Do Mobile Data Plans Affect Usage? Results from a Pricing Trial with ISP Customers

Carlee Joe-Wong[1], Sangtae Ha[2], Soumya Sen[3], and Mung Chiang[1]

[1] Princeton University
[2] University of Colorado-Boulder
[3] University of Minnesota
{cjoe,chiangm}@princeton.edu,sangtae.ha@colorado.edu,ssen@umn.edu

Abstract. The growing amount of traffic in mobile data networks is causing concern for Internet service providers (ISPs), especially smaller ISPs that need to lease expensive links to Tier 1 networks. Large amounts of traffic in "peak" hours are of especial concern, since network capacity must be provisioned to accommodate these peaks. In response, many ISPs have begun trying to influence user behavior with pricing. Time-dependent pricing (TDP) can help reduce peaks, since it allows ISPs to charge higher prices during peak periods. We present results from the first TDP trial with a commercial ISP. In addition to analyzing application-specific mobile and WiFi traffic, we compare changes in user behavior due to monthly data caps and time-dependent prices. We find that monthly data caps tend to reduce usage, while TDP can increase usage as users consume more data during discounted times. Moreover, unlike data caps, TDP reduces the network's peak-to-average usage ratio, lessening the need for network over-provisioning and increasing ISP profit.

1 Introduction

Mobile data usage is growing at unprecedented rates, with Cisco estimating that global mobile data traffic grew 81% in 2013 and projecting a compound annual growth rate of 61% over the next five years [1]. This trend has significantly increased ISPs' capital expenses, as they must provision their network to accommodate peak usage during the day [3,16]. Smaller ISPs are particularly affected, as their network capacity is limited by middle mile links to Tier 1 operators, which are leased at rates based on peak usage [20]. Many ISPs are therefore trying to reduce their peak mobile data traffic [18,22]. In this paper, we focus on the use of pricing as an incentive for users to reduce their peak usage.

Most U.S. ISPs charge fixed fees for limited monthly data caps. Yet data caps may not effectively limit usage peaks, as users can remain under their caps by using less data at off-peak times and not changing their peak-time usage. *Time-dependent pricing* (TDP) allows the ISP to effectively target network peaks by offering higher prices at those times, incentivizing users to consume data at other times. Yet TDP's effectiveness depends on users' willingness to shift their data usage in exchange for reduced prices, which can vary for different users and applications: business users, for instance, might not wait to download email attachments, but teenagers might wait to download video purchases [8]. To the best of our knowledge, there are no systematic studies of these price-delay tolerances, and no works on TDP have yet accounted for the effect of displaying usage statistics to users: showing users these statistics would make them more aware of their usage and might affect their usage behavior. Such trials have also focused only on university populations [2,8].

In this paper, we present *results from the first TDP trial with a commercial ISP*. We recruited 27 customers of a local U.S. ISP, dividing users into time-independent pricing (TIP) and TDP groups. The TIP users used a data usage monitoring application with their regular pricing plan. We show that this monitoring induced them to reduce their usage below their monthly data caps, but that they still had very high peak usage. The TDP users both monitored their data usage and received time-dependent prices; we show that the prices induced TDP users to increase their usage at discounted times. Thus, *simple data caps do not effectively reduce ISPs' peak network usage, but TDP does.* Our work makes the following contributions:

– An analysis of the results of the first TDP trial with a commercial ISP, including:
– A study of temporal and per-app WiFi and cellular usage data.
– An analysis of the impact of data usage monitoring apps on cellular and WiFi usage behavior.
– An evaluation of real customers' price sensitivity and delay tolerance for different applications.
– An examination of TDP's cost benefits with empirical price sensitivity and delay tolerance estimates.

In the next section, we give an overview of related work. We then describe the trial structure and our analysis methodology in Section 3. We analyze users' pre-trial data usage in Section 4 before presenting the trial results in Section 5. We conclude in Section 6.

2 Related Work

Previous trials in a university setting demonstrated TDP's effectiveness in changing mobile data usage patterns [8]. Others have suggested that data usage and user responses to incentives depend on psychological [2] or socioeconomic [14] factors. Another work on price elasticities for wireline broadband speeds considers a wider population [7]. These trials, however, do not analyze TDP's effects on different apps or account for the effect of simply displaying usage statistics to users. We find that displaying usage statistics generally decreases usage volume, but when combined with TDP can result in increased usage at low-price times.

