Submitted to *Information Systems Research*
manuscript (Version: August 18, 2018)

The Voice of the Customer: Managing Customer Care in Twitter

We investigate digital customer care in the Telecommunications industry. The primary goal is to determine an optimal strategy to manage customer sentiment on social media sites such as Twitter. We also aim to identify factors and external events that can influence the effectiveness of customer care. Recently, managing customer sentiment – particularly on social media – has become crucial as more customers have started to use social media to seek help from firms.

Our study uses data consisting of sentiments expressed by customers directed at Twitter's service accounts of four major U.S. telecommunication- service providers: AT&T, Verizon, Sprint and T-Mobile. To understand the antecedents of digital customer care, we model a diffusion process of customer sentiment over time. This diffusion process is influenced (or *controlled*) by the firm through the strategy employed to respond to customer tweets. The main benefit of our methodology is that it can be used in a prescriptive sense (specifically, to improve digital customer care), rather than merely for prediction (e.g., to forecast customer sentiment).

The parameters of the controlled diffusion process are estimated to shed several insights into digital customer care in this industry. First, we find a clear separation among the firms in terms of digital customer care effectiveness. Second, we find that good customer care is not merely a matter of responding to customer tweets: T-Mobile and Sprint have high response rates, but are low on effectiveness. Third, the quality of digital customer care that customers expect varies across firms: Customers of higher priced firms (e.g., Verizon and AT&T) expect better customer care. Fourth, seemingly unrelated events (such as signing an exclusive contract with a celebrity) can impact digital customer care. These events can be firm-initiated or exogenous. Our study has important implications for managers as it can help firms determine the optimal strategy to influence customer sentiment. The study also helps firms anticipate the impact of external events on digital customer care and adjust the response strategy to accommodate these events.

Key words: Digital Customer Care; Customer Sentiment; Stochastic Differential Equations

1. Introduction

Given the proliferation of Web 2.0 technology, customers are able to communicate with firms through new channels of communication such as social media websites. According to J.D.Power (2013), almost two-thirds of customers have used a company's social media site to receive service. Such digitally provisioned service is attractive to customers due to its fast response, and is being increasingly preferred to more traditional service channels such as phone or email (Frumkin 2017).

Firms too have an incentive to embrace service provision through social media. According to Hyken (2016), the average cost of a service-related response in Twitter is \$1, while the average cost of interacting with a customer through a traditional call center can be close to \$6. Firms that use Twitter as a social care channel are seeing a large 19% increase in customer satisfaction. Given that digital customer care has mutual benefits for firms and customers, it is not surprising to observe a 250% increase in customer service interactions over Twitter during the past few years (Frumkin 2017).

Although there are significant benefits of digital customer care management, it imposes certain challenges. The communication between firms and customers is public. Therefore, a firm's response to a customer's query could impact not only the focal customer but also other existing and potential customers. This, of course, can also be regarded as a key benefit of digital customer care, namely, that a response to one customer could potentially benefit other customers. However, it is important for firms to devise a response strategy that adequately addresses customers' queries in social media. Gunarathne et al. (2017) show some evidence that firms strategically respond to customers' queries. In particular, their study suggests that firms tend to provide a more satisfactory response to customers who are more influential in social networking websites.

Numerous examples in Twitter provide evidence that a failure to promptly respond to customers on social media may prove disastrous. For instance when British Airways had an eight-hour delay in responding to a dissatisfied customer, it made the customer's tweet go viral, rapidly disseminating negative sentiment about the airline. The ineffective response not only aggravated the focal customer, but also made the treatment observable to other social media users (social media bystanders). In another example, American Airline responded with an automated *"Thank You for Your Support"* tweet to a Twitter user's negative tweet, giving the impression that the airline's responses are perfunctory and automatic, rather than carefully constructed to resolve customers' problems (Burke 2016).

These examples, along with many others, indicate that customers will continue to use social media channels to communicate with firms, and it is up to firms to respond effectively. According to J.D.Power (2013), "consumer expectations for social interactions vary across industries, although quality content and responsive service representatives are keys to higher satisfaction levels." Therefore, we explore the question: *What is an optimal response strategy for digital customer care management?*

To answer this question, we develop a controlled diffusion model of customer sentiment. This model is a Stochastic Differential Equation (SDE) that describes the dynamics of customer sentiment driven by the response strategy used by the firm. The control objectives of the firm are explicitly modeled to capture how a firm should optimally react to customer sentiment. We recover the parameters of this controlled diffusion process using Maximum Likelihood Estimation (MLE) and data on social media customer care. This data was collected from Twitter's service accounts of the *Big Four* telecommunications firms in the U.S. (AT&T, Verizon, Sprint, and T-Mobile). Sentiment analysis was used to measure customer sentiment over time for each firm during a four-month period. To lend credibility to our structural model, we compared the predictive performance of the diffusion model with state-of-the-art forecasting methods and found that our proposed model outperforms these models.

The main advantage of our approach is not just in its predictive ability, but that it can prescribe an optimal response strategy for effective digital customer care management. There are several useful insights gleaned from this study. First, we find that there is a clear separation among the firms in terms of digital care effectiveness. The top two firms in the industry (AT&T and Verizon) do better – in terms of the effectiveness of care support – than Sprint and T-Mobile. Second, for all firms, we find that good digital care consists not merely of responding to tweets, but an effort-intensive activity where customer tweets need to be carefully examined and adequately addressed. Third, customers expect better quality of care from firms that charge more for similar cellular plans. Fourth, because of its structural nature, our methodology allows us to estimate the impact of seemingly unrelated events that could significantly affect customer care (e.g. signing an exclusive contract with a celebrity can improve how customers perceive the quality of care). Our findings in that regard are that events external to the care platform (whether firm initiated or exogenous) can profoundly influence different aspects of digital customer care. Finally, we develop a closed-form approximation of the optimal response strategy to effectively manage customer sentiment and recover useful structural properties of such a strategy.

The rest of the paper is organized as follows. We first provide a review of related work, followed by a discussion of a controlled diffusion model of how customer sentiment evolves over time. We next estimate the parameters of this controlled diffusion process using Twitter data from the service handles of four major telecommunications service providers in the U.S. (AT&T, Verizon, T-Mobile, and Sprint). The estimated model is validated by comparing its out-of-sample predictive performance with that of state-of-the-art prediction models. We next provide several insights into the nature of digital customer care in the telecommunications industry using a number of policy simulations and real events. Then we further explore the structure of the optimal digital customer care policy. We conclude the paper with a discussion and summary of the findings.

2. Literature Review

In this section, we review the literature on firm-level impacts of customer sentiment as well as the main methodology employed in this study.

2.1. Customer Sentiment

A basic premise of this study is that customer sentiment can be measured and is important to manage for a firm. Customer sentiment can be gauged using sentiment analysis which is defined as "the field of study that analyzes people's opinions, sentiments, evaluations, appraisals, attitudes, and emotions towards entities such as products, services, organizations, individuals, issues, events, topics, and their attributes." (Liu 2015). A variety of studies in business-related domains find a strong relationship between customer sentiment and firm-level outcomes. Luo et al. (2013) report that social media-based metrics such as web blogs and consumer ratings are significant indicators of firm equity value. Goh et al. (2013) find that compared to user-generated content in brand communities, marketer-generated content has a weaker impact on consumer purchasing behavior. An effective way to influence customer sentiment is to use expert testimonies. In a study by Luo et al. (2017) sentiments expressed by experts on blogs are found to positively influence consumer perceptions of the focal brand. Lau et al. (2012) report that domain-specific sentiment analysis can be used to make more informed decisions about business mergers and acquisitions. Customer sentiment is not just an individual expression, but also a social phenomenon. Forman et al. (2008) report that online community members rate reviews containing identity-descriptive information more positively, and the prevalence of reviewer disclosure of identity information is associated with increases in subsequent online product sales. Lak and Turetken (2017) provide support for the positive impact of sentiment scores on the efficiency of purchase decisions (time to make a purchase decision), but do not consider its role to be influential in increasing decision effectiveness (customer's confidence in the purchase decision). Finally, Archak et al. (2011) report that textual data can be used for the predictive modeling of future sales.

Table 1 summarizes several studies, topics, and main findings in this area. These studies suggest that customer sentiment could influence firm equity value, branding activities, merger and acquisition decisions, and product sales and purchase decisions. Although these studies suggest the importance of customer sentiment for firm success, they do not prescribe any systematic strategy to manage or influence customer sentiment. Our study intends to fill the gap, namely, to study response strategies to manage customer sentiment within social media sites such as Twitter.

2.2. Digital Customer Care

Like any other business function, customer service operations have also been impacted by the digital transformation. According to Bianchi et al. (2014), roughly 70% of the customers use online platforms to purchase telecommunications-related products and services, and almost 90% of the customers use online platforms for customer service inquiries. Digital care platforms including social media websites, have empowered customers to express their opinions about firms and their

Table 1 Summary of Representative Studies Related to Customer Sentiment.

Study	Topic	Main Findings
Luo et al. (2013)	Social Media & Firm Equity Value	Social media-based metrics such as web blogs and consumer ratings are significant, leading indicators of firm equity value.
Goh et al. (2013)	Social Media Brand Community & Consumer Behavior	Compared to user-generated content in brand communities, marketer-generated content has a weaker impact on consumer purchasing behavior.
Luo et al. (2017)	Expert Blogs' Impacts on Consumer Perceptions of Competing Brands	Expert blog sentiments and volume on a focal brand have a positive relationship with consumer perceptions of the focal brand.
Lau et al. (2012)	Adaptive Decision Support for Business Mergers and Acquisitions	Domain-specific sentiment analysis can be used to make more informed decisions about business mergers and acquisitions.
Forman et al. (2008)	The Relationship Between Consumer Reviews and Product Sales	Online community members rate reviews containing identity-descriptive information more positively, and the prevalence of reviewer disclosure of identity information is associated with increases in subsequent online product sales.
Lak and Turetken (2017)	The Impact of Sentiment Scores on Purchase Decisions	Sentiment scores improve the efficiency (speed) of purchase decisions without significantly affecting decision effectiveness (confidence).
Archak et al. (2011)	Product Features and Consumer Reviews	Textual data can be used for predictive modeling of future changes in sales.

products and services. Customers can take advantage of these platforms to seek better and faster response from firms. An important survey published by The Institute of Customer Service (2013), reported that one undeniable impact of the digital customer service was that customers expect better service from firms. According to the same report, although "the percentage of customers who experienced a problem has decreased from 17% in January 2008 to 11.7% in July 2012, the proportion of those went on to make a complaint rose from 72% to 76% in the same period." Ma et al. (2015) suggest that customers decry or complement a firm on social media depending on their relationship with the firm (negative, neutral, or positive). Furthermore, interventions made by firms to modify firm-customer relationships often encourage further complaints in future.

The majority of academic and practitioner publications assert that digital care platforms have a lot to offer to firms. For instance, firms can lower their cost of providing customer service through the effective use of digital care platforms (Bianchi et al. 2014). Firms can also encourage experienced customers to support more novice customers (Lithium Technologies Inc. 2017). Finally, customer satisfaction on digital platforms appears to be higher than the same on traditional channels such as phone and mail (Bianchi et al. 2014).

