

Device-Oriented Telecommunications Customer Call Center Demand Forecasting

by

Ashish Koul

S.B. Electrical Engineering and Computer Science, MIT, 2001
M.Eng. Electrical Engineering and Computer Science, MIT, 2003

Submitted to the MIT Sloan School of Management and the Engineering Systems Division
in Partial Fulfillment of the Requirements for the Degrees of

Master of Business Administration

and


Master of Science in Engineering Systems

In conjunction with the Leaders for Global Operations Program at the
Massachusetts Institute of Technology

June 2014

© 2014 Ashish Koul. All rights reserved.

ARCHIVES



The author hereby grants to MIT permission to reproduce and to distribute publicly paper and
electronic copies of this thesis document in whole or in part in any medium now known or
hereafter created.

Signature redacted

Signature of Author

MIT Sloan School of Management, MIT Engineering Systems Division
May 9, 2014

Signature redacted

Certified by

Stephen Graves, Thesis Supervisor
Professor of Management Science, MIT Sloan School of Management

Signature redacted

Certified by

Bruce Cameron, Thesis Supervisor
Lecturer, Engineering Systems Division

Signature redacted

Accepted by

Richard C. Larson
Mitsui Professor of Engineering Systems
Chair, Engineering Systems Division Education Committee

Signature redacted

Accepted by

Maura Herson, Director of MIT Sloan MBA Program
MIT Sloan School of Management

This page intentionally left blank.

Device-Oriented Telecommunications Customer Call Center Demand Forecasting

by

Ashish Koul

Submitted to the MIT Sloan School of Management and the MIT Systems Engineering Division on May 9, 2014 in Partial Fulfillment of the Requirements for the Degrees of Master of Business Administration and Master of Science in Engineering Systems

Abstract

Verizon Wireless maintains a call center infrastructure employing more than 15,000 customer care representatives across the United States. The current resource management process requires a lead time of several months to hire and train new staff for the customer rep position. To ensure that call center resources are balanced with customer demand, an accurate forecast of incoming call volume is required months in advance.

The standard forecasting method used at Verizon relies on an analysis of aggregate call volume. By analyzing the growth trend of the customer base and the month-upon-month seasonal fluctuations within each year, the total incoming call volume is predicted several months in advance. While this method can yield solid results, it essentially groups all customers into a single category and assumes uniform customer behavior. Given the size of the Verizon customer base, forecast inaccuracy in the current process can lead to resource allocation errors on the order of tens of thousands of labor hours per month.

This thesis proposes a call forecasting model which segments customers according to wireless device type. By taking into consideration customer behavior on a per device basis and accounting for the continuous churn in mobile devices, there is the potential to create a forecasting tool with better accuracy. For each device model, future call volumes are estimated based upon projected device sales and observed customer behavior. Aggregate call volume is determined as the sum across all device models. Linear regression methods are employed to construct forecast models for each of the top 20 mobile devices (those which generate the most customer service calls) using historical device data.

The aggregate call volume forecast for these top 20 devices is benchmarked against the standard forecast currently in use at Verizon to validate the reliability of the new approach. Furthermore, device-oriented analytics processes developed for this project will enable Verizon to build a rich library of device data without additional staff or resource investments. By incorporating device-oriented data analysis into the call volume forecasting process, Verizon Wireless hopes to improve forecast accuracy and staff planning, effectively maintaining service levels while reducing overall staffing costs at call centers.

Thesis Supervisor: Stephen Graves

Title: Professor of Management Science, MIT Sloan School of Management

Thesis Supervisor: Bruce Cameron

Title: Lecturer, Engineering Systems Division

This page intentionally left blank.

Acknowledgments

This project was completed under the auspices of both the Leaders for Global Operations (LGO) program at the Massachusetts Institute of Technology and the Verizon Corporation. The research would not have been possible without generous support from these two organizations.

At Verizon there are several individuals whom I wish to acknowledge. First, I would like to offer thanks to my project supervisors at Verizon, Christine Wong and Michael O'Brien. They were both tireless advocates for the goals of the research and guided my work during my time at Verizon. Furthermore, I would like to thank Rudy Cardenas, the director in charge of Verizon Wireless call center resource planning, for providing me with the opportunity and resources to work on this demand forecasting project. Lastly, I would like to acknowledge my colleague Christi Marret for helping me to navigate the information technology infrastructure at Verizon.

I would also like to thank my MIT thesis advisors, Stephen Graves and Bruce Cameron, for their valuable input over the course of the project. Additionally my classmate in the LGO program, Yanai Golany, offered advice which was of help in refining some of my research ideas.

Lastly, I'd like to thank my fiancée Deepali for her willingness to deal with my awkward and shifting work schedule.

This page intentionally left blank.

Table of Contents

Abstract	3
Acknowledgments	5
Table of Contents	7
List of Figures	9
1 Introduction	11
1.1 Verizon Wireless	11
1.2 Customer Service Overview	12
1.2.1 Call Center Organization	12
1.2.2 Call Center Staff Management	13
1.2.3 Call Volume Forecasting Process.....	15
1.3 Device-Oriented Call Volume Forecasting	17
2 Device Call Volume Data Analysis	20
2.1 Data Collection	20
2.1.1 Call Volume and Device Data Recording	20
2.1.2 Abandonments and Unidentified Calls	21
2.2 Call Volume Trend Analysis	22
2.2.1 Device Activation Age Factor	22
2.2.2 Transient Factor	25
2.2.3 Call Volume Seasonality Factor.....	27
3 Forecast Model Proposal and Calibration	30
3.1.1 Proposed Model Structure	30
3.1.2 Model Fitting and Calibration	33
3.2 Potential Model Performance	39
4 Results	42
4.1 Device Sales Forecast Data	42
4.2 Forecast Model Performance	45
5 Discussion	47
5.1 Further Development	47
5.2 Implementation and Deployment	48
5.3 Potential Value for Customer Service Processes	50
6 Bibliography	53

This page intentionally left blank.

List of Figures

Figure 1: Call Center Staffing Decision Lead Time	14
Figure 2: Current Forecasting Process Components	16
Figure 3: Current Forecast Method Accuracy	16
Figure 4: Call Volume for Representative VzW Device, Beginning with Launch	18
Figure 5: Device-Oriented Aggregate Call Volume Breakdown	18
Figure 6: Weekly CIR_D for Representative VzW Devices	23
Figure 7: Call Volume for Devices Activated in Two-Week Time Frame	24
Figure 8: Active Base for Representative VzW Device	24
Figure 9: Transient Events, Software Updates	25
Figure 10: Transient Events, National Holidays	26
Figure 11: Seasonality in Aggregate Call Volume	27
Figure 12: Flagship Device Activation Seasonality	28
Figure 13: Seasonality in Device CIR_D	29
Figure 14: Calling Behavior Model for New Activations	31
Figure 15: Forecast Average Prediction % Error with Varying Parameters	35
Figure 16: Accounting for Transient Events	36
Figure 17: Example of Complete Forecast for Sample VzW Device	39
Figure 18: Aggregate Call Volume Forecast Potential Performance	40
Figure 19: Pareto Chart of Top 20 VzW Device Models	41
Figure 20: Weekly Active Base for VzW Device Models	43
Figure 21: Complete Forecasting Process Steps	45
Figure 22: Aggregate Call Volume Forecast Performance	46
Figure 23: Deliverables for VzW	49
Figure 24: Device-Specific Data Results	51

List of Equations

Equation 1: Initial Model	31
Equation 2: Incorporating Transients	32
Equation 3: Complete Model	33

This page intentionally left blank.

1 Introduction

This thesis report is built upon the results of a six month research project conducted within the Supply Chain group at Verizon Wireless corporate headquarters in Basking Ridge, New Jersey between June and December of 2013. The project was done in collaboration with the Verizon Wireless Customer Care group with the intention of improving call center management processes.

This chapter details the background and objectives for this research project.

