Cellular Service Demand:
Biased Beliefs, Learning, and Bill Shock*

Michael D. Grubb†and Matthew Osborne‡

February 27, 2012

Abstract

By April 2013, the FCC's recent bill-shock agreement with cellular carriers requires consumers be notified when exceeding usage allowances. Will the agreement help or hurt consumers? To answer this question, we estimate a model of consumer plan choice, usage, and learning using a panel of cellular bills. Our model predicts that the agreement will lower average consumer welfare by \$2 per year because firms will respond by raising monthly fees. Our approach is based on novel evidence that consumers are inattentive to past usage (meaning that bill-shock alerts are informative) and advances structural modeling of demand in situations where multi-part tariffs induce marginal-price uncertainty. Additionally, our model estimates show that an average consumer underestimates both the mean and variance of future calling. These biases cost consumers \$42 per year at existing prices. Moreover, absent bias, the bill-shock agreement would have little to no effect.

1 Introduction

Cellular phone companies frequently offer consumers contracts with included allowances of voice minutes, text messages, and data usage that are followed by overage charges for higher usage. Consumers are often unaware that they are incurring overage charges during the month, which leads to *bill shock* at the end of the month. On October 17^{th}, 2011 President Barack Obama declared:

> Far too many Americans know what its like to open up their cell-phone bill and be shocked by hundreds or even thousands of dollars in unexpected fees and charges. But we can put an end to that with a simple step: an alert warning consumers that they're about to hit their limit before fees and charges add up (CTIA - The Wireless Association 2011a).

President Obama made this statement at the announcement of a new bill-shock agreement between the FCC and cellular carriers. By April 2013, this agreement commits cellular service providers to inform consumers when they approach and exceed their included voice, text, and data allowances (CTIA - The Wireless Association 2011a). Prior to the agreement, the FCC had proposed a similar regulation which was strongly supported by consumer groups but opposed by the industry (Deloney, Sherry, Grant, Desai, Riley, Wood, Breyault, Gonzalez and Lennett 2011, Altschul, Guttman-McCabe and Josef 2011).[1]

Will the new bill-shock agreement help or hurt consumers? If carriers held their prices fixed after implementing the agreement then it would weakly help consumers. Such prices-fixed logic likely lies behind consumer groups' strong advocacy for bill-shock alerts. However, the bill-shock agreement could hurt consumers once endogenous price changes are taken into account. Moreover, complementary theoretical work by Grubb (2011) shows that the answer is theoretically ambiguous. Therefore, to address this question, we develop and estimate a dynamic model of plan choice and usage that makes use of detailed cellular phone data. Given our parameter estimates, counterfactual simulations show that the net effect of the bill-shock agreement and endogenous prices changes is an overall annual reduction in consumer welfare of $2 per consumer.

En route to making our prediction about the bill-shock agreement's effect on consumer welfare we make two additional contributions. First, we provide new evidence on how consumers make consumption choices under marginal-price uncertainty and estimate a tractable model incorporat-

---

[1]The wirelesss industry trade group, C.T.I.A. - The Wireless Association, argued that proposed bill-shock regulation "violates carriers' First Amendment protections. . . . against government compelled speech" (Altschul et al. 2011).

ing such realistic behavior. In particular, we find that consumers are inattentive to their remaining minute balance. Given such inattention, we assume that consumers optimally respond to exogenously arising calling opportunities by choosing a calling threshold and making only those calls more valuable than the threshold. Unlike standard models, this approach allows for consumers to endogenously adjust their calling behavior in response to bill-shock alerts in our counterfactual simulations. (Attentive consumers would never find new information in a bill-shock alert.) Second, we relax the standard rational expectations assumption and infer consumers' beliefs about their future calling opportunities from plan choices. By comparing these beliefs to actual usage at the population level, systematic differences identify consumer biases such as overconfidence. Identifying consumer biases is important for our endogenous-price counterfactual simulations because firm pricing decisions are strongly influenced by overconfidence and other biases (Grubb 2009).

Our primary data were obtained from a major US university that acted as a reseller for a national cellular phone carrier, and covers all student accounts managed by the university from 2002 to 2004. We begin by documenting five stylized facts in our data that shape our modelling approach. First, a sharp increase in calling when free off-peak calling begins shows that consumers' usage choices are price sensitive. Second, absence of bunching at tariff kink points and other evidence show that consumers are uncertain about the ex post marginal price when making calling choices. Third, novel evidence from call-level data suggests consumers are inattentive to their remaining balance of minutes. Fourth, consumers are uncertain about their own average taste for usage when first choosing a calling plan, which leads to frequent ex post plan choice mistakes. However, consumers learn about their own tastes over time and switch plans in response. Finally, consumers make ex ante mistakes that are predictable given information held by a carrier.

The first three stylized facts suggest that the arrival of a bill-shock alert will be informative and cause a consumer to reduce calling. The second stylized fact, marginal-price uncertainty, naturally arises whenever consumers make a series of small purchase choices that are aggregated and billed under a multipart tariff, as in cellular phone service, electricity, and health care. Addressing such marginal-price uncertainty represents a challenge for the literature which has typically side-stepped the issue by assuming that consumers can perfectly predict their future usage (Cardon and Hendel 2001, Reiss and White 2005, Lambrecht, Seim and Skiera 2007), or that consumers believe they can perfectly predict their usage up to an implementation error which they ignore (Iyengar, Ansari and Gupta 2007). (Notable exceptions are Yao, Mela, Chiang and Chen (2011) and Jiang (2011).) By recognizing that consumers are inattentive, our modeling approach incorporates marginal-price uncertainty realistically and tractably and allows consumers to endogenously respond to bill-shock alerts. Our consumers behave optimally given their inattention, by choosing a calling threshold each

month (related to expected marginal price) and accepting only calls valued above the threshold. This approach has been proposed in earlier work (Saez 2002, Borenstein 2009), but has not been implemented in a structural model.[2] An advantage of our structural approach is that we can estimate the consumer beliefs required to calculate calling thresholds.

To account for the last two stylized facts concerning plan choice, we model consumer beliefs and learning. We call a consumer's average taste for calling his true type. A consumer's plan choices are determined not by his true type but by his beliefs about his true type. We assume that each consumer's prior consists of a point estimate of her own true type and a level of perceived uncertainty about this point estimate. We assume that consumers are Bayesian learners, following Erdem and Keane (1996), Ackerberg (2003), Crawford and Shum (2005), and Goettler and Clay (2011) and therefore learn their true types in the long run. At the same time, to account for the predictable nature of plan choice mistakes in the short run, we allow consumers' initial beliefs to be biased.

Our data are informative both about consumers' actual average tastes for cellular phone usage and about their prior beliefs about their own tastes. Consumers' usage choices identify the distribution of consumers' true types, while consumers' initial plan choices and subsequent switching decisions identify beliefs. The joint distribution of beliefs and true types determines whether beliefs are biased in the population. For instance, suppose that we consider the subset of consumers that all share a particular prior belief about their own types. A common assumption (often labeled rational expectations) is that this belief coincides with the distribution of true types within this subset of the population. We relax this assumption, separately identify both beliefs and the distribution of true types conditional on beliefs, and then compare the two distributions. We label differences between these distributions as biases.[3] Moreover, we allow consumers to over- or under-estimate the monthly volatility in their tastes.

We identify two substantial biases causing predictable mistakes. The first we label *overconfidence*, which arises when a consumer underestimates her own uncertainty surrounding her point estimate of her true type. We find that consumers underestimate their own uncertainty about their true type by 84%. Overconfident consumers initially choose plans that are too risky. Moreover,

---

[2] In the context of electricity demand, Borenstein (2009) independently proposes that consumers choose *behavioral rules*, such as setting the thermostat, similar to our calling threshold. Borenstein (2009) uses the behavioral rule assumption to motivate using expected marginal price rather than realized marginal price in reduced form estimates of electricity price elasticities. Saez (2002) also suggests a very similar model for labor choice by income tax filers.

[3] An alternate interpretation is that unmeasurable prior beliefs were unbiased at some previous time, but are now measurably and systematically different from reality at the population level (although consistent with rational expectations) due to the arrival of a correlated shock or signal at the population level. The distinction is pedantic as it does not matter for optimal firm pricing, consumer welfare, policy counter-factuals or other issues of interest.

they place too much weight on their prior point estimates when updating beliefs and will be slow to learn and switch plans based on experience.

The second bias that we focus on is *volatility bias*, which arises when consumers underestimate the monthly volatility in their tastes for usage.[4] We estimate that consumers underestimate the volatility in their taste for usage by 10%. Similar to overconfidence, volatility bias causes consumers to underestimate the uncertainty in their usage predictions when making plan choices, and choose plans that are too risky. However, volatility bias has the opposite effect of overconfidence on the rate of learning: volatility bias causes consumers to underweight their priors relative to past usage when updating their beliefs about their average tastes for usage. This leads to faster learning and more frequent plan switching. Thus the rate of plan switching allows us to separate the two biases.[5] Because we find that overconfidence is stronger than volatility bias, consumers overweight their prior beliefs relative to new information and learn and switch plans relatively slowly. Thus initial plan choice mistakes are especially costly. Holding observed prices constant, we find that overconfidence and volatility bias jointly reduce annual consumer welfare by $34 per student.

There are other biases in beliefs which could result in consumer behavior that is similar to that caused by overconfidence and volatility bias. To ensure we do not misattribute other errors as overconfidence or volatility bias, we estimate a flexible distribution of initial beliefs which captures (at least) two other potential sources of bias. We are able to separately identify these biases due to the rich choice set of plans in our data that importantly include both three-part tariffs and a two-part tariff. Holding observed prices constant, these biases reduce annual consumer welfare by an additional $8 per student, for a total annual cost of all biases of $42 per student.

Turning back to the recent FCC agreement, we conduct a counterfactual simulation where we allow firms to adjust prices in response to bill-shock alerts. To do so, we add additional supply side structure to our model and add a parameter λ measuring the amount of differentiation across firms. This firm differentiation parameter λ is omitted from our estimated demand model because our demand data are from a single carrier and do not identify λ. To complete our endogenous price counterfactual simulations, we therefore first calibrate the firm differentiation parameter λ conditional on our demand estimates using observed prices. (We use EconOne data on the prices of all cellular-phone plans offered during 2002-2004 in the vicinity of the university that provided our primary data.) We find that firms respond to bill-shock regulation by raising fixed fees, reducing

---

[4]Overconfidence could more broadly be interpreted to include volatility bias, however we seek to draw a distinction between two different biases and define overconfidence more narrowly to do so.

[5]Our model includes a price consideration parameter that plays a similar role to a switching cost. This is separately identified from the learning rate by the rate at which consumers fail to switch away from strictly dominated plans.

included minute allowances, and reducing overage rates on three-part tariffs. By doing so, firms maintain annual profits close to unregulated levels (falling by just $0.20 per person).[6] This means that consumers are approximately residual claimants on total welfare, which falls by $2.21 per person, and hence annual consumer welfare drops by $2.01 per person. The social welfare loss results from consumers' reduced calling. Absent consumer biases, we find that firms offer two-part tariffs but not three-part tariffs, which means that bill-shock regulation has no effect.

Section 2 discusses related literature. Section 3 describes our data and documents five stylized facts that shape our modeling approach. Sections 4 and 5 describe our model and explain identification. Sections 6, 7, and 9 discuss estimation, present results and conclude. Additional details are in the Online Appendix available at \url{www.mit.edu/~mgrubb/GrubbOsborneAppendix.pdf}.

2   Related Literature

Complementary work by Jiang (2011) also evaluates the recent bill-shock agreement via counterfactual simulation, predicting a $370 million welfare improvement. In contrast to our own approach, Jiang (2011) imposes rational expectations rather than estimating consumer beliefs and has cross-sectional data so cannot address learning. Finally, Jiang's (2011) bill-shock counterfactual corresponds to removing a taste shock from the model. In contrast, a strength of our approach is that consumers endogenously change calling behavior in response to information in bill-shock alerts. (A strength of Jiang's (2011) data is that they are nationally representative and cover all carriers.)

Related work provides evidence that individual labor choices (Liebman and Zeckhauser 2004) and electricity consumption choices (Ito 2010) respond to average prices rather than marginal prices. This is not surprising because electricity tariffs and the income tax code are both very complex and often not well understood by consumers. A typical consumer may not realize electricity pricing is nonlinear, in which case average price is a good estimate of marginal price. However, this model is not appealing in the context of cellular service because consumers are fully aware that contracts include an allowance of 'free' minutes.

A significant body of experimental evidence shows that individuals are overconfident about the precision of their own predictions when making difficult forecasts (e.g. Lichtenstein, Fischhoff and Phillips (1982)). In other words, individuals tend to set overly narrow confidence intervals relative to their own confidence levels. A typical psychology study might pose the following question to a group of subjects: "What is the shortest distance between England and Australia?" Subjects would then be asked to give a set of confidence intervals centered on the median. A typical finding

---

[6]Annual profits fall by $3.69 per consumer for any single firm that independently chooses to offer bill-shock alerts.

is that the true answer lies outside a subject's 98% confidence interval about 30% to 40% of the time. Consumers who exhibit volatility bias underestimate the extent to which their tastes will change over time. This is closely related to projection bias, a prevalent behavior that has been documented in a variety of experiments, surveys, and field studies (Loewenstein, O'Donoghue and Rabin 2003, Conlin, O'Donoghue and Vogelsang 2007). Via mean biases, we allow for overestimation of demand, which is one of the causes of flat-rate bias documented by Lambrecht and Skiera (2006) in internet service choice.

A small number of empirical papers relax rational expectations for consumer beliefs and estimate mean biases (Crawford and Shum 2005, Goettler and Clay 2011). Most similar to our work is Goettler and Clay (2011), which estimates mean biases. Goettler and Clay (2011) cannot identify higher moments of beliefs because the choice set in online grocery-delivery service is limited to two-part tariffs. In contrast, the rich tariff choice-set in our setting enables us to measure (rather than assume away) volatility bias and overconfidence in addition to mean biases.

To identify beliefs from plan choices, we assume consumers are risk neutral.[7] In contrast, related work on health insurance markets often does the reverse and imposes rational expectations to identify risk preferences from plan choices (Cardon and Hendel 2001, Handel 2011, Einav, Finkelstein, Pascu and Cullen Forthcoming). Following a third approach, Ascarza, Lambrecht and Vilcassim (2012) impose rational expectations and risk neutrality but estimate preferences for cellular phone usage that depend directly on whether contracts are two or three-part tariffs.

Our results are consistent with a related sequence of papers about Kentucky's 1986 local telephone tariff experiment (Miravete 2002, Miravete 2003, Miravete 2005, Narayanan, Chintagunta and Miravete 2007, Miravete and Palacios-Huerta 2011). First, although the standard model of consumer choice does well at explaining behavior in the Kentucky experiment, our estimates of negative aggregate mean bias and positive conditional mean bias are consistent with evidence in Miravete (2003) which documents that on average all consumers who chose a small metered plan would have saved money on a larger flat rate plan.[8] Second, as in the Kentucky experiment we find that most consumers (55 to 71 percent) initially choose the tariff that turns out to be optimal ex post. Moreover, consumers switch plans and most switches appear to be in the right direction to lower bills (Section 3.2).

Our counterfactual simulations with endogenous prices relate to the literatures with standard

---

[7]If consumers are risk averse then our estimates of overconfidence and volatility bias are lower bounds on bias.

[8]Interestingly, in Miravete (2003) the bias that can be inferred from elicited expectations differs from that inferred from choices. Consumers were not offered three-part tariffs in the Kentucky experiment so their choices do not shed light on overconfidence or volatility bias.

consumers on monopoly sequential-screening (surveyed by Rochet and Stole ((2003), Section 8), including Baron and Besanko (1984), Riordan and Sappington (1987), Miravete (1996), Courty and Li (2000), Miravete (2005), and Grubb (2009)) and competitive static-screening (surveyed by Stole (2007), including Armstrong and Vickers (2001) and Rochet and Stole (2002)). Moreover, it is related to the growing literature on optimal contracting with non-standard consumers (for which Spiegler (2011) provides a good guide). Of particular relevance are DellaVigna and Malmendier (2004), Uthemann (2005), Eliaz and Spiegler (2006), Eliaz and Spiegler (2008), Grubb (2009), Herweg and Mierendorff (Forthcoming), and Grubb (2011).

Finally our paper is about the cellular phone industry, about which there is a small literature. Beyond work already mentioned, other work on the cellular phone industry examines risk while driving (Bhargava and Pathania 2011), carrier switching costs (Kim 2006), the effect of entry on pricing (Seim and Viard 2010, Miravete and Röller 2004), the effect of number portability regulation on competition (Park 2009), the role of multi-market contact in competition (Busse 2000), and demand (Iyengar, Jedidi and Kohli 2008, Huang 2008).

