Marketing Science Institute Working Paper Series 2020
Report No. 20-120

Timing Customer Reactivation Initiatives

Niels Holtrop, Jaap E. Wieringa

Timing Customer Reactivation Initiatives

Niels Holtrop[1]

Jaap E. Wieringa[2]

July 1, 2020

[1] Niels Holtrop, Assistant Professor of Marketing, School of Business and Economics, Maastricht University, Tongersestraat 53, 6211 LM Maastricht, The Netherlands; n.holtrop@maastrichtuniversity.nl.

[2] Jaap E. Wieringa, Professor of Research Methods in Business, Department of Marketing, Faculty of Economics and Business, University of Groningen, PO Box 800, 9700 AV Groningen, The Netherlands; J.E.Wieringa@rug.nl.

Acknowledgements: This work is part of the first author's dissertation, written at the University of Groningen. The authors would like to thank the anonymous retailer for its cooperation in this study. Furthermore, the authors would like to thank Dylan Beijers, Arend-Jan Hengeveld, Aurélie Lemmens, and Thomas Reutterer for their comments on earlier versions of this manuscript. Participants at the 2016 and 2019 EMAC Conferences, 2016 and 2019 Marketing Dynamics Conferences, and 2018 Marketing Science Conference as well as seminar participants at Goethe University, KU Leuven-Groningen Research Seminar, Maastricht University, University of Amsterdam, and University of Groningen also provided valuable feedback on this work.

Timing Customer Reactivation Initiatives

Abstract

Firms operating in non-contractual settings use customer reactivation initiatives to stimulate inactive customers to resume purchasing. Customer reactivation requires these firms to identify which customers have become inactive at what point in time, and then approach them through direct mail or e-mail. Existing approaches struggle to separate active from inactive customers, and do not provide calendar time estimates of when to send a reactivation mailing. Addressing these shortcomings, we develop an approach to target and time the sending of reactivation mailings. Building on control chart methods, we introduce a gamma–gamma control chart, modeling the average customer interpurchase time and the variation therein to determine inactivity boundaries. Crossing these boundaries generates an inactivity signal, which should trigger reactivation. Comparing our control chart approach to multiple competing models establishes the improved ability of our chart to predict customer activity. Additionally, a field test in the greetings and gifts industry illustrates that this approach increases overall activity by 1.9–3.5 percentage points, and leads to incremental activity and revenue gains of 111% and 38%, respectively.

Keywords: customer reactivation, control chart, customer base analysis, field test, non-contractual setting

Consider a marketing professional charged with managing the customer base of an organization operating in a non-contractual setting (e.g., retailing, catalogs, charity donations) faced with the following challenges. First, she observes that many customers interact (e.g., purchase or donate) regularly with her organization (i.e., active customers). However, over time some customers reduce their interaction frequency while others stop interacting entirely (i.e., inactive customers). Foremost, she would like to identify *which* individual customers are active and which are inactive. Second, she would like to know *when* they become inactive. As customers can interact at any time with the organization, are heterogeneous in their interaction frequency, and do not interact at set time intervals, detecting such a change in their interaction pattern is not straightforward. Third, given information on which customers (are likely to) become inactive at what time, she can then proactively reach out to these customers (e.g., through direct or electronic mail) in an attempt to restore their interaction frequency to prior levels or to avoid permanent discontinuation of their interactions. We term such attempts *customer reactivation*. While customer reactivation is addressed frequently in professional publications[1], and numerous practical tips on how to execute customer reactivation are available (e.g., Pokorynik, 2017; Stevens, 2017), academic guidance on customer reactivation and its execution is limited. In this article, we address this question with the aim of implementing effective customer reactivation initiatives.

The implementation of customer reactivation initiatives depends on 1) identifying inactive customers, 2) the time at which they become inactive, and 3) the initiation of a (direct or electronic) mailing. Identifying inactive customers in non-contractual settings has been studied extensively in the existing literature. Model-based approaches are most common, giving rise to stochastic models such as the Pareto/NBD model (Schmittlein. Morrison & Columbo, 1987) and its variants (e.g., the BG/NBD model) (Fader, Hardie & Lee

[1] A Google search returns over two million hits.

2005a). These models provide accurate customer base-level predictions of quantities including future purchases and customer lifetime value (CLV). Importantly, however, Wübben and Von Wangenheim (2008) find that such models perform poorly when predicting individual-level future purchases (e.g., as required for determining customer-level inactivity). The authors recommend that simple business rules, such as the time-since-last-purchase heuristic, be used. However, intuitively, business rules may be too generic, especially in highly heterogeneous non-contractual settings, where customers vary wildly in their time between interactions. For example, in Figure 1 we show that the classification performance of the business rule considered by Wübben and Von Wangenheim (2008) depends on the purchase frequency of customers, with either high- or low-frequency customers being inaccurately classified in two of the three datasets considered in this paper. Further supporting this intuition, we later show that more specialized models not considered by Wübben and Von Hangenheim (2008) outperform heuristic approaches.

Research is more scarce regarding the time at which customers become inactive and the consequences of sending a mailing in response. Studies have acknowledged that intervention timing is important, as prolonged inactivity increases recency and thereby decreases future purchase likelihood (Bult & Wansbeek, 1995; Fader, Hardie & Lee 2005b; Neslin et al., 2013). Furthermore, the amount of time between mailings has been shown to affect overall customer equity (Drèze & Bonfrer, 2008). Finally, guidance on intervention timing should be based on date estimates rather than on activity probabilities produced by stochastic models, as managers make decisions—including those on customer reactivation—in calendar time (Dew & Ansari, 2018; Korkmaz, Kuik & Fok 2013). Indeed, the managerial question is *when* an action should be taken; for marketing scheduling purposes, this is more easily communicated in terms of a calendar date. Since e-mail—which has a short lead time—is used in 90% of customer reactivation initiatives (McGee, 2016), swift and efficient scheduling is a requirement. The question then is how to provide such calendar time estimates linked to the moment at which a customer becomes inactive and whether sending a mailing in response to this event can restore customers' interaction frequency.

In this article, we develop an integrated approach to customer reactivation, taking into account the abovementioned considerations. We rely on techniques developed in the statistical quality control literature (Montgomery, 2009). This stream of research is concerned with monitoring and controlling (industrial) processes in the face of process variability, making it ideally suited for our goals. In particular, we focus on a technique called control charts (Shewhart, 1931), which is used to monitor a process variable (e.g., the average concentration of tin in a chemical bath) and to detect situations in which the process does not obey its requirements (Wieringa, 1997). We fuse this control chart approach with existing models for purchasing in non-contractual settings (e.g., Colombo & Jiang, 1999; Fader, Hardie & Lee 2005b) to develop a customer monitoring system that provides frequent updates on the (in)activity of customers in calendar time. The system is calibrated to give a signal when a customer is deemed inactive for too long given her regular purchase pattern, which is used as an indicator to send a reactivation mailing.

We demonstrate the validity of our approach in two empirical studies. First, we compare our control chart approach to a wide range of competing models identified from the literature as well as the business rule described in Wübben and Von Wangenheim (2008). Across three different datasets, we establish that the control chart approach is better able to predict customer (in)activity. Second, to assess validity in the field and to determine whether sending mailings to inactive customer is efficacious, we perform a field test using data and systems provided by a firm in the greetings and gifts industry. We demonstrate the efficacy of our approach at the customer base level, as customer activity increases by 1.9 percentage points compared to the current firm policy. We also investigate the incremental activity and economic value of our approach and show that it increases these measures by 111% and 38% respectively, compared to the current firm policy.

The remainder of this article proceeds as follows: In the next section, we review the existing literature related to customer reactivation and provide an introduction to control charts. Next, we outline the development of our approach, elaborating how we can adapt existing stochastic purchase models to be combined with control charts. We then show the results of our comparison study, followed by an overview of our field test and its results. We conclude with a discussion on the implications of our work and provide directions for future research.

Research Background

Customer reactivation aims to stimulate customers who have reduced or ceased purchasing for some time to resume their purchases by sending them a mailing (Blömeke, Clement & Bijmolt, 2010). The successful implementation of customer reactivation initiatives in a non-contractual setting requires firms to 1) determine which customers have become inactive, 2) the specific time this occurs, and 3) contact these customers with a reactivation mailing at that point in time.

Intuitively, firms can observe (the time between) customer purchases occurring through the purchase records they keep of their customers. When they observe long spells with no purchases, the question naturally arises of whether a customer will purchase ever again. Given the non-contractual setting of the firm, this question cannot be answered with certainty. Existing research, which we review below (see Table 1), has focused on identifying which customers are (in)active in non-contractual settings using stochastic models of the underlying customer behavior. We draw on this stream of research to develop our method for identifying potentially inactive customers, to which we add the notion of customer monitoring. As purchase processes are inherently dynamic, and as firms need to act the moment a customer becomes inactive, our approach provides a way to frequently infer a customer's activity level.

Second, given that firms wish to approach these identified customers with a reactivation mailing, we draw on studies in the direct mailing domain that consider this problem. These studies have focused on how and how often to target customers but have provided limited insights on when to target customers (see Table 2). We add to this stream of research by using customer activity as a decision rule for targeting decisions.

Third, we review the statistical quality control literature on control charts and show which aspects can help us develop a method for customer reactivation. Combining insights from these three research streams, we next show how control charts can be adapted for reactivation timing and how this adds to findings from prior studies while addressing the shortcomings of current approaches.

Studies on Determining Inactivity in Non-contractual Settings

Modeling purchase behavior in non-contractual settings is challenging, as customer attrition is unobserved, making forecasts for customer lifetime value (CLV) and its components less straightforward than in contractual settings. Building on the Pareto/NBD model (Schmittlein, Morrison & Columbo 1987), many new variants of these 'buy-till-your-die' models were developed in previous decades, such as the BG/NBD model (Fader, Hardie & Lee 2005a). Such models use recency, frequency and monetary value (RFM) variables as inputs and provide probabilistic forecasts, for example, for the number of future purchases, customer lifetime, or CLV as well as for the probability of making future purchases, or

P(Alive)[2]. While this probability is useful for customer base valuation (e.g., Fader, Hardie & Lee 2005b), the fact that it is a probability instead of a deterministic number makes it difficult to translate into a binary active/inactive decision usable in managerial decision-making. Prior work has thus suggested to 1) determine a cut-off value for *P(Alive)* to aid in the translation to a binary decision (Reinartz & Kumar, 2000; Wübben & Von Wangenheim, 2008) or 2) estimate a reactivation probability with an appropriate cut-off (Ma, Tan & Shu, 2015). However, Wübben and Von Wangenheim (2008) show that such cut-off approaches only capture *aggregate* behavior accurately, making them less suited for our intended goal of identifying *individual* inactive customers. Moreover, beyond determining *which* customers are inactive, it is also important to determine *when* they are inactive (Neslin et al., 2013). While cut-off-based methods provide this information when a customer's probability is below the cut-off value, expensive continuous model updates are required to produce up-to-date forecasts to detect the moment at which the cut-off value is reached. A different approach is thus required.

Alternatives for accurately determining customer activity have been explored, aimed at modeling the time between purchases (interpurchase time, IPT). Platzer and Reutterer (2016) introduced the Pareto/GGG model, showing that including purchase timing using IPT provides additional information to improve the performance of stochastic purchase models. Building on this insight, we develop a reactivation model with IPT as the underlying measure, where we consider IPT in calendar time. The advantage of considering IPT in calendar time is that there is no need to apply cut-off transformations to our model outcomes, as the model produces calendar time rather than probabilistic predictions (e.g., Allenby, Leone & Jen 1999). In so doing, we 1) incorporate recent advances showing the good

[2] Ascarza, Netzer, and Hardie (2018) use a Hidden Markov model to separate inactive from churned customers in a hybrid contractual/non-contractual setting. We do not consider it here, as their approach does not use RFM statistics and does not consider customer reactivation, only churn.

performance of this measure for determining customer activity but 2) also provide a direct estimate of the time at which reactivation should be initiated. Later, we discuss how we translate this continuous measure into a binary active/inactive decision to aid managerial interpretation and marketing scheduling.

