

Consumer Dynamic Usage Allocation and Learning under Multi-part Tariffs

Arun Gopalakrishnan
Raghuram Iyengar
Robert J. Meyer *

The Wharton School, University of Pennsylvania

Abstract

Multi-part tariffs are widely favored within service industries as an efficient means of mapping prices to differential levels of consumer demand. Whether they benefit consumers, however, is far less clear as they pose individuals with a potentially difficult task of dynamically allocating usage over the course of each billing cycle. In this paper we explore this welfare issue by examining the ability of individuals to optimally allocate consumption over time in a stylized cellular-phone usage task for which there exists a known optimal dynamic utilization policy. Actual call behavior over time is modeled using a dynamic choice model that allows decision makers to both discount the future (be myopic) and be subject to random errors when making call decisions. Our analysis provides a “half empty, half full” view of intuitive optimality. Participants rapidly learn to exhibit far-sightedness, yet learning is incomplete with some level of allocation errors persisting even after repeated experience. We also find evidence for an asymmetric effect in which participants who are exogenously switched from a low (high) to high (low) allowance plan make more (fewer) errors in the new plan. The effect persists even when participants make their own plan choices. Finally, interventions that provide usage information to help participants eradicate errors have limited effectiveness.

Keywords: Multi-part tariffs, dynamic allocation, consumer learning, dynamic decision making, intertemporal discounting

*Authorship order is alphabetical. Arun Gopalakrishnan is a doctoral candidate; Raghuram Iyengar is an Associate Professor of Marketing; and Robert J. Meyer is the Gayfryd Steinberg Professor of Marketing, all at the Wharton School of the University of Pennsylvania. Correspondence concerning this manuscript should be addressed to Arun Gopalakrishnan, agop@wharton.upenn.edu, The Wharton School, 700 Jon M. Huntsman Hall, 3730 Walnut Street, Philadelphia, PA 19104. The authors thank Alessandro Arlotto, Eric Bradlow, Jeffrey Cai, Yupeng Chen, Pete Fader, Christophe Van den Bulte, seminar participants at the 2012 Wharton Marketing Lunch Seminar series, Cornell University, and University of Washington Marketing Camp, and session participants at the 2012 INFORMS Marketing Science, 2013 UT Dallas and 2013 TPM conferences for their helpful comments. The authors thank Amanda Gulick, Tony Kutovoy, Robert Botto, and Rich Fritzon for data collection assistance and support. The authors thank the Marketing Science review team for their thoughtful comments throughout the review process. Arun Gopalakrishnan acknowledges the generous support of the Russell Ackoff Doctoral Student Fellowship and the Wharton Risk Management and Decision Processes Center for this research. This work was supported by an AWS in Education Grant award.

1. Introduction

The multi-part tariff is one of the most widely used approaches to pricing in service industries such as cellular phones, electricity distribution, vehicle leasing, retail banking, and online retailing. By charging consumers different prices for different levels of consumption, multi-part tariffs have the advantage of allowing firms both to price discriminate among consumers who have different levels of demand and fairly allocate limited capacity to those who need it the most (Bagh and Bhargava 2013; Oren et al. 1985). Whether such pricing mechanisms offer comparable welfare benefits for consumers, however, is far less clear.

The primary source of concern is that multi-part tariffs potentially complicate otherwise simple decisions about service utilization and tariff choice. Under a typical pay-per-use plan, for example, a consumer who wishes to decide whether to consume a unit of a service (such as place a phone call) simply needs to assess whether the marginal utility of usage exceeds its marginal cost. Under multi-part plans, however, the decision is rendered more complex by the need to also consider how consumption in the present may affect the marginal cost of future consumption. If consumers find it difficult to make present-future consumption tradeoffs, or lose track of accumulated usage under a given plan (e.g., Grubb and Osborne 2013), multi-part tariffs could distort consumption patterns in a way that departs from what would be optimal for the consumer in the long term. An example would be cell phone consumers who too quickly use up their “free minutes” (the monthly allowance), exposing them to otherwise avoidable overage charges (which represent an increase in marginal cost) later in the month.

How pervasive are such consumer errors within the context of multi-part tariffs? The empirical evidence, based largely on aggregate (typically monthly) usage, is somewhat mixed. Bar-Gill and Stone (2009), for example, find that cell-phone consumers are often poor judges of

the calling plans that best match their actual usage. In their work consumers exceeded their call allowance 17% of the time (by an average of 33%), and those who did not exceed it used, on average, just 47% of their call allowance. Likewise, Ascarza et al. (2012) report that consumers increase usage when switching from two-part (fixed access fee and pay-per-use rate) to three-part tariffs (fixed access fee, usage allowance, and overage rate), an inflation they attribute to a biased tendency to increase usage when consumption within the usage allowance is perceived as “free”. Finally, Lambrecht and Skiera (2006), Lambrecht et al. (2007), Goettler and Clay (2011), and Ater and Landsman (2013) all offer suggestive evidence that consumers themselves sense the difficulties of making usage decisions under multi-part tariffs, with many preferring flat-rate plans to multi-part tariffs even when total *ex post* cost is higher.

Two sets of factors limit the strength of conclusions about optimality that can be drawn from data on aggregate usage. First, because such data do not reveal the full set of beliefs held by consumers when they were making decisions--such as their expectations of future usage, consumption utility, and temporal planning horizon --it is impossible to draw strong conclusions about the optimality of these decisions. For example, a consumer who is observed incurring overage charges might well have been behaving optimally if s/he was strategically responding to an unexpected surge in consumption opportunities providing high utility. Second, as noted by Miravete (2003), even if consumers do err when choosing plans and/or making consumption decisions, such mistakes may be transient, diminishing over time as consumers become more skilled in forecasting their usage patterns and disciplining consumption during a billing cycle.

The purpose of this paper is to take a first step towards empirically assessing the ability of individuals to make optimal usage decisions under multi-part pricing schemes. We do so by studying how consumers make call-by-call usage allocation decisions under three-part tariffs in a

controlled laboratory task for which there is a known optimal policy for consumption over time. In this task participants make a series of decisions about whether to accept or reject a series of incoming phone calls that have a known stochastic arrival rate. As in natural settings, when calls arrive, participants observe the utility of the call (simulating caller ID), decide whether to accept each call, and can terminate calls. Participants play for a monetary incentive with the goal of maximizing the total net utility of their phone usage decisions over a time horizon.

We emerge from the work with four key findings. First, we show how participants' call-acceptance decisions in the task can be captured by a heterogeneous dynamic utility model that assumes that they are forward-looking when making call-acceptance decisions, but subject to both differential rates of temporal discounting and random errors in usage. Estimates of the structural parameters of the model suggest that while participants quickly learn to avoid myopia (excessive discounting) when making decisions in the task, random errors are more persistent, which forms the main source of departures from optimality among experienced users. Second, while aggregate performance improves relative to normative benchmarks with game experience, this convergence is incomplete and displays several sustained departures from optimality. Among these is a tendency for behavior to be "noisier" as the free allowance is increased in the three-part tariff, and participants who are exogenously switched from a low (high) to high (low) allowance plan make more (fewer) errors in the new plan. Third, this asymmetry persists even when allowing participants to self-select into plans. Finally, interventions such as providing subjects with full information about talk time usage and/or creating an alert when reaching 80% of free allowance has almost no effect on performance, suggesting that *inattention* to usage is unlikely to be a major driver of errors.

The remainder of this paper is organized as follows. In Section 2, we describe the general form of the decision task faced by participants in our laboratory setting, and develop a dynamic choice model of consumer decision making that incorporates temporal discounting, random errors, learning, and heterogeneity. In Section 3, we explain our design choices and describe the experimental setting in which we collect data. In Section 4, we report the findings from four studies based on our experimental paradigm along with robustness checks. We conclude in Section 5 with a discussion of the work’s implications for researchers and regulators.

2. Theory and Model Development

2.1 Overview

In this paper we study how individuals solve a series of cell-phone calling tasks that have the following structure:

A consumer i faces a series of incoming calls in game g (out of a total of G games) and must decide whether to accept or decline each opportunity. The costs of calls are governed by a three-part-tariff (TPT) plan \wp_{ig} in which the consumer pays a fixed up-front fee F_{ig} for the first \bar{W}_{ig} seconds of “free” talk time, with overage charges of c_{ig} units per second. Upon receiving a call at time t the consumer first observes its marginal utility per second u_{igt} , which is an independent and identically distributed (i.i.d.) draw from a known distribution $h(u_{igt})$. If he or she accepts the call it lasts for up to L seconds. Once a call is accepted, the consumer can terminate the call at any time or accept another incoming call whose utility can be observed even while an accepted call is in progress. Time is discrete, and calls arrive as a Bernoulli process with a fixed, known, probability p in each second. The consumer’s goal is to make a series of accept, switch,

or reject decisions that maximize total net call utility over a finite time horizon of T seconds.

Although abstract in form, this stylized task captures many of the common features of real-world calling problems. As in natural settings the consumer must decide how best to utilize “free” talk time under a budget that mimics monthly billing cycles, s/he faces potential overage charges that can exceed the average unit time cost (fixed fee divided by “free” talk time), the likely utility of an incoming call is known (e.g., the consumer is assumed to have “caller ID”), call duration is under the control of the consumer, and s/he has the ability to switch to a new incoming call while on a current call. Likewise, because the task will be played repeatedly, there are opportunities to learn, and the consumer faces the possibility of plan changes that s/he must adapt to. That said, we do not accommodate all features of real world calling in order to retain a form that yields a clear optimal policy. We do not, for example, allow subjects to make outgoing calls, allow for rollover of unused call allowance, or distinguish between off-peak and on-peak usage, and every call’s duration has the same upper bound. In the discussion we address the degree to which the current findings might be expected to generalize to a broader range of dynamic consumption problems.

2.2 Normative Theory and Empirical Model

In this section we describe how call acceptance decisions would be made by a rational decision maker who seeks to maximize net call utility over a finite time horizon. Because one of our central research goals is to examine how actual acceptance decisions compare to this optimal benchmark, the model allows for two behavioral features that might change over games as participants gain task experience: intertemporal discounting, or the tendency to give more weight to the value of calls that are immediately at hand versus those that may lie in the future, and

time-varying random response errors. Following Salisbury and Feinberg (2010), these errors reflect unobserved drivers of utility that lie outside the rational policy, such as momentary lapses of attention and transient errors that arise from making decisions via non-normative heuristics.

In Section 2.2.1 we first describe how a utility-maximizing consumer would solve our focal decision problem given that s/he discounts future utilities absent random choice errors, and then extend this model in 2.2.2 to allow for such errors. In Section 2.2.3 we discuss the form of the model that we estimate empirically, which includes heterogeneity in discounting and error propensities.

2.2.1 The optimal policy without random choice errors

Consider the optimal policy for consumer i who wishes to maximize utility over a discounted finite time horizon in game g with discount factor β_{ig} under a TPT plan \wp_{ig} (with fixed fee F_{ig} , “free” allowance \bar{W}_{ig} , and per-second overage charge c_{ig}) by making a series of decisions $\{\delta_{igt}\}$.

The consumer is assumed to be intertemporally risk neutral in assessing the expected total utility of future calls.¹ Note that the discount factor is allowed to be consumer- and game- specific. Let $t \in \{1, 2, \dots, T\}$ be the decision times within game g , where T is the total number of seconds.

Assume that in each second there is a constant probability p of an arriving call with strictly positive and stochastic utility u_{igt} with p.d.f. $h(u_{igt})$. For notational convenience, we set $u_{igt} = 0$ to represent the event of no-call arrival at time t (which occurs with probability $1 - p$),

yielding the mixed probability distribution $f(u_{igt}) = (1 - p) \cdot I(u_{igt} = 0) + p \cdot h(u_{igt}) \cdot$

$I(u_{igt} > 0)$.

¹ In our experimental work we will examine whether observed departures from optimality might accrue due to risk aversion among consumers. We impose intertemporal risk neutrality here both for model identification and to make the analysis consistent with prior dynamic consumer utility models, which also make the same assumption (e.g., Erdem and Keane 1996).

At time t , a consumer derives current-period utility $K(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}, \delta_{igt})$ from new call acceptance ($\delta_{igt} = 1$) or existing call continuation ($\delta_{igt} = 0$) and obtains zero utility from termination or rejection of a call ($\delta_{igt} = -1$). As shown in Figure 1, the space of call decisions $\Delta_{igt}(u_{igt}, \underline{u}_{igt})$ depends on available new and existing calls and period utility is a function of u_{igt} – the value of an incoming call (could be 0 if none), \underline{u}_{igt} – the value of an ongoing call (could be 0 if none), D_{igt} – the number of seconds remaining on an existing call, and w_{igt} – the number of free seconds (or talk time) used.

Insert Figure 1 Here

The consumer's expected utility is given by $E[\sum_{t=1}^T \beta_{ig}^{t-1} \cdot K(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}, \delta_{igt} \in \Delta_{igt}(u_{igt}, \underline{u}_{igt}))]$, and can be formulated as a dynamic programming problem as shown in equation (1) where $V_{igt}(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}, \delta_{igt})$ is the conditional value function and $\bar{V}_{ig,t+1}(\underline{u}_{ig,t+1}, D_{ig,t+1}, w_{ig,t+1}) \equiv E_{u_{ig,t+1}}[\max_{\delta_{ig,t+1}} V_{ig,t+1}(u_{ig,t+1}, \underline{u}_{ig,t+1}, D_{ig,t+1}, w_{ig,t+1}, \delta_{ig,t+1})]$ is the expected continuation value conditional on future state parameters. State variables include $u_{igt}, \underline{u}_{igt}, D_{igt}$, and w_{igt} but only u_{igt} is unknown prior to the start of time t .

$$\begin{aligned}
 &V_{igt}(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}, \delta_{igt}) = \\
 &\left\{ \begin{array}{l} [u_{igt} - c_{ig}I(w_{igt} \geq \bar{W}_{ig}) + \beta_{ig}\bar{V}_{ig,t+1}(u_{igt}, L - 1, w_{igt} + 1)]I(u_{igt} > 0), \delta_{igt} = 1 \\ [u_{igt} - c_{ig}I(w_{igt} \geq \bar{W}_{ig}) + \beta_{ig}\bar{V}_{ig,t+1}(u_{igt}, D_{igt} - 1, w_{igt} + 1)]I(\underline{u}_{igt} > 0)I(D_{igt} > 0), \delta_{igt} = 0 \\ \beta_{ig}\bar{V}_{ig,t+1}(0, 0, w_{igt}), \delta_{igt} = -1, \end{array} \right\} \tag{1}
 \end{aligned}$$

In Web Appendix A, we lay out formal propositions that characterize the properties of the above value function and the optimal call-acceptance policy that follows from solving this

dynamic program. Most have a strong intuitive flavor. For example, if the consumer has an incoming call but no existing call at time t ($u_{igt} > 0, \underline{u}_{igt} = 0$), the optimal decision of whether to accept ($\delta_{igt} = 1$) or reject ($\delta_{igt} = -1$) the call involves a tradeoff between immediate call utility (less any overage charge) and its associated opportunity cost: the difference in expected future utility from saving this unit of talk time for future consumption and taking the call. Once a call is accepted, a rational consumer then needs to decide how long to stay on it. In the absence of another incoming call, the optimal call continuation decision has a similar intuition to call acceptance. A rational consumer might, however, choose to switch calls if a new competing one arrives. In this situation three decisions are possible ($\delta_{igt} = \{1, 0, -1\}$): continue the current call, take the new call, or terminate the current call without taking the new one. In this case the consumer would compute the same tradeoff between present utility and future opportunity costs as above, but here via two steps, first deciding which of the two calls offers the higher level of long-run utility, and then deciding whether the better of the two is worth continuing or taking (which could cause both calls to be rejected).