3 Background: Data and Evidence for Stylized Facts

3.1 Data

Our primary data are a panel of individual monthly billing records for all student enrollees in cellular-phone plans offered by a national cellular carrier in conjunction with a major university from February 2002 to June 2005. During this period, cellular phones were a relatively new product in the US, having 49% penetration in 2002 compared to 98% in 2010.[9] This data set includes both monthly bill summaries and detailed call-level information for each subscriber.[10] We also acquired EconOne data on the prices and characteristics of all cellular-phone plans offered at the same dates in the vicinity of the university. The price menu offered to students differed from that offered by the carrier directly to the public: university plans included a two-part tariff, a limited three-month contractual commitment, different monthly promotions of *bonus* minutes, and a $5 per month surcharge on top of carrier charges to cover the university's administrative costs.

---

[9]This feature makes our data ideal for studying consumer beliefs about new products. Penetration rates are calculated as estimated total connections (CTIA - The Wireless Association 2011b) divided by total population (U.S. Census Bureau 2011).

[10]Students recieved monthly phone bills, mailed by default to their campus residence. The sample of students is undoubtedly different than the entire cellular-phone-service customer-base. However, a pricing manager from one of the top US cellular phone service providers made the unsolicited comment that the empirical patterns of usage, overages, and ex post "mistakes" documented in Grubb (2009) using the same data were highly consistent with their own internal analysis of much larger and representative customer samples.

The bulk of our work makes use of the monthly billing data. We exclude individuals who are left censored (those who are existing subscribers at the start of the panel). For most analysis, including our structural estimation, we also restrict attention to the period August 2002 to July 2004. (This is the period for which we can reliably infer university prices from billing data. See Appendix A). We focus on customer choice between four *popular* local plans, that account for 89% of bills in our data. We group the remaining price plans (including national and free long distance plans) with the outside option, and hence drop the 11% of bills with unpopular price plans.[11] Finally, rate plan codes are frequently miscoded as a default value on a customers initial bill, in which case we remove the first bill. Our final data set contains 1366 subscribers and 16,283 month-subscriber observations. Note that for much of our analysis, we also exclude pro-rated bills during months of partial service, or customer switching between plans (however, pro-rated bills are included in the sample we use to estimate the structural model).

Figure 1 shows the four popular plans, which we label as plans 0 through 3. Plan 0 is a two-part tariff that charges $14.99 per month and 11 cents per minute. Plans 1-3 are three-part tariffs that charge monthly fees (M_j) of 34.99, 44.99, and 54.99 respectively, include an allowance (Q_j) of 280 to 1060 free peak-minutes, and charge an overage rate (p_j) of 35 to 45 cents per additional peak minute. We say that one plan is larger than another if it coincides with the lower envelope of the tariff menu at a higher interval of usage. Plans are numbered in order of size, smallest to largest. We say that one plan is riskier than another if it yields higher expected bills for sufficiently high usage uncertainty. Plan 0 is the safest plan, plan 1 is the riskiest, and plans 1-3 are numbered in order of decreasing risk.

All four plans include surcharges of 66 to 99 cents per minute for roaming outside a subscriber's tri-state area and 20 cents per minute for long distance. Plans 1-3 always offer free off-peak calling but Plan 0 does so only prior to fall 2003. Plan 0 includes free in-network calling, while plans 1-3 do not with the exception of plan 2 in 2004. Once a customer chooses a plan, the plan terms remain fixed for that customer, regardless of any future promotions or discounts, until they switch plans or terminate service. However, the terms of any given plan, such as the included allowances and overage rates for plans 1-3, vary according to the date a customer chooses the plan.

Shares of plans 0-3 are 44, 28, 15, and 2 percent of bills, respectively. Plan prices are shown for Spring 2003 in Figure 1 and are described for all dates in Appendix A Table 7. This price series was inferred from billing data rather than directly observed, as discussed in Appendix A.

---

[11]In fact, we treat switching to an unpopular plan the same as quitting service, hence we also drop all remaining bills once a customer switches to an unpopular plan, even if they eventually switch back to a popular plan.

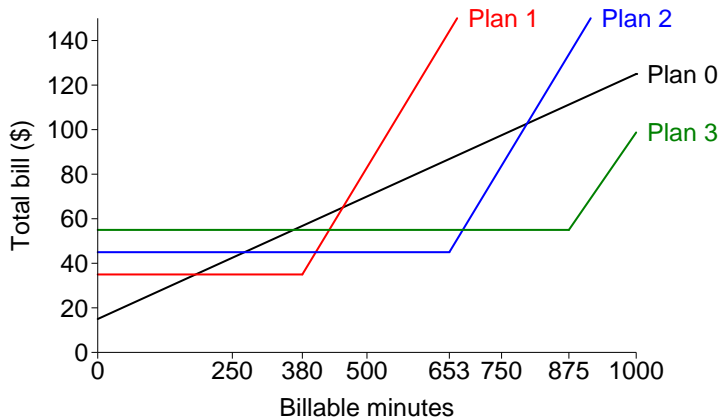| Plan	M_j	Q_j	p_j
Plan 0	$14.99	0	$0.11
Plan 1	$34.99	380	$0.45
Plan 2	$44.99	653	$0.40
Plan 3	$54.99	875	$0.40

Figure 1: Popular Plan Prices, Spring 2003.

3.2 Evidence for Stylized Facts

3.2.1 Three stylized facts relevant to modeling usage choices

Three features of the data are important to accurately model usage choices by customers of cellular phone service. First, consumers' usage choices are price sensitive. Second, consumers' usage choices are made while consumers are uncertain about the ex post marginal price. Third consumers are inattentive to the remaining balance of included minutes during the course of a billing cycle. These three stylized facts motivate our assumption that, rather than choosing a precise quantity, consumers choose calling thresholds and proceed to make all calls valued above the threshold.

Consumer price sensitivity is clearly illustrated by a sharp increase in calling volume on weekday evenings exactly when the off-peak period for free night and weekend calling begins (Figure 2). This is not simply a 9pm effect, as the increase occurs only on weekdays, and at 8pm for plans with early nights-and-weekends.[12]

Two pieces of evidence demonstrate consumer uncertainty about ex post marginal price. First, given clear sensitivity to marginal price, if consumers could anticipate whether they would be under their allowance (zero marginal price ex post) or over their allowance (35 to 45 cents per minute marginal price ex post) we would expect to see substantial bunching of consumers consuming their entire allowance but no more or less. Figure 3 shows there is no bunching, which is consistent with similar findings in the contexts of electricity consumption (Borenstein 2009) and labor supply

---

[12]For plans with free weeknight calling starting at 8pm, there is still a secondary increase in usage at 9pm (Figure 2 panel C). Restricting attention to outgoing calls made to land-lines (recipients for whom the cost of receiving calls was zero) almost eliminates this secondary peak (Figure 2 panel D). This suggests that the secondary peak is primarily due to calls to and from cellular numbers with 9pm nights (the most common time for free evening calling to begin) rather than a 9pm effect.
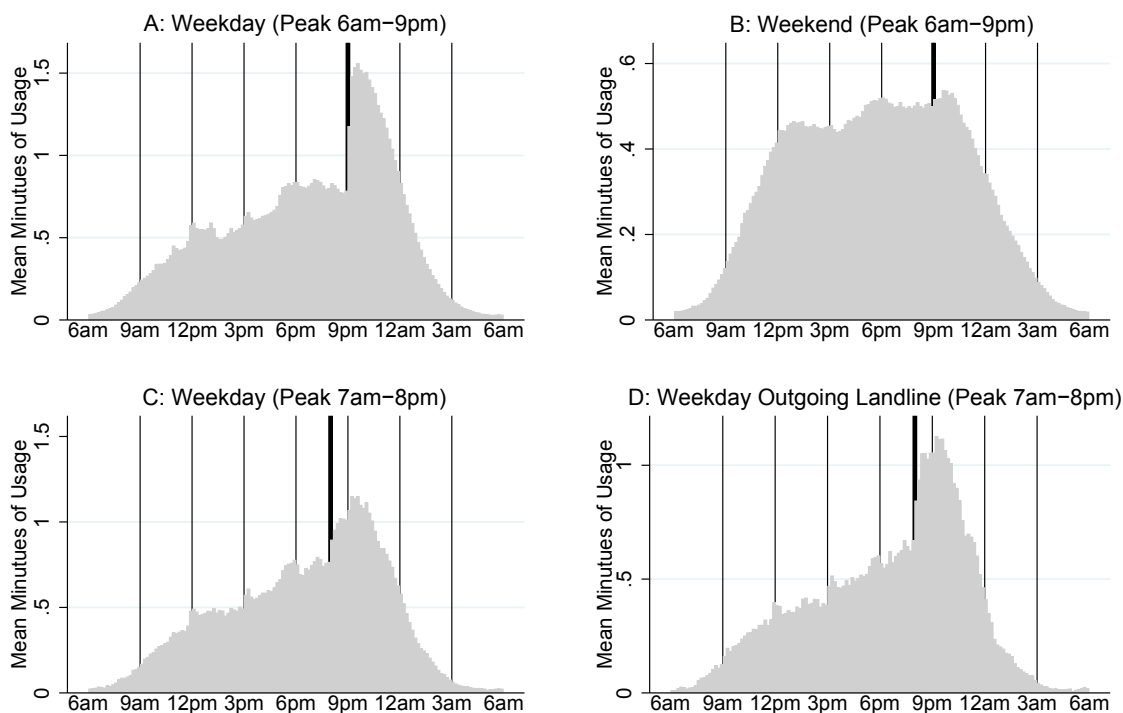
9

Figure 2: Daily usage patterns for subscribers with free nights and weekends. Top row: weekday (Panel A) and weekend (Panel B) usage patterns for subscribers with 6am-9pm peak hours. Bottom row: weekday usage patterns for subscribers with 7am-8pm peak hours. Panel C shows all weekday calling, while Panel D is restricted to outgoing calls to land-lines.

(Saez 2010). Second, consumers who anticipate being strictly under their allowance (zero marginal price ex post) should exhibit no price response at the commencement of off-peak hours. However, the sharp increase in calling at 9pm shown in Figure 2 persists even in months for which the peak allowance is under-utilized. These are natural consequences of usage choices made under uncertainty about ex post marginal price.

Now we turn to evidence that consumers are inattentive. If consumers are attentive to the remaining balance of included minutes during the billing cycle they should use this information to continually update their beliefs about the likelihood of an overage and a high marginal price ex post. Following an optimal dynamic program, an attentive consumer should (all else equal) reduce her usage later in the month following unexpectedly high usage earlier in the month. This should be true for any consumers who are initially uncertain whether they will have an overage in the current month. For these consumers, the high usage shock early in the month increases the likelihood of an overage, thereby increasing their expected ex post marginal price, and causing them to be more selective about calls. If calling opportunities arrived independently throughout the month, this
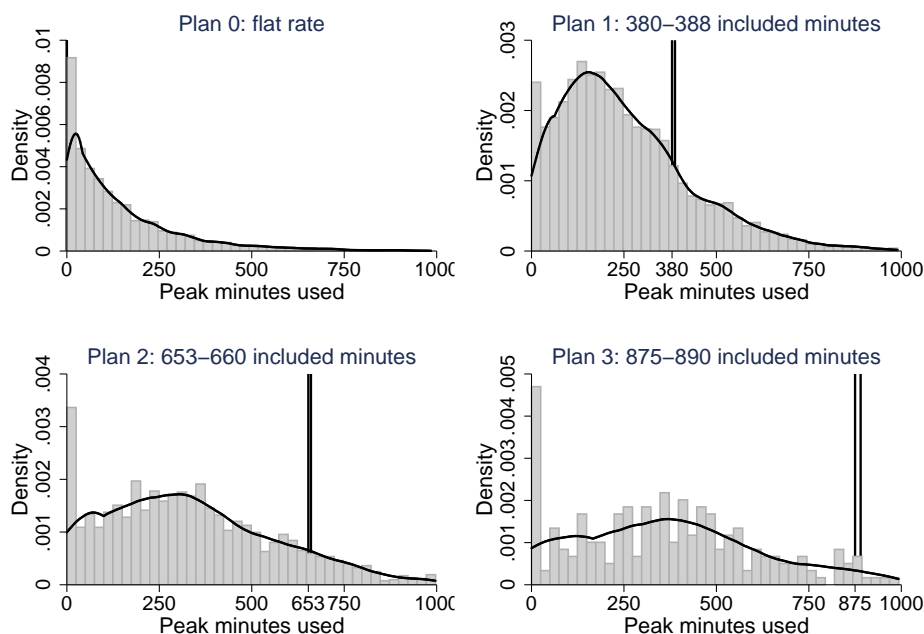
Figure 3: Usage densities for popular plans are constructed with 9,080, 5,026, 2,351, and 259 bills for plans 0-3 respectively. The sample for plans 1-3 is selected to only include bills for which in-network calls were costly and for which included peak minutes were within a narrow range, as indicated above each plot. Vertical lines bound the range of included free minutes for each plan.

strategic behavior by the consumer would lead to negative correlation between early and late usage within a billing period. However, looking for negative correlation in usage within the billing period is a poor test for this dynamic behavior because it is likely to be overwhelmed by positive serial correlation in taste shocks.

To test for dynamic behavior by consumers within the billing period, we use our data set of individual calls to construct both fortnightly and weekly measures of peak usage.[13] A simple regression of usage on individual fixed effects and lagged usage shows strong positive serial correlation. However, we take advantage of the following difference: Positive serial correlation between taste shocks in periods t and $(t-1)$ should be independent of whether periods t and $(t-1)$ are in the same or adjacent billing cycles. However, following unexpectedly high usage in period $(t-1)$, consumers should cut back usage more in period t if the two periods are in the same billing cycle. Thus by including an interaction effect between lagged usage and an indicator for the lag being in the same billing cycle as the current period, we can separate strategic behavior within the month from serial

---

[13]We divide each month into four weeks or two fortnights, and drop the extra 2-3 days between weeks 2 and 3.

correlation in taste shocks.

Table 1 shows a regression of log usage on lagged usage and the interaction between lagged usage and an indicator equal to 1 if period $(t-1)$ is in the same billing cycle as period t. We also include time and individual fixed effects and correct for bias induced by including both individual fixed effects and lags of the dependent variable in a wide but short panel (Roodman 2009). Reported analysis is for plan 1, the most popular three-part tariff. As expected, positive serial correlation in demand shocks leads to a positive and significant coefficient on lagged usage in the full sample (column 1) and most subsamples (columns 2-6). If consumers adjust their behavior dynamically within the billing cycle in response to usage shocks, then we expect the interaction effect to be negative. In the full sample (column 1) the interaction effect has a positive point estimate, but is not significantly different from zero. This suggests that consumers are not attentive to past usage during the course of the month.

Table 1: Dynamic usage pattern at fortnightly level.

	(1)	(2)	(3)	(4)	(5)	(6)
Overage Percentage	0-100%	0	1-29%	30-70%	71-99%	100%
$\ln(q_{t-1})$	0.649***	0.607***	0.535***	0.499***	-1.046	0.958***
	(0.0258)	(0.0529)	(0.0431)	(0.0683)	(1.065)	(0.0441)
SameBill*$\ln(q_{t-1})$	0.0133	0.0245	0.0193	-0.0149	-0.0837	3.685
	(0.0107)	(0.0183)	(0.0181)	(0.0222)	(1.180)	(4.745)
Observations	9068	3727	3218	1830	217	76
Number of id	386	167	130	87	11	6

Dependent variable $\ln(q_t)$. Standard errors in parentheses. Time and individual fixed effects.
Key: *** p<0.01, ** p<0.05, * p<0.1

Consumers who either never have an overage (43% of plan 1 subscribers) or always have an overage (3% of plan 1 subscribers) should be relatively certain what their ex post marginal price will be, and need not adjust calling behavior during the month. For instance, consumers who always make overages may only make calls worth more than the overage rate throughout the month. For such consumers we would expect to find no interaction effect, and this may drive the result when all consumers are pooled together as in our first specification. As a result, we divide consumers into groups by the fraction of times within their tenure that they have overages. We repeat our first specification for different *overage-risk groups* in Columns 2-6 of Table 1. The interaction effect is indistinguishable from zero in all overage risk groups. Moreover, in unreported analysis, more

flexible specifications that include nonlinear terms[14] and a similar analysis at the weekly rather than fortnightly level all estimate an interaction effect indistinguishable from zero. There is simply no evidence that we can find that consumers strategically cut back usage at the end of the month following unexpectedly high initial usage. We conclude that consumers are inattentive to their remaining balance of included minutes during the billing cycle.[15]

Our evidence for inattention is supported by Leider and Şahin's (2011) experimental work, which suggests that consumers who receive feedback about past usage do not follow an optimal dynamic program but instead use a constant calling threshold until all included minutes are used up and then adjust to the overage rate. This finding is consistent with our model of consumer behavior under our bill-shock counterfactual in which consumers are alerted when exceeding their allowance. In contrast, Yao et al. (2011) reject our static calling threshold model in favor of attentive dynamic behavior using Chinese cellular phone data.[16] The discrepancy between Yao et al.'s (2011) finding and our own may be due in part to the fact that, unlike consumers in our data, the Chinese consumers could check their minute balance. Moreover, results in all three papers can be reconciled by the fact that the financial incentives to pay attention were likely stronger for Chinese consumers than for American consumers and lab subjects.

