Reshaping Bank Branch Networks due to Mobile Banking

Francisco O. Cisternas-Vera, Alan L. Montgomery and Willem-Jan van Hoeve

September 2016
Working Draft, Do not cite

Francisco Cisternas is a doctoral candidate, Alan L. Montgomery is an Associate Professor of Marketing, and Willem van Hove is an Associate Professor of Operations Research at the Tepper School of Business, Carnegie Mellon University, 5000 Forbes Ave., Pittsburgh, PA 15213. Francisco Cisternas is the corresponding author and can be reached through email at fcisternas@cmu.edu. We would like to thank the anonymous financial service that provided the data. Francisco Cisternas also thanks Dipankar and Sharmila Chakravarti for their generous fellowship. Any opinions, findings, and conclusions or recommendations expressed in this research are those of the author(s) and not of our sponsors.

Reshaping Bank Branch Networks due to Mobile Banking

Abstract:

Mobile banking has grown from 29% in 2012 to 43% in 2015 among US bank customers and is likely to keep growing due to a further increase in the adoption of smart phones, improvement in the quality of mobile banking apps, and increasing awareness of these apps. Mobile banking is changing the way consumers interact with their banks, displacing many banking functions performed through other channels, such as automated teller machines (ATM), telephone banking, and online banking. In response banks are testing new branch formats, developing improved ATMs, and reducing the number of branches. Multi-channel management in the financial industry is critical to attract and retain customers. Using geo-coded transaction data from a large consumer bank, we develop a dynamic structural model that represents consumers' preferences for online and physical channels. Our demand model takes into account consumer banking behavior as a function of the branch network structure as well the mobile channel. We use this model to optimize the branch network in terms of capacities, amenities, location, and number of branches. Counterfactuals are constructed to evaluate potential levels of channel adoption and consider its effect on banking transactions and, more important, on customer loyalty. Our model shows that all channels remain relevant after mobile banking adoption; moreover, we find complementarity between the physical and digital channels. Our conjecture is that the importance of physical channels to banks are lessened in the presence of digital channels but is not replaced entirely. Our findings suggest that instead of reducing the number of branches, banks may want to aim to adjust current branch capacities and have physical branches specialize in those transactions that cannot be served with digital channels.

Keywords: Mobile Banking, Multi-channel Management, Dynamic Structural Model.

1. Introduction

Mobile banking has the potential to fundamentally change how, when, and where consumers bank. Mobile banking began in 2010 with specialized access to banks' web pages for mobile devices. Since then the type and number of bank services offered through the mobile channel has increased, now ranging from simple balance inquiries or personal funds transfers to mobile payments and check deposits. As a consequence, mobile banking is changing the traditional role of physical bank branches[1]. In fact, some consumers do not visit any branches, except perhaps to open their accounts. However, this does not mean that branches can be closed for several reasons. First, some consumers continue to be heavily reliant upon traditional branch services and ignore other channels. Second, some functions like cash withdrawals or accessing a safety deposit boxes cannot be provided digitally. Third, branches provide more than just banking services, i.e., branches are the closest point of interaction and are strongly valued by customers, even if they are never used. This close interaction is also valued by banks, as an industry marketing research report about US consumers (Celent, June 2013) found that "[physical] branches are the best opportunity to cultivate strong relationships such as new customer acquisition and opening of accounts."

Currently, banks themselves are introducing new technologies that automate many functions performed by branch employees. Some options for redesigning branches include adding more sophisticated ATMs, self-service tablets, and videoconferencing services, while decreasing traditional teller services. As banks encourage their consumers to migrate transactions to self-service and less costly channels, the question bank managers are asking is: "What should the branch network look like?" The answer to this question is complex because there is an interdependence between the branch network and consumer behavior. For example, if a new ATM is more conveniently located near a consumer then a weekly visit to withdraw cash at the branch may transform into two smaller withdrawals from the ATM. In turn the bank must predict demand for banking services to appropriately design the branch, so it is important for the bank to anticipate how consumer banking changes as a function of the branch network design. Our research proposes a

[1] We use the term bank branch to denote a physical brick-and-mortar location and use mobile or online to refer to virtual bank channels. For clarity we only use the term branch to refer to physical locations. Additionally we think of ATM's as a type of branch with extremely limited capabilities (typically only cash withdrawal).

solution to this problem through an analytical model of consumer behavior and leverages this model to predict how banks can optimize their existing branch networks.

Our framework addresses two primary objectives. The first objective is to construct a model of consumer attraction to a branch that predict which branches a customer is likely to visit and which services are demanded at each location. This model needs to consider both choices of branches and ATMs as well as the online and mobile bank channels. To better understand consumer behavior, we use the locations of debit and credit card transactions to predict where consumers will shop and relate this to branch location usage. For example, a customer who commutes from the suburbs to downtown may be more amenable to using a branch location downtown, whereas a retired customer in the suburbs may have a strong preference for a location near their home. A challenge of this modeling process is to understand how consumer branch usage will change over time as consumers adopt online and mobile bank services. Additionally, the bank may even take steps to encourage consumers to make these changes more quickly.

The second objective is to combine our model of consumer and business demand with operational considerations of the branches to design the optimal branch network. Operational considerations include transaction capacity, customer satisfaction, branch expertise, neighborhood potential, and competitive characteristics. The goal is to provide insight into the design of an optimal branch network based on these considerations. In particular, how should services and resources be allocated across the different branches in the network? Should the number of branches be expanded or reduced in order to reach a certain level of customer satisfaction (or any other objective)? Our conjecture is that physical branches will continue to play an important role in banking, but that their role, as well as their quantity and size, will change dramatically in the next decade. We estimate our model with anonymized data from a large US bank and construct counterfactuals that suggest that in the next decade we should observe that each branch becoming more specialized with a limited number of services. These services will vary from branch to branch and are determined by customers' preferences as predicted by models like the one presented in this paper. These models are possible only because of new sources of data, advances in econometric algorithms, and higher computing capacity.

1.1 Technological evolution of bank Industry

Modern banks can trace their origins to the rich cities of northern Italy during the early Renaissance periods14th century[2] with fractional reserve banking and banknotes appearing between the 17th and 18th centuries. The industry evolved slowly until the 20th century, when new technology fostered new types of interactions between consumers and bank. In the 1960s, the first call center and the precursors of ATMs appeared. One of the earliest call centers was created in the UK in 1964, where it was known as a "Private Automated Business Exchange" or PABX[3]. The precursor of the ATM was the Bankograph, which was installed in New York City in 1961 by the City Bank of New York[4,5]. This automated envelope deposit machine was removed after six months due to lack of customer acceptance. It is widely accepted that the first modern ATM was installed by Barclays Bank in London in 1967; since then, the technology has continued to advance with new generations featuring touch screens, video conferencing, biometrics, coin handling, scanning of individual checks without envelopes, and offering non-bank related services like dispensing movie tickets, phone cards, or traveler's checks.

Another leap in banking technology came with the advent of the Internet. The early 1980s precursors of online banking used phone lines and a keyboard to access account information; later, banks began to use the World Wide Web mainly as a way to advertise their services. In 1995, Wells Fargo was the first bank to add account services to its website, and many banks followed. Many users were reluctant to adopt online banking in the early years, with only 0.4% of households in the US using online banking at the end of 1999[6]. This number has grown to 31% in 2004, 47% in 2009[7], and 51% by 2013.