Many studies have found a significant time-of-day pattern in cellular network traffic [11]. Others have analyzed LTE network performance [9] and compared the performance of different network interfaces (e.g., LTE and WiFi) [19]. Papers focusing on individual users' data consumption show a large diversity in the amount of data used by different users and different apps on mobile and WiFi networks [5,6,13,21]. These lead to distinct temporal usage patterns, which [12] showed can be leveraged to improve users' experience with intelligent WiFi offloading. Similarly, [10] shows that delaying mobile off-screen traffic, which is assumed to be delay-tolerant, can improve energy usage. Another work on Super Bowl traffic shows that short-term delays can be leveraged to eliminate congestion [4]. Our work provides a more nuanced estimation of delay tolerances and examines their monetary value to users by offering price incentives.

3 Methodology

We designed the trial to determine the effects of data usage monitoring and a combination of TDP with usage monitoring. We first outline the trial structure and then describe the data collected and apps distributed to trial participants. We finally present a model for users' price-delay tolerances that allows us to evaluate TDP's benefits for ISPs.

3.1 Trial Participants and Structure

We recruited 27 active trial participants from an ISP's customer base. While our sample size is small, the number of participants was limited by the fact that we changed some of their mobile data plans to TDP, broadening the trial's financial implications beyond those of simply measuring usage. All participants used their own Android devices. They did not use data monitoring apps before the trial, but did have monthly data caps.

All active participants downloaded custom-built apps for the trial, which we describe in more detail in the next section. These participants were divided into two groups: time-independent pricing (TIP) and TDP users. The TIP users installed data monitoring apps, allowing us to estimate the effect of usage monitoring with data caps. The TDP users' app both monitored data and offered time-dependent prices. Thus, their behavior is affected by both data monitoring and TDP. We additionally collected passive network data on more than 5000 "control" users, who did not install any apps. Table 1 summarizes the three groups of users.

The control and TIP users' data caps, which are not shared among devices, ranged from 1 to 10 GB and were the same as before the trial. TDP users were charged hourly time-dependent prices, e.g., \$10/GB from 12 to 1am and \$15/GB from 1 to 2am. The prices offered ranged from \$10/GB to \$20/GB, and were chosen to be no higher than the ISP's most popular data plan: a monthly 1 GB cap for \$19.99. Prices were randomly determined and shown to the TDP users 24 hours in advance, allowing them to plan their usage over the next day.

	Recruitment	Data Collection	Data Plan
Control	Random	RADIUS logs	Unchanged
TIP	Volunteer	Trial app & RADIUS	Unchanged
TDP	Volunteer	Trial app & RADIUS	TDP rates

Table 1: Three groups of trial participants.

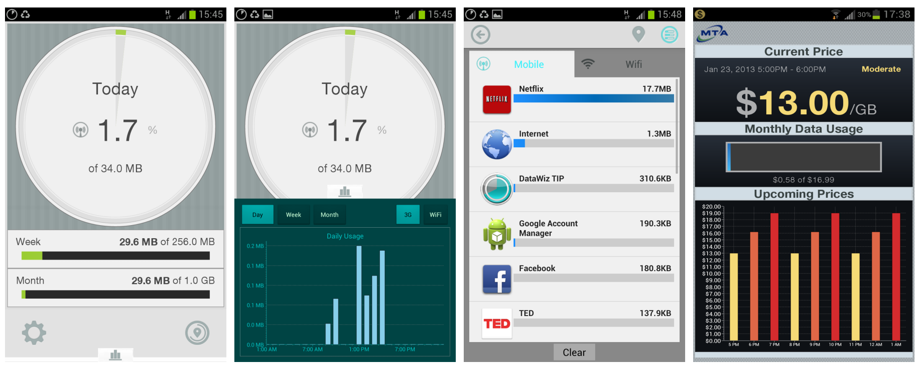

(a) TIP app home. (b) Usage graphs. (c) Per-app usage. (d) TDP app home.

