http://pubsonline.informs.org/journal/isre/

INFORMATION SYSTEMS RESEARCH
Vol. 30, No. 2, June 2019, pp. 616–635
ISSN 1047-7047 (print), ISSN 1526-5536 (online)

Forward-Looking Behavior in Mobile Data Consumption and Targeted Promotion Design: A Dynamic Structural Model

Lizhen Xu,[a] Jason A. Duan,[b] Yu Jeffrey Hu,[a] Yuan Cheng,[c] Yan Zhu[c]

[a] Scheller College of Business, Georgia Institute of Technology, Atlanta, Georgia 30308; [b] McCombs School of Business, University of Texas at Austin, Austin, Texas 78712; [c] School of Economics and Management, Tsinghua University, 100084 Beijing, China

Contact: lizhen.xu@scheller.gatech.edu, https://orcid.org/0000-0002-8603-9891 (LX); duanj@mccombs.utexas.edu, https://orcid.org/0000-0002-3621-7055 (JAD); jeffrey.hu@scheller.gatech.edu, https://orcid.org/0000-0001-8482-4899 (YJH); chengyn@sem.tsinghua.edu.cn (YC); zhuyan@sem.tsinghua.edu.cn (YZ)

Received: July 20, 2015
Revised: April 3, 2017; March 6, 2018
Accepted: July 9, 2018
Published Online in Articles in Advance: June 5, 2019

https://doi.org/10.1287/isre.2018.0820

Abstract. This paper examines the dynamic consumption behavior of individual mobile data users by employing a unique data set on individual-level daily usage over multiple months. Whether and which individual mobile data users are forward looking by dynamically balancing present and future usage and how to design profitable promotions targeting these users are questions of both academic and managerial interest. By developing a dynamic structural model and formally proving its theoretical properties, we discover distinct temporal usage patterns that can distinguish forward-looking users from myopic ones. An empirical test is constructed to test for individual forward-looking behavior by matching the observed usage patterns with the theoretical results. We find a considerable proportion of users (about 40%) are indeed forward looking and also find empirical evidence of individual consumer myopia. Our approach enables us to apply the dynamic model only to those exhibiting forward-looking behavior. It hence serves as a feasible option to control for consumer myopia in estimating dynamic structural models given the inherent limitation that individual discount factors are generally unidentifiable. Our structural model is shown to accurately capture the dynamic trends observed in the actual usage data. It enables sophisticated counterfactual simulations incorporating various factors (e.g., consumer anticipation, plan switch) to deliver rich implications for targeted promotion design. As we find, promotions targeting only forward-looking consumers could be significantly more profitable than blanket promotions uniformly applied to all. Properly designed end-of-month promotions targeting forward-looking users could help mobile carriers fully utilize the otherwise excess network bandwidth and increase revenue at little extra cost.

History: Il-Horn Hann, Senior Editor; Prasanna Tambe, Associate Editor; Alok Gupta, Editor-in-Chief.
Funding: L. Xu acknowledges the support from the National Natural Science Foundation of China [Grants 71572029 and 71671027]. Y.J. Hu and Y. Cheng acknowledge the support from the National Natural Science Foundation of China [Grant 71628102].
Supplemental Material: The online appendix is available at https://doi.org/10.1287/isre.2018.0820.

Keywords: mobile data • dynamic consumption pattern • forward-looking behavior • consumer myopia • targeted promotion • individual heterogeneity

1. Introduction

As consumers worldwide increasingly access digital content via the mobile broadband on their smart devices, mobile carriers around the globe have been actively offering various finely designed mobile data plans to consumers (e.g., Segan 2014). Such mobile data plans typically adopt a three-part tariff price structure, which consists of a prespecified data quota in a billing cycle for a fixed fee and a unit rate for extra usage exceeding the quota. Under such a price structure, as various activities ranging from web browsing to multimedia streaming can quickly exhaust the allotted quota, users need to carefully manage their data usage over time within a billing cycle in order to avoid excessive surcharges while achieving the highest utility from their plan quota. This reflects the so-called *forward-looking* consumption behavior. As economic theories posit, consumers ought to be forward looking and take into account future utilities when determining their amount of mobile data usage for the present day. As a result, consumers' daily consumption, rather than maximizing the present utility alone (i.e., consuming as needed), needs to balance between the present utilities (i.e., consuming now) and the future ones (i.e., saving for the future). This leads to dynamic consumption strategies that constantly adjust the daily usage based on the amount of quota remaining and the number of days left. However, common wisdom suggests that not all consumers are forward looking and some can be *myopic* in their consumption strategies (e.g., Urminsky and Zauberman 2015). These consumers simply consume as needed every day without considering future days, resulting in static consumption patterns independent of the amount of quota or the number of days remaining. Therefore, *whether* and *which* individual consumers are forward looking in mobile data consumption are interesting questions of both managerial and academic significance.

Identifying individual consumers with forward-looking consumption behavior is highly valuable to business practitioners. First, forward-looking consumers are likely to be more price sensitive and responsive to various promotions, whereas myopic consumers are likely less so, which our empirical findings will show is indeed the case. Therefore, designing promotions targeting forward-looking consumers only could yield significant revenue gains compared with the common industry practice of sending blanket promotions uniformly applied to all consumers. Second, in the context of mobile data consumption, because forward-looking users typically reduce their data usage when the remaining quota diminishes near the end of a billing cycle (usually a month), total mobile data usage volume tapers off, leaving excess bandwidth capacity within the mobile networks. If mobile carriers could properly design end-of-month promotions targeting forward-looking users to boost usage, they could fully utilize the available network bandwidth and increase revenue at little extra cost. Therefore, the ability to identify individual forward-looking consumers and design targeted promotions especially for these consumers can open up enormous revenue opportunities for mobile carriers and many other industries alike.

It is also of academic interest and importance to deepen the understanding of individual forward-looking consumption behavior. Whether consumers are forward looking in various contexts has long been a question with important business and policy implications intriguing researchers. Previous studies using aggregate data have found empirical evidence that consumers *collectively* are forward looking in purchasing durable goods such as textbooks (Chevalier and Goolsbee 2009) and cars (Busse et al. 2013). On the other hand, behavioral theories suggest that some consumers may display myopic behavior in making consumption decisions (e.g., Hoch and Loewenstein 1991, Frederick et al. 2002, Worthy et al. 2012, and Urminsky and Zauberman 2015)). Theoretical models often explicitly consider some consumers as forward looking and others as myopic (e.g., Pashardes 1986 and Gabaix and Laibson 2006). However, there is a lack of empirical studies that formally test for forward-looking consumption behavior at the *individual* consumer level along with theoretical results regarding the associated temporal consumption patterns. Questions such as *how many* and *which* consumers are forward looking and what *temporal consumption patterns* can enable us to identify these individuals remain largely unanswered. We aim to fill this gap by studying the dynamic consumption behavior of individual consumers in the emerging context of mobile data usage. As data at the individual consumer level with sufficient observations both cross-sectionally and longitudinally are becoming increasingly available in this big-data era, we believe that the theories and methods developed in this paper will also generate a wide range of applications and implications in various contexts.

Being able to distinguish individual forward-looking consumption behaviors from myopic ones can also help mitigate an inherent limitation in the estimation of dynamic structural models. In dynamic structural models, where consumer behaviors in observed data are typically modeled based on a utility-maximization framework, whether consumers are forward looking or not would essentially be reflected in the value of a model parameter known as the *discount factor*, which quantifies the present value of future utilities. A successful estimation of individual-level discount factors, if possible, would answer all the questions pertinent to the forward-looking behavior of individual consumers. Unfortunately, it is well known, in both theory and application, that the discount factor is deeply confounded with other utility parameters and cannot be identified by observational data in general (Rust 1987, Manski 1993). The common approach in the literature is to assume that all individuals are forward looking and pick a fixed value as the common discount factor for all individuals (e.g., Rust 1987, Misra and Nair 2011, and Huang et al. 2015). Although some recent studies attempt to estimate a common discount factor for all individuals utilizing special structures in their data, as we discuss in more detail in Section 2, it has not been possible to estimate individual-specific discount factors along with all the other utility parameters based on observational data only. In the face of such an inherent obstacle in dynamic structural model estimation, if we can identify forward-looking consumers based on certain dynamic consumption patterns consistent with the structural model, we will be able to apply the dynamic structural model to these users only rather than the entire sample. In this way, we advance the current practice by avoiding assuming all consumers as forward looking and hence controlling for possible consumer myopia, which largely mitigates this notorious limitation in estimating dynamic structural models.

Motivated by these considerations, in this paper we seek to address two sets of unanswered questions with regard to the forward-looking consumption behavior of mobile data users. First, is there any notable dynamic pattern in mobile data consumption that enables us to identify evident individual forward-looking behavior? We aim to answer this question with both theory and evidence: We first seek to establish a solid theoretical foundation rooted in dynamic economics to discover such a pattern, upon which we develop an empirical test for individual forward-looking behavior and apply it to real data. The findings can thus shed light on whether and which individual consumers are indeed forward looking in mobile data consumption. Second, being able to identify forward-looking consumers, how can mobile carriers design effective promotions targeting these consumers, and how can such promotions help increase revenue? To appropriately answer this question calls for a well-developed structural model that properly captures the dynamic consumption decisions of mobile data users. The model needs to account for the apparent heterogeneity among users in various aspects such as usage needs and price sensitivity. The structural parameters capturing these aspects are invariant to policy changes and can thus deliver meaningful counterfactual implications for various proposed promotion designs.

Employing a unique data set on the daily mobile data usage over a nine-month period for a large number of subscribers from a leading mobile carrier in China, we are able to construct and estimate a dynamic structural model for daily mobile data usage. We model an individual user's daily data usage in a framework of maximizing the sum of present and expected future utilities facing day-to-day random utility shocks. Hence, our dynamic structural model captures the intertemporal substitution between current and future consumption. Taking advantage of the richness of our individual-level and multiple-month daily usage data, we fully capture the heterogeneity among users by specifying the model at the individual level so that the model parameters are user specific.

To identify individual forward-looking behavior, we first derive and prove formal analytical results from our structural model, which uncover distinct usage patterns for forward-looking and myopic users. As we show through rigorous proofs, forward-looking users' daily usage is positively correlated with the remaining data quota, whereas myopic users' usage is uncorrelated with the remaining quota. Such distinct patterns enable us to develop a reduced-form empirical test for forward-looking behavior of individual users by matching their observed usage patterns and the theoretical results. We apply the test to each individual user in our data sample and find that about 40% of the users exhibit forward-looking usage patterns, whereas the rest demonstrate myopic patterns.

We then apply our dynamic structural model to these forward-looking users. In estimating the model, we adopt a two-stage estimation approach similar to Bajari et al. (2007). The estimation method is computationally efficient and can be easily parallelized, which facilitates speedy estimation for hundreds of users in our sample and opens the feasibility for a larger-scale implementation by mobile carriers. By examining the relationship between estimated structural parameters and observed consumer demographics, we find that older users generally have lower price sensitivity and less daily variation in usage. To validate the estimation results, we develop a simulation method by solving the dynamic programming problem given the estimated parameter values. Simulated usage well captures the dynamic trends observed in the actual data, further validating our model development.

To demonstrate how our model can help design various promotions targeting forward-looking users, we conduct counterfactual simulations to evaluate the design of two types of promotions: discounting the price rate for overage and offering unlimited-use passes with a fixed fee in the last several days of a month. We find a stark contrast in profitability between offering price discounts to forward-looking users and offering them to myopic users. An optimal discount of 40%–50% in the overage charge targeting forward-looking users can increase the carrier's revenue by about 27%, whereas offering discounts to users with myopic usage patterns can only hurt the revenue simply because a large number of myopic users are not price sensitive and do not respond to the price reduction actively. We further show that blanket promotion uniformly applied to all users does not help improve the carrier's revenue because the revenue gain from forward-looking users are offset by the revenue loss from myopic users. This result highlights the revenue potential of targeted promotions focusing on the segment of forward-looking users, which makes our proposed method especially valuable for business practice. We also find that selling unlimited-use passes at a proper price level and at the right time (e.g., a ¥10 pass for the last four days of a month) can increase the expected revenue by as much as 20%. As we further show, even when users may anticipate the coming of the promotion offers and hence adjust their usage behavior or even plan choice in the first place, the promotion effects change little, and the relative magnitude remains robust.