Related methods using time as a measure exist, but it is unclear how these methods perform in predicting individual-level customer activity. We therefore include them in our model comparison later on. Specifically, the Pareto/GGG model (Platzer & Reutterer, 2016) could potentially be adapted for reactivation timing. Second, duration models provide a natural way of modeling data with a time-to-event nature (e.g., Seetharaman & Chintagunta, 2003).

Approaching Inactive Customers with Reactivation Mailings

The problem of which customers to target with direct or electronic mailings has been investigated extensively in the direct mailing literature (e.g., Bult & Wansbeek, 1995). Studies have shown that consideration of customers' mailing and purchase histories is important for capturing dynamics (e.g., Gönül & Shi, 1998; Gönül, Kim & Shi, 2000; Gönül & Ter Hofstede, 2006; Simester, Sun & Tsitsiklis, 2006), as they might influence the effectiveness of such actions (Van Diepen, Donkers & Franses, 2009). Thus, our approach will also leverage customers' historical data to inform mailing decisions.

However, the main aim of existing studies is to enhance overall profitability by balancing the benefits of increased purchase value with the costs of sending mailings. Most studies offer recommendations about the amount of mailings; for example, after how many mailings you should stop, and the optimal number of mailings (e.g., Neslin et al., 2013; Van Diepen, Donkers & Franses, 2009). Drèze and Bonfrer (2008) show at an aggregate level that the time between mailings affects overall customer equity. However, this does not inform managers faced with the scheduling of mailings about *when* to send these direct mails in calendar time to individual customers (Dew & Ansari, 2018; Korkmaz, Kuik & Fok. 2013). This question is increasingly relevant as firms switch from physical to digital marketing spending (CMO Survey 2018), resulting in e-mail being used in 90% of customer reactivation initiatives and direct mail in only 55% (McGee, 2016). Given the lower cost and shorter lead-time of this instrument, its deployment is more flexible than, for example, catalogs. However, such flexibility raises questions about the timing of direct mails at the individual customer level, which has not been considered previously in this setting (Ansari & Mela, 2001; Khan, Lewis & Singh 2016; Neslin et al., 2013; Zhang, Kumar & Cosguner, 2017). Our approach addresses this issue. Accordingly, we contribute to recent studies investigating the deployment of marketing actions in calendar time (Ascarza & Hardie, 2013; Dew & Ansari, 2018; Korkmaz, Kuik & Fok, 2013), specifically that of reactivation e-mails. Importantly, our approach incorporates the idea that firms should launch such interventions at the point in time when the marketing action is most likely to be effective (Ascarza, 2018).

Statistical Quality Control and Control Charts

Determining the time at which individual customers should be approached with a reactivation mailing requires knowledge about the activity status of *each* customer at the (any) point in time when the reactivation decision is made; this could be daily, weekly (as in our field test), biweekly, or monthly. Given the typically large number of customers in a customer base, this requires continuously *monitoring* a substantial amount of units (i.e., customers). This problem of monitoring many units has been investigated in the statistical quality control literature.

This stream of literature is concerned with monitoring industrial processes and acting on any disturbances to ensure their continuation (Montgomery, 2009). Statistical quality control approaches have been used to monitor inventory stocks (Ernst, Guerro & Roshwalb, 1993), wafer stepper production (Does et al., 1999), and the tin-plating of surface-mounted diodes (Wieringa, 1997). Applications in the marketing field include the selection of marketing test panel members (Marcuse, 1945), monitoring market shares and promotions (Crespy Stearns & Krebhiel, 1995), and yearly monitoring of aggregate customer satisfaction scores (Sharma, Niedrich & Dobbins, 1999) . Different from these prior studies, we monitor the purchase behavior of many individual customers near-continuously instead of monitoring aggregate-level variables at limited time points (e.g., monthly, yearly). Furthermore, we show how to generate individual-level predictions of future purchase behavior.

In this research stream, various methods have been developed to efficiently monitor many units simultaneously. One of the earliest and most prominent examples is the control chart (Shewhart, 1931). A control chart monitors the process performance through a target variable, such as the average concentration of a chemical substance. By measuring this target variable at different points in time and plotting the resulting time series, a chart is created. The chart also includes predetermined bounds within which the process is allowed to fluctuate. If the target variable crosses one of these bounds, a signal is given, and an intervention can occur to bring the process back within its bounds. A fictitious example control chart is shown in Figure 2.

Situations where one of the bounds is crossed indicate that the process is not performing according to its normal operation, and so it is said to be *out of control*. This excess variation is due to what Shewhart (1931) calls special causes (e.g., the chemical concentration deviates strongly from the mean (target) concentration).

However, some variability is allowed, as each process suffers from normal variation that is inherent to the process (common causes; Shewhart, 1931). The bounds of the control chart are determined, such as separating common from special causes of variation.

Importantly, as long as it is in control (i.e., moves within the predetermined boundaries), one should not interfere with the process. Doing so would only lead to additional variation and potentially destabilize the process ("tampering with the process"; Deming, 1982).

Using Control Charts for Customer Reactivation

Intuitively, control charts can also be used to monitor the purchasing behavior of individual customers, which provides a suitable solution for our problem of deciding which customers to approach with a reactivation action. By design, control charts intend to monitor a variable over time. Adapting this to the problem of determining whether a customer is active or not, we track the time between purchases (IPT). A time series of IPT observations is formed by tracking this variable at the individual customer level across purchase occasions.

Thus, we derive customer-specific trajectories of purchase behavior with associated bandwidths around them. Combined with the predetermined boundaries inherent to a control chart, we can convert the IPT from a calendar time measure to a binary active/inactive decision. This occurs at the point where the time since the customer's last purchase is outside the bandwidth of normal customer purchase behavior. The control chart boundary is crossed, and a reactivation intervention should occur. Earlier intervention is not advisable because a purchase could have been absent due to common cause variation in purchase behavior. In contrast, crossing one of the boundaries indicates that the process is affected by a special cause. This is a strong enough signal to warrant a reactivation action.

This separation of signals is important for effective customer management. As Ascarza, Iyengar and Schleicher (2016) show, firms should not react to all behaviors. They found that recommending mobile phone price plans to decrease customer churn could instead increase churn. The root cause was that customers who were unaware of their phone plan (not inactive in our setting) were made aware of their unfit phone plan. This is a typical example of “tampering with the process” (Deming, 1982), which may lead to destabilization of the process and negative consequences (e.g., customer churn). The control chart approach naturally deals with this situation by separating the common and special causes of variation. Firms should not interfere in a stable process and only target customers who are deviating from their normal behavior. Interventions that are initiated too soon will likely not target inactive customers, making them either ineffective or even causing adverse effects.

In sum, given the natural fit between control charts and the problem of determining when to target which customers with a reactivation intervention, we will outline how to operationalize such a control chart in the next section.

Model Development

Given our objective of developing a control chart approach for reactivation purposes, we first need to select a variable that characterizes a customer’s purchase behavior. The control chart can then monitor this variable to guide the timing of reactivation actions (Montgomery, 2009). As noted above, we use IPT, which has been used as a measure of purchase timing (Platzer & Reutterer, 2016; Zhang, Bradlow & Small, 2014). Individual-level IPT is a variable that is widely available in many non-contractual settings, making it ideally suited for a general approach to manage reactivation initiatives.

We start our model development by considering the characteristics of the statistical process underlying our variable of interest for the situation when the customer is active. Our model development is based on the following assumptions, which we adapt from Fader, Hardie and Lee (2005b):

- For a given purchase occasion, the time since the last purchase varies around the average customer-specific IPT

- The average IPT varies across customers but not over time for a specific customer (given that we model the characteristics of the customer’s behavior in a situation where she is active)

The latter assumption implies that there is a true process mean to be estimated, with noisy variation around the mean. To accurately model the time between purchases, we need to consider the properties of the process distribution (i.e., IPT in our case). In Figure 3, we summarize this distribution based on the data from our empirical applications. We observe that the distributions appear to be right skewed, and the time between purchases is always non-negative. These characteristics indicate that a normal distribution is not suitable for modeling IPT. We therefore adopt the gamma distribution to model this process, as suggested previously (e.g., Platzer & Reutterer, 2016; Zhang et al., 2007).

When considering the purchase process of a single customer, this assumption would have been sufficient. However, as we are seeking to model the purchase processes of many customers, it seems unreasonable to assume a similar distribution for each customer (see assumption 2 above). We therefore assume that heterogeneity across the population exists, and this heterogeneity can be modeled with another gamma distribution. In so doing, we follow Fader, Hardie and Lee (2005b), and adapt the gamma–gamma model of Colombo and Jiang (1999) as our model of IPT. An important advantage of this approach is that we can obtain closed-form solutions for our key expressions later on.

To develop a control chart that is usable to monitor the IPT process, we first consider the general form of such a chart. Here, we focus on the simplest form following Shewhart (1931). Such a control chart plots the individual observations along with the mean and an upper control limit (Montgomery, 2009), which we can write as

$$(1)\ CC(n_i) = \mu(n_i) + c \times \sigma(n_i)^3$$

Here, μ and σ can be replaced with suitable estimators of this quantity depending on the underlying process distribution, and n_i represents a customer-specific input (IPT in our case). The width of the control chart *c* is often taken as 3, based on the assumptions that the successive observations of the process are identically and normally distributed. This corresponds to a false alarm probability (i.e., giving a signal when no signal should be given) of 1–0.9973 = 0.0027. That is, a false alarm is highly unlikely, and the occurrence of one should be cause for concern (Shewhart, 1931). Given our earlier exposition on the non-normality of our underlying distribution, we will seek to replace each of these quantities with forms suitable to our desired application. Thus, we will derive expressions for μ and σ based on the gamma–gamma model to deal with non-normality, and we will propose a procedure to determine *c* given this non-normal data structure.

Deriving μ and σ for the Gamma–Gamma Model

First, we derive expressions characterizing the mean and standard deviation for the gamma–gamma model. We use these expressions to determine the average IPT of a customer after x_i purchases as well as the variation therein, which determines whether an out-of-control situation has occurred.