Fig. 1: Screenshots of the TIP and TDP apps. The TIP app's small pie chart indicator on the upper left of the screen indicator bar (1a) shows the approximate portion of a user's monthly data cap used so far. The TDP app's colored price indicator on this bar (1d) indicates the current price range.

3.2 Data Collection

Our dataset consists of two separate types of data: one 21.5 GB set of RADIUS network data, and one 10.5 GB set of application usage data. The RADIUS data was collected from March 2012 to June 2013 for all TIP, TDP, and control group users and contains 140 million session records, including input and output byte counts and start and end timestamps.

The second dataset was collected by TIP and TDP trial participants' apps during the June 2013 trial. This data consists of uplink and downlink cellular and WiFi byte counts for every application, collected every ten minutes, as well as the hourly prices offered to TDP participants.[4] We developed separate TDP and TIP apps for the trial, which collect usage information and display it to users.

The TIP app is a usage monitoring application with screens shown in Figs. 1a–1c. Users could view their monthly, weekly, and daily usage as a fraction of their data cap (Figs. 1a and 1b), as well as their per-app usage (Fig. 1c). Daily and weekly data caps were calculated based on the monthly cap and number of days left in the month. Users could quickly see the remaining fraction of their monthly cap by looking at the pie chart icon on the bar at the top of the screen.

[4] We did not collect more detailed data, e.g., packet traces, to maintain users' privacy. Participants fully consented to the data collection, but complete anonymity was not possible as we had to calculate how much to charge the TDP users.

The TDP app allows users to monitor their spending on data and see the future prices. As with the TIP app, users can see their per-app usage (Fig. 1c). However, the main screen has been modified (Fig. 1d) to show the future prices and the amount the user has spent during the month. On the top left of the home screen bar, we show a color-coded price indicator that is visible both inside and outside our app; the indicator lets users easily see the current price, making it easier for them to decide whether or not to consume data at a given time [15]. It is colored red, orange, yellow, or green for high, medium, low, and very low prices respectively.

3.3 Estimating Price-Delay Tolerances and Optimizing Prices

We quantify users' price-delay tolerances by fitting their observed usage with TDP to a model of users' expected usage volume given the prices offered and their price-delay tolerances. We then calculate the ISPs' expected profit and users' expected traffic patterns with these user parameters. We use the following process:

Establish baseline usage: We establish the average amount of data used in each hour of the day by extrapolating from TDP users' pre-trial RADIUS data. We divide the usage into different apps using the fraction of data used by each app in each hour by TIP users.[5]

Model users' price-delay tolerances: We use a model adapted from our previous work [8,17]. We define "waiting functions" $w_\beta(d,t)$ that give the probability that a user will wait for time t, given a savings d on the usage price. The waiting functions have the form $w_\beta(d,t) = C(\beta)\max(d,0)(t+1)^{-\beta}$, where $C(\beta)$ is a normalization constant and the β parameter controls the user's "willingness-to-wait:" w_β decreases faster with t for larger β, making users less likely to wait for longer amounts of time. The value of β differs for different applications, e.g., a user is more likely to delay a software update than checking email. We can compare apps' delay tolerances by comparing their β parameters.

Estimate the model parameters: We choose the model parameters that provide the best fit between observed TDP trial usage and the usage predicted by our model, given the prices offered during the trial.

To predict TDP usage, we identify two types of changes in usage relative to the baseline: first, users may shift some usage from higher- to lower-priced times. We use the waiting functions above to calculate the expected amounts shifted for each app. Second, price discounts can induce users to increase their overall usage [15,17]. Since the amount of the increase depends on the app and time of the day (e.g., users are unlikely to increase their usage while sleeping), we parameterize the usage increase with $\alpha_a(t)$, which depends on the app a and hour t. We use the form $V_a(t)\left((1+d(t))^{\alpha_a(t)}-1\right)$, where $V_a(t)$ is the pre-trial (baseline) usage for app a and $d(t)$ the discount offered (i.e., the maximum price, normalized to 1,

[5] We use per-app data for the TIP users since TDP can skew the app distribution [8], and we have no pre-trial per-app data. RADIUS logs do not have per-app data, and distributing apps before the trial would have skewed users' behavior.

