulatory issues in the telecommunications industry.

Keywords: Steering, Bundling, Nonlinear Pricing, Telecommunications Industry, Cord Cutting, Broadband

JEL Codes: L11, L13, L96.

*We are grateful for the generosity of the North American ISP that provided the data used in this paper. We thank Shane Greenstein and Jeff Prince for comments, as well as audiences at Penn State and the NBER's IT and Digitization workshop. Zach Nolan and Jonathan Williams thank Cable Television Laboratories, Inc., and its member organizations, for their support of this work.

†Department of Economics, University of North Carolina at Chapel Hill, mcmanusb@unc.edu
‡Department of Economics, University of Pennsylvania, nevo@wharton.upenn.edu.
§Department of Economics, Duke University, zach.nolan@duke.edu.
¶Department of Economics, University of North Carolina at Chapel Hill, jonwms@unc.edu.

1 Introduction

Firms in a variety of industries provide consumers access to the products or services of other firms. In many cases, the intensity of use of these third-party services determine costs borne by the access provider. For example, internet service providers (ISPs) offer access to endless online applications that generate traffic requiring network investment. In health care markets, insurers may offer patients access to primary care physicians who can affect the insured consumers' use of other medical services covered by the same insurance contract. In these settings, firms face complex tradeoffs – access creates value but unlimited access can lead to economic losses from serving a customer. This tension leads to concerns that access providers may foreclose certain third-party firms or design complex contracts or pricing schedules that accomplish this effect by steering consumers to avoid more costly services.

These concerns are particularly acute in the telecommunications industry, and they are at the forefront of policy and regulatory discussions due to the advent and popularity of "over the top TV" (OTTV) services (e.g., Netflix). Specifically, ISPs use mixed bundling to sell both internet access and TV services, which, respectively, provide access to and compete with OTTV services. Further, OTTV services generate substantial network costs, in contrast to traditional TV services for which costs are essentially fixed.[1] The cost and revenue consequences create an incentive to steer customers or even foreclose some services. While steering may impose utility costs on consumers, it can be socially beneficial if the steering strategies (e.g. price premia or discounts) help align consumers' choices with the asymmetry in the firm's costs. These salient features of the modern telecommunications industry are not captured well by standard bundling models. In this paper, we develop a mixed-bundling framework to formalize the tradeoffs faced by ISPs, and we use unique panel data to provide the first empirical evidence on steering and relevant elasticities for the ongoing "net-neutrality" debate and other antitrust concerns.

Our analysis begins with a simple theoretical mixed bundling model that includes TV and internet subscriptions in the presence of competitively supplied third-party OTTV services. We show that many of the important phenomena described above can be captured by adding two elements to the standard bundling framework: quantity choices by consumers for the services they receive, and a parameter that describes internet subscribers' access to OTTV services. The model predicts that an increase in streaming video access, all else fixed, leads to increased "cord-cutting" behavior in which consumers drop their TV subscriptions and increase OTTV use. This reduces the ISP's revenue and increases its costs. When the firm can use only subscription

[1] The differences in the two services' costs are related to OTTV's "unicast" nature, which can have a unique transmission for each consumer, in contrast to the "broadcast" nature of traditional TV.

prices to steer consumers' subscription choices and usage, it raises the internet-only price and reduces its bundle price to encourage video use through conventional TV rather than OTTV. The firm's opportunities to steer consumers increase when we expand the model to include usage allowances and internet tiers. An internet usage allowance encourages some consumers to keep their video subscriptions while limiting internet use. When the ISP adds internet tiers which associate greater allowances with higher prices, the firm can extract more revenue from high demand consumers. While some high demand consumers experience a reduction in surplus when their internet use is capped or more expensive, some low demand consumers may gain access to the market because subscription prices can fall when set jointly with an allowance. These benefits, together with increased profit to the firm, suggest ambiguous welfare impacts of this nonlinear pricing strategy.

At the core of our empirical analysis is unique household-level panel data from a North American ISP which sells TV and internet access as stand-alone services and in a discounted bundle. The sample includes information on each household for nearly one year and is drawn from a representative sample of markets. During our sample period, the ISP implemented usage-based pricing (UBP) in a subset of markets, which offers a rare opportunity to measure consumers' responses to steering strategies in this industry. The data include detailed information on each household's internet data usage at an hourly frequency, and whether they also have a TV subscription. Central to our analysis, and in contrast to previous academic work on the telecommunications industry, we observe which OTTV services a household uses and the volume of traffic generated by these services.

In our sample, OTTV accounts for nearly two thirds of all data usage on the internet. Engagement varies substantially across households, but a large majority of households received at least one subscription OTTV service. Households with subscriptions to the ISP's TV service use much less data and engage less frequently with OTTV. These differences are not merely due to selection into different types of subscriptions; we document that when a household cuts the cord, it increases its internet usage by nearly one quarter. Overall, cord-cutting accelerated during the sample period, tripling in frequency between the start and end of the sample.

We exploit temporal and cross-sectional variation in UBP policy to investigate how consumers responded to the ISP's introduction of three-part tariffs (i.e., an access fee, a usage allowance, and an overage price) for internet service. We find that the new usage-based prices were able to steer a significant share of internet-only households into the bundle, where they can receive video entertainment through traditional TV. New bundle subscriptions were most likely for households which used OTTV services with content that is very similar to traditional TV. Following the addition of a TV subscription, a household's internet usage fell significantly. In addition to reflecting an intuitive substitution pattern, the reduction in these particular OTTV services can provide especially large cost savings for the ISP.

The UBP strategy had an even larger impact on consumers' choices to upgrade their internet service to tiers with greater usage allowances (to avoid paying overage charges). The consumers most likely to upgrade were the heaviest users prior to UBP, with only modest differences between households that engage with OTTV services and those with similar internet usage but less connection to streaming video. Households who were likely to receive overage charges and upgraded their tiers typically increased their internet usage during the sample period, while similar households that did not upgrade reduced their usage, on average. In other markets without UBP, households increased their usage regardless of whether they upgraded their internet tiers.

These steering outcomes – increased bundle subscriptions and upgraded internet tiers – represent an effective reallocation of OTTV surplus from households to the ISP. While these immediate outcomes appear unfavorable for consumers, ISPs' incentives to create new content and expand their networks may be stronger when ISPs capture a greater share of total surplus from OTTV. The steering outcomes also have an impact on third-party providers of video content, e.g. Netflix. Consumers who add TV subscriptions, for example, may be less willing to pay for subscription OTTV services.

Our results contribute to a growing literature on the rapidly changing telecommunications industry. The emergence of OTTV providers has dramatically reshaped the industry. Between 2012 and 2017, per-person average daily use of online video increased twenty-fold from 3.5 minutes to 72 minutes.[2] This usage increase coincides with the emergence and rapid growth of several prominent firms that offer OTTV services. In 2015 about 40% of US households subscribed to a streaming service such as Netflix (40M subscribers), Amazon Prime (14.5M), or Hulu (7M). While on-demand services like Netflix offer a variety of original programming along with a library of previously distributed movies and television programs, other streaming services like Sling TV offer live TV over the internet.[3] Due to data availability, much of the research on video programming has focused on traditional TV (Crawford et al. (2017), Crawford and Yurukoglu (2012), Crawford and Shum (2007), and Crawford and Shum (2015)). Our data's richness allows us to examine ISPs' incentives to interact with these emerging online sources of video-programming competition.

Relationships between ISPs and internet content providers is an active area for

[2]The industry facts in this paragraph are reported by Neilsen and Statista.

[3]Other similar services include: Sony Play Station Vue, DirecTV Now, YouTubeTV, and Hulu Live.

public policy. The FCC's 2015 Open Internet Order prevented ISPs from discriminating among various online applications. This order limited ISPs' ability to reduce usage of video services from some third-party providers. The FCC voted in 2017 to roll back the order, and future policy in this area continues to be debated, including antitrust scrutiny of recently-proposed vertical mergers between media companies and ISPs. The literature on these issues largely began with Wu (2003), who introduced the term "net neutrality" and provides an excellent summary of the issues. Lee and Wu (2009) and Greenstein et al. (2016) discuss and review the literature on the topic. However, most of the existing economic analysis of the topic is theoretical: Economides and Hermalin (2012), Armstrong (2006), Bourreau et al. (2015), Choi et al. (2015), Choi and Kim (2010), Economides and Tag (2012), Gans (2015), Economides and Tag (2016), Reggiani and Valletti (2016), and Sidak (2006). Our empirical analysis on steering incentives complements these theoretical studies by providing insight into relevant elasticities for the debate.

Our work is related to several studies on firm strategies and demand in the telecommunications industry. We build on Nevo et al. (2016), Malone et al. (2016), and Malone et al. (2014), who use high-frequency data to study subscriber behavior on residential broadband networks. However, our access to detailed information on consumers' subscriptions and OTTV usage offers a unique opportunity to analyze the complex incentives faced by ISPs due to OTTV's role in the interaction between internet and television services. Other studies of the demand for and consumers' value from broadband services include Goetz (2016), Tudon (2018), Goolsbee and Klenow (2006), Dutz et al. (2009), Rosston et al. (2013), Greenstein and McDevitt (2011), Goolsbee and Klenow (2006), Edell and Varaiya (2002), Varian (2002), and Hitte and Tambe (2007).

More generally, our model and empirical analysis also contribute to literatures on firms' strategic efforts to steer and sort heterogeneous consumers across product menus. Ho and Lee (2017), Liebman (2017), and Raval and Rosenbaum (2017) study how insurers influence patients' choices across medical providers. Barwick et al. (2017) examine conflicts of interest and steering by residential real-estate brokers. Crawford et al. (2017) consider similar incentives in cable TV markets, and estimate the value to cable distributors of including vertically integrated versus non-integrated sports networks in their channel bundles. The incentive to degrade product quality for discriminatory or steering purposes, as is present in our model, is related to the classic work of Mussa and Rosen (1978), which Crawford and Shum (2007) apply in the context of the telecommunications industry. In the bundling literature, Armstrong (2013) and Gentzkow (2007) study how the consumption of one product in a bundle affects utility from other products, which is similar to the relationship between OTTV and TV that we study. Chu et al. (2011) and Crawford and Yurukoglu (2012) empirically explore how variations on bundling and other pricing strategies can affect firms' profit and consumer welfare. Nonlinear-pricing strategies similar to those we examine have been studied in broadband markets (Economides and Hermalin (2015), Lambrecht et al. (2007)), phone service contracts (Miravete (2003), Grubb (2015), and Grubb and Osborne (2015)), and other markets (Hagemann (2017), McManus (2007)).

2 Model

We develop a mixed bundling model to capture some of the central features of how ISPs have responded to the advent of OTTV. The emergence and popularity of OTTV services has coincided with a trend in cord cutting. Between 2015 and 2017, the number of US households that received video exclusively through a broadband connection almost doubled to about 5.4M. ISPs have generally used two strategies to manage these changes. First, they have offered TV at a substantial discount when bundled with broadband, in an attempt to persuade consumers to receive video though traditional TV services that do not carry the same network costs as broadband OTTV delivery. For example, Comcast (the largest ISP in the US) offers some bundled internet and cable TV packages at a discount of 30%.[4] Second, ISPs have created internet services tiers which vary in speed and usage allowances, which can curtail streaming video use, provide an additional incentive for consumers to continue their TV subscriptions, and allow the ISP to capture surplus from heavy internet users. The US GAO reports that seven of the thirteen largest ISPs had monthly data allowances in 2014. These strategies are of interest to regulators and policymakers because they have complex effects on consumers' choices and welfare. Increased bundling discounts are appealing in that they reduce the marginal price of a valued service, but other subscription prices may increase as part of the same strategy. Similarly, ISP efforts to reduce use of OTTV services through prices and usage allowances may appear harmful to consumers who value OTTV strongly, but the ISP may reduce broadband's subscription price for other consumers when its costs can be controlled.

In specifying our theoretical model, we make minimal changes to the standard monopoly model of bundling to highlight the modifications necessary to capture the salient features of the modern telecommunications industry in which internet and TV services are no longer distinct services.[5] First, we incorporate an intensive margin to

[4] Comcast offers packages with 200+ TV channels and 100Mbps internet at \$60 and \$85 per month, respectively, when either is selected individually, or a bundle of both services for an average of \$100 per month over the term of the contract. Spectrum subscribers save about \$20 (18%) when they subscribe to bundle with 125+ TV channels and 100Mbps internet.

[5] Most previous bundling models with complementarity include assumptions that make the utility of the bundle greater or less than the sum of the individual utilities. This type of between-product relationship is distinct from what we consider here. Indeed a model could be constructed that includes both features, but we choose not to do so for the sake of (relative) simplicity. In our

consumers' choices in addition to the discrete choice made over subscription services. This is important because costs associated with building and maintenance of networks are increasing in total traffic. Second, we model internet service as a gateway to online applications, including OTTV. This captures the current tradeoff for ISPs: access to OTTV increases the value of internet subscriptions but access increases costs and makes TV and bundles relatively less attractive to consumers.

