Does Pricing of Internet Usage Steer Consumers or Meter Usage? Evidence from a Pricing Experiment^{*}

Brian McManus Aviv Nevo Zachary Nolan Jonathan W. Williams

March 7, 2024

Abstract

Regulatory agencies have expressed concerns that usage-based pricing (UBP) of internet service steers consumers from streaming video to traditional TV subscriptions. We study this issue with household-level panel data from an internet service provider's UBP experiment, capturing the pricing strategy's effects on internet and TV subscriptions, application-specific internet usage, and payments to the firm. UBP served largely to meter internet usage by high-demand households rather than steering them toward TV. Households' payments increased due to usage-related overage charges and internet subscription upgrades to avoid overages. Households that avoided internet-related payments reduced their internet usage rather than adding TV subscriptions.

Keywords: Usage-based Pricing, Steering, Metering, Bundling, Telecommunications Industry, Broadband Internet, Net Neutrality

JEL Codes: L11, L13, L96.

^{*}We thank the North American internet service provider that provided the data used in this paper. This paper is one of two that derive from a now-superseded earlier manuscript, "Steering Incentives of Platforms: Evidence from the Telecommunications Industry." The other current paper, "The Steering Incentives of Gatekeepers in the Telecommunications Industry," is available on the authors' websites. McManus: Department of Economics, University of North Carolina at Chapel Hill, mcmanusb@unc.edu; Nevo: Department of Marketing, University of Arizona, znolan@arizona.edu; Williams: Department of Economics, University of North Carol. Nevo is on leave from the University of Pennsylvania while working at the Federal Trade Commission (FTC). The views expressed in this paper are those of the authors and do not necessarily represent the views of the FTC or any individual Commissioner.

1 Introduction

The pricing of residential broadband is an ongoing focus of government policy. One specific concern is the incentive of Internet Service Providers (ISPs) to disadvantage, through pricing and other means, third-party internet-based services that compete with the ISP's own services. The Federal Communications Commission's (FCC's) 2010 Open Internet Order asserts that "broadband providers have incentives to interfere with the operation of third-party Internet-based services that compete with the providers' revenue-generating telephony and/or pay-television services."¹ To address this, the Order provides a set of rules ISPs should follow, including transparency in network management practices and a ban on "unreasonable discrimination" against lawful internet traffic.²

A primary way in which regulators have claimed ISPs could disadvantage third parties is by implementing usage-based pricing (UBP). In the US, Comcast, AT&T, Spectrum, Cox, and numerous other ISPs have implemented UBP in the form of multi-part tariffs. These tariffs typically consist of a fixed monthly access fee, a monthly usage allowance, and an overage schedule that applies additional fees to internet usage in excess of the allowance. Regulators have raised concerns that UBP offers ISPs a mechanism to raise the cost of over-the-top video (OTT) services, which currently account for the majority of internet traffic. ISPs that sell both internet access and TV service (also referred to as multiple-system operators, or MSOs) may benefit if raising OTT's price causes consumers to substitute to TV service. Indeed, concerns over potential harm to consumers and OTT providers motivated the restrictions on UBP in Charter's acquisition of Time Warner Cable in 2016.³ On the other hand, metering usage ensures that users who consume more also pay more, avoiding a regressive subsidy between users. A marginal price may also result in more efficient usage

 $^{^1\}mathrm{Preserving}$ the Open Internet, Report and Order, 25 FCC Rcd 17905, 17916, para. 22 (2010 Open Internet Order.

²The FCC's approach to regulating internet service providers is still evolving. After rolling-back Obamaera guidelines in 2017, in 2023 the FCC voted to reinstate the policy (https://docs.fcc.gov/public/ attachments/DOC-397827A1.pdf).

³https://docs.fcc.gov/public/attachments/FCC-16-59A1.pdf.

of networks by eliminating low-value traffic.⁴

The contrasting arguments on UBP's effects are based on theoretical considerations that make implicit assumptions about how consumers behave if faced with UBP. The concern over disadvantaging OTT assumes that, in response to UBP, consumers will be steered towards the MSO's own services and away from OTT. The argument that UBP meters usage and will lead to more efficient network use assumes that consumers will sort to plans and usage patterns that align with their (marginal) willingness to pay for usage. Metering's potential sorting effect suggests another motivation for an MSO to introduce UBP: to capture more of the surplus from internet usage and streaming video. In this paper, we use a unique dataset to measure how UBP's introduction actually affects consumers' subscription choices, internet usage, and payments to the MSO. We find that UBP's main impacts are on consumers' internet usage and the associated payments. Consumers who opted to continue relatively high internet usage did so by increasing payments to the MSO, while those who avoided usage payments did so by reducing internet usage rather than adding TV subscriptions.

To build intuition for our empirical analysis and the interpretation of our results, we begin by offering a simple model to illustrate an MSO's incentives and consumers' potential responses to UBP. We augment the standard mixed bundling model to accommodate several relevant characteristics of the TV-internet bundles offered by MSOs. Specifically, we allow "internet-only" subscribers to receive a share of TV's video entertainment through OTT streaming services, and we account for intensive-margin decisions about internet usage. As OTT improves, internet subscriptions offer greater consumption opportunities and become closer substitutes for TV subscriptions. This provides a dual motivation for an MSO to introduce UBP for internet usage. To the extent consumers continue using OTT, UBP can meter usage and direct some of streaming video's surplus back to the MSO. On the other hand, UBP may also steer some consumers with internet-only subscriptions to add television subscriptions, so their video consumption occurs through traditional broadcast TV rather

⁴See, for example, FCC's Open Internet Advisory Committee's 2013 Report https://transition.fcc.gov/cgb/oiac/Economic-Impacts.pdf.

than OTT. Our model provides a simple framework to describe which consumers are on the margins of key subscription and usage choices, and therefore the model provides guidance for investigating UBP's effects empirically.

Our empirical analysis measures a variety of consumer responses to UBP. We use novel household-level panel data from a pricing experiment implemented by an MSO that operates in multiple North American markets. In one (treated) market, the MSO shifted its pricing model to UBP partway through our sample period. Prior to this change, internet access was provided in exchange for fixed access fees. There were multiple tiers of service, with different prices associated with different connection speeds, but none of the tiers assigned prices based on usage level. After UBP was introduced, each tier's price schedule included a usage allowance and overage fees. In a second (control) market that we observe contemporaneously, the MSO held its pricing fixed throughout the sample period. Other service attributes like connection speeds were held constant in both markets. Our data include monthly subscription decisions and daily information on internet usage volume by category (e.g., Web Browsing, OTT, etc) and for several large applications within the OTT category (e.g., Netflix, Hulu, etc).

Several features of the data provide challenges and opportunities for our empirical analysis. First, the treated and control markets have non-negligible differences in the market shares of the MSO's subscription options and in the level and composition of internet usage. Second, UBP's nonlinear structure implies that consumers with different internet usage levels have different exposure to the price change, so any market-level analysis will miss the opportunity to study how households' UBP responses vary with their exposure to the policy. To overcome these challenges and provide a more precise analysis of how UBP affects households' choices and payments to MSOs, we employ the penalized synthetic control approach in Abadie and L'Hour (2021). The methodology creates a matched sample of control market households for each household in the treated market based on a specified set of pre-treatment characteristics. These control households allow us to identify which differences in post-treatment behaviors are the result of the treatment rather than differences in composition of the two markets. The technique also provides a straightforward way to measure heterogeneity in treatment intensity and response. In particular, each treated household's matched control households provide a "counterfactual" distribution of usage in the absence of UBP, which we use to calculate a household-specific measure of the price increase associated with UBP's implementation. We use the usage distribution of matched control households to calculate the counterfactual overage fees that would have been incurred by each treated household in the absence of behavioral changes. We use this dollar-valued measure of treatment intensity like a price index to examine heterogeneity in the treatment effects. On average, treated households would pay \$2.53 more per month for their internet usage if they did not alter their behavior under UBP. The distribution of price effects is heavily skewed, however, with 84% of households incurring an expected price increase of less than \$0.30, while the 95th percentile household has a price increase of \$16.39.

We find that treated households with non-negligible predicted price increases took several actions to limit UBP's impact on their monthly bills, primarily by changing their internet subscriptions and usage patterns. Subscription changes were largely concentrated in internet tier upgrades, which came with greater data allowances and speeds. Treated households in the top ten percent of the treatment intensity measure were 3.6 times more likely than matched control households to upgrade their internet tier, and households in the top two percent were 6.4 times more likely to upgrade. We interpret these upgrades as a form of metering, in which high-demand consumers are sorted into higher-priced options for internet service. By contrast, we find little impact on households' adoption of the MSO's TV service, implying that UBP did not serve as an instrument to steer consumers toward TV subscriptions.

UBP also had a meaningful impact on internet usage for some households. The average treated household's daily usage level decreased by about 0.25 Gigabytes (GBs), or 6%. This decrease was concentrated among households with larger treatment intensity, who generally

had much higher baseline usage levels. Among households with a predicted price increase of less than \$1, the average usage reduction was less than 1%. Households with an expected price increase in the top ten percent decreased their daily internet usage by 2.5 GBs, a 19% reduction. Among those households that upgraded to a higher-allowance internet tier, UBP had less effect on usage levels. Households in the top 10% of treatment intensity who upgraded their internet tier decreased their daily usage by less than 1%.