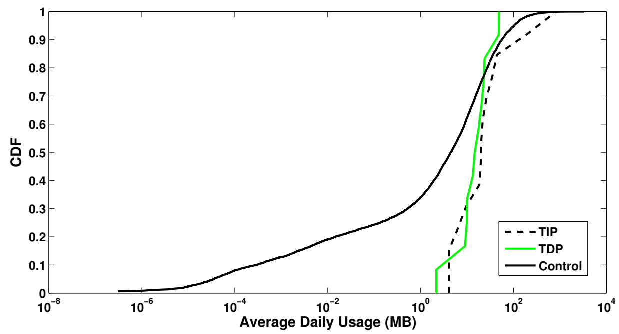

Fig. 2: Average daily usage (March 2012–June 2013).

minus the offered price) in hour t. In accordance with the economic principle of diminishing marginal utility, we constrain $\alpha_a(t) \in [0, 1]$. Note that if $\alpha_a(t) = 0$, the usage does not increase with $d(t)$. We add this term to the amount of traffic shifted to find the total traffic for each app in each hour as a function of the discounts offered and model parameters β and $\alpha_a(t)$.
Calculate profit-maximizing prices: Given the parameter estimates, we can optimize the prices offered over the day so as to maximize ISPs' profit with TDP, i.e., revenue minus cost. The revenue is simply the sum of the time-dependent prices multiplied by the expected usage under TDP. We model the cost as a piecewise-linear function, with zero marginal cost below a fixed capacity C and a constant marginal cost γ for usage above this capacity. Thus, ISPs will choose time-dependent prices so as to maximize their profit

$$\sum_{t=1}^{T} (1 - d(t))\, X(t) - \gamma \max\left(X(t) - C, 0\right), \tag{1}$$

where $X(t)$ is the expected usage at time t after TDP. By continually re-estimating the price-delay tolerances and re-optimizing the prices offered accordingly, the ISP can adapt its prices to changes in user behavior.

4 Traffic Characteristics

In this section, we first construct baseline usage information for TIP, TDP, and control users from our pre-trial RADIUS dataset. We then characterize the major apps used by TIP and TDP users. In all figures, hours given are in local time.

4.1 How much data do users consume?

Figure 2 shows the cumulative distribution function (CDF) of all users' average daily usage. We see that the TIP and TDP users use similar amounts of data, ranging from 2 to 100MB, i.e., a few hundred MB to 3 GB per month. While a substantial minority (34.1%) of control users use less than 1MB per day, none of these users volunteered for our TIP or TDP trial groups.

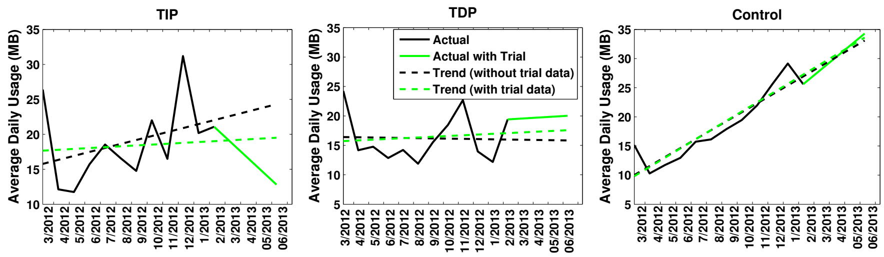

Fig. 3: Average monthly usage for the TIP, TDP, and control users.

Rank	App (Mobile)	%	App (WiFi)	%
1	`com.facebook.katana`	15.24	`com.facebook.katana`	18.93
2	`android.process.media`	11.39	`android.process.media`	17.83
3	`com.pandora.android`	9.05	`com.android.browser`	11.64
4	`com.android.browser`	8.71	`com.android.email`	7.37
5	`com.android.email`	7.26	`mobi.ifunny`	6.75
6	`mobi.ifunny`	3.19	`com.android.chrome`	4.20
7	`com.motorola.motoemail`	2.27	`com.pandora.android`	3.64
8	`com.datawiz.tip`	2.01	`com.rhythmnewmedia.android.e`	3.00
9	`com.motorola.blur.service.main`	1.99	`com.alphonso.pulse`	2.51
10	`com.motorola.contacts`	1.99	`com.datawiz.tip`	2.11

Table 2: Usage fraction of the top 10 apps, comprising 63.1 and 78.0% of total mobile and WiFi usage respectively.