2.1 Firm, Consumers, and Offerings

Consider a market in which a monopoly ISP offers consumers access to a pair of distinct services: non-video internet (1) and video entertainment (2). An individual consumer's taste for services 1 and 2 is $v = (v_1, v_2)$. We normalize the consumer population to one and assume that consumers' tastes are distributed independently and uniformly on $[0,1] \times [0,1]$.

The ISP offers subscriptions (also called "plans" or "packages") that allow consumers to gain access to these services. To begin, we assume the firm offers three subscription options: broadband internet access (i), TV (t), or a bundle (b) that includes both i and t. The firm's mixed bundling pricing strategy includes prices for the stand-alone subscriptions (p_i and p_t) and a price for the bundle (p_b). A consumer can subscribe to one of firm's three plans, $\{i, t, b\}$ or an outside option denoted by 0 that provides zero utility. To capture the presence of OTTV, we assume that consumers can access some video content (service 2) through an internet-only subscription (i).

In contrast to standard bundling models which assume that usage is binary, we assume that usage quantity matters to consumers and firms. Let $q_1(v)$ and $q_2(v)$ represent utility-maximizing quantities for a consumer with preferences v. The consumer's choice of q in general depends on his choice of $k \in \{i, t, b, 0\}$ and perhaps other aspects of the firm's packages, which are made clear in context below. We measure q in units that are meaningful to a consumer, e.g. hours of entertainment. The quantity choice for video services on a given package, $q_2(v)$, can include both traditional TV and OTTV, so we separate q_2 into q_2^i and q_2^t, with $q_2 = q_2^i + q_2^t$. The consumer receives no utility from a service when the quantity is zero. For simplicity, we assume that marginal utility for a service is constant at one until the consumer reaches a satiation level (determined by v), and then marginal utility is zero for any greater quantity. This results in rectangular demand.

Given these assumptions on preferences, in Columns 1 and 2 of Table 1 we summarize the relationship among v, quantity, and utility across the ISP's plans. If the consumer selects the outside option, 0, quantities are zero for both services and utility

context, subadditive utility from the bundle could be justified, for example, by time constraints that effectively raise the utility cost of consuming both video and internet services.

is zero. A subscription to t results in usage for video services up to the satiation level, $q_2^t(v) = v_2$, and there is zero non-video internet usage given the lack of access. This consumer receives utility equal to $U_t = v_2 - p_t$. Internet-only subscribers receive utility of $U_i = v_1 + \delta v_2 - p_i$, where the first and second terms captures utility from consuming non-video internet and OTTV applications, respectively. The parameter $\delta \in [0, 1)$ allows the consumer to use OTTV to receive a fraction of his preferred video quantity. We assume that OTTV is available at no additional expense to the consumer. The restriction $\delta < 1$ has several possible interpretations, including limited available OTTV content and diminished video quality, which could be due to transmission (e.g. congestion and buffering) or hardware limitations.[6] We assume that consumers who subscribe to the bundle receive all video entertainment through traditional TV, with $q_2^t = v_2$ and $q_2^i = 0$.[7] These consumers receive utility equal to $U_b = v_1 + v_2 - p_b$.

Table 1: Utility and Quantity Choice by Plan

Offering	Net Utility	Quantities: $q(v)$, Hours	Quantities: $g(v)$, GBs
i	$U_i = v_1 + \delta v_2 - p_i$	$q_1 = v_1$ $q_2 = q_2^i = \delta v_2$	$v_1 + \beta\delta v_2$
t	$U_t = v_2 - p_t$	$q_1 = 0$ $q_2 = q_2^t = v_2$	0
b	$U_b = v_1 + v_2 - p_b$	$q_1 = v_1$ $q_2 = q_2^t = v_2$	v_1
0	$U_0 = 0$	$q_1 = 0$ $q_2 = 0$	0

Notes: This table summarizes the net utility and usage quantities associated with each subscription choice.

The ISP's profit from serving the market depends on the number of subscribers to each plan, the quantities they select, and the costs of providing these services. On the TV side of the market, the ISP pays content providers a fixed fee, γ_t, for each consumer who subscribes to the service, plus the ISP pays a fixed cost, Γ_t, to maintain the service network. Once the firm pays γ_t there is no additional cost of supplying q_2^t. Internet service, by contrast, requires the ISP to transmit bytes through its network, and we assume that transmission costs are increasing in quantity. This relationship

[6]While we treat δ as an exogenous parameter, alternative models with endogenous δ could consider a range of additional interesting strategies. For example, an ISP might choose to increase or decrease δ through network investment or internet terms of service. Alternatively, the ISP might set different δ values for different internet tiers, or different δs for different OTTV content depending on whether the ISP owns the content.

[7]This assumption represents the best-case/least-cost outcome for the ISP and strengthens incentives to steer consumers to the bundle, but the assumption does not qualitatively change any of the model's predictions.

between quantity and costs captures the firm's need to build-out its network to provide the demanded content.[8] We assume that the firm's transmission costs are expressed in gigabytes, denoted g, and costs increase in the quantities q_1 and q_2^t. The firm has marginal cost of γ_i to transmit one unit of internet content. We normalize units so that one gigabyte (GB) provides one unit (e.g. hour) of non-video internet services, while $\beta > 1$ GBs are required for a unit of OTTV consumption. This captures the greater bandwidth requirements and immediacy needs of OTTV relative to other internet applications.[9] In addition to the variable cost of transmitting content, the ISP has fixed cost of Γ_i to provide its internet infrastructure.

In Columns 2 and 3 of Table 1, we provide the relationship between the units that consumers value (q hours) and the unit relevant for the firm's cost (g GBs). Subscribers to the outside option (0) or stand-alone TV (t) demand zero GBs, i.e., $g_0(v) = g_t(v) = 0$. For a consumer with preferences v who subscribes to stand-alone internet (i) or the bundle (b), we denote these plan-specific mappings (hours to GBs) as $g_i(v)$ and $g_b(v)$, respectively. An internet-only customer who consumes q_1 and q_2^i requires $q_1 + \beta q_2^i$ GBs to be transmitted through the firm's internet infrastructure. Given the relationship between q and v, this implies $g_i(v) = v_1 + \beta\delta v_2$. When a consumer chooses the bundle, $g_b(v) = v_1$ because all video demand is satisfied through the TV service ($q_2^i(v) = 0$).

Together, these assumptions on preferences and costs imply that the ISP's profit function is

$$\pi = \int\limits_{v \in S_i} [p_i - \gamma_i g_i(v)]\, dv + \int\limits_{v \in S_t} [p_t - \gamma_t]\, dv + \int\limits_{v \in S_b} [p_b - \gamma_i g_b(v) - \gamma_t]\, dv - \Gamma_i - \Gamma_t.$$

where S_i, S_t, and S_b are the sets of consumers choosing internet, TV, and the bundle, respectively. In our numerical examples below, we find it useful to set to zero the parameters γ_t, Γ_i, and Γ_t.

Our model allows an intuitive analysis of how OTTV affects consumers' choices and ISP profit, as well as illustrating the firm's benefits from using bundle discounts, nonlinear pricing, or allowances to steer consumers among the subscription options. In Panel (a) of Figure 1 we demonstrate the effect of an increase in δ, i.e., OTTV becomes more attractive, holding prices fixed. Consumers substitute along two margins. The stand-alone internet plan attracts consumers who had selected the outside option because of the increased value associated with access to the internet ($0 \rightarrow i$).

[8]The ISP maintains a service network that must grow with total usage and can experience content interruptions or slowdowns during times of high total usage. In the present model we assume-away direct affects of congestion on consumers' utility or usage.

[9]The assumption $\beta > 1$ can also be interpreted as implying a lower per-gigabyte marginal utility of OTTV than other internet activity.

Figure 1: *Consumer Response to Variation in δ*

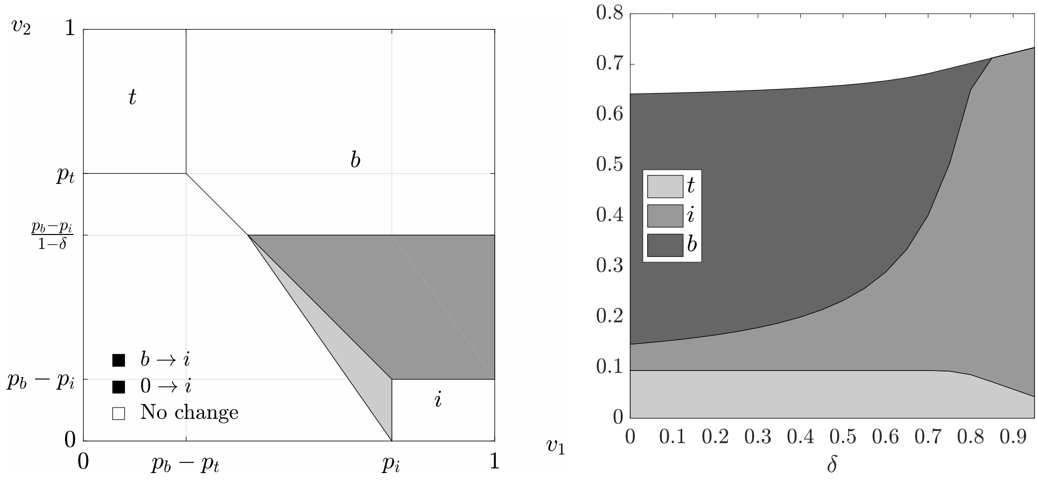

(a) Substitution (b) Market Shares

Notes: These figures illustrate consumer substitution when δ increases and prices remain fixed. In panel (a), prices are fixed at $(p_i, p_t, p_b) = (0.75, 0.65, 0.9)$ and delta is allowed to increase from 0 to 0.7. In panel (b), prices are set at the $\delta = 0$ optimum. Market shares for each service are computed as δ increases with prices held fixed.

The new consumers increase ISP profit. However, some fraction of bundle subscribers now prefer stand-alone internet service ($b \to i$), which diminishes revenue ($p_i \leq p_b$) and increases costs ($g_i(v) \geq g_b(v)\ \forall v$). In Panel (b) of Figure 1, we plot market shares for each plan at different values of δ, holding prices at the optimal level for $\delta = 0$ and $\gamma_i = 0.2$.[10] The i share grows monotonically in δ, and this largely comes from a reduction in the bundle market share, which is completely eliminated around $\delta = 0.85$. When OTTV is sufficiently beneficial, some consumers even convert their t subscriptions to i, as i offers services very similar to the original bundle. Thus, if the ISP does not respond to the emergence and improvement of OTTV, the model predicts a reallocation of market shares among plans that is consistent with the recent and rapid trend in cord cutting.

2.2 ISP Strategies

ISPs currently have a variety of strategies available to mitigate the effects of cord cutting, or to simply acquire a greater share of the surplus from OTTV's rise. A repeal of net neutrality would expand these opportunities. We use our model to explore the implications of some of these strategies, and we discuss their importance for allocating surplus. We focus on the strategies observed in our data to provide testable implications, which we study empirically in Section 4.

[10] These prices are $p_i \simeq 0.721$, $p_t \simeq 0.642$, and $p_b \simeq 0.906$

Optimal Prices

If prices are not altered in response to an increase in δ, consumer movement from b to i reduces the ISP's profit. The consumers trade a more expensive subscription for one with a lower price, and they generate streaming costs ($\gamma_i > 0$) from video usage that was previously delivered at no cost ($\gamma_t = 0$). If the firm increases p_i, reduces p_b, or both, it can dampen these effects. Adjustments to these prices can reverse the cord-cutting choices illustrated in Figure 1. Whether changes to p_i, p_b, or both make a greater positive impact on profit will depend on the changes along each plan's extensive margin and profit from inframarginal consumers.

Figure 2: *Optimal Prices for Varying Values of* δ

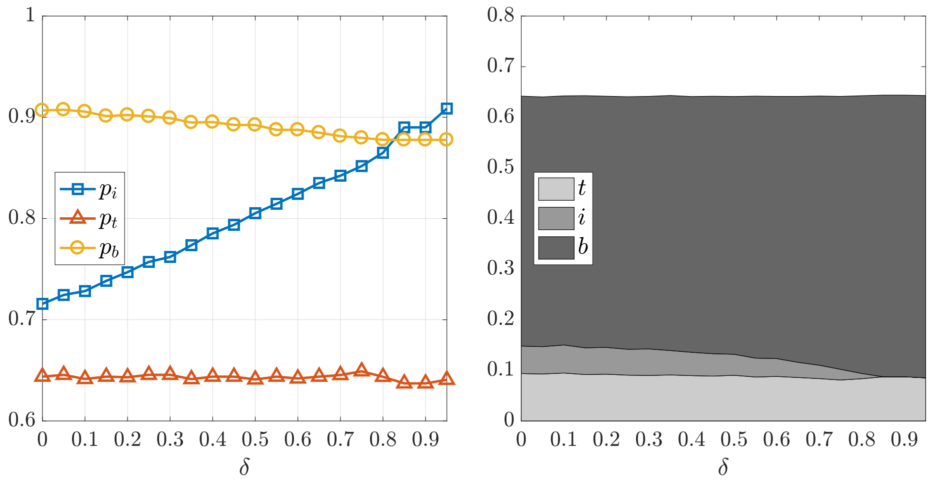

(a) Prices (b) Market Shares

Notes: Panel (a) provides optimal prices for each level of δ. Panel (b) provides the market shares resulting from these optimal prices.