Decomposing these changes in internet usage across applications, we find reductions in usage that are generally proportional to the pre-UBP level. In particular, OTT consumption accounts for the majority of the usage and therefore the majority of the usage decline among households that did not upgrade their internet tiers. Within the online video category, we find different responses by content provider. Notably, despite Netflix's position as the mostused OTT provider by volume, reductions in Netflix usage were smaller in magnitude than YouTube and other OTT applications among households that did not upgrade their internet tier. These differences in usage effects are suggestive of differences in the value consumers place on types of content, both among third-party internet services and between OTT and conventional TV.

We conclude by documenting how UBP affected MSO revenue. Consumers with limited expected exposure to UBP, under our measure of treatment intensity, made payments to the MSO that were identical to their control groups' payments, on average. Consumers in the top 10% of estimated price exposure, on the other hand, paid \$10.90 more (8.5%) to the MSO following the price change. Realized payments increased monotonically with our measure of treatment intensity. Consumers' increased payments were largely due to overage charges for internet usage and upgrades to internet subscriptions; treated consumers' payments for TV service fell slightly relative to their matched control households.

Related literature Our analysis contributes to the literatures on Net Neutrality and the telecommunication industry's pricing practices. A central focus of the Net Neutrality debate, introduced by Wu (2003), is the interaction of upstream content creators and the distribution networks that deliver content to consumers. This literature, which is primarily theoretical and surveyed by Lee and Wu (2009) and Greenstein et al. (2016), generally deals with relationships between content and distribution firms.⁵ Our focus, by contrast, is on how content-neutral prices paid by consumers influence choices over content type, quantity, and distribution channel. We build on this analysis in a companion paper, McManus et al. (2024), where we use structural estimates of subscription demand and internet usage to study an MSO's incentives to charge premium or discounted prices for OTT usage.

Content-neutral usage-based broadband prices, like those we study, were the subject of a 2013 report by the Open Internet Advisory Committee and a growing theoretical and empirical literature.⁶ Important contributions to the theoretical literature in this space include MacKie-Mason and Varian (1994), Bauer and Wildman (2012), Odlyzko et al. (2012), and Chillemi et al. (2020). Related empirical research on usage-based pricing for residential broadband, e.g., Malone et al. (2014), Nevo et al. (2016), and Malone et al. (2021), focus on how prices affect usage volume. In addition, several related studies analyze nonlinear pricing for wireless telecommunications.⁷ Relative to these studies, our empirical analysis has the advantage of an experimental setting and includes richer subscription information and application-specific usage data that allows us to measure the impact of UBP on TV subscriptions and third-party content providers.

There are empirical studies of nonlinear pricing in many other markets. Like the demand for the Internet, electricity demand exhibits substantial heterogeneity across households and intra-day variation, which makes it a natural candidate for congestion or real-time pricing. Some recent work in this area includes Wolak (2007, 2010, 2016), Strapp et al. (2007), Ito

⁵The extensive theoretical literature includes contributions from 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), Sidak (2006). Recent empirical work in this area includes Goetz (2019), who examines how consumer welfare is affected by bargaining between between content providers and ISPs over network investment

⁶For the Open Internet report, see https://transition.fcc.gov/cgb/oiac/oiac-2013-annual-report.pdf.

⁷Lambrecht et al. (2007); Miravete (2003); Grubb (2015); Grubb and Osborne (2015).

(2014), and Anderson et al. (2017); Newsham and Bowker (2010) and Fauqui and Sergici (2010) provide a comprehensive review of the literature. Einav et al. (2015a) and Einav et al. (2015b) study health insurance contracts featuring deductibles and caps on out-ofpocket expenditures that are similar to the multi-part tariffs in our study. Performance incentives in labor contracts often also take a nonlinear form with thresholds for bonuses.⁸

More broadly, our analysis contributes to the rich literature on demand for telecommunications services. In addition to studies of demand for internet services,⁹ there is a well-developed literature on demand for TV services.¹⁰ Similar to our analysis, Malone et al. (2023) study the substitutability of the two services using data on households that drop TV service. Our contribution to this literature is a better understanding of the trade-offs households are willing to make between the services offered by MSOs, and whether usagebased pricing, as empirically implemented, is likely to have a meaningful negative impact on households and third-party content producers.

2 Model

In this section we introduce a model to describe how UBP for internet service affects consumers' choices in a setting where video entertainment is available through both OTT and traditional TV. We begin with a standard mixed bundling model in which an MSO sells standalone TV and internet subscriptions as well as in a bundle, and we augment the model in a few ways. First, we allow consumers to make usage decisions in addition to subscription choices, so we capture activity at both the intensive and extensive margins. Second, we allow consumers with internet subscriptions to access some TV video content through the internet. This increases the value that consumers receive from the MSO's internet service, which affects consumers' subscription and usage decisions. We demonstrate that OTT's

⁸Copeland and Monnet (2009); Chung et al. (2010); Misra and Nair (2011); Duflo et al. (2012).

⁹Prince and Greenstein (2017); Goolsbee and Klenow (2006); Dutz et al. (2009); Rosston et al. (2013); Greenstein and McDevitt (2011); Edell and Varaiya (2002); Varian (2002); Hitte and Tambe (2007)

¹⁰Crawford and Shum (2007); Crawford and Yurukoglu (2012); Crawford et al. (2018, 2019)

availability may create incentives for an MSO to introduce UBP or similar nonlinear prices for internet service. UBP can serve as a metering instrument, capturing some of the surplus consumers receive from OTT, and it can affect both the total amount and composition of internet data usage. In raising the price of OTT and other internet content, UBP may also steer some consumers toward TV subscriptions. We use our model to highlight the margins at which these effects occur and discuss the factors which may contribute to whether UBP's effects are primarily to meter or steer consumers' choices.

2.1 The Setup

Consider a market in which a monopolist MSO offers consumers access to two types of content. Type 1 is available on the internet only and type 2 is video entertainment available through TV.¹¹ An individual consumer's taste for "units" (e.g. hours) of content 1 and 2, relative to the outside option, is given by $v = (v_1, v_2)$. We normalize the consumer population to one and assume that consumers' tastes are distributed on $[0, 1] \times [0, 1]$.

The MSO offers subscription plans which allow the consumers to access the content. We begin by assuming that the MSO offers three plans: broadband internet access (i), TV (t), and a bundle (b) that includes both i and t. The firm's mixed bundling pricing strategy includes prices for the stand-alone plans $(p_i \text{ and } p_t)$ and a price for the bundle (p_b) . A consumer can subscribe to one of the firm's three plans, $\{i, t, b\}$, or an outside option denoted by 0 that provides utility normalized to zero. To capture the presence of OTT, we assume that consumers can access some fixed fraction, $\delta \in [0, 1]$, of type 2 content through an internet-only plan (i). We assume that OTT is available at no additional expense to the consumer. The restriction $\delta \leq 1$ has several possible interpretations, including limited available OTT content and diminished video quality, which could be due to transmission (e.g. congestion and buffering) or hardware limitations.

¹¹In this stylized model video content available only through the internet and not TV, e.g. Netflix, is part of type 1 content. In practice, different varieties of video content may have complex substitution relationships with each other. In our empirical analysis, we highlight some of these relationships.

An individual consumer receives utility from consuming q_1 units of content type 1 and q_2 units of content type 2. The quantity choice for type 2 content, q_2 , can include both traditional TV, $q_{2,t}$, and OTT, $q_{2,i}$, with $q_2 = q_{2,t}+q_{2,i}$. The consumer decides on consumption based on his tastes (v) and subscription plan. For simplicity, we assume that a consumer has marginal utility equal to one for content j up to a satiation level equal to the taste parameter v_j , and then marginal utility is zero for any greater quantity. For example, a consumer with taste v_2 and a TV-only subscription consumes $q_{2,t} = v_2$ units of video entertainment through his TV and receives surplus of v_2 from this activity. We integrate OTT into this framework by assuming that when the consumer uses OTT, his marginal utility from video hours remains equal to one up to δv_2 , where it falls to zero. This can be viewed as a scenario where a consumer enjoys v_2 distinct shows available on TV, but only the fraction δ of the shows are available through OTT. To simplify our descriptions of bundle subscribers' consumption choices, we assume that bundle subscribers receive all of type 2 content through TV, with $q_{2,t} = v_2$ and $q_{2,i} = 0$.

Putting this all together, subscribers in internet-only plans receive utility of $U_i = v_1 + \delta v_2 - p_i$, where the first and second terms capture utility (and quantities) from consuming internet content and OTT applications, respectively. A subscription to the TV service, t, results in TV content consumption of $q_{t,2} = v_2$, zero internet usage given the lack of access, and net utility equal to $U_t = v_2 - p_t$. Bundle subscribers consume quantities of content types 1 and 2 up to their satiation levels and receive utility equal to $U_b = v_1 + v_2 - p_b$. Finally, if the consumer selects the outside option, 0, quantities are zero for both content types and utility is zero.

