Telecommunications Policy 38 (2014) 1035–1045

Contents lists available at ScienceDirect

Telecommunications Policy

URL: www.elsevier.com/locate/telpol

Do three-part tariffs improve efficiency in residential broadband networks?

Jacob B. Malone *, John L. Turner, Jonathan W. Williams

Department of Economics, University of Georgia, United States

ARTICLE INFO

Available online 16 August 2014

Keywords:
Broadband
Demand
Peak
Usage-based pricing

ABSTRACT

We analyze subscriber usage data from an Internet Service Provider that sells service using three-part tariff and unlimited plans. Subscribers facing three-part tariffs have lower average usage than subscribers on unlimited plans, and differences among heavy users explain nearly all the overall difference. Hence, the three-part tariff saves network costs and narrows the gap, between light and heavy users, in price per Gigabyte used. However, subscribers facing three-part tariffs cut usage similarly during peak and off-peak hours. Since off-peak usage adds no network costs, these off-peak usage reductions lower welfare. Differentially pricing peak usage could further enhance efficiency.

1. Introduction

The Internet is now an integral part of modern life. Recent research estimates the average American spends over 3 h online daily.[1] Moreover, with the proliferation of faster access speeds and bandwidth intensive applications like online video, which now comprise over 60% of peak usage, subscribers use more data than ever before.[2] This growth in subscriber demand puts pressure on Internet Service Providers (ISPs) to manage demands placed on their networks.

In response, a sizable number of US providers now sell service via usage-based pricing plans. The most typical "three-part tariff" plan specifies an access fee, a usage allowance and an overage price. Subscribers who use less data than the allowance pay just the access fee for service, while subscribers who use more pay the overage price for each additional GB used.[3]

ISPs typically argue that such plans lower overall and peak usage, helping to reduce network costs, and reduce the level of cross-subsidization between light and heavy users, more closely linking subscriber costs to usage. Government agencies such as the US Federal Communications Commission (FCC) are closely watching to see how usage-based pricing affects costs and efficiency (OIAC, 2013). Yet, there is little empirical evidence to date.

In this paper, we analyze subscriber-specific usage data from a North American ISP during May 2011–May 2013. Importantly, this provider sells service via a menu of three-part tariff plans but also sells service via unlimited plans to a group of grandfathered customers. In comparing usage behavior across groups of subscribers, we identify key effects of three-part tariff pricing. We also analyze how subscriber usage changes across the day and across the month for a single billing cycle.

* Corresponding author.
E-mail addresses: jbmalone@uga.edu (J.B. Malone), jlturner@uga.edu (J.L. Turner), jonwms@uga.edu (J.W. Williams).

[1] See http://www.emarketer.com/Article/Social-Usage-Involves-More-Platforms-More-Often/1010019.

[2] See https://www.sandvine.com/downloads/general/global-internet-phenomena/2013/sandvine-global-internet-phenomena-report-1h-2013.pdf.

[3] Some plans bill for overage more crudely. For example, Comcast's XFINITY plan bills $10 for additional 50 GB blocks of data, and also gives subscribers three warnings before imposing the fees.

http://dx.doi.org/10.1016/j.telpol.2014.06.004
0308-5961/

First, we analyze monthly data. Despite enjoying faster connection speeds, the top ten (one) percent of users on unlimited plans use about 56% (74%) more data in May 2013 than the top ten (one) percent of users facing three-part tariffs. In contrast, the median user who faces a three-part tariff uses just 10% less data than the median user on an unlimited plan. Other than the top few deciles of users, distributions of usage are quite similar across subscriber groups. Thus, allowances and overage prices are mostly effective at reining in heavy users. Since such users account for a disproportionate share of total usage, three-part tariffs save significant network costs for the ISP.

By linking payments to usage, three-part tariffs also reduce the amount of cross-subsidization between the low- and high-volume users. Among subscribers on unlimited plans, the average subscriber in the top ten percent of users consumes 21.5 times more data per month but pays just 1.1 times more than the average subscriber in the bottom half of the usage distribution. Among subscribers facing three-part tariffs, however, the average subscriber in the top ten percent of users consumes just 15 times more data per month than the average bottom-half-usage subscriber, but pays about 2.3 times more for service.

Three-part tariffs do not strongly affect how much subscribers favor downloading (e.g., streaming Netflix) versus uploading (e.g., moving a file from local storage to Dropbox). Subscribers facing three-part tariffs download an average of 90.3% of data per month. While this is statistically smaller than the 90.9% fraction for subscribers to unlimited plans, we do not view a 0.6 percentage point difference as economically significant.

Second, we more closely examine data aggregated up to daily peak hours and daily off-peak hours for a one-month billing cycle during May and June of 2012. We find that subscribers facing three-part tariffs change their usage throughout the month in ways consistent with forward-looking behavior. Subscribers on pace to exceed their allowance cut back on usage in statistically and economically significant ways. For example, subscribers with cumulative usage at between 80% and 100% of the allowance face high implicit overage prices and reduce usage by 28% over the last six days of the billing cycle. In contrast, those who have consumed under 40% of the usage allowance face low implicit overage prices and increase usage by 13% over the last six days of the month.

Most importantly, the percentage effects described above—indeed, all effects of cumulative usage and time of month on current usage—are *statistically the same during peak and off-peak hours*. Intuitively, the way that subscribers cut back does not depend on time of day. Hence, while three-part tariffs do curtail overall usage by heavy users, these users lower their peak usage by the same proportion as they lower their off-peak usage.

Because total off-peak usage is (by definition) below network capacity constraints, any reduction in overall usage that occurs during off-peak hours has no effect on ISP welfare but lowers welfare for both subscribers and content providers (e.g., Netflix), who lose value-creating transactions. This suggests service plans that differentially price peak and off-peak usage may improve upon the three-part tariff schedule's ability to enhance consumer welfare and increase ISP profitability. Such plans would also provide incentives for content providers to make content more portable across the day. For example, since Netflix offers no “live” content, much of this traffic (currently over 30% of peak traffic) could be downloaded during off-peak hours.

