

A Case Study of Traffic Demand Response to Broadband Service-Plan Upgrades

Sarthak Grover, Roya Ensafi, and Nick Feamster

Department of Computer Science, Princeton University.
{sgrover, rensafi, feamster}@cs.princeton.edu

Abstract. Internet service providers are facing mounting pressure from regulatory agencies to increase the speed of their service offerings to consumers; some are beginning to deploy gigabit-per-second speeds in certain markets, as well. The race to deploy increasingly faster speeds begs the question of whether users are exhausting the capacity that is already available. Previous work has shown that users who are already maximizing their usage on a given access link will continue to do so when they are migrated to a higher service tier.

In a unique controlled experiment involving thousands of Comcast subscribers in the same city, we analyzed usage patterns of two groups: a control group (105 Mbps) and a randomly selected treatment group that was upgraded to 250 Mbps without their knowledge. We study how users who are already on service plans with high downstream throughput respond when they are upgraded to a higher service tier without their knowledge, as compared to a similar control group. To our surprise, the difference between traffic demands between both groups is higher for subscribers with moderate traffic demands, as compared to high-volume subscribers. We speculate that even though these users may not take advantage of the full available capacity, the service-tier increase generally improves performance, which causes them to use the Internet more than they otherwise would have.

1 Introduction

With the large impact of broadband Internet on our daily lives and its rapid increase in bandwidth-intensive services, policymakers and service providers (ISPs) are trying to determine how much bandwidth consumers need. With the proliferation of high quality video content, and the recent boom in Internet-enabled consumer device, it is worth studying—and continually re-evaluating—whether (and how) users consume the capacity that ISPs offer. Up to a certain point, users will exhaust available capacity, and they will also adapt when more capacity becomes available; this increased demand in turn drives provisioning. Above certain speeds, however, the typical user no longer exhausts the available capacity. At what speed does this inflection point occur? How do users adapt their demands when an ISP offers faster speed tiers? Answers to these questions will ultimately help inform policymakers and ISPs determine how to make investments in infrastructure, and when to make them.

In the United States, the Federal Communications Commission (FCC) is interested in the relationship between demand and capacity for several reasons. First, the FCC recognizes the need to define broadband benchmarks based on traffic demand and is considering doing so [9]. It has defined a “typical” household traffic demand to enable concurrent broadband use, such as video streaming, web browsing, and VoIP. The FCC has also asked for comments and suggestions on how to define such a demand-based benchmark for future planning [7,8]. Second, recent research shows that diurnal Internet usage patterns are correlated with GDP, Internet allocations, and the electrical consumption of a region [11], which makes the study of usage potentially relevant to the regulatory bodies responsible for development. Finally, the FCC is responsible for increasing broadband deployment throughout the US, and it recently decided to aggressively increase the broadband threshold benchmark to 25 Mbps downstream and 3 Mbps upstream. Yet, a survey conducted by NCTA (for the FCC) showed that the largest deterrent to deployment of faster speed tiers is that consumers do not *want* the faster speeds (the second largest deterrent is the price) [8]. Clearly, this question deserves both rigorous and continuous study.

Previous work discovered that users who are already maximizing their usage on a given access link will continue to do so when they are migrated to a higher service tier [1]. In this paper, we study how the traffic demands of subscribers who are *already* on service plans with high downstream throughput respond to an undisclosed service plan upgrade as part of a randomized control trial (RCT). This experiment offers the unique opportunity to explore the effects of a service-tier upgrade on user traffic demand, while mitigating the cognitive bias of the service-tier upgrade by withholding that information from subscribers. To the best of our knowledge, this is the first such comparative study of usage behavior in a controlled experiment to study responses to service upgrades.

Our study is based on data from the residential home gateways of Comcast subscribers in Salt Lake City, Utah. To measure traffic demand, Comcast collects aggregate byte counts every 15 minutes from two types of users: *control*, or users who pay and use a high capacity access link (105 Mbps); and *treatment*, or users who pay for 105 Mbps but were actually offered a 250 Mbps access link *without their knowledge*. We evaluate three months of traffic demand for more than 6,000 Comcast subscribers, 1,519 of whom were in the treatment group. We find that subscribers who are already using most of their available capacity at the 105 Mbps and the 250 Mbps service tiers do not show a significant difference in traffic demand. On the other hand, subscribers who exhibit moderate traffic demands in the both groups often exhibit a large relative difference in their traffic demands. This result suggests that even users who are not fully exhausting the available capacity at one service tier may increase usage at higher service tiers, perhaps because the improved performance at the higher tier may cause these subscribers to use the Internet more than they otherwise would. We also observed that the most significant difference in per-subscriber traffic demand occurred during non-prime-time hours on weekdays, suggesting that this demographic of consumer may disproportionately include users who work from home. Such a

phenomenon is also consistent with our observation that traffic demands at these higher service tiers consistently rises throughout the course of the day, with no mid-afternoon drop in traffic volume, as is evident in other studies.