3.2.2 Two stylized facts relevant to modeling plan choices

Two important features of the data are important to accurately model plan choice by cellular customers. First, while 29%-45% of contract choices are suboptimal ex post, consumers learn about their own usage levels over time and switch plans in response. Second, consumers' prior beliefs are biased: in the short run, before learning and switching plans, consumer plan-choice mistakes are predictable and can be exploited for profit. (We assume that consumers always make optimal plan-choices conditional on beliefs. When initial choices are suboptimal in a predictable way, we refer to consumers' prior beliefs as biased.)

Consumers switch plans. This may be in response to changes in tastes or prices but the pattern of switches shows that they are also made in response to learning. There are 1366 customers in our

---

[14]Average q_t will vary with expected marginal price, which is proportional to the probability of an overage. The probability of an overage in a billing period which includes periods t and $(t-1)$ increases nonlinearly in q_{t-1}. In one specification, we first fit a probit on the likelihood of an overage as a function of the first fortnights usage, and then used the estimated coefficients to generate overage probability estimates for all fortnights. We then included these (lagged) values as explanatory variables. In an alternative specification we added polynomial terms of lagged q_{t-1}.

[15]The finding is perhaps not surprising because service was resold by a university and, as a result, consumers could not contact the carrier to check minute balances.

[16]Yao et al. (2011) show that a scatter plot of cumulative weekly usage within a billing cycle against its lag is concave. In contrast, the relationship is linear in our data, which is consistent with our constant calling threshold.

data set, who we observe for an average of 12 months before either the data set ends or the customer quits.[17] Among all customers, 207 (15%) switch plans at least once, and 28 (2%) switch plans more than once, leading to a total of 246 plan switches. Of these switches, 85 (35%) are to plans that have either dropped in price or been newly introduced since the customer chose their existing plan. These switches could be motivated by price decreases rather than learning. However, the remaining 161 (65%) switches are to plans that are weakly more expensive than when the customer chose his or her existing plan. These switches must be due to learning or taste changes.

Not only do consumers switch plans, but they switch in the "right" direction. To substantiate this claim we make two calculations. First we calculate how much the customer would have saved had they signed up for the new plan initially, holding their usage from the original plan fixed. By this calculation, 60 to 61 percent of switches which can not be explained by price decreases saved customers money. (Switches that can not be explained by price decreases are those to plans which are weakly more expensive at the switching date than at the initial choice date.) Average savings, across money saving and money losing switches, are $11.03 to $15.44 per month.[18]

The savings estimates of $11.03 to $15.44 per month are underestimates because they do not take into account the fact that consumers can re-optimize usage choices upon switching plans. For instance, when switching to a plan with more included minutes consumers may optimally choose to talk more in response to the lower marginal price. An upper bound on the value of these additional calls is their price under the old plan. Hence our second calculation is the money that would have been lost had the customer not switched plans and remained on their original plan, again holding usage fixed. By this calculation average savings for switching are $24.42 to $31.84 per month, and 68 to 75 percent of switches saved money.[19] Hence consumers' expected benefit is between $11.03 and $31.84 per month when switching to plans that have not decreased in price since their previous choice, and 60 to 75 percent of switches are in the "right" direction.

---

[17]In our sample, 31 percent of customers are observed for more than 12 months. Standard cellular phone contracts often include switching costs (such as extension of commitment and delay of new phone subsidy) for switching plans prior to the expiry of one or two year contracts. In such a setting, more than 12 months of data would be needed to observe switching and learning. The students in our sample, however, could switch plans at any time and cancel after only three months, without any cost except hassle costs. As a result, we are able to observe active switching and learning over shorter time periods.

[18]We calculate bounds because we cannot always distinguish in-network and out-of-network calls. Both figures are statistically greater than zero at the 99% level. The 60-61 percent rates of switching in the "right" direction are statistically greater than 50 percent at the 95% level. This calculation is based on 98 of the 161 switches which can not be explained by price decreases. The remaining 63 switches occur so soon after the customer joins that there is no usage data prior to the switch that is not from a pro-rated bill.

[19]This calculation is based on 157 of the 161 switches which can not be explained by price decreases. The calculation cannot be made for the remaining 4 switches since there is no usage data following the switch that is not from a pro-rated bill. Figures are significant at the 99% confidence level.

In unreported analysis, additional evidence of learning is that: (1) the likelihood of switching declines with tenure, and (2) the likelihood of switching to a larger plan increases after an overage. Narayanan et al. (2007) estimate that consumers in the Kentucky experiment learn to switch up from overuse faster than they learn to switch down from underuse. In the context of retail banking, Ater and Landsman's (2011) results suggest that the asymmetry could be large enough that banking customers' tendency too choose overly large plans grows overtime through switching. For simplicity, we implement symmetric learning in our structural model.

Table 2: Predictable Customer Mistakes Yield Savings Opportunities

	First Opportunity	Second Opportunity
Dates	10/02-8/03	9/03 onwards
Enrollment Change	plan 1-3 $\rightarrow$ plan 0	plan 1 $\rightarrow$ plan 2
Affected Customers	251 (34%)	445 (55%)
Savings		
Total	$20,840 (47%)	$7,942 (28%)
Per Affected Bill	$8.76	$2.64
Per Affected Cust.	$83.03 (149%)	$17.85 (46%)

The University acts as a reseller and could bill students for their chosen plan, sign them up for an alternative plan, and save the difference in charges. These plan-level savings opportunities indicate that consumers choose overly risky plans (overconfidence or projection bias). Savings estimates are a lower bound because we cannot always distinguish in and out-of-network calls.

The presence of ex post mistakes alone shows only that consumers face uncertainty ex ante at the time of plan choice. However, ex post mistakes are not only present, they are also predictable given an individual's initial plan choice and population usage data. This implies that consumers' prior beliefs are biased and differ from average posteriors. Two plan-level savings opportunities demonstrate that customer mistakes are predictable and show how such predictability can be exploited by firms. (The first savings opportunity is an extension of that documented in Grubb (2009).) The university acts as a reseller and charges students a fixed five dollar fee per month to cover administrative costs. Although the university did not do so, they could have billed students based on the terms of their chosen calling plan, but signed them up for a predictably cheaper plan and saved the difference in charges. Table 2 illustrates two substantial opportunities. In the 2002-2003 academic year, when plan 0 offered free off-peak calling, by signing the 248 students who selected plans 1-3 up for plan 0, the university would have saved at least $20,731, or $83.59 per affected student. In the following year, the cellular company closed this opportunity by ending free off-peak calling on plan 0. However, an alternative was to sign up the 439 students who chose

plan 1 onto plan 2, which would have saved at least \$7,934, or \$18.07 per affected student. These plan-level savings opportunities indicate that consumers choose overly risky plans (overconfidence or volatility bias).[20]

4 Model

At each date t, consumer i first chooses a plan j and then chooses peak and off-peak quantities summarized by the vector $\mathbf{q}_{it} = (q_{it}^{pk}, q_{it}^{op})$. (The text suppresses the distinction between in-network and out-of-network calling, which is covered in Appendix C.) Total billable minutes for plan j are

$$q_{itj}^{billable} = q_{it}^{pk} + OP_j q_{it}^{op},$$

where OP_j is an indicator variable for whether plan j charges for off-peak usage. At the end of period t, consumer i is charged

$$P_j(\mathbf{q}_{it}) = M_j + p_j \max\{0, q_{itj}^{billable} - Q_j\},$$

where pricing plan j has monthly fee M_j, included allowance Q_j, and overage rate p_j.

We assume consumers are risk neutral, consumers have quasi-linear utility, and peak and off-peak calls are neither substitutes nor complements.[21] Consumer i's money-metric utility in month t from choosing plan j and consuming $\mathbf{q}_{it}$ units is

$$u_{itj} = \sum_{k \in \{pk, op\}} V\left(q_{it}^k, \theta_{it}^k\right) - P_j(\mathbf{q}_{it}) + \frac{1}{\alpha}\eta_{itj},$$

where

$$V\left(q_{it}^k, \theta_{it}^k\right) = \frac{1}{\beta}\left(\theta_{it}^k \ln\left(q_{it}^k/\theta_{it}^k\right) - q_{it}^k\right)$$

is the value from category $k \in \{pk, op\}$ calling, which depends on a pair of non-negative taste-shocks $\boldsymbol{\theta}_{it} = (\theta_{it}^{pk}, \theta_{it}^{op})$, and η_{itj} is an i.i.d. logit error.[22] The marginal value of a dollar is normalized to

---

[20]Aggregate and conditional mean biases could explain one or other plan-level savings opportunity but only overconfidence and volatility bias can simultaneously explain both savings opportunities. Note that the first savings opportunity is robust to dropping the top 30 percent of customers with the highest average savings, while the second savings opportunity is robust to dropping the top 2 percent of customers.

[21]In reality, consumers likely do delay calls until off-peak periods. Our assumption ruling out such substitution should not bias our final results. In particular, as off-peak calling is typically free and is exogenously so in our counterfactual simulations, whether peak calls are foregone entirely or shifted off-peak does not effect firm revenues or peak-pricing. Moreover, in either case, foregone peak calls carry a social cost captured in our welfare estimates.

[22]We model consumers' choice between the four most popular pricing plans (plans 0-3), comparable plans from other

one, $1/\alpha$ scales the logit error variance, and β is a price coefficient that determines how sensitive calling choices are to the marginal price of an additional minute of calling time. Our choice of functional form for $V\left(q_{it}^k, \theta_{it}^k\right)$ implies that the taste shock θ_{it}^k enters demand multiplicatively, as discussed below.

4.1 Quantity Choices

Recognizing that consumers are uncertain about the ex post marginal price when making usage choices from three-part tariffs is a key feature of our model and where we take a new approach (also suggested independently by Borenstein (2009)). We assume that at the start of billing period t, consumer i is uncertain about her period t taste shock $\boldsymbol{\theta}_{it}$. She first chooses a plan j and then chooses a calling threshold vector $\mathbf{v}_{itj}^* = (v_{itj}^{pk}, v_{itj}^{op})$ based on chosen plan terms and her beliefs about the distribution of $\boldsymbol{\theta}_{it}$. During the course of the month, the consumer is inattentive and does not track usage but simply makes all category-k calls valued above v_{itj}^k. Over the course of the month, for $k \in \{pk, op\}$ this cumulates to the choice:

$$q_{it}^k = q(v_{itj}^k, \theta_{it}^k) = \theta_{it}^k \hat{q}(v_{itj}^k), \tag{1}$$

where $\hat{q}(v) = 1/(1 + \beta v)$ and $\hat{q}(0) = 1.$[23]

The interpretation is that θ_{it}^k is the volume of category-k calling opportunities that arise and $\hat{q}(v)$ is the fraction of those calling opportunities worth more than v per minute. Timing is summarized in Figure 4. Figure 5 shows the calling threshold v_{itj}^{pk} and resulting consumption choice $\theta_{it}^{pk} \hat{q}(v_{itj}^{pk})$ in relation to a consumer's realized inverse demand curve for calling minutes, $V_q(q_{it}^{pk}, \theta_{it}^{pk})$.
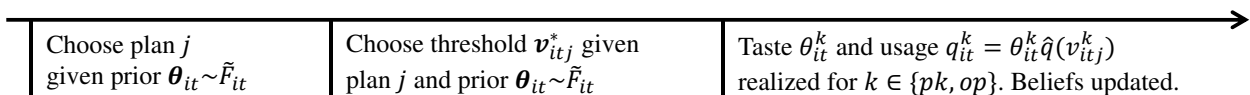
Choose plan j given prior $\boldsymbol{\theta}_{it} \sim \tilde{F}_{it}$	Choose threshold $\boldsymbol{v}_{itj}^*$ given plan j and prior $\boldsymbol{\theta}_{it} \sim \tilde{F}_{it}$	Taste θ_{it}^k and usage $q_{it}^k = \theta_{it}^k \hat{q}(v_{itj}^k)$ realized for $k \in \{pk, op\}$. Beliefs updated.

Figure 4: Model Time Line

---

carriers, and an outside option. For plans other than the four popular university plans, the logit error η_{itj} has a clear economic interpretation: it includes all unmodeled plan heterogeneity including network quality, available phones, and roaming charges. Within the four popular plans, the logit error η_{itj} has no satisfactory economic interpretation, as these plans only differ in price, and in the complete model we capture all the dimensions on which prices differ. All initial plan choices could be explained without including the logit error, but they are required to explain switches that appear to be in the "wrong" direction.

[23]The fact that demand is multiplicative in θ_{it}^k follows from the assumption that $V\left(q_{it}, \theta_{it}^k\right)$ can be expressed as $V\left(q_{it}, \theta_{it}^k\right) = \theta_{it}^k \hat{V}\left(q_{it}/\theta_{it}^k\right)$ for some function $\hat{V}$. In this case, $\hat{V}(x) = (\ln x - x)/\beta$. The fact that $\hat{q}(0) = 1$ simply reflects the chosen normalization of θ_{it}^k.
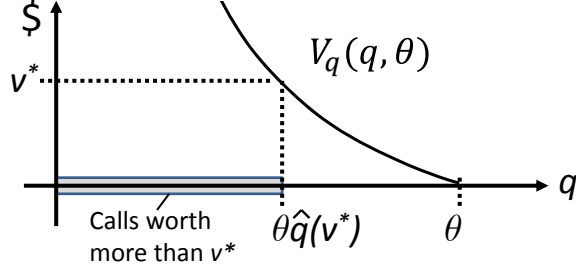
Figure 5: Inverse Demand Curve and Calling Threshold

Making all peak calls valued above the constant threshold v_{itj}^* is the optimal strategy of an inattentive consumer who does not track usage within the current billing cycle and hence cannot update his beliefs about the likelihood of an overage within the current billing cycle. (It is analogous to an electricity consumer setting a thermostat rather than choosing a quantity of kilowatt hours.)

When marginal price is constant, a consumer's optimal calling threshold is simply equal to the marginal price. Thus for plan zero, which charges 11 cents per minute for all billable calls, $\mathbf{v}_{itj}^* = (0.11, 0.11 OP_j)$. Further, $v_{itj}^{op} = 0$ for plans 1-3 because they offer free off-peak calling.

Conditional choosing one of plans 1-3, which include free off-peak calling and an allowance of peak minutes, consumer i chooses her period t peak-calling threshold v_{itj}^{pk} to maximize her expected utility conditional on her period t information $\Im_{it}$. Given allowance Q_j, overage rate p_j, and multiplicative demand (equation (1)), the optimal threshold (derived in Appendix B.1) is uniquely characterized by equation (2):

$$
v_{itj}^{pk} = p_j \Pr\left(\theta_{it}^{pk} \geq Q_j/\hat{q}(v_{itj}^{pk}) \mid \Im_{it}\right) \frac{E\left[\theta_{it}^{pk} \mid \theta_{it}^{pk} \geq Q_j/\hat{q}(v_{itj}^{pk}); \ \Im_{it}\right]}{E\left[\theta_{it}^{pk} \mid \Im_{it}\right]}. \tag{2}
$$

The threshold v_{itj}^{pk} will be between zero and the overage rate p_j.[24]

Note that choosing threshold v_{itj}^{pk} is equivalent to choosing a target peak-calling quantity $q_{it}^T \equiv E[\theta_{it}^{pk}]\hat{q}(v_{itj}^{pk})$, which is implemented with error $(\theta_{it}^{pk} - E[\theta_{it}^{pk}])\hat{q}(v_{itj}^{pk})$. Importantly, consumers are aware of their inability to hit the target precisely and take this into account when making their threshold/target choice.

---

[24]Equation (2) may seem counter-intuitive, because the optimal v_{itj}^{pk} is greater than the expected marginal price, $p_j \Pr(q(v_{itj}^{pk}, \theta_{it}^{pk}) > Q_j \mid \Im_{it})$. This is because the reduction in consumption from raising v_{itj}^{pk} is proportional to θ_{it}^{pk}. Raising v_{itj}^{pk} cuts back on calls valued at v_{itj}^{pk} more heavily in high demand states when they cost p_j and less heavily in low demand states when they cost 0.

4.2 Plan Choices

We model consumers' choice between the four most popular pricing plans (plans 0-3), comparable AT&T, Cingular, and Verizon plans (Sprint offered no local plans), and an outside option which incorporates all other plans. We adopt Ching, Erdem and Keane's (2009) consideration set model by assuming that consumers make an active choice with exogenous probability P_C and keep their current plan with probability $(1 - P_C)$. We use the frequency of failures to switch away from dominated plans to identify P_C.[25]

Customer i's perceived expected utility from choosing plan j at date t is

$$U_{itj} = E\left[\sum_{k \in \{pk, op\}} V\left(q(v_{itj}^k, \theta_{it}^k), \theta_{it}^k\right) - P_j\left(\mathbf{q}(\mathbf{v}_{itj}^*, \boldsymbol{\theta}_{it})\right) \mid \Im_{it}\right] + \frac{1}{\alpha}\eta_{itj}, \qquad (3)$$

and from choosing the outside option is $U_{it0} = O + \eta_{it0}$. The parameter O will be identified from the frequency at which consumers leave the data set. Conditional on making an active choice, a consumer's consideration set includes plans offered by her current provider, the outside option, and plans from a randomly selected alternative carrier.[26] Consumers myopically[27] choose the plan (or outside option) from their consideration set that maximizes expected utility in the current period.