We also show that usage is highly persistent, in the sense that a subscriber's usage in one month almost perfectly predicts usage in the next month. In the monthly data, more than 90 percent of subscribers use an amount that places them within two deciles of their position in the prior month's usage distribution. In the intra-month data, over half of the variation in usage is explained by individual-specific fixed effects. Hence, heavy users in one period tend to be heavy users in the next period. This is important, because it makes it relatively easy for ISPs to target such users with three-part tariffs while limiting the impact to a small number of subscribers.

Our study contributes to the policy debate surrounding appropriate network management. Due to uncertainty about how government agencies such as the FCC will enforce net-neutrality, ISPs have been reluctant to negotiate contracts with content providers to pay for the costs of delivering their traffic to end users. However, usage-based pricing of subscriber service plans appears to be (for the moment) a safe harbor from such scrutiny. For example, the 2010 FCC Open Internet Report and Order states:

> “… prohibiting tiered or usage-based pricing and requiring all subscribers to pay the same amount for broadband service, regardless of the performance or usage of the service, would force lighter end users of the network to subsidize heavier end users. It would also foreclose practices that may appropriately align incentives to encourage efficient use of networks. The framework we adopt today does not prevent broadband providers from asking subscribers who use the network less to pay less, and subscribers who use the network more to pay more.” (FCC, 2010, Paragraph 72)

While three-part tariff pricing is common practice by U.S. cellular providers and foreign residential broadband companies, it and other forms of usage-based pricing are controversial in the United States. Numerous consumer groups argue usage-based pricing is unfair and an unnecessary use of market power,[4] and Senator Ron Wyden has proposed legislation restricting it.[5] To inform this debate, the Federal Communications Commission tasked its Open Internet Advisory Committee to study the economics of usage-based pricing. The committee's recent report (OIAC, 2013) includes information on the proliferation of usage-based pricing and sets up a useful framework for thinking about the efficiency of usage-based

[4] See, e.g., StopTheCap.com.

[5] Sen. Wyden's bill, the so-called Data Cap Integrity Bill, can be found at http://www.wyden.senate.gov/news/press-releases/wyden-data-cap-legislation-will-protect-consumers-and-promote-innovation.

pricing and its potential effects on access. However, the report has limited access to usage data and leaves most questions unanswered.

Beyond policy insights, our analysis contributes to multiple literatures. The first studies access to and demand for broadband usage. A number of theoretical studies consider usage-based pricing (Bauer & Wildman, 2012; Mackie-Mason & Varian, 1995, Odlyzko, Stallman, & Weinberg, 2012), primarily to focus on its welfare implications. On the empirical side, Hitt and Tambe (2007) show that broadband adoption significantly increases usage, while several papers estimate the economic value of broadband internet (Dutz, Orszag, & Willig, 2009; Greenstein & McDevitt, 2011; Rosston, Savage, & Waldman, 2010). A small number of papers use highly detailed and disaggregated usage data. Edell and Varaiya (2002) and Varian (2002) use data from the INDEX experiments, where subjects were offered different prices for different (and very slow, by today's standards) broadband speeds, to estimate demand for broadband and opportunity costs of time. Goolsbee and Peter (2006) use time-use data and hourly wages to estimate the value of broadband usage. More recently, Lambrecht, Seim, and Skiera (2007) and Nevo, Turner, and Williams (2013) study choices made by subscribers facing three-part tariffs. Lambrecht et al. (2007) use monthly usage data to study how uncertainty influences choices of three-part tariff plans, while Nevo et al. (2013) exploit higher frequency data to estimate demand for broadband services.

A separate literature examines whether consumers are forward-looking and able to make rational decisions when faced with complicated pricing schedules. These studies span a wide range of topics from telecommunications (Grubb & Osborne, 2012; Miravete, 2003), to stockpiling behavior in response to sales (Hendel & Nevo, 2006), to optimally using health insurance (Aron-Dine, Einav, Finkelstein, & Cullen, 2012; Handel, 2013).

The remainder of the paper is as follows. Section 2 describes the data in detail and presents statistics on plan features. Section 3 then provides descriptive statistics on subscriber usage at the monthly and hourly levels. We test for how subscribers respond to changes in implicit prices of usage, and for how such responses depend on peak or off-peak hours, in Section 4. Section 5 concludes.

2. Data

The data used in this paper are from a North American ISP and contain information on subscriber usage from May 2011 to May 2013.[6] All usage data originate from 15-minute level Internet Protocol Detail Records (IPDR), a trusted data source used for subscriber billing. We analyze data at a 15-minute frequency for May 12, 2012 to June 30, 2012. Outside of this window, we have access to monthly aggregates of a subscriber's downstream and upstream usage. We focus on this ISP's largest markets to ensure the sample is representative of other North American communities.

IPDR records are collected by a Cable Modem Termination System (CMTS), which converts Internet traffic to a coaxial signal for home delivery. Fig. 1 illustrates where a CMTS device is located on the network. The CMTS is marked by an oval.

Our data also include variables on subscribers and plan choices. Each subscriber has a cable modem with a unique Media Access Control (MAC) address. Monthly plan-specific variables include a subscriber's base price, usage allowance, and overage fees. We merge plan-specific variables with the IDPR data using the MAC address.

This ISP notably serves subscribers on both three-part tariff and unlimited usage plans. These unlimited plans are legacy options; new subscribers are offered just three-part tariff plans. We refer to subscribers remaining on unlimited plans as being grandfathered.

Table 1 presents statistics on plan details. From May 2011 to May 2013, the download speeds on three-part tariff plans almost double to 15.1 Mb/s, whereas speeds for unlimited plans are steady at 6.5 Mb/s. With these faster speeds come a 21% increase in the fixed access price for three-part tariff subscribers. Unlimited subscribers experience no such increase in prices. The number of unlimited subscribers drops almost 60% while the number of three-part tariff subscribers grows by 42%, indicating a strong preference for speed.