In addition to the managerial implications, our study makes several contributions on the methodological fronts. First, our analytical results establish a theoretical understanding of the temporal patterns associated with the class of dynamic consumption problem as we study. Theoretically uncovering such patterns with formal proofs is nontrivial given the complex nature of dynamic problems and has been absent in the literature. Second, the empirical test we develop, along with its theoretical foundation, offers a practical and reliable tool for identifying individual forward-looking behavior that is useful to both future researchers and practitioners. To the best of our knowledge, we are the first in the literature to conduct such a formal test at the individual consumer level and to provide evidence for individual forward-looking behavior. Third, our empirical findings also show that not all consumers are forward looking and thus underscore the necessity of controlling for possible consumer myopia in dynamic model estimation. Our approach, which first derives the analytical properties of a dynamic structural model followed by the reduced-form test for individual forward-looking behavior, serves as a novel and feasible method for this purpose. It enables us to apply the dynamic structural model only to those individuals exhibiting behaviors consistent with the dynamic model and to avoid the strong assumption that all consumers are forward looking. It therefore helps mitigate the inherent limitation in dynamic model estimation that individual discount factors cannot be identified in general.

The rest of the paper is organized as follows. In the next section, we briefly review the related literature and discuss the distinction of our work. We describe our data in Section 3 and lay out our dynamic structural model in Section 4. We test for individual forward-looking behavior in Section 5 by first deriving formal theoretical results and then developing the empirical test based on these results and applying it to our data. In Section 6, we discuss the estimation and simulation methods for our model, followed by the estimation results. We conduct counterfactual simulations of various promotion strategies to provide policy implications in Section 7. Section 8 concludes the paper.

2. Literature Review

In this section, we briefly review the literature on traditional wireless phone services and consumer usage of mobile data and smartphone apps, which are relevant to the substantive context of our study. We also discuss the literature on dynamic structural models involving consumer forward-looking behavior, followed by how our study contributes to this stream of literature.

Previous studies have examined different aspects of traditional wireless phone services, including voice call service and short message service (SMS). For example, Iyengar et al. (2007) study consumer learning over months in choosing their voice call plans. Kim et al. (2010) develop a static structural model to study the substitution between voice and SMS demands. Yao et al. (2012) estimate consumers' weekly discount rate in using voice call service. Recently, increasing research interest has been attracted to consumer behavior related to mobile data and smartphone apps. For example, Ghose and Han (2011) investigate the correlation between uploading and downloading for users using the mobile internet. Niculescu and Whang (2012) examine the codiffusion process of the adoption of wireless voice and mobile data services in Japan. Xu et al. (2014) explore the complementary effect between the introduction of a mobile app and website visits for news media. By contrast, our study differs from these papers in the fundamentals, such as research context (i.e., mobile data consumption), research methodology (i.e., dynamic structural model), and research focus (i.e., dynamic patterns of daily usage and targeted promotion design).

Dynamic structural models and their estimation have long been applied to a large variety of economics and business problems involving forward-looking agents—including, for example, bus engine replacement (Rust 1987), sales-force compensation (e.g., Misra and Nair 2011and Chung et al. 2014), online grocery service (e.g., Goettler and Clay 2011), and social media contribution (e.g., Huang et al. 2015). In estimating dynamic structural models, however, an inherent limitation persists: the discount factor, a key structural parameter characterizing forward-looking behavior, is notoriously difficult to identify.

Theoretically, it is shown that the discount factor in dynamic structural models generally cannot be separately identified from other structural parameters by observational data (Manski 1993, Magnac and Thesmar 2002). Theoretical results on strict identification conditions for discount factors have yet to be established. Magnac and Thesmar (2002, proposition 4) theoretically derive an "exclusion restriction" condition in the discrete choice context, which is in general not easily verifiable or economically interpretable in practice (Yao et al. 2012, Abbring and Daljord 2017). Recent results (Abbring and Daljord 2017) show that the discount factor may be identifiable only up to a finite set of distinct values under a more economically interpretable exclusion condition, which is still not generally satisfied in applications. In applied studies with observational data, the common approach is to assume all individuals are forward looking and set a fixed value as the common discount factor for all of them (e.g., Rust 1987, Misra and Nair 2011, and Huang et al. 2015). Some studies attempt to identify the discount factor by imposing restrictions on the model structure. For example, Goettler and Clay (2011) assume homogeneity across all consumers with different pricing structures, which enables the identification of the common price sensitivity through cross-sectional variation and thus the common discount factor through intertemporal variation. A few recent studies manage to identify a common discount factor for all individuals utilizing special structures in their data. For example, Yao et al. (2012) use a data set spanning before and after a natural experiment where there is a shift from linear pricing to nonlinear pricing; Chung et al. (2014) exploit particular sales-force compensation structures to find variables that do not affect current utility but affect future utility so as to form exclusion restrictions. Nevertheless, even under these situations, it is still possible that a common discount factor cannot be reliably estimated in practice. For example, despite the identifiability by their data in theory, Chung et al. (2014, p. 178) find in practice that the objective function is "relatively flat with respect to changes in the discount factor," and hence they have to rely on grid search to find the most appropriate discount factor instead of the usual estimation methods. Therefore, it is not surprising that successful estimation of individual-specific discount factors along with all utility parameters using observational data only has not been possible in the existing literature.

In this regard, the approach we propose, in which theoretical results are first derived from the structural model and a reduced-form test is then developed based on these results to identify individual forward-looking behavior, contributes to this stream of literature by offering a novel and feasible method to control for possible consumer myopia in estimating dynamic structural models. Our approach shares the same spirit of previous nonstructural studies that develop empirical tests for collective forward-looking behavior using aggregate data (e.g., Chevalier and Goolsbee 2009). To the best of our knowledge, we are the first to conduct formal tests for consumer forward-looking behavior at the individual level and to provide empirical evidence for both forward-looking and myopic behavior of individual consumers.

Another distinction of our study from the existing literature on dynamic structural analysis is that our paper is among the few that perform individual-specific analysis. Panel data with sufficient observations longitudinally are not often available, and the computational burden for dynamic models is typically high. Therefore, most existing studies of dynamic models usually specify and estimate the key structural parameters uniformly either across all individuals (e.g., Goettler and Clay 2011) or in segments (e.g., Arcidiacono and Miller 2011 and Chung et al. 2014), forgoing individual heterogeneity to a certain extent. In our context of individual usage of mobile data, significant heterogeneity exists among users, and a small number of segments cannot capture the rich variation in multiple heterogeneous parameters. Therefore, we take advantage of the fine granularity of our data and develop computationally efficient estimation based on Hotz and Miller (1993) and Bajari et al. (2007) so as to specify and estimate our dynamic model individual specifically. In this respect, our estimation strategy is similar to Misra and Nair (2011), who perform individually separate estimation for 87 sales agents. Our successful implementation of speedy estimation for hundreds of users serves as a meaningful example for future studies when significant individual heterogeneity needs to be taken full account of and/or large-scale parallel implementation is important.

3. Data

We obtain the data for this study from one of the leading telecommunications companies that provide mobile services in China.[1] This carrier has nationwide mobile network coverage and has hundreds of millions of subscribers. We collaborate with the company's subsidiary in a certain province and obtain the mobile data usage records of its subscribers in the capital city of that province.[2] The data furnished to us span a nine-month period from January to September 2013 and include all subscribers in the focal city who had been using the company's mobile data service with monthly data plans throughout the nine months.

A mobile data user with the focal carrier usually has a main service package, which typically consists of a combination of voice call minutes, text messages, and other value-added mobile services (e.g., free music downloads, real-time stock market information), marketed under different brand series with different cost structures. On top of the main service package, a user may also have a mobile data plan, which adopts a typical three-part tariff price structure: After paying an up-front monthly fee for a limited data quota (e.g., ¥50.00 for 500 MB), a user can consume mobile data up to the plan quota without any additional cost within a month.[3] If the total data usage within a month exceeds the plan quota, overage charges will be levied as a fixed unit price rate per megabyte of additional usage. The billing cycle is based on calendar months and hence the same for all users. At the beginning of each month, the unused data quota from the previous month is automatically cleared, and the accumulated data usage is reset to zero. For our study period, there are seven different data plans available, with the plan quota ranging from 30 MB to 5,120 MB and the monthly fee from ¥5.00 to ¥200.00. The extra rate, however, is the same across all these data plans,[4] which simplifies our estimation approach as we discuss in Section 6.

For this study, we select from the original data a random sample of 1,000 users who have not changed their data plan within the nine-month period. We focus on these users because of practical consideration. Note first that our model can be easily extended to incorporate users who changed their data plan. We can simply allow the mean and variance of the daily usage to differ before and after the plan change, because it reflects their rational choices based on varying consumption needs. Nevertheless, the increased number of model parameters results in much fewer data points applicable to the identification of each parameter. Given that we have only nine months of data, for the purpose of better identification and more efficient estimation, we choose to focus on the users who have not changed their data plan within the study period. We also exclude from our sample those newly joined users. Because we focus on the general mobile data users instead of new users for this study, we do not model users' initial learning process. For this reason, we only include in our sample those experienced users who have been with the focal carrier for at least one year, which account for over 85% of all users in our original data.

For each individual user in our sample, we include the following information for the analysis: (i) the monthly plan information, including the chosen main service package and data plan each month, the plan quota, and the monthly fees and overage price rates; (ii) the daily mobile data usage for each day throughout the entire study period, resulting in 273 data points for each user; and (iii) the user profile information, including age, gender, and customer history (i.e., how long, measured in months, the user has been with this carrier at the beginning of the study period). Table 1 provides the summary statistics of the user profile information.

Figure 1 shows the histograms of the daily mobile data usage for four representative users. As we can see, there exists variance in a user's mobile data consumption from day to day, and the distribution of the daily usage can be well approximated by a normal distribution truncated at zero. Notice that with all negative values replaced by zero, there is a point mass at zero accounting for the mass of the negative values of a normal distribution. The data patterns justify our model setting with regard to the distribution of the random utility shocks.

4. Model

In this section, we construct our dynamic structural model of daily mobile data consumption. To take full account of the heterogeneity among users, we exploit the granularity of our individual-level daily usage data over multiple months and specify our model such that the model parameters are user specific and individually estimated.

Consider an individual user i with monthly mobile data plan j, which adopts a typical nonlinear pricing scheme of three-part tariff. With a fixed monthly fee F_j (in Chinese yuan, CNY or ¥, as we use throughout this paper), the plan includes a monthly data quota Q_j (in megabytes, or MB, as we use throughout this paper). The user can consume mobile data up to Q_j without incurring any extra cost. If the total data usage within a month exceeds Q_j, the user needs to pay a unit price p_j per MB of additional usage. Denote user i's daily mobile data usage as a_{it} ($a_{it} \geq 0$) for any given day t within a month, $t = 1, \ldots, T$, where T is the total number of days in the month.[5] At the beginning of day t, the remaining data quota user i has is thus $q_{it} = \max\{Q_j - \sum_{\tau=1}^{t-1} a_{i\tau}, 0\}$.