The appeal of using prices to steer consumers is illustrated in Panel (a) of Figure 2, which presents the relationship between optimal mixed bundling prices and δ. Greater values of δ are generally associated with greater values of p_i and greater bundle discounts. Together, these prices act to steer consumers from i to b, while also reducing some i subscriptions at the extensive margin. The steering effects are especially strong for $\delta > 0.8$, when the firm uses prices to completely eliminate broadband-only subscriptions, and consequently forecloses OTTV usage completely.[11] Figure 2 Panel (b) illustrates the impact of optimal prices on market shares. The bundle market share grows in δ while the stand-alone internet share decreases.

When price are fixed (as in Figure 1), an increase in δ reallocates surplus from the firm to consumers. When prices are re-optimized as δ increases, by contrast, the

[11]It is notable that at approximately the same value of δ with fixed prices, the bundle is shut out of the market.

firm's profit does not necessarily decline. An expansion in OTTV's utility increases consumers' benefits from i and therefore their willingness to pay. In our model it is easy to find cases in which re-optimized prices result in roughly constant profit while consumer surplus increases modestly.

Optimal Prices with Allowances and Tiering

Our model provides stark predictions for market outcomes when δ is sufficiently large and firms are restricted to alter only prices. In particular, OTTV is foreclosed through an increase in p_i and decrease in p_b. If firms are permitted to use more sophisticated strategies, their incentives to foreclose OTTV may be dampened. We explore the implications of usage allowances and tiering, similar to those we observe in our data.

Table 2: Internet Subscriber Consumption Quantities and Utility with Usage Allowances

Offering	Net Utility	Quantities	
		$q(v)$, Hours	$g(v)$, GBs
i ($v_1 \geq \kappa$)	$U_i = \max\{v_1, \kappa\} - p_i$	$q_1 = \max\{v_1, \kappa\}$ $q_2 = q_2^i = 0$	$\max\{v_1, \kappa\}$
i ($v_1 < \kappa$)	$U_i = v_1 +$ $\max\{\delta v_2, (\kappa - v_1)/\beta\} - p_i$	$q_1 = v_1$ $q_2^i = \max\{\delta v_2, (\kappa - v_1)/\beta\}$	$v_1 + \beta \times$ $\max\{\delta v_2, (\kappa - v_1)/\beta\}$
b	$U_b = \max\{v_1, \kappa\} + v_2 - p_b$	$q_1 = \max\{v_1, \kappa\}$ $q_2 = q_2^t = v_2$	0

Notes: This table summarizes the net utility and usage quantities associated with each subscription choice when the ISP sets a usage allowance κ.

While prices can be a useful way for the ISP to deal with rising values of δ, a usage allowance can always generate further improvements in profit when $\gamma_i > 0$. Let κ denote the usage allowance in GBs. We assume this is a strict allowance that applies to all types of usage, so $g_i(v) \leq \kappa$.[12] Consider the situation illustrated in Panel (a) of Figure 1 with $\delta > 0$. If the firm sets κ slightly above p_i, it changes no consumers' choices on the margin between i and the outside option, but it reduces internet usage for all inframarginal consumers who strictly prefer i to all other options. The usage reduction also applies to b subscribers with large v_1 that do not change their subscription decisions despite a reduction in surplus. We summarize the impact of an allowance on a consumers' choices and outcomes in Table 2 (excluding unaffected subscribers to t and 0).

Panel (a) of Figure 3 illustrates optimal prices and allowances for a range of δ values,

[12]In practice, these types of limits are set through "acceptable-use" policies or complemented with overage prices that permit the consumer to exceed the allowance for a fee.

Figure 3: *Optimal Prices and Allowance for Varying Values of δ*

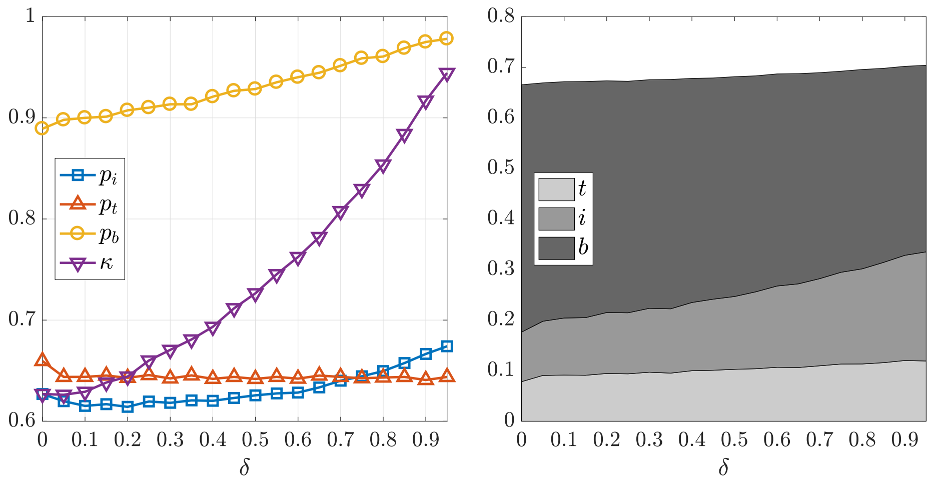

(a) Prices and Allowance (b) Market Shares

Notes: Panel (a) depicts optimal prices when the ISP sets a usage allowance, as δ varies from 0 to 1. Panel (b) summarizes the change in market shares at each optimal price combination.

and Panel (b) of Figure 3 illustrates the resulting market shares. Several important differences exist between the outcomes with and without an allowance. The allowance becomes the firm's primary steering tool between i and b, as p_i is largely unchanged by the increase in δ while p_b rises. The firm allows the i market share to grow with δ, and the ISP serves a greater share of the full consumer population. Some broadband subscribers (with large v_1 and small v_2) use no OTTV, while others consume both internet services and streaming video. Thus, the introduction of the allowance can permit the ISP to balance providing access to more consumers while simultaneously steering customers with greater values of $g_i(v)$ to low-cost offerings (i.e., b).

In addition to using base prices (p_t, p_i, p_b) and allowances (κ) to steer consumers among offerings, ISPs often also separate their internet services into quality tiers. The tiers may vary in terms of usage allowances, overage prices, or connectivity speeds. In Figure 4 we demonstrate the effect of one such strategy where a "low usage" tier is subject to the usage allowance κ while a "high usage" tier has no allowance. The same tiers are available to both internet-only and bundle subscribers (i_L and b_L for low-usage tiers, and i_H and b_H for high usage). We impose a simple set of prices in the market to facilitate discussion of how the introduction of internet tiers shifts subscription and usage choices. This resembles the implementation of the tiering and UBP strategy that has been implemented by numerous North American ISPs.

Former subscribers to the initial unlimited internet service (i) may update their subscription and usage choices in several ways. First, consumers may opt to subscribe to i_L, but in some cases their internet usage will be limited by κ. In our example

Figure 4: *Effect of Tiers and Allowances*

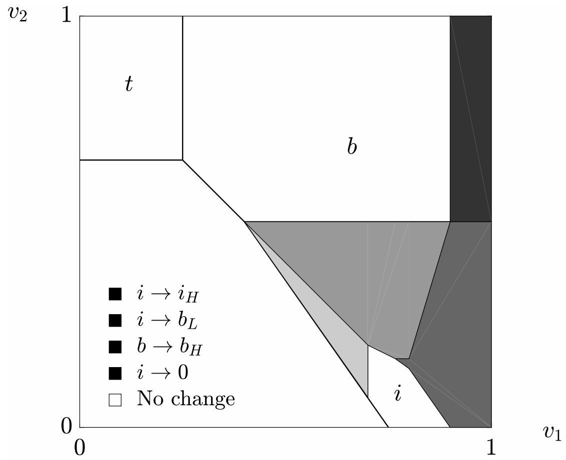

Notes: This figure shows the effect of the introduction of a tier with a usage allowance on subscription choices. Market shares are first plotted for prices $(p_i, p_t, p_b) = (0.75, 0.65, 0.9)$. Next, a usage allowance is placed on the original internet tier, and a new premium internet tier is introduced with no allowance. The new prices are $(p_{i,L}, p_{i,H}, p_t, p_{b,L}, p_{b,H}) = (0.75, 0.85, 0.65, 0.9, 0.7)$. Each shaded region depicts a set of consumer types that makes a particular subscription change. The usage allowance is set at 0.8. Throughout, δ is fixed at 0.7.

this occurs for relatively low-usage consumers. Second, consumers may convert their subscription to b_L. While the prices are no different from the initial case with no internet allowance, the limits on internet use are sufficient to steer consumers toward b_L. These consumers receive their video entertainment through the ISP's television services, and their internet and OTTV use fall. Relative to the consumers who remain in i_L, the consumers who switch to b_L have a stronger taste for video entertainment. Third, consumers may "upgrade" their internet subscription to i_H, although this is really an opportunity to pay a greater price for the same unlimited service that the consumers used previously. Consumers who switch from i to i_H continue to use the internet and OTTV at the same quantities as under the initial menu. Finally, some i consumers switch to the outside option. In addition to these margins for former i subscribers, some bundle subscribers with strong internet tastes will opt for b_H so that they can continue unlimited internet use, while other b subscribers convert their subscriptions to b_L and reduce internet use.[13]

[13]Some additional transitions are possible although we do not consider them in our simple example. For example, some consumers with strong video and internet tastes may choose to add video while also upgrading their internet tier; this could occur with price and tier assumptions different than the ones we display. In addition, continuing increases in δ will create countervailing incentives for consumers to drop video.

3 Data overview and descriptive analysis

3.1 Empirical Setting and Data Sources

Our data come from an anonymous North American ISP. The ISP offers three residential services: internet, conventional TV, and telephone. Like many telecommunication firms in North America, the ISP sells TV and internet services via mixed bundling, giving discounts off of stand-alone prices when consumers subscribe to both services. We exclude the firm's telephone service from the present study because it lacks the policy relevance and interrelated nature of TV and internet services.[14] Like many ISPs, the firm offers tiers of internet service that are differentiated by speed and in some cases usage allowance.

The sample is drawn from a handful of markets during a one-year period in the mid-2010s, and it is nationally representative in terms of demographics, service offerings, and usage patterns.[15] The scope and length of the data make it unique. During our one-year sample period we observe about 350,000 consumers' billing information, subscriptions, and internet usage. Most consumers in the sample (72%) are active subscribers for the full sample period; a small share of consumers terminate their subscriptions (16%) and some consumers (14%) start subscriptions after the sample begins.

We organize the subscription data at the household-month level, to reflect the ISP's regular billing cycle. 30% of ISP customers have internet-only subscriptions, and the remaining 70% subscribe to an internet-TV bundle; no ISP customers subscribe to TV alone. Similar bundling patterns occur at other ISPs. While the ISP offered several internet service tiers, over half of all households (57%) subscribed to one "middle" tier in the ISP's menu. The bundling frequency varied somewhat across tiers, from a low of 57% in the highest-speed tier to a maximum of 73% in the most popular tier.

For each household in the sample, we observe internet usage data information from two linked collection platforms, Internet Protocol Detail Records (IPDR) and Deep Packet Inspection (DPI).[16] IPDR data offer precise details on the ISP's byte counts flowing into and out of the household, i.e. download and upload volume. DPI data provides information on the application (e.g., Netflix) or protocol (e.g., File Transfer Protocol or FTP) generating each byte used by a consumer, but not the specific content (e.g., particular movie title or website).[17] In some instances the DPI data are missing

[14] About 40% of ISP customers subscribe to its telephone service.

[15] Our non-disclosure agreement prevents us from identifying the firm and any details that could be used to infer their identity. This includes the specific markets served and some details about the implementation of UBP.

[16] The data sources are connected by a hashed household-level code which preserves the anonymity of consumers.

[17] Given the vast number of applications, we develop a taxonomy to group the applications by utility or function following industry standards like Sandvine's reporting of aggregate internet traffic. We

byte counts for a consumer's applications, and in those cases we use the IPDR data to adjust usage volume. Adjustments of this sort account for a very small share of all download volume that we describe and analyze. Thus, we are able to infer which OTTV services are utilized by a customer and the intensity of usage for each. We organize the usage data at the household-day-hour level, which generates about 8,000 observations per household on internet usage.