During the first year, the average usage allowance doubles in size to 91.6 GB/month. As a result, the percentage of people exceeding the usage allowance falls by 12% to 8.3%. Increasing the usage allowances also improves matching between subscribers and plans. Let a plan be “dominated” if there is another plan such that the subscriber would have paid less, given their usage, and the advertised download speed is just as great. In May 2011, 26.4% of three-part tariff subscribers are on a dominated plan, but this falls to just 4.6% in May 2012, a drop of about 80%.

These statistics reveal that the average three-part tariff subscriber rarely exceeds the usage allowance. While this is partly due to higher usage allowances, subscribers are also better informed. This ISP now offers a notification system that alerts subscribers once common usage thresholds have been surpassed.[7] For the entire panel, around 15% of subscriber-month observations exceed the usage allowance which translates to around 3.5 occurrences per subscriber over two years.

Subscribers who exceed the usage allowance nearly always use a consistently high amount over the entirety of the month. That is, overage charges rarely result from a small number of high-usage days. For subscribers who exceed the allowance in May 2012, only 0.3% of observations are for days where the subscriber uses more than 50% of the usage allowance in one day. This percentage increases to just 11.5% when 10% of the usage allowance is used as the cutoff. These results are further proof that subscribers are knowledgeable of the implications on overages by various Internet activities.

[6] Due to data corruption issues, there are no data for March and April 2012.
[7] This is now mandated for the US cellular industry.

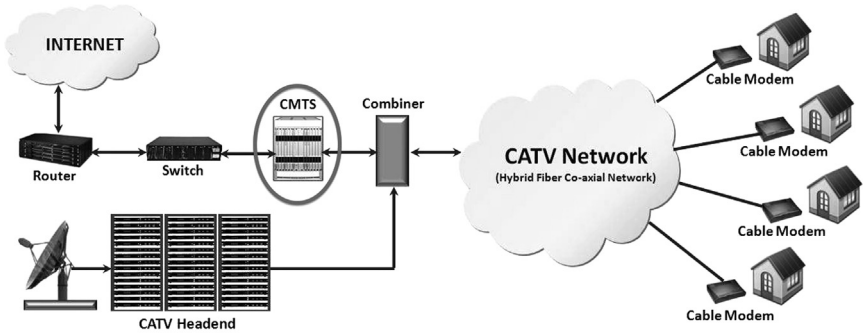

Fig. 1. CMTS relationship to subscriber. *Note:* This figure illustrates a typical cable-broadband network. It comes from http://www.commverge.com/Solutions/BroadbandAccessNetworks/CableCMTS/tabid/165/Default.aspx.

Table 1
Summary statistics of plan details.

	May 2011	May 2012	May 2013
Three-part tariff			
Fixed access fee ($)	61.5	73.8	77.9
Per GB overage fee ($)	3.7	3.1	3.0
Down speed (Mb/s)	8.9	13.7	15.1
Allowance size (GB)	44.1	91.6	104.3
Over allowance (%)	20.5	8.3	8.2
Allowance used (%)	85.1	42.3	44.6
On dominated plan (%)	26.4	4.6	4.4
Subscribers	48,894	59,550	69,600
Unlimited			
Fixed access fee ($)	44.9	44.7	44.4
Per GB overage fee ($)	0.0	0.0	0.0
Down speed (Mb/s)	6.5	6.5	6.4
Allowance size (GB)	∞	∞	∞
On dominated plan (%)	0.0	0.0	0.0
Subscribers	28,075	17,426	11,761

Note: These statistics reflect plan characteristics and usage by subscribers to a single ISP during May 2011–May 2013. Across plans, download speed is non-decreasing in the access fee and the overage price is non-increasing in the access fee. Usage is based upon Internet Protocol Detail Record (IPDR) data, captured in 15-minute intervals and aggregated to the monthly level. We say a plan is dominated if the subscriber could have chosen another plan and (holding usage constant) paid less and enjoyed advertised download speed no lower. Except for the count of subscribers, all reported values are averages at the subscriber level.

It is clear from Table 1 that there is significant movement between plans during 2011–2013. Much of this movement is to higher-speed, higher-allowance plans. Fig. 2 highlights how this reflects forward-looking behavior by subscribers. The Figure shows median growth rates and the percentage of subscribers who upgrade to plans with larger usage allowances, conditional on the portion of the usage allowance the subscriber consumed in May 2011 or 2012. Subscribers near the usage allowance experience the slowest growth and also show the highest rates of plan switching in the following year, which suggests subscribers are forward-looking and can effectively manage usage. Both ways in which these subscribers respond to the possibility of overage fees are beneficial to the ISP. Slower growth rates save costs, and plan switching generates more revenue.

3. Monthly usage

Differences among subscribers at the top of the usage distribution account for nearly all of the differences in average usage in May 2013. The cumulative distribution functions shown in Fig. 3 are nearly identical at the median but diverge sharply for the top few deciles of usage. Table 2 shows that this divergence is statistically significant. A Kolmogorov–Smirnov (K–S) test, where the null hypothesis is that the distributions are equal, returns a highly significant test statistic of 0.074.[8] By reining in these extreme subscribers, three-part tariffs also reduce the level of inequality. The Gini coefficient for usage among three-part tariff subscribers, 0.548, is nearly 9% lower than the Gini coefficient for usage among unlimited subscribers (0.599).[9]

[8] This statistic is well above the 1% critical value for this test, 0.006 (Massey, 1951). Results are qualitatively identical for May 2011 and May 2012.
[9] These Gini coefficients are calculated using May 2013 data. For comparison, the Gini coefficient for U.S. income in 2011 is 0.48.

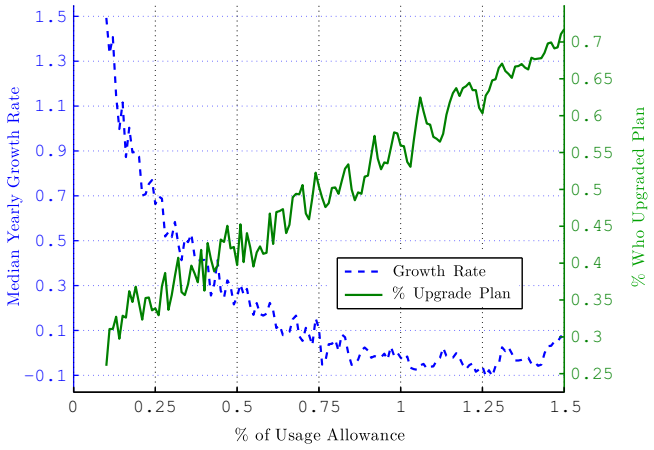

Fig. 2. Yearly growth and plan upgrade incidence. *Note:* These statistics reflect usage and plan choices by subscribers to a single ISP during May 2011–May 2013. Usage is based upon IPDR data, captured in 15-minute intervals and aggregated to the monthly level.