For any given day t, user i's per-period utility from consuming an a_{it} amount of mobile data for that particular day is modeled as follows:

$$u_i(a_{it}) = (\mu_i + \xi_{it})a_{it} - \frac{1}{2}a_{it}^2 - \eta_i p_j \max\{a_{it} - q_{it}, 0\}. \quad (1)$$

Table 1. Summary Statistics of the User Profile Information

	Mean	Std. dev.	Max	Min
Age	32.8	9.87	67	16
Male	0.697	0.460	1	0
Customer History	58.7	35.4	192	12

The first two terms of the above utility function (i.e., $(\mu_i + \xi_{it})a_{it} - (1/2)\,a_{it}^2$) capture the direct utility from mobile data consumption. We adopt the typical quadratic functional form, which is commonly used for utility functions in the literature related to mobile communication (e.g., Iyengar et al. 2007 and Kim et al. 2010). The quadratic functional form properly models the diminishing marginal utility as the consumption volume increases, where the linear term can be viewed as reflecting the utility gain from the data consumption and the quadratic term as the disutility associated with the time and effort spent as well as the opportunity cost. To account for the variation in an individual user's mobile data usage need from day to day, in addition to the mean level μ_i, we incorporate a daily random shock, ξ_{it}, into the utility function (1). The private shock ξ_{it} is observed only to each individual user but not to the researchers. We assume ξ_{it} follows a normal distribution with individual-specific variance, $N(0, \sigma_i^2)$ and is independently and identically distributed across time. As we discuss in Section 3, a normal distribution of the daily random utility shocks well approximates the distribution of daily usage observed in the data. Notice that the daily mobile data usage a_{it} is bounded below at zero, and we allow both μ_i and ξ_{it} to take either positive or negative values. Therefore, a negative value of $\mu_i + \xi_{it}$ indicates no need for consuming mobile data on that particular day and will result in the optimal usage volume of zero, whereas a positive and large value of $\mu_i + \xi_{it}$ indicates the opposite and can lead to a high usage volume on that particular day.

The last term in the per-period utility function (1) (i.e., $\eta_i p_j \max\{a_{it} - q_{it}, 0\}$) accounts for the cost of extra usage exceeding the plan quota.[6] Here, η_i captures the individual-specific price sensitivity, measuring the utility cost per unit of monetary spending. Notice that the coefficients in the utility function (1) is only identifiable up to a constant scale. Therefore, for identification purpose, we normalize the coefficient before the second term (i.e., a_{it}^2) to a fixed constant 1/2. It is also worth noting that the focal mobile carrier deducts overage charges daily from users' account balances, so we include the associated utility cost in the per-period utility function rather than till the end of a month.

When determining the amount of mobile data to consume each day, rational users who are forward looking not only consider the utility for the present day but also take into account the expected utility for the future days in the rest of the month. To properly model such a decision process, we assume that at the beginning of each day t, each individual user observes her private daily shock ξ_{it} and *then* decides her data usage for that day, a_{it}. A user optimizes her usage in any day by weighing her present utility (given the observed shock for that day) against her expected future utility (by taking expectation over all the future utility shocks). The

Figure 1. Histogram of Daily Data Usage

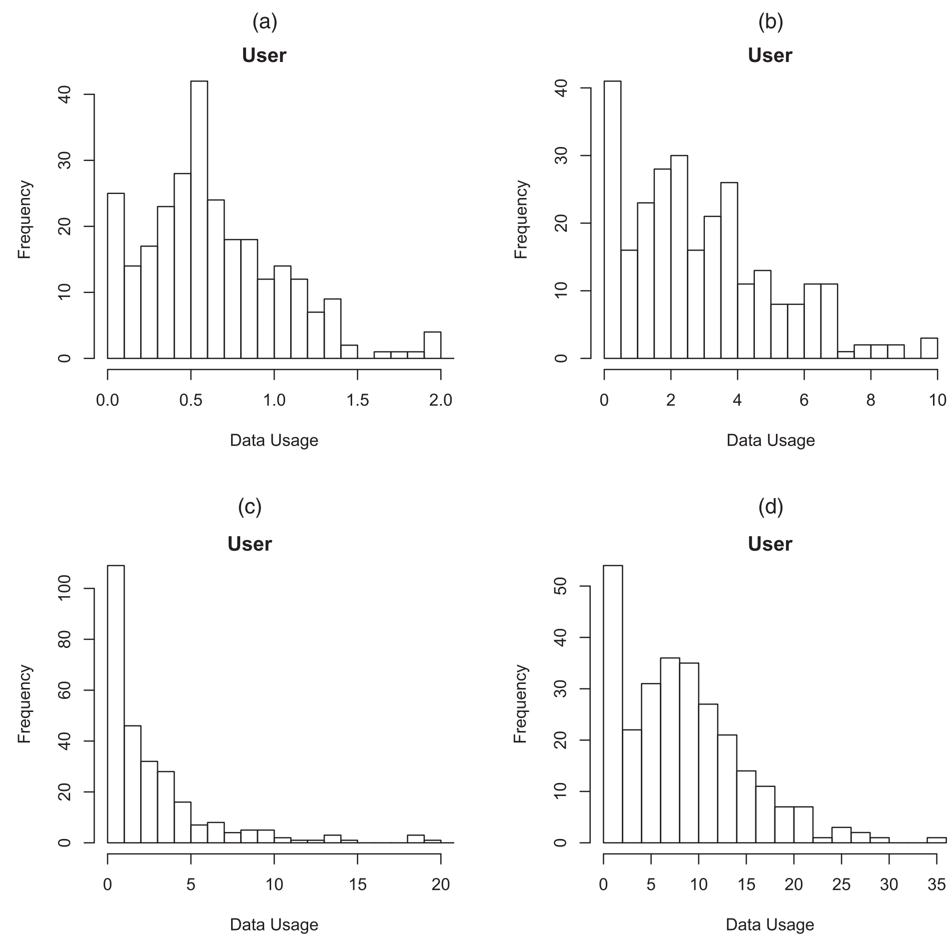

optimization problem for each user is thus to maximize the sum of the present utility and the expected future utility, which can be formulated as

$$u_i(a_{it}; q_{it}, \xi_{it}) + E_{\{\xi_{i\tau}\}_{\tau=t+1}^{T}}\left[\sum_{\tau=t+1}^{T} \beta^{\tau-t} u_i(a_{i\tau}; q_{i\tau}, \xi_{i\tau})\right], \qquad (2)$$

where β is the discount factor for future utility. Notice that our dynamic model assumes users are forward looking, and $0 < \beta < 1$. For users who behave myopically, however, $\beta = 0$, and the dynamic model degenerates into a simple static model that only optimizes the per-period utility in (1), as we explore further in Section 5.1. The expectation in (2) is taken with respect to all the future shocks $\{\xi_{i\tau}\}$ for $\tau = t+1, \ldots, T$, which, in turn, determine the optimal future consumption path $\{a_{i\tau}\}$ and the associated utilities $\{u_i(a_{i\tau})\}$.

The problem described above can be formulated as a dynamic programming problem with T periods (i.e., days). At the beginning of each period t, a user i decides her optimal usage based on three state variables: the remaining unused data quota from her monthly plan, q_{it}; the number of days left in the month, d_t; and the private utility shock for the current period, ξ_{it}. Among them, q_{it} and d_t are evolving states transitioning from period to period deterministically given the previous usage; that is,

$$q_{i,t+1} = \max\{q_{it} - a_{it}, 0\} \quad \text{and} \quad d_{t+1} = d_t - 1. \qquad (3)$$

Note that here we assume each user knows her remaining data quota each day. This is a reasonable assumption considering the wide availability of built-in tools in most smartphones and various third-party apps that can easily track the cumulative mobile data usage with a single finger tap.

The value function, $V(q_{it}, d_t, \xi_{it})$, can be defined recursively backward by the Bellman equation:

$$V_i(q_{it}, d_t, \xi_{it}) = \max_{a_{it}\geq 0}[u_i(a_{it}; q_{it}, \xi_{it}) + \beta E_{\xi_{i,t+1}} V_i(q_{i,t+1}, d_{t+1}, \xi_{i,t+1})]. \qquad (4)$$

Starting backward from the last period (i.e., the last day in a month so that $t = T$ and $d_t = 1$), where $V_i(q_{it}, d_t = 1, \xi_{it}) = \max_{a_{it}\geq 0} u_i(a_{it}; q_{it}, \xi_{it})$, (4) determines the maximized sum of the present utility and the expected future utility as a function of the state variables in any period. The solution to (4) yields the policy function $a_i^*(q_{it}, d_t, \xi_{it})$, which determines the optimal mobile data usage in any day as a function of the state variables q_{it}, d_t, and ξ_{it}. Because the policy function directly links to the daily usage observed in our data, it hence plays a central role in our estimation strategies.

The structural model parameters to be estimated for each individual user are therefore μ_i, σ_i^2, and η_i. Note that we assume users have full information about their structural parameter values and do not model users' initial learning processes, because for this study, we examine the general population of mobile data users rather than the newly acquired users. As we discuss in Section 3, the vast majority of the users in our data set are experienced users who have been with the carrier for a long time. Therefore, we focus on experienced users and abstract away from the learning problem in the scope of this study. It is also worth mentioning that our study focuses on the consumption of mobile data through the mobile network of the focal carrier. When users have access to alternative data connections (e.g., Wi-Fi), it diminishes their need of using mobile data, which will be reflected in our model in either the mean usage need μ_i (for constant Wi-Fi access) or the changing utility shocks ξ_{it} (for occasional Wi-Fi access).

5. Test for Forward-Looking Behavior

In this section, we first derive important theoretical properties from our dynamic structural model; we then develop a reduced-form empirical test based on these theoretical results and apply the test to our data sample. The theoretical results themselves characterize notable and undocumented temporal consumption patterns for the class of dynamic problem as we study. The empirical test offers a useful tool to identify individual forward-looking behavior. The test results provide empirical evidence that a certain proportion of users are indeed forward looking and dynamically optimize their usage over time. They also enable us to control for possible consumer myopia and apply our dynamic model only to those users exhibiting the forward-looking usage pattern consistent with the model.

5.1. Theoretical Properties

We are particularly interested in exploring any theoretical results on the consumption patterns that will differentiate those users who are forward-looking from those who are not. We start with a benchmark case in which users are myopic, who determine their daily mobile data usage only based on the present utility (as in (1)) without considering the future utility (as in (2)). Myopic behavior can be viewed as degeneration of the dynamic model with $\beta = 0$. In such a case, the dynamic programming problem reduces to a static optimization problem, which can be fully solved in a closed form, leading to the following result.

Proposition 1. *For myopic users, given any utility shock ξ_{it} in any given day before the day when the data plan quota is fully expended (i.e., $q_{it} > a_i^*(q_{it}, d_t, \xi_{it})$), their daily mobile data usage $a_i^*(q_{it}, d_t, \xi_{it})$ is independent of the remaining data plan quota q_{it}.*

Proof. The proof is detailed in the online appendix. □

For myopic users, in maximizing the per-period utility (1) in a given day, as long as there is enough plan quota remaining (i.e., q_{it} is large enough), they simply maximize the direct utility from data consumption (i.e., the first two terms in (1) as the last term equals 0 when $q_{it} > a_{it}^*$). In other words, they just use as much as they need, which maximizes the direct consumption utility given the realized utility shock for that day and does not depend on q_{it} at all.

Next, we turn to our dynamic model and consider users who are forward looking. They optimize their daily usage to maximize the sum of the present utility and the expected future utility as in (2). As a result, their optimal usage typically changes dynamically over time, depending on the state in each period. In particular, we are interested in whether and how the policy function $a_i^*(q_{it}, d_t, \xi_{it})$ depends on the remaining plan quota q_{it}, so as to contrast with the result for myopic users.

To prove the result formally and generally, we first prove two lemmas with regard to some important properties of the expected value function $\bar{V}_i(q_{it}, d_t)$, which is the value function $V_i(q_{it}, d_t, \xi_{it})$ taken expectation over the random utility shock ξ_{it},

$$\bar{V}_i(q_{it}, d_t) = E_{\xi_{it}} V_i(q_{it}, d_t, \xi_{it}). \tag{5}$$

The two lemmas and their proofs are presented in the online appendix because of their technical nature. in brief, Lemma A.1 shows that the expected value function in the last period (i.e., when the number of remaining days $d_t = 1$) is *strictly concave* in the remaining quota q_{it}; Lemma A.2 further shows that if the next-period expected value function $\bar{V}_i(\cdot, d-1)$ is strictly concave, then the current-period expected value function $\bar{V}_i(\cdot, d)$ is also strictly concave. By backward induction utilizing both lemmas, we can conclude that $\bar{V}_i(q_{it}, d_t)$ is strictly concave in the remaining quota q_{it} in any period. To illustrate this result, in Figure 2, we present two examples of the expected value function numerically calculated under different sets of parameter values. The algorithm we developed to numerically solve the value function will be discussed in detail in Section 6.3.