1.1 Verizon Wireless

Verizon Wireless (VzW) is currently the largest provider of cellular communications services in the United States, serving a customer base of approximately one-third of the United States population. The company was originally created as a joint venture between Verizon Communications and Vodafone, but after a fourteen year partnership between the two firms, Vodafone agreed to sell its stake in the business to Verizon Communications for a total of \$130 billion in September of 2013. This marked the third largest corporate acquisition in history.

Verizon Wireless offers cellular voice and data services across the United States, operating the largest deployed 4G LTE (Long-Term Evolution) wireless network in the world. Alongside its service offerings, the company sells an array of wireless devices from many major electronics vendors including Apple, Samsung, Motorola, LG and HTC. Devices offered by Verizon Wireless tend to be among the most sophisticated handsets available, including many of the latest smartphones.

1.2 Customer Service Overview

To provide live support to its customers, Verizon Wireless manages a call center infrastructure employing more than 15000 customer care representatives at more than 20 sites across the United States. These representatives provide telephone support to customers addressing billing issues, technical problems and other general questions. Given the size of the Verizon customer base, the call centers handle a volume of well over one hundred million calls every year.

1.2.1 Call Center Organization

To address the diversity of customer service issues which may come up, the call center workforce is organized into a number of distinct work groups, each with primary responsibility for a specific subset of service calls. The major work groups are:

1. *Care* – Representatives are responsible for dealing with customer billing issues, basic tech support and general questions.
2. *Care Bilingual* – Representatives have the same responsibilities as normal Care reps, but cater to Spanish speaking customers.
3. *Tech* – Representatives are responsible for advanced tech support issues and address problems which normal Care reps are unable to resolve.

When a customer service call is placed, initially the customer is greeted by an automated response system that is intended to determine the nature of the call. The system first attempts to determine the identity of the caller. Given that the customer service call center network is intended to provide support to current Verizon Wireless customers, it is assumed that all callers have an account. For potential new customers looking to ask questions, there is an independent sales call center system.

If the customer calls from a Verizon Wireless line, the system will automatically recognize the phone and retrieve the customer account information from an internal database. For devices which cannot be recognized, the customer is asked by the system to provide the phone number associated with the account in question.

Once the account information has been established, the customer is asked a series of questions regarding their reason for calling. This information is used to determine the work group to which the call will be delivered for live support. The large majority of calls (more than 80%) are directed to the Care group. The Bilingual and Tech support groups are more specialized and have far less employees than the Care group.

1.2.2 Call Center Staff Management

The objective of call center staff management at Verizon Wireless customer service is first to ensure that call centers have adequate staff to guarantee an acceptable service level. Verizon has consistently earned a high rating amongst cellular service providers for customer service according to J.D. Power and Associates, and maintaining this reputation is a continuing corporate objective. At the same time, Verizon leadership seeks to minimize costs associated with support. Thus the staff planning process seeks to balance service level with service cost by ensuring that resources are deployed judiciously.

The first mechanism applied to meet customer demand is the management of shift timing. Over the course of each day, incoming call volume levels are variable. In general, call center demand starts off slow in the morning, grows to a maximum by mid-afternoon and declines into the evening. To meet this demand, the call center resource management team at Verizon collects statistics regarding daily call volume to determine the patterns in customer

demand. This data is used to set working shifts for the call center representatives such that the available staff is proportional to the expected demand. The Verizon customer service system provides 24-hour phone support, though not all call centers operate on a 24-hour schedule.

The second mechanism applied to meet customer demand is the management of call center staffing levels. The number of representatives employed by the call center system must be sufficient given the gross level of demand. This can be a difficult issue to contend with given the requirements of the job. To be an effective call center representative, one must have solid communication skills and be able to quickly learn the technical details of Verizon service plans and mobile devices. Furthermore, one must be able to handle frustrated and sometimes combative customers. Due to the demanding nature of these jobs, adding staff to a call center takes time. The current process requires a lead time of several months to hire and train new staff for the customer rep position. (See Figure 1 Below)




Figure 1: Call Center Staffing Decision Lead Time

The goal of the staff resource planning team is to ensure that the hiring decisions made today will yield the appropriate level of staffing in the ensuing months. As a result, to ensure that call center resources are balanced with customer demand, an accurate projection of incoming call volume is required months in advance. The methods used for the demand projection are the focus of this research project.

1.2.3 Call Volume Forecasting Process

Due to the rapid pace of change in the wireless telecommunications industry, demand for live customer support can be difficult to predict. The current call forecasting process seeks to predict the aggregate number of incoming calls on a monthly basis. Towards the end of every calendar year, the Call Center Resource Planning Team creates a twelve month forecast, providing an estimate of incoming call volume for each month of the coming year for each of the major call center work groups (*Care, Tech, Bilingual Care*). This long-term forecast yields an initial estimate of the resources that will be required within the call center system – It is used primarily as a financial and accounting guideline by Verizon Wireless management.

Apart from the twelve month forecast, the team also generates a short-term four month forecast of incoming call volume per month, updated every two weeks with the latest data. This short-term forecast is considerably more accurate and is used for operational decisions. Given the expected level of demand four months out, staff hiring decisions are made to ensure adequate capacity. Improving on this short-term forecast is the focus of this research project.

The forecasting method used by the resource planning team is largely based upon a single metric known as the monthly Call-In-Rate (*CIR*), defined below:

$$CIR = \frac{\text{Incoming Calls (one month)}}{\text{Active Devices}}$$

Essentially, *CIR* is the number of calls recorded in a month for each active device (an active device can be a phone, tablet or wireless hotspot device). Note that the denominator of this expression, number of active devices, is not the same as the number of customers, given that a single customer may have multiple active devices. The *CIR* metric is recorded every month and analyzed to determine gross trends as well as seasonal trends. (See Figure 2 Below)

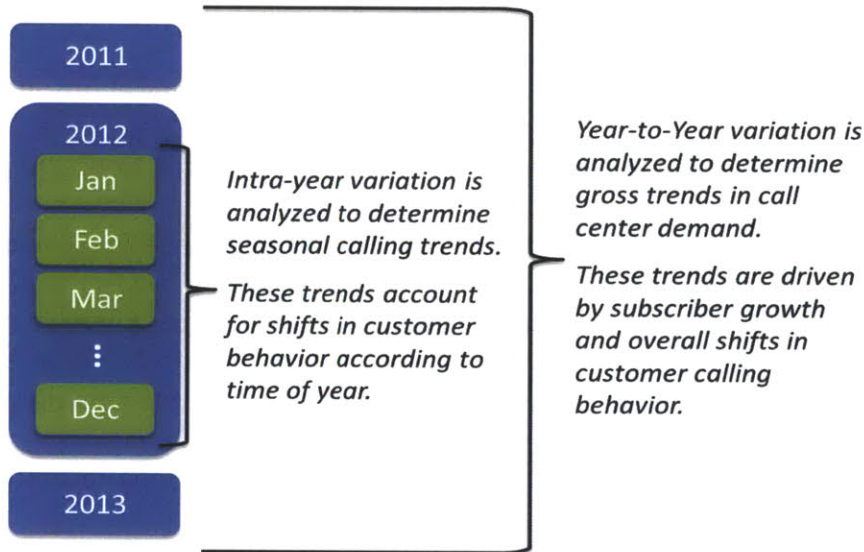


Figure 2: Current Forecasting Process Components

This analysis of trends is used to project an expected CIR level for the subsequent months. In addition, the Verizon Wireless marketing division generates a forecast for the number of active lines of service. Using these forecast values, the expected monthly call volume for each of the coming months is calculated as follows:

$$\text{Forecast Calls (one month)} = \text{Forecast CIR} * \text{Forecast Active Lines of Service}$$

The absolute percent error of the current forecasting method in the case of the short-term four month projection is summarized in the table below. Approximately every two weeks,

Monthly Forecast Error

Bi-Monthly Forecast Date	Feb	Mar	Apr	May	Jun	Jul	Aug	Average
	01/22 Re-Forecast	0.02479	0.00288	0.04090	0.03159			
02/08 Re-Forecast	0.02479	0.04026	0.07893	0.06471				0.05217
02/26 Re-Forecast		0.00478	0.04283	0.03017	0.02383			0.02540
03/11 Re-Forecast		0.00478	0.04726	0.03487	0.02949			0.02910
03/25 Re-Forecast			0.04726	0.02508	0.02456	0.08136		0.04457
04/09 Re-Forecast			0.04726	0.01316	0.01264	0.04880		0.03047
04/23 Re-Forecast				0.01316	0.01264	0.04880	0.02813	0.02568
05/09 Re-Forecast				0.01316	0.01087	0.03716	0.09783	0.03976

Figure 3: Current Forecast Method Accuracy

a new four month projection is generated, and the table shows the percent error for each of the monthly projections.

