

USAGE-BASED PRICING AND DEMAND FOR RESIDENTIAL BROADBAND

BY AVIV NEVO, JOHN L. TURNER, AND JONATHAN W. WILLIAMS¹

We estimate demand for residential broadband using high-frequency data from subscribers facing a three-part tariff. The three-part tariff makes data usage during the billing cycle a dynamic problem, thus generating variation in the (shadow) price of usage. We provide evidence that subscribers respond to this variation, and we use their dynamic decisions to estimate a flexible distribution of willingness to pay for different plan characteristics. Using the estimates, we simulate demand under alternative pricing and find that usage-based pricing eliminates low-value traffic. Furthermore, we show that the costs associated with investment in fiber-optic networks are likely recoverable in some markets, but that there is a large gap between social and private incentives to invest.

KEYWORDS: Demand, broadband, dynamics, usage-based pricing.

1. INTRODUCTION

THE TELECOMMUNICATIONS SECTOR IS UNDERGOING MAJOR CHANGES and is the focus of several important public policy debates. A driving force behind these changes is the growing importance of data services and the proliferation of online activities, especially the popularity of (over the top) video providers such as Netflix and YouTube. Cable companies, which once mainly delivered video, are shifting their focus to broadband services. The same is true for cellular carriers, whose networks are increasingly used to deliver data. Traditional telecom companies are trying to keep up with this trend and offering their version of data delivery services.² In this paper, we contribute an important ingredient for studying the economics of this industry: demand for residential broadband. In particular, we estimate demand using a unique data set and (shadow) price variation created by usage-based pricing. We demonstrate the

¹We are grateful to the North American Internet Service Providers that provided the data used in this paper. We thank participants in several seminars, Gautam Gowrisankaran, Shane Greenstein, Tom Holmes, Nathan Miller, Mar Reguant, Ron Reuss, Scott Savage, and Tommaso Valletti for insightful comments, as well as the co-editor and four anonymous referees. Jim Metcalf and Craig Westwood provided expert IT support for this project and Jacob Malone provided excellent research assistance. We gratefully acknowledge support from NSF Grants SES-1324851 and SES-1324717.

²Major recent public policy discussions in this sector include the proposed mergers between Comcast and Time Warner Cable, which was called off, ATT and DirecTV, which was approved with some conditions (one of which was a limit on the use of usage-based plans), and Charter and Time Warner (currently under review). Another proposed merger, between Sprint and T-Mobile, was canceled but was rumored to be driven by incentives to invest in next-generation wireless technology that will deliver faster data services. Other policy issues facing the FCC and Congress are the equal treatment of content on the Internet (so-called net neutrality) and encouragement of municipal broadband networks.

implications of the demand estimates by computing consumers' plan choices and usage when faced with a variety of contracts including unlimited plans and high-speed fiber-to-the-premise (FTTP) options.

In order to estimate demand for residential broadband, we rely on two sources of variation. We use variation in prices and attributes of the plans subscribers choose, but more importantly, we rely on variation created by three-part tariff plans we observe. Subscribers pay a monthly fee, which provides them a monthly data allowance. If they exceed the allowance, they pay a price per gigabyte (GB). When facing a three-part tariff, the marginal price paid for usage is zero until the subscriber exceeds their allowance. However, a forward-looking subscriber realizes that the shadow price of usage depends on how many days are left in the billing cycle and the fraction of the allowance already used. To exploit this variation, we build a dynamic model of utility-maximizing subscribers' intertemporal decision making throughout a billing cycle.

At the core of the paper is a data set we secured from an Internet Service Provider (ISP). The data include information on hour-by-hour Internet usage for roughly 55,000 subscribers facing different price schedules. We also know plan-specific variables (speed, prices, etc.) for the plan the household subscribes to and for the alternatives not chosen. The ISP has in place three-part tariff plans in addition to subscribers who are grandfathered in to unlimited plans. Download speed, a key attribute of a plan, is measured in megabits per second (Mb/s).³ With a 15 Mb/s connection, roughly the mean in our data for usage-based plans, one can download 900 megabits, or 112 megabytes, per minute. A 1 GB file will take about 9 minutes to download. A standard-definition movie is approximately 2.5 GB and a high-definition movie twice that. The speed we use in the analysis is advertised speed, but realized speeds might vary due to congestion. This does not seem to be an issue in our data, since at the time the ISP's network was rarely impacted by congestion.

Using these data, we provide descriptive evidence that consumers respond to variation in the shadow price of usage. We then estimate a (finite horizon) dynamic choice model, by adapting the techniques of [Akerberg \(2009\)](#), [Bajari, Fox, and Ryan \(2007\)](#), and [Fox, Kim, Ryan, and Bajari \(2011\)](#). Specifically, we solve the dynamic problem for a large number of subscriber types, once for each type. We then estimate the distribution over these types by matching moments recovered from the data to those predicted by a weighted average of the optimal behavior of the types.

The estimates allow us to quantify several aspects of consumer demand for broadband. First, we find that consumers' willingness to pay for speed is het-

³As a quick reminder, a megabit is equal to 2^{20} , or 1,048,576, bits. Speed is measured in megabits per second. A megabit should not be confused with a megabyte, which is the standard for measuring file size. A byte is equal to 8 bits.

erogeneous, which is intuitive given the different ways consumers use the Internet. The estimates suggest that consumers will pay between \$0 and \$5 per month for a 1 Mb/s increase in connection speed, with an average of \$2. The mean download speed in our data is roughly 15 Mb/s, and therefore a 1 Mb/s increase is roughly 6.7% faster. The median household uses 24 GB of data per month. The increased speed shortens the download time by roughly 14 minutes per month, and therefore the average of \$2 willingness to pay amounts to roughly \$8 per hour. For households with higher usage, the time savings are greater, potentially explaining the higher willingness to pay. With the availability of more content and applications, consumers will likely increase their usage, therefore implying greater time savings and accordingly a greater willingness to pay for speed.

Second, we find a large difference between the marginal and infra-marginal value of content. The average consumer's willingness to pay for a 1 GB increase in usage allowance is \$0.36 per month, which suggests that marginal content has relatively low value. On the other hand, the infra-marginal value of content is high. At a price of zero and with average download speed, we find that the surplus of the average consumer would be \$165 per month, and average usage about 66 GB per month (compared to \$85 per month and 48 GB per month in the existing setting). The large difference between the value of marginal and infra-marginal content lines up with our intuition of the value of online content.

To check our estimated model's ability to predict usage growth, as well as to check how well the estimates fit out of sample, we take the menu of plans offered by a second provider. We do not use data from this provider in estimation, because it does not offer usage-based pricing. We are able to accurately predict the usage levels for this provider both in June 2012 and in June 2015. We do especially well at matching the growth rate, of roughly 102%, during this period.

We demonstrate the implications of the demand estimates by calculating usage and welfare under several alternative scenarios. We start by evaluating the implications of usage-based pricing. Usage-based pricing has been proposed as one way to manage congestion in the current bandwidth-intensive environment. The term typically refers to nonlinear pricing based on the quantity of usage, not the type or composition of usage, which is at the heart of the net-neutrality debate. Usage-based pricing is popular for broadband service outside the United States, and for cellular plans in the United States. However, it has generated a policy discussion in the U.S. when proposed as the standard for pricing broadband service ([Open Internet Advisory Committee \(2013\)](#)). Much of the debate on the usage and welfare implications of usage-based pricing has been theoretical ([Mackie-Mason and Varian \(1995\)](#), [Bauer and Wildman \(2012\)](#), [Odlyzko, Arnaud, Stallman, and Weinberg \(2012\)](#)), and has not been

informed by data. See [De Fontenay, Shugard, and Sibley \(1990\)](#) for a discussion of similar issues with telecommunications services in the 1980s.⁴

We find that usage-based pricing is effective at lowering usage without reducing consumer welfare significantly, relative to a world with just unlimited plans. This is driven directly by the finding that marginal content is not very valuable and that consumer welfare is mainly driven by infra-marginal usage. Generally, usage-based pricing shifts surplus from consumers to providers. The magnitude, as well as the effect on total welfare, depends on the prices of the unlimited plans in the counterfactual setting. We explore several scenarios for these prices, but we do not solve for the equilibrium prices.

Next we evaluate adoption, usage, and welfare when consumers are presented with an unlimited plan with a connection speed of one gigabit per second (Gb/s). We find that surplus generated from usage is substantial. Yet at a fee of \$70, which is what Google charges for Google Fiber in Kansas City, the ISP captures only a small portion of this surplus. A general finding is that there is a big gap between the social return and the private return. For example, using cost estimates from [Kirjner and Parameswaran \(2013\)](#), we estimate it takes 27 months to recover the capital expenditures from a social perspective, relative to typical cable offerings, and only 12 months relative to not having any broadband service. On the other hand, a typical ISP that upgrades to gigabit speeds would recover these costs only after about 149 months. The exact recovery time depends on the competitive environment, but the general point—that there is a large gap between social and private incentives to invest—is quite robust. This gap has been recognized by policy makers, who have pushed for municipal broadband networks. We evaluate adoption and usage of these plans under several combinations of monthly fees and speeds.

Instead of using the variation in the shadow price, as the dynamic model does, one could estimate a static, or myopic, model using the price variation from overage charges, as well as plan choice data. This will lead to an underestimate of the response to price. The intuition is simple. We show that in the data, consumers reduce their consumption as they near the allowance. The dynamic model interprets this as a response to the shadow price. In a static framework, the price is unchanged, so the lowered consumption implies a lower (average) usage at a price of zero, which would lead to an underestimate of the price sensitivity. This intuition is confirmed empirically: we find that the price response from static estimation is, on average, 38.6% lower than the dynamic estimates.

Our paper is related to a literature that studies demand for broadband service. [Varian \(2002\)](#) and [Edell and Varaiya \(1999\)](#) ran experiments where users

⁴Another way to manage congestion is by setting peak-load prices and giving users incentives to shift usage to off-peak times. Without a change in some of the major applications that currently require real-time streaming, this strategy is unlikely to be effective.

faced varying prices for different allowances and speeds. Goolsbee and Klenow (2006) used data on individuals' time spent on the Internet and earnings to estimate consumer benefit from residential broadband, assuming an hour spent on the Internet is an hour of forgone wages. Lambrecht, Seim, and Skiera (2007) used monthly consumption data from a German ISP to study the role of uncertainty in consumers' selection of usage-based plans. Several additional papers (Dutz, Orszag, and Willig (2009), Rosston, Savage, and Waldman (2010), Greenstein and McDevitt (2011)) estimated the economic value of broadband Internet using plan choice data. Hitt and Tambe (2007) showed that broadband adoption increases Internet usage by 1,300 minutes per month, suggesting a strong preference for content that requires high bandwidth.