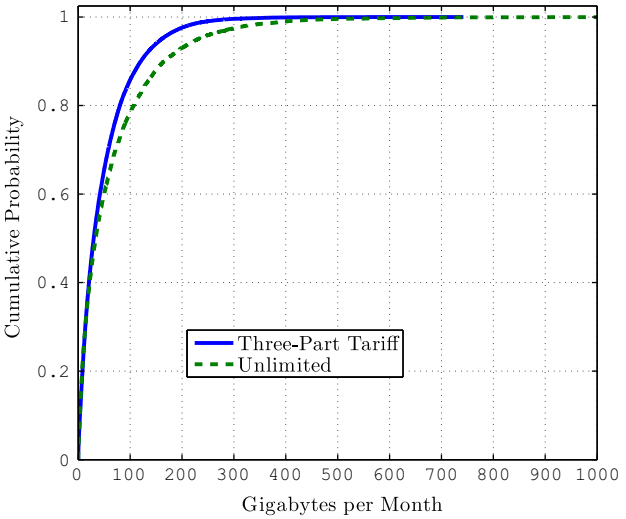

Fig. 3. Usage distributions. *Note:* These cumulative distribution functions reflect usage by subscribers to a single ISP during May 2013. Usage is based upon IPDR data, captured in 15-minute intervals and aggregated to the monthly level.

Despite these differences in the usage distributions, the composition of monthly usage is very similar among subscribers. Although a two-sample *t*-test rejects the null hypothesis that the downstream percentage of traffic is the same for three-part tariff and unlimited subscribers, the difference is not economically significant. For example, in May 2013, the unlimited subscribers' average download percentage (90.9%) is less than one percentage point greater than for subscribers facing a three-part tariff (90.3%). This relationship is consistent across the entire panel. It is not surprising that downstream content dominates, because popular downstream applications (like video streaming) are bandwidth intensive.

The three-part tariff is also effective at reducing usage by the heaviest subscribers given the growth rates observed in Table 2. While the annualized growth rate for the 99th percentile is 22.8%, usage grows by just 3% between May 2012 and May 2013. This is significantly less than growth in usage by the unlimited subscribers' (34%). The smaller growth rates at the top of the distribution suggest that three-part tariffs are pushing low-valued content off the network.

The growth rates of the bottom half of three-part tariff subscribers are notably higher than any others. However, these subscribers consume substantially less data per month, so the lower growth rates for the top percentile subscribers still imply much higher yearly increases in GBs because of the higher initial usage. Therefore, these high-usage subscribers are still the primary drivers of network costs.

The differences in growth rates between three-part tariff subscribers and unlimited subscribers are also sizable. There are a couple of explanations for this. First, three-part tariff subscribers enjoy speeds (15.1 Mb/s) more than twice as fast as unlimited subscribers (6.4 Mb/s), which makes it easier to access content. Second, many applications are bandwidth-adaptive. That is,

Table 2
Summary statistics of monthly usage.

	May 2011	May 2012	May 2013	Annualized growth rate
Three-part tariff				
Mean	21.4	40.1	48.9	0.511
p25	2.8	7.1	10.7	0.955
p50	8.9	20.7	29.8	0.830
p75	25.0	52.4	68.0	0.649
p90	55.7	103.1	120.7	0.472
p95	83.9	145.7	160.6	0.384
p99	164.8	241.1	248.6	0.228
Usage by downloading (%)	88.7	89.9	90.3	
Subscribers	48,894	59,550	69,600	
Unlimited				
Mean	32.8	46.6	65.4	0.412
p25	4.7	7.3	10.1	0.466
p50	14.6	22.7	33.4	0.513
p75	39.3	60.6	87.1	0.489
p90	83.1	119.9	165.2	0.410
p95	127.0	173.1	230.8	0.348
p99	243.8	299.0	401.0	0.282
Usage by downloading (%)	89.7	90.8	90.9	
Subscribers	28,075	17,426	11,761	
K–S test statistic	0.123**	0.037**	0.074**	
Download % DIM *t*-statistic	11.3**	8.9**	5.4**	

Note: These statistics reflect usage by subscribers to a single ISP during May 2011–May 2013. Usage is based upon IPDR data, captured in 15-minute intervals and aggregated to the monthly level. We denote percentile with "p," hence the numbers next to p25 denote usage by the subscriber whose usage is at the 25th percentile. All usage statistics are in gigabytes. K–S ≡ Kolmogorov–Smirnov and DIM ≡ difference in means.
∗ denotes significance at the 0.05 level.
** denotes significance at the 0.01 level.

even if the subscriber does nothing behaviorally that would alter their monthly usage, the application will adapt on its own to use whatever bandwidth is available.[10]

Because heavy users generate the majority of network costs, ISPs would ideally generate more revenue from them accordingly. This is possible if usage is persistent. If usage is persistent, only heavy usage users will be affected by the three-part tariff and the amount of cross-subsidization across users will decrease.

To study persistence, we group subscribers by decile of usage for each month. We then create a transition matrix of movement between deciles and plot it in Fig. 4. The ridge in the figure shows a subscriber's current decile is his most likely decile in the subsequent month and that the vast majority of transitions are to a "local" region of the distribution. If a ± 2 decile window is used, for example, more than 90% of subscriber transitions are captured.[11] The saddle shape also implies that subscribers are least likely to transition to a new decile when they start in a more extreme decile. In fact, we find that the tenth decile is the most persistent of all with users remaining there about 65% of the time.