The strict concavity of the expected value function leads to the following properties of the policy function for forward-looking users.

Proposition 2. *For forward-looking users, given any utility shock ξ_{it} in any given day (except the last day of a month) before the day when the data plan quota is fully expended (i.e., $q_{it} > a_i^*(q_{it}, d_t, \xi_{it})$), the optimal daily usage $a_i^*(q_{it}, d_t, \xi_{it})$ is strictly increasing in q_{it} if $0 < a_i^* < q_{it}$.*

Figure 2. Illustration of the Expected Value Function $\bar{V}(q_{it}, d_t)$

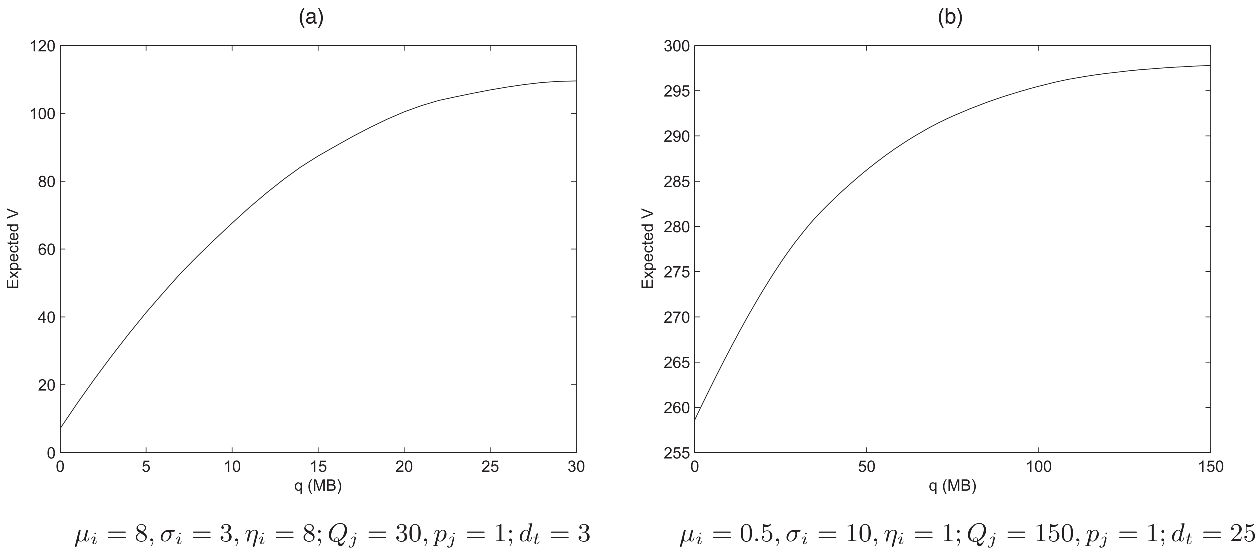

Proof. The proof is detailed in the online appendix. □

To understand the reasoning leading to Proposition 2, recall the optimization problem facing forward-looking users as in (4), which can be written more simply as

$$a_i^*(q_{it}, d_t, \xi_{it}) = \underset{a_{it} \geq 0}{\arg\max}\,[u_i(a_{it}) + \beta \bar{V}_i(\max\{q_{it} - a_{it}, 0\}, d_{i,t+1})]. \tag{6}$$

From (6), it is clear that users face the trade-off between their immediate utility from current consumption, $u_i(a_{it})$, and their future utility given the quota remaining after the current-period consumption, $\bar{V}_i(\max\{q_{it} - a_{it}, 0\}, d_{i,t+1})$. Users will increase their usage amount if the marginal benefit from current consumption exceeds the marginal benefit from saving for the future and decrease their usage amount if the opposite. If the remaining quota at the beginning of the current period q_{it} increases, with the same amount of data usage a_{it}, the quota left for the next period ($\max\{q_{it} - a_{it}, 0\}$) also increases in general. Because $\bar{V}_i(\cdot, d_{i,t+1})$ is strictly concave as we have shown, the marginal utility gain from saving for the future will then decrease, which gives users the incentive to increase their consumption in the current period. As a result, the optimal usage a_{it}^* strictly increases with q_{it} in general. Exceptions arise after the plan quota is used up or in the last day of the month. In these cases, there is simply nothing to save for the future, so the dynamic programming problem reduces to the same static optimization problem facing myopic users as is discussed in Proposition 1.

The distinct usage patterns that distinguish forward-looking users from myopic ones are testable given our observed data. We can use each individual user's daily usage before the plan quota is used up or the last day of a month to test for any positive relationship with the remaining quota so as to discover evidence of forward-looking consumption behavior, as we develop next. It is worth clarifying that the dependence of the policy function on the remaining quota shown in the two propositions holds *given* the random utility shock ξ_{it} and a particular day d_t. The said relationship with the remaining quota may not hold across different days. Moreover, even when d_t and q_{it} are given, the optimal usage a_{it}^* is still a random variable depending on the realization of the unobservable utility shock ξ_{it}. Therefore, what we should test for is the average trend between the daily usage and the remaining quota in terms of expectation rather than each realization per se.

5.2. Empirical Test

From the theoretical results in Section 5.1, we now develop a reduced-form empirical test and apply it to our data sample. According to the two propositions above, to identify forward-looking behavior, we can examine each individual user's daily mobile data usage over the nine months for the days before the plan quota is used up (excluding the last day in each month). For users who are forward looking and dynamically optimize their usage, their daily usage given the number of remaining days, which is a random variable depending on the unobservable demand shock, should demonstrate positive correlation with the remaining quota at the beginning of each day. By contrast, for myopic users, their daily usage should be uncorrelated with the remaining quota. To test for such usage patterns, we apply the following reduced-form empirical model:

$$a_{it} = \begin{cases} \tilde{a}_{it} & \text{if } \tilde{a}_{it} > 0, \\ 0 & \text{otherwise;} \end{cases}$$
$$\tilde{a}_{it} = \alpha_{i0m} + \alpha_{i1} q_{it} + \alpha_{i2} d_t + \varepsilon_{it}, \; \varepsilon_{it} \sim N(0, \omega_i^2). \tag{7}$$

This is a typical Type I Tobit model, where the dependent variable a_{it}, which is the observed usage for user i on day t of a month, reflects a latent variable $\tilde{a}_{it}$ such that a_{it} equals $\tilde{a}_{it}$ if it is positive and equals 0 otherwise. The latent variable $\tilde{a}_{it}$ has a regression structure regarding the remaining quota q_{it}, controlling for the number of remaining days d_t. The error term ε_{it} follows a normal distribution, which is consistent with the actual distribution of daily usage observed in our data as is discussed in Section 3 (Figure 1). We allow the coefficients (α_i's) and the variance of the error term (ω_i^2) to be individually different for different users. We also account for the monthly fixed effects by allowing the intercept α_{i0m} to differ across different months. The key parameter of interest here is α_{i1}, which is expected to be significantly greater than zero for forward-looking users and insignificantly different from zero for myopic users.

We use maximum likelihood estimation to estimate the model and obtain the p-values for coefficient estimates. We apply the model for each individual user within our data sample one by one and hence need to run 1,000 separate estimations. Among all 1,000 users in our sample, we found 392 users with a positive estimate of α_{i1} that is significant at the 1% level. Among the remaining 608 users, 597 users have α_{i1} estimates insignificantly different from zero. There are only 11 users (i.e., 1.1%) with a significantly negative α_{i1}, which can be considered as singular outliers. It is worth noting that 99% of the users in our data are shown to be explainable by our model, which provides empirical validation for our model and the associated theoretical results. Figure 3 depicts the histogram of the estimated α_{i1}'s.

To ensure the reliability of the test results, we perform further robustness checks. Notice that the model parameters in (7) are individually specific, meaning that the model already accounts for all kinds of individually specific fixed effects. To further take account of possible time-varying effects, such as weekend or national holidays, we add additional dummies into the regression specification of $\tilde{a}_{it}$ in (7) as control variables. We found that both weekends and national holidays do not have significant effects on daily usage. Furthermore, after controlling for the weekend and holiday effects, the test results (in terms of the significance of α_{i1}) remain unchanged for 98% of all users. We further test whether users' consumption behaviors are stable over time (e.g., whether users may display myopic usage patterns at the beginning of the month and fall more in line with the forward-looking behavior near the month end, whether users may accelerate consumption in the last several days of a month). For such tests, we allow users to have different regression coefficients for different time periods within a month and do not find statistically significant differences among these coefficients. In addition, we also check whether a possible nonlinear functional relationship between daily consumption amount a_{it} and number of remaining days d_t could affect the estimated coefficient of q_{it} in our model. We find our test results are robust when using fixed effects for each day of a month or including a nonlinear (quadratic) term of d_t. Moreover, we further verify that the observed temporal variation in daily usage is unlikely to be caused by systematic demand changes within a month by finding no significant relationship between the daily consumption a_{it} and the day within a month d_t for myopic users.

To further corroborate the proposed test, we explore outside of the commonly used hypothesis-testing framework adopted here and conduct two alternative methods of segmentation, which both lead to consistent results. First, we seek to identify forward-looking individuals based on model comparison. We estimate a constrained version of model (7) by assuming α_{i1} to be zero along with the original unconstrained version of model (7) and then compare the resulting Bayesian information criteria (BIC). We find that all the 392 forward-looking users according to our proposed test yield higher BIC under the unconstrained model than

Figure 3. Histogram of the Test Results (Estimated α_{i1})

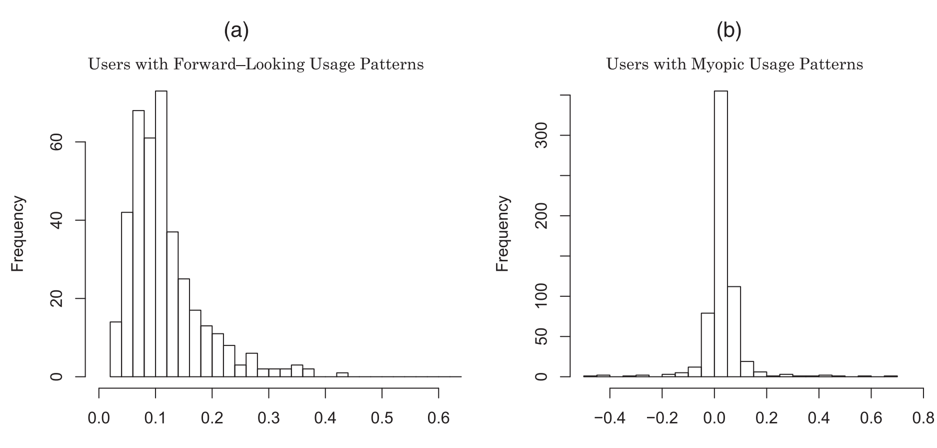

the constrained one, suggesting consistent correlation between the daily usage and the remaining quota. Second, we further apply the clustering analysis from unsupervised machine learning. In particular, we calculate the partial correlations between the daily usage and the remaining quota (or the remaining days) controlling the remaining days (or the remaining quota) for each individual and then perform the k-means clustering based on these two correlations. The gap statistics (Tibshirani et al. 2001) clearly suggest that the two-cluster segmentation strongly dominates any other number of clusters. The correlations in the usage patterns for these two clusters are consistent with our theoretical model predictions of the consumption behaviors of forward-looking and myopic users. The cluster membership results achieve about 90% overlap with our original test results. It is noteworthy that the model-free clustering analysis is independent of our theoretical or empirical models and endogenously segments users solely based on the data. The consistent findings further corroborate the robustness and validity of our proposed framework.

In summary, we can confidently conclude that the reduced-form test developed in this section, rooted in the classical framework of hypothesis testing with a clear interpretation of the underlying theoretical foundation, does reliably capture a fundamental distinction among mobile data users' dynamic consumption behaviors. The empirical results provide formal evidence that there is a considerable proportion of users (i.e., about 40%) who are forward looking and dynamically plan their mobile data usage in a way consistent with our structural model. Meanwhile, the findings also suggest the necessity of controlling for possible consumer myopia in studying dynamic consumption behavior at the individual level.