The modeling in this paper is related to several literatures. First is a literature that focuses on estimating demand in dynamic settings (e.g., Crawford and Shum (2005), Hendel and Nevo (2006a), Gowrisankaran and Rysman (2012), and others). Like our analysis, Yao, Mela, Chiang, and Chen (2012) exploited intra-month (weekly) variation in the shadow price of usage under three-part tariffs to identify consumers' discount factors. Second is a literature studying incentives under nonlinear contracts. In labor contracts, a nonlinear compensation structure based on performance during a fixed period of time makes the worker's decision regarding the optimal level of effort a dynamic one, in much the same way usage is under a three-part tariff (e.g., Copeland and Monnet (2009), Chung, Steenburgh, and Sudhir (2010), Duflo, Rema, and Ryan (2012), Misra and Nair (2011)). In deciding on health care expenditures with an annual deductible and an out-of-pocket maximum, consumers also face a similar trade-off (Einav, Finkelstein, and Schrimpf (2015a, 2015b)). Finally, our paper is related to a literature that examines if consumers are forward-looking (Aron-Dine, Einav, Finkelstein, and Cullen (2015), Chevalier and Goolsbee (2009), Grubb (2015), Grubb and Osborne (2015), and Hendel and Nevo (2006b)).

2. DATA

The data come from a North American ISP that offers several plans. Features of a plan include maximum download speed, an access fee, usage allowance (if any), and overage price per GB of data (if any).⁵ Unlimited plans, where subscribers do not face overage prices, are only available to subscribers who previously had them. Usage in GBs is recorded for both uploads and downloads, but for billing purposes, and consequently our purposes, the direction of the traffic is ignored. For each subscriber, we observe usage at the monthly level from May 1st, 2011 to May 31st, 2012, and for 15-minute inter-

⁵Subscribers are not on long-term contracts, only incurring a disconnection fee if service is canceled.

vals during May 10th to June 30th, 2012. We also know the plan chosen by the subscriber.⁶

2.1. *Sample and Descriptive Statistics*

The sample includes 54,801 subscribers in four different markets served by the ISP. The residents of these four markets had per-capita income of \$47,592 in 2011, relative to \$45,222 for residents in all U.S. metropolitan markets.

The data demonstrate a sharp increase in usage. The median subscriber, to our ISP, more than doubles usage, from 9 GBs in May 2011 to over 20 GBs in May 2012, an increase equivalent to roughly four movies. The average subscriber's usage increases from 23 to over 40 GBs. This trend is not unique to our provider or to the period 2011–2012. In Table I, we present the change in usage and composition of traffic from 2012 to 2015 for another provider.⁷ Over the course of three years, average usage doubles. The increase is almost completely due to increased video usage: in 2012, video accounts for roughly a third of usage, while in 2015, video was almost two thirds of total traffic. A question we will return to later is whether our estimates of preferences, using 2012 data, can match 2015 usage.

TABLE I
COMPOSITION AND USAGE GROWTH, 2012 VERSUS 2015^a

Source	% of Usage		% Usage Growth
	2012	2015	2012 to 2015
Video	34.1	61.1	260.1
Web Browsing	31.9	21.5	36.2
File Sharing	8.3	0.2	−95.2
Gaming	1.3	3.1	357.1
Music	0.4	3.4	1,650.0
Backup	0.2	0.5	400.0
Other	23.7	10.3	−12.4
Total	100.0	100.0	101.5

^aThis table presents the percent of usage from different uses in 2012 and 2015, as well as the percent growth in GBs between 2012 and 2015 from each source and overall. The data are for June 2012 and June 2015 and come from a North American ISP, different from our main provider. In these data, we do not observe usage-based pricing, but we do see more information about the content.

⁶Unfortunately, we do not have information on bundling of other services. So we do not know if the subscriber is paying for the Internet service as part of a bundle.

⁷The data we have from this provider do not include usage-based plans and therefore we do not use these data in the rest of this paper. On the other hand, for the provider we use below, we do not have 2015 data or the breakdown by source, so we cannot generate the equivalent of Table I using its data.

Since a large proportion of traffic consists of bandwidth-intensive online video offerings (e.g., Netflix, YouTube, etc.), it can be argued that Internet consumption is increasingly “lumpy.” For these types of content, consumption often involves a single download or upload of substantial size. One way to measure lumpiness in usage is the ratio of daily usage relative to the household’s mean. If the distribution was not lumpy, we would commonly observe a ratio near 1. Instead, we see that 25% (50%) of the daily usage observations are less than 23% (65%) of their mean, and 5% of observations are over three times the household’s average daily consumption. In Section 5, we discuss the model’s ability to match this lumpiness in usage.

For the main analysis, we use disaggregated data, which include one complete billing cycle for each subscriber during May and June, 2012. Usage shows a cyclical pattern during the day. Peak usage occurs during 10 pm–11 pm, when the average user consumes over 4.5 GBs each month. This is almost six times the amount of traffic generated during 5 am–6 am. Throughout the day, approximately 90% of the usage is for download. Aggregating usage (uploads and downloads) to a daily level results in 1,644,030 subscriber-day observations.

Table II reports summary statistics on monthly usage and plan characteristics for the May–June 2012 billing cycle. These statistics highlight the corre-

TABLE II
DESCRIPTIVE STATISTICS OF SUBSCRIBER PLAN CHOICES AND
USAGE, MAY–JUNE 2012^a

	Unlimited Plans	Usage-Based Plans
<i>Number of Subscribers</i>	12,316	42,485
<i>Plan Characteristics</i>		
Mean Access Fee (\$)	44.33	74.20
Mean Download Speed (Mb/s)	6.40	14.68
Mean Allowance (GB)	∞	92.84
Mean Overage Price (\$/GB)	–	3.28
<i>Usage</i>		
Mean (GB)	50.39	43.39
Median (GB)	25.60	23.63
Median Price per GB (\$)	1.68	3.02
<i>Overages</i>		
Mean Share of Allowance Used (%)	–	49.02
Subscribers Over Allowance (%)	–	9.45
Median Overage (GB)	–	17.03
Median Overage Charges (\$)	–	51.19

^aThese statistics reflect characteristics of plans chosen and usage by subscribers to a single ISP, in four markets during May–June 2012. Usage is based upon Internet Protocol Detail Record (IPDR) data, captured in 15-minute intervals and aggregated to the monthly level. Means and medians are at the subscriber level.

lations between allowances, overage prices, and usage. An average subscriber to an unlimited plan pays \$44.33 for a month of service, enjoys a maximum download speed of 6.40 Mb/s and uses just over 50 GB. In contrast, an average subscriber to a usage-based plan pays nearly \$30 more per month to enjoy faster download speed (14.68 Mb/s), but uses under 44 GB. The median subscriber to an unlimited plan pays about \$1.68 per GB, less than 60% of what a usage-based subscriber pays. Of course, the subscribers to the grandfathered unlimited plans were early adopters and could have high usage for that reason (although one would think that they would also want to shift to the higher speed of the usage-based plans).

In Table II, we also report summary statistics for overage charges incurred by subscribers on plans with usage-based pricing. During the May–June 2012 billing cycle, about ten percent of subscribers on plans with usage-based pricing exceed their allowance. This is important, as our identification strategy relies on having enough subscribers with a positive probability of incurring overage charges during the month. On average, subscribers use slightly less than half of their usage allowance. For those subscribers exceeding their allowance, the median overage is 17 GBs and the corresponding charge is \$51.19. Figure 1 presents the variation in the fraction of the allowance used by consumers for the May–June 2012 billing cycle.

Our data-use agreement prevents us from disclosing the actual plan features, but we can show an approximate relative ranking of the plan features and costs. For each plan, Figure 2 presents how the total cost of subscribing to a plan

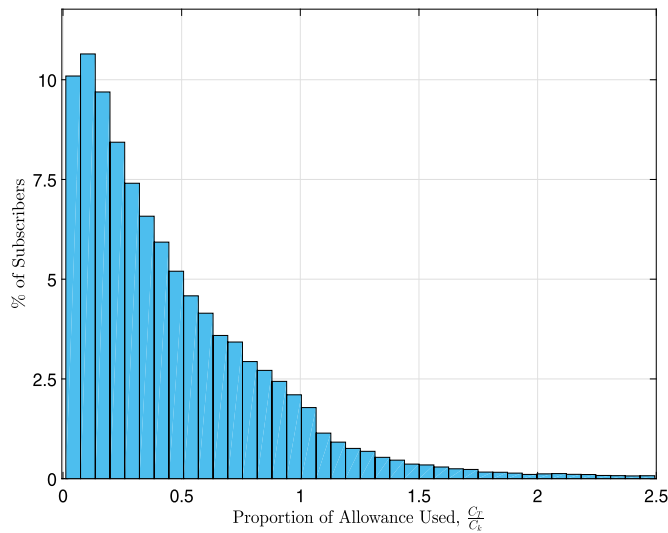


FIGURE 1.—Proportion of allowance used. *Note:* This figure presents a histogram where each observation represents a consumer’s monthly usage relative to their allowance, both in GBs.

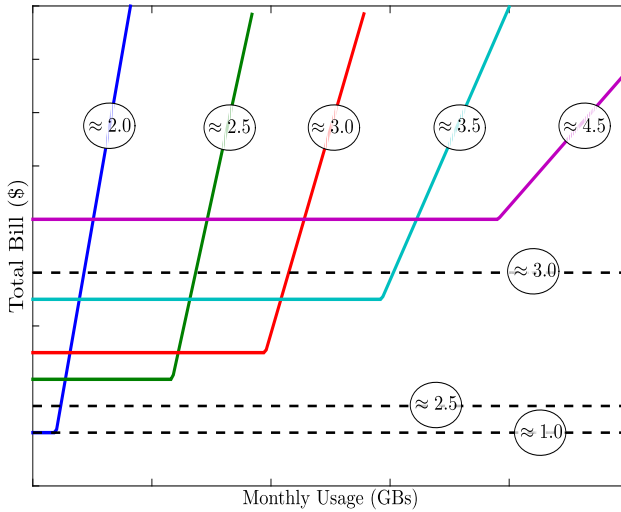


FIGURE 2.—Plan features and billing. *Note:* This figure illustrates the relationship between monthly usage and cost for the usage-based (solid lines) and grandfathered unlimited plans (dashed lines). The approximate relative speed for each plan, normalized by the slowest plan, is indicated by the circle intersecting each line.

changes with monthly usage. Each plan is also labeled with a relative speed ranking, with the slowest plan normalized to 1.

