Portfolio Dynamics for Customers of a Multiservice Provider

David A. Schweidel

Wisconsin School of Business, University of Wisconsin–Madison, Madison, Wisconsin 53706, dschweidel@bus.wisc.edu

Eric T. Bradlow, Peter S. Fader

The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104
{ebradlow@wharton.upenn.edu, faderp@wharton.upenn.edu}

Multiservice providers, such as telecommunication and financial service companies, can benefit from understanding how customers' service portfolios evolve over the course of their relationships. This can provide guidance for managerial issues such as customer valuation and predicting customers' future behavior, whether it is acquiring additional services, selectively dropping current services, or ending the relationship entirely. In this research, we develop a dynamic hidden Markov model to identify latent states that govern customers' affinity for the available services through which customers evolve. In addition, we incorporate and demonstrate the importance of separating two other sources of dynamics: portfolio inertia and service stickiness. We then examine the relationship between state membership and managerially relevant metrics, including customers' propensities for acquiring additional services or terminating the relationship, and customer lifetime value. Through a series of illustrative vignettes, we show that customers who have discarded a particular service may have an increased risk of canceling all services in the near future (as intuition would suggest) but also may be more prone to acquire more services, a provocative finding of interest to service providers. Our findings also emphasize the need to look beyond the previous period, as in much current research, and consider how customers have evolved over their *entire* relationship in order to predict their future actions.

Introduction

Among the many questions that multiservice providers frequently ask about their customers are which customers are most at risk for terminating (or enhancing) their relationships and which customers are the most valuable. In addressing these questions, firms commonly consider the length and breadth of customers' relationships (e.g., Bolton et al. 2004) as good "rules of thumb" from which appropriate CRM (customer relationship management) decisions can be made. Though such measures have seen widespread use in practice, they fail to account for the full nature of how customers have evolved during their relationships.

Consider three hypothetical customers of a multiservice provider (e.g., a telecommunication or financial services company) who began service at the same time and *currently* subscribe to the same set of services. The first customer has maintained this portfolio since beginning service, and the second has gradually acquired services over time to get to this same portfolio. The third customer, in contrast, started with a large portfolio and has slowly discarded services. Though all three currently subscribe to the same

portfolio, is it fair to assume that they have the same future value to the firm? Probably not. The firm's valuation of these customers, as well as related diagnostics such as predictions of who is most likely to defect in the next month, should consider the entire path of portfolio decisions to date. Though it seems logical that customers' entire past behavior would yield insight into expected future behavior, simple heuristics and much extant research ignore this information.

Although methods have been developed to value customers in contractual settings and to estimate the likelihood that they will cancel or acquire individual services (e.g., Fader and Hardie 2010; Schweidel et al. 2008a, b), there are additional considerations unique to the multiservice context that we consider here. First, correlation patterns may exist among the likelihoods of adopting and discarding the available services (e.g., Li et al. 2005). Second, tendencies to adopt and discard services may change over the course of a customer's relationship. A telecommunications customer may "ramp up" early in his relationship to sample the services, whereas a veteran customer may progress through a "staged death" and slowly discard services (Smurl 2002). Rather than devoting resources

to these waning customers, it may be more profitable to focus efforts on new prospects, providing them with incentives to begin (and broaden) their relationship. More generally, customers may progress through a set of latent states as their relationships evolve, over which time their affinities for services shift.

In addition to underlying changes in customers' service affinities over time, portfolio choices may be subject to additional dynamics. Customers may exhibit *portfolio inertia*, as they may not deliberate about their service portfolio each month and simply maintain it for a lengthy duration without considering any changes. The extent of portfolio inertia may depend on the duration for which the current portfolio has been maintained, as customers may be more prone to keep a portfolio without considering alternatives the longer they have already subscribed to it. When customers reevaluate their portfolios, another source of state dependence may arise, as current subscribers of a particular service may be more likely to retain it than a nonsubscriber would be to acquire it, thus exhibiting *service stickiness*—selectively adding or dropping particular services while maintaining most of the services they previously owned.[1]

Though customers might maintain all or most of their services from one month to another (due, perhaps, to portfolio inertia or service stickiness), they may still move among the different latent states. In other words, customers' baseline service affinities may change although the observed portfolios do not. Taken together, these dynamics influence customers' monthly portfolio choices and ultimately their value to the service provider. There is lore that "broader" customers are more valuable, yet we contend that "how you got there" also matters; that is, one must look back beyond the current period to see how a customer's portfolios have evolved over time and may continue to evolve going forward.

To this end, we develop a dynamic hidden Markov model (HMM) to capture the evolution of customers' subscription propensities over time. This lets us update our beliefs about customers' propensities of adding or dropping services based on their observed portfolio history, an application of importance to managers interested in determining the value of their customers. The remainder of this paper proceeds as follows. We review the related literature and describe the data used in our empirical application. We then develop a portfolio choice model that allows for evolution over the course of customers' relationships, portfolio inertia and service stickiness. This is followed by a discussion of the empirical results and a demonstration of the managerial insights into future behavior and customer value afforded by our framework. We conclude with a discussion of our contributions and directions for future research.

Previous Research and Managerial Relevance

The growing literature on customer equity proposes that marketing decisions should be linked to financial metrics, such as customer lifetime value (CLV) (e.g., Rust et al. 2004, Bolton et al. 2004, Blattberg et al. 2001). In applying the customer equity framework to multiservice providers, the duration of the relationship and the future acquisition of additional services (cross-buying) must be considered, as they represent central components of CLV. As mentioned earlier, previous research has posited that customer value increases with both the duration of the relationship and the number of services that a customer has acquired. We believe this proposition deserves further scrutiny and should not be accepted at face value.

Rust et al. (2004), for example, link increased customer retention to higher CLV. Although the authors do not study the effects of cross-selling on customer equity, they note its importance to multiservice providers. Bolton et al. (2004) present a framework that links antecedents such as price, relationship marketing instruments, and service quality to the length and breadth of the relationship, which in turn influences CLV. Though the authors account for the impact of acquiring multiple services, they do not fully explore how the acquisition and retention of services may relate to customers' decisions to terminate the relationship. For example, some customers may be in the process of gradually discarding services but have a low risk for ending their relationship, whereas customers who have recently acquired more services may temporarily be at greater risk. Such patterns may change throughout a customer's relationship, necessitating a dynamic model of customer behavior that can be linked to customer valuation and related marketing decisions (Hogan et al. 2002).

Donkers et al. (2007) compare different methods that can be used to determine the CLV of customers in a multiservice setting. The authors find that customers are more prone to purchase a service if they had it in the previous year (i.e., the aforementioned service stickiness). Although the authors find evidence of service stickiness, they do not consider the effects of portfolio inertia or latent changes in customers' baseline service affinities, which will impact

---

[1] This may occur for a number of reasons, such as current subscribers learning about the service (e.g., Iyengar et al. 2007) or other forms of state dependence (e.g., Erdem 1996, Roy et al. 1996). Although we account for this effect, we do not attempt to distinguish between these potential explanations. We discuss this as a direction for future research with an appropriate data set.

a customer's value to the firm and should hence be incorporated into estimates of CLV.

We extend previous literature by developing a dynamic model of customers' portfolio choices at a multiservice provider, which we then link to metrics including CLV and the likelihood of ending the relationship. We employ an HMM (e.g., Netzer et al. 2008, Fader et al. 2004, Schmittlein et al. 1987) to capture latent shifts in customers' baseline preferences for services, as well as their risk for ending their relationship. In addition to generalizing univariate HMMs to a multivariate context (e.g., Du and Kamakura 2006), our framework incorporates additional sources of dynamics that may influence customers' subscription decisions and consequently their value to the service provider.

HMMs have been used extensively in CLV calculations, though primarily in single-product transactional exchanges. Much research in customer value stems from the Pareto/NBD (Schmittlein et al. 1987), an HMM that assumes that customers are either active or inactive. Reinartz and Kumar (2003) use the Pareto/NBD to determine the most valuable customers. Fader et al. (2005) employ the Pareto/NBD to compute CLV and generate isovalue curves to examine how different purchase histories can lead to similar valuations. Venkatesan and Kumar (2004) demonstrate how CLV can be employed to develop resource allocation strategies. Although they find that cross-buying is associated with increased purchase frequency (and hence CLV), the authors consider cross-buying as exogenous and do not consider how such behavior may shift over time.

Accounting for the aforementioned dynamics, we estimate CLV for new customers and examine how customers' remaining value may change over time. Using a series of vignettes, we demonstrate how customers' full histories can be used to identify those who have more cross-selling potential or are at a greater risk for ending their relationships—even among customers who currently subscribe to the same portfolio. Such metrics can be used to customize marketing efforts and for tracking purposes in managerial dashboards (e.g., O'Sullivan and Abela 2007).