The average error rate seen in the above table is 3.5%, and the historical average of forecast accuracy for the current process is approximately 4%. This may seem quite good, but in the context of Verizon call center operations, it still leaves much to be desired. Given the sheer volume of the Verizon Wireless customer base, an error rate of 3.5% can yield a substantial impact on customer service costs. Assuming approximately 10 million incoming calls per month and an average customer service call length of 10 minutes, this error equates to more than 58 thousand man hours of work on the part of customer service representatives. Further assuming an average employee wage rate of 20 dollars per hour, this equates to more than 1.1 million dollars per month. Thus, any improvement in forecast accuracy could yield significant value for the Verizon Wireless Customer Service organization.

1.3 Device-Oriented Call Volume Forecasting

The current forecasting method used at Verizon relies on an analysis of aggregate call volume. By taking into account growth projections of total customer base and projections of monthly Call-In-Rate, the total incoming call volume is predicted several months in advance. While this method has been able to yield solid results, it essentially groups all customers into a single category and assumes a standard customer behavior.

This research project aims to create a call forecasting model which segments customers according to wireless device type. The sometimes rapid shifts in customer device mix, as customers swap out old phones for the latest products, can dramatically influence calling patterns. In the figure below, weekly customer service call volume for users of a popular

smartphone is given, beginning with device launch. As more devices are purchased and activated by users over time, call volume grows rapidly.




Figure 4: Call Volume for Representative VzW Device, Beginning with Launch

By taking into consideration customer behavior on a per device basis and accounting for the continuous churn in mobile devices, there is the potential to create a forecasting tool with better accuracy. For each device model, future call volume is to be estimated based upon projected device sales and previous customer behavior. Aggregate call volume can then be determined as the sum of call volume projections across all device models. (See Figure Below)

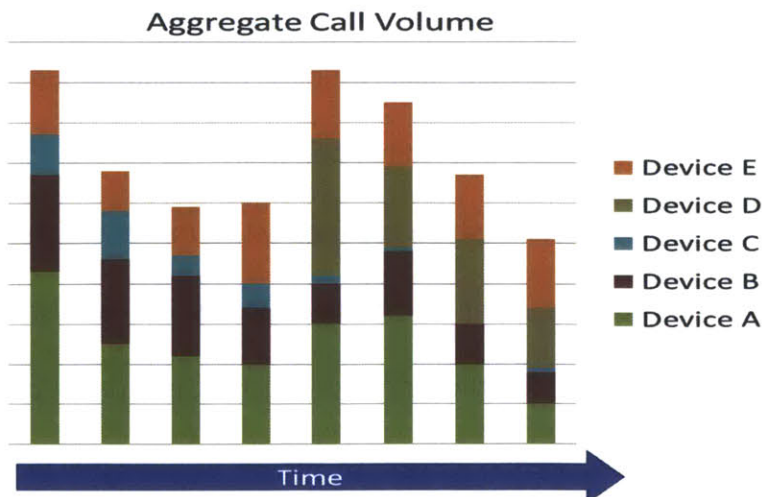


Figure 5: Device-Oriented Aggregate Call Volume Breakdown

The intended result of this project is a new call volume forecasting process for use by the Verizon Wireless Call Center Resource Planning Team. The purpose of this process is to help the team refine their short-term operational projections so as to enable staff hiring decisions which better achieve a balance between customer service level and service cost. The proposed process is designed to be integrated into the current workflow of Resource Planning Team employees. It will utilize resources which are available and straightforward to access and will not require employees to learn new tools beyond those which are standard in the enterprise (Microsoft Excel and Access).

2 Device Call Volume Data Analysis

This chapter provides an overview of the data analysis methods applied to examine device calling patterns and details key observations used in the development of the device-oriented forecasting model.

2.1 Data Collection

2.1.1 Call Volume and Device Data Recording

Verizon Wireless maintains an extensive information technology infrastructure for the call center system, recording information for every incoming call including:

1. *Customer Identity*
2. *Device Model*
3. *Time Call Received*
4. *Duration of Call*
5. *Call Type*

The *Call Type* data parameter indicates the intended call center work group (*Care, Tech, Bilingual Care*) for each call. Calling data from approximately the 16 month period from May, 2012 to August 2013 was utilized for this project.

In addition, for all active device models in the Verizon Wireless portfolio, the Supply Chain Group maintains a dataset of the weekly *active base* count per device. *Active base* count is defined as the total number of devices of a particular model in use by Verizon Wireless customers at a given time. This information was available per week over the same time frame as the calling data.

Finally, the marketing division at Verizon generates weekly sales forecasts per device for all active devices in the Verizon Wireless portfolio. This data was utilized in the final formulation for call volume forecasting.

Data was collected for those 20 device models used on the Verizon Wireless network which yielded the highest number of customer service calls as of August 2013. A device model was defined to be a product platform that is an aggregation of all SKU variations for the device such as color and storage capacity.

2.1.2 Abandonments and Unidentified Calls

In the case of call abandonments, a record is created in which unavailable parameters are set to null. When a customer calls from a Verizon Wireless phone, the customer identity and device information can be determined immediately from the source phone number. Even if the call is abandoned before a service representative answers the phone, this data will still be recorded in the database.

There are some calls for which no customer or device data can be determined. These are calls which originate from a non-Verizon device and are abandoned before the customer is asked to provide any identifying information. In addition, some calls do get routed into the customer service system from people who are not actual Verizon customers and are looking to ask general questions. These types of unidentified calls cannot be attributed to any particular device model and are thus not considered in the scope of this project. In general, it was observed that unidentified calls tend to comprise a small fixed percentage of total calls over time. Thus, such calls could be incorporated into the predictive model relatively easily.

2.2 Call Volume Trend Analysis

Using call volume and active base count data per device, the weekly Call-In-Rate was calculated for analysis of device level calling behavior:

$$CIR_D = \frac{\text{Incoming Calls for Device } D \text{ (one week)}}{\text{Active Base Count of Device } D}$$

This value equates to the number of calls received within the call center system per active device within a given week for device model D. It is similar to the *CIR* used for the current Verizon forecasting method, except that CIR_D is device specific and is calculated on a weekly basis rather than a monthly basis. This value is used extensively for device calling behavior analysis in the remainder of this report.

In the sections that follow, patterns observed from call volume data analysis are described in detail.

2.2.1 Device Activation Age Factor

Examination of CIR_D data over time for a number of Verizon devices reveals a consistent pattern of calling behavior for all devices. In the following figure, the CIR_D data for several representative devices is plotted on a weekly basis over the course of 72 weeks between the spring of 2012 and the fall of 2013. This data reveals that for any device, the customer calling rate is highest at the time of market introduction for the device. Over time, the calling rate decays to a steady state level and remains fairly consistent over time. The steady-state level of CIR_D varies by device model, but for all observed data the device calling behavior seems to follow this pattern.