We abstract away from specifying the MSO's cost structure and profit function. While the firm's costs are an important part of how it sets prices with and without OTT, along with its incentive to introduce UBP, in this paper focus on tracking consumers' responses to UBP. For related discussion of the MSO's incentives see McManus et al. (2024).

Figure 1: Consumer choices in simplified model



Notes: The Figure shows consumer subscription choices for $\delta = 0$ and $\delta = 0.7$, holding prices fixed. At $\delta = 0$, consumers with v values in i choose internet-only subscriptions, those in t choose TV subscriptions, and those in b and II select the bundle. Consumers with v values in regions 0 and I do not purchase from the MSO.

2.2 Consumer Choice

We now turn to the choices consumers make in this setup. In Panel (a) of Figure 1 we present choices different consumers make for a fixed set of prices. When no OTT is available (i.e., $\delta = 0$), consumers in the areas labeled '0' and 'I' select the outside option, and those in areas 'b' and 'II' select the bundle. Consumers in areas *i* and *t* select the stand-alone internet and TV subscription plans. The split is intuitive: consumers with high valuations for both content types choose the bundle, and consumers with high valuation for one content type and not the other choose the plan with just the appropriate stand-alone subscription. Consumers with low valuations of both content types choose the outside option. The locations of the margins between the regions depend on the prices of the various options.

We next consider the effect of OTT becoming more attractive (i.e., the effect of an increase

in δ) on subscription choices. In Figure 1 we show the effect of δ increasing to 0.7, holding prices fixed at their original levels. Two types of consumers change their choices. First, some consumers (in area I) who did not purchase, despite moderately high valuation for content 1 or 2, will subscribe to *i* because it became more attractive. These new consumers increase the MSO's revenue and are one reason the MSO has an incentive to encourage OTT. Second, some consumers (in area II) decide to "cut the cord." These consumers choose a bundle when $\delta = 0$ but have relatively low tastes for TV content among bundle subscribers. As δ increases, they prefer stand-alone internet service because they can consume OTT using the internet service. The cord-cutting by these consumers diminishes the MSO's revenue, as the bundle price is higher than the internet-only price. Whether the consumers in area II reduce the firm's profit, however, depends on several factors which may work in opposite directions. The relative costs of providing TV and internet affect the bundle and internet-only profit margins, and consumers in area II may be moving to a higher- or lower-margin service.

2.3 UBP's Impact

We illustrate the impact of introducing UBP with a stylized tiered internet service in which consumers must pay a premium for greater usage. As in the UBP policy we observe empirically, consumers have the option to upgrade to a higher internet tier if they desire more content than is permitted in their initial tier's allowance.¹² We illustrate this strategy with a simple menu of two internet plans, with the low-usage plan (i_L) available for price $p_{i,L}$ and usage cap κ , plus a high-usage plan (i_H) with unlimited usage. The usage cap and tiers serve two purposes in an internet subscription. First, the high-usage tier extracts a premium from high-demand individuals who are willing to pay a premium $(p_{i,H} - p_{i,L})$ for extra usage $((v_1 + \delta v_2) - \kappa_L)$. Second, the usage cap prevents additional internet usage by inframarginal consumers on the low tier whose tastes would lead them to consume in excess of the cap; this may be valuable to the MSO if internet costs increase with usage. To the extent that

¹²Unlike the empirical UBP in our data, the stylized example does not allow the consumer to pay an overage charge for extra internet usage while remaining on his initial tier.

the cap and tiers limit OTT usage in some cases while charging a premium for it in others, they act as metering instruments for the MSO. In a setting with $\delta > 0$, the caps and tiers reduce consumers' net benefit from video entertainment over the internet, and therefore may steer consumers from i to b.

To illustrate the impact of tiers on consumers' choices, we consider a case in which the MSO introduces tiers to internet service while keeping other prices fixed. This scenario resembles the situation we see in our data, in which the MSO introduced tier allowances and overage charges without adjusting other subscription prices. When the MSO introduces caps and tiers to a setting which had neither, several types of consumers are affected in distinct ways. We illustrate these changes in Figure 2. We impose a set of prices that facilitate reading the different regions of Figure 2, including setting $p_{i,L}$ equal to the initial p_i .¹³ Panel (a) of Figure 2 provides an initial distribution of consumers across subscription options with with $\delta > 0$, before any tiers or cap are offered.

The introduction of usage tiers splits i and b subscribers into a new collection of actions, illustrated in Figure 2 Panel (b). Former subscribers to the initial unlimited internet service (i) may update their subscription and usage choices in several ways. Some consumers, in Figure 2 Panel (b)'s area I, accept the usage cap κ and remain internet-only with a low allowance (i_L) . This tier's usage cap will cause some i_L subscribers to reduce their internet usage relative to their pre-UBP levels. Consumers with a stronger taste for internet usage, whether for video- or non-video content, may "upgrade" their internet subscription to i_H ; these consumers are in Panel (b)'s area II. This is the central metering channel in our simple model. From the MSO's perspective, the tier upgrade is a way to collect a greater price from consumers who receive internet service that is equivalent to their pre-UBP outcome. Consumers with relatively strong values of v_2 switch from i into the bundle (areas III and IV). Of these consumers, those with high values of v_1 pay for a tier upgrade in addition to TV

¹³In panel (a), we plot market shares for prices $(p_i, p_t, p_b) = (0.75, 0.65, 0.9)$. In panel (b), the MSO places a usage allowance $\kappa = 0.8$ on the original internet tier, now i_L , and a new premium internet tier, i_H is introduced with no allowance. The new prices are $(p_{i,L}, p_{i,H}, p_t, p_{b,L}, p_{b,H}) = (0.75, 0.85, 0.65, 0.9, 1.0)$.

service (area IV). When consumers switch to the bundle, they receive video entertainment through TV, so their OTT usage falls. In our empirical setting, the flow of consumers into III and IV, which represents the steering effects of UBP, depends on the substitutability of OTT and traditional TV, as well as the density of consumers in these regions. Finally among initial *i* subscribers, some cancel their subscriptions completely (area V) because capped internet, at the present price, is worth less than the outside option. In addition to these margins for former *i* subscribers, some bundle subscribers with strong internet tastes (in area VI) opt for b_H so that they can consume internet without a usage limit.



Figure 2: Effect of Tiers and Allowances

Notes: This figure shows the effect of the introduction of internet tiers on subscription choices. Throughout, δ is fixed at 0.7. In panel (a), the MSO offers a single internet tier with unlimited usage. In panel (b), consumers must select a high-usage internet-only or bundle tier for internet usage greater than κ .

3 Data and Descriptive Analysis

In this section, we first describe our data sources and discuss the implementation of the MSO's UBP experiment. We then present detailed summary statistics on households' sub-

scription choices and internet usage. Finally, we provide a descriptive analysis of how household behavior changed following the introduction of UBP.

3.1 Data Sources and the Usage-Based Pricing Experiment

Our data come from a North American MSO; our data use agreement with the MSO prevents us from identifying the firm or any details that could be used to infer its identity.¹⁴ We observe nine months of billing information, subscriptions, and application-specific internet usage data for 70,500 households in two large markets. The data were collected during the latter half of the 2010s. The MSO that provided this sample, like most other MSOs during this period, offered a menu of internet tiers and TV service. These services could be purchased as standalone subscriptions or as bundles including both an internet tier and TV service.¹⁵ Subscription contracts and billing periods were about one month long, so we use the terms "month" and "billing period" interchangeably below. The internet tiers are differentiated by speed, while TV service was available in several tiered channel packages (basic, premium, etc.). Adding a TV plan to an existing internet subscription increased a household's bill by about \$100, on average. For each household, we observe the internet tier chosen, the presence or absence of a TV-service subscription (but not consumption or the set of available channels), and monthly payments to the MSO. We also observe internet usage in several distinct categories, including real-time entertainment (RTE), web browsing, gaming, and peer-to-peer traffic. The RTE category contains online/streaming video, and within the RTE category we can identify usage of some major applications like Netflix.¹⁶

One important feature of our data is that the MSO introduced UBP in one of the two

¹⁴To maintain the MSO's anonymity, we cannot provide details on the specific markets served, the exact dates and details of the implementation of UBP, and the detailed characteristics of the MSO's product menus.

¹⁵The MSO also offers telephone service, which about 40% of its customers subscribe to. We do not use the telephone service information in this paper.

¹⁶Information on internet usage comes from two sources: internet protocol data records (IPDR) and a deep packet inspection (DPI) platform. IPDR is considered the most reliable source for high-frequency customerspecific upload and download byte counts and is used by MSOs for usage-based billing purposes. The DPI platform (e.g., Sandvine) provides detailed information on the composition of bytes used by a household. In these DPI data, we observe household-specific byte counts at an hourly frequency for each traffic category.

markets during our sample period. Prior to introducing UBP in the "treated" market, the MSO offered identical product menus in the treated and control markets, with the control market's menu remaining unchanged throughout the sample period. Under the original product menu, internet tiers were differentiated by price and connection speed, with faster connection speeds associated with higher prices. Under UBP, each internet tier received a monthly usage allowance, with more expensive tiers associated with higher allowances, while the baseline access fees and connection speeds associated with each tier remained the same. Internet usage in excess of the UBP allowance triggered the automatic purchase of a "top-up" quantity of additional data. The top-up quantity was smaller than the allowance difference between any two adjacent tiers, and its price was approximately equal to the subscription price difference between adjacent tiers. Households could purchase an unlimited number of top-ups.