2.2.2 Introducing Random Choice Errors

It is unlikely, of course, that consumers would always trade-off utility and costs in an optimal manner. Even if they understood the optimal policy, its application may occasionally be disrupted by lapses of attention or transient mood effects that are unobserved by the analyst. Likewise, in the course of learning how best to accept calls consumers might randomly try out a mix of heuristic policies that sometimes lead them to accept calls they should decline and sometimes reject those that should be accepted. To capture the effect of such errors in the decisions of consumer i in game g , we introduce a time-varying additive i.i.d. Extreme Value

Type I shock term $\epsilon_{igt}(\delta_{igt})$ with scale parameter λ_{ig} in the model formulation (Rust 1987; Salisbury and Feinberg 2010) such that the consumer chooses the decision

$\delta_{igt}^* = \operatorname{argmax}_{\delta_{igt}} [V_{igt}(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}, \delta_{igt}) + \epsilon_{igt}(\delta_{igt})]$. We non-parametrically capture learning that may reduce transient errors across games by allowing λ_{ig} to be game-specific.

In the extreme case, as $\lambda_{ig} \rightarrow 0$, a consumer's actions are perfectly characterized by the optimal policy for her given discount factor β_{ig} . On the other hand, a larger λ_{ig} indicates a noisy policy that deviates with increasing probability from what is optimal given the consumer's β_{ig} (cf. McKelvey and Palfrey 1995). As $\lambda_{ig} \rightarrow \infty$, consumer i makes completely random decisions in game g . The expected continuation value

$\bar{V}_{ig,t+1}(\underline{u}_{ig,t+1}, D_{ig,t+1}, w_{ig,t+1}) \equiv E_{u_{ig,t+1}}[\bar{V}_{ig,t+1}(u_{ig,t+1}, \underline{u}_{ig,t+1}, D_{ig,t+1}, w_{ig,t+1})]$ is now a function of ex ante value function $\bar{V}_{ig,t+1}(u_{ig,t+1}, \underline{u}_{ig,t+1}, D_{ig,t+1}, w_{ig,t+1})$, which requires integrating over a vector of future shock terms $\epsilon_{ig,t+1}$, unknown to both the consumer and researcher. Note that current-period error terms ϵ_{igt} are assumed known to the consumer at time t and only realized for $\delta_{igt} \in \Delta_{igt}(u_{igt}, \underline{u}_{igt})$. However, the researcher only knows the distribution of these shocks and must integrate over them to arrive at the probability of making decision δ_{igt} .

Given that the error terms are distributed Extreme Value Type I, closed form solutions are available for conditional choice probabilities and the ex ante value function (Arcidiacono and Ellickson 2011; Rust 1987) as shown in equations (2) and (3) respectively. Here, the term

$\delta_{igt}^{max} = \operatorname{argmax}_{\delta_{igt}} \Pr(\delta_{igt} | u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt})$, and γ is the Euler-Mascheroni constant.

$$\Pr(\delta_{igt} | u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}) = \frac{\exp(V_{igt}(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}, \delta_{igt})/\lambda_{ig})}{\sum_{m \in \Delta_{igt}(u_{igt}, \underline{u}_{igt})} \exp(V_{igt}(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}, m)/\lambda_{ig})} \quad (2)$$

$$\begin{aligned} \bar{V}_{igt}(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}) = \\ -\ln(\Pr(\delta_{igt}^{max} | u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt})) + V_{igt}(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}, \delta_{igt}^{max}) + \gamma \lambda_{ig} \end{aligned} \quad (3)$$

The *ex-ante* value function can be interpreted as the value function for the decision with the highest probability of being chosen plus an upside coming from possible shocks that may make other decisions more attractive. The term involving γ does not affect conditional choice probabilities. As noted before, in the limiting case of $\lambda_{ig} \rightarrow 0$, the ex ante value function becomes equivalent to the conditional value function given δ_{igt}^{max} as the decision rule becomes deterministic (i.e., $\Pr(\delta_{igt}^{max} | u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}) \rightarrow 1$).

Given a vector of observed consumer decisions $Y_{ig} \equiv \{\delta_{ig1}, \delta_{ig2}, \dots, \delta_{igT}\}$ and call arrivals $X_{ig} \equiv \{u_{ig1}, u_{ig2}, \dots, u_{igT}\}$ for each game, we can construct the likelihood function for consumer i over G games using the standard assumption of conditional independence of the temporal error terms (e.g., Rust 1987) as shown below.

$$\prod_{g=1}^G p(Y_{ig} | \beta_{ig}, \lambda_{ig}, X_{ig}, \wp_{ig}) = \prod_{g=1}^G \prod_{t=1}^T \Pr(\delta_{igt} | u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}) \quad (4)$$

State variables at each decision node can be constructed from X_{ig} and Y_{ig} . The data therefore constitute discrete choices made at different values of state parameters. An important point to note is that X_{ig} is independent of Y_{ig} by design – arriving calls have nothing to do with the decisions the consumer made in the past.

2.2.3 Model Identification and Heterogeneity

Our dynamic choice model has two parameters to capture decision making by consumer i in game g : an intertemporal discount factor (β_{ig}) and error scale parameter (λ_{ig}). The scale parameter is usually set to a constant (e.g., Arcidiacono and Ellickson 2011) to estimate the

slopes of a latent utility function. Since call utilities are assumed known to both the consumer and researcher (we provide them in the experimental task), the slope for current period utility is essentially fixed to a value of 1, which allows us to estimate λ_{ig} .² We also assume that state transition probabilities are known to both the consumer and researcher (since they are explicitly provided in the experimental task) unlike models estimated on field data (e.g., Erdem and Keane 1996), which allows us to pin down the discount factor based on observed choices across a large variation of state parameters at different points in game time (cf. Dubé et al. 2013) as it is the only unknown driving the systematic portion of conditional choice probabilities in equation (2).

The estimation of individual- and game- specific parameters will, however, be challenging due to sampling variation in call arrivals such that some games' data may be less informative than others for a particular pair of parameters (see Web Appendix E). To circumvent this issue of data sparseness and address heterogeneity across consumers as well, we propose a latent class framework in which the model parameters for each segment are allowed to vary across games non-parametrically (Kamakura and Russell 1989).³ Note that an important substantive benefit of allowing for multiple segments with game-specific parameters is that it will capture heterogeneity in learning within the subject pool – for example, some segments may learn to become optimal quickly, while others may not despite repeated experience. Further, our non-parametric specification (which nests the case of allowing a linear trend over games for each parameter) of how the segment level parameters (the discount factor and scale parameter) vary over the six games imposes no a priori restriction on the direction and speed of learning, which is

² As long as the slope for one term in the period utility function is fixed, the scale parameter is identified along with slopes of other terms. We exploit this in one of our robustness checks in which we estimate curvature in the per-period utility.

³ Consistent with typical latent class modeling framework, we assume homogeneity within a segment. While there may be heterogeneity even within a latent class, since the objective of this paper is to characterize the overall effects of three-part tariffs on learning and usage, we do not introduce further heterogeneity into the model. Further, several individual-level posterior distributions closely resemble the “segment average” given the sparseness of information in data from any one person-game.

important since patterns may not be monotonic. The likelihood function for a dataset of N subjects playing G games each is as follows:

$$\prod_{i=1}^N \sum_{s=1}^S p(i \in s) \prod_{g=1}^G p(Y_{ig} | \beta_{gs}, \lambda_{gs}, X_{ig}, \rho_{ig}) \quad (5)$$

where $p(i \in s)$ is the probability that subject i belongs to segment s and $p(Y_{ig} | \beta_{gs}, \lambda_{gs}, X_{ig}, \rho_{ig})$ the likelihood function for a person-game formed using conditional choice probabilities with segment level parameters (equation 2). Note that a homogenous model would have 12 parameters (two game-specific parameters by six games), a two-segment model would have 25 parameters (including the probability of segment membership), and so forth.

3. Empirical Setting

In this section we describe our experimental design, starting with the selection of task parameters and the associated normative predictions in Section 3.1. In Section 3.2, we provide an overview of the experimental method and studies that we ran.

3.1 Design Considerations and Normative Predictions

One of the major challenges in designing the experiment was to choose task parameters (e.g., the free allowance \bar{W} , the overage cost per second c , the time horizon T , and the probability distribution of call events $f(u)$) that would allow us to efficiently identify the parameters of the empirical model described above while minimizing the task burden on participants.

For example, while it would be desirable to observe a large number of call choices made in different contexts from each participant, pilot work suggested that participants became fatigued when the task lasted longer than 25-30 minutes. We also found that participants typically took twice the amount of time as the number of game-time seconds across all games to complete the task (including the time for decision making whenever a call arrives). As a result, we determined that six games with $T = 120$ seconds would provide a sufficiently large number

of games to assess learning and adaptation to plan changes, while also providing a long enough time horizon to create sufficient present/future tradeoffs.

Given $T = 120$, it was important to choose an allowance \bar{W} that was neither too high nor too low in order to facilitate identification of individuals' intertemporal discount rates. In particular, to identify these rates we had to choose \bar{W} such that participants faced consequential trade-offs in most periods of the game when deciding whether to accept a call now versus saving the talk time for future calls. Based on pilot work we found that allowances of 20 seconds (one sixth the total time) and 40 seconds (one-third) captured scenarios that were perceived by participants as reflecting sparse and ample free-talk time budgets, where opportunity costs were more and less salient, respectively. We also chose maximum call length $L = 8$ to allow for the real-world feature of call termination (which requires $L > 1$) and potentially multiple opportunities to switch to other calls (in this case, up to 7 opportunities since at most one call can arrive per second in our task). Though the fixed fee was sunk, we choose values of \$80 (\$20) for the plan with a 40-second (20-second) allowance to ensure task realism and set the expected surplus under both plans for a consumer who does not discount future call values ($\beta=1$) to be roughly equal.

The choices for the distribution of call arrivals $f(u)$ and overage rates c were crucial in determining the call acceptance threshold. We choose $f(u = 0) = 0.8$ so that the average probability of a call arriving was 0.2, with an expected number of 24 calls in a 120-second time horizon, providing a reasonable number of decision opportunities to participants within a game. A higher $f(u = 0)$ decreased the likelihood of overage by rendering call opportunities too sparse.

For participants to make tradeoffs about present and future call consumption, the overage charge should exceed some of the call values. We chose three call values (low: $u_L = \$0.25$, medium: $u_M = \$0.75$, and high: $u_H = \$6$), which arrive with known probabilities ($f(u_L) = 0.1, f(u_M) = 0.05, f(u_H) = 0.05$) and set the overage charge c to \$4 (falling between u_L and u_H)⁴. The large value differential between the high- and medium-valued calls simplified the consumer's decision calculus while allowing for efficient recovery of our model parameters.

In Figures 2 and 3, we present the optimal policies for when to accept low-valued (u_L) and medium-valued (u_M) calls under the above set of plan and game characteristics, for both a consumer who does not discount future call values ($\beta = 1$) and one who is more myopic. High-valued (u_H) calls should always be accepted regardless of temporal discounting.

Insert Figures 2 and 3 here

From Figures 2a and 2b, the higher the amount of free allowance consumed, the more judicious a consumer should be in taking low-valued calls (solid line) and medium-valued calls (dotted line). These lower-valued calls can be accepted earlier under a plan with a larger allowance since there is a larger pool of “free” seconds before the consumer starts approaching the overage phase. As the prospect of overage nears, both types of calls should generally be avoided. Under discounting for the 40-second allowance plan (Figures 3a and 3b) consumers would accept lower-valued calls much sooner than a consumer who does not discount future call values ($\beta=1$), and then shift to a higher acceptance threshold as the overage phase approaches. The timing of these inflection points varies by discount factor; a more myopic consumer will continue accepting lower-valued calls for a longer period of time before applying more stringent standards. Note that a low-valued call becomes unattractive more quickly than a medium-valued

⁴ A c that is too low makes overage essentially inconsequential while a c that is too high would preclude some calls from ever being rational to accept.

call with increasing usage. The intuition is similar for the 20-second allowance plan. Plan and game characteristics other than the ones we chose can also lead to meaningful differences in optimal policies which can be exploited for estimating structural parameters. An interesting direction for future research is to understand the boundary conditions of how these characteristics interact with consumer behavioral patterns.

3.2 Overview of Method and Studies

Each of the experimental studies described below had a similar structure. In all cases participants first played a 30-second practice game to familiarize themselves with the task and interface, and then proceeded to play six games for compensation. After each game, the net utility was displayed to the participant, breaking down the total value of calls consumed less fixed and overage costs. For incentive compatibility, we informed participants that all amounts shown above were in game dollars and revealed the exchange rate to convert net accumulated game dollars over the six games to a bonus payment in US dollars (over and above a show-up fee). The show-up fee and exchange rate varied depending on the platform we used to run each study. Web Appendix B contains illustrative call arrival patterns and decisions and how these help identify model parameters. In Web Appendix C, we include the full set of instructions to participants as they undertook the task.

The overall experiment comprised four main studies and several follow-up studies that allowed us to examine call behavior under different plan characteristics. In Study 1, we examine consumer learning as a function of plan allowance by exposing each subject to one of two TPT plans for six games. In Study 2, we examine consumer learning and adaptation to imposed plan change by exposing each subject to one TPT plan for three games and the other for the remaining three games. In Study 3, we replicate the structure of Study 2 but allow for initial plan choice and the option to switch after three games, as a test of whether self-selection impacts within-game

behavior. In Study 4, we examine the effect of risk attitude on intertemporal discounting. Two follow-up studies then tested the robustness of the earlier findings to manipulations that might act to improve the efficiency of call decisions: (1) providing a pop-up alert when subjects use up 80% of their free allowance, and (2) providing the alert and a visible real-time counter of talk time used. Additional follow-up studies tested for sunk cost effects arising from the fixed fee of the TPT and curvature in current period utility. For each study, we present model-free evidence of how subjects performed as a function of experience and/or plan changes using the natural benchmark of a consumer who does not discount future call values with no random errors in decision making.