Potentially smartphones have the ability to fundamentally alter how consumers interface with banks, being always on and every present. Mobile banking was first introduced in 1999 through SMS[8], and later with

[2] Hoggson, N. F. (1926) Banking Through the Ages, New York, Dodd, Mead & Company.
[3] Science and invention in Birmingham#cite note-45
[4] Machine Accepts Cash Deposits, The New York Times, 12 April 1961
[5] "From punchcard to prestaging: 50 years of ATM innovation". ATM Marketplace. 31 July 2013. Last retrieved 15 February 2016
[6] Online Banking Report
[7] Survey by Gartner Group
[8] Short Message Service (SMS) is a text messaging service component of phone, web, or mobile communication systems

the introduction of smart phones with WAP[9] technology that allowed consumers to access web platforms. In 2010 banks began to widely introduce special client programs (apps) for smartphones, but it was not until 2013 that they truly began to take advantage of unique mobile features, like location-based services. According to the last Survey of Consumers and Financial Services in 2015 conducted by the Board of Governors of the Federal Reserve System, the ubiquity of mobile phones is changing the way consumers access financial services: 39% of all mobile phone owners and 52% of smart phone owners with bank accounts have used mobile banking in 2014 (up from 29% in 2012 and 33% in 2013). And this rapid growth is expected to continue since 11% of phone owners with bank accounts who do not currently use mobile banking expect that they will probably or definitely use it during the next year. The most common use of mobile banking is to make inquiries about account balances (94% of mobile banking users); the second and third most used services are money transfers and receiving alerts (61% and 57% of mobile banking users, respectively). Also during 2014, 51% of mobile banking users performed at least one inquiry using this channel, up from 38% during 2013. The median frequency of use of mobile banking is five times per month. The main impediments to the adoption of mobile banking are the preference for other banking methods and security concerns: 86% of consumers who do not use mobile banking believe that their banking needs are being met without the use of mobile banking, whereas 62% cite concerns about security.

1.2 Literature Review

Our research relates to other work in multichannel customer management (MCM). We follow the definition first proposed by Neslin et al. (2006): "the design, deployment, and evaluation of channels to enhance customer value through effective customer acquisition, retention, and development." In MCM, the first open question is whether firms should add more channels to the traditional ones. In the banking industry, the introduction of a mobile channel is almost compulsory for medium to large banks. Blattberg, Kim, and Neslin (2008, Chapter 25) suggests that firms should encourage multiple channel adoption if the strategy increases loyalty or marketing response, but should discourage it if adoption decreases loyalty, has no

[9] Wireless Application Protocol (WAP) is a technical standard for accessing information over a mobile wireless network

impact on marketing response, or just offers customers greater convenience without increasing the firm's share of customers' wallets.

In terms of loyalty, many studies (Kumar and Venkatesan 2005; Venkatesan, Kumar, and Ravishanker 2007) show that increasing the number of channels can yield to higher customer satisfaction, and with the rapid increase in mobile adoption, banks should add this channel to prevent customer attrition. Empirical evidence suggests that multichannel availability may enhance loyalty (Shankar, Smith, and Rangaswamy 2003; Hitt and Frei 2002; Danaher, Wilson, and Davis 2003; Wallace, Giese, and Johnson 2004), although some studies suggest that increased Internet usage may erode loyalty (Ansari, Mela, and Neslin 2008).

In our paper we describe and support three new reasons banks may adopt the mobile channel and encourage their customers to do likewise. First, we show that mobile banking and physical channels are complementary and that mobile adoption can increase bank usage, which can be beneficial for the business in all the channels. When consumers choose to use the digital channels, it releases capacity in traditional channels. This allows firms to potentially reduce capacity without affecting their service level. The capacity reduction can be achieved by eliminating services from branches that are available on digital channels, allowing branches to specialize in their unique offerings which in turn may increase their efficiency. Second, the use of mobile banking affects customer loyalty by increasing switching costs. Finally, mobile banking is still a differentiating factor, because not all banks fully support this channel yet.

The MCM literature shows that multiple channel customers are not necessarily more profitable (Kushwaha and Shankar 2013); however, a study of the banking industry (Cambra, et al 2015) shows an improvement in profit with multiple channels in cases where customers were encouraged to high-margin channels, in dual channel combinations. We find that although the mobile channel increases the frequency of interaction, at the same time it decreases the bank's opportunity to cross-sell or up-sell products because digital channels are less efficient in this respect, which in the long term may yield lower profits.

Our research focuses on understanding how customers' channel decisions are affected by the introduction of the mobile channel in the banking industry, an issue that has not been studied previously. Channel choices became popular with the introduction of online channels, but this research is usually constrained to retailing. Chintagunta, Chu, and Cebollada (2012) are the closest to our work in terms of methodology; using a hierarchical Bayes model they found significant transaction costs to purchase in-store versus online, but they did not consider the mobile channel. Laukkanen (2007) found significant differences in value perception between mobile and online banking, which is consistent with our findings. Moreover, we find that value perception is related to location and other branch characteristics when compared with safety and awareness of the digital channels.

2. Modeling Consumer Financial Transactions

We obtain anonymized data from a large US bank. Our data comprises bank transaction data for a sample of more than 500,000 accounts with more than 1.7 billion transactions across all channels for over a ninety six-month period (June 2007 and Jun 2015). For each transaction, an anonymized account identifier, date, channel, amount, and type were provided along with associated customer information. Additionally, the location associated with branches is known and the locations of some debit card transactions through the merchant's postal code can be inferred (not all transactions give postal codes). **Table 1** provides a simulated example of the raw information for a consumer. The description in this example is given for illustrative purposes and is not actual data.

Date	Description	Channel	Location	+/-	Amount	Balance
11/14/12	ATM withdrawal	ATM	15213	-	$80.00	$63.15
11/15/12	Check deposit	Branch	15213	-	$130.41	-$67.26
11/16/12	Salary from direct deposit	ACH	15213	+	$287.42	$189.16
11/17/12	Check balance	ATM	15213	o		$189.16
11/17/12	Debit card purchase at supermarket	Debit	15213	-	$97.84	$91.32

Table 1. Example of simulated transaction data for an individual customer.

In our dataset, withdrawals and deposits represent more than 98% of transactions that consumers perform at the bank, and as a consequence, these two types of transactions form the focus of our study[10]. Additionally, most consumers visit a branch about once a month on average. Specifically, we observe that almost 60% of consumers interact with the bank one time per month on average. Again, most branch visits

[10] There are other types of transactions, like safety deposit box access or opening an account with a sales' representative that can be important ones for the bank. For expositional purposes, we want to keep our model as simple as possible and, therefore, we ignore these types of events. However, our model can be extended for these other types of activities.

result in a single transaction but when more than one transaction is performed it is typically done as a single type of transaction (e.g., cash deposit and cash withdrawal).

Following the bank's nomenclature, consumers can operate with the bank through ten channels that can perform twenty types of transactions within sixteen service types. Among the transaction types, the top five account for almost 90% of the transaction types. The most popular are Inquiry and Purchase; these transactions do not use significant bank capacity. The next set of transaction types by frequency are withdrawals, deposits, and credits, which are more demanding on the bank's infrastructure. **Figure 1** summarizes the channels and their usage across the industry (Capgemini 2015). We recode the channels as those with branch interaction (Branches and ATMs) and those from digital channels (phone, Internet, Mobile and Social Media).

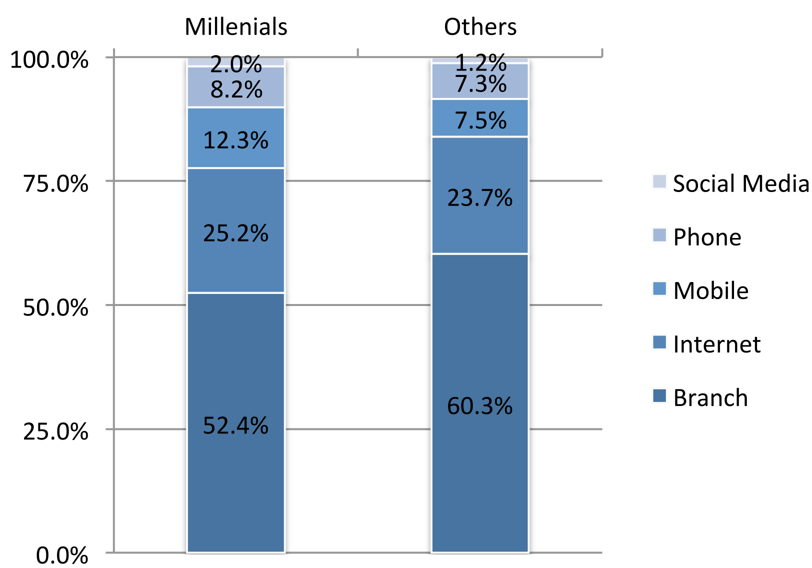

Figure 1. Channel preferences when performing current accounts and credit cards transactions.