Data

Data were provided by a major telecommunications firm that chose to remain anonymous. Monthly subscription information was provided from one geographic region of the United States for January 2002 through May 2004, indicating the services to which customers subscribed each month. The services include basic cable, a digital cable package, premium channels (HBO, Showtime, Starz, Cinemax, and TMC), and high-speed Internet service. We constructed a calibration data set by taking a random

sample of customers who began service between February and September 2002 and subsequently tracked them through January 2004, yielding a sample of 3,393 customers. The first 24 months of observations (February 2002 through January 2004) were used for calibration. Depending on when they began service, customers were under calibration observation for 17–24 months. The remaining four months were used for out-of-sample validation.

Exploratory Analyses

During the calibration period, we observe a general trend of services being discarded (with the exception of high-speed Internet service), with approximately 60% of customers not subscribing to any services at the end of this period. Examining the aggregate number of subscribers offers a logical summary, but it yields limited insights into both copurchasing patterns and acquisition and retention processes.

To address these limitations, we also examined retention of services by initial subscribers. Of the initial subscribers to premium channels, approximately 10% discarded services during the three-month introductory period, and approximately 50% discarded them within five months and approximately 70% discarded them within one year. Original subscribers to basic cable, the digital package, and Internet service also discarded services, which is potentially indicative of the relationship being terminated, though at a slower rate.

To examine copurchasing behavior and as part of a strong empirical benchmark in our analyses, we present the 10 most popular portfolios (containing at least one service) in descending popularity order in Table 1.

These portfolios account for 86.9% of the "customer-month" combinations in the calibration period with at least one active service. It appears that Internet service is not copurchased commonly with any of the add-ons, whereas Showtime and TMC, and digital cable service and HBO, are commonly copurchased. Table 1 offers a snapshot of copurchasing behavior

Table 1   Most Popular Portfolios

Portfolio popularity	Proportion observed (%)	Basic cable	Digital service	Internet service	HBO	Cinemax	SHO	Starz	TMC
1	36.1	✓							
2	12.3	✓	✓		✓				
3	8.8	✓	✓						
4	6.7	✓	✓		✓	✓	✓	✓	✓
5	6.4	✓			✓				
6	4.1	✓		✓					
7	3.6	✓	✓		✓		✓		✓
8	3.4			✓					
9	2.9	✓	✓		✓	✓			
10	2.4	✓	✓		✓			✓	

during the calibration period, but it does not reveal how these trends may change over time.

During the calibration period, we also observe that subscribers made between 0 and 10 portfolio changes (average = 1.38, s.d. = 1.24), with approximately 21% of subscribers choosing to maintain their initially chosen portfolio and 46% making a single portfolio change. Of those who did not change their portfolio during the calibration period, 49% only subscribed to basic cable and 6% subscribed to the portfolio containing basic cable, digital cable, and HBO. Contrasting these proportions with those observed in Table 1 highlights (and portends the need for models that account for) the potential dynamics that may influence subscribers' portfolio choices and ultimately their value to the service provider.

In a longitudinal holdout analysis, we see that many subscribers had ceased subscribing to any services prior to the start of the holdout period. Among the customers who still subscribed to at least one service (40% of our sample), 79% made no changes during the holdout period and 18% made a single change. Given the infrequency of portfolio changes during this period, a simple heuristic "model" that assumes customers maintain the last portfolio to which they subscribed during the calibration period provides a second strong benchmark for model validation (i.e., the "hold and keep" model).

Given the large number of portfolios from which customers can choose and the limited number of portfolio decisions made during the calibration period, it can be difficult to identify dynamics in the choice decisions through purely exploratory methods (e.g., windowing and counting) that rely solely on the observed transitions between portfolios. To overcome this limitation, we next develop a model to identify latent states that are characterized by service affinities and discuss (and later demonstrate) how it can be used to derive managerially relevant measures including CLV and defection risk.

Model Development

To capture temporal changes in customers' choice tendencies, we employ a dynamic HMM. The HMM consists of three parts: the state-specific portfolio choice model, transition matrix, and initial distribution. We next describe each component of the customer-level model.

State-Specific Portfolio Choice Model

In Figure 1, we first provide an overview of the state-specific portfolio choice models we considered that arise from the inclusion or omission of portfolio inertia and service stickiness.

In the absence of portfolio inertia and service stickiness (cell I in Figure 1), a customer's portfolio choice

Figure 1 State-Specific Portfolio Choice Models

Portfolio inertia	Service stickiness	
	Omitted	Included
Omitted	**I** • Multivariate choice every month • State-specific choice probabilities *are not* affected by service subscriptions in the previous month	**II** • Multivariate choice every month • State-specific choice probabilities affected by service subscriptions in the previous month
Included	**III** • Multivariate choice *only when the portfolio is reconsidered* • State-specific choice probabilities *are not* affected by service subscriptions in the previous month	**IV** • Multivariate choice *only when the portfolio is reconsidered* • State-specific choice probabilities affected by service subscriptions in the previous month

is not influenced by the services to which he subscribed in the previous month. By incorporating service stickiness (without portfolio inertia), we allow customers to be more (or less) prone to retaining the individual services to which they subscribed in the previous month (cell II). Although service stickiness offers a potential explanation for the low frequency with which individual services are changed, it does not explicitly recognize customers' tendencies to maintain *all* the same services (i.e., portfolio inertia, cell III), which may be common in many multiservice contexts (e.g., financial services and insurance providers). Allowing for both service stickiness and portfolio inertia (cell IV) recognizes that customers may infrequently reconsider their portfolios and make selective changes when they do so. The difference in the level at which these two dynamics operate (individual service versus portfolio) allows us to distinguish between them conceptually and empirically identify their effects.[2] The distinction between service stickiness and portfolio inertia in a multiservice setting has not been discussed previously and is one contribution of this research.

To describe how we incorporate these components into the state-specific portfolio choice model, we

---

[2] Simulations were conducted in which we varied the number of active states and the extent of portfolio inertia to create scenarios with differing frequencies of portfolio changes while also accounting for service inertia. We considered scenarios with one and two active states and medium and high levels of portfolio inertia. In all cases, the true parameter values were contained in the 90% highest posterior density intervals. Further details are available from the authors upon request.

explicate the model in a sequential manner. We begin by deriving the model for portfolio utility, which is comprised of the utilities of the services that constitute it. These service utilities may be affected by previous subscription to the service (i.e., service stickiness) and a customer's latent affinity for the service, the latter of which depends on his current state. We then incorporate portfolio inertia into the portfolio choice model.

In many contractual relationships, including our empirical context, customers must subscribe to a base service (i.e., basic cable) to subscribe to "add-on" services (i.e., premium channels and a digital package). We let $k = 1$ denote basic cable service, A the set of add-on services, and B the set of services that can be chosen regardless of subscription to basic cable. Letting z index all available portfolios, accounting for the service nesting structure, we let $V_k(z) = 1$ if service k is included in portfolio z and equal 0 otherwise. Conditional on being in state s at time t, the utility of portfolio z for customer h is given by $Q_{hts}(z) + \varepsilon_{htsz}$, where $Q_{hts}(z)$ is

$$Q_{hts}(z) = (V_1(z)) \cdot \left(\underbrace{U_{hts1}}_{\text{base service utility}} + \underbrace{\sum_{a=2}^{1+A} U_{htsa} \cdot V_a(z)}_{\text{add-on utility}} \right)$$
$$+ \underbrace{\sum_{b=2+A}^{1+A+B} U_{htsb} \cdot V_b(z)}_{\text{independent service utility}} \qquad (1)$$

and ε_{htsz} are independent and identically distributed random components of the portfolio utility that follow a Gumbel distribution, and U_{htsk} is customer h's utility for service k at time t conditional on being in state s. The first term of $Q_{hts}(z)$ captures the utility of basic cable; the second term captures the utility from the add-on services in z that require basic cable, and the final term reflects the utilities of services that do not require basic cable. Although we opt for the parsimonious additive structure, Equation (1) can be generalized to include additional utility components that "turn on" for portfolios with certain service combinations, similar to a bivariate logit model (e.g., Niraj et al. 2008), thereby allowing for a flexible covariance structure among the services.