Unsurprisingly, three-part tariffs reduce the amount of subsidization occurring between the bottom and the top of the distribution. Table 3 shows the distribution of relative prices paid per GB. In May 2013, the average subscriber in the top 10% consumed 15 times more data than the average subscriber in the bottom 50%, but paid only 2.3 times more. For unlimited subscribers the results are more extreme. The average subscriber in the top 10% used 21.5 times more and paid 1.1 times more. From this perspective, three-part tariffs are effective at reducing the burden of subsidization on the low-usage subscribers by generating more revenue from high-usage subscribers.

Interestingly, the per-GB cost to a subscriber falls substantially across time. For example, the median user on a three-part tariff plan pays 61.5% less in May 2013 relative to May 2011. This pattern is consistent across all quantiles of the usage distribution for subscribers on both types of plans.

4. Intra-day and intra-month usage

In the three-part tariff plans offered by our ISP, allowances and overages impose implicit prices on users that vary depending upon cumulative usage during that month and on the remaining number of days in the month. Ideally, such prices reduce usage in a way that efficiently saves on network costs. However, allowances that apply to *total* usage only, such

[10] Note also that content providers actively seek to accommodate bandwidth constraints driven by usage allowances. For example, in response to relatively low allowances by some Canadian ISPs, Netflix lowered the default bitrate limit of video from 4800 kb/s to 625 kb/s to help users stay under their allowances (OIAC, 2013, p. 10).

[11] Fig. 4 includes all subscribers, but if the same analysis is performed separately for subscribers on unlimited and three-part tariff plans, the results do not change.

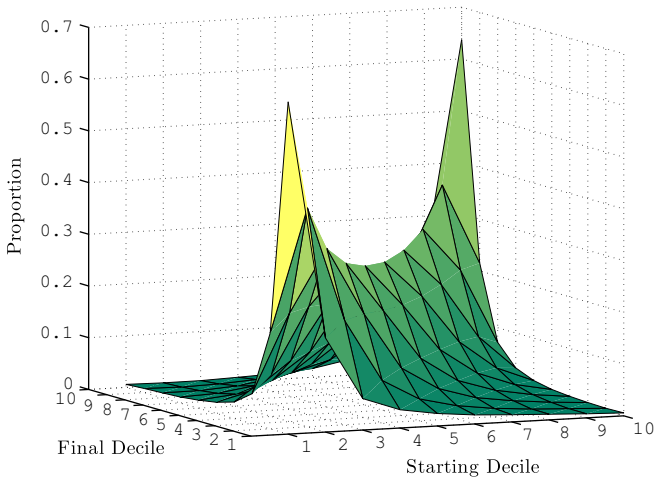

Fig. 4. Usage transition surface. *Note:* The level of the surface is the likelihood a subscriber's usage is in a particular (final) decile of the usage distribution during a particular month, conditional on which (starting) decile that subscriber's usage was in during the previous month. These statistics reflect usage by subscribers to a single ISP during May 2011–May 2013. Usage is based upon IPDR data, captured in 15-minute intervals and aggregated to the monthly level.

Table 3
Summary statistics of monthly revenue.

	May 2011	May 2012	May 2013
Three-part tariff			
Mean revenue ($)	83.79	81.44	84.74
p75 ($/GB)	18.56	8.82	6.21
p50 ($/GB)	6.39	3.34	2.46
p25 ($/GB)	4.27	1.57	1.28
p10 ($/GB)	2.00	1.00	0.86
p5 ($/GB)	1.33	0.79	0.74
p1 ($/GB)	0.82	0.65	0.64
Subscribers	48,894	59,550	69,600
Unlimited			
Mean revenue ($)	44.91	44.65	44.39
p75 ($/GB)	9.21	5.93	4.24
p50 ($/GB)	3.01	1.92	1.29
p25 ($/GB)	1.12	0.73	0.49
p10 ($/GB)	0.53	0.36	0.26
p5 ($/GB)	0.35	0.26	0.19
p1 ($/GB)	0.19	0.15	0.11
Subscribers	28,075	17,426	11,761

Note: These statistics reflect usage and expenditures by subscribers to a single ISP during May 2013. Usage is based upon IPDR data, captured in 15-minute intervals and aggregated to the monthly level. Expenditures per GB reflect subscriber plan choices for that month. We denote percentile with "p," hence the numbers next to p25 denote the subscriber whose $/GB is at the 25th percentile.

as those used by our ISP, provide no direct incentive by users to reduce peak usage (OIAC, 2013, p. 16) and may inadvertently reduce off-peak usage as well. Because peak usage alone drives network costs, it is important to know whether such allowances lower peak usage by more than off-peak usage.

To study this issue, we consider usage data for a complete billing cycle during May–June 2012. We start by discussing the data we use to test how three-part tariff allowances affect peak and off-peak usage. We then present direct tests.

Fig. 5 plots usage for each hour of the day for the average three-part tariff and unlimited subscriber. We report statistics in Megabits per second (Mb/s) to provide perspective on how demanding usage is relative to the maximum potential usage at provisioned speeds. Just as in the monthly data, the average unlimited subscriber consumes more across the entire day. Notice that the pattern of usage observed for both three-part tariff and unlimited subscribers is very similar: early morning hours are the slowest, while the late evening's are the busiest.