Following Fader and Hardie (2013), let x_i denote the number of transactions of a customer *i*, and let $t_{i1}, t_{i2}, \dots, t_{ix_i}$ denote the time between each transaction[4]. Define $\bar{t}_i = \sum_{j=1}^{x_i} \frac{t_{ij}}{x_i}$ as an estimate of the (unobserved) true average IPT of customer *i*, denoted as ξ_i. We are interested in two quantities related to ξ_i: its conditional mean denoted as $\mathrm{E}(T_{ij}|\bar{t}_i, x_i)$ and

[3] In general, control charts also have a lower bound, defined as $CC(n_i) = \mu(n_i) - c \times \sigma(n_i)$. However, as purchase time is naturally bounded at 0, we only consider charts with an upper control limit.
[4] We focus on repeat transactions here; hence, t_{1i} denotes the time between the first and second purchase. Customers with only one purchase are excluded from the analysis, as it is not certain that they will make another purchase.

its conditional variance denoted as $\mathrm{Var}(T_{ij}|\bar{t}_i, x_i)$. These two quantities will serve as estimates for μ and σ as desired. To arrive at these expressions, we formalize the model of Colombo and Jiang (1999) as follows:

1. We assume that $t_{ij} \sim gamma(p, \nu_i)$ with shape parameter p and rate parameter ν_i. This implies $E(T_{ij}|p, \nu_i) = \xi_i = \frac{p}{\nu_i}$, which gives $\bar{t}_i \sim gamma(px_i, \nu_i x_i)$.
2. We assume that $\nu_i \sim gamma(q, \gamma)$

Under these conditions, Fader, Hardie and Lee (2005b) show that

$$(2)\ \mathrm{E}(T_{ij}|p, q, \gamma, \bar{t}_i, x_i) = \frac{p(\gamma + x_i\bar{t}_i)}{px_i + q - 1} = \left(\frac{q-1}{px_i + q - 1}\right)\frac{p\gamma}{q-1} + \left(\frac{px_i}{px_i + q - 1}\right)\bar{t}_i,$$

which gives our desired expression for μ. Estimation of the parameters p, q, and γ will be discussed subsequently. Note that, as x_i increases, more weight is placed on the actual observed average IPT $\bar{t}_i$, and less weight is placed on the population mean. Hence, when we observe few purchases from a customer, we use the population mean as our expected IPT for customer *i*. The more purchases we observe, the more importance is given to the customer-specific average IPT $\bar{t}_i$. We next derive a closed-form expression for $\mathrm{Var}(T_{ij}|\bar{t}_i, x_i)$, the full derivation of which is provided in Appendix A, where we show that

$$(3)\ \mathrm{Var}(T_{ij}|p, q, \gamma, \bar{t}_i, x_i) = \frac{p(p + px_i + q - 1)(\gamma + x_i\bar{t}_i)^2}{(px_i + q - 2)(px_i + q - 1)^2}.$$

The square root of this expression then provides our estimate for σ. To obtain maximum likelihood estimates of the parameters p, q, and γ, we require the marginal distribution of $\bar{t}_i$. Fader, Hardie and Lee (2005b) derive that this marginal distribution is equal to

$$(4)\ f(\bar{t}_i|p, q, \gamma, x_i) = \frac{\Gamma(px_i + q)}{\Gamma(px_i)\Gamma(q)} \frac{\gamma^q \bar{t}_i^{\,px_i - 1} x_i^{px_i}}{(\gamma + \bar{t}_i x)^{px_i + q}},$$

where Γ is the gamma function. The likelihood to optimize for *N* customers given their observed number of purchases x_i and average IPT $\bar{t}_i$ is thus

$$(5)\ L(p,q,\gamma|x_i,\bar{t}_i)=\prod_{i=1}^{N} f\ (\bar{t}_i|p,q,\gamma,x_i).$$

Based on the parameter estimates obtained by maximizing Equation 5, we can compute our other quantities of interest, $\mathrm{E}\left(T_{ij}\middle|p,q,\gamma,\bar{t}_i,x_i\right)$ and $\mathrm{Var}\left(T_{ij}\middle|p,q,\gamma,\bar{t}_i,x_i\right)$.

Estimating c for the Gamma–Gamma Model

Having obtained estimates for μ and σ in Equation 1, we also need to obtain the width of the control chart boundaries *c*. To do so, we developed a simulation approach that calibrates *c* based on simulated customer purchase trajectories. We outline this approach in Appendix B.

Estimation of Control Charts & Comparison to Existing Approaches

In this section, we compare the control chart approach outlined above with a variety of competing models studied in the literature. First, however, we provide some background information on the firm that provided the focal dataset for this study and the corresponding data.

Background of the Focal Firm

We cooperated closely with an anonymous European online retailer in the greetings and gifts industry. The firm allows customers to send printed greeting cards and larger gifts to other people. We cannot reveal more about the firm for reasons of anonymity. Similar to many firms, the company is concerned with stimulating customers to repurchase after their initial purchase. The non-contractual setting the retailer is operating in makes this a challenge, as customers are free to purchase from any other retailer at any time.

Currently, the retailer tries to prevent customers from becoming inactive by sending a so-called reactivation mailing after a customer has not purchased for two months (eight weeks).[5] Such an e-mail primarily serves as a reminder to the customer and offers free postage on the next order, aiming to move the customer to purchase from the retailer again. This is a typical example of a reactivation action (e.g., Blömeke, Clement & Bijmolt, 2010; Pokornyik, 2017).

However, the retailer's current approach ignores the particularities of purchase patterns. Sending a reactivation e-mail to all customers after two months does not account for the fact that the purchase patterns of individual customers differ widely (see Figure 3). Potentially, the control chart approach outlined above could better accommodate these patterns. To test this assertion, we performed a field test with the retailer, which is described later. Before doing this, we needed to establish the ability of the control chart approach to accurately predict future activity. We therefore first used the data provided by the firm to compare the control chart approach to a broad range of competing models. The full purchase histories of 16,790 customers before the start of the field test were available to us. Limited demographic information was also available. Table 3 provides pre field-test summary statistics for these variables.

Calibrating Control Charts

To illustrate the estimation and usage of control charts, we use the IPT data described above to calibrate the control charts. We estimate control charts for 8,112 customers, whose purchase timing distribution adheres to the one described in Figure 3. First, we obtain the parameters for the gamma–gamma control chart by optimizing Equation 5. These estimates correspond to $\hat{p}$ = 15.73, $\hat{q}$ = 1.62, and $\hat{\gamma}$ = .57. Next, we determine the boundaries for the control chart using the procedure outlined before, obtaining the estimate $\hat{c}$ = 4.9. Compared to the case of the normal distribution ($\hat{c}$ = 3, Shewhart 1931), our boundaries in this case are

[5] Discussions with firm managers revealed that two months was chosen because it seemed "good" and was not based on actual purchase data. Hence, it is independent of the number and timing of purchases, and therefore does not affect the observed IPTs for individual customers.

slightly wider. This is because Shewhart (1931) assumed that the sequence of observations is independent, which is not the case in our situation, so we need wider control limits (Wieringa, 1999). These four parameters fully specify the control chart, which we can compute for every customer by plugging Equations 2 and 3 into Equation 1.

In Figure 4, we plot two control charts, one for a customer with below-average IPT (customer A) and one for a customer with above-average IPT (customer B). The former illustrates what happens for a customer who has short purchase cycles, while the latter illustrates what happens for a customer with long purchase cycles.

Both plots clearly illustrate the salient characteristics of the control chart. First, they show the time since the last purchase, which increases from week to week until a new purchase is made; it then drops to zero, leading to a saw tooth-like pattern. Second, both plots show the estimated average time between purchases (Equation 2) as a solid line, which evolves as new information becomes available. Finally, and most importantly, the upper bound for the control chart is visible. It is here that the plots show a marked difference: While the time between purchases for customer B (lower panel) never crosses the upper boundary of the control chart, the chart for customer A crosses the boundary several times.

Thus, while for customer B we would never attempt reactivation, we would do so for customer A—specifically, at the point in time at which the boundary is crossed. It is at that point that we are unsure whether the customer will purchase again[6], and reactivation might stimulate a purchase.

Intuitively, this difference between both customers makes sense. Given the long time between customer B's purchases, we could still reasonably expect a purchase to occur in the future; reactivation would likely not be effective given this customer's normal purchasing

[6] In hindsight, we know this customer repurchased. However, considering the decision of whether to approach this customer around week 75, the available information suggests that this customer was at risk of not purchasing again.

pattern. In contrast, given the short purchase cycles of customer A, a longer time between purchases (e.g., around week 75) could be indicative of the customer becoming inactive, and reactivation could be warranted.

Comparison Against Existing Approaches

How well does the control chart approach outlined above fare in separating active and inactive customers compared to existing approaches identified in the literature? To answer this question, we followed the protocol outlined in Wübben and Von Wangenheim (2008, p. 86–88) to compare the ability of various models to identify active and inactive customers. Briefly, we separate our dataset into two equal periods of 76 weeks, one calibration sample and one holdout sample. We calibrate each model on the estimation sample, and determine whether a customer is active or not at the end of this period. We use the holdout period to determine actual activity, classifying a customer as active if at least one purchase is made, and inactive otherwise. We then compare predicted with observed activity to determine the model's ability to separate active and inactive customers.

Based on the literature review, we include the following competing models next to the control chart approach outlined above:

1. Pareto/NBD (Schmittlein, Morrison & Columbo, 1987)
2. BG/NDB (Fader, Hardie & Lee, 2005a)
3. Pareto/GGG (Platzer & Reutterer, 2016)
4. Pareto/NBD with reactivation (Ma, Tan & Shu, 2015)
5. Hidden Markov Model (Schwartz, Bradlow & Fader, 2014)
6. Generalized Gamma model (Allenby, Leone & Jen, 1999)
7. Survival model, i.e., heterogeneous discrete-time proportional hazard model with expo-power baseline hazard (Seetharaman & Chintagunta, 2003)

All these models were specifically developed to model heterogeneous customer purchase behavior, accounting for heterogeneity in their own way. A key difference is the way the models distinguish active and inactive customers. Models 1–5 provide an activity probability, P(Active), as output. To classify customers as active or inactive, a cut-off c_{active} needs to be selected such that when P(Active) $\geq c_{active}$ a customer is classified as active and when P(Active) $< c_{active}$ she is classified as inactive. While it is natural to consider .5 as a cut-off (Reinartz & Kumar, 2000), Wübben and Von Wangenheim (2008) argue and show that this might not be the optimal value. We follow their approach and use a grid search to determine the optimal cut-off for these models, optimizing the overall classification performance[7].

Models 6 and 7 provide direct estimates of the time until the next purchase, thus allowing the direct classification of a customer as active or inactive if a purchase has not occurred by that time. For Model 6, beyond the mean time-to-next-purchase estimate, we also consider the upper limit of the 95% (Bayesian) prediction interval as an estimate, which in spirit is similar to the control chart upper bound. As the control chart selects its bound by design, no predictive optimization is required beyond setting *c* as outlined previously.

Beyond these model-based approaches, we also add a heuristic (business rule) approach (Wübben & Von Wangenheim, 2008). The primary reason for this is that the data-providing firm currently uses a two-month hiatus heuristic, to which we will compare our approach in the field test. Additionally, including this approach has the benefit of comparability to Wübben and Von Wangenheim (2008), who find that such heuristics perform favorably compared to model-based approaches. In line with their work, we also report results of an optimized hiatus heuristic, where we select the hiatus (i.e., time since last purchase) to optimize overall classification. We also use these two months as the reactivation period in Model 4.

Table 4 shows the classification performance when we do not optimize each model. The control chart performs favorably on most classifications, particularly the overall (mis)classification of customers. The Generalized Gamma model performs best with respect

[7] While it might seem tempting to optimize the classification of inactive customers in line with our goal, the importance lies in separating active from inactive customers, for example, to avoid negative consequences of sending reactivation mailings to customers that were not inactive (Ascarza, Iyengar & Schleicher, 2016).

to classifying active customers, although this is due to classifying very few customers as active, resulting in poor classification of inactive customers. Finally, in contrast to Wübben and Von Wangeheim (2008), we find that the hiatus heuristic does not perform that well, potentially due to the unguided selection of the threshold of two months (see footnote 5).

In Table 5 we provide the results for the models in which we optimized the holdout classification of active and inactive customers. First, we observe, consistent with Wübben and Von Wangenheim (2008), that many cut-offs differ from the naïve cut-offs in Table 4. For the BG/NBD, Pareto/NBD, and Pareto/GGG models, the cut-offs are all higher than .5 (.57, .70, and .77, respectively); for the Pareto/NBD variant of Ma et al. (2015) and HMM, the optimized cut-offs are lower (.02 and .38). The optimized hiatus length is also substantially longer (13 months).

