Mobile Internet Usage and Usage Based Pricing

Jeffrey Prince and Shane Greenstein*

January 2019

Abstract

We provide both a theoretical and empirical analysis of mobile Internet usage in the presence of usage based pricing in the form of data caps. We begin by building a simple model of mobile Internet usage, grounded in prior empirical work on home Internet usage. Within the model, we generate empirical predictions for mobile usage, particularly in response to usage based pricing, and we identify conditions where usage based pricing (via a data cap) is not revenue enhancing for providers. Using data on mobile Internet usage of thousands of individuals, we provide some of the first analyses linking mobile usage to key demographics such as income. We find a reverse-U relationship between mobile Internet usage and income – notably different than the monotonically declining relationship found on home devices. We also find a largely monotonically increasing relationship between income and usage intensity, measured by number of page views in a session. Combined with our model, these findings suggest data caps are particularly binding on low-income users and that the use of caps by providers is more likely driven by cost than by revenue enhancement. They further suggest that price discrimination strategies may be more effective in terms of revenue generation if tied to usage intensity rather than duration.

---

* Indiana University, Department of Business Economics and Public Policy, Kelley School of Business (jeffprin@indiana.edu); and Harvard Business School, Department of Technology and Operations Management (sgreenstein@hbs.edu). We thank the Harvard Business School for funding. Yanhao Wang provided excellent research assistance. We are responsible for all errors.

1. Introduction

Little more than a decade after their introduction into the commercial mainstream, mobile usage has claimed a significant share of Americans' time and attention. For example, as of 2019, time spent on mobile devices (70% of which is on smartphones) exceeded time spent watching television for the first time (MarketWatch, 2019). Further, a large proportion of time spent on a smartphone is in the form of Internet traffic; in fact, by 2017 over half of global Internet traffic was generated through mobile phones (Statista, 2019). This rapid emergence of mobile Internet usage as a prominent share of consumer attention makes the determinants and composition of such usage a subject of keen interest for a wide and diverse set of entities, including advertisers, Internet service providers (ISPs), content generators, and policymakers. Nonetheless, possibly due to limited data access, there still exists little empirical analysis of smartphone Internet usage beyond broad measures, such as total usage.

In this paper we apply a simple but classic question to smartphone pricing, namely, whether supply behavior is sustainable in the face of demand behavior by smartphone owners. In particular, we investigate whether user actions support the strategic use of price discrimination via a data cap, as most straightforward models would imply. Though it is well known that certain demand elasticities must take on certain values along some observable dimension to support price discrimination, most economic analysis merely assumes those conditions hold rather than checks whether they do. There is insight to be gained from a conversation between theory and data. With the aid of extensive data on user demographics and usage, we can look for numerous features consistent with demand supporting price discrimination in mobile usage.

A summary of our approach is as follows. After sketching prior findings on adoption and usage for smartphones and the Internet, and the basics of usage-based-pricing (UBP), we construct a simple theoretical model grounded in the prior findings. We note two interesting implications of this model – one theoretical and one empirical. On the theory side, the model has two key elements, the demand for a device to use data, and the use of it, where the latter can have novel features. In particular, we consider the possibility that high income users have a high willingness to pay for mobile access, but get less value out of high usage levels. That leads to a variance in the willingness to pay for usage between high- and low-income users. In the presence of such demand heterogeneity, we provide conditions for which second degree price discrimination – in the form of UBP via a data cap – is not profit/revenue enhancing. On the empirical side, we provide conditions for which UBP via a data cap would result in downward pressure on low-income usage, hence bending the usage/income curve "downward at the left end." We then take the insights of our model to ComScore data on thousands of smartphone users over the last several months of Internet usage in 2015.

Our empirical results show that smartphone users do not display a monotonic relationship between Internet usage and income. That finding sharply contrasts with the pattern previously seen in home usage of dial-up and broadband (e.g., Goldfarb and Prince, 2008). Instead, mobile Internet usage and income have a hill-shaped relationship. That is, as income rises, mobile Internet usage increases, reaches a peak, and then falls for the highest-income users. Hence, the greatest usage is among middle-income users. This finding – specifically the lack of relatively high usage by low-income users – is consistent with the empirical prediction of our model, where data caps have the most substantial impact on low-income users. To the extent it corroborates our model, this result also suggests that UBP via data caps is not reflective of demand, and is likely driven mostly by cost considerations among carriers, since market conditions are likely consistent with those for which we've shown data caps are not revenue enhancing.

We consider two notable alternatives that would be consistent with our finding but yield different conclusions. First, it could be that the relatively low mean usage of high-income individuals masks their relative prevalence among the highest-volume users. If this were the case, carriers may be using UBP to enhance revenue via high-priced, unlimited data plans meant for high-income, heavy users. However, our quantile analysis – we particularly focus on the 90^{th} percentile – indicates this is not happening.

Second, we consider the possibility that productivity from mobile usage takes a considerable jump when moving from low to middle income, but then tapers when moving into the higher income brackets. If this were the case, it would suggest that the relatively low mobile usage we see by low-income users, as compared to home usage, is not the result of more restrictive UBP but rather differences in relative productivity between home and mobile usage. Again, we argue this alternative is unlikely. We do this by comparing usage rates of a basic "productivity-oriented" usage activity, email, across income levels. This comparison shows no unusual increase in email usage from low to middle income, relative to other income changes, suggesting that the increase in usage from low to middle income is not driven by added productivity. We also note that such a mechanism makes the decline in usage by users with the highest income quite unintuitive. Specifically, it would imply that smartphone usage yields increased productivity as income rises, but relatively less so at the highest income levels.

A notable additional finding is a largely monotonically increasing relationship between intensity of use (measured as page views) and income. We go on to show that this relationship is not being driven by differences in the use of online video. This additional finding suggests that price discrimination strategies may be more effective in terms of revenue generation if tied to usage intensity rather than duration.

Taken together, our results constitute one of the first large scale assessments of online smartphone usage as pertains to basic demographics. The patterns we find are suggestive of a similar relationship compared to what has been found for recent usage of internet home devices (Boik, Greenstein and Prince, 2019), but with a key difference: data caps binding more on low income users. Under these circumstances and through the lens of our theoretical model, it is unlikely that data caps are an activity designed to enhance revenue by exploiting market power over demand. The more likely alternative motivation is cost-driven constraints.

This finding complements the insight of Rogerson (2016) that competition will prevent suppliers from exploiting consumers using price discrimination; our finding suggests that demand is such that, even with market power, price discrimination via data caps is unlikely to be utilized to exploit consumers. In addition, our findings indicate that low-income users are likely the ones being most notably affected by UBP. Lastly, the fact that usage intensity, and not duration, tracks income in a largely monotonic way, suggests that price discrimination according to intensity rather than duration may hold more promise for revenue enhancement – a finding of potential interest to carriers and regulators alike.

2. Internet and Smartphone Adoption and Usage

The growth in usage of smartphones directs attention at a long standing question about usage pricing of wireline and wireless services. Until the most recent decade wireline data access rarely used any forms of usage pricing. In contrast, from the late 1990s onward cellular carriers priced calls on the basis of minutes, and from the early 2000s onward sold data contracts for wireless data based on total usage per month. In other words, though the pricing started in different places, it has acquired more similar forms over time. Yet, it is an open question whether the motivation is similar.

In the present era, wireline carriers have moved to using data caps, sometimes with tiered pricing for bandwidth. This has all the symptoms of price discrimination among different types of demand in a classic Mussa-Rosen sense (Mussa & Rosen, 1978). This behavior was so novel it generated brief attention from the Federal Communications Commission (See FCC, 2013). In contrast, in an era in which 4G supports the smartphone, all wireless devices and access modes are oriented towards data usage. Often this has tiered pricing with overage charges, a form of UBP where price goes up after crossing a usage limit, which is a similar but less severe constraint compared to a hard cap. Is this something different, possibly motivated by cost considerations? Questions about smartphone usage pricing are ungrounded in any data, and, to our knowledge, have not received attention from systematic analysis.

With limited direct analyses of mobile Internet usage, a natural starting point for understanding this type of usage is home Internet usage. Analysis has largely focused on home (i.e., static) devices, such as the home personal computer. A pervasive finding in that literature is a negative relationship between home Internet usage and income, first noted by Goldfarb and Prince (2008), who examined dial-up home usage in the United States. This result has been replicated for broadband (Boik et al., 2019) as well. However, it has yet to be determined whether the negative relationship between usage and income extends to mobile usage as well.

Prior work also has examined home usage in the face of data caps. Nevo et al. (2016) analyze usage data for a set of customers of a single ISP. These users face distinctive three-part tariffs, which impose a shadow value on the price of data as users approach their monthly allowances. Users are sensitive to the charges affiliated with reaching a data cap, but they also endogenously select into capacity consistent with their own use. Variation in user behavior permits an analyst to recover variation in the willingness to pay for broadband. Such willingness-to-pay estimates provide insight into the gaps between private and social incentives to build or upgrade broadband. Nevo et al.'s (2016) estimates suggest the gap is substantial.[1]

Possibly due to limited data access, there still exists little similar empirical analysis of smartphone Internet usage beyond broad measures, such as total usage. Some recent studies have focused on smartphone usage, but have been quite limited in scope (Andone et al., 2016 & Christensen et al., 2016), utilizing small samples and/or limited information on users. Nonetheless, findings from these studies include: females and the young use smartphones more than males and older users (Andone et al., 2016), and longer usage is associated with poorer sleep (Christensen et al., 2016). Smartphone users act in ways that indicate higher search costs due to the smaller size of the screen, and with more sensitivity to geographic considerations, prioritizing local matters (Ghose et al., 2013). Using a panel of Korean smartphone users, Lee (2017) shows that usage substitutes for activities done on other devices where the activities were difficult, making them much easier. Aguiar (2019) utilizes a relatively larger sample, comparing web browsing behavior for desktops and mobile devices for a set of German consumers during a similar time frame to ours. Consistent with smartphones having relatively higher search costs, he finds more concentrated consumption on mobile devices, with less depth of attention. He also

[1] Malone, Nevo, and Williams (2106) also examine the willingness to pay for more bandwidth, based on usage data from one ISP. They focus on the tradeoffs for different ways to approach congestion of networks. They show that peak load pricing along with caching more effectively deals with congestion than does throttling of traffic. (This is McManus, Brian, Aviv Nevo, Zachary Nolan, and Jonathan W. Williams (2018). Steering Incentives and Bundling Practices in the Telecommunications Industry (October 1). NET Institute Working Paper No. 18-12. https://ssrn.com/abstract=3267060).

finds nontrivial limits on substitutability between content on the desktop and mobile, suggesting mobile Internet consumption may have fundamental differences from desktop consumption[2].

There are reasons to question whether and how home Internet usage patterns would extend to mobile. In particular, the very nature of mobile usage allows for usage at a wider range of locations and times, and the smaller screen limits the amount of data, both of which can impact its usefulness for many purposes. Further, compared to home usage, mobile usage in the U.S. has generally experienced substantially greater constraints from wireless suppliers, in the form of stricter data caps on subscription plans and throttled data for heavy users. Hence, one might expect mobile Internet usage to change relative to home usage, but whether it does, and in what direction, remains largely an open question.

Other factors could shape behavior at the time we observe it. By 2016 smartphones are in 75% of households, and 4G networks support use in most of the urban areas of the country. Unlike the 3G networks, which mixed legacy voice with novel data traffic, the 4G networks are entirely data. By this time, the majority of users have gotten far past the initial point of adoption, and novel uses generally are not the determinant of most usage. If anything, by 2016 the key factor driving usage is growth on the intensive margin by existing users, typically for video and streaming. In this era there are many reports of growing data traffic necessitating more antennae and alleviation of congestion as a major cause of pricing.

