

Digital Twins: A Generative Approach for Counterfactual Customer Analytics

Samuel Levy

Carnegie Mellon University, silevy@andrew.cmu.edu, levysamuel.com

July 31, 2023 [Click for Latest Version](#)

Abstract

This research provides a novel methodology, Digital Marketing Twins, that automatically extracts latent features from individual-level brand survey responses to inform a statistically-principled, deep generative model of customer-side brand affinity and firm-side performance factors. The proposed model enables marketers to find drivers of individual-level brand affinity, as opposed to traditionally observed metrics that must be analyzed in aggregation. The framework serves a counterfactual purpose at the customer level. The generative part of the model *completes* the distribution of survey responses over time, and across firms – thereby addressing the archetypal missing data problem – by imputing customer responses in counterfactual regimes. The proposed prescriptive framework also proposes policy optimization through customer surveys, using Bayesian optimization, which efficiently identifies “paths of least resistance” among customer responses to service-quality questions – a search that otherwise would represent a complexity of $\mathcal{O}(n^d)$.

This research applies Digital Marketing Twins methodology to the competitive landscape of the U.S. wireless telecommunications retail market, leveraging a unique dataset of large-scale quarterly brand surveys from all three major carriers (AT&T, T-Mobile, and Verizon) from 2020 to 2022. It optimizes over the learned generative model from the multi-firm brand surveys to provide marketing policy recommendations according to individual-level counterfactual responses and different carriers. Empirically, this approach reveals latent asymmetries in competition in terms of brand affinity, together with a nonlinear increase in brand affinity for certain types of drivers, such as satisfaction with network speed, but a nonlinear decrease in brand affinity for customers who report greater likelihoods of changing plans, providers, or devices, relative to their current wireless services.

Key words: Customer Satisfaction; Competitive Environment; Survey Research; Deep Generative Models; Counterfactual Reasoning; Bayesian Optimization.

1. Introduction

Customer surveys are ubiquitous tools. Marketers leverage them to fuel brands and boost corporate growth, as well as determine the causes of customer satisfaction and of customer churn. Marketing researchers adopt them to learn customer preferences, gauge customer satisfaction, identify competitive offers, improve existing products and services, tailor marketing strategies, and innovate personalized services. Due to the ever growing complexity, and frequency of customer surveys and

survey touchpoints though, firms today mostly rely on third-party platforms (e.g., Salesforce, HubSpot) or survey companies (e.g., Ipsos and Kantar) to execute large-scale customer surveys, known commercially as *brand surveys*.

The increasing scale and scope of such brand surveys also is part of a broader trend of marketers embracing data-driven approaches, with an emphasis on the use of behavioral metrics to compute customer lifetime value (CLV) (e.g., Venkatesan and Kumar 2004, Fader et al. 2005). Behavioral metrics help identify which customers are at risk of churning, though recent research also calls for efforts to distinguish the causes from the predictors of churn (Braun and Schweidel 2011, Ascarza et al. 2018, Ascarza 2018). That is, predictors of churn include demographics and behavioral patterns that statistically indicate a high likelihood of discontinuing products or services, but the causes of churn might include poor customer service, high prices, product quality issues, or more attractive offerings from competitors. Customer satisfaction targets the root issues that lead to customer attrition (Gustafsson et al. 2005). Moreover, customer satisfaction and dissatisfaction drive companies' stock prices asymmetrically: dissatisfaction harms returns far more than a one-unit increase enhances them (Malshe et al. 2020). Accordingly, a new methodology is needed to carve out the "paths of least resistance" to individual-level customer satisfaction; this research proposes a prescriptive framework that pioneers such policy optimization by relying on customer surveys.

Despite the numerous benefits of customer surveys for marketers, there are two issues – one theoretical and the other practical – that plague most large-scale surveys carried out for customer relationship management (CRM) purposes. First, the surveys are difficult to integrate into prescriptive frameworks. Linear functional forms, strong parametric assumptions, and limited consideration of customer heterogeneity as a result of limited or incomplete individual-level data across brands makes it difficult for marketers to understand customer churn and retention in mature, competitive environments. Which marketing action should be recommended when customer satisfaction declines? Although one obvious lever is promotional offers, there are others; for example, in the wireless telecommunications industry, managers can increase customer satisfaction by improving network quality and network speeds, providing better data plans, strengthening brand perception, and solving problems that customers encounter when using providers' devices and services. However, the question remains: which aspects should be prioritized at the customer level? Second, in practice, customer surveys often are repeated cross-sectional. Because repeated cross-sectional surveys represent *different* sets of customers at various points in time, they cannot track individual changes. Unlike longitudinal surveys, they take "snapshots" and thereby provide less depth of information about individual respondents. If the sampled population changes significantly over time, comparisons between different cross-sectional surveys become challenging, if not impossible. For

this reason, the nature of data variations in the repeated cross-sectional format of brand surveys is described as *pseudo-longitudinal*.

In this paper, I propose a novel methodology – Digital Marketing Twins – that leverages large-scale brand surveys conducted by a focal firm and its competitors in the U.S. wireless telecommunications retail market. This methodology finds paths of least resistance to individual-level customer satisfaction, in a statistically principled way. It uses a unique dataset from a representative sample of customers of AT&T, T-Mobile, and Verizon, the three major players in the U.S. telecommunications market. Using quarterly cross-sectional survey responses that span ten quarters – from 2020 to 2022 – the methodology overcomes both the substantive and technical limitations previously mentioned. The framework builds a generative model of customer preferences by flexibly mapping individual-level surveyed characteristics to various dimensions of customer satisfaction. Generative models capture the joint probability distribution between observed and latent variables of interest. In practice, they provide the steps that explain how the data are assumed to be generated, allowing marketers and researchers to incorporate domain expertise into their models. The generative aspect not only supports the forecasting of customers’ responses in the next quarter, but also provides counterfactual responses according to different scenarios, such as customer responses as if they were using a different wireless carrier, all else being equal.

The digital twins approach, already an established method for counterfactual simulations in the realm of manufacturing, presents a novel and previously unexplored avenue for application within the marketing field. Until now, this innovative approach has, to the best of my knowledge, remained unapplied to this context. Digital twins integrate data from different sources to mimic the behavior of physical objects or systems; they can be used to test hypotheses, simulate scenarios, and optimize the performance of the systems. The use of generative models as a basis for digital twins is not novel; for example, generative adversarial networks (GANs) and conditional GANs can learn distributions of interest in structures with material nonlinearities and uncertainties (Tsiamanis et al. 2021). In marketing contexts, digital twins serve a counterfactual purpose at the customer level. The generative part of the model *completes* the distribution of survey responses over time, and across firms, such that it can address the archetypal missing data problem. The proposed Digital Marketing Twins methodology offers solution to the missing data problem that arises from the pseudo-longitudinality of brand surveys – in which individual respondents can be only observed in one time period and at one company – by imputing customer responses in the next time period and in a counterfactual regime in which all individually observed characteristics remain constant. The goal is to infer customer satisfaction under counterfactual regimes of their experiential, engagement, and usage characteristics, identifying the potential causes of satisfaction.

Table 1 provides a summary of the conceptual benefits of the Digital Marketing Twins framework. Quasi-experimental methods can be particularly useful for understanding the effect and strategic value of an intervention on an outcome of interest from recorded CRM data. A combination of propensity score matching and a flexible Bayesian parametric or nonparametric model such as a GAM, a Gaussian process (GP), or a mixture-of-normals can be useful for making counterfactual predictions. However, these methods are seldom scalable out-of-the box; they require careful selection of kernel functions and/or hyperparameters, which in turn requires expert knowledge. The quasi-experimental methods’ inference (especially of natural experiments) is necessarily that of what *has* happened, rather than what *could* happen. Although state-of-the-art predictive CRM methods provide forward-looking analytical insights, often with the help of rich machine learning techniques, they lack the ability to provide counterfactuals by focusing on predictors of customer churn, neglecting complex structures such as competitive effects, and relying on flexible statistics of customer behavior.

	Predictions / Tactical	Counterfactuals / Strategic
Retroactive	Basic CRM Reports	Quasi-Experimental Methods
Proactive	Predictive CRM Models	Digital Twins

Table 1 Digital Marketing Twins as a proactive and strategic framework for customer analytics.

From a technical standpoint, the goal of this research is to develop a novel deep generative and probabilistic latent factor model, as well as to leverage Bayesian optimization to find the best marketing actions to recommend at the individual level, from a large-scale survey. The Digital Marketing Twins model captures customer-side brand affinity at the individual level, for each brand, in each time period, controlling for observed heterogeneity and firm-side factors. The inference model mapping from data to the latent space is parameterized by a neural network, for high flexibility. Latent, customer-side brand affinity provides an interpretable layer that maps to a latent utility model that in turn yields an ordinal logit structure for brand survey questions pertaining to customer satisfaction. To generate counterfactual responses and missing quarters, it relies on amortized inference, learning a set of parameters that can map any data point to the latent space. I train the model using stochastic variational inference with mini-batching, for high scalability and uncertainty quantification. After training, a Bayesian optimization method maximizes individual-level, latent, customer-side brand affinity for customers of the focal firm, to discover the marketing actions most likely to increase customer satisfaction. This maximization leads to a “path of least resistance” at the individual level, enabling marketers to use surveys to identify causes of satisfaction. Applications of Bayesian optimization on a latent space appear on other contexts, outside

marketing (e.g., Gómez-Bombarelli et al. 2018, Griffiths and Hernández-Lobato 2020), and is useful in situations in which gradients are not accessible.