Letting $Y_{htk} = 1$ if customer h subscribes to service k at time t, we specify U_{htsk} as

$$U_{htsk} = \alpha_{sk} + \beta_k \cdot \text{Promo}(t) + \theta_k \cdot 1(Y_{h(t-1)k} = 1), \qquad (2)$$

where α_{sk} is a state-by-service-specific intercept for the affinity of service k in state s, β_k reflects the impact of an introductory three-month promotional offer, and

θ_k accounts for the shift in service utility associated with subscribing to service k at time $t-1$.[3]

Let z_{ht} denote the portfolio that customer h subscribes to at time t. When customer h is in state s and reconsiders the portfolio to which he subscribes, the state-specific portfolio choice probabilities follow a multinomial logit model:

$$r(z_{ht} \mid s) = \frac{\exp(Q_{hts}(z_{ht}))}{\sum_{z' \in Z} \exp(Q_{hts}(z'))}. \qquad (3)$$

Note that $Q_{hts}(z)$ can also be generalized to account for firm-specific service nesting structures.[4]

Customers in a contractual relationship may not actively reconsider their portfolio subscriptions each month but often maintain the portfolio to which they previously subscribed. To account for this notion of portfolio inertia, we assume that the likelihood of customer h subscribing to portfolio z_{ht} in state s is given by

$$P(z_{ht} \mid s) = \begin{cases} (1 - \psi_{ht}) r_{ht}(z_{ht} \mid s) & z_{ht} \neq z_{h(t-1)}, \\ \psi_{ht} + (1 - \psi_{ht}) r_{ht}(z_{ht} \mid s) & z_{ht} = z_{h(t-1)}, \end{cases} \qquad (4)$$

where ψ_{ht} accounts for portfolio inertia.[5]

The extent of portfolio inertia may depend on the duration for which customer h has most recently subscribed to portfolio $z_{h(t-1)}$, which we denote m_{ht}. The decision to reconsider the portfolio may also depend on the firm's marketing activity. To incorporate the impact of the introductory promotion into the extent of portfolio inertia and to account for duration dependence, we model ψ_{ht} in a fashion similar a discrete-time hazard model with a Weibull baseline distribution (e.g., Seetharaman and Chintagunta 2003, Schweidel et al. 2008b):

$$\psi_{ht} = \exp[-\lambda(m_{ht}^c - (m_{ht} - 1)^c) \exp(\gamma_1 \text{Promo}(t)$$
$$+ \gamma_2 I(\text{Promo}(t-1) \neq \text{Promo}(t)))]. \qquad (5)$$

The coefficients γ_1 and γ_2 allow for the extent of portfolio inertia during the promotional period and in

---

[3] We considered the addition of household-by-service-specific random effects to address the concern of spurious state dependence in estimates of θ_k. Between the two models, we found no differences in the estimated signs of θ_k. As the model with these random effects performed worse during the holdout period, we present the results based on the more parsimonious specification in Equation (2).

[4] The multinomial portfolio choice probability in Equation (3) can also be derived as the product of multiple binary service subscription decisions, which we detail in the online supplement available on the first author's website (http://www.dschweidel.com). We favor the portfolio utility specification because it facilitates incorporating marketing actions that promote combinations of services.

[5] Our model does not allow us to statistically identify portfolio inertia at the state-specific level. We discuss this issue and potential remedies in our future research section toward the end of the paper.

the month immediately following it ($\text{Promo}(t-1) \neq \text{Promo}(t)$) to differ from the extent of portfolio inertia in the remaining months. The probability with which customer h reconsiders his portfolio can also increase ($c > 1$) or decrease ($c < 1$) with the duration for which it has been maintained to date (m_{ht}). If $c = 1$, portfolio inertia does not vary with m_{ht}.

If customer h changed his portfolio at $t(z_{ht} \neq z_{h(t-1)})$, he must have reconsidered the portfolio, with a probability of $1 - \psi_{ht}$. In contrast, there are two possible ways he may subscribe to the same portfolio ($z_{ht} = z_{h(t-1)}$). First, he may not have reconsidered his portfolio, which has a probability of ψ_{ht}. Second, though customer h did reconsider the portfolio (with a probability of $1 - \psi_{ht}$), he chose the same services, which may be due in part to service stickiness.

Evolution of Latent Relationship States

To incorporate evolution in customers' affinities for different services and their propensities to end their relationships, we use an HMM with multiple active states (e.g., Netzer et al. 2008). Each latent state is characterized by the baseline affinity for services, which may shift over time. The latent states serve three critical roles in our model. First, they offer a parsimonious means of allowing for customers' service affinities to vary over time. Second, as customers in one state may be more (or less) prone to subscribe to multiple services than customers in other states, they allow for correlations in subscriptions to different services at the margin. Last, as the states vary in the probability with which customers move to the "End" state (described next), the latent states allow us to relate observed subscriptions to defection risk, which we will illustrate.

In contrast to service stickiness and portfolio inertia, movement among the latent states is not directly affected by a customer's previous service subscriptions. Although a customer's decision to retain a service will be influenced by the baseline affinity of his current state, portfolio inertia, and service stickiness, the latter two dynamics do not influence a customer's decision to acquire a new service. This distinction allows us to distinguish the effects of changes in customers' underlying service affinities resulting from movement among the latent states, from portfolio inertia and service stickiness.

Consistent with prior CLV models, we incorporate an absorbing End state in which customers have terminated their relationship. Although customers may subscribe to no services in any active state *or* the End state, resuming service is not an option from the End state. We thus distinguish between a "temporarily cancelled" relationship (no services in an active state) and a terminated relationship (the End state), a key managerial insight. Though the relationship states are

Figure 2   Illustration of Movement Through the States
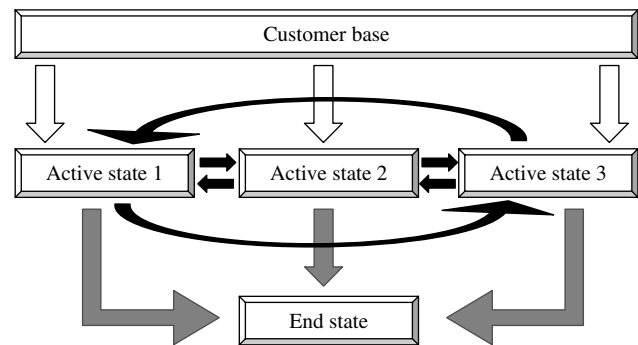
latent, they can be inferred probabilistically from the sequence of portfolios to which a customer has subscribed. We illustrate the possible movements through a model with three active states in Figure 2 and detail the transition matrix specification in the appendix.

Analogous to the Pareto/NBD and prior HMMs that have appeared in the literature, we allow customers to move among the latent states at different rates (e.g., Schmittlein et al. 1987, Netzer et al. 2008, Fader and Hardie 2010, Schweidel and Fader 2009). We also allow the monthly transition matrix $W_h(t)$ to vary over time, accounting for the firm's marketing actions (e.g., Montoya et al. 2010). This general specification nests a number of extant models. Setting $W_h(t)$ to the identity matrix results in a latent class model with no movement among the states. Such a model, though, may be inappropriate for estimating CLV, because it does not permit movement to the End state. Our model also nests the case in which there is no evolution while a customer is active, which is akin to the Pareto/NBD (Schmittlein et al. 1987). If $W_h(t) = W(t)$, the model would account for nonstationarity but would assume that all customers change service affinities at the same rate. Permitting movement among the states at different rates provides increased flexibility without requiring additional states. Our specification of $W_h(t)$ also nests a zero-order Markov process (i.e., a renewal model), thus allowing for inertia among the latent states. Viewing portfolio choice as akin to brand choice, our model also extends the "lightning bolt" models (e.g., Chintagunta 1999, Roy et al. 1996) by incorporating latent shifts in customers' service affinities.

Initial Distribution

We assume that customer h may begin in any active state, following an initial distribution $\pi_{h\cdot}$ in which the likelihood of beginning in the End state is zero ($\pi_{hE} = 0$). If information were available regarding customers' past experiences with the firm or the industry, such information could influence the initial

distribution or possibly be set exogenously. The likelihood of the observed portfolio decisions can then be expressed as (MacDonald and Zucchini 1997):

$$L_h = \pi_h' \times \tilde{L}_{h1} \times W_h(1) \times \tilde{L}_{h2} \times W_h(2)$$
$$\times \cdots \times W_h(M_h - 1) \times \tilde{L}_{hM_h} \times \mathbf{1}, \quad (6)$$

where π_h is the initial distribution, M_h denotes the number of months for which customer h is under observation, $\mathbf{1}$ is an $(S+1) \times 1$ vector of ones, $\tilde{L}_{hm}$ is an $(S+1) \times (S+1)$ diagonal matrix in which the sth diagonal element is $P(z_{ht} \mid s)$ from Equation (4), and the $(S+1)$th element is the likelihood of being in the End state (equal to 1 if z_{ht} consists of no services and 0 otherwise).

The Role of Promotional Activity

Our dynamic HMM allows for the firm's introductory promotion—and marketing activity more generally—to influence subscription decisions in three ways. First, the promotion affects the state-specific choice process when customers reconsider their portfolios (β_k in Equation (1)). Second, the promotion may influence the frequency with which customers reconsider their portfolios through the portfolio inertia probability ψ_{ht} (γ_1 and γ_2 in Equation (5)). Last, the transition matrix $W_h(t)$ incorporates the effects of the promotion by allowing the probability with which customers move to the End state to potentially differ during the promotion and in the month immediately following it.[6] Thus, the promotion may affect which portfolios are chosen, when customers consider changing their portfolios, and how long they remain as active subscribers, providing a rich picture of how marketing may influence customer value, which the firm can use to compare the related costs to their expected returns.

Empirical Application and Results

We compare nested versions of the full model to empirically assess the appropriate number of states, presented in Table 2, panel (a).[7]

In Table 2, panel (a), we present the calibration and holdout log marginal densities (LMD), which are calculated as the log of the harmonic mean of the likelihoods across iterations (Newton and Raftery 1994) for each model specification. We also calculate the hit rate for the entire portfolio, and the sum of the prediction

---

[6] Given the limited portfolio changes observed in our data, we specify a transition matrix $W_h(t)$ in which the transitions to the End state vary with the promotional activity. In contexts where more variation is observed in the data, time-varying covariates could be incorporated into the other elements of the transition matrix.