We note how these observations can motivate this study at an intuitive level. On the one hand, mobile usage has become a common part of the household budget, and has all the typical

[2] Although inconclusive, possibly due to a smaller sample size, the relationship Aguiar (2019) finds between duration and income for mobile devices is somewhat similar to the one we show for the U.S., providing some corroborating evidence of our findings outside the U.S.

properties of a normal good. Higher income households adopt with higher propensity, consistent with more usage, and greater willingness to pay for data.[3] On the other hand, "more usage" in the modern context means more time-intensive activities. Higher income households generally have a higher opportunity cost of time, and, therefore, if mobile usage resembles wireline usage (Boik et al., 2019), high income-consumers will consume less time-intensive activities. Two margins will shape use, the pricing of adoption and usage, and the value of time, and these do not necessarily push in the same direction. Empirical data can illuminate the interaction of these margins for mobile use and allow us to contrast their roles vis a vis wireline use.

3. A Model of Mobile Internet Usage

In this section, we present a simple model of mobile usage. Motivated by prior empirical work on online usage, we consider the case of two user types, H and L, where H types generally get higher usage value but reach a satiation point relatively more quickly compared to L types. Using this model, we solve for conditions under which a monopoly provider, facing constant marginal costs, would find price discrimination via a data cap to be profit enhancing. We also provide a simple empirical prediction for the case where the provider offers a capped plan that is purchased by L types but not H types[4].

[3] https://www.pewInternet.org/fact-sheet/Internet-broadband/, "Who has Home Broadband?" sort by income.

[4] Note that our model allows for data caps, which implicitly have the price of overage being infinite. As noted above, mobile plans with usage limits generally allow usage in excess of the limit with overage charges. However, these overage charges generally are high enough that, for a given amount of usage in excess of a plan's limit, consumers face lower charges by switching to a plan with a higher limit rather than staying in the more limited plan and paying the overage charge (e.g., $15/MB over the limit for Verizon in 2019: https://www.verizonwireless.com/plans/).

The model is as follows. A monopolist provider sells mobile subscription usage plans to a set of consumers at constant marginal cost, assumed to be zero for ease of exposition. As they don't impact the solution, we ignore fixed costs, thus eliminating the distinction between profit and revenue maximization. Each consumer is one of two possible types, H or L, where there are $N_H > 0$ consumers that are type H and $N_L > 0$ consumers of type L. Let the inverse demand for mobile usage of H types be represented by continuous, differentiable function f(.) and the inverse demand of L types be represented by continuous, differentiable function g(.). We make the following assumptions concerning each type's inverse demand functions:

1. $f(0) > 0$ & $g(0) > 0$
2. $f' < 0$ & $g' < 0$
3. $f^{-1}(0) < \infty$ & $g^{-1}(0) < \infty$
4. $g^{-1}(0) > f^{-1}(0)$
5. $\int_0^{f^{-1}(0)} f(x)dx > \int_0^{g^{-1}(0)} g(x)dx$

The above assumptions imply the following. First, both types get positive surplus from at least some usage (Assumption 1). Next, the inverse demand for both types is downward sloping (Assumption 2). Assumption 3 implies both types have finite satiation points, that is, both types have a level of usage such that, beyond this level, they get no additional value. Next, we have that the satiation point of usage for L types is greater than the satiation point of usage for H types (Assumption 4). Put another way, this assumption implies that, with no constraints and no usage price, L types would have more mobile usage than H types. Lastly, we have H types attaining a higher total surplus at their satiation point than L types (Assumption 5). This means that the area under the inverse demand for H types (and above the X-axis of usage) is greater than the area under the inverse demand for L types. Figure 1 provides an illustration of linear inverse demands that satisfy the above assumptions. This figure is only illustrative for the general case we initially consider, but has a more literal application when we subsequently consider the specific case of linear inverse demands[5].

[Figure 1 about here]

We note here that the first three assumptions are quite standard for inverse demands. The last two, Assumptions 4 and 5, are rooted in observed differences between high-income and low-income Internet users. Specifically, they are rooted in the basic observation that for online usage, high-income individuals/households tend to be higher on the extensive margin (adoption of home Internet, smartphone) but lower on the intensive margin, at least for home Internet (limited evidence regarding mobile usage prior to this paper).

We consider two possible offerings by the monopolist. First, it can offer two plans, one with a cap and one without, at possibly differing prices (price discrimination option). Second, it just offers one unlimited plan at a single price.

We begin by solving for the optimal prices and maximum revenue the firm can attain in both scenarios, where in the price discrimination scenario, the firm sells a capped plan to low types and an unlimited plan to high types[6]. Consider first the price discrimination scenario. The firm solves the following problem:

---

[5] Note that our model materially differs from that of Anderson and Dana (2009) with endogenous quality choice. Specifically, it is the case in our model that the increase in value with improved quality (higher cap) is lower for H types than L types; they assume the opposite. They also assume that H types always get more incremental value from increased quality, whereas our model has that true only for low cap levels. Our differing assumptions mean it is not a guarantee that firms can sort H types from L types, but we limit our solution and insights to the cases where they can.

[6] Note that, under our assumptions, the comparison is trivial if instead we have the monopolist sell a capped plan to H types and unlimited plan to L types. In such a scenario, the monopolist cannot charge more than the maximum willingness-to-pay by L types for unlimited usage for the unlimited plan (or else L types won't buy it) or for the

$$Max_{P_U,P_C,C} N_H P_U + N_L P_C$$

Subject to:

1) $P_U \leq \int_0^{f^{-1}(0)} f(x)dx$ (High types buy)
2) $P_C \leq \int_0^{C} g(x)dx$ (Low types buy)
3) $\int_0^{f^{-1}(0)} f(x)dx - P_U \geq \int_0^{C} f(x)dx - P_C$ (High types buy unlimited)
4) $\int_0^{C} g(x)dx - P_C \geq \int_0^{g^{-1}(0)} g(x)dx - P_U$ (Low types buy capped)

where P_U is the price of an unlimited plan, P_C is the price of a capped plan, and C is the cap, or maximum allowed usage, for a capped plan.

Solving this problem leads to the following proposition:

Proposition 1: The solution to the price cap problem with H types buying unlimited and L types buying a capped plan is one of the following: A) $P_C = \int_0^{C} g(x)dx$, $P_U = \int_0^{g^{-1}(0)} g(x)dx$, with C implicitly defined by $\int_0^{f^{-1}(0)} \mathrm{f(x)dx} - \int_0^{g^{-1}(0)} \mathrm{g(x)dx} = \int_0^{C} \mathrm{f(x)dx} - \int_0^{C} \mathrm{g(x)dx}$. Revenue is $N_L \int_0^{C} g(x)dx + N_H \int_0^{g^{-1}(0)} g(x)dx$. This is the solution when constraints 3 and 4 bind. B) $P_C = \int_0^{C} g(x)dx$, $P_U = \int_0^{f^{-1}(0)} f(x)dx + \left[\int_0^{C} g(x)dx - \int_0^{C} f(x)dx\right]$, with C implicitly defined by $(N_H + N_L)g(C) = N_H f(C)$. Revenue is $N_L \int_0^{C} g(x)dx + N_H \left(\int_0^{f^{-1}(0)} f(x)dx + \left[\int_0^{C} g(x)dx - \int_0^{C} f(x)dx\right]\right)$. This is the solution when constraint 3 binds but constraint 4 does not.

capped plan (or H types will buy the unlimited plan). Hence, for this scenario, the monopolist at best can attain revenue equal to the maximum willingness-to-pay of L types multiplied by the total number of consumers. It can always attain this revenue level selling just an unlimited plan at the same price, meaning this type of price discrimination is trivially weakly inferior, in terms of revenue, to selling a single unlimited plan.

Proof: See Appendix.

In words, our proposition states that, if the solution is such that incentive compatibility for both types (constraints 3 and 4) is binding, the price for the capped plan is the area under inverse demand for L types up to the cap (i.e., the most L types would pay for the capped plan), and the price for the unlimited plan is the total surplus of L types with unlimited usage. As shown in the proof, the only other viable solution is one where incentive compatibility for H types (constraint 3) is binding, but not for L types. Here, we get the same price for the capped plan (area under inverse demand for L types up to the cap), but now the price for the unlimited plan is the total surplus of H types with unlimited usage, plus the difference in surplus between L types and H types when usage is at the capped level (and this difference must be negative).

Now, consider the single offering scenario, where the monopolist only offers an unlimited plan and must set its price. The monopolist then simply picks the larger revenue generator between:

i) Sell only to high types, price at $P = \int_0^{f^{-1}(0)} f(x)dx$ and earn revenue of $N_H \int_0^{f^{-1}(0)} f(x)dx$

ii) Sell to both types, price at $P = \int_0^{g^{-1}(0)} g(x)dx$ and earn revenue of $(N_H + N_L) \int_0^{g^{-1}(0)} g(x)dx$

In words, scenario i) has the monopolist charging the maximum that H types would pay with unlimited usage, which would price out the L types from the market. In scenario ii), the monopolist prices as high as possible so that both types will buy, which is the maximum that L types would pay with unlimited usage.

With solutions to both the single offering and two-offering problems in hand, we can assess when offering two plans, one capped and one unlimited, will enhance revenues. First, note that if the solution to the two-offering problem is A) in the above proposition, revenues are less than or equal to what materialize in option ii) for the single-offering problem. Hence, the monopolist can always do at least as well in terms of revenues by offering just the single unlimited plan.

The more complicated comparison occurs when the solution to the two-offering problem is B) in the above proposition. In this circumstance, it is possible for the price discrimination strategy to improve revenue over a single-plan offering. However, we can characterize conditions for the market and/or preferences that preclude this from happening. Specifically, price discrimination will *not* enhance revenues when: a) the number of L types is large and/or b) H types get a relatively high proportion of their total possible surplus after a relatively low (compared to their satiation point) level of usage[7].

We provide a more formal discussion of the above claim in the Appendix; however, we highlight its intuition here. Consider the case where the number of L types is large. First, note that as N_L increases, a monopolist making a single offering will choose to sell to both groups, since revenue is strictly increasing in N_L for option ii) but flat for option i). Next, under the price discrimination offerings, note that L types pay a lower price than under the single unlimited offering. Hence, revenue on L types is lower under price discrimination than under a single

[7] We note that this second condition aligns with Anderson and Dana's (2009) increasing percentage differences condition. The first condition expands on Anderson and Dana's (2009) findings; they show that for high numbers of L types, the firm will offer only one product, but they do not assess the impact on revenues from increasing L types conditional on still offering two products as we do here.

offering, meaning the relative revenue attained under price discrimination is less as the number of L types grows.

Consider now the case where H types get a relatively high proportion of their total possible surplus after a relatively low (compared to their satiation point) level of usage. In this scenario, it becomes difficult for a price discriminating monopolist to prevent H types from buying a cheaper, capped plan. The optimal price for the capped plan is the surplus L types receive from a capped plan; therefore, to prevent H types from also buying the capped plan, the monopolist must lower the price for the unlimited plan. In sum, this second scenario puts downward pressure on the price for the unlimited plan for a price discriminating monopolist, and thus puts downward pressure on revenues.

While there are conditions on the market sizes and preferences such that price discrimination can be revenue enhancing under our assumptions, we note that such conditions vanish when inverse demand is linear (and we make a standard second-order sufficiency assumption), the case we illustrate in Figure 1. We state this fact in Proposition 2 below.

Proposition 2: In addition to Assumptions 1-5 above, assume differentiable and continuous functions f(x) and g(x) are both linear, and that $(N_H + N_L)g'(x) < N_H f'(x)$. Then, revenue from selling a single, unlimited plan is always at least as great as revenue from selling two plans, one with a cap and one without, at possibly different prices.

Proof: See Appendix.