The remainder of the paper is organized as follows: Section 2 presents a literature review. In Section 3, contains a description and exploration of the data, using simple descriptive techniques. Section 4 provides the methodology. Section 5 provides fit metrics, benchmarks the Digital Marketing Twins model against nested baseline models, analyzes the probabilistic latent factors, and provides counterfactual results. Section 6 introduces a novel prescriptive framework to optimize customer satisfaction with Bayesian optimization. Section 7 concludes.

2. Related Literature

This work contributes to academic literature on digital twins using generative modeling, as well as to literature on machine learning methods in marketing for competitive environments; it also shows that early models for customer satisfaction are special cases of the proposed framework.

2.1. Digital Twins and Generative Modeling

Tsialiamanis et al. (2021) suggest how to advance system simulation by creating digital twins for specific systems, referring to fields such as manufacturing, control systems, the Internet of Things (IoT), smart cities, social networks, and management. Digital twins can help predict the behavior of structures in different situations, thus maximizing the operational lives of the structures and minimizing costs. However, the construction of digital twins is inherently complex and uncertain. Aleatory uncertainty, related to random events, and epistemic uncertainty, related to a lack of knowledge, are key considerations. To address these issues, Tsialiamanis et al. (2021) propose the use of generative models as the foundation for a digital twin, providing estimations of aleatory and epistemic uncertainty. They study two types of generative models: the Stochastic Finite Element (SFE) method – a physics-based, white-box model – and the Conditional Generative Adversarial Network (cGAN) – a data-driven, black-box model. Each has strengths and limitations. For example, SFE models excel in predefined conditions but struggle with unknown scenarios, whereas cGANs can perform across a wide range of conditions but cannot extrapolate beyond available data. With a hybrid, grey-box approach, incorporating both models to overcome these limitations, generative models might better accommodate uncertainty in digital twins. By combining a generative white box (SFE) and a generative black box (cGAN), they propose a fully generative grey box that they assess in relation to other existing models, such as variational auto-encoder (VAE) and Gaussian processes (GP). In this paper, on the other hand, I use a deep generative model building on a variational auto-encoding neural network to mirror the competitive environment, and use a Gaussian process to optimize the latent customer-side brand affinity.

Kapteyn et al. (2021) propose a new mathematical foundation for digital twins, that is, computational models that mirror structures, behaviors, and contexts of physical assets. Because digital twin applications usually require extensive resources and expertise for implementation, the authors propose a unifying mathematical model that uses dynamical systems theory and probabilistic graphical models, with the digital twin and the physical asset modeled as coupled dynamical systems that evolve over time, and the digital twin constantly updating its internal models according to observational data. By demonstrating this approach with a digital twin of an unmanned aerial vehicle (UAV), they show how the model aids in calibration by updating internal models, and facilitating decision-making. They conclude with the presentation of an abstract state-space formulation for digital twins, describing a realized dynamic decision network based on this mathematical model and illustrating its application to a UAV's structural digital twin.

Finally, Yu et al. (2021) propose a health monitoring solution for complex systems in smart manufacturing, applying a digital twin approach with a nonparametric Bayesian network model. With advancements in sensor technology and artificial intelligence, modern manufacturing systems need to be intelligent, visual, and capable of self-assessing their health throughout their life cycle. The Prognostic and Health Management (PHM) process is crucial in this context, and the proposed model offers an innovative solution for tracking the health states of such systems. The model collects sensor data from the physical world, updating its simulated physical model in real time and providing optimization and decision support. Their nonparametric Bayesian network model can adapt in real-time too, thus reducing model uncertainty. Yu et al. (2021) also include model validation experiments on electro-optical systems, and provide more accurate health monitoring than a traditional data-driven Convolution Neural Network (CNN) approach.

2.2. Machine Learning Methods in Marketing for Competitive Environments

Among the proposals for machine learning techniques to study market structures and competitive landscapes, Netzer et al. (2012) systematically analyze online user-generated content to “listen” to what customers write about a focal firm's and competitors' products; they use text mining to overcome the difficulties involved in extracting and quantifying the wealth of online data that customers generate and network analysis tools to convert the mined relationships into co-occurrences among brands or between brands and terms. (Tirunillai and Tellis 2014) extract latent dimensions of customer satisfaction with quality, using an unsupervised latent Dirichlet allocation model, and Lee and Bradlow (2011) automatically elicit product attributes and extract brands' relative positions from online customer reviews, providing both predictive and descriptive support for managerial decision making.

Brand competition occurs not only in single markets, but also in different sub-markets and structured markets. The hypothesis of multiple structured markets (Kannan and Wright 1991)

helps us understand how brands compete by including marketing mix variables. In type-primary markets, “switchers” are highly responsive to changes in marketing mix variables whereas in brand-primary markets, the “loyal” segment remains relatively unresponsive to marketing programs (e.g., in contexts of ground coffee purchases or store panel records). Ringel (2023) recently proposed the visualization of brand competition in a multimarket membership product (MMP) context, in which products that compete in multiple submarkets that are each characterized by distinct competitors and customer preferences, with competitive relationships inferred from customers’ online searches using bootstrapped neural network product embeddings in the digital camera market.

2.3. Customer Satisfaction and Survey Research

Using an ordered probit model, Kekre et al. (1995) study determinants of customer satisfaction for software products and service support for mainframes and workstations. Their main dependent variable is an overall satisfaction score, measured on an ordered categorical scale. The authors can explain how certain features of the software, such as reliability, capability and usability, affected overall satisfaction. They consider other explanatory variables, such as the type of product and the type of user, allowing for interaction effects. If an ordered logit were substituted for their ordered probit, their model can be nested within the proposed framework, by replacing amortized neural networks with linear functions, assuming that business key performance indicators (KPI) have no impact on customer satisfaction and considering only overall satisfaction as a unique target variable.

3. Exploring the Data

The data for this study consist of repeated cross-sectional responses from a brand survey for all major U.S. telecom carriers (AT&T, T-Mobile, and Verizon) between the third quarter of 2019 and the third quarter of 2022. Responses are recorded quarterly. For each carrier – not necessarily in the order previously mentioned, for confidentiality – I observe responses from a sample of 8770 customers, 7129 customers, and 4370 customers, in every quarter. The number of customers is the same across quarters for a given carrier, but the sample of customers for each carrier differs between quarters (i.e., repeated cross-sectional data).

According to managerial sources from one of the three major U.S. telecom carriers, the objectives of this survey were threefold: (1) gain an understanding of how wireless, internet, and pay TV customers view and rate customer experience with their carrier or provider, (2) determine the driving factors of customer satisfaction, and (3) determine what the focal firm does well and where it falls behind competitors, according to not only customer satisfaction (i.e., net promoter score) but also specific drivers and attributes.

3.1. Inputs and Outputs

The questions in the survey data fall into three categories, representing three different goals. The first group of questions provide customer characteristics; their characteristics have a predictive function, because they cannot be manipulated by managers (e.g., age and ethnicity cannot be influenced by any marketing action). The second group of perceptual questions offer immediate strategic value to managers, in that they ask customers about their feeling toward competing carriers. The third group of questions relates directly to customer satisfaction and form the basis for the proposed digital twin approach, because I assume that managers aim to maximize customer satisfaction. Therefore, the survey questions reflect three categories:

- **Predictive Variables:** during the inference phase, and at test time, these variables are fixed. In the optimization phase, they remain fixed. A key assumption of the model is that invariant predictors completely characterize customer heterogeneity. For the empirical application, I use large numbers of socio-demographic and usage questions, including age, gender, race or ethnicity, annual household income before taxes, devices at home, name of the wireless service provider, type of plan, tenure with provider, dollar amount paid per month for the plan, data usage, 5G usage, and rewards program.
- **Strategic Variables:** at training time and test time, these variables are fixed. At optimization time, these variables are the arguments of the optimization problem, and are assumed to be manipulated by the marketing analyst. They include:
 - Importance of any of the following according to likelihood to recommend: network in rating, price / value; billing; customer service; general feeling; plans; rewards and benefits; other factors.
 - Satisfaction with network speed; network reliability; data plans that meet my needs; value of price paid; accuracy of billing; rewards and recognition; ease of doing business; solving problems for the first time; “brand for me”; total cost of wireless service; device selection.
- **Target Variables:** these variables are reconstructed at inference time, and predicted at test time, and optimized at optimization time. They include¹:
 - Likelihood to recommend (LTR) (0-10);
 - Likelihood to recommend current provider’s phone to a friend or a colleague (Phone LTR) (0-10);
 - Likelihood to switch wireless service providers within the next 12 months (Intention to Switch) (0-4);
 - Overall satisfaction with current provider (0-9);

¹ The complete list of strategic and target variables is available in Appendix A; because there are more than 300 one-hot encoded predictors, they are not listed here. The complete list remains available upon request.