Users' average daily usage changes over time. Figure 3 shows the average daily usage in each month over one year (March 2012–February 2013), fitted with a linear trendline. We see that usage generally increases for TIP and control users, as is consistent with the growing amounts of mobile data traffic, but remains anomalously flat for TDP users. Usage observed during the June 2013 trial period fits this trend for the control group. However, the TIP users see a large decrease and the TDP users a slight increase in usage compared to that predicted by the trendlines. Thus, *TIP users decrease their usage and TDP users increase their usage* during the trial. We examine these findings and their psychological causes in Section 5.

4.2 How is usage distributed among apps?

Table 2 shows the fraction of mobile (cellular) and WiFi usage corresponding to the top 10 apps. Many of the same apps appear for mobile and WiFi, with Facebook and Android's media process the number 1 and number 2 apps for both interfaces. Pandora, web browsing, email, and iFunny also appear in the top 7 apps for both WiFi and mobile usage.[6] Mobile usage is more evenly distributed among apps than is WiFi usage, with the top 10 apps comprising 63.1% of mobile and 78.0% of WiFi usage. Apps outside the top 10 each accounted for less than 2% of usage.

[6] Larger sample sizes with a broader population may yield different top apps.

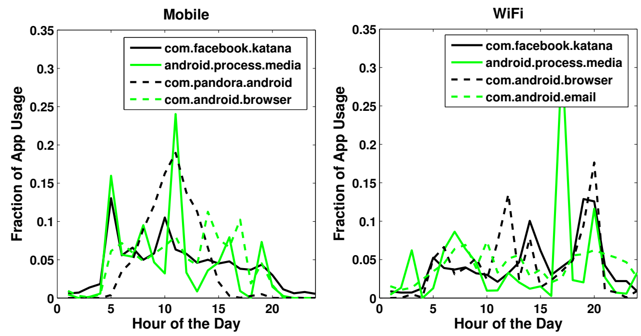

Fig. 4: Mean hourly usage for the top 4 apps.

Figure 4 shows the hourly usage of the top four apps for mobile and WiFi. We see that WiFi is generally used more in the evenings, likely because people are at home then and have WiFi connectivity there. While most apps have generally similar usage patterns, there are some differences: Pandora, for instance, is only used between 5am and 3pm on mobile and peaks around 10am. Android's media process, which is used by other apps to stream videos, shows high peaks for mobile and WiFi usage, likely due to its high bandwidth requirements.

5 Pricing Effects

In this section, we present the trial results. Throughout the discussion, we use the peak-to-average ratio (PAR) of hourly usage over a day to measure the degree to which ISPs must over-provision their network. A higher PAR indicates that the ISP's network has more idle capacity, as it is provisioned for higher peak capacity than is needed on average. Before the trial, TIP and TDP users had an average PAR of 1.88, indicating that the peak hourly traffic was almost twice the mean.

We first show that TIP users decrease their total usage to remain below their data caps, but increase their mobile usage's PAR and may increase their overall WiFi usage. TDP users increase their usage in response to price discounts, allowing ISPs to reduce their PAR by up to 31.4% with profit-maximizing prices.

5.1 Do TIP users decrease their usage?

Most TIP users decrease their usage in order to remain below their data caps. However, their PAR increases to 2.67 from 1.88 before the trial.