We also present model estimates obtained by running an MCMC algorithm that is detailed in Web Appendix D. In addition to typical model selection criteria such as Log Marginal Likelihood (Newton and Raftery 1994) and DIC (Spiegelhalter et al. 2002) for determining the appropriate number of latent classes, we also considered segment size for model selection. We find that a two-segment model considerably improves fit over a homogeneous model, as there appear to be two highly heterogeneous sub-groups in the data. Additional segments continue to sub-divide one of these sub-groups to better fit the data but do not add new diagnostic insights (see Web Appendix F for fit statistics for a range of models with up to four latent classes and various nested cases) and in some instances result in segment sizes that are too small for robust parameter recovery (see Web Appendix E). To permit intuitive and robust like-for-like comparisons across studies and conditions, we present results for all studies using a two-segment latent class model.

4. Results

4.1 Study 1 – Learning under the same plan

4.1.1 Motivation and Procedure

The purpose of this study is to quantify how consumer temporal discounting and random response errors change with game experience as a function of tariff allowance. Consumers might be expected to consider a larger planning horizon (consistent with a higher discount factor) and reduce the propensity for random errors as they receive feedback from bills. In particular, it has been widely noted that “as-if” optimality often arises by the recurrent application of simple trial-and-error (reinforcement) learning rules based on repeating behaviors that yield positive outcomes (low bills and high utility) and avoiding those that do not (e.g., Meyer and Hutchinson 2001; Kunreuther et al. 2009; Young 2009). But whether and with what speed reinforcement rules will allow calling behavior to converge to optimality in this task is unclear. A potential barrier is the noisy nature of feedback that cellphone customers receive; if one ends a game facing overage charges and wants to learn from this experience, one has to imagine the counterfactual policy that would have yielded lower charges but still provided high levels of call utility.

In this study, participants were 125 members of a subject pool of a large Northeastern university. All studies were run in a computer lab. Subjects were paid \$10 to participate and could earn a bonus (typically between \$0 and \$2.50) based on an exchange rate of US\$1 for every 400 game dollars). Upon being seated participants first read the overall task instructions and were randomly allocated to a plan with an allowance of either 20 seconds (55 subjects) or 40 seconds (70 subjects). Participants then made decisions under this plan for all six games. We used a pre-defined criterion of task comprehension (accepting 33% or fewer of high-value calls across six games implies lack of task understanding) to screen out subjects not meeting this criterion.⁵ After screening, there remain 44 subjects in the 20-second condition and 63 subjects in

⁵ The qualitative features of the aggregate descriptive statistics do not change with the exclusion of these subjects.

the 40-second condition.⁶ The exact experimental instructions used and screen shots of the game are provided in Web Appendix C.

4.1.2 Findings: Overall Performance

In Table 1, we show descriptive statistics for number of call arrivals by type and how many were accepted on average across the six games in each condition. As expected, most high-value calls are accepted. A smaller number of low- and medium- value calls (lumped together) are accepted, though there is high variance across subjects.⁷

Insert Table 1 Here

In Figures 4 and 5 we plot how participants' call-acceptance decisions corresponded to those prescribed by a fully optimal ($\beta=1$) error-free policy in the 20-second and 40-second plans, respectively.

Insert Figures 4 and 5 Here

The figures plot the incidences of three types of potential decision errors: failing to correctly accept a new call, failing to correctly continue an existing call, or failing to correctly reject a call. In Figures 4a and 5a we see that participants were quite skilled at accepting calls that they should (which would usually be any high-value call) almost from the start, and this achievement increased with play experience. In the first game for the 40-second allowance condition, for example, participants correctly accepted 81% of all optimal calls, and this increased to 92% by the sixth game. In addition, as shown in Figures 4b and 5b, participants were even more skilled at knowing when to continue calls that should be continued; across all

⁶ A two-sample test of proportions shows that the proportions of subjects screened out in each condition are not statistically different (p-value = 0.11).

⁷ Most subjects exceed their free allowance since high-value calls continue to arrive in the overage phase and are worth more than overage cost. In our setting therefore, exceeding overage is not a diagnostic statistic.

games in both plan conditions participants exhibited normative continuation behavior in 99% of all accepted calls.

In contrast, participants were far less skilled at knowing when to *decline* the low- and medium-value calls that should be rejected (especially early on in a game). As shown in Figure 4c (5c) for the 20-second (40-second) allowance condition, participants successfully rejected or terminated only 17% (14%) of all calls that should have been declined, including both new and ongoing ones in game one. With experience decision-making abilities increased, but learning was far from complete, with 62% (53%) of correct rejections and terminations for the 20-second (40-second) plan condition in game six. Participants in aggregate seem to do worse in the larger allowance plan even after experience.

4.1.3 Model Results

While the descriptive results show that participants were able to learn to adopt more efficient call-acceptance policies with experience in the task (albeit to a limit), they are silent on the *source* of the learning; specifically, whether it might have been due to participants being more far-sighted as the task went on, or becoming more consistent in the application of an otherwise far-sighted policy. We can explore this by examining the across-game properties of the two key parameters of our two-segment behavioral-structural model that we fit to participants' behavior in the task: implied discount rates and scale parameters. In Figures 6 and 7 we plot the posterior mean and 95% posterior interval of the parameters obtained from our MCMC implementation.

Insert Figure 6 and Figure 7 here

Figure 6 shows the game-by-game variation for the model parameters for the two segments (that we label as “error-prone” and “learner” respectively) for the $W=20$ plan. The across-game variation of the parameters for the error-prone segment does not show any trend

while that for the learner segment clearly shows a pattern. For the latter, the discount parameter (which determines the level of far-sightedness) increases over games while the scale parameter (which determines the level of errors) decreases over games. We test for consumer learning by examining whether the difference between two structural parameters in adjacent games includes 0 in the 95% posterior interval, as shown in Table 2 for the scale parameter of the learner segment.

Insert Table 2 here

The improvements are not statistically significant beyond game five. The pattern in Figure 6 and Table 2 are both indicative of consumer learning – consumers in the learner segment are becoming more far-sighted and are reducing their errors. The learner segment constitutes 84% of the subjects (with the posterior probability of segment membership practically 100% indicating clear separation into latent classes). In Figure 7, we show that the larger allowance plan ($W = 40$) also features a learner segment (Segment 2, 67% of the subjects). Notably, the scale parameter (λ) under the larger allowance plan is significantly larger than that under the smaller allowance plan for each of the games except the first one (for the learner segment), suggesting that subjects learn to reduce errors to a greater extent in the smaller allowance plan.

Note that under both plans, a residual amount of error remains after six games of experience even in the learner segment, since some subjects do not learn to eliminate errors. Though even the error-prone segment learns to reduce errors, residual error variance is significantly higher than the learner segment in both conditions. Further, consumer learning does not continue beyond game five (see Table 2).

4.2 Study 2 – Imposed Plan Change

4.2.1 Motivation and Procedure

In natural settings consumers are often faced with changes in plans, either due to voluntary choices or policy changes by imposed by their service provider. There is anecdotal evidence that plan changes have at least a short-term effect in degrading the efficiency of calling decisions as heuristic policies learned under the previous regime may have little applicability to the new one. For example, consider a consumer who has become skilled at making call decisions under a plan with a low free allowance who then transitions to a plan with a larger allowance. If the consumer exhibits rule carry-over (e.g., Restle 1962), s/he would, at least for a while, maintain higher standards for call quality as if s/he were in a low-allowance regime, to avoid incurring overage charges for which there was little real risk. On the other hand, work by Lichtenstein and Fischhoff (1977) and Alba and Hutchinson (2000) suggest that we might see the opposite bias by virtue of what they term the “hard-easy” effect: the larger allowance plan is objectively a harder regime since it involves a larger set of policies to ponder over prior to overage, yet the consumer may misconstrue the regime as one where heuristics that insure vigilance are no longer required, resulting in lower standards for the quality of accepted calls. The “hard-easy” effect would predict worse (better) performance when transitioning to a higher (lower) allowance plan while the rule carry-over effect’s predictions would be in the opposite direction.

The purpose of this study is to examine consumer learning and adaptation to imposed plan change by exposing each subject to one TPT plan for three games and the other for the remaining three games. Participants were 156 members of a subject pool recruited on Amazon Mechanical Turk, an online platform for running tasks which offers access to a stable and diverse subject pool (Mason and Suri 2012). We put in place the following requirements for subjects to qualify for our study: a historical task approval rate greater than or equal to 95%, at least 50 tasks

approved to date, and based in the United States. Subjects who qualified could see our task posted and undertake it by clicking on the URL that led to our web-based study (ran on Mozilla Firefox browser). Subjects were paid \$3 to participate and could earn a bonus (typically between \$0-\$1) based on an exchange rate of US\$1 for every 1000 game dollars.

Subjects were randomly assigned to start with a 40-second allowance TPT plan (cell 1) or a 20-second allowance TPT plan (cell 2), and experienced an imposed change to the other plan after three games. From the original pool, 47 were screened out based on a pre-defined criterion of task comprehension (accepting 33% or fewer of all high value calls across the six games). Of the remaining 109 subjects, 43 were in cell 1 and 66 were in cell 2.⁸

4.2.2 Findings: Overall Performance

Similar to Study 1, participants were skilled in correctly continuing calls once taken with an average of 98% accuracy across all games in both conditions. We plot the relative incidence of call acceptance and rejection decision errors for cell 1 and cell 2 (Figure 8). Participants accept on average 80% of the calls they should rationally take (Figures 8a and 8c) across six games in both cells, and while there is learning in the first two games, we see little evidence of a disruptive effect of plan change on this error type.

In contrast, we observe a tangible and asymmetric effect of plan changes on the kind of decision errors that were most prominent in Study 1: correct rejections of low- and medium-value calls. Participants who first became familiar with a large plan allowance and then were switched to a smaller plan allowance improved decision accuracy after the plan change (Figure 8b). In contrast, participants who first became familiar with a small plan allowance and were then switched to a larger plan allowance performed worse after the plan change, and only

⁸ A two-sample test of proportions shows that the proportions of subjects screened out in each condition are not statistically different (p-value = 0.18).

partially recovered from this deterioration by game six (Figure 8d). As can be seen in Figure 8d, increased acceptance of medium-value calls that should be rejected drives much of the deterioration.

Insert Figure 8 here

4.2.3 Model Results

The results suggest that, similar to Study 1, there are two very different types of subjects. Subjects in the error-prone segment are erratic in their temporal discounting and exhibit a large amount of randomness. Learner subjects, who represent over 60% of the subjects in both cells, exhibit a learning pattern. The latent classes clearly distinguish learners from error-prone subjects as evidenced by the virtually 100% posterior probability of each subject belonging to one of the segments. To examine how the task-change effects were manifested in the structural parameters of the behavioral model, in Figure 9a we plot the estimates for the scenario in which participants in the learner segment started with a 40-second allowance and were then switched to a 20-second allowance.

Insert Figure 9 here

Subjects in the learner segment (63% of subjects) become more far-sighted from game one to game two and from game two to game three. The learning effect is much more pronounced from game one to game two for the discount parameter (β). However, learning asymptotes after game three, and does not seem affected by plan change. For the scale parameter (λ), learning is significant between games 2 and 3, and games 3 and 4 (see Table 3). This reduction in errors when switched to the lower allowance plan largely explains the aggregate effect in Figure 9a.

Insert Table 3 here

For the case where participants started with a 20-second allowance and were then switched to a 40-second allowance, we again see evidence of two different segments. 60.7% of subjects consist of learners who quickly learn to become far-sighted while reducing noisy deviations within the first three games (Figure 9b). We find that even these learners become noisier when shifted to the higher allowance plan (starting from game four), but recover after some experience in the new plan. Thus there is clear asymmetry in the learning behavior across the two experimental conditions for subjects in the learner segment, in the direction suggested by the hard-easy effect.

4.3 Study 3 – Performance under volitional plan choice

4.3.1 Motivation and Procedure

One potential limitation of the findings of incomplete learning in Study 1 and the asymmetric effect of plan switches in Study 2 is that both were uncovered in settings where plan were exogenously imposed on decision makers. In natural settings, consumers can self-select into the plans that best fit their own needs and decision skills. Hence, if participants are aware of their differential abilities to make efficient calling decisions under different plan types, we might expect both higher achievement overall and the asymmetric biases observed in Studies 1 and 2 to vanish—or at least be mollified—under plan choice.

To test this, we recruited a new panel of 140 participants from Amazon Mechanical Turk, using the same recruitment criteria as study 2. In this new study, participants were initially offered a choice of a 40-second allowance plan at a fixed fee of \$80 or a 20-second allowance plan at a fixed fee of \$20, after playing a practice game. Participants were given the opportunity to change plans after three games. The fixed fees were chosen such that the expected surplus for a consumer who does not discount future call values is about the same across the two plans.

From the original pool, 33 were screened out based on a pre-defined criterion of task comprehension (accepting 33% or fewer of all high-value calls across the six games). Of the remaining 107 subjects, 29 started in the 40-second allowance plan and 78 started in the 20-second allowance plan. 12 out of 29 subjects starting in the 40-second plan switched to the 20-second plan after three games, whereas 32 out of 78 starting in the 20-second plan switched to the 40-second plan after three games.

4.3.2 Findings: Overall Performance

In Figure 10 we plot the key measure of performance—percentage of calls that were incorrectly pursued when in fact optimal to reject—for each of the four possible plan-choice combinations: those who started with the 40-second plan and stayed with it (N=17), those who stayed with the 40 and switched to a 20 (N=12), those who started with 20-second plan and stayed (N=46), and those who started with a 20 and switched to a more generous plan (N=32). The other performance measures (correct acceptance and continuance of calls) were similar to the corresponding imposed plan conditions in Study 1 and 2.

Insert Figure 10 here

By comparing the panels of Figure 10 with those from Studies 1 and 2 (Figures 5, 8 and 9) we see an immediate answer to the question of whether having free choice helped: it did not. Similar to Study 1, those who stuck to the same plan throughout displayed steady learning (Figures 10a, 10b). Likewise, the self-selection data show an asymmetric effect of learning for those who switched plans. As shown in Figure 10c, switchers from the higher to lower plan allowance learn at a gentle rate (the solid line) while correct rejection of low- and medium-valued calls improves mildly. In contrast, switchers from the lower to higher plan allowance

perform worse (solid line in Figure 10d), which is especially driven by increased usage allocation to medium-valued calls.