Figure 6: Weekly CIR_D for Representative VzW Devices

This pattern can be attributed to the logical behavior one would expect from any customer purchasing a new device. Upon activation of a new cell phone, it is rational to believe that a customer would have a high probability of calling into customer service. The customer is not yet familiar with the device, and any technical problems or questions regarding service plans are likely to come up when the device has just been activated. This concept is apparent from Figure 7 below, which shows weekly customer call volume for a representative device, limited to customers who activated the device within a specific two-week time period.

The plot shows a burst of calls during the initial few weeks. Then as customers become more comfortable with the new phone and service issues are resolved, the call volume decays to a consistent level.

This behavior explains the decay pattern observed in the CIR_D level over time in Figure 6. Initially, when a device is launched into the market, all of the device users are essentially new

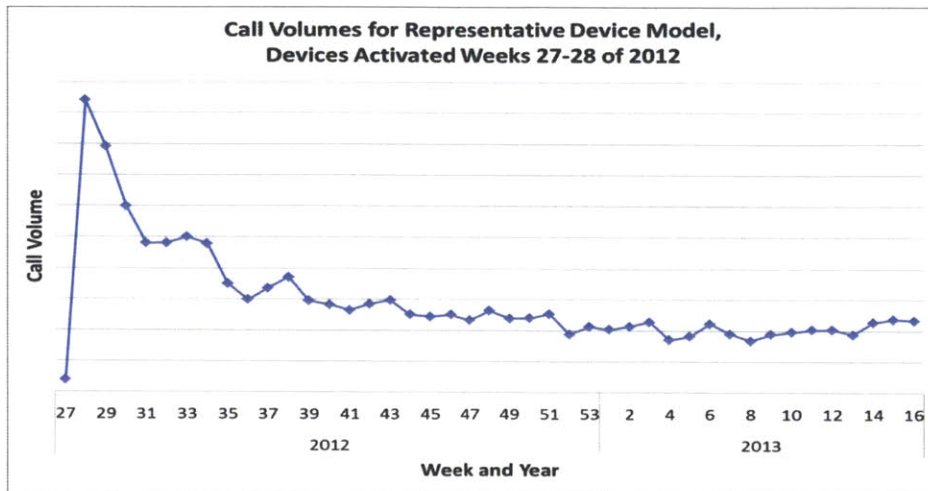


Figure 7: Call Volume for Devices Activated in Two-Week Time Frame

customers, for whom the device was just activated. In Figure 8 below, weekly active base over time for a representative device is presented, beginning with device launch. As more devices are sold into the marketplace, the proportion of experienced device users to new device users




Figure 8: Active Base for Representative VzW Device

increases. Given that the propensity of experienced users to call into customer service is low, the overall CIR_D level decreases with time as the proportion of experienced device users increases.

At a given moment in time, the active base for any device model can be divided up into those devices which were recently activated and those which have been active for some time. The differing probability of placing a call into customer service for these groups forms the basis of the call forecasting model, to be discussed in a subsequent section.

2.2.2 Transient Factor

In addition to the device activation factor, observed CIR_D levels over time tend to exhibit short-term fluctuations not related to device activation. These fluctuations are driven by transient market events which either stimulate or depress call center service demand for a brief period of time. In the plot below, the CIR_D levels for a subset of Verizon Wireless devices are shown to demonstrate this transient behavior.




Figure 9: Transient Events, Software Updates

In the case of the above plot, the transient behavior is driven by device software updates. For most modern devices, both basic and smartphone devices, manufacturers have the ability to deliver software updates to add new features or correct problems with devices

currently in use. Software updates affect the operation of devices, typically upgrading device functionality but also potentially creating new problems. For many devices, software updates yield a noticeable transient shift in CIR_D level.

In addition to software updates, transient events which can affect calling behavior include national holidays, wireless network outages and severe weather events. In the figure below, the transient effect of holidays is visible. The plot displays a weekly aggregate CIR level for the top twenty devices on the Verizon network (as of August 2013). In the weeks containing a major national holiday, the CIR level dips momentarily due to reduced demand for customer service. In addition, the week following a holiday tends to exhibit a slight increase in CIR level, due to customers delaying their service requests till after the holiday.




Figure 10: Transient Events, National Holidays

For any forecasting model to be accurate, the effect of transient events must be taken into account. In the case of holidays and software updates, prior calling behavior can be used as a model for prediction. Since the schedules for both holidays and software updates are generally known in advance, past impacts can be used to predict future impacts. However,

both network outages and severe weather events are not predictable. Without advance knowledge of when such events may take place, it is difficult to account for them in any forecast.

For the forecast model implemented for this project, only predictable transient events (holidays and software updates) are taken into account.

2.2.3 Call Volume Seasonality Factor

Incoming call volume to the call center system exhibits a noticeable seasonality, with the call volume peaking in the latter half of each calendar year and then decreasing into the first half of the following year. The figure below shows the seasonality pattern in aggregate call volume over the course of three years.




Figure 11: Seasonality in Aggregate Call Volume

This behavior is to be expected given the routine seasonal device activation pattern for Verizon Wireless and the wireless industry in general. Verizon tends to release the majority of its flagship devices into the marketplace in the latter half of the calendar year to coincide with the increase in consumer spending into the holiday season. The figure below plots the sum of

weekly activations of flagship device models on the Verizon network in a representative year. (Flagship devices are those which are heavily marketed and attract the most customers. Such devices constitute a significant proportion of total active device count.)

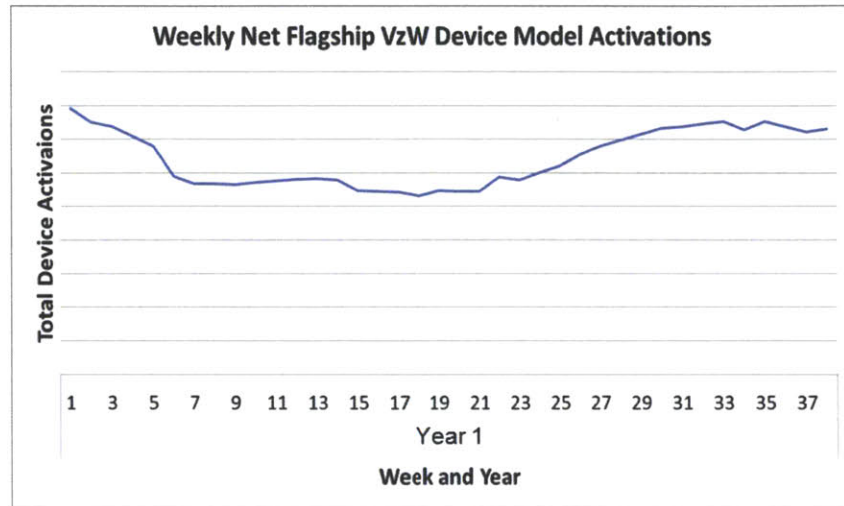


Figure 12: Flagship Device Activation Seasonality

Given that new device users have a higher propensity to call for customer service, the increase in new activations is accompanied by an increase in call volume. Thus, the seasonal behavior of gross call volume should primarily be explained by the device activation age factor, as previously described in section 2.2.1. However, when examining the CIR_D level for several device models, there appears to be a slight seasonal increase in call volumes for some individual device models as well. (See Figure Below)

This seasonality contradicts the device activation age factor, given that the long term CIR_D level does not simply converge to a static level, but rather seems to be influenced by the seasonality of gross device activations on the Verizon Wireless network. The devices which exhibit this behavior tend to be relatively older models which have been on the market for at least a full year. Most users of such devices have been using their devices for a long time, and




Figure 13: Seasonality in Device CIR_D

perhaps may be considering upgrading to a newer model. Furthermore, given that the bulk of device activations take place in the latter half of the calendar year, many customers are coming to the end of their contracts within this time period and are able to upgrade at reduced prices. Given that the long-term CIR_D levels tend to move with gross device activation levels, it is assumed that these users are calling into customer service to inquire about upgrading to the latest model devices. Such behavior would explain the observed slight seasonal shifts. This idea of seasonality influencing long term CIR_D is the final component of the forecasting model developed for this project.