To conclude this section, we make a simple empirical prediction for the case where firms offer capped plans that are generally purchased by L types, who are predominantly low income. The simple prediction is that the monotonic, declining relationship between income and usage we observe in scenarios where usage is largely unlimited (i.e., home Internet) will be weaker in magnitude, and possibly even disappear altogether. The intuition here is straightforward: if caps are binding on low-income consumers, we expect to observe a suppression of their usage volume; it will be less than it otherwise would have been. Consequently, whereas usage among low-income consumers would otherwise be higher than high-income consumers, with this constraint in place, the difference will be smaller and possibly even disappear. It could even become negative. The outcome is an empirical question.

4. Data

We obtained data on the online activity of individual smartphones, and some iPads, from ComScore for the last three months of 2015. We observe each device for one month, meaning we have three pooled one-month panels. An observation is a session, consisting of a continuous visit to a website (via an app or browser) on the smartphone. The information collected includes the sites visited on the device, how much time was spent at each site, and the number of pages visited within the site. We also observe several corresponding demographic measures for the device user, including income, sex, ethnicity, age, whether the user has children, and household size.

We first define a unique session by *device id*×*log-in time*×*duration*×*website id*. There are 25,846,679 sessions in the raw data with 242 duplicates showing up in pairs (92 in Oct, 86 in

Nov and 58 in Dec). There is no significant difference in browsing pattern or user demographics with the rest of the data (i.e., duplicates probably caused by a data recording error), so we simply drop them. We then proceed by excluding outlier sessions. Specifically, we drop any session by users who are over 100-years old or live in unknown region, and any session with duration of over 6 consecutive hours or involving more than 100 web pages. This leaves us with 24,955,632 unique and cleaned sessions left in the sample. These are sessions by 12,270 distinct devices visiting 112,359 distinct domains. There are 8,822, 8,170 and 8,466 distinct devices in Oct, Nov and Dec, respectively. Devices enter and exit the data each of the three months, such that there are: 6,674 devices showing up both in Oct and Nov, 6,242 both in Nov and Dec, 5,533 both in Oct and Dec, and 5,315 in all three months.

For some of our supplemental analysis, we augment the ComScore data with data on website categories. We create three categories relevant to that analysis: Email; News and Weather; and News, Weather and Sports. We then manually went through each of the top 3,000 sites and apps in our data (by usage) and identify which, if any, of these categories each site or app belongs to. There are 53 sites and apps categorized as Email (e.g., Gmail, Outlook, YahooMail), 103 sites and apps categorized as News and Weather (e.g., CNN, BBC, NYTIMES, WSJ, YahooWeather, Weather.com), and 152 sites and apps categorized News, Weather, and Sports (in addition to News and Weather domains, we have 49 Sports sites and apps, e.g., ESPN, NFL.com, NBA.com).

Table 1 summarizes the session data by device platform and access method. Table 2 summarizes the demographic information of our device users.

(Tables 1 and 2 about here)

As can be seen from the above tables, our data skew toward Apple devices (iPhone and iPad), and toward younger, female, and lower income individuals. The access method is quite balanced between apps and browsers.

In Tables 3 and 4, we break down monthly usage and page views by our demographic measures. Here, we see only minor differences across ethnicity (Hispanic vs. not Hispanic), sex, and household size, and somewhat greater usage by individuals with children. We see largely declining usage with age, and a roughly hill-shaped usage pattern with income. With regard to intensity, as measured by page views, we see little difference across household size and the presence of children. We see modestly higher intensity by women vs. men and by Hispanics vs. non-Hispanics. Lastly, we see an increasing pattern with income and a U-shaped pattern with age.

[Tables 3 and 4 here]

The statistics in Tables 3 and 4 are unconditional. As a point of reference, in Figures 2a and 2b, we compare the unconditional relationship between usage and income from Boik et al., 2019, which uses weekly home device data from 2008 and 2013, to our unconditional relationship between usage (and page views) and income on a smartphone or iPad in 2015. The Boik et al. (2019) relationships mirror results from prior work, including Goldfarb and Prince (2008); however, the relationship we find concerning usage for smartphones and iPads stands in contrast, with an interior maximum point and a relatively flat shape until the highest incomes.

(Figures 2a and 2b about here)

5. Empirics and Findings

The summary statistics from Section 4 provide suggestive unconditional patterns, particularly with regard to income. In this section, we build on these findings with a battery of regression analyses. For our analyses, the dependent variables are either usage time or page views, and the independent variables consist of our demographic measures. Hence, our regression model is some variation of:

$$(1)\ Weekly\_Usage_i = \beta_0 + \beta_1 Income_i + \beta_2 Age_i + \beta_3 FamilySize_i + \beta_4 Children_i + \beta_5 Hispanic_i + \beta_6 Gender_i + \varepsilon_i$$

where Income, Age, and Family Size are categorical variables.

The parameter estimates for (1) for both usage and page views are in Table A1 in the Appendix. Figure 3 plots the parameter estimates for each income category from both regressions. Here we see that the unconditional patterns from Section 4 still hold after controlling for a range of other demographic variables. Namely, usage and income have a hill-shaped relationship, and page views increase with income, controlling for age, household size, sex, ethnicity and children[8].

[Figure 3 about here]

[8] Our data do not distinguish mobile usage via WiFi vs. the mobile network. However, our conclusions about the impact of data caps only rely on enough mobile network usage for caps to conceivably bind; otherwise, one must consider an alternative explanation for the change in pattern with income relative to home broadband, two of which we rule out here.

Our robust finding of usage eventually declining as income drops into the lowest categories is consistent with our basic empirical prediction from Section 3, namely that data caps pose a relatively stronger restriction for these income groups[9]. Furthermore, as suggested by smartphone adoption patterns and relatively unbounded home usage patterns, if it's the case that preferences for high vs. low income users resemble those presented in Section 3, our analysis in that section indicates these data caps likely are not a price discrimination scheme to enhance revenue. As we've shown, revenue enhancement via price discrimination would require non-linear inverse demands and either relatively quick exhaustion of value by high (-income) types or low numbers of (low-income) types, which seems unlikely in this market. Hence, the data caps are likely cost-driven UBP that prove binding for low-income mobile usage.

Of course, it could be the case that the simplified expression of preferences we use in Section 3 misses key features that could generate our empirical findings and lead to different conclusions. We consider two in particular. First, it could be that our analysis, which focuses on means, is missing important information in the tails. Specifically, while high-income individuals may have relatively low mean usage levels on mobile devices, this may mask their relative prevalence among the highest-volume users. If this were the case, carriers may be using UBP to enhance revenue via high-priced, unlimited data plans meant for high-income, heavy users.

We test this first alternative possibility by running quantile regressions for monthly usage, using the 50th, 80th, and 90th percentiles. The coefficients on income from these three quantile regressions are in Figure 4 (full regression results are in Table A2). All three quantile

[9] We note that a more direct approach toward demonstrating differences across income in caps binding would be to examine whether there are differences in end-of-month usage patterns. However, our data do not contain information on contract parameters, such as the start and end of usage metering. We did try looking at end-of-calendar-month patterns, but unsurprisingly (since contracts need not align with calendar months), there were no clear patterns.

regressions show a quite similar pattern to what we saw in our OLS regressions – usage quantiles that are hill-shaped in income, with the highest income group having the lowest cutoff for each of the three quantiles we consider. Particularly given our findings for the 90th percentile, these results indicate that it is unlikely that UBP for mobile is targeting heavy users with high income for revenue enhancement. We also note that our finding with regard to page views (increasing with income) also largely persists for these quantiles.

[Figure 4 about here]

The second variant of preferences we consider is the possibility that productivity from mobile usage takes a considerable jump when moving from low to middle income, but then tapers when moving into the higher income brackets. If this were the case, it would suggest that the relatively low mobile usage we see by low-income users, as compared to home usage, is not the result of more restrictive UBP but rather differences in relative productivity between home and mobile usage.

To test this second possibility, we compare usage rates of a basic "productivity-oriented" usage activity, email, across income levels, both on the extensive and intensive margins. We report the results Figure 5a (extensive) and 5b (intensive, in terms of usage and intensity). Here we see no unusual increase in email usage from low to middle income, relative to other income changes. This relationship indicates that the increase in usage from low to middle income is not driven by added productivity. Beyond this basic finding, we note that this alternative "productivity-oriented" mechanism makes the decline in usage by users with the highest income quite unintuitive. Specifically, it would imply that smartphone usage yields increased productivity as income rises, but relatively less so at the highest income levels.

[Figures 5a and 5b about here]

We conclude our findings by highlighting an alternative to UBP as a potential vehicle for revenue-enhancing price discrimination. Our findings pertaining to intensity of use (measured as page views), both OLS and the upper quantiles, largely indicate a monotonic relationship with income. To the extent that intensity of use can be easily tracked and priced, its close movement with income makes it a perhaps more promising means of detecting and pricing for higher income compared to UBP.

A possible caveat concerning intensity as an effective price discrimination tool is that the relationship we find is driven by video consumption. That is, low-income individuals simply consume more video on mobile, which may involve fewer page views on average, and this basic difference could be driving the positive relationship we find between page views and income. As a simple test of this caveat, we re-ran our analysis, considering only usage and page views for news and weather, and news, weather, and sports, applications in which we expect video not to be a dominant feature of the application. The OLS and quantile regression results are in Figures 6 and 7, respectively (regression results in Table A2). Here we see that intensity of use continues to have a strong, positive relationship with income for both categories, strongly suggesting that video is not driving this result.

[Figures 6 and 7 about here]

6. Discussion and Conclusions

Using an empirically grounded theoretical model, we identify conditions where mobile service providers would find no revenue benefits from imposing data caps. Then, using data on actual mobile usage, we provide evidence suggestive of data caps largely binding low-income users and an overall relationship between usage and income that aligns with our model basic empirical prediction. Taken together, our analysis suggests that data caps on mobile usage are not motivated by demand. They are likely motivated more by costs than as a revenue-enhancement mechanism.

Our additional analysis helped rule out other alternative mechanisms. We showed that high-income individuals not only have lower average usage but a lower 90th percentile of usage, casting doubt on a firm strategy of catering to a subset of wealthy, very high-volume mobile users. We also showed that a comparison of basic "productivity-oriented" usage, as measured by email, indicates no unusual increase in email usage from low to middle income, relative to other income changes. This additional finding suggests that the increase in usage from low to middle income we find is not driven by added productivity.

Overall, our analyses shed new light on mobile usage patterns and the strategic use of UBP. While it's certainly possible data caps can enhance revenues in other ways, our results suggest that their use as a means to separate purchases of high-income and low-income users likely doesn't serve that purpose but are instead driven by cost considerations. Although there are key differences across the markets in terms of costs, etc., this finding can serve as a useful input as firms and regulators continue to weigh possibilities concerning data caps on fixed line connections. In addition, our finding of a more persistent, monotonic relationship between usage intensity and income provides an interesting new dimension of potential price discrimination for both providers and regulators to consider.

References

Aguiar, L. 2019. “Going Mobile: The Effects of Smartphone Usage on Internet Consumption.” JRC Digital Economy Working Paper.

Anderson, E. and J. Dana. 2009. “When Is Price Discrimination Profitable?” *Management Science*, 55, 980-989.

Andone, I., Blaszklewicz, K., Elbes, M., Trendafilov, B., Markowetz, A., and C. Montag. 2016. “How Age and Gender Affect Smartphone Usage.” *UbiComp/ISWC ’16 Adjunct*.

Bauer, J. and S. Wildman. 2012. “The Economics of Usage-Based Pricing in Local Broadband Markets.” Michigan State University.

Boik, A., Greenstein, S. and J. Prince. 2019. “The Persistence of Broadband User Behavior: Implications for Universal Service and Competition Policy.” Forthcoming at *Telecommunications Policy*.