2.3. Stochastic Differential Equations

In this study, we develop a Stochastic Differential Equation (SDE) model of the dynamics of the online customer sentiment in the presence of a response strategy used by the firm. Stochastic Differential Equation has been applied in finance to model the time series of stock price movements, where randomness is inevitable (Black and Scholes 1973). A stochastic process models the random variable of interest that varies continuously (or almost continuously) and stochastically through time. We model customer sentiment as a Markov process, where the probability distribution of the future value depends only on its current value, which subsumes the effects of past values of the process. Randomness is captured by a Weiner process (the continuous limit of Random Walk), a fundamental building block for randomness in stochastic processes. A Stochastic process, is defined to be a Wiener process if it has independent increments that are normally distributed with mean zero and variance equal to the time interval (dt) between the increments (Ross 2014). We next develop an SDE model for the evolution of customer sentiment over time.

3. Model Development

In this section, we develop a diffusion model of customer sentiment followed by a stochastic control model of how this diffusion model is driven by the control objectives of the firm.

3.1. Diffusion Model of Customer Sentiment

In a typical digital customer care setting (such as the one in Twitter), customer comments (or tweets) arrive in a chronological sequence forming time series data. We perform sentiment analysis on each tweet and generate a sentiment score corresponding to each tweet. The sentiment scores are normalized to range from zero and a maximum value M.

On social media sites, customer opinions arrive randomly and can be very different across customers. Firms usually respond to every single tweet posted by the customer, but all tweets may not need the same amount of effort to be addressed. For instance, some tweets could be replied to using a template response such as "Thank you @username for contacting our support team." On the other hand, some of the customer tweets may require additional effort by the firm. Therefore, firms often use a ticket-generation system to be able to detect the tweets that require more effort (zendesk.com 2018). However, the effort spent to respond to different tweets is usually not publicly obsvervable. These systems usually use the sentiment of the customer's post (tweet) as one of the indicators for generating tickets (zendesk.com 2018) and hence effort allocation decisions. A dissatisfied customer with a negative sentiment in her/ his tweet has a higher likelihood of requiring additional effort (ticket generation). In our setting, the aggregate customer sentiment at any point in time is affected by the arrival of new customer sentiment as well as response effort exerted by the firm in order to address customer needs. If the firm is able to properly handle the tickets, the sentiment will likely be adjusted in future periods. The sentiment (denoted by the state variable x_t) is derived from tweets posted by customers during each time period t and then computing its moving average over a chosen time window. In time series data, a moving average is commonly used as the measure of the state (Gooijer and Hyndman 2006). For example, moving average of a stock's price measures how that stock is trending; whereas in economics, moving average is commonly used to study the gross domestic product, employment rates, and other macroeconomic variables (Brown 2004).

The goal in this section is to model a stochastic process that captures the evolution of the state variable (x_t) over time. To model this process, we start with a simplified micro structure: the *change* in the state (dx_t) in a small interval from t to $t + dt$. The change in the state consists of a deterministic component, that is influenced by the effort the firm exerts to manage customer sentiment. In addition, it consists of a stochastic component that captures the randomness in the change in sentiment in a short time interval.

First, we discuss the deterministic component, or the expected change in the state, $\mathbb{E}(dx_t)$. The deterministic component consists of a positive term and a negative term. The negative term ($-k_2 x_t$) can be interpreted as the rate of decrease in the sentiment when the firm exerts no effort to manage it. Customers often have beliefs or expectations about the quality of customer care. The negative drift in the sentiment depends on the quality of digital customer care that customers expect from the firm (k_2) and the current sentiment (state) x_t. According to Parasuraman et al. (1991), "[...] customers want service companies to play fair. Customers are paying good money, and the company should provide good service in exchange." Given that the expectation of customer service for a similar product/ service would mainly depend on the price, firms that charge more for a similar product/ service would raise expectations of customer care. Hence, the negative impact of zero customer care effort is higher when customer expectation of care (k_2) is also high. Furthermore, for higher values of x_t, the negative impact is also higher. Given that we normalize the values of the sentiment such that the lowest sentiment score is zero, we ensure in the SDE that the sentiment cannot drop below zero. Next, we discuss the positive term within the deterministic component. The positive impact of the effort exerted by the firm depends on the difference between the highest possible sentiment score (M) and the current sentiment (x_t). When the sentiment is already high, the impact of effort diminishes. In the extreme case when $x_t = M$, there is no further gain from additional effort. The basic idea is that when the sentiment is low, there is more *room to grow* and hence, there is a higher positive impact of the effort exerted by the firm. The response effort is given by $u(x_t, t)$. Thus, we allow for the firm's response effort to be dynamic, and a function of the current sentiment, x_t. The impact of this effort on the state depends on the effectiveness of the response effort (k_1). The value of k_1 is firm-specific and depends on the knowledge level, reliability, responsiveness and assurance of the customer care team (Pitt et al. 1995). In addition, k_1 could represent the quality of the ticket-generation system (zendesk.com 2018). Thus, the positive impact of the response effort on the state is modeled as, $k_1 u(x_t,t)(M-x_t)$. Combining the effects on the firm's sentiment discussed above we can write:

$$\mathbb{E}(dx_t) = (k_1 u(x_t,t)(M-x_t) - k_2 x_t)dt \tag{1}$$

The stochastic component (diffusion) of the change in the firm perception is modeled as $k_3 d\omega_t$, where $d\omega_t$ is the Wiener process used to capture the white noise or randomness; $d\omega_t \overset{d}{=} N(0,dt)$. The parameter k_3 influences the magnitude of the random component. To summarize, we model the change in the state using the following SDE:

$$dx_t = (k_1 u(x_t,t)(M-x_t) - k_2 x_t)dt + k_3 d\omega_t \tag{2}$$

3.2. Stochastic Control Problem

The firm uses knowledge of the current customer sentiment (the state variable) to determine the response effort. The cost rate $q(x_t,t)$, measured in dollars (say) per unit time, is defined as below:

$$q(x_t,t) = \kappa u(x_t,t)^2 + c(M-x_t)^2 \tag{3}$$

The first term in the cost rate is the direct cost associated with response effort. The cost of response effort is assumed to be a quadratic function of the effort. This is consistent with other similar models (Sethi 1973), where effort is considered to be an input into a production function for producing the state (x_t). The parameter κ is the cost rate per unit effort.

The second term in the cost rate is the indirect cost associated with less than perfect customer sentiment. The indirect cost depends on a damage parameter c and the current sentiment. We assume the indirect cost to be increasing and convex with the deficit in customer sentiment from its perfect value. This implies that the damage from low customer sentiment gets increasingly higher.

We write the firm's objective as one of minimizing the total discounted cost over a planning horizon T, by determining the optimal trajectory for the control, namely, the firm's response effort, $u(x_t,t)$. Moreover, customer sentiment (x_t) is influenced by the control in a manner described by the stochastic differential equation (2) that acts as a constraint to the control problem.

Assuming a discount rate ρ, the above control problem can be expressed as shown below:

$$\underset{u(x_t,t)}{Min}\ \mathbb{E}\int_0^T q(x_t,t)e^{-\rho t}dt \equiv \underset{u(x_t,t)}{Min}\ \mathbb{E}\int_0^T \left(\kappa u(x_t,t)^2 + c(M-x_t)^2\right)e^{-\rho t}dt \tag{4}$$

subject to:

$$dx_t = (k_1 u(x_t,t)(M - x_t) - k_2 x_t)dt + k_3 d\omega_t$$

$$0 \le x_t \le M, u(x_t,t) \ge 0$$

Where M is the maximum sentiment score, k_1 is the effectiveness of response effort, k_2 is the expectation of customer care quality, k_3 is the magnitude of the random component, κ is the unit cost of effort, c is the damage parameter, ρ is the discount rate, and T is the planning horizon. Table 2 lists the main parameters used in this study for easy reference.

Table 2 Parameter Definitions

Parameter Name	Definition
M	Maximum sentiment score
k_1	Effectiveness of response effort
k_2	Expectation of customer care quality
k_3	Magnitude of the random component
κ	Unit cost of effort
c	Damage parameter
ρ	Discount rate
T	Planning horizon

To solve this problem, let $V(x_t,t)$, known as the value function, be the expected value of the objective from t to T, when an optimal policy is followed from t to T. Then, by the principle of optimality,

$$V(x_t,t) = \underset{u(x_t,t)}{Min} \; \mathbb{E}\left[(\kappa u(x_t,t)^2 + c(M - x_t)^2)dt + (1-\rho dt)V(x_t + dx_t, t+dt)\right] \tag{5}$$

Using Taylor's expansion of $V(x_t + dx_t, t+dt)$ in equation 5 and properties of stochastic calculus, we can derive the Hamilton-Jacobi-Bellman (HJB) equation for stochastic optimal control as below:

$$\rho V(x_t,t) = \underset{u(x_t,t)}{Min} \left[\kappa u(x_t,t)^2 + c(M - x_t)^2 + V_x(k_1 u(x_t,t)(M - x_t) - k_2 x_t) + \frac{1}{2}k_3^2 V_{xx} + V_t\right] \tag{6}$$

Differentiating with respect to the control, yields the following optimal form for the firm's response effort,

$$u^*(x_t,t) = -\frac{k_1}{2\kappa}(M - x_t)V_x(x_t,t) \tag{7}$$

Note that the optimal response effort in this expression depends both on the current state (x_t) as well as the marginal value function (V_x). In order to evaluate the value function, we substitute the optimal response effort into the HJB. This yields a Partial Differential Equation (PDE) as shown below, that needs to be solved to obtain the value function $V(x_t,t)$,

$$-\frac{k_1^2 V_x^2 (M-x_t)^2}{4\kappa} + c(M-x_t)^2 - k_2 x_t V_x + \frac{1}{2} k_3^2 V_{xx} + V_t = \rho V(x_t, t) \tag{8}$$

Equation 8 is a non-linear, second-order partial differential equation that does not lend itself to an analytical, closed-form solution. Therefore, we proceed with a numerical solution for $V(x_t, t)$. There is a large body of literature that proposes the use of Method of Lines (MOL) for numerical solutions of PDEs (Sadiku and Obiozor 2000). The first step is discretization along the x_t (state variable). The region is divided into strips by N dividing straight lines (hence the name method of lines) parallel to t (time)-axis. Since we are discretizing along x_t, we replace the first and second derivative with respect to x_t with its finite difference equivalent. In other words, with only one remaining independent variable (t), we have a system of Ordinary Differential Equations (ODE) that approximate the original PDE. Thus, one of the salient features of the MOL is the use of existing, and generally well-established, numerical methods for ODEs for solving PDEs. In our case, we use Backward Diffusion Formula (BDF), which is a well-known linear multi-step method for numerical integration of individual discretized ODEs. Ultimately, in this fashion, the PDE for $V(x_t, t)$ is solved along discrete lines in the x_t (state variable) direction, and the solution between lines is found by numerical interpolation. This allows us to numerically evaluate the optimal response effort, $u^*(x_t, t)$.

In Section 5, we use the values of the optimal response effort to estimate the parameters in the controlled diffusion process described in equation 2. Note that the controlled diffusion process (the process obtained after substituting for the optimal control), becomes a function of the parameters (k_1, k_2, k_3), the state variable x_t, and time t. The parameter M – the highest value of the sentiment – depends on the range of sentiment values and does not need to be estimated. The values of x_t and t are, of course, directly observed in the data.