4.3 Distribution of Tastes

We assume that the non-negative taste-shocks which determines usage are latent taste shocks censored at zero:

$$\theta_{it}^k = \begin{cases} 0 & \tilde{\theta}_{it}^k < 0 \\ \tilde{\theta}_{it}^k & \tilde{\theta}_{it}^k \geq 0 \end{cases}, \ k \in \{pk, op\}.$$

---

[25]When prices fall consumers often do not switch away from their existing plans even when they are now dominated by plans on the current menu. For instance, most consumers paying $54.99 for 890 minutes on plan 3 do not switch to plan 2 during the one month promotion in April 2004 when it offered 1060 minutes for only $44.99. We believe this is because consumers who are not actively making a plan choice do not find out about the price cuts.

[26]We avoid including all plans in the consideration set to reduce computational time.

[27]We assume learning is independent of plan choice, so there is no value to experimentation with an alternative plan. Nevertheless, myopic plan choice is not optimal for several reasons. First, when a consumer is currently subscribed to a plan that is no longer offered (and is not dominated) there is option value to not switching, since switching plans will eliminate that plan from future choice sets. Second, if $P_C < 1$, a forward looking consumer would tend to discount her current period logit-error η_{it}. Third, if $P_C < 1$, a forward looking consumer should anticipate that her current plan choice may persist in the future but her future calling threshold choices v^* will improve as she learns about her type μ_i. This consideration makes plans 1 and 2 marginally more attractive relative to plans 0 and 3 but the effect is not large. We ignore these issues for tractability.

We assume that the latent shock $\tilde{\theta}_{it}^k$ is normally distributed and that consumers observe its value even when censored. This adds additional unobserved heterogeneity to the model but preserves tractable Bayesian updating. Censoring makes zero usage a positive likelihood event, which is important since it occurs for 10% of plan 0 observations.

Usage choices in the data are strongly serially-correlated conditional on customer-plan and date fixed effects. We therefore incorporate simple serial-correlation into our model by assuming that the latent shock $\tilde{\boldsymbol{\theta}}_{it}$ follows a stationary AR1 process with a bivariate normal innovation,

$$\tilde{\boldsymbol{\theta}}_{it} = \boldsymbol{\mu}_i + \varphi\tilde{\boldsymbol{\theta}}_{i,t-1} + \boldsymbol{\varepsilon}_{it},$$

where $\boldsymbol{\mu}_i$ is customer i's true type, φ is the common serial coefficient, and $\boldsymbol{\varepsilon}_{it} \sim N(0, \boldsymbol{\Sigma}_\varepsilon)$ is the normally-distributed mean-zero innovation with variance-covariance matrix

$$\boldsymbol{\Sigma}_\varepsilon = \begin{bmatrix} (\sigma_\varepsilon^{pk})^2 & \rho_\varepsilon \sigma_\varepsilon^{pk} \sigma_\varepsilon^{op} \\ \rho_\varepsilon \sigma_\varepsilon^{pk} \sigma_\varepsilon^{op} & (\sigma_\varepsilon^{op})^2 \end{bmatrix}.$$

(We assume AR(1) rather than AR(k) for simplicity.) Consumers' true types, $\boldsymbol{\mu}_i = (\mu_i^{pk}, \mu_i^{op})$, are normally distributed across the population as described below.

4.4 Beliefs and Learning

Estimation of consumer beliefs and learning is focused on a single dimension of usage: total peak-calling. We make this restriction because plans 1-3 always offer free off-peak calling and hence the choice data are not rich enough to allow us to identify beliefs about off-peak calling. For simplicity, we assume that while consumers are learning about their peak type μ_i^{pk} over time, there is no learning about off-peak demand because consumers know their off-peak types μ_i^{op}.[28]

We assume the serial-correlation coefficient φ is known by all consumers. While taste innovations $\boldsymbol{\varepsilon}_{it}$ have variance-covariance $\boldsymbol{\Sigma}_\varepsilon$, consumers believe the variance-covariance matrix is

$$\tilde{\boldsymbol{\Sigma}}_\varepsilon = \begin{bmatrix} (\tilde{\sigma}_\varepsilon^{pk})^2 & \rho_\varepsilon \tilde{\sigma}_\varepsilon^{pk} \sigma_\varepsilon^{op} \\ \rho_\varepsilon \tilde{\sigma}_\varepsilon^{pk} \sigma_\varepsilon^{op} & (\sigma_\varepsilon^{op})^2 \end{bmatrix},$$

where $\tilde{\sigma}_\varepsilon^{pk} = \delta_\varepsilon \sigma_\varepsilon^{pk}$ and $\delta_\varepsilon > 0$. If $\delta_\varepsilon = 1$, then consumers' perceptions match reality. If $\delta_\varepsilon < 1$, then consumers underestimate the volatility of their peak tastes from month-to-month and exhibit

---

[28]This assumption does not effect our endogenous-price counterfactual simulations because we assume free off-peak calling.

volatility bias. If $\delta_\varepsilon < 1$, then consumers will predictably choose too risky plans and overreact to past usage when deciding whether or not to switch plans.[29] Consumer beliefs about the variance of off-peak tastes and the correlation between peak and off-peak tastes are both correct.

Consumers learn about their own peak-type μ_i^{pk} over time. At date t, consumer i believes that μ_i^{pk} is normally distributed with mean $\tilde{\mu}_{i,t}^{pk}$ and variance $\tilde{\sigma}_t^2$: $\mu_i^{pk}|\Im_{i,t} \sim N(\tilde{\mu}_{i,t}^{pk}, \tilde{\sigma}_t^2)$. At the end of each billing period, usage q_{it}^{pk} is realized and consumers can infer $\theta_{it}^{pk} = q_{it}^{pk}/\hat{q}(v_{itj}^{pk})$. When $q_{it}^{pk} = \theta_{it}^{pk} = 0$, we assume that consumers can observe the latent taste shock $\tilde{\theta}_{it}^{pk}$. The latent shock provides an unbiased normal signal about μ_i^{pk} and consumers update beliefs according to Bayes rule (see Appendix B.3).[30] Over time consumers learn their own types: $\tilde{\mu}_{i,t}^{pk}$ converges to μ_i^{pk} and $\tilde{\sigma}_t^2$ converges to zero.

Consumers' plan choices and threshold choices depend on beliefs about the distribution of tastes $\boldsymbol{\theta}_{it}$. When choosing a plan and a usage threshold for the first time, consumers believe:

$$\tilde{\theta}_{i1}^{pk} \sim N\left(\frac{\tilde{\mu}_{i1}^{pk}}{1-\varphi}, \tilde{\sigma}_{\theta 1}^2\right), \tag{4}$$

where

$$\tilde{\sigma}_{\theta 1}^2 = \frac{\tilde{\sigma}_1^2}{(1-\varphi)^2} + \frac{(\tilde{\sigma}_\varepsilon^{pk})^2}{1-\varphi^2}. \tag{5}$$

In all later periods $t > 1$, when consumers can condition on $\tilde{\theta}_{i,t-1}^{pk}$, beliefs are:

$$\tilde{\theta}_{it}^{pk} \mid \Im_{it} \sim N\left(\tilde{\mu}_{it}^{pk} + \varphi\tilde{\theta}_{it-1}^{pk}, \tilde{\sigma}_t^2 + (\delta_\varepsilon \sigma_\varepsilon^{pk})^2\right).$$

Following a month with surprisingly high usage, consumer i's belief about the distribution of demand in the following month increases for two reasons. First the consumer increases his estimate of his type ($\tilde{\mu}_{i,t+1}^{pk} > \tilde{\mu}_{it}^{pk}$), and second he knows that his demand is positively correlated over time. In the standard model the only behavior change that might result is a switch to a larger plan. In our model, a consumer might also switch to a larger plan but, conditional on not switching, would cut back on usage by choosing a higher calling threshold ($v_{i,t+1}^{pk} > v_{i,t}^{pk}$) and being more selective about calls.

---

[29]For tractability, we assume that consumers learn about means but not variances, so volatility bias is persistent.

[30]In fact, given our assumption that consumers know μ_i^{op}, consumers can also infer ε_{it}^{op} from off peak usage which is informative about μ_i^{pk} because it is correlated with ε_{it}^{pk}. We assume consumers only update beliefs using θ_{it}^{pk} and not ε_{it}^{op}. This choice is conservative in the sense that our finding that consumers respond to data too little is biased downwards. It is also realistic for two reasons. First, consumers are unlikely to pay attention to off-peak usage when they are on contract with free off-peak calls. Second, we only assume consumers know μ_i^{op} for simplicity as we cannot identify off-peak beliefs. In reality, consumers are unlikely to know μ_i^{op} so cannot actually infer ε_{it}^{op}.

4.5 Priors

Each customer is characterized by the individual specific triple $\{\tilde{\mu}_{i1}^{pk}, \mu_i^{pk}, \mu_i^{op}\}$. Together with the population parameter $\tilde{\sigma}_1^2$, this triple specifies each customer's true type $\boldsymbol{\mu}_i$ and prior belief $\mu_i^{pk} \sim N(\tilde{\mu}_{i1}^{pk}, \tilde{\sigma}_1^2)$. (Consumers are assumed to know their own off-peak types.) The population is described by the joint distribution of $\{\tilde{\mu}_{i1}^{pk}, \mu_i^{pk}, \mu_i^{op}\}$, which we assume is a trivariate normal distribution. We outline the nine parameters of the distribution at the end of the section. By not restricting any of these parameters, our formulation allows for consumer beliefs about their peak types to be biased in three ways, as described below.

First, aggregate mean bias can arise when the average point estimate is too low or too high. We define μ_0^{pk} and $\tilde{\mu}_0^{pk}$ to be the population averages of true peak types μ_i^{pk} and prior point estimates $\tilde{\mu}_{i1}^{pk}$, respectively. A typical assumption (perhaps labeled rational expectations) is that $\mu_0^{pk} = \tilde{\mu}_0^{pk}$, which implies that an average individual's initial point estimate is an unbiased estimate of her true type. We do not impose this assumption. If $b_1 \equiv \tilde{\mu}_0^{pk} - \mu_0^{pk} \neq 0$, then there is aggregate mean bias and consumers will predictably choose plans which are too small ($b_1 < 0$) or too large ($b_1 > 0$).

Second, overconfidence can arise when the precision of consumers' beliefs about their type is miscalibrated. In our notation, $\sigma_{\mu^{pk}}^2$ is the conditional variance of true peak types in the population, $Var(\mu_i^{pk} \mid \tilde{\mu}_{i1}^{pk})$, and $\tilde{\sigma}_1$ is consumers' uncertainty about their peak type. A typical assumption (perhaps labeled rational expectations) is that $\tilde{\sigma}_1 = \sigma_{\mu^{pk}}$.[31] We do not impose this assumption either. If $\delta_\mu \equiv \tilde{\sigma}_1 / \sigma_{\mu^{pk}} < 1$ then consumers exhibit overconfidence: they underestimate their own uncertainty about their type μ_i^{pk}. Overconfident consumers, like those with volatility bias, will predictably choose overly risky plans. However, in contrast to those with volatility bias, they will under-react to past usage when making plan switching decisions. Grubb's (2009) analysis is static, so could not distinguish between overconfidence and volatility bias, but found that customers do choose overly risky plans, so exhibit either overconfidence, volatility bias, or both.

Third, conditional mean bias can arise when consumers over or under react to their own private information, forming individual point estimates, $\tilde{\mu}_{i1}^{pk}$, that differ from the population average, $\tilde{\mu}_0^{pk}$, too much or too little. Conditional on $\tilde{\mu}_{i1}^{pk}$, we write the population average of μ_i^{pk} as $E\left[\mu_i^{pk} \mid \tilde{\mu}_{i1}^{pk}\right] = \mu_0^{pk} + \psi^{pk}(\tilde{\mu}_{i1}^{pk} - \tilde{\mu}_0^{pk})$.[32] Letting $b_2 \equiv 1 - \psi^{pk}$, we have:

$$\tilde{\mu}_{i1}^{pk} - E\left[\mu_i^{pk} \mid \tilde{\mu}_{i1}^{pk}\right] = b_1 + b_2(\tilde{\mu}_{i1}^{pk} - \tilde{\mu}_0^{pk}). \tag{6}$$

---

[31] An alternative rational-expectations benchmark discussed in Appendix B.4 would be $\tilde{\sigma}_1^2 = Var(\mu_i^{pk} \mid \tilde{\mu}_{i1}^{pk}, \mu_i^{op})$. Using this benchmark would affect the description of bias but not its economic consequences: it would not alter our evaluation of bill-shock regulation or any welfare results outside of our de-biasing counterfactual simulations.

[32] Implicitly this defines ψ^{pk} as $Cov(\mu_i^{pk}, \tilde{\mu}_{i1}^{pk}) / Var(\tilde{\mu}_{i1}^{pk})$.

A typical assumption is that $\tilde{\mu}_{i1}^{pk} = E[\mu_i^{pk} \mid \tilde{\mu}_{i1}^{pk}]$, or $b_1 = b_2 = 0$, which implies that all individuals' initial point estimates are unbiased estimates of their true types.[33] (Assuming $b_1 = 0$ only makes the weaker restriction that the average individual's point estimate is unbiased.) The parameter b_2 measures the amount of conditional mean bias in the population. If $b_2 > 0$ then point estimates differ too much from the population average and consumers predictably choose plans which are too extreme. If $b_2 < 0$ then point estimates differ too little from the population average and consumers predictably choose plans which are too moderate.

For completeness, we finish by describing the joint normal distribution of $\{\tilde{\mu}_{i1}^{pk}, \mu_i^{pk}, \mu_i^{op}\}$ in two parts. First, the marginal distribution of initial point estimates $\tilde{\mu}_{i1}^{pk}$ is normal with mean $\tilde{\mu}_0^{pk}$ and variance $\tilde{\sigma}_{\mu^{pk}}^2$. Second, conditional on the point estimate $\tilde{\mu}_{i1}^{pk}$, the population distribution of true types $\boldsymbol{\mu}_i$ is normal with mean $\boldsymbol{\mu}_0 + \boldsymbol{\psi}(\tilde{\mu}_{i1}^{pk} - \tilde{\mu}_0^{pk})$ and variance matrix $\boldsymbol{\Sigma}_\mu$. The vectors $\boldsymbol{\mu}_0$ and $\boldsymbol{\psi}$ are defined as (μ_0^{pk}, μ_0^{op}) and (ψ^{pk}, ψ^{op}), respectively; the diagonal elements of $\boldsymbol{\Sigma}_\mu$ are $\sigma_{\mu^{pk}}^2$ and $\sigma_{\mu^{op}}^2$, while the off-diagonal is $\rho_\mu \sigma_{\mu^{pk}} \sigma_{\mu^{op}}$, where ρ_μ is the conditional correlation of μ_i^{pk} and μ_i^{op}. Note that the joint distribution of true types and priors we describe can naturally be generated from the marginal distribution of true types, a common prior, and an unbiased signal that accounts for aggregate uncertainty. This is the presentation adopted by Goettler and Clay (2011).

5 Identification

Parameters can be categorized into four groups: (1) parameters governing beliefs ($\tilde{\mu}_0^{pk}$, $\tilde{\sigma}_{\mu^{pk}}$, $\tilde{\sigma}_1$, and $\tilde{\sigma}_\varepsilon^{pk}$), (2) the true (conditional) distribution of tastes ($\boldsymbol{\mu}_0$, $\boldsymbol{\psi}$, $\boldsymbol{\Sigma}_\mu$, $\boldsymbol{\Sigma}_\varepsilon$, and φ), (3) the price coefficient β, and (4) parameters related to switching and quitting (P_C, α, and O). Broadly speaking, plan choices identify beliefs, the distribution of actual usage identifies the distribution of true tastes, and changes in usage in response to the discontinuous change in marginal price between peak and off-peak hours identify the price coefficient β. Finally, the rate of switching away from dominated plans, the rate of switching in the "wrong" direction, and the rate of quitting identify, respectively, the active choice probability P_C, the logit error weight $1/\alpha$, and the outside option O.

5.1 Price Coefficient

If consumers' chosen thresholds ($\mathbf{v}_{it}^*$) were known, the price coefficient β could be inferred from marginal price variation and the induced variation in $\hat{q}(v_{it}^k)$.[34] Unfortunately, we require β to

---

[33] An alternative benchmark discussed in Appendix B.4 would be $\tilde{\mu}_{i1}^{pk} = E[\mu_i^{pk} \mid \tilde{\mu}_{i1}^{pk}, \mu_i^{op}]$. See foonote 31.

[34] For instance, there is one clean experiment in the data in which existing plan 1 subscribers were automatically upgraded from 280 free minutes to 380 free minutes and increased their usage in response by an average of 53 minutes.

calculate $\mathbf{v}^*_{it}$. We circumvent this problem by relying on a source of marginal price variation for which v^*_{it} is known. Prior to fall 2003, v^*_{it} is 11 cents during peak hours and 0 cents during off-peak hours for plan 0 subscribers.