Although this gives us an extensive set of information from which to make inferences about consumer behavior, there are important deficiencies. Clearly, consumers use cash and other payment mechanisms (e.g., credit cards from competing banks) that we do not observe. Therefore, we abstract away from our observed transactions and assume that consumers have an unobserved demand for cash that is described by a stochastic process. Additionally, consumers may receive checks at any moment, but we only observe when the checks are deposited, presumably when visiting the branch is convenient. We assume that there is an unobserved arrival of checks described by another stochastic process. Additionally, consumers experience an opportunity or holding cost when not depositing a check, which encourages them to deposit checks earlier.

Consumers are represented as rational economic agents who incur transactions costs for visiting a branch or making mobile or online transactions. These transactions costs vary by branch and time and can explain why a consumer who works close to an ATM may have a very different pattern of branch usage than one who works far away from an ATM. Every time the bank modifies its branch network by opening/closing a branch/ATM or by enabling more transaction types on its mobile channel, it is changing the transaction costs associated with the interaction, and these changes affect the frequency and intensity with which customers interact with the bank. For example, if a branch is opened in the same building where a customer works, it is likely that the consumer will visit this branch more frequently than the branch previously visited, thus decreasing the average number of transactions per trip or even the amount of cash withdrawn.

In summary, consumers make decisions in order to minimize current and future transaction and waiting costs. Therefore, these sequential decisions are the solution of a stochastic dynamic programing (SDP) model. The SDP timing is described as follows:

1. The consumers realize their needs for financial transactions (cash withdrawal and deposits).
2. Transaction costs for each alternative are realized.
3. Consumers decide whether to perform pending transactions or to wait. This decision is made by comparing the costs of waiting one more period with the transaction costs of the "cheapest" alternative.
4. The cash balance and amount of check not deposited are updated according to consumers' decisions.

The timing of each period is described in **Figure 2**.

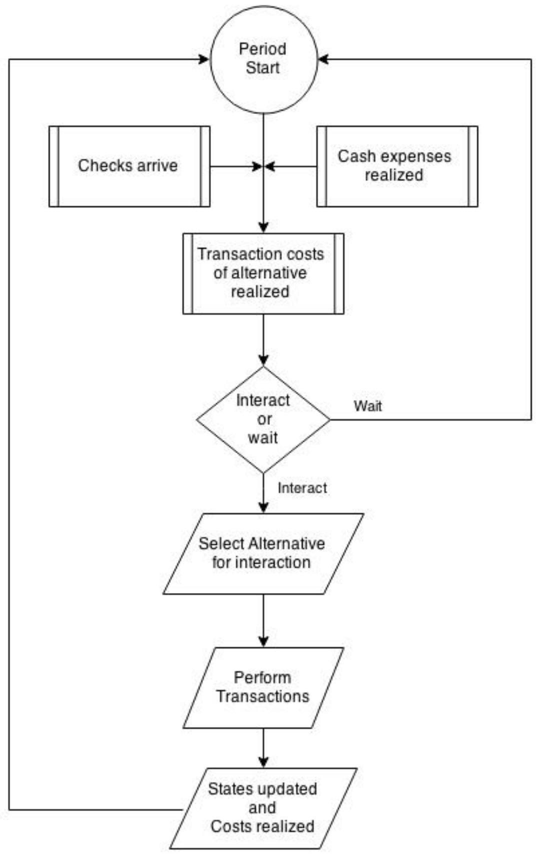

Figure 2. Model Timing

2.1 Model Specification

Our model has three main components: 1) waiting cost, which includes the costs associated with postponing interactions with the bank; 2) transaction costs, which are implicit costs associated with a consumer interacting with the bank; and 3) consideration set, where we assume that consumers consider only a subset of the available alternatives. We allow the set of alternatives to be updated over time which allows variation in the consideration set. At the end of this subsection we explain how we handle heterogeneity in the model.

2.1.1 Demand for money and its waiting costs

Cash demand. The cash demand for consumer i at period t is represented by D_{it}. This demand is not observed by the researcher, but we model it as a random term with a Poisson distribution with rate λ_{D_i}:

$$D_{it} \sim Pois\left(\lambda_{D_i}\right) \tag{1}$$

The balance of cash held by consumers is denoted as k_{it}. If a consumer does not visit the bank and make a cash withdrawal then their cash balance is $k_{it+1} = k_{it} - D_{it}$. Alternatively if consumers choose to visit the bank then they choose their cash balance as the quantity Q_{it}, which is net amount of cash after withdrawals or deposits.

Check deposits (h_{it}). We denote the dollar amount of checks at the beginning of period *t* that need to be deposited by consumer *i* as h_{it}. If a consumer visits the bank during period *t* (if a visit is made then $V_{it} = 1$ otherwise $V_{it} = 0$) then we assume the consumer deposits all checks; otherwise, the total amount of checks that are not deposited is updated by adding the checks that arrived this period:

$$h_{i,t-1} = \begin{cases} h_{it} + \sum_{s=1}^{N_{it}} A_{its} & \text{if } V_{it} = 0 \\ 0 & \text{if } V_{it} = 1 \end{cases} \tag{2}$$

N_{is} is the number of checks that arrived during the current period *t*, and A_{its} is the amount of check *s*. We do not observe the arrival of the checks, only deposits are observed. Therefore, we model these values as a compound Poisson process, where N_{it} follows Poisson arrival timing and A_{its} is uniformly distributed within the range of checks amount deposited by the consumer:

$$\begin{gathered} N_{it} \sim Pois\left(\lambda_{N_i}\right) \\ A_{its} \sim Unif\left(Am_i\right) \end{gathered} \tag{3}$$

Holding Costs. When a consumer decides to postpone their interaction with the bank for a future period, the consumer incurs a cost of not performing pending transactions. For example, not depositing a check can cause an overdraft, or not withdrawing cash might cause a suboptimal consumption. Using the notion of opportunity costs or interest rate, we use a linear representation of waiting as a function of the amounts involved. The waiting cost for consumer *i* at period *t* is described by ω_{it} and defined as:

$$\omega_{it} = \underbrace{\eta_i^h \left(k_{it} - D_{it}\right)}_{\text{Holding Costs}} \cdot 1_{\left[k_{it} - D_{it} > 0\right]} + \underbrace{\eta_i^p \left(D_{it} - k_{it}\right)}_{\text{Penalizing Costs}} \cdot 1_{\left[k_{it} - D_{it} < 0\right]} + \underbrace{\gamma_i \cdot h_{it}}_{\text{Undeposited Check Costs}} \tag{4}$$

Where η_i^h is the holding cost associated with having excess cash, perhaps due to the risk of losing the cash or the opportunity cost of not having it one deposit, and η_i^p is the penalty associated with not having the right amount of cash at a given period which might force a consumer to forgo consumption or borrow funds. The parameter associated with not depositing a check, γ_i, represents the opportunity costs or a greater risk of overdrafting associated with not depositing in the current period and is proportional to the amount of undeposited checks (h_{it}).

2.1.2 Transaction Costs

A consumer chooses a branch from a large set of branches denoted by the subscript *b*, one of the branches represents the mobile and online channel. We model this choice through an implicit transaction cost incurred by the consumer that represents the effort and time to complete the transaction. The transaction cost is influenced by characteristics of the alternative chosen and the customers' individual characteristics and preferences. In the case of the physical branches and ATMs, our exploratory analysis showed that location is a driving factor in determining how attractive a given branch location is with respect to others. The location of the branch is fixed, but customers do travel during a period. For example, consumers travel from home to work, or work to home, or from home to a shopping mall and these movements affects the branch attractiveness.