The MSO introduced UBP in two phases. First, during an "announcement period" that began several months prior to UBP implementation, the MSO publicized the details of the new policy to its customers in the treated market. The MSO provided the UBP starting date, the menu of tier usage allowances, and the price and quantity of an allowance topup. During each billing period in the announcement period, the MSO also informed treated subscribers about how their monthly usage compared to the data allowance that would be associated with their current tier under UBP. In the second phase, which we call the "treatment period," the MSO enforced its UBP policy, including assessing overage charges on households with internet usage greater than their monthly allowance. In all, our sample period includes multiple months of the "pre-policy period" prior to UBP's announcement, the full announcement period, and several months of the treatment period. To our knowledge, the MSO selected the treated market for UBP based on its technological capability to charge subscribers for usage there. In both the treated and control markets, the MSO's competitors did not change their subscription menus meaningfully during the sample period, including in response to UBP's introduction in the treated market. Satellite TV was available in each market, as was internet service via DSL at substantially slower speeds than the MSO's broadband service.

3.2 Descriptive Statistics

In Table 1, we describe household-level monthly internet usage and plan choices in each market during the pre-policy period. Table 1's left side contains statistics on internet-only households who had no TV subscription, and the right side describes households with an internet-TV bundle. The distributions of internet usage are displayed in Panel 1. Among internet-only subscribers, the treated and control distributions are quite similar. For bundle subscribers, internet usage is greater in the control market. These differences motivate our empirical approach, which constructs household-specific control groups to aid in measuring treatment effects. In both the treated and control markets, internet-only households use more internet data than bundle subscribers.¹⁷

Table 1's Panel 2 provides a breakdown of subscription choices. About 28% of the MSO's customers have internet-only subscriptions, and the remaining 72% subscribe to an internet-TV bundle; no MSO customers in our sample subscribe to TV alone. We aggregate the internet tiers into three categories by speed: low, medium and high. There are some minor differences across subscription choices, but in general the medium-speed tiers are the most popular, followed by the low-speed tiers.

Panels 3 and 4 of Table 1 provide a breakdown of internet usage by category, including several major streaming video applications. More than half of the internet usage we observe is OTT consumption. Netflix, which offers a variety of original programming along with a library of previously distributed movies and television programs, is the most-used subscription service. Engagement with YouTube generates the second-largest level of network usage. Other observable applications include Hulu, which emphasizes opportunities to stream-on-

 $^{^{17}}$ Despite the rapid usage growth within our sample, we see little evidence that congestion affected internet use. 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. (2021) for a study of the impact of congestion on broadband networks.

Household Type	Has T	V: NO	Has TV: YES				
Household Type	Treated	Control	Treated	Control			
Panel 1: GBs Total Monthly U	Jsage						
Mean	154.47	156.48	73.62	102.94			
Standard Deviation	169.49	170.57	114.27	122.40			
1st Percentile	1.49	1.90	0.21	0.34			
25th Percentile	43.31	60.10	9.08	22.57			
Median	105.50	117.76	30.66	65.65			
75th Percentile	211.27	208.62	93.38	139.79			
99th Percentile	732.85	655.40	515.91	560.41			
Panel 2: Subscription Choices							
Households (N)	7,330	$13,\!374$	$23,\!439$	$28,\!278$			
Total Bill (\$)	74.92	78.56	175.46	181.47			
Low Speed Internet $(\%)$	0.32	0.23	0.25	0.12			
Median Speed Internet (%)	0.53	0.56	0.66	0.63			
High Speed Internet (%)	0.15	0.21	0.10	0.25			
Panel 3: Mean Usage by Cates	gory						
Video	100.14	98.11	43.19	61.58			
Browsing	39.43	42.30	22.05	29.79			
Other	14.91	16.06	8.38	11.57			
Panel 4: Mean Usage of OTT Services							
Netflix	54.82	40.27	22.59	27.83			
YouTube	20.64	29.64	12.25	18.28			
Hulu	4.38	3.08	0.86	1.11			
Sling TV	0.66	0.40	0.04	0.06			

Table 1: Usage and Plan Choice during the Pre-Policy Period

Notes: Summary statistics at a household-month level of observation. Usage levels are totals within monthly billing cycles. Tier choice shares are the fraction of households that choose each speed tier group.

demand TV shows currently airing on network TV, and Sling TV, which offers live TV over the internet. Internet-only households use each of these applications more intensely than bundle households.

3.3 Descriptive Analysis of Responses to UBP

UBP's price effects, if relevant to households' choices, would create differences in subscription patterns between the treated and control markets. In Figure 3, we report how the



Figure 3: UBP Response: TV Adds and Drops



(b) Drops

Notes: Panels (a) and (b) of this figure plot the propensity to add and drop video, respectively, in the treatment and control market. Only households eligible to take each action are considered (e.g., only households who begin the sample without a TV subscription may add TV).

propensities to add and drop video service (panels (a) and (b), respectively), changed in the treatment and control markets during our panel.¹⁸ Panel (a) shows that, prior to UBP's introduction, the flows of households from internet-only to bundle subscriptions were virtually identical in the treated and control markets. After UBP was introduced, treated households were substantially more likely to add TV subscriptions. While this difference in new TV subscriptions is consistent with arguments about UBP's potential steering effects, in panel (b) we show that the treated market maintained a slightly greater cord-cutting rate than the control market throughout the sample period. Combining the effects in the two panels, the share of households with TV subscriptions fell by more in the treated market than the control, as the difference in cord-cutting frequency applies to a larger starting population (initial bundle subscribers) than the add-TV frequency.

In Figure 4 we display the cumulative shares of households upgrading or downgrading

¹⁸To construct Figure 3 and similar figures below, we use all households in the sample, i.e., without matching on households' pre-policy characteristics as in our main empirical results.



Figure 4: UBP Response: Tier Upgrades and Downgrades



(b) Downgrades

Notes: Panels (a) and (b) of this figure plot the propensity to upgrade and downgrade tiers, respectively, in the treatment and control market. Only households eligible to take each action are considered (e.g., households who begin the sample on the highest internet tier may not upgrade their tier).

their internet tiers. In panel (a) we show that the upgrade rates were identical during the pre-policy months of the sample period, and then they diverged significantly. Treated households upgraded their internet tiers more frequently than control households starting during the announcement period, and this continued after UBP's full introduction. In panel (b) we show that the treated market also had a greater rate of internet downgrades than the control market, but the difference is relatively small in magnitude. As we discuss below, this may be due to UBP informing some low-usage households that their consumption could be accommodated in a lower tier.

In addition to UBP's impact on subscriptions, we investigate how its prices affected internet usage quantities. In Table 2, we provide an initial summary of the difference-indifferences impact of UBP on usage. Internet usage grew in both markets during the sample period, but growth was substantially slower under UBP. In the control market, average GBs per household per month grew by 36% (from 119.68GB to 162.09GB) during the sample period, but growth in the treated market was only 20% (from 91.77GB to 110.21GB). Just under half of the difference between the two markets is due to slower growth in video usage, and a similar share is due to reduced growth in browsing, which itself includes embedded video. The major video streaming services we highlight in Table 2 account for considerably more than half of all pre-UBP video usage, but, in relative terms, video usage reductions occurred primarily outside of these services in the treated market.

	Treated: YES		Treated	Diff_in_Diff	
	Pre-UBP	UBP	Pre-UBP	UBP	
Video	55.98	64.17	73.01	92.53	-11.33
Netflix	29.83	34.97	31.72	40.60	-3.74
YouTube	14.13	16.46	21.83	25.39	-1.24
Hulu	1.65	2.70	1.73	2.73	0.04
Sling TV	0.18	0.70	0.17	1.19	-0.51
Browsing	25.95	32.48	33.71	51.32	-11.09
Other	9.85	13.56	12.97	18.24	-1.55
Total	91.77	110.21	119.68	162.09	-23.97

Table 2: UBP Response: Usage

Notes: Average monthly internet usage by category, treatment period, and treatment status.

To provide some preliminary context on UBP's incidence in the treated markets, we describe briefly how pre-announcement internet usage levels compared to the allowances that were eventually introduced to a household's tier. In the most popular internet tier, 5% of households had average pre-UBP monthly usage that exceeded the tier's allowance under UBP. Average usage in excess of the allowance was more common in lower tiers than higher tiers. Across all household-month pre-UBP observations in the treated market, 3.9% would have generated an overage charge, had UBP been in place. Overall, 5.2% of all treated households would have received at least one bill with a positive overage charge. Due to substantial growth in internet usage during our sample period, UBP would affect more households during the treatment period than in the pre-announcement months.