Christensen, M., Bettencourt, L., Kaye, L., Moturu, S., Nguyen, K., Olgin, J., Pletcher, M., and G. Marcus. 2016. “Direct Measurements of Smartphone Screen-Time: Relationships with Demographics and Sleep.” *PLOS One*, 1-14.

Falaki, H., Mahajan, R., Kandula, S., Lymberopoulos, D., Govindan, R. and D. Estrin. 2010. “Diversity in Smartphone Usage.” *MobiSys ’10*.

FCC. 2013. “Policy Issues in Data Caps and Usage-Based Pricing.” *Open Internet Advisory Committee – 2013 Annual Report.*

Ghose, A., Goldfarb, A., and S.P. Han. 2016. “How is the Mobile Internet Different? Search Costs and Local Activities,” ISR, 24, 3, 613-631.

Goldfarb, A. and J. Prince. 2008. “Internet Adoption and Usage Patterns Are Different: Implications for the Digital Divide.” *Information Economics and Policy*, 20, 1, 2-15.

Hitt, L. and P. Tambe. 2007. “Broadband Adoption and Content Consumption. *Information Economics and Policy*, 19, 3-4, 362-378.

Lee, S. 2017. “Quantifying the Effects of Smartphone Adoption: Digital Device Substitution and Digital Device Expansion,” working paper, University of Washington.

Lyons, D. 2013. “Internet Policy’s Next Frontier: Data Caps, Tiered Service Plans, and Usage-Based Broadband Pricing.” *Federal Communications Law Journal*, 66, 1, 1-44.

MarketWatch. 2019. “For the First Time Ever, Americans Will Spend More Time on Mobile Devices than Watching TV.”

Moore, G. 2002. Crossing the Chasm. HarperCollins Publishers, New York, NY.

Mussa, M. and S. Rosen. 1978. "Monopoly and Product Quality," *Journal of Economic Theory*, 18, 301-317.

Nevo, A., Turner J., and J. Williams. 2016. "Usage Based Pricing and Demand for Residential Broadband," *Econometrica*, 84, 2, 411- 443

Pew Research Center. 2013. "Smartphone Ownership – 2013 Update." Pew Report.

Public Knowledge. 2013. "The Wrong Tool for the Job: Data Caps, Price Discrimination, and Bandwidth Pricing."

Rogerson, W. 2016. "The Economics of Data Caps and Free Data Services in Mobile Broadband." Working paper, Northwestern University.

Statista. 2019. "Percentage of All Global Web Pages Served to Mobile Phones from 2009 to 2018." https://www.statista.com/statistics/241462/global-mobile-phone-website-traffic-share/, accessed June 6, 2019.

Wagner, D., Rice, A. and A. Beresford. 2014. "Device Analyzer: Understanding Smartphone Usage."

Wallsten, S. 2013. "What Are We Not Doing When We're Online," NBER working paper.

Figures

Figure 1
WTP for Mobile Usage by Type

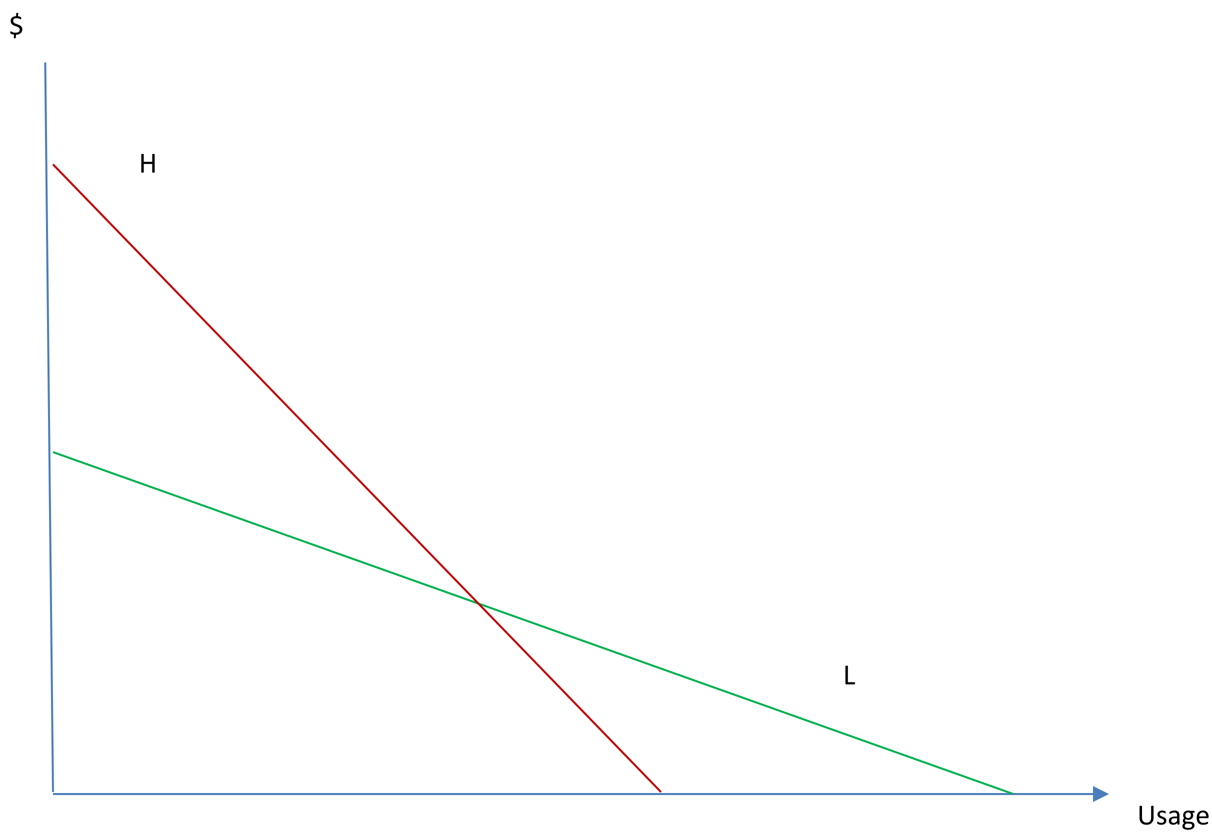

Figure 2a
Weekly Home Internet Usage by Income (Boik et al. 2019)

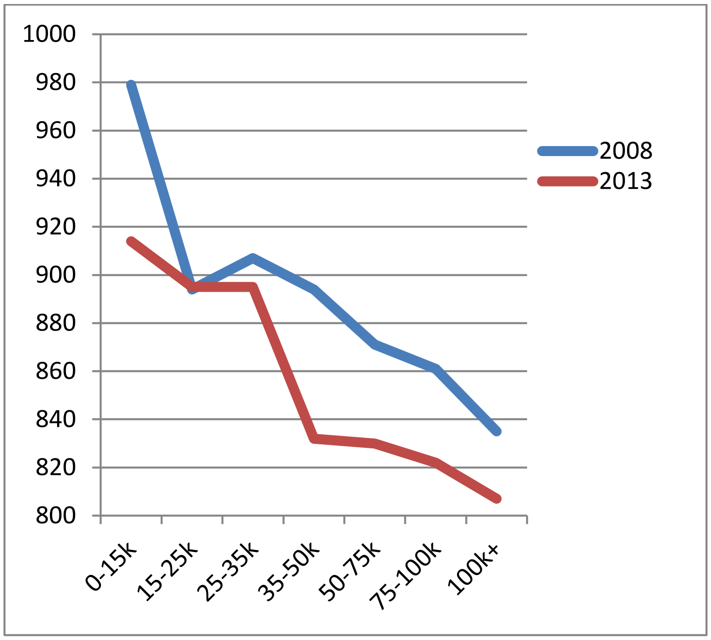

Figure 2b
Average Weekly Usage by Income

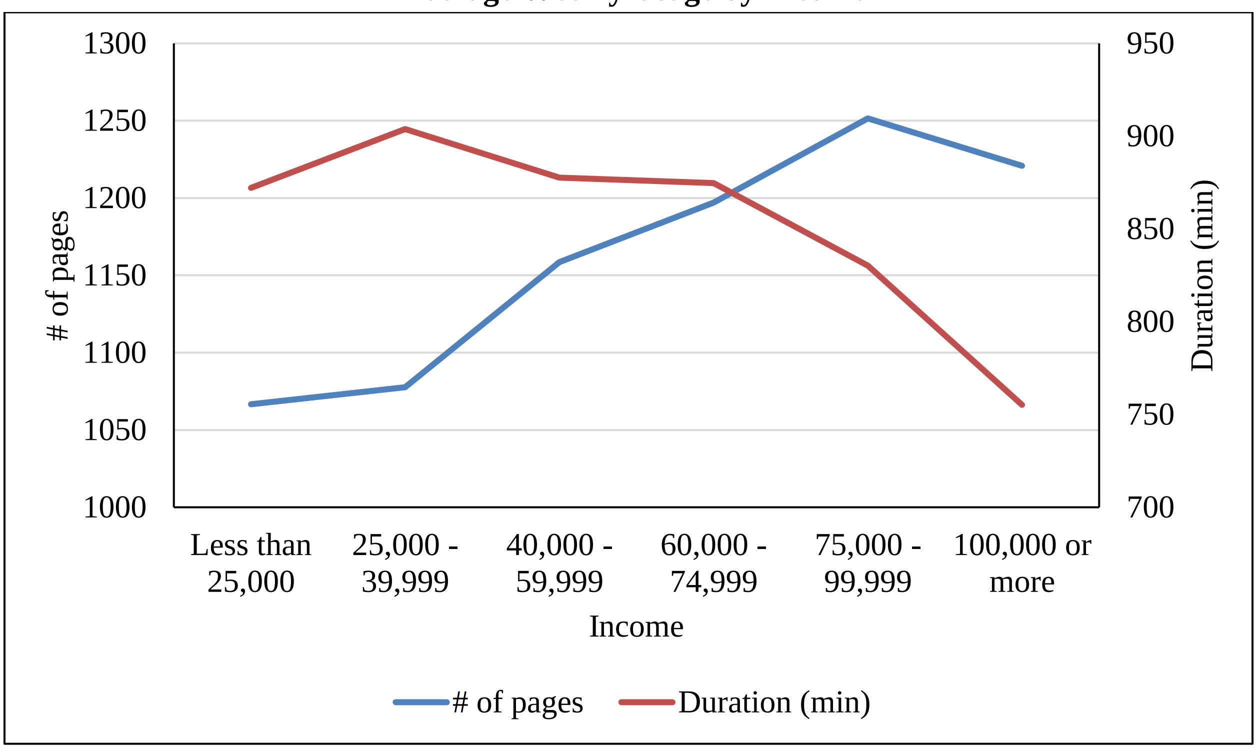

Note: Data are from table 3.