The natural logarithm of transaction costs T_{ibt} for customer i choosing alternative *b* at period *t* are:

$$Ln\left[T_{ibt}\right] = \begin{cases} -\tau_i \cdot B_{bt} - \delta_i \cdot DW_t - \pi_i \cdot 1_{[b \in \text{Used}_t]} - \varepsilon_{ibt} & \text{if } b \in CS_{it} \\ -\theta_i - \varepsilon_{ibt} & \text{if } b \notin CS_{it} \\ -\psi_i - \varepsilon_{ibt} & \text{outside option} \end{cases} \tag{5}$$

Notice that there are three specific cases in evaluating transactions costs: branch alternatives that are in the consideration set (e.g., those branches visited recently), branch alternatives outside the consideration set (e.g., a consumer unexpectedly visits are branch outside their usual areas), and an outside option which captures the decision to not visit any branch (e.g., not visiting a branch or choosing to visit a competitor).

If the alternative is in the consideration set then the transaction cost is the sum of three components plus a random error. The first component is the inner product $\tau_i \cdot B_{bt}$ where B_{bt} is a vector of transformed attributes of alternative *b* at time *t*—for example, size or capacity of the alternative chosen. We also include in this vector the number of competitor branches in the same ZIP code, and the number of other branches of the same bank in the same ZIP code. The second component represents an individual preference for a day of week and is measured as the inner product of $\delta_i \cdot DW_t$, where DW_t is a vector of zeros with a one in the place of the day of the week of t and δ_i is its respective sensitivity. The third component is meant to capture persistence in usage above and beyond the consideration set and scales an indicator, $1_{[b \in Used]}$, that detects if alternative b has been used in the past three months. If the alternative is not in the consideration set then the transaction cost is the sum of a constant and random component to give a small probability that any branch can be chosen.

We assume that consumers choose the alternative with the lowest transaction costs:

$$\tilde{v}_{it} = \min_b T_{ibt} \tag{6}$$

The unobserved idiosyncratic random shock ε_{ibt} in (5) is assumed to follow a Type I extreme value distribution and is i.i.d.. This assumption yields a logit form for the probability of choosing alternative *b*.

2.1.3 Consideration Sets

The data shows that most consumers visit a small number of branches, which suggests that consumers focus on branches close to areas that they live, work or shop. Also, most consumers in our dataset, as of 2012 and early 2013, have not used the mobile channel or online channel. This suggests that consumers lack smartphones or might not be aware of all the available alternatives. Therefore, we employ a consideration set model (Schoker et al. 1991, Roberts and Lattin 1991). We allow the consideration sets to vary across consumers and time.

Consideration sets are not observed directly in the data, so we have to infer them from our data set using transactions, demographic, and geographical data. To simplify our analysis we propose the following structure to decide upon a consumer's consideration set. For the digital channels, the mobile and online channels are assumed to be included in the consideration set if a consumer performed a transaction on one of these channels in the past 12 months, otherwise these channels are not in the consideration set.

In the case of physical channels—branches and ATMs—consumers tend to concentrate their transactions in few geographic areas with some apparently random deviations. To capture the geographic influence on consideration sets we use the geographic boundaries imposed by ZIP codes and assume that consumers will form their consideration sets based upon a set of ZIP codes that they frequent. Choosing consideration sets based upon geographic location is important so that we can assess the effect of opening new branches. We design the consideration sets based upon focal points of activity, such as their home ZIP code, and other ZIP codes that are likely to be visited based upon their connectedness.

Typically consumers have high concentration of activity in certain geographic areas like their home or work. We call these areas focal points. High activity areas are found for each consumer. Specifically, we designate the first focal location as the ZIP code area associated with the model value of the number of customer transactions from all sources: debit cards, credit cards, visits to ATMs or branches. A consumer specific second focal location is defined as the home address ZIP code.

Consumers often use shop or bank near their focal points or along routes between focal points. For example, a consumer might stop at a branch en route to work from home. Connectedness between geographic areas are related to physical proximity, but limiting ourselves to contiguous zones does not adequately reflect the visitation patterns that we observe in our exploratory analysis. For example, a consumer may never visit a branch that is in the zone south of their home, but may visit several branches to the north according to the route that they take from their home in the south to the office in the north. Therefore, we decided to construct a measure of connectivity between ZIP codes.

The connectedness of ZIP codes is measured by constructing a matrix of co-occurrences of pairs of ZIP codes being visited. We use the data from customer transactions from the previous year to avoid biasing our connectedness matrix. In each cell, the number of customers who performed transactions in both ZIP codes was tallied. The diagonal measures the total number of consumers who performed transactions in the specific location. Finally, the matrix rows were normalized to reflect the percentage in each cell (*i,j*), which is the percentage of consumers of ZIP code *i* who also performed transactions in *j*. We call this matrix was the matrix of connectivity. Note that this matrix is not symmetric, reflecting the fact that connectivity is not symmetric between locations. For instance, a suburb is well "connected" with downtown, but not vice versa; this is because most of the people who bank in the suburb also bank in downtown, but not a large percentage of people who bank in downtown also bank in that specific suburb. An interesting aspect is of this matrix is that connected ZIP locations are not necessarily contiguous. For example, this matrix shows a suburban area connected with a shopping area that are not contiguous, another example is that when two non-contiguous areas may be connected due to a major highway between these areas, or a third example is two contiguous areas due to a physical boundary like a river or mountain. Our assumption is that branch placement does not affect our connectivity matrix, and this connectivity matrix is stable through time. Our connectivity matrix correlates well with our knowledge of shopping and travel patterns.

Consideration sets may change over time, perhaps due to a branch being opened or closed. In the dynamic programing problem, we assume that consumers do not anticipate future consideration set changes. This is to simplify an already complex optimization problem. However, consumers do anticipate variations in the transaction costs because of shopping or travel patterns. For example, a consumer who lives in the suburbs may have lower transaction costs associated with a downtown branch during weekdays or workweek and higher transaction costs on the weekend.

Based upon our formulation of the consideration set we can compute the probability for consumer i to visit branch b during period t to perform a banking transaction:

$$P_{ibt} = \begin{cases} \dfrac{\exp\{T_{ibt}\}}{\exp\{-\psi_i\} + \exp\{-\theta_i\} + \sum_{j \in CS_i} \exp\{T_{ijt}\}} & j \in CS_i \\ \dfrac{\exp\{-\theta_i\}}{\exp\{-\psi_i\} + \exp\{-\theta_i\} + \sum_{j \in CS_i} \exp\{T_{ijt}\}} & j \notin CS_i \\ \dfrac{\exp\{-\psi_i\}}{\exp\{-\psi_i\} + \exp\{-\theta_i\} + \sum_{j \in CS_i} \exp\{T_{ijt}\}} & \text{Outside Option} \end{cases} \tag{7}$$

Notice that whenever a new branch is added to the consideration set the probability of the outside option is decreased, while whenever a new branch is added to the consideration set from perhaps opening a new branch there is a decrease in the probability of the outside option. Conversely closing a branch always decreases the service level for consumers that include that branch in their consideration set.

2.1.4 Heterogeneity

We have two sources of heterogeneity in the model. First, through parameter sensitivities and transaction costs parameters. We use a hierarchical Bayesian framework where for each parameter $\phi_i \in \{\eta_i^h, \eta_i^p, \gamma_i, \tau_i, \delta_i, \pi_i, \theta_i\}$ a normal prior distribution is defined $\phi_i \sim N(\bar{\phi}, \Sigma)$, with $\bar{\phi}$ a linear function of demographics. We specify a diffuse hyper-prior for these parameters. The second source of heterogeneity is based on customer location through the consideration set. Model testing with and without consideration sets show large improvement in fit and prediction.