With average usage in Mb/s, Fig. 5 shows that average usage demands are only a small fraction of a subscriber's potential (advertised) bandwidth. For example, during the 10 PM hour, the busiest of the day, the average three-part tariff subscriber uses around 0.25 Mb/s, or roughly 0.12 GB per hour. This level of utilization represents around 2% of the maximum possible

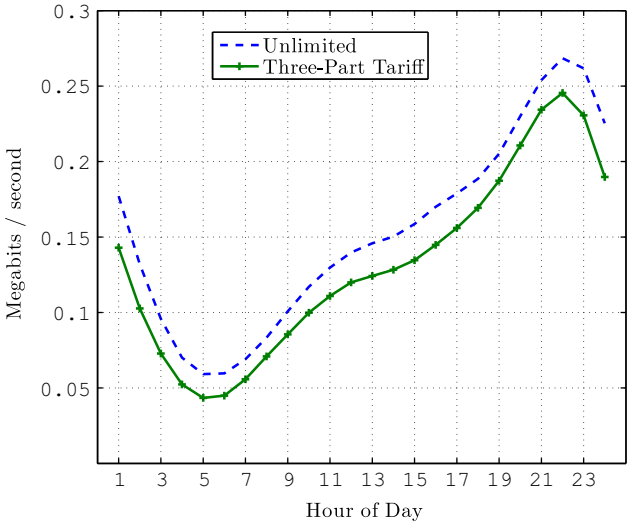

Fig. 5. Average intra-day subscriber bandwidth usage. *Note:* These statistics reflect usage by subscribers to a single ISP during a single billing cycle during May–June 2012. Usage is based upon IPDR data, captured in 15-minute intervals and aggregated to the hourly level.

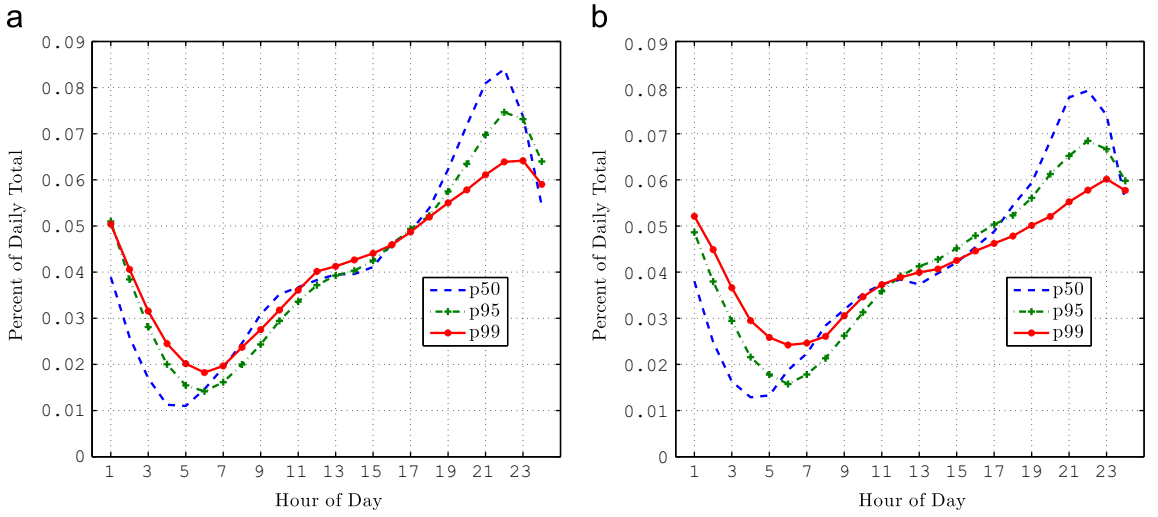

Fig. 6. Hourly usage as a percent of daily total. (a) Three-part tariff. (b) Unlimited. *Note:* These statistics reflect usage by subscribers to a single ISP during a single billing cycle during May–June 2012. Usage is based upon IPDR data, captured in 15-minute intervals and aggregated to the monthly level. We denote percentile with "p," hence the numbers next to p50 denote the subscriber whose usage is at the 50th percentile.

given the average provisioned connection speed. The low level of utilization is not surprising as many online activities require the subscriber to take time to consume the content even if it arrives very quickly.

Fig. 6(a) and (b) presents the proportion of a subscriber's daily usage at each hour of the day for subscribers facing a three-part tariff and for unlimited subscribers, respectively. The figures present usage for different quantiles of the distribution (50th, 95th and 99th percentile users). Patterns of usage across the day are nearly identical for three-part tariff and unlimited subscribers. This suggests that those subscribers facing three-part tariffs reduce usage in a proportional manner across the day, since the proportion of daily usage at each hour of the day is the same, but the levels are lower for those facing three-part tariffs as shown in Fig. 5.

Notably, the heaviest-using subscribers have slightly flatter profiles across the day. This is consistent with such subscribers using more file-sharing applications, like BitTorrent, that generate traffic even if the subscriber is absent. This pattern holds for subscribers on both types (three-part tariff and unlimited) of plans.

To test directly for how the three-part affects peak and off-peak usage, we modify and extend the framework used by Nevo et al. (2013, Section 2.3) to demonstrate forward-looking behavior by subscribers. Intuitively, subscribers face a dynamic problem within a given billing cycle, regarding how to allocate usage across that billing cycle. A subscriber well below the pace to exceed the usage allowance before the end of the billing cycle has a low probability of exceeding the

allowance and consequently a low implicit price of usage. Conversely, a subscriber on pace to exceed the usage allowance has a higher implicit price of usage, but one that is less than the overage price because there is some chance the subscriber's usage will not reach the allowance. Finally, subscribers that have exceeded the usage allowance at any point in the billing cycle face a constant price, the overage fee.

Hence, when a subscriber's cumulative usage increases relative to the pace to exceed the usage allowance, then the implicit price of usage increases.[12] If subscribers reduce usage in a way that lowers the chance of paying overage fees when the probability of overage fees is high, it suggests that they are forward-looking and understand how to manage usage across a billing cycle.

In contrast to Nevo et al. (2013), we disaggregate usage into peak (6:00 PM to 11:59 PM) and off-peak hours (otherwise). We then modify the model to capture "peak effects" of implicit prices on usage. If peak effects are non-zero, then subscribers have different sensitivity to implicit prices during these hours. If peak effects are zero, then subscribers respond the same during peak and off-peak periods. If the latter holds, then three-part tariffs do not efficiently target peak usage.