Figure 3: Income Dummy Coefficients from OLS Regressions in Table A1

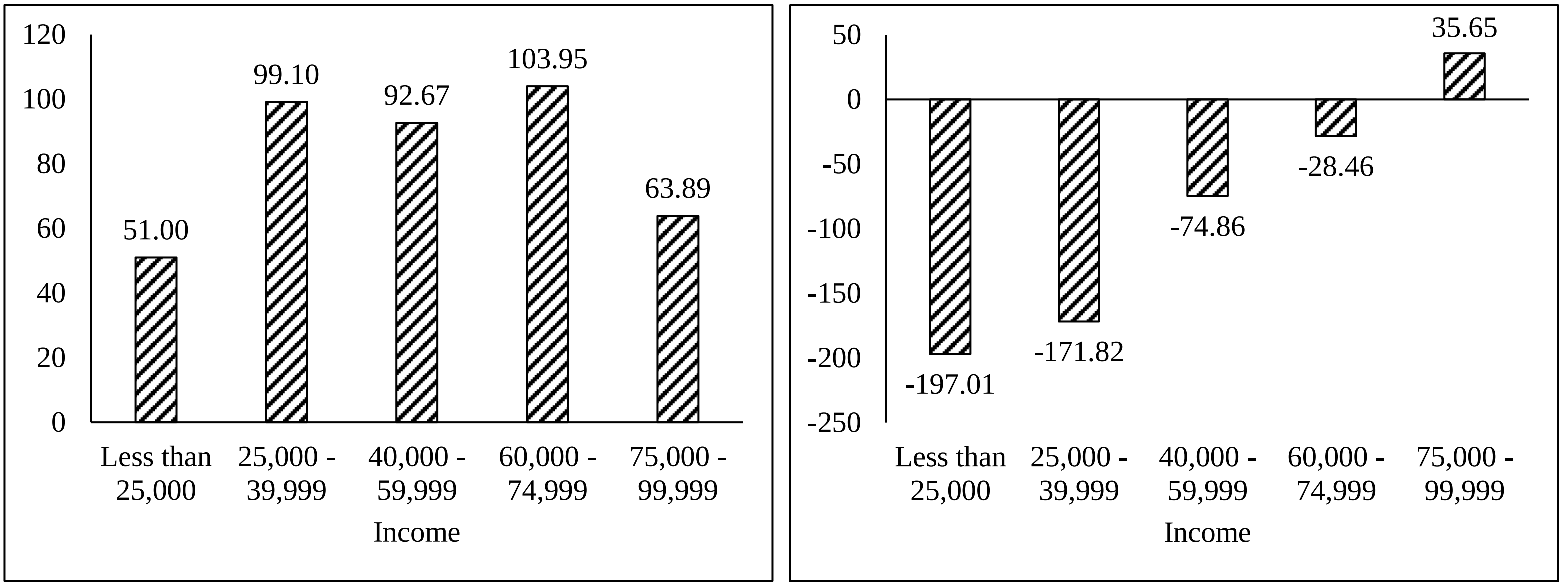

(i) Duration

(ii) # of pages visited

Note: Dummy coefficients are estimated from regression:

$$Avg.Weekly\ Usage = \beta_0 + \beta_1 Income + \beta_2 Age + \beta_3 FamilySize + \beta_4 Children + \beta_5 Ethnicity + \beta_6 Gender + \varepsilon,$$

with both region and week fixed effects controlled. Usage is either duration (min) or #of pages visited. The benchmark dummy is with income of $100,000 or more.

Figure 4
Income Dummy Coefficients from Quantile Regressions

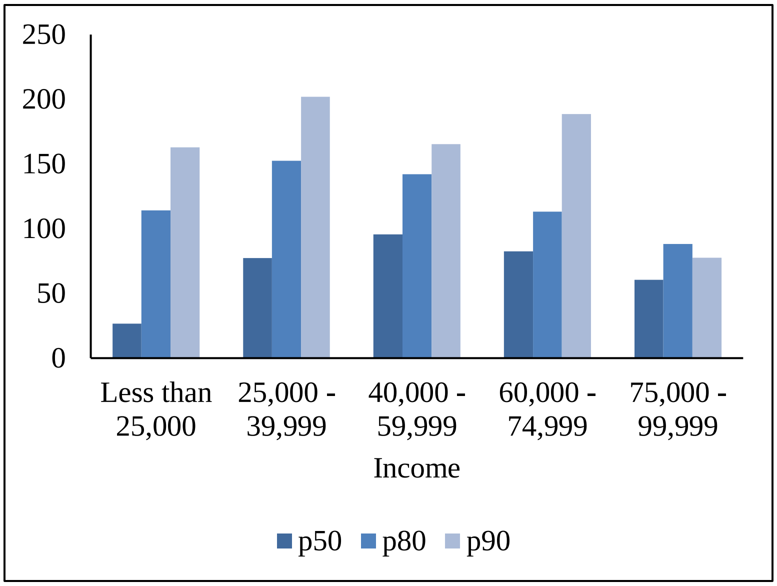

(i) Duration

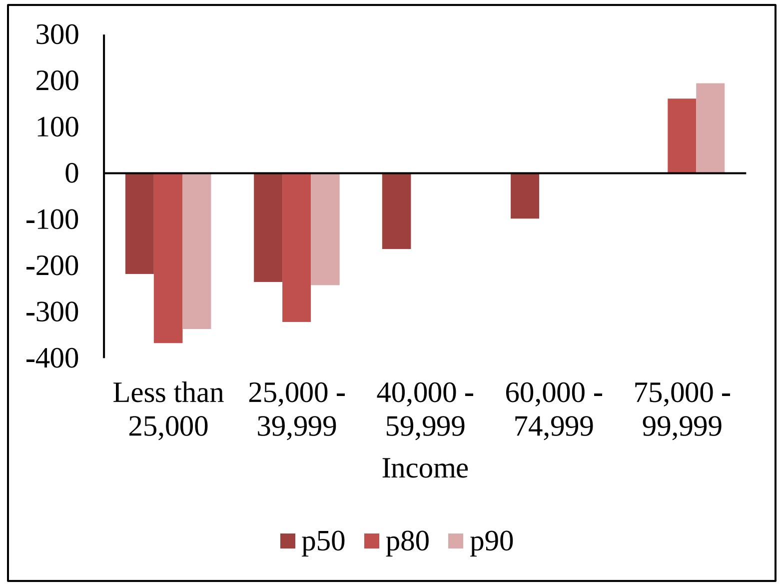

(ii) # of pages visited

Note: Dummy coefficients are estimated from quantile regression:

$$Avg.Weekly\ Usage = \beta_0 + \beta_1 Income + \beta_2 Age + \beta_3 FamilySize + \beta_4 Children + \beta_5 Ethnicity + \beta_6 Gender + \varepsilon,$$

with both region and week fixed effects controlled. Usage is either duration (min) or #of pages visited. Usage is evaluated at 50%, 80% and 90% quantiles. The benchmark dummy is with income of $100,000 or more.

See Table A2 for estimated coefficients.

Figure 5a: **Probability of visiting email websites by income**

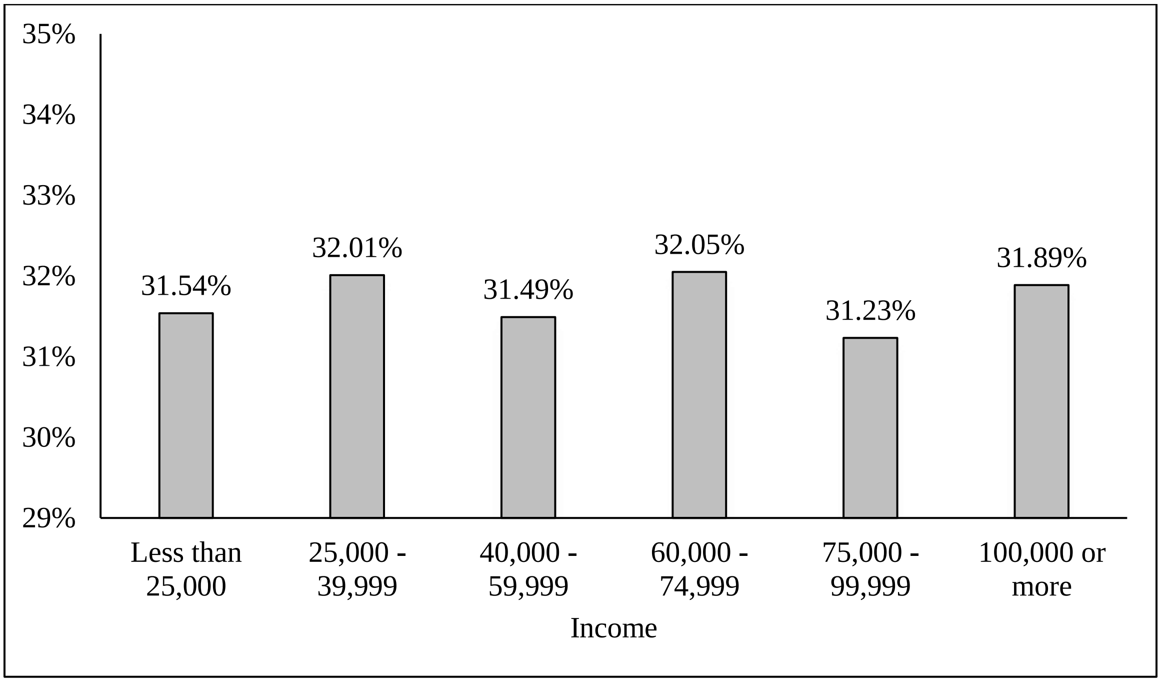

Figure 5b: Average email usage and intensity by income, conditional on using

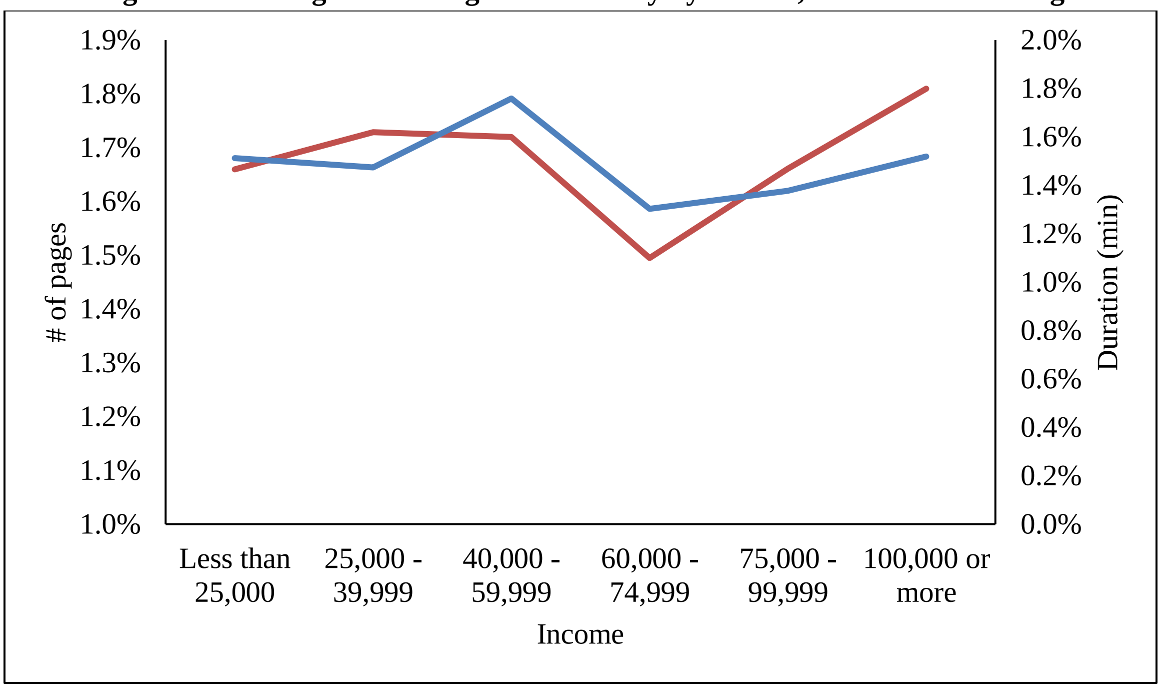

Note: Statistics are calculated for the whole sample. Probability of using email service is measured by the percentage of devices that ever visited email domains throughout the period. Email usage and intensity are measured by the fractions of time spent and pages visited on an email domain by an average device, conditional on the device ever using email service. A list of email service providers are available upon request (e.g., gmail, outlook, etc.).