The rest of the paper is organized as follows. In Section 2 we overview some previous studies of traffic demand and service capacity. In Section 3, we offer details about our data, sanitization, and characterization. We then proceed by describing our evaluation criteria and analyze traffic demand in response to a service tier upgrade in Section 4. We summarize our findings in Section 5.

2 Related Work

The measurement community has produced a plethora of studies of broadband performance analysis, yet has performed relatively fewer studies of traffic demand in broadband access networks. The increasing availability of high-bandwidth Internet services and the FCC’s recent interest in exploring traffic demand as a broadband benchmark [8] now calls for increased attention to the relationship between user traffic demand and broadband capacity.

Our work complements an earlier study by Bischof *et al.* [1], who used *natural* experiments to investigate causal relationships between the traffic demand (which they refer to as “user demand”, or “usage” in their paper) and factors such as service capacity, performance, and price. Bischof *et al.* showed that demand increases with capacity, but “follows a law of diminishing returns”; in other words, increases in capacity for an already high tier results in a lower increase in demand. Our work presents complementary results from a large-scale *controlled* experiment and examines in particular a high service tier (105 Mbps) that has not been studied before. Our dataset mitigates the affect of price, performance, and other potential biases (such as regional [2, 4], capped usage [3], and “geek-effect” [1]) by limiting the dataset to a large number of users selected randomly from the same service tier and location.

Zheleva *et al.* present a case study of the effects of an Internet service upgrade, from 256 kbps satellite to 2 Mbps terrestrial wireless, in rural Zambia [15]. This work observed that the stark change in traffic demand three months after the upgrade caused a performance bottleneck. In contrast, our case study focuses on traffic demands of subscribers from much higher service tiers who are not continuously bottlenecked by their access link; additionally, we study how users adjust their traffic demands without informing them of the upgrade, thus eliminating potential cognitive bias.

Other efforts such as [10, 12] study the characteristics of residential broadband, and report the contributions of the most popular web applications to the total usage. The bi-annual Sandvine reports [13, 14] provide an overview of overall Internet traffic demand from fixed lines and mobile carriers as well as an updated analysis of the most popular Internet applications. They showed that video accounts for 63% of traffic usage overall, and traffic demand peaks during the peak evening hours, possibly due to increasing video content consumption. Our work does not concern with the applications responsible for most traffic,

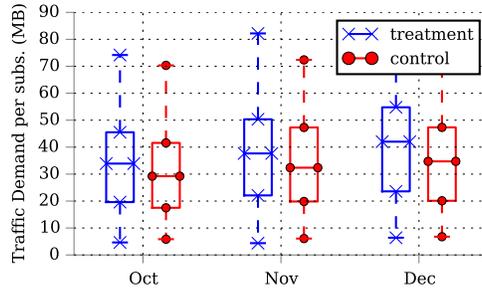


Fig. 1: Distribution of downstream demand averaged across subscribers in the control and treatment groups from October through December 2014.

but only with the peak period during which an individual subscriber’s traffic demand is high.

3 Method and Data

We describe the design of our randomized control experiment and the dataset that we used for this experiment.

3.1 Method

In designing our controlled experiment, we follow the popular statistical convention of experimental designers to refer to the service upgrade as *factor*, the group of users without the upgrade as *control* and the upgraded users as *treatment* [5].

Controlled experiments are difficult to do on the Internet scale. Our work involves a randomized control experiment on the scale of a large urban city. This enables us to study the effect of just one factor, *the service plan upgrade*, while other factors, such as price, performance, or regional differences between users, are controlled. We believe the effects observed on this dataset will also be observed in others collected from urban cities and high tiers.