Although prices in the model depend only on total peak and total off-peak calling, we additionally break out the share of calling demand for weekday outgoing-calls to landlines immediately before and after 9pm to help identify the price coefficient. The shock $\mathbf{r}^{9pm}_{it} = (r^{9pk}_{it}, r^{9op}_{it}) \in [0,1]^2$ captures the share of peak and off-peak calling demand that is within 60 minutes of 9pm on a weekday and is for an outgoing call to a landline. The distribution of r^k_{it} for $k \in \{9pk, 9op\}$ is a censored normal,

$$
\begin{aligned}
\tilde{r}^k_{it} &= \alpha^k_i + e^{r,k}_{it} \\
r^k_{it} &= \begin{cases} 0 & \text{if } \tilde{r}^k_{it} \leq 0 \\ \tilde{r}^k_{it} & \text{if } 0 < \tilde{r}^k_{it} < 1 \\ 1 & \text{if } \tilde{r}^k_{it} \geq 1 \end{cases},
\end{aligned}
$$

where α^k_i is unobserved heterogeneity and $e^{r,k}_{it}$ is a mean-zero shock normally distributed with variance $(\sigma^k_e)^2$ independent across i, t, and k. We assume that α^{9pk}_i is normally distributed in the population with mean μ^{9pk}_α and variance $(\sigma^{9pk}_\alpha)^2$.

Our identifying assumption for the price coefficient is that consumer i's expected outgoing calling demand to landlines on weekdays is the same between 8:00pm and 9:00pm as it is between 9:00pm and 10:00pm:

$$
E\left[r^{9pk}_{it}\right] E\left[\theta^{pk}_{it}\right] = E\left[r^{9op}_{it}\right] E\left[\theta^{op}_{it}\right]. \tag{7}
$$

In other words, we assume that the increase in observed calling to landlines on weekdays immediately after off-peak begins at 9pm is a price effect rather than a discontinuous increase in demand at 9pm.[35] As a result, equation (7) implicitly defines α^{9op}_i as a function of α^{9pk}_i and other parameters.

Given plan 0 pricing prior to fall 2003, $\theta^{op}_{it} = q^{op}_{it}$ and $\theta^{pk}_{it} = q^{pk}_{it}(1 + 0.11\beta)$. Moreover, the pre and post 9pm calling shares are always observed because calling thresholds are constant within peak and within off-peak hours: $r^{9op}_{it} = q^{9op}_{it}/q^{op}_{it}$ and $r^{9pk}_{it} = q^{9pk}_{it}/q^{pk}_{it}$. Thus equation (7) can be

---

(The 95% confidence interval on this increase is 26-81 minutes.) However, without knowing how consumer thresholds were affected by the price change, this does not identify β.

[35]We focus on calls to landlines because the other party to the call pays nothing both before and after 9pm. The assumption would be unreasonable for calls to or from cellular numbers since such calling opportunities increase at 9pm when the calls become cheaper for the other party and the other party is more likely to call or answer.

solved for β as a function of moments of the data:

$$\beta = \frac{100}{11} \left(\frac{E\left[q_{it}^{9op}/q_{it}^{op}\right] E\left[q_{it}^{op}\right]}{E\left[q_{it}^{9pk}/q_{it}^{pk}\right] E\left[q_{it}^{pk}\right]} - 1 \right).$$

5.2 Serial Correlation

Data prior to fall 2003 identifies the AR1 coefficient φ. During this period, all plans offered free nights-and-weekends so that we observe

$$q_{it}^{op} = \theta_{it}^{op} = \mu_i^{op} + \varphi\theta_{it-1}^{op} + \epsilon_{it}^{op}. \tag{8}$$

The argument follows the identification argument for the parameters of a linear regression model with person level fixed effects and a lagged dependent variable. By taking the first difference of equation (8), we remove the impact of the fixed effect μ_i^{op}. Then φ can be estimated using past values of θ_{it}^{op} as instruments, as in Blundell and Bond (1998).

5.3 Beliefs

Next, consider identification of consumers' prior beliefs from plan choices. Choice data are quite informative about beliefs about peak usage, as illustrated by Figure 6, but relatively uninformative about beliefs about off-peak usage. Hence we assume consumers know their own off-peak taste distribution (including μ_i^{op} and σ_ε^{op}). Prior to fall 2003, when off-peak calling is free, an individual consumer's initial plan choice depends only on β, η_{i1}/α, and her beliefs about θ_{i1}^{pk} described by $\tilde{\mu}_{i1}^{pk}/(1-\varphi)$ and $\tilde{\sigma}_{\theta 1}$. Thus initial plan-choice shares depend only on α, β, φ, $\tilde{\sigma}_{\theta 1}$, and the population distribution of $\tilde{\mu}_{i1}^{pk}$, described by $\tilde{\mu}_0^{pk}$ and $\tilde{\sigma}_{\mu^{pk}}^2$. Parameters φ and β are already identified. For transparency of the argument, we begin by considering a restricted model that excludes logit errors ($1/\alpha = 0$). Initial plan choice shares identify the remaining parameters $\tilde{\mu}_0^{pk}$, $\tilde{\sigma}_{\mu^{pk}}^2$, and $\tilde{\sigma}_{\theta 1}$. Finally, the learning rate separately identifies $\tilde{\sigma}_1$ and $\tilde{\sigma}_\varepsilon$ from $\tilde{\sigma}_{\theta 1}$. Initial choice shares in post fall 2003 data also aid identification, but require a more complicated argument involving beliefs about off-peak tastes.

Absent the logit-error, initial plan choices place bounds on each individual's prior beliefs about the mean ($\tilde{\mu}_{i1}^{pk}/(1-\varphi)$) and variance ($\tilde{\sigma}_{\theta 1}^2$) of their first taste shock, $\tilde{\theta}_{i1}^{pk}$. (Recall $\tilde{\sigma}_{\theta 1}^2$ is related to model parameters by equation (5).) Based on October-November 2002 pricing data (ignoring free in-network calling), Figure 6 (top panel) shows plan-choice as a function of prior beliefs $\{\tilde{\mu}_{i1}^{pk}, \tilde{\sigma}_{\theta 1}^2\}$ given $\beta = 4$ and $\varphi = 1/2$. Consumers joining in October-November 2002 with beliefs in the gray
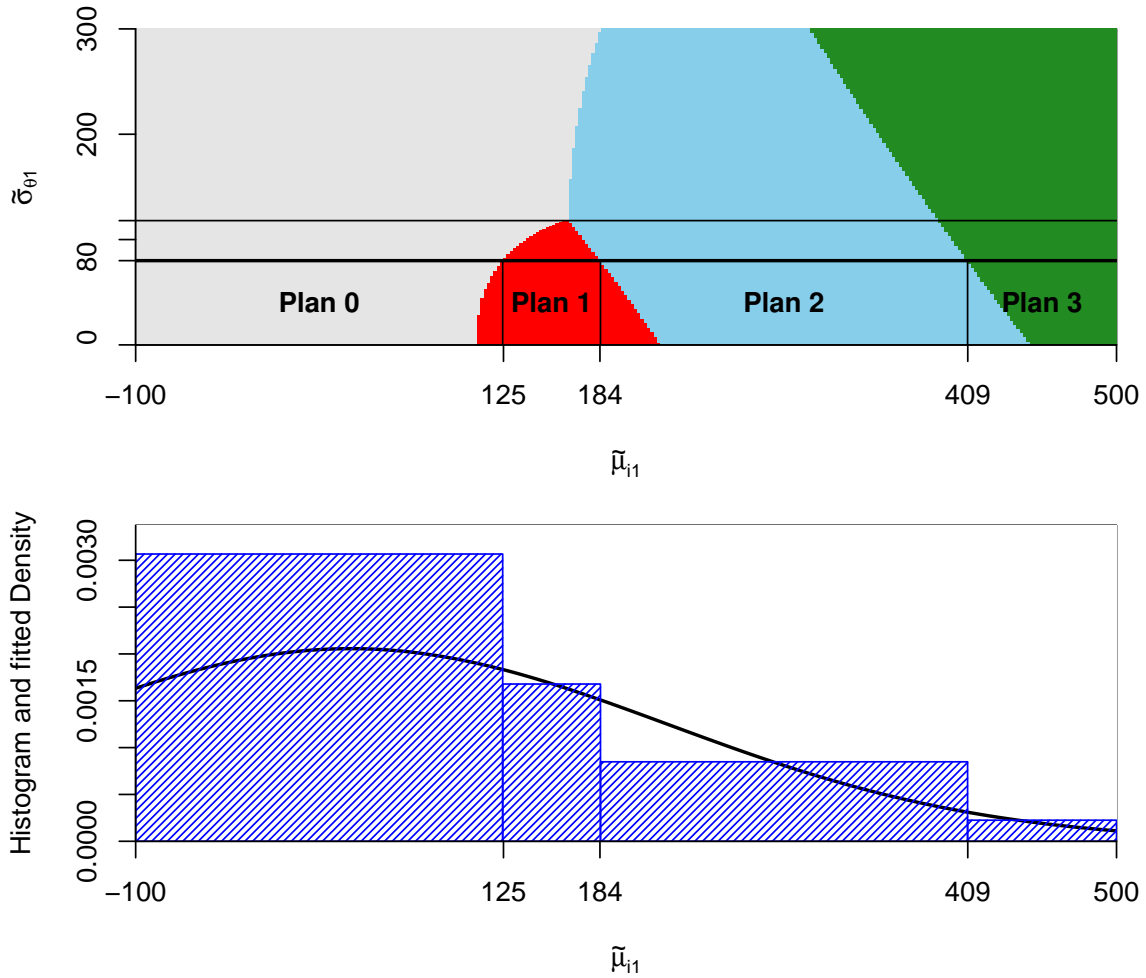
Figure 6: Top panel: Plan choice as a function of initial beliefs $\{\tilde{\mu}_{i1}, \tilde{\sigma}_{\theta 1}\}$ implied by the model evaluated at October-November 2002 prices given $\beta = 4$ and $\varphi = 1/2$. Bottom panel: Histogram and fitted normal distribution over $\tilde{\mu}_{i1}$ implied by the assumption $\tilde{\sigma}_{\theta 1} = 80$ and October-November 2002 new subscriber plan choice shares of 69%, 10%, 19%, and 2% for plans 0 to 3 respectively.

region choose plan 0, those with beliefs in the red region choose plan 1, those with beliefs in the blue region choose plan 2, and those with beliefs in the green region choose plan 3. This means that observing a new customer in October-November 2002 choose plan j will bound her beliefs to be within the relevant colored region.

Figure 6 shows that plan 0 is chosen both by individuals with low expectations of usage (low $\tilde{\mu}_{i1}^{pk}$), as it has the lowest fixed fee, and by individuals with high uncertainty about usage (high $\tilde{\sigma}_{\theta 1}$), as it never charges more than 11 cents per minute and is therefore a safe option. Figure 6 shows that for any $\tilde{\sigma}_{\theta 1}$ larger than 118, plan 1 is never chosen. Thus the assumption that $\tilde{\sigma}_{\theta 1}$ is common across individuals and the fact that a sizable fraction of individuals chose plan 1 in

October-November 2002 puts an upper bound on $\tilde{\sigma}_{\theta 1}$ of 118.

If we were to fix $\tilde{\sigma}_{\theta 1}$ at any level below 118, individual i's plan choice bounds $\tilde{\mu}_{i1}^{pk}$ to an interval. For instance, if overconfidence and volatility bias were complete ($\tilde{\sigma}_1 = \delta_\mu = \delta_\varepsilon = 0$) so that consumers believed they could predict their usage perfectly ($\tilde{\sigma}_{\theta 1} = 0$) and consumers were inelastic ($\beta = 0$), then consumers would choose from the lower envelope of the tariff menu, and initial choice of plan j would imply the following bounds on the prior point estimate $\tilde{\mu}_{i1}^{pk}$:

$$(M_j - M_{j-1})/p_{j-1} + Q_{j-1} \leq \frac{\tilde{\mu}_{i1}^{pk}}{1 - \varphi} \leq (M_{j+1} - M_j)/p_j + Q_j.$$

For $\tilde{\sigma}_{\theta 1}$ and β strictly positive, the bounds do not have an analytical solution but can be read from the corresponding horizontal slice of Figure 6. For example, the bounds are given for $\tilde{\sigma}_{\theta 1} = 80$ by the vertical lines in Figure 6. Combining plan share data from customers who join in October-November 2002 with these bounds generates of histogram over $\tilde{\mu}_{i1}^{pk}$ with four bins, one for each of the four pricing plans. Since we assume that $\tilde{\mu}_{i1}^{pk}$ is normally distributed with mean $\tilde{\mu}_0$ and standard deviation $\tilde{\sigma}_\mu$, this histogram would then (over) identify the distribution. The resulting histogram and fitted normal distribution, are both shown in the lower panel of Figure 6 for the case $\tilde{\sigma}_{\theta 1} = 80$, $\beta = 4$, and $\varphi = 1/2$.

The model identifies $\tilde{\sigma}_{\theta 1}$ as the value between 0 and 118 that generates the best fit between the histogram and the fitted normal distribution. Choosing a larger value for $\tilde{\sigma}_{\theta 1}$ implies a higher mean and a lower variance for the distribution of $\tilde{\mu}_{i1}^{pk}$.[36] Given $\beta = 4$, the overall best fit is at $\tilde{\sigma}_{\theta 1} = 83.5$.

The preceding argument for identifying $\tilde{\sigma}_{\theta 1}$, $\tilde{\mu}_0^{pk}$, and $\tilde{\sigma}_{\mu^{pk}}$ clearly bounds $\tilde{\sigma}_{\theta 1} \leq 118$ (given $\beta = 4$) but then relies heavily on the functional form assumption that $\tilde{\mu}_{i1}^{pk}$ is normally distributed for point identification. Nevertheless, there is additional information in the data which reduces reliance on the functional form assumption: As prices change over time, the bounds depicted in Figure 6 change, so that plan share data from later dates provide additional restrictions on $\tilde{\sigma}_{\theta 1}$ and the distribution of $\tilde{\mu}_{i1}^{pk}$.

The exercise described above identifies consumer uncertainty about initial tastes ($\tilde{\sigma}_{\theta 1}$) but it still remains to separate out uncertainty about own type ($\tilde{\sigma}_1$) from perceived taste volatility ($\tilde{\sigma}_\varepsilon^{pk}$), which in turn will distinguish overconfidence (δ_μ) from volatility bias (δ_ε). By equation (5), $\tilde{\sigma}_{\theta 1}^2$ is a weighted sum of $\tilde{\sigma}_1^2$ and $(\tilde{\sigma}_\varepsilon^{pk})^2$. The two parameters are distinguished by the rate of learning

---

[36]This is because higher uncertainty (higher $\tilde{\sigma}_{\theta 1}$) leads individuals who choose plans 1-3 to insure themselves by choosing plans with more included minutes. They are willing to choose plan 2 over plan 1 and plan 3 over plan 2 at lower values of $\tilde{\mu}_{i1}^{pk}$. However, they are only willing to choose plan 1 over plan 0 at higher values of $\tilde{\mu}_{i1}^{pk}$.

and plan switching, which is decreasing in $\tilde{\sigma}_\varepsilon^{pk}/\tilde{\sigma}_1$. This is apparent from the expression for $\tilde{\mu}_{i,t+1}^{pk}$ derived from Bayes rule in Appendix B.3 equation (16), which shows that a consumer's updated beliefs are a weighted average of her prior and her signals, where the weight placed on her prior is proportional to $(\tilde{\sigma}_\varepsilon^{pk}/\tilde{\sigma}_1)^2$. Recall that we identify the probability of an active choice based on the rate at which consumers switch away from dominated plans. Thus we can distinguish slow learning from a failure to actively consider switching.

5.3.1 Logit Error Weight

The preceding discussion ignores logit-errors, which the model does incorporate into plan choice. As a result, plan choices do not actually give sharp bounds on prior beliefs, but rather smooth likelihoods over priors, since beliefs outside the bounds described by Figure 6 can be explained by the logit error. Without logit-errors, all initial plan choices could be rationalized by prior beliefs. However, the model requires logit-errors to rationalize switches that appear to be in the 'wrong' direction. For example, suppose a customer with high average usage chooses a small plan and subsequently experiences a string of overage charges. A low prior belief ($\tilde{\mu}_{i1}^{pk}$ small) could rationalize the initial choice of a small plan. However, given the assumption of Bayesian learning, no prior can simultaneously rationalize the initial choice and a subsequent switch to an even smaller plan. The degree to which switching is in the wrong direction identifies the logit error weight $1/\alpha$.