2.2. Are Subscribers Forward-Looking?

We now examine whether consumers in our data are forward-looking. This is interesting for two reasons. First, the evidence we provide adds to a growing literature demonstrating that consumers are forward-looking when making economic choices. Second, our identification relies on consumers responding to changes in the shadow price of usage over a billing cycle. It is therefore useful to know that consumers are indeed responding to this variation before proceeding to the model.

Our data are from an ISP that allows subscribers to carefully track their usage, by receiving text messages and emails at regular intervals after they exceed one-half of their allowance. Consumers may also log into the provider's web site at any time. We therefore have confidence that subscribers are aware of cumulative usage during the month.

If subscribers are forward-looking, we expect certain patterns in usage throughout a billing cycle. The heaviest-volume subscribers who know they have high probability of exceeding their allowance should behave as though the shadow price is equal to the overage price: there should be little change in average usage throughout the billing cycle. Similarly, for subscribers with

only a small probability of exceeding their allowance, behavior should not vary throughout the billing cycle. The only exception would be a small increase in usage towards the end of the billing cycle when the probability of exceeding the usage allowance approaches zero. For subscribers between these two extremes, usage should vary significantly depending on both the day in the billing cycle and a subscriber's cumulative usage up until that day.

To test whether consumers respond to the price variation introduced by past usage within a billing cycle, we estimate the following regression:

$$(1) \quad \ln(c_{jkt}) = \sum_{m=1}^{M=4} \sum_{n=1}^{N=5} \alpha_{nm} \mathbb{1} \left[pct_n \leq \frac{C_{jk(t-1)}}{\bar{C}_k} < pct_{n+1} \right] \\ \times \mathbb{1}[day_m \leq t < day_{m+1}] + \mathbf{x}_t \boldsymbol{\psi} + \mu_j + \varepsilon_{jkt},$$

where $\ln(c_{jkt})$ is the natural logarithm of subscriber j 's usage on day t , on plan k . The ratio $\frac{C_{jk(t-1)}}{\bar{C}_k}$ is the proportion of the usage allowance used up until day t , or the subscriber's total usage in the previous $(t-1)$ days of the billing cycle, $C_{jk(t-1)} = \sum_{\tau=1}^{t-1} c_{jk\tau}$, divided by the usage allowance on plan k , \bar{C}_k . The first set of indicators, $\mathbb{1}[pct_n \leq (\frac{C_{jk(t-1)}}{\bar{C}_k}) < pct_{n+1}]$, equals 1 when the proportion of a subscriber's usage allowance that has been used to date is in a particular range, such that $pct_1 = 0$, $pct_2 = 0.40$, $pct_3 = 0.60$, $pct_4 = 0.80$, $pct_5 = 1.00$, and $pct_6 = \infty$. The other set of indicators, $\mathbb{1}[day_m \leq t < day_{m+1}]$, equals 1 when the day is in a particular range, such that $day_1 = 10$, $day_2 = 15$, $day_3 = 20$, $day_4 = 25$, and $day_5 = 31$. We omit the interactions for the first ten days of the billing cycle. We include dummy variables for the days of the week and a time trend to account for any organic growth in usage over the course of the billing cycle. Since different households have different billing dates, we can separate time trends and demand shocks on certain days from the dynamics of our model. We also include subscriber fixed effects, μ_j , to remove persistent forms of heterogeneity across subscribers.

The estimates of Equation (1) are reported in Table III. Each cell reports the estimate for the coefficient on the interaction between the indicators in the respective row and column. The patterns in the table are consistent with forward-looking behavior: at each point in the billing cycle, current usage is lower the closer the consumer is to the allowance (and hence the shadow price is higher). Furthermore, subscribers who are near the allowance early in the billing cycle reduce usage less than subscribers near the allowance later in the billing cycle (i.e., coefficients decrease monotonically from left to right within rows four and five of Table III). This is consistent with consumers near the allowance later in the billing cycle reducing usage proportionally more, relative to their own mean, than consumers nearing the allowance early in the billing cycle. For subscribers well below the allowance late in the billing cycle, we observe a small increase in usage, consistent with these subscribers becoming confident that they will not exceed the allowance.

TABLE III
FORWARD-LOOKING BEHAVIOR, WITHIN-MONTH REGRESSION^a

	$\mathbb{1}[10 \leq t < 15]$	$\mathbb{1}[15 \leq t < 20]$	$\mathbb{1}[20 \leq t < 25]$	$\mathbb{1}[25 \leq t < 31]$
$\mathbb{1}[0 \leq \frac{C_{jk(t-1)}}{C_k} < 0.40]$	-0.04** (0.01)	-0.04** (0.01)	0.03** (0.01)	0.08** (0.01)
$\mathbb{1}[0.40 \leq \frac{C_{jk(t-1)}}{C_k} < 0.60]$	-0.02 (0.02)	-0.12** (0.01)	-0.12** (0.01)	-0.04** (0.01)
$\mathbb{1}[0.60 \leq \frac{C_{jk(t-1)}}{C_k} < 0.80]$	-0.07** (0.03)	-0.12** (0.02)	-0.20** (0.02)	-0.16** (0.01)
$\mathbb{1}[0.80 \leq \frac{C_{jk(t-1)}}{C_k} < 1.00]$	-0.19** (0.05)	-0.26** (0.03)	-0.39** (0.02)	-0.42** (0.02)
$\mathbb{1}[1.00 \leq \frac{C_{jk(t-1)}}{C_k}]$	-0.12** (0.05)	-0.35** (0.03)	-0.41** (0.02)	-0.47** (0.02)
Adjusted R^2	0.46			

^aThis table presents OLS estimates of Equation (1) using 1,644,030 subscriber-day observations. The dependent variable is natural logarithm of daily usage. Each cell in the table gives the coefficient on the interaction between the indicators in the respective row and column. Controls include a constant, time trend, indicators for the day of the week, and subscriber fixed effects. Asterisks denote statistical significance: **1% level, *5% level.

In addition to the within-month variation in price, subscribers also encounter a change in the shadow price when their usage allowance is refreshed at the beginning of a new billing cycle. A forward-looking subscriber near the allowance at the end of a billing cycle knows that the shadow price decreases at the beginning of the next billing cycle. Conversely, a subscriber well below the allowance likely experiences an increase in the shadow price as the new billing cycle begins. Subscribers well over the allowance at the end of the billing cycle, who expect to go over the allowance again next month, should behave as though the price always equals the overage price and not respond at all.

For most subscribers, we observe at least one day of usage beyond the full billing cycle used for the rest of our analysis, allowing for a test of whether subscribers respond to this across-month price variation. To do so, we first calculate the percentage change in usage from the final day of the billing cycle ($t = 30$) to the first day of the next billing cycle ($t = 31$) for each subscriber, $\frac{C_{jk(30)} - C_{jk(31)}}{C_{jk(30)}}$. We then calculate the mean percentage change for groups of subscribers that used various fractions of the allowance by the end of the month, $\frac{C_{30}}{C_k}$. Figure 3 presents the results. Subscribers facing a price increase at the beginning of the next month consume relatively more at the end of the current month, while those expecting a price decrease consume relatively less. We observe little change in usage for those well above the allowance in the cur-

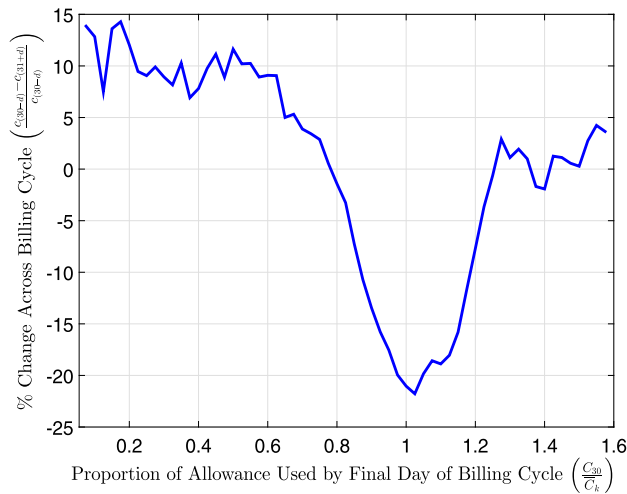


FIGURE 3.—Across-month dynamics. *Note:* This figure presents how the percentage change in usage from the last day of a billing cycle to the first day of the next varies with the proportion of the allowance consumed by a subscriber at the end of the billing cycle.

rent month.⁸ Collectively, our results provide support for the hypothesis that subscribers are forward-looking. Consumers are responsive, in an economically meaningful way, to variation in the shadow price of usage both within and across billing cycles.

3. MODEL

We model the subscriber's problem in two stages. The subscriber first chooses a plan anticipating future demand for content, and then chooses usage given the chosen plan.

3.1. Utility From Content

Subscribers derive utility from content and a numeraire good. To consume content, each subscriber chooses a plan, indexed by k . Each plan is characterized by the speed content is delivered, s_k , by a usage allowance, \bar{C}_k , by a fixed fee, F_k , that pays for all usage up to the allowance, and by an overage price, p_k per GB of usage in excess of the allowance. For any plan, the number of days in the billing cycle is T .

⁸In the Supplemental Material Section S.1.2 (Nevo, Turner, and Williams (2016)), Figure 8, we provide further analysis to demonstrate that the result holds for different time windows.

