Intern. J. of Research in Marketing 26 (2009) 119–124

Contents lists available at ScienceDirect

Intern. J. of Research in Marketing

journal homepage: www.elsevier.com/locate/ijresmar

Dynamic changepoints revisited: An evolving process model of new product sales

David A. Schweidel [a,*,1], Peter S. Fader [b]

[a] University of Wisconsin-Madison School of Business, United States
[b] The Wharton School of the University of Pennsylvania, United States

ARTICLE INFO

Article history:
First received in August 28, 2007
and was under review for 4 months

Area Editor: Russell Winer

Keywords:
New products
Nonstationarity
Changepoint models
Duration models

ABSTRACT

This paper posits a new framework to model the trial of and repeat purchasing of a new product. While much research has examined underlying shifts in consumer purchasing patterns, the typical assumption has been that the underlying purchasing process remains the same although the purchasing rate may change over time. Motivated by Fader, Hardie, and Huang's development of a dynamic changepoint model [Fader, P. S., Hardie, B. G. S., & Huang, C. -Y. (2004). A Dynamic Changepoint Model for New Product Sales Forecasting. *Marketing Science*, 23 (1), 50–65], we consider an *evolving process* as consumers gain more experience with a new product.
Our framework assumes that consumers progress through two purchasing states, becoming more regular in their inter-purchase times as they gain experience with the product through repeat purchases. More specifically, they move from an initial state of exponential purchasing to a "steady state" that is characterized by a more regular Erlang-2 timing distribution. The proposed model is very flexible and nests a number of existing models, enabling it to explain a wide range of observed behavioral patterns. We apply our evolving process model to the same datasets used by Fader, Hardie, and Huang [Fader, P. S., Hardie, B. G. S., & Huang, C. -Y. (2004). A Dynamic Changepoint Model for New Product Sales Forecasting. *Marketing Science*, 23 (1), 50–65] and compare our results to a number of competing models. We find empirical evidence to support the use of a two-state model, since it yields relevant insights as well as improved empirical performance. We discuss the implications as they relate to forecasting new product sales.

1. Introduction

In assessing the likely long-run success of new products, managers must understand both the trial (Fader, Hardie, & Zeithammer, 2003; Hardie, Fader, & Wisniewski, 1998) and repeat purchasing (Kalwani & Silk, 1980) behavior of consumers. While consumers may be quick to try a new product, this behavior may not be mirrored in their repeat purchasing. Though they may continue to purchase it over time, they may do so at a slower (or faster) rate. Additionally, they may demonstrate increased regularity over time. Rather than examining trial and repeat purchasing in isolation, Fader, Hardie, and Huang (2004; hereafter denoted as FHH) developed a flexible model to capture and explain a systematic shift observed in purchase behavior as a panel of consumers gains experience with a new product. FHH posited a "dynamic changepoint" model in which consumers' purchasing rates can be updated with a decreasing probability over time in order to capture the (presumed) evolution toward "steady state" purchasing behavior; thus, reflecting an individual's acquisition of experience with the product and the associated decrease in novelty. While the authors explicitly tested the exponential and Erlang-2 distributions for the underlying inter-purchase timing process and allowed individuals' purchasing rates to change over time, they assumed that the underlying process itself (i.e., the exponential or Erlang-2 distribution) remains unchanged.

In this short paper, we propose an alternative framework that is rooted in consumer psychology literature. Rather than assuming that consumers continually "trade in" their purchasing rates for a particular timing model and renew them from the same distribution, we consider a simpler behavioral explanation in which there are two purchasing states. Consistent with the spirit of FHH, we believe that consumers' repeat purchasing of a new product is highly variable in its infancy, but will tend to stabilize over time as consumers gain experience through additional repeat purchases. We account for this stabilization as a shift from the "trial regime" to a "steady state."

In developing a model of consumer behavior that allows for a "steady state," we consider *process evolution* as an explanation for this movement toward stability. In contrast to FHH, we propose that consumers may move toward a more regular timing process as they determine how the new product fits into their usual purchasing routines. In this more regular purchasing state, the frequency of purchasing may vary significantly across individuals. Some individuals may decide to purchase the product with a higher frequency than they did while they were acquiring knowledge of the product, while others

* Corresponding author. 975 University Avenue, 4191B Grainger Hall, Madison, WI 53706, United States. Tel.: +1 608 265 2736.
E-mail address: dschweidel@bus.wisc.edu (D.A. Schweidel).

[1] The first author acknowledges the financial support of the Center for Brand and Product Management at the Wisconsin School of Business.