4 Measuring the Impact of Usage-Based Pricing

In this section, we describe our empirical approach, which adapts the penalized synthetic control (PSC) framework of Abadie and L'Hour (2021). We construct a household-specific measure of treatment intensity that captures the cost of inaction in response to UBP, which we relate to treatment effects which capture the impact of UBP on household usage and subscription decisions.

4.1 Empirical Strategy

Our data exhibit both market-level and intertemporal variation in exposure to treatment. The PSC estimator provides a powerful tool to identify household-specific measures of treatment intensities and effects while controlling for observable differences between treated and control households. Specifically, the PSC estimator yields two useful objects for each household: weights for a matched sample of control households and an estimate of the treatment effect for each outcome of interest, in our case UBP's impact on usage and subscriptions. In a treated household's matched sample, the control households' usage distribution provides a measure of treatment intensity in the absence of behavioral changes brought on by treatment. The treatment intensity is the expected price change from UBP given the synthetic control's usage distribution in the post-treatment period. This measure captures the relative cost of inaction for treated households in response to UBP. We relate this treatment intensity measure to the treatment effect estimates to characterize the heterogeneity in households' subscription and usage responses to UBP.

The PSC approach has particular advantages for constructing the measure of treatment intensity. In particular, it explicitly accounts for the trade-off between (a) minimizing behavioral discrepancies between a given treated household and its synthetic control, and (b) minimizing discrepancies between a given treated household and each of the individual control households included in its synthetic control. Favoring option (a) will lead to the best overall fit between treated households and their corresponding synthetic control, but accounting for option (b) can reduce bias by ensuring that each treated household's matched sample includes only control households with similar observable characteristics. Given that our measure of treatment intensity relies on the synthetic control's distribution of usage in the post-treatment period, it is important that each matched control household have similar usage patterns to the treated household in the pre-treatment period. The large number of control households and disaggregated nature of our data is ideal for this purpose, because it allows us to account for (b) with little to no impact on (a).

Formally, the PSC framework solves the following quadratic program:

$$\min_{W_i \in \mathbb{R}^{n_c}} \|X_i - \sum_{j=1}^{n_c} W_{i,j} X_j \|^2 + \lambda \sum_{j=1}^{n_c} W_{i,j} \|X_i - X_j \|^2 \tag{1}$$
s.t. $W_{i,j} \ge 0 \quad \forall j = 1, \dots, n_c$

$$\sum_{j=1}^{n_c} W_{i,j} = 1.$$

The solution, $W_i^*(\lambda)$, is a set of weights that describe the matched sample among the $j = \{1, \ldots, n_c\}$ control households for each treated household $i = \{1, \ldots, n_t\}$. The tuning parameter, λ , balances the trade-off described above between considerations (a) and (b), described above. Larger values of λ correspond to more weight on ensuring that the characteristics of treated household $i(X_i)$ are similar to the characteristics of each matched control household (X_j) . We include eleven covariates that describe the level, variance, and composition of internet usage during the M_0 months in the pre-UBP period, as well as plan choice. Specifically, X includes: total internet usage in each of three pre-UBP months, variance in daily usage during the pre-UBP months, share of total pre-UBP usage in each of four categories (online video, web browsing, Netflix, and Youtube), whether the household subscribed to TV service, and which internet tier they chose at the start of the pre-UBP period. We follow the leave-one-out cross-validation of post-intervention outcomes for the untreated approach of Abadie and L'Hour (2021) to identify the λ penalty values that minimize mean squared prediction error and bias. We find a value of $\lambda = 0.1$ is optimal, so this is our focus in the discussion of our results presented in Section 5. We describe results for alternative values of λ in the Appendix; results generated using alternative λ parameters are qualitatively similar to the case $\lambda = 0.1$.

4.2 Measuring Treatment Intensity and Effect

We use the solution to Equation 1 to estimate household-specific measures of treatment intensity for each treated household, which we relate to treatment effect estimates for a variety of household behaviors.

To construct the measure of treatment intensity, we first calculate a counterfactual cumulative distribution function describing usage in the absence of treatment. For each treated household $(i = 1, ..., n_t)$, we calculate

$$\hat{F}_{i}(z;\lambda) = \frac{1}{M_{1}n_{c}} \sum_{j=1}^{n_{c}} \sum_{m=1}^{M_{1}} \mathbb{1}\left[c_{jm} < z\right] W_{i,j}^{*}(\lambda),$$
(2)

where c_{jm} refers to the realized monthly usage of control household j during month m of the treatment period ($m = 1, ..., M_1$). The expected monthly price increase from UBP in the absence of behavioral changes for household i, our measure of treatment intensity, is then

$$\hat{\rho}_i(\lambda) = \int_0^\infty \mathcal{O}(z) d\hat{F}_i(z;\lambda) \tag{3}$$

where $\mathcal{O}(z)$ are the overages associated with z GBs of usage on household *i*'s internet tier. This measure is like a price index because it places (probability) weights on the prices associated with different quantities, conditional on choices remaining fixed.

We calculate estimates of household-specific treatment effects for an outcome of interest

 (y_i) for household *i* during the months of the post-treatment period $(m = 1, \ldots, M_1)$ as

$$\hat{\tau}_i(\lambda) = \frac{1}{M_1} \sum_{m=1}^{M_1} \sum_{j=1}^{n_c} \left(y_{im} - y_{jm} W_{i,j}^*(\lambda) \right).$$
(4)

This measure captures the average deviation in behavior of household i from its synthetic control in the post-treatment period.

In Section 5, we examine the relationship between the household-specific measures of treatment intensity and effect, $\hat{\rho}_i(\lambda)$ and $\hat{\tau}_i(\lambda)$, respectively. This offers a systematic way of providing insight into the heterogeneous actions taken by households in response to the market-level introduction of UBP. To calculate standard errors for these measures and the relationships between them, we use block re-sampling with 200 draws. In particular, we sample with replacement n_t treated households (the entire block of a household's data), and similarly n_c control households.¹⁹

5 Results

In this section, we describe our estimates of treatment intensity from UBP $(\hat{\rho}_i(\lambda))$, and then we relate these measures to changes in treated households' subscription and usage choices under UBP $(\hat{\tau}_i(\lambda))$. The relative magnitudes of households' subscription and usage responses allows us to break down UBP's potential steering and metering effects.

5.1 UBP Treatment Intensity

The average expected UBP overage charge, calculated with Equation (3), is \$2.53. The distribution of $\hat{\rho}_i(\lambda)$ is highly skewed, with 84% of treated households with expected charges of less than \$0.30 if they remain on their pre-UBP subscriptions and usage trends. At

¹⁹We use the entire population of treated households in this process, but each re-sampling of size n_c represents about 20% of the number of control households. This has no measurable impact on the estimates, but makes the PSC calculations computationally tractable.

the 95th percentile, the expected overage charge is \$16.31. We display the distribution of expected overages, conditional on the charge being greater than \$0 (16.7% of households), in Figure 5. While most households can meaningfully reduce overages by upgrading, the top few percentiles would incur overages even after upgrading if usage does not decrease.²⁰



Figure 5: Distribution of Expected Overages

Notes: Histogram of household-level estimated expected overages, $\hat{\rho}_i(\lambda)$. The 83.3% of households with zero expected overage are not included in the figure.

Our measure of treatment intensity is not a price that is actually paid by a household, but rather it is the additional cost associated with maintaining current behaviors in the post-UBP period as captured by the household's synthetic control. For example, two high-usage households may have substantially different expected overage charges, but each could avoid its charge by paying the same fee to upgrade to the next-highest internet tier or decreasing usage. For our purposes, it is worthwhile to describe the two households as ex-ante different (in expected overage, conditional on no upgrades or changes to usage behavior) rather as equivalent through realized payments under UBP. With this in mind, we describe expected

²⁰If we calculate $\hat{\rho}_i(\lambda)$ using the allowance on the next highest tier instead of on the chosen tier, the 95th percentile is \$0, the 96th percentile is \$0.64, the 97th percentile is \$3.26, the 98th percentile is \$8.13, and the 99th percentile is \$18.54.

overage charges in bins defined by overage percentiles rather than dollar values alone. We group together the 84% of households with overages equal to \$0.30 or below, and then we separate the remaining households into 8 bins with equal numbers of treated households in each. For example, 2% of treated households have an expected overage between \$0.30 and \$1.30, and the next-greater 2% have expected overages between \$1.30 and \$2.80.

5.2 UBP's Impact on Subscriptions

A prominent concern about UBP is that it steers consumers toward bundles that include TV subscriptions, which might be more profitable for the MSO than internet-only subscriptions. In particular, the TV subscription may allow the MSO to capture consumers' surplus from content delivered through their TV service, while similar OTT content generates surplus that is captured, in part, by third-party OTT providers. If UBP is successful at steering consumers toward TV subscriptions, it would come at the cost of the otherwise-preferred OTT delivery and potentially harm third-party OTT providers.