3 Forecast Model Proposal and Calibration

The following chapter describes the model formulation based upon previously described patterns discovered through data analysis. All data for call volume analysis was available on a per week basis, so the forecasting model was designed to forecast weekly call volume per device.

The model structure is presented in steps, with each step expanding the model to account for each of the three factors outlined in the previous chapter. After the model structure has been presented, the model calibration and fitting process is described.

3.1.1 Proposed Model Structure

The basic model formulation is inspired by mobile device return forecasting techniques developed for Verizon Wireless [1]. In prior research, the return rate for wireless devices was shown to be driven primarily by date of purchase. The warranty policy for Verizon cell phones covers the cost of replacement for a fixed period (generally one year from purchase). Thus, the device return rate was found to be relatively constant during the warranty coverage period, after which it would immediately fall to a minimal level. For this research project, the equivalent notion of customer service demand being driven by device age is applied in the forecasting model.

Given the device activation age factor described in section 2.2.1, the following initial model is proposed to describe user calling behavior. (See Figure Below) At the time of device activation, the propensity to call for customer service is greatest. After some time it decreases to a mid-level, and ultimately it settles to a steady-state level. These three temporal periods are represented by windows A, B and C in the below diagram.

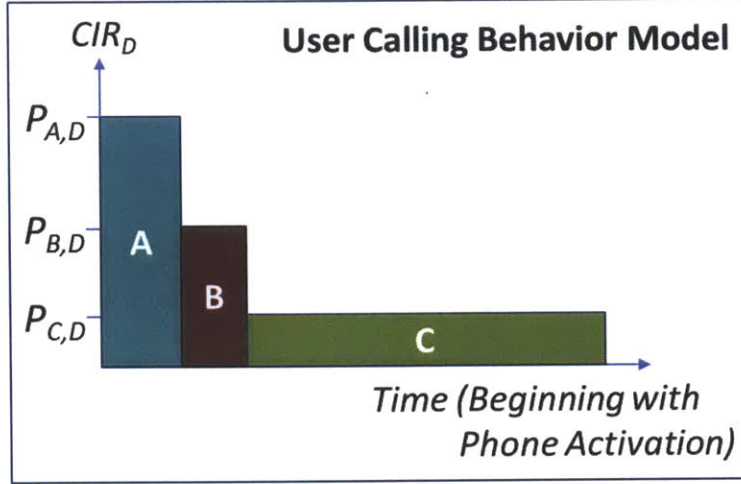


Figure 14: Calling Behavior Model for New Activations

Given this model, when assessing user calling behavior for a given week, active devices can be split into three groups – new activations (activated within window A), recent activations (window B) and mature activations (window C). Thus, over the lifecycle of a device model D , for a given week n the incoming call volume $C_D(n)$ for the device is initially calculated as follows.

$$C_D(n) = P_{A,D}S_{A,D}(n) + P_{B,D}S_{B,D}(n) + P_{C,D}S_{C,D}(n)$$

Equation 1: Initial Model

In this expression, $P_{A,D}$, $P_{B,D}$ and $P_{C,D}$ are the number of calls received per device per week for device model D for users in temporal groups A , B and C , respectively. The terms $S_{A,D}(n)$, $S_{B,D}(n)$ and $S_{C,D}(n)$ denote the number of active devices in temporal groups A , B and C , respectively. Thus, for each device model, the initial calculation of incoming calls for a week is accomplished by dividing up users according to device activation date and applying the above linear expression.

Note that in the context of this study, the only available statistic regarding device activation is the weekly active base count per device. Thus, the number of device activations which takes place in a week is assumed to be the increase in active base over that week. In

reality this is not the case, as device de-activations will also take place, as in the following expression.

$$\begin{aligned} \text{Active Base Current Week} &= \text{Active Base Prior Week} \\ &+ \text{Activations in Prior Week} \\ &- \text{Deactivations in Prior Week} \end{aligned}$$

Due to the nature of the available data, for simplicity in this study the de-activation count is effectively ignored. Implications of this choice will be addressed in the discussion chapter. For parameters $S_{A,D}(n)$ and $S_{B,D}(n)$, the values computed for analysis will be equivalent to the active base growth within windows A and B, respectively. In the case where a device active base is in decline, as for an older device model, the values of $S_{A,D}(n)$ and $S_{B,D}(n)$ will be set to zero.

At this point the proposed model must be expanded to account for transient events. In the context of this project, both major holidays and device software updates are included as transient events. These events are assumed to affect all users of a device model, irrespective of device activation age. Thus, to include transients, Equation 1 above is expanded as follows.

$$\begin{aligned} C_D(n) &= P_{A,D}S_{A,D}(n) + P_{B,D}S_{B,D}(n) + P_{C,D}S_{C,D}(n) \\ &+ T_D(n)S_{Total,D}(n) \end{aligned}$$

Equation 2: Incorporating Transients

In this expression, $T_D(n)$ is the transient function, which reflects the total transient impact in terms of number of calls per device for week n for device model D . The term $S_{Total,D}(n)$ denotes the total active base for device model D in week n . (This term is equal to the sum of $S_{A,D}(n)$, $S_{B,D}(n)$ and $S_{C,D}(n)$) The transient function is to be constructed to account for both major holidays and software updates. This process will be discussed in a subsequent section.

The final factor to be taken into account for the forecast model is the seasonality effect. The hypothesis drawn from data observations was that this effect was the result of users of older generation devices calling into customer service to inquire about the latest devices. The observed CIR_D levels tend to increase with the seasonal increase in gross activations for new devices. Thus, the forecast model is further extended to account for this seasonal shift as follows.

$$\begin{aligned}
C_D(n) &= P_{A,D}S_{A,D}(n) + P_{B,D}S_{B,D}(n) + P_{C,D}S_{C,D}(n) \\
&+ T_D(n)S_{Total,D}(n) \\
&+ H_D G(n)S_{Total,D}(n)
\end{aligned}$$

Equation 3: Complete Model

In this expression, $G(n)$ denotes the total number of device activations of all models for week n . H_D denotes the impact of seasonal device activations on call volume for users of device D . The product term $H_D G(n)$ reflects the increase in call volume in terms of calls per device for week n for device model D . The value of H_D is distinct for each device and is expected to be greater for older device models, whose users may be looking to upgrade.

The completed model above was fit to input calling data for all devices using linear regression. The fitting process is detailed in the following section.

3.1.2 Model Fitting and Calibration

The first step required for model fitting is selecting appropriate data window lengths. The goal of the model is to construct a four month call volume forecast. Thus, calling data for all devices was divided into training and testing data, whereby a set number of weeks of calling data was used to fit model parameters, while the subsequent 16 weeks of data was used to

evaluate the model prediction accuracy. The number of weeks of calling data to use for model fitting was selected in order to yield an optimal forecast.

Given the pace of change in the mobile device marketplace, consumer behavior is likely to shift with time, and thus the model should try to capture recent information. However, the model must also capture enough information to ensure that the model fit is robust in the presence of data variability.

In addition to length of the fitting data window, the model as given in Equation 1 requires the choice of appropriate lengths for temporal windows A and B, as shown in figure 14. Window C encompasses all mature devices (those activated prior to windows A and B). Defining window lengths a and b for windows A and B respectively, the active base counts referenced in Equation 1 are defined below.