[7] Details of the estimation procedure are provided in the online supplement available on the first author's website (http://www.dschweidel.com).
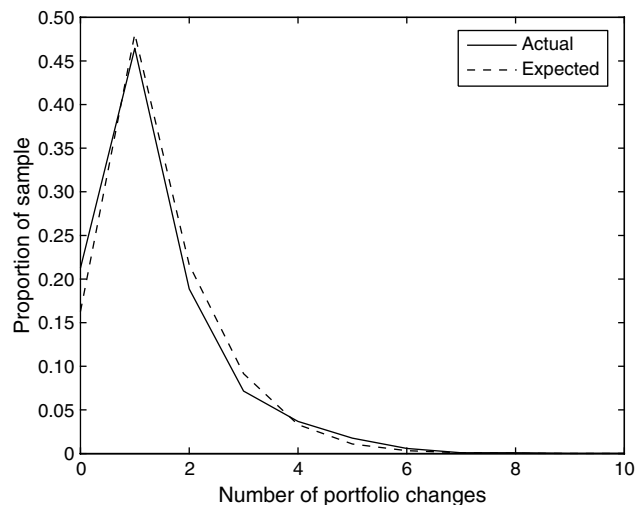
errors for individual services. At first glance, it might seem strange that the portfolio- and service-level fit statistics get better as we move from the calibration to the holdout period. This unusual pattern reflects the fact that many services (and complete portfolios) are dropped as customers "age" with the firm, and hence it is an easier prediction target.

In choosing among the models, we see that although the calibration LMD continues to improve as we add states, the holdout LMD diminishes slightly with the addition of the fourth active state. In moving from three to four active states, though predictive performance improves, we do not observe substantial improvement in either of the error measures during the holdout period. Given the marginal improvement in error measures and the additional complexity of the four-state model, we present the results from our model with three active states.

To ensure that the three-state model reflects the observed data, we compare the expected portfolio choices from the model to that observed in our data. For the 10 most popular portfolios, which account for 86.9% of portfolio choices with at least one active service in the calibration period, the average absolute difference between the observed proportion of portfolio choices with at least one service (Table 1) and the corresponding expectation from our model is 1.49%. Examining the distribution of portfolio choices across all portfolios, including the portfolio with no services, the average absolute difference is 0.13%, suggesting that we are capturing the distribution of portfolio choices. We also compare the observed and expected distributions of portfolio changes, as shown in Figure 3, which appear very similar.

Taken together, these results suggest that the model reflects observed portfolio choice behavior quite well. To further assess model identification, we simulate a

Figure 3 **Distribution of Portfolio Changes**

Table 2 **Model Comparison**

	Calibration LMD	Holdout LMD	Calibration portfolio hit rate (%)	Holdout portfolio hit rate (%)	Calibration service MAE	Holdout service MAE
(a) Model performance, varying states						
States						
1	−37,382	−2,830	83.6	94.2	0.360	0.105
2	−32,930	−2,611	83.8	94.2	0.360	0.106
3	−31,496	−2,521	84.0	94.3	0.357	0.104
4	−29,540	−2,527	84.5	94.3	0.344	0.104
(b) Three-state model performance, varying choice model components						
Description						
Full model	−31,496	−2,521	84.0	94.3	0.357	0.104
Portfolio inertia only	−32,256	−2,535	84.1	94.2	0.364	0.114
Service stickiness only	−35,088	−2,970	81.0	92.7	0.407	0.132
Multinomial choice	−84,258	−9,685	56.0	74.4	0.913	0.485

Note. The hit rate and mean absolute error (MAE) calculations were performed at each iteration of the Markov chain Monte Carlo (MCMC) sampler and averaged across iterations.

data set based on the posterior means of the model parameters. Fitting our model to the simulated data reveals that the posterior means of the parameters (i.e., the "true" values) are all within the 99% interval.

We next assess the need to include different model components by varying the inclusion of portfolio inertia and service stickiness in a 2×2 manner, as illustrated in Figure 1. The resulting LMD and fit statistics are presented in Table 2, panel (b).

The full model performs better in terms of portfolio predictions during the holdout period than the model with only portfolio inertia, though only to a slight extent. Both of these models yield more accurate predictions than the models that omit portfolio inertia, suggesting its importance in this empirical application and suggesting also that service stickiness alone is insufficient to capture the degree of inertia at the portfolio level. Not surprisingly, the model that lacks both portfolio inertia and service stickiness yields the worst performance, despite allowing for preferences to evolve through the latent states. Although service stickiness appears to be only marginally useful for this data set, we retain it for our subsequent analyses to demonstrate the full range of implications and inferences associated with our proposed model.

In addition to these results, we also consider alternative specifications for $W_h(t)$. We separately estimate a zero-order Markov process and a model with a homogeneous transition matrix. The model that omits heterogeneity performs slightly worse during holdout at both the portfolio level (hit rate = 93.8%) and the service level (MAE = 0.106), whereas the zero-order Markov model performs similarly at the service level (hit rate = 94.3%, service MAE = 0.106), suggesting that there is some predictive value to our more general specification of $W_h(t)$.

For comparative purposes, we consider a strong benchmark based on the observed data for the most common portfolios. We assume 12 observable

"states": one for each of the 10 portfolios from Table 1, another for all other nonempty portfolios, and a final one for the portfolio that consists of no services, resulting in a 12×12 monthly switching matrix.[8] This benchmark is discussed by Donkers et al. (2007) as an approach to estimate CLV in a multiservice context, as it captures (marginally) the most common co-occurrences and hence is not trivial to beat, given the high concentration of observed portfolios and inertia present in our data. The portfolio hit rate (treating all "other" portfolios as the same) during the holdout period is 93.4%. At the service level, assuming that the likelihood of subscribing to a particular service in the "other" state is equal to the proportion of observations in the calibration period in the "other" state that contained that service, the holdout error was 0.29. Thus, we outperform this benchmark at both the portfolio and service level.

We also examine the aforementioned "hold and keep" heuristic to predict subscriptions during the holdout period. The portfolio hit rate during the holdout period is 93.7%, and the service MAE is 0.136. The performance of this naïve heuristic and the previously described empirical benchmark provide us with perspective, suggesting that we should not read much into the sheer magnitude of these metrics because of the extent of inertia observed in our data. On a relative basis, our model reduces the portfolio error by approximately 8.5%, compared to this heuristic, as well as allowing us to estimate CLV and evaluate the impact of marketing actions.

We next present detailed results from the model with three active states that incorporates portfolio

[8] Discussions with managers from another multiservice provider revealed that this benchmark is used in practice. We also tested a similar benchmark based on the five most common portfolios, an "other" portfolio, and the null portfolio, which performed worse out-of-sample than the described 12×12 switching matrix.

inertia and service stickiness. We then demonstrate how our framework can be used for valuation and discuss its use for related activities of interest to managers.

Model Estimates

We first examine the parameters governing the portfolio choice process when customers reconsider their current portfolio, presented in Table 3.

As expected, the introductory promotion (β) generally increases the service utilities. We see from the state-specific service affinities (α) that State 1, in which customers begin with a mean probability of 16%, is marked by high attraction to cable-related add-ons. Compared to the service affinities from the other states, we colloquially refer to it as the "full-size" state. State 2, in which customers begin with a mean probability of 46%, is marked by high affinities for the digital package and HBO (hence this is called the "mid-size" state). Customers begin in State 3 with a mean probability of 38% and have a low affinity for all add-on services (the "economy" state). Given the variation in service affinities across states, found ex post, it is not surprising that a multistate model that allows service affinities to evolve is required.

In addition to the promotion and state-specific service affinities, service stickiness influences the portfolio choice process when customers reconsider their portfolios. For all services except basic cable, the service stickiness term θ is positive, suggesting that customers are more prone to keep a service they already have than they would be to add it. The most sizeable effects are observed for high-speed Internet and digital cable, which both require hardware and may therefore require greater effort to discard, compared to premium channels that can be canceled with a

phone call. It would be of future research interest to see how general this phenomenon is for add-on services for which technology investment may be an initial barrier to acquisition but later serves as a "barrier to exit." Although the service stickiness term for basic cable is negative, this term is not interpreted in isolation as easily, as the probability of subscribing to basic cable depends on the affinities for basic cable and the add-on services that require it. Thus, though a customer may be less attracted to basic cable on its own, his attraction to the add-on services may drive his decision to subscribe to basic cable.

We turn now to the parameter estimates associated with portfolio inertia (Equation (5)). The posterior mean of $\lambda = 0.23$ reflects a slow rate at which customers will reconsider their portfolios. The tendency for portfolio inertia increases the longer the portfolio has been kept, reflected by the posterior mean of $c = 0.73$. With a posterior mean of $\gamma_1 = -0.55$, we see that the promotion slows the rate of portfolio changes (i.e., increasing portfolio inertia). On the other hand, the posterior mean of $\gamma_2 = 0.45$ reflects a temporary rise in portfolio reconsideration as soon as the promotion ends. Thus, customers are more prone to choose larger portfolios (because of the β parameters) and maintain them with minimal reconsideration during the promotional period (γ_1). As one might expect, once the promotion ends, customers are more likely to reconsider their portfolios (γ_2), but service stickiness makes them somewhat inclined to keep many of the services that they hold.