The control chart shows the best overall predictive performance. The Hidden Markov Model is best at classifying inactive customers, while the Generalized Gamma model is best at classifying active ones. As before, both findings should be interpreted in light of worse classification performance in the other classification group. Finally, in this optimized scenario we find, similar to Wübben and Von Wangenheim (2008), that a heuristic approach is highly competitive compared to many model-based approaches.

These results are based on our focal dataset. To generalize them further, we replicated our analyses using two other datasets: the canonical CDNOW dataset (Fader & Hardie, 2001) and a grocery retail dataset (Platzer & Reutterer, 2016). The results are available on request from the authors, and confirm the ability of the control chart to separate active from inactive customers better than other approaches (except for one instance: for the optimized grocery dataset, the Generalized Gamma model performs best overall).

Overall, we can conclude that the control chart approach has better predictive ability than all of the alternative models considered. Thus, we argue that it provides a good basis to guide reactivation decisions. However, the current analysis has some shortcomings, which we address in the next section.

Field Test

Having established the favorable performance of control charts in distinguishing active from inactive customers compared to existing approaches, we next address two shortcomings of our earlier analysis. First, our assessment was based on historical data, reflecting only customer behaviour based on existing policies of these firms. In contrast, we are interested in influencing customer behaviour based on our predictions, which cannot be measured using historical data. Second, we focused on the models' ability to distinguish active and inactive customers and not on whether inactive customers also increased their activity after receiving a reactivation mailing. A field test addresses these shortcomings.

Description of the Reactivation Field Test

We performed a randomized field test comparing our proposed approach (the model group) to the current status quo at the retailer (the business rule group, sending an e-mail after two months[8]) and a control condition, in which no e-mails were sent. Comparisons of the model and business rule groups to the control condition allowed us to investigate the relative effect of reactivation targeting on customer activity. Figure 5 presents a graphical overview of the field test and its main outcomes.

The field test was run for two months from October 1 to November 30, 2016. We chose this period due to the absence of any major holidays (e.g., Valentine's Day, Christmas) or events (e.g., high school graduation ceremonies) that are strongly associated with products

[8] We did not add an optimized business rule group to our field test, as the firm was not yet convinced of its merit.

from the focal firm and could interfere with the field test. These months thus reflect normal business months for the focal firm[9].

The firm selected 16,790 customers from their customer database and assigned them randomly to the three experimental groups (see Figure 5). To be included in the field test, customers should have made at least one purchase in the previous year and at least two purchases during their tenure at the firm. The firm uses the former cut-off to determine which customers are considered potentially active customers, while the latter ensures that we only considered customers who have the intention to repurchase from the focal firm. Moreover, the final purchase for these customers should have occurred at the latest two months before the start of the field test, and the model-predicted purchase time should be within the two months of the field test (i.e., reactivation should occur during the field test period). These restrictions ensured that, when customers were assigned to either the normal or the model group, they would actually be approached during the field test. All customers included in the field test were approached only once during the test. Customers did not receive any other targeted marketing actions from the firm in the six months before or during the field test to avoid any confounding effects.

Model-Free Analysis

Our outcome variable of interest is activity: whether or not the customer made a purchase during the field test. This is in line with our and the firms' objective to reactivate customers who did not recently purchase and could thus be at risk of not purchasing again in the future. Given the randomized nature of our field test, we can directly compare the activity

[9] This assertion was validated by applying time-series outlier detection using a LOESS seasonal and trend decomposition on aggregate daily sales for the years 2012–2015. No outliers were found during October and November, but outliers were concentrated in February (Valentine's) and December (Christmas).

of customers across the three conditions to gain model-free insights on the results of the field test (see Figure 5).

We find that, on average, the model-based targeting is more effective than the business rule targeting method: 63.5% of the customers in the model group made a purchase, compared to 61.6 % of the customers in the business rule group and 60.0% of the customers in the control group. A three-way chi-square test indicates significant differences between groups ($X^2(2) = 15.77$, $p < .01$). A follow-up pairwise proportion test (corrected for multiple comparisons using Benjamini and Hochberg's (1995) method) confirms that the difference between the model and the business rule group is significant ($p = .049$), and the difference between the model and control group is significant ($p < .01$). There is no significant difference between the normal and control groups ($p = .148$). We can thus conclude that, on average, the model-based targeting approach outperforms the business rule approach while also increasing the total number of active customers, compared to not taking any action.

Aggregate Effects of Improved Reactivation Policy

While these results provide some initial evidence of the effectiveness of our reactivation approach, they do not control for potential confounding factors influencing activity, nor do they consider customer heterogeneity. Therefore, we also estimate a formal model relating activity to a set of control factors. We control for variables typically used in the CRM literature: relationship length and IPT, RFM variables[10], and gender (Blattberg, Kim & Neslin, 2008). Furthermore, we explore heterogeneity in the treatment effects by adding the interactions between the dummies indicating the groups and the control factors to our model. We estimate the following linear probability model:

---

[10] Only information on *cumulative* net sales is available. Hence, we use a categorical variable with the categories greetings-only, gifts-only, and mixed customers for monetary value to avoid collinearity problems with relationship length and frequency. The latter two categories have higher value customers, given the price difference between greetings and gifts.

(6) $\text{Activity}_\text{i} = \beta_0 + \beta^m TM_i + \beta^b TB_i + \beta^c X_i^c + \beta^{mc} TM_i X_i^c + \beta^{bc} TB_i X_i^c + \varepsilon_i,$

where TM_i and TB_i are binary variables indicating the model and business rule groups, respectively, and X_i^c is a vector containing the control variables (relationship length, IPT, RFM, and gender). The vector β contains the intercept β_0, the average treatment effects for the model and business rule group β^m and β^b, the main effects for the control variables β^c, and the interaction effects β^{mc} and β^{bc}. Finally, the error term ε_i is assumed to follow a normal distribution. Continuous variables are mean-centered to represent the effects for the 'average' customer.

Table 6 presents the results for the linear probability models. The results in the first column confirm that, compared to the control group, activity is significantly higher for the model group but not for the business rule group. Using the business rule group as the baseline reveals that the difference between the model and business rule group is significant (b = .024, p = .007). The second column shows the robustness of these effects when controlling for pre-test factors potentially influencing activity. This is reassuring and suggests that our result is not an artifact of a failure of randomization. The difference between the model and business rule group remains significant (b = .031, p = .003).

Some pre-test factors also influence activity directly. We find that customers who have been with the firm longer and those with a higher purchase frequency have a higher probability of purchasing during the period of the field test. In contrast, customers with a higher average IPT, greetings-only customers, customers who purchased longer ago, and male customers have a lower probability of purchasing during the field test. Finally, an exploratory analysis of interaction effects shows a marginally significant negative interaction between recency and the business rule group, indicating that targeting customers after two months is even less effective for customers who purchased longer ago. We do not find any other heterogeneous effects related to either of the treatments. Overall, we confirm the findings of the previous section that, on average, the model-based targeting method significantly increases the purchase probability of customers.

Incremental and Economic Value of Improved Reactivation Targeting

While we have established a statistically significant average increase in purchase probability due to the proposed targeting method, we should also consider whether this estimate has incremental and/or economic significance. From the firm perspective, the questions are 1) whether to target customers at all and 2) which targeting method to use. The answer to question one should be affirmative when the incremental benefit of targeting a customer is positive compared to not targeting the customer (e.g., Bodapati, 2008; Gönül, Kim, & Shi, 2000). Note that targeting does not necessarily lead to positive effects. As noted by Ascarza (2018), a targeted message "could make customers realize their (latent) need to churn (Berson, Smith, & Thearling, 2000) or could break the inertia that prevented them from churning (Ascarza, Iyengar, & Schleicher, 2016)". When targeting is found to be beneficial, the choice of targeting method should be based on which method generates the highest incremental revenue (Lemmens & Gupta, 2020).

One limitation of our field test is that each customer was assigned to one experimental group only. We therefore cannot determine what would have happened when the customer was not targeted or was targeted by another method. Ascarza (2018) provides a solution to this conundrum. She introduces the LIFT metric, or the incremental change in probability of targeting versus not targeting a customer. This metric can be computed using uplift random forests (Guelman, Guillén, & Pérez-Marín, 2015), a machine learning technique that provides estimates of the two probabilities of interest for *each* customer: the probability of activity in case of targeting and in case of not targeting. By subtracting these probabilities, we can derive the incremental effect (or LIFT) of targeting. In our case, for customer i, LIFT_i =

P(Active$_i$ =1|X_i, T_i=1) – P(Active$_i$=1|X_i, T_i =0), where T_i is a variable indicating whether or not a customer was targeted.[11]

We apply uplift random forests to our field test data, using the *uplift* package in the R language. Specifically, we estimate two uplift random forest models, one to identify the effects of both treatments (i.e., LIFT) and another aimed at identifying the monetary impact of both treatments (i.e., Value LIFT). Both models use activity as the dependent variable, the same independent variables as before, and one of the two targeting methods (model or business rule) as the treatment. The models are estimated on the relevant subsets of data (comparing model to control and business rule to control), after which predictions for *all* 16,790 customers were made based on the respective models. Consequently, we obtain the incremental change in probability of each customer being active when targeted with either method, even if that customer did not receive the specific treatment. Table 7 presents the results of this analysis. We find that our proposed model-based approach gives an average LIFT of .019, while it is only 0.009 with the business rule approach. This difference is significant. The cause for the difference is that, compared to the normal approach, the model-based approach generates stronger positive lifts overall (50th/75th/90th percentile), even though this effect is also attenuated by stronger negative lifts (10th/25th percentile). Thus, on average, reactivation increases the probability of being active during the field test, and this effect is significantly stronger for the model-based approach.

A supplemental analysis (available from the authors) shows that the firm can avoid targeting negative LIFT customers by more carefully selecting the targets for reactivation. Specifically, this analysis shows that greetings-only customers and customers with a higher-than-average purchase frequency are more likely to have negative LIFTs. Therefore,

[11] For further details on the method, refer to Ascarza (2018) and Guelman, Guillén, & Pérez-Marín (2015).

excluding such customers from the reactivation action might further increase the efficacy of the reactivation action.

Next, to quantify the economic value of both targeting methods, we compute a Value LIFT metric (Ascarza, 2018). Specifically, we multiply the probabilities with and without targeting for each customer with the average net spending (i.e., corrected for costs) per order prior to the treatment (i.e., cumulative net sales divided by cumulative purchase value; Table 3) to assess the incremental change in net sales due to targeting. Table 7 shows that this leads to an average incremental net sales increase of 5.29% for the model approach and a 3.82% incremental net sales increase for the business rule approach. Again, this difference is significant. Similar to the LIFT case, the overall positive Value LIFT for the model-based approach is higher (50th–90th percentile), even though the negative Value LIFT is also larger (10th percentile). Thus, targeting customers has positive economic value, and this value is significantly larger with the model-based approach.

Discussion

Extending Extant Research

While many marketing practitioners struggle with the challenge of restoring customers from an inactive to an active purchasing state (customer reactivation), the marketing literature has paid scant attention to approaches for doing so. Therefore, in this article we proposed a method to guide the timing of reactivation mailings in non-contractual settings. In this setting, customers become inactive, and firms cannot determine whether they will still purchase in the future. Customer reactivation aims to revert these inactive customers back to a state where they purchase again. The implementation of such reactivation initiatives requires firms to 1) determine which customers have become inactive, 2) when this occurs, and 3) contact these customers with a reactivation mailing at that point in time. Existing models of purchase behavior in non-contractual settings could be used for this goal (Ma, Tan, & Shu 2015; Reinartz & Kumar, 2000; Wübben & Von Wangenheim, 2008). However, their estimates of customer activity show poor performance (Wübben & Von Wangenheim, 2008). To remedy this, we developed a gamma–gamma control chart model, combining insights from customer management in non-contractual settings with those from statistical quality control theory. We demonstrated improved performance of this method compared to a wide range of competing methods on three different datasets. Moreover, we validated these findings in a field test using data from the greetings and gifts industry.