$S_{A,D}(n) = \text{Devices activated between week } (n - a) \text{ and week } n$

$S_{B,D}(n) = \text{Devices activated between week } (n - a - b) \text{ and week } (n - a)$

$S_{C,D}(n) = \text{Devices activated before week } (n - a - b)$

Along with fitting data window length, appropriate values for a and b had to be selected such that the model would have optimal predictive value. To accomplish this, an experiment was conducted to test a number of parameter combinations. Input data window length (the number of weeks of data used to fit the regression model) was varied between 20 and 40 weeks while the parameters a and b were varied between 0 and 3 weeks. For the purpose of this experiment, only the initial model given in Equation 1 was considered. Linear regression was applied to training data to determine the coefficient values $P_{A,D}$, $P_{B,D}$, $P_{C,D}$ for each device (using independent variables $S_{A,D}(n)$, $S_{B,D}(n)$ and $S_{C,D}(n)$ and dependent variable $C_D(n)$). These model coefficients were then applied to the subsequent 16 weeks of testing data to determine




Figure 16: Accounting for Transient Events

For those circumstances with noticeable transient impact, the nominal level of CIR_D was inferred by looking at levels before and after the transient and interpolating the points between. This nominal level is assumed to be the level of CIR_D which would have occurred if the software update had not taken place. The transient impact for a given software update, measured in calls per device for week n , was calculated to be the difference between CIR_D and the nominal level.

$$T_{D,Update}(n) = CIR_D(n) - Nominal_D(n) \quad [\text{For the given software update}]$$

This method is analogous to how the impact of promotions is often calculated for sales forecasting. The impact of distinct promotions can be combined by superposition to find the full impact for sales over time. Such methods have also been used to assess the impact of promotions on call center demand [2].

For those software updates without any noticeable impact, it was assumed that this was a minor software change which did not yield any customer issues. Such updates were

effectively considered as non-events and removed from the database. Transient impact for all software updates for all devices was determined in this manner.

Prominent national holidays were assumed to affect all Verizon customers. To find these impacts, rather than analyze CIR_D for all devices, the aggregate weekly CIR level for the top twenty devices (See Figure 10) was used. For the major holidays which exhibit a noticeable transient shift, the impact was measured in the same manner used for software updates. Note that in the case of holidays, transient shifts only occur for a single week, rather than several weeks as in the case of software updates. This is because software updates tend to have staggered distribution such that it takes several weeks for the update to be applied to all customers.

The final transient impact for device D for all weeks n is computed as the sum of impacts for all holidays and software updates.

$$T_D(n) = T_{Holidays}(n) + T_{D,Updates}(n)$$

This transient function can be incorporated into the forecasting model shown in Equation 2 to improve model performance. In order to use the transient data, the linear regression process must be modified. Prior to the model fitting, the training data must be adjusted to remove the impact of the transient events. Essentially, this requires subtracting call volume attributed to transient events from the full call volume $C_D(n)$.

$$C_{D,Adjusted}(n) = C_D(n) - T_D(n)S_{Total,D}(n) \text{ [For Training Data]}$$

After this correction, the linear regression is applied as before to the adjusted call volume data to determine the coefficients $P_{A,D}$, $P_{B,D}$, $P_{C,D}$ for each device. These can then be applied in linear combination to the 16 weeks of testing data to determine an initial call

volume. Finally, the 16 week call volume estimate is completed by adding the transient impact for those weeks back into the call volume.

$$C_D(n) = C_{D,Adjusted}(n) + T_D(n)S_{Total,D}(n) \text{ [For Testing Data]}$$

Accounting for transient events has the additional benefit of improving the potential accuracy of the regression coefficients, given that the adjusted data used in the linear regression is free of transient artifacts.

The final piece needed to complete device model fitting is the seasonality factor as given in Equation 3. The observed behavior revealed that there was an impact in CIR_D related to the sales cycle for devices. In the case of customers who have been using their phones for some time, the release and sale of new devices appears to boost call volume.

Given the timeframe of the available data, a subset of device models was selected to build function $G(n)$ in Equation 3.

$$G(n) = S_{1,D1}(n) + S_{1,D2}(n) + S_{1,D3}(n) + \dots$$

In the above expression, $S_{1,D1}(n)$ is the number of activations of device model $D1$ which have taken place within the 1 week prior to week n . Device models $D1$, $D2$, $D3$ and so on are selected to be the current flagship devices within the timeframe of the forecast fitting data. These are devices which will spur the most interest amongst current customers. $G(n)$ is calculated as the sum in the expression above and is incorporated into the forecasting model as given in Equation 3. This element completes the model – The plot below displays an example of the application of the complete model to sample data.




Figure 17: Example of Complete Forecast for Sample VzW Device

Note that the procedure for fitting and forecasting is the same as before. The training data must be corrected in advance to remove transient impacts, but now the regression is applied to find an additional coefficient: $P_{A,D}$, $P_{C,D}$ and H_D (Note that $P_{B,D}$ is no longer required, as the B window has been removed from the model). The parameter H_D represents the impact of seasonal sales on call volume for device D . Using the window parameters already selected, the completed model yields an average percent error of 5.2% across all device models, an improvement from the previous value of 6.7% (See Figure 15).

3.2 Potential Model Performance

Applying the completed model, the aggregate forecast performance can now be analyzed. In previous sections, the forecast model structure has been developed to independently address device models. Now, the forecast output for all of the top 20 devices is summed over the 16 week span of testing data. This yields the results in the figure below.

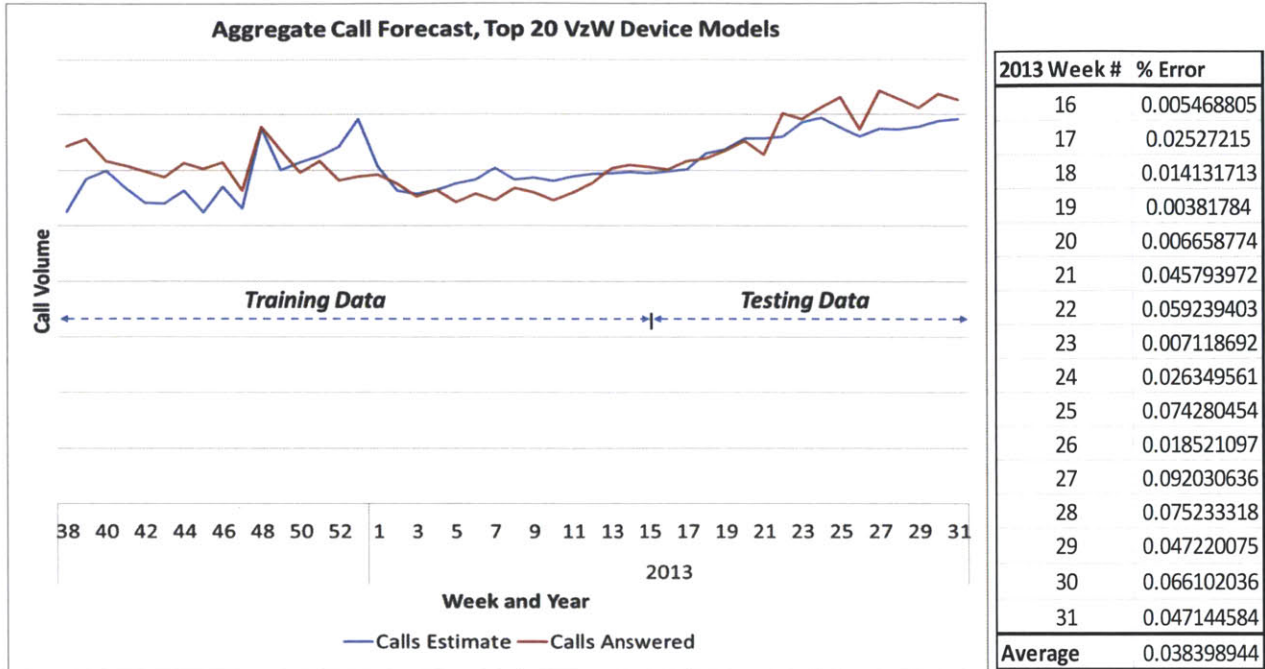


Figure 18: Aggregate Call Volume Forecast Potential Performance

The average percent error for the weekly forecast is 3.8%, and the maximum absolute error is 9.2%. The logical expectation for this aggregate forecast would be that the average percent error for the aggregate would be lower than that of the individual device models, assuming that the forecast error for all device models is mutually independent [3].