Figure 5a shows TIP usage as a fraction of users' data caps before and during the trial. Each circle represents a user, and the circle size is proportional to the user's data cap. The dashed line represents equal usage fractions before and during the trial. In general, users' usage amounts are closer to their data cap during the trial. A few users' data points lie above the dashed line, indicating that they used less of their data caps during the trial. These users, all with relatively small 1GB caps, exceeded their data caps before the trial, but no

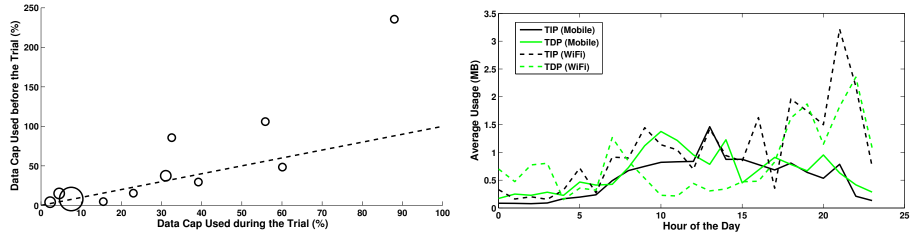

(a) Mean TIP monthly usage. (b) Mean hourly mobile and WiFi usage.

Fig. 5: Monthly and hourly usage volumes.

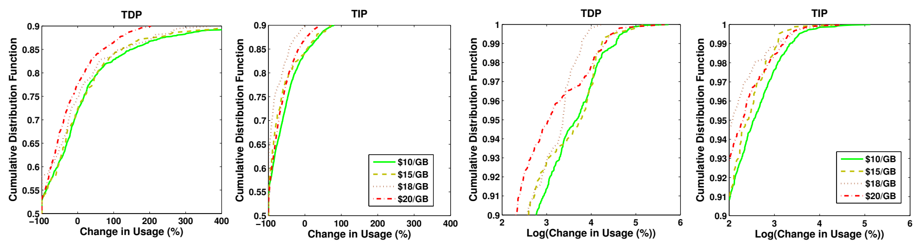

(a) Bottom 90th percentile. (b) Top 10th percentile.

Fig. 6: Change in hourly usage relative to pre-trial usage for the same user in the same hour of the day.

users did so during the trial. Other users' data points lie below the dashed line, indicating that they used more of their data caps during the trial than before. The data monitoring app ensured that they did not have to worry about hitting their data caps.

We conjecture that users reduce their monthly data usage by shifting some of their usage to WiFi. While we do not have pre-trial WiFi statistics (WiFi data was not collected by the network), 65.42% of TIP users' data was consumed over WiFi, versus 55.39% of TDP users'. Figure 5b shows the hourly mobile and WiFi usage patterns for TIP and TDP users. WiFi is used more than mobile in the evening, and spikes at these times for TIP users. This spike may indicate unusually large WiFi usage due to users' not using mobile data.

5.2 Do TDP users respond to price discounts?

TDP users increase their usage more in discounted hours. ISPs' profit-maximizing prices can decrease their peak-to-average hourly traffic ratio by up to 31.4%.

Price-delay tolerances: We offered four different prices during the trial: \$10 (green price indicator), \$15 (yellow), \$18 (orange), and \$20 (red) per GB. Figure 6 shows the % change in usage in different hours for each price, compared to usage in the same hour (e.g., 12 to 1am) for the same user before the trial. While the TIP usage changes are similar for all prices, TDP users have more positive changes for \$10/GB versus \$20/GB, in both the bottom 90th (Fig. 6a) and top 10th (Fig. 6b) percentiles of usage changes. The difference

App	Estimated β	Mean α
`com.facebook.katana`	2.326	0.503
`android.process.media`	1.341	0.234
`com.pandora.android`	0.479	0.141
`com.android.browser`	0	0.212
`com.android.email`	3.000	0.979

Table 3: Price-delay tolerance for the top 5 mobile apps.

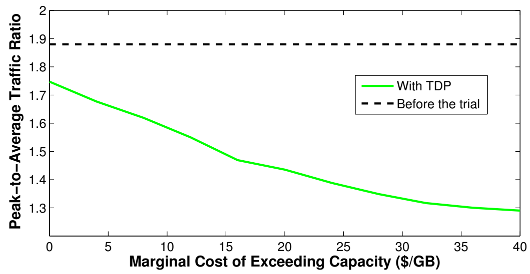

Fig. 7: Peak-to-average hourly traffic ratio with TDP.

is less pronounced for the intermediate \$15/GB and \$18/GB prices, but is still apparent, especially around the 80th percentile. TDP users thus distinguished between very low, moderate, and high prices, perhaps using the colored price indicators. For all prices, TDP users had more positive usage changes than TIP users, likely because they were saving money on some of their usage and felt they could use more data overall. TDP changes above the 97th percentile are less price-dependent, but these are likely outliers occurring when usage increases during hours of very small pre-trial usage.