A central feature of our data is the ISP's choice to introduce UBP in some markets partway through the sample period. We observe one "treated" market that eventually had UBP and several additional "control" markets without the prices. The ISP's introduction of UBP came in two phases. First, the ISP announced that UBP would be arriving within a few months, and it provided households with information about how their monthly usage compared to the data allowance within their current internet tier. During this phase, which we call the "Announcement Period," households were not billed if their usage exceeding their tier's allowance. In the next phase, which we call the "Treatment Period," the ISP assessed overage charges on households that exceeded their allowances. In addition to the Announcement and Treatment periods, which each last a few months in our sample, we observe two months of activity (a "Pre-Policy Period") prior to any announcement of the pricing program. Markets without UBP had no announcements or policy changes linked to households' internet usage. Other than the introduction of UBP, the ISP's internet tiers were identical across markets in our sample.

The ISP faced similar competitors in the treated and control markets during the sample period. Satellite TV was available in each market, as was internet via DSL lines. During our sample period in the markets we study, the firm from which we received data offered internet service on a substantially higher-speed and higher-capacity network than its competitors. The competitors' subscription offerings did not change meaningfully during the sample period, including in response to the UBP policy's introduction.

It is informative to separate households by internet usage intensity, so we partition customers by how their usage compares to the allowance in the treated market for their internet tier. For each household, we calculate average internet usage during sample period's first two months (Pre-Policy). For the household's initial internet tier and its corresponding usage allowance under UBP, let φ represent the fraction of the allowance used by the household. This fraction is irrelevant for many households in our data sample – they were never exposed to UBP and did not need to consider whether they exceeded a usage allowance. Nevertheless, the fraction φ provides a convenient

present this classification in Appendix Tables 10 and 11. This simple grouping captures nearly 99% of all traffic, with the "other" category capturing the remainder.

way to describe all households as heavy or light internet users within their tier, and to examine the impact of UBP across households.

3.2 Consumers' Internet Usage Choices

Households in the sample use an average of 4.7GB of internet data per day; median internet usage is 1.37GB per day. Internet usage differs substantially across households, and those who use more data (including OTTV) tend to select higher-speed tiers. Mean usage in the highest-priced tier is nearly seven times as great as in the lowest-priced tier. The usage dispersion within a tier is also substantial, as the interquartile range differs by an order of magnitude for each, and the within-tier coefficient of variation ranges from 1.67 to 2.05. There are also considerable differences between households with and without bundled service.[18] The median unbundled subscriber uses 137% more internet data than the median bundle subscriber. The percentage difference in data use is larger below the median; the difference in mean usage is 61%.

In Table 3 we present statistics on the distribution of daily usage, by application type, for bundled and unbundled households. Mean usage for each application type is greater for unbundled customers. This suggests that OTTV services are not the only online applications that are used more intensively by households without TV subscriptions. Video, browsing, and streaming services generate the large majority (nearly 90%) of all internet traffic, and within these services unbundled households' average usage is 81% greater than bundled customers. In Table 4, we provide a breakdown of OTTV streaming across individual applications, separately for bundled and unbundled households. Both adoption and use intensity of OTTV applications are greater for internet-only households.

Table 3: Daily Usage Statistics by Traffic Type

Type	Average Usage (GBs)		
	Bundled	*Unbundled*	*Overall*
Browsing	1.86	1.25	1.44
Gaming	0.17	0.13	0.15
Music	0.19	0.15	0.16
Video	3.86	2.25	2.75
Other	0.31	0.18	0.22
Total	6.39	3.96	4.72

Notes: This table reports distributional statistics of subscriber-day observations in our panel broken out by online activity category.

[18] Both types of households have the same general patterns: substantial heterogeneity within tiers and strong selection of higher-volume users into more expensive tiers.

Table 4: OTTV Service Adoption and Daily Usage

Application	Adoption (%)	Unconditional Use (GBs)	Conditional Use (GBs)
Unbundled			
HBO Go	11.3	0.03	0.25
Hulu	22.6	0.14	0.39
Netflix	72.4	2.00	2.66
SlingTV	2.93	0.05	1.59
YouTube	92.5	0.95	1.02
Bundled			
HBO Go	10.1	0.02	0.20
Hulu	12.0	0.05	0.19
Netflix	62.6	1.17	1.81
SlingTV	0.53	0.01	1.15
YouTube	84.5	0.65	0.76

Notes: This table reports the share of subscribers that engage with the OTTV services observed in the DPI data and the average daily consumption of these services.

There was substantial internet usage growth during our sample period. Overall, average internet usage increased at a 44% annualized rate, with the largest growth rates associated with consumers who began the sample at lower usage levels. This may follow from lower-usage consumers finding more opportunities to begin or expand OTTV streaming during the sample period. Nevertheless, internet usage growth was also substantial for high-use households, and because of large usage differences across households, the choices of high-use consumers still drive overall growth.[19] Despite the rapid usage growth within our sample, there is very little evidence that congestion affected internet use.[20] The ISP invested heavily in its network prior to our sample period. Across the industry in general, innovation and network investments have played a critical role in stimulating and accommodating internet growth of about 40% per year during the past decade.

3.3 Plan and Usage Transitions

During the sample period many ISP customers switched between subscription packages, which could include switching to a different internet tier or adding or dropping TV. Some transitions were largely unprompted by the ISP's actions. For example, changes to household composition, available time for entertainment, or the acquisition

[19]We observe growth rates that are more similar across tiers than those reported by Nevo et al. (2016).

[20]Packet loss, which is a quality disruption often caused by congestion, averaged less than 0.01% during peak hours in our sample. See Malone et al. (2016) for a study of the impact of congestion on broadband networks.

Table 5: Conditional Probabilities of Subscription Changes

	Final State			
Market	Internet-only		Bundled	
Initial State	*Downgrade/Same*	*Upgrade*	*Downgrade/Same*	*Upgrade*
Internet-only				
Treated Market				
$\varphi < 0.5$	93.66	2.58	3.38	0.38
$\varphi \geq 0.5$	79.64	17.03	2.34	0.99
Control Markets				
$\varphi < 0.5$	94.55	2.60	2.48	0.37
$\varphi \geq 0.5$	92.45	4.09	2.95	0.51
Bundled				
Treated Market				
$\varphi < 0.5$	1.82	0.32	92.86	5.00
$\varphi \geq 0.5$	2.29	2.10	77.01	18.59
Control Markets				
$\varphi < 0.5$	2.04	0.21	94.33	3.42
$\varphi \geq 0.5$	3.25	0.55	91.12	5.08

of new internet-ready hardware could prompt a household to change subscriptions. Other transitions coincided with the ISP's introduction of UBP.

In Tables 5 and 6 we provide summary statistics on households' transitions across plans and their subsequent usage choices. It is informative to separate consumers by whether they are relatively intense internet users, so we partition the users by how their usage compares to the treated market's allowance for their internet tier. We separate the households by whether $\varphi \geq 0.5$ at the start of the sample period. To provide a preliminary exploration of how UBP's introduction affects households' decisions, we also separate consumers by whether they were in the treated market.

In Table 5, we summarize the transitions between services made by customers during our entire sample. Specifically, the rows and columns provide the consumer's initial and final states, respectively, and the entry in each cell is the probability of transition conditional on the initial state (i.e., each row adds to one). Most households keep the same overall collection of services (internet-only or the bundle) and remain in their initial internet tier, but a significant share change some aspects of their subscriptions. For example, consumers with $\varphi \geq 0.5$ are more likely to upgrade their internet tier during the sample period. This difference is relatively strong among high-usage households in the treated market.

In Table 6 we summarize consumers' internet usage changes. The usage patterns provide preliminary evidence that is consistent with our model's predictions. High-usage consumers in the treatment market generally reduced their internet use if they did not upgrade their internet tier; the exception is usage among cord-cutters who did

Table 6: Consumption Change (%) Conditional on Subscription Changes

Initial State	Final State: Internet-only		Final State: Bundled	
Market	*Downgrade/Same*	*Upgrade*	*Downgrade/Same*	*Upgrade*
Internet-only				
Treated Market				
$\varphi < 0.5$	30.82	114.91	21.58	42.25
$\varphi \geq 0.5$	-8.78	7.16	-25.76	-13.72
Control Markets				
$\varphi < 0.5$	43.69	92.19	29.36	70.88
$\varphi \geq 0.5$	8.16	28.56	-12.17	0.78
Bundled				
Treated Market				
$\varphi < 0.5$	80.33	188.58	25.42	74.56
$\varphi \geq 0.5$	7.88	22.34	-11.77	9.29
Control Markets				
$\varphi < 0.5$	78.50	137.58	36.24	65.90
$\varphi \geq 0.5$	11.27	31.51	1.45	15.76

Notes: This table describes consumption changes conditional on starting characteristics and plan transitions during the sample. The leftmost column describes initial states (internet-only vs. bundled, treated vs. control, and usage level φ) while the upper columns describe plan state transitions (internet-only vs. bundled, downgrade/same internet allowance vs. upgrade internet allowance). Changes in consumption are expressed as a percent difference between initial and final level.

not upgrade internet. In addition, heavy-usage treated consumers who added video reduced their internet usage, on average, even if they upgraded their internet tier. Overall, the treated market's internet usage grew more slowly than in the control markets. For consumers with $\varphi < 0.5$ during the pre-policy period, this could be due to general increased caution about internet use following UBP's introduction, or because fewer consumers in this low-usage group allowed their usage to grow beyond their tier thresholds.

3.4 Cord-Cutting Motivations and Implications

ISPs' challenges with OTTV are at their sharpest with households that cut the cord and rely on the internet for all video entertainment. We provide a detailed description of cord-cutters' changes to internet use in order to offer a more complete description of the ISP's motivation to introduce UBP.

Our illustrative model predicts that a consumer can be prompted to cut the cord if δ shifts while all prices are held fixed. In practice, this can occur because an individual acquires new internet-ready hardware, or if the quality or quantity of OTTV improves. The model suggests that cord-cutting is accompanied by an increase in the amount of internet usage, especially through OTTV.

If consumers' internet usage choices remained fixed after cutting the cord, this subscription change alone would have a major impact on ISP revenue. Consumers' behavioral changes following cord-cutting, however, amplify the negative impact of cord-cutting on ISP profits. Specifically, once TV service is dropped, consumers often re-allocate their time toward online applications that require greater data transmission and therefore greater network costs. We quantify this shift by analyzing the daily internet usage of cord-cutters during the eight weeks before and after dropping their TV subscriptions. Relative to the statistics on Table 3, this analysis provides a tighter focus on activity immediately before and after the cord-cutting decision.

In Figure 5 Panel (a) we plot the shift in average internet usage within cord-cutting households, and, for comparison, include the average usage of consumers who kept their TV subscriptions. Each "Fitted Line" is an estimate of the mean per-household usage on each day of the 112-day sample window featured in the Figure, with estimates presented separately for cord-cutting (CC) and bundle (non-CC) households. We center each cord-cutting observation on the date of the subscription change ("Day 0"), and we present mean internet usage levels for the same date for households that do not cut the cord. We supplement the fitted lines with estimates of the average levels of internet usage in the weeks before and after cord-cutting. Together, these estimates show that the average (eventual) cord-cutter's internet usage begins approximately 25% higher than other subscribers, and increases by 23% after cutting the cord (4.9 GB/day to 6.0 GB/day). Thus, there is strong selection among the individuals who drop TV in our sample (i.e., heavier internet users on average), and following the drop these consumers substantially increase internet usage.