The results here do show a reduction in error, but this improvement is limited primarily because of the distribution of call volume amongst devices. While this analysis incorporates all of the top 20 Verizon Wireless device models, the actual call volume tends to be dominated by only the top few. Furthermore, there tends to be correlation in call volume in across devices in the case of software updates. For many device manufacturers, software updates for multiple smart phone models are released together. Thus, the transient impact of a software update can affect several device models at once. Any error in accounting for the update is magnified across more than one device.

The pareto chart below shows the distribution of call volume for the top 20 devices. The bulk of call volume is generated by the top few device models. Given this distribution, the error-cancelling effects that one would expect in the aggregate forecast are reduced.




Figure 19: Pareto Chart of Top 20 VzW Device Models

This performance analysis shows potential for the aggregate forecasting technique, but the last step in analysis is to assess performance in the case of true forecast device activations. Up to this point, all device calling and activation data has been drawn from the available dataset of prior values. Specifically, device activation data used to compute the independent variables $S_{A,D}(n)$, $S_{B,D}(n)$ and $S_{C,D}(n)$ and $G(n)$ (See Equation 3) for forecast testing data has been taken from the true dataset, ie., the actual data for activations. In an actual forecasting scenario, prior values would be used for training, but for the actual forecast the independent variables would themselves have to be estimated values.

This estimation is done using the device model sales forecast data, as described in the results chapter which follows.

4 Results

4.1 Device Sales Forecast Data

For the purpose of testing the live performance of the forecasting model, after fitting the model coefficients using prior data, the independent variables $S_{A,D}(n)$, $S_{B,D}(n)$ and $S_{C,D}(n)$ and $G(n)$ for the 16 week forecast period must be estimated. Active base count data for the future is of course not available, but in the previous chapter it was established that the Verizon Wireless marketing division generates a 16 week sales forecast for all active device models. This data is used as a proxy for the active base data upon which the model was constructed and evaluated.

The sales forecast figures cannot be utilized directly. As discussed in Section 3.1.1, the calculated values used for $S_{A,D}(n)$ and $S_{B,D}(n)$ are equivalent to the net increase in device active base within windows A and B, respectively. The sales figures account for all projected device sales each week via retail channels. They do not account for potential device activations which may occur when people obtain devices through non-standard channels, such as eBay or Craigslist. Furthermore, they do not account for device de-activations as do the active base statistics used to compute $S_{A,D}(n)$ and $S_{B,D}(n)$. Despite these distinctions, the projected sales figures still represent a significant portion of net active base change and are used for prediction as follows.

The 16 week sales forecast data is used to project the device active base for the same 16 weeks. This is accomplished by first making a simple observation. The plot below shows the active base count for several device models. Unlike the weekly call volume, the active base

count varies smoothly, with minimal variability from week to week. Furthermore, changes in the growth trend are gradual, with the slope of the curve shifting slowly over time.




Figure 20: Weekly Active Base for VzW Device Models

To project the active base for the subsequent 16 weeks, first a weekly shift trend R_D is established by applying a weighted moving average to the prior five weeks of available active base data for each device in the following manner. (In the equations that follow, the most recent available week is referenced as week 0, and previous weeks are referenced with negative indices.)

$$\begin{aligned}
 R_D = & \left(\frac{3}{18}\right) * (\text{Active Base Change between week } (-4) \text{ and week } (-3)) + \\
 & \left(\frac{4}{18}\right) * (\text{Active Base Change between week } (-3) \text{ and week } (-2)) + \\
 & \left(\frac{5}{18}\right) * (\text{Active Base Change between week } (-2) \text{ and week } (-1)) + \\
 & \left(\frac{6}{18}\right) * (\text{Active Base Change between week } (-1) \text{ and week } (0))
 \end{aligned}$$

This trend value serves as the baseline expected weekly shift in active base for device D over the next 16 weeks. The value is biased such that the highest weight is applied to the most

recent weekly change. This manner of moving weighted average is a common basic technique used for univariate forecasting. While more sophisticated techniques were considered [4], this method was chosen for its simplicity.

The sales forecast data is now used to project deviations from this baseline trend. This is done by taking a similar weighted average of the first four weeks of sales forecast data, with the greatest weight applied to the first week. This average value L_D is then subtracted from all sales forecast values to find the deviations $V_D(n)$.

$$V_D(n) = (\text{Week } n \text{ sales forecast for Device } D) - L_D$$

This function essentially captures how the sales forecast is deviating from the beginning trend over the course of 16 weeks. The assumption here is that the ending trend R_D of the active base count shift corresponds to the beginning trend L_D of the sales forecast. This is based upon the observation that active base for all devices varies gradually. The values of function $V_D(n)$ capture how the sales forecast is effectively varying from the expected level, and thus these values can be incorporated into the active base projection to account for changes in the expected level of sales over the next 16 weeks.

The final projection $F_D(n)$ is calculated as follows.

$$F_D(n) = F_D(0) + \sum_{k=1}^n (R_D + V_D(k))$$

$F_D(n)$ is the projected active base for week n for device D over the course of 16 weeks. (Note again that week 0 is the latest prior week, thus $F_D(0)$ is the initial active base prior to the 16 week projection.) This projection is used to calculate the model input parameters $S_{A,D}(n)$, $S_{B,D}(n)$ and $S_{C,D}(n)$ and $G(n)$ for the forecast.

4.2 Forecast Model Performance

Having established projected data parameters for all devices, model performance can now be evaluated. For the case of the projection, the true impact on transient events is estimated based upon past behavior. National holidays are assumed to have to the same impact in the future as was the case in the previous year. Software updates are assumed to have the same impact in the future as similar software updates in the past. Note that since new software updates are different each time, the impact will likely not be the same, but for the purposes of forecasting this is an appropriate assumption.

The full forecasting process is summarized in the following, with the model fitting and forecasting processes separated.

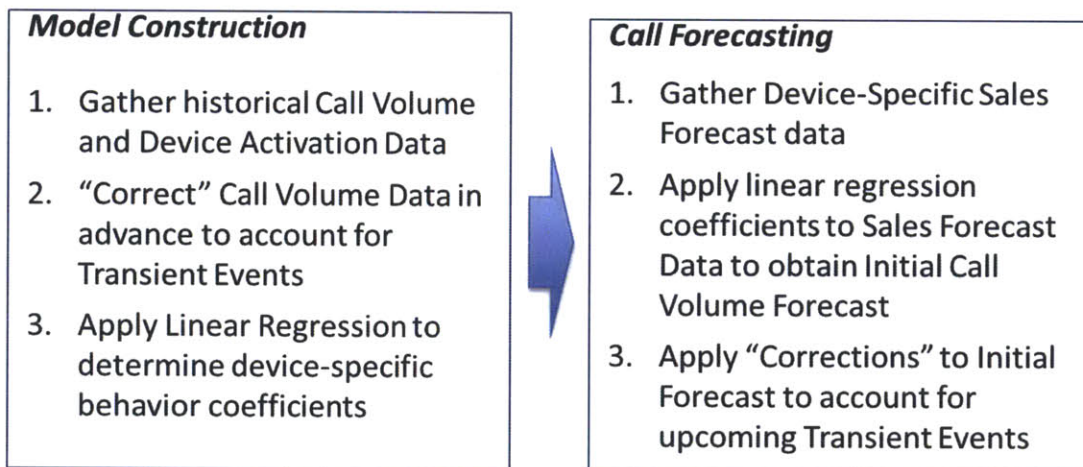


Figure 21: Complete Forecasting Process Steps

Applying this process with the latest 32 week window of device data and the 16 week device sales forecast yields the following aggregate call volume forecast for the top 20 Verizon device models. In the below diagram, weeks 4 through 35 of Year 1 are used as training data, while the weeks labeled *F0* to *F15* are forecast.