Figure 6: Income Dummy Coefficients Using OLS for Restricted Subsamples

Top 3000 News & weather

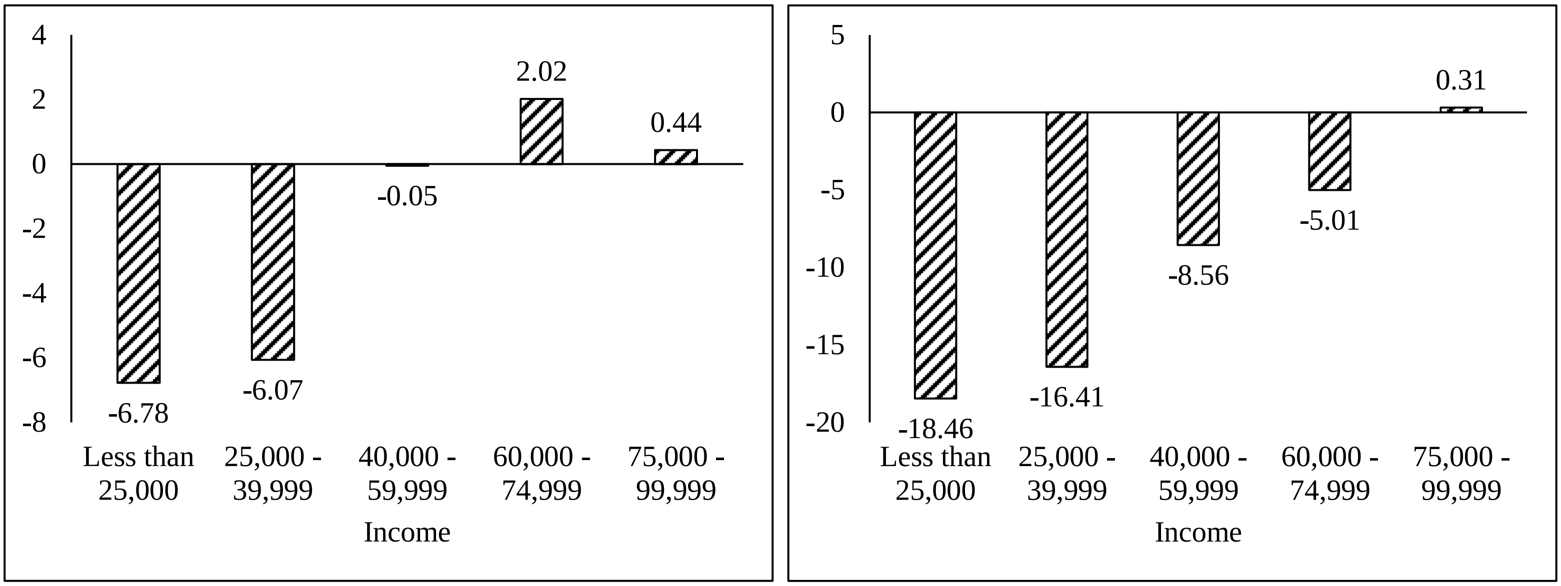

(i) Duration

(ii) # of pages visited

Top 3000 News, weather, and sports

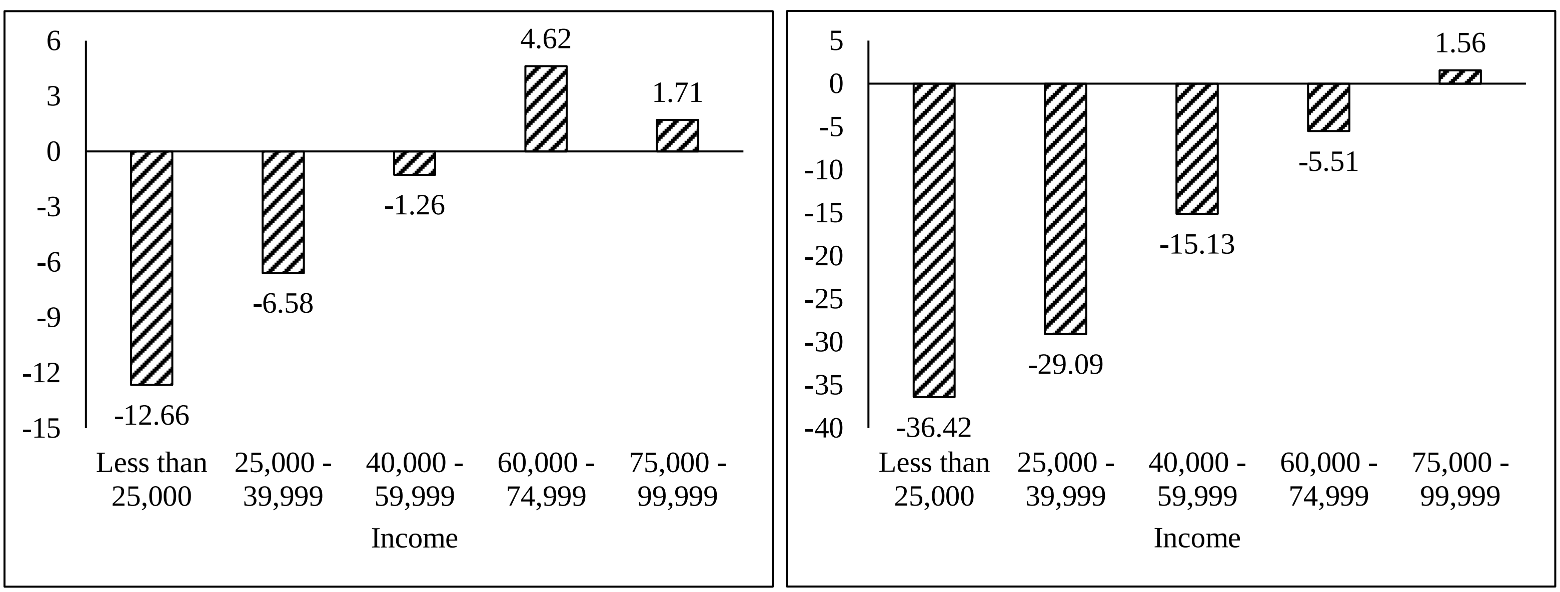

(iii) Duration

(iv) # of pages visited

Note: Based on OLS regressions as in Figure 3, but for sub-samples. See Table A1 for estimated coefficients.

Figure 7: Income Dummy Coefficients Using Quantile Regression for Restricted Subsamples

Top 3000 News & weather

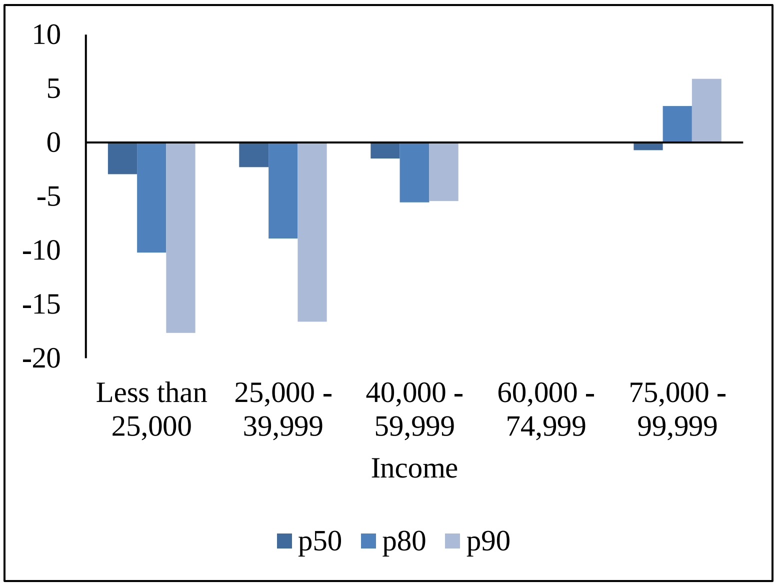

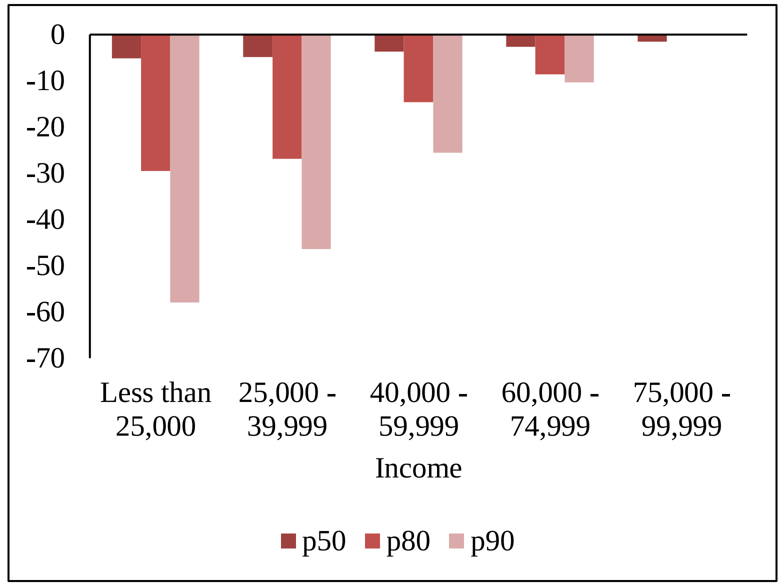

(i) Duration

(ii) # of pages visited

Top 3000 News, weather, and sports

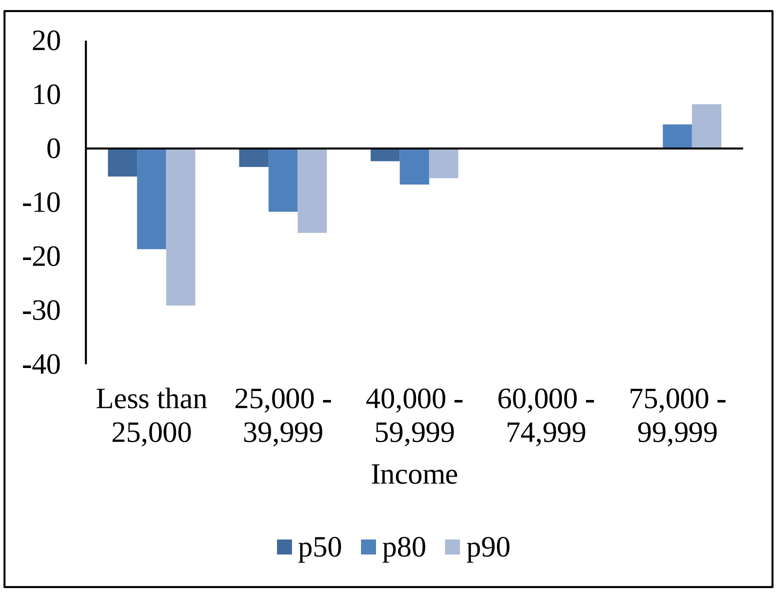

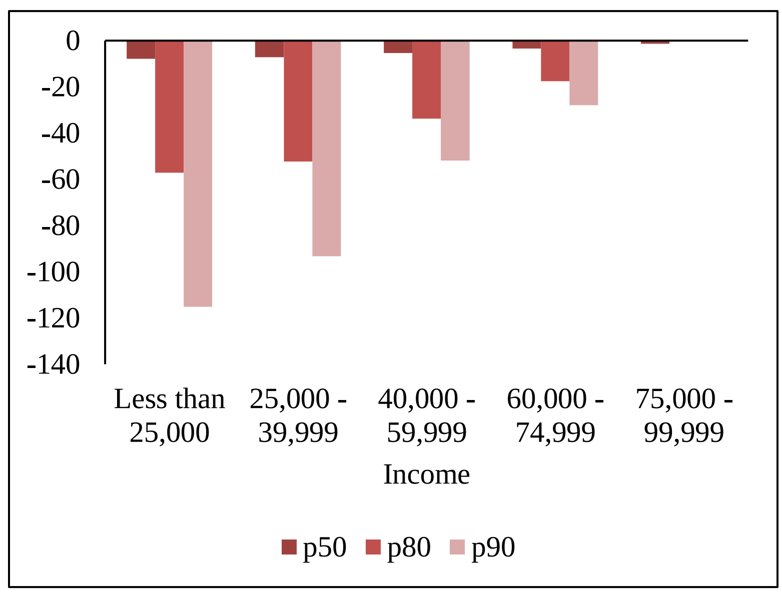

(iii) Duration

(iv) # of pages visited

Note: Based on quantile regressions as in Figure 4, but for sub-samples. See Table A2 for estimated coefficients.