6. Dynamic Model Estimation

Being able to identify users with evident forward-looking usage patterns, we can apply our dynamic model to these users and estimate their structural parameters $\{\mu_i, \sigma_i, \eta_i\}$. As discussed in Section 2, the discount factor cannot be identified by general observational data such as ours. Therefore, we follow the common approach in the literature and set $\beta = 0.9$ for these forward-looking users.[7]

We adopt the two-stage estimation strategy proposed by Bajari et al. (2007; henceforth referred to as BBL). We choose this estimation approach for its advantage in computational efficiency. The general idea of the two-stage estimation strategy is as follows. In the first stage, we empirically estimate the policy function by fitting a distribution of the daily usage from the observed data. In the second stage, we estimate the structural parameters by ensuring that the policy function estimated in the first stage is indeed optimal. As shown by BBL, such a two-stage estimation yields consistent estimates of the structural model parameters.

6.1. First-Stage Estimation

In the first stage, we empirically estimate the policy function $a_i^*(q_{it}, d_t, \xi_{it})$ actually adopted by users in the observed data. Recall that without observing the private utility shock ξ_{it}, given the two observed states (q_{it} and d_t) only, the daily usage $a_i^*(q_{it}, d_t, \xi_{it})$ is observed as a random variable. Therefore, to estimate the policy function, we first empirically estimate the conditional distribution of the daily usage a_{it} given the two observed state variables, defined by its cumulative distribution function (cdf) $F(a_{it}|q_{it}, d_t)$. We then back out the policy function $a_i^*(q_{it}, d_t, \xi_{it})$ based on the estimated $F(a_{it}|q_{it}, d_t)$ and the distribution of the private shock ξ_{it}.

Following BBL, we specify a flexible functional form to approximate the conditional distribution $F(a_{it}|q_{it}, d_t)$, as follows:

$$a_{it} = \begin{cases} \breve{a}_{it} & \text{if } \breve{a}_{it} > 0, \\ 0 & \text{otherwise;} \end{cases} \qquad (8)$$
$$\breve{a}_{it} \sim N(\theta_i(q_{it}, d_t), \tau_i^2(q_{it}, d_t)).$$

We let daily usage a_{it} be the censored observation of a latent variable $\breve{a}_{it}$ so that a_{it} equals $\breve{a}_{it}$ if it is positive and equals 0 otherwise. The latent variable $\breve{a}_{it}$ follow a normal distribution whose mean and variance are flexible functions of (q_{it}, d_t). We specify $\theta_i(q_{it}, d_t)$ and $\log \tau_i^2(q_{it}, d_t)$ in flexible parametric forms as polynomial functions of q_{it} and d_t. Specifically, we let $\theta_i(q_{it}, d_t) = [1, q_{it}, q_{it}^2, d_t, d_t^2, q_{it}d_t] \cdot \boldsymbol{\Theta_i}$ and $\tau_i^2(q_{it}, d_t) = \exp([1, q_{it}, d_t] \cdot \boldsymbol{\Gamma_i})$, where $\boldsymbol{\Theta_i}$ and $\boldsymbol{\Gamma_i}$ are a series of coefficients to be estimated. We use the maximum likelihood estimation to estimate the coefficients $\boldsymbol{\Theta_i}$ and $\boldsymbol{\Gamma_i}$ based on the observed data $\{a_{it}, q_{it}, d_t\}$ for all days throughout the entire observation period (excluding the days with no quota left and the last day of a month) for each individual user i.

After obtaining the estimated coefficients, we determine $F_i(a_{it}|q_{it}, d_t)$ according to (8) and then derive the policy function given $F_i(a_{it}|q_{it}, d_t)$. We utilize the fact that $a_i^*(q_{it}, d_t, \xi_{it})$ is increasing in ξ_{it}, so $F_i(a_i^*(q_{it}, d_t, \xi_{it})|q_{it}, d_t) = G_i(\xi_{it})$,[8] where $G_i(\cdot)$ is the cdf of ξ_{it}. As a result, the policy function can be written as

$$a_i^*(q_{it}, d_t, \xi_{it}) = F_i^{-1}(G_i(\xi_{it})|q_{it}, d_t). \qquad (9)$$

Equation (9) determines the policy function estimated from the observed data, which is the outcome of the first-stage estimation.

6.2. Second-Stage Estimation

In the second stage, we estimate the structural parameters such that the policy function estimated in the first stage and the observed data plan choice are both optimal under these parameter values.

We first define the expected utility given a usage strategy $a_i(q_{it}, d_t, \xi_{it})$ at any state (q_0, d_0) *before* the realization of random shock ξ_0 as

$$\bar{U}_i(q_0, d_0; a_i) = E_{\{\xi_{it}\}_{t=0}^{d_0-1}} \sum_{t=0}^{d_0-1} \beta^t u_i(a_i(q_{it}, d_t, \xi_{it})). \tag{10}$$

Under the true model parameters, the policy function estimated in the first stage, $a_i^*(q_{it}, d_t, \xi_{it})$, should be optimal such that for any state, the expected utility given the policy function $a_i^*(q_{it}, d_t, \xi_{it})$ is no less than that given any alternative $a_i'(q_{it}, d_t, \xi_{it})$. Following BBL, we construct the alternative policy functions by adding arbitrary perturbations to the observed policy function; that is,

$$a_i'(q_{it}, d_t, \xi_{it}; e) = \max\{a_i^*(q_{it}, d_t, \xi_{it}) + e, 0\}. \tag{11}$$

Note that the perturbation e can be positive or negative, whereas $a_i'(q_{it}, d_t, \xi_{it})$ is bounded below at zero because the data usage cannot be negative. Thus, the optimality of $a_i^*(q_{it}, d_t, \xi_{it})$ implies that the following inequality holds for any set of (q, d, e):

$$g_i(q, d, e) = \bar{U}_i(q, d; a_i^*) - \bar{U}_i\left(q, d; a_i'(e)\right) \geq 0. \tag{12}$$

For estimation, we generate $n_I = 500$ different sets of states and policy function perturbation, $\{q_k, d_k, e_k\}_{k=1}^{n_I}$, and evaluate (12) at all of these 500 different sets of values.

In addition to $a_i^*(q_{it}, d_t, \xi_{it})$ being optimal, under the true model parameters, users' data plan choices should also be optimal. Given the chosen data plan j, let F_j be its monthly fee, and let Q_j be the total plan quota accordingly. A user's rational selection of data plan j implies that the expected utility at the beginning of a month under this particular plan is no less than that under any other data plan j'. Therefore, the following inequalities should also hold:

$$g_i'\left(Q_j, Q_{j'}, T; F_j, F_{j'}\right) = \bar{U}_i(Q_j, T; a_i^*) - \bar{U}_i\left(Q_{j'}, T; a_i^*\right) - \eta_i\left(F_j - F_{j'}\right) \geq 0, \quad \forall j' \neq j. \tag{13}$$

Note that $\bar{U}_i(Q_j, T; a_i^*)$ is evaluated at the beginning of a month, and T is the total number of days in a month, which can be 31, 30, or 28. In formulating (13), we exploit the fact that the price rate for overage remains unchanged across different data plans for any given user, as discussed in Section 3. For this reason, even if a user were to switch to an alternative data plan, her policy function would remain the same as the one observed in the data. Therefore, we can use the same $a_i^*(q_{it}, d_t, \xi_{it})$ estimated in the first stage to compute $\bar{U}_i(Q_{j'}, T; a_i^*)$, $\forall j' \neq j$.

We obtain the estimators of the structural parameters $\{\mu_i, \sigma_i, \eta_i\}$ by minimizing the overall violation of the two sets of inequalities in (12) and (13)—that is, by minimizing the following objective function:

$$\sum_{k=1}^{n_I} (\min\{g_i(q_k, d_k, e_k), 0\})^2 + \sum_{j'}\left(\min\left\{g_i'\left(Q_j, Q_{j'}, T; F_j, F_{j'}\right), 0\right\}\right)^2. \tag{14}$$

In evaluating (14), we apply Monte Carlo methods to compute the expected utility $\bar{U}_i(q, d; a)$, which cannot be derived in a closed form. For each expected utility, we simulate $n_S = 10{,}000$ random shock paths of $\{\xi_{it}\}_{t=0}^{d-1}$, calculate the usage path $\{a_{it}\}$ according to the given policy function, and generate the state path $\{q_{it}, d_t\}$. We then calculate the sum of discounted utility in each simulation round and average over all n_S simulation paths to approximate the expected utility $\bar{U}_i(q, d; a)$,[9] based on which we can compute the objective function and obtain the structural parameter estimates.

6.3. Simulation Method

In this section, we briefly describe the method we develop to simulate each individual user's daily usage based on the estimated parameter values. We use this method to validate our estimation results by comparing the simulated usage with the actual usage, as we report in Section 6.4. It will also be used in the policy experiments in Section 7, where we explore various promotion designs.

Given the estimated parameter values, to simulate users' daily usage according to the dynamic model, we first numerically solve the dynamic programming problem. In particular, we need to calculate the expected value function $\bar{V}_i(q, d)$ (defined by (5)) at any state (q, d). We use backward recursion to solve $\bar{V}_i(q, d)$. We start from the last period (i.e., the number of remaining days $d = 1$), where we can analytically derive the optimal data usage, and hence the value function $V_i(q, d = 1, \xi)$, in a closed form. We then compute $\bar{V}_i(q, d = 1)$ by simulating from the distribution of ξ and taking the average.

Note that we need to solve $\bar{V}_i(q, d)$ as a function of q continuously and obtain the value of $\bar{V}_i(q, d)$ at any possible q given d. To do so in a computationally feasible way, we first compute $\bar{V}_i(q, d)$ at multiple different values of q and then interpolate the function values at other points. Recall that we have shown $\bar{V}_i(q, d)$ is increasing in q (Lemmas A.1 and A.2 in the online appendix). Therefore, the interpolation algorithm needs to preserve the monotonicity of the original function. For this reason, we choose to use the monotone piecewise cubic Hermite interpolation (Fritsch and Carlson 1980), which is a variant of spline interpolation that preserves monotonicity of the original data points being interpolated while ensuring smoothness of the interpolated function. Thus, our method introduces a useful alternative to the dynamic structural model literature, which is different from the commonly used interpolation methods (e.g., Erdem and Keane 1996 and Judd 1998) and especially relevant in the context when monotonicity is clearly established.

Knowing $\bar{V}_i(q, d)$ at any possible q given d, we can recursively derive $V_i(q, d+1, \xi)$ by solving the optimization problem in (4). Repeating the steps described above, we can compute $\bar{V}_i(q, d+1)$ at multiple points of q and then interpolate the function values at other points. By backward recursion, we numerically solve the expected value functions day by day till the first day of the month. Figure 2 presents two examples of the expected value functions solved under different sets of parameter values.

After solving all expected value functions, the optimal data consumption $a_i^*(q_{it}, d_t, \xi_{it})$ can be obtained by solving the maximization problem in (6) given any simulated utility shock ξ_{it}.

6.4. Results

We follow the two-stage estimation strategies described in Sections 6.1 and 6.2 and estimate our dynamic model individual by individual for the 392 users who exhibit forward-looking usage patterns consistent with the model based on the test results in Section 5.2. We write the estimation programs in MATLAB and run them on a Windows workstation computer with multicore processors, which facilitates parallel computing and significantly reduces the computing time despite the hundreds of separate estimations. We are able to obtain the estimates of the structural model parameters for each of the 392 users individually. Figure 4 shows the histogram of the estimates of $\{\mu_i, \sigma_i, \xi_i\}$ across all these users.

As Figure 4 shows, significant variation exists in the estimates of all three structural parameters across users, reflecting considerable heterogeneity in users' dynamic behavior of mobile data consumption. Note that although the mean value of the daily data consumption need, μ_i, is positive for most users, a small portion of users have negative μ_i's, indicating that they oftentimes do not need to use mobile data, which is consistent with the observations in our data (e.g., Figure 1). The parameter of particular interest is η_i, which measures each individual user's price sensitivity in mobile data consumption. As we can see, although the less price-sensitive users have η_i close to zero, some users' price sensitivity levels could be considerably high. Such variation reaffirms the necessity of fully capturing individual heterogeneity and avoiding the bias of assuming a common parameter value or just a small number of segment values for all users. It is also worth noting that our dynamic model is able to identify a user's price sensitivity through intertemporal substitution reflected in daily usage, even if she has never exceeded her plan quota and thus has never actually incurred any overage charge. By contrast, a static model would be unable to identify the price sensitivity for the users who have never exceeded their plan quota.