Utility from content is additively separable over all days in the billing cycle.⁹ Let consumption of content on day t of the billing cycle be c_t and the consumption of the numeraire good on day t be y_t . The utility for a subscriber of type h on plan k is given by

$$u_h(c_t, y_t, v_t; k) = v_t \left(\frac{c_t^{1-\beta_h}}{1-\beta_h} \right) - c_t \left(\kappa_{1h} + \frac{\kappa_{2h}}{\ln(s_k)} \right) + y_t.$$

The first term captures the subscriber's gross utility from usage. Quickly declining marginal utility seems natural: a subscriber's first email sent, favorite website to check, first Netflix/YouTube video, etc., should bring higher marginal utility than subsequent usage. The specification—where the *curvature* of the utility function can vary from log ($\beta_h \rightarrow 1$) to linear ($\beta_h = 0$)—permits an interpretation based on a close link to price elasticity of demand. The utility from consumption is scaled by a time-varying shock, v_t , which captures randomness in utility from consumption. The shock is observed by the subscriber only in period t . For type h , each v_t is independently and identically distributed $\text{LN}(\mu_h, \sigma_h)$, truncated at point \bar{v}_h to exclude the top 0.5% of the distribution. For simplicity, we denote type h 's distribution of v_t as G_h , and refer to μ_h and σ_h as the mean and standard deviation of the distribution. A more general model would relax the i.i.d. assumptions in a variety of ways. For example, the distribution could vary by day of the week, exhibit serial correlation, or be more “lumpy” (to account for different usage generated, say, by email and movies). However, as we show in Sections S.1.1–S.1.3 of the Supplemental Material, it seems like none of these is a major concern in our data.

The second term captures the subscriber's non-price cost of consuming online content. Marginal cost is constant, at $\kappa_{1h} + \frac{\kappa_{2h}}{\ln(s_k)}$. The parameter $\kappa_{1h} > 0$ captures the consumer's *opportunity cost of content*. The ratio $\frac{\kappa_{2h}}{\ln(s_k)}$, where $\kappa_{2h} > 0$ is the subscriber's *preference for speed*, captures the waiting cost of transferring content. This specification implies that the subscriber has a satiation point absent overage charges: the marginal utility from content is decreasing while the opportunity cost of time is constant. This explains why consumers on unlimited plans consume a finite amount of content. Our specification assumes that the cost of consuming content is linear, but the marginal utility is convex. Alternatively, we could allow for curvature in cost. However, it is not clear how the data would separate these two and it seems to us that a linear cost, reflecting time cost, and declining marginal utility is the natural assumption.

⁹In this way, we assume content with a similar marginal utility is generated each day or constantly refreshed. This may not be the case for a subscriber who has not previously had access to the Internet. Below, we will assume decreasing marginal utility within a time period, but additive across periods. This seems to us as a natural assumption here, but it does potentially create sensitivity to how the time period is defined.

The vector of parameters, $(\beta_h, \kappa_{1h}, \kappa_{2h}, \mu_h, \sigma_h)$, describes a subscriber of type h . Conditional on choosing plan k , this subscriber’s problem is

$$\begin{aligned} & \max_{\{c_1, \dots, c_T\}} \sum_{t=1}^T E[u_h(c_t, y_t, v_t; k)] \\ \text{s.t. } & F_k + p_k \text{Max}\{C_T - \bar{C}_k, 0\} + Y_T \leq I, \\ & C_T = \sum_{j=1}^T c_j, \quad Y_T = \sum_{j=1}^T y_j. \end{aligned}$$

We do not discount future utility since we model daily decisions, over a finite and short horizon. Uncertainty involves the realizations of v_t . We assume that wealth, I , is large enough so that it does not constrain consumption of content.

3.2. Optimal Usage

We now solve for the optimal usage implied by the model. Denote the unused allowance at the beginning of period t , for a subscriber on plan k , as $\bar{C}_{kt} \equiv \text{Max}\{\bar{C}_k - C_{t-1}, 0\}$. Similarly, denote period- t overage as $\mathcal{O}_{tk}(c_t) \equiv \text{Max}\{c_t - \bar{C}_{kt}, 0\}$.

In the terminal period (T) of a billing cycle, there are no intertemporal tradeoffs. The subscriber solves a static utility maximization problem, given cumulative usage up until period T , C_{T-1} , and the realization of the preference shock, v_T . For a subscriber well below the allowance (i.e., \bar{C}_{kT} is high) and without a high draw of v_T , it is optimal to consume content up to the point where $\frac{\partial u_h(c_t, y_t, v_t; k)}{\partial c_t} = 0$. If marginal utility at $c_t = \bar{C}_{kT}$ is positive but below p_k , then it is optimal to consume exactly the remaining allowance. For a subscriber who is already above the allowance (i.e., $\bar{C}_{kT} = 0$) or who draws a high v_T , it is optimal to consume up to the point where $\frac{\partial u_h(c_t, y_t, v_t; k)}{\partial c_t} = p_k$. Denoting this optimal level of consumption by $c_{hkT}^*(C_{T-1}, v_T)$, the subscriber’s utility in the terminal period is then

$$\begin{aligned} V_{hkT}(C_{T-1}, v_T) &= v_T \left(\frac{(c_{hkT}^*)^{1-\beta_h}}{1-\beta_h} \right) - c_{hkT}^* \left(\kappa_{1h} + \frac{\kappa_{2h}}{\ln(s_k)} \right) \\ &+ y_t - p_k \mathcal{O}_{tk}(c_{hkT}^*). \end{aligned}$$

For any other day in the billing period $t < T$, usage adds to cumulative consumption and affects the next period’s state, so the optimal policy function for a subscriber incorporates this. We therefore solve for the optimal usage recur-

sively. Specifically, type h on plan k solves

$$c_{hkt}^*(C_{t-1}, v_t) = \arg \max_{c_t} \left\{ v_t \left(\frac{c_t^{1-\beta_h}}{1-\beta_h} \right) - c_t \left(\kappa_{1h} + \frac{\kappa_{2h}}{\ln(s_k)} \right) + y_t - p_k \mathcal{O}_{tk}(c_t) + E[V_{hk(t+1)}(C_{t-1} + c_t)] \right\}.$$

Define the *shadow price* of consumption

$$\tilde{p}_k(c_t, C_{t-1}) = \begin{cases} p_k, & \text{if } \mathcal{O}_{tk}(c_t) > 0, \\ \frac{dE[V_{hk(t+1)}(C_{t-1} + c_t)]}{dc_t}, & \text{if } \mathcal{O}_{tk}(c_t) = 0. \end{cases}$$

Then the consumer's optimal choice in period t satisfies

$$(2) \quad c_{hkt}^* = \left(\frac{v_t}{\kappa_{1h} + \frac{\kappa_{2h}}{\ln(s_k)} + \tilde{p}_k(c_{hkt}^*, C_{t-1})} \right)^{1/\beta_h}.$$

Equation (2) implies that a type with parameter β_h has demand elasticity equal to $-\frac{1}{\beta_h}$ with respect to changes in the *total disutility* of content, $\kappa_{1h} + \frac{\kappa_{2h}}{\ln(s_k)} + \tilde{p}_k(c_t, C_{t-1})$. The demand elasticity with respect to changes in $\tilde{p}_k(c_t, C_{t-1})$ does not equal $-\frac{1}{\beta_h}$. Intuitively, a subscriber with curvature β_h will be less sensitive to changes in $\tilde{p}_k(c_t, C_{t-1})$ than an elasticity of $-\frac{1}{\beta_h}$ implies.

The value functions are given by

$$V_{hkt}(C_{t-1}, v_t) = v_t \left(\frac{(c_{hkt}^*)^{1-\beta_h}}{1-\beta_h} \right) - c_{hkt}^* \left(\kappa_{1h} + \frac{\kappa_{2h}}{\ln(s_k)} \right) + y_t - p_k \mathcal{O}_t(c_{hkt}^*) + E[V_{hk(t+1)}(C_{t-1} + c_{hkt}^*)]$$

for each ordered pair (C_{t-1}, v_t) . Then for all $t < T = 30$, the expected value function is

$$E[V_{hkt}(C_{t-1})] = \int_0^{\bar{v}_h} V_{hkt}(C_{t-1}, v_t) dG_h(v_t),$$

and the mean of a subscriber's usage at each state is

$$(3) \quad E[c_{hkt}^*(C_{t-1})] = \int_0^{\bar{v}_h} c_{hkt}^*(C_{t-1}, v_t) dG_h(v_t).$$

The solution to the dynamic program for each type implies a distribution for the time spent in particular states (t, C_{t-1}) over a billing cycle.

3.3. Plan Choice

Subscribers select plans before observing any utility shocks. Specifically, entering the first period with $C_0 = 0$, the subscriber selects plan $k \in \{1, \dots, K\}$ to maximize expected utility. The subscriber may also choose no plan at all, $k = 0$. Formally, the plan choice is given by

$$k_h^* = \arg \max_{k \in \{0, 1, \dots, K\}} \{E[V_{hk1}(0)] - F_k\},$$

where the value $E[V_{h01}(0)]$ and the fixed access fee F_0 for $k = 0$, the outside option, are normalized to 0. Note that we assume that there is no error, so consumers choose the plan that is optimal.

Previous work has studied whether consumers make what look like suboptimal choices *ex post* in a variety of settings.¹⁰ In our data, it is not obvious that subscribers systemically make mistakes. One way to measure these mistakes is to ask how many subscribers could switch to a plan that costs less and is no slower. Using this definition, if we look at the complete billing cycle (June 2012) in isolation, we find that 7.24% of subscribers used a dominated plan. However, the frequency of this type of mistake goes down to 0.13% if we ask how many subscribers could have paid less and used service that is no slower during the 13 months from May 2011 to May 2012. This is a weak test of the optimality of plan choice since, as we see in Figure 2, speed and usage allowances are positively correlated. However, nothing in our data suggests that consumers are consistently making what seems *ex post* like clear mistakes.

4. ESTIMATION AND IDENTIFICATION

We estimate the parameters of the model using a method-of-moments approach similar to the two-step algorithms proposed by [Akerberg \(2009\)](#), [Bajari, Fox, and Ryan \(2007\)](#), and [Fox et al. \(2011\)](#). First, we solve the dynamic program for a wide variety of subscriber types. Second, we estimate a weight for each of the types by matching the weighted average of optimal behavior, calculated in the first stage, to the equivalent moments observed in the data. This yields an estimated distribution of types. In this section we outline the main steps, and provide more details and sensitivity analysis in Section S.2 of the Supplemental Material.