4.3.3 Model Results

For the cases where participants switched plans - either from 40 to 20 seconds or 20 to 40 seconds, we find that learning appears to come more from participants reducing response errors (λ) than from overcoming myopia (β), in line with the earlier studies. We find that participants significantly improve learning after choosing a lower allowance while those choosing a higher allowance plan regress to some extent⁹ as shown in Table 4.

Insert Table 4 here

4.4 Study 4 – Effect of risk attitude on discounting

Note that in our task there is no uncertainty in current period utility. There is, however, uncertainty in the distribution of future payoffs, which raises the possibility that our estimates of time discounting may also be capturing individual differences in risk aversion. Specifically, respondents may appear to be giving more weight to immediate (and certain) call values less because they are myopic and more because they prefer certain over uncertain call values. To test this we recruited a sample (N=41) of participants to perform the 40-second calling task described in Study 1, and also took measures of their risk attitudes using the lottery task described by Holt and Laury (2002). Details of the lottery task are in Web Appendix G. We test whether this risk measure (R_i) moderates each individual's discount factor (β_{ig}) by incorporating it as a covariate in our structural model and allowing slopes (ω_s) to be latent class-specific (i.e. $\beta_{ig} = function(\beta_{gs}, \omega_s, R_i)$).

⁹ Though we observe an increase in λ when participants choose a higher allowance plan, the effect is not significant perhaps due to the lower sample size resulting from self-selection in this condition.

The mean raw risk aversion score of the 41 participants was 6.22 (the risk-neutral score is 4) indicating that subjects on average tended towards being risk-averse in line with Holt and Laury (2002). We subtract 4 from the raw score to yield the risk attitude covariate (higher than zero indicates risk aversion, less than zero indicates risk seeking).

We find that the error-prone segment has a risk attitude slope that is not statistically different from zero while the learner segment has a slope that is positive (0.11) whose 95% posterior interval excludes zero. The result suggests individuals who are more risk averse as measured by the Holt and Laury (2002) task were, surprisingly, also more forward-looking in our simulation task. Whether this result might generalize to other settings is uncertain, as there have been relatively few prior attempts to examine the relationship between risk and intertemporal discounting (e.g., Andersen et al. 2008; Coble and Lusk 2010). As the learning curves for discounting and random error propensities remain similar to those in Study 1, this result suggests that our model is capturing intertemporal preference effects that go beyond the effects of relative risk aversion, which appear empirically small.

4.5 Robustness Checks

In this section, we describe additional studies and analyses that test the robustness of our findings from the main set of studies. First, decision errors might be reduced if participants were given more information about their usage of time in the task—something the Federal Communications Commission has negotiated with telecom companies to implement (Jiang 2013). To test usage-awareness effects, we modified the 40-second plan condition of Study 1 such that participants were provided with either a pop-up alert when they used up 80% of their free allowance (call alert study), or an alert and visible real-time counter of talk time used (usage display study). Participants learned more quickly early on to achieve a higher correct call rejection percentage of

29% (35%) in the first three games of the call alert (usage display) study as compared to the Study 1 40-second plan condition (23% average over first three games). However, participants achieve approximately the same performance in the final three games (original study: 49%, call alert study: 46%, usage display study: 47%) indicating that these interventions do not lift overall performance after participants build task experience.

Second, we reduced the size of the fixed fee (from \$80 to \$20) in the 40-second plan condition of Study 1 to understand if the larger up-front fee of the 40-second plan relative to the 20-second plan may be driving our results, since participants may be driven by a need to use as many seconds as possible to justify the larger fixed fee (a sunk cost effect akin to Arkes and Blumer 1985). The reduced fixed-fee study causes a slight lift in correct call rejection performance (58% average in final three games) but does not markedly change the learning curve compared to the corresponding Study 1 condition.

Finally, we tested for curvature effects in current period utility using Study 1 data by running a model that included a quadratic net current utility (call value less any cost) term and allowed the slope for the quadratic term to vary by latent class membership.¹⁰ We find that the slope coefficient is very close to the lower bound of zero for both segments and the DIC for this model is worse than the original model for the 40-second plan (17,210 vs 17,200), reflecting that the two additional curvature slope parameters did not improve fit. This result makes sense in our setting where call utilities are explicitly given to subjects (hence no uncertainty in current period utility), who have an incentive proportional to maximizing net utility.

5. Discussion

¹⁰ We restrict this slope to be between zero and an upper bound computed to ensure that the resulting utility function is concave and increasing. The slope is identified because we set the slope of the linear net utility term to 1.

The purpose of this paper was to investigate the ability of individuals to make optimal call-usage decisions under multi-part pricing schemes. Our findings are based on data drawn from a series of controlled laboratory experiments in which participants made a series of decisions whether to accept simulated incoming phone calls in a setting for which there was a known optimal usage policy. Although the laboratory environment lacked the ecological realism that would have come from studying call behavior in the field (e.g., Grubb 2009), it had the advantage of allowing us to draw strong inferences about intuitive optimality that would be otherwise difficult to achieve. For example, by providing participants with the distribution of arriving calls we were able to tease apart variance in usage behavior that accrued to varying rational expectations about future calls and how participants traded off the present versus future value of those calls.

One of our major findings was that mean usage behavior approached that which would be optimal under different plans given repeated experience, as evidenced by increasing far-sightedness and reduced errors in responses over time. However, we also found learning to be incomplete, with participants still making the error of incorrectly accepting too many lower-valued calls despite multiple rounds of experience. The experiments also revealed an asymmetric bias in achievement when participants were forcibly switched between plans that varied in the amount of free talk time. Specifically, transitioning from a large to small allowance plan accelerated learning whereas the converse resulted in regression. The direction of this asymmetric effect is in line with the hard-easy effect (Lichtenstein and Fischhoff 1977) rather than a uniform tendency to carry over policies learned in one context to another (Restle 1962). Perhaps most surprisingly, these results proved to be robust to such manipulations as allowing participants to choose their own plans, providing warnings when their free talk time was almost used up, and by making talk time visible, as well as to risk preferences. We also find systematic

evidence for a small but highly error-prone segment whose learning asymptotes at much higher levels of error variance compared to the “learners.”

In line with past research that has investigated how plan characteristics impact consumption (e.g., Ascarza et al 2012; Leider and Sahin 2014) we find that increasing free allowance results in overconsumption of lower-valued calls that ought to be rejected. Our work suggests this is due to increased errors under plans with large “free” allowances, rather than a utility bias for “free” minutes or sub-optimal threshold policies. Other work has considered how consumers’ inattention to past usage may impact their future consumption under multi-part tariffs (e.g., Grubb 2014; Grubb and Osborne 2013; Jiang 2013). For instance, Jiang (2013) and Grubb and Osborne (2013) cite evidence of a lack of bunching at points when the free allowance is used up and attribute this to consumer inattention to usage. Interestingly, we find little evidence that inattention drives sub-optimal usage allocation in our experimental setting. Consumers accept too many lower-valued calls and deplete their free allowance, therefore resulting in overage when they take the higher-valued calls that they should rationally accept.

Subjects in the “learner” segment who did not largely eliminate errors fell prey to a flat maximum around the objective function as shown in Table 5, as the lift from eliminating game six errors is about 5% for Study 1 conditions. Cognitive miserliness can explain this result as additional effort to improve performance has diminishing payoffs for learner subjects. Yet, error-prone subjects could have substantially improved earnings relative to learners by updating their decision-making strategies with game feedback. In Study 1, subjects in the 20-second (40-second) plan allowance who were classified as error-prone earned an average bonus of \$0.27 (\$0.60) compared to an average of \$1.57 (\$1.69) for those in the learner segment. One possibility is that these subjects simply could not exercise the discipline needed to avoid lower-valued

calls—these types of consumers may be especially prone to incurring overage charges and racking up higher monthly bills (whether it comes as a shock or not).

Insert Table 5 here

While observed in a laboratory setting, we suggest that the findings may aid recent attempts to inform telecom regulators (such as the Federal Communications Commission) seeking to enhance consumer welfare under multi-part tariffs (e.g., Grubb and Osborne 2013; Jiang 2013). To date most actions that have been introduced to help consumers avoid “bill shock” have been based on the assumption that usage errors accrue primarily due to usage inattention. If true, providing consumers with more complete information about usage (such as limit alerts) should allow them to properly balance short-term and long-term costs when making usage decisions. One of the key findings of our work, however, is such interventions may be less successful than widely presumed; the usage biases we observed were robust to usage-limit alerts and posting elapsed usage. The reason is that the errors we observed were *not* primarily caused by inattention, but rather by participants being overly myopic when first exposed to a plan and then making too many errors later—two sources not remedied by making them more informed about time. Knowing this, it might be more useful to communicate projected usage costs in the future rather than how much talk time has been used. Likewise, our data also show that consumers cannot be counted on to be aware of their own biases and choose the plans to which they are cognitively best suited, as allowing free choice of plans also did little to mitigate the pattern of errors.

But while there may be no simple fix either for usage or plan choice errors, the positive message here is that, with experience these errors diminish; hence, in some cases the best consumer healing tool may just be time itself. In other cases where errors may not diminish

quickly (such as for the error-prone segment), more detailed interventions may be needed. For instance, providing a planning tool that projects usage and overage costs at the current run-rate may draw more attention to the future consequences of current decisions. Tools such as intelligent virtual assistants (e.g., Siri in the Apple iOS system) may be able to learn user preferences and biases over time and provide recommendations when it may help the user.

One of the limitations of our work is that our data were drawn from behavior in an experimental task that simplified the real-world problem of cell-phone usage. In our setting participants were not tempted to make outgoing calls, call value and duration were known upon arrival and uncorrelated, and there was no data usage of the phone. Given these simplifications, the view of consumer optimality that we offered may, if anything, be an optimistic one relative to that which we might observe in a more complex task setting. In field settings, consumers may have significantly lower attention due to the longer time span of the billing cycle and distractions due to other activities. Further, the distribution of call utilities is likely much more complex than the design choices we implemented, which can slow down learning. As a result, the value of alerts and other interventions in the field may have benefits beyond what we demonstrate in our experiments.

Future experimental work can probe the boundaries of the biases uncovered here by changing the task design. On one hand, task difficulty could be increased by providing a larger set of call values (and probabilities), extending the calling task over a longer time period, and adding more distractions to mimic real-world task attention spans, all of which would be expected to exacerbate decision errors. On the other hand, learning effects may be stronger if participants played a larger set of games – though they would almost certainly need breaks during the task to avoid mental depletion.

In this same light, researchers can explore the effects uncovered in this paper using field data by combining data sources, whether through consumer “call diaries,” recalled distribution of calls on a suitable importance scale, or by exploiting natural variation in the data. Using either natural experiments (e.g., Yao et al. 2012), or imposing exclusion restrictions to separate current and future payoffs (e.g., Chung et al. 2014) can help pin down discount rates, which are essential to understand temporal usage tradeoffs, along with patterns of learning and error variance. Akin to Dubé et al. (2013), our approach demonstrates that experiments can help with identification of structural parameters that are challenging to recover using field data. Augmenting field data with experimental data could yield further insights on consumer primitives.

We hope that our work in the cell-phone consumption context encourages future research in several other related domains such as diet management (e.g., Guthrie et al 1995), consumption of goods using credit cards (e.g., Gross and Souleles 2002), the use of banking products (e.g., Stango and Zinman 2009), and health insurance (e.g., Aron-Dine et al. 2012). Understanding what assumptions typically used in empirical models hold in terms of underlying dynamic usage behavior can better aid both researchers and practitioners in their attempts to better anticipate the consequences of policy actions on consumer behavior and welfare.

References

- Alba, J. W., & Hutchinson, J. W. (2000). Knowledge calibration: What consumers know and what they think they know. *Journal of Consumer Research*, 27(2), 123-156.
- Andersen, S., Harrison, G. W., Lau, M. I., & Rutström, E. E. (2008). Eliciting risk and time preferences. *Econometrica*, 76(3), 583-618.
- Arcidiacono, P., & Ellickson, P. B. (2011). Practical methods for estimation of dynamic discrete choice models. *Annu. Rev. Econ.*, 3(1), 363-394.
- Arkes, H. R., & Blumer, C. (1985). The psychology of sunk cost. *Organizational behavior and human decision processes*, 35(1), 124-140.
- Aron-Dine, A., Einav, L., Finkelstein, A., & Cullen, M. R. (2012). Moral hazard in health insurance: How important is forward looking behavior? (No. w17802). National Bureau of Economic Research.
- Ascarza, E., Lambrecht, A., & Vilcassim, N. (2012). When talk is “free”: The effect of tariff structure on usage under two- and three-part tariffs. *Journal of Marketing Research*, 49, 882-899
- Ater, I., & Landsman, V. (2013). Do customers learn from experience? Evidence from retail banking. *Management Science*, 59(9), 2019-2035
- Bar-Gill, O., & Stone, R. (2009). Mobile misperceptions. *Harv. JL & Tech.*, 23, 49.
- Bagh, A., & Bhargava, H. K. (2013). How to price discriminate when tariff size matters. *Marketing Science*, 32(1), 111-126.
- Chung, D. J., Steenburgh, T. & Sudhir K. (2014). Do bonuses enhance sales productivity? A dynamic structural analysis of bonus-based compensation plans. *Marketing Science*, 33(2), 165-187.
- Coble, K. H., & Lusk, J. L. (2010). At the nexus of risk and time preferences: An experimental investigation. *Journal of Risk and Uncertainty*, 41(1), 67-79.
- Dubé, J-P., Hitsch, G.J., & Jindal, P. (2013). The joint identification of utility and discount functions from stated choice data: An application to durable goods adoption, *Quantitative Marketing and Economics (forthcoming)*
- Erdem, T., & Keane, M. P. (1996). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Science*, 15(1), 1-20.
- Goettler, R.L. & Clay, K. (2011). Tariff choice with consumer learning and switching costs. *Journal of Marketing Research*, 48, 633-652

- Gross, D. B., & Souleles, N. S. (2002). Do liquidity constraints and interest rates matter for consumer behavior? Evidence from credit card data. *The Quarterly Journal of Economics*, 117(1), 149-185.
- Grubb, M.D. (2009). Selling to overconfident consumers. *American Economic Review*, 99(5), 1770-1807.
- Grubb, M.D. (2014). Consumer inattention and bill-shock regulation. *Review of Economic Studies* (forthcoming)
- Grubb, M.D & Osborne, M. (2013). Cellular service demand: Biased beliefs, learning, and bill shock. *American Economic Review* (forthcoming)
- Guthrie, J. F., Fox, J. J., Cleveland, L. E., & Welsh, S. (1995). Who uses nutrition labeling, and what effects does label use have on diet quality?. *Journal of Nutrition education*, 27(4), 163-172.
- Holt, C. A., & Laury, S. K. (2002). Risk aversion and incentive effects. *American Economic Review*, 92(5), 1644-1655.
- Iyengar, R., Jedidi, K., & Kohli, R. (2008). A conjoint approach to multipart pricing. *Journal of Marketing Research*, 45(2), 195-210.
- Jiang, L. (2013). The welfare effects of “bill shock” regulation in mobile telecommunications markets. *Working paper*
- Kamakura, W. A., & Russell, G. J. (1989). A probabilistic choice model for market segmentation and elasticity structure. *Journal of Marketing Research*, 379-390.
- Kunreuther, H., Silvasi, G., Bradlow, E. T., & Small, D. (2009). Bayesian analysis of deterministic and stochastic prisoner’s dilemma games. *Judgment and Decision Making*, 4(5), 363-384.
- Lambrecht, A., Seim, K., & Skiera, B. (2007). Does uncertainty matter? Consumer behavior under three-part tariffs. *Marketing Science*, 26(5), 698-710.
- Lambrecht, A., & Skiera, B. (2006). Paying too much and being happy about it: Existence, causes and consequences of tariff-choice biases. *Journal of Marketing Research*, 43, 212-223.
- Leider, S., & Sahin, O. (2014). Contracts, biases and consumption of access services. *Management Science* (forthcoming).
- Lichtenstein, S., & Fischhoff, B. (1977). Do those who know more also know more about how much they know? *Organizational Behavior and Human Performance*, 20(2), 159-183.