The findings of our broad model comparison show that existing stochastic models for customer base analysis (e.g., Fader, Hardie, & Lee, 2005a; Platzer and Reutterer 2016; Schmittlein, Morrison & Columbo 1987) struggle with predicting future activity at the customer level. Thus, while such models provide good performance at the customer base level—their most common use—individual-level predictions might suffer. Models incorporating a richer individual-level heterogeneity structure (e.g., Allenby, Leone & Jen 1999; Schwartz, Bradlow & Fader 2014) show better performance on this task, although they were not developed specifically for predicting future activity. Our control chart approach was developed with this goal in mind, and thus it exhibits better overall predictive performance for this specific task.

The results from our field experiment show that our approach also works well in an actual reactivation setting. We show that sending reactivation mailings can increase customer activity, but timing is crucial. We found that the current firm business rule of sending mailings after two months of inactivity did not increase activity, whereas sending the mailings following our control chart approach did increase customer activity. Drèze and Bonfrer (2008) established the importance of timing at the aggregate level; we extend this finding to the individual customer level. Firms that aim to introduce such micro-marketing strategies (Zhang & Krishnamurthy, 2004) should thus consider where in the purchase cycle of customers such actions are initiated.

Not every customer may benefit from improved targeting. Indeed, some customers might be annoyed or even become aware of the possibility to churn when receiving firms' mailings (Ascarza, 2018). However, improved timing could alleviate such negative effects by making the mailings more relevant to customers since they are closer to their purchase occasion (Chen, Narasimhan & Zhang 2001; Goldenberg, 2008). Our results support this notion, showing that the LIFT (Ascarza, 2018) generated by our control chart approach is higher, on average, compared to the firm's heuristic approach. However, the results also show substantial heterogeneity in the LIFT distribution, with stronger positive and negative targeting effects associated with the control chart approach compared to the firm's business rule. Improved selection of customers based on demographic and behavioral factors could alleviate these stronger negative effects.

While not the focus of our approach, we also show a small increase in incremental net sales when using a control chart, beyond increased activity. This contrasts the findings of Drèze and Bonfrer (2008), who reported that activity and revenue behave asymmetrically under adapted timing. Here, we show that it possible to increase both simultaneously.

Corroborating Extant Research

Some of our findings also corroborate existing research. First, we corroborate the work of Wübben and Von Wangenheim (2008) by showing that their finding that heuristic approaches (i.e., business rules) to individual-level customer base management work quite well pertains to a larger set of models and data sets than they originally included. Even though we show that the control chart approach is able to provide better individual-level predictions of future purchases, the heuristic approach is competitive in all three datasets used and often performs better than many standard models. Still, heuristic approaches could be improved by considering the underlying data-generating process. This is illustrated by the observation that the optimal heuristic approach in our greetings and gifts data would use a thirteen-month (instead of two-month) hiatus. Such “smart heuristic” approaches could reduce the negative difference in the (value)LIFT of the heuristic approach (compared to the control chart) we observed in our field test, which gives a slight upper hand to model-based approaches.

Second, we also extend the finding of Ascarza, Iyengar & Schleicher (2016) that targeting customers might increase instead of decrease churn. In a supplementary analysis available on request, we computed control chart-based predictions for customers in the business rule group in our field experiment. As these customers were targeted at a fixed time (two months), by comparing these dates to the predictions of the control chart, we can separate customers into three groups: targeted too early, targeted on time, and targeted too late. We show that targeting customers before their expected purchase data (i.e., too early) reduces their probability to purchase. Thus, not only can proactive churn prevention be less effective than believed but also it matters *when* the proactive action is taken. If the action is taken after a customer’s predicted next purchase date, there is no significant difference compared to targeting on time. This echoes the finding of Drèze and Bonfrer (2008), who showed that long intercommunication times are less problematic than short ones.

Limitations and Future Research

Within the scope of this article, there are some limitations that could not be addressed. Given limits on the runtime of the field test, customers were sent only one reactivation mailing. Assessing the performance of our approach across multiple mailings and over a longer period would provide further insights into the efficacy of our approach.

Other data sources could provide additional information on whether a customer is active or inactive. For example, customer complaints (Van Oest & Knox, 2011), web site visits, or response to e-mail (Ascarza, Netzer & Hardie, 2018) could indicate customer activity outside of purchasing. Such variables could be incorporated in stochastic models as covariates (Fader & Hardie, 2007; Van Oest & Knox, 2011) and thus could theoretically be added to our gamma–gamma control chart model as well. We did not consider this extension here to keep the exposition brief. We do note that incorporating covariates might be more easily facilitated using Bayesian approaches (e.g., Allenby, Leone, & Jen, 1990; Platzer & Reutterer, 2016) rather than the frequentist methods we considered in the current article due to the increased complexity of the gamma distribution compared to the exponential distribution considered in earlier work (Fader & Hardie, 2007; Van Oest & Knox, 2011).

Finally, we only focused on modeling purchase timing to improve customer retention. However, it is possible to model purchase value using the same framework, with the aim of identifying customer development opportunities. An additional extension could be to consider both purchase timing and purchase value simultaneously in one control chart. Such an approach is motivated by the fact that substantial correlations between purchase timing and purchase value can exist, depending on the dataset under consideration (Glady, Lemmens & Croux, 2015). One appropriate choice would be the Hotelling T^2 control chart, which can consider the mean and variation therein for two focal processes (Wierda, 1994).

References

Allenby, G. M., Leone, R. P., & Jen, L. (1999). A dynamic model of purchase timing with application to direct marketing. *Journal of the American Statistical Association*, *94*(446), 365–374.

Ascarza, E. (2018). Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research*, *55*, 80–98.

Ascarza, E., & Hardie, B. G. S. (2013). A joint model of usage and churn in contractual settings. *Marketing Science*, *32*(4), 570–590.

Ascarza, E., Iyengar, R., & Schleicher, M. (2016). The perils of proactive churn prevention using plan recommendations: Evidence from a field experiment. *Journal of Marketing Research*, *53*(1), 46–60.

Ascarza, E., Neslin, S. A., Netzer, O., Anderson, Z., Fader, P. S., Gupta, S., Hardie, B. G. S., Lemmens, A., Libai, B., Neal, D., Provost, F., & Schrift, R. (2018). In pursuit of enhanced customer retention management: Review, key issues, and future directions. *Customer Needs and Solutions*, *5*, 65–81.

Ascarza, E., Netzer, O., Hardie, B. G. S. (2018). Some customers would rather leave without saying goodbye. *Marketing Science*, *37*(1), 54–77.

Benjamini, Y., & Hochberg, Y. (1995). Controlling the false discovery rate: A practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society Series B*, *57*, 289–300.

Berson, A., Smith, S., & Thearling, K. (2000). *Building data mining applications for CRM*. McGraw-Hill.

Bodapati, A. V. (2008). Recommendation systems with purchase data. *Journal of Marketing Research*, *45*(1), 77–93.

Blattberg, R. C., Kim, B.-D., & Neslin, S. A. (2008). *Database marketing: Analyzing and managing customers*. Springer Science + Business Media.

Blömeke E., Clement, M., & Bijmolt, T. H. A. (2010). *Should they stay or should they go? Reactivation and termination of low-tier customers: Effects on satisfaction, word-of-mouth, and purchases*. SOM Research Reports (Vol. 10008). University of Groningen, SOM Research school.

Bult, J. R., & Wansbeek, T. (1995). Optimal selection for direct mail. *Marketing Science*, *14*(4), 378–394.

CMO Survey (2018). https://cmosurvey.org/results/february-2018/

Colombo, R., & Jiang, W. (1999). A stochastic RFM model. *Journal of Interactive Marketing*, *13*, 2–12.

Crespy, C. T., Stearns, J. M., & Krehbiel, T. C. (1995). Integrating analytic methods into marketing research education: Statistical control charts as an example. *Marketing Education Review*, *5*(1), 11–23.

Deming, W. E. (1982). *Out of the crisis*. Cambridge University Press.

Dew, R., & Ansari, A. (2018). Bayesian nonparametric customer base analysis with model-based visualizations. *Marketing Science*, *37*(2), 216–235.

Does, R. J. M. M., Roes, K. C. B., & Trip, A. (1999). *Statistical process control in industry: Implementing and assuring SPC*. Kluwer Academic.

Drèze, X., & Bonfrer, A. (2008). An empirical investigation of the impact of communication timing on customer equity. *Journal of Interactive Marketing*, *22*(1), 36–50.

Ernst, R., Guerrero, J.-L., & Roshwalb, A. (1993). A quality control approach for monitoring inventory stock levels. *Journal of the Operational Research Society*, *44*(11), 1115–1127.

Fader, P. S., & Hardie, B. G. S. (2001). Forecasting repeat sales at CDNOW: A case study. *Interfaces*, *31*, S94–S107.

Fader, P. S., & Hardie, B. G. S. (2007). *Incorporating time-invariant covariates into the Pareto/NBD and BG/NBD models*. http://brucehardie.com/notes/019

Fader, P. S., & Hardie, B. G. S. (2013). *The Gamma–Gamma model of monetary value*. http://brucehardie.com/notes/025

Fader, P. S., Hardie, B. G. S., & Lee, K. L. (2005a). Counting your customers the easy way: An alternative to the Pareto/NDB model. *Marketing Science*, *24*(2), 275–284.

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

Glady, N., Lemmens, A., & Croux, C. (2015). Unveiling the relationship between the transaction timing, spending and dropout behaviors of customers. *International Journal of Research in Marketing*, *32*, 78–93.

Goldenberg, B. J. (2008). *CRM in real-time: Empowering customer relationships*. CyberAge Books.

Gönül, F., & Shi, M. Z. (1998). Optimal mailing of catalogs: A new methodology using estimable structural dynamic programming models. *Management Science*, *44*(9), 1249–1262.

Gönül, F., Kim, B.-D., & Shi, M. Z. (2000). Mailing smarter to catalog customers. *Journal of Interactive Marketing*, *14*(2), 2–16.

Gönül, F., & ter Hofsede, F. (2006). How to compute optimal catalog mailing decisions. *Marketing Science*, *25*(1), 65–74.

Guelman, L., Guillén, M., & Pérez-Marín, A. M. (2015). Uplift random forests. *Cybernetics and Systems*, *46*(3/4), 230–248.

Romana, K., Lewis, M., & Singh, V. (2009). Dynamic customer management and the value of one-to-one marketing. *Marketing Science*, *28*(6), 1063–1079

Korkmaz, E., Kuik, R., & Fok, D. (2013). *Counting your customers: When will they buy next? An empirical validation of probabilistic customer base analysis models based on purchase timing*. ERIM Report Series ERS–2013–001–LIS. https://ssrn.com/abstract=2198260

Lemmens, A., & Gupta, S. (2020). Managing churn to maximize profits. *Marketing Science*, in press

Ma, S., Tan, H., & Shu, F. (2015). When is the best time to reactivate your inactive customers? *Marketing Letters*, *26*, 81–98.

Marcuse, S. (1945). An application of the control chart method to the testing and marketing of foods. *Journal of the American Statistical Association*, *40*(230), 214–222.