As explained in Section 3.3, we compare the delay tolerance for different apps by fitting our waiting function model to the trial usage. Table 3 shows the resulting β parameters and average α parameters over time for the top five mobile apps (Table 2). We see that while Pandora has a lower value of β, corresponding to higher delay tolerance, email has the lowest delay tolerance (highest value of β). Web browsing, however, has the highest delay tolerance, perhaps reflecting users' use of the web for looking up non-urgent information. Surprisingly, email has the highest α value (i.e., increase at low-price times independent of shifting), likely because users downloaded more email attachments and images when the price was low.

Maximizing ISP profit: Finally, we use the parameters in Table 3 and app usage fractions in Table 2 to calculate the optimal time-dependent prices offered by the ISP, which maximize (1) for different marginal costs of exceeding capacity (γ). To measure TDP's effect on usage peaks, we calculate the PAR with these optimized prices. Figure 7 shows the achieved PAR for a range of γ values, compared to that before the trial. Even when $\gamma = 0$, the PAR improves due to discounts in less congested hours, which induce an increase in usage and revenue. Thus, TDP can more effectively increase ISP profit and reduce the network's PAR than can simple data caps.

6 Discussion and Conclusion

Pricing is a unique way of controlling network usage in that it explicitly relies on user attitudes and responses to incentives. Thus, to supplement our measurement results, we conducted three opinion surveys with the TIP and TDP participants before, during, and after the trial.[7] As part of the survey, users were asked their opinions on TDP's viability. Most users–especially TDP users in the mid-trial survey–expressed some concern over TDP's possible complexity. However, nearly all users preferred TDP to forced usage throttling in the mid- and post-trial surveys. Combined with our measurement results, we see that TDP can be more effective than capping or throttling usage, but must be implemented carefully to avoid undue complexity. One possible strategy is to use binary prices, e.g., charging either \$10/GB or \$20/GB in any given hour.

Our work shows that users do change their behavior in response to changes in their pricing plans; in particular, TIP users reduce their usage in response to data caps, possibly increasing their WiFi usage. However, data caps are not sufficient to prevent ISPs' need to over-provision networks according to their peak usage. Time-dependent pricing allows ISPs to reduce their peak-to-average traffic ratio, yet requires more sophisticated understanding from users than monthly data caps. While customers are willing to shift their usage in response to time-dependent prices, a full implementation and deployment of TDP will require more experimentation with a wider range of users.

7 Acknowledgments

We gratefully acknowledge the assistance of our colleagues at the Matanuska Telephone Association. Part of the work was supported by NSF CNS-1117126.

References

1. Cisco Visual Networking Index: Global mobile data traffic forecast update, 2013 - 2018, Feb 2014. `http://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white_paper_c11-520862.pdf`.
2. J. M. Dyaberi, B. Parsons, V. S. Pai, K. Kannan, Y. Chen, R. Jana, D. Stern, A. Varshavsky, and B. Wei. Managing cellular congestion using incentives. *IEEE Communications Magazine*, 50(11):100–107, 2012.
3. M. El-Sayed, A. Mukhopadhyay, C. Urrutia-Valdés, and Z. J. Zhao. Mobile data explosion: Monetizing the opportunity through dynamic policies and QoS pipes. *Bell Labs Technical Journal*, 16(2):79–100, 2011.
4. J. Erman and K. K. Ramakrishnan. Understanding the super-sized traffic of the Super Bowl. In *Proc. of ACM IMC*, pages 353–360. ACM, 2013.

[7] Pre-trial TIP and TDP, mid-trial TIP, mid-trial TDP, post-trial TIP, post-trial TDP surveys: `https://www.surveymonkey.com/s/`{`LPYDGWG, 63PVQCW, ZLLLQ86, CPPBH92, CPZP57Q`}.