In step 1 of the estimation, we solve the dynamic program for 16,807 types (seven points of support for each of the five parameters), where each type is

¹⁰In particular, some research has highlighted what seem like suboptimal choices made by consumers facing nonlinear pricing, similar to ours, in cell phone usage ([Grubb and Osborne \(2015\)](#)) and health care ([Abaluck and Gruber \(2011\)](#), [Handel \(2013\)](#)). On the other hand, several papers (e.g., [Miravete \(2003\)](#), [Economides, Seim, and Viard \(2008\)](#), [Goettler and Clay \(2011\)](#), [Ketcham, Lucarelli, Miravete, and Roebuck \(2012\)](#)) highlight circumstances where individuals make choices that are rational *ex post*.

defined by a value of the parameter vector $(\beta_h, \kappa_{1h}, \kappa_{2h}, \mu_h, \sigma_h)$. For a plan, k , and subscriber type, h , we solve the finite-horizon dynamic program described in the previous section recursively, starting at the end of each billing cycle. To do so, we discretize the state space. Because the subscriber does not know v_t prior to period t , we can integrate over its support, and the solution to the dynamic programming problem for each type of subscriber can be characterized by the expected value functions, $E[V_{hkt}(C_{t-1})]$, and policy functions, $E[c_{hkt}^*(C_{t-1})]$. Having solved the dynamic program for a subscriber of type h , we generate the transition process for the state vector implied by the solution.

In step 2 of the estimation, we choose a weight for each subscriber type to match moments we recover from the data to the (weighted) average of the behavior predicted by the model.¹¹ Formally, we choose weights to satisfy

$$\hat{\theta} = \arg \min_{\theta} \mathbf{m}_k(\theta) \hat{\mathbf{V}}^{-1} \mathbf{m}_k(\theta),$$

$$\text{subject to } \sum_{h=1}^{H_k} \theta_h = 1 \quad \text{and} \quad \theta_h \geq 0 \quad \forall h.$$

The plan-specific vector $\mathbf{m}_k(\theta)$ is given by $\mathbf{m}_k(\theta) = \hat{\mathbf{m}}_k^{\text{dat}} - \mathbf{m}_k^{\text{mod}} \theta$, where $\hat{\mathbf{m}}_k^{\text{dat}}$ is the vector of moments recovered from the data, $\mathbf{m}_k^{\text{mod}} \theta$ is a weighted average of the equivalent type-specific moments predicted by the model, and $\hat{\mathbf{V}}^{-1}$ is a weighting matrix. Note that type weights, θ_h , are chosen to match the moments for each plan, and they sum up to 1 for each plan. After the estimation, we rescale the weights by the probability that each plan is chosen, and therefore we also match the share of each plan in the data.

In choosing which moments to match, we focus on two considerations: identification and computational ease. Estimation is much simpler if the moments are linear in the weights.¹² We therefore choose the following two sets of moments. First, we use the mean usage at each state $\sum_{h=1}^H E[c_{hkt}^*(C_{t-1})] \times \gamma_{hkt}(C_{t-1}) \theta_h$, where $E[c_{hkt}^*(C_{t-1})]$ is the mean usage of type h in time t given plan k and past usage of C_{t-1} , and $\gamma_{hkt}(C_{t-1})$ is the probability that this type reaches the state. Note that the average is taken across all types on the plan, not just those that arrive at the state with positive probability, which keeps the moment linear in the parameters. The second set of moments is the mass of subscribers at a particular state, $\sum_{h=1}^H \gamma_{hkt}(C_{t-1}) \theta_h$, which, like the first set of moments, is easy to calculate and linear in the weights.

¹¹We allow for 16,807 types, but we estimate weights for only 8,626 types. Roughly half the types are ruled out by the plan choice, since they choose none of the offered plans and there are no usage moments.

¹²As pointed out by Bajari, Fox, and Ryan (2007) and Fox et al. (2011), least squares minimization subject to linear constraints, and over a bounded support, is a well-defined convex optimization problem. Even though the optimization is over a potentially large number of weights, it is quick and easy to program as long as the moments are linear in the weights.

We calculate standard errors using a block-resampling methodology (Lahiri (2003)). Specifically, we sample the data by consumer with replacement, keeping all 30 days for each of 54,801 consumers drawn, which results in 1,000 samples of size 1,644,030. For each sample, we recalculate the moments and then re-estimate the weights. We calculate standard errors for subsequent statistics and counterfactual analyses by repeating the calculation using the 1,000 different estimates of the weights.

The estimation procedure recovers the weights of each type by choosing the mixture of types that best matches the data. The data we use to identify the parameters includes plan choice and usage. The logic of identification follows, in some ways, that of Bayesian estimation. The selection of the boundaries of the initial grid amounts to putting a uniform prior on the distribution of types over the grid (and zero probability elsewhere). Plan selection—specifically, the share of consumers who choose each plan—provides information on the type distribution. In the model, each type has an optimal plan, so plan selection splits the type space into distinct groups but does not provide any information about the relative importance of types within a group. In other words, it puts a weight on each group of types equal to the share of the plan that group chooses. As the number of plans increases, and the attributes and prices of the plans vary, the type distribution can be recovered.

The usage moments allow us to distribute the weight among types within a group who choose each plan. Since the objective function is linear in the weights, the intuition for how the weights are identified is similar to that of a linear regression: the weights are identified as long as the behavior predicted by different types is not collinear over all the moments and all states used in estimation. Thus, the key to identification is to understand how each parameter impacts the variation in predicted behavior across moments and states.¹³

We rely on two types of usage variation. The first is variation in mean usage across different states (i.e., day in the billing cycle and fraction of allowance used). The second is the probability of reaching various states. Since cumulative usage (at different times of the month) is a state variable, the probability of reaching a state depends on higher-order moments of the usage distribution and efficiently summarizes the information in these moments. Note that the first set of moments we use are the unconditional mean usage, that is, taking into account usage conditional on getting to a state and the probability of reaching that state.

Each parameter impacts behavior in these states in a different way. For example, roughly speaking, the parameter μ_h determines average usage across days for each type. To see this, consider Equation (2), but fix $\beta_h = 1$ and as-

¹³The logic behind the identification follows closely the formal argument in Kasahara and Shimotsu (2009).

sume that the shadow price does not vary with the parameters. In this case, average consumption is given by

$$E[c_{hkt}^*] = \frac{E(v_t)}{\kappa_{1h} + \frac{\kappa_{2h}}{\ln(s_k)} + \tilde{p}}.$$

An increase in the average of the shocks impacts usage in all states. On the other hand, changes in the disutility of consuming content, κ_{1h} and κ_{2h} , have a different effect for different shadow prices. With $\beta_h \neq 1$ and with the shadow price varying with the parameters, the same idea holds: the change in mean usage across states is affected in a different way by the different parameters.

In addition to variation in average usage across states, we also rely on changes to higher-order moments of usage. A change in σ_h impacts the variance of usage and therefore the likelihood of reaching certain states. Other parameters also impact the probability of reaching certain states. For example, while different combinations of the curvature parameter β_h and the disutility parameters, κ_{1h} and κ_{2h} , may imply similar average usage, these parameters affect higher-order moments differently. For example, a high-curvature (high β_h), low-disutility (low κ_{1h}) subscriber could have similar average usage as a low-curvature (low β_h), high-disutility (high κ_{1h}) subscriber. But the latter subscriber is relatively less sensitive to movements in the shadow price than the former. Hence, mean usage across states and the probability of reaching states will be different for these two types. Figures 10 and 11 in the Supplemental Material illustrate this point.

The nonlinearity of the parameters makes it hard to precisely define what variation in the data identifies each parameter. However, loosely speaking, plan selection is critical for pinning down κ_{2h} . The mean usage across all states pins down μ_h and the variance in usage across all states pins down σ_h . Finally, as we discussed in the previous paragraph β_h and κ_{1h} are identified from the variation in usage across states with different shadow prices and the probability of reaching different states. In Section S.2.4 of the Supplemental Material, we further demonstrate how types behave differently by looking at the behavior of the types we estimate.

To demonstrate the importance of plan selection and the usage moments in pinning down the distribution of types, we present below the estimated distribution when we only use plan choice and when we use both plan choice and usage.

5. RESULTS

We estimate a weight greater than 0.01% ($\theta_h > 0.0001$) for 53 types. The most common type accounts for 28% of the total mass, the top 5 types account

for 65%, the top 10 for 78%, and the top 20 for 90%.¹⁴ No plan has more than 20 types receiving positive weights, while the average number of types across plans is only 6.6.¹⁵ To get an idea of what the results imply: the most common type is predicted to use 29 GB per month and have a willingness to pay of about \$72 for a plan with a speed of 14.68 Mb/s, the mean in our data, and no overage charges.

Figure 4 demonstrates the importance of plan selection and information on usage (during the billing cycle) for identifying the distribution of types. As we noted above, in the model, each type has an optimal plan, so plan selection splits the type space into distinct groups but does not provide any information about the relative importance of types within a group. Figure 4(a) presents the joint distribution of the utility curvature, β_h , and the mean of the distribution of random shocks, μ_h , if only plan selection information is used and uniform weights are applied to all types that select a particular plan, that is, if we put equal weight on each of the types within a group (so that the weights sum up to the share of the plan), and then integrate over the other parameters to recover the joint distribution of β_h and μ_h . Figure 4(b) presents our estimates of the joint distribution when we use both plan choice and usage information. The two distributions are very different, suggesting that the usage moments are a crucial source of information.¹⁶


The estimated joint distribution in Figure 4(b) is highly irregular and looks quite different from a normal distribution. For the highest-volume subscribers (high μ_h), there is substantial variation in the elasticity of demand. In fact, for high- μ_h subscribers, the distribution of β_h is clearly multi-peaked (unconditional on values of other parameters). The majority of high-volume subscribers have highly elastic demand, a value of β_h less than or equal to 0.3, including the most common type of subscriber. Most of the remainder of the high- μ_h subscribers have less elastic demand, or a value of β_h greater than or equal to 0.7. The distribution of the other parameters is similarly irregular and multi-peaked. For example, a relatively small group of individuals places high value on increased connection speeds (high κ_{2h}), but the majority of high- μ_h subscribers have a relatively low preference for speed.

Overall, our model fits the data quite well. For all plans, the correlation between the empirical moments and the fitted moments is very high. The model also fits patterns in the data not explicitly used in estimation. For example, bandwidth-intensive activities, like online video and cloud-based services, will generate a lumpiness in usage. To show that our model can match this behavior, we calculate in the data, and simulate from the model, the ratio of daily


¹⁴Further information about the estimated distribution is available upon request.

¹⁵Expanding the grid of types to allow for two additional values of each parameter results in estimates that assign no weight greater than 0.01% to any of the additional types.