Beyond the promotional period, portfolio inertia may lead a customer to keep the same portfolio while his baseline service affinities change as he moves among the latent states. Thus, a customer's current service affinities may not manifest immediately while

Table 3 Posterior Means and 95% Intervals for Portfolio Choice Model Parameters

	α_1 (Full-size)	α_2 (Mid-size)	α_3 (Economy)	β (Promotion effect)	θ (Service stickiness)
Basic	−1.05 (−3.87, 1.99)	−3.47 (−4.01, −3.00)	4.51 (3.76, 5.22)	1.24 (0.78, 1.72)	−1.79 (−2.24, −1.34)
Internet	−2.48 (−2.71, −2.26)	−2.47 (−2.66, −2.29)	−2.71 (−2.88, −2.56)	0.17 (0.00, 0.35)	3.70 (3.45, 3.96)
Digital	1.70 (1.41, 1.99)	1.46 (1.13, 1.81)	−2.53 (−2.72, −2.39)	0.17 (−0.07, 0.41)	2.95 (2.73, 3.17)
HBO	0.77 (0.57, 0.98)	0.74 (0.53, 0.97)	−2.57 (−2.77, −2.39)	0.93 (0.74, 1.10)	1.05 (0.86, 1.25)
Showtime	5.35 (4.26, 7.02)	−4.75 (−5.49, −4.06)	−6.65 (−7.54, −5.87)	1.19 (0.52, 1.95)	1.76 (0.94, 2.54)
Starz	−0.54 (−0.72, −0.37)	−2.30 (−2.50, −2.09)	−6.17 (−7.15, −5.42)	0.49 (0.29, 0.69)	1.66 (1.37, 1.95)
Cinemax	−0.40 (−0.57, −0.23)	−1.97 (−2.17, −1.78)	−5.50 (−6.20, −4.94)	0.38 (0.19, 0.56)	1.43 (1.16, 1.70)
TMC	4.19 (3.32, 5.28)	−6.03 (−7.36, −4.91)	−6.69 (−7.74, −5.80)	−0.41 (−1.49, 0.67)	0.85 (−0.35, 2.02)

he maintains the same portfolio. As we will show, in estimating a customer's risk for defecting or potential for acquiring more services, the firm must look beyond the most recent decision and consider the entire portfolio path to date.

To explore how customers evolve over the course of their relationship with the firm, we next examine the posterior mean transition matrix (across customers) during the promotion, immediately after the promotion, and for all subsequent months. These data are presented in Table 4.

We first consider movement among the active states. The economy state is always fairly sticky, with customers exhibiting the highest degree of inertia among the latent states both during and after the promotional period. Should customers move out of this state and remain active, they will most likely move to mid-size. Customers in mid-size are also prone to remain in the same state, but when they move, the average propensities to move to economy or full-size are fairly similar. But unlike economy and mid-size, the full-size state is relatively ephemeral, with customers more likely to transition away from it and move into the mid-size. There is a high degree of asymmetry when it comes to broad jumps between

full-size and economy: few customers will move up from economy to full-size in a single change; it is far more likely for a full-size customer to suddenly drop down to economy.

Examining the tendencies to move to the End state, we observe variation both across the latent states and in relation to the promotional period. During the promotion, customers are unlikely to end their relationships. When the promotion ends, we see a sharp increase in the probability that customers in full-size and mid-size move to the End state, whereas the probability of moving from economy to the End state increases only slightly. Thus, although customers in economy are unlikely to acquire more services, limiting their short-term cross-buying potential, they may have longer tenures because of this low transition probability to the End state that will fuel high customer valuations (as we will see). In contrast, full-size and mid-size customers are at greater risk of ending their relationships, which may limit their long-term value.

Updating Beliefs of State Membership

Though portfolio inertia and service stickiness can delay the manifestation of a customer's service affinities in its current state, we can infer the current state from the sequence of portfolios to date. Using the HMM, the posterior probability of being in state s at time t is

$$P(X_t = s \mid z_{h1}, \ldots, z_{ht})$$
$$= \frac{\pi_h \times \tilde{L}_{h1} \times W_h(1) \times \tilde{L}_{h2} \times W_h(2) \times \cdots \times \tilde{L}_{h(t-1)} \times W_{h \cdot s}(t-1) \times L_{hts}}{\pi_h \times \tilde{L}_{h1} \times W_h(1) \times \tilde{L}_{h2} \times W_h(2) \times \cdots \times \tilde{L}_{h(t-1)} \times W_{h \cdot s}(t-1) \times \tilde{L}_{ht} \times \mathbf{1}'},$$
(7)

where L_{hts} is the sth diagonal element of $\tilde{L}_{ht}$, the likelihood of choosing portfolio z_{ht} conditional on being in state s, given in Equation (4). To highlight the managerial relevance of using the full portfolio histories, we return to the initial example of three customers currently subscribing to the same portfolio. The first customer has *maintained* the portfolio that has the digital (and basic) cable and HBO. The second customer began by subscribing only to basic and digital cable and *added* HBO after three months of service. The third customer initially subscribed to digital cable with HBO, Showtime, and TMC but *dropped* Showtime and TMC after three months.[9] Although these customers all end up with the same portfolio (so that we may draw direct comparisons), the *path to get there* differs. This will reinforce the importance of using customers' portfolio histories in examining their value and expected future behavior.

Table 4   Transition Probabilities

From	To			
	Full-size	Mid-size	Economy	End
(a) Posterior means and 95% intervals for transition probabilities during promotion				
Full-size	0.48 (0.33, 0.69)	0.34 (0.25, 0.39)	0.16 (0.05, 0.30)	0.02 (0.00, 0.04)
Mid-size	0.18 (0.09, 0.31)	0.68 (0.56, 0.81)	0.12 (0.04, 0.29)	0.03 (0.00, 0.10)
Economy	0.04 (0.01, 0.11)	0.11 (0.03, 0.24)	0.83 (0.62, 0.97)	0.02 (0.00, 0.05)
End	0	0	0	1
(b) Posterior means and 95% intervals for transition probabilities immediately after promotion				
Full-size	0.43 (0.30, 0.65)	0.31 (0.22, 0.37)	0.15 (0.04, 0.29)	0.11 (0.04, 0.20)
Mid-size	0.14 (0.08, 0.26)	0.58 (0.41, 0.77)	0.11 (0.03, 0.26)	0.17 (0.03, 0.39)
Economy	0.04 (0.01, 0.10)	0.11 (0.03, 0.23)	0.83 (0.62, 0.97)	0.02 (0.00, 0.05)
End	0	0	0	1
(c) Posterior means and 95% intervals for transition probabilities after promotion				
Full-size	0.44 (0.31, 0.66)	0.31 (0.22, 0.38)	0.16 (0.04, 0.29)	0.10 (0.03, 0.18)
Mid-size	0.13 (0.08, 0.25)	0.55 (0.38, 0.75)	0.10 (0.03, 0.26)	0.22 (0.04, 0.46)
Economy	0.04 (0.01, 0.10)	0.10 (0.02, 0.22)	0.80 (0.57, 0.96)	0.06 (0.01, 0.13)
End	0	0	0	1

---

[9] We selected the exemplar portfolios for our vignettes from the 10 most popular portfolios (Table 1).