We estimate the following regression model:

$$\ln(c_{itp}) = \alpha_i + \alpha_i^{peak}\mathbf{1}\{p=1\} + [\psi + \psi^{peak}\mathbf{1}\{p=1\}]\mathbf{x}_{it} + \sum_{m=1}^{M=4}\sum_{n=1}^{N=5}[\beta_{nm} + \beta_{nm}^{peak}\mathbf{1}\{p=1\}]\mathbf{z}_{itnm} + \epsilon_{itp}. \quad (1)$$

The dependent variable, $\ln(c_{itp})$, is the log transformation of subscriber i's peak or off-peak usage at day t in the billing cycle, where $p=1$ indicates peak usage.[13] Subscriber fixed effects are denoted by α_i and α_i^{peak}; the latter is the "peak" fixed effect. The vector of controls, $\mathbf{x}_{it}$, includes day-of-week dummy variables and a billing cycle time trend,[14] and ψ^{peak} captures the peak effect of controls on usage.

The terms inside the double summation capture a series of interactions that describe a subscriber's state, such that

$$\mathbf{z}_{itnm} = \mathbf{1}[pct_n \le \tilde{c}_{i,t-1} < pct_{n+1}]\mathbf{1}[day_m \le t < day_{m+1}],$$

where $\tilde{c}_{i,t-1}$ is the proportion of the usage allowance consumed through $t-1$ days of the billing cycle. The thresholds for the usage allowance indicators are $pct_1 = 0$, $pct_2 = 0.40$, $pct_3 = 0.60$, $pct_4 = 0.80$, $pct_5 = 1$, and $pct_6 = \infty$, while those for the billing day are $day_1 = 10$, $day_2 = 15$, $day_3 = 20$, $day_4 = 25$, and $day_5 = 31$.

We also include a peak indicator here, and interact it with each of the variables. This permits a direct test of whether the response to variation in the possibility of overages is different during peak hours. If $\beta_{nm}^{peak} = 0$ for each ordered pair, (n, m), then we can conclude that subscribers respond similar to the possibility of overages during each part of the day.

Table 4 reports, for each (n,m) ordered pair, the coefficient estimate for the usage allowance interaction in Eq. (1), β_{nm} and for the peak effect, β_{nm}^{peak}. First, note that none of the peak-effect estimates are statistically or economically significant. This implies that subscribers respond to expected overages by changing behavior uniformly across the day, consistent with the descriptive evidence in Fig. 6. Thus, the discussion of the impact of three-part tariffs on usage during all hours can focus on the top-half of Table 4.

We find a sharp reduction in usage by subscribers with a high probability of exceeding the usage allowance. Subscribers with between 80% and 100% of the usage allowance reduce usage by about 28% over the last six days of the billing cycle.[15] This effect is slightly stronger for those who have already exceeded the usage allowance (33%). Second, subscribers who have consumed only a small proportion of their usage allowance near the end of their billing cycle actually accelerate usage with the assurance that overages are very unlikely. For example, subscribers who have consumed under 40% of the usage allowance increase usage by about 13% over the last six days of the month. Collectively, the results from Table 4 show that subscribers are quite sophisticated in their ability to respond to within-month variation in the possibility of overages, and that the reduction in usage is quite large on average and similar during peak and off-peak hours.

At the bottom of Table 4, we report the proportion of variation explained by the subscriber fixed effect and its interaction with the peak indicator, α_i and α_i^{peak}, respectively. Consistent with our earlier discussion, the fixed effects explain about 55% of the total variation in usage, despite the substantial subscriber-specific variation in day-to-day online activities.

By pricing usage similarly at all times of the day, a three-part tariff saves costs by reducing low-value traffic during peak hours but eliminates costless welfare-enhancing usage during off-peak hours. This latter effect harms subscribers and content providers, who lose transactions, without helping the ISP. This suggests that usage-based pricing that specifically incorporates time of day may be more effective at improving overall welfare. The magnitude of any improvement would depend critically on how effectively subscribers can redistribute traffic to off-peak hours (i.e., elasticity across the day). We view this as a crucial topic for future research and strongly encourage experiments that permit clean identification of such elasticities.

[12] Notably, our ISP offers a notification system to subscribers regarding the proportion of their allowance that has been consumed up until each point in the billing cycle. Hence, forward-looking behavior is feasible.

[13] Only subscribers for which a complete billing cycle is observed are used in estimation.

[14] This time trend is identified because subscribers start the billing cycle on different calendar days.

[15] Note that we interpret these coefficient estimates, due to the log dependent variable, according to Halvorsen and Raymon (1980). Specifically, $100*(e^{-0.325}-1) = -27.74$.

Table 4
Main regression estimates.

	$10 \le t < 15$	$15 \le t < 20$	$20 \le t < 25$	$25 \le t < 31$
Interactions (β_{nm})				
$0 \le \tilde{c}_{i,t-1} < 0.40$	−0.041** (0.008)	−0.026** (0.009)	0.039** (0.010)	0.123** (0.010)
$0.40 \le \tilde{c}_{i,t-1} < 0.60$	0.002 (0.023)	−0.076** (0.020)	−0.084** (0.018)	−0.001 (0.017)
$0.60 \le \tilde{c}_{i,t-1} < 0.80$	−0.008 (0.046)	−0.156** (0.031)	−0.160** (0.024)	−0.080** (0.021)
$0.80 \le \tilde{c}_{i,t-1} < 1$	−0.181* (0.078)	−0.257** (0.048)	−0.325** (0.035)	−0.371** (0.028)
$1 \le \tilde{c}_{i,t-1}$	−0.079 (0.080)	−0.324** (0.061)	−0.406** (0.042)	−0.460** (0.033)
Peak effects (β_{nm}^{peak})				
$0 \le \tilde{c}_{i,t-1} < 0.40$	−0.015 (0.012)	−0.003 (0.014)	0.023 (0.016)	−0.017 (0.016)
$0.40 \le \tilde{c}_{i,t-1} < 0.60$	0.018 (0.039)	−0.025 (0.031)	0.039 (0.029)	0.007 (0.026)
$0.60 \le \tilde{c}_{i,t-1} < 0.80$	−0.021 (0.074)	0.055 (0.049)	0.050 (0.038)	−0.031 (0.034)
$0.80 \le \tilde{c}_{i,t-1} < 1$	0.006 (0.117)	−0.004 (0.077)	−0.005 (0.055)	−0.007 (0.045)
$1 \le \tilde{c}_{i,t-1}$	−0.003 (0.130)	−0.019 (0.095)	−0.001 (0.067)	−0.057 (0.053)
Proportion of variation Explained by α_i and α_i^{peak}	0.550			

Note: This table reflects estimates of selected parameters from equation (1). The rows reflect the proportion of a subscriber's allowance consumed prior to the given period, while the columns reflect the day of the month. We use data from a complete billing cycle during May–June 2012. Usage is based upon IPDR data, captured in 15-minute intervals and aggregated to peak (6:00 PM–11:59 PM) and off-peak (12:00 AM – 5:59 PM) categories. Robust standard errors are in parentheses.