Figure 22: Aggregate Call Volume Forecast Performance

Within the 16 week forecast window, a major software update which affected several Verizon Wireless devices occurred along with a major national holiday. In the case of the software update, the forecast underestimates the call volume. As discussed above, this is a potential issue when trying to account for such events in the forecast, given that software updates are always different.

The average percent error resulting from the forecast is about 4.5%, with a maximum of 8.8%. In comparison to the historical performance of the current forecasting process in use by the Customer Service Resource Planning Team (See Section 1.2.3), this result has less accuracy. In spite of this, with further development it is likely that the device-oriented forecasting method can be improved to yield better results. This will be addressed in the discussion section.

5 Discussion

5.1 Further Development

While the current performance results for the device-oriented forecasting method do not yield tangible improvement over the current process, the potential for improvement exists with further study. The active base data utilized in the tracking of device models for this study is a reasonable metric for tracking changes in the user base for each device; however the separate details of device activations versus de-activations are not accounted for. This is a potential issue, as device de-activations are likely to generate customer service calls as well. Early in the lifecycle of a device, the level of activations is significantly higher and the de-activations can likely be ignored. However, when a device reaches maturity and the active base begins to move into decline, users will likely be calling into customer service to inquire about switching devices. The forecast attempts to capture this using the device seasonality trend parameter in the regression model, but actual deactivation statistics for the device model would likely be a much better metric.

One of the primary objectives of this study was implementation practicality. The forecasting process was to be designed for direct use by the Staff Planning Team. Using the current information technology infrastructure at Verizon Wireless, obtaining accurate detail information for each device model can be difficult. However, as processes such as device-oriented forecasting are implemented, there will likely be a stronger focus on extracting value from device-level data. Thus, it is likely in the future that building a more sophisticated model incorporating complete activation and de-activation data will become feasible.

In addition, another avenue of improvement for the forecast would be to expand the analysis to include a broader range of transient events. As of now, only software updates and holidays are accounted for because they can be predicted. The current forecasting method in use by Verizon actually attempts to account for two other types of transients: severe weather events and network outages. These events tend to affect large numbers of Verizon customers and can be incorporated for the current aggregate forecasting method. The device-oriented method developed for this project currently ignores these transients, as they are not predictable events. However, accounting for them within the fitting data will help to improve the model regression accuracy.

5.2 Implementation and Deployment

Constructing the device-oriented forecasting model requires the compiling of device data from several different sources. Weekly call volume information is recorded in a Verizon Wireless database known internally as the Enterprise Data Warehouse (EDW). Extracting the device calling data from this database required crafting a set of custom SQL scripts. Active Base count information for all devices is maintained by a specific set of analysts within the Supply Chain group in Microsoft Excel format. Device sales forecast information is maintained by another set of analysts within the Marketing Group. Finally, software update schedule information is maintained on a website on the Verizon Wireless internal network.

The full set of device data was collected into a set of Excel spreadsheets and the forecast model was implemented using the regression functionality within Microsoft Excel. Excel was utilized for simplicity, such that transferring and teaching use of the model to personnel from the Call Center Resource Planning Team would be straightforward.

The final deliverable to Verizon stakeholders was the set of custom SQL scripts for querying the database, a set of Excel files containing the functional implementation of the predictive model and a detailed guide document describing the full forecasting process. The document was created to lead an analyst step-by-step through the process. (See Figure Below)

In addition, toward the end of the six month research project, the mid-level management of the Call Center Resource Planning Team was invited to Verizon headquarters in Basking Ridge, NJ where they participated in a training session. Over the course of two and a half days, each step of the forecasting process was demonstrated live. The current forecasting model

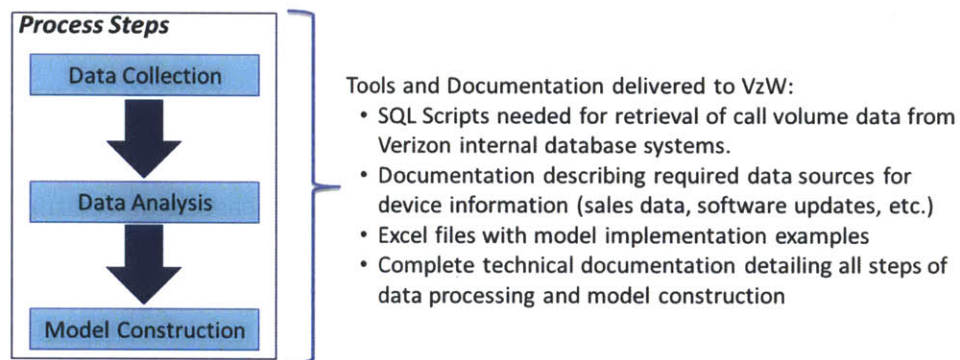


Figure 23: Deliverables for VzW

used by the team does not require data from as diverse a set of sources as the device-oriented model. Thus, during the live training session the required data sources were emphasized. Since the data sources are maintained by analysts in different groups, appropriate personal introductions were made to ensure that the Resource Planning Team would know where to go with data requests.

The ultimate goal of the live training and documentation was to ensure that analysts on the team would know where to find the information required for device-oriented forecasting and how to process that information correctly. Currently, the Resource Planning Team planning

is in the early stages of incorporating device-oriented analysis into their regular processes. They are working to build a centralized database of device calling data using Microsoft Access which should help to automate data processing and model construction. The plan is to expand the forecasting model to incorporate all active Verizon Wireless devices.

The new methodology was designed to be completed alongside the current processes without the need for additional personnel. While implementing the process will require an initial up-front investment of time and effort to establish the necessary information systems and procedures, in the long term constructing the device-oriented forecast should require little additional effort beyond the current workflow.

5.3 Potential Value for Customer Service Processes

The current processes in place for call center demand forecasting at Verizon Wireless are actually quite accurate, and in its current form the device-oriented approach does not appear to yield an improvement in forecast accuracy. However, there are still tangible benefits in incorporating the new approach. The main insight derived from the analysis of this project is that customer calling behavior is heavily influenced by device activation patterns. At the moment when a new device is activated, the customer is most likely to call customer service with questions. This likelihood decays over time to a steady state value. Propensity to call for service also tends to vary with the buying season, with customers more likely to call as new device activations increase. The quantitative relationship between device activation and demand for support can serve to guide decision making regarding operations within the larger Verizon organization.

Up to this point, the Customer Service Resource Planning Team has been relatively limited with regard to the data they incorporate into forecast processes.




Figure 24: Device-Specific Data Results

Examining data at the aggregate level hides details which may be of interest for resource planning. Incorporating a device-oriented approach provides a more precise quantitative and qualitative understanding of the nature of customer service demand. In the long term this data can potentially be used to customize call center services to better meet customer needs.

Furthermore, since the data required for device-oriented analysis comes from several distinct sources within Verizon Wireless, adopting the new process will expand the internal network of business relationships for members of the Resource Planning Team. The resulting connections will yield a bidirectional flow of information which can enhance collaboration and efficiency within Verizon Wireless.

The appeal of having a rich source of device-oriented data was apparent to the team from the onset of the project, and implementation of the new data analytics processes will certainly yield benefits in the long term.

This page intentionally left blank.

6 Bibliography

- [1] B.J. Petersen, "Reverse Supply Chain Forecasting and Decision Modeling for Improved Inventory Management," master's thesis, Eng. Systems Division, Massachusetts Inst. Of Technology, 2013.
- [2] R. Soyer and M.M. Tarimcilar, "Modeling and Analysis of Call Center Arrival Data: A Bayesian Approach," *Management Science*, vol. 54, no. 2, 2008, pp. 266-278.
- [3] G.G. Vining and S.M. Kowalski, *Statistical Methods for Engineers*, Brooks/Cole, 2011.
- [4] J.W. Taylor, "A Comparison of Univariate Time Series Methods for Forecasting Intraday Arrivals at a Call Center," *Management Science*, vol. 54, no. 2, 2008, pp. 253-265.