5.4 Tastes

Having identified beliefs it is straightforward to identify taste process parameters. Given the AR1 coefficient φ, the price coefficient β, and consumer beliefs, we can calculate v_{it}^k for $k \in \{$pk-in,pk-out,op-in,op-out$\}$ and infer taste-shocks $\boldsymbol{\theta}_{it}$ and $\tilde{\mathbf{r}}_{it}^{9pm}$ from usage. Observing r_{it}^k for $k \in \{$9pk,9op$\}$ (a censoring of $\tilde{r}_{it}^k = \alpha_i^k + e_{it}^{r,k}$) identifies $E\left[\alpha_i^k\right]$, $Var(\alpha_i^k)$, and $Var(e_{it}^{r,k})$.[37] Correlation between observed usage and initial plan choices identifies $\boldsymbol{\psi}$, which determines the correlation between beliefs and true types. Given φ and $\boldsymbol{\theta}_{it}$, we can calculate the composite error $(\boldsymbol{\mu}_i + \boldsymbol{\varepsilon}_{it}) = \boldsymbol{\theta}_{it} - \varphi\boldsymbol{\theta}_{i,t-1}$, which is joint-normally distributed conditional on $\tilde{\mu}_{i1}^{pk}$, so unconditionally is the mixture of joint normals. The argument for identifying this distribution is then similar to that for identifying the error structure in a random effects distribution. This delivers the parameters $\boldsymbol{\mu}_0$, $\boldsymbol{\Sigma}_\mu$, and $\boldsymbol{\Sigma}_\varepsilon$. Finally, bias measures δ_μ, δ_ε, b_1, and b_2 can be computed from their definitions.

---

[37]Without censoring, these would simply be $E\left[\alpha_i^k\right] = E\left[r_{it}^k\right]$, $Var\left(\alpha_i^k\right) = Cov(r_{it}^k, r_{it-1}^k)$, and $Var(e_{it}^{r,k}) = Var(r_{it}^k) - Var(\alpha_i^k)$.

6   Estimation Procedure

Before describing our estimation procedure, we outline the parameters to be estimated. First are those associated with beliefs: the parameters governing the distribution of consumer beliefs, $\tilde{\mu}_0^{pk}$ and $\tilde{\sigma}_{\mu^{pk}}$, consumers' initial uncertainty about their peak type, $\tilde{\sigma}_1$, and consumers' estimate of taste volatility, $\tilde{\sigma}_\varepsilon$. The parameters associated with actual tastes for usage are the means of the μ_{it}^k's, μ_0^{pk} and μ_0^{op}, their variances and correlation, $\sigma_{\mu^{pk}}^2$, $\sigma_{\mu^{op}}^2$, and ρ_μ, and the variances and correlation of the idiosyncratic errors, $(\sigma_\varepsilon^{pk})^2$, $(\sigma_\varepsilon^{op})^2$, and ρ_ε, as well as ψ^{pk} and ψ^{op}, which capture correlation between beliefs and actual usage. There are four parameters which govern the shares of outgoing landline calls occurring between 8:00 pm and 10:00 pm: the average peak share μ_α^{9pk}, the individual specific variance $(\sigma_\alpha^{9pk})^2$, and the two idiosyncratic variances $(\sigma_e^{9k})^2$ for $k \in \{pk, op\}$.[38] The final set of parameters that are discussed in the text include the price coefficient β, the logit error weight $1/\alpha$, the active choice probability P_C, and the outside good utility O. Finally, we estimate an additional six parameters that govern the share of in-network usage and a parameter that reflects consumer beliefs about the share of in-network usage. We discuss these parameters further in Appendix C. We denote the vector of all parameters as Θ, which is 30 dimensional.

We begin this section by describing the structure of the likelihood function which arises from our model. As discussed below, the likelihood function for our model does not have a closed form expression due to the presence of unobserved heterogeneity. We therefore turn to Simulated Maximum Likelihood to approximate the likelihood function (Gourieroux and Monfort 1993).

An observation in our model is a usage plan-choice pair for a consumer at a given date. At each observation, we must evaluate the joint likelihood of observed usage and plan choice conditional on observed prices and the consumer's usage and choice history. The likelihood for an observation arises naturally from the distributional assumptions on our model's unobservables. To facilitate the exposition, we divide the unobservables into two groups. The first group consists of random variables that are independent across individuals, but are not independent across time within an individual. This consists of the unobservables $\tilde{\mu}_{i1}^{pk}$, μ_i^{pk}, μ_i^{op}, α_i^{9pk}, two normally distributed individual specific effects which govern the share of in-network usage for peak and off peak, α_i^{pk} and α_i^{op}, and latent $\tilde{\theta}_{it}^k$ when $\theta_{it}^k = 0$ for $k \in \{pk, op\}$. (When category $k \in \{pk, op\}$ usage is zero, we can infer that the censored taste shock θ_{it}^k is zero but the latent taste shock $\tilde{\theta}_{it}^k \leq 0$ is unobserved.) We group these random variables together into a vector denoted $\boldsymbol{u}_i$. The second group of error terms consist of structural shocks that are independent across time and individuals: the logit plan

---

[38]Recall that we do not need to estimate a mean or individual specific variance for off peak 9:00 pm to 10:00 pm usage because we restrict average peak and off-peak tastes for 8:00 pm to 10:00 pm usage to be equal in equation (7).

choice error η_{itj}, the errors in the stochastic process of $\tilde{\theta}_{it}^k$ when $\tilde{\theta}_{it}^k > 0$, ε_{it}^k, idiosyncratic errors for the 8:00 pm to 10:00 pm shares, r_{it}^{9k} and e_{it}^k for $k \in \{9pk, 9op\}$, as well as two normally distributed idiosyncratic errors governing in-network usage, which we denote e_{it}^{pk} and e_{it}^{op}, respectively.

For individual i at time period t, we observe a plan choice j as well as a vector of usage, $\boldsymbol{q}_{it}$, where $\boldsymbol{q}_{it} = \{q_{it}^{pk,in}, q_{it}^{pk,out}, q_{it}^{op,in}, q_{it}^{op,out}, q_{it}^{9pk}, q_{it}^{9op}\}$, and the in and out superscripts refer to in-network and out-of-network usage. Conditional on $\mathbf{u}_i$, the likelihood of an observation will simply be the product of the choice probability and the likelihood of the observed usage.

First, consider the choice probability. Conditional on information set $\Im_{it}$ and an active choice in period t,[39] an individual will choose plan j when that plan has the highest utility according to equation (3). Let J_{it} denote the set of plans available to consumer i in period t. Conditional on an active choice, our assumption of logit errors gives rise to the following choice probability:

$$P_{it}(j'|C; \Im_{it}, \boldsymbol{u}_i, J_{it}) = \frac{\exp(U_{ijt}(\Im_{it}, \boldsymbol{u}_i))}{\sum_{k \in J_{it}} \exp(U_{ikt}(\Im_{it}, \boldsymbol{u}_i))}.$$

Unconditional on an active choice, the probabilities that an existing customer switches to plan j' in period t (where j' could be the outside good) or keeps the existing plan j are $P_C P_{it}(j'|C; \Im_{it}, \boldsymbol{u}_i, J_{it})$ and $P_C P_{it}(j|C; \Im_{it}, \boldsymbol{u}_i, J_{it}) + (1 - P_C)$ respectively:

$$P(\text{Choose } j'|\Im_{it}, \boldsymbol{u}_i, J_{it}) = \begin{cases} P_C P_{it}(j'|C; \Im_{it}, \boldsymbol{u}_i, J_{it}) & \text{if } j' \neq j \\ P_C P_{it}(j'|C; \Im_{it}, \boldsymbol{u}_i, J_{it}) + (1 - P_C) & \text{if } j' = j \end{cases}. \tag{9}$$

The consumer's information set in period t will contain some of the random draws, as well as past $\boldsymbol{q}_{it}$'s which impact the Bayesian updating process. The consumer's choice set J_{it} depends on the plan choices drawn from the non-university plans and the consumer's past plan choices. For a new customer, the initial choice set J_{i1} includes plans currently offered through the university but does not include the outside option or any other plans, and does not vary with the simulation draw. Other options are not included for new customers because we only observe consumers who sign up; hence the probability of plan choice for these customers is the probability of choosing plan j conditional on signing up. For existing customers, the choice set J_{it} also includes the customer's existing plan, those currently offered by the other provider considered, and the outside good. We assume that the consumer considers only one outside provider (AT&T, Cingular, or Verizon), in addition to the possibility of quitting each month. The option considered is drawn from a discrete distribution which assigns probability 1/3 to each of the three providers.

---

[39]Notation: conditioning on C means conditioning on an active choice.

As there are three possible choice sets, we index each choice set by J_{it}^k, $k = 1, ..., 3$.

Next, consider the likelihood of observed usage. A consumer's observed usage, $\boldsymbol{q}_{it}$, will be a function of $\boldsymbol{u}_i$, the idiosyncratic errors ε_{it}^k and e_{it}^k, and past values of $\boldsymbol{q}_{it}$ for $t > 1$. Conditional on $\boldsymbol{u}_i$ and $\boldsymbol{q}_{i1}, ..., \boldsymbol{q}_{i,t-1}$, the distributions of ε_{it}^k and e_{it}^k will generate a distribution for $\boldsymbol{q}_{it}$. We denote this density function as $f_{\boldsymbol{q}}(\boldsymbol{q}_{it}|\boldsymbol{u}_i, \boldsymbol{q}_{i1}, ..., \boldsymbol{q}_{i,t-1}, \Theta)$. Given assumed distributions of ε_{it}^k and e_{it}^k, we derive the distribution of $\boldsymbol{q}_{it}$ using a change of variables. We describe the exact form of $f_{\boldsymbol{q}}$ in Appendix D.

The likelihood of a sequence of observed usages and plan choices will be the product of the individual usage and choice likelihoods with the unobservable $\boldsymbol{u}_i$ integrated out:

$$L_i(\Theta) = \int_{\boldsymbol{u}_i} \prod_{t=1}^{T_i} \left[\left(\sum_{k=1}^{3} \frac{1}{3} P(\text{Choose } j'|\Im_{it}, \boldsymbol{u}_i, J_{it}^k) \right) f_{\boldsymbol{q}}(\boldsymbol{q}_{it}|\boldsymbol{u}_i, \boldsymbol{q}_{i1}, ..., \boldsymbol{q}_{i,t-1}) \right] f_{\boldsymbol{u}}(\boldsymbol{u}_i) d\boldsymbol{u}_i. \quad (10)$$

Because it has no closed form solution, we approximate the integral over $\boldsymbol{u}_i$ in equation (10) using Monte Carlo Simulation. For each individual, we take S draws on the random effects from $f_{\boldsymbol{u}}(\boldsymbol{u}_i)$ and the choice sets J_{it}^k, and approximate the likelihood using

$$\hat{L}_i(\Theta) = \frac{1}{S} \sum_{s=1}^{S} \left[\left(P(\text{Choose } j'|\Im_{it,s}, \boldsymbol{u}_{is}, J_{it,s}^k) \right) f_{\boldsymbol{q}}(\boldsymbol{q}_{it}|\boldsymbol{u}_{is}, \boldsymbol{q}_{i1}, ..., \boldsymbol{q}_{i,t-1}) \right].$$

The model log-likelihood is the sum of the logarithms of the individual likelihoods:

$$\hat{LL}(\Theta) = \sum_{i=1}^{I} \log(\hat{L}_i(\Theta)). \quad (11)$$

It is well-known that the value of Θ which maximizes $\hat{LL}$ is inconsistent for fixed S due to the logarithmic transformation in equation (11). However, it is consistent if $S \to \infty$ as $I \to \infty$, as discussed in Hajivassiliou and Ruud (1994). We chose $S = 300$; to arrive at this value we conducted some simple artificial data experiments where we simulated our model and attempted to recover the parameters, finding that 300 draws was sufficient to recover the true parameter draws to roughly 5% accuracy. We also found in our experiments that we were able to reduce simulation bias significantly by using a deterministic Sobol sequence generator to create the random draws, rather than canonical random number generators. Goettler and Shachar (2001) describe some of the advantages of this technique in detail. We use the algorithm provided in the R package randtoolbox to create the draws (Dutang and Savicky 2010).

A second issue that arises in the formulation of equation (11) is due to the censoring of serially correlated taste shocks. As noted above, whenever peak or off-peak usage is zero, the corresponding

latent taste shock $\tilde{\theta}_{it}^k$ is negative and unobserved. In such cases we substitute the probability that $\tilde{\theta}_{it}^k$ is censored for f_q. However, we always need a value of $\tilde{\boldsymbol{\theta}}_{it}$ to calculate period $(t+1)$'s likelihood because we assume both that $\tilde{\boldsymbol{\theta}}_{it}$ follows an AR1 process and that consumers observe $\tilde{\boldsymbol{\theta}}_{it}$ when updating their $(t+1)$ beliefs. Hence, when censored, we also draw a value of $\tilde{\theta}_{it}^k$ using an importance-sampling procedure and include it in $\boldsymbol{u}_i$ to be integrated out. Our approach, which results in a smooth likelihood, is an adaptation of Lee's (1999) procedure for integrating out serially correlated latent unobservables in dynamic Tobit models.

Additional details about the likelihood function, including a treatment of in-network calling, are in Appendix D. We describe the computational procedures we use to evaluate and maximize the likelihood function in Section E. The computational difficulties in the estimation arise primarily from two sources: one is the high dimensional unobserved heterogeneity, which requires many evaluations of the likelihood function. The second is the computation of $\mathbf{v}^*$. Because there is no closed-form solution for $\mathbf{v}^*$, we use a nonlinear equation solver to solve for it numerically. We must do this for each simulation draw, at each time period, for every individual, at every choice that is not the outside good or the two-part tariff, plan 0.

7 Results

7.1 Parameter Estimates

Our parameter estimates are shown in Table 3. The first three columns show the coefficients, estimates, and standard errors for the first 15 parameters, while the fourth through sixth columns show the same for the next 15 parameters. The calling price coefficient β is 4.02, which indicates that a price increase from 0 cents to 11 cents per minute decreases usage by 31%.

The next 10 parameters characterize the distribution of $\tilde{\mu}_{i1}^{pk}$ as well as the perceived and true distributions of μ_i^{pk} and μ_i^{op} conditional on $\tilde{\mu}_{i1}^{pk}$. On average, consumers believe their mean draw of $\tilde{\theta}_{it}^{pk}$ to be negative 19, while the actual mean of $\tilde{\theta}_{it}^{pk}$ is 103 minutes. (Accounting for censoring of the latent shock, the average consumer believes the mean of θ_{it}^{pk} is 72 minutes and the true mean is 175 minutes.) The average off-peak draw $\tilde{\theta}_{it}^{op}$ is slightly below the peak value at 97 minutes. (The model predicts higher off-peak usage due to consumer price sensitivity.)

The standard deviation in consumers' initial belief $\tilde{\mu}_{i1}^{pk}$ is 143 minutes. Conditional on $\tilde{\mu}_{i1}^{pk}$, the standard deviation of consumer uncertainty about true type, $\tilde{\sigma}_1$, is 12.9 minutes. In contrast, conditional on $\tilde{\mu}_{i1}^{pk}$, the population standard deviations of μ_i^{pk} and μ_i^{op} are 78 and 162 minutes respectively. Thus consumers are overconfident, underestimating uncertainty about true type by 84%. The estimates of $\boldsymbol{\psi}$ indicate that initial beliefs are slightly negatively correlated with μ_i^{pk}

Table 3: Parameter Estimates

Coefficient	Estimate	Std. Err	Coefficient	Estimate	Std. Err
β	4.024	(0.02)	μ_α^{9pk}	-0.004	(0.001)
$\tilde{\mu}_0^{pk}$	-19.275	(1.33)	$(\sigma_\alpha^{9pk})^2$	0.06	(0.001)
μ_0^{pk}	103.327	(1.896)	$(\sigma_e^{9pk})^2$	0.104	(0.001)
μ_0^{op}	97.512	(3.418)	$(\sigma_e^{9op})^2$	0.116	(0.001)
$\tilde{\sigma}_{\mu^{pk}}$	143.196	(0.646)	φ	0.579	(0.008)
$\tilde{\sigma}_1$	12.876	(0.046)	α	0.096	(0.163)
$\sigma_{\mu^{pk}}$	78.584	(1.462)	Price Consideration	0.063	(0.046)
$\sigma_{\mu^{op}}$	161.784	(2.235)	Outside Good Utility	-71.191	(30.312)
ψ^{pk}	-0.043	(0.012)	δ_r	0.003	(0.037)
ψ^{op}	0.228	(0.017)	μ_α^{pk}	0.35	(0.002)
ρ_μ	0.981	(0.002)	μ_α^{op}	0.401	(0.002)
$\tilde{\sigma}_\varepsilon^{pk}$	163.574	(1.746)	$(\sigma_\alpha^{pk})^2$	0.035	(0.001)
σ_ε^{pk}	182.521	(0.452)	$(\sigma_\alpha^{op})^2$	0.039	(0.001)
ρ_ε	0.407	(0.004)	$(\sigma_e^{pk})^2$	0.03	(0)
σ_ε^{op}	306.574	(0.66)	$(\sigma_e^{op})^2$	0.025	(0)
Log-likelihood	264684.7				

and are positively correlated with μ_i^{op}. Finally, conditional on $\tilde{\mu}_{i1}^{pk}$ the correlation between peak and off-peak μ_i^k is high at 98%. The unconditional correlation between peak and off-peak μ_i^k is somewhat lower at 94%; the unconditional standard deviations of the peak and off-peak μ_i^k are slightly higher than their conditional values at 79 minutes and 165 minutes respectively.

The last four rows of column 1 describe the perceived and true distributions of the error term ε. The perceived standard deviation of ε_{it}^{pk}, $\tilde{\sigma}_\varepsilon^{pk}$, is 164 minutes. In contrast, the true standard deviation is 183 minutes, meaning consumers underestimate volatility by 10%. The variances of peak and off peak errors are higher than the unconditional variances of μ_i^{pk} and μ_i^{op}, indicating that more of the variation in usage can be attributed to monthly volatility than the consumer-level fixed effect; additionally, their correlation is much lower.