McGee, T. (2016). *2016 customer acquisition, retention and the best ROI*. https://www.targetmarketingmag.com/article/2016-customer-acquisition-retention-and-the-best-roi/

Montgomery, D. C. (2009). *Introduction to statistical quality control*. John Wiley & Sons Inc.

Neslin, S. A., Taylor, G. A., Grantham, K. D., & McNeil, K. R. (2013). Overcoming the "recency trap" in customer relationship management. *Journal of the Academy of Marketing Science*, *41*, 320–337.

Platzer, M., & Reutterer, T. (2016). Ticking away the moments: Timing regularity helps to better predict customer activity. *Marketing Science*, *35*(5), 779–799.

Pokornyik, R. (2017). *Customer reactivation marketing: Why and how to do it.* https://www.omnisend.com/blog/customer-reactivation-marketing-why-and-how-to-do-it/

Reinartz, W., & Kumar, V. (2000). On the profitability of long-life customers in a noncontractual setting: An empirical investigation and implications for marketing. *Journal of Marketing*, *64*(4), 17–35.

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

Schwartz, E. M., Bradlow, E. T., & Fader, P. S. (2014). Model selection using database characteristics: Developing a classification tree for longitudinal incidence data. *Marketing Science, 33*(2), 188–205.

Seetharaman, P. B., & Chintagunta, P. (2003). The proportional hazard model for purchase timing. *Journal of Business and Economics Statistics*, *21*(3), 368–382.

Sharma, S., Niedrich, R. W., & Dobbins, G. (1999). A framework for monitoring customer satisfaction: An empirical illustration. *Industrial Marketing Management*, *28*(3), 231–243.

Shewhart, W. A. (1931). *Economic control of quality of manufactured product*. Van Nostrand Company.

Simester, D., Sun, P., & Tsitsiklis, J. N. (2006). Dynamic catalog mailing policies. *Management Science*, *52*(5), 683–696.

Stevens, R. (2017). *Ten tips for customer reactivation*. Retrieved September 25, 2018, from https://biznology.com/2017/06/ten-tips-customer-reactivation/

Van Diepen, M., Donkers, B., & Franses, P. H. (2009). Dynamic and competitive effects of direct mailings: A charitable giving application. *Journal of Marketing Research*, *46*, 120–133.

Van Oest, R., & Knox, G. (2011). Extending the BG/NBD: A simple model of purchases and complaints. *International Journal of Research in Marketing*, *28*, 30–37.

Wierda, S. J. (1994). *Multivariate statistical process control* (Unpublished dissertation). University of Groningen.

Wieringa, J. E. (1997). *The case of tin-plating of surface mounted glass diodes*. SOM Research Report 97A18. University of Groningen.

Wieringa, J. E. (1999). *Statistical process control for serially correlated data* (Unpublished dissertation). SOM Research School, University of Groningen.

Wübben, M., & Von Wangenheim, F. (2008). Instant customer base analysis: managerial heuristics often “get it right”. *Journal of Marketing, 72*(3), 82–93.

Zhang, Y., Bradlow, E. T., & Small, D. S. (2014). Predicting customer value using clumpiness: From RFM to RFMC. *Marketing Science, 34*(2), 195–208.

Zhang, J., & Krishnamurthi, L. (2004). Customizing promotions in online stores. *Marketing Science*, *47*, 561–578.

Zhang, X. , Kumar, V., & Cosguner, K. (2017). Dynamically managing a profitable email marketing program. *Journal of Marketing Research*, *54*(6), 851–866.

Zhang, C. W., Xie, M., Liu, J. Y., & Goh, T. N. (2007). A control chart for the Gamma distribution as a model of time between events. *International Journal of Production Research*, *45*(23), 5649–5666.

Appendix A: Derivation of the variance component

To develop our control charts, we require an expression for the variance component of the chart. In the model development section, we provided an expression for the mean component of the control chart (Equation 2), given in Fader, Hardie, and Lee (2005b) and Fader and Hardie (2013). In Equation 3 we stated that

$$(A1)\ \mathrm{Var}\left(\mathrm{T}_{\mathrm{ij}} \middle| p, q, \gamma, \bar{\mathrm{t}}_i, \mathrm{x}_\mathrm{i}\right) = \frac{p(p + px_i + q - 1)(\gamma + x_i \bar{\mathrm{t}}_i)^2}{(px_i + q - 2)(px_i + q - 1)^2}$$

To arrive at this expression, we start from first principles and note that

$$\text{(A2) } \mathrm{Var}\left(\mathrm{T}_{\mathrm{ij}} \middle| p, q, \gamma, \bar{\mathrm{t}}_i, \mathrm{x}_\mathrm{i}\right) = E\left(T_{ij}^2 \middle| p, q, \gamma, \bar{\mathrm{t}}_i, \mathrm{x}_\mathrm{i}\right) - E\left(T_{ij} \middle| p, q, \gamma, \bar{\mathrm{t}}_i, \mathrm{x}_\mathrm{i}\right)^2$$

The latter part of this expression can be derived by taking the square of the expression in Equation 2, yielding

$$(A3)\ E\left(T_{ij} \middle| p, q, \gamma, \bar{\mathrm{t}}_i, \mathrm{x}_\mathrm{i}\right)^2 = \frac{p^2(\gamma + x_i\, \bar{\mathrm{t}}_i)^2}{(px_i + q - 1)^2}$$

What remains is a derivation of the first part of Equation A2. Under the assumption that T_{ij} follows a gamma distribution,

$$E\left(T_{ij}^2\right) = \int_0^\infty x_i^2 \frac{\nu^p}{\Gamma(p)} x^{p-1} e^{-\nu x_i} dx_i = \frac{(p + 1\,)p}{\nu^2}$$

Given the stochastic nature of ν, we can use the change of variables technique (Fader & Hardie, 2013) to derive the distribution $f(\tau | p, q, \gamma, \bar{\mathrm{t}}_i, x_i)$ if we let $\tau = h(\nu)$. Then, we have

$$f_{\tau(\tau) = \left|\frac{d}{d\tau} h^{-1}(\tau)\right| f_\nu(h^{-1}(\tau))}$$

which can be evaluated by realizing that, when $\tau = \frac{(p+1\,)p}{\nu^2}$, it implies that $\nu = \frac{\sqrt{(p+1\,)p}}{\sqrt{\tau}}$ and $\frac{d\nu}{d\tau} = -\frac{\sqrt{(p+1\,)p}}{2\tau^{3/2}}$. Then, applying the change of variables we get

$$f(\tau|p,q,\gamma,\bar{\mathrm{t}}_i,x_i) = -\frac{\sqrt{(p+1)p}}{2\tau^{\frac{3}{2}}}\frac{(\gamma + x_i\,\bar{\mathrm{t}}_i)^{px_i+q}\left(\frac{\sqrt{(p+1)p}}{\sqrt{\tau}}\right)^{px_i+q-1} e^{\frac{-(\gamma+x_i\,\bar{\mathrm{t}}_i)\sqrt{(p+1)p}}{\sqrt{\tau}}}}{\Gamma(px_i+q)}$$

$$= -\frac{[\sqrt{(p+1)p}(\gamma + x_i\,\bar{\mathrm{t}}_i)]^{px_i+q}\left(\tau^{-\frac{1}{2}}\right)^{px_i+q-1} e^{\frac{-(\gamma+x_i\,\bar{\mathrm{t}}_i)\sqrt{(p+1)p}}{\sqrt{\tau}}}}{2\tau^{\frac{3}{2}}\Gamma(px_i+q)}$$

The expectation of this distribution can be computed as follows:

$$E\big(f(\tau)\big) = \int_0^\infty -\tau\frac{[\sqrt{(p+1)p}(\gamma + x_i\,\bar{\mathrm{t}}_i)]^{px_i+q}\left(\tau^{-\frac{1}{2}}\right)^{px_i+q-1} e^{\frac{-(\gamma+x_i\,\bar{\mathrm{t}}_i)\sqrt{(p+1)p}}{\sqrt{\tau}}}}{2\tau^{\frac{3}{2}}\Gamma(px_i+q)}d\tau$$

$$= -\frac{\left[\sqrt{(p+1)p}(\gamma + x_i\bar{\mathrm{t}}_i)\right]^{px_i+q}}{\Gamma(px_i+q)}\int_0^\infty \frac{(\tau^{-\frac{1}{2}})^{px_i+q-1}e^{\frac{-\sqrt{(p+1)p}(\gamma+x_i\bar{\mathrm{t}}_i)}{\sqrt{\tau}}}}{2\sqrt{\tau}}d\tau$$

If we define $u = \frac{1}{\sqrt{\tau}}$ (so that $\tau = \frac{1}{u^2}$ and $\frac{d\tau}{du} = -\frac{2}{u^3}$), we can rewrite the expression as

$$(A4)\ E\big(f(\tau)\big) = -\frac{\left[\sqrt{(p+1)p}(\gamma + x_i\bar{\mathrm{t}}_i)\right]^{px_i+q}}{\Gamma(px_i+q)}\int_0^\infty u\frac{(u^{px_i+q-1}e^{-\sqrt{(p+1)p}(\gamma+x_i\bar{\mathrm{t}}_i)\,u})}{2}\frac{-2}{u^3}du$$

$$= \frac{p(1+p)(\gamma + x_i\bar{\mathrm{t}}_i)^2}{(px_i+q-2)(px_i+q-1)}$$

Substituting A3 and A4 into Equation A2, we find that

$$\mathrm{Var}\big(\mathrm{T_{ij}}\big|p,q,\gamma,\bar{\mathrm{t}}_i,\mathrm{x_i}\big) = \frac{p(1+p)(\gamma + x_i\bar{\mathrm{t}}_i)^2}{(px_i+q-2)(px_i+q-1)} - \frac{p^2(\gamma + x_i\,\bar{\mathrm{t}}_i)^2}{(px_i+q-1)^2}$$

$$= \frac{p(p+px_i+q-1)(\gamma + x_i\bar{\mathrm{t}}_i)^2}{(px_i+q-2)(px_i+q-1)^2},$$

which is the expression in Equation 3.

Appendix B: Estimating c for the Gamma–Gamma Model

Having obtained estimates for μ and σ in Equation 1, we also need to obtain the width of the control chart boundaries c. In general, this width can be set to a fixed number, such as 3 (e.g., Shewhart, 1931), but it can also be selected based upon a sample of observations known to be in-control in combination with a suitable false alarm rate (e.g., P(false alarm) = 0.0027; Montgomery, 2009). As the number 3 is based on normal distribution theory and we are dealing with non-normal data in our setting, we will use an in-control sample to determine the value of c. A new challenge arises, as the notion of in-control is difficult to define in our case; based on historical purchase data, we do not know which customers (or observations from these customers) are behaving as 'normal' and which are not. Furthermore, customers vary purchase frequency, and we cannot reliably determine a false alarm rate due to short time horizons for many customers. Using only customers with long time horizons would bias our estimates, as we would ignore a substantial portion of customers.

We develop a simulation approach whereby we generate our own in-control observations based on the distribution of the observed purchase behavior of customers in the data and determine the value for c based on this simulated sample. In particular, for a given set of simulated customers we need to generate a series of purchase occasions, which have IPTs that align with the time between purchases observed in the data. To achieve this, we first generate the number of purchases according to the negative binomial model (NBD), in the tradition of the Pareto/NBD and BG/NBD models (Fader, Hardie & Lee, 2005a; Schmittlein & Morrison, 1987).

We obtain the parameters of the NBD model by fitting it to the customer purchase data. Subsequently, we use this estimated model to simulate new, long purchase trajectories for 5,000 simulated customers with IPTs that match the behavior of customers in the data. We calibrate control charts on these simulated purchase trajectories and select c using a grid search such that most of these observations are within the bounds of the control chart (i.e., are in-control). We follow existing theory by requiring that c is chosen such that 99.73% of the observations are within the bounds of the control chart (Shewhart, 1931).