- Meyer, R.J., & Hutchinson, J.W. (2001). Bumbling geniuses: The power of everyday reasoning in multi-stage decision making. *Wharton on Making Decisions*, New York: John Wiley and Sons, 37-61
- McKelvey, R. D., & Palfrey, T. R. (1995). Quantal response equilibria for normal form games. *Games and Economic Behavior*, 10(1), 6-38.
- Mason, W., & Suri, S. (2012). Conducting behavioral research on Amazon's Mechanical Turk. *Behavior Research Methods*, 44(1), 1-23.
- Miravete, E. J. (2003). Choosing the wrong calling plan? Ignorance and Learning. *The American Economic Review*, 93(1), 297-310.
- Newton, M. A., & Raftery, A. E. (1994). Approximate Bayesian inference with the weighted likelihood bootstrap. *Journal of the Royal Statistical Society: Series B*, 3-48.
- Oren, S., Smith, S., & Wilson, R. (1985). Capacity pricing. *Econometrica*, 545-566.
- Restle, F. (1962). The selection of strategies in cue learning. *Psychological Review*, 69(4), 329.
- Rust, J. (1987). Optimal replacement of GMC bus engines: An empirical model of Harold Zurcher. *Econometrica*, 999-1033.
- Salisbury, L. and Feinberg, F. (2010). Alleviating the constant stochastic variance assumption in decision research: Theory, measurement, and experimental test. *Marketing Science*, 29(1), 1-17.
- Spiegelhalter, D. J., Best, N. G., Carlin, B. P., & Van Der Linde, A. (2002). Bayesian measures of model complexity and fit. *Journal of the Royal Statistical Society: Series B*, 64(4), 583-639.
- Stango, V., & Zinman, J. (2009). What do consumers really pay on their checking and credit card accounts? Explicit, implicit, and avoidable costs. *The American Economic Review*, 99(2), 424-429.
- Yao, S., Mela, C.F., Chiang, J., & Chen, Y. (2012). Determining consumers' discount rates with field studies. *Journal of Marketing Research*, 49, 822-841.
- Young, P.H. (2009). Learning by trial and error. *Games and Economic Behavior* (65), 626-643

TABLES AND FIGURES

| Statistics across participants over the set of six games | Expected call arrivals | 20-second plan allowance | 40-second plan allowance |
|--|------------------------|--------------------------|--------------------------|
| Total calls | 144 | 145 (12) | 145 (11) |
| Total low- and medium-valued calls | 108 | 110 (12) | 107 (10) |
| Number of accepted low-and medium- valued calls | n/a | 16 (17) | 23 (15) |
| Total high-valued calls | 36 | 35 (6) | 38 (6) |
| Number of accepted high value calls | n/a | 32 (7) | 35 (8) |

Table 1: Descriptive statistics for Study 1 – the values shown are averages across the six games played by each participant, with across-participant standard deviations in parentheses.

| Conditions (imposed) | $\lambda_{2L} - \lambda_{1L}$ | $\lambda_{3L} - \lambda_{2L}$ | $\lambda_{4L} - \lambda_{3L}$ | $\lambda_{5L} - \lambda_{4L}$ | $\lambda_{6L} - \lambda_{5L}$ |
|------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|
| Study 1, 20-second allowance | (-6.2,-4.5) | (-1.7,-0.9) | (-0.7,-0.3) | (-0.3,0.1) | (-0.3,0.0) |
| Study 1, 40-second allowance | (-2.0,-0.0) | (-2.4,-1.3) | (-1.0,-0.0) | (-0.8,-0.1) | (-0.4,0.1) |

Table 2: Study 1 - 95% posterior interval for difference in learner segment's error scale parameter between game pairs. Learning is indicated if the posterior interval excludes non-negative values (in bold).

| Conditions (imposed) | $\lambda_{2L} - \lambda_{1L}$ | $\lambda_{3L} - \lambda_{2L}$ | $\lambda_{4L} - \lambda_{3L}$ | $\lambda_{5L} - \lambda_{4L}$ | $\lambda_{6L} - \lambda_{5L}$ |
|--------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|
| Study 2, Low to high allowance | (-5.2,-2.9) | (-0.7,-0.2) | (0.1,0.6) | (-0.5,0.0) | (-0.3,0.1) |
| Study 2, High to low allowance | (-1.4,0.1) | (-1.9,-0.9) | (-1.1,-0.4) | (-0.3,0.1) | (-0.3,0.1) |

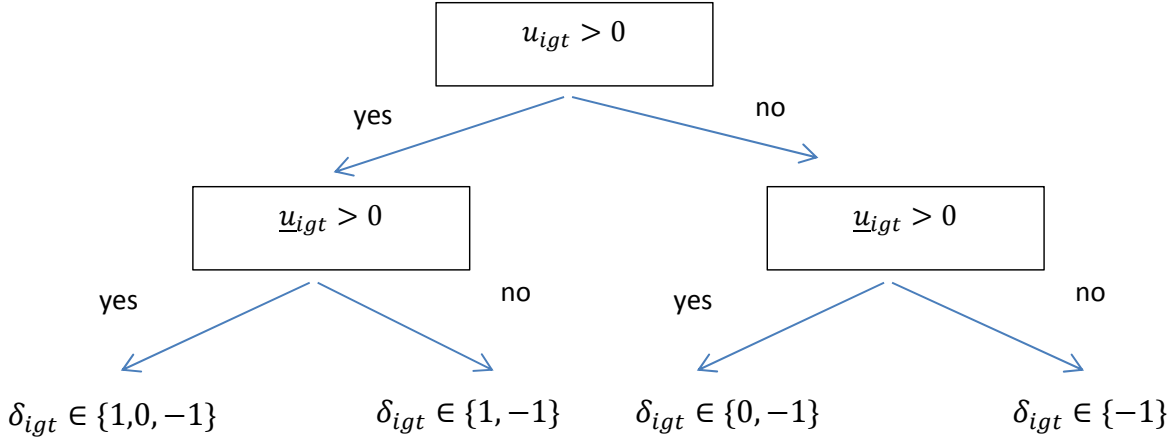
Table 3: Study 2 - 95% posterior interval for difference in learner segment's error scale parameter between game pairs. Learning or deterioration is indicated if the posterior interval excludes zero (in bold). The shaded cells indicate transition due to plan change.

| Conditions (self-selected) | $\lambda_{2L} - \lambda_{1L}$ | $\lambda_{3L} - \lambda_{2L}$ | $\lambda_{4L} - \lambda_{3L}$ | $\lambda_{5L} - \lambda_{4L}$ | $\lambda_{6L} - \lambda_{5L}$ |
|--------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|-------------------------------|
| Study 3, Low to high allowance | (-2.4,-0.9) | (-1.0,-0.4) | (-0.0,0.6) | (-0.3,0.3) | (-0.5,0.0) |
| Study 3, High to low allowance | (-3.7,-1.1) | (-0.8,0.5) | (-1.6,-0.7) | (-0.4,0.3) | (-0.4,0.3) |

Table 4: Study 1 - 95% posterior interval for difference in learner segment's error scale parameter between game pairs. Learning or deterioration is indicated if the posterior interval excludes zero (in bold). The shaded cells indicate transition due to plan change.

| | Game 1 | Game 2 | Game 3 | Game 4 | Game 5 | Game 6 |
|---------------------|--------|--------|--------|--------|--------|--------|
| 20-second allowance | 254% | 42% | 20% | 7% | 6% | 4% |
| 40-second allowance | 195% | 54% | 20% | 12% | 9% | 5% |

Table 5: Lift in average bonus for learner segment by using error-free optimal policy with no discounting of future call values for Study 1.



$$K(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}, \delta_{igt}) = \begin{cases} [u_{igt} - c_{ig}I(w_{igt} \geq \bar{W}_{ig})]I(u_{igt} > 0), & \delta_{igt} = 1 \\ [\underline{u}_{igt} - c_{ig}I(w_{igt} \geq \bar{W}_{ig})]I(\underline{u}_{igt} > 0)I(D_{igt} > 0), & \delta_{igt} = 0 \\ 0, & \delta_{igt} = -1 \end{cases}$$

Figure 1: Current period utility K as a function of state variables and decision. For a respondent i , game g and time t , u_{igt} is the value of an incoming call, \underline{u}_{igt} is the value of an ongoing call, D_{igt} is the number of seconds remaining on an existing call, w_{igt} is talk time used and δ_{igt} denotes the decision being new call acceptance ($\delta_{igt} = 1$), existing call continuation ($\delta_{igt} = 0$) or termination or rejection of a call ($\delta_{igt} = -1$).




Figure 2: Earliest time (in seconds) for a consumer following the optimal policy to accept call with given utility under each plan when (a) $\beta = 1, \bar{W} = 40$, (b) $\beta = 1, \bar{W} = 20$. Low-value call decisions are represented by the solid line while medium-value call decisions are represented by the dotted line.

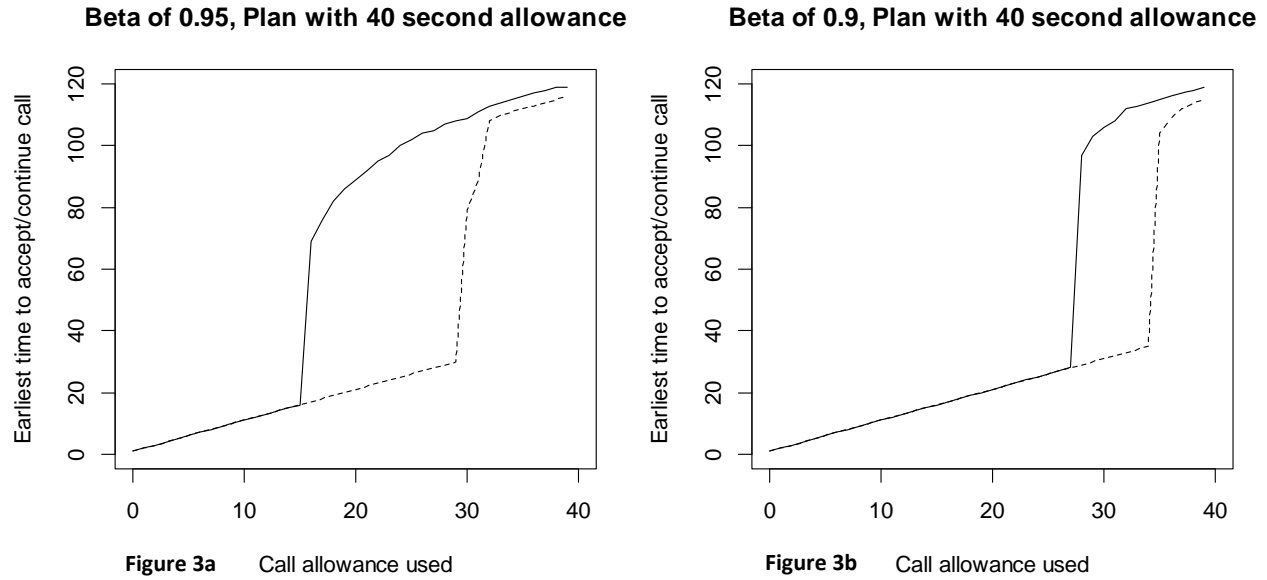


Figure 3: Earliest time (in seconds) for a consumer following the optimal policy to accept call with given utility under each plan when (a) $\beta = 0.95, \bar{W} = 40$, (b) $\beta = 0.90, \bar{W} = 40$. Low-value call decisions are represented by the solid line while medium-value call decisions are represented by the dotted line. Note that it is never optimal to accept low- or medium-valued calls once the call allowance is fully used up in our experimental design.

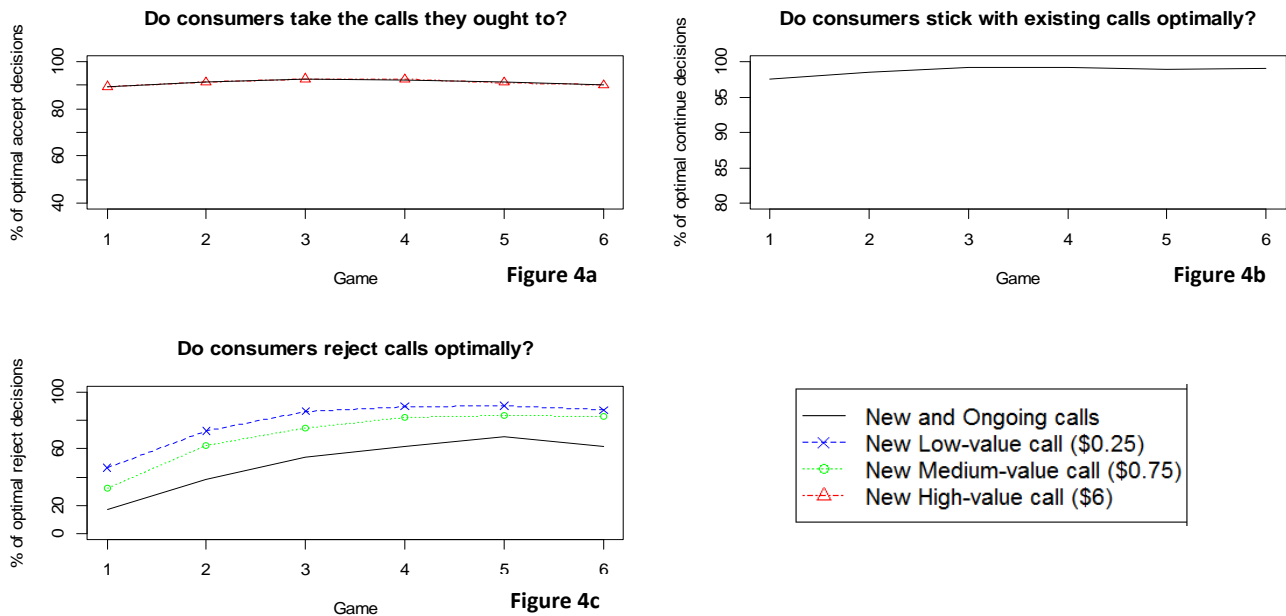


Figure 4: Study 1 (20-second allowance). Most of the calls that should be accepted are high-value ones, hence the overall trajectory is virtually identical to the high-value call acceptance trajectory (top left graph). Subjects learn to reject new low- and medium-value calls with experience but performance asymptotes after three games (bottom left graph). The overall call rejection trajectory is lower than for new call rejection since this includes incorrect continuation of ongoing calls that should be rejected.