4. Data

We collected data on customer service-related queries and firm responses from the Twitter service handles of the *Big Four* telecommunications firms (AT&T, Verizon, Sprint, and T-Mobile) from June 16 of 2016 to October 15 of 2016. According to statista.com, these four carriers secure almost 95% of the wireless communications market share in the United States.[1]

To collect this data, we made two types of API calls[2]: API calls to obtain all of the tweets *directed at* these firms, and, API calls to obtain all of the tweets *from* these firms. Figure 1 shows

[1] https://www.statista.com/statistics/199359/market-share-of-wireless-carriers-in-the-us-by-subscriptions/

[2] Application programming interface (API) is a set of procedures, protocols, and tools for building software applications. APIs make it easier for programmers to develop an application by providing the building blocks. An API call or request is simply a request sent to the API for a particular transaction (e.g. logging in or posting a comment or retrieving content)

an example of a tweet posted by a customer as well as the firm's response to that tweet. Appendix A describes our data collection process. Given our focus on customer service, we only focused on tweets through support accounts (e.g. "@ATTCares" rather than "@ATT"). Table 3 reports the counts of tweets "directed at" and "from" each firm's support account.[3]

Table 3 **Counts of Tweets "directed at" and "from" Support Accounts**

Firm	Count of tweets "directed at" firm's support account	Count of tweets "from" firm's support account
AT&T	82,667	83,186
Verizon	56,629	41,572
Sprint	58,726	89,956
T-Mobile	80,051	119,113

We measure customer sentiment by performing sentiment analysis on the tweets posted by customers on a firm's customer service account in Twitter. The Package "sentimentr" in R was used to perform this task. Before performing sentiment analysis, the following text pre-processing steps were applied. First, the entire corpus was converted into lowercase. Next, stop-words were removed from the corpus and SnowBall stemming was applied. Finally, punctuations and numbers were removed from the corpus.

To verify the accuracy of the sentiment analysis, we hired two graduate students to manually determine the sentiment of a random sample of 1,000 tweets. We then compared the human-generated sentiments with those generated by "sentimentr." A high match of 83.49% between the human-generated sentiments and sentimentr-generated sentiments was observed. Appendix B reports the details of our sentiment analysis approach. Since "sentimentr" returns the scores that range from a negative minimum to a positive maximum, we used Min-Max transformation to transform these sentiment scores to range from 0 to 1. Scores close to 0 are from very negative tweets while scores close to 1 are from very positive tweets.

After determining the sentiment of each individual tweet, the mean sentiment score during each hour can be calculated. Although our approach can work with any unit of time (day, hour, minute, etc.), we use one hour as the unit of time. There are several reasons for this choice. Mainly: First, if the unit of time is too small, there are many periods with no customer tweet. On the other hand if the unit of time is too large, there is too much aggregation of customer sentiment. According to Table 3, the number of tweets directed at firms range form 56,629 to 82,667. Since we collected data from June 16 of 2016 to October 15 of 2016, there are 2,928 hours in this period,

[3] T-Mobile and Sprint often respond to a single customer query with more than one tweet. This explains why the number of tweets from these firm are more than the number of tweets from customers. In Appendix A, we discuss how we identified these types of tweets.

Figure 1 **Customer Care Transaction in Twitter**

resulting in tweets that range between 19.341 and 28.233 tweets per hour. Second, according to Hutchinson (2017), 72% of customers who complain on Twitter expect to receive a response within one hour. Therefore, firms need to prepare an hourly planning scheme to be able to keep up with the customers' expectations.

The state variable x_t in our model is the hourly moving average (two-hour window) sentiment score at time t. During certain hours, there were no tweets posted. To address this, we used Predictive Mean Matching (PMM) to impute the missing sentiment scores for these hours. Less than 10% of hours are imputed for the four firms. Our data set includes both original tweets and retweets.

Table 4, reports the mean and Gini's Mean Difference (GMD), which is a measure of variability.[4] Here, T-Mobile has the highest mean sentiment score, followed by Verizon, AT&T, and Sprint. The Gini's Mean Difference is similar across the four firms. Figure 2 visualizes the hourly sentiment scores for the *Big Four* telecommunications firms for the period of our study.

Table 4 **The Mean and Variation of Sentiment Scores for the Big Four**

Firm	Mean	Gini's Mean Difference (GMD)
AT&T	0.468	0.040
Verizon	0.472	0.040
Sprint	0.463	0.043
T-Mobile	0.480	0.040

[4] According to Yitzhaki (2003), GMD is a superior measure of variability as it is more informative than variance if the data is not normally distributed.

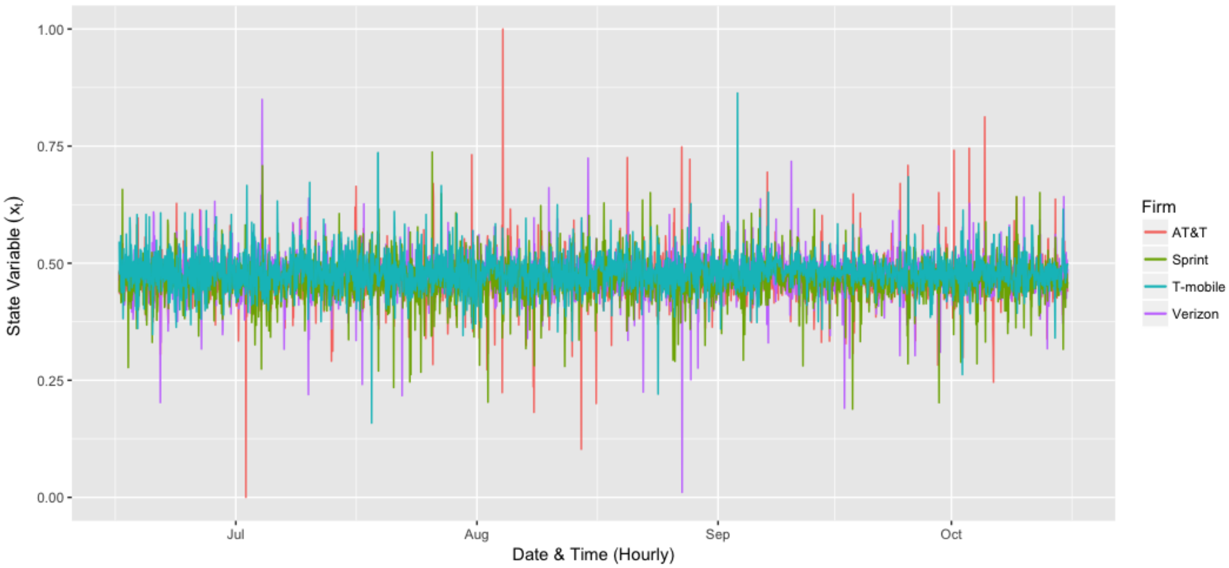

Figure 2 The State Variable (Sentiment Score) for the Big Four

5. Model Estimation

We apply the Maximum Likelihood Estimation (MLE) procedure to recover the parameters in our model. Typically, however, MLE requires knowledge of the probability density function of x_t. However, deriving the value function $V(x_t, t)$ discussed in section 3 is intractable. Hence, it is infeasible to obtain an analytical solution for the probability density of x_t. Even so, the change in customer sentiment (dx_t), has some useful properties that can be exploited. The density of dx_t is a Gaussian, with a mean $(k_1 u(x_t, t)(M - x_t) - k_2 x_t)dt$ and standard deviation $k_3\sqrt{dt}$. In particular, the probability density function of dx_t, that represents the change in state between t and $t + dt$ can be written as:

$$p(dx_t|x_t, k_1, k_2, k_3, \rho, c, \kappa) = \frac{1}{\sqrt{2\pi(k_3 dt)^2}} \exp^{\frac{-(dx_t - \mu(x_t,t))}{2(k_3 dt)^2}} \tag{9}$$

where:

$$\mu(x_t, t) = (k_1 u(x_t, t)(M - x_t) - k_2 x_t)dt$$

For sentiment data with T observations, we approximate the infinitesimal quantities (namely, dx_t, dt) with their numerical analogues. Thus, we write:

$$dx_{t_i} \approx \Delta x_{t_i} = x_{t_{i+1}} - x_{t_i}$$
$$dt \approx \Delta t = t_{i+1} - t_i.$$

If we define $\boldsymbol{\theta} \equiv (k_1, k_2, k_3, \rho, c, \kappa)$, the log-likelihood function for Δx_t with T observations is:

$$\ln L(\boldsymbol{\theta}) = \sum_{i=1}^{T-1} \ln p(\Delta x_{t_i}|x_{t_i}; \boldsymbol{\theta}) \tag{10}$$

By substituting the density function and simplifying we get:

$$\ln L(\boldsymbol{\theta}) = -\frac{T-1}{2}\ln 2\pi - (T-1)\ln k_3\Delta t - \frac{1}{2(k_3\Delta t)^2}\sum_{i=1}^{T-1}(\Delta x_{t_i} - \mu(x_{t_i}, t_i))^2 \tag{11}$$

where $\mu(x_t, t) = (k_1 u(x_{t_i}, t_i)(M - x_{t_i}) - k_2 x_{t_i})\Delta t$.[5]

For every observation (x_{t_i}, t_i) in the data and parameter set $\boldsymbol{\theta}$, the optimal response (control) $u(x_{t_i}, t_i)$ is obtained by using

$$u(x_{t_i}, t_i) = -\frac{k_1}{2\kappa}(M - x_{t_i})V_x(x_{t_i}, t_i) \tag{12}$$

In the above, $V_x(x_t, t)$ is obtained as a numeric value by numerically solving

$$-\frac{k_1^2 V_x^2 (M - x_{t_i})^2}{4\kappa} + c(M - x_{t_i})^2 - k_2 x_{t_i} V_x + \frac{1}{2}k_3^2 V_{xx} + V_t = \rho V(x_{t_i}, t_i), V(x_{t_i}, T) = 0$$

The Maximum Likelihood estimate $\hat{\boldsymbol{\theta}}$ is obtained by maximizing the log-likelihood function described in equation 11 over its parameter space:

$$\hat{\boldsymbol{\theta}} \equiv (\hat{k_1}, \hat{k_2}, \hat{k_3}, \hat{\kappa}) = \arg\max_{\boldsymbol{\theta}} \ln L(\boldsymbol{\theta})$$

In order to estimate $\hat{\boldsymbol{\theta}}$, we need to numerically solve $V_x(x_t, t)$ using the MOL approach. For this analysis we considered a total of 10^{12} possible value sets for $\boldsymbol{\theta}$. Four parameters appear in the Log-likelihood expression. These are the effectiveness of effort by the firm (k_1), the expectation of customer care quality (k_2), the magnitude of the random component (k_3), and the unit cost of effort (κ). However, we cannot separately estimate k_1 and κ, since these parameters appear together as a ratio ($\frac{k_1}{\kappa}$). Since we are interested in the effectiveness of response effort (k_1), we fixed the unit cost of effort ($\kappa = 1$). We also verified the robustness of the estimation results over a range values of κ. The values of c and ρ do not affect the estimates and were fixed at 1 and 0.0005.