Our 23% estimate of cord-cutting's impact is likely to be conservative. The data suggest that many consumers "prepare" to drop their TV subscriptions by adding and using online video services during the weeks before cutting the cord.[21] In Figure 5 Panel (a) the cord-cutters' increased usage begins prior to the subscription change date, and consumers who do not alter their subscriptions do not have a corresponding usage change. To decompose this activity by cord-cutters, we examine changes in consumer engagement with three popular subscription-based video services: Netflix, Hulu, and SlingTV. In Figure 5 Panel (b) we provide the change in the number of active users for each video service in the eight weeks before consumers drop TV. Among cord cutters, Netflix adds the most active users in absolute terms, although this is a relatively small percentage increase due to Netflix's high existing penetration rate. Sling TV, arguably the closest substitute for pay-TV, adds a very similar number to Hulu in absolute terms but dramatically more in percentage terms.[22]

We explore the nature of cord-cutters' changing internet usage by decomposing the

[21] On the other hand, some growth could be due to underlying organic growth in internet use that

Figure 5: *Average Daily Usage and OTTV Adoption for Cord-Cutters*

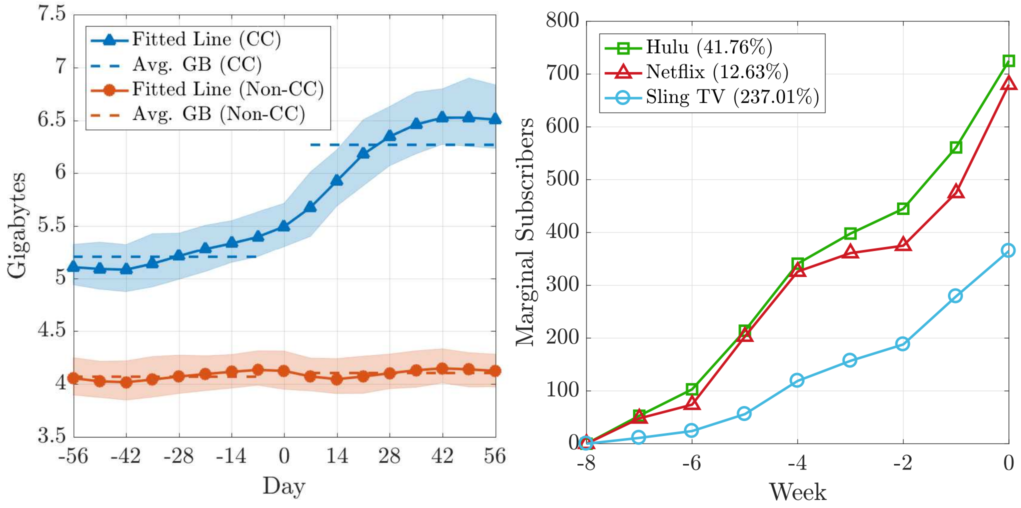

(a) Usage (b) OTTV Adoption

Notes: Panel (a) summarizes the average daily usage (GBs) of internet subscribers in the eight weeks before and after dropping pay-TV services (upper plots) in comparison to all other internet subscribers (lower plots). The points are average daily usage, and the dotted lines are average usage across days before and after the change. One week on either side of the cord-cut date is omitted. Panel (b) shows the change in active engagement with select OTTV services by cord-cutters in the eight weeks prior to dropping pay-TV services. As a baseline, a household is considered a subscriber to every OTTV service engaged with prior to eight weeks before the cord-cut date. The total count of active subscribers is then determined in each week prior to the cord-cut date, and the differences between these weekly counts and the baseline count are plotted.

23% increase in traffic across types of applications. In Figure 6 Panel (a) we provide the change in usage across Browsing, Gaming, Music, Video Streaming, and Other applications. The online activities that increase the most are those that seem to provide the closest substitute for watching TV. Online video and webpage browsing have the largest absolute and proportional gains. Usage of each service grows by about a third, and together they account for about 90% of an average internet-only customer's data usage. (Some online video that is delivered through websites, e.g. video on Facebook, appears in our data as browsing rather than video streaming.) Other entertainment applications (Gaming and Music) which are less similar to TV have smaller increases in usage after cord-cutting; cloud-based backup services (e.g. Dropbox) also have more modest usage increases. In Figure 6 Panel (b) we extend the results in Panel (a) to decompose the 30% increase in online video across individual OTTV services. Among cord cutters, Hulu usage increased by 92%, Netflix by 32%, and SlingTV by 695%.[23]

is not related to dropping TV.

[22]The number of YouTube actually decreases slightly, but this is not surprising as this free service has nearly 100% penetration prior to cord-cutting.

[23]Use of other relatively less-popular streaming services, like Amazon and HBO Go, actually de-

Figure 6: *Average Daily Usage by Application Type for Cord-Cutters*

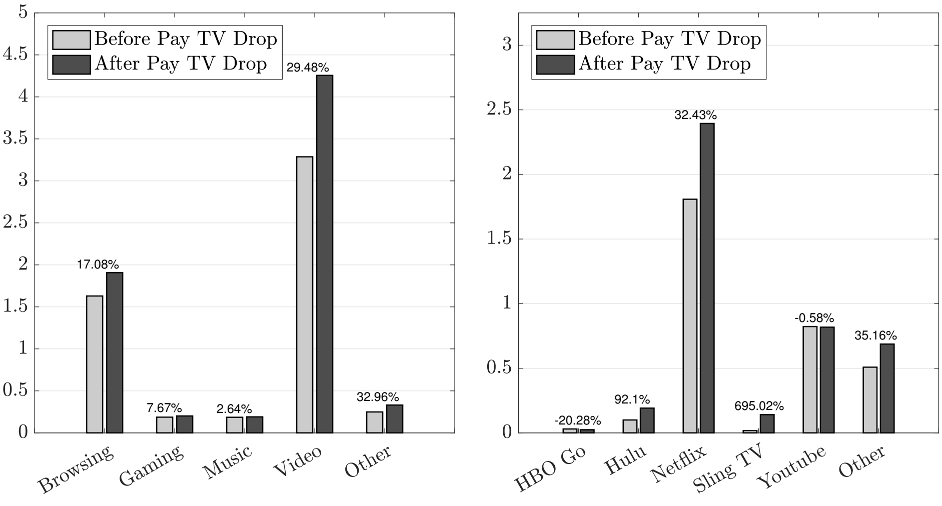

(a) Total Usage by Type (b) OTTV by Type

Notes: Panel (a) summarizes the average daily usage (GBs) of cord-cutters in the eight weeks before and after dropping pay-TV services. Total usage is broken up into traffic types. For each traffic type, the percent change in usage level between the two periods is shown above the two bars. Panel (b) summarizes the average daily usage (GBs) of online video applications by cord-cutters in the eight weeks before and after dropping pay-TV services. For each application, the percent change in usage level between the two periods is shown above the two bars.

SlingTV's impact is particularly noteworthy because it provides essentially the same live programming that ISPs broadcast across their networks for their own TV customers, but it imposes a substantial marginal burden on broadband networks because SlingTV is transmitted separately to each customer (potentially multiple streams simultaneously per household).

Consumers who cut the cord sacrifice a portion of their available video entertainment, and in exchange they pay substantially less in total for the entertainment they receive. We analyze this change in payments by focusing on consumers who begin the sample with a bundle subscription and, at some point, drop TV but remain internet customers. Average payments from these consumers to the ISP fall by about 50%, from \$132 to \$68. We also calculate estimated total out-of-pocket spending by consumers to the ISP plus all subscription services using 2015 prices. In Figure 7 we plot the densities of consumers' total payments before and after cutting the cord. Average total spending per customer drops by half when the consumer drops a TV subscription. The variation around each distribution's mode is due to differences in consumers' subscription details (i.e., channel packages and internet tiers) and third-party OTTV subscriptions. The expenditure change is reversed, of course, if the ISP can induce an

crease over this period, due in part to the end of seasonal programming. In addition, at the very end of our sample period we encountered encryption issues that interfered with our ability to observe these particular services.

internet-only subscriber to add video.

Figure 7: *Estimated Total Subscriber Monthly Payments with OTTV*

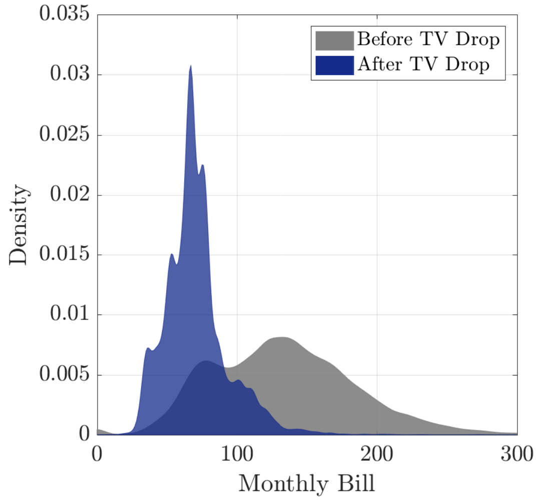

Notes: This figure shows the distribution of estimated monthly payments by cord-cutters in the month before and after dropping pay TV. The total monthly operator bill is observed in the ISP billing data. Monthly OTTV expenditures are estimated using 2015 subscription fees for each of the video services observed in the usage data.

4 Empirical Analysis of Steering

In our main empirical analysis, we exploit our sample's cross-sectional and temporal variation in UBP policy to estimate the impact of these prices on households, the ISP, and OTTV providers. We divide our empirical analysis into two main pieces. First, we examine the impact of UBP on households' subscription choices. Depending on a household's pre-policy subscription portfolio, this may include opportunities to add a video subscription, cut the cord, and upgrade its internet tier. Second, we analyze UBP's impact on internet usage choices, which are typically different conditional on households' subscription decisions. Finally, we synthesize these analyses in a discussion of the overall impact of UBP on households and firms, including how our results inform contemporary telecommunications policy issues such as net neutrality and vertical merger approvals in the industry.

4.1 Descriptive evidence

In our analysis of video additions, we focus on households that begin the sample period with an internet-only subscription. We illustrate some critical features of the data in Figure 8, where we plot the cumulative share of households that transition from internet-only to the bundle during the full sample period. There is a regular trickle of consumers into the bundle during the pre-policy and announcement periods.

Figure 8: *Cumulative New Video Subscriptions*

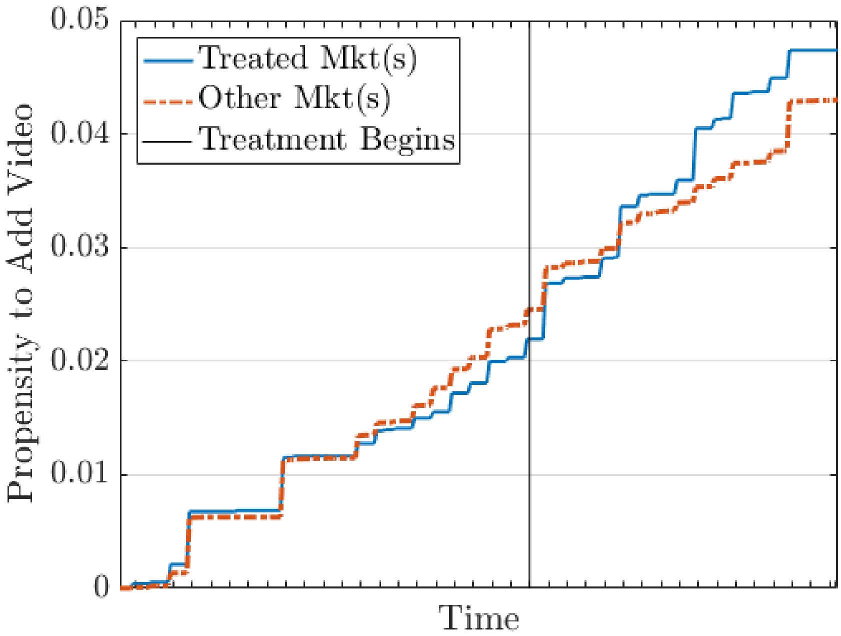

Notes: This figure provides the cumulative share of consumers who chose to switch from internet-only to bundle subscriptions. The hashmarks on the "Time" axis indicate weeks during the sample period.

For the first several months, the treated and control markets accumulate new bundle subscribers at nearly identical rates; in the weeks immediately before the treatment period, the control markets have a greater rate of video additions. Once the treatment period begins, however, households exposed to UBP appear to be more likely to switch to the bundle. By the end of the sample period, a greater share of households in the treated market has added video subscriptions.

We also provide a preliminary examination of differences in video-addition behavior across households with different characteristics. In Figure 9 we plot the difference between treatment and control markets in the propensity to add a video subscription as a function of the household's usage intensity (φ) during the pre-policy period. (We estimate this difference with a kernel regression to account for variation in φ.) The overall difference is positive but quantitatively small, and Netflix users with large values of φ appear somewhat more likely than average to add video if they are in the treated market. The largest differences exist for heavy-usage linear-OTTV subscribers. These households appear to be substantially more likely to add video under the treatment. In our empirical analysis we evaluate whether this difference is statistically significant and economically meaningful.

We use a similar approach to analyze households' decisions to upgrade their internet tiers in response to UBP. We begin by describing some key patterns in the data.