Figure 1 Classification performance of hiatus heuristic business rule according to customer purchase frequency (high versus low, median split) across three datasets. In the greetings & gifts dataset, all high customers are naively classified as active, while in the CDNOW dataset all low customers are naively classified as inactive.

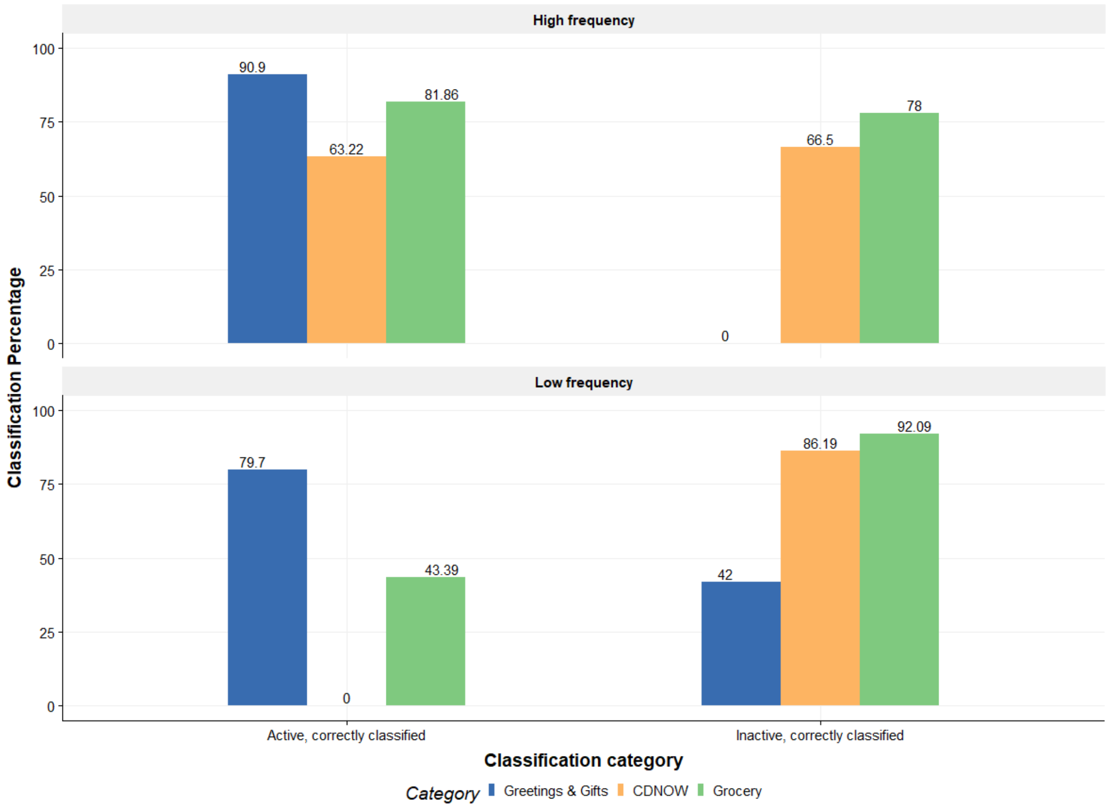

Figure 2 A fictitious control chart with upper and lower control limits

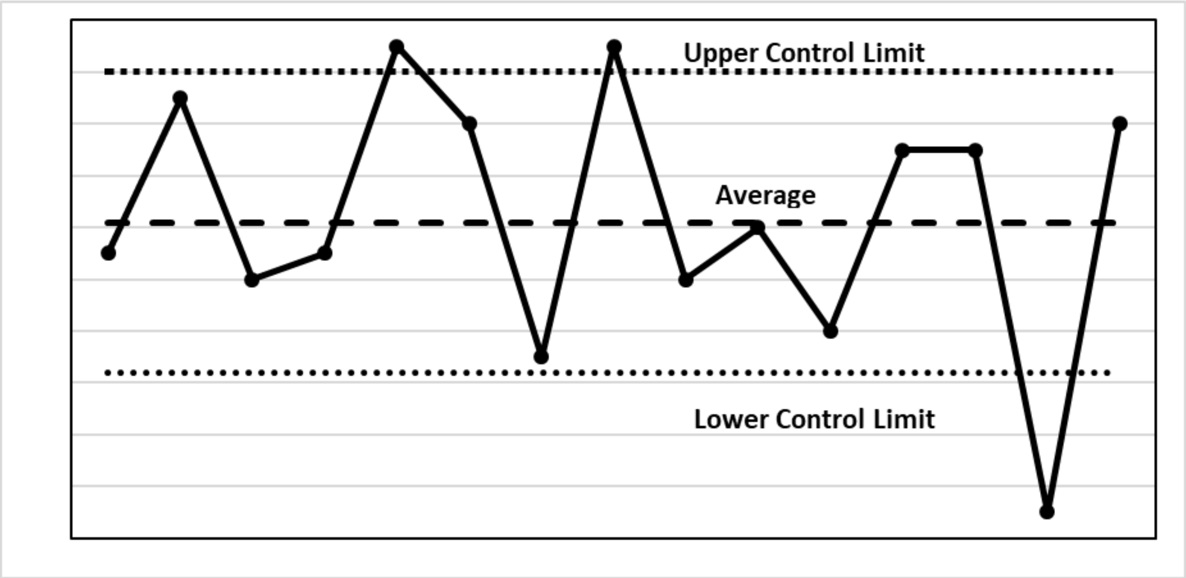

Figure 3 Distribution of interpurchase time of repeat purchases per customer. The values for greetings and gifts have been indexed for confidentiality reasons.