In Figure 6 panel (a) we describe UBP's impact on TV service addition choices, separated by percentiles of UBP incidence. Consumers in the lowest expected overage bin (\$0.30 or below), who have virtually no price exposure to UBP, have a statistically significant but small positive response in TV subscriptions. The nonzero effect may be due to increased salience of prices for internet usage, or households perceiving greater risk of encountering future UBP charges than we are able to capture with our approach to expected overages. Treated households with greater expected UBP charges have no significant difference from matched control households. The point estimates of UBP effects on TV additions are similar in magnitude for all expected overage levels, but the smaller sample sizes in the bins with nontrivial overages generates fairly imprecise estimates. The average effect across all treated households is a 1.1 percentage point (67%) increase in the probability of adding TV.

In Figure 6 panel (b), we display decisions to drop TV subscriptions across the percentiles of $\hat{\rho}_i(\lambda)$. Among households with near-zero UBP exposure, treated households are



Figure 6: TV Subscription Changes by Expected Overage Level



(b) Cancelled Subscriptions

Notes: Heterogeneity in the effect of UBP on the take-up of TV subscriptions. Households that began the sample without a TV subscription are in panel (a); households with a TV subscription are in panel (b). Households are grouped into bins using percentiles of the $\hat{\rho}_i(\lambda)$ distribution. All bins except for the left-most bin contain 2% of households; the left-most bin contains all remaining households. Confidence intervals are calculated using block re-sampling with 200 permutations.

significantly more likely to drop their TV subscriptions than control households, but this difference is quantitatively small. Per household, this effect is similar in magnitude to the corresponding difference in TV additions among internet-only treated households with nearly zero price exposure. For treated households with greater expected overage charges, there is generally no difference from the control market. Together, the small but positive estimated effects for additions and drops result in no overall impact on TV subscriptions for the MSO. This is inconsistent with concerns that commonly-implemented forms of UBP in the US have significant impacts as steering mechanisms.

While UBP had only small effects on TV subscription choices, it had an economically meaningful impact on internet tier subscription decisions. In Figure 7, we display the changes in the propensity to upgrade and downgrade tier for different levels of UBP price exposure. Panel (a) shows that UBP had a positive but very small impact on upgrading decisions for treated households with near-zero price exposure, and as exposure increased the propensity



Figure 7: Internet Tier Changes by Expected Overage Level



(b) Tier Downgrades

Notes: Heterogeneity in the effect of UBP on the internet tier choice decision. Households that began the sample on the highest internet tier are omitted from panel (a); households on the lowest internet tier are omitted from panel (b). Households are grouped into bins using percentiles of the $\hat{\rho}_i(\lambda)$ distribution. All bins except for the left-most bin contain 2% of households; the left-most bin contains all remaining households. Confidence intervals are calculated using block re-sampling with 200 permutations.

to upgrade did as well. Where expected overage charges were nontrivial, treated households responded meaningfully. Treated households in the 84^{th} to 90^{th} percentiles had a 7.7 percentage point (117%) increase in the probability of upgrading tiers. In the higher percentiles of expected UBP exposure, the effects were substantially larger. Households in the 92^{nd} percentile had a 13.6 percentage point (207%) increase in probability of upgrading, while those in the top 2% were 39.4 percentage points (642%) more likely to upgrade. These outcomes are evidence that UBP acted as a form of metering, sorting higher-demand households into higher-priced tiers. We also examine tier-downgrade decisions in response to UBP. In Figure 7 panel (b) we show that there was generally no economically or statistically significant response to UBP on this dimension.



Figure 8: Internet Usage by Expected Overage Level

Notes: Heterogeneity in the effect of UBP on overall daily internet usage. Households are grouped into bins using percentiles of the $\hat{\rho}_i(\lambda)$ distribution. All bins except for the left-most bin contain 2% of households; the left-most bin contains all remaining households. Confidence intervals are calculated using block re-sampling with 200 permutations.

5.3 UBP's Impact on Internet Usage

In addition to impacts on subscription choices, UBP may affect households' internet usage in several ways. If a household with relatively high usage stays with its initial internet tier under UBP, it may decrease usage so it does not incur overage charges. This would imply some (perhaps modest) utility loss for the household, and a traffic reduction for third-party firms. If a similar household upgrades its tier to maintain its existing usage, this represents a transfer from the household to the MSO, but the household and third-parties may benefit from a greater speed. We display the overall effect of the price change on usage levels in Figure 8, then decompose the overall effect in the two panels of Figure 9, which displays usage changes separately for households that did not (panel a) or did (panel b) upgrade their internet tier.

The average treatment effect across all households is a 0.24 GB reduction in daily usage



Figure 9: Internet Usage by Expected Overage Level and Upgrade Decision

(a) No Tier Upgrade

(b) Tier Upgrade

Notes: Heterogeneity in the effect of UBP on overall daily internet usage. Households that did not upgrade to a higher internet tier are described in panel (a); households that upgraded their internet tier are described in panel (b). Households are grouped into bins using percentiles of the $\hat{\rho}_i(\lambda)$ distribution. All bins except for the left-most bin contain 2% of households; the left-most bin contains all remaining households. Confidence intervals are calculated using block re-sampling with 200 permutations.

(a 6% reduction from the baseline). For households that did not upgrade their internet tier, responses to UBP are zero or significantly negative, as would be expected. For example, households in the 90^{th} to 92^{nd} overage percentile, with expected overages between \$4.90 and \$7.90, reduced their internet usage by 1.37 GB (16%) if they did not upgrade. Households with expected overages at or above the 90th percentile reduced their internet usage by 3.47 GBs (26.8%) on average if they did not upgrade.

In panel (b), we show the responses of households that move to a higher-allowance tier. These households valued continuing their internet usage above the incremental price of increasing their tier, and most households display a significant increase in total GB used. A prominent mechanism for usage increases is the greater speed of a higher tier, which typically generates an automated response from bandwidth-adaptive applications. This can increase usage in GBs even if the household spends no additional time using the internet. In addition, the increased bit rate may be valued by households because of increased video resolution or reduced download times. For households with the greatest price exposure, we find a usage reduction for households that upgrade. These households have predicted usage so great that, without further reduction, they risk exceeding the usage allowance of their new internet tier.

We decompose these usage changes by application in Figures 10 and 11. In Figure 10, we display usage changes, again separated by tier upgrading choices, for video, browsing, and other internet usage. In panel (a), we show that video usage reductions accounted for most of the usage declines among households that did not upgrade their internet tiers. These impacts of UBP, measured as daily GB changes, reflect both the pre-policy usage levels and magnitudes of reduction. In percentage terms, the reductions in usage types are quite similar to each other. The reduction in video usage is consistent with some concerns about how an MSO may use UBP to steer consumers toward its TV service, but our results in Figure 6 suggest that consumers were not responsive on the extensive margin of TV subscriptions.²¹ Our results in Figure 10 panel (b) show that the usage increases in Figure 9 panel (b) were largely concentrated in video applications. This is consistent with a mechanism for increased usage we described above, as video applications have both greater benefits from increased bit rates (through higher-definition video) and are likely to be bandwidth-adaptive.

In Figure 11, we provide a further decomposition of usage changes. We separate video usage into four categories: Netflix, YouTube, Hulu or SlingTV (combined), and all other video. Despite Netflix's position as the largest category of video usage (see Table 1), its usage reductions are smaller in magnitude than YouTube or other video among households that did not upgrade their internet tier. This suggests that Netflix has a relative high value among video categories: when treated households perceived a need to reduce video usage, they chose to focus their reductions on video categories other than Netflix. Changes to Hulu and SlingTV, which are near zero in Figure 11 panel (a), are partially due to their small average usage in the population, but households' utility for the services may play a role, as in the case

²¹We do not observe households' TV viewing activity. Households substitute on the intensive margin from streaming video to broadcast TV, this could raise the value of MSOs' TV offerings. Households may also have chosen to reduce bit rates for video applications to reduce usage. We defer this issue to future research.



Figure 10: Composition of Internet Usage by Expected Overage Level and Upgrade Decision

(a) No Tier Upgrade

(b) Tier Upgrade

Notes: Heterogeneity in the effect of UBP on daily internet usage by category. The 3 categories are mutually exclusive and collectively exhaustive. Households that did not upgrade to a higher internet tier are described in panel (a); households that upgraded their internet tier are described in panel (b). Households are grouped into bins using percentiles of the $\hat{\rho}_i(\lambda)$ distribution. All bins except for the left-most bin contain 2% of households; the left-most bin contains all remaining households. Confidence intervals are calculated using block re-sampling with 200 permutations.

of Netflix. In Figure 11 panel (b), we display usages changes by video category for households that upgraded their internet tiers. Netflix usage increased significantly for households in all expected-overage bins, including those with near-zero price exposure. Netflix's applications automatically adapt their data usage to provisioned speeds, so this channel is one likely explanation for increased Netflix usage. The reduction in YouTube and the other video categories, which is statistically significant in some of the greater expected-overage bins, and could be due partially to substitution toward higher quality (resolution) Netflix video.

5.4 UBP's Impact on Payments to the MSO

Several impacts of UBP – subscription changes and overage charges – affect households' payments to the MSO. We conclude our analysis by describing the distribution and magnitudes of changes to subscribers' payments due to UBP.