2.2 Consumer's Dynamic Programming Problem

In each period, consumers decide whether to visit a bank (if a visit is made then $V_{it} = 1$ otherwise it is 0) and how much to withdraw (Q_{it}) in order to minimize total costs of waiting, transactions, and expected future costs. The state space is defined as $S_{it} = \{k_{it}, r_{it}, D_{it}, A_{sit}, N_{it}, \varepsilon_{ibt}\}$. The optimal sequence of decisions at time origin τ can be found solving the following dynamic programing problem with discount factor β as:

$$\min_{\{V_{it},Q_{it}\}_{t=\tau}^{\infty}} E_{\{S_t\}_{t=\tau}^{\infty}}\left[C_{i\tau}\left[V_{i\tau},Q_{i\tau};\mathrm{S}_{i\tau}\right] + \sum_{t=\tau+1}^{\infty} \beta^t C_{it}\left[V_{it},Q_{it};\mathrm{S}_{it}\right] \middle| \{S_t\}_{t=0}^{\tau-1} \right] \tag{8}$$

Where

$$C_{it}\left[V_{it},Q_{it};S_{it}\right] = \begin{cases} \omega_{it}\left[V_{it},Q_{it};S_{it}\right] & \text{if } V_{it}=0 \\ \nu_{it}\left[V_{it},Q_{it};S_{it}\right] & \text{if } V_{it}=1 \end{cases} \tag{9}$$

We use the term $\nu_{it} = \tilde{\nu}_{it} + \eta_i \cdot k_{it}^*$ to represent the smallest transaction cost among all alternatives given in (6) plus the holding cost of the optimal cash amount to be held at the end of the period, ω_{it}, as defined in (4).

We define the value function $\zeta_{i\tau}$ as follows:

$$\zeta_{i\tau}\left[S_{i\tau}\right] = \min_{\{V_{it},Q_{it}\}_{t=\tau}^{\infty}} E_{\{S_t\}_{t=\tau+1}^{\infty}}\left[C_{i\tau}\left[S_{i\tau}\right] + \sum_{t=\tau+1}^{\infty} \beta^{t-\tau} C_{it}\left[S_{it}\right] \middle| \{S_t\}_{t=0}^{\tau} \right] \tag{10}$$

Since this is an infinite horizon dynamic problem in equilibrium the policy function is independent of time and we can write the problem using the Bellman equation:

$$\zeta_i\left[S_{i\tau}\right] = \min_{V_{i\tau},Q_{i\tau}} E_{S_{i,\tau+1}}\left[C_i\left[V_{i\tau},Q_{i\tau};S_{i\tau}\right] + \beta\ \zeta_i\left[S_{i,\tau+1}\right] \mid S_{i\tau} \right] \tag{11}$$

No analytic solution is known, so we rely on a numerical approach which we describe in the following section.

2.2.1 Identification and data limitation

The model as stated before cannot be completely identified with the data we have available. First, the outside option is normalized to unity for every consumer $(\psi_i = 1)$. Second, the frequency of interaction of consumers with the bank is observed which allows us to infer the relative tradeoffs between waiting and transaction costs. However, our main interest is in the transaction costs, since it drives decisions about branch choice. Therefore we choose to fix the parameters associated with holding costs and chose to set the cost of postponing a check deposit and the cost of having excess cash to be unity $(\eta^h{}_i = \gamma_i = 1, \forall i)$ for every consumer. As a consequence the penalty for not having enough cash equals two. Third, we calibrate a constant rates of consumption and check arrivals between bank visits using the first year of data which is not used for estimating the parameters.

2.2.2 Understanding the Model

To illustrate our model we simulate two customers and depict their cash withdrawls, deposits and visits to branches in **Figure 3**. For each customer we plot three time series: the top plot represents their available cash per period; the middle plot are the amount of undeposited checks, and the bottom plot depicts which branch is visited during each period. The branch choices are indexed by an integer from one to five, and zero represents the choice to not visit any branches. The consumer on the left (panel a) depicts a customer with low holding costs relative to their transaction costs (i.e., the ratio ρ). In other words the customer is willing to wait quite a while to deposit checks and experiences a relatively high transaction cost, and as a consequence chooses to only visit the bank eight times during 100 periods. In contrast the consumer on the right (panel b), visits branches more frequently since their holding costs relative to their transactions costs are greater. In other words the consumer does not want to wait to deposit a check and visits a branch 23 times during the same 100 periods. Remember that consumers make a decision about whether to visit a branch or not based on expected future costs, so consumers are trading off the cost of holding a check against the transaction cost of visiting a branch and running out of cash.

(a) Low Holding/Transaction ratio (b) High Holding/Transaction ratio

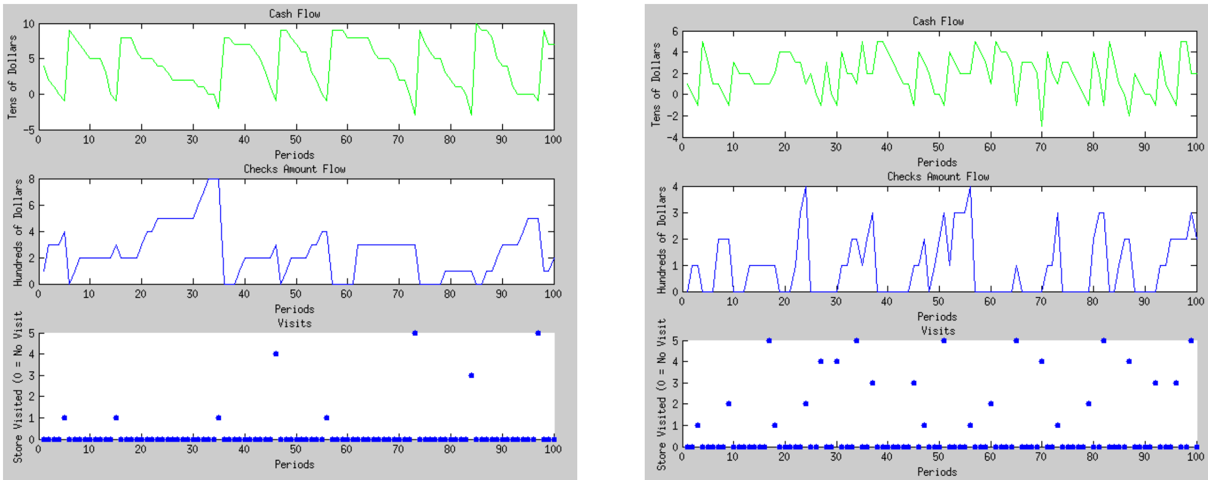

Figure 3. Model workings, bank activities vs periods.

2.3 Model Estimation

In order to estimate our model, it is necessary to solve a discrete choice dynamic programing (DDP) problem. In the marketing literature, many models with dynamic decisions have been estimated—for example, dynamic brand choice (Erdem and Keane 1996; Gönül and Srinivasan 1996), dynamic quantity choice (Sun 2005), or new product adoption (Song and Chintagunta 2003). However, the techniques employed in these papers cannot be used in our case due to the size and complexity of our problem.

Our problem shares many similarities with dynamic inventory problems found in the operations research literature, which is quite extensive. Harris (1913) proposed the classic model of economic order quantity (EOQ). Clark and Scarf (1960) proved the optimality of the (s, S)-policy for the stochastic demand model under very general conditions. The (s, S)-policy is a closed-form solution for the dynamic problem, where an order is placed when inventory level reaches s and the order quantity is S-s. Since 1984, efficient algorithms have provided fast computation of the policy (Ferdeguren and Zipkin). Our problem is an extension of the inventory problem to a multiple product inventory problem known as a Stochastic Joint Replenishment Problem (SJRP). Many policies have been proposed to solve this DDP, but none of them have been shown to be optimal (Ozkaya et al 2006). However, there are heuristics which can yield approximate solutions (Viswanathan 1997, Johansen, and Melchiors 2003).

To solve our DDP and simultaneously estimate the parameters of the model, we adopt the technique proposed by Imai, Jain, and Ching (2009) (which we abbreviate as IJC). This approach uses the Metropolis-Hastings algorithm to estimate the model's parameters, and within each MCMC step it iterates the value function once to improve our solution but avoid the computational cost of solving the value function. Solutions of the value function can be approximated from previous evaluations using a non-parametric approach. This method allows estimating the model without explicitly solving for the optimal policy function at every iteration. This procedure reduces our computation burden making it possible to estimate our model for a large number of consumers such as our problem.

However, even using the more efficient approach from IJC it is not possible to apply it directly to our data set due to the scale of our dataset. To speed up the computation, we employ a parallel computing approach suggested by Neiswanger, Wang and Xing (2014). Exploiting the properties of the likelihood function they propose a parallel MCMC algorithm in which subsets of the data are processed independently. The samples can then be combined using a semiparametric mixture to make inferences about the posterior distribution. This combination of IJC and parallel computation was suggested by Liu et al (2016). Running the IJC method in parallel with Neiswanger at al. draw from the true posterior reduced the estimation time to a few days of computation. In our problem we use 30% of the consumers in our sample and employ 50 CPUs to estimate in 1 week what would have taken 3.4 years of computation on a single processor.