Figure 4   Illustration of Posterior State Predictions

(a) Posterior state membership

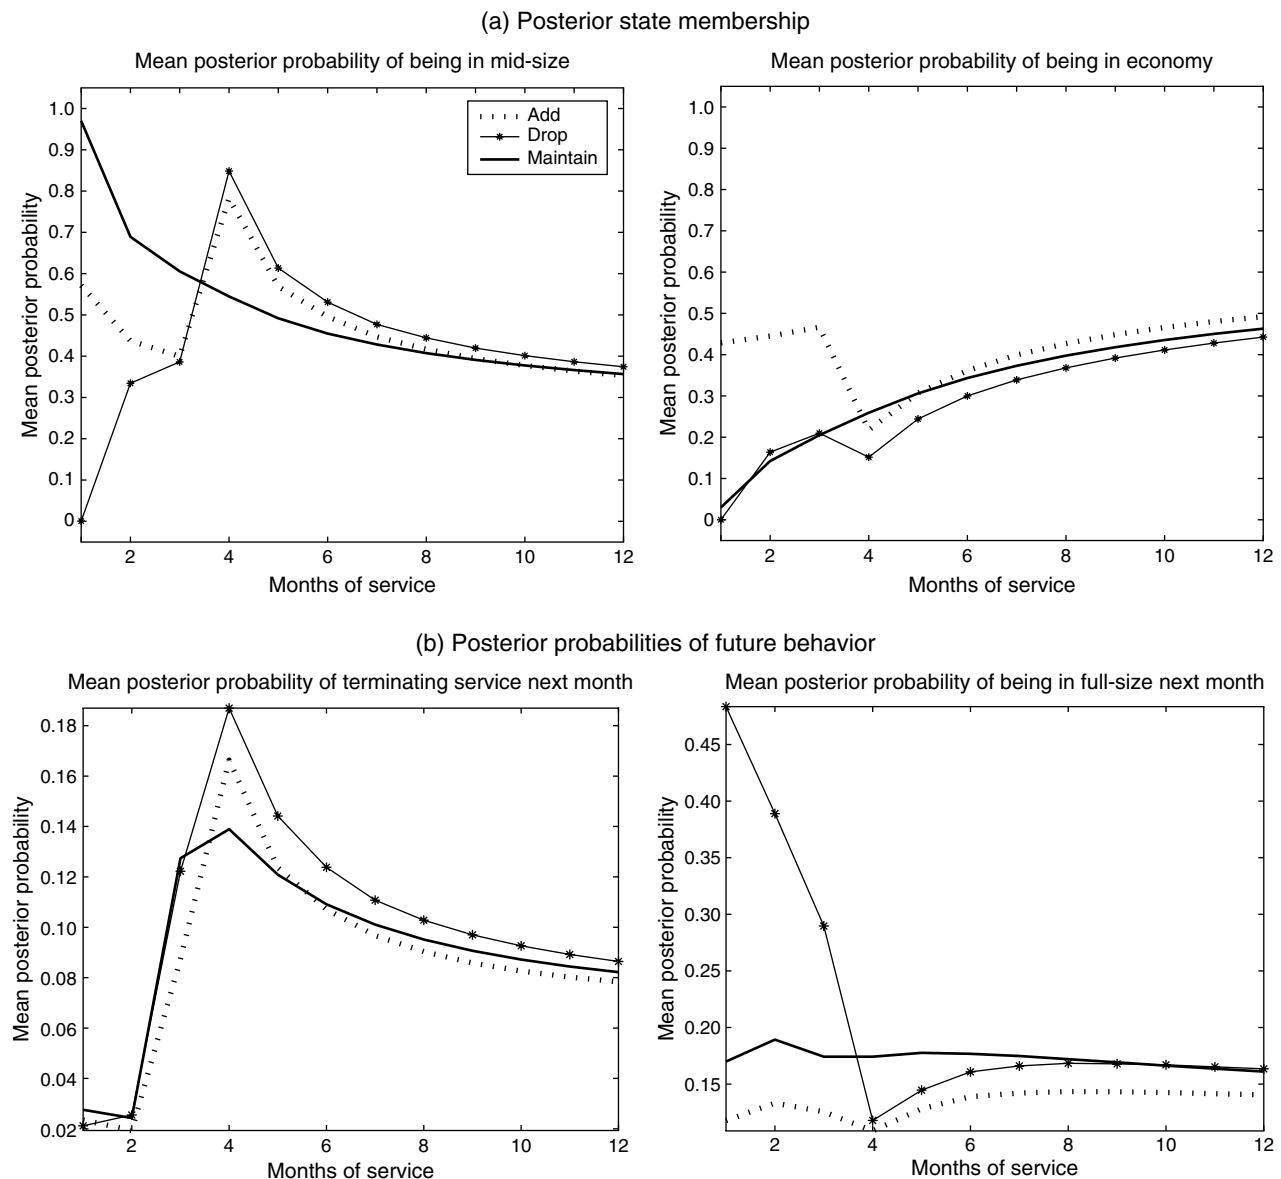(b) Posterior probabilities of future behavior

Figure 4(a) illustrates the mean posterior probability of being in mid-size and economy through the first 12 months of service for each of these customers.

The "maintain" customer most likely began service in mid-size, whereas the "add" customer may have begun in either mid-size or economy. The "drop" customer, who started with the largest portfolio and discarded services, most likely began in full-size. It gradually becomes more likely that the maintain customer is in economy, as he may have transitioned to economy but not yet changed his portfolio because of portfolio inertia and/or service stickiness. When the portfolio changes are observed in the third period for the add and drop customers, both are initially more likely to be in mid-size, given the chosen

portfolio. Over time, though, their likelihoods of being in economy gradually increase like the maintain customer.

In these illustrations, the customers have the same portfolio by the fourth month but have arrived at it through different paths, resulting in varying posterior state memberships. For several months after the changes take place, the state-membership probabilities are quite different, before these "ripple effects" start to fade away and the posterior beliefs converge. During this period of instability, the portfolio paths clearly point to differing posterior beliefs of state membership, yielding expected future actions that vary across these illustrative customers. The firm may leverage this variation by identifying those who

are most at risk for ending service, as well as those who are most likely to be in a state in which they may acquire more services.

To demonstrate this, we consider customers' propensities to terminate the relationship by moving to the End state in the next month. This probability is given by

$$p(\text{switch to End state in month } t)$$
$$= \sum_{s=1}^{S} P(X = s \mid z_{h1}, \dots, z_{ht}) w_{hs_s E}(t), \qquad (8)$$

where the first term reflects the updated state membership (Equation (7)) and the second term is the probability of transitioning to the End state. A customer's portfolio path influences the firm's expectation of the probability with which a customer terminates service through the updated state membership of the customer. The left panel of Figure 4(b) presents the probability of ending service next month, based on information available to date, for the three vignettes.

Although all three customers have a low probability of ending the relationship during the promotion, this probability rises for all three customers as the promotion concludes. When the add and drop customers change their portfolios in month 4, we see a sharp temporary rise in the risk that they will end service in the next month. It may seem somewhat surprising that these opposite behaviors will yield similar probabilities of termination; one might naively think that the add customer has become more engaged (and therefore less likely to terminate his relationship) whereas the drop customer is starting to disengage. Our model offers an alternative explanation that may help prevent the analyst from reading too much into these marginal changes. Specifically, both customers may have moved into mid-size despite starting in different states, and their future trajectories are fairly similar.

This analysis counters the conventional wisdom that a customer who has selectively discarded a service likely has "one foot out the door." Though the drop customer is at a (slightly) greater risk for ending the relationship altogether, as might be expected given his previous decision to discard services, such a customer still has plenty of upside potential. This finding, which was unexpected a priori, again underscores the need to use full portfolio histories in predicting customers' future actions.

Moving from customer termination to customer development, we perform a similar analysis about maximizing each customer's cross-buying potential. We derive the posterior probability of being in full-size in the next month in the same fashion that the defection probability was calculated in Equation (8),

which we illustrate in the right panel of Figure 4(b). Interestingly, following the portfolio changes made at $t = 4$, it is the maintain customer who is most likely to move to full-size. As noted above, the selective changes made by the add and drop customers at $t = 4$ suggest that it is unlikely that these customers are in full-size. It is particularly surprising that it is the add customer who has the lowest probability of moving to full-size (although we acknowledge that the differences are not very substantial). Again, this runs counter to the usual conventional wisdom about customer development, but it makes sense in light of the richer portfolio dynamics that our model captures.

Customer Valuation

Through posterior predictive simulation, we can estimate each customer's overall CLV (i.e., from the time of initial acquisition), as well as their residual value at any point in time after acquisition. Estimating such metrics at the individual level is particularly important for contractual service providers that wish to manage their spending across customers. To compute CLV, at each iteration of the Markov chain Monte Carlo (MCMC) sampler, we simulated the behavior of 1,000 customers, including their initial state and the sequence of portfolio choices made over a lengthy time horizon (to ensure that all customers have dropped service), to assess their lifetime value. The CLV of customer h is calculated as

$$\text{CLV}_h = \sum_{t=1}^{\infty} \sum_{k \in K} V_k(z_{ht}) \times (P_k - C_k) \times \left(\frac{1}{1+d} \right)^{(t-1)/12}, \quad (9)$$
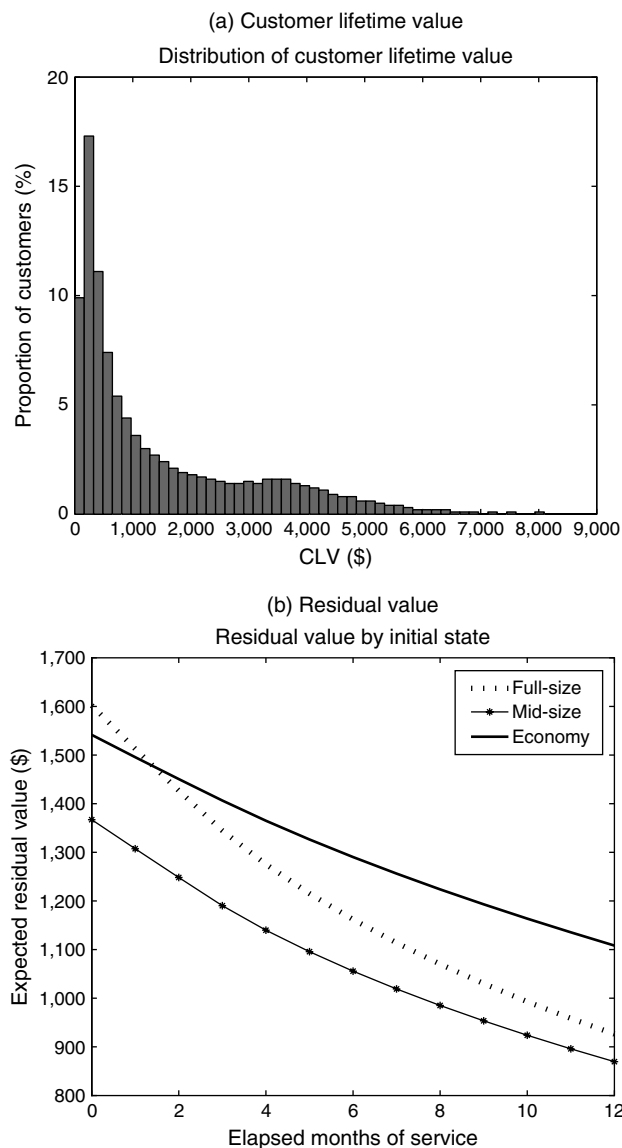
where z_{ht} indicates the portfolio to which h subscribes at time t, $V_k(\cdot)$ denotes the inclusion of service k in portfolio z_{ht}, P_k is the price of service k, C_k is the cost of service k, and d is the monthly discount rate. We assume that $C_k = 0$; that P_k is constant over time, which is consistent with the company's observed prices; and that d reflects an annual discount rate of 15%.