- Overall feeling about current provider (0-4);
- Overall feeling about competitions' providers² (0-4).

The model also controls for different aspects of firm performances, using Generally Accepted Accounting Principles (GAAP) and non-GAAP measures published quarterly by AT&T, T-Mobile, and Verizon between the second quarter of 2020 and the fourth quarter of 2022. For brevity, I denote these variables as business key performance indicators (KPIs). They include total revenue, operating revenue, cost of revenue, gross profit, operating expense, churn, and average revenue per user (ARPU). The measures are standardized. Tables 4 and 5 (in the Appendix) lists all questions included as target variables and strategic variables. Figure AP.1 includes summary statistics at the question level, per carrier.

3.2. Multivariate Analysis

Before providing a generative model of the target variables (LTR, Phone LTR, Satisfaction, Overall Feeling about Carrier {1, 2, 3}, Intention to Switch), it is helpful to understand how they are associated with one another. Therefore, I undertake a correlation analysis of the target variables, at the carrier level.

First, I recode the Intention to Switch as Retention Likelihood, applying the formula $f(x) = 4 - x$. Intuitively, Satisfaction, Likelihood to Recommend and Overall Feeling about Own Provider should correlate positively with Retention Likelihood; an empirical analysis verifies these correlations (Figure 1). For each carrier, the correlation between retention likelihood and satisfaction ranges from 0.28 to 0.31. The correlation between Retention Likelihood and Feeling about Current Provider also is positive (respectively, 0.50, 0.43, and 0.46 for Carriers 1, 2, and 3). Phone LTR is also positively correlated with Retention Likelihood. The strong positive correlations between LTR and Feeling about Own Provider and Satisfaction suggest that marketers at least indirectly capture a measure of satisfaction when they record the popular Net Promoter Scores³ (Reichheld 2003).

It is more challenging to understand the relationship between Feeling about Competitions' Providers and other target variables. More positive feelings are associated with a lower Retention Likelihood (correlations from -0.08 to -0.18, Figure 1). However, more positive feelings are also associated with higher LTR, Phone LTR, and Satisfaction, suggesting that customers may simply be "happier" about the telecommunications industry in general. This suggestion is corroborated

² For example, an AT&T customer is asked about their overall feeling about T-Mobile and Verizon, as two separate questions.

³ The Net Promoter Score is the measure that transforms LTR by assigning a +1 to respondents who indicate a LTR of 0 to 6, 0 to those who indicate a LTR between 7 and 8, and +1 to those who provide a LTR of 9 and 10, then taking the average across all respondents.

by the slightly positive correlation between all Feeling measures about Own and Competitor’s providers.

Finally, it is interesting to notice symmetries and asymmetries between the three carriers in terms of correlations across target variables. In terms of symmetries, customers from all carriers tend to have more positive or more negative feeling about all competitions’ carriers simultaneously; moreover, the variable that correlates most with Retention Likelihood is Feeling about Own carrier, whereas Overall Satisfaction comes second. In terms of asymmetries, customers from Carrier 3 do not express positive or negative associations with their Feeling about Own Carriers and Competitions’ Carriers (correlations of 0.01 and 0.03, in Figure 1) whereas customers from Carrier 1 and especially Carrier 2 tend to have a stronger associations (correlations of 0.14 and 0.16 for Carrier 2, Figure 1).

3.3. Linear Modeling is Limited for Analyzing Brand Surveys

The first step in gauging whether a nonlinear model is needed to investigate the relationship between explanatory variables (invariant predictors) is to compare two simple discriminative machine learning models. For simplicity, I use multiple output linear regression as a benchmark for the discriminative linear model, and multiple output random forest as a benchmark for the discriminative nonlinear model.

Multiple output linear regression is a generalization of the simple linear regression model when more than one output variables is considered. The model learns a linear relationship between the input variables and each of the output variables. Each output variable is modeled as a linear combination of the input variables plus an error term. In contrast, random forests are ensemble learning methods that operate by constructing a multitude of decision trees and outputting the mean prediction for a regression task. A multiple output random forest is an extension of this technique to handle multiple output variables. This method is useful when the output variables are not independent of each other and share some correlation, as indicated in Subsection 3.2.

Fitting the multiple-output linear regression and the multiple-output random forest on the pooled data shows that the coefficient of determination score R^2 is higher for all target variables in the random forest analysis (Table 2). This higher goodness-of-fit metric indicates that the nonlinear model better captures the relationships between input (invariant and strategic) variables and target variables.

Although linear regression and random forest are useful models, they are fundamentally predictive and do not necessarily provide meaningful interpretations of the relationships between inputs and outputs. For counterfactual reasoning, other techniques may be more desirable, such as structural equation modeling (SEM). Yet, even if SEM may be more interpretable, it also yields