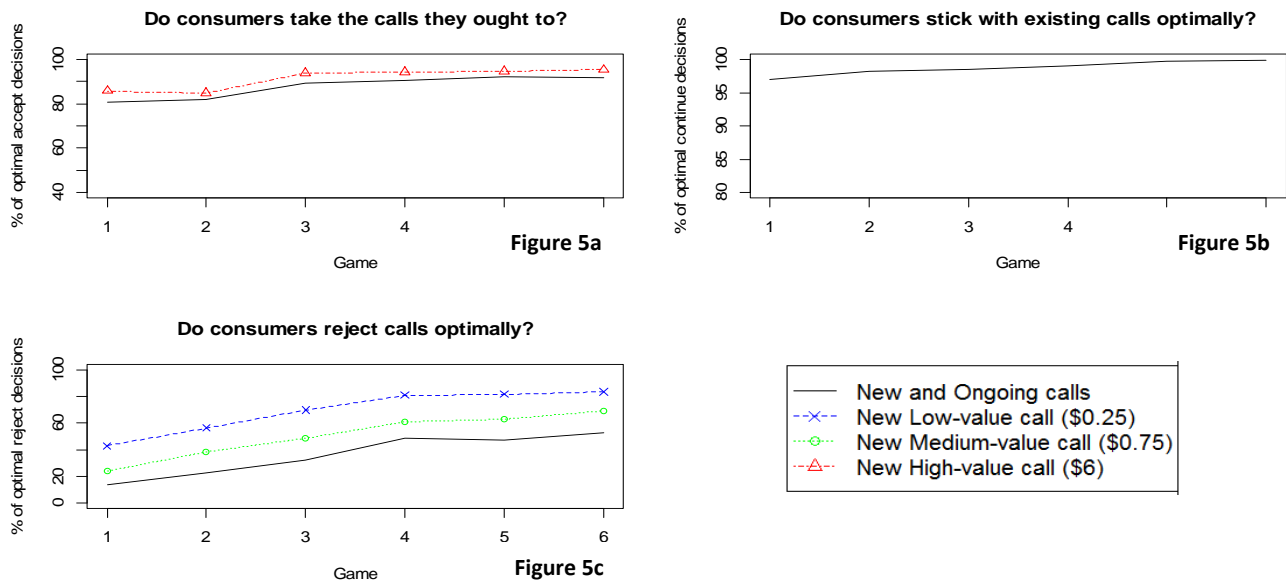


Figure 5: Study 1 (40-second allowance). Similar to Figure 4, learning to reject new low- and medium-value calls asymptotes after three to four games.

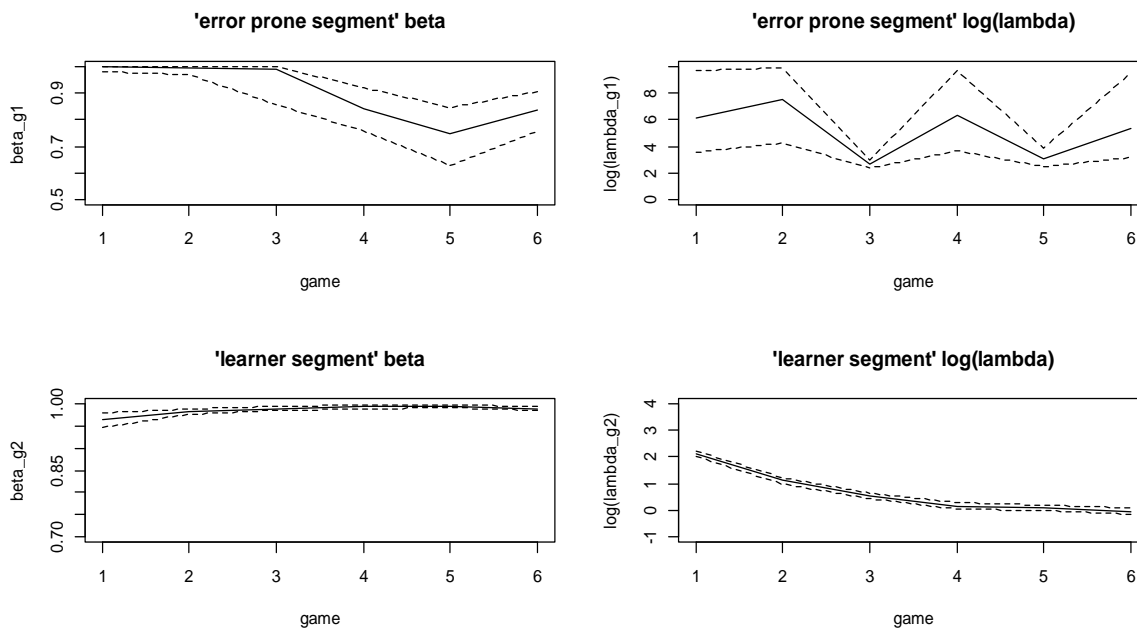


Figure 6: Study 1: Posterior means and 95% credible interval of parameters by segment and game for study 1's 20-second allowance plan. Note that intervals are wider for large values of log(lambda) and low values of beta, since identification from the data is more challenging for such values. Each graph has a different y-axis range to emphasize variations of interest. Learners improve mostly in the first 3 games.




Figure 7: Study 1: Posterior means and 95% credible interval of parameters by segment and game for study 1's 40-second allowance plan. Each graph has a different y-axis range to emphasize variations of interest. Learners improve mostly in the first 3 games.

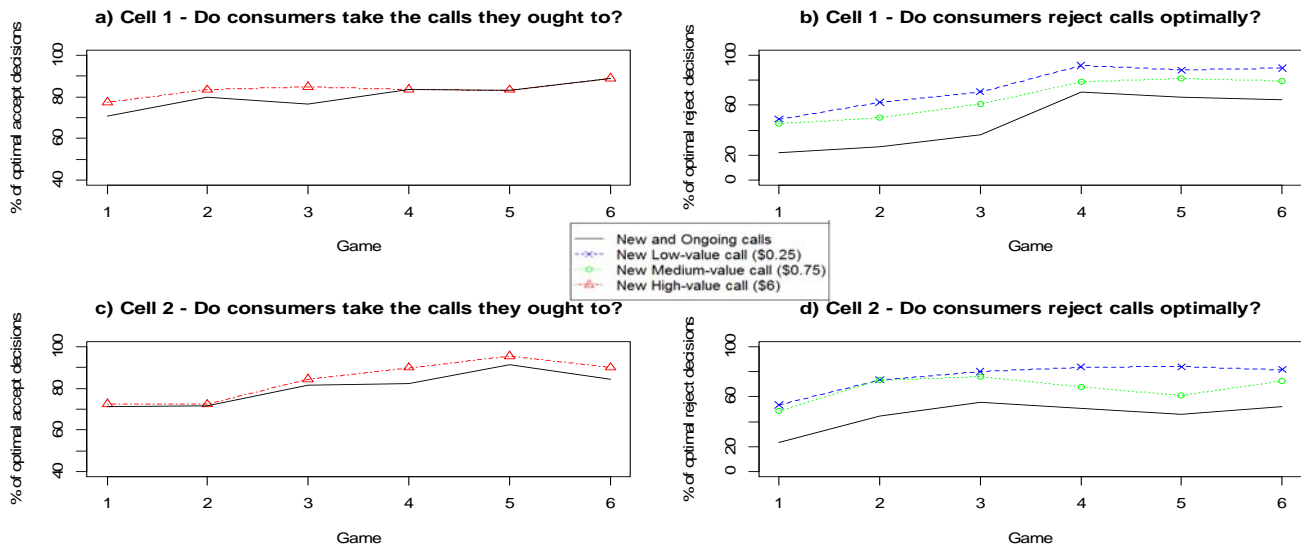


Figure 8: Study 2: Cell 1 (40-second allowance for first three games, 20-second allowance for last three games) and cell 2 (20-second allowance for first three games, 40-second allowance for last three games). Subjects correctly reject *more* low- and medium-value calls after transitioning to the lower allowance plan. Subjects correctly reject *fewer* medium-value calls after transitioning to the higher allowance plan.




Figure 9: Study 2: Posterior means and 95% credible interval of parameters for learner segment by game for (a) study 2's imposed 40 to 20-second plan transition, and (b) study 2's imposed 20 to 40-second plan transition. The dotted vertical line demarcates the plan transition point.

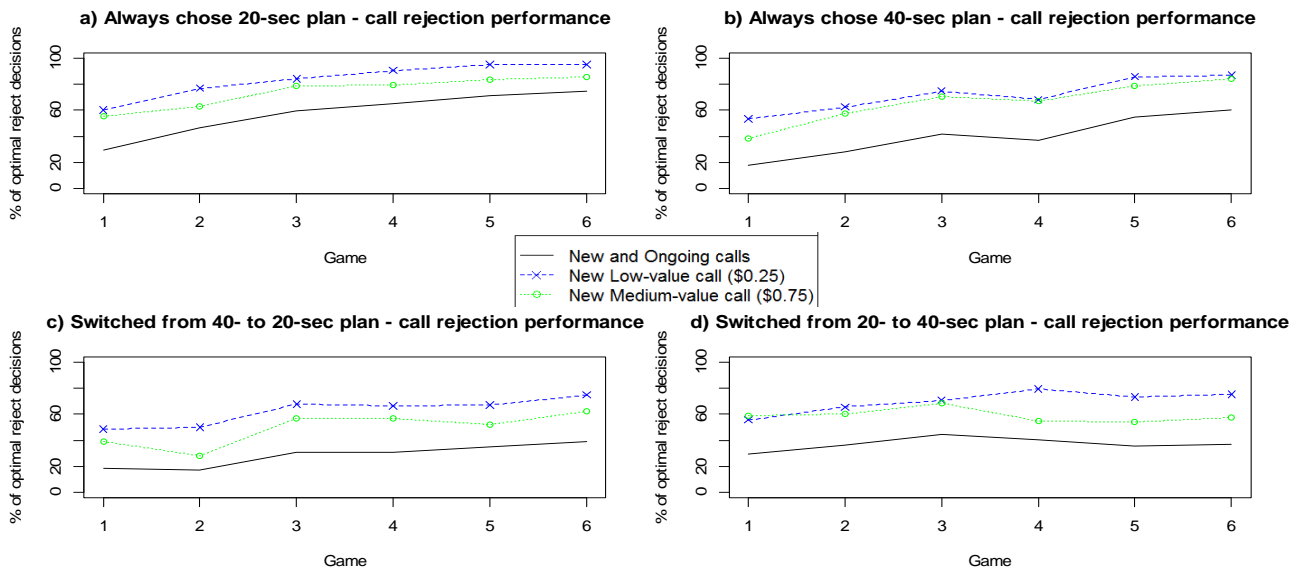


Figure 10: Study 3: Call rejection performance for all four self-selected cells that participants could end up in, based on their first and second plan choice.

Web Appendix

accompanying

Consumer Dynamic Usage Allocation and Learning under
Multi-part Tariffs

Arun Gopalakrishnan

Raghuram Iyengar

Robert J. Meyer

Appendix A: Propositions relating to optimal policy and threshold rules

In this Appendix, we first characterize the properties of the value function, dropping the subscripts for individual i and game g for notational convenience.

Proposition 1: The expected continuation value function is non-decreasing in \underline{u}_t at given values of $1 \leq D_t < L$ and $0 \leq w_t < t$.

$$\text{For } t = T, \bar{V}_T(\underline{u}'_T, D_T, w_T) = \int \max\{u_T - \text{cl}(w_T \geq \bar{W}), \underline{u}'_T - \text{cl}(w_T \geq \bar{W}), 0\} f(u_T) du_T$$

$$\text{and } \bar{V}_T(\underline{u}''_T, D_T, w_T) = \int \max\{u_T - \text{cl}(w_T \geq \bar{W}), \underline{u}''_T - \text{cl}(w_T \geq \bar{W}), 0\} f(u_T) du_T$$

If $\underline{u}''_T \geq \underline{u}'_T$, $\max\{u_T - \text{cl}(w_T \geq \bar{W}), \underline{u}''_T - \text{cl}(w_T \geq \bar{W}), 0\} \geq \max\{u_T - \text{cl}(w_T \geq \bar{W}), \underline{u}'_T - \text{cl}(w_T \geq \bar{W}), 0\} \forall u_T$.

Hence $\bar{V}_T(\underline{u}''_T, D_T, w_T) \geq \bar{V}_T(\underline{u}'_T, D_T, w_T)$.

Assume $\bar{V}_{t+1}(\underline{u}''_{t+1}, D_{t+1}, w_{t+1}) \geq \bar{V}_{t+1}(\underline{u}'_{t+1}, D_{t+1}, w_{t+1})$ whenever $\underline{u}''_{t+1} \geq \underline{u}'_{t+1}$.

$$\text{Given } \bar{V}_t(\underline{u}_t, D_t, w_t) = \int \max\{u_t - \text{cl}(w_t \geq \bar{W}) + \beta \cdot \bar{V}_{t+1}(\underline{u}_t, L - 1, w_t + 1), \underline{u}_t - \text{cl}(w_t \geq \bar{W}) + \beta \cdot \bar{V}_{t+1}(\underline{u}_t, D_t - 1, w_t + 1), \beta \cdot \bar{V}_{t+1}(0, 0, w_t)\} f(u_t) du_t,$$

$\bar{V}_t(\underline{u}''_t, D_t, w_t) \geq \bar{V}_t(\underline{u}'_t, D_t, w_t)$ for $\underline{u}''_t \geq \underline{u}'_t$, since $\underline{u}_t - \text{cl}(w_t \geq \bar{W}) + \beta \cdot \bar{V}_{t+1}(\underline{u}_t, D_t - 1, w_t + 1)$ is a non-decreasing function of \underline{u}_t (by assumption for \bar{V}_{t+1} term and obvious for \underline{u}_t linear term). By induction, the proposition holds true $\forall t$.