The first four parameters of column 2 describe consumers' tastes for 8:00 pm to 10:00 pm usage. The low value of μ_α^{9pk} indicates that outgoing 8:00 to 9:00 pm landline usage is small as a fraction of total peak usage, which is consistent with the data. The φ value of 0.58 indicates strong serial correlation in tastes from month to month. The logit error scaling parameter, α, is estimated to be a little less than 0.1. The price consideration parameter is 0.063, indicating that consumers seldom look at prices, but it is not precisely estimated. The imprecision is consistent with our artificial data experiments, where we found that this parameter was difficult to identify. The outside good utility

is estimated to be -71. Compared to average utilities of about -20, this implies that consumers prefer inside goods to the outside good by a large margin.

The last seven parameters relate to in-network usage. We describe the modifications to the model needed to distinguish in and out-of-network usage in Appendix C. Loosely, the parameter δ_r measures consumers' underestimation of the fraction of calls that are in-network. Since our estimate of δ_r is close to zero, consumers believe that almost all usage is out of network.[40] The next two parameters govern the shares of $\boldsymbol{\theta}_{it}$ which can be apportioned to peak and off-peak in-network usage, respectively, while the final four govern the variances of in-network usage.

7.2 Biases and Learning

Returning to consumer beliefs, the parameters which summarize consumer biases are functions of our estimated parameters. We display estimates of these parameters in Table 4. Our estimates of δ_μ and δ_ε indicate strong overconfidence and mild volatility bias, respectively. Consumers underestimate their uncertainty about their own average tastes by 84% and underestimate the monthly volatility in their tastes by 10%. Together, overconfidence and volatility bias imply that the standard deviation of consumers' initial uncertainty about $\tilde{\theta}_{i1}^{pk}$, $\tilde{\sigma}_{\theta 1}$, is 203 minutes rather than the correctly calibrated 292 minutes. (Note that if consumers are risk averse rather than risk neutral then these estimates are lower bounds on the magnitudes of overconfidence and volatility bias.) Aggregate mean bias is negative, indicating that the average consumer underestimates her initial $\tilde{\theta}_{it}^{pk}$ draw by 123 minutes. Finally, the positive estimate of b_2 reflects strong positive conditional mean bias.[41]

Table 4: Estimates of Consumer Beliefs

Coefficient	Estimate	Std. Err
δ_μ	0.164	(0.003)
δ_ε	0.896	(0.01)
b_1	-122.602	(2.19)
b_2	1.043	(0.012)

The fact that overconfidence is stronger than volatility bias ($\delta_\mu < \delta_\varepsilon$) implies that consumers

---

[40]Plan 0 always offered free in-network usage and plan 2 did so as well near the end of our sample period. We incorporated this parameter to help explain the high share of Plan 1 relative to Plan 0, as plan 0 dominates plan 1 for anyone with a median in-network usage share.

[41]In the context of grocery home delivery service, Goettler and Clay (2011) also find $b_1 < 0$ and $b_2 > 0$.

overweight their priors relative to new experience and hence learn slowly. This is illustrated in Figure 7, which plots an average of consumers' evolving point-estimates $\tilde{\mu}_{it}^{pk}$ for consumers whose true value is $\mu_i^{pk} = \mu_0^{pk} \approx 103$. A consumer's time t point-estimate $\tilde{\mu}_{it}^{pk}$ is a function of her initial belief $\tilde{\mu}_{i1}^{pk}$ and her past taste shocks $\boldsymbol{\theta}_i^{t-1}$. The dotted lines in the figure show the average of $\tilde{\mu}_{it}^{pk}$ for 1000 simulated consumers, where each consumer's $\tilde{\mu}_{i1}^{pk}$ and $\boldsymbol{\theta}_i^T$ are drawn from their estimated distributions. The thick red line shows how consumers' beliefs evolve given estimated overconfidence and volatility bias. An average consumer whose true μ_i^{pk} is roughly 103 minutes and who enters the sample believing $\tilde{\mu}_{i1} = \tilde{\mu}_0 = -19$ increases her belief to $\tilde{\mu}_{i,13} = -9.6$ after one year (an 8% reduction in aggregate mean bias). The blue dashed line shows how beliefs evolve when overconfidence and volatility bias are removed. Debiasing consumers speeds up learning and after 1 year an average consumer's belief about μ_i^{pk} will be $\tilde{\mu}_{i,13} = 66$ (a 70% reduction in aggregate mean bias).
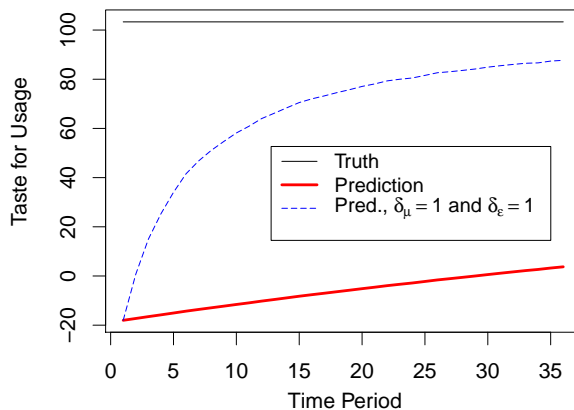

Figure 7: Population average $\tilde{\mu}_{it}^{pk}$ (point estimate of μ_i^{pk}) for those with true value $\mu_i^{pk} = \mu_0^{pk} \approx 103$.

7.3   Fixed-Price Counterfactual: Impact of Biased Beliefs

Before proceeding to simulate endogenous price changes in Section 8, we briefly simulate the change in firm profits, consumer welfare, and total welfare that results from debiasing consumers while holding observed prices fixed (Table 5). We construct these counterfactual simulations at our data in the sense that we hold fixed the number of consumers, and when consumers enter and exit the data set. Surplus changes are measured in dollars per student over the two year period that they are observed and assume marginal cost is zero. The first three columns of Table 5 show the welfare effects when students face university prices, while the last three columns show the welfare effects when consumers face publicly available prices. The consequences of debiasing are larger in the latter case because the university's plan 0, which is absent from the public menu, tended to protect biased consumers.

Table 5: Counterfactual: Per student change in surpluses from bias elimination (fixed prices)

Beliefs	University Plans			Public Plans		
	Profits	Cons. Welf.	Total Welf.	Profits	Cons. Welf.	Total Welf.
$\delta_\mu = 1$	−23.09	16.05	−7.04	−48.91	33.09	−15.82
$\delta_\varepsilon = 1$	−6.59	4.86	−1.73	−12.41	9.07	−3.34
$\delta_\mu = 1$ and $\delta_\varepsilon = 1$	−23.56	15.21	−8.35	−51.88	34.44	−17.44
No Biases	−44.01	36.95	−7.07	−58.54	42.21	−16.34

Changes in surpluses (profits, consumer welfare, and total welfare) are measured in dollars per student over the 2 year sample period. Changes are relative to surpluses at estimates.

The first row of Table 5 shows the impact of removing overconfidence, which raises consumer surplus but lowers profits and total welfare. Row two shows similar effects of removing volatility bias and row three shows the combined effects of removing both biases, which raises consumer surplus by \$34 given public prices. Finally, row four shows the total effect of removing all biases, including mean biases and underestimation of in-network calling (discussed in Appendix C), which raises consumer surplus by \$42 given public prices. On average, debiased consumers are less likely to choose plan 1 and make fewer calls because they are more aware of overage risk. The reduction in plan 1 share is more pronounced for public prices because its initial share is higher without plan 0. The reduction in calling reduces total welfare because marginal costs are approximately zero. Thus gains in consumer surplus are overshadowed by profit losses.

8 Endogenous-Price Counterfactual: Bill-Shock Regulation

8.1 Nested-Logit Specification

To predict the effect of bill-shock regulation on equilibrium prices it is important to correctly capture the degree of competition between carriers. Hence, we modify the error structure of the demand model to be a two level nested logit, rather than logit. In our nested-logit specification, we assume that each inside nest contains the plans offered by a carrier (the option of shutting off cellular-phone service is also put in its own nest). The outside nest consists of all the carriers (including no service) in a particular consumer's consideration set.[42] We assume that the inclusive value parameter, denote by λ, is the same for each option.

The more restrictive logit specification implies that if consumers choose plans within carrier

---

[42] A new consumer chooses among three carriers and no-serivce whereas an existing consumer who considers switching chooses among her current carrier, a randomly chosen outside carrier, and no service.

primarily based on price then carriers are close substitutes. Thus the logit specification leads to unrealistically high competition and low prices in counter-factual simulations. We choose the more flexible nested-logit specification because it allows consumers to have strong idiosyncratic carrier preferences (due to network coverage or phone availability) that create market power, while at the same time making within carrier plan choices primarily based on price.

Ideally, we would like to estimate λ jointly with the other demand parameters using demand side choice data. Unfortunately, we observe neither carrier market shares on campus nor the alternate carriers chosen by students quitting university plans. Hence only the quitting rate is available to identify utility of the outside good, average utility of university plans relative to other carriers, and λ. In our demand estimates we assume $\lambda = 1$ (logit specification) and carrier symmetry to identify the outside good utility.[43] To address this identification problem, we calibrate λ using supply-side price data: We select the value of λ that best rationalizes observed prices conditional on our demand estimates. Our algorithm, which is described in Appendix F, calibrates λ to be 0.2.

Before proceeding, we make two comments on our calibration approach. First, one potential problem is that our demand estimates were made conditional on $\lambda = 1$ (which generates the logit model), but different values of λ might produce different demand estimates. Fortunately, our demand estimates are relatively insensitive to λ, which we show in Appendix F. Second, in principle we could have estimated λ and the other parameters jointly by using constrained maximum likelihood and constraining observed prices to be optimal at the estimated parameters. We avoided this approach because we prefer only to impose our supply-side structural assumptions (that competition is symmetric static Nash in prices and that our student population is representative[44]) only when they are necessary in the endogenous-price counterfactual simulations.

8.2 Modeling bill-shock alerts and endogenous prices.

In our bill-shock regulation counterfactual, consumers are informed when their usage reaches Q, their allotment of free minutes.[45] In response to this new policy, a consumer's usage rule changes: A consumer will accept all calls valued above v^* until she exhausts her included minutes. After that

---

[43]Outside price variation is too limited to separately identify λ. Moroever, in an unreported specification, we instead chose a natural normalization for the outside good and estimated λ. We rejected this alternative, however, because the resulting estimate of λ was zero, an implausible number that implies carriers have monopoly power. This may have been due to the fact that our normalization of the outside good was too low, that carrier symmetry is a bad assumption when including the university plans, or the fact that forced quits due to graduations (outside the model) biased the estimate downwards.

[44]In reality, university plans are not symmetric to other carrier offerings and our population of students is likely overweighted towards new and low-volume users relative to the overall population.

[45]Alerts are not applicable to two-part tariffs with constant marginal prices.

point, she only accepts calls valued above p. Because the consumer adjusts her calling threshold upon making Q calls, the optimal initial threshold v^* differs from that characterized by equation (2). Appendix B.2 describes expected utility and characterizes v^* under bill-shock regulation. To calculate endogenous equilibrium prices, we assume that there are three symmetric carriers, equilibrium is symmetric static Nash in prices, marginal costs are zero, overage rates are at most fifty cents,[46] and each carrier offers a menu of three plans.

8.3 Counterfactual Simulation Results

Table 6 shows the results of our endogenous-price counterfactual simulations. Column 1 shows predicted plan prices and welfare outcomes under our estimated demand parameters. The model predicts that firms offer a two-part tariff at fifty cents per minute for $26.70 per month, a three-part tariff with 349 included minutes for $60.17 per month, and an unlimited plan for $75.58 per month.

Table 6: The Impact of Bill Shock Regulation and Removing Biases on Equilibrium Prices

		Est	Est, Bill Shock (fixed prices)	Est, Bill Shock	$\delta_\mu = 1$ and $\delta_\varepsilon = 1$	No Biases
		(1)	(2)	(3)	(4)	(5)
Plan 1	M	26.70	26.70	26.42	26.15	74.43
	Q	0	0	0	0	∞
	p	50	50	50	50	N/A
	Share	54	53	56	48	33
Plan 2	M	60.17	60.17	61.73	56.27	74.43
	Q	349	349	211	0	∞
	p	50	50	12	8	N/A
	Share	28	29	26	27	33
Plan 3	M	75.58	75.58	76.52	77.15	74.43
	Q	∞	∞	∞	∞	∞
	p	N/A	N/A	N/A	N/A	N/A
	Share	18	17	19	25	33
Outside Good Share		0	0	0	0	0
Δ Annual Profit			-16.47	-0.2	8.18	32.44
Δ Annual Cons Welfare			8.34	-2.01	20.28	168.23
Δ Annual Total Welfare			-8.13	-2.21	28.47	200.67

All welfare and profit numbers are expressed in thousands of dollars. Because the counterfactuals in columns 4 and 5 produced two part tariffs, under bill shock regulation equilibrium prices are unchanged. We simulate 1000 consumers for 12 months.

---

[46]Otherwise the combination of biased beliefs and inattention lead to implausibly high overage-rate predictions.

Column 2 of Table 6 holds constant the predicted prices from column 1 but imposes bill-shock regulation. Holding prices constant, bill-shock alerts help plan 2 customers avoid an average of $54.97 in overage payments annually (reducing profits by the same amount). Avoided overage charges correspond to reduced calling, so average consumer surplus rises by only $27.00 per year for an average plan 2 customer. Because marginal costs are zero, the difference, $27.97, reflects the annual decrease in total welfare per plan 2 customer due to reduced calling.[47] (Column 2 of Table 6 reports these figures averaged across all customers.)

Column 3 of Table 6 imposes bill-shock regulation but allows firms to adjust prices. In equilibrium, markups are determined primarily by the calibrated inclusive value parameter, which at $\lambda = 0.2$ implies markups of about $74 per month. As a result, following bill-shock regulation, firms adjust the price of plan 2 to compensate for lost overage revenue and maintain a stable markup. Hence annual profits are stable (falling by only $0.20 per consumer per year) and consumers lose (an average of $2.01 per consumer per year) because they essentially become residual claimants of total welfare.[48]

As before, plan 2 customers are most affected by bill-shock regulation. Now, however, there are two distinct groups to consider – those who continue to choose plan 2 and those who switch to plan 1 in response to price changes. Those who continue to choose plan 2 after bill-shock regulation is implemented only benefit by an average of $15.53 per year due to changes in plan 2 pricing, including the $1.56 monthly fee increase and the 138 minute included allowance reduction. However, their contribution to total welfare actually increases because they make more calls in response to the lower overage rate of 12 cents per minute.

The preceding good news, that bill-shock alerts raise consumer surplus and improve efficiency for those who continue to choose plan 2, is only part of the story. Although average consumer surplus rises by $15.53 on plan 2, consumers do not appreciate this and actually perceive a $31.76 drop in annual surplus. This is because much of the benefit comes from the ability to respond to bill-shock alerts and take advantage of a low overage rate in months with high usage. Unfortunately, biased consumers underestimate the incidence of such high usage and therefore underestimate both the value of bill-shock alerts the value of a lower overage rate. In contrast, even biased consumers fully appreciate the cost of a $1.56 monthly fee increase. Thus, on the margin, more consumers

---

[47]Holding average total calling constant, bill-shock alerts reduce welfare by inducing consumers to call more in low demand months (by choosing a lower v^*) and to call less in high demand months (when receiving an alert). This reduces the average value of placed calls. In addition, average total calling is reduced because biased consumers choose too low a calling threshold v^*, but correct their behavior following a bill-shock alert.

[48]An additional counterfactual simulation shows that a single firm which introduced bill-shock alerts on its own would lose $3.69 per customer annually.

choose plan 1 over plan 2. Moreover, this is a bad decision for those marginal consumers because they talk more than they anticipate. As a result, average consumer surplus falls by $309.60 annually for those who switch to plan 1. Moreover, their contribution to total welfare falls as well because they substantially reduce calling in response to the 50 cent per minute rate.

In sum, we find that bill-shock alerts are neutral or beneficial for most consumers but severely harm a minority. Averaging across all consumers yields the negative results in column 3: average annual losses of $2.21 to social welfare and $2.01 to consumer surplus due to bill-shock regulation.

Turning to columns 4 and 5 of Table 6, we investigate the consequences of debiasing consumers. Column 4 shows the effect of eliminating overconfidence and projection bias. In column 4, plans 1 and 3 are similar to those in column 1 but plan 2 becomes a two-part tariff charging $56 per month and 8 cents per minute. Moreover, consumers tend to choose larger contracts and total welfare increases by $28 per consumer annually. Column 5 shows the effect of eliminating all biases. In this case, firms offer three identical plans that each charge $74 for unlimited calling and total welfare increases by $201 per consumer annually. In both columns 4 and 5, debiasing increases firm profits but most of the increase in total surplus accrues to consumers.