The distribution of CLV across customers, averaged across iterations of the MCMC sampler, is shown in Figure 5(a).

We find considerable heterogeneity in customer value, which is fortunate for the firm because it can be used as the basis for targeting (e.g., Venkatesan and Kumar 2004). To calculate the expected CLV (ECLV), one need just compute the mean of this distribution, leading to an overall ECLV of $1,471 with a median lifetime of 15.2 months. From the firm's perspective, this can provide sharp guidelines for necessary or maximal allocated spending.

To illustrate how this can be used for targeting, we examine the expected value and lifetime of customers by initial state, which may be related to past experience or acquisition channel. The ECLV for customers

Figure 5   Valuing Customers

(a) Customer lifetime value

Distribution of customer lifetime value

(b) Residual value

Residual value by initial state

beginning in full-size is $1,602 (with a median lifetime of 12.0 months), and the ECLV for those who begin in mid-size is $1,367 (with a median lifetime of 11.5 months), and that of those starting in economy is $1,542 (with a median lifetime of 23.4 months). The ECLV of customers starting in full-size, where they are likely to choose many add-on services, enjoys a premium of almost 20% compared with those starting in mid-size, which is consistent with common promotional activity in the telecommunications industry to provide incentives to subscribers to begin with many services. However, because of the differences in expected lifetimes, the value of customers starting in full-size only exceeds that of those starting in economy by 4%, despite a considerable difference in the size of the portfolios initially chosen.

Interestingly, customers starting in economy are more valuable than those starting in mid-size. Although

customers initially in mid-size are prone to subscribe to more services, they tend to have shorter tenures. As a result, the expected lifetime value of economy customers ends up being greater. By deriving CLV from our model, thus accounting for the relationship between portfolio choices and the likelihood of ending the relationship, we see that the link between portfolio size and customer value is not as clear as previously thought. One might think that "bigger is better" in regard to the size of the portfolio and its relation to CLV, yet a "slow and steady" customer may ultimately deliver greater value.

We further examine these differences by calculating the expected residual value after t months of service. We illustrate the residual value over a one-year period in Figure 5(b).

Though customers beginning in full-size are more valuable initially, the residual value of customers starting in economy exceeds that of those beginning in full-size after just a few months. Based on the residual value of current customers and the ECLV of new customers, the firm may consider allocating resources toward prospects and newer customers rather than investing in older customers, depending on the mix of customers in the market that remains.

To further demonstrate how customer valuation can be employed by multiservice firms, we repeated the same simulation procedure but varied the length of the promotional period from 0 to 6 months. With no promotion, we find that the ECLV is $1,232. Increasing the length of the promotion one month at a time, the ECLV increases by $116, $68, and $55 for each of the first three months, yielding an increase of $239 per customer in expected revenue during the promotion. If the firm were to extend the promotion, for the next three months the incremental increases in revenue would be $50, $46, and $43, respectively, continuing the trend of diminishing returns from extending the promotion. Comparing the incremental changes in ECLV from extending the promotion to the associated costs, the firm can evaluate the appropriate length for the promotion, which may vary across customers. If the data contained sufficient variation in promotional activity, the same approach could be used to evaluate alternative promotions, such as offering a promotion at different times or only on certain services.

Conclusions and Directions for Future Research
We have developed a dynamic model of portfolio choice for use in multiservice contractual settings. Through our framework, we see the relevance of using customers' full portfolio paths to make inferences regarding managerial metrics such as the likelihood of acquiring additional services and terminating

the relationship. We demonstrate how our model can be used to estimate CLV and residual value, which can offer guidance for decisions such as the appropriate compensation for customer complaints, targeting customers with specific marketing initiatives, and the amount to invest in customer acquisition and retention efforts.

Although we recognize that our results are data set dependent, our analysis reveals a number of key insights. Our illustrative vignettes highlight the importance of leveraging the full portfolio histories rather than just the most recent portfolio. Interestingly, those who have selectively dropped add-on services may have a higher likelihood of enhancing their relationship by acquiring additional services. This contradicts the notion that selectively dropping services is indicative of a weakening relationship. To the contrary, it may reflect an engaged customer who is considering his options—either pruning or enhancing his portfolio.

In addition to being informative of future service acquisition, we see that the portfolio history can provide insight into those customers most at risk for ending the relationship. Informal discussions with managers from other industries have provided strong support for the notion of customers "on life support" in contractual contexts such as magazine subscriptions, extended warranty plans, and financial programs. In such cases (among others), attempts to reinvigorate the customer "relationship" may backfire, as doing so may serve as the catalyst for a customer to end the relationship (i.e., moving to the End state). Future research may shed light on the circumstances in which attempts should be made to re-engage such customers and those cases where the best strategy is to "let sleeping dogs lie."

There are limitations to our work that provide open avenues for future research. We consider variation across states in the transition probabilities and baseline service affinities. Variation across states could also be introduced in the effects of marketing activities, service stickiness, and portfolio inertia. Given the infrequency with which portfolio changes occur, though, there may be methodological challenges to identifying these distinct effects with the limited variation observed in subscription contexts. To overcome such limitations, certain modeling assumptions may be necessary. For example, restricting state transitions to occur with different periodicities than portfolio changes may allow for state-specific portfolio inertia (e.g., Jerath et al. 2010). Another way our model can be further generalized would be to relax the first-order Markov process, commonly made in HMMs (e.g., Netzer et al. 2008, Montgomery et al. 2004) and allow for the transitions among states to depend on

the duration in the latent state (i.e., a hidden *semi-*Markov model; see Guédon 2003). It may also be worth exploring how changes to the service portfolio may directly impact movement among the latent states. We leave these extensions as areas for future research.

Our data provider did not have detailed information on customer-firm interactions, which can provide insight into customers' future behavior (e.g., Berger et al. 2002). The inclusion of such data would allow service providers to determine the best way to respond to each customer. Additionally, the availability of service usage behavior may allow researchers to distinguish between customers who are maintaining services because of inertia (e.g., subscribers with diminished or little usage) and those who have learned about the service's value subsequent to acquisition (e.g., subscribers with increased or heavier usage). The analyses resulting from the inclusion of such data may reveal the optimal strategies for firms to employ with different customers, based on their subscription *and* usage histories.

Though our results suggest doubt about common heuristics of customer value such as portfolio size, managers may find it useful to identify others that may aid in forecasting the best (and worst) customers in a multiservice setting, as well as those who are at the greatest risk for defection (e.g., Wübben and Wangenheim 2008). Additionally, it might be worthwhile to examine the diffusion of new services and their effect on the overall portfolio choice process, as well as to learn how customers' total portfolios (i.e., their relationships *across* providers) evolve over time. Last, as this research is intended to provide managers with a tool with which to understand their customers, additional real-world testing in other industries is warranted, though we expect our model to perform well in other settings.

Acknowledgments

Appendix

Transition Matrix Specification

To parameterize the transition matrix $W_h(t)$, we assume that each state s' has a relative attractiveness (gravity of attraction, if you will) from state s, which may vary across customers to allow for heterogeneous evolution. To place structure on the transition matrix and identify the states, we assume that customers in state j, where $j \in \{A_1, \ldots, A_s\}$, transition to the End state with a probability of $w_{hjE}(t)$ such that

$$\text{logit}(w_{hjE}(t)) = \tau_{jE} + \kappa_{j1} \cdot \text{Promo}(t)$$
$$+ \kappa_{j2} I(\text{Promo}(t-1) \neq \text{Promo}(t)) + \varphi_{hjE}, \quad (10)$$

where $\tau_{A_1E} < \tau_{A_2E} < \cdots < \tau_{A_SE}$. As such, we assume that $\tau_{A_1E} = \eta_1$, $\tau_{A_2E} = \tau_{A_1E} + \exp(\eta_2)$, $\tau_{A_3E} = \tau_{A_2E} + \exp(\eta_3)$, etc.