Proposition 2: The expected continuation value function is non-increasing in w_t at given values of \underline{u}_t and $1 \leq D_t < L$.

$$\text{For } t = T, \bar{V}_T(\underline{u}_T, D_T, w'_T) = \int \max\{u_T - \text{cl}(w'_T \geq \bar{W}), \underline{u}_T - \text{cl}(w'_T \geq \bar{W}), 0\} f(u_T) du_T$$

$$\text{and } \bar{V}_T(\underline{u}_T, D_T, w''_T) = \int \max\{u_T - \text{cl}(w''_T \geq \bar{W}), \underline{u}_T - \text{cl}(w''_T \geq \bar{W}), 0\} f(u_T) du_T$$

If $w''_T \geq w'_T$, $\max\{u_T - \text{cl}(w'_T \geq \bar{W}), \underline{u}_T - \text{cl}(w'_T \geq \bar{W}), 0\} \leq \max\{u_T - \text{cl}(w''_T \geq \bar{W}), \underline{u}_T - \text{cl}(w''_T \geq \bar{W}), 0\} \forall u_T$, with equality holding when both w'_T and w''_T are strictly less or greater than \bar{W} .

Hence $\bar{V}_T(\underline{u}_T, D_T, w''_T) \leq \bar{V}_T(\underline{u}_T, D_T, w'_T)$.

Assume $\bar{V}_{t+1}(\underline{u}_{t+1}, D_{t+1}, w''_{t+1}) \leq \bar{V}_{t+1}(\underline{u}_{t+1}, D_{t+1}, w'_{t+1})$ whenever $w''_{t+1} \geq w'_{t+1}$.

$$\text{Given } \bar{V}_t(\underline{u}_t, D_t, w_t) = \int \max\{u_t - \text{cl}(w_t \geq \bar{W}) + \beta \cdot \bar{V}_{t+1}(\underline{u}_t, L - 1, w_t + 1), \underline{u}_t - \text{cl}(w_t \geq \bar{W}) + \beta \cdot \bar{V}_{t+1}(\underline{u}_t, D_t - 1, w_t + 1), \beta \cdot \bar{V}_{t+1}(0, 0, w_t)\} f(u_t) du_t,$$

$\bar{V}_t(\underline{u}_t, D_t, w_t'') \leq \bar{V}_t(\underline{u}_t, D_t, w_t')$ for $w_t'' \geq w_t'$, since $-cI(w_t \geq \bar{W}) + \beta \cdot \bar{V}_{t+1}(\underline{u}_{t+1}, D_{t+1}, w_t + 1)$ is a non-increasing function of w_t (by assumption for \bar{V}_{t+1} term and obvious for $-cI(w_t \geq \bar{W})$ term). By induction, the proposition holds true $\forall t$.

Proposition 3: The expected continuation value function is non-decreasing in D_t at given values of \underline{u}_t and $0 \leq w_t < t$.

$1 \leq D_t < L$ in our setting, since $D_t = 0$ implies no ongoing call, and $D_t \geq L$ is not possible since this implies that the call could exceed the maximum call duration of L .

For $t = T$, $\bar{V}_T(\underline{u}_T, D_T'', w_T) = \bar{V}_T(\underline{u}_T, D_T', w_T)$ for any admissible value of D_T'', D_T' . With only one time period left, it does not matter how many additional periods of talk time are left on an ongoing call as long as there is at least one period. It is therefore also true that $\bar{V}_T(\underline{u}_T, D_T'', w_T) \geq \bar{V}_T(\underline{u}_T, D_T', w_T)$ for $D_T'' \geq D_T'$.

Assume $\bar{V}_{t+1}(\underline{u}_{t+1}, D_{t+1}'', w_{t+1}) \geq \bar{V}_{t+1}(\underline{u}_{t+1}, D_{t+1}', w_{t+1})$ whenever $D_{t+1}'' \geq D_{t+1}'$.

Given $\bar{V}_t(\underline{u}_t, D_t, w_t) = \int \max\{u_t - cI(w_t \geq \bar{W}) + \beta \cdot \bar{V}_{t+1}(\underline{u}_t, L - 1, w_t + 1), \underline{u}_t - cI(w_t \geq \bar{W}) + \beta \cdot \bar{V}_{t+1}(\underline{u}_t, D_t - 1, w_t + 1), \beta \cdot \bar{V}_{t+1}(0, 0, w_t)\} f(u_t) du_t$,

$\bar{V}_t(\underline{u}_t, D_t'', w_t) \geq \bar{V}_t(\underline{u}_t, D_t', w_t)$ since $\bar{V}_{t+1}(\underline{u}_t, D_t - 1, w_t + 1)$ is non-decreasing in D_t by assumption and is the only term that involves D_t . By induction, the proposition holds true $\forall t$.

A consequence of Proposition 1 and 2 is that $\bar{V}_{t+1}(0, 0, w_t) - \bar{V}_{t+1}(\underline{u}_t, D_t - 1, w_t + 1)$, which features in both equations (2) and (3) is not generally greater than or less than zero, since the expected continuation value function $\bar{V}_{t+1}(\underline{u}_t, D_t - 1, w_t + 1)$ has both a weakly higher continuing call utility and greater used talk time.

Proposition 4: Switching to a call of equal or higher utility than an existing call is always better than continuing on the existing call.

Switching makes sense when $u_t + \beta \cdot \bar{V}_{t+1}(\underline{u}_t, L - 1, w_t + 1) \geq \underline{u}_t + \beta \cdot \bar{V}_{t+1}(\underline{u}_t, D_t - 1, w_t + 1)$. If $u_t \geq \underline{u}_t$, and given $L \geq D_t$,

$\bar{V}_{t+1}(\underline{u}_t, L - 1, w_t + 1) \geq \bar{V}_{t+1}(\underline{u}_t, D_t - 1, w_t + 1)$ by Proposition 1 and 3. Hence, switching weakly dominates continuation. Note that it can, in some circumstances, be optimal to switch to a call of *lower* utility than an existing call. Hence, Proposition 4 establishes a sufficient but not necessary condition for switching.

We now describe the threshold rules that apply based on the consumer's decision space at each time period. In equation (1), the value function at time t depends on which of the four decision scenarios manifests in that time period. The continuation value of accepting a new call

$\bar{V}_{ig,t+1}(u_{igt}, L - 1, w_{igt} + 1)$ reflects the extra second of talk time used at time t while having an ongoing call with utility u_{igt} and remaining duration $L - 1$ (since L is the maximum duration of a call in our decision problem). Similarly, the continuation value of staying on an existing call $\bar{V}_{ig,t+1}(\underline{u}_{igt}, D_{igt} - 1, w_{igt} + 1)$ reflects the extra second of talk time used at time t and call progress. If the consumer rejects or terminates a call, or has no decision to make, the continuation value $\bar{V}_{ig,t+1}(0, 0, w_{igt})$ reflects that talk time stays the same and no ongoing call will be available in the next period.

The optimal policy depends on the available decisions. If the only decision is to accept or reject a new call (no ongoing call), a threshold rule applies. As shown in equation (A1), the optimal decision rule ($\delta_{igt,ACC-REJ}^*$) is either 1 (accept) or -1 (reject). The intuition is straightforward: overage cost raises the bar for accepting calls, a higher continuation value for taking a call lowers the bar, and a higher continuation value for call rejection raises the bar. The difference in continuation values is critical to the optimal policy as this weighs the saving of another second of usage against the utility from continuing the call in the future. The effect of the tradeoff on the decision rule depends crucially on the amount of talk time used (w_{igt}).

$$\delta_{igt,ACC-REJ}^*(u_{igt}, w_{igt}) = 2 \cdot I(u_{igt} > c_{ig}I(w_{igt} \geq \bar{W}_{ig}) + \beta_{ig} \cdot \bar{V}_{ig,t+1}(0, 0, w_{igt}) - \beta_{ig} \cdot \bar{V}_{ig,t+1}(u_{igt}, L - 1, w_{igt} + 1)) - 1 \quad (\text{A1})$$

Similarly, another threshold rule applies when the only option is to continue ($\delta_{igt,CON-TER}^* = 0$) or terminate ($\delta_{igt,CON-TER}^* = -1$) an existing call as shown in equation (A2). The intuition is analogous to the new call acceptance decision, except that the utility and remaining duration are based on an existing call.

$$\delta_{igt,CON-TER}^*(\underline{u}_{igt}, D_{igt}, w_{igt}) = I(\underline{u}_{igt} > c_{ig}I(w_{igt} \geq \bar{W}_{ig}) + \beta_{ig} \cdot \bar{V}_{ig,t+1}(0, 0, w_{igt}) - \beta_{ig} \cdot \bar{V}_{ig,t+1}(\underline{u}_{igt}, D_{igt} - 1, w_{igt} + 1)) - 1 \quad (\text{A2})$$

When both a new and existing call are available, the consumer's optimal decision involves two steps. The first step determines whether switching to the new call ($\delta_{igt,SWITCH}^* = 1$) or staying with the existing call ($\delta_{igt,SWITCH}^* = 0$) is more attractive (equation A3). The more attractive option is then weighed against the continuation value of declining all calls, resulting in $\delta_{igt,SWITCH-CON-TER}^*$ being a composite of the decision rules in equations (A1), (A2), and (A3).

$$\delta_{igt,SWITCH}^*(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}) = I(u_{igt} + \beta_{ig} \cdot \bar{V}_{ig,t+1}(u_{igt}, L - 1, w_{igt} + 1) \geq \underline{u}_{igt} + \beta_{ig} \cdot \bar{V}_{ig,t+1}(\underline{u}_{igt}, D_{igt} - 1, w_{igt} + 1)) \quad (\text{A3})$$

$$\begin{aligned} \delta_{igt,SWITCH-CON-TER}^*(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}) \\ = \delta_{igt,CON-TER}^*(\underline{u}_{igt}, D_{igt}, w_{igt}) \cdot [1 - \delta_{igt,SWITCH}^*(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt})] + \\ \delta_{igt,ACC-REJ}^*(u_{igt}, w_{igt}) \cdot \delta_{igt,SWITCH}^*(u_{igt}, \underline{u}_{igt}, D_{igt}, w_{igt}) \end{aligned} \quad (\text{A4})$$

Appendix B: Examples of call choice data collected from participants and what in the data identify structural parameters

In Figures B1 and B2, we show illustrations of call arrivals (marked as circular points), usage (indicated by the solid lines – 0 denotes not being on a call at that time either due to lack of calls or deciding against taking/staying on a call), and when usage hits talk time allowance (dotted vertical line) from our first study (in the 40-second talk time condition). The subject in Figure B1, for example, rejects several high-value calls in the first two games and experiences overage for the first time in game three. The subject then allocates usage to all high-value calls but also takes lower-valued calls in the remaining three games (which may or may not be optimal). In contrast, another subject from the same experimental condition (Figure B2) always allocates usage to high-value calls and tends to be conservative with using talk time allowance for lower-valued calls after game one. In game six, this subject turns down 13 consecutive lower-valued calls and starts using talk time close to the 50-second mark.




Figure B1: Illustrations of call arrival and usage patterns for a subject with large errors under a 40-second allowance plan for all six games. The dotted vertical line is when free call allowance was used up (if at all).

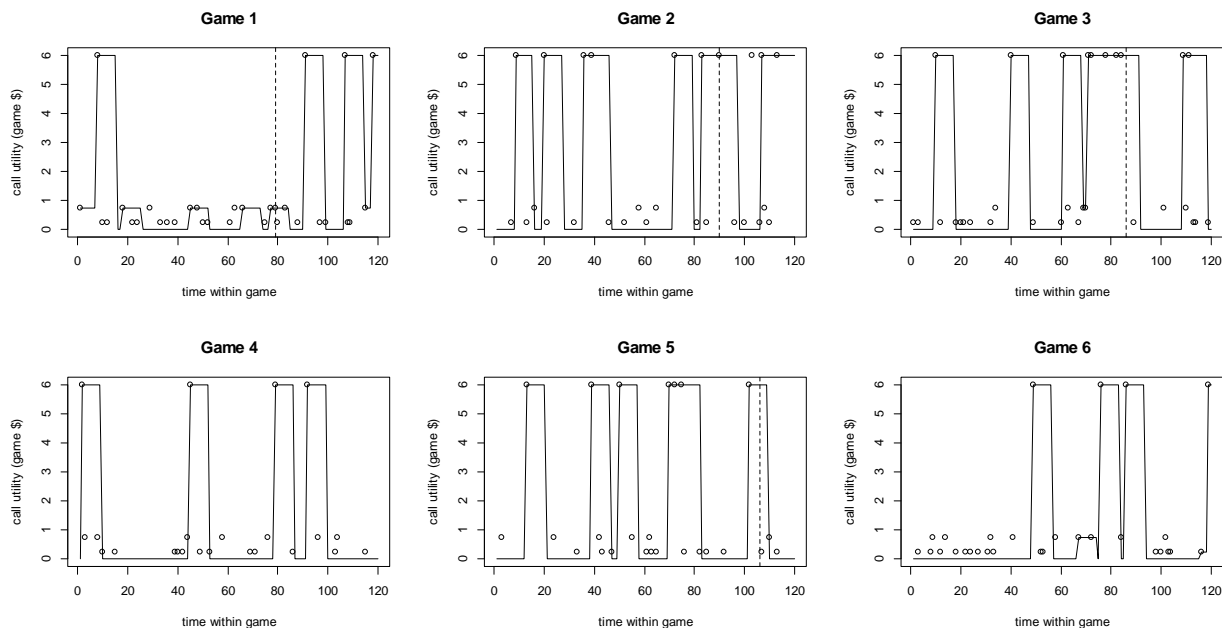


Figure B2: Illustrations of call arrival and usage patterns for a subject who learns to allocate usage to mostly high-valued calls under a 40-second allowance plan for all six games. The dotted vertical line is when free call allowance was used up (if at all).

Though plan and game characteristics have been chosen to allow for empirical recovery of structural parameters, estimator properties also depend on the “signal” contained in each game about an individual’s decision process due to the stochastic nature of call arrivals. Imagine that a person received *only* high-value calls and accepted all of them. This would be consistent with many models of behavior (including a myopic consumer). However, the timing of acceptance or rejection of low- and medium-value calls (e.g., Figures B1 and B2) is informative of how forward looking a person may be. Specifically, if a person did take these lower-valued calls early on, the inflection points along consumed talk time (as illustrated in Figure 3) at which the person stops taking such calls helps identify the discount factor. Rejecting high-value calls is informative of a noisy policy (with large idiosyncratic shocks) as is exhibiting deviations from the optimal policy for any given discount factor (moderate idiosyncratic shocks).