Tables

Table 1
Device Characteristics

	Platform			*Access Method*	
	Android Phone	iPad	iPhone	Mobile App	Mobile Browser
October	2,368,150	1,315,617	5,140,395	4,530,842	4,293,320
(%)	26.84	14.91	58.25	51.35	48.65
November	2,103,145	1,213,548	4,571,053	3,996,959	3,890,787
(%)	26.66	15.39	57.95	50.67	49.33
December	2,494,867	1,238,476	4,510,381	4,151,481	4,092,243
(%)	30.26	15.02	54.71	50.36	49.64
Combined	6,966,162	3,767,641	14,221,829	12,679,282	12,276,350
(%)	27.91	15.1	56.99	50.81	49.19

Note: The unit of observations here is a session to an app or website by a device.

Table 2
Demographic Distributions

	October	November	December	Combined
Income				
Less than 25000	2,605	2,280	2,388	7,273
25000 - 39999	1,861	1,731	1,785	5,377
40000 - 59999	1,710	1,634	1,708	5,052
60000 - 74999	807	777	806	2,390
75000 - 99999	850	818	831	2,499
100000 or more	975	919	939	2,833
Sex				
Female	5,827	5,398	5,708	16,933
Male	2,981	2,761	2,749	8,491
Ethnicity				
Hispanic	1,149	1,025	1,031	3,205
Non-Hispanic	7,659	7,134	7,426	22,219
Age				
18-24	1,779	1,550	1,652	4,981
25-34	2,884	2,612	2,764	8,260
35-44	1,897	1,827	1,878	5,602
45-54	1,329	1,297	1,303	3,929
55-64	666	634	619	1,919
65 and over	253	239	241	733
With children				
No	4,587	4,209	4,318	13,114
Yes	4,221	3,950	4,139	12,310
Family size				
1	1,362	1,221	1,235	3,818
2	2,325	2,156	2,277	6,758
3	1,783	1,650	1,715	5,148
4	1,781	1,689	1,739	5,209
5+	1,557	1,443	1,491	4,491
Total	8,808	8,159	8,457	25,424

Note: The unit of observations here is a distinct device.

Table 3
Average weekly usage by income and age

		# of Pages	Duration (min)			# of Pages	Duration (min)
Income	***Less than 25,000***	7273		***Age***	***18-24***	4981	
	mean	1066.63	872.11		mean	1298.78	937.94
	std.dev	1539.85	747.35		std.dev	1692.23	708.80
	min	0.00	0.00		min	0.00	0.03
	medium	374.00	706.22		medium	495.00	824.28
	max	15184.00	5638.02		max	14999.00	6394.37
	25,000 - 39,999	5377			***25-34***	8260	
	mean	1077.56	903.82		mean	1095.43	958.17
	std.dev	1602.46	743.99		std.dev	1551.79	748.54
	min	0.00	0.07		min	0.00	0.00
	medium	349.00	745.31		medium	379.00	813.26
	max	19697.00	5839.25		max	17059.00	6433.68
	40,000 - 59,999	5052			***35-44***	5602	
	mean	1158.61	877.62		mean	1084.34	860.36
	std.dev	1611.60	720.61		std.dev	1555.04	739.15
	min	0.00	0.13		min	0.00	0.02
	medium	415.00	724.64		medium	387.00	695.60
	max	16737.00	6394.37		max	16737.00	5729.90
	60,000 - 74,999	2390			***45-54***	3929	
	mean	1196.98	874.67		mean	1030.82	728.17
	std.dev	1614.19	746.09		std.dev	1495.56	713.85
	min	0.00	0.15		min	0.00	0.17
	medium	473.00	705.81		medium	423.00	524.73
	max	13992.00	5666.93		max	19697.00	9584.33
	75,000 - 99,999	2499			***55-64***	1919	
	mean	1251.49	830.29		mean	1153.56	632.93
	std.dev	1590.95	690.56		std.dev	1589.03	642.03
	min	0.00	0.10		min	0.00	0.07
	medium	566.00	678.00		medium	509.00	434.10
	max	13860.00	6433.68		max	13625.00	5561.47
	100,000 or more	2833			***65 and over***	733	
	mean	1220.83	755.20		mean	1350.41	613.31
	std.dev	1540.68	694.70		std.dev	1647.55	629.97
	min	0.00	0.03		min	0.00	0.48
	medium	588.00	588.67		medium	716.50	387.65
	max	12210.00	9584.33		max	11302.00	3524.07

Note: Average usage is calculated at the *device*×*week* level.

Table 4
Average weekly usage by sex, ethnicity, family size and structure

	# of Pages	Duration (min)			# of Pages	Duration (min)
Female	16933		***Family size***	***1***	3818	
mean	1149.79	876.88		mean	1111.63	836.23
std.dev	1572.65	719.09		std.dev	1556.73	759.41
min	0.00	0.00		min	0.00	0.00
medium	428.00	727.95		medium	421.00	642.16
max	16737.00	9584.33		max	15184.00	6394.37
Male	8491			***2***	6758	
mean	1104.53	835.51		mean	1134.27	831.56
std.dev	1597.93	754.52		std.dev	1598.74	727.40
min	0.00	0.02		min	0.00	0.07
medium	410.00	644.80		medium	432.00	663.09
max	19697.00	6433.68		max	19697.00	9584.33
Hispanic	3205			***3***	5148	
mean	1178.55	884.56		mean	1176.16	861.76
std.dev	1578.78	745.47		std.dev	1623.67	714.00
min	0.00	0.02		min	0.00	0.02
medium	455.00	724.54		medium	449.00	707.05
max	13750.00	6433.68		max	15061.00	5666.93
Non-Hispanic	22219			***4***	5209	
mean	1128.40	860.01		mean	1122.21	889.23
std.dev	1581.51	729.24		std.dev	1555.85	727.64
min	0.00	0.00		min	0.00	0.04
medium	416.00	699.66		medium	409.00	734.20
max	19697.00	9584.33		max	13860.00	5651.58
No	13114			***5+***	4491	
mean	1138.80	829.20		mean	1121.89	904.51
std.dev	1598.06	735.30		std.dev	1554.36	733.89
min	0.00	0.00		min	0.00	0.03
medium	430.00	653.87		medium	395.00	766.74
max	19697.00	9584.33		max	12823.00	5729.90
Yes	12310		***Total***		25424	
mean	1130.37	899.15		mean	1134.71	863.10
std.dev	1563.18	725.38		std.dev	1581.25	731.34
min	0.00	0.02		min	0.00	0.00
medium	412.00	751.42		medium	421.00	702.54
max	15061.00	5729.90		max	19697.00	9584.33

Note: Average usage is calculated at the *device*×*week* level.

Appendix

Table A1: Regression coefficients for duration and page views

		All websites		Top 3000 News & weather		Top 3000 News, weather, and sports	
		Duration (unit: minute)	# of pages	Duration (unit: minute)	# of pages	Duration (unit: minute)	# of pages
Mean Weekly Usage		863.10	1134.71	23.31	39.25	34.60	60.10
Income	income{<25000}	51.00***	-197.0***	-6.781***	-18.46***	-12.66***	-36.42***
		(7.718)	(17.03)	(0.865)	(1.587)	(1.109)	(2.317)
	income{25000 - 39999}	99.10***	-171.8***	-6.066***	-16.41***	-6.583***	-29.09***
		(7.924)	(17.64)	(0.901)	(1.563)	(1.174)	(2.418)
	income{40000 - 59999}	92.67***	-74.86***	-0.0513	-8.556***	-1.259	-15.13***
		(7.798)	(17.66)	(1.224)	(1.586)	(1.373)	(2.526)
	income{60000 - 74999}	104.0***	-28.46	2.020	-5.010***	4.621***	-5.513*
		(9.405)	(20.95)	(1.319)	(1.876)	(1.549)	(3.010)
	income{75000 - 99999}	63.89***	35.65*	0.436	0.314	1.711	1.558
		(8.886)	(20.46)	(1.015)	(1.843)	(1.336)	(2.921)
Family size	2	10.71	3.920	-3.482***	-4.829***	-3.509***	-4.553**
		(7.235)	(15.43)	(0.902)	(1.437)	(1.065)	(2.170)
	3	-3.956	34.01*	-3.688**	-10.18***	-5.755***	-16.44***
		(8.393)	(18.53)	(1.580)	(1.661)	(1.608)	(2.312)
	4	11.32	-29.37	-6.052***	-14.54***	-6.582***	-19.73***
		(8.965)	(19.28)	(1.235)	(1.827)	(1.388)	(2.541)
	5+	20.59**	-27.42	-5.687***	-14.01***	-9.636***	-23.57***
		(9.492)	(20.12)	(1.382)	(1.949)	(1.494)	(2.749)
Children	Yes	47.90***	-20.61	0.291	0.720	0.893	3.153*
		(5.878)	(13.10)	(1.038)	(1.201)	(1.069)	(1.726)
Age	age	-8.267***	-4.198***	0.0985***	0.275***	-0.107***	-0.176***
		(0.167)	(0.398)	(0.0268)	(0.0353)	(0.0283)	(0.0531)
Gender	male	-35.94***	-69.51***	0.614	-2.900***	20.01***	35.45***
		(4.651)	(10.21)	(0.591)	(0.912)	(0.741)	(1.600)
Ethnicity	non-Hispanic	16.93**	-44.35***	-3.290***	-10.83***	-2.249*	-15.13***
		(6.720)	(14.51)	(1.059)	(1.593)	(1.222)	(2.537)
Constant	cons	1097.7***	1533.4***	27.19***	55.39***	42.18***	97.65***
		(15.81)	(35.62)	(2.054)	(3.222)	(2.410)	(4.743)
Controls	Region FE	YES	YES	YES	YES	YES	YES
	Week FE	YES	YES	YES	YES	YES	YES
	Obs	112344	112344	58290	58290	64311	64311
	Adj R2	0.038	0.010	0.006	0.011	0.021	0.023

Note: Regression of equation (1) in section 5. Dependent variable is weekly usage. Observations are at *device*×*week* level. Robust standard errors in the parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Table A2: Quantile Regressions for whole sample and sub-samples

Panel A: All websites

	Duration			Pages		
Income	p50	p80	p90	p50	p80	p90
income{<25000}	26.62***	114.1***	162.9***	-218.1***	-367.3***	-337.3***
	(8.545)	(12.67)	(20.11)	(15.40)	(33.75)	(51.95)
income{25000 - 39999}	77.32***	152.4***	202.0***	-235.1***	-321.6***	-242.4***
	(8.659)	(13.52)	(20.10)	(15.57)	(35.00)	(52.10)
income{40000 - 59999}	95.66***	142.1***	165.4***	-164.1***	-54.87	-5.856
	(8.799)	(13.35)	(20.39)	(16.50)	(35.89)	(52.82)
income{60000 - 74999}	82.60***	113.1***	188.6***	-98.45***	-49.49	10.63
	(10.52)	(16.51)	(24.77)	(20.74)	(40.64)	(68.15)
income{75000 - 99999}	60.61***	88.29***	77.66***	-6.588	161.4***	194.5***
	(9.931)	(15.44)	(23.01)	(20.24)	(42.68)	(58.46)
Obs	112344	112344	112344	112344	112344	112344