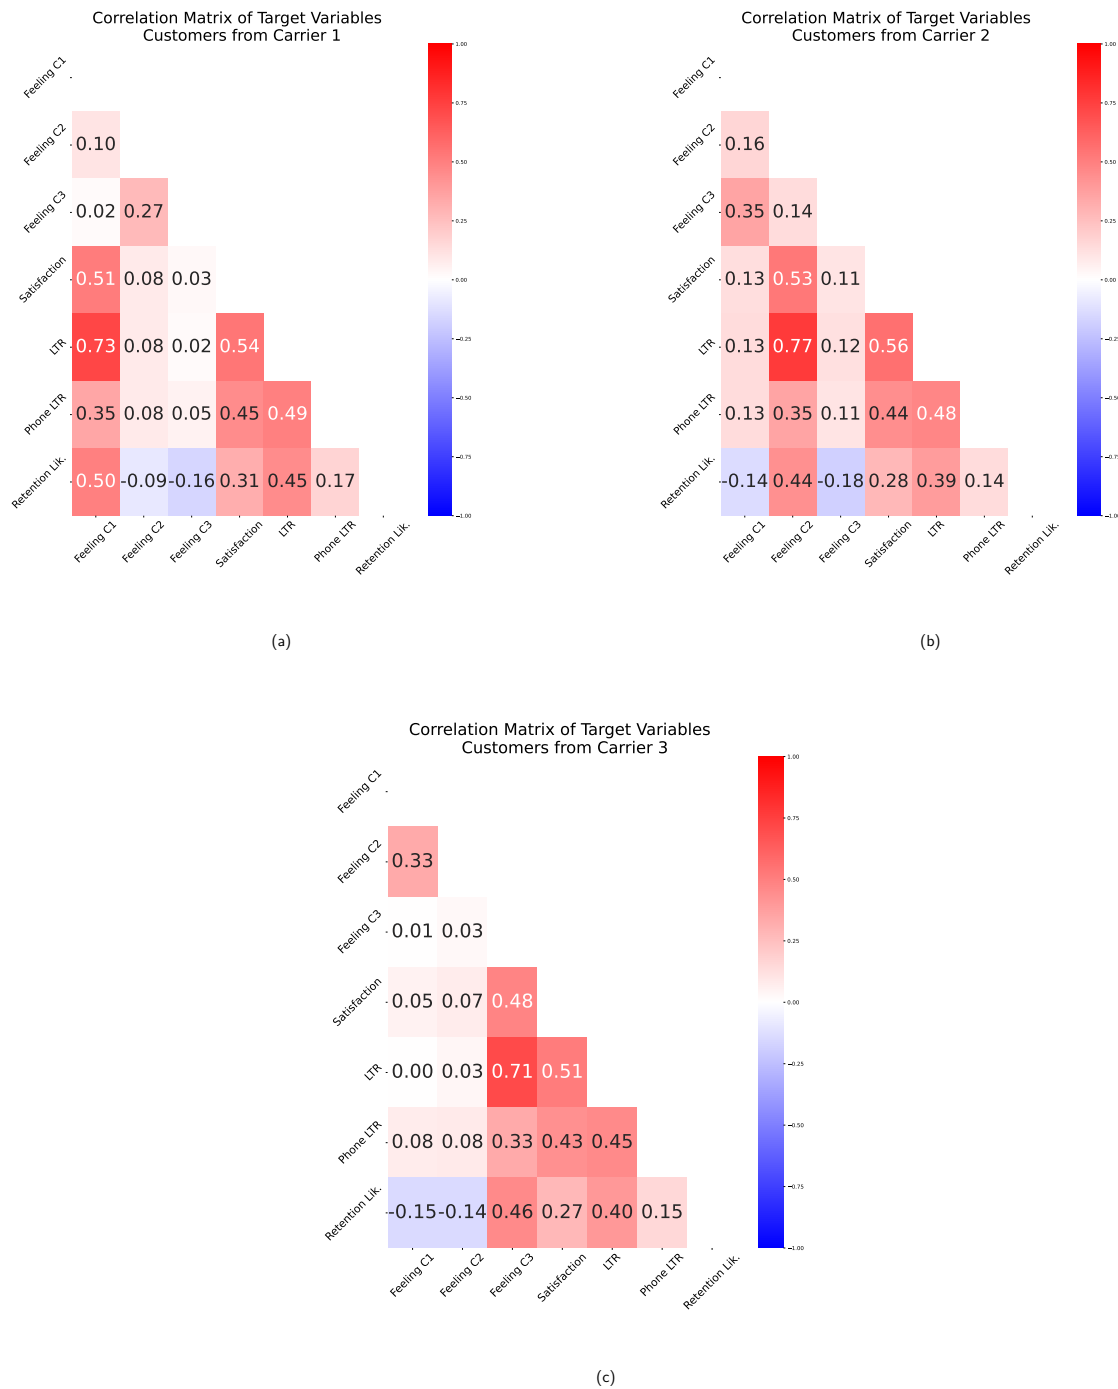


Figure 1 Correlation matrices for target variables in the survey, for each carrier.
 Notes: Feeling C1, Feeling C2, and Feeling C3 respectively refer to Overall Feeling about Carrier 1 (a), Carrier 2 (b), and Carrier 3 (c), respectively.

questionable assumptions, because it is difficult to know *a priori* how the different input variables relate to latent constructs of interest that explain the various target variables that managers care about and try to optimize.

Data Source	Multiple Output Discriminative Models	
	Linear Regression (Test R^2)	Random Forest (Test R^2)
Q12 (Likelihood-to-recommend, LTR)	0.78	0.82
Q27 (Phone LTR)	0.30	0.42
Q18 (Satisfaction)	0.58	0.67
Q20Ar1 (Overall Feeling Carrier 1)	0.38	0.58
Q20Ar2 (Overall Feeling Carrier 2)	0.40	0.57
Q20Ar3 (Overall Feeling Carrier 3)	0.39	0.55
Q20 (Intention to Switch)	0.52	0.60

Table 2 Goodness-of-fit of two discriminative models on a test sample of the data.

Notes: Multi-output Linear regression is a standard benchmark, while multi-output random forest is a nonlinear benchmark. A nonlinear model better explains the variation in the data.

4. Modeling Framework

Subsection 4.1 details the model architecture and training, based on mapping customer-side and firm-side input variables to latent variables using amortized neural networks. After documenting the latent variables and their marketing interpretations (Subsection 4.2), Subsection 4.3 presents the model layer for ordered categorical variables. Subsection 4.4 details the inference procedure, with implementation details. Section 5 outlines the predictions tasks, and Section 6 offers a description of the optimization phase.

4.1. Mapping Customer-Side and Firm-Side Input Variables to Latent Parameters through Amortized Neural Networks

The survey data refer to N respondents and K carriers. In a single quarter t , for a given firm k , a subset N_{kt} of these N respondents are surveyed, such that $\sum_{k=1}^K \sum_{t=1}^T N_{kt} = N$. The number of respondents within a firm remains constant over time, but individuals are not surveyed more than once, such that $N_{k1} = \dots = N_{kT}$ for all $k = 1, \dots, K$. Throughout this paper, $K = 3$, referring to the three largest carriers: AT&T, T-Mobile, and Verizon⁴. The survey data has J questions. I label the J^{pred} predictive variables \mathbf{x}_{it}^{pred} , the J^{str} strategic variables \mathbf{x}_{it}^{str} and the $J^{targets}$ targets \mathbf{y}_{ijtk} , where $i = 1, \dots, N$ customers, $j = 1, \dots, J$ questions, $k = 1, 2, 3$ firms. In summary, $J^{str} + J^{pred} + J^{targets} = J$. The training phase does not make a conceptual distinction between invariant predictors and

⁴ Sprint was also a major carrier but merged with T-Mobile U.S. on April 1, 2020. It was the fourth-largest telecommunications carrier in the United States before the merger. Since the data starts in Q2 2020, I do not consider Sprint customers in the analysis.

strategic variables, because they enter the same neural network. Therefore, $\mathbf{x}_{it} = [\mathbf{x}_{it}^{str}, \mathbf{x}_{it}^{pred}]^T$. For a given quarter t and a given firm k , business KPI are \mathbf{x}_{kt}^{KPI} , which is a vector of size H .

Amortized inference refers to inference over variational parameters that are parameterized by a function of the data, instead of approximating separate variables for each data point (Zhang et al. 2018). For this research, the parameterized function is the neural network $f(\cdot)$ that represents the parameters of the variational distribution across all data points from the J questions in the survey. An alternative would be to separately learn a set of parameters for each data point, rather than learning a set of mean and location parameters for each customer, at each time period, and each firm. The word ‘‘amortized’’ herein means that the cost of learning the variational parameters is ‘‘amortized’’ over all the data points.

Amortized inference is a powerful way to infer the posterior over customer-level and firm-level latent variables according to \mathbf{x}^{KPI} , \mathbf{x}^{str} and \mathbf{x}^{pred} . Using variational inference to approximate the posterior distribution of customer-side and firm-side latent variables implies replacing the locational variational parameters with a function of the data where parameters – weights and biases of neural networks – are shared across all data points, for all firms and at all quarters. The neural network parameters automatically learn a complex representation of the inputs across firms and over time, and this representation is mapped to latent variables that are the building blocks of the target variables. As a major methodological contribution, the current study proposes amortized inference as a way to augment repeated cross-sectional data.

The feed-forward neural networks $f(\cdot)$ and $g(\cdot)$ map customer-side and firm-side, respectively, input variables to a set of latent location and scale parameters that generatively model the target variables. In such feed-forward neural networks, hidden layers are dense and sequentially connected. Consider the feed-forward neural network function $f(\mathbf{x}; \theta_f)$ with D hidden layers; is detailed as follows. The input layer is $d = 0$, hidden layers are $d = 1, 2, \dots$, and the output layer is D . The weights connecting layer d and layer $d + 1$ can be referred to as $W^{(d)}$, and the biases in layer $d + 1$ are indicated by $b^{(d)}$. The pre-activation at layer $d + 1$ can be denoted as $a^{(d+1)}$, and the post-activation is $h^{(d+1)}$. The activation function is the hyperbolic tangent (\tanh), such that:

$$\mathbf{a}^{(1)} = W^{(0)}\mathbf{x}_{it} + \mathbf{b}^{(0)} \tag{1}$$

$$\mathbf{h}^{(1)} = \tanh(\mathbf{a}^{(1)}) \tag{2}$$

$$\mathbf{a}^{(2)} = W^{(1)}\mathbf{h}^{(1)} + \mathbf{b}^{(1)} \tag{3}$$

⋮

$$\mathbf{a}^{(D)} = W^{(D-1)}\mathbf{h}^{(D-1)} + \mathbf{b}^{(D-1)} \tag{4}$$

$$\mathbf{h}^{(D)} = \tanh(\mathbf{a}^{(D)}) \tag{5}$$

$$\mu_{ikt} = W_{\mu}^{(D)} \mathbf{h}^{(D)} + \mathbf{b}_{\mu}^{(D)} \quad (6)$$

$$\nu_{ikt} = \exp \left(W_{\nu}^{(D)} \mathbf{h}^{(D)} + \mathbf{b}_{\nu}^{(D)} \right) \quad (7)$$

Here, \mathbf{x}_{it} is the batch input to the network. Because ν_{ikt} is a variance and must be non-negative, I apply an exponential function to obtain it from $\mathbf{a}^{(D)}$. The weights and biases (collectively referred to as θ_f) are learned by training the network. These weights and biases are parameters of amortized neural networks.

For the feed-forward neural network $g(\mathbf{x}^{KPI}; \theta_g)$ with D' hidden layers, the inputs are the KPI for the three major carriers in the U.S. market (AT&T, T-Mobile, and Verizon), published quarterly over the corresponding 10 quarters of survey data. The neural network's output at a batch level is a concentration parameter γ_{ktl} and a rate parameter ω_{ktl} . The function g also relies on amortization to learn a shared representation across all quarters and firms instead of learning individual γ_{ktl} and ω_{ktl} . The dimension l refers to a set of L latent dimensions summarizing the various aspects of performance across firms and over time. These L latent dimensions provide dimensionality reduction, similar to principal components in principal component analysis.

Finally, the use of the *tanh* activation function introduces non-linearities between layers, allowing the network to learn complex mappings from inputs to outputs. The *tanh* function is particularly well-suited to the empirical application, due to its differentiability and its output range of -1 to 1 , which helps with the normalization of the outputs.