To understand these pricing results, we consider scenarios in reverse order, beginning with column 5. The unlimited calling plans in column 5 achieve first-best surplus via marginal cost pricing and earn a markup of $74 from each consumer. (Three identical such contracts do better than a single contract due to the red-bus/blue-bus problem.) When consumers are unbiased this is optimal because: (1) we assume that the inclusive value parameter is the same for all consumers, and (2) we estimate a low outside good value that predicts full market coverage. Thus, firms have no incentive to price discriminate and charge unbiased consumers different markups.

In column 4, positive conditional mean bias implies that consumers choose overly extreme plans. Thus consumers who choose plan 3 overestimate their usage while those who choose plan 1 underestimate their usage. (Negative aggregate mean bias means plan 1 consumers outnumber plan 3 consumers.) Plan 3 is optimal for over-estimators because it sells calls to consumers up front and then offers no refunds. Plan 1 is optimal for under-estimators because it extracts value ex post when consumers realize their true value for calls. Plan 2 caters to those consumers in between whose under or over estimation is mild.

Returning to columns 1-3, plan 1 and plan 3 pricing is driven by conditional mean bias following the same logic as in column 4. However, in columns 1-3 overconfidence and volatility bias predominate for intermediate plan 2 consumers and hence a three-part tariff is optimal (Grubb 2009).

Because three-part tariff pricing is driven by overconfidence and volatility bias, eliminating these biases also eliminates three-part tariff pricing. Thus bill-shock regulation has no effect without bias.

Moreover, the fact that bill-shock regulation is less important for unbiased consumers does not depend entirely on the elimination of three-part tariffs. In a final counterfactual, we simulate the effect of bill-shock regulation while holding observed public prices constant. In this simulation, bill-shock regulation benefits consumers with estimated biases by $22 but benefits debiased consumers by only $6. Debiased consumers are affected less by bill-shock alerts because (even holding prices constant) they make better plan choices that lead to lower incidence of overages.

9   Conclusion

We specify and estimate a model of consumer cellular-phone plan and usage choices. We identify the distribution of consumer tastes from observed usage and consumers' beliefs about their future usage from observed plan choices. Comparing the two we find that consumers underestimate their average taste for calling, underestimate their own uncertainty about their average tastes, and underestimate the volatility of their tastes from month-to-month. Because the magnitude of overconfidence is substantially larger than that of volatility bias, consumers correct initial plan choice mistakes more slowly than would unbiased consumers.

We conduct counterfactual simulations in which we (a) eliminate biases and (b) quantify the welfare impact of bill-shock regulation. We find that eliminating biases significantly increases consumer welfare, by $42 annually per consumer holding observed public prices fixed, and $201 annually per consumer accounting for firms' endogenous pricing response. If observed prices do not respond to bill-shock regulation, then the average consumer will benefit by $22 annually. This finding is reversed when firms optimally respond to bill-shock regulation. Although consumers avoid overage fees, firms raise monthly fees and average consumer surplus falls by $2 annually. In either case, bill-shock regulation lowers total welfare. Finally, we find that bill-shock regulation would have little to no effect if consumers were unbiased.

Our evaluation of bill-shock regulation could be insightful in other relevant contexts as well. For instance, in 2009 US checking overdraft fees totalled more than $38 billion and have been the subject of new Federal Reserve Board regulation (Martin 2010, Federal Reserve Board 2009). Convincing evidence of consumer inattention (Stango and Zinman 2009, Stango and Zinman 2010) suggests that this fee revenue would be dramatically curtailed if the Fed imposed its own bill-shock regulation by requiring debit card processing terminals to ask users "$35 overdraft fee applies, continue Yes/No?" before charging fees. Our counterfactual shows that in the cellular context consumers are nevertheless made worse off after accounting for endogenously higher fixed fees.

References

Ackerberg, Daniel A., "Advertising, Learning, and Consumer Choice in Experience Good Markets: An Empirical Examination," *International Economic Review*, 2003, *44* (3), 1007–1040.

Altschul, Michael F., Christopher Guttman-McCabe, and Brian M. Josef, "Comments of CTIA - The Wireless Association," January 10th 2011. `http://fjallfoss.fcc.gov/ecfs/document/view?id=7021025497`.

Armstrong, Mark and John Vickers, "Competitive Price Discrimination," *RAND Journal of Economics*, 2001, *32* (4), 579–605.

Ascarza, Eva, Anja Lambrecht, and Naufel J. Vilcassim, "When Talk is 'Free': The Effect of Tariff Structure on Usage under Two- and Three-Part Tariffs," *SSRN eLibrary*, 2012.

Ater, Itai and Vardit Landsman, "The Role of Overage Payments in Shaping Customers' Choice Among Three-Part Tariff Plans," 2011.

Bar-Gill, Oren and Rebecca Stone, "Pricing Misperception: Explaining Pricing Structure in the Cellular Service Market," *SSRN eLibrary*, 2009. `http://ssrn.com/paper=1425046`.

Baron, David P. and David Besanko, "Regulation and Information in a Continuing Relationship," *Information Economics and Policy*, 1984, *1* (3), 267–302.

Bhargava, Saurabh and Vikram Pathania, "Driving Under the (Cellular) Influence," Working Paper August 27 2011.

Blundell, Richard and Stephen Bond, "Initial conditions and moment restrictions in dynamic panel data models," *Journal of Econometrics*, 1998, *87* (1), 115–143.

Borenstein, Severin, "To What Electricity Price Do Consumers Respond? Residential Demand Elasticity Under Increasing-Block Pricing," Preliminary Draft April 30 2009.

Busse, M. R., "Multimarket Contact and Price Coordination in the Cellular Telephone Industry," *Journal of Economics and Management Strategy*, 2000, *9* (3), 287–320.

Cardon, James H. and Igal Hendel, "Asymmetric Information in Health Insurance: Evidence from the National Medical Expenditure Survey," *The RAND Journal of Economics*, 2001, *32* (3), 408–427.

Ching, Andrew, Tülin Erdem, and Michael P. Keane, "The Price Consideration Model of Brand Choice," *Journal of Applied Econometrics*, 2009, *24* (3), 393–420.

Chintagunta, Pradeep, Puneet Manchanda, and S. Sriram, "Empirical Investigation of Consumer Adoption, Consumption, and Termination of a Video on Demand Service," Work in Progress 2009.

Conlin, Michael, Ted O'Donoghue, and Timothy J. Vogelsang, "Projection Bias in Catalog Orders," *American Economic Review*, 2007, *97* (4), 1217–1249.

Courty, Pascal and Hao Li, "Sequential Screening," *The Review of Economic Studies*, 2000, *67* (4), 697–717.

Crawford, Gregory S. and Matthew Shum, "Uncertainty and Learning in Pharmaceutical Demand," *Econometrica*, 2005, *73* (4), 1137–1173.

CTIA - The Wireless Association, "CTIA-The Wireless Association, Federal Communications Commission and Consumers Union Announce Free Alerts to Help Consumers Avoid Unexpected Overage Charges," October 17 2011. `http://www.ctia.org/media/press/body.cfm/prid/2137`.

＿, "Year-End 2010 Top-Line Survey Results," Technical Report 2011. `http://files.ctia.org/pdf/CTIA_Survey_Year_End_2010_Graphics.pdf`.

DeGroot, Morris H., *Optimal Statistical Decisions*, New York: McGraw-Hill, 1970.

DellaVigna, Stefano and Ulrike Malmendier, "Contract Design and Self-Control: Theory and Evidence," *The Quarterly Journal of Economics*, 2004, *119* (2), 353–402.

Deloney, Amalia, Linda Sherry, Susan Grant, Parul P. Desai, Chris M. Riley, Matthew F. Wood, John D. Breyault, Jessica J. Gonzalez, and Benjamin Lennett, "Comments of the Center for Media Justice, Consumer Action, Consumer Federation of America, Consumers Union, Free Press, Media Access Project, National Consumers League, National Hispanic Media Coalition and New America Foundation Open Technology Initiative in response to notice of proposed rulemaking.," January 10th 2011. `http://fjallfoss.fcc.gov/ecfs/document/view?id=7021025418`.

Dutang, Christophe and Petr Savicky, "randtoolbox: Generating and Testing Random Numbers.," 2010. R package version 1.10 `http://cran.r-project.org/web/packages/randtoolbox/`.

Einav, Liran, Amy Finkelstein, Iuliana Pascu, and Mark Cullen, "How General Are Risk Preferences? Choices under Uncertainty in Different Domains," *American Economic Review*, Forthcoming.

Eliaz, Kfir and Ran Spiegler, "Contracting with Diversely Naive Agents," *The Review of Economic Studies*, 2006, *73* (3), 689–714.

＿ **and** ＿, "Consumer Optimism and Price Discrimination," *Theoretical Economics*, 2008, *3* (4), 459–497.

Erdem, Tülin and Michael P. Keane, "Decision-Making under Uncertainty: Capturing Dynamic Brand Choice Processes in Turbulent Consumer Goods Markets," *Marketing Science*, 1996, *15* (1), 1–20.

Federal Reserve Board, "Federal Reserve announces final rules prohibiting institutions from charging fees for overdrafts on ATM and one-time debit card transactions," Press Release November 12 2009.

fei Lee, Lung, "Estimation of dynamic and ARCH Tobit models," *Journal of Econometrics*, 1999, *92* (2), 355–390.

Gaynor, Martin S., Yunfeng Shi, Rahul Telang, and William B. Vogt, "Cell Phone Demand and Consumer Learning - An Empirical Analysis," *SSRN eLibrary*, 2005.

Goettler, Ronald L. and Karen B. Clay, "Tariff Choice with Consumer Learning and Switching Costs," *Journal of Marketing Research*, 2011, *48* (4), 633–652.

⎯⎯ **and Ron Shachar**, "Spatial Competition in the Network Television Industry," *The RAND Journal of Economics*, 2001, *32* (4), 624–656.

Gourieroux, Christian and Alain Monfort, "Simulation-based inference : A survey with special reference to panel data models," *Journal of Econometrics*, 1993, *59* (1-2), 5–33.

Grubb, Michael D., "Selling to Overconfident Consumers," *American Economic Review*, 2009, *99* (5), 1770–1807.

⎯⎯ , "Bill Shock: Inattention and Price-Posting Regulation," mimeo 2011.

Hajivassiliou, Vassilis A. and Paul A. Ruud, "Classical Estimation Methods for LDV Models Using Simulation," in R.F. Engle and Daniel L. McFadden, eds., *Handbook of Econometrics, Volume IV*, Amsterdam: North-Holland, 1994.

Handel, Benjamin, "Adverse Selection and Switching Costs in Health Insurance Markets: When Nudging Hurts," 2011.

Herweg, Fabian and Konrad Mierendorff, "Uncertain Demand, Consumer Loss Aversion, and Flat-Rate Tariffs," *Journal of the European Economic Association*, Forthcoming.

Huang, Ching-I., "Estimating Demand for Cellular Phone Service Under Nonlinear Pricing," *Quantitative Marketing and Economics*, 2008, *6* (4), 371–413.

Ito, Koichiro, "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing," 2010.

Iyengar, Raghuram, Asim Ansari, and Sunil Gupta, "A Model of Consumer Learning for Service Quality and Usage," *Journal of Marketing Research*, 2007, *44* (4), 529–544.

⎯⎯ , **Kamel Jedidi, and Rajeev Kohli**, "A Conjoint Approach to Multipart Pricing," *Journal of Marketing Research*, 2008, *45* (2), 195–210.

Jiang, Lai, "The Welfare Effects of ''Bill Shock" Regulation in Mobile Telecommunication Markets," Working Paper July 2011.

Kim, Jiyoung, "Consumers' Dynamic Switching Decisions in the Cellular Service Industry," Working Paper, SSRN November 2006.

Lambrecht, Anja and Bernd Skiera, "Paying Too Much and Being Happy About It: Existence, Causes, and Consequences of Tariff-Choice Biases," *Journal of Marketing Research*, 2006, *43* (2), 212–223.

⎯⎯ , **Katja Seim, and Bernd Skiera**, "Does Uncertainty Matter? Consumer Behavior under Three-Part Tariffs," *Marketing Science*, 2007, *26* (5), 698–710.

Leider, Steve and Özge Şahin, "Contracts, Biases and Consumption of Access Services," Ross School of Business Paper July 14 2011. `http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1895468`.

44

Lichtenstein, Sarah, Baruch Fischhoff, and Lawrence D. Phillips, "Calibration of Probabilities: The State of the Art to 1980," in Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment under uncertainty : heuristics and biases*, Cambridge ; New York: Cambridge University Press, 1982, pp. 306–334.

Liebman, Jeffrey B. and Richard Zeckhauser, "Schmeduling," Working Paper October 2004.

Loewenstein, George, Ted O'Donoghue, and Matthew Rabin, "Projection Bias in Predicting Future Utility," *The Quarterly Journal of Economics*, 2003, *118* (4), 1209–1248.

Martin, Andrew, "Bank of America to End Debit Overdraft Fees," *The New York Times*, March 10 2010.

Miravete, Eugenio J., "Screening Consumers Through Alternative Pricing Mechanisms," *Journal of Regulatory Economics*, 1996, *9* (2), 111–132.

———, "Estimating Demand for Local Telephone Service with Asymmetric Information and Optional Calling Plans," *The Review of Economic Studies*, 2002, *69* (4), 943–971.

———, "Choosing the wrong calling plan? Ignorance and learning," *American Economic Review*, 2003, *93* (1), 297–310.

———, "The Welfare Performance of Sequential Pricing Mechanisms," *International Economic Review*, 2005, *46* (4), 1321–1360.

——— **and Ignacio Palacios-Huerta**, "Inattention, Choice Dependence, and Learning from Experience in a Repeated Decision Problem," mimeo 2011.

——— **and Lars-Hendrik Röller**, "Estimating Price - Cost Markups Under Nonlinear Pricing Competition," *Journal of the European Economic Association*, 2004, *2* (2-3), 526–535.

Narayanan, Sridhar, Pradeep K. Chintagunta, and Eugenio J. Miravete, "The role of self selection, usage uncertainty and learning in the demand for local telephone service," *Quantitative Marketing and Economics*, 2007, *5* (1), 1–34.

Osborne, Matthew, "Consumer Learning, Switching Costs and Heterogeneity: A Structural Examination," *Quantitative Marketing and Economics*, 2011, *9* (1), 25–70.

Park, Minjung, "The Economic Impact of Wireless Number Portability," 2009.

Reiss, Peter C. and Matthew W. White, "Household Electricity Demand, Revisited," *The Review of Economic Studies*, 2005, *72* (3), 853–883.

Riordan, Michael H. and David E. M. Sappington, "Awarding Monopoly Franchises," *American Economic Review*, 1987, *77* (3), 375–387.

Rochet, Jean-Charles and Lars A. Stole, "Nonlinear Pricing with Random Participation," *The Review of Economic Studies*, 2002, *69* (1), 277–311.

— **and Lars Stole**, "The Economics of Multidimensional Screening," in M. Dewatripont, Lars Peter Hansen, and Stephen J. Turnovsky, eds., *Advances in economics and econometrics: theory and applications - eighth world Congress*, Vol. 36 of <u>Econometric Society Monographs</u>, New York: Cambridge University Press, 2003.

Roodman, David, "How to do xtabond2: An introduction to difference and system GMM in Stata," *Stata Journal*, 2009, *9* (1), 86–136.

Saez, Emmanuel, "Do Taxpayers Bunch at Kink Points?," Working Paper June 2002.

— , "Do Taxpayers Bunch at Kink Points?," *American Economic Journal: Economic Policy*, 2010, *2* (3), 180–212.

Seim, Katja and V. Brian Viard, "The Effect of Market Structure on Cellular Technology Adoption and Pricing," *American Economic Journal: Microeconomics*, 2010.

Spiegler, Ran, *Bounded Rationality and Industrial Organization*, Oxford University Press, 2011.

Stango, Victor and Jonathan Zinman, "What do Consumers Really Pay on Their Checking and Credit Card Accounts? Explicit, Implicit, and Avoidable Costs," *American Economic Review Papers and Proceedings*, 2009, *99* (2).

— **and** — , "Limited and Varying Consumer Attention: Evidence from Shocks to the Salience of Overdraft Fees," 2010.

Stole, Lars A., "Price Discrimination and Competition," in Mark Armstrong and Robert K. Porter, eds., *Handbook of Industrial Organization*, Vol. Volume 3, Elsevier, 2007, chapter 34, pp. 2221–2299.

Train, Kenneth, *Discrete Choice Methods with Simulation*, 2nd ed., Cambridge University Press, 2009.

U.S. Census Bureau, "Table 1. Preliminary Annual Estimates of the Resident Population for the United States, Regions, States, and Puerto Rico: April 1, 2000 to July 1, 2010," Technical Report NST-PEST2010-01 February 2011. `http://www.census.gov/popest/states/tables/NST-PEST2010-01.xls`.

Uthemann, Andreas, "Competitive Screening of Customers with Non-Common Priors," 2005. `http://www.ucl.ac.uk/~uctpaut/research/uthemann_ncp.pdf`.

Yao, Song, Carl F. Mela, Jeongwen Chiang, and Yuxin Chen, "Determining Consumers' Discount Rates with Field Studies," 2011.