5. H. Falaki, D. Lymberopoulos, R. Mahajan, S. Kandula, and D. Estrin. A first look at traffic on smartphones. In *Proc. of ACM IMC*, pages 281–287. ACM, 2010.
6. H. Falaki, R. Mahajan, S. Kandula, D. Lymberopoulos, R. Govindan, and D. Estrin. Diversity in smartphone usage. In *Proc. of ACM MobiSys*, pages 179–194. ACM, 2010.
7. V. Glass, S. Stefanova, and R. Dibelka. Customer price sensitivity to broadband service speed: What are the implications for public policy? In S. Sen, C. Joe-Wong, S. Ha, and M. Chiang, editors, *Smart Data Pricing*. Wiley, 2014.
8. S. Ha, S. Sen, C. Joe-Wong, Y. Im, and M. Chiang. TUBE: Time-dependent pricing for mobile data. *Proc. of ACM SIGCOMM*, 42(4):247–258, 2012.
9. J. Huang, F. Qian, Y. Guo, Y. Zhou, Q. Xu, Z. M. Mao, S. Sen, and O. Spatscheck. An in-depth study of LTE: Effect of network protocol and application behavior on performance. In *Proc. of ACM SIGCOMM*, pages 363–374. ACM, 2013.
10. J. Huang, F. Qian, Z. M. Mao, S. Sen, and O. Spatscheck. Screen-off traffic characterization and optimization in 3G/4G networks. In *Proc. of ACM IMC*, pages 357–364. ACM, 2012.
11. J. Huang, Q. Xu, B. Tiwana, Z. M. Mao, M. Zhang, and P. Bahl. Anatomizing application performance differences on smartphones. In *Proc. of ACM MobiSys*, pages 165–178. ACM, 2010.
12. Y. Im, C. Joe-Wong, S. Ha, S. Sen, T. T. Kwon, and M. Chiang. AMUSE: Empowering users for cost-aware offloading with throughput-delay tradeoffs. In *Proc. of IEEE INFOCOM*, pages 435–439. IEEE, 2013.
13. G. Maier, F. Schneider, and A. Feldmann. A first look at mobile hand-held device traffic. In *Passive and Active Measurement*, pages 161–170. Springer, 2010.
14. A. Rahmati, C. Tossell, C. Shepard, P. Kortum, and L. Zhong. Exploring iPhone usage: the influence of socioeconomic differences on smartphone adoption, usage and usability. In *Proc. of MobileHCI*, pages 11–20. ACM, 2012.
15. S. Sen, C. Joe-Wong, S. Ha, J. Bawa, and M. Chiang. When the price is right: Enabling time-dependent pricing of broadband data. In *Proc. of SIGCHI*, pages 2477–2486. ACM, 2013.
16. S. Sen, C. Joe-Wong, S. Ha, and M. Chiang. Incentivizing time-shifting of data: A survey of time-dependent pricing for Internet access. *IEEE Communications Magazine*, 50(11):91–99, 2012.
17. S. Sen, C. Joe-Wong, S. Ha, and M. Chiang. Smart data pricing (SDP): Economic solutions to network congestion. In H. Haddadi and O. Bonaventure, editors, *Recent Advances in Networking*, pages 221–274. ACM SIGCOMM, 2013.
18. S. Sen, C. Joe-Wong, S. Ha, and M. Chiang. A survey of smart data pricing: Past proposals, current plans, and future trends. *ACM Computing Surveys*, 46(2):15, 2013.
19. J. Sommers and P. Barford. Cell vs. WiFi: On the performance of metro area mobile connections. In *Proc. of ACM IMC*, pages 301–314. ACM, 2012.
20. O. Tipmongkolsilp, S. Zaghloul, and A. Jukan. The evolution of cellular backhaul technologies: Current issues and future trends. *IEEE Communications Surveys & Tutorials*, 13(1):97–113, 2011.
21. Q. Xu, J. Erman, A. Gerber, Z. Mao, J. Pang, and S. Venkataraman. Identifying diverse usage behaviors of smartphone apps. In *Proc. of ACM IMC*, pages 329–344. ACM, 2011.
22. J. Zander and P. Mähönen. Riding the data tsunami in the cloud: Myths and challenges in future wireless access. *IEEE Communications Magazine*, 51(3):145–151, 2013.