¹⁶In Section S.2.3 of the Supplemental Material, we show the marginal distribution of all five parameters when we use different moments for estimation.



(a) Only plan selection



(b) Plan selection and usage

FIGURE 4.—Sources of identification: plan selection and usage. *Note:* (a) presents the joint distribution of the utility curvature parameter, β_h , and the mean of shocks, μ_h , when only information on optimal plan selection is used and uniform weights within group are applied. (b) presents the distribution when information on optimal plan selection is used and the weights are chosen to match usage moments from the data.

consumption to a consumer's mean usage over a billing cycle. In the model, this ratio will vary both because of variation in the state and because of random draws to the consumption shocks, v_t . We therefore simulate usage predicted by the model 1,000 times for each consumer type over a 30-day period,

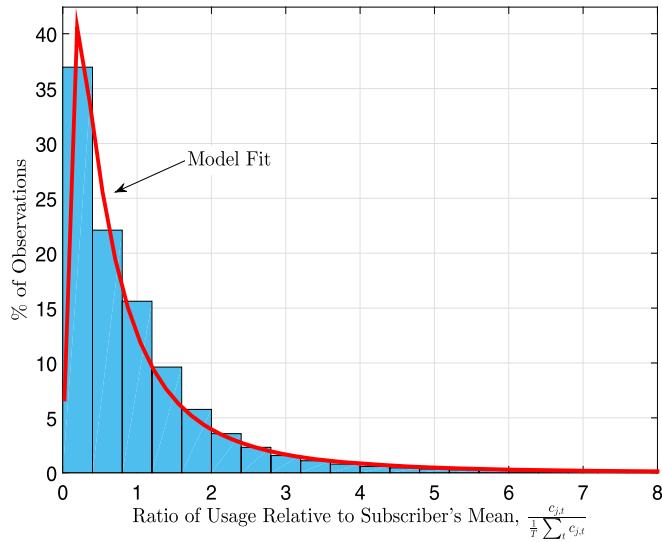



FIGURE 5.—Model fit: distribution of usage relative to a subscriber's mean. *Note:* This figure presents the ratio of daily usage, c_{jt} , to a subscriber's monthly average, $\frac{1}{T} \sum_i c_{jt}$, from the data and simulations from the model.


perform the same calculation for each of the simulations, and then aggregate across types accounting for the relative mass of each. Figure 5 presents the distribution of this ratio for the data and the model simulations.

The model and estimates also produce accurate out-of-sample predictions. To demonstrate this, we use the estimated type distribution to predict usage on the plans of the ISP, whose data we presented in Table I. This is a completely different ISP, with different plans, than the ISP whose data we use to estimate the model. First, we predict the usage in both June 2012 and June 2015. Our predictions are within 6% and 5% of the observed usage, respectively. If we focus on the growth rate between 2012 and 2015, the predictions match the data almost exactly. The actual growth rate is 102% and the prediction is 104%. The growth rate in the data can be driven by the change in plan characteristics (such as price and speed), or by a change in preferences or availability of content. The fact that the estimated model fits the growth pattern suggests that it is the former that is driving the change and not the latter. Additionally, the ability of the estimates to provide accurate out-of-sample predictions gives us confidence that the high dimensionality of the problem, that is, the large number of weights we estimate, does not lead us to over fit within sample.

To demonstrate the implications of the estimated type distribution, Figure 6 shows the distribution of willingness to pay to increase the usage allowance by 1 GB on the first day of the billing cycle, $\frac{dEV_{hk1}(0)}{d\bar{C}_k}$ (in panel (a)), and speed by 1 Mb/s for the entire billing cycle, $\frac{dEV_{hk1}(0)}{ds_k}$ (in panel (b)). Figure 6(a) shows



(a) Value of increasing usage allowance by 1 GB



(b) Value of increasing speed by 1 Mb/s

FIGURE 6.—Distributions of value of increasing usage allowance by 1 GB and speed by 1 Mb/s. *Note:* (a) and (b) show the distribution of willingness to pay to increase usage allowance by 1 GB and to increase speed by 1 Mb/s, respectively.

that approximately eighty percent of subscribers have a positive probability of incurring overage charges and would be willing to pay to increase their allowance if given the opportunity. The average (median) willingness to pay for a 1 GB increase is \$0.36 (\$0.09), and the distribution is left-skewed with a small number of subscribers who are willing to pay substantial amounts. Figure 6(b) shows there is substantial variation in the preference for speed across consumers. The willingness to pay to improve the speed by 1 Mb/s ranges from nearly zero to just over \$5; the average is \$2.02 and the median is \$2.48. As we discuss in the [Introduction](#), the mean number translates to roughly \$8 per hour in saved time due to faster downloads.

To further visualize what our estimates imply for demand, we consider subscriber behavior under a linear tariff. Suppose the ISP eliminates access fees and instead sets a price p per GB, and offers just one download speed s . Because there is no fixed fee, every subscriber type consumes something under this plan. Conditional on v_t , a subscriber of type h chooses consumption according to Equation (2), with $s_k = s$ and $\tilde{p}_k(c_t, C_{t-1}) = p$. Taking expectations over G_h for each type, and averaging across subscriber types, expected daily demand for content is then

$$D(p) = \sum_{h=1}^H \hat{\theta}_h \int_0^{\bar{v}_h} \left(\frac{v}{\kappa_{1h} + \frac{\kappa_{2h}}{\ln(s)} + p} \right)^{1/\beta_h} dG_h(v).$$

in Table IV we present expected demand for five different speeds: (1) 2 Mb/s, a relatively slow speed by most standards; (2) 14.68 Mb/s, the average speed

TABLE IV
EXPECTED DAILY USAGE UNDER A LINEAR TARIFF^a

Price (\$)	Speed (Mb/s)				
	2	14.68	50	100	1,024
0.00	0.97 (0.005)	2.20 (0.010)	2.97 (0.015)	3.42 (0.018)	4.62 (0.034)
1.00	0.50 (0.002)	1.14 (0.001)	1.50 (0.005)	1.70 (0.006)	2.31 (0.009)
2.00	0.29 (0.001)	0.66 (0.002)	0.86 (0.003)	0.96 (0.003)	1.24 (0.004)
3.00	0.18 (0.001)	0.42 (0.001)	0.54 (0.001)	0.59 (0.001)	0.74 (0.001)
4.00	0.12 (0.001)	0.29 (0.001)	0.36 (0.001)	0.39 (0.001)	0.48 (0.001)
5.00	0.09 (0.001)	0.21 (0.001)	0.25 (0.001)	0.28 (0.001)	0.33 (0.001)

^aThis table presents the expected daily usage averaged across all subscriber types when facing a linear tariff. Standard errors, in parentheses, are calculated using a block-resampling methodology as described in the text.

for subscribers in our data; (3) 50 Mb/s, a mid-tier speed offered currently; (4) 100 Mb/s, a top-tier speed offered currently; and (5) 1,024 Mb/s, the highest speed currently offered in North America. We present expected demand for prices of 0 to \$5 per GB.

For the average speed in our data, subscribers facing a zero price would consume an average of 2.20 GB per day, or roughly 66 GB per month.¹⁷ This usage is significantly reduced as the price increases. At a price of just \$1 per GB, usage is cut by half, and at \$5 per GB, usage is more than 10 times lower. The impact of speed can be seen by comparisons across columns. At a price of zero, expected usage at a speed of 2 Mb/s is more than 60% lower than usage at the average speed in our data, while at speed of 1,024, expected usage more than doubles. Demand is much less price-sensitive at lower speeds, since waiting costs form a much greater part of the subscriber's overall costs from consuming content, reducing the effect of price.

The relatively modest increase in usage with large increases in speed is intuitive and provides a reality check on our estimates. The applications used by the average consumer in 2012 (or even 2015) do not require speeds that are substantially above the average speed in the data. Therefore, the significantly higher speeds do not induce much higher usage, on average. However, even with 2012 applications, these estimates demonstrate that current usage is limited by offered connection speeds. As the average user adopts more bandwidth-intensive applications, the volume of Internet traffic is likely to increase following investment in high-speed next-generation networks.

6. IMPLICATIONS FOR MANAGING INTERNET TRAFFIC

To further illustrate the implications of our estimates, we conduct several exercises to study different solutions proposed to managing the growth in Internet traffic.

6.1. *The Impact of Usage-Based Pricing*

We consider the impact of usage-based pricing (UBP) on usage, subscriber welfare, and the ISP's costs and revenues by comparing behavior under UBP, to consumers' plan and usage choices predicted by the demand model in various counterfactual settings. The analysis is not an equilibrium analysis, since we do not solve for the optimal offering of plans in the counterfactual setting. Instead, we use the model to simulate consumer behavior in a variety of settings.

¹⁷We note that the computed standard errors are relatively small. This is driven by the aggregate moments we use for estimation and the smoothing across states described in Section S.2.2 of the Supplemental Material, which leads to very low variance in the moments used for estimation. Thus, in the resampling algorithm we use to calculate standard errors, there is little variance in the moments to generate imprecision in the estimates.

TABLE V
USAGE-BASED PRICING VERSUS UNLIMITED PLANS^a

	(1)	(2)	(3)	(4)
<i>Scenario Description</i>				
UBP/Unlimited Plan Attributes	UBP current	Unlim current	Unlim typical US	Unlim rev-max F_k
<i>Usage and Surplus</i>				
Usage (GBs)	48.2 (0.203)	60.2 (0.261)	62.0 (0.264)	65.4 (0.322)
Speed (Mb/s)	14.2 (0.021)	10.3 (0.010)	10.8 (0.018)	12.6 (0.069)
Consumer Surplus (\$)	84.7 (0.810)	111.9 (0.791)	113.5 (0.789)	97.1 (0.810)
Revenue (\$)	69.4 (0.132)	42.1 (0.044)	44.8 (0.068)	64.3 (0.209)

^aThis table presents estimates of usage, surplus, and revenue information for several scenarios. Standard errors, in parentheses, are calculated using a block-resampling methodology as described in the text.