To validate that the estimates of model parameters properly reflect the true dynamics in the actual data, we simulate each user's daily usage paths given the estimated structural parameter values by following the approach described in Section 6.3. For each user, we first simulate the daily utility shocks for a 30-day period and compute the optimal usage for each day given these shocks. We simulate such a 30-day usage path 100 times and calculate the average usage for each day. We then contrast the simulated data with the actual data (i.e., the average usage for each day of a month over the nine-month observation period). Figure 5 depicts the simulation results for six representative users.

In Figure 5, the solid curves represent the average daily usage over 100 rounds of simulation outputs,

Figure 4. Dynamic Model Estimation Results

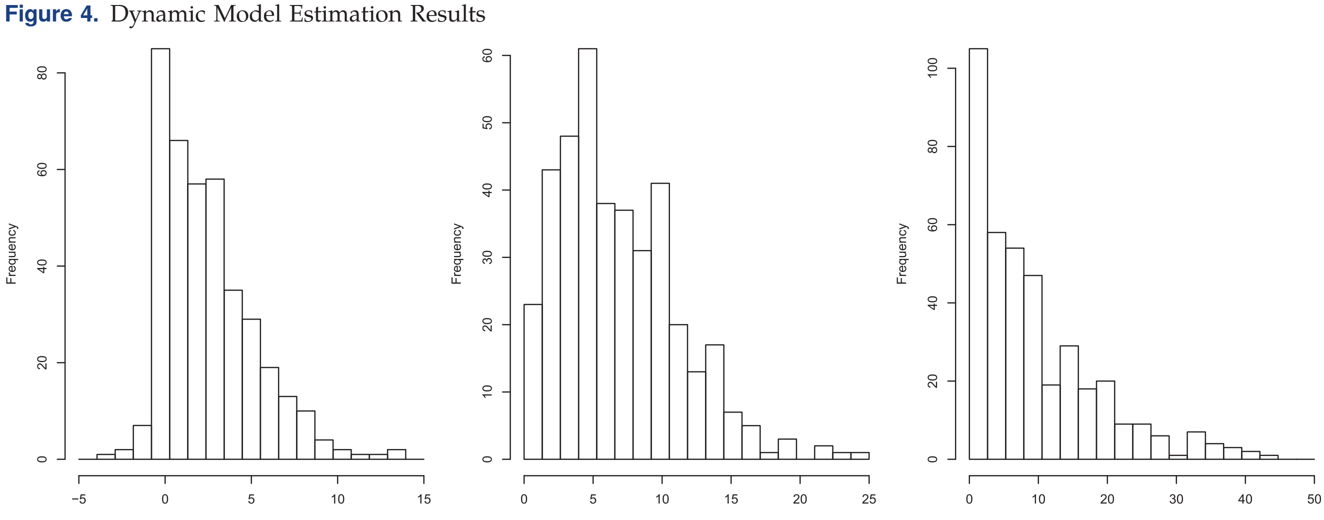

Figure 5. Simulation Output vs. Actual Data

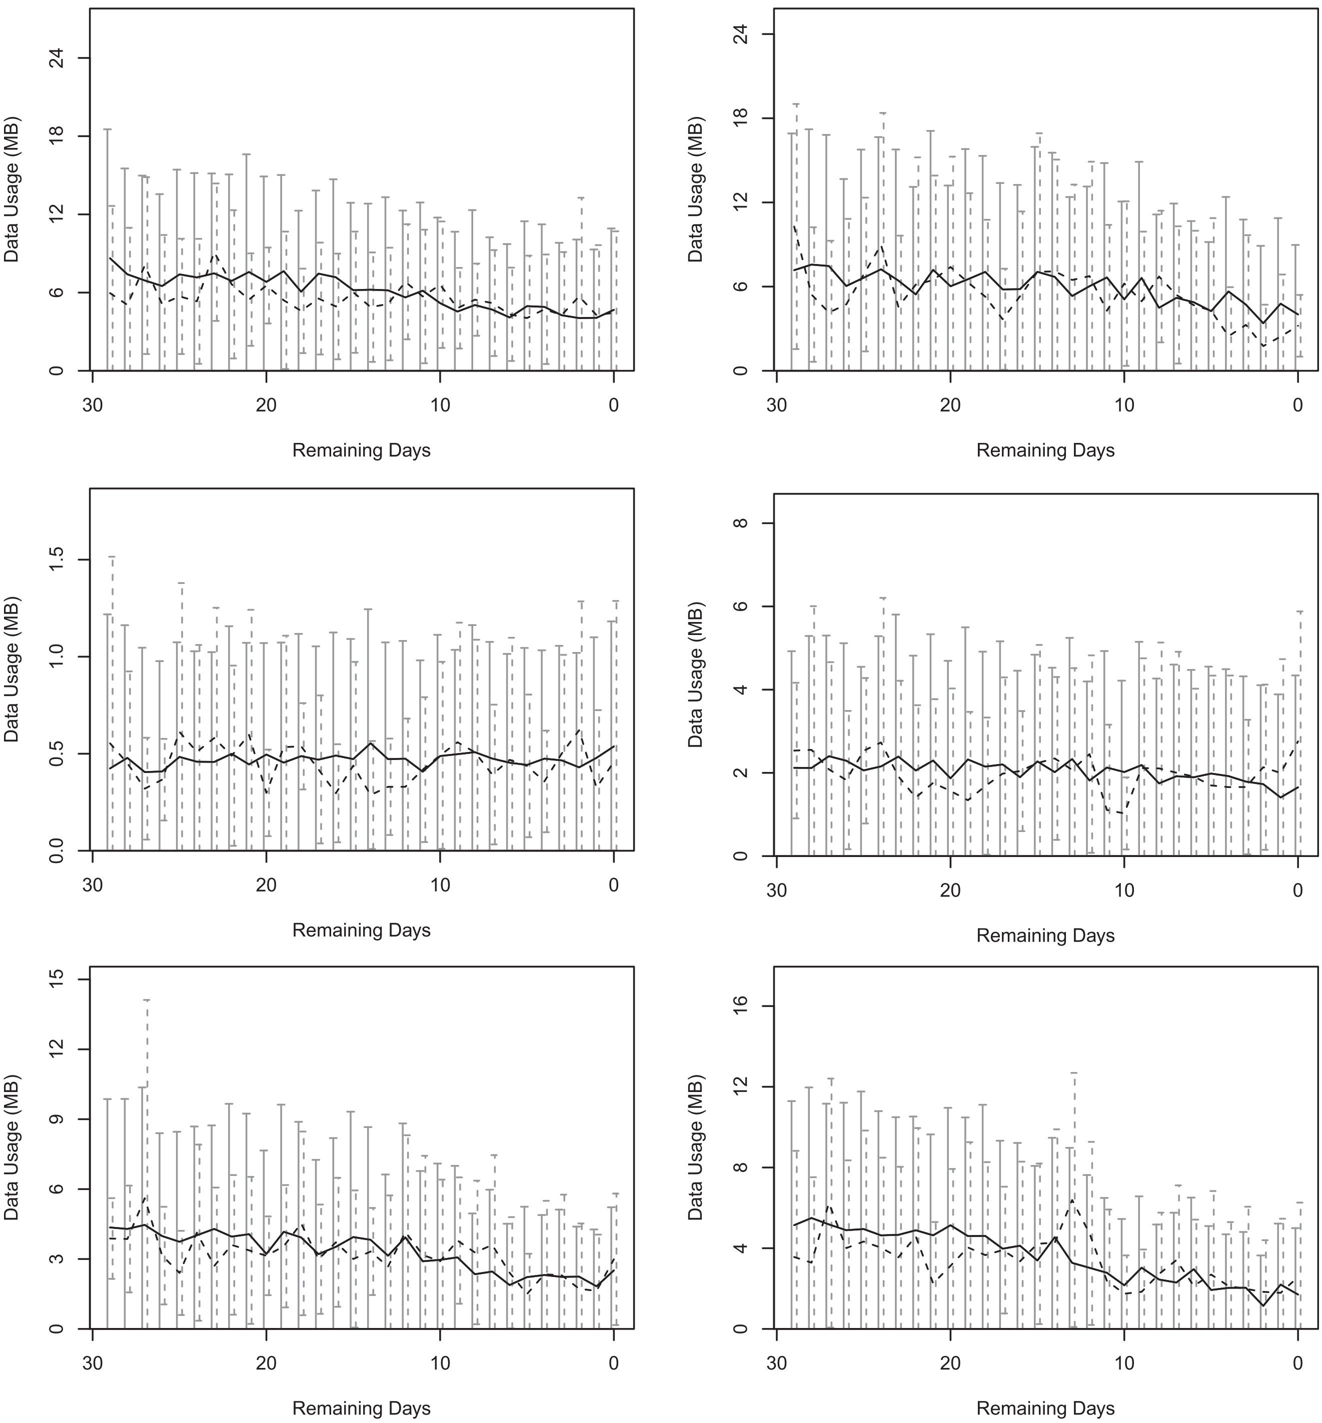

whereas the dashed curves represent the average daily usage over the nine months from the actual data. The solid vertical bars in gray are the 95% intervals constructed from 100 rounds of simulated data (using the 2.5% and 97.5% empirical quantiles), and the dashed vertical bars represent the variation of daily usage in the actual data (constructed using two standard deviations above and below the mean to approximate a 95% interval). As we can see, the average simulated usage falls into the intervals of actual data usage, and the average actual usage also falls in the intervals of the simulated data for almost all days. The trajectories of the average simulated daily usage and actual daily usage closely approximate each other, indicating goodness of fit of the estimation in general. It is especially noteworthy that the simulated usage paths capture the dynamic trends observed in the actual data very well. Recall that by weighing immediate utility against future utility, forward-looking users restrain their data usage when the remaining quota is less because the marginal benefit of saving for the future becomes higher. Such intertemporal substitution results in a smooth and gradual decrease in daily usage over time toward the end of the month, as we observe from the actual usage in Figure 5. The simulated usage well captures such intertemporal trends, which validates our proposed model for the dynamic consumption of mobile data.

We are also interested in how the estimated structural parameters relate to users' demographic information. Thus, we regress the estimated $\{\mu_i, \sigma_i, \eta_i\}$ on user demographics, including age, gender, and customer history (which is converted into the unit of 10 months). To account for the possible correlation among the three estimates for each individual, we apply the seemingly unrelated regression (SUR) so that a system of equations is estimated with the estimated $\{\mu_i, \sigma_i, \eta_i\}$ as the vector dependent variable. Table 2 shows the results. Somewhat consistent with the general notion that demographic information is not most informative in predicting user behavior, we find that neither gender nor customer history has any significant effects on the estimated mean usage level μ_i, the standard deviation of the daily random shock σ_i, or the price sensitivity η_i. On the other hand, age is found to have significantly negative effects on η_i and σ_i, which indicates that older users generally have lower price sensitivity and less daily variation in usage. The results can be explained by the fact that income increases with age in general, so the price sensitivity is lower for older users accordingly. Meanwhile, older users are likely to have more stable usage habit, resulting in lower usage variance in general.

7. Policy Implications

Having estimated individual-specific structural parameters, which are invariant to policy changes, we are able to predict each individual user's rational response to various new pricing policies that mobile carriers may be interested in implementing. In particular, the focal carrier in our study is especially interested in exploring promotions near the end of a month, when users are likely to have expended their quotas, and the overall data usage typically drops. Effective month-end promotions can help the carrier make full use of the excess bandwidth capacity in their mobile network to increase both revenue and customer satisfaction at little extra cost.

In this section, we conduct counterfactual simulations to examine the effectiveness of month-end promotions. As an illustration of how our model can flexibly account for different scenarios to deliver widely applicable implications, we study two types of promotions, each focusing on different aspects. In the first counterfactual study, we demonstrate how offering the same promotion to forward-looking and myopic users could lead to completely different outcomes. In the second study, we further illustrate the impacts on the promotion outcomes when forward-looking users anticipate the coming of the promotion offer.