3. Empirical Results

In this section, we show our estimation results and discuss their managerial implications[11]. The bank which provided our data is for a single metropolitan area, which includes the city and its surrounding suburbs, and is located in the midwest. Our target bank is the leading bank in this region with approximately 200 branches and more than 2,000 ATMs. In order to understand consumer behavior, we use daily transaction data from all bank channels during one calendar year (2012). To estimate the connectivity among ZIP codes, we use the prior year of transactions (2011).

[11] The numbers shown in this version are computed with 10% of the customers available to us in our database. A sample of the customers were chosen to reduce the computational burden.

Digital channels—online and mobile—are becoming more and more important, but the vast majority of digital transactions, 85%, are inquiries to check account balances. In this analysis, however, we ignore balance inquiries and focus on banking services that involve transfers of money between accounts, check deposits and cash withdrawals. We chose these types of transactions because they represent more than 99% of all transactions performed at a branch or ATM, and they are the main reason customers choose when and how to visit the bank.

If a consumer visits a branch they choose an alternative in the consideration set 91% of time. We conjecture that consumers are using branches near their home, work, or shopping locations. Among alternatives within the consideration set, the probability of choosing an alternative increases by 15% when it was previously used. Consumers' likelihood of banking on a non-favorite day drops 66% on average. We found that branches in high income ZIP codes are visited 22% less than other branches in the other ZIP codes. Further, we found that consumers prefer branches in suburban regions. Urban branches are, in fact, chosen 11% less than suburban ones.

The finding that high income regions are visited 22% less than other branches in other regions suggests that high income consumers perform fewer transactions, which is consistent with the lack of urgency to perform transactions based on need for cash or to deposit a check. Wealthier customers also have other means to perform transactions, thus alleviating the burden for branches in high income areas. Another effect that explains this behavior is that high income locations are populated with high branch density; thus, increased competition among branches reduces the attractiveness of individual branches. Similarly, a preference for suburban branches can be a consequence of branch density, because urban regions tend to have higher branch density than suburban and rural areas; therefore, the competition for consumers is stronger in those regions.

Our model shows that consumers prefer large branches. We were surprised to find that medium sized branches are far less preferred than small branches. Specifically their attractiveness drops by more than 15%. In addition to branch size, we use the log of the square footage of the branch as a proxy for a branch's capacity. The size of a branch is a proxy for its capability to perform more types of transactions, for its capacity to perform each service (e.g., more tellers), and for its increased amenities. Customers show that they prefer these attributes in the branches but, surprisingly, not in high proportions. We further determined that an increase of 10% in the surface area makes the branch only 1% more attractive.

ATMs are more preferred than branches but to a lesser extent than expected. We found that ATMs are approximately 8% more attractive than branches. The mild consumer preference for ATMs was counter-intuitive to us, but can be explained by three major effects. First, the number of ATMs is more than ten times higher than that of branches, so each ATM receives less attention. Second, at ATMs, consumers cannot perform all the type of transactions that are possible at branches. Third, ATMs present more security concerns than branches, so people tend to perform high value transactions at branches. Although ATMs can be more convenient, many customers still prefer branches.

We find few consumers use mobile channels, but more than 70% of those who try mobile banking at least once continue using it later on. Since the mobile channel is relatively new, we consider consumers to be aware of this channel after they try it once, thus incorporating it within their consideration set. When the mobile channel is within the consideration set, the alternative is 12% more attractive than the average branch, whereas the online channel is 22% more attractive than the average branch.

We find that consumers favorite day to bank is Friday, followed by Monday and then Wednesday. This strong desire to perform transactions on given days can be explained by the time compression caused by the limited weekend hours. Because of this lack of availability customers tend to use more branches located near where they work, as opposed to branches near where they live.

4. Branch Optimization

To find the optimal branch networks we have to make decisions about opening and closing branches within our target area. The number of permutations of potential locations leads to a combinatorial explosion, so an efficient optimization tool is required. To make our problem more tractable we assume that only a single new branch can be opened in each ZIP code. Obviously there are many possible locations within each ZIP code to locate a branch, but we lack detailed information about location availability, building costs, rental and leasing costs, government regulations, and which locations the bank is actually considering for expansion.

Therefore, we have restricted our simulations to new "average" branches with a medium size and average capacity. However, this not a limitation of our technique just our information set[12].

Using the posterior of the parameters associated with our demand model we simulate transaction costs T_{ijt} for a one month period of time for a sample of customers[13] S, using all the current branches, available channels and the outside option represented as set B. We average the transaction costs over our time frame and set the average for our alternatives in the consideration set as μ_{ij}, we let μ_{i0} be the average of the transaction costs of the alternatives outside the consideration set, and let $\mu_{i,-1}$ be the transaction cost of the outside option. We define x_j as binary variable, taking value one when we decide to keep branch j open and zero if the branch is closed, C_{ij} is binary parameter matrix that takes the value one if the alternative j is in the consideration set of i and zero otherwise.

The logistic form of our demand model yields a non-linear optimization problem. Given the size and complexity of our model, a direct optimization of our model is not feasible. Instead, we re-formulate the problem as a linear programing model. Taking advantage of the IIA property of the logistic probabilities we notice that the ratio between branches in the consideration set should remain constant for each customer:

$$p_{ij} \le \frac{\exp(\mu_{ij})}{\exp(\mu_{ik})} p_{ik} + (1 - C_{ik}) \qquad \begin{array}{l} \forall i \in S \\ \forall j,k \in B \\ j \neq k \end{array} \tag{12}$$

We need to impose a constraint to ensure that probabilities add to one:

$$\sum_j p_{ij} = 1 \qquad \forall i \tag{13}$$

We also impose that consumers can only choose alternatives that are available:

$$p_{ij} \le x_j \qquad \begin{array}{l} \forall i \in S \\ \forall j \in B \end{array} \tag{14}$$

Additionally, alternatives outside the consideration set have equal probability:

---

12 The number of branches per ZIP code can be extended. However, the complexity of the optimization grows multiplicatively with the number of branches per ZIP code.

13 We generate a random sample of 10,000 customers to speed up these computations.

$$p_{ij} \le p_{ik} + C_{ik} \qquad \begin{matrix} \forall i \in S \\ \forall j,k \in B \\ j \ne k \end{matrix} \tag{15}$$

Following this formulation we can now use different objective functions based on the customer's service level, which we define as $\left(1-p_{i,-1}\right)$. This measure's the probability that a consumer visits a branch. Our objective is to choose a branch format such that the probability that consumers choose to visit our bank is maximized:

$$\max_{x_j} \sum_i \left(1-p_{i,-1}\right) \tag{16}$$

In this simulation we assume that the number of branches desired by the bank is fixed at the value *NB*, which yields the constraint:

$$\sum_j x_j = NB \tag{17}$$

In other words, our objective is to maximize the service level or minimize customer attrition.

4.1 Counterfactuals

Using our model estimates and the linear programming approximation just described we perform three counterfactual simulations. First, we analyzed customer attrition when the bank closes a branch. Second, we analyzed location for opening new branches. Third, we analyzed the impact of increased mobile channel adoption by consumers.

4.1.1 Branch Closure

When the bank decides to close a branch, customers' utility decreases, as customers need to replace the branch with another alternative that has higher transaction costs. After closure, customers' probability of attrition increases or remains the same. In order to determine which branch to close, we can use the model to find branch to close that minimizes the negative impact on customers. To determine optimal branch closure we use the following objective function:

$$\min_b \sum_{i,b} \Delta P_{ib} \cdot V_i \quad (18)$$

Where ΔP_{ib} is the variation on the probability of attrition of customer *i* when closing branch *b*. V_i is the value of customer *i*. The customer value is computed as the average balance in the account for the prior year[14].