6. Estimation Results

Table 5 presents the estimation results. Based on our findings, Verizon, with the largest value for effectiveness of response effort (k_1), is the most effective in responding to customer sentiment and Sprint, which has the smallest value for k_1, is the least effective. We note that while T-Mobile and Sprint have higher response rates (the number of firm-generated tweets divided by the number of customer-generated tweets), they are low on effectiveness. Our findings concerning the effectiveness

[5] The Gaussian property of $dx_t|x_t$ may not always hold for $\Delta x_t|x_t$, and hence needs to be verified in the data. Conceptually, if the data is sufficiently granular, implying that data points are not too far apart in time, the Normal assumption is likely to be upheld. In Appendix D, the Normal assumption is verified in the Twitter feed collected for all four firms.

of the response effort (k_1) are aligned with the inverse of average monthly blended customer churn for the second quarter of 2016. This period overlaps with most of the time period of our data set which was from June 16 of 2016 to October 15 of 2016. According to a research study by Strategy Analytics (Dano 2016), Verizon's blended customer churn[6] during this time period was 1.19%, the lowest among the *Big Four*. The second spot in the report belonged to AT&T with a blended churn rate of 1.35%. The third and fourth spots belonged to T-Mobile and Sprint with blended churn rates of 2.22% and 2.81% accordingly. This alignment between the inverse of customer churn during the second quarter of 2016 and the effectiveness of the response effort during an overlapping period lends support to our findings.

Table 5 **Maximum Likelihood Estimation Results for the Entire Data**[7]

Firm	**Effectiveness of response effort (k_1)**	**Expectation of customer care quality (k_2)**	**Magnitude of random component (k_3)**
AT&T	0.396*** [0.011]	0.164*** [0.007]	0.028*** [<0.001]
Verizon	0.408*** [0.013]	0.165*** [0.006]	0.025*** [<0.001]
Sprint	0.358*** [0.010]	0.151*** [0.005]	0.027*** [<0.001]
T-Mobile	0.393*** [0.015]	0.153*** [0.009]	0.025*** [<0.001]

The parameter k_2 measures the quality of digital customer care that customers expect from a firm. Based on the estimates for k_2, we can clearly divide the *Big Four* into two groups of firms: AT&T and Verizon, and, Sprint and T-Mobile. AT&T and Verizon have very similar values for k_2 which are greater than those of Sprint and T-Mobile. As indicated previously, a greater value of k_2 is associated with a larger downward drift in customer sentiment when the firm's effort (u) is held constant. Our findings therefore suggest that AT&T and Verizon's customers expect better quality of care than the customers of Sprint and T-Mobile. This is consistent with the plan-price rankings of these firms, with Verizon and AT&T being more expensive than Sprint and T-Mobile.[8] The values of the random component (k_3) for the four firms are similar, implying that the firm's effort does not appear to influence the underlying variability of the sentiment data.

[6] Blended churn rate, as reported by U.S. wireless carriers, is a churn rate figure that is based on both pre-paid and contract customer losses. This figure is one of the main figures in evaluating wireless carriers performance.

[7] Table notes: *** significant at 0.001; standard errors are reported in brackets.

[8] According to a Business Insider post (Dunn 2016), during the time period of our study, the cheapest unlimited plans (unlimited talk, text, and data) offered by Sprint, T-Mobile, and AT&T were, respectively, $60, $70, and $100 per month for the first line. Verizon did not offer an unlimited data plan at the time, but offered something similar with unlimited talk and text plus 24GB of data for $110 per month (komando.com 2016).

6.1. Comparison with Structure-free Models

The best validation of our structural model would have been through an experiment that allowed us to alter a firm's response effort and compare the observed outcomes with our predictions. Unfortunately, we did not have the luxury to conduct experiments in any of the firms. According to Keane (2010), another rigorous approach to examine the validity of an structural model is to compare the predictions obtained from the model with those obtained from state-of-the-art, but structure-free models (e.g. using time series analysis). This is the approach taken in our study.

To measure the predictive performance of a model, we use the first two-thirds of the data for building the predictive model, and the remaining one-third of the data for testing the model. Furthermore, we use one-step ahead forecasting to generate a forecast. In one-step ahead forecasting, we use previous forecasts to generate future forecasts. We compared the predictions from the structural model with five structure-free models: ARIMA, ARCH, GARCH, ARIMA+GARCH, and ARIMA+apARCH. Additional details of these comparisons are provided in Appendix C. We next summarize the results of this comparison.

The results in Tables 6 and 7 show that the SDE model performs slightly better than advanced structure-free models such as ARIMA+GARCH, and ARIMA+apARCH. Figure 3 visualizes the changes in the actual data as well as the forecasts obtained from the SDE and structure-free models.

Table 6 Mean Absolute Error (MAE) Comparisons

Model	AT&T	Verizon	Sprint	T-Mobile
SDE	0.011	0.011	0.010	0.009
ARIMA	0.014	0.014	0.013	0.012
ARCH	0.019	0.018	0.017	0.018
GARCH	0.019	0.018	0.017	0.018
ARIMA+GARCH	0.013	0.013	0.012	0.012
ARIMA+apARCH	0.013	0.013	0.012	0.012

Table 7 Symmetric Mean Absolute Percent Error (SMAPE) Comparisons

Model	AT&T	Verizon	Sprint	T-Mobile
SDE	2.944	2.913	2.903	2.271
ARIMA	3.002	2.999	2.960	2.436
ARCH	4.113	3.954	3.817	3.752
GARCH	4.107	3.933	3.829	3.752
ARIMA+GARCH	2.969	2.935	2.789	2.491
ARIMA+apARCH	2.964	2.939	2.798	2.500

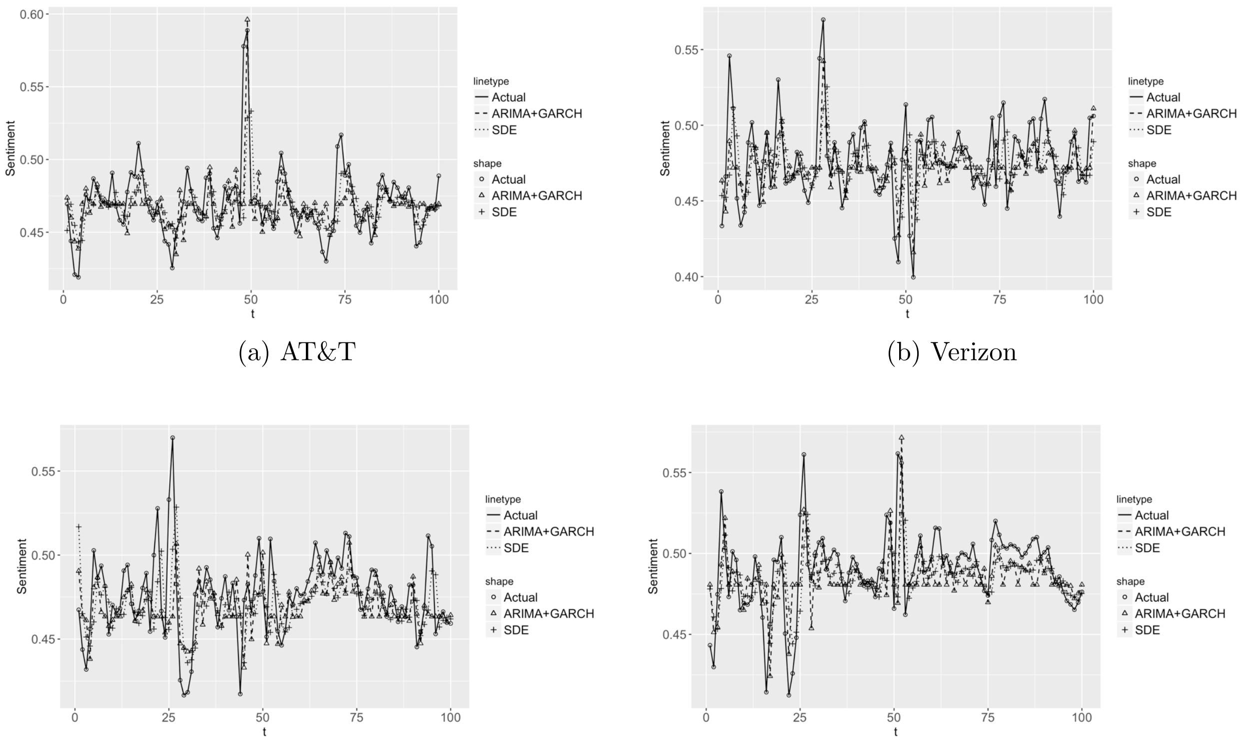

(a) AT&T (b) Verizon

(c) Sprint (d) T-Mobile

Figure 3 Forecasts of 100 Hours of Test Data (Actual x_t VS SDE Predections VS ARIMA+GARCH Predictions)

7. The Effects of Major Events

Although structure-free models work relatively well to forecast customer sentiment, these models do not help firms determine the optimal response effort needed to manage customer sentiment. In addition, the forecast from a structure-free model will only be accurate if the environmental conditions underlying the data do not change. A structural model, on the other hand, allows the firm to conduct policy simulations, i.e., determine the potential impact of environmental events (such as a price change, introduction of a new phone plan, a new marketing campaign, etc.) on the parameters that influence customer sentiment (e.g., the effectiveness of care (k_1), or the quality of care that the customers expect from the firm (k_2)). Because the directional change (up or down) in these parameters can be anticipated by the firm, the change in the optimal response effort can also be anticipated. In what follows, we consider different events that could potentially influence customer care and the optimal response effort. All these events can be considered to be external to the customer care platform, i.e., we deliberately chose events that are typically not controlled by customer care managers. Furthermore, these events can be classified into two kinds: Firm-initiated events, and Exogenous events (not initiated or controlled by the firm).

7.1. Firm-initiated Events

These events are chosen such that one could expect them to have a positive or a negative impact on customer sentiment. For instance, a price increase for popular plans is likely to draw a negative reaction from customers. The question arises: does this reaction spill over to create negative sentiment in the care platform? If so, what is the mechanism behind the drop in sentiment? For example, the sentiment could fall if the parameter k_2 increases as a result of the price increase. This is plausible since customers would expect better care from a firm that charges more. If customer care managers can anticipate the impact of a price increase, they could adapt accordingly, e.g., increase the response effort. On the other hand, when an event is perceived positively by customers, they might cut the firm some slack regarding service-related issues. For instance, if a firm offers more data download for the same price, its customers may be willing to tolerate a drop in the quality of customer care. Once again, it would be beneficial for customer care managers to anticipate the change in the parameters that affect customer sentiment so that they determine a new level of response effort.

To examine the potential effects of firm-initiated events, we focus on the impact on k_2, the parameter in our model that represents the quality of care that customers expect from a firms. We compare the value of k_2 for the three-day period before the event with its value for the three-day period after the event. To identify the events during the study period, we used Google search with the name of the firm (e.g. "AT&T") as the search keyword and the study period as the interval of time. Next, we used the first Google result that was related to an action initiated by the firm as an event to study for each firm. These events are discussed below.

7.1.1. AT&T: Exclusive Contract with Taylor Swift Our first example is an event that, at first glance, appears to be unrelated to customer care. On October 4th 2016, AT&T announced an exclusive, multi-year deal with Taylor Swift.[9] The deal resulted in a new service called "Taylor Swift Now", an exclusive AT&T service that would allow AT&T customers access to exclusive content and events related to Taylor Swift. Importantly, there was no corresponding price increase. Given that the customers would react positively to this event, we tested to see if this event had a spillover effect on customer care. As reported in Table 8, the value of k_2 decreased from 0.174 to 0.167 as a result of this event. This was a positive outcome for AT&T customer care because customer sentiment would decrease at a slower pace for the same level of response effort.