* denotes significance at the 0.05 level.

** denotes significance at the 0.01 level.

5. Conclusion

We analyze usage data from a North American ISP to study the effects of three-part tariffs, currently the most popular form of usage-based pricing. We find that three-part tariffs are effective at reducing overall usage, generating more revenue from high-usage subscribers and reducing inequality in usage and fees per GB. We also find subscriber usage to be highly persistent, which means that usage-based pricing plans such as three-part tariffs can effectively target high-usage subscribers.

Yet, three-part tariffs have minimal effects on the timing of usage. Thus, such plans reduce off-peak usage, causing pure deadweight losses. This suggests that ISPs may be able to further enhance welfare by differentially pricing peak and off-peak usage. Notably, at least one ISP (Exede) has experimented with time-sensitive caps (OIAC, 2013, p. 16, footnote 23).[16] It will be interesting to see whether additional ISPs follow suit.

Our results leave many important related questions unanswered. For example, we do not identify granular details of what applications subscribers use. While it seems intuitive that three-part tariffs would tend to remove more high-bandwidth usage like video streaming, further research will be necessary to quantify these effects.

In addition, when ISPs upgrade networks they avoid (or at least mitigate) network congestion, but it is not well known exactly what costs they are avoiding. To quantify completely the welfare effects of reducing network upgrades, it would be useful to estimate costs from congestion. A key challenge is to find data where subscribers face usage prices and congestion simultaneously.

Finally, broadband access is typically bundled with television service from many large cable companies. Given recent moves toward consolidation in the cable industry, such as the recently proposed Comcast-Time Warner merger, it appears that cable companies are under pressure to improve their competitive positions. Understanding the welfare implications of such changes, and how they might affect decisions to use usage-based pricing of broadband, are unexplored and important to understand. We look forward to further progress in the area.

Acknowledgments

We thank the National Science Foundation (Grant SES-1324717) and Cable Television Laboratories, Inc., and its member organizations, for their support of this work. We are grateful for the generosity of the cable operators that provided the data used in this paper. This paper has benefited from discussions with Ike Elliott, Shane Greenstein, Ron Reuss, Scott Savage,

[16] http://www.dslreports.com/shownews/Exede-Caps-Lifted-For-Overnight-Use-120776.

Terry Shaw, and Steven Wildman. Jim Metcalf and Craig Westwood provided expert IT support for this project. All remaining errors are our own.

References

Aron-Dine, A., Einav, L., Finkelstein, A., & Cullen, M. (2012). *Moral hazard in health insurance: How important is forward looking behavior?* NBER working paper no. 17802.
Bauer, J., & Wildman, S. (2012). *The economics of usage-based pricing in local broadband markets*. Michigan State University working paper.
Dutz, M., Orszag, J., & Willig, R. (2009). *The substantial consumer benefits of broadband connectivity for US households*. Internet intervention alliance working paper.
Edell, R., & Varaiya, P. (2002). Providing internet access: What we learn from INDEX. *Broadband: Should we regulate high-speed internet access?* Brookings Institution.
FCC, Federal Communications Commission (2010). *Open internet report and order*.
Goolsbee, A., & Peter, K. (2006). Valuing products by the time spent using them: An application to the internet. *American Economic Review P&P*, *96*(2), 108–113.
Greenstein, S., & McDevitt, R. (2011). The broadband bonus: Estimating broadband internet's economic value. *Telecommunications Policy*, *35*(7), 617–632.
Grubb, M., & Osborne, M. (2012). *Cellular service demand: Biased beliefs, learning, and bill shock*. MIT working paper.
Halvorsen, R., & Palmquist, R. (1980). The interpretation of dummy variables in semilogarithmic equations. *American Economic Review*, *70*(3), 474–475.
Handel, B. (2013). Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review*, *103*(7), 2643–2682.
Hendel, I., & Nevo, A. (2006). Measuring the implications of sales and consumer inventory. *Econometrica*, *74*(6), 1637–1673.
Hitt, L., & Tambe, P. (2007). Broadband adoption and content consumption. *Information Economics and Policy*, *19*(3–4), 362–378.
Lambrecht, A., Seim, K., & Skiera, B. (2007). Does uncertainty matter? Consumer behavior under three-part tariffs. *Marketing Science*, *26*(5), 698–710.
Mackie-Mason, J., & Varian, H. (1995). Pricing congestible network resources. *IEEE Journal on Selected Areas in Communications*, *13*(1141–1149).
Massey, F. (1951). The Kolmogorov–Smirnov test for goodness of fit. *Journal of the American Statistical Association*, *46*(253), 68–78.
Miravete, E. (2003). Choosing the wrong calling plan? Ignorance and learning. *American Economic Review*, *93*(1), 297–310.
Nevo, A., Turner, J., & Williams, J. (2013). *Usage-based pricing of broadband Internet service*. Working paper.
Odlyzko, A., St. Arnaud, B., Stallman, E., & Weinberg, M. (2012). *Know your limits: Considering the role of data caps and usage-based billing in internet access service*. Public knowledge whitepaper.
OIAC, Open Internet Advisory Committee (2013). Policy issues in data caps and usage-based pricing. *OIAC annual report*.
Rosston, G., Savage, S., & Waldman, D. (2010). Household demand for broadband internet in 2010. *The B.E. Journal of Economics and Policy Analysis*, *10*(1).
Varian, H. (2002). The demand for bandwidth: Evidence from the INDEX experiment. *Broadband: Should we regulate high-speed internet access*? Brookings Institution.