4.2. Interpreting the Probabilistic Latent Factors Generating Digital Twins

For the latent parameter layer of the digital twin architecture, which includes the latent variables and their prior distributions, recall that i indexes customer identifiers from 1 to N ; k indexes firms from 1 to K ; t indexes time from 1 to T . Customer-side factors include the following latent variables:

- $z_{ikt} \sim \mathcal{N}(\mu_{ikt}, \nu_{ikt})$: this latent factor has a Normal prior distribution. Because μ_{ikt} and ν_{ikt} are functions of an amortized neural network, this prior is highly flexible and encodes a wide range of customer characteristics, automatically accounting for interactions and nonlinearities. This parameter is interpreted as the *customer-side brand affinity factor*; it represents, for a given customer at a given time, their affinity with brand k .
- $\alpha_{jkt} \sim \mathcal{N}(0, 1)$: The parameters α_{jkt} represent the *baseline* for question j for firm k at time t . It has a standard Normal prior distribution for simplicity.
- $\beta_{jl} \sim \mathcal{N}^+(0, 1)$ The parameter β_j is interpreted as the *polarization* of question j in the l -th dimension of service quality, that is, how much question j elicits a response on the l -th service characteristic.

The firm-side factors include the following latent variables:

- $\phi_{ktl} \sim \mathcal{G}(\gamma_{ktl}, \omega_{ktl})$: The parameterization of a prior on ϕ_{ktl} , a firm-side latent factor on dimension l for firm k at time t , has a Gamma prior distribution. Because γ_{ktl} and ν_{ktl} are functions of an amortized neural network, this prior also is highly flexible; it encodes a wide range of firm characteristics, automatically accounting for interactions and nonlinearities.

Support for both ϕ_{ktl} and β_{jt} is the real positive line, for identification. The sign of z_{ikt} becomes then immediately interpretable, as explained in Subsection 4.3.

4.3. Model Layer for Ordered Categorical Outcomes

Because the target variables are all ordered categorical variables, I use an ordered logit specification. Let y_{ijkt}^* denote the latent response of respondent i to the entire set of J questions. Questions have different numbers of scale points: some questions have five scale points (0-4) whereas others have 10. For $M + 1$ common and ordered cut points $\{c_m : c_{m-1} \leq c_m, m = 1, \dots, M\}$ where $c_0 = -\infty$ and $c_M = +\infty$, latent utility values y_{ijkt}^* depend linearly on α_{jkt} , which are baseline values for question j at firm k ; the customer-side brand affinity z_{ikt} ; the polarization of question j in firm-side latent dimension l , β_{jl} ; and the firm-side factors ϕ_{ktl} :

$$y_{ijkt}^* = \alpha_{jkt} + z_{ikt} \sum_{l=1}^L (\beta_j \phi_{kt})_l + \varepsilon_{ijkt} \quad \text{where } \varepsilon_{ijkt} \underset{i.i.d.}{\sim} EV(0, 1) \quad (8)$$

The individual responses y_{ijkt} for a customer $i = 1, \dots, N$ of firm $k = 1, \dots, K$ at question $j = 1, \dots, J_{targets}$, at time $t = 1, \dots, T$ take the following values: $m = 1, 2, \dots, M$ where M is the maximum number of scale points for question j . Because questions differ in their total number of scale points, m and M should have a subscript j , but I omit it for simplicity. The following holds:

$$y_{ijkt} = m \quad \text{if } c_{j,m-1} \leq y_{ijkt}^* \leq c_{j,m} \quad (9)$$

where a Dirichlet prior model applies to ordinal probabilities, which serves to induce cut points indirectly. This approach enables a proper, principled prior on the cut points, which is useful when some categories are not strongly separated due to their data sparsity in some categories (Betancourt 2020).

By marginalizing out the latent utilities y_{ijkt}^* , it is possible to write the probability of observing category m for question j in customer i of firm k at time t :

$$p(y_{ijkt} | c_{j,1}, \dots, c_{j,M}) = \begin{cases} \Pi(c_{j,1} - \alpha_{jkt} - z_{ikt} \sum_{l=1}^L (\beta_j \phi_{kt})_l) & \text{if } m = 1 \\ \Pi(c_{j,m} - \alpha_{jkt} - z_{ikt} \sum_{l=1}^L (\beta_j \phi_{kt})_l) \\ - \Pi(c_{j,m-1} - \alpha_{jkt} - z_{ikt} \sum_{l=1}^L (\beta_j \phi_{kt})_l) & \text{if } 1 < m < M \\ 1 - \Pi(c_{j,m-1} - \alpha_{jkt} - z_{ikt} \sum_{l=1}^L (\beta_j \phi_{kt})_l) & \text{if } m = M \end{cases} \quad (10)$$

where $\Pi(\cdot)$ is the cumulative distribution function of the Type I extreme value distribution, that is, the logistic function.

4.4. Inference and Implementation

The set of latent variables to infer is $\tilde{\mathbf{z}} = [\mathbf{z}, \boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\phi}, \boldsymbol{\mu}, \boldsymbol{\nu}, \boldsymbol{\gamma}, \boldsymbol{\omega}]$. The set of parameters to learn is $\boldsymbol{\theta} = [\boldsymbol{\theta}_f, \boldsymbol{\theta}_g]$. By writing \mathbf{y} as the set of all observations (survey data and KPI), it is possible to approximate the posterior distribution $p_{\boldsymbol{\theta}}(\tilde{\mathbf{z}}|\mathbf{y})$.

Because of the size of the data, and the use of neural networks to parameterize the latent variables, exact inference (e.g., Markov chain Monte Carlo algorithms) is not feasible. Therefore, it is necessary to rely on approximate Bayesian inference; stochastic variational inference (SVI) aims at determining a variational distribution $q_{\boldsymbol{\lambda}}(\tilde{\mathbf{z}})$ that is as close as possible to the posterior $p_{\boldsymbol{\theta}}(\tilde{\mathbf{z}}|\mathbf{y})$ as measured by Kullback-Leibler (KL) divergence. Minimizing the KL divergence is equivalent to maximizing the evidence lower bound (ELBO) on the log marginal probability of the data $\log p_{\boldsymbol{\theta}}(\mathbf{y})$, with $\log p_{\boldsymbol{\theta}}(\mathbf{y}) \geq \text{ELBO}$ and $\log p_{\boldsymbol{\theta}}(\mathbf{y}) - \text{ELBO} = \text{KL}(q_{\boldsymbol{\lambda}}(\tilde{\mathbf{z}})||p_{\boldsymbol{\theta}}(\tilde{\mathbf{z}}|\mathbf{y}))$.

The evidence lower bound (ELBO) is:

$$\mathcal{L}(\boldsymbol{\lambda}) = \mathbb{E}_{q_{\boldsymbol{\lambda}}(\tilde{\mathbf{z}})} [\log p_{\boldsymbol{\theta}}(\mathbf{y}, \tilde{\mathbf{z}})] - \mathbb{E}_{q_{\boldsymbol{\lambda}}(\tilde{\mathbf{z}})} [\log q_{\boldsymbol{\lambda}}(\tilde{\mathbf{z}})] \quad (11)$$

The ELBO creates two expectations with respect to the variational distribution. The first expectation, $\mathbb{E}_{q_{\boldsymbol{\lambda}}(\tilde{\mathbf{z}})} [\log p_{\boldsymbol{\theta}}(\mathbf{y}, \tilde{\mathbf{z}})]$, represents the expected log-likelihood of the data given the model parameters, which encourages densities that place their mass on configurations of the latent variables that explain the observed data (Blei et al. 2017). The second expectation, $\mathbb{E}_{q_{\boldsymbol{\lambda}}(\tilde{\mathbf{z}})} [\log q_{\boldsymbol{\lambda}}(\tilde{\mathbf{z}})]$, is the negative divergence between the variational density and the prior. Maximizing the ELBO is akin to finding a balance between encouraging the model to fit the data well (maximizing the first term) and encouraging densities close to the prior (maximizing the second term).

In line with standard practice, this study uses mean-field variational approximation. The model implementations relies on the machine learning framework Google JAX for fast computation on Graphics Processing Unit (GPU), and the NumPyro probabilistic programming language (Phan et al. 2019). With the Adam optimization algorithm (Kingma and Ba 2014), a Monte Carlo version of the loss function is optimized in Equation (11) and a test set is used to determine all model hyperparameters, namely, the number of hidden layers per neural network, number of hidden units per neural network, and number L of latent firm-side dimensions.