The results are presented in Table V. In column (1), we present the results when consumers face the UBP plans in the data, where we use the model to allocate the consumers who were on unlimited grandfathered plans to UBP plans (and calculate their usage). In column (2), we examine the case when consumers face the same set of plans except that allowances are unlimited (and monthly fees are held constant). Not surprisingly, since the marginal price of usage is zero and monthly fees are unchanged, usage increases substantially, as does consumer welfare. The ISP's revenue decreases mainly because consumers select cheaper plans that previously had a lower allowance but now differ only in speed. The sum of consumer surplus and revenue is slightly lower than column (1), and since the ISP's costs are higher (due to increased traffic on the network), total surplus is lower. Relative to identical unlimited plans, UBP increases total welfare, but mainly shifts surplus from consumers to the ISP.

In column (2), we hold constant the prices of plans, the speeds, and the number of plans. It is therefore likely that we overstate the surplus lost by subscribers and the revenue gains to the ISP from UBP. We explore two ways to relax this assumption. In column (3), we present results when consumers face a typical set of unlimited plans offered by a U.S. cable provider in 2012.¹⁸ Interestingly, the results are very similar to those from column (2).

Next, in column (4), we return to the set of plans we see in the data, but calculate the fees that maximize revenue associated with unlimited service. The

¹⁸In particular, the monthly fixed fees for the plans are \$34.99, \$47.99, \$59.99, and \$79.99 with speeds of 8, 12, 15, and 18 Mb/s.

ISP significantly raises fixed fees, collecting an average fee of \$64.3, excluding about 7% of those who would subscribe under UBP plans. The effect is as expected, consumer surplus decreases relative to column (2) due to the higher fees, and ISP revenues increase. As in the case of column (2), UBP shifts surplus to the ISP. However, unlike the results of column (2), where the direction of the effect relative to UBP is known in advance, here consumer surplus could be higher or lower when UBP is introduced. Unlimited plans have lower marginal prices but the fixed fees are likely higher, and more importantly, subscribers may switch plans, get higher speeds, and change usage. The results in Table V show that, on net, consumer surplus is lower with UBP.¹⁹ The calculation in column (4) might be overestimating the ability of the ISP to raise prices because we do not allow for entry or for a competitive response. To check this, we explore introducing a DSL option, which is slower than the plans in the data, or introducing a telecom with a fiber-to-the-node technology (such as ATT's U-verse). The results (not presented in the table) only strengthen the patterns we observe: UBP generally shifts surplus from consumers to the ISP.

6.2. *Economic Viability of Next-Generation Networks*

We now evaluate the economic viability of various network expansions. Recently, several ISPs and some municipalities have considered network expansions. These expansions could involve 1 Gb/s networks. This type of speed can be provided by next-generation high-speed broadband networks, such as FTTP, which has been introduced by, among others, Google.²⁰ DOCSIS 3.1 for cable-broadband networks is expected to be capable of similar performance. In addition to expansion plans by private providers, there is also a push to provide fast broadband by municipalities, which we evaluate at the end of this section.

As in the previous section, our analysis is not an equilibrium analysis: we do not have a supply model to determine the price of the FTTP plan. Furthermore, the speed is a significant extrapolation from what we see in the sample, and we (explicitly) hold the availability of content constant. One could imagine that when fast speeds become widely available, content and consumers' willingness to pay for speed could change (and probably be higher than our estimates).

The results are presented in Table VI. In column (1), we display adoption and usage when FTTP is offered for free. This provides a benchmark of consumer surplus generated by the availability of a fast connection. The higher

¹⁹In a working paper version of this paper we provided the consumer welfare numbers by types and showed that some consumers were better off under UBP. Indeed, the net result aggregating across consumers can be sensitive to the exact estimates.

²⁰See <https://fiber.google.com/about/> for current offerings and expansion plans.

TABLE VI
ADOPTION OF FTTP AND USAGE^a

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Scenario Description</i>						
Fee (F_k)	0	70	70	70	rev-max	rev-max
Competition			KC-cable	U-verse	KC-cable	U-verse
<i>Usage and Surplus</i>						
Usage (GBs)	138.8 (0.855)	136.6 (0.857)	134.5 (0.856)	134.4 (0.871)	133.1 (0.901)	132.0 (0.897)
Speed (Mb/s)	1024.0 (0.000)	977.9 (1.481)	687.0 (3.597)	673.0 (4.022)	596.4 (3.482)	592.8 (3.461)
Consumer Surplus (\$)	279.4 (1.025)	212.9 (1.014)	213.2 (0.968)	215.5 (0.981)	194.3 (0.922)	175.0 (0.889)
Revenue (\$)	0.00 (0.000)	66.8 (0.101)	55.3 (0.125)	58.5 (0.133)	77.7 (0.197)	95.3 (0.231)
FTTP Share (%)	100.0 (0.000)	95.5 (0.145)	64.7 (0.359)	67.1 (0.397)	57.2 (0.348)	57.2 (0.351)

^aThis table presents estimates of average usage, speed consumer surplus, revenue, and adoption for pricing options of FTTP as well as other broadband offerings. The adoption rates are for the population we estimated, namely, the users of broadband subscribers. Standard errors, in parentheses, are calculated using a block-resampling methodology as described in the text.

speed does indeed generate substantial surplus. However, due to a declining marginal value of speed implied by our utility function, speeds of more than 10 times those offered by the typical cable plans imply only 2 times the surplus.

In column (2), we evaluate usage and adoption when a fee of \$70 is charged, which is what Google charges in Kansas City for Google Fiber. At this fee, but with no alternatives, over 95% of households in our population, of broadband users, are predicted to adopt FTTP, and as a result usage decreases only slightly relative to a fee of zero. The next four columns examine the effect of alternative broadband offerings and different prices. In columns (3) and (4), the fee for FTTP is still \$70, but either cable or U-verse options are introduced as competitors.²¹ Adoption rates of FTTP fall significantly to roughly two thirds, with usage and speed being lower because some consumers choose the alternative plans. In the last two columns, when the FTTP provider charges revenue-maximizing fees, adoption of FTTP, usage, and speeds fall even further. The optimal FTTP fee in columns (5) and (6) are \$106.4 and \$134.4, respectively.

To address how long it will take to return the cost of investment in FTTP, we draw on estimates of the capital costs associated with the fiber network

²¹The cable offerings are similar to those available in Kansas City, prices of \$29.99, \$39.99, \$49.99, \$59.99, at speeds of 3, 25, 50, 105 Mb/s. U-verse offerings are \$29.95, \$34.95, \$44.95, \$64.95 for speeds of 3, 6, 18, 45.

built in Kansas City by Google Fiber (Kirjner and Parameswaran (2013)).²² If we compare FTTP to no availability of broadband at all, then from a social welfare point of view, the estimates from column (1) suggest that these capital costs can be recovered in approximately 12 months ($\$3,284/\279.4). If, alternatively, broadband competes with typical cable options, column (3) of Table V, then the social costs are recovered in approximately 27 months ($\$3,284/(\$279.4 - \$113.5 - \$44.8)$). From the ISP's perspective, the time to recover the capital costs is much higher. For example, for an existing cable service, it will take approximately 149 months ($\$3,284/(\$66.8 - \$44.8)$) to recover the investment cost if the revenue from FTTP, priced at \$70, is compared to the revenue of a typical U.S. cable plan (column (3) of Table V). The exact number is sensitive to what we assume about competition before and after FTTP is introduced and whether we consider a new entrant or an existing ISP. However, the general point—that there is a large gap between social and private incentives to invest—is quite robust.

This gap between social and private returns has been noted by policy makers and has led to a push for municipality-based broadband.²³ The advocates for muni-broadband make several arguments as to why the social return of a broadband network differs from the private return. There is a claim that because firms have limited ability to price discriminate (basically only through speed and possibly the monthly allowance), they cannot capture the infra-marginal gains from the network. There are other arguments for muni-broadband including (1) serving underserved populations, and (2) stimulating business relocations and formations, to spur growth. We can say something about the first argument but cannot speak to the latter two arguments.

In Table VII, we present results for adoption and usage of muni-broadband network for different speeds and fixed fees that are in the range of what is commonly seen. In each scenario in Table VII, there is a single offering with a fixed fee, F_k , and speed, s_k . Since municipal-broadband offerings are often introduced in more rural communities where innovative wifi solutions are necessary to reach subscribers (rather than fiber or coaxial networks), the speeds offered are typically lower than high-end cable offerings or FTTP. To fully address the optimal network speed, we would need cost data, which we do not have. Therefore, we focus only on adoption and usage. Our results in Table VII support the conclusions we saw above: the gap between revenue and social return is substantial.

²²The authors estimated that it will cost \$84 million dollars to “pass” 149,000 homes, or approximately \$564 per household. To actually connect each home, the authors estimated it will cost an additional \$464 per subscriber. If one assumes a 20% penetration rate for the service, this equates to capital costs of \$3,284 ($5 \times \$564 + \464) per household served.

²³See, e.g., a white paper issued by the White House https://www.whitehouse.gov/sites/default/files/docs/community-based_broadband_report_by_executive_office_of_the_president.pdf.

TABLE VII
MUNICIPAL BROADBAND^a

	(1)	(2)	(3)	(4)	(5)	(6)
<i>Plan Description</i>						
Fee (\$)	50	75	50	75	50	75
Speed (Mb/s)	25	25	50	50	100	100
<i>Usage and Surplus</i>						
Usage (GBs)	75.4 (0.344)	72.9 (0.339)	88.9 (0.432)	87.9 (0.429)	102.2 (0.532)	101.1 (0.530)
Speed (Mb/s)	24.6 (0.029)	21.0 (0.047)	49.2 (0.058)	46.2 (0.098)	98.4 (0.110)	92.5 (0.195)
Consumer Surplus (\$)	132.1 (0.827)	108.4 (0.818)	153.7 (0.855)	129.8 (0.846)	173.9 (0.888)	149.9 (0.878)
Revenue (\$)	49.2 (0.058)	63.0 (0.141)	49.2 (0.058)	69.3 (0.147)	49.2 (0.055)	69.3 (0.146)
Muni-BB Share (%)	98.4 (0.124)	84.0 (0.178)	98.4 (0.121)	92.4 (0.142)	98.4 (0.123)	92.4 (0.140)

^aThis table presents estimates of average usage, speed, consumer surplus, and revenue for pricing options of typical municipal broadband offerings. The adoption rates are for the population we estimated, namely, broadband subscribers. Standard errors, in parentheses, are calculated using a block-resampling methodology as described in the text.