Figure 9: *Differences Across Markets in Video Additions*

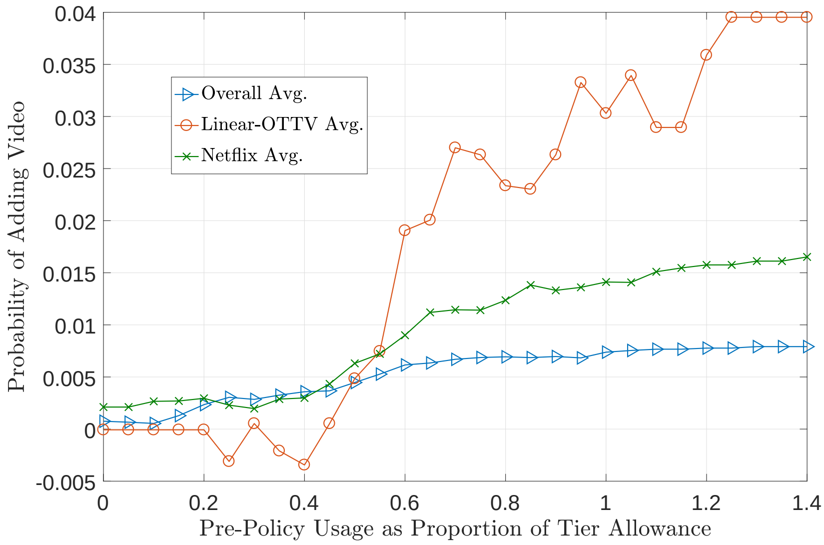

Notes: This figure displays between-market differences in video addition probabilities. The horizontal axis provides the household's internet usage intensity during the pre-policy period. The plotted lines are the differences between treatment and control markets for all households overall, linear-OTTV subscribers, and Netflix subscribers.

In Figure 10 we display the cumulative share of tier-upgrading actions in the treatment and control markets. (These shares are smoother than the video addition shares because nearly all households have the opportunity to upgrade their internet tier, and more choose to do so.) We ignore the option to downgrade internet because very few households take this action.[24] The Figure shows clearly that upgrading choices were more common in the treatment market than the control markets, and this started in the weeks before the UBP policy was enforced. The rate of upgrading choices diverged during the announcement period, when households in the treated market were informed that overage charges would be coming soon.

In Figure 11 we provide a preliminary analysis of whether tier-upgrading responses to UBP varied by households' internet usage intensity or their subscriptions to streaming video services. The differences between treatment and control markets grow steadily with φ, and there is little difference among the overall population, Netflix subscribers, and linear-OTTV subscribers. The plotted differences suggest that households whose pre-policy usage was near their tier allowances ($\varphi = 1$) have an upgrade probability that is 20 percentage points greater in the treated market than in the controls.

[24]Low-usage households in the treated market had about a 1% chance of reducing their internet tier. While the downgrading frequency is fairly small, it is slightly more frequent in the treated market during the Full-UBP period. This could be due to an informational effect of the pricing program, in some cases alerting them that their usage could be reduced to a lower-priced tier.

Figure 10: *Cumulative Tier Upgrades*

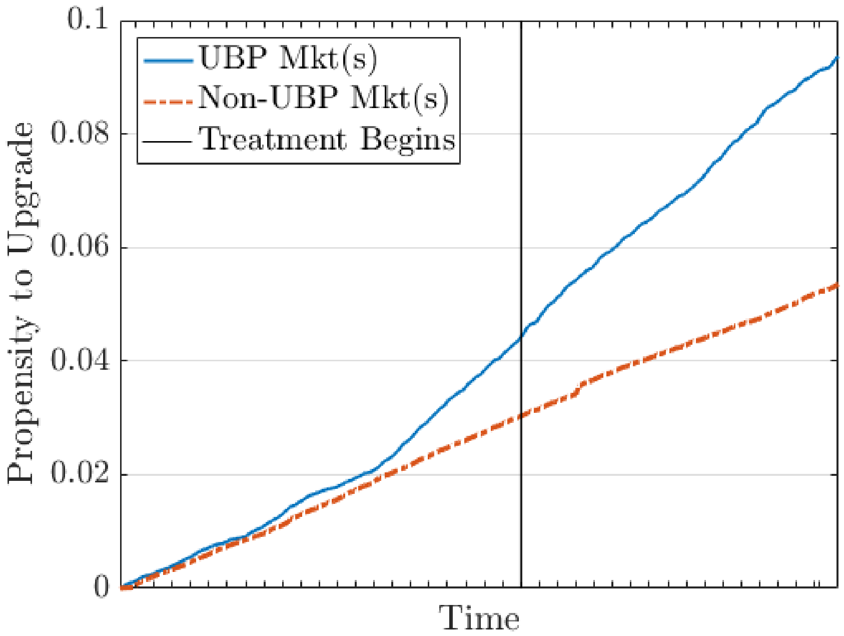

Notes: This figure provides the cumulative share of consumers who chose to upgrade their internet tier. The hashmarks on the "Time" axis indicate weeks during the sample period.

Figure 11: *Differences Across Markets in Tier Upgrades*

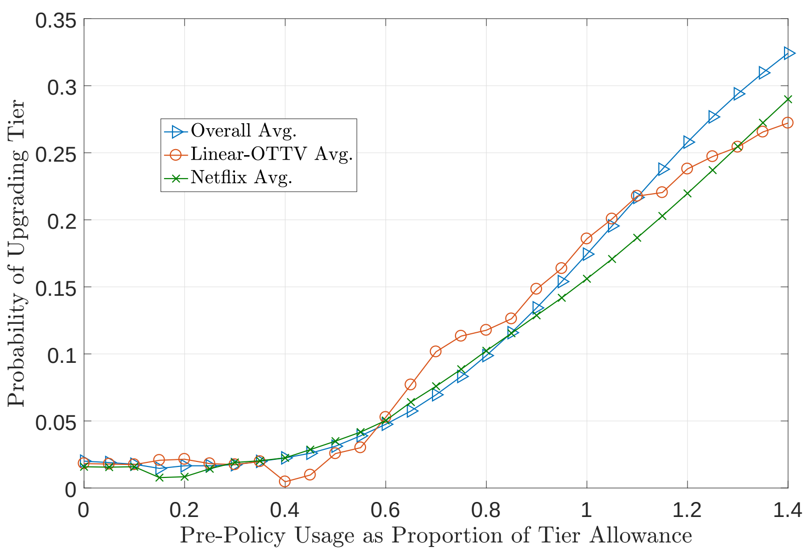

Notes: This figure displays between-market differences in internet tier upgrade probabilities. The horizontal axis provides the household's internet usage intensity during the pre-policy period. The plotted lines are the differences between treatment and control markets for all households overall, linear-OTTV subscribers, and Netflix subscribers.

4.2 Analysis of subscription choices

To test steering's impact on households' subscription choices, we specify a model of transitions across services (TV and internet) and internet tiers. Each household begins the sample period as internet-only or with the bundle, and each has an observed internet tier. We ask whether and how a household changes its subscription following the ISP's initial announcement of its UBP policy in the treated market. For internet-only households, we investigate whether they add a TV subscription, upgrade their internet tier, or both. For bundle subscribers, we examine the likelihood that they cut the cord, upgraded internet tier, or both. Very few households downgrade internet during the sample period, and very few start the period in the highest internet tier, so we do not consider these options in our analysis.[25] We use j to index the individual actions that a household might take. In grouping together transitions that occur during the announcement and treatment periods, we acknowledge that information about upcoming UBP policy may cause some households to update their subscriptions.

We use a multinomial logit (MNL) model to estimate the probability that a household makes each possible change to its subscription portfolio, given its choice set. We allow these probabilities to be functions of the household's exposure to the treatment, its pre-policy usage intensity, and some of its pre-policy OTTV subscriptions. We use h to index individual households, and we define *Treat* as an indicator for whether h resides in the treated market. We incorporate households' characteristics in two ways. First, we define *Heavy* as an indicator for whether a household's pre-policy usage intensity $\varphi > 0.5$. Second, we create the dummy variables *Netflix* and *Linear* to indicate whether h is an active user of Netflix or a linear-OTTV service (i.e. Sling TV or Hulu), respectively, during the pre-policy period. Additional household heterogeneity is in $\epsilon_{h,j}$, which is distributed type I extreme value.

We specify that a household's utility from selecting transition option j is:

$$\begin{aligned} u_{h,j} =& \beta_{j,0} + \beta_{j,T} Treat + \beta_{j,H} Heavy + \beta_{j,HT} Heavy \times Treat + \\ & \beta_{j,HL} Heavy \times Linear + \beta_{j,HLT} Heavy \times Linear \times Treat + \\ & \beta_{j,HN} Heavy \times Netflix + \beta_{j,HNT} Heavy \times Netflix \times Treat + \epsilon_{h,j}. \end{aligned} \tag{1}$$

We normalize to zero the mean utility from retaining the original subscription, so $u_{h,0} = \epsilon_{h,0}$ for this option. This model allows for an overall effect ($\beta_{j,T}$) of UBP on all treated households, the effect ($\beta_{j,HT}$) of UBP on all treated heavy-usage households, and effects ($\beta_{j,HLT}$ and $\beta_{j,HNT}$) of UBP on heavy-usage households who also subscriber to linear OTTV or Netflix during the pre-policy period.

In reporting the estimation results, we focus on predicted transition probabilities in the treatment and control markets. We report the estimated parameter values in

[25] We treat households that downgrade internet as if they did not change their tier. We exclude households from the estimation sample if they were in the highest internet tier during the pre-policy period. We also drop all households that change their subscription portfolio during the pre-policy period.

Appendix Tables 12 and 13. We begin by describing the overall impact of UBP on households that were internet-only during the pre-policy period; see the "All households" results of Table 7. Treated households were more likely to add TV alone, upgrade their internet tier alone, or take both actions together. New TV subscriptions were 16% more common in the treated market, and the internet upgrade rate was more than twice as large among households subject to UBP. These differences are statistically significant at $p < 0.01$. Transitions in both TV subscription and internet tier were relatively rare in the overall population (about 0.5%), and the treatment effect is not statistically significant. Our findings for heavy-internet households are somewhat different. Treated households with $Heavy = 1$ were less likely to add TV alone ($p < 0.05$), but they were significantly more likely to add TV while also upgrading internet. The magnitudes of these differences are roughly equal. The greatest difference between treatment and control markets is in the probability of upgrading internet tier, which was four times greater among households subject to UBP.[26]

The remaining results in Table 7 provide the treatment effects for heavy-usage internet-only households that subscribed to linear OTTV or Netflix. Consistent with the results in Figure 9, we find that treated households with $Heavy \times Linear = 1$ are more likely add video, although this difference is entirely concentrated among households who also upgrade their internet tier. In the control markets, the probability of this outcome (0.6%) is about equal to its value in the overall population (0.4%), while in the treated market we estimate that 1.7% of households make this transition. This provides a substantial benefit to the ISP, as the households pay increased subscription prices for both TV service and upgraded internet, plus the subsequent reduction in internet usage (described below) represents a cost reduction for the ISP. An additional significant feature of this result is that it demonstrates nontrivial substitution between streaming video and conventional TV. In addition to new TV subscriptions, UBP stimulated a significant share of households with $Heavy \times Linear = 1$ to upgrade their internet tiers.

The impact of UBP is generally weaker for internet-only households with $Heavy \times Netflix = 1$. These households were more likely to upgrade their internet tier alone or jointly with adding TV, but the differences between treatment and control markets are smaller than for otherwise similar linear OTTV subscribers. In addition, the share of households that add TV alone is smaller in the treated market than in the control markets, and the magnitude of this difference is slightly greater than the positive difference for joint TV and internet upgrades. The differences between the results for internet-only Netflix and linear OTTV households supports the intuition that the

[26]Our results on the overall treatment effect and among households with $Heavy = 1$ are essentially identical if we estimate a simpler MNL model that allows for these effects only.

latter service is a closer substitute for conventional TV, and therefore substitution along this margin is more sensitive to changes in the internet pricing policy.

Table 7: Treatment effect on internet-only subscriber transition probabilities

Treated Population		Add TV	Upgrade Tier	Add and Upgrade
All households	Control	2.62	3.05	0.41
	Treatment	3.06	7.00	0.56
	Difference	0.44	3.95	0.15
	(SE diff.)	(0.19)	(0.28)	(0.09)
$Heavy = 1$	Control	2.95	4.08	0.52
	Treatment	2.33	17.04	0.98
	Difference	-0.62	12.96	0.47
	(SE diff.)	(0.32)	(0.75)	(0.20)
$Heavy \times Linear = 1$	Control	3.57	4.03	0.60
	Treatment	3.47	21.81	1.72
	Difference	-0.11	17.78	1.12
	(SE diff.)	(0.70)	(1.57)	(0.49)
$Heavy \times Netflix = 1$	Control	3.25	4.11	0.52
	Treatment	2.55	17.34	1.08
	Difference	-0.70	13.23	0.56
	(SE diff.)	(0.36)	(0.83)	(0.23)

In Table 8 we provide a parallel set of results for households who subscribed to the bundle at the end of the pre-policy period. The possible actions for these households included cord-cutting and upgrading internet tier. In this analysis we omit the effects that capture linear OTTV subscribers the pre-policy period. Fairly few households subscribed to linear OTTV and conventional TV simultaneously, and some of those who did may not have viewed the portfolio as a long-term choice, but were instead experimenting with linear OTTV or in the midst of switching away from conventional OTTV.[27] Our estimates of the overall treatment effect, provided in the rows labeled "All households," show that treated households were 15% less likely ($p < 0.01$) to drop their TV subscriptions without changing their internet tier. Internet tier upgrades (alone) were almost twice as likely in the treated market (6.54% vs. 3.68%). More households in the treated market dropped TV and upgraded internet jointly, but the mass of consumers in this outcome is smaller than those steered away from cord-cutting alone. When we focus on bundle households with $Heavy = 1$, the differences between treatment and control markets are all larger. Cord-cutting alone was 29% less common in the treated market, and tier upgrades were more than three times as likely (5.08% vs. 18.61%). The share of households that both dropped TV and

[27]While linear OTTV provides some access and content options unavailable through conventional TV, the services were very similar during our sample period.

upgraded internet is also over three times as large in the treated market, and for these heavy-usage households the net effect of UBP on cord cutting is a slight increase in this outcome, although it is often accompanied by the (partially) countervailing outcome of a tier upgrade. In the final set of results in 8 we show that bundled households with $Heavy \times Netflix = 1$ are very similar to heavy-usage households without Netflix subscriptions: they are less likely to cord-cut alone, more likely to upgrade internet alone, and slightly more likely to both drop TV and upgrade internet jointly.