4.5. Digital Twin Generative Process

To summarize, the specification is such that for all $m = 1, \dots, M$, $i = 1, \dots, N$, $j = 1, \dots, J$, $k = 1, \dots, K$ and $t = 1, \dots, T$:

$$\begin{bmatrix} \mu_{ikt} \\ \nu_{ikt} \end{bmatrix} = f(\mathbf{x}_{it}, \boldsymbol{\theta}_f) \quad \text{where } f \text{ is a feed-forward neural network}$$

$$\begin{aligned}
 \begin{bmatrix} \gamma_{kt} \\ \omega_{kt} \end{bmatrix} &= g(\mathbf{x}_{kt}^{(\text{KPI})}, \boldsymbol{\theta}_g) && \text{where } g \text{ is a feed-forward neural network} \\
 \beta_j &\sim \mathcal{N}^+(0, 1) \\
 z_{ikt} &\sim \mathcal{N}(\mu_{ikt}, \nu_{ikt}) \\
 \phi_{ktl} &\sim \text{Gamma}(\gamma_{ktl}, \omega_{ktl}) \\
 y_{ijkt}^* &= \alpha_{jkt} + z_{ikt} \sum_{l=1}^L (\beta_j \phi_{kt})_l + \varepsilon_{ijkt} && \text{where } \varepsilon_{ijkt} \underset{i.i.d.}{\sim} EV(0, 1) \\
 \pi(\mathbf{c}_j | \kappa, \lambda) &= \mathcal{D}(p(\mathbf{c}_j, \lambda) | \kappa) \cdot |J(\mathbf{c}_j, \lambda)| \\
 y_{ijkt} &= m && \text{if } c_{j,m-1} \leq y_{ijkt}^* \leq c_{j,m}
 \end{aligned}$$

where \mathcal{D} is the Dirichlet probability density function, and J is the Jacobian matrix. A uniform Dirichlet prior $\kappa = (1, 1, \dots, 1)$ and the anchor point $\lambda = 0$ identify the model, without loss of generality. Figure 2 provides an illustration of the full digital twin architecture. The hybrid deep learning and probabilistic generative framework allow the best of both worlds: high flexibility and representation power of the amortized neural network on the left hand side of the illustration, in orange, and high interpretability and theory-based for counterfactual reasoning on the right hand side, in green.

4.6. Relation to Prior Literature

4.6.1. Variational Autoencoders My model is novel in its use of customer-level predictors and strategic variables that parameterize an amortized neural network for high flexibility, that output structured latent variables that can be subsequently interpreted by marketers, and generate a coherent model of customer satisfaction. However, neither the use of amortized neural networks nor the use of Bayesian optimization in marketing is novel.

The model resembles a Variational Autoencoder (VAE) (Kingma and Welling 2013), which is also a generative model that also uses variational inference for learning. The differences lie primarily in the specific structure of the model and the form of the decoder. In the VAE, the encoder is a neural network that takes the observed data as inputs, then outputs parameters of a distribution over the latent variables. The proposed model has two such “encoders”, f and g , each of which produces parameters for different distributions over subsets of the latent variables, such that f encodes μ_{ikt} and ν_{ikt} for customer-side brand affinity z_{ikt} , and g encodes γ_{kt} and ω_{kt} for firm-side factors ϕ_{kt} .

In a VAE, the latent variables capture unobserved factors of variation in the data, whereas z_{ikt} and ϕ_{kt} capture observed factors of variation in the data, because they are parameterized by survey and KPI inputs. The only unobserved factors of variations come through the type I extreme value that affects latent utilities.

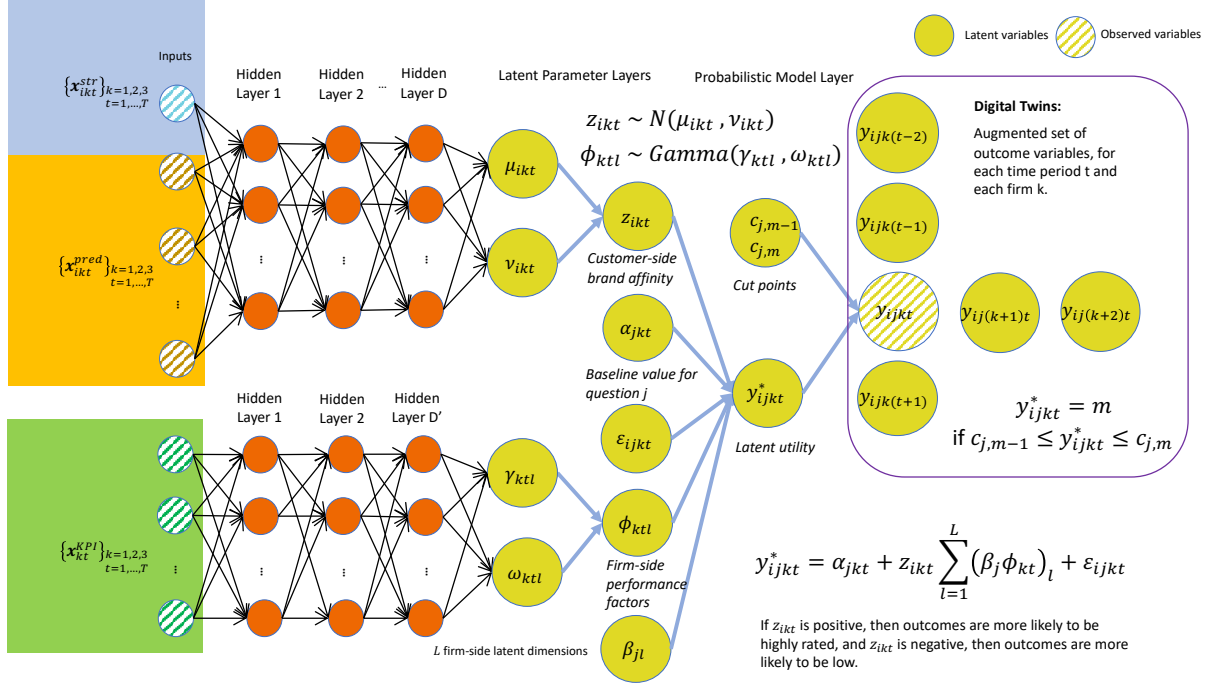


Figure 2 A deep probabilistic architecture of the modeling framework. An amortized inference neural networks (in orange, top) take survey data question as inputs and parameterize latent customer-side brand affinity z_{ikt} . Another amortized inference neural network (in orange, bottom) takes KPI data as input and parameterize firm-side performance factors ϕ_{ktl} . These two types of latent variables are then combined in the probabilistic model layer according to the Equation (8) as a latent customer-level utility. This latent utility is then evaluated against latent question-level cut points to present ordered categorical variables for target customer satisfaction questions.

The decoder in a VAE takes the latent variables and generates parameters for the distribution over the observed data. The equations involving y_{ijkt} and y_{ijkt}^* can be interpreted as part of a kind of decoder that uses the latent variables, together with an ordered logit model, gives a distribution over the observed variable y_{ijkt} . However, unlike a VAE, this decoder does not involve a neural network but is determined by an ordered logit model and a latent factor model that decomposes firm-side and customer-side effects.

4.6.2. Bayesian Models in Political Science The proposed model connects loosely with political science literature, through the notion of Bayesian ideal points.⁵ In political science, a latent factor model quantifies lawmakers' political preferences using roll-call votes (Jackman 2001, Clinton et al. 2004). Lawmakers yea or nay voters on a shared set of bills can be coded as a binary

⁵ Note that the ideal point in political science has a different meaning than in marketing; in marketing, it refers to the hypothetical product attributes or characteristics that a consumer would perceive as perfect, indicating their absolute preference. Market research and product development often use the ideal point concept to tailor offerings to align closely with consumer desires, increase product appeal, and ensure market success.

variable y_{ij} for lawmaker i voting for bill j . Each lawmaker is assumed to have a latent variable z_i , also known as the ideal point, and per-bill latent variables (α_j, β_j) . The vote can be modeled as a Bernoulli distribution $y_{ij} \sim \mathcal{N}(\sigma(\alpha_j + z_i\beta_j))$ where $\sigma(x) = \frac{1}{1+\exp(-x)}$. The customer-side brand affinity functions as an ideal point, such that each customer’s measurement on target variables is akin to rating each company.

5. Model Results

5.1. Fit and Benchmarks

The comparison of the proposed model with three benchmarks confirms its validity, as the goodness-of-fit metrics in Table 3 reveal. The three competing specifications are as follows: Model (1) is a linear version of the proposed model, in which the neural networks have been replaced by a linear layer. It is akin to a traditional SEM, with observed inputs and outputs, but has latent variables parameterizing the relationship between inputs and outputs. A traditional SEM is therefore a special case. Model (2) omits individual predictive and strategic variables from the full specification, but retains individual level customer-side brand affinity z_{ikt} for all $i = 1, \dots, N$, $k = 1, \dots, K$, $t = 1, \dots, T$ as “random effects”. Model (3) omits KPI, but retains firm level performance factors ϕ_{klt} for all $l = 1, \dots, L$, $k = 1, \dots, K$, $t = 1, \dots, T$ as “random effects”.

Figure 3 shows evidence of convergence, in that the average training and test loss (i.e., negative evidence lower bound – ELBO) decrease rapidly before stabilizing. The average test loss is close to the training loss, suggesting good generalizability. The greater variance of the test loss results from the test sample being smaller than the training sample.