7.1.2. Verizon: Price Hike On July 6th 2016, Verizon announced an increase in the price that also increased the control customers had over their data plan.[10] Although more control over data plans is a positive change, the increase in price is clearly not so.[11] As T-Mobile's CEO John

[9] http://fortune.com/2016/10/04/taylor-swift-att-deal-super-bowl/

[10] http://money.cnn.com/2016/07/06/technology/verizon-data-plan/

[11] http://time.com/money/4394575/verizon-price-increase/

Legere summarized it in Twitter: "And OMG, my favorite part is that @verizon is going to charge you $5, to promise not to charge you for overages?" Thus, the overall impact of this event is not clear.

Table 8 reports the results of the changes in k_2 due to this firm-initiated event. During the three-day period before the announcement, k_2 had a value of 0.161. During the three-day period after the announcement k_2 increased to 0.174. This increase in k_2 implies that the event was negatively perceived by Verizon's customers. While it may be argued that a price increase is likely make customers demand better quality of care, we substantiate this intuition as well as quantify the magnitude of the effect.

7.1.3. Sprint: "Unlimited Freedom" Plan On August 18th 2016, Sprint launched its new "Unlimited Freedom" plan.[12] This plan was claimed to allow customers to use more data with lower cost (the lowest among the *Big Four*). Therefore, this is perceived as a positive change from the customer's point of view. As per results reported in Table 8, the value of k_2 decreased from 0.159 to 0.152. This decrease in k_2 is expected as this was a positive event for Sprint's customers.

7.1.4. T-Mobile: "Unlimited One" Plan On the same day that Sprint announced its new "Unlimited Freedom" plan, T-Mobile announced its "Unlimited ONE" plan that gave customers the privilege to download more data without any increase in price.[13] Similar to the effect of the event on Sprint's customer care, we believe that this event should have a positive influence on T-mobile's customers care. According to the results reported in Table 8, k_2 decreased to a value of 0.153 from a value of 0.16 after the announcement of the new plan.

Table 8 Change in Expectation of Customer Care Quality (k_2) due to Firm-initiated Events

Firm	k_2 Before the Event	k_2 After the Event	k_2 Change
AT&T	0.174	0.167	**-0.007**
Verizon	0.161	0.174	**+0.013**
Sprint	0.159	0.152	**-0.007**
T-Mobile	0.160	0.153	**-0.007**

7.2. External Event: Release of iPhone 7

This is an example of an external event that is not initiated by firms. On September 9 2016, iPhone 7 was released by Apple. Given the popularity of iPhones in the U.S., many customers ordered the new iPhone immediately after its release. Did this cause a significant increase in customer care requests? For instance, the requests could increase if customers experience service issues when

[12] https://www.engadget.com/2016/08/18/sprints-unlimited-freedom-plan/

[13] https://newsroom.t-mobile.com/news-and-blogs/rip-data-plans.htm

ordering the new iPhone, if they need assistance to complete the order, or if they have questions about delivery date and other logistical issues. Figure 4 depicts the count of customer tweets during each hour from the beginning of September 8 2016 to the end of September 10 2016. The plots (a) through (d) in Figure 4 clearly show that all four firms experienced a surge of customer tweets on September 9 2016. Therefore, it is reasonable to expect that the release of the new iPhone would impact the customer care workload, and could perhaps, reduce the quality of customer care.

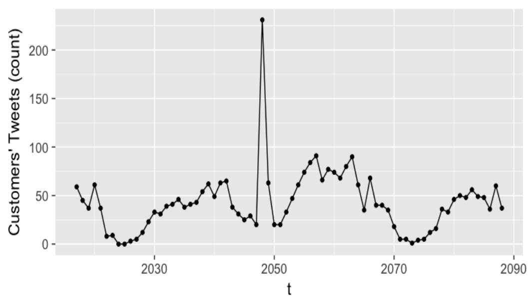

(a) AT&T

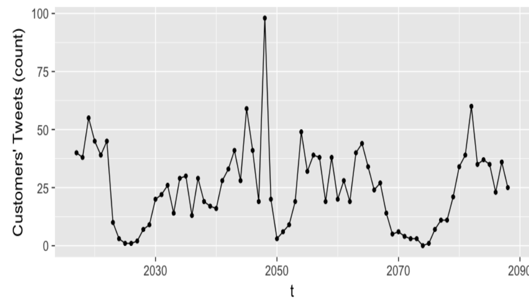

(b) Verizon

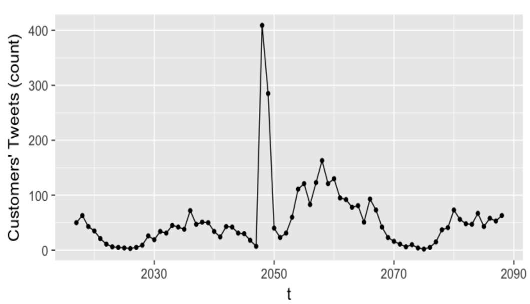

(c) Sprint

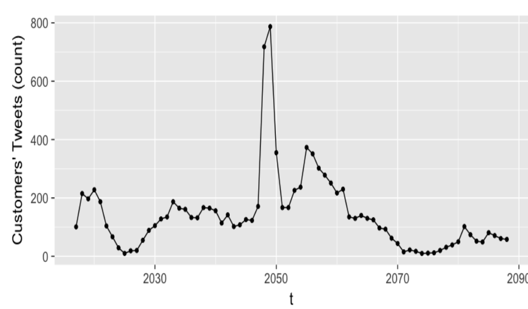

(d) T-Mobile

Figure 4 The Surge in Customer Tweets during iPhone 7 Release on September 9th, 2016

Given that all four firms experienced a sudden increase in customers tweets caused by the release of iPhone 7, we ask: How prepared were these firms to respond to customer queries during this time? Also, did the event adversely affect the effectiveness of customer care? We estimated k_1 for a three-day period before and a three-day period after the release of iPhone 7 for all four firms. Table 9 reports the changes in k_1 before and after the release of iPhone 7. As expected, the estimate for k_1 dropped for all four firms within the few days after the release of iPhone 7. Among the *Big Four*, Verizon had the smallest change in k_1, and AT&T had the second smallest change. The other two firms (T-Mobile and Sprint) had the biggest drop in the effectiveness of response effort. It is also worth noting that T-Mobile received more tweets during this period when compared to other firms. The impact the iPhone release had on k_1 perfectly matches the order of the number of tweets each firm received during this time period (i.e. T-Mobile, Sprint, AT&T, and Verizon).

Table 9 **Changes in the Effectiveness of Response Effort (k_1) during the Release of iPhone 7**

Firm	k_1 Before the Event	k_1 After the Event	k_1 Change
AT&T	0.329	0.322	**-0.007**
Verizon	0.348	0.345	**-0.003**
Sprint	0.321	0.299	**-0.022**
T-Mobile	0.361	0.337	**-0.024**

8. Structural Properties: Optimal Response and Cost

Our goal in this section is to study the optimal response strategy and the optimal cost obtained from solving the stochastic control problem in section 3.2. In particular, we develop a closed-form approximation of the steady state optimal response strategy and cost. This task is achieved using the following approach. We numerically solve the HJB to evaluate the value function over a large range of parameter values and make an informed guess about the form of the value function. The form of the value function provides the form of the optimal response strategy.

8.1. Numerical Solution of HJB

An illustrative solution for the value function and the response effort as a function of the customer sentiment (x_t) at time t is shown in the 3D plots below.

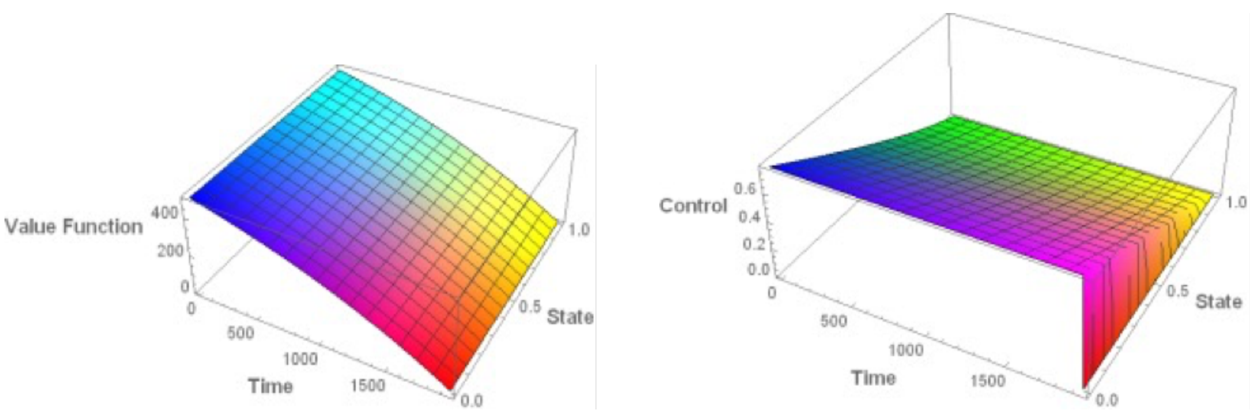

(a) Value Function (b) Response Effort

Figure 5 **Value Function and Response Effort**

Figure 5 (a) indicates that the value function is linear in the state variable. Figure 5 (b) illustrates that the optimal effort reduces with current state of the customer sentiment. This is intuitive in that the firm should increase the response effort when customer sentiment is low. Towards the end of the planning horizon, there is less incentive to incur any additional cost of response effort. Thus, the optimal response effort has the extreme value $u^*(x_t, T) = 0$. Based on these observations, we propose the following functional approximation for the optimal response strategy:

$$u^*(x_t) = \alpha(M - x_t)(1 - \exp^{-\beta(T-t)}) \tag{13}$$

If the planning horizon is relatively large, we can study the behavior of the optimal effort in steady state or independent of time, i.e., $u^*(x) = \alpha(M - x)$. From (7) we know that $u^*(x_t, t) = -\frac{k_1}{2\kappa}(M - x_t)V_x(x_t, t)$. Thus, under steady state, the marginal value function should be constant, i.e., $V_x(x) = -\frac{2\kappa}{k_1}\alpha$. This is consistent with the observation that the value function in Figure 5 (a) is linear in the state variable.

As a robustness check, we compare the marginal value function ($V_x(x_t, t)$) approximated as $-\frac{2\kappa}{k_1}\alpha$ with that obtained using the standard MOL approach for our data set. Figure 6 plots the numerical value of V_x (obtained using MOL) versus the analytical approximation ($V_x(x_t, t) = -\frac{2\kappa}{k_1}\alpha$) obtained above. As can be seen, the approximation performs well (less than 5 percentage point difference) over most of the planning horizon. However, as expected, we observe an increase in the percentage difference towards the end of the planning horizon, since we do not consider the exponential term involving time. We also study the average percentage difference over a range of parameter values across all four firms and find that the marginal value function is, on average, within 2% of the actual value obtained using the MOL approach.