By examining a single ISP’s high-capacity tier with an unannounced upgrade, our dataset mitigates several biases that previous studies may have suffered. Studying the behavior of users who opt for buying a higher service plan (unsatisfied subscribers) will naturally show an increase in demand on upgrading service [1]. Similarly users who have been offered an upgrade in service may change their behavior to utilize the upgraded capacity (cognitive bias) [15]. Studying datasets with these biases are prone to positive high correlation between demand and capacity.

3.2 Data

Our raw dataset consists of network usage byte counters reported every 15 minutes from October 1, 2014 to December 29, 2014 from about 20,000 Comcast

Dataset	Hourly traffic per 1,000 subscribers (GB)				per subscriber
	Total	95% Traffic	PT	Non-PT	Daily Demand (GB)
control down	2.67×10^5	234.5	205.1	108.5	2.97
treatment down	2.95×10^5	244.42	209.5	122.3	3.30
control up	2.98×10^4	21.39	18.942	12.80	0.33
treatment up	4.27×10^4	31.48	22.81	19.02	0.48

Table 1: Overview of the control (4,845 subscribers) and treatment (1,519 subscribers) datasets for upstream and downstream traffic. The 95 percentile traffic is the peak of total demand. PT traffic is the average traffic demand per hour during prime-time hours. We normalize traffic for both groups to 1,000 subscribers for comparison. The daily demand is the average traffic demand per subscriber over a single day. All values are in gigabytes (GB).

residential broadband gateways in Salt Lake City, Utah. Each dataset contains the following fields: Device ID (household identifier), the 15-minute time interval (end time), service direction ({downstream, upstream}), anonymized IP address, and the bytes transferred in each 15-minute interval.

The data consisted of two groups: a *control* set, consisting of 18,354 households with a 105 Mbps access link; and a *treatment* set, consisting of 2,219 households that were paying for a 105 Mbps access link, yet were receiving 250 Mbps instead. Subscribers in the treatment group were selected randomly and were not told that their access bandwidth had been increased for the three months of our analysis. Our initial analysis of the data from more than 20,573 households showed that not all gateways were reporting their traffic counters every 15 minutes over the whole three-month period: 32% of the *treatment* dataset and 72% of the *control* dataset gateway devices were responsive less than 80% of the time throughout the measurement period. For the analysis in Section 4, we present results based on the accepted group of subscribers that contributed to the three-month dataset more than 80% during their lifetime. Our ultimate dataset consists of 4,845 subscribers in the *control* dataset and 1,519 subscribers in the *treatment* dataset.

Figure 1 shows the distribution of downstream traffic demand, averaged across subscribers (average bytes per 15-minute sample period), for the three months in our measurement period for both groups. Table 1 compares the total demand for subscribers in the *control* and *treatment* datasets, scaled to a thousand households. The downlink 95th percentile traffic demand over an hour is 234.5 GB for the lower tier control group, and 244.42 GB for the higher tier treatment group. Table 1 also shows that an average subscriber in the control group would download 2.97 GB in a day, and 3.30 GB if they belonged to the treatment group. As for the uplink, an average subscriber would transfer 0.33 GB over a day in the control group, and 0.48 GB over a day in the treatment group.

4 Results

Table 2 shows the metrics that we use to evaluate how user demand responds to service-tier upgrades. The *traffic demand* for a subscriber is defined as the

Parameter	Definition
Traffic Demand per Subscriber (§ 4.1)	total bytes transferred in measurement int. number of contributing subscribers
Peak Demand (§ 4.1)	Daily 95th percentile of bytes transferred in any 15-minute interval
Prime-Time Ratio (§ 4.2)	$\frac{\text{avg usage in peak (prime-time) hour}}{\text{avg usage in off-peak hour}}$
Peak-to-Average Ratio (§ 4.3)	$\frac{95\text{-ile of daily traffic demand}}{\text{mean of daily traffic demand}}$

Table 2: Evaluation Metrics

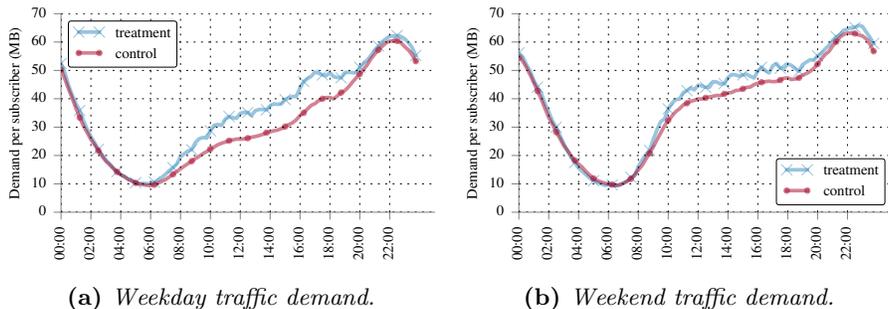


Fig. 2: Mean subscriber demand (bytes per 15-minute interval).