Table 8: Treatment effect on bundle subscriber transition probabilities

Treated Population		Drop TV	Upgrade Tier	Drop and Upgrade
All households	Control	2.23	3.68	0.27
	Treatment	1.87	6.54	0.53
	Difference	-0.36	2.86	0.26
	(SE diff.)	(0.09)	(0.15)	(0.05)
$Heavy = 1$	Control	3.25	5.08	0.55
	Treatment	2.29	18.61	2.11
	Difference	-0.96	13.53	1.55
	(SE diff.)	(0.28)	(0.70)	(0.26)
$Heavy \times Netflix = 1$	Control	3.25	4.11	0.52
	Treatment	2.55	17.34	1.08
	Difference	-0.70	13.23	0.56
	(SE diff.)	(0.36)	(0.83)	(0.23)

4.3 Analysis of usage choices

The ISP's UBP policy has two types of benefit for the firm. First, it may induce switches in households' subscriptions which increase revenue; this is the impact we studied above. Second, it may lead consumers to reduce their internet usage, which can reduce pressure on the ISP to make costly investments in network capacity. We now consider whether the ISP was able to induce reductions in internet usage following the introduction of UBP.

We begin by examining the case of initially internet-only households that added a conventional TV subscription in the treated market. In Figure 12 we compare internet usage changes for two groups of households who added video to their subscriptions. In Panel (a) we illustrate customers' internet choices in the pre-policy period. These households may have had no TV subscription prior to their switch, or they may have received TV channels from an unaffiliated satellite provider like DirecTV. For this group, the choice to add video coincides with no significant change in internet usage. By contrast, the households captured in Panel (b) add video after UBP is present in their home market, and this group significantly reduces internet usage. This change is consistent with the usage change described in Section 2's illustrative model. The reduction's magnitude (25%) is about equal in magnitude but opposite in sign to the change in internet usage that occurs after a household cuts the cord. Households induced by UBP to add video, therefore, appear to be steered away from heavy internet usage.

Figure 12: *Usage Before and After Addition of Video Service*

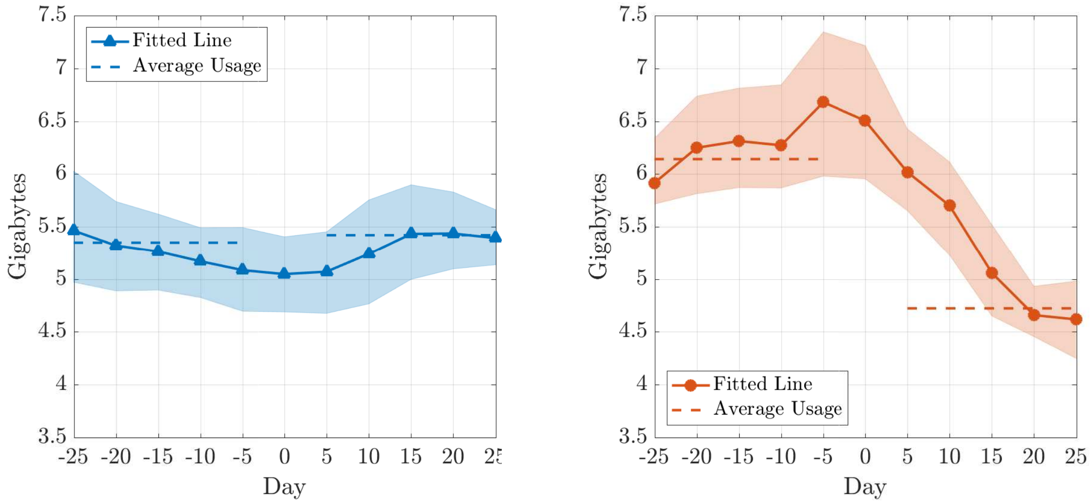

(a) Pre-UBP (b) Post-UBP

Notes: This figure presents statistics summarizing average daily usage (GBs) of customers on the days before and after adding pay-TV services. The dots are average daily usage, the hashed line is average usage across days before and after, and the solid line is a global local-linear regression fit of the data.

In addition to the revenue and cost implications for the ISP, there are also third-party services that may be impacted by steering. To examine this, in Table 9 we decompose by application the decrease in usage in Figure 12(b). Average usage of some types of applications like Music, Gaming, and Sharing increase slightly during our sample, while Streaming and Browsing remain at similar levels. The large and substantial usage decrease that we observe in Figure 12(b) is almost entirely due to a reduction in OTTV usage. The largest absolute reduction in usage is Netflix, while the largest percentage reduction is SlingTV. This mirrors the application-level usage changes by cord-cutting households that we describe in Section 3.4. The closest substitutes for pay-TV, linear-OTTV services, are most heavily impacted by steering, while other on-demand services are impacted to a lesser degree.

To complement our analysis of internet usage change by households that add video, we also examine how usage changes among households that upgraded (or did not upgrade) their internet tier. We explore this outcome by decomposing usage choices by households' decisions to upgrade. For all households that subscribe to the ISP's

Table 9: Change in Usage After Pay TV Add During UBP Implementation

Traffic Type	Average Usage (GBs)	
	Pre-Add	*Post-Add*
Browsing	1.67	1.43
Gaming	0.20	0.15
Music	0.18	0.16
OTTV	2.911	2.018
HBO Go	0.03	0.02
Hulu	0.17	0.07
Netflix	2.52	1.40
Sling TV	0.05	0.02
YouTube	1.00	0.74
Other Video	0.56	0.49
Other	0.26	0.21

Notes: This table reports average usage by traffic type for subscribers that add pay TV after UBP is implemented. Each average is taken across subscriber days for one month before and one month after adding pay TV.

most popular tier at the start of the sample period, we record their value of φ, whether they were exposed to UBP, their decision to upgrade or not, and the difference in their internet usage (in GBs) during the sample period.[28] We plot these usage differences in Figure 13.

Across all markets, households that upgraded their internet tier generally had more usage growth than households that did not upgrade, as expected. The usage growth patterns among upgrading households are very similar for low values of φ across markets. Around $\varphi = 0.6$, upgrading households begin to diverge between the treated and control markets. This could be due to a combination of selection (less intense users nevertheless upgrade under UBP) and a behavioral response (despite upgrading, households become more cautious in using data). A more striking difference exists among households that did not upgrade. While usage growth is smaller for all non-upgrading households in the treated market, the difference is especially large for $\varphi > 0.8$. These households reduce their internet usage during the sample period, which is striking considering the general trend toward usage growth in the U.S. during recent years. The introduction of UBP, therefore, appears successful at both stimulating tier upgrades and slowing internet usage growth.

[28]We focus within this tier so that all households usage intensities (φ) are calculated for the same allowance, plus usage changes can be expressed in quantities. The Figure is essentially the same if we construct it using all households regardless of tier.

Figure 13: *Changes in Internet Usage across Markets and Tier Choices*

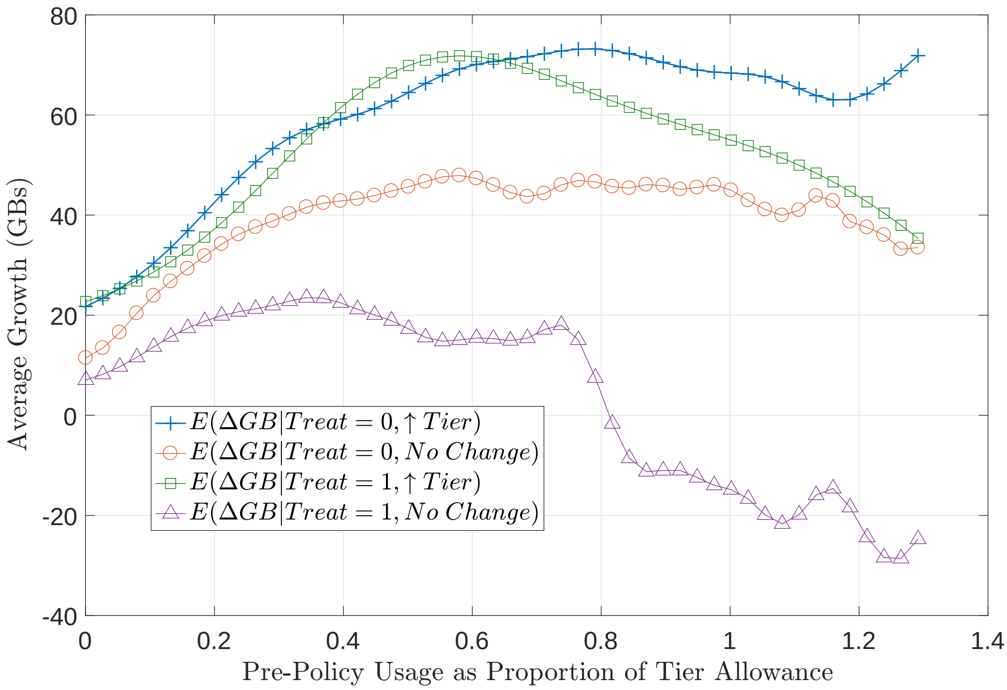

Notes: This figure displays the difference in internet usage (ΔGB) during the sample period for households divided by their pre-policy usage intensity (φ), their exposure to UBP (*Treat*), and their tier upgrade decision ($\uparrow$ *Tier* or *No Change*).

5 Conclusions

We study the impact of increasingly popular online video services on the telecommunications industry. OTTV represents an opportunity for ISPs in that streaming video increases the demand for internet subscription services, but the new applications also present several important challenges. OTTV improvements reduce demand for ISPs' television services, and delivering internet content can substantially increase ISPs' costs. The firms have responded with nonlinear pricing strategies to steer consumers across subscription options and possibly curtail internet and OTTV usage.

We provide an illustrative model that describes the some of the central incentives behind ISPs' prices, and we use panel data on household-level subscriptions and internet usage to evaluate the model's predictions. Our model shows that bundle discounts, usage limits, and usage tiers are valuable instruments for steering consumers. In our empirical analysis, we find that telecommunications firms successfully steer consumers to higher-priced subscriptions by using nonlinear pricing for internet services. In addition to changes in consumers' subscriptions, steering curtails growth in internet usage – in some cases reducing usage despite widespread trends toward positive growth. These general outcomes, along with more nuanced results regarding the heterogeneity in steering's impact, have a number of important implications for ongoing policy debates.

Steering's impact is relevant for antitrust policy related to the telecommunications industry. In particular, evaluation of mergers between content and distribution firms present a number of challenges. First, market boundaries may be difficult for regulators and antitrust authorities to identify because little evidence exists on consumers' willingness to substitute across conventional TV, streaming video, and other non-video internet applications. Our results show that consumers are willing to substitute among online entertainment sources and with conventional TV. Specifically, we find that cord-cutters increase their usage of most online applications after dropping an ISP's TV service, and these increases are roughly proportional to usage levels prior to the subscription change.[29] Thus, telecommunications antitrust analysis might consider broad market definitions that encompass many forms of digital entertainment (e.g., Facebook and NBC compete for a consumer's time), as well as the central role of ISPs in shaping how content is distributed and surplus is allocated.

Second, antitrust authorities need to assess how existing or new vertical relationships may affect a ISPs' incentives to introduce restrictive cross-licensing agreements or use price instruments to favor its own content over competitors'. The impact of these strategies depends on consumers' sensitivity to steering strategies, which we show can be significant. An ISP that is vertically integrated with a content-producing firm may foreclose some content from availability to consumers via a competing ISP. A price-based steering strategy with similar effects is "zero rating," which favors certain content by not counting its usage against a monthly allowance. Our elasticity estimates show that even blunt mechanisms like usage-based pricing can have important allocative consequences among consumers and various firms. Firms may perceive even greater benefits of zero rating than UBP because a targeted mechanism is likely to be more profitable.