Figure 11: Online Video Usage by Expected Overage Level and Upgrade Decision

(a) No Tier Upgrade

(b) Tier Upgrade

Notes: Heterogeneity in the effect of UBP on daily online video usage. The 4 categories are mutually exclusive and collectively exhaustive. Households that did not upgrade to a higher internet tier are described in panel (a); households that upgraded their internet tier are described in panel (b). Households are grouped into bins using percentiles of the $\hat{\rho}_i(\lambda)$ distribution. All bins except for the left-most bin contain 2% of households; the left-most bin contains all remaining households. Confidence intervals are calculated using block re-sampling with 200 permutations.

In Figure 12, we show how treated households with varying exposure to UBP changed their payments to the MSO, relative to control households. We find that the 84% of households with near-zero expected overages had no change in their payments to the MSO. For greater levels of $\hat{\rho}_i(\lambda)$, additional payments to the MSO increase monotonically. Households in the 85th through 92nd percentiles make additional payments of \$2 to \$5, while those in the next two percentile bins pay \$7 to \$8 more per month, relative to control households. Households with the greatest expected overages make payments to the MSO greater than those in lower bins, but actual payments are much less than our estimated cost without behavioral changes, i.e., $\hat{\rho}_i(\lambda)$. This is because treated households' actual payments allow for re-optimization, which can include upgrading the internet tier or reducing usage, while the expected overage values assume that behavior continues according to their matched sample of control households.



Figure 12: Change in Payments by Expected Overage Level

Notes: Heterogeneity in the effect of UBP on monthly bill level. Households are grouped into bins using percentiles of the $\hat{\rho}_i(\lambda)$ distribution. All bins except for the left-most bin contain 2% of households; the left-most bin contains all remaining households. Confidence intervals are calculated using block re-sampling with 200 permutations.

In Figure 13, we show how treated households' additional payments are divided among realized overage charges, internet tier upgrades, and changes to TV subscriptions. In most percentile bins, overage charges and upgrade fees each contribute about half of the overall change in payments reported in Figure 12. Changes to TV subscription payments are indistinguishable from zero in all but one percentile bin, which is line with the estimates of subscription changes reported in Figure 6. Households with the greatest values of $\hat{\rho}_i(\lambda)$, i.e., expected overages of \$36.50 or more, are the exception to the patterns described above. We find that these households pay significantly more in additional usage fees versus internet tier upgrades. This could be due to especially heavy internet usage being a transient phenomenon for households. In these cases, it may be worthwhile to pay overage charges while demand is high rather than making a commitment to an internet tier that would accommodate all of their usage or already being on the tier with the greatest allowance.

During UBP's initial implementation, we observe some notable changes to the composi-



Figure 13: Change in Payments by Expected Overage Level

Notes: Heterogeneity in the effect of UBP on monthly bill level. Households are grouped into bins using percentiles of the $\hat{\rho}_i(\lambda)$ distribution. All bins except for the left-most bin contain 2% of households; the left-most bin contains all remaining households. Confidence intervals are calculated using block re-sampling with 200 permutations.

tion of the additional revenue. In Figure 14, we show initial trends in payments of the three types of charges we highlight above. Additional payments by treated households for internet tier upgrades, displayed on the left, rise steadily through the announcement and treatment periods. Households with greater expected overages make greater additional payments for tier upgrades throughout this portion of the sample period. Overage charges, displayed on the right, fall over time for households in the two highest expected overage bins, while charges for treated households in lower bins rise slightly over time or remain flat. Differences in levels and changes are generally small among these households. In general, consistent with optimizing behavior by consumers, overage charges are less than the typical price to upgrade the internet tier.



Figure 14: Change in Payments by Month



(b) Usage Fees

Notes: Change in internet subscription fees and usage fees during the sample. Pre-policy change usage fees (dotted line) are the implied usage fees had the policy change been in effect. No usage fees were billed until month 0. Households are grouped using percentiles of the $\hat{\rho}_i(\lambda)$ distribution. All groups except for the smallest $\hat{\rho}_i(\lambda)$ bin contain 2% of households; the smallest bin contains all remaining households.

6 Conclusions

Usage-based pricing of internet access has attracted scrutiny because of its potential to shift consumers' choices on several margins and reallocate surplus from consumers and thirdparty content providers to MSOs. Despite interest from policymakers and other interested parties, there has been little prior empirical research to evaluate UBP's steering and metering impacts. We use novel panel data on an MSO's introduction of UBP to measure its effects on consumers, third-party content sources, and MSO revenue. We exploit highly detailed subscription and usage data on treated households, to whom UBP was introduced, and matched control households to construct measures of heterogeneous treatment intensity and effects.

We find that consumers facing nontrivial charges under UBP responded meaningfully to the policy, largely through their internet subscriptions and usage. A significant share of consumers facing a high cost of inaction upgraded their internet tiers to accommodate their internet usage, while other "in the money" consumers reduced their usage in order to limit overage charges while remaining in their original tiers. Thus, UBP served as an effective instrument for metering households' internet demand, prompting greater payments to the MSO from households that value usage most. Notably, we uncover two results that are contrary to some warnings about UBP. First, despite concerns that UBP will disproportionately affect OTT, we find no meaningful difference between OTT and other internet content when consumers reduced usage under UBP. Second, UBP was ineffective in steering consumers toward the MSO's TV service. In total, UBP primarily served to transfer surplus from consumers to the firm, which we observe through increases in the MSO's revenue from overage fees and tier upgrades. This highlights the potential for distributional issues that follow from UBP, but a full consideration of welfare effects should account for the relationship between firm revenue and its incentives for network investment, which benefits consumers.

There are a number of issues that remain for future research. Despite the richness of our data, more detail is required to understand how the increasingly complex relationships between MSOs and content providers impact their pricing and steering incentives. MSOs that are vertically integrated with content producers may have a greater incentive to use pricing to steer customers towards their TV service, or they may implement non-neutral pricing that favors their content. The recent growth of stand-alone streaming services (e.g., Peacock, Disney+, HBO Max, ESPN+, etc) has also altered MSOs' incentives. Given OTT's popularity, MSOs may focus on strategies to capture surplus from these services rather than steer customers toward their TV services.

References

- Abadie, Alberto and Jérémy L'Hour (2021). "A Penalized Synthetic Control Estimator for Disaggregated Data." Journal of the American Statistical Association, 116(536): 1817– 1834.
- Anderson, Laura M, Lars Gårn Hansen, Carsten Lynge Jensen and Frank A Wolak (2017). "Using Real-Time Pricing and Information Provision to Shift Intra-Day Electricity Consumption: Evidence from Denmark." Department of Economics, Stanford University.
- Armstrong, Mark (2006). "Competition in Two-Sided Markets." RAND Journal of Economics, 37(3): 668–691.
- Bauer, Johannes M. and Steven S. Wildman (2012). "The Economics of Usage-Based Pricing in Local Broadband Markets." Technical report, Michigan State University.
- Bourreau, Marc, Frago Kourandi and Tommaso Valletti (2015). "Net Neutrality with Competing Internet Platforms." *Journal of Industrial Economics*, 63(1): 30–73.
- Chillemi, Ottorino, Stefano Galavotti and Benedetto Gui (2020). "The impact of data caps on mobile broadband Internet access: A welfare analysis." *Information Economics and Policy*, 50: 100843.
- 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.
- Chung, Doug, Thomas Steenburgh and K Sudhir (2010). "Do Bonuses Enhance Sales Productivity? A Dynamic Structural Analysis of Bonus-Based Compensation." HBS Working Paper 11-041.
- Copeland, Adam and Cyril Monnet (2009). "The Welfare Effects of Incentive Schemes." *Review of Economic Studies*, 76(1): 93–113.
- 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 Ali Yurukoglu (2012). "The Welfare Effects of Bundling in Multichannel Television Markets." *American Economic Review*, 102(2): 643–685.

- Crawford, Gregory S., Robin S. Lee, Michael D. Whinston and Ali Yurukoglu (2018). "The Welfare Effects of Vertical Integration in Multichannel Television Markets." *Econometrica*, 86(3): 891–954.
- Crawford, Gregory S., Oleksandr Shcherbakov and Matthew Shum (2019). "Quality Overprovision in Cable Television Markets." *American Economic Review*, 109(3): 956–95.
- Duflo, Esther, Rema Hanna and Stephen P. Ryan (2012). "Incentives Work: Getting Teachers to Come to School." *American Economic Review*, 102(4): 1241–78.
- Dutz, Mark, Jonathan Orszag and Robert Willig (2009). "The Substantial Consumer 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 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.
- Einav, Liran, Amy Finkelstein and Paul Schrimpf (2015a). "Bunching at the Kink: Implications for Spending Responses to Health Insurance Contracts." *Working Paper*.
- Einav, Liran, Amy Finkelstein and Paul Schrimpf (2015b). "The Response of Drug Expenditure to Non-Linear Contract Design: Evidence from Medicare Part D." *Quarterly Journal* of Economics, 130(2): 841–899.
- Fauqui, Ahmad and Sanem Sergici (2010). "Household Response to Dynamic Pricing of Electricity: A Survey of 15 Experiments." *Journal of Regulatory Economics*, 38(2): 193– 225.
- Gans, Joshua (2015). "Weak Versus Strong Net Neutrality." Journal of Regulatory Economics, 47(2): 183–200.