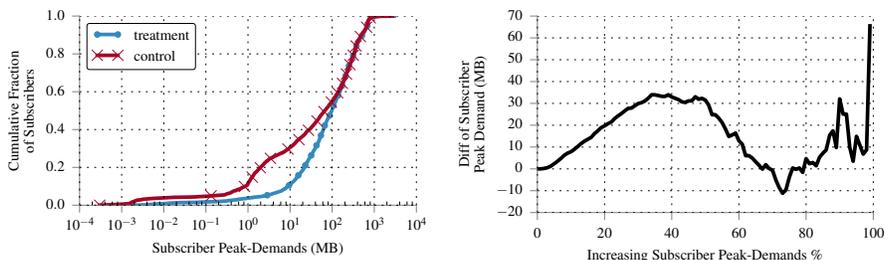
total bytes transferred, in upstream or downstream, during a single sample measurement (15 minutes). We use traffic demand to calculate the total demand per hour, and the average and 95th percentile peak demand over a day. To compare the total traffic of the control and treatment groups, we scale to a thousand subscribers wherever applicable. We define *prime time* as 8:00 p.m. to 12:00 a.m., when Internet usage tends to be highest. Indeed, we observed that the total daily traffic consistently falls within 90th percentile during this four-hour period. We define the *prime-time ratio* as the ratio of traffic during an average prime-time hour, to the average hourly traffic outside the prime-time hour. This ratio conveys the disparity between demand during the prime-time and the rest of the day. The rest of this section explores the effects of a service-tier upgrade on user traffic demand in the context of these metrics.

4.1 Traffic Demand Per Subscriber

We first explore how an upgrade to a higher service tier affected the average traffic demand per subscriber, for different times of the day and days of the week. Figure 2 shows the average downlink traffic demand across subscribers for a week, for both the treatment and control groups. We observe that subscriber behavior differs significantly on weekdays and weekends. The average per-subscriber demand over a weekday is 35.6 MB, and the 95th percentile peak demand is 61.12 MB for subscribers in the treatment group (Table 3). Over a

		median	mean	95%
Weekday	treatment	35.97	35.58	61.12
	control	28.06	31.12	58.78
Weekend	treatment	45.27	40.10	64.27
	control	41.15	37.66	62.23

Table 3: Weekday and weekend traffic demands (MB) per measurement window.



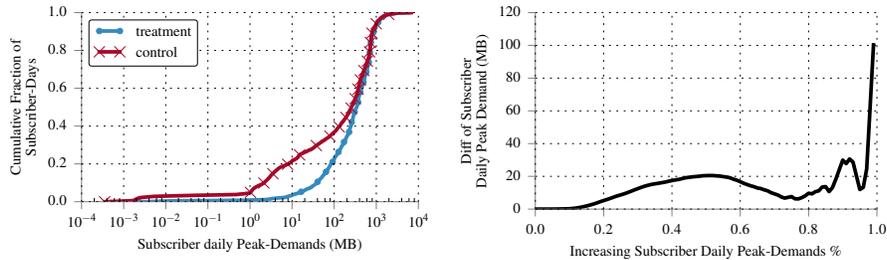
(a) Peak (95%) traffic demand per subscriber. (b) Change in overall peak (95%) demand per subscriber.

Fig. 3: 95th percentile traffic demand (bytes per 15 minutes) per subscriber for the control and treatment groups over the three-month measurement period. Subscribers were considered at every 5% in each group. The y-axis units are bytes transferred in the peak 15-minute interval, in MB.

weekend, the average demand is 40.1 MB, and the 95th percentile demand is 64.3 MB for treatment, but the median is 45.27 MB due to consistent use in the major part of the day. On weekdays, traffic demand increases monotonically from morning until prime-time hours in the evening. On weekends, we observed a sharp rise in demand in the early morning period, from 8:00 a.m. to 10:00 a.m. Then, the demand plateaued until the next rise before evening prime-time hours. Previous reports indicate that the aggregate traffic volume for US fixed access link providers usually troughs during mid-afternoon hours (between 2:00 p.m.–6:00 p.m.) [13]. In contrast to these previous reports, we do not observe such troughs in subscriber demand.