More broadly, our results are also relevant for the net neutrality debate, in which empirical evidence is rare. Net neutrality's repeal provides ISPs more latitude to discriminate across types of internet traffic. While we do not observe source-specific discrimination in our analysis, our results are informative about ISPs' incentives to discriminate when they have the opportunity. For example, we find that usage-based pricing's primary impact is in inducing consumers to upgrade their internet tiers and continuing using their preferred online applications (e.g., Netflix), although there are some exceptions like linear OTTV, where bundle subscriptions shift significantly. If

[29] Consumers who are steered into the bundle appear to reduce only their OTTV usage, but it is difficult to untangle the policy's effect from general trends toward greater internet usage in every type of application.

an ISP can successfully use tier premiums to extract some of the rents associated with OTTV innovations, it may not seek more targeted mechanisms to foreclose or diminish the attractiveness of OTTV. These incentives may change, however, as firms diversify and vertically integrate into media production, so more research is needed in this area.

There are several issues that our model and empirical results do not address, and that we leave for future research. While our model provides a useful framework for formalizing the steering incentives of ISPs, a richer specification is required to quantify the welfare implications of steering. Similarly, the model makes simplifying assumptions on the interaction between firms, for example competitive OTTV supply. Given the increasingly complex relationships between content providers and ISPs, and the evolving regulatory and antitrust environment, modeling and evaluating these policy issues is a fruitful area for future research.

References

Armstrong, Mark (2006). "Competition in Two-Sided Markets." *RAND Journal of Economics*, 37(3): 668–691.

Armstrong, Mark (2013). "A More General Theory of Commodity Bundling." *Journal of Economic Theory*, 148(2): 448–472.

Barwick, Panle Jia, Parag A. Pathak and Maisy Wong (2017). "Conflicts of Interest and Steering in Residential Brokerage." *American Economic Journal: Applied Economics*, 9(3): 191–222.

Bourreau, Marc, Frago Kourandi and Tommaso Valletti (2015). "Net Neutrality with Competing Internet Platforms." *Journal of Industrial Economics*, 63(1): 30–73.

Choi, Jay Pil, Doh-Shin Jeon and Byung-Cheol Kim (2015). "Network Neutrality, Business Models, and Internet Interconnection." *American Economic Journal: Microeconomics*, 7(3): 104–141.

Choi, Jay Pil and Byung-Cheol Kim (2010). "Net Neutrality and Investment Incentives." *RAND Journal of Economics*, 41(3): 446–471.

Chu, Chenghuan, Phillip Leslie and Alan Sorensen (2011). "The Welfare Effects of Vertical Integration in Multichannel Television Markets." *American Economic Review*, 101(1): 263–303.

Crawford, Greg, Robin Lee, Mike Whinston and Ali Yurukoglu (2017). "The Welfare Effects of Vertical Integration in Multichannel Television Markets." *Working Paper.*

Crawford, Gregory and Matthew Shum (2007). "Monopoly Quality Degradation and Regulation in Cable Television." *Journal of Law and Economics*, 50(1): 181–219.

Crawford, Gregory and Matthew Shum (2015). "The Welfare Effects of Monopoly Quality Choice: Evidence From Cable Television Markets." *Working Paper.*

Crawford, Gregory and Ali Yurukoglu (2012). "The Welfare Effects of Bundling in Multichannel Television Markets." *American Economic Review*, 102(2): 643–685.

Dutz, Mark, Jonathan Orszag and Robert Willig (2009). "The Substantial Consumber Benefits of Broadband Connectivity for US Households." *Internet Intervention Alliance Working Paper.*

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

Economides, Nicholas and Benjamin Hermalin (2015). "The Strategic Use of Download Limits by a Monopoly Platform." *The RAND Journal of Economics*, 46(2): 297–327.

Economides, Nicholas and Joacim Tag (2012). "Network Neutrality on the Internet: A Two-Sided Market Analysis." *Information Economics and Policy*, 24(2): 91–104.

Economides, Nicholas and Joacim Tag (2016). "Internet Regulation, Two-Sided Pricing, and Sponsored Data." *Working Paper*.

Edell, Richard and Pravin Varaiya (2002). *Providing Internet Access: What We Learn from INDEX*, volume Broadband: Should We Regulate High-Speed Internet Access? Brookings Institution.

Gans, Joshua (2015). "Weak Versus Strong Net Neutrality." *Journal of Regulatory Economics*, 47(2): 183–200.

Gentzkow, Matthew (2007). "Valuing New Goods in a Model with Complementarity: Online Newspapers." *American Economic Review*, 97(3): 713–744.

Goetz, Daniel (2016). "Competition and Dynamic Bargaining in the Broadband Industry." *Princeton University Working Paper*.

Goolsbee, Austan and Peter Klenow (2006). "Valuing Products by the Time Spent Using Them: An Application to the Internet." *American Economic Review P&P*, 96(2): 108–113.

Greenstein, Shane and Ryan McDevitt (2011). "The Broadband Bonus: Estimating Broadband Internet's Economic Value." *Telecommunications Policy*, 35(7): 617–632.

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

Grubb, Michael (2015). "Consumer Inattention and Bill-Shock Regulation." *Review of Economic Studies*, 82(1): 219–257.

Grubb, Michael and Matthew Osborne (2015). "Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock." *American Economic Review*, 105(1): 234–271.

Hagemann, Garrett (2017). "Upstream Quantity Discounts and Double Marginalization in the New York Liquor Market." *Working Paper*.

Hitte, Loran and Prasanna Tambe (2007). "Broadband Adoption and Content Consumption." *Information Economics and Policy*, 74(6): 1637–1673.

Ho, Kate and Robin Lee (2017). "Equilibrium Insurer-Provider Networks: Bargaining and Exclusion in Health Care Markets." *Working Paper.*

Lambrecht, Anja, Katja Seim and Bernd Skiera (2007). "Does Uncertainty Matter? Consumer Behavior Under Three-Part Tariffs." *Marketing Science*, 26(5): 698–710.

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

Liebman, Eli (2017). "Bargaining in Markets with Exclusion: An Analysis of Health Insurance Networks." *Working Paper.*

Malone, Jacob, Aviv Nevo and Jonathan Williams (2016). "A Snapshot of the Current State of Residential Broadband Networks." *NET Institute Working Paper No. 15-06.*

Malone, Jacob, John Turner and Jonathan Williams (2014). "Do Three-Part Tariffs Improve Efficiency in Residential Broadband Networks?" *Telecommunications Policy*, 38(11): 1035–1045.

McManus, Brian (2007). "Nonlinear Pricing in an Oligopoly Market: the Case of Specialty Coffee." *RAND Journal of Economics*, 38: 512–532.

Miravete, Eugenio (2003). "Choosing the Wrong Calling Plan? Ignorance and Learning." *American Economic Review*, 93(1): 297–310.

Mussa, Michael and Sherwin Rosen (1978). "Monopoly and Product Quality." *Journal of Economic Theory*, 18(2): 301–317.

Nevo, Aviv, John Turner and Jonathan Williams (2016). "Usage-Based Pricing and Demand for Residential Broadband." *Econometrica*, 84(2): 411–443.

Raval, Devesh and Ted Rosenbaum (2017). "Payer Steering and Patient Welfare: Evidence from Medicaid." *Working Paper.*

Reggiani, Carlo and Tommaso Valletti (2016). "Net Neutrality and Innovation at the Core and at the Edge." *International Journal of Industrial Organization*, 45(1): 16–27.

Rosston, Gregory, Scott Savage and Bradley Wimmer (2013). "Effect of Network Unbundling on Retail Price: Evidence from the Telecommunications Act of 1996." *Journal of Law and Economics*, 56(2): 487–519.

Sidak, Gregory (2006). "A Consumer-Welfare Approach to Network Neutrality Regulation of the Internet." *Journal of Competition Law & Economics*, 2(3): 349–474.

Tudon, Jose (2018). "Congestion v. Content Provision in Net Neutrality: The Case of Amazon's Twitch.tv." *Working Paper*.

Varian, Hal (2002). *The Demand for Bandwidth: Evidence from the INDEX Experiment*, volume Broadband: Should We Regulate High-Speed Internet Access? Brookings Institution.

Wu, Tim (2003). "Network Neutrality, Broadband Discrimination." *Journal of Telecommunications and High Technology Law*, 1(2): 141–178.

Appendix

Table 10: Description of Internet Traffic Groups

Group	*Description (Examples)*	*% of All Usage*
Admin	System administrative tasks (STUN, ICMP)	1.19
Backup	Online storage (Dropbox, SkyDrive)	0.58
Browsing	General web browsing (HTTP, Facebook)	26.70
CDN	Content delivery networks (Akamai, Level3)	2.95
Gaming	Online gaming (Xbox Live, Clash of Clans)	3.06
Music	Streaming music services (Spotify, Pandora)	3.40
Sharing	File sharing protocols (BitTorrent, FTP)	0.20
Streaming	Generic media streams (RTMP, Plex)	6.26
Tunnel	Security and remote access (SSH, ESP)	0.07
Video	Video streaming services (Netflix, YouTube)	55.47
Other	Anything not included in above groups	0.13

Table 11: Daily Usage Statistics by DPI Group, 2015

	count	mean	std	min	25%	50%	75%	95%	max
Tunnel	119,878,359	0.00	0.20	0.00	0.00	0.00	0.00	0.00	360.01
Music	119,878,359	0.12	0.66	0.00	0.00	0.00	0.07	0.54	836.33
Gaming	119,878,359	0.12	1.13	0.00	0.00	0.00	0.00	0.29	789.63
Admin	119,878,359	0.05	0.41	0.00	0.00	0.00	0.01	0.18	246.95
Backup	119,878,359	0.03	0.58	0.00	0.00	0.00	0.00	0.05	694.27
CDN	119,878,359	0.11	0.77	0.00	0.00	0.00	0.01	0.38	542.75
Other	119,878,359	0.03	0.45	0.00	0.00	0.00	0.00	0.07	593.60
Sharing	119,878,359	0.01	0.32	0.00	0.00	0.00	0.00	0.00	408.40
Browsing	119,878,359	1.19	3.49	0.00	0.12	0.43	1.16	4.39	1,191.10
Streaming	119,878,359	0.32	1.40	0.00	0.00	0.01	0.14	1.46	841.76
Video	119,878,359	1.93	4.29	0.00	0.00	0.10	1.86	9.91	582.08

Note: This table reports distributional statistics of subscriber-day observations in our panel broken out by online activity category.

Table 12: Responses of Unbundled Subscribers to UBP

	(1)			(2)		
	Add	Upgrade	Both	Add	Upgrade	Both
Heavy User	0.20***	0.47***	0.35***	-0.42***	0.43***	0.33
	(0.043)	(0.038)	(0.104)	(0.110)	(0.074)	(0.201)
Treated Market	0.32***	0.00	0.04	0.32***	0.00	0.04
	(0.077)	(0.086)	(0.223)	(0.077)	(0.086)	(0.223)
Heavy × Treat	-0.40***	1.58***	0.76**	-0.58	1.42***	-0.13
	(0.157)	(0.106)	(0.311)	(0.473)	(0.178)	(0.775)
Heavy × Linear				0.18**	-0.02	0.22
				(0.078)	(0.071)	(0.188)
Heavy × Treat × Linear				0.50*	0.50***	0.78*
				(0.281)	(0.134)	(0.449)
Heavy × Netflix				0.66***	0.06	-0.03
				(0.114)	(0.078)	(0.214)
Heavy × Treat × Netflix				-0.03	-0.00	0.64
				(0.490)	(0.170)	(0.778)
Constant	-3.64***	-3.59***	-5.54***	-3.64***	-3.59***	-5.54***
	(0.025)	(0.024)	(0.063)	(0.025)	(0.024)	(0.063)
Observations	104790			104790		

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 13: Responses of Bundled Subscribers to UBP

	(1)			(2)		
	Cut	Upgrade	Both	Cut	Upgrade	Both
Heavy User	0.50***	0.43***	0.99***	0.36***	0.41***	0.76***
	(0.035)	(0.028)	(0.090)	(0.085)	(0.065)	(0.215)
Treated Market	-0.10*	0.39***	0.44***	-0.10*	0.39***	0.44***
	(0.050)	(0.032)	(0.123)	(0.050)	(0.032)	(0.123)
Heavy × Treat	-0.08	1.07***	1.07***	-0.26	1.00***	0.38
	(0.133)	(0.061)	(0.189)	(0.370)	(0.142)	(0.558)
Heavy × Netflix				0.16*	0.02	0.27
				(0.090)	(0.069)	(0.224)
Heavy × Treat × Netflix				0.19	0.08	0.75
				(0.389)	(0.149)	(0.565)
Constant	-3.83***	-3.32***	-6.09***	-3.83***	-3.32***	-6.09***
	(0.017)	(0.013)	(0.051)	(0.017)	(0.013)	(0.051)
Observations	240370			240370		

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$