8.2. Mode of the State Variable

Using the expression for the steady state effort and equation (3), the optimal steady state cost per unit time is given by $q^*(x) = (\kappa\alpha^2 + c)(M - x)^2$. Because the expressions for the optimal effort and cost both involve the state variable x, we propose studying the behavior of the optimal effort at the most frequent value of the state, i.e., its mode. Unfortunately, an exact expression for the mode cannot be found. An intuitive approximation for the mode can be found by finding the value of the state (say, $\hat{x}$) such that $\mathbb{E}(dx_t) = 0$. As the value of x_t approaches $\hat{x}$, the magnitude of the change in the state variable can be expected to diminish, implying that the process should spend maximum time in the vicinity of $\hat{x}$. We therefore have:

$$k_1\alpha(M - \hat{x})^2 - k_2\hat{x} = 0$$

The above equation has two roots, but the positive root is infeasible since it is greater than M. Therefore,

$$\hat{x} = \frac{1}{2}\left\{2M + \frac{k_2}{k_1\alpha} - \sqrt{\left(\frac{k_2}{k_1\alpha}\right)^2 + \frac{2Mk_2}{k_1\alpha}}\right\} \tag{14}$$

To check the accuracy of the above approximation, we simulated 10,000 sample paths – each consisting of 20,000 time steps – of the stochastic process in equation (2). The above process was performed for 1,000 combinations of the independent parameters viz., k_1, k_2, k_3 and α. The mode of x_t found from simulation was found to agree, very closely, to the mode provided by the expression in equation (14). The MSE of the approximation was less than 0.01% with a maximum error of 0.38%. We also performed a Welch Two Sample t-test to confirm that the means of the two populations (simulated versus predicted using equation (14)) are not different (p-value = 0.265).

8.3. Steady State Effort and Cost

The approximate expressions for steady state optimal effort and cost are:

$$u^*(\hat{x}) = \alpha(M - \hat{x})$$

$$q^*(\hat{x}) = (\kappa\alpha^2 + c)(M - \hat{x})^2.$$

The value of $\hat{x}$ is obtained from equation (14). The value of α is a function of the model parameters: k_1, k_2, k_3, ρ, c, and κ. For a given set of parameter values, we use singular value decomposition (SVD) with an objective to minimize the least square error to determine the values of α that best fits the proposed form of the steady state marginal value function. Appendix E provides the details of this study. The following form for α best fits the data:

$$\alpha = a_0 + \frac{a_1}{k_1} + \frac{a_2}{k_2} + \frac{a_3}{\kappa} + a_4 c + a_5 \rho \tag{15}$$

To recall, k_1, k_2 and κ – firm dependent parameters – correspond respectively to the effectiveness of response effort, customer expectation of care, and the cost of response effort. The parameter c corresponds to the marginal cost of a reduction in customer sentiment (damage parameter) and ρ is the discount rate. This coefficient estimates in (15) are: $a_1 = -0.107$, $a_2 = 0.035$, $a_3 = 0.361$, $a_4 = 0.344$, and $a_5 = -5.080$. The approximation performs well with a MSE of about 3%.

8.4. Structural Properties

Table 10 provides the estimated values of the derivative $\frac{\partial u^\star}{\partial \zeta}$ and $\frac{\partial q^\star}{\partial \zeta}$, where $u^\star$ is the optimal effort, $q^\star$ is the optimal cost per unit time, and ζ is a problem parameter. The problem parameters k_1 and k_2 are directly related to customer care, whereas the parameters κ and c are financially related, and indirectly affect customer care. The derivative of the optimal response effort with respect to financial parameters does not vary much across the firms. That is, while financial parameters clearly affect the optimal response effort, there is not much *differential* impact across the firms. This is expected since factors like the cost of response effort or the marginal benefit of increasing customer sentiment should affect any firm in more or less the same manner. Concerning κ, the unit cost of effort, it is intuitive that as this cost increases, the effort should decrease ($\frac{\partial u^*}{\partial \kappa} < 0$). Finally,

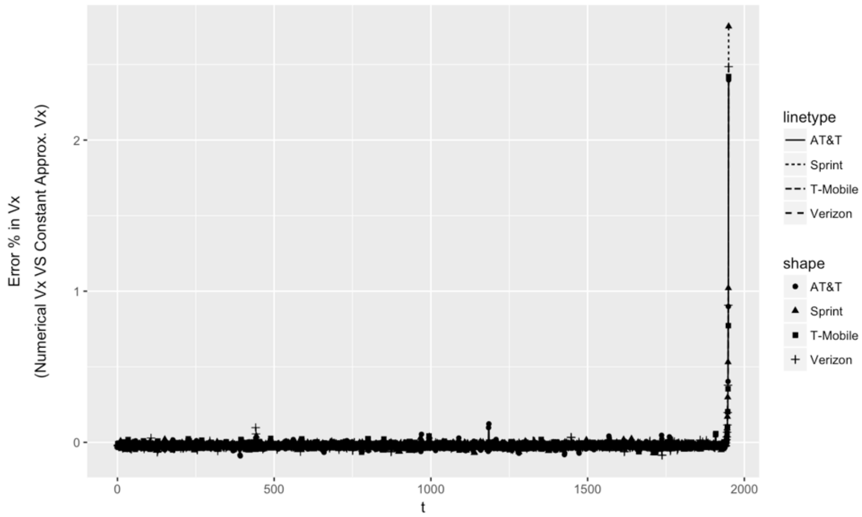

Figure 6 **Percentage Error between the Linear Approximated versus Optimal Marginal Value Function**

$\frac{\partial u^*}{\partial c} > 0$ simply means that as the damage of low customer sentiment increases, it pays to put more effort into responding to customer tweets.

On the other hand, the two digital care parameters, k_1 and k_2, do exhibit differential impacts (Table 10). First, there is a separation between the premium firms (AT&T and Verizon) and the bargain firms (T-Mobile and Sprint). The premium firms exhibit lower marginal benefit of the efficiency of response effort, implying that there are diminishing returns to efficiency improvement. The derivative of the optimal response effort with respect to k_2 shows a similar trend: (a) the premium firms separate from the bargain firms, and (b) the customers of the premium firms, being already more demanding of better service, have a lower marginal impact on the optimal response effort.

Turning to the marginal impact on cost, the two premium firms cluster well on the digital care parameters, k_1 and k_2. There is greater benefit for the bargain firms to improve the efficiency of response effort. This is likely because the bargain firms currently operate at a lower level of response effort efficiency. On the other hand, the cost for the bargain firms is more sensitive to an increase in k_2, the care customer's expect from a firm. The cost impacts of the financial parameters are more or less similar across the firms; these derivatives can be used by the firm to find how the increase in the price of a productive input (such as the salary of employees that provide customer care) can affect the cost of the digital care unit.

Not reported in Table 10, is information on how the different marginal impacts are affected by the level of another parameter, i.e., other than the one that is being varied. Since there are many such effects, we report only a few interesting ones here. For example, we find that optimal effort decreases with k_1 for low values of k_2, but changes direction later, i.e., for large values of k_2, the optimal effort increases with k_1. This finding suggests that when customer expectation of care is high (k_2 is high), as response effort gets more efficient (k_1 increases) more effort must be applied. However, when customer expectation is low, it is beneficial for the firm to decrease the effort and benefit from the reduction in the cost of effort. As can be observed in Table 10, response effort decreases with an increase in efficiency for all four firms. This implies all firms can be expected to reduce effort (and hence, incur lower personnel costs) if the tools supporting the care platform become more sophisticated and efficient.

Similarly, from equations (14) and (15), it is easy to see that the change in optimal response effort with an increase in k_2 can be positive or negative depending on the values of the other model parameters. For example, for relatively low values of k_2, we find that the response effort increases with k_2. This is intuitive as the response effort must increase as customers become more demanding of better customer care. However, for large values of k_2, there are decreasing returns of any further increase in effort and the impact is negative. Table 10 shows that the impact of k_2 on the optimal response effort is positive for all four firms, implying that all four firms can be expected to increase response effort if their customers become more demanding of care.

Table 10 **Sensitivity Analysis with Respect to Model Parameters**

	Optimal Response Effort				**Optimal Cost**			
Firm	$\frac{\partial u^\star}{\partial k_1}$	$\frac{\partial u^\star}{\partial k_2}$	$\frac{\partial u^\star}{\partial \kappa}$	$\frac{\partial u^\star}{\partial c}$	$\frac{\partial q^\star}{\partial k_1}$	$\frac{\partial q^\star}{\partial k_2}$	$\frac{\partial q^\star}{\partial \kappa}$	$\frac{\partial q^\star}{\partial c}$
AT&T	-0.218	0.145	-0.135	0.129	-1.143	1.581	0.138	0.142
Verizon	-0.214	0.149	-0.134	0.128	-1.091	1.554	0.140	0.149
Sprint	-0.263	0.164	-0.136	0.129	-1.396	1.774	0.137	0.116
T-Mobile	-0.244	0.184	-0.133	0.127	-1.199	1.693	0.144	0.138

9. Conclusion

In this study, we develop a controlled diffusion model of the dynamics of customer sentiment in the presence of a response strategy used by firms. We used the model to predict customer sentiment as well as to study the structural properties of the optimal response strategy. The optimal response strategy is a function of several firm-specific parameters and the current customer sentiment. We apply Maximum Likelihood Estimation to recover the parameters of the stochastic process describing the evolution of customer sentiment using the data from the *Big Four* telecommunications firms in the United States. We also compare our model's predictive performance on customer sentiment with the state-of-the-art structure-free models. These comparisons indicate that the predictive performance of our controlled diffusion model is better than the predictive performances of several well-known, structure-free models.

Given the availability of data from Twitter and other digital platforms, firms can easily capture, process, and monitor customer sentiment within these platforms (Fan and Gordon 2014). Firms can also (and do) deploy a variety of analytical tools to forecast customer sentiment over time. Although there is mounting evidence that the leading B2C firms monitor and respond to customers' service-related queries in social networking websites (Twitter for Customer Service Team 2015), there is little information about how customer sentiment can be managed using an optimal response strategy. For instance, it is known that some firms prioritize customers' tweets to efficiently allocate available resources in providing responses (Gunarathne et al. 2017, Muralidhar et al. 2015). However, there is limited research that prescribes a response strategy to respond to customer sentiment within digital platforms. The concern is that, even though firms capture, monitor, and forecast customer sentiment, many firms may not be equipped well to develop a response strategy that yields the desired outcomes. Therefore, our SDE model could help practitioners go beyond the existing methods used in practice.

Furthermore, our model can be used to compare firms with respect to their overall effectiveness in responding to customer sentiment as well as their customers' expectation of care quality. With respect to the effectiveness of digital customer care management, Verizon topped the list during the period of our study followed by AT&T, T-Mobile, and Sprint, in that order. The ranking of the effectiveness of digital customer care management of the four firms in our study perfectly aligns with their blended customer churn rates during the same time period. That is, Verizon had the lowest blended churn among the *Big Four*. Verizon's blended churn was followed by those of AT&T, T-Mobile, and Sprint accordingly. Interestingly, while T-Mobile has the highest response rate (measured by the number of tweets generated by the firm divided by the number of tweets generated by customers in the care platform), it ranks low in terms of the effectiveness of care effort. This indicates that good care quality comes from carefully evaluating customer requests and responding to them in an appropriate manner. Such quality care may require a well-designed ticket generation system that would accurately detect the tweets that require further follow-up as well as a good customer service team that would be able to resolve the issues once the tickets are generated. Merely sending out automated tweets is not sufficient for achieving good quality care.

With respect to the expectation of customer care quality, our findings suggest that AT&T and Verizon's customers expect better care than those of Sprint and T-Mobile. This finding perfectly aligns with the prices that these firms charge (for comparable plans), with Verizon and AT&T being the most expensive followed by Sprint and T-Mobile. With respect to the net effort exerted by the firms, we find that Verizon has the best rank, followed by AT&T, T-Mobile and Sprint, in that order.