Appendix C: Full set of Screenshots

In this appendix, we reproduce the exact instructions that subjects were exposed to, as they undertook the simulation.

Screen 1: Welcome to the Call Study Simulator

On the next screen you will be introduced to a cell-phone call study in which you will make a series of decisions. Bonus money can be earned based on your performance.

To get started, enter your MTurk worker ID in the field below and press the Start button below. Upon completing the study, which will take 15 to 20 minutes, you will get a code to enter in the MTurk window from which you launched this study, to get paid. No payment will be made if the study is not completed.

Subject id:

Screen 2: Introduction to the calling game

This is a study about phone-call management. You will be playing a game where the objective is to maximize the amount of pleasure you get from calls while minimizing the cost of those calls. It works like this. You will play a game where you have to make a series of decisions about whether to accept or reject a series of incoming calls in a 120-second window. You will play this game 6 times.

At the start of the game, you will be randomly assigned to a payment plan which provides a certain amount of "free" calling time for a fixed up-front charge. If you go over this allowance you are charged by the second for all the calls you take. Unused free seconds do not roll over to subsequent games and the payment plan may change from game to game. Check the plan details which will appear at the top of the screen before playing each game.

There are three types of calls – low-value (\$0.25/sec), medium-value (\$0.75/sec) and high-value (\$6.00/sec). At any given moment, you may receive one of these call types or no call. If you do receive a call, you will be asked if you want to accept or reject the call.

Each call lasts for a maximum of 8 seconds. If you accept a call, you can hang up anytime, or switch to another arriving call.

The probability of receiving a call at any point is 20%. If a call is received, 50% of the time it will be low-valued, 25% of the time it will be medium-valued, and 25% of the time, it will be high-valued.

All amounts shown during the task are in Game dollars, and your accumulated net value after the 6 games will be converted into real US currency using an exchange rate of \$0.10 for every 100 Game dollars. The bonus range would be roughly in the range of US\$0.40 to US\$1, and is contingent on Game dollars accumulated. If your net value is negative, no bonus will be paid.

The image on the next page shows the main features of the game screen.

After the next screen, you will first play a 30 second practice game which does not count towards your net value. You will then start the first of the six games.

Screen 3:

Your calling plan is described here.

Mobile Phone Game #1

Calling Plan: Fixed Fee of \$20, free allowance of 40 secs, \$4/sec overage

Call Values: low (\$0.25/sec), medium (\$0.75/sec), high (\$6.00/sec)

You will play six games in total.

This clock counts down from 2 minutes.

Calls

01:58

Active Call: Value \$6/sec
Time Left: 7 secs

You are either "idle" or on an "Active Call"

Status

Total Value: \$6.00

Display the value earned from your calls so far.

Incoming Call with value of \$0.75/sec
Take this call?

Yes No

You can terminate a call before it is finished with this button.

Incoming calls are presented to you here.

Screen 4:

Mobile Phone Game - Practice Game

Calling Plan: Fixed Fee of \$80, free allowance of 40 secs, \$4.00/sec overage charges.

Call Values: low (\$0.25/sec), medium (\$0.75/sec), high (\$6.00/sec)

| | |
|---|---|
| <p>Calls</p> <div style="border: 1px solid #ccc; background-color: #4a86e8; color: white; padding: 10px; border-radius: 10px;"><p>0:28 Idle</p><p>Incoming Call with value of \$0.25/sec Take this call?</p><p style="text-align: center;">Yes No</p></div> | <p>Status</p> <div style="border: 1px solid #ccc; background-color: #4a86e8; color: white; padding: 10px; border-radius: 10px;"><p>Total Value: \$0.00</p></div> |
|---|---|

Next Game

Screen 5:

The following payment plan will apply for your call usage in game 1:

A fixed fee of **\$80**,

an allowance of **40 free seconds**, and

a **\$4.00 per second** overage charge.

You will be notified if the payment plan changes.
Please check the top of the game screen to verify the payment plan for each game.

Click here to begin next game

Screen 6:

Mobile Phone Game #1

Calling Plan: Fixed Fee of \$80, free allowance of 40 secs, \$4.00/sec overage charges.

Call Values: low (\$0.25/sec), medium (\$0.75/sec), high (\$6.00/sec)

| | |
|---|---|
| <p>Calls</p> <div style="border: 1px solid #ccc; background-color: #4a86e8; color: white; padding: 10px; border-radius: 10px;"><p>1:53 Active Call. Value \$0.25/sec Time Left: 3 secs</p><p>Incoming Call with value of \$6/sec Take this call?</p><p style="text-align: center;">Yes No</p></div> | <p>Status</p> <div style="border: 1px solid #ccc; background-color: #4a86e8; color: white; padding: 10px; border-radius: 10px;"><p>Total Value: \$1.25</p></div> |
|---|---|

Next Game

Appendix D: MCMC Algorithm for 2-segment latent class model

Markov Chain Monte Carlo (MCMC) procedures are used for obtaining posterior distributions of the following parameters: β_{gs} – the discount factor for segment s in game g , λ_{gs} – the scale parameter for segment s in game g , α – the probability of belonging to segment 1, Z_i – the latent class membership for each subject i . We transform the segment and game specific parameters as $\begin{bmatrix} \text{logit}(\beta_{gs}) \\ \log(\lambda_{gs}) \end{bmatrix} = \mu_{gs}$.

The joint distribution of all model parameters is given by:

$$p(\{\mu_{g1}, \mu_{g2} \forall g\}, \alpha, \{Z_i \forall i\} | Y) \\ \propto p(\alpha) \prod_{g=1}^6 p(\mu_{g1}) p(\mu_{g2}) \prod_{i=1}^N p(Z_i | \alpha) \left(\prod_{g=1}^6 p(Y_{ig} | \mu_{g1}) \right)^{I(Z_i=1)} \left(\prod_{g=1}^6 p(Y_{ig} | \mu_{g2}) \right)^{I(Z_i=2)}$$

The following steps will implement a hybrid Gibbs/Metropolis sampler.

1. Draw from full condition distribution of $\mu_{gs} | \{Z_i \forall i\}, Y \forall g, s$:

$p(\mu_{gs}) \prod_{i=1}^N \left(p(Y_{ig} | \mu_{gs}) \right)^{I(Z_i=s)}$ using a Metropolis step since the distribution has no closed form. $p(Y_{ig} | \mu_{gs})$ requires a backward sweep to compute expected value functions that are used to form conditional choice probabilities. $p(\mu_{gs})$ is uniformly distributed in the region where $\text{logit}(\beta_{gs}) \in [-10, 10]$ and $\log(\lambda_{gs}) \in [-10, 10]$ and has zero support outside this region. Due to the logit and log transformations, the range of discount factors covered is $[0.00005, 0.99995]$ and the range of error scale parameters is $[0.00005, 22,026.5]$. Values outside of this range would be extremely unlikely to materialize in this setting.




2. α conditional on $\{Z_i \forall i\}$ is deterministic and is computed as $\frac{\sum_{i=1}^N I(Z_i=1)}{N}$.
3. Draw $Z_i | \{\mu_{g1}, \mu_{g2} \forall g\}, \alpha, Y \sim \text{Bernoulli}\left(\frac{\alpha(\prod_{g=1}^6 p(Y_{ig} | \mu_{g1}))}{\alpha(\prod_{g=1}^6 p(Y_{ig} | \mu_{g1})) + (1-\alpha)(\prod_{g=1}^6 p(Y_{ig} | \mu_{g2}))}\right)$.

To generalize to S segments, steps 2 and 3 simply need to be modified to account for the additional segments. The most computationally intensive process (within step 1) is to calculate the *ex ante* value function at each point in state space and time for use in the conditional choice probabilities. We parallelize this computation for efficiency using Amazon Web Services (AWS) elastic cloud computing resources. We implemented this using the R library “parallel” which allows dynamic creation and allocation of tasks to worker processes spawned as new threads. Computationally intensive tasks were written in the C programming language and linked to the R interface as a shared library. The C and R source code for implementing this algorithm are available upon request from the authors. 100,000 iterations of the algorithm were run for each data set.


Appendix E: Simulation and Parameter Recovery for Model

We test parameter recovery for varying amounts of data for a grid of true $\beta \in \{0.785, 0.907, 0.989\}$ and $\lambda \in \{0.21, 1.01, 2.81, 6.00\}$ under the 40-second plan allowance. This set of parameter values covers the typical range of estimates that we find in our studies. For each true parameter set, we simulate 30 datasets of three, six, and twelve subjects (N) each by generating random call arrivals and applying the decision rule given by the behavioral model for the true parameters for each call decision. For the purposes of this simulation, each dataset is composed of homogenous subjects with one game each. The finite-sample bias is computed by estimating the posterior mean of the parameters for each of the data sets and computing the average deviation from the true parameters. We observe that for $N = 3$, recovery of parameters is more robust for higher true β . $N = 6$ provides a much tighter parameter recovery by comparison and for $N = 12$, recovery is fairly tight around the true parameters. The implication for latent-class analysis is that any class should have at least 6 subjects to allow for robust parameter recovery.


The legend below applies for all plots.




Bias in beta for N = 6




Bias in lambda for N = 6



Bias in beta for N = 12



Bias in lambda for N = 12



Appendix F: Model Comparison Results for Study 1

In Study 1, we estimate a set of models for the 20-second and 40-second plan allowance conditions to characterize fit. We find that a 2-segment latent class model allowing for learning in both β (discount factor) and λ (scale parameter representing deviations from optimal behavior) fits well compared to a homogenous model or 2-segment models in which learning is turned off for either parameter. While increasing the number of latent classes beyond two segments does improve fit, the largest improvement occurs when comparing a homogenous model to one that assumes two segments. The model fit results for the other studies are similar. Since our goal is to characterize the effects of three-part tariffs on overall task performance, we use 2 latent classes for ease of comparison across studies, and these capture much of the distinct differences between two highly heterogeneous sub-groups.

| | 1 segment | 2 segments | 3 segments | 4 segments |
|-------------------------------|-----------|------------|------------|------------|
| Total Parameters | 12 | 25 | 37 | 51 |
| Log Marginal Likelihood (LML) | -6,264 | -4,967 | -4,568 | -4,283 |
| DIC | 12,528 | 9,930 | 9,142 | 8,574 |

Table F1: Study 1 20-second plan allowance condition (N = 44) – comparing fit across different numbers of latent classes

| | No Learning on β , No Learning on λ | Learning on β , No Learning on λ | No Learning on β , Learning on λ | Learning on β , Learning on λ |
|-------------------------------|--|---|---|--|
| Total Parameters | 5 | 15 | 15 | 25 |
| Log Marginal Likelihood (LML) | -5,739 | -5,532 | -5,006 | -4,967 |
| DIC | 11,477 | 11,050 | 10,003 | 9,930 |

Table F2: Study 1 20-second plan allowance condition (N = 44) – comparing fit across nested 2-segment models.

| | 1 segment | 2 segments | 3 segments | 4 segments |
|-------------------------------|-----------|------------|------------|------------|
| Total Parameters | 12 | 25 | 37 | 51 |
| Log Marginal Likelihood (LML) | -9,592 | -8,608 | -8,223 | -7,943 |
| DIC | 19,185 | 17,200 | 16,395 | 15,896 |

Table F3: Study 1 40-second plan allowance condition (N = 63) – comparing fit across different numbers of latent classes

| | No Learning on β , No Learning on λ | Learning on β , No Learning on λ | No Learning on β , Learning on λ | Learning on β , Learning on λ |
|-------------------------------|--|---|---|--|
| Total Parameters | 5 | 15 | 15 | 25 |
| Log Marginal Likelihood (LML) | -9,428 | -9,185 | -8,811 | -8,608 |
| DIC | 18,855 | 18,366 | 17,626 | 17,200 |

Table F4: Study 1 40-second plan allowance condition (N = 63) – comparing fit across nested 2-segment models.

Appendix G: Details of Study 4 procedure

In Study 4, we administered the Holt and Laury (2002) lottery task prior to our main cellphone simulation task as shown in the screenshot below.

Part I - Lottery Task

For each of the pairs of choice below, choose the option you prefer (either Option A or B). There is no right or wrong answer.

For example, 5%: \$1.00, 95%: \$0.50 would mean that in a hypothetical lottery, you would have a 5% chance of winning \$1.00 and a 95% chance of winning \$0.50.

Only pick one option per choice. Once you are done with all 10 choices, click the submit button.

Choice 1: 10%: \$2.00, 90%: \$1.60 10%: \$3.85, 90%: \$0.10
Choice 2: 20%: \$2.00, 80%: \$1.60 20%: \$3.85, 80%: \$0.10
Choice 3: 30%: \$2.00, 70%: \$1.60 30%: \$3.85, 70%: \$0.10
Choice 4: 40%: \$2.00, 60%: \$1.60 40%: \$3.85, 60%: \$0.10
Choice 5: 50%: \$2.00, 50%: \$1.60 50%: \$3.85, 50%: \$0.10
Choice 6: 60%: \$2.00, 40%: \$1.60 60%: \$3.85, 40%: \$0.10
Choice 7: 70%: \$2.00, 30%: \$1.60 70%: \$3.85, 30%: \$0.10
Choice 8: 80%: \$2.00, 20%: \$1.60 80%: \$3.85, 20%: \$0.10
Choice 9: 90%: \$2.00, 10%: \$1.60 90%: \$3.85, 10%: \$0.10
Choice 10: 100%: \$2.00, 0%: \$1.60 100%: \$3.85, 0%: \$0.10

For this study, we recruited a new panel of 48 participants from Amazon Mechanical Turk, using the same recruitment criteria as Study 2. From the original pool, 5 were screened out based on a pre-defined criterion of task comprehension (accepting 33% or fewer of all high-value calls across the six games). Of the remaining 43 subjects, two participants' lottery task responses displayed inconsistent preferences. Specifically, one participant accepted option A for choice 10 which is inconsistent with any level of risk aversion since a certain \$3.85 dominates a certain \$2.00. Another participant switched from option A to B, and then back again, which Holt and Laury (2002) also designate as violating the basic premise of the task. These two participants were removed from the dataset to yield a final pool of 41 participants.

A risk neutral participant would pick option A for choices 1 through 4 and switch to option B from choice 5 onwards. We code risk attitude as the number of option A choices less 4 (such that 0 corresponds to risk neutral, > 0 corresponds to risk averse, and < 0 corresponds to risk seeking). On average, subjects are risk averse in line with Holt and Laury (2002).