Figure 3a shows the distribution of the the 95th percentile downlink traffic demand over the three-month measurement period. The highest peak demand per 15-minute interval amongst subscribers in the control group was 2.97 GB; in the treatment group, the highest peak demand was 3.0 GB. The average peak traffic demand was 169.8 MB for control and 186.6 MB for treatment. Given the 105 Mbps service-tier capacity, this means that users rarely utilize their links, even on averaging the 95th percentile demand (average utilization was 1.43% for control and 1.5% for treatment).

We suspected that the subscribers who downloaded most bytes in the higher service tier would be the ones causing the largest difference in mean demand, as previous studies have observed such a phenomenon. In fact, we observed that the more moderate (median) subscribers actually seemed to exhibit larger dif-



(a) Daily peak (95%) traffic demand per subscriber. (b) Change in daily peak (95%) demand per subscriber.

Fig. 4: 95th percentile traffic demand (bytes per 15 minutes) per day per subscriber for the control and treatment groups. Subscribers were considered at every 5% in each group. y-axis units are bytes transferred in the peak 15-minute interval, in MB.

ferences in traffic demand: The median peak demand was 66.7 MB for the lower service tier, and 98.4 MB for the higher tier. This result indicates that the more moderate subscribers who received a service-tier upgrade exhibit significantly higher peak demand than comparable users in the control group.

We also observed a significant difference in the mean peak demand was present in the 50% of subscribers in the control group with the lowest traffic demand when compared to the same set of subscribers of the treatment group. (This disparity appears as a large gap under the 50% tick in Figure 3a.)

Figure 3b shows another way of looking at this phenomenon: it explores the difference between the distribution of users with particular traffic demands in the control and treatment groups. For each group, we sort the subscribers according to increasing demand. Then we compute the difference in peak demand for each percentile in the group. For example, the plot shows the median user (50% on the x-axis) increased their peak demand by about 25% in response to the service tier upgrade. Comparing the 70% subscribers of both groups with the least demand, we see that peak demand in the treatment group is higher than the peak demand in the control group, indicating that in fact even moderate users may increase their demand as a result of the service-tier increase, even though they are not using the full capacity in either case. When we combine this analysis with that in Figure 3a, we find that these subscribers who respond with increased usage have a peak demand less than 200 MB. Naturally, the small number of users with the highest demand (closer to 100%) also show a substantially larger usage for the higher service-tier.

Further investigation revealed that users with moderate peak traffic demands not only exhibit a large difference in their traffic demands in aggregate, but also on a daily basis. Figure 4 shows that when subscribers on the lower tier had a daily peak demand under 600 MB, 70% of subscribers in the treatment group had 15-minute demands that were 5–20 MB higher. The ratio of the differences in demand across percentiles also shows that the 40% of subscribers with lowest

		Hourly traffic in PT	Hourly traffic in Non-PT	Prime-time Ratio
Weekday	treatment	233.12	124.18	1.88
	control	225.40	104.30	2.16
Weekend	treatment	246.93	143.08	1.73
	control	238.15	133.16	1.79

Table 4: Hourly traffic demand during (and outside) prime-time hours per subscriber, in MB. Prime-time traffic demand is defined as the average traffic demand during a prime-time hour.

peak demands in the treatment group demonstrate more than double the daily peak traffic demand of comparable users in the control group.

One possible explanation for why moderate users might increase their usage in response to a service-tier upgrade is that the higher service tier not only affords more capacity, but also a better user experience (*e.g.*, faster downloads). Thus, even though users may not be exhausting the capacity of the higher service tier, they nonetheless seem to respond to the service tier upgrade by using the Internet more than they had before the service-tier upgrade.