Overall, the proposed SDE model prescribes a dynamic response strategy that allows the firm to effectively react to external events that might have a significant spillover on customer sentiment. Using real events that occurred during the time period of our study, we revealed that an event that is perceived positive by the customers would lower their expectations of care quality. On the other hand, an event that is perceived as negative by the customers (such as a price hike) would increase their expectation of care quality. Such analysis enables the firms to tailor their response efforts in anticipation of potentially influential events, such as a marketing campaign, a new product release or even a data security breach. To our knowledge, quantifying the spillover effects of seemingly unrelated events on digital customer care management has not been rigorously studied in the literature.

Our study is not without limitations. We did not obtain closed-form results for the optimal response strategy. While the current formulation of the Stochastic Control Model does not lend itself to a closed-form solution, it is possible that other reasonable formulations could yield an analytically tractable form. We also did not consider any noise in the specification of the control problem. A noisy formulation would consider the possibility that a firm does not act optimally. This could happen for a variety of reasons. A firm may be acting optimally, but could be using signals (state variables) that the researcher is unaware of. Alternatively, the firm may be committing optimization errors. This could explain the differences between the optimal actions predicted by the model and the firm's actual actions. Including noise in model specification can lead to a more comprehensive treatment of the problem being addressed in this study.

References

Archak, Nikolay, Anindya Ghose, Panagiotis G. Ipeirotis. 2011. Deriving the pricing power of product features by mining consumer reviews. *Management Science* **57**(8) 1485–1509.

Bianchi, Raffaella, Davide Schiavotto, Daniel Svoboda. 2014. *Why Companies Should Care about E-care*. URL `www.mckinsey.com/business-functions/marketing-and-sales/our-insights/why-companies-should-care-about-ecare`.

Black, Fischer, Myron Scholes. 1973. The pricing of options and corporate liabilities. *The Journal of Political Economy* **3**(81) 637654.

Brown, Robert Goodell. 2004. *Smoothing, Forecasting and Prediction of Discrete Time Series*. Prentice-Hall, Englewood Cliffs, N.J.

Burke, Sarah. 2016. *Examples Of The Good, The Bad & The Ugly Of Customer Service On Social Media!*. URL `www.getspokal.com/examples-of-the-good-the-bad-the-ugly-of-customer-service-on-social-media`.

Dano, Mike. 2016. *How Verizon, AT&T, T-Mobile, Sprint and More Stacked up in Q2 2016: The top 7 Carriers*. URL `http://www.fiercewireless.com/wireless/how-verizon-at-t-t-mobile-sprint-and-more-stacked-up-q2-2016-top-7-carriers`.

Dunn, Jeff. 2016. *Mobile Carriers Are Going in on "Unlimited" Data Plans, But They're All Full of Red Flags*. URL `http://www.businessinsider.com/tmobile-unlimited-data-plan-sprint-att-verizon-2016-8/#att-4`.

Fan, Weiguo, Michael D. Gordon. 2014. The power of social media analytics. *Communications of the ACM* **6**(57) 74–81.

Forman, Chris, Anindya Ghose, Batia Wiesenfeld. 2008. Examining the relationship between reviews and sales: The role of reviewer identity disclosure in electronic markets. *Information Systems Research* **19**(3) 291–313.

Frumkin, Tamar. 2017. *The 7 Most Important Customer Service Stats for 2017*. URL `http://www.conversocial.com/blog/the-7-most-important-customer-service-stats-for-2017`.

Goh, Khim Yong, Cheng-Suang Heng, Zhijie Lin. 2013. Social media brand community and consumer behavior: Quantifying the relative impact of user- and marketer-generated content. *Information Systems Research* **24**(1) 88–107.

Gooijer, Jan G. De, Rob J. Hyndman. 2006. 25 years of time series forecasting. *International Journal of Forecasting* **22**(3) 443–473.

Gunarathne, Priyanga, Huaxia Rui, Abraham Seidmann. 2017. Whose and what social media complaints have happier resolutions? evidence from twitter. *Journal of Management Information Systems* **34**(2) 314–340.

Hutchinson, Andrew. 2017. *Consumer Expectations Rising on Social Customer Care*. URL `https://www.socialmediatoday.com/social-business/consumer-expectations-rising-social-customer-care-report/`.

Hyken, Shep. 2016. *How To Use Twitter For Customer Service*. URL `https://www.forbes.com/sites/shephyken/2016/04/30/how-to-use-twitter-for-customer-service/`.

J.D.Power. 2013. *Poor Social Media Practices Can Negatively Impact a Businesses' Bottom Line and Brand Image*. URL `https://www.sharelatex.com/project/5853150a70e2cc5e5efc80a3`.

Keane, Michael. 2010. Structural vs. atheoretic approaches to econometrics. *Journal of Econometrics* **156**(1) 3 – 20.

komando.com. 2016. *2016 Mobile Data Plans - How Do They Compare?*. URL `https://www.komando.com/tips/370499/2016-mobile-data-plans-how-do-they-compare/all`.

Lak, Parisa, Ozgur Turetken. 2017. The impact of sentiment analysis output on decision outcomes: An empirical evaluation. *AIS Transactions on Human-Computer Interaction* **9**(1) 1–22.

Lau, Raymond Y. K., Stephen S. Y. Liao, K. F. Wong, Dickson K. W. Chiu. 2012. Web 2.0 environmental scanning and adaptive decision support for business mergers and acquisitions. *MIS Quarterly* **36**(4) 1239–1268.

Lithium Technologies Inc. 2017. *Discussion Forums: Connect Product Experts and Brand Enthusiasts through Peer-to-Peer Discussions*. URL `www.lithium.com/products/online-communities/forums`.

Liu, Bing. 2015. *Sentiment Analysis: Mining Opinions, Sentiments, and Emotions*. Cambridge University Press, Cambridge, UK.

Luo, Xueming, Bin Gu, Jie Zhang, Chee Wei Phang. 2017. Expert blogs and consumer perceptions of competing brands. *MIS Quarterly* **41**(2) 371–395.

Luo, Xueming, Jie Zhang, Wenjing Duan. 2013. Social media and firm equity value. *Information Systems Research* **24**(1) 146–163.

Ma, Liye, Baohong Sun, Sunder Kekre. 2015. The squeaky wheel gets the grease-an empirical analysis of customer voice and firm intervention on twitter. *Marketing Science* **34**(5) 627–645.

Muralidhar, S., Chinmay Sardesai, Amit Morya. 2015. *Power to the People: Customer Care and Social Media*. URL `cognizant.com/whitepapers/power-to-the-people-customer-care-and-social-media-codex1226.pdf`.

Parasuraman, A., Leonard L. Berry, Valarie A. Zeithaml. 1991. Understanding customer expectations of service. *Sloan Management Review* **32**(3) 39–48.

Pitt, Leyland F., Richard T. Watson, C. Bruce Kavan. 1995. Service quality: A measure of information systems effectiveness. *MIS Quarterly* **19**(2) 173–187.

Ross, Sheldon. 2014. Variations on brownian motion. *Introduction to Probability Models (11th ed.)*. Elsevier, Amesterdam, 612–614.

Sadiku, Matthew N.O., C.N. Obiozor. 2000. A simple introduction to the method of lines. *International Journal of Electrical Engineering Education* **3**(37) 282–296.

Sethi, Suresh. 1973. Optimal control of the vidale-wolfe advertising model. *Operations Research* **21**(4) 998–1013.

Stephens, M. A. 1974. Edf statistics for goodness of fit and some comparisons. *Journal of the American Statistical Association* **69**(347) 730–737.

The Institute of Customer Service. 2013. *Handle with Care: An Analysis and Toolkit to Improve Complaint Handling*. URL `www.instituteofcustomerservice.com/research-insight/research-library/handle-with-care-an-analysis-and-toolkit-to-improve-complaint-handling`.

Twitter for Customer Service Team. 2015. *Customer Service on Twitter*. URL `cdn.cms-twdigitalassets.com/content/dam/marketing-twitter/downloads/customerservice-playbook.pdf`.

Yitzhaki, Shlomo. 2003. Ginis mean difference: A superior measure of variability for non-normal distributions. *METRON - International Journal of Statistics* **61**(2) 285–316.

zendesk.com. 2018. *Providing great social media customer service*. URL https://www.zendesk.com/resources/customer-service-through-social-media/.

Appendix A: Twitter Data Collection

To obtain all of the tweets "directed at" and "from" the telecommunications firms, we used Gnip, an official vendor of Twitter data. We created a JSON file that contained information about the data we desired to collect from Twitter. In the file, we set "fromDate" to "201606160000" and"toDate" to "201610160000". We used the following keywords as the search terms in the JSON file to collect any tweet directed at the four telecommunications firms: @TMobile OR @TMobileHelp OR #tmobile OR #t-mobile OR t-mobile OR tmobile OR @ATT OR @ATTCARES OR #ATT OR #AT&T OR AT&T OR @verizon OR #verizon OR @VerizonSupport OR #VerizonSupport OR @sprint OR @sprintcare OR #sprint OR #sprintcare

It is worth noting that the matching in Gnip is not case sensitive.[14]

Using the keywords above, we collected all of tweets "directed at" the firms. Particularly, the query returned 1,421,697 tweets "directed at" AT&T, 386,835 tweets "directed at" Verizon, 301,633 tweets "directed at" Sprint, and 1,162,252 tweets "directed at" T-Mobile (a total of 3,272,417 tweets). Given that the focus of this study is on care-related interactions between firms and customers, we only used the tweets that were within the corresponding care accounts: ATTCares, VerizonSupport, sprintcare, and TMobileHelp. This filtering process resulted in the tweets reported in Table 3.

We then used the following steps to collect all of the tweets posted "from" firms' official Twitter accounts:

from: TMobile OR from: TMobileHelp OR from: att OR from: attcares OR from: verizon OR from: verizonsupport OR from: sprintcare OR from: sprint

This resulted in collecting 84,624 tweets "from" AT&T, 58,076 tweets "from" Verizon, 96,892 tweets "from" Sprint, and 193,765 tweets "from" T-Mobile. Since our focus in this study is on the customer service-related tweets, we only used the tweets that were posted "from" firms' support accounts. This filtering process resulted in the tweets that are reported in Table 3. It is worth noting that T-Mobile's support account (@TMobileHelp) not only responded to customers' service-related tweets, but also posted a variety of public-relations tweets. Primarily, given that T-Mobile's CEO John Legere (@JohnLegere) has a very active Twitter account (he has posted over 36k tweets as of December 2017), T-Mobile's support team retweets or responds to general tweets directed at him. Therefore, we decided to remove any tweet that is directed at John Legere. We also noticed that T-Mobile often posts several tweets in response to a single tweet posted by customers. In such cases, T-Mobile uses the following format at the end of the tweet to show the sequence of the tweets that were directed at the customer: order of tweet/number of tweets in the sequence. For instance, 2/3 at the end of a tweet means that the focal tweet is the second tweet out of three tweets in response to the customer's original tweet (see Figure A.1). In case of T-Mobile, we found 57,990 tweets that were in a sequence of tweets in response to a single tweet by the customer. Sprint uses the same format. We found 4,990 tweets were in a sequence of tweets in response to a single tweet by the customer. These types of sequences of tweets were rare in AT&T and Verizon's data.