In order to test the impact of making decisions using this approach, we look at actual branch closures made by the bank. During the period of analysis, the bank closed several branches in order to comply with competitive regulations. One closure was for a moderate size, medium-income level branch located in a suburban region that is near a commercial district. In the data, we observed a high attrition level[15] among customers who visited the branch at least once in the six months before the branch closed. The data showed close to 8.5% attrition, while the model predicted an expected attrition level of just over 6.5% among the same customers during the next six months. One reason for this discrepancy is that when banks close a branch, they tend to close more than one at a time; this combined effect might cause an increase in the attrition level.

When limiting the branch selection to branches within the same ZIP code, or ZIP codes located in the immediate neighboring ZIP codes, the model suggests the closure of a different type of branch. Although this branch was of moderate size, it is a high income level branch located in an urban region with higher branch density, highly connected to a region with a high density of branches, and not highly connected to suburban neighborhoods.

We expected the attrition level to be 3.2%, and the expected weighted value loss ($\Delta P_{ib} \cdot V_i$) was 27% of the value loss computed closing the bank choice this implies a 73% reduction in value loss. For reference, we can convert attrition level difference into dollars, using customer lifetime value (LTV). We assume that LTV can be computed using the following formula:

$$LTV = \sum_{t=1}^{T} m_i \frac{r_i^t}{(1+d)^t} \quad (18)$$

[14] An example of the optimization problem can be found in Appendix 1.
[15] Attrition is defined as customer inactivity during the next six months.

We used the following parameters: average profitability per consumer \$300 per year, retention rate of 98.5% annually, and an annual discount factor of 3% with an infinite time horizon. According to this calculation, the bank would have saved \$530,000 with a single branch closure. During that year, the bank closed more than 20 branches in different locations. Thus, if we expect similar savings per each closure, in one year the bank would have saved more than \$10 million using this tool. The objective function is not minimizing attrition, but attrition weighted by customer value. If we were to minimize attrition without weights and using the same numbers as before for the same branch, the annual decision would have saved the bank more than \$12 million in lifetime value during that year.

4.1.2 Branch Opening

When a bank opens a new branch, it is important to determine the type and location of the branch. We can use our model to predict by how much the probability of attrition decreases when a branch is added. The model was used to rank regions in terms of potential attrition gain when adding a branch. We found that locations with low connectivity and low branch density are good candidates for a branch opening. Alternatively, we found that in locations where branch presence is low, mobile channel adoption should be encouraged. Conversely, in regions with high competition, a large branch with more amenities is recommended.

4.1.3 Mobile Adoption

The model assumes that consumers who have never used a mobile channel are not aware of that possibility; therefore, it is not included in the consideration set. To represent an increase in the adoption of mobile channels, we increase the mobile channel adoption among consumers who were aware of the option but who had never used it before. In other words, we include this alternative in the consideration set, so consumers who find this channel more attractive than other alternatives will begin to use it. When we add a mobile channel to the consideration set, customers' utility can potentially increase, since they now can substitute an alternative channel for this one if it gives them higher utility.

In our data set, the percentage of consumers who use a mobile channel is 18%. We ran counterfactuals for mobile adoption at 30%, 50%, and 70% levels of adoption. Specifically we random chose individuals in our sample and made them aware of the channel to complete the required adoption level. This can be interpreted as a campaign to educate consumers about mobile banking. Many banks are already using these campaigns to encourage and accelerate the transition of customers to the digital world. For each level of adoption, we then reevaluate the attrition probabilities. We look at small regions between one and three ZIP codes, and evaluate whether it is possible to eliminate a branch without increasing the attrition level above the level where it was before increasing the mobile adoption. When consumers adopt a mobile channel, the probability of attrition decreases more than 37% on average. For example, if a consumer had an attrition probability of 6%, then after the mobile channel is adopted this probability drops to 3.8%.

As we expected an increase in mobile adoption generates a substitution effect with transactions performed at branches, causing an excess capacity in that channel. However, even with high levels of mobile adoption, branches are always needed. This is not only because there are services that cannot be done through other channels, but also because for many customers' transactions performed using digital channels are far less preferred and lack of branches might make these customers switch to competitors.

It was surprising to find that an increase in mobile adoption can lead to an increase in demand for branches. We found two mechanisms that yield this effect. First, there is an attraction effect. Mobile channels may create the perception that the bank as a whole is more attractive. In response customers respond by switching transactions from competitors to the bank through mobile and branch channels. Second, there is a branch switching effect. Mobile channels create a distortion within the bank, switching transactions from one branch to another. Consider, for instance, a consumer who deposits large checks at a branch every Saturday. For convenience, the consumer also simultaneously performs other transactions at this branch. When a mobile channel becomes available, the consumer can perform the urgent deposit using their smartphone on Saturday but postpone the rest of their transactions until Monday when they go to work and are near to a physical branch location. In this case, the consumer's preferred branch is different than the one that is typically used on Saturdays. In this way, mobile channels alter the location of some transactions.

The overall branch usage depends on the summation of these effects. The substitution effect seems to dominate in most scenarios. Attraction and branch switching effects tend to be weaker when mobile adoption very high or very low. However, banks are in a transition period now, since mobile adoption is not too low, nor too high, and in this scenario, these effects can have a substantial impact. Using the knowledge gleaned from our methodology, we can summarize our findings in the following table:

		Branch Density	
		High	**Low**
Competition	**Leader**	Mobile ↑ => Branch Usage ↓	Mobile ↑ => Branch Usage ↓
	Follower	Mobile ↑ => Branch Usage ↓	Mobile ↑ => Branch Usage ↑

If the bank is dominant and there is high branch density, then the substitution effect dominates due to the bank's role as leader. Also, since the market share of competitors is small there are not many new customers to attract. Conversely, when the bank is a follower and does not have many branches, there are many customers who can migrate to the bank from competitors. Because branches are sparse, the mobile channel makes the bank very attractive to new customers, so an increase in mobile channel adoption can lead to an increase in branch usage. In the other quadrants, both effects coexist, and the substitution effect tends to dominate.

5. Conclusion

Our analysis suggests that it is essential to make branch opening and closure decisions using the predictions of the adoption of online and mobile channels. Making decisions without considering these other channels can lead to expensive decisions and market share loses. From the model, we conclude that rural areas and low density regions can be well served with mobile channels. When facing strong competition, the mobile channel is not enough and branches are needed. Given the current industry trend to reduce the number of branches we would recommend against closing branches in highly competitive regions.

An unexpected finding was that an increase in mobile adoption may lead to an increased demand for branch transactions. In regions where the bank is not a leader and there is low branch density, we recommend that the bank be prepared for an increase in demand at branch locations when mobile adoption increases. Depending on the situation, the bank should consider opening a branch in the region. This can be a good strategy to grow, attracting customers from competitors. As mobile usage continues to grow, our model will be increasingly useful for banks that need to determine optimal branch configurations. And with small changes, our model can be used to make decisions about other existing and new channels.

References

Ansari A., Mela C., & Neslin, S. (*2008*) Customer Channel Migration. Journal of Marketing Research: February 2008, Vol. 45, No. 1, pp. 60-76.

Arrow, K, Harris. T., & Marschak, J. (1951). Optimal Inventory Policy. Econometrica.

Blattberg, R., Kim, B., & Neslin, S. (2008). Database Marketing: Analyzing and Managing Customers. Springer-Verlag New York, eBook ISBN 978-0-387-72579-6.

Board of Governors of the Federal Reserve System. (2014) "Consumers and mobile Financial Services". *Federal Reserve Board Publications*.

Cambra-Fierroa, J., Wagner, K., Melero-Poloc, I., & Sesec, F. (2016). "Are multichannel customers really more valuable? An analysis of banking services" IJRM 1140.

Capgemini (2015), "World Retail Banking Report", http://www.worldretailbankingreport.com.

Clark, A. & Scarf, H., (1960). Optimal Policies for a Multi-Echelon Inventory Problem. Management Science.

Accenture Consulting (2013). "Celent Model Bank 2013: Case Studies of Effective use of Technology in Banking", web: https://www.accenture.com/us-en/insight-celent-model-bank-2013-banking-summary.aspx

Chintagunta, P. K., Chu, J., & Cebollada, J. (2012). Quantifying transaction costs in online/off-line grocery channel choice. Marketing Science, 31(1), 96-114.