7.1. Discounts in Overage Charges

Offering discounts in price is perhaps the most common strategies to promote usage or sales. Thus, we first investigate how the focal carrier could increase revenue by offering discounts in the price rates for overage when users are about to deplete their quotas. Such a promotion can be applied to both forward-looking and myopic users, which allows us to explore the optimal promotion designs targeting these two segments of users separately.

We start with targeting the forward-looking users and conduct the following counterfactual simulation study based on our dynamic structural model. We consider the carrier notifies users about a discount, as a percentage δ $(0 < \delta < 1)$ off the price rate for overage, when users are about to deplete their plan quota. Under such a promotion strategy, we simulate each user's usage individual by individual for all of the 392 users estimated using the dynamic model. For each user, we simulate their usage during a 30-day period for 100 rounds. In each round, the user starts using mobile data given the price structure of her chosen data plan. We first solve the expected value functions $\bar{V}(q, d)$ for $d = 1, \ldots, 30$, given the structural parameter estimates, following the backward recursion method described in Section 6.3. We can then simulate this user's usage path by simulating the daily utility shocks and

Table 2. SUR Analysis for Estimated Model Parameters and Demographics

	Estimated μ_i		Estimated $\log \sigma_i$		Estimated $\log \eta_i$	
Demographic var.	Estimate	Std. dev.	Estimate	Std. dev.	Estimate	Std. dev.
Intercept	3.33**	0.48	2.18**	0.17	2.39**	0.30
Age	−0.016	0.015	−0.013*	0.0053	−0.027**	0.0096
Gender	0.21	0.28	0.024	0.10	−0.20	0.18
Customer History	−0.082	0.043	−0.017	0.015	0.019	0.027

$^{**}p < 0.01$; $^{*}p < 0.05$.

computing the optimal usage for each day according to (6). Once the remaining quota falls below a certain threshold (i.e., 10% of each user's total monthly quota as we set for this analysis), the user is notified of the discounted rate. Notice that forward-looking users start to adjust their usage immediately upon being notified of the discounted rate, even if they still have some quota left and do not incur overage charges right away. Therefore, we need to recalculate the expected value functions $\bar{V}(q, d)$ using the new discounted rate $p'_j = (1 - \delta)p_j$ and find the new optimal usage for the rest days till the end of the month. We can then calculate the total usage simulated over the 30-day period and the total extra charge (beyond the fixed plan fees) incurred. Averaging the extra charges simulated over the 100 rounds for each user and then over all 392 users, we can obtain the expected average (extra) revenue under a certain promotion strategy. The leftmost panel of Table 3 shows the results for different discount levels when we target forward-looking users only.

As Table 3 shows, the expected average revenue increases as discount deepens, indicating that forward-looking users, especially those with high price sensitivity, do respond actively to price reduction by increasing their usage significantly. Facing the usual trade-off between profit margin and sales volume, the revenue follows a typical inverted U curve, reaching its peak around $\delta = 50\%$ and starting to decrease as price goes further down. As the counterfactual analysis shows, a month-end promotion targeting forward-looking users offering 40%–50% off the original overage charges could increase the expected average revenue by as much as 27%.

Next, we examine the optimal promotion design targeting myopic users. We first estimate the structural parameters of myopic users based on a static model, in which each user simply maximizes her per-period utility as in (1). The optimal daily usage can be solved explicitly, and the parameter estimates can be obtained via maximum likelihood estimation. Notice that unlike our proposed dynamic model, which identifies individual user's price sensitivity through intertemporal variation in usage, the static model cannot identify the price sensitivity for those users who have never exceeded their quotas during the observation period. As a result, we are able to estimate the three structural parameters for 340 of the 597 users with myopic usage patterns as discussed in Section 5.2. We can then simulate the usage paths for each of these 340 users according to the static model. Each user is offered the discount rates when she is about to deplete her quota, but myopic users start to adjust their usage only when overage charges actually occur. The expected average revenues when targeting myopic users can be calculated in a similar fashion as described above and are summarized in the central panel of Table 3.

In stark contrast to the promotion outcomes targeting forward-looking users, offering discounts in overage rate to myopic users turns out to adversely affect the carrier's revenue. This is because a large number of myopic users are not price sensitive and do not increase their usage significantly in response to the price reduction. As a result, the greater the price discount, the more revenue loss the carrier incurs, which explains the revenue decreases shown in the central panel of Table 3.

We further examine the revenue implications of offering discounts to both forward-looking and myopic users at the same time. As the rightmost panel of Table 3 shows, offering the promotion to the two user segments altogether cannot help increase the overall revenue, simply because the revenue gain from the forward-looking users is canceled out by the revenue loss from the myopic users. Combining the three sets of results, we provide sharp policy implications for mobile carriers. On one hand, the little room of revenue improvement by adjusting the price uniformly charged to all users justifies the current price setting of the focal carrier. On the other hand, we demonstrate remarkable revenue potential by targeting the segment of users who are likely to actively respond to month-end promotions. In this regard, the method we proposed in this paper to identify such a user segment provides mobile carriers with a tremendously valuable tool that is both rigorous in theory and easy to apply in practice.

7.2. Month-End Unlimited-Use Passes

The mobile carrier is also interested in another type of month-end promotion, which offers unlimited-use

Table 3. Results of Counterfactual Analysis 1: Discounts in Overage Charges

	Forward looking		Myopic		Both segments	
Discount (δ)	Avg. rev.	% increase	Avg. rev.	% increase	Avg. rev.	% increase
0%	5.66	n/a	6.21	n/a	5.92	n/a
10%	6.12	8.1	5.74	–7.6	5.94	0.3
20%	6.56	15.9	5.24	–15.6	5.95	0.5
30%	6.93	22.4	4.72	–24.0	5.90	–0.3
40%	7.18	26.9	4.18	–32.7	5.78	–2.4
50%	7.21	27.4	3.60	–42.0	5.53	–6.6
60%	6.92	22.3	2.98	–52.0	5.09	–14.0

passes that grant users unlimited mobile data usage during the last several days of a month with a fixed fee. Notice that only through a dynamic model can the effectiveness of such a promotion strategy be properly assessed. This is because we need to evaluate how likely each user will accept the promotion offer, and users make their acceptance or rejection decisions by comparing their expected future utility in both cases, which can be evaluated only by properly modeling users' forward-looking decision processes. We therefore conduct the second counterfactual simulation study focusing on forward-looking users. In addition, if the carrier repeatedly offers the promotion, users may form expectations about it and adjust their usage and even plan choices accordingly. Therefore, we would also like to illustrate how the possibility that users may anticipate the coming of the promotion offer could impact the promotion effects. In what follows, both scenarios in which users do not and do anticipate the promotion are studied.

We consider the mobile carrier sends out a promotion on the nth day to the end of a month. The promotion provides users with the option to purchase a month-end unlimited-use pass such that after paying a fixed fee P, users can freely use mobile data till the end of the month without any additional charge, regardless of the remaining quota from their original data plan. If a user chooses to decline the offer, she continues with her original data plan and pays for any extra usage according to the plan rate.

Starting with the case when users do not anticipate the promotion, we simulate the behavior of each of the 392 forward-looking users following a similar approach as in Section 7.1. For each user, we simulate a 30-day usage path for 100 rounds. In each round, we simulate a user's daily usage for the first $(30-n)$ days according to the original dynamic programming problem given her chosen data plan and the structural parameter estimates. At the beginning of the $(30-n+1)$th day, the user is notified of the promotional offer and decides whether to accept it. If she chooses to accept the offer, the expected utility can be derived as

$$\bar{V}^{Accept}(n,P) = \sum_{\tau=0}^{n-1} \beta^{\tau} E_{\xi_{it}} \frac{1}{2}(\max\{\mu_i + \xi_{it}, 0\})^2 - \eta_i P. \quad (15)$$

The first term of (15) represents the sum of discounted expected utility from unconstrained data usage over the last n days of a month; the second term represents the disutility associated with the fee of the pass. A user accepts the promotional offer if and only if $\bar{V}^{Accept}(n,P) > \bar{V}(q_{it}, n)$, where $\bar{V}(q_{it}, n)$ is the expected utility from continuing with the original plan as in (5). Notice that whereas $\bar{V}^{Accept}(n,P)$ is independent of the remaining quota q_{it}, $\bar{V}(q_{it}, n)$ does depend on q_{it}. As a result, whether a user will accept the promotional offer depends on her usage history in the earlier part of the month and is hence probabilistic. Thus, we can determine a user's acceptance decision and continue to simulate her usage for the last n days, which yields the total surcharge (beyond the monthly plan fee) a user pays in each simulation round. We can then compute the average expected revenue and the average probability of accepting the offer by averaging over all simulation rounds and all users. The left panel of Table 4 shows the results under different values of n and P.

As Table 4 (the "No anticipation" panel) shows, the average acceptance rate of the promotional offer increases as the fee P reduces or the number of days n increases. For the same n, the average revenue follows an inverted U curve as the fee increases, which can be easily explained by the usual trade-off between price and demand. In comparison with the average revenue without any promotion (i.e., ¥5.83), selling the unlimited-use pass at a price too low (e.g., a five-day pass for ¥5.00) can result in a high acceptance rate but may not help increase the revenue at all. Instead, a properly priced promotion (e.g., a four-day pass for ¥10.00) can increase the expected revenue by as much as 20%.

Table 4. Results of Counterfactual Analysis 2: Month-End Unlimited-Use Passes

		No anticipation			Anticipation		
Last n days	Fee P	Accpt. rate	Avg. rev.	% increase	Accpt. rate.	Avg. rev.[†]	% increase
3	5.00	0.4148	6.46	10.8	0.4345	6.49	11.3
3	10.00	0.1511	6.62	13.4	0.1604	6.66	14.1
3	15.00	0.0384	6.20	6.3	0.0393	6.21	6.5
4	5.00	0.4569	6.17	5.7	0.4960	6.36	9.0
4	10.00	0.2569	7.02	20.3	0.2698	7.11	21.9
4	15.00	0.0766	6.48	11.1	0.0809	6.51	11.5
5	5.00	0.4893	5.85	0.19	0.5361	6.06	3.8
5	10.00	0.3055	6.98	19.6	0.3240	7.08	21.4
5	15.00	0.1421	6.88	18.0	0.1521	6.94	19.0

[†]Extra revenue net of the revenue loss in monthly plan fee as a result of a plan switch.

Next, we further extend our counterfactual analysis to the scenario when users are assumed to anticipate the promotional offer.[10] Note that the dynamic programming problem is different in this scenario, and hence we need to solve for a new optimal solution. The key difference arises in the day preceding the arrival of the promotion—that is, the $(n+1)$th day to the end of the month (i.e., $d_t = n+1$). The new value function, $V^{Ant}(q_{it}, n+1, \xi_{it})$, incorporates the possible acceptance of the promotion and thus can be reformulated as

$$\begin{aligned} V^{Ant}&(q_{it}, n+1, \xi_{it}) \\ &= \max_{a_{it}\geq 0}[u(a_{it}; q_{it}, \xi_{it}) \\ &\quad + \beta \max\{\bar{V}(\max\{q_{it}-a_{it}, 0\}, n), \bar{V}^{Accept}(n, P)\}]. \end{aligned} \tag{16}$$

Consequently, the expected value function can be obtained by taking expectation over ξ_{it} such that $\bar{V}^{Ant}(q_{it}, n+1) = E_{\xi_{it}} V^{Ant}(q_{it}, n+1, \xi_{it})$. This revised expected value function then enters the dynamic programming problem for all the days preceding $d_t = n+1$. In this way, for each proposed promotion design (n, P), we re-solve the associated new dynamic programming problem backward recursively for each user. Before simulating the daily usage, we compare the expected utility from choosing each of the available data plans at the beginning of a month to allow the possibility of plan switching. A user chooses data plan j^* that yields the highest expected utility such that

$$j^* = \arg\max_j [\bar{V}^{Ant}(Q_j, T) - \eta_i F_j]. \tag{17}$$

We then simulate each user's daily usage based on $\bar{V}^{Ant}(q, d)$ under the chosen plan. The expected promotion outcomes can be obtained in a similar way as previously done and are summarized in the right panel of Table 4.