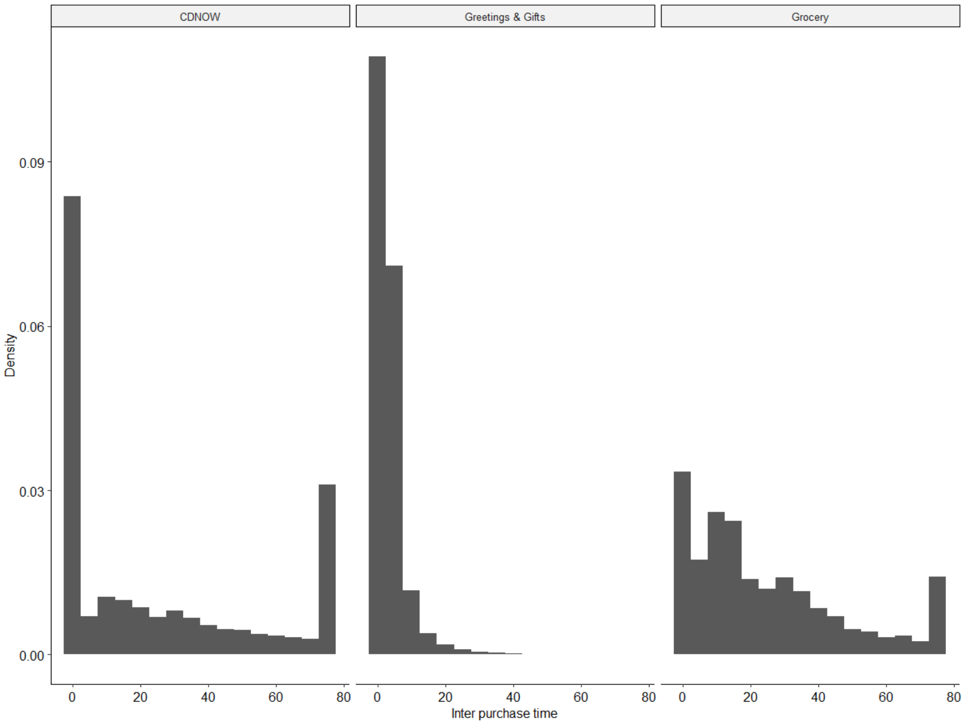

Figure 4 Control charts for two randomly selected customers A and B

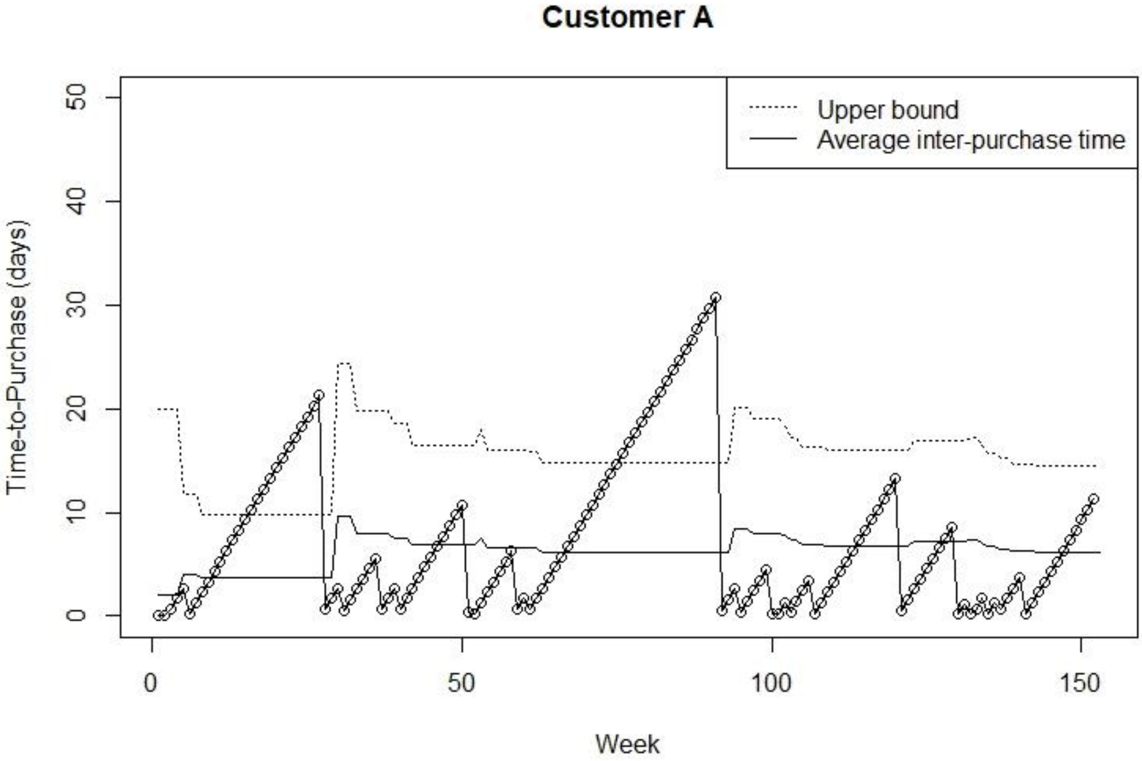

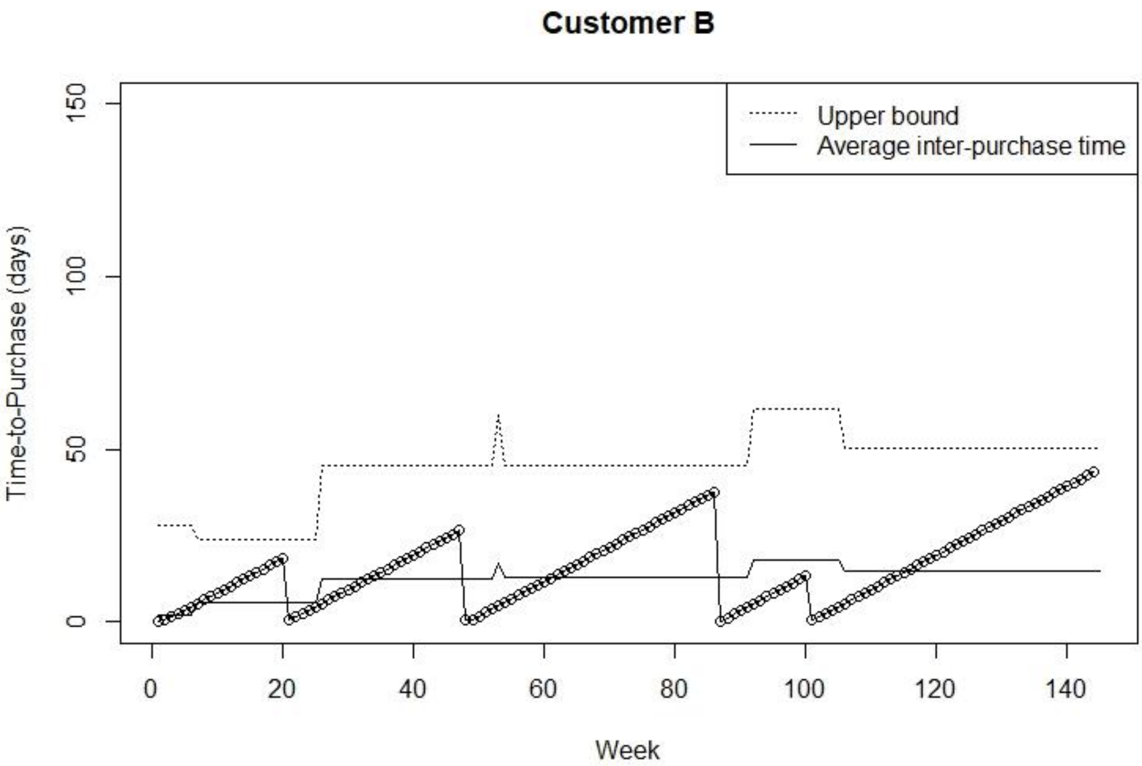

Figure 5 Overview of the groups in the field test

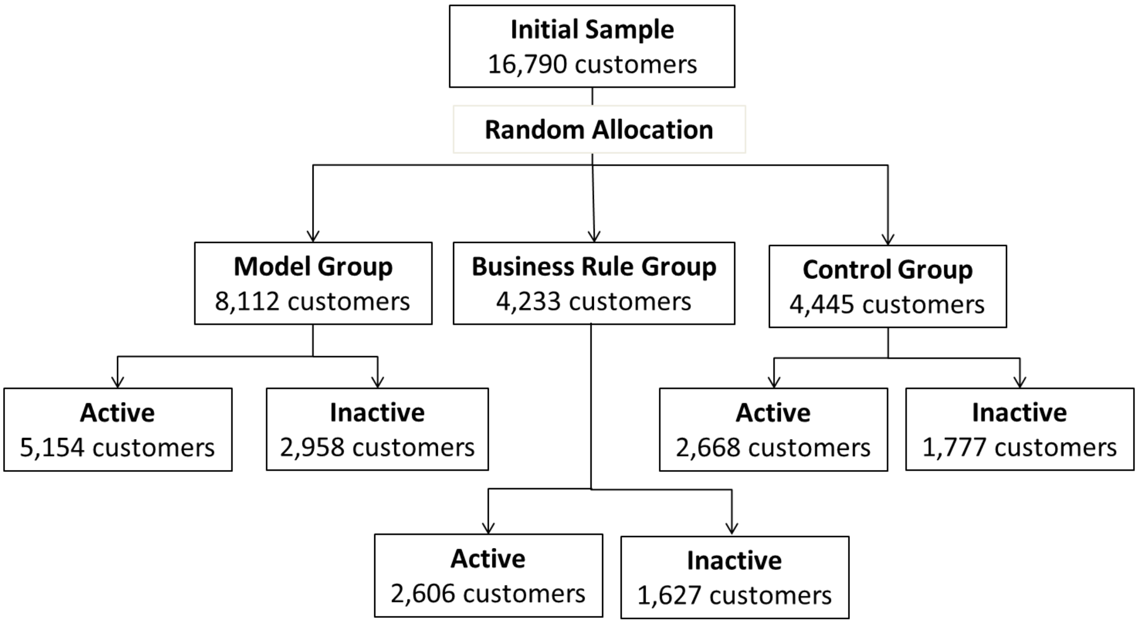

Table 1 Overview of studies related to determining inactivity in non-contractual settings

	Focus on reactiva-tion	**Binary activity decision**	**Calendar time decisions**	**Online communi-cation**	**Valida-tion in the field**	**Accounts for customer heterogeneity**	**Accounts for past behavior/ dynamics**
Allenby, Leone, and Jen (1999)			✔			✔	✔
Ascarza and Hardie (2013)			✔			✔	✔
Blömeke et al. (2010)	✔				✔	✔	
Fader, Hardie, and Lee (2005a)						✔	✔
Korkmaz, Kuik, and Fok (2013)			✔			✔	✔
Neslin et al. (2013)				✔			✔
Ma et al. (2015)	✔					✔	✔
Platzer and Reutterer (2016)			✔			✔	✔
Reinartza and Kumar (2000)		✔	✔				✔
Schmittlein et al. (1987)						✔	✔
Seetharaman and Chintagunta (2003)			✔			✔	✔
Wübben and Von Wangenheim (2008)		✔				✔	✔
This study	✔	✔	✔	✔	✔	✔	✔

Table 2 Overview of studies related to customer management using direct mailing

	Focus on reactiva-tion	Binary activity decision	Calendar time decisions	Online communi-cation	Valida-tion in the field	Accounts for customer heterogeneity	Accounts for past behavior/ dynamics
Ansari and Mela (2001)				✔		✔	
Bult and Wansbeek (1995)							✔
Dew and Ansari (2018)			✔			✔	✔
Drèze and Bonfrer (2008)			✔	✔		✔	✔
Gönül et al.(2000)			✔			✔	✔
Gönül, Shi (1998)						✔	✔
Gönül and Ter Hofstede (2006)			✔			✔	✔
Khan et al. (2009)				✔		✔	✔
Simester et al. (2006)					✔	✔	✔
Van Diepen et al. (2009)					✔	✔	✔
Van Diepen et al. (2011)						✔	✔
Zhang et al. (2017)				✔		✔	✔
This study	✔	✔	✔	✔	✔	✔	✔

Table 3 Pre-test sample characteristics

	Mean	Standard deviation	Minimum	Maximum
Interpurchase time (weeks)[a]	3.57	4.73	0	100
Cumulative purchase volume[a]	3.74	3.29	.13	100
Cumulative net sales[a,b]	9.73	6.43	0	100
Relationship length (years)	4.57	2.03	.07	7.16
Recency (weeks)	6.51	1.18	.29	8
Frequency	42.74	29.10	4	423
Gender (male)	.09	.83	0	1
Greetings-only customer	.47	.50	0	1
Gifts-only customer	.001	.008	0	1

[a] This number has been transformed into an index for confidentiality reasons. The index was set to 0 at the 0-point and to 100 at the maximum value.
[b]Net sales are corrected for costs incurred by the firm (e.g., discounts)

Table 4 Classification of active and inactive customers for various models using standard thresholds

	Inactive, correctly classified (%)	Active, correctly classified (%)	Overall, correctly classified (%)	Inactive but classified active (%)	Active but classified inactive (%)	Overall, incorrectly classified (%)
Two-month hiatus	50.82	81.11	66.84	49.18	18.89	33.16
BG/NBD (.5)	71.12	78.37	76.57	28.88	21.63	23.43
Pareto/NBD (.5)	68.74	73.91	72.96	31.26	26.09	27.04
Pareto/GGG (.5)	63.05	71.24	70.00	36.94	28.76	30.00
Pareto/NBD Ma et al. (.5)	54.62	89.34	70.96	45.38	10.66	29.04
Generalized Gamma (mean)	39.55	**100**	48.14	60.45	**0**	51.86
Survival Model	34.35	66.99	44.55	65.65	33.01	55.45
HMM (.5)	33.27	65.52	50.92	66.73	34.48	49.08
Control Chart	**78.60**	90.50	**86.30**	**21.39**	9.50	**13.70**

Notes: Best classification performance per column in boldface.

Table 5 Classification of active and inactive customers for various models using optimized thresholds

	Inactive, correctly classified (%)	Active, correctly classified (%)	Overall, correctly classified (%)	Inactive but classified active (%)	Active but classified inactive (%)	Overall, incorrectly classified (%)
13-month hiatus	78.78	77.07	77.41	21.22	22.93	22.59
BG/NBD (.57)	70.92	82.20	78.77	29.08	17.80	21.23
Pareto/NBD (.70)	68.90	83.38	78.59	31.10	16.62	21.41
Pareto/GGG (.77)	62.40	84.18	75.71	37.60	15.82	24.29
Pareto/NBD; Ma et al. (.02)	72.45	82.27	79.36	27.55	17.73	20.64
Generalized Gamma (upper)	65.64	**99.83**	82.20	34.36	**.17**	17.80
Survival Model	34.35	66.99	44.55	65.65	33.01	55.45
HMM (.38)	**100**	66.08	66.09	**0**	33.92	32.91
Control Chart	78.60	90.50	**86.30**	21.39	9.50	**13.70**

Notes: Best classification performance per column in boldface.

Table 6 Linear probability model results for drivers of customer activity

	(1) Activity (main effect)	(2) Activity (controls)	(3) Activity (heterogeneity)
Intercept	.600 (.007) ***	.613 (.007) ***	.613 (.008) ***
Model group	.035 (.009) ***	.042 (.010) ***	.042 (.010) ***
Business rule group	.016 (.010)	.011 (.010)	.011 (.10)
Gender (male)		-.023 (.013)	-.024(.013) *
Relationship length[a]		.005 (.002) *	.005 (.003) *
Interpurchase time[a]		-.003 (.002) *	-.003 (.001) *
Greetings-only customer		-.029 (.007) ***	-.030 (.007) ***
Gifts-only customer		-.607 (.482) *	-.608 (.482)
Recency[a]		-.004 (.0004) ***	-.004 (.0001) ***
Frequency[a]		.001 (.0002) ***	.001 (.0001) ***
Recency: Model group			.001 (.001)
Recency: Business rule group			-.002 (.001) *
N	16,790	16,790	16,790
Adjusted R^2	.0001	.016	.015

Notes: [a] All continuous variables are mean-centered, so the main effects refer to the average customer. Standard errors in parentheses. *** p < 0.01; ** p < 0.05; * p < 0.10

Table 7 Incremental and economic value of different targeting methods

	LIFT : Model	LIFT: Business rule	Mean difference[b]	p-value	Value LIFT: Model	Value LIFT: Business rule	Mean difference[b]	p-value
Mean	.019	.009	.010	.000	.053	.038	.015	.000
10th percentile[a]	-.101	-.056	-.056	.000	-.150	-.109	-.100	.000
25th percentile	-.045	-.059	-0.44	.000	-.070	-.071	-.077	.000
50th percentile	.018	.003	-.006	.000	.029	.013	-.019	.000
75th percentile	.083	.026	.023	.000	.151	.054	.032	.000
90th percentile	.139	.104	.075	.000	.278	.254	.146	.000

Notes: [a] Percentiles are computed using the model-based lift/value, after which the corresponding percentile average for the *same* customers in the normal group is computed and reported.
[b] We report the mean difference given by a paired samples t-test, comparing the same customer under model and normal conditions. This is generally not equal to the difference of the means of the preceding two columns.