4.2 Prime-Time Ratio

ISPs design networks to handle peak demand, which is usually observed during prime-time hours, when subscribers heavily consume real-time entertainment traffic, such as video. The FCC defines prime-time as the local time from 7:00–11:00 p.m. [6]. To measure the concentration of network usage during prime-time, we use Sandvine’s definition of the *prime-time ratio*: the ratio of the average (hourly) traffic demand during prime-time hours to the average demand in non-prime-time hours [13,14]. We measured the prime-time ratio of the subscribers in the control and treatment groups considering each contiguous four-hour period in each day. Our experiment shows that, in fact, the evening hours with the largest prime-time ratio are 8:00 p.m.–12:00 a.m., so we use this time interval for our definition of prime time.

Table 1 shows that the average hourly prime-time downstream traffic per 1,000 subscribers is 209.5 GB for the treatment group, compared to 205.1 GB for the control group, which is about a 2% difference. In contrast, during an average hour *outside* of prime time, the traffic per 1,000 subscribers is 122.3 GB for the treatment group, compared to 108.5 GB for the control group, amounting to about a 12% difference. The more significant difference in demand during hours outside of the daily prime-time is also apparent from the weekly usage patterns in Figure 2.

We also calculated the prime-time ratio per day over weekends and weekdays, as shown in Table 4. On weekends, the prime-time ratios for the treatment and control groups are 1.73 and 1.79 respectively. On the weekdays, the prime-time ratio for the control group is 2.16 compared to 1.88 for the treatment group. In terms of absolute demand, the prime-time demand on weekdays in the treatment group is within 4% of that in the control group. In contrast, the demand in

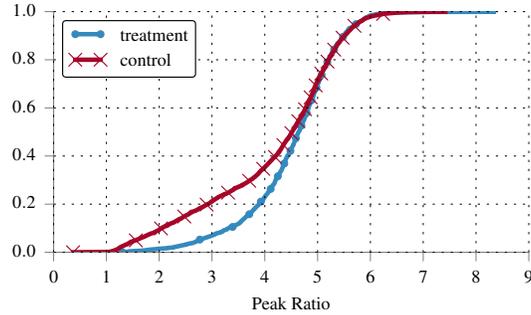


Fig. 5: Distribution of the daily peak-to-average ratio per subscriber, averaged for each subscriber over the measurement period in the treatment and control groups.

non-prime-time hours is 19% higher for the treatment group on weekdays, and only 7.5% higher on weekends. The large difference in non-prime-time demand between the control and treatment groups on weekdays suggests that many of the users in the treatment group may in fact be subscribers who work from home may adjust their behavior during non-prime-time hours and weekdays in response to a higher service tier.

Although 6% of the subscribers in both groups had a prime-time ratio over 100, we also observed that 9% of the control group and 14% of the treatment group had prime-time ratios *less than one*, indicating that these users actually had higher demand during the day than they did during prime time. Similarly, these users may be small home businesses or subscribers who work at home.

4.3 Peak-to-Average Ratio

In addition to examining traffic demands across the entire four-hour prime-time window, we also explored how subscribers in the treatment group exhibited different behavior for the 15-minute interval of highest (95th percentile) demand in a day, regardless of prime-time hours. We measure the disparity between a subscriber’s daily 95th percentile and the mean usage as the *peak-to-average ratio* (PAR). This standard metric shows the ratio of peak values to the effective value and extends those used in conventional studies of user traffic patterns, such as the Sandvine’s peak traffic analysis [13].

Figure 5 shows the PAR for each subscriber in the treatment and control groups. The median PAR for subscribers from the treatment group is 4.64, compared to 4.51 for the control group. We found that 40% of the subscribers in both groups have PAR greater than five; the PAR of subscribers in the treatment group is higher than those in the control group, perhaps indicating that users in both higher service tiers do in fact use the additional capacity for short periods of time. The notable difference occurs for peak-to-average ratios of less-than 5: as we observed in Section 4.1, subscribers with more moderate (median) traffic

demands tend to increase their peak demand more in response to the increased service tier. Again, we believe these trends appear not because users are necessarily eager to fill the additional capacity of a higher service tier, but rather may be occurring because the upgrade results in better performance, and that this improved user experience in turn causes these subscribers to make more use of the Internet.

The lower prime-time ratio by volume, and a consistently higher peak-to-average ratio per subscriber indicates the following: subscribers in the treatment group have higher peak-to-average ratio than those in the control group. However, these subscribers tend to still have low absolute demand, so the relatively higher PAR for the treatment group does not significantly affect total traffic during prime-time and, when it is high, the demand tends to be in non-prime-time hours. Consistent with the results in Section 4.2, we also found that on weekdays, the peak-to-average ratios in the treatment group are higher than the control group, whereas on weekends peak-to-average ratios for both the control and treatment groups are similar.