In terms of goodness-of-fit, the proposed model (Model (4)) consistently performs better than other models across all carriers and metrics. The average training and test losses are lowest in this model, indicating that it offers the best fit to the data. For example, the Average Training Loss for Model (4) is 11.54, lower than the corresponding values for other models. The same trend is apparent in the Average Test Loss, in which Model (4) outperforms other models with a loss of 12.81. Furthermore, with regard to the Mean Absolute Error (MAE) for each carrier across different tasks, Model (4) generally exhibits the smallest error, suggesting it the most capable of accurately reconstructing the data. Some exceptions involve the Phone LTR and Carrier Satisfaction for Carriers 1 and 2, and the Retention Likelihood for Carrier 3, for which Model (4) does not perform best. However, the overall performance of Model (4) remains superior.

Thus, Table 3 suggests that a neural network model that includes individual predictors and KPI performs best among the presented models across a variety of reconstruction tasks. It also indicates the complexity of the relationships in the data and the ability of neural networks to better capture these complex relationships and extract predictive latent features.

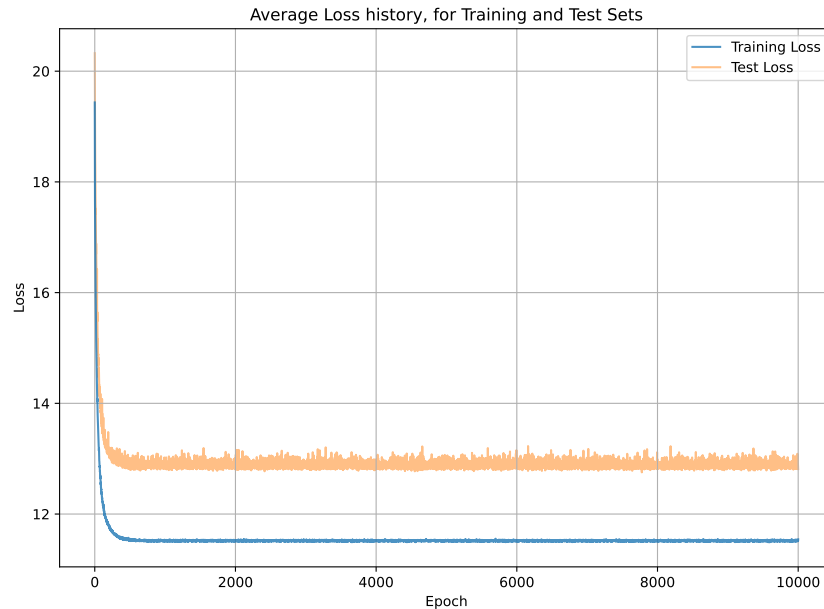


Figure 3 Average training and test loss (ELBO), over all individuals in training and test sets, respectively. **Notes:** The plot suggests that convergence is reached after about 1000 epochs, but the model was trained for 10,000 epochs in total. Test loss is slightly greater than training loss, but remains constant after convergence, as expected.

5.2. Analyzing and Interpreting the Probabilistic Latent Factors

A crucial aspect of the proposed framework is its ability to analyze and interpret the estimated probabilistic latent factors while relaxing the functional form between various predictors and these factors for maximum flexibility.

Figure 4 plots the counterfactual customer-side brand affinities z_{ikt}^* using posterior means across all customers. A darker color signifies higher density. Each dot represents the triplet $(z_{i1t}, z_{i2t}, z_{i3t})$ where $t = 0$, after transformation through a softmax function, which is then converted in a barycentric coordinate system to obtain values that live in a three-dimensional simplex. For example, a customer in Carrier 2 with a dot near the “Carrier 1” vertex reflects that the brand affinity level that customer would obtain if they were assigned to Carrier 1, i.e., the *digital twin* of that customer under the counterfactual regime that this customer’s carrier is now Carrier 1.

One interesting phenomenon to notice is the higher density of customers toward Carrier 1, for all carriers’ customer bases. This digital twin representation suggests a large group of customers who would have a high brand affinity with Carrier 1, if they were ever assigned to be its customers. This latent asymmetry in brand affinity could not be identified without a rigorous counterfactual analytical framework. Carrier 1 likely should target this group of prospective customers to steal them from Carrier 2’s and Carrier 3’s customer bases.

Models		(1)	(2)	(3)	(4)
ϕ		Linear	NN	w/o KPI	NN
z		Linear	w/o indiv. predictors	NN	NN
No. epochs		10000	10000	10000	10000
Avg. Training Loss		12.22	14.06	11.57	11.54
Avg. Test Loss		13.49	15.81	12.94	12.81
Test Mean Absolute Error (MAE)					
Carrier 1	LTR	0.84	1.59	0.67	0.64
	Phone LTR	1.39	1.57	1.24	1.27
	Carrier Satisfaction	1.10	1.51	0.95	1.00
	Overall Feeling Carrier 1	0.53	0.71	0.49	0.42
	Overall Feeling Carrier 2	0.68	0.61	0.68	0.64
	Overall Feeling Carrier 3	0.55	0.47	0.56	0.50
	Retention Likelihood	0.74	0.90	0.69	0.71
Carrier 2	LTR	0.88	1.75	0.69	0.67
	Phone LTR	1.31	1.48	1.15	1.17
	Carrier Satisfaction	1.11	1.53	0.97	1.01
	Overall Feeling Carrier 1	0.70	0.70	0.69	0.68
	Overall Feeling Carrier 2	0.51	0.75	0.47	0.43
	Overall Feeling Carrier 3	0.61	0.57	0.60	0.57
	Retention Likelihood	0.85	1.01	0.81	0.82
Carrier 3	LTR	0.84	1.54	0.63	0.62
	Phone LTR	1.35	1.60	1.20	1.21
	Carrier Satisfaction	1.11	1.54	0.95	1.00
	Overall Feeling Carrier 1	0.75	0.69	0.75	0.70
	Overall Feeling Carrier 2	0.76	0.70	0.77	0.71
	Overall Feeling Carrier 3	0.49	0.68	0.47	0.38
	Retention Likelihood	0.74	0.89	0.68	0.69

Table 3 Goodness-of-fit metrics. Training and Testing Loss, and Mean Absolute Error for Reconstruction Tasks. Model (4) is benchmarked against nested versions (1,2,3). Model (1) assumes a linear link between inputs (predictive variables, strategic variables, key performance indicators) and latent variables. Model (2) omits individual predictors. Model (3) omits key performance indicators.

Figure 5 plots baseline values α_{jkt} for each target variable over time, showing posterior mean and 95% credible intervals. These values represent the base utility for a given question, at a given time, and for a given carrier, after accounting for nonparametric variations in individual predictors or KPI. Baseline values for LTR seem to be lower for Carrier 2, though they seem higher in terms of recommending their carrier’s phone. Unsurprisingly, baseline values for Overall Feeling about Carriers 1, 2, and 3 are higher for corresponding customer bases. Finally, baseline values for Retention Likelihood for Carrier 1 and 3 are higher than for Carrier 2; customers of Carrier 2 are more likely to switch to the competition, on average.

Figure 6 plots the carrier performance loading on each target variable over time. The evidence is more mixed, because reflecting considerable uncertainty, as also indicated by the credible interval. This uncertainty is propagated into the model’s predictive performance, explaining why the model