As Table 4 (the "Anticipation" panel) indicates, the expected promotion outcomes remain qualitatively unchanged compared with the case without user anticipation. With user anticipation, two counteracting effects come into play. On one hand, users tend to slightly increase their usage in the earlier part of a month because they anticipate the option of the unlimited pass if they use up the quota before the month ends. On the other hand, a small portion of the users may switch to a smaller plan in the very beginning because the availability of the promotion reduces the necessity of a larger plan. The two effects impact the carrier's revenue in the opposite directions and mostly cancel each other out. The overall effect is a slight increase in the expected revenue, whereas the relative magnitude of the promotion outcomes (e.g., acceptance rate, revenue increase) across various promotion designs remains unchanged. Therefore, the policy implications on the promotion effectiveness as we derived are robust regardless of whether users may anticipate the promotion.

8. Conclusion

With the proliferation of modern mobile devices and the expansion of fast-speed data networks nowadays, our paper addresses the important yet underexamined issue of dynamic consumption behavior of mobile data users and provides implications for both researchers and practitioners.

As we show, forward-looking users, who value future utilities in determining present consumption and hence dynamically manage their usage over time, are found to be more responsive to various promotions. Therefore, promotions designed to target these users can be remarkably more profitable than the blanket promotions uniformly applied to all consumers. Properly identifying forward-looking consumers is therefore an imperative issue of not only academic importance but also tremendous business value.

Our study develops a novel and practical method to identify forward-looking consumers. Building on a rigorous theoretical foundation, we uncover distinct usage patterns for forward-looking and myopic behaviors, which have been absent in the literature. These patterns are easily testable with observational data by examining the presence of a positive correlation between the daily consumption and the remaining amount of the prespecified quota. In this sense, our study not only deepens the theoretical understanding of dynamic consumption behaviors but also provides a useful instrument for business practitioners to easily implement.

Enriching the academic literature, to our knowledge we are the first in the literature to conduct a formal test for forward-looking behavior at the individual consumer level. Our empirical findings provide evidence of individual forward-looking behavior and show that a considerable portion of all consumers (i.e., about 40% in our data) are indeed forward looking. Meanwhile, we also find evidence that not all users are forward looking, although consumers collectively are found to be forward looking in the previous literature. Moreover, we propose a novel approach to control for possible consumer myopia in estimating dynamic structural models by matching the temporal patterns with the theoretical properties derived from the structural model. This approach offers a valuable option in the face of the inherent difficulty in identifying individual discount factors in dynamic optimization problems.

In addition, we illustrate the great potential of dynamic promotions to be sent to the right group at the right time. As we show, forward-looking users usually restrain their usage toward the end of a month, leaving excess bandwidth capacity in carriers' mobile data network. Properly designed month-end promotions targeting forward-looking users can increase carriers' revenue at little extra cost. The dynamic structural model we construct and the associated estimation and simulation methods furnish a helpful tool to identify such revenue opportunities. Our model is shown to well capture the dynamic trends observed in the actual usage data. Our estimation method is computationally efficient and can be easily parallelized, which enables large-scale implementation. The counterfactual simulations can account for complex factors (e.g., consumer anticipation, plan switch) to deliver various reliable implications in assessing the design of different dynamic promotion strategies. It thus opens a new front of marketing opportunities that offer customized prices and products to targeted consumers based on dynamic past consumption behaviors, which goes beyond the prevailing industry practice of static targeting using demographic data and simple usage statistics.

For future research, opportunities abound in the flourishing area related to dynamic mobile data consumption. For example, if we could have the right type of data from field or natural experiments (e.g., if we had data on consumer usage before and after a carrier changes its plan and fee structures), it might be possible to design a method to estimate individual discount factors, which would be a more precise way to quantify forward-looking behaviors. Moreover, beyond identifying the segment of forward-looking consumers through easy-to-implement tests as we propose, if one could further develop simple heuristics to identify within the forward-looking segment smaller groups or even individuals who are more responsive to promotions than others, it would open up greater potential for revenue improvement through individual targeting in business practice. We hence believe our study serves as a start for a variety of promising future research directions.

Acknowledgments

The authors thank the department editor, the associate editor, and the two anonymous reviewers for their constructive and insightful comments. The authors also thank the seminar participants at Purdue University, the University of Notre Dame, Georgia State University, the University of Georgia, Emory University, the City University of Hong Kong, Erasmus University Rotterdam, the Georgia Institute of Technology, the University of Florida, the University of Texas at Dallas, George Washington University, Johns Hopkins University, the University of Minnesota, the University of Texas at Austin, Indiana University, the University of California at San Diego, Shanghai Jiao Tong University, the University of Electronic Science and Technology of China, Dalian University of Technology, Nanjing University, Beijing University of Posts and Telecommunications, the Workshop on Information Systems and Economics (2014), the China Summer Workshop on Information Management (2015), the Conference on Information Systems and Technology (2015), and the ISMS Marketing Science Conference (2016) for their helpful feedback. Y. Cheng is the corresponding author.

Endnotes

[1] This company is among the largest mobile carriers in the world. We are unable to reveal the identity of the company because of a nondisclosure agreement.

[2] This city has a population of over 10 million and an annual gross domestic product (GDP) of over US$140 billion. It is ranked among the top 20 cities in China by either population or GDP.

[3] Some main service packages come with an allowance of free data as well. In these cases, a user can consume up to the combined total of the quotas included in both the main package and the data plan before incurring any extra cost. Note that we do include the data quota contained in the main service package in our analysis, but we take users' main service packages as exogenously given and do not model users' choice of main service package. This is because most components of a main service package are beyond the scope of this study, and there is a labyrinthine selection of more than 100 different packages.

[4] The rate is ¥1.00 per extra megabyte regularly, whereas some users' main service packages offer discounts in the extra rate (e.g., ¥0.29/MB). Therefore, depending on their main service packages, different users may have different effective extra rates. In any case, for a given user, the extra rate stays the same across different data plans.

[5] We explicitly account for different numbers of days in different months (e.g., 28 days for February and 31 days for March).

[6] We have also estimated a model with a quadratic cost function to account for potential risk aversion. We find the coefficient of the quadratic term to be statistically insignificant for almost all users and therefore choose to keep our model parsimonious.

[7] We also tried alternative values of β (e.g., 0.99, 0.95, 0.85, 0.80). It is found that the estimation results are generally insensitive to the value of β, and $\beta = 0.9$ generates the best model fit relatively.

[8] Note that $F_i(a_{it}|q_{it}, d_t)$ has a mass point at $a_{it} = 0$. Therefore, $\forall \xi_{it}$ such that $G_i(\xi_{it}) \leq F_i(0|q_{it}, d_t)$, we define $F_i^{-1}(G_i(\xi_{it})|q_{it}, d_t) = 0$ in (9).

[9] We exploit the fact that the utility function (1) is linear in all model parameters (we can rewrite ξ_{it} as $\sigma_i \xi_{0t}$, where $\xi_{0t} \sim N(0, 1)$). Therefore, in computing each $\bar{U}_i(q, d; a)$, we only need to simulate the n_S paths once independent of the parameter values, which significantly expedites the estimation.

[10] We consider complete information here so that users fully anticipate both the day n and the fee P of the promotion being implemented.

References

Abbring J, Daljord O (2018) Identifying the discount factor in dynamic discrete choice models. Becker Friedman Institute for Economics Working Paper 2017-17, University of Chicago, Chicago.

Arcidiacono P, Miller RA (2011) Conditional choice probability estimation of dynamic discrete choice models with unobserved heterogeneity. *Econometrica* 79(6):1823–1867.

Bajari P, Benkard CL, Levin J (2007) Estimating dynamic models of imperfect competition. *Econometrica* 75(5):1331–1370.

Busse MR, Knittel CR, Zettelmeyer F (2013) Are consumers myopic? Evidence from new and used car purchases. *Amer. Econom. Rev.* 103(1):220–256.

Chevalier J, Goolsbee A (2009) Are durable goods consumers forward-looking? Evidence from college textbooks. *Quart. J. Econom.* 124(4):1853–1884.

Chung DJ, Steenburgh T, Sudhir K (2014) Do bonuses enhance sales productivity? A dynamic structural analysis of bonus-based compensation plans. *Marketing Sci.* 33(2):165–187.

Erdem T, Keane MP (1996) Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. *Marketing Sci.* 15(1):1–20.

Frederick S, Loewenstein G, O'Donoghue T (2002) Time discounting and time preference: A critical review. *J. Econom. Literature* 40(2): 351–401.

Fritsch FN, Carlson RE (1980) Monotone piecewise cubic interpolation. *SIAM J. Numer. Anal.* 17(2):238–246.

Gabaix X, Laibson D (2006) Shrouded attributes, consumer myopia, and information suppression in competitive markets. *Quart. J. Econom.* 121(2):505–540.

Ghose A, Han SP (2011) An empirical analysis of user content generation and usage behavior on the mobile internet. *Management Sci.* 57(9):1671–1691.

Goettler RL, Clay K (2011) Tariff choice with consumer learning and switching costs. *J. Marketing Res.* 48(4):633–652.

Hoch SJ, Loewenstein GF (1991) Time-inconsistent preferences and consumer self-control. *J. Consumer Res.* 17(4):492–507.

Hotz VJ, Miller RA (1993) Conditional choice probabilities and the estimation of dynamic models. *Rev. Econom. Stud.* 60(3):497–529.

Huang Y, Singh PV, Ghose A (2015) A structual model of employee behavioral dynamics in enterprise social media. *Management Sci.* 61(12):2825–2844.

Iyengar R, Ansari A, Gupta S (2007) A model of consumer learning for service quality and usage. *J. Marketing Res.* 44(4):529–544.

Judd K (1998) *Numerical Methods in Economics* (MIT Press, Cambridge, MA).

Kim Y, Telang R,Vogt WB, Krishnan R (2010) An empirical analysis of mobile voice service and SMS: A structural model. *Management Sci.* 56(2):234–252.

Magnac T, Thesmar D (2002) Identifying dynamic discrete decision processes. *Econometrica* 70(2):801–816.

Manski CF (1993) Dynamic choice in social settings: Learning from the experiences of others. *J. Econom.* 58(1–2):121–136.

Misra S, Nair HS (2011) A structural model of sales-force compensation dynamics: Estimation and field implementation. *Quant. Marketing Econom.* 9(3):211–257.

Niculescu MF, Whang S (2012) Codiffusion of wireless voice and data services: An empirical analysis of the Japanese mobile telecommunications market. *Inform. Systems Res.* 23(1):260–279.

Pashardes P (1986) Myopic and forward looking behavior in a dynamic demand system. *Internat. Econom. Rev.* 27(2):387–397.

Rust J (1987) Optimal replacement of gmc bus engines: An empirical model of Harold Zurcher. *Econometrica* 55(5):999–1033.

Segan S (2014) The death of unlimited mobile data. *PC Magazine* (March 10), https://www.pcmag.com/article2/0,2817,2454764,00.asp.

Tibshirani R, Walther G, Hastie T (2001) Estimating the number of clusters in a data set via the gap statistic. *J. Roy. Statist. Soc. Ser. B* 63(2):411–423.

Urminsky O, Zauberman G (2015) The psychology of intertemporal preferences. Keren G, Wu G, eds. *The Wiley Blackwell Handbook of Judgment and Decision Making* (John Wiley & Sons, Chichester, UK), 141–181.

Worthy DA, Otto AR, Maddox WT (2012) Working-memory load and temporal myopia in dynamic decision-making. *J. Experiment. Psych.: Learn. Memory Cognition* 38(6):1640–1658.

Xu J, Forman C, Kim JB, van Ittersum K (2014) News media channels: Complements or substitutes? Evidence from mobile phone usage. *J. Marketing* 78(4):97–112.

Yao S, Mela CF, Chiang J, Chen Y (2012) Determining consumers' discount rates with field studies. *J. Marketing Res.* 49(6):822–841.