0167-8116/$ – see front matter. Published by Elsevier B.V.
doi:10.1016/j.ijresmar.2008.12.005

may exhibit longer inter-purchase cycles. However, whether the new rate is faster or slower than the old one, the degree of variability from one purchase cycle to the next will be reduced over time. The key factor distinguishing this purchasing state from the earlier one is increased regularity in inter-purchase times, which we incorporate via the Erlang-2 distribution.

Our model is also motivated by an empirical observation seen in FHH and earlier research in the general area of multi-event timing models. Consistent with past research (e.g., Fader et al., 2003; Moe & Fader, 2004), FHH found that the exponential distribution offered better empirical performance than the more regular Erlang-2 timing model. However, FHH frequently referred to the work of Gupta (1991) who, working with data for an established product, found evidence favoring the use of the Erlang-2 timing distribution. Other researchers (e.g., Chatfield & Goodhardt, 1973; Herniter, 1971; Jeuland, Bass, & Wright, 1980; Kahn & Morrison, 1989), have also provided strong support (and logical arguments) in favor of more regular timing models such as the Erlang-2. To the best of our knowledge, no one in the marketing literature has provided a synthesis or justification for these seemingly conflicting results. By considering the evolution toward regularity as more repeat purchases are made, the process evolution model provides a parsimonious explanation that allows both timing models to co-exist, but at different stages of a consumer's experience with a new product.

We further review related literature in Section 2. In Section 3, we formulate the proposed evolving process model and discuss the salient features of the model, as well as the extant models that it nests. We also put forth a two-state purchasing model which, while not accounting for process evolution, allows for a separate "trial regime" and "steady state." We then apply the evolving process and the two-state purchasing models to the same datasets used by FHH and discuss the results in Section 4. In Section 5, we conclude by discussing the managerial implications of the model as well as some salient limitations and future research directions.

2. Related research

FHH estimated a dynamic changepoint model of customer inter-purchase times and provided a thorough demonstration of this model's empirical performance. Their model is based on the following three assumptions:

1. Each household's purchasing follows an exponential distribution with parameter λ_h.
2. Each household's purchasing rate is drawn from a gamma distribution with parameters r and α.
3. Following a household's j^{th} purchase, it obtains a new purchasing rate as an independent draw from the same gamma distribution with a probability of $\gamma_j = 1 - \psi(1 - e^{-\theta(j+1)})$.

While these assumptions are plausible and the empirical performance of the FHH model is very good, it assumes that the core purchase process remains unchanged. More specifically, it assumes that each consumer's purchasing rates are drawn from the same exact underlying distribution over time.

There are, however, alternative methods by which changes in consumers' purchasing patterns can be taken into account. In this research, we consider a model with two purchasing states, allowing households in the first state to latently transition to the second one after each purchase. This is similar in spirit to extant marketing models that have considered customers' progression from an "active" purchasing state to an "inactive" state (e.g., Fader et al., 2004; Schmittlein, Morrison, & Colombo, 1987). Unlike such models, though, we allow for purchasing to occur in the second state, making it akin to a "steady state." Moreover, we explicitly consider changes in the underlying process rather than just in the parameters, per se.

The consumer psychology literature supports the notions of distinct purchasing states and underlying process changes over time in decision-making. Howard and Sheth (1967) described a model of buyer behavior in which customers progress through a series of stages, from extensive problem solving behavior to limited problem solving behavior, and finally to routinized response behavior. This would appear to support FHH's assumption that the probability of a household changing its purchasing rate decreases over time, but calls into question whether such behavior is consistent with a renewal of the purchasing rate from the same distribution or with drawing a purchasing rate from a different distribution. In addition, under Howard and Sheth's (1967) framework, it is the underlying purchasing process itself that differs in each of the stages, suggesting that process evolution may be warranted.

Meyer (1987) found that decision-making evolves from the use of episodic knowledge toward the use of alternative judgment strategies that draw on generalized knowledge, into which past experience has been incorporated. This would suggest that consumers' purchasing of a product with which they have experience may differ systematically from their purchasing of a product with which they have more limited experience (or none at all), mirroring the shift from a "new" to an "established" product, further supporting a two-state model with process evolution.

West, Brown, and Hoch (1996) demonstrated that preferences become more consistent as respondents gain a greater depth of knowledge. Hoeffler and Ariely (1999) found similar results using different experimental methods. The authors assert that when first "encountering a new domain, consumers are more likely to be constructing their preferences. Eventually, as consumers gain experience in a domain, stable preferences can develop." Through repeated purchases, it stands to reason that a consumer's preference for a new product will eventually stabilize in relation to alternative products. For example, when a new beverage is launched, consumers may be uncertain as to how it compares to those products that they are purchasing on a regular basis. After learning about the product through trial and repeat purchases, stable preferences would develop upon which the product can be incorporated into a regular cycle.

These papers support the notion of greater stability in preferences (and consequently behavior) over time, as posited by FHH. They go one step further, however, in suggesting that the underlying decision process may become fundamentally different as consumers gain experience. It is therefore reasonable to develop and test a model that allows for separate purchasing states, as well as one in which customers become more regular in their purchasing as they gain experience with the product through repeat purchases. If this is the case, then consumers' purchasing patterns, once their preferences for the product stabilizes and they are in a "steady state," should be more consistent with Erlang-2 inter-purchase times than with exponential inter-purchase times.

3. Model development

3.1. Model formulation

We first present the development of the household-level evolving process model, followed by an alternative two-state specification that does not incorporate process evolution. Under the evolving process story, we assume that households begin purchasing under a "trial regime." Subsequent to the trial purchase, a household may move to the second, more regular "steady state" (which includes the possibility that it will never buy the product again), or it may need additional purchases in the first state to determine how the product fits into its purchasing routine. When the household has acquired sufficient knowledge through repeat purchases, it moves into the “steady state” regime, which is characterized by more regular inter-purchase times. This reflects the product's transition from “new” to “established” for this particular household. In contrast to FHH, who assumed homogeneity in the changepoint process, we allow consumers to shift to more regular purchasing at different rates. Once this transition has occurred, the evolutionary process is over and the household remains in the “steady state” for the duration of the observation (or forecast) period.

More formally, the household-level model is based on the following assumptions:

1. There are two underlying purchasing states, one characterized by an exponential distribution (with rate λ_{h1}) and the other by an Erlang-2 distribution (with rate λ_{h2}).
2. If a household has not yet transitioned to the Erlang-2 state, it will do so after its next purchase with a probability of p_h.
3. The purchasing rates (λ_{h1} and λ_{h2}) and the transition probability p_h are independent.[2]

Assume that household h has made x purchases at $T_h=\{t_1, t_2,\ldots, t_x\}$during the time period $(0,T]$, k of which were made in the exponential purchasing state before transitioning. The likelihood of household h making its k^{th} purchase at t_k is:

$$L1(t_k|\lambda_{h1},k)=\lambda_{h1}^{k}e^{-\lambda_{i1}t_k}. \tag{1}$$

It should be noted that, for those households who do not make any purchases during the observation period, there is no opportunity to exit the exponential purchasing state. As such, $L1(T|\lambda_{h1},k=0)$ appropriately yields the survival function of the exponential distribution. More generally, $L1(T|\lambda_{h1},k)$ accounts for the likelihood that household h makes k purchases in the exponential purchasing state, given that it never transitions to the Erlang-2 purchasing state.

If a household does move to the more regular purchasing state, the inter-purchase times of the remaining $x-k$ purchases follow an Erlang-2 distribution. The likelihood of observing these purchases at times $t_{k+1}, t_{k+2},\ldots, t_x$ over the interval (t_k,T) is:

$$L2\left(t_{k+1},\ldots,t_x|\lambda_{h2},k\right)=\lambda_{h2}^{2(x-k)}e^{-\lambda_{h2}(T-t_k)}(\lambda_{h2}(T-t_x)+1)\left(\prod_{j=k}^{x-1}\left(t_{j+1}-t_j\right)\right)^{I(k<x)}. \tag{2}$$

Note that when $k=x$, the household transitions to the Erlang-2 purchasing state after making its last observed purchase and the likelihood contribution is simply the survival function of the Erlang-2 distribution for a duration of $T-t_x$.

Of course, we do not know exactly when the state transition occurs for each household (i.e., the precise value of k). A natural way to approach this issue is to assume that after each purchase, household h transitions to the Erlang-2 state with a probability of p_h. This results in a geometric process governing transition to the regular state, where the probability of transitioning after k purchases is given by:

$$P(k|p_h)=p_h(1-p_h)^{k-1}\text{ for }k=1,2,\ldots \tag{3}$$

Combining each of the pieces above, we have the household-level likelihood function based on household h's purchasing pattern T_h, conditional on the parameters λ_{h1}, λ_{h2}, and p_h:

$$L(\lambda_{h1},\lambda_{h2},p_h|T_h)=(1-p_h)^{x}L1(T|x,\lambda_{h1})+\sum_{k=1}^{x}P(k|p_h)L1(t_k|\lambda_{h1},k)L2\left(t_{k+1},\ldots,t_x|\lambda_{h2},T,k\right). \tag{4}$$

The first term on the right hand side of Eq. (4) accounts for the possibility that household h never transitions to the Erlang-2 state. The summation in the second term allows for the transition to occur after the k^{th} purchase with a probability of $P(k|p_h)$.

Eqs. (1)–(4) provide the household-level model for our evolving process story. In accounting for cross-sectional heterogeneity, we consider two different sources of heterogeneity in households' purchasing behavior. The first is in the household-level parameters governing the inter-purchase times in each state (λ_{h1} and λ_{h2}). The second is in the likelihood with which households transition from the exponential state to the Erlang-2 state (p_h). Though heterogeneity in the latter has been incorporated into previous research (e.g., Schmittlein et al., 1987), it is one of the key areas where our model specification diverges from FHH. We begin by assuming that purchasing rates for the exponential stage follow a gamma distribution with shape parameter r_1 and scale parameter α_1, and the rates for the Erlang-2 stage follow a gamma distribution with shape parameter r_2 and scale parameter α_2:

$$g_s(\lambda_{hs})=\frac{\alpha_s^{r_s}\lambda_{hs}^{r_s-1}e^{-\alpha_s\lambda_{hs}}}{\Gamma(r_s)}\text{ for }s=1,2. \tag{5}$$

To allow for heterogeneity in the transition probabilities for each household, we assume that the probabilities are drawn from a beta distribution with parameters a and b:

$$q(p_h)=\frac{p_h^{a-1}(1-p_h)^{b-1}}{\beta(a,b)}. \tag{6}$$

[2] We considered an alternative hierarchical Bayesian model specification in which this assumption was relaxed and found that performance in the holdout period suffered compared to our proposed specification and other benchmarks. Details of this specification are available from the authors upon request.

After accounting for heterogeneity in purchasing rates and the transition probability, the likelihood is given by:

$$L(r_1,\alpha_1,r_2,\alpha_2,a,b\,|\,T_h) = \int_0^1\int_0^\infty\int_0^\infty L(\lambda_{h1},\lambda_{h2},p_h\,|\,T_h)g_1(\lambda_{h1})g_2(\lambda_{h2})q(p_h)d\lambda_{h1}d\lambda_{h2}dp_h$$
$$= \left(\frac{\beta(a,b+x)}{\beta(a,b)}\frac{\Gamma(r_1+x)}{\Gamma(r_1)}\left(\frac{\alpha_1}{\alpha_1+T}\right)^{r_1}\left(\frac{1}{\alpha_1+T}\right)^{x}\right)$$
$$+\,1(x>0)\sum_{k=1}^{x}\left\{\begin{array}{l}\dfrac{\beta(a+1,b+k-1)}{\beta(a,b)}\left(\dfrac{\Gamma(r_1+k)}{\Gamma(r_1)}\left(\dfrac{\alpha_1}{\alpha_1+t_k}\right)^{r_1}\left(\dfrac{1}{\alpha_1+t_k}\right)^{k}\right)\times\\ \left(\dfrac{\Gamma(2(x-k)+r_2)}{\Gamma(r_2)}\left(\dfrac{\alpha_2}{\alpha_2+T-t_k}\right)^{r_2}\left(\dfrac{1}{\alpha_2+T-t_k}\right)^{2(x-k)}\left(1+\dfrac{(2(x-k)+r_2)(T-t_x)}{\alpha_2+T-t_k}\right)\left(\prod_{i=k}^{x-1}(t_{i+1}-t_i)\right)^{1(k<x)}\right)\end{array}\right\}. \quad (7)$$

The first term accounts for the likelihood that a household never leaves the exponential state. As such, all x purchases occur in this state and, conditional on remaining in this state, the likelihood of making x purchases in an observation period of length T is given by the exponential/gamma model. For those households who make at least one purchase ($x>0$), the summation term allows for the probability of a transition after the k^{th} purchase. In such a case, the first k purchases (that are made by time t_k) follow an exponential/gamma timing model, while purchases $k+1,\ldots,x$ (which are made between t_k and T) follow an Erlang-2/gamma model.

The calibration log-likelihood is then given by:

$$LL(r_1,\alpha_1,r_2,\alpha_2,a,b) = \sum_{h=1}^{H} ln[L(r_1,\alpha_1,r_2,\alpha_2,a,b\,|\,T_h)]. \quad (8)$$

3.2. Features of the evolving process model

The inclusion of process evolution and heterogeneity in the likelihood of transitioning to the regular purchasing state makes the proposed model extremely flexible and capable of explaining a number of different buying patterns. In the event of stationary repeat purchasing ($p_h=0$), the model collapses into the standard exponential/gamma model, the timing equivalent of the NBD count model (Gupta & Morrison, 1991).

Our framework also nests extant changepoint models (e.g., Howard, 1965; MacDonald & Zucchini, 1997). It may be the case that consumers will cease buying the product after making a few purchases. This shift to an "inactive" state was first modeled using the Pareto/NBD (Schmittlein et al., 1987). When $\lambda_{h2}\rightarrow 0$, our model becomes the BG/NBD, which has been shown to serve as a very close approximation to the Pareto/NBD (Fader, Hardie, & Lee, 2005). Households may exhibit distinct trial and repeat purchasing behavior. If $p_h=1$, a household will transition to the more regular purchasing state after its trial purchase, in which it will make all of its repeat purchases. When $p_h<1$, our framework allows for the possibility that households will require additional purchases in the "trial regime" before their behavior stabilizes and they move to the "steady state."

In its most general form, the proposed model allows for the transition of a product from "new" to "established," the speed of which can vary by household. We assume that inter-purchase times during the "established" stage of the relationship are more regular than in the initial stage. One interesting (and managerially useful) feature of the proposed model is its ability to estimate the parameters of the "steady state" – the Erlang-2 state in our model, in which households reside after transitioning – even when many households may not yet be in that state. Based on the behavior of those households that have made a number of repeat purchases, we can estimate the distribution of steady-state purchasing rates for all households as well as how quickly they will transition to the steady state. This is an advantage over the FHH model, as our formulation allows us to estimate the distribution of purchasing rates in the steady state; it is not easy to obtain an equivalent expression for FHH, and the authors made no attempt to provide it.

In addition to examining the process evolution model outlined in Eqs. (1)–(8), we also consider a two-state model where the purchasing in both states follows an exponential distribution. The household-level model can be derived in a straightforward manner by substituting $L1(T-t_k|\lambda_{h2},k)$ for $L2(t_{k+1},\ldots,t_x|\,\lambda_{h2},T,k)$ in Eq. (4), ultimately yielding a two-state exponential/gamma model. Like our proposed evolving process model, this alternative formulation has two distinct purchasing states and nests extant changepoint models, thus having the same advantages compared to FHH. As it does not consider process evolution, it offers a middle ground and lets us determine whether process evolution is empirically supported, or if multiple purchasing states (both of which follow an exponential timing distribution) is sufficient.

4. Empirical analysis

We examine the ability of the evolving process model using sales data for two products that underwent year-long tests in two markets prior to a national launch. For a more detailed description of the data, we refer readers to FHH. Using the first 26 weeks as a calibration period for each dataset, we estimate a series of models: (1) the stationary exponential/gamma (EG), (2) the stationary Erlang-2/gamma (ErlG), (3) the two-state EG, (4) FHH's non-stationary EG (NSEG), and (5) the proposed evolving process model. To assess the predictive ability of the proposed models, we compare models based on their holdout log-likelihoods during the following 26 weeks. This provides a rigorous assessment of model performance, as aggregated errors will not mask each other.

We first present the model results for the Kiwi Bubbles dataset in Table 1.

As the EG model is nested by each of the non-stationary models (the two-state EG, the NSEG, and the evolving process model), the value of a dynamic model becomes clear. While the dynamic models perform fairly similarly during the calibration period, we see that the evolving process model yields the best performance in the holdout period. In looking at the mixing distribution for the exponential purchasing state of each of the dynamic models (as given by estimates of r_1 and α_1), we also see that they yield similar parameter estimates. Interestingly, while the two-state EG model suggests that the second state is akin to an inactive state, the evolving process model yields a "steady state" that reflects a shorter expected inter-purchase time, compared to the "trial regime." Thus, inferences about the timing and

Table 1
Model comparison for the kiwi bubbles data.

	Parameters								Log likelihood		
	r_1	α_1	r_2	α_2	a	b	ψ	θ	Calibration	BIC	Holdout
EG	0.08	71.38	–	–	–	–	–	–	−3812	7641	−2118
ErlG	0.10	33.09	–	–	–	–	–	–	−3973	7963	−2323
Two-State EG	0.05	25.15	0.00	44.73	0.87	1.59	–	–	−3772	7593	−2077
FHH NSEG	0.05	24.06	–	–	–	–	0.85	1.14	−3772	7575	−2066
Evolving process	0.05	24.69	0.05	7.68	1.83	2.88	–	–	−3771	7589	−2063

extent of customer "death" may be highly sensitive to the model specification; the typical "buy 'til you die" frameworks (Fader et al., 2005; Schmittlein et al., 1987) may be severely overstating the magnitude of the "death" component.

To further explore the differences across these models, we next examine the distribution of transition probabilities from the exponential to the Erlang-2 purchasing state. In both the two-state EG and evolving process models, we see that there is approximately a 1-in-3 chance that households will transition to the second state after making a purchase ($\mu = a/(a+b) = 0.36$ and 0.39 for the two models, respectively). The estimated beta distribution under the two-state EG model, however, is more polarized than the evolving process model ($\varphi = 1/(1+a+b) = 0.29$ and 0.18, respectively). Thus, while some households will be fairly likely to transition quickly after a purchase, a greater proportion will be reluctant to do so. In contrast to the "backwards-J" shape of the mixing distribution under the two-state EG model, when we incorporate process evolution, the estimated beta distribution features an interior mode, implying a relatively high degree of homogeneity in the transition process. It is hard to say for sure which of these "stories" about the transition to a steady state is more valid, but this comparison suggests that there is value in developing and comparing alternative model specifications.

To further assess the predictive ability of the evolving process model, we turn our attention to the Four Seasons Biscuits dataset. As FHH mention, this dataset is characterized by a greater degree of sparseness (fewer buyers, fewer purchases per buyer) than the Kiwi Bubbles dataset. As such, it offers a challenge for any model that is attempting to capture evolving behavior. It is a particular challenge for the evolving process model, as a more regular purchasing pattern may be difficult to infer with limited purchases. In Table 2, we summarize the performance of the five models for the Four Season Biscuits dataset.

Again we find that the non-stationary models outperform the stationary EG model that they nest, as well as the stationary Erlang-2/gamma model. In this case, while the NSEG model has the lowest BIC, it is the two-state EG model that yields the best out-of-sample performance. Whereas the evolving process model yielded better performance compared to the other dynamic models for the Kiwi Bubbles dataset, as anticipated, it struggles with this sparser dataset.

Comparing parameter results across the models, as with the Kiwi Bubbles data, the initial purchasing states for the three dynamic models yield similar results despite the different mechanism of accounting for non-stationarity (i.e., a dynamic changepoint vs. a single but heterogeneous changepoint). The "steady state" in both the two-state EG and evolving process models point toward an increased purchasing rate once this state is reached by households. In contrast, while the NSEG model finds that a "steady state" will ultimately be reached for this product (because $\psi = 1$), it assumes the underlying distribution remains unchanged.

Taken together, both datasets provide support for the use of a flexible two-state model that allows one to distinguish the "trial regime" from the "steady state." Additionally, the Kiwi Bubbles dataset supports the notion of process evolution, which has not previously been considered in the extant models in marketing. While both datasets ultimately reveal faster purchasing in the "steady state," the empirical results may call for different managerial actions. While purchasers of Four Season Biscuits have a high expected probability of transitioning to the "steady state," managers must overcome the low expected purchasing rate in the "trial regime" and spur initial trial. In contrast, households are slow to transition from the "trial regime" to the "steady state" when purchasing Kiwi Bubbles. As such, marketers may need to employ tactics that will increase repeat sales to drive households into the "steady state." These stories serve to demonstrate both the flexibility and applicability of the two-state purchasing models. They also provide a stark contrast between the renewal model proposed by FHH and our two-state purchasing models, the latter of which indicates that the optimal decision for marketers may depend on the relative speed of purchasing in each state, as well as the rate at which the "steady state" is reached.

5. Conclusion

The intent of this research has been to indicate a direction in modeling that, despite being behaviorally sound, has been methodologically overlooked. Unlike previous models that account for evolution in customer preferences by updating purchasing rates for a given process, we present a model with two distinct states and empirically assess the need to incorporate evolution in the underlying purchase process. Both the two-state EG and the evolving process models account for changes that occur in customer purchasing as a product goes from "new" to "established." In addition to nesting a number of multiple-event timing models that are used within marketing and yielding managerially relevant insights, this formulation offers improved performance compared to the FHH model.

While the improvement in empirical performance compared to FHH's modeling framework is small, as revealed in Tables 1 and 2, two-state purchasing models afford managers a complete picture of the behavior that can be expected during both the "trial regime" and "steady state" phases of the relationship as well as the speed with which customers move into the "steady state." We find some support for incorporating process evolution and show that it can offer better performance and substantially different insights, while also resolving

Table 2
Model comparison for the four season biscuits data.

	Parameters								Log likelihood		
	r_1	α_1	r_2	α_2	a	b	ψ	θ	Calibration	BIC	Holdout
EG	0.05	160.90	–	–	–	–	–	–	−1865	3746	−1245
ErlG	0.06	65.56	–	–	–	–	–	–	−1960	3936	−1316
Two-State EG	0.02	18.09	0.18	54.66	24.84	5.84	–	–	−1842	3734	−1228
FHH NSEG	0.02	23.57	–	–	–	–	1.00	0.49	−1840	3713	−1232
Evolving process	0.02	17.79	0.18	23.10	39.80	11.00	–	–	−1835	3720	−1233

some of the inconsistent findings regarding the use of exponential and Erlang-2 models that have existed within the literature.

The notions of multiple purchasing states and evolving process are not only consistent with existing behavioral literature, but are also applicable to other settings besides new product sales. Future research may consider its performance with regard to newly acquired customers for a service provider. It may also be of merit to consider the inclusion of additional purchasing states in the customer relationship. Extending the original "buy 'til you die" framework of Schmittlein et al. (1987), one can envision a combination of exponential and Erlang-2 purchasing states followed by an inactive state. Another direction for future research would be to incorporate the behavioral work that explains changes in preferences. For instance, research examining changes in behavior based on underlying motivations (e.g. Yang, Allenby, & Fennel, 2002) could be incorporated into these multi-stage models to provide a more complete picture of customer behavior. Additionally, future research may address how individuals' decision-making strategies (e.g., Gilbride & Allenby, 2004) evolve over time.

There are limitations in our research that should be recognized. For instance, we did not consider the impact of covariates such as pricing and promotion, as our interest was purely focused on understanding the distinct states and the possibility of process evolution. A natural extension would be to include covariates in the model following the method prescribed by Gupta (1991) and applied by FHH. We may see systematic changes in covariate effects as customers move from one purchasing process to another. Just as there may be differences in the long-run effectiveness of marketing activities across households (e.g., Lim, Currim, & Andrews, 2005), activities may also differ in their effectiveness for the same household from one state to the next. Future research may also explore the impact of new product launches on the preferences and purchase processes for all brands within a category (Ramaswamy, 1997; Wagner & Taudes, 1991).

While we find that the best-performing model varies by dataset, the "winning" model in both cases features two purchasing states. Future work may examine the scenarios under which each of these alternative specifications (the two-state EG and evolving process) is best used. Such a meta-analysis across different product categories and across customers in different markets (e.g., Andrews & Currim, 2002; Gielens & Steenkamp, 2007) may reveal systematic patterns in evolutionary behavior. For instance, some categories may exhibit the "sleeper" pattern of slow trial and increased subsequent purchasing (as we see in both of the datasets used here), while the purchasing in other categories may mirror a "flash in the pan" pattern of frequent early purchasing that gives way to slower "steady state" purchasing.

Such a comparison can also provide insight into the differences in buying behavior exhibited at different types of stores. For example, purchases of new products at grocery and other brick-and-mortar stores may follow a relatively regular purchase cycle (e.g., Kahn & Morrison, 1989) and may be best explained by the evolving process model. It should be noted that, while an evolving process model may best explain the observed behavior in this scenario, it may be the case that demand arises under an exponential process but manifests under regular shopping trips. As such, while buying appears more regular, this may simply be due to increased purchase frequency under an exponential timing process.

While trips to brick-and-mortar stores may follow a regular purchase cycle, online purchases may be made only when the demand manifests and thus may therefore call for the two-state EG model. It remains a relevant empirical question if customers of an online store will gravitate toward a more regular purchase cycle or if this modeling framework is more appropriate for traditional stores. It may also be worthwhile to investigate timing models beyond the exponential and Erlang-2.

But perhaps the key, at least at this early stage of the research process, should not be to focus too much on identifying the timing distributions that provide the best possible fit; it is to get a better understanding of the evolutionary process itself. Perhaps it is appropriate to go back to lab experiments to better understand the behavioral theory on which this research draws (such as those conducted by Hoeffler and Ariely (1999), Meyer (1987), and West et al. (1996)) to try to better match the mathematical assumptions discussed here with the implicit process(es) that subjects follow as they go through their evolutionary transitions. We hope that quantitative and behavioral researchers will find common interest in this type of investigation.