- Goetz, Daniel (2019). "Dynamic Bargaining and Size Effects in the Broadband Industry." 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.
- Hitte, Loran and Prasanna Tambe (2007). "Broadband Adoption and Content Consumption." Information Economics and Policy, 74(6): 1637–1673.
- Ito, Koichiro (2014). "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." American Economic Review, 104(2): 537–563.
- 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.
- MacKie-Mason, Jeffrey K. and Hal Varian (1994). "Economic FAQs About the Internet." The Journal of Economic Perspectives, 8(3): 75–96.
- Malone, Jacob, Aviv Nevo and Jonathan W. Williams (2021). "The Tragedy of the Last Mile: Economic Solutions to Congestion in Broadband Networks." *Working Paper*.
- Malone, Jacob, John Turner and Jonathan Williams (2014). "Do Three-Part Tariffs Improve Efficiency in Residential Broadband Networks?" *Telecommunications Policy*, 38(11): 1035–1045.
- Malone, Jacob B., Aviv Nevo, Zachary Nolan and Jonathan W. Williams (2023). "Is OTT Video a Substitute for TV? Policy Insights from Cord-Cutting." *The Review of Economics* and Statistics, 105(6): 1615–1623.

- McManus, Brian, Aviv Nevo, Zachary Nolan and Jonathan W. Williams (2024). "The Steering Incentives of Gatekeepers in the Telecommunications Industry." *Working Paper*.
- Miravete, Eugenio (2003). "Choosing the Wrong Calling Plan? Ignorance and Learning." American Economic Review, 93(1): 297–310.
- Misra, Sanjog and Harikesh Nair (2011). "A Structural Model of Sales-Force Compensation Dynamics: Estimation and Field Implementation." *Quantitative Marketing and Economics*, 9(3): 211–257.
- Nevo, Aviv, John Turner and Jonathan Williams (2016). "Usage-Based Pricing and Demand for Residential Broadband." *Econometrica*, 84(2): 411–443.
- Newsham, Guy R and Brent G Bowker (2010). "The Effect of Utility Time-Varying Pricing and Load Control Strategies on Residential Summer Peak Electricity Use: A Review." *Energy Policy*, 38(7): 3289–3296.
- Odlyzko, A., B. St. Arnaud, E. Stallman and M. Weinberg (2012). "Know Your Limits: Considering the Role of Data Caps and Usage Based Billing in Internet Access Service." Technical report.
- Prince, Jeffrey and Shane Greenstein (2017). "Measuring Consumer Preferences for Video Content Provision via Cord-Cutting Behavior." Journal of Economics & Management Strategy, 26(2): 293–317.
- 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.
- Strapp, James, Chris King and Sharon Talbott (2007). "Ontario Energy Board Smart Price Pilot Final Report." IBM Global Business Services and eMeter Strategic Consulting for the Ontario Energy Board.
- Varian, Hal (2002). The Demand for Bandwidth: Evidence from the INDEX Experiment, volume Broadband: Should We Regulate High-Speed Internet Access? Brookings Institution.

- Wolak, Frank A (2007). "Residential Customer Response to Real-Time Pricing: The Anaheim Critical Peak Pricing Experiment." *Center for the Study of Energy Markets.*
- Wolak, Frank A (2010). "An Experimental Comparison of Ciritcal Peak and Hourly Pricing: The PowerCentsDC Program." *Department of Economics, Stanford University.*
- Wolak, Frank A (2016). "Designing Nonlinear Price Schedules for Urban Water Utilities to Balance Revenue and Conservation Goals." *National Bureau of Economic Research*.
- Wu, Tim (2003). "Network Neutrality, Broadband Discrimination." Journal of Telecommunications and High Technology Law, 1(2): 141–178.

Appendix A Penalty Selection

To select an appropriate value of the tuning parameter λ , we follow the leave-on-out crossvalidation approach of post-intervention outcomes for the untreated approach described in Abadie and L'Hour (2021). Specifically, for each candidate λ and each control household j, we implement the PSC procedure to construct a synthetic control for household j from a donor pool comprised of all other control households (i.e., the donor pool for household j is $\{1, \ldots, n_c) \setminus j$). We then calculate the "treatment effect,"

$$\tilde{\tau}_{j}(\lambda) = \sum_{k \neq j, \, k=1}^{n_{c}} \left[\frac{1}{M_{1}} \sum_{m=1}^{M_{1}} \left(y_{jm} - y_{km} W_{j,k}^{*}(\lambda) \right) \right],$$

where the outcome of interest y_{im} is the monthly usage of household *i* during UBP period month *m*.

We use these estimates to identify the values of λ that minimize mean squared prediction error and bias, where prediction error and bias are defined as follows:

$$RMSE(\lambda) = \left(\frac{1}{n_c} \sum_{j=1}^{n_c} \tilde{\tau}_j(\lambda)^2\right)^{1/2}$$
$$Bias(\lambda) = \left|\frac{1}{n_c} \sum_{j=1}^{n_c} \tilde{\tau}_j(\lambda)\right|.$$

In Table 3, we describe the results of the exercise for 6 values of λ , including prediction error, bias, and the density of the estimated synthetic control weights. For a given λ , each density statistic describes the distribution across synthetic controls of the count of units in the donor pool that receive positive (non-zero) weight.

				Density	
λ	$ \text{Bias}(\lambda) $	$\mathrm{RMSE}(\lambda)$	Min	Median	Max
$\rightarrow 0$	29.77	106.91	1	1130	3474
0.001	1.05	95.65	1	7	11
0.01	1.19	95.37	1	7	10
0.1	0.26	95.97	1	5	9
1	0.02	103.32	1	3	7
10	0.93	113.52	1	1	4

Table 3: Cross-Validation Results

Appendix B Robustness

The estimates reported in the main text are obtained by fixing the λ parameter at 0.1. In this section, we describe the sensitivity of pre-treatment fit and the robustness of the estimates reported in the main text to alternative values of λ .

In Panel I of Table 4, we report the mean levels of the 11 pre-treatment covariates used in the construction of the synthetic controls across treated households, untreated households, and five alternative synthetic controls.

In Panel II of Table 4, we report the distribution of household-level usage treatment effects $(\hat{\tau}_i(\lambda))$ and expected overages $\hat{\rho}_i(\lambda)$ for each set of alternative synthetic controls.

Panel I: Pre-treatment Fit			Synthetic Control				
	Treated	Untreated	$\lambda = 0.001$	$\lambda = 0.01$	$\lambda = 0.1$	$\lambda = 1$	$\lambda = 10$
Usage Month 1	92.26	130.36	92.38	92.39	92.18	91.28	91.00
Usage Month 2	90.90	130.84	91.19	91.27	91.37	90.70	90.39
Usage Month 3	101.90	140.55	102.42	102.39	101.75	100.01	99.56
Share Video	0.45	0.52	0.45	0.45	0.45	0.45	0.45
Share Browsing	0.43	0.36	0.43	0.43	0.43	0.43	0.43
Share Netflix	12.20	5.37	10.36	10.33	9.98	9.47	9.31
Share YouTube	8.30	3.65	7.05	7.03	6.79	6.44	6.34
Share Linear OTT	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Usage Variance	17.40	23.07	17.27	17.30	16.36	14.39	13.75
Has TV	0.78	0.70	0.78	0.78	0.78	0.78	0.78
Internet Tier	2.82	3.07	2.82	2.82	2.82	2.82	2.82

Table 4: Fit and Estimates for Alternative λ Penalties

Panel II: Estimates

Panel II: Estimates			Synthetic Control				
	Treated	Untreated	$\lambda = 0.001$	$\lambda = 0.01$	$\lambda = 0.1$	$\lambda = 1$	$\lambda = 10$
Usage Treatment Effect							
Mean		-53.60	-0.05	-0.08	-0.11	-0.11	-0.11
SD			3.38	3.40	3.46	3.66	4.03
5th Ptile			-4.45	-4.55	-4.77	-5.25	-5.99
10th Ptile			-2.78	-2.87	-3.05	-3.33	-3.71
Median			-0.11	-0.12	-0.12	-0.08	-0.03
90th Ptile			2.82	2.80	2.85	3.04	3.35
95th Ptile			4.85	4.88	4.93	5.18	5.56
Expected Overages							
Mean			3.32	3.04	2.65	2.26	2.16
SD			13.95	12.70	11.45	10.14	10.85
Median			0.00	0.00	0.00	0.00	0.00
90th Ptile			7.22	6.47	4.92	1.99	0.00
95th Ptile			20.17	19.00	16.66	13.65	10.00

Notes: Panel I: household-level averages of pre-treatment matching variables for treated households, untreated households, and five sets of synthetic controls. Panel II: distribution of household-level estimated usage treatment effects and expected overages for five sets of synthetic control. Untreated treatment effect is a simple average difference between treated and untreated households.