Danaher, P. J., Wilson, I. W., & Davis, R. A. (2003). A comparison of online and offline consumer brand loyalty. *Marketing Science*, *22*(4), 461-476.

Erdem, T. & Keane, M. P. (1996). Decision-making under uncertainty: Capturing dynamic brand choice processes in turbulent consumer goods markets. Marketing science, 15(1), 1-20.

Federguren, A. & Zipkin, P. (1984). A Combined Vehicle Routing and Inventory Allocation Problem. Operations Research.

Gönül, F. & Srinivasan, K. (1996). Estimating the impact of consumer expectations of coupons on purchase behavior: A dynamic structural model. Marketing Science, 15(3), 262-279.

Harris, F. (1913). How many parts to make at once. The Magazine of Management.

Hitt, L. M. & Frei, F. X. (2002). Do better customers utilize electronic distribution channels? The case of PC banking. Management Science, 48(6), 732-748.

Hoehle, H., & Venkatesh, V. (2015). Mobile application usability: conceptualization and instrument development. Mis Quarterly, 39(2), 435-472.

Imai, S., Jain, N., & Ching, A. (2009). Bayesian estimation of dynamic discrete choice models. Econometrica.

Johansen, S.G. & Melchiors, P. (2003). Can-order policy for the periodic-review joint replenishment problem. Journal of the Operational Research Society.

Kumar, V. & Venkatesan, R. (2005). Who are the multichannel shoppers and how do they perform?: Correlates of multichannel shopping behavior. Journal of Interactive Marketing, 19(2), 44-62.

Kushwaha, T. & Shankar, V. (2013) Are Multichannel Customers Really More Valuable? The Moderating Role of Product Category Characteristics. Journal of Marketing: July 2013, Vol. 77, No. 4, pp. 67-85.

Liu, X., A.L. Montgomery and K. Srinivasan (2016). Optimizing Bank Overdraft Fees with Big Data. Carnegie Mellon University, Working Paper.

Neiswanger, W., Wang, C., & Xing, E. (2013). Asymptotically exact, embarrassingly parallel MCMC. arXiv preprint arXiv:1311.4780.

Neslin, S., Grewal, D., Leghorn, R., Shankar, V., Teerling, M., Thomas, J., & Verhoef, P. (2006). Challenges and Opportunities in Multichannel Customer Management. Journal of Service Research, Volume 9, No. 2.

Laukkanen, T. (2007),"Internet vs mobile banking: comparing customer value perceptions", Business Process Management Journal, Vol. 13 Iss 6 pp. 788 – 797.

Ozkaya, B.Y., Gurler, U., & Berk, E. (2006). The stochastic joint replenishment problem: A new policy, analysis, and insights. Naval Research Logistics.

Roberts, J. H. & Lattin, J. M. (1991). Development and testing of a model of consideration set composition. Journal of Marketing Research, 429-440.

Shankar, V., Smith, A. K., & Rangaswamy, A. (2003). Customer satisfaction and loyalty in online and offline environments. International journal of research in marketing, 20(2), 153-175.

Shocker, A. D., Ben-Akiva, M., Boccara, B., & Nedungadi, P. (1991). Consideration set influences on consumer decision-making and choice: Issues, models, and suggestions. Marketing letters, 2(3), 181-197.

Song, I., & Chintagunta, P. K. (2003). A micromodel of new product adoption with heterogeneous and forward-looking consumers: Application to the digital camera category. Quantitative Marketing and Economics, 1(4), 371-407.

Sun, B. (2005). Promotion Effect on Endogenous Consumption. Marketing Science 24(3) 430–443.

Taha, A., Jahed, D. H., Ahmad, M. N., & Zakaria, N. H. (2013, November). Antecedents of customer satisfaction in mobile commerce: A systematic literature review. In 2013 International Conference on Research and Innovation in Information Systems (ICRIIS) (pp. 554-558). IEEE.

Takeyasu, K. (2008). Mobile marketing. Information Systems Research Methods, Epistemology, and Applications, 328.

Valentini, S., Montaguti, E., Neslin, S. (2011) Decision Process Evolution in Customer Channel Choice. Journal of Marketing: November 2011, Vol. 75, No. 6, pp. 72-86.

Venkatesan, R., Kumar, V., & Ravishanker, N. (2007). Multichannel shopping: causes and consequences. Journal of Marketing, 71(2), 114-132.

Viswanathan, S. (1997) Periodic review (s, S) policies for joint replenishment inventory systems. Management Science.

Wallace, D. W., Giese, J. L., & Johnson, J. L. (2004). Customer retailer loyalty in the context of multiple channel strategies. Journal of retailing, 80(4), 249-263.

Yang, B., & Ching, A. T. (2013). Dynamics of consumer adoption of financial innovation: The case of ATM cards. Management Science, 60(4), 903-922.

Appendix

1. Summary of variables

Name	Description
C_{it}	Total costs incurred by consumer i from period t to the future
k_{it}	Money consumer i held at the beginning of period t
h_{it}	Amount of checks to deposit by consumer i at the beginning of period t
η_i	Holding cost parameter for consumer i
Q_{it}	Quantity that consumer i will withdraw at period t
W_{it}	Waiting time of consumer i from period t to the next transaction at a branch.
γ_i	No deposit cost parameter for consumer i
N_{it}	Number of checks that arrived for consumer i at period t
A_{its}	Amount to deposit from check s
a_s	Arrival time of check s
β	Discount factor
τ_{it}	Lowest transaction cost i have at period t given its branch consideration set
Q^*_i	Optimal quantity to withdraw in the future
d_i	Consumer i average money requirement per period
ε_{wit}	Random cost shock
α_i^1, α_i^2	Consumer sensitivity to alternative/branch characteristics
B_{bt}^1, B_b^2	Branch attributes and location characteristics
L_i	Vector of location probabilities

2. Other Optimization procedures

In this paper, we want to answer questions about proper branch network structure and how the bank should react to the introduction of the mobile channel and continue growth of alternative channels. There are many objectives we could use to improve the branch network; here, we use consumer attrition for two reasons. First, the lifetime value of a consumer represents $X that is $Y proportion of an average branch operating costs. Second, it is an objective measure that can be computed directly from observed data.

We compute regional maps with consumers' activity with current configurations and shifts in consumer activities due to branch network configurations or channel adoption. Based on this shifting and variation in transaction/holding costs caused by this shift, we compute what variation generates the minimal attrition level or lower decrease in service level. Based on the current branch network configuration, the optimization model can suggest what branches could be closed without increasing consumer attrition more than a certain level. We also can recommend openings or increase capacity in a region that shows high demand for branch activity. With these tools we can run counterfactuals with different levels of adoption of alternative channels, what branches would be not necessary and in what regions would be profitable to encourage consumers to move to these alternative channels. We define consumer attrition percentage as the proportion of inactive consumers for more than 3 months. A consumer is considered inactive when she doesn't perform any interaction with the bank in a given period.

3. Branch closure optimization example

The output of the consumer demand model can be used to compute customers' probabilities of using other alternatives or quitting the bank. This optimization can be done with any number of branches in the analysis—for example a neighborhood, a ZIP code, or a city. To demonstrate the optimization workings, we use an example of five branches. We first use the model to compute all the transition probabilities for a given consumer when the bank closes each branch, as depicted in **Table 1**.

S(b,b')	B1	B2	B3	B4	B5	Attrition
B1	-	0.15	0.25	0.30	0.25	0.05
B2	0.20	-	0.35	0.12	0.30	0.03
B3	0.40	0.15	-	0.19	0.20	0.06
B4	0.20	0.30	0.29	-	0.17	0.04
B5	0.25	0.35	0.18	0.20	-	0.02

Table 1. Example of transition probability after closing a branch

In this example, the customer would be least likely to quit the bank if B5 were closed, because its probability of attrition is the smallest (0.02). Following the same procedure for multiple customers and weighting the attrition probability by their value, it is possible to choose the optimal branch to close.