5 Conclusion

In this paper, we study how subscribers respond to an increase in their ISP’s service tier. To do so, we use a randomized control trial to compare per-subscriber traffic volumes between two groups of Comcast subscribers in the same city during the same time period: a control group, with Comcast’s 105 Mbps service offering; and a treatment group of subscribers who were upgraded to the 250 Mbps service tier without their knowledge. We observed that subscribers with more moderate traffic demands exhibit a relatively higher usage for the upgraded service-tier as compared to subscribers who were already sending relatively high traffic volumes in both groups.

Initially, we were surprised by this result: after all, both intuition and previous work suggest that when users experience service-tier upgrades, they immediately exhaust the available capacity (particularly the high-volume subscribers). At higher tiers, however, we observe a completely different phenomenon: in general, users are not exhausting the available capacity, but a service tier upgrade may simply result in a better user experience that causes subscribers with more moderate traffic demands to use the Internet more than they otherwise would. The fact that the most significant difference that we observed between the two service-tier groups occurred during non-prime-time hours on weekdays also suggests that these higher service tiers may generally be disproportionately used by subscribers who work from home. Future research should aim to repeat our experiment for different cohorts (*i.e.*, different subscribers, geographies, service tiers, and ISPs), and could also strive to obtain more fine-grained traffic statistics to explore exactly which applications are responsible for the behavioral changes that we have observed.

Acknowledgments

This research was supported by the Comcast Tech Research Fund and by NSF Awards CNS-1539902 and CNS-1540066. We thank Jason Livingood and James Moon from Comcast for helpful discussions and access to the data that we used for this study.

References

1. Z. S. Bischof, F. E. Bustamante, and R. Stanojevic. Need, want, can afford: Broadband markets and the behavior of users. In *Proceedings of the 2014 Conference on Internet Measurement Conference*, IMC '14, pages 73–86, New York, NY, USA, 2014. ACM.
2. J. C. Cardona, R. Stanojevic, and R. Cuevas. On weather and internet traffic demand. In *Proceedings of the 14th International Conference on Passive and Active Measurement*, PAM'13, pages 260–263, Berlin, Heidelberg, 2013. Springer-Verlag.
3. M. Chetty, R. Banks, A. Brush, J. Donner, and R. Grinter. You're capped: Understanding the effects of bandwidth caps on broadband use in the home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pages 3021–3030, New York, NY, USA, 2012. ACM.
4. K. Cho, K. Fukuda, H. Esaki, and A. Kato. The impact and implications of the growth in residential user-to-user traffic. In *Proceedings of the 2006 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, SIGCOMM '06, pages 207–218, New York, NY, USA, 2006. ACM.
5. V. Easton and J. McColl. *Statistics Glossary*. STEPS, 1997.
6. Federal Communications Commission. Measuring Broadband America - 2014, April 2014.
7. Federal Communications Commission. Tenth Broadband Progress Report No 14-113, February 2014.
8. Federal Communications Commission. Eleventh Broadband Progress Report No 15-10A1, February 2015.
9. Federal Communications Commission. International Broadband Data Report (Fourth), February 2015.
10. S. Grover, M. S. Park, S. Sundaresan, S. Burnett, H. Kim, B. Ravi, and N. Feamster. Peeking behind the nat: An empirical study of home networks. In *Proceedings of the 2013 Conference on Internet Measurement Conference*, IMC '13, pages 377–390, New York, NY, USA, 2013. ACM.
11. Lin Quan, and Heidemann, John and Pradkin, Yuri. ANT Evaluation of the Diurnal Internet, October 2014.
12. G. Maier, A. Feldmann, V. Paxson, and M. Allman. On dominant characteristics of residential broadband internet traffic. In *Proceedings of the 9th ACM SIGCOMM Conference on Internet Measurement Conference*, IMC '09, pages 90–102, New York, NY, USA, 2009. ACM.
13. Sandvine. Global Internet Phenomena Report - 1H, April 2014.
14. Sandvine. Global Internet Phenomena Report - 2H, November 2014.
15. M. Zheleva, P. Schmitt, M. Vigil, and E. Belding. The increased bandwidth fallacy: Performance and usage in rural zambia. In *Annual Symposium on Computing for Development (DEV)*, 2013.