Conditional on not transitioning to the End state, with a probability of $(1 - w_{hjE}(t))$, customers transition among the active states with a probability of $n_{hjj'}$, where $j, j' \in \{A_1, \ldots, A_s\}$. This corresponds to the multinomial logit probabilities based on the relative attractiveness of the active states, $g_{hjj'}$, where $j, j' \in \{A_1, \ldots, A_s\}$ (e.g., Nakanishi and Cooper 1974, Cooper and Nakanishi 1983). We assume that the relative attractiveness of A_1 from any of the other active states is 0. The relative attractiveness $g_{hjj'}$ is given by

$$g_{hjj'} = \begin{cases} g_{jj'} + \varphi_{hjj'} & j \in \{A_1, \ldots, A_S\}, \; j' \in \{A_2, \ldots, A_S\}, \\ 0 & j' = A_1. \end{cases} \tag{11}$$

The resulting transition probabilities are then given by

$$w_{hjj'}(t) = \begin{cases} (1 - w_{hA_sE}(t))n_{hA_sA_{s'}} & j \neq E, \; j' \neq E, \\ w_{hA_sE}(t) & j = A_s, \; j' = E, \\ 1 & j = j' = E, \\ 0 & j = E, \; j' = A_s, \end{cases} \tag{12}$$

where $\varphi_{hj\cdot} = [\varphi_{hjA_2} \cdots \varphi_{hjA_S} \; \varphi_{hjE}]^T$ and $\phi_{hj\cdot} \sim \text{MVN}(\mathbf{0}, \Sigma)$, $\mathbf{0}$ is an $S \times 1$ vector of zeros, and Σ is an $S \times S$ covariance matrix. This allows the customer-specific effects to vary by state, as well as for the attraction of a state (both active and End states) from other active states to be correlated.

To identify the active states, as described in Equations (10)–(12), we order them by increasing likelihood of transitioning to the End state. For expositional ease, we reorder the states in our discussion of the results based on the expected portfolio size.

Initial Distribution Specification

To allow for variation across customers in the relationship state in which they begin service, we assume a customer-level initial distribution. In a model with S active relationship states, we parameterize the probability with which customer h begins in relationship state s as

$$\pi_{hs} = \begin{cases} \dfrac{1}{1 + \sum_{s'=2}^{S} \exp(\upsilon_{s'} + \delta_{hs'})} & s = 1, \\ \dfrac{\exp(\upsilon_s + \delta_{hs})}{1 + \sum_{s'=2}^{S} \exp(\upsilon_{s'} + \delta_{hs'})} & s = 2, \ldots, S, \end{cases} \tag{13}$$

where $\delta_{h\cdot} \sim \text{MVN}(\mathbf{0}, T)$, $\mathbf{0}$ is an $(S-1) \times 1$ vector of zeros, and T is an $(S-1) \times (S-1)$ covariance matrix.

References

Berger, P. D., R. N. Bolton, D. Bowman, E. Briggs, V. Kumar, A. Parasuraman, C. Terry. 2002. Marketing actions and the value of customer assets: A framework for customer asset management. *J. Service Res.* **5**(1) 39–54.

Blattberg, R. C., G. Getz, J. S. Thomas. 2001. *Customer Equity: Building and Managing Relationships as Valuable Assets*. Harvard Business School Press, Boston.

Bolton, R. N., K. N. Lemon, P. C. Verhoef. 2004. The theoretical underpinnings of customer asset management: A framework and propositions for future research. *J. Acad. Marketing Sci.* **32**(3) 1–20.

Chintagunta, P. K. 1999. Variety seeking, purchase timing, and the "lightning bolt" brand choice model. *Management Sci.* **45**(5) 486–498.

Cooper, L. G., M. Nakanishi. 1983. Standardizing variables in multiplicative choice models. *J. Consumer Res.* **10**(1) 96–108.

Donkers, B., P. C. Verhoef, M. G. de Jong. 2007. Modeling CLV: A test of competing models in the insurance industry. *Quant. Marketing Econom.* **5**(2) 163–190.

Du, R., W. A. Kamakura. 2006. Customer lifecycles and life styles in America. *J. Marketing Res.* **43**(1) 121–132.

Erdem, T. 1996. A dynamic analysis of market structure based on panel data. *Marketing Sci.* **15**(4) 359–378.

Fader, P. S., B. G. S. Hardie. 2010. Customer-base valuation in a contractual setting: The perils of ignoring heterogeneity. *Marketing Sci.* **29**(1) 85–93.

Fader, P. S., B. G. S. Hardie, C.-Y. Huang. 2004. A dynamic change-point model for new product sales forecasting. *Marketing Sci.* **23**(1) 50–65.

Fader, P. S., B. G. S. Hardie, K. L. Lee. 2005. RFM and CLV: Using iso-value curves for customer base analysis. *J. Marketing Res.* **42**(4) 415–430.

Guédon, Y. 2003. Estimating hidden semi-Markov chains from discrete sequences. *J. Computational and Graphical Statist.* **12**(3) 604–639.

Hogan, J. E., K. N. Lemon, R. T. Rust. 2002. Customer equity management: Charting new directions for the future of marketing. *J. Service Res.* **5**(1) 4–12.

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

Jerath, K., P. S. Fader, B. G. S. Hardie. 2010. New perspectives on customer "death" using a generalization of the Pareto/NBD model. Working paper, Carnegie Mellon University, Pittsburgh.

Li, S., B. Sun, R. T. Wilcox. 2005. Cross-selling sequentially ordered products: An application to consumer banking services. *J. Marketing Res.* **42**(2) 233–239.

MacDonald, I. L., W. Zucchini. 1997. *Hidden Markov and Other Models for Discrete-Valued Time Series*. Chapman and Hall, London.

Montgomery, A. L., S. Li, K. Srinivasan, J. C. Liechty. 2004. Modeling online browsing and path analysis using clickstream data. *Marketing Sci.* **23**(4) 579–595.

Montoya, R., O. Netzer, K. Jedidi. 2010. Dynamic allocation of pharmaceutical detailing and sampling for long-term profitability. *Marketing Sci.* **29**(5) 909–924.

Nakanishi, M., L. G. Cooper. 1974. Parameter estimation for a multiplicative competitive interaction model: Least squares approach. *J. Marketing Res.* **11**(3) 303–311.

Netzer, O., J. M. Lattin, V. Srinivasan. 2008. A hidden Markov model of customer relationship dynamics. *Marketing Sci.* **27**(2) 185–204.

Newton, M. A., A. E. Raftery. 1994. Approximate Bayesian inference by the weighted likelihood bootstrap. *J. Roy. Statist. Soc. B* **56**(1) 3–48.

Niraj, R., V. Padmanabhan, P. B. Seetharaman. 2008. A cross-category model of households' incidence and quantity decisions. *Marketing Sci.* **27**(2) 225–235.

O'Sullivan, D., A. V. Abela. 2007. Marketing performance: Measurement ability and firm performance. *J. Marketing* **71**(2) 79–93.

Reinartz, W. J., V. Kumar. 2003. The impact of customer relationship characteristics on profitable lifetime duration. *J. Marketing* **67**(1) 77–99.

Roy, R., P. K. Chintagunta, S. Haldar. 1996. A framework for investigating habits, "the hand of the past," and heterogeneity in dynamic brand choice. *Marketing Sci.* **15**(3) 280–299.

Rust, R. T., K. N. Lemon, V. A. Zeithaml. 2004. Return on marketing: Using customer equity to focus marketing strategy. *J. Marketing* **68**(1) 109–127.

Schmittlein, D. C., D. G. Morrison, R. Colombo. 1987. Counting your customers: Who are they and what will they do next? *Management Sci.* **33**(1) 1–24.

Schweidel, D. A., P. S. Fader. 2009. Dynamic changepoints revisited: An evolving process model of new product sales. *Internat. J. Res. Marketing* **26**(2) 119–124.

Schweidel, D. A., P. S. Fader, E. T. Bradlow. 2008a. A bivariate timing model of customer acquisition and retention. *Marketing Sci.* **27**(5) 829–843.

Schweidel, D. A., P. S. Fader, E. T. Bradlow. 2008b. Understanding service retention within and across cohorts using limited information. *J. Marketing* **72**(1) 82–94.

Seetharaman, P. B., P. K. Chintagunta. 2003. The proportional hazard model for purchase timing: A comparison of alternative specifications. *J. Bus. Econom. Statist.* **21**(3) 368–382.

Smurl, P. 2002. Subscriber retention: From hook-up to catch-up. *Multichannel News* (November 3), http://www.multichannel .com/article/CA256568.html.

Venkatesan, R., V. Kumar. 2004. A customer lifetime value framework for customer selection and resource allocation strategy. *J. Marketing* **68**(4) 106–125.

Wübben, M., F. v. Wangenheim. 2008. Instant customer base analysis: Managerial heuristics often "get it right." *J. Marketing* **72**(3) 82–93.