[14] For more information about the matching rules in Gnip, please refer to:support.gnip.com/enrichments/matching_rules.html

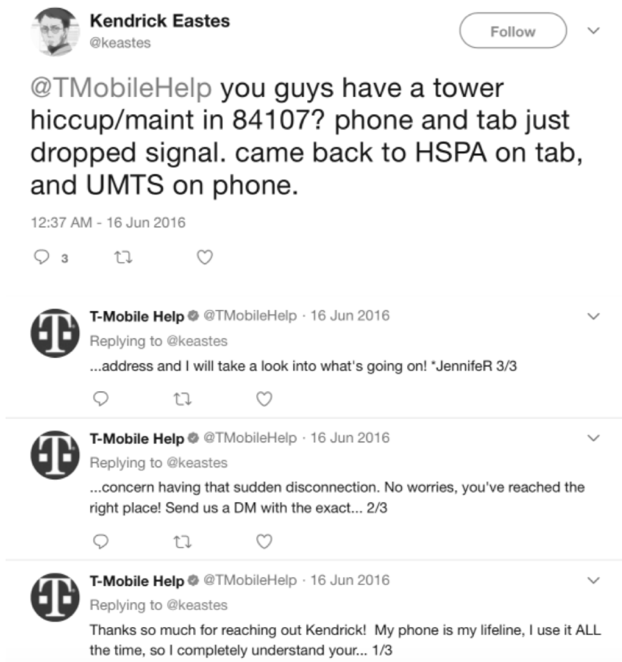

Figure A.1 A sample Sequence of Tweets Posted by T-Mobile's Support Account in Response to a Customer's Tweet

Appendix B: Sentiment Analysis

To perform sentiment analysis on the tweets, we deployed package "sentimentr" in R [15]. To evaluate the accuracy of this method, we also hired two graduate students to independently determine the sentiment of each tweet from a randomly drawn sample of 1,000 tweets. The two students allocated labels -1 for negative tweets, +1 for positive tweets, and 0 for neutral tweets. The students initially rated 100 tweets for checking the inter-rater agreement. The initial agreement was 84%, meaning that they both chose the same label for 84% of the tweets. Then, they labeled the rest of the sample of tweets resulting in 1,000 manually labeled tweets. The inter-rater agreement on the entire sample was 87%.

We then used these tweets (870 tweets that received the same label by both students) to compare with the predictions of the package sentimentr. Package sentimentr returns a sentiment score for each tweet. To compare these scores with the manual labels, we transformed the sentiment scores by using cut-off points -0.1 and 0.1. This means that a tweet with sentiment score of less than or equal to -0.1 is labeled -1 and a tweets with sentiment score of greater than or equal to 0.1 is labeled +1. The rest are labeled 0. Then we compared the labels predicted by the package with the manual labels. This resulted in 84.94% match.

Appendix C: SDE VS Structure-free Models

We fit a model and generate a forecast for the next step (hour). Then we use this forecast to generate the next forecast. We continue this process until we generate the forecasts for the entire period of the test data. Although we roll the training data one step at each iteration, we use the same model specification that

[15] https://github.com/trinker/sentimentr

we obtained by fitting the model using the training data. For instance, we do not change the values of the ARIMA parameters p, d, q in every time we roll ahead the training data. Rather, we use the same values of these parameters that we obtained by fitting the initial training data for ARIMA(p, d, q).

To test if the time series is a Random Walk, we used the Phillips-Perron Unit Root Test. The test resulted in a p-value of 0.01 that recommends that the null hypothesis be rejected, implying that the time series is not a Random Walk Process. Using KwiatkowskiPhillipsSchmidtShin (KPSS) test, we also tested if the time series is stationary. The p-value was greater than 0.05 suggesting that the time series is stationary around a deterministic trend.[16]

The first structure-free model that we used was ARIMA with the general form

$$ARIMA(p,d,q) \times (P,D,Q)_S$$

where p is non-seasonal Autoregressive (AR) order, d is non-seasonal differencing, q is non-seasonal Moving Average (MA) order, P is seasonal AR order, D is seasonal differencing, Q is seasonal MA order, and S is time span of repeating seasonal pattern. To determine the orders in ARIMA, we deployed "auto.arima" function in package "forecast" in R. We used Bayesian Information Criterion (BIC) to compare different alternatives in the ARIMA time series analysis. The function "auto.arima" conducts a search over possible models within the constraints provided. It then will return the best possible model based on the chosen criterion (BIC in our case). In terms of the constraints in "auto.arima", we set the maximum orders p and q at 24, P and Q at 5, and d and D at 2. Since our data set is hourly data, we used 24 as the value for "frequency" when we transformed the data to time series data. For all four firms, "auto.arima" returned a non-seasonal, non-differenced AR model; AR(15) for AT&T, AR(9) for Verizon, AR(15) for Sprint, and AR(11) for T-Mobile.[17]

We also ran ARCH(1) and GARCH(1,1) as additional examples of structure-free models. Autoregressive Conditionally Heteroscedastic (ARCH) and Generalized Autoregressive Conditionally Heteroscedastic (GARCH) models are widely used in financial sector to estimate the volatility of returns. In general, these models are used to describe the variance of a time series; particularly a time series with a changing and volatile variance. Figure C.1 illustrates the changing in the variance for the four time series. It appears that Sprint and T-Mobile have a higher variation in sentiment than do AT&T and Verizon.

Given the existence of volatility in the sentiment, ARCH and GARCH might better predict the changes in the sentiment. Furthermore, since according to "auto.arima" the process seems to be autoregressive, therefore we tried ARIMA(p,0,0)+GARCH(1,1) where p is the order we obtained from "auto.arima" (e.g. $p = 15$ for AT&T). Finally we tried ARIMA(p,0,0)+apARCH(1,1). Because a shock at time $t-1$ could also impact the variance at time t, the volatility is more likely to be high at time t if it was also high at time $t-1$. This yields clustering of volatility. Therefore, we used apARCH to capture the asymmetry in the volatility of the sentiment.

[16] We created a time series for each firm. The results of the tests were very similar for all four time series.

[17] We still refer to these models as ARIMA models due to the fact that we allowed the algorithm to determine the values for d (differencing) and p (MA).

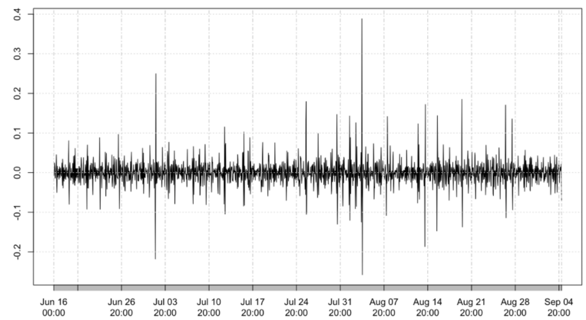

(a) AT&T

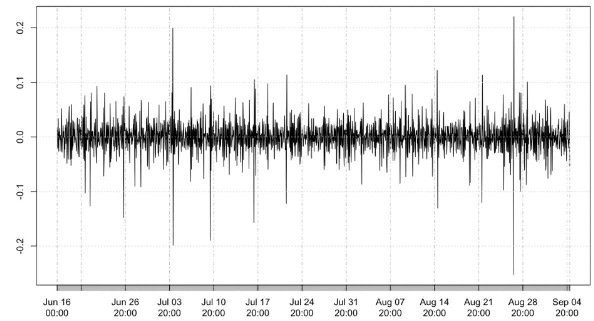

(b) Verizon

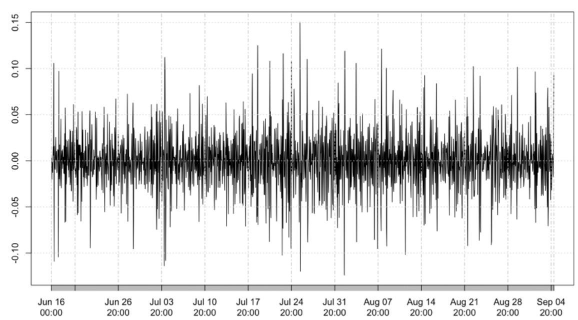

(c) Sprint

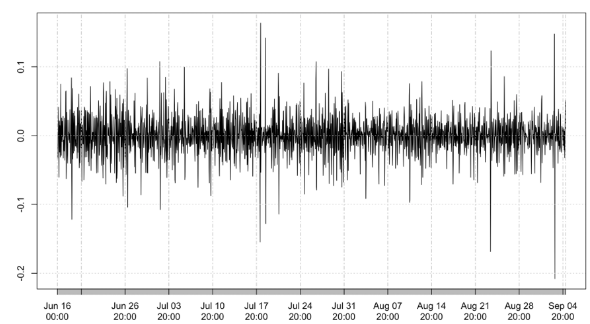

(d) T-Mobile

Figure C.1 The Change in the Variance of Customer Sentiment over Time

Appendix D: Testing the Normality Assumption

For our parameter estimates to be correct, $\Delta x_t|x_t$ needs to be drawn from a normal distribution. That is, for every given x_t in our data, Δx_t should be drawn from a normal distribution. To test this assumption, we use the Anderson-Darling test. This test has been claimed to be one of the most powerful tools for testing whether a normal distribution adequately describes a set of data (Stephens 1974). To test the normality assumption, we need to run Anderson-Darling test over Δx_t for every given x_t. Given that x_t ranges from 0 to 1 in our data set, we cannot systematically test the normality assumption for every single x_t in our data set. Therefore, we decided to find the most frequent value of x_t (mode of x_t- we call it x_M) and run the test over $\Delta x_t|x_M$. This way we would have a bigger sample size for Anderson-Darling test. Before finding x_M for each firm, we rounded the values of x off to four decimal places. x_M turned out to be 0.4510 for all four firms. To add more observations for the normality test, we also included 0.4508, 0.4509, 0.4511, and 0.4512 (0.0002 above and below 0.4510). This process resulted in sufficient sample size for all four firms for testing the normality assumption.

According to the results reported in Table 11, the p-value of the Anderson-Darling test is greater than 0.05 for all four firms, suggesting that $\Delta x_t|x_M$ is drawn from a normal distribution for all four firms.

Appendix E: Numerical Experiment to Derive α

Our objective here is to derive an accurate approximation of α with respect to our model parameters (k_1, k_2, ρ, c, and κ). To do this, we performed a numerical study by allowing these parameters to vary across a

Table 11 **Anderson-Darling Test of Normality**

Firm	Sample Size	Anderson-Darling Test Statistic	p-value
AT&T	23	0.137	0.971
Verizon	34	0.663	0.102
Sprint	17	0.516	0.163
T-Mobile	20	0.344	0.450

range and calculated the value of α for each combination of these value sets. In particular, we let k_1 to range from 0.2 to 1.8 in steps of 0.2, k_2 to range from 0.1 to 0.9 in steps of 0.1, k_3 to range from 0.01 to 0.05 in steps of 0.02, ρ to range from 0.005 to 0.021 in steps of 0.002, c to range from 0.4 to 2.0 in steps of 0.2, and finally κ to range from 0.4 to 2.0 in steps of 0.2. This resulted in 100,000 combinations of these parameters. Then we used these values to solve for $V_x(x_t, t)$ in equation 7 using Method of Lines (MOL) approach. Out of 100,000 combinations, in 56,000 cases the algorithm converged. For each one of these 56,000 combinations, we calculated the value for α. We observed that the relationships between α and k_1, k_2, κ are non-linear and decreasing. Hence, we approximated the function for α (equation 15) based on the parameters in our SDE model. We used Mean Squared Error (MSE) to evaluate the accuracy of this approximation. We found that this approximation performed well with a percentage MSE of 3.1. This process resulted in the following coefficient estimates in equation (15): $a_1 = -0.107$, $a_2 = 0.035$, $a_3 = 0.361$, $a_4 = 0.344$, and $a_5 = -5.080$.