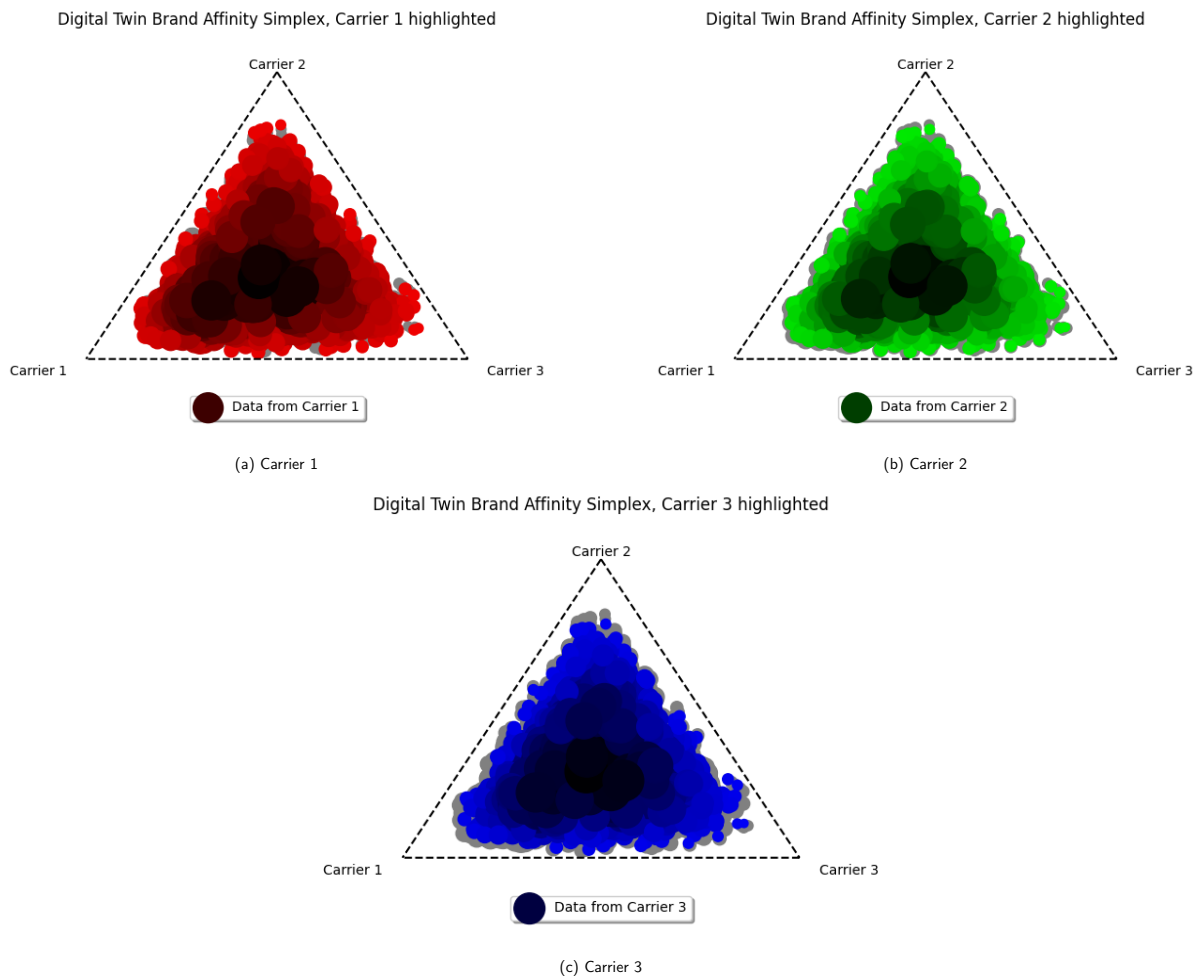


Figure 4 Plotting counterfactual customer-side brand affinities z_{ikt}^* , using posterior means across all customers. **Notes:** A higher density is signified by darker colors. Brand affinities have been transformed using a softmax to fit into a simplex. Each dot represents a customer from a given carrier, and can be projected onto the edges of the triangle to reveal the manifest digital twins, i.e., counterfactual brand affinities summarizing target variables.

with KPI performs only marginally better than the model without them. An interesting aspect to notice is that carrier performance loading on Overall Feeling about Own Carrier and Competition's Carrier are well informed, as is also reflected in the lower test MAE shown above in Model (4) compared with Model (3).

5.3. Application: Personalized Marketing Communication Campaigns using Grid Search on Strategic Variables

After the model is trained, digital twins can be used to optimize personalized marketing communications campaigns. Certain features of telecommunications services are complex and not well understood or known by customers. The proposed framework allows marketers to automatically

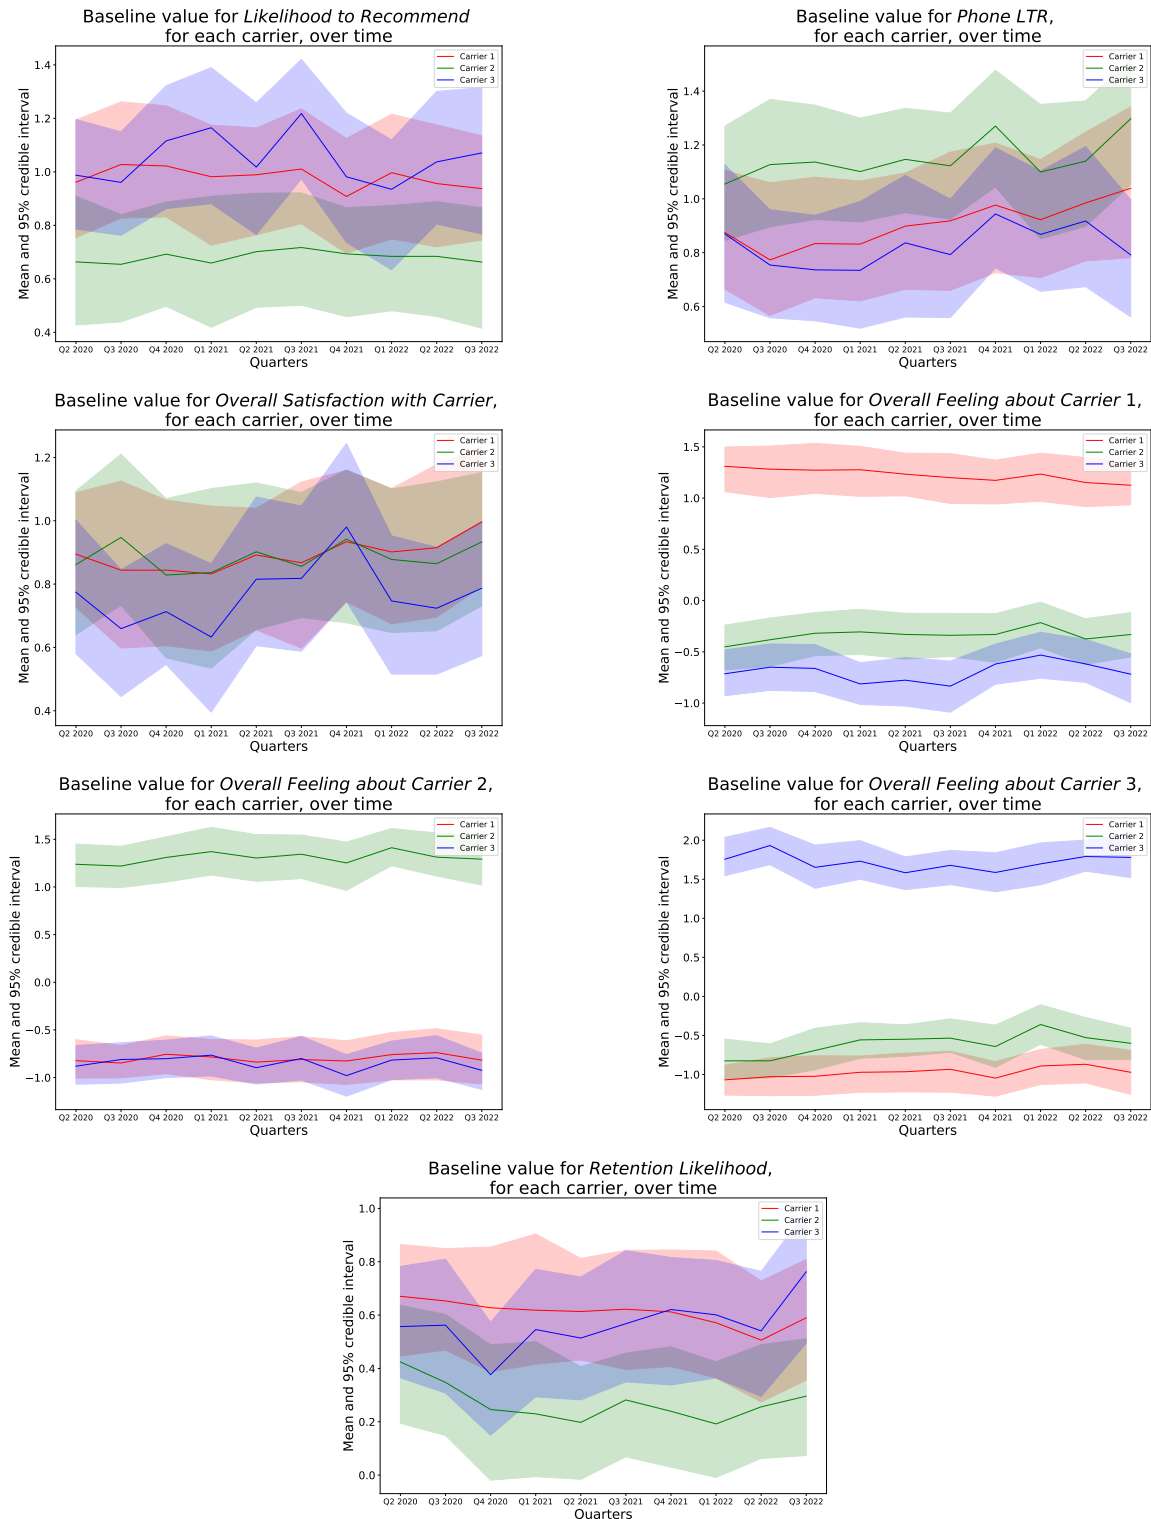


Figure 5 Plot of the baseline values α_{jkt} for each target variable over time.

Notes: These values represent the base utility for a given question, at a given time, and for a given carrier, abstracting away from individual predictors or key performance indicators.