7. CONCLUSION

We estimate demand for residential broadband using plan choices and high-frequency usage data. The three-part tariff plans make the usage problem dynamic and generate variation in the shadow price of usage. We show that consumers respond to this variation. Next, we use the variation in shadow price to estimate a dynamic choice model and then use the estimates to evaluate the usage and welfare implications of alternatives proposed to dealing with growth in Internet usage: usage-based pricing and high-speed next-generation networks. Our results suggest that usage-based pricing is an effective means to remove low-value traffic from the Internet. We find that FTTP generates significant consumer surplus but that there is a large gap between the private and social incentives for investment in such networks. Thus, without subsidization, these investments will come much later than is socially optimal.

Our estimates accurately predict usage in 2012, the time of our data, in a network of a different provider. The estimates do remarkably well in predicting the growth rate of usage between 2012 and 2015 in this provider's network. We do not know if this is by chance or if it is a more general result. Indeed, the only way to verify this finding is to apply the model to additional settings and time periods, which we hope will be done by future work.

There are several issues that our model does not address, and that we leave for future research. First, network congestion, which is argued to be a driver of the move towards usage-based pricing, was not necessary to model because

the ISP providing our data operated an overly provisioned network. An interesting question for future research is to measure the size and impact of congestion externalities among subscribers. Second, our analysis, because of data limitations, focused on GBs but not on the type of content viewed. In future work, we hope to have more detailed data on the actual content, which, coupled with information on TV viewing, would let us explore several questions on how consumers actually use broadband. Finally, our analysis aggregates to the daily level. However, as we noted, usage is significantly higher during peak periods. This suggests that a natural way to deal with congestion is to introduce peak-load pricing. To be effective, this will require changes in popular applications, such as Netflix. We leave it to future work to explore substitution between peak and non-peak periods, which is key to the effectiveness of peak-load pricing.

REFERENCES

- ABALUCK, J., AND J. GRUBER (2011): "Choice Inconsistencies Among the Elderly: Evidence From Plan Choice in the Medicare Part D Program," *American Economic Review*, 101 (4), 1180–1210. [426]
- ACKERBERG, D. (2009): "A New Use of Importance Sampling to Reduce Computational Burden in Simulation Estimation," *Quantitative Marketing and Economics*, 7 (4), 343–376. [412,426]
- ARON-DINE, A., L. EINAV, A. FINKELSTEIN, AND M. CULLEN (2015): "Moral Hazard in Health Insurance: How Important Is Forward-Looking Behavior?" *Review of Economics and Statistics*, 97 (4), 725–741. [415]
- BAJARI, P., J. FOX, AND S. RYAN (2007): "Linear Regression Estimation of Discrete Choice Models With Nonparametric Distributions of Random Coefficients," *American Economic Review P&P*, 97 (2), 459–463. [412,426,427]
- BAUER, J. M., AND S. WILDMAN (2012): "The Economics of Usage-Based Pricing in Local Broadband Markets," Working Paper, Michigan State University. [413]
- CHEVALIER, J., AND A. GOOLSBEE (2009): "Are Durable Goods Consumers Forward-Looking? Evidence From College Textbooks," *Quarterly Journal of Economics*, 124 (4), 1853–1884. [415]
- CHUNG, D., T. STEENBURGH, AND K. SUDHIR (2010): "Do Bonuses Enhance Sales Productivity? A Dynamic Structural Analysis of Bonus-Based Compensation Plans," Working Paper 11-041, HBS. [415]
- COPELAND, A., AND C. MONNET (2009): "The Welfare Effects of Incentive Schemes," *Review of Economic Studies*, 76 (1), 93–113. [415]
- CRAWFORD, G., AND M. SHUM (2005): "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 73 (4), 1137–1173. [415]
- DE FONTENAY, A., M. SHUGARD, AND D. SIBLEY (1990): *Telecommunications Demand Modelling: An Integrated View*. Amsterdam: North-Holland. [414]
- DUFLO, E., H. REMA, AND S. P. RYAN (2012): "Incentives Work: Getting Teachers to Come to School," *American Economic Review*, 102 (4), 1241–1278. [415]
- DUTZ, M., J. ORSZAG, AND R. WILLIG (2009): "The Substantial Consumer Benefits of Broadband Connectivity for US Households," Report. [415]
- ECONOMIDES, N., K. SEIM, AND B. VIARD (2008): "Quantifying the Benefits of Entry Into Local Phone Service," *The RAND Journal of Economics*, 39 (3), 699–730. [426]
- EDELL, R., AND P. VARAIYA (1999): "Providing Internet Access: What We Learn From INDEX," *IEEE Network*, 13 (5), 18–25. [414]
- EINAV, L., A. FINKELSTEIN, AND P. SCHRIMPF (2015a): "The Response of Drug Expenditure to Non-Linear Contract Design: Evidence From Medicare Part D," *Quarterly Journal of Economics*, 130 (2), 841–899. [415]

- (2015b): “Bunching at the Kink: Implications for Spending Responses to Health Insurance Contracts,” Working Paper. [415]
- FOX, J., K. I. KIM, S. RYAN, AND P. BAJARI (2011): “A Simple Estimator for the Distribution of Random Coefficients,” *Quantitative Economics*, 2 (3), 381–418. [412,426,427]
- GOETTLER, R., AND K. CLAY (2011): “Tariff Choice With Consumer Learning and Switching Costs,” *Journal of Marketing Research*, 48 (4), 633–652. [426]
- GOOLSBEE, A., AND P. J. KLENOW (2006): “Valuing Products by the Time Spent Using Them: An Application to the Internet,” *American Economic Review P&P*, 96 (2), 108–113. [415]
- GOWRISANKARAN, G., AND M. RYSMAN (2012): “Dynamics of Consumer Demand for Durable Goods,” *Journal of Political Economy*, 120 (6), 1173–1219. [415]
- GREENSTEIN, S., AND R. MCDEVITT (2011): “The Broadband Bonus: Estimating Broadband Internet’s Economic Value,” *Telecommunications Policy*, 35 (7), 617–632. [415]
- GRUBB, M. (2015): “Consumer Inattention and Bill-Shock Regulation,” *Review of Economic Studies*, 82 (1), 219–257. [415]
- GRUBB, M., AND M. OSBORNE (2015): “Cellular Service Demand: Biased Beliefs, Learning, and Bill Shock,” *American Economic Review*, 105 (1), 234–271. [415,426]
- HANDEL, B. (2013): “Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts,” *American Economic Review*, 103 (7), 2643–2682. [426]
- HENDEL, I., AND A. NEVO (2006a): “Measuring the Implications of Sales and Consumer Inventory Behavior,” *Econometrica*, 74 (6), 1637–1673. [415]
- (2006b): “Sales and Consumer Inventory,” *The RAND Journal of Economics*, 37 (3), 543–561. [415]
- HITT, L., AND P. TAMBE (2007): “Broadband Adoption and Content Consumption,” *Information Economics and Policy*, 19 (3–4), 362–378. [415]
- KASAHARA, H., AND K. SHIMOTSU (2009): “Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices,” *Econometrica*, 77 (1), 135–175. [428]
- KETCHAM, J. D., C. LUCARELLI, E. J. MIRAVETE, AND M. C. ROEBUCK (2012): “Sinking, Swimming, or Learning to Swim in Medicare Part D,” *American Economic Review*, 102 (6), 2639–2673. [426]
- KIRJNER, C., AND R. PARAMESWARAN (2013): “Google Fiber—How Much Will Google Spend to Overbuild Kansas City? How Much Would It Cost to Expand?” Bernstein Research. [414,439]
- LAHIRI, S. N. (2003): *Resampling Methods for Dependent Data*. New York: Springer. [428]
- LAMBRECHT, A., K. SEIM, AND B. SKIERA (2007): “Does Uncertainty Matter? Consumer Behavior Under Three-Part Tariffs,” *Marketing Science*, 26 (5), 698–710. [415]
- MACKIE-MASON, J., AND H. VARIAN (1995): “Pricing Congestible Network Resources,” *IEEE Journal on Selected Areas in Communications*, 13, 1141–1149. [413]
- MIRAVETE, E. (2003): “Choosing the Wrong Calling Plan? Ignorance and Learning,” *American Economic Review*, 93 (1), 297–310. [426]
- MISRA, S., AND H. NAIR (2011): “A Structural Model of Sales-Force Compensation Dynamics: Estimation and Field Implementation,” *Quantitative Marketing and Economics*, 9 (3), 211–257. [415]
- NEVO, A., J. L. TURNER, AND J. W. WILLIAMS (2016): “Supplement to ‘Usage-Based Pricing and Demand for Residential Broadband’,” *Econometrica Supplemental Material*, 84, <http://dx.doi.org/10.3982/ECTA11927>. [422]
- ODLYZKO, A., B. S. ARNAUD, E. STALLMAN, AND M. WEINBERG (2012): “Know Your Limits: Considering the Role of Data Caps and Usage Based Billing in Internet Access Service,” Whitepaper, Public Knowledge. [413]
- OPEN INTERNET ADVISORY COMMITTEE (OIAAC) (2013): “Policy Issues in Data Caps and Usage Based Pricing,” OIAC 2013 Annual Report, 14–38. Available at <http://transition.fcc.gov/cgb/oiaac/oiaac-2013-annual-report.pdf>. [413]
- ROSSTON, G., S. SAVAGE, AND D. WALDMAN (2010): “Household Demand for Broadband Internet in 2010,” *The B.E. Journal of Economic and Policy Analysis (Advances)*, 10 (1), Article 79. [415]

VARIAN, H. (2002): "The Demand for Bandwidth: Evidence From the INDEX Experiment," in *Broadband: Should We Regulate High-Speed Internet Access?* ed. by J. Alleman and R. Crandall. Washington, DC: Brookings Institution. [414]

YAO, S., C. MELA, J. CHIANG, AND Y. CHEN (2012): "Determining Consumers' Discount Rates With Field Studies," *Journal of Marketing Research*, 49 (6), 822–841. [415]

Dept. of Economics, Northwestern University, 2001 Sheridan Rd., Evanston, IL 60208, U.S.A. and NBER; nevo@northwestern.edu,

Dept. of Economics, University of Georgia, Brooks Hall, Athens, GA 30602, U.S.A.; jlturner@uga.edu,

and

Dept. of Economics, University of North Carolina–Chapel Hill, Gardner Hall, Chapel Hill, NC 27599, U.S.A.; jonwms@unc.edu.

Co-editor Liran Einav handled this manuscript.

Manuscript received September, 2013; final revision received October, 2015.