Panel B: Top 3000 News & weather

	Duration			Pages		
Income	p50	p80	p90	p50	p80	p90
income{<25000}	-2.929***	-10.20***	-17.64***	-5.168***	-29.48***	-57.97***
	(0.229)	(0.915)	(2.327)	(0.339)	(1.967)	(4.584)
income{25000 - 39999}	-2.287***	-8.916***	-16.61***	-4.888***	-26.87***	-46.40***
	(0.236)	(0.897)	(2.285)	(0.341)	(2.000)	(5.021)
income{40000 - 59999}	-1.494***	-5.545***	-5.420**	-3.710***	-14.62***	-25.58***
	(0.243)	(0.954)	(2.418)	(0.358)	(2.231)	(4.980)
income{60000 - 74999}	-0.494	-1.759	-2.096	-2.626***	-8.602***	-10.35*
	(0.313)	(1.103)	(2.839)	(0.410)	(2.695)	(6.000)
income{75000 - 99999}	-0.712**	3.371***	5.907**	-1.514***	2.758	8.074
	(0.301)	(1.274)	(2.933)	(0.428)	(2.848)	(6.104)
Obs	58291	58291	58291	58291	58291	58291

Panel C: Top 3000 News, weather, and sports

	Duration			Pages		
Income	p50	p80	p90	p50	p80	p90
income{<25000}	-5.219***	-18.69***	-29.13***	-7.821***	-57.14***	-115.2***
	(0.345)	(1.405)	(2.805)	(0.455)	(2.785)	(6.740)
income{25000 - 39999}	-3.393***	-11.70***	-15.64***	-7.166***	-52.33***	-93.32***
	(0.360)	(1.497)	(2.965)	(0.459)	(2.914)	(6.893)
income{40000 - 59999}	-2.374***	-6.701***	-5.471*	-5.321***	-33.69***	-51.83***
	(0.371)	(1.534)	(3.195)	(0.484)	(3.184)	(7.572)
income{60000 - 74999}	0.390	2.736	2.995	-3.422***	-17.57***	-27.98***
	(0.496)	(1.816)	(3.520)	(0.546)	(4.044)	(8.370)
income{75000 - 99999}	-0.0408	4.458**	8.220**	-1.349**	0.661	-4.964
	(0.469)	(1.759)	(3.475)	(0.630)	(3.905)	(9.578)
Obs	64310	64310	64310	64310	64310	64310

Note: Quantile regression of equation (1) in section 5. Dependent variable is weekly usage. Observations are at *device*×*week* level. Robust standard errors in the parentheses. * p<0.10, ** p<0.05, *** p<0.01.

Proof of Proposition 1

This price cap problem yields the following Lagrangian:

$$
\begin{aligned}
L = N_H P_U + N_L P_C + \lambda_1 \left(\int_0^{f^{-1}(0)} f(x)dx - P_U \right) + \lambda_2 \left(\int_0^{C} g(x)dx - P_C \right) \\
+ \lambda_3 \left(\int_0^{f^{-1}(0)} f(x)dx - P_U - \int_0^{C} f(x)dx + P_C \right) \\
+ \lambda_4 \left(\int_0^{C} g(x)dx - P_C - \int_0^{g^{-1}(0)} g(x)dx + P_U \right)
\end{aligned}
$$

with $\lambda_1, \dots, \lambda_4 \geq 0$

The Kuhn-Tucker conditions are as follows:

1. $L_{P_U}: N_H - \lambda_1 - \lambda_3 + \lambda_4 = 0$
2. $L_{P_C}: N_L - \lambda_2 + \lambda_3 - \lambda_4 = 0$
3. $L_C: \lambda_2 g(C) - \lambda_3 f(C) + \lambda_4 g(C) = 0$
4. $\lambda_1 \left(\int_0^{f^{-1}(0)} f(x)dx - P_U \right) = 0$
5. $\lambda_2 \left(\int_0^{g^{-1}(0)} g(x)dx - P_C \right) = 0$
6. $\lambda_3 \left(\int_0^{f^{-1}(0)} f(x)dx - P_U - \int_0^{C} f(x)dx + P_C \right)$
7. $\lambda_4 \left(\int_0^{C} g(x)dx - P_C - \int_0^{g^{-1}(0)} g(x)dx + P_U \right)$

The steps to solving the problem using the Kuhn-Tucker conditions are as follows:

1) The inequality in constraint #1 must be strict. To see this, suppose it is not, i.e., that constraint #1 is an equality. Then, high types get 0 net surplus from buying unlimited. The most you can charge to get low types to buy a capped plan is $\int_0^{C} g(x)dx < \int_0^{f^{-1}(0)} f(x)dx$ by Assumption 6. Hence, the High types would get positive surplus from buying the capped plan, meaning they would buy it, and violating constraint #3.
2) At least one of constraints #1 and #2 must bind. Otherwise, you could raise both prices by ε, make more revenue, and still satisfy all constraints. Since we know constraint #1 is not binding, it must be that constraint #2 is. Therefore, $\boldsymbol{P_C = \int_0^{C} g(x)dx}$.
3) Combine the first two Kuhn-Tucker conditions to get $N_H + N_L = \lambda_1 + \lambda_2$. Then, since $\lambda_1 = 0$, we have $\lambda_2 = N_H + N_L$.
4) From the second Kuhn-Tucker condition, and step #3 above, we have $N_H = \lambda_3 - \lambda_4$. Therefore, we can't have the last two constraints (#3 and #4) both be non-binding (or else we'd have N_H be zero, and it is assumed positive).

5) Consider the three remaining possibilities for constraints #3 and #4 (both bind, #3 only binds, #4 only binds). Start by supposing both constraint #3 and #4 are binding. Looking at constraint #4, the left side is zero (plugging in for our solution to P_C). Hence, we have $\boldsymbol{P_U = \int_0^{g^{-1}(0)} g(x)dx}$. The implicit solution then for C is the value such that (looking at constraint #3): $\int_0^{f^{-1}(0)} f(x)dx - \int_0^{g^{-1}(0)} g(x)dx = \int_0^C f(x)dx - \int_0^C g(x)dx$. While we don't explicitly solve for C we note that it is some point between 0 and $g^{-1}(0)$, since the difference in surplus is 0 at usage = 0 and is decreasing past $f^{-1}(0)$. Hence, by the Intermediate Value Theorem, there exists a C between 0 and $f^{-1}(0)$ such that the difference in surplus up to C is equal to the difference in total surplus.
6) Suppose now that only constraint #4 binds, meaning $\lambda_3 = 0$. Then, according to the first Kuhn-Tucker condition, we have $\lambda_4 = -N_H < 0$, which is a violation.
7) Suppose that only constraint #3 binds, meaning $\lambda_4 = 0$. Then, from Kuhn-Tucker condition #1, we have $\lambda_3 = N_H$. Then, from Kuhn-Tucker condition #3, we have $\boldsymbol{(N_H + N_L)g(C) = N_H f(C).}$ From constraint #3 binding, we have $\boldsymbol{P_U = \int_0^{f^{-1}(0)} f(x)dx + \left[\int_0^C g(x)dx - \int_0^C f(x)dx\right].}$

Hence, we have either:

A) $\boldsymbol{P_C = \int_0^C g(x)dx, P_U = \int_0^{g^{-1}(0)} g(x)dx}$, with C implicitly defined by $\mathbf{\int_0^{f^{-1}(0)} f(x)dx - \int_0^{g^{-1}(0)} g(x)dx = \int_0^C f(x)dx - \int_0^C g(x)dx.}$ Revenue is: $\boldsymbol{N_L \int_0^C g(x)dx + N_H \int_0^{g^{-1}(0)} g(x)dx}$

B) $\boldsymbol{P_C = \int_0^C g(x)dx, P_U = \int_0^{f^{-1}(0)} f(x)dx + \left[\int_0^C g(x)dx - \int_0^C f(x)dx\right]}$, with C implicitly defined by $\boldsymbol{(N_H + N_L)g(C) = N_H f(C)}$. Revenue is: $\boldsymbol{Revenue = N_L \int_0^C g(x)dx + N_H \left(\int_0^{f^{-1}(0)} f(x)dx + \left[\int_0^C g(x)dx - \int_0^C f(x)dx\right]\right)}$

QED

Proof of Proposition 2

Define: $f(x) = A_H - B_H x$ and $g(x) = A_L - B_L x$.

Assumptions #4 and #5 imply that $A_H > A_L$ and that $\frac{A_H}{B_H} < \frac{A_L}{B_L}$

Consider the solution for the price discriminating firm. We know the optimal cap satisfies:

$$(N_L + N_H)(A_L - B_L C) = N_H(A_H - B_H C)$$

Consider possible solution $\frac{A_H}{B_H}$ = C*. Then, we have:

$$(N_L + N_H)(A_L - B_L C^*) > N_H(A_H - B_H C^*) = 0$$

If we consider a lower value for C* to satisfy the inequality, the LHS increases by $(N_L + N_H)B_L$. The RHS increases by $N_H B_H$. We've assumed $(N_H + N_L)g'(x) < N_H f'(x)$, so the LHS increases by more than the RHS. Hence $C^* \geq \frac{A_H}{B_H}$. That is, the data cap is at least as high as the maximum usage by high types. This means, even in the previously ambiguous case #2 above, $\boldsymbol{P_U = \int_0^{f^{-1}(0)} f(x)dx + \left[\int_0^{C} g(x)dx - \int_0^{C} f(x)dx\right] = \int_0^{C} g(x)dx}$ (first and third terms cancel). Therefore, the price charged to the high and low types is no more than the maximum willingness-to-pay of the low types, which is what a firm selling just the unlimited plan could charge and get everyone to buy.

QED

Examination of Revenues in Scenario B of Proposition

To help understand what influences revenues in scenario B of the price discrimination case, recall that revenues in scenario B are: $\boldsymbol{Revenue = N_L \int_0^C g(x)dx + N_H \left(\int_0^{f^{-1}(0)} f(x)dx + \left[\int_0^C g(x)dx - \int_0^C f(x)dx\right]\right)}$. We can rewrite these revenues as:

$$\begin{aligned}\boldsymbol{Revenue} &= \boldsymbol{(N_L + N_H)\int_0^{g^{-1}(0)} g(x)dx} \\ &\quad \boldsymbol{- N_H\left[\left(\int_0^{g^{-1}(0)} g(x)dx - \int_0^C g(x)dx\right) - \left(\int_0^{f^{-1}(0)} f(x)dx - \int_0^C f(x)dx\right)\right]} \\ &\quad \boldsymbol{- N_L\left(\int_0^{g^{-1}(0)} g(x)dx - \int_0^C g(x)dx\right)}\end{aligned}$$

Note that the above expression has three terms. The first is the revenue the monopolist can attain when selling just the unlimited plan to both types. Consider next the third term. The expression in parentheses is weakly positive, since C must be less than the satiation point of L types. As the number of L types increases, we know two things. First, the single-offering monopolist will want to sell to both types (revenue when selling to both is increasing in N_L while revenue when selling to just H types is not changing with N_L). Second, the difference in the above revenue and the revenue when selling to both types is decreasing in N_L, due to the last expression above. Consequently, the difference in price discrimination revenue and single-offering revenue is weakly declining in the number of L types. This is our first claim in the text.

Now consider the second term in the above revenue equation. The first expression in parentheses is the surplus for L types when increasing usage from the cap (C) to their satiation point, and the second expression is the analog for H types. The difference in these expressions may be positive or negative. However, the second expression is decreasing as H types experience more surplus with smaller amounts of usage (leaving less remaining surplus to be had when using beyond any cap). As the second expression decreases, the revenue equation above declines, leading to our second claim in the text.