References

Andrews, R. L., & Currim, I. S. (2002). Identifying segments with identical choice behavior across product categories: An intercategory logit mixture model. *International Journal of Research in Marketing*, *19*(1), 65—79.

Chatfield, C., & Goodhardt, G. J. (1973). A consumer purchasing model with Erlang inter-purchase times. *Journal of the American Statistical Association*, *68*(December), 828—835.

Fader, P. S., Hardie, B. G. S., & Huang, C. -Y. (2004). A dynamic changepoint model for new product sales forecasting. *Marketing Science*, *23*(1), 50—65.

Fader, P. S., Hardie, B. G. S., & Lee, K. L. (2005). "Counting your customers" the easy way: An alternative to the Pareto/NBD model. *Marketing Science*, *24*, 275—284 (Spring).

Fader, P. S., Hardie, B. G. S., & Zeithammer, R. (2003). Forecasting new product trial in a controlled test market environment. *Journal of Forecasting*, *22*, 391—410 (August).

Gielens, K., & Steenkamp, J. -B. E. M. (2007). Drivers of consumer acceptance of new packaged goods: An investigation across products and countries. *International Journal of Research in Marketing*, *24*(2), 91—111.

Gilbride, T. J., & Allenby, G. M. (2004). A choice model with conjunctive, disjunctive, and compensatory screening rules. *Marketing Science*, *23*(3), 391—406.

Gupta, S. (1991). Stochastic models of interpurchase time with time-dependent covariates. *Journal of Marketing Research*, *28*, 1—15 (February).

Gupta, S., & Morrison, D. G. (1991). Estimating heterogeneity in consumers' purchase rates. *Marketing Science*, *10*(3), 264—269.

Hardie, B. G. S., Fader, P. S., & Wisniewski, M. (1998). An empirical comparison of new product trial forecasting models. *Journal of Forecasting*, *17*, 391—410 (June-July).

Herniter, J. (1971). A probabilistic market model of purchase timing and brand selection. *Management Science*, *18*(12), 102—113.

Hoeffler, S., & Ariely, D. (1999). Constructing stable preferences: a look into dimensions of experience and their impact on preference stability. *Journal of Consumer Psychology*, *8*(2), 113—139.

Howard, R. A. (1965). Dynamic inference. *Operations Research*, *13*(5), 712—733.

Howard, J.A., & Sheth, J.N. (1967). Theory of Buyer Behavior. In R. Moyer (Ed.), Changing Marketing Systems: Consumer, Corporate and Government Interfaces (pp. 253-262), Proceedings, Winter Conference, American Marketing Association, 253-262.

Jeuland, A. P., Bass, F. M., & Wright, G. P. (1980). A multibrand stochastic model compounding heterogeneous Erlang timing and multinomial choice processes. *Operations Research*, *28*(2), 255—277.

Kahn, B. E., & Morrison, D. G. (1989). A note on "random" purchasing: additional insights from Dunn, Reader and Wrigley. *Applied Statistics*, *38*(1), 111—114.

Kalwani, M., & Silk, A. J. (1980). Structure of repeat buying for new packaged goods. *Journal of Marketing Research*, *17*(August), 316—322.

Lim, J., Currim, I. S., & Andrews, R. L. (2005). Consumer heterogeneity in longer-term effects of price promotions. *International Journal of Research in Marketing*, *22*(4), 441—457.

MacDonald, I. L., & Zucchini, W. (1997). *Hidden Markov models and other models for discrete-valued time series*. London: Chapman & Hall.

Meyer, R. J. (1987). The learning of multiattribute judgment policies.*Journal of Consumer Research*, *14*, 155—173 (September).

Moe, W. W., & Fader, P. S. (2004). Capturing evolving visit behavior in clickstream data. *Journal of Interactive Marketing*, *18*(1), 5—19.

Ramaswamy, V. (1997). Evolutionary preference segmentation with panel survey data: An application to new products. *International Journal of Research in Marketing*, *14*(1), 57—80.

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

Wagner, U., & Taudes, A. (1991). Microdynamics of new product purchase: A model incorporating both marketing and consumer-specific variables. *International Journal of Research in Marketing*, *8*(3), 223—249.

West, P. M., Brown, C. L., & Hoch, S. J. (1996). Consumption vocabulary and preference formation. *Journal of Consumer Research*, *23*(2), 120—135.

Yang, S., Allenby, G. M., & Fennel, G. (2002). Modeling variation in brand preference: The roles of objective environment and motivating conditions. *Marketing Science*, *21*(1), 14—31.