Marketing Science

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Commentaries and Reply on “Predicting Customer Value Using Clumpiness: From RFM to RFMC” by Yao Zhang, Eric T. Bradlow, and Dylan S. Small

To cite this article:
(2015) Commentaries and Reply on “Predicting Customer Value Using Clumpiness: From RFM to RFMC” by Yao Zhang, Eric T. Bradlow, and Dylan S. Small. Marketing Science 34(2):209-217. https://doi.org/10.1287/mksc.2015.0904

MARKETING SCIENCE
Vol. 34, No. 2, March–April 2015, pp. 209–217
ISSN 0732-2399 (print) | ISSN 1526-548X (online)

http://dx.doi.org/10.1287/mksc.2015.0904

Commentaries and Reply on "Predicting Customer Value Using Clumpiness: From RFM to RFMC" by Yao Zhang, Eric T. Bradlow, and Dylan S. Small

This series of discussions presents commentaries and a reply on Zhang et al. [Zhang Y, Bradlow ET, Small DS (2015) Predicting customer value using clumpiness: From RFM to RFMC. *Marketing Sci*. 34(2):195–208].

Keywords: customer lifetime value; RFM; clumpiness
History: Preyas Desai served as the editor-in-chief for this article.

http://dx.doi.org/10.1287/mksc.2015.0904a

Commentary on "Predicting Customer Value Using Clumpiness"

Vineet Kumar, Yale School of Management, New Haven, Connecticut 06520, vineet.kumar@yale.edu; and Kannan Srinivasan, Tepper School of Business, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213, kannans@cmu.edu

In this commentary, we provide our perspective on clumpiness, how it might be caused in practical settings, the implications of clumpy behavior for firm strategy, and detail methodological and measurement issues. In the process we point to open questions that would be well served by further research. Finally, we demonstrate how well the notion of customer lifetime value based on a hidden Markov model (HMM) is captured by measures of clumpiness suggested by the authors.

The paper by Zhang, Bradlow, and Small, in this issue of *Marketing Science*, builds on their work in Zhang et al. (2013), by using the idea of clumpiness and demonstrating that it has predictive power over and beyond the three commonly used variables of RFM: recency, frequency, and monetary value. We reiterate the authors' suggestion that these companion papers are highly complementary, and Zhang et al. (2013) provides highly insightful ideas regarding the idea of clumpiness. The theoretical grounding of these measures is derived from a set of desirable criteria (convergence, continuity, maximum/minimum, symmetry, and convexity) leading to a family of measures that have better statistical power and properties than commonly used measures (e.g., serial correlation or runs). In Figure 2 (p. 199), the authors provide a highly helpful framework that illustrates how clumpiness can play a role both as a dependent variable of interest and as a predictive variable for customer value and RFM variables. They demonstrate the empirical validity and importance of the clumpiness construct across a wide range of data sets, including both traditional and digital business settings.

Foundations and Types of Clumpiness

At the outset, we note that clumpiness is orthogonal to monetary value or the size of the transaction, i.e., five visits to a store in quick succession to buy one product on each visit is different conceptually than buying the exact same set of five products in a single shopping trip. A natural question the reader might have is whether clumpiness is caused merely by intertemporal substitution, so that consumers just bunch their visits or purchases together, which they might have made in a more spread out manner. The authors clearly demonstrate that clumpiness does have predictive power for a range of outcomes of interest over and beyond frequency, so the construct is independently important in a predictive sense, further suggesting that there are likely to be causal mechanisms that make clumpiness important from a practical perspective.

What causes a pattern of clumpiness in customer behavior? It is helpful to understand this because in most settings, we might expect consumers to intertemporally smooth their consumption patterns, and clumpiness is the exact opposite of that. We suggest a few mechanisms by which consumers might demonstrate clumpiness, and it would be helpful to have more, even though these might be specific to the setting. First, consider customers in a search mode (perhaps measured by visit clumpiness). For example, a customer might visit a car dealership multiple times before making a purchase; a house hunter might look at a number of houses in short order before making a purchase, and then displaying no search for several years, until she begins the process again. In such cases, higher variance in price or quality might lead to more clustered visits or clumpiness. Second, when consumers have serial correlation with respect to occasion availability, that would result in clumpiness. Customers often make more visits to stores in the holiday season than during the rest of the year. An amusement park might similarly find more customer visits during the summer than the rest of the year. Finally, there might be complementarity between consumption occasions leading to clumpiness. Consider the notion of binge-watching TV shows online (e.g., on Hulu), where consumers might watch a whole season in a few days, followed by no activity for several weeks.

We would distinguish between exogenous and endogenous (or induced) clumpiness. Consumers who are inherently or exogenously clumpy would be more likely to have periods of activity or events clumped together, independent of marketing interventions made by the firm.

Connection to heterogeneity/state dependence: The idea of clumpiness is also connected to the well-known literature on heterogeneity versus state dependence (see Heckman 1991 for an overview and Dubé et al. 2010 for recent research in this area). The idea that the conditional probability of an event is higher given that it has occurred recently within the same individual is indicative of state dependence, rather than heterogeneity. However, heterogeneity would lead to different degrees of clumpiness across consumers, as the authors find in their results (see Figure 3, p. 200), and have implications for why clumpiness might predict the monetary amount of lifetime value. Note that heterogeneity in preference is unlikely to appear as clumpiness, but would rather appear as a difference in frequency. As clumpiness is related to the monetary value (see results in Table 7, p. 205), this might permit profitable segmentation strategies based on the clumpiness variable. Further, since the authors find that clumpiness is highly impacted by the product category (see Tables 8 and 9, p. 205), it might be worthwhile to undertake an in-depth examination of how this is driven by differences in the buying process.

Clumpiness in Practice

A connection to clumpiness might involve "superconsumers," a small number of buyers who contribute a large proportion of revenues or profits (Yoon et al. 2014). Firms are increasingly focused on these consumers because despite their already high consumption levels, across a range of settings, these consumers also are more prone to increasing their consumption even further, because they use the product across a wider range of usage scenarios. Given that marketing to superconsumers would require specialized marketing messages beyond regular consumers, it would be worthwhile to investigate whether these superconsumers potentially demonstrate highly clumpy consumption patterns, and whether they are exogenously and/or endogenously clumpy.

Clumpiness in Digital Settings: The authors point out that in digital settings (e.g., Hulu), clumpiness becomes more important for two reasons. First, we learn that a larger proportion of consumers exhibit clumpy behavior, as illustrated in Figure 7 (p. 203). Second, clumpiness seems to have more impact on future consumer behavior that we care about as marketers, e.g., future visit frequency. It would be very useful to understand why we find more clumpy behavior in digital settings. Is it because it is easier to visit websites repeatedly rather than physical stores, due to lower transaction costs? We might also expect that consumer data is captured in finer granularity in digital settings, i.e., events are effectively specified in continuous time rather than discrete. An advantage with the present approach is that it allows us to determine the appropriate interval of inter-event time (IETs). For example, if a consumer watches two episodes of a show on Netflix or Amazon Instant Video during a day, perhaps one in the morning and one in the evening, it is possible to characterize these as two separate visits, and determine the period length to perhaps be half a day rather than a day. More broadly, the characterization of clumpiness in the paper allows the flexible specification of visits to be customized to the setting.

Clumpiness and Firm Strategy

Zhang et al. (2015) demonstrates across a range of settings that it might be profitable to increase (or decrease, in certain cases) clumpiness. It is therefore worthwhile to investigate the implications for firms.

Consumers who are endogenously clumpy would respond to firm choices by displaying more clumpy behavior. More concretely, consider again the notion of binge-watching several episodes (or even a whole season of a TV show) in a small number of viewings. With a traditional network or cable TV, such consumers could only watch the show once a week. Thus, now since consumers have a choice, we might consider binge watchers to be those who are more clumpy, whereas those who space out their viewings of episodes across weeks are less clumpy. This represents the notion of exogenous clumpiness. However, now the firm, say Amazon, could induce clumpiness by making the next episode of the show more prominent and easy to access (or even play automatically after the previous episode ends), or in the case of a paid show, offer a discount for the next episode if watched within a time frame. Thus, the strategic marketing choices made by firms could induce clumpiness, and cause consumers to intertemporally substitute their consumption. Zhang et al. (2015) specifically discusses this in §4, and details this in Figure 2 (p. 199).

Thus, consumers could be further classified as: (a) non-clumpy, both exogenous and induced, (b) exogenously non-clumpy, but induced clumpiness, and (c) both exogenous and induced clumpiness.

Promotions: A further example of endogenous clumpiness is present between clumpiness (especially visit clumpiness) and retargeting (or behavioral targeting), which involves showing the consumer advertisements for products based on their prior browsing patterns. For example, a consumer might visit an online seller, e.g., Zappos.com and view a pair of shoes, which then follow her as an advertisement as she visits the *New York Times* website, attempting to cause her to return to Zappos and complete a purchase. This might also have different implications for visit clumpiness and purchase clumpiness, and it would be helpful to build on the results in Lambrecht and Tucker (2013), who demonstrate the importance of targeting based on a consumer's search and decision process.

A few other promotional vehicles have strong potential to induce clumpiness (if designed well). Loyalty programs are often designed with rewards that have a specified expiration policy, and might induce clumpy behavior. Catalina's coupon printers at the checkout lane in supermarkets have the potential to induce clumpiness by incentivizing the consumer to make another shopping trip in the near future.

Product Complementarity: Firms that create complementary products might be more able to induce and take advantage of clumpiness, e.g., consecutive episodes of a TV show for a season are often designed to be complementary. Retailers might also invite consumers to purchase complementary products related to purchases they have already made, e.g., providing a targeted coupon for softener to a consumer who purchases detergent. Some online retailers like Amazon increasingly do this by prominently featuring complementary products during the next visit of the customer, potentially leading to more purchase clumpiness.

Operational Implications: Finally, clumpy customer behavior could have a significant operational impact on firms, e.g., a smaller number of clumpy customers requiring higher service levels in shorter timeframes. Are there settings where it might make sense to use clumpiness across different companies, so we capture the full range of customer behavior in a particular context, e.g., a consumer shopping online for auto insurance across multiple providers? The issue of behavior targeting also becomes increasingly important in such settings.

Measuring Clumpiness

Practical Appeal and Applicability: The appeal of the method in the present setting of customer lifetime value (CLV) is along several dimensions. RFM has long been a simple way to classify or segment customers, and is used extensively in practice. First, the authors define the clumpiness construct thoughtfully in consumer settings. Although the practitioner has flexibility in using potentially any of a number of functions to characterize clumpiness, the authors provide further guidelines to which measures perform better in practice. Second, the measure is clearly scalable in big data settings with a large *volume*, and its simplicity makes it easy to recompute in real-time for high *velocity* settings. A simple way to develop an intuition for clumpiness might be to think of the variance in the frequency of the events (the authors suggest second moment in Zhang et al. 2013), and this aspect becomes intuitively apparent when thinking of consumer visits as following a HMM (see below).

Clumpiness across a different number of events: Intuitively, the notion of clumpiness is that given there are n out of N possible events ($n < N$), a sequence is more clumpy when the events are bunched together (IET), compared to the case when they are evenly spaced. However, it might not make sense to compare clumpiness across a different number of events n for a fixed N, since IETs would be expected to decrease. Rather, it is best to think of the notion of clumpiness as having both fixed n and N, which leads to the question of whether there might be a natural way to compare clumpiness across different values of n.

Measures of clumpiness: Like the authors detail in Zhang et al. (2013), the measure of clumpiness can vary significantly over time, and it is nonmonotonic in events, implying a new event can either increase or decrease the measure of clumpiness. However, it might be possible to provide lower and upper bounds on the measure of clumpiness expected at a future time, which remains an issue for further research. Such an approach might allow for a better prediction because future clumpiness is likely to have an impact on lifetime value.

Clumpiness Elasticity: When consumers are induced to be clumpy, it would naturally lead to the issue of what interventions might induce most clumpiness, and in which consumers, perhaps leading to the notion of clumpiness elasticity, i.e., how a small proportional change in a marketing action might lead to a corresponding change in clumpiness. This elasticity might become an important metric over and beyond the traditional ones in assessing how response to marketing interventions might be driven in the form of clumpiness.

Multiple types of events: It might be helpful to explore how clumpiness could be extended to incorporate multiple types of events. In Zhang et al. (2015), the authors deal with this in the form of visit and purchase clumpiness, and effectively treat them as two separate sequences. Is there a way to define an integrated notion of clumpiness that captures the distinctions between different event types but also recognizing the similarities between them, and using that information in computing an overall clumpiness?

Clumpiness Predicts Customer Lifetime Value in a Hidden Markov Model

We next attempt to determine how clumpiness might be useful in predicting customer behavior in a HMM. We use a HMM since that is a parsimonious way to generate serial state dependence, and like the authors have suggested, it has been used in a wide variety of consumer settings (Netzer 2008).

This is connected most closely with §4.6 in Zhang et al. (2015) where the authors examine how in-sample RFM and C can predict out of sample monetary value.

Note that since the authors intended their approach to be simple and used RFMC, rather than inferring the latent state and making predictions conditional on the latent state, we attempt to do the same. Following the companion article (Zhang et al. 2013), we illustrate a consumer HMM of purchase behavior with the following parameters, where $y_t \in \{0, 1\}$ denoting purchase or no purchase, and $Z_t \in \{1, 2\}$ denoting the latent state of the consumer in period t.

We define two types of customers, A and B, with different purchase probabilities across the two states. The proportion of type A consumers is Φ. For a customer of type A, we have

$$y_t^A \sim \begin{cases} Bernoulli(p_1^A), & Z_t = 1, \\ Bernoulli(p_2^A), & Z_t = 2. \end{cases}$$

For both customer types, the transition matrix is specified as

$$\Theta = \begin{pmatrix} 1-\theta_{12} & \theta_{12} \\ \theta_{21} & 1-\theta_{21} \end{pmatrix}.$$

The parameters are set to baseline values: $p_1^A = p_2^A = 0.5$, $p_1^B = 0.25$, $p_2^B = 0.5$, $\theta_{12} = 0.1$, and $\theta_{21} = 0.1$. Thus, type A consumers do not vary their purchase behavior based on their latent state whereas type B consumers have a different probability of purchasing based on their state. We set each purchase to be a constant \$1 profit margin, and predict CLV based on RFM and RFMC, using a discount factor of $\beta = 0.99$ between periods. We use a simulation of $M = 10{,}000$ simulations over $T = 100$ periods. We use the first 50 periods to compute the RFC variables, and aim to predict the CLV for the next 50 periods based on this.

We can include both persistent heterogeneity and true state dependence with the HMM. Persistent heterogeneity is modeled by consumer type (A) or (B). The regression is specified with CLV as the dependent variable, with a subset of RFC as the explanatory variables (monetary value is fixed). The population proportion of Type A consumer is varied as $\Phi \in \{0, 0.25, 0.5\}$. We detail two measure of fit in Tables 1 and 2. In Table 1, we list the sum of squared residuals for different degrees of consumer heterogeneity (i.e., the proportion of type A customers). We find that the sum of squared residuals is lowest when clumpiness is included along with frequency, suggesting that these two variables are capturing different types of variation. However, the model with just recency and clumpiness has the highest sum of squared residuals, suggesting that these variables might not be sufficient to capture the dynamics of consumer behavior in the data. It is also interesting to note that the models with just frequency and clumpiness, but excluding recency, does almost as well as the full RFC model. Tables 2 and 3 detail the $\mathbf{R}^2$ corresponding to these models, and the patterns here are also consistent with the same logic we have previously described.

Table 1 Sum of Squared Residuals

Consumer heterogeneity/model	$\Phi = 0$	$\Phi = 0.25$	$\Phi = 0.5$
RF	195,398.3	186,484.4	171,931.7
RC	374,627.2	613,108.9	766,364.5
FC	142,429.1	152,866.4	148,549.3
RFC	141,467.1	152,656.1	147,796.4

Table 2 Adjusted R^2

Consumer heterogeneity/model	$\Phi = 0$	$\Phi = 0.25$	$\Phi = 0.5$
RF	0.9133	0.9298	0.9443
RC	0.8339	0.7691	0.7518
FC	0.9368	0.9424	0.9519
RFC	0.9373	0.9425	0.9521

Table 3 Multiple R^2

Consumer heterogeneity/model	$\Phi = 0$	$\Phi = 0.25$	$\Phi = 0.5$
RF	0.9134	0.9298	0.9448
RC	0.8339	0.7692	0.7589
FC	0.9368	0.9425	0.9529
RFC	0.9373	0.9425	0.9529

Conclusion

Overall, Zhang et al. (2015) makes a strong case for including clumpiness as an additional construct over and beyond RFM, and managers would be well served to use RFMC in making decisions, especially in settings where consumers display variation in clumpiness. Clumpiness thus captures an important degree of variation in customer behavior that is increasingly important in a wide variety of settings. Like the authors convincingly suggest, we believe that there is likely to be a strong potential in managing (and in most cases increasing) customer clumpiness, and practitioners would be well advised to seek out active ways of measuring and managing this important new construct.

References

Dubé JP, Hitsch GJ, Rossi PE (2010) State dependence and alternative explanations for consumer inertia. *RAND J. Econom.* 41(3):417–445.

Heckman JJ (1991) Identifying the hand of past: Distinguishing state dependence from heterogeneity. *Amer. Econom. Rev.* 81(2):75–79.

Lambrecht A, Tucker C (2013) When does retargeting work? Information specificity in online advertising. *J. Marketing Res.* 50(5):561–576.

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

Yoon E, Carlotti S, Moore D (2014) Make your best customers even better. *Harvard Bus. Rev.* (March). http://hbr.org/2014/03/make-your-best-customers-even-better/ar/1.

Zhang Y, Bradlow ET, Small DS (2013) New measures of clumpiness for incidence data. *J. Appl. Statist.* 40(11):2533–2548.

Zhang Y, Bradlow ET, Small DS (2015) Predicting customer value using clumpiness: From RFM to RFMC. *Marketing Sci.* 34(2): 195–208.

---

http://dx.doi.org/10.1287/mksc.2015.0904b

Comments on "Predicting Customer Value Using Clumpiness from RFM to RFMC"

Vithala R. Rao, S.C. Johnson Graduate School of Management, Cornell University, Ithaca, New York 14853, vrr2@cornell.edu

I appreciate the opportunity to prepare a few comments on this technically sophisticated paper (Zhang et al. 2015 or ZBS) on predicting CLV using information on the density of purchase events in a time period (measured as clumpiness index or C). It augments the three summary indexes of Recency (R), Frequency (F), and Monetary value (M), often used in practice. This paper builds on the seminal work by Fader and his colleagues (2005) and a recent publication by the authors (Zhang et al. 2013) cited in the ZBS paper. Because the model captures the variability in the inter-visit purchase time in a unique sense (i.e., the C-index), it may improve the predictability of CLV beyond the RFM framework. I expect the augmented approach of RFMC to have a similar impact as the work by Fader et al. (2005).

In this note, I will take a pragmatic view and identify some issues that may need to be addressed if the paper's approach were to be widely employed in practice by less sophisticated analysts or managers. I will cover four major items: salient features of this paper, computation of the C-index, some missing aspects, some ideas to make this method more useful including some suggestions for future work.

Salient features of the paper: The paper relies on analysis of large data sets of customer purchases (or visits) to a website and utilizes known probability distributions of time-dependent purchases (or visits) to draw inferences on CLV. The novel feature is the ease of computation of the clumpiness index (or C), which is simple and elegant. The C-index is computed at the individual level and its variability across selected individual characteristics is shown in the paper. The paper considers the advantage of using C both as a predictor and an outcome variable. The method relies on the availability of large samples of past observations to compute the C-index. But, the details of how the C-index enhances the computation of CLV are not easy to decipher in the paper.

In my opinion, the probabilistic formulation for CLV and RFM in Fader et al. (2005) is somewhat intricate for practitioners. It may therefore be worth focusing on the simple formula for $\text{CLV} = mr/(1+r-d)$, where m is the monetary value of a purchase, r is the retention rate, and d is the discount rate (Gupta and Lehmann 2005) to gain insights on the value of the C-index. While this simple formula does not capture the longitudinal variability in the three parameters (r, m, and d), which can be random, it can provide a good approximation to CLV. Further, managers can update customer valuations as additional information on customer behavior for computing the C-index (which affects m, r, and d) is acquired.

Computation of the C-index: ZBS utilize the starting time and end time of the observation period as events in computing IETs for a given n (number of events) and N (length of observation period). Such starting and end points of the observation period are necessarily arbitrary. If the start is at the time of the first event in the sequence and the end is at the last event in the observation period, the clumpiness measure may yield a different value.

For example, consider the two customers B and C shown in Figure 1 of ZBS (p. 198) with purchase events at (2, 3, 4, 27, 28, and 29) and (13, 14, 15, 16, 17, and 18). According to the computation used in the paper, Customer B is clumpier with an H_p-value of 0.48 than Customer C with an H_p value of 0.34.

Recomputed C-indexes with the starting time as that of the first event and end time as the time of the last event (or effectively dropping the first and last term in the summation term of Equation (1), p. xx), the C-index for B and C will be H_p $(B) = 0.66$ and H_p $(C) = 0.72$. The implication is that Customer C will now be considered clumpier relative to B. Such a difference can have implications for segmentation and other uses of the clumpiness index.

A related point on the computation of the C-index is whether the measure is well defined if the number of events is small or if the length of the observation period is short. It is also worth thinking about whether the measure captures the duration of clumped observations well.

Some missing aspects: I will now identify some aspects of ZBS that may lead to future research possibilities and perhaps increase the usability of the C-index to managers. First, ZBS do not delve into underlying behavioral processes of customers. Given that the paper's focus is on past data on purchases (or visits) to a particular firm's website, the method does not show how competition affects the results. Also, it is unclear the level at which the model is specified (is it the firm, product category, or brand?); presumably the level is the same as the level at which observations are made.

Further, it is unclear if the analyst can get data for customer activities offline or on competitive websites to capture the full customer purchase history in the computation of the C-index. This will enable revealing a holistic picture of customer behavior, which can be very valuable to managerial decisions. As with most of the probability models, the outcome is for an average customer with details on the variability; however, it appears that the individual characteristics are not directly included in the calculations.

A potentially fruitful avenue of research is to develop a direct relationship between the RFM or RFMC variables and the variables (m and r) of the simple CLV formula. Presumably, R and F relate directly to the retention rate (r). One may conjecture that r will increase as purchases are more recent and more frequent. The C-index modifies the computed value of r (by adjusting it up or down in some fashion). If this is possible, the revised r in the simple CLV formula can be expressed as $r' = rf(C)$, where $f(C)$ is the modifier to r for a customer with a clumpiness index of C. Then, the revised CLV for this customer will be $\text{CLV} = mr'/(1 + r' - d)$. I realize that this interpretation is not technically complex as in the published works on the relationship of RFM to CLV, but it may express the essential ideas of ZBS, which can be further developed.

While the authors show the relationship of the C-index to several individual characteristics and marketing activities (such as email and direct marketing), it is not clear how these are included in the ultimate calculation of CLV.

It may behoove the authors to seek additional information from customers (via small-scale survey) to understand the conditions under which purchase or visit activities are clumped. Such information may enable making adjustments in the segmentation of customers based on the value of their C-index and therefore their CLVs.

In a similar vein, it may be useful to apply the theory of runs and state dependence in the calculation and interpretation of the C-index.

It is not so clear whether firms would prefer customers who exhibit clumpy behavior versus customers who purchase items from the firms on a regular basis. Clumped purchases would imply severe costs for the sellers both in marketing and operations/logistics. Once such costs are evaluated, it may be necessary for firms to encourage customers to purchase regularly. This aspect of benefits and costs of clumpy behavior calls for further exploration.

Some Ideas to Enhance the Utility to Managers: In this section, I am thinking of managers who wish to use current advances in customer relationship management (CRM) methodologies in their decision making and who also appreciate the import of the kind of work in ZBS but perhaps do not possess highly advanced math/stat skills. Under such conditions, I wonder how they can implement the ideas of the paper. I believe that the authors can suggest ways to enable increasing use of this concept of RFMC (particularly the C-index) by clearly spelling out how the C-index can be employed for various managerial actions based on CLV.

Based on the results presented in ZBS, it is quite obvious that the C-index requires a large series of purchase (or visit) observations. Naturally, this can limit its use for situations when the purchase histories are quite small or when there are no past data altogether. One important context is managing new products. It will be beneficial to describe how these ideas can be utilized in such contexts.

One possibility is to develop an adaptive approach to measuring CLV; first a measure of CLV will be computed for every customer using prior guesses or judgments on (m, r, and d). As new customer behavior data become available, the new values of m, r, and possibly d can be used in refined values of CLV. In this process, one can conceptualize using the RFMC approach to obtain more refined CLV numbers. Of course, such an adaptive approach implies that some direct connections can be developed between m, r, and possibly d and the RFMC data. The adaptive approach also will ensure measurements to be quite current.

Related to this, the authors can suggest ways to determine the smallest sample size necessary for using the approach of RFMC. While we now have access to large sets of data with a significantly large number of observations, it may be worth looking into the issue of minimum sample size to draw reasonable inferences.

Given that the RFM to CLV model relies so heavily on advanced probability distributions, I wonder whether new approaches with distribution-free methods can be developed. Going back to my theme, managers may find distribution free methods more appealing to understand.

Having observed that the implications of the *C*-index can differ depending on the way the periods are accounted for (as shown above), it will be useful to determine the robustness of managerial actions for differing ways of computing the *C*-index. Such sensitivity analysis will enhance the confidence of managers in using the RFMC approach. In general, what is the effect on managerial decisions for customers if the *C*-index is missing or if it is miscalculated? What is the consequence on outcomes measured in terms of profit or number of customers, etc., over time? Analyses required to answer such questions may require large scale simulations.

I also think that managers will benefit by showing how the *C*-index (and in general *R* and *F* measures) affects the retention rate of customers. The underlying model needs to be clarified as discussed earlier.

A comprehensive study of how the *C*-index varies by product category, firm, type of retail setting (offline versus online), etc., and will assist managers in benchmarking their own product situation against some norms.

Conclusion: I admire the authors in bringing a new thinking to CRM by implicitly asking researchers to look at IETs as well as RFM and for measuring CLV of customers with the emphasis on the *C*-index. I think that the authors can add more value by positing a sound theory for understanding the *C*-index and including several effects such as competition. While it is tempting to restrict the analysis to available data (*albeit* a large data set or several large data sets), it may be useful to consider the general problem related to customer buying products/services over time and look for the appropriate data. Perhaps additional data can be collected to understand the clumpiness behavior and incorporate such understanding in managerial decisions.

The managerial consequences of omitting the *C*-index need to be better understood. Several people will benefit if the inclusion of the *C*-index can be brought to the level of a simple formula for CLV such as the one mentioned before (Gupta and Lehmann 2005). Attention needs to be paid to situations with limited or no purchase (or visit) history data.

References

Fader PS, Hardie BGS, Lee KK (2005) RFM and CLV: Using iso-value curves for customer base analysis. *J. Marketing Res.* 42(4):415–430.

Gupta S, Lehmann DR (2005) *Managing Customers as Investments: The Strategic Value of Customers in the Long Run* (Wharton School Publishing, Upper Saddle River, NJ).

Zhang Y, Bradlow ET, Small DS (2013) New measures of clumpiness for incidence data. *J. Appl. Statist.* 40(11):2533–2548.

Zhang Y, Bradlow E, Small DS (2015) Predicting customer value using clumpiness: From RFM to RFMC. *Marketing Sci.* 34(2): 195–208.

http://dx.doi.org/10.1287/mksc.2015.0904c

Empirically Testable Sources and Implications of Clumpiness

Yao Zhang, Fixed Income Research Division, Credit Suisse, New York, New York 10010, yao.a.zhang@gmail.com; Eric T. Bradlow, and Dylan S. Small, The Wharton School, University of Pennsylvania, Philadelphia, Pennsylvania 19104 {ebradlow@wharton.upenn.edu, dsmall@wharton.upenn.edu}

We would like to thank Professors Kumar, Srinivasan, and Rao for three striking aspects of their discussions of our work on clumpiness that a posteriori make us feel even better about the potential practical and academic impact of our work than immediately after Zhang et al. (2015) was accepted.

First, each of these discussants actually computed *C* (noting how easy and fast it was to do) and applied it! In the Kumar and Srinivasan discussion, it was utilized to provide a more nuanced understanding of the relationship between HMM-simulated CLV and *R*, *F*, *M*, and *C* (extending the work of Zhang et al. 2013—ZBS 2013, who related the parameters of the HMM to *C* but not CLV), whereas in the case of Rao's discussion it was to show the sensitivity of *C* to how one treats the start and end IETs. Simply put, we were happy that the first goal of this research, come up with a summary measure that people could easily compute and apply to different questions, was accomplished.

Second, these two discussions highlight exactly what our paper is and what our paper is not. What the paper does, as summarized by Kumar and Srinivasan, is to apply ZBS 2013 who establish a new statistically powerful measure of clumpiness, to assess whether *C* adds predictive power to out-of-sample CLV (and its subcomponents) beyond that of R, F, M, and marketing action. Thus, more broadly stated and framing this work within extant literature, our paper is about data compression and data sufficiency. Is RFM sufficient to compress a customer history? The answer is no in a broad set of cases. Can we assess by simply computing a statistic on each consumer when RFM-based segmentation and probability models derived using these summaries (e.g., buy-to-you-die (BTYD) models, Fader et al. 2005), before fitting the models, are sufficient? Yes, compute the fraction of clumpy customers in your data set. Thus, Zhang et al. (2015) is a measurement paper and one that shows that marketers should care about the value of C because predicting CLV matters to marketers. What this paper is not (and generates many of the discussants questions) is a paper that posits a theory of clumpy behavior nor a paper that does more than provide a first exploration of its potential drivers (product category, demographics). It is for this reason that we thank the discussants because we believe that each of the questions they raised (exogenous versus not, its usage for new products, etc.) are all viable research questions that suggest that understanding the clumpy customer is not a single paper research area.

Last, and this is reflected in the title of our rejoinder, each of the points raised by the discussants is empirically testable and has managerial implications. That is the beauty of C: compute it, test your question, and apply it! In particular, we list some of the questions raised by reviewers and how a researcher armed with C and the RIGHT data (beyond that which was available to us) can answer these questions using simple methods and then firms can apply the answers.

- *Intertemporal Substitution*: Whereas it is possible that clumpy behavior could be generated by customers who simply forward buy, our analyses show that conditional on F, people with high C purchase more out-of-sample. Thus C is likely to be generated by a different mechanism which longer time-series might be able to tease out. If in fact, though, higher C does mean higher value, then maybe a fraction of high C customers should be part of the marketer's metric dashboard (Farris et al. 2010).
- *Consumer Experience*: If one had a data set that linked transactions/visits to birth date in a product category, one could simply correlate time-varying clumpiness with product category experience. Understanding the coincidence between binge consumption and information search has implications for new product launches and the optimal amount of information firms should provide.
- *High variance in prices/quality*: In Zhang et al. (2015), we informally linked clumpiness to digital (more clumpy) versus non-digital goods. A more formal analysis, if clumpiness could be computed across a large number of product categories, where each category was described by a vector of attributes, would be to relate the fraction of clumpy customers to variables like price and quality dispersion. This could provide a roadmap for firms to where more valuable long-term CLV customers exist.
- *Stockouts/Product Availability*: While Zhang et al. (2015) provided mainly customer-level explanations for clumpiness, the discussants both point out that the firm is not a passive player and that firm-level actions may inadvertently or intentionally drive clumpiness (i.e., people cannot consume a product that is stocked out). Optimizing the pattern of consumption via product supply is an interesting managerial area for future study.
- *Clumpy, is it good or bad*?: While as marketers, we assumed that high C which leads to higher CLV is a good thing, this may not be the case. Both reviewers point out the operational costs associated with uneven consumption and thus since C provides better estimates of future consumption patterns, de facto, it also provides better estimates of potential operational costs due to uneven demand.
- *Clumpy with You, but how about with Others*?: Both discussants discuss this issue from a different perspective. Kumar and Srinivasan question whether clumpiness is endogenous (an inherent trait if you will) versus something that can be manipulated exogenously. If the former, clumpy with you may imply clumpy everywhere. Rao discusses this issue from a different perspective, competition. For one to get a full sense of clumpy behavior, the site-centric data used in Zhang et al. (2015) does, in hindsight, seem woefully inadequate to get a full view of the customer's clumpiness propensities. Better data, likely panel-centric, would help to address this issue and would allow firms to understand if the customer is clumpy or is the customer clumpy in this product category?

We would like to conclude our rejoinder by noting a few future research areas that seem quite promising beyond the empirically-oriented ones described here. First, when ZBS 2013 was written, we tried to reverse engineer Fader et al. (2005) and find the probability model for which RFMC were the sufficient statistics. The Weibull counting model of McShane et al. (2008) yields a model where all of the IETs are sufficient, but its connection to clumpiness is not direct. Second, as pointed out by both discussants, relating clumpiness to work on state dependence is a fruitful area of study and is closely aligned with papers that try to tease apart heterogeneity from state dependence (Dubé et al. 2010). This may unlock the key to finding a probability model with RFMC as sufficient. Finally, a closer look at the relationship between goal achievement (e.g., completing a series of a recorded show) and clumpiness using behavioral theories as explanations (e.g., goal-gradient, Kivetz et al. 2006) is likely to provide greater insights for those looking to apply our clumpiness measures.

References

Dubé JP, Hitsch GJ, Rossi PE (2010) State dependence and alternative explanations for consumer inertia. *RAND J. Econom.* 41(3):417–445.

Fader PS, Hardie BGS, Lee KK (2005) RFM and CLV: Using iso-value curves for customer base analysis. *J. Marketing Res.* 42(4): 415–430.

Farris PW, Bendle NT, Pfeifer PE, Reibstein DJ (2010) *Marketing Metrics: The Definitive Guide to Measuring Marketing Performance*, 2nd ed. (Pearson Education, Upper Saddle River, NJ).

Kivetz R, Urminsky O, Zheng Y (2006) The goal-gradient hypothesis resurrected: Purchase acceleration, illusionary goal progress, and customer retention. *J. Marketing Res.* 43(1): 39–58.

McShane BB, Adrian M, Bradlow ET, Fader PS (2008) Count models based on Weibull interarrival times. *J. Bus. Econom. Statist.* 26(3):369–378.

Zhang Y, Bradlow ET, Small DS (2013) New measures of clumpiness for incidence data. *J. Appl. Statist.* 40(11):2533–2548.

Zhang Y, Bradlow ET, Small DS (2015) Predicting customer value using clumpiness: From RFM to RFMC. *Marketing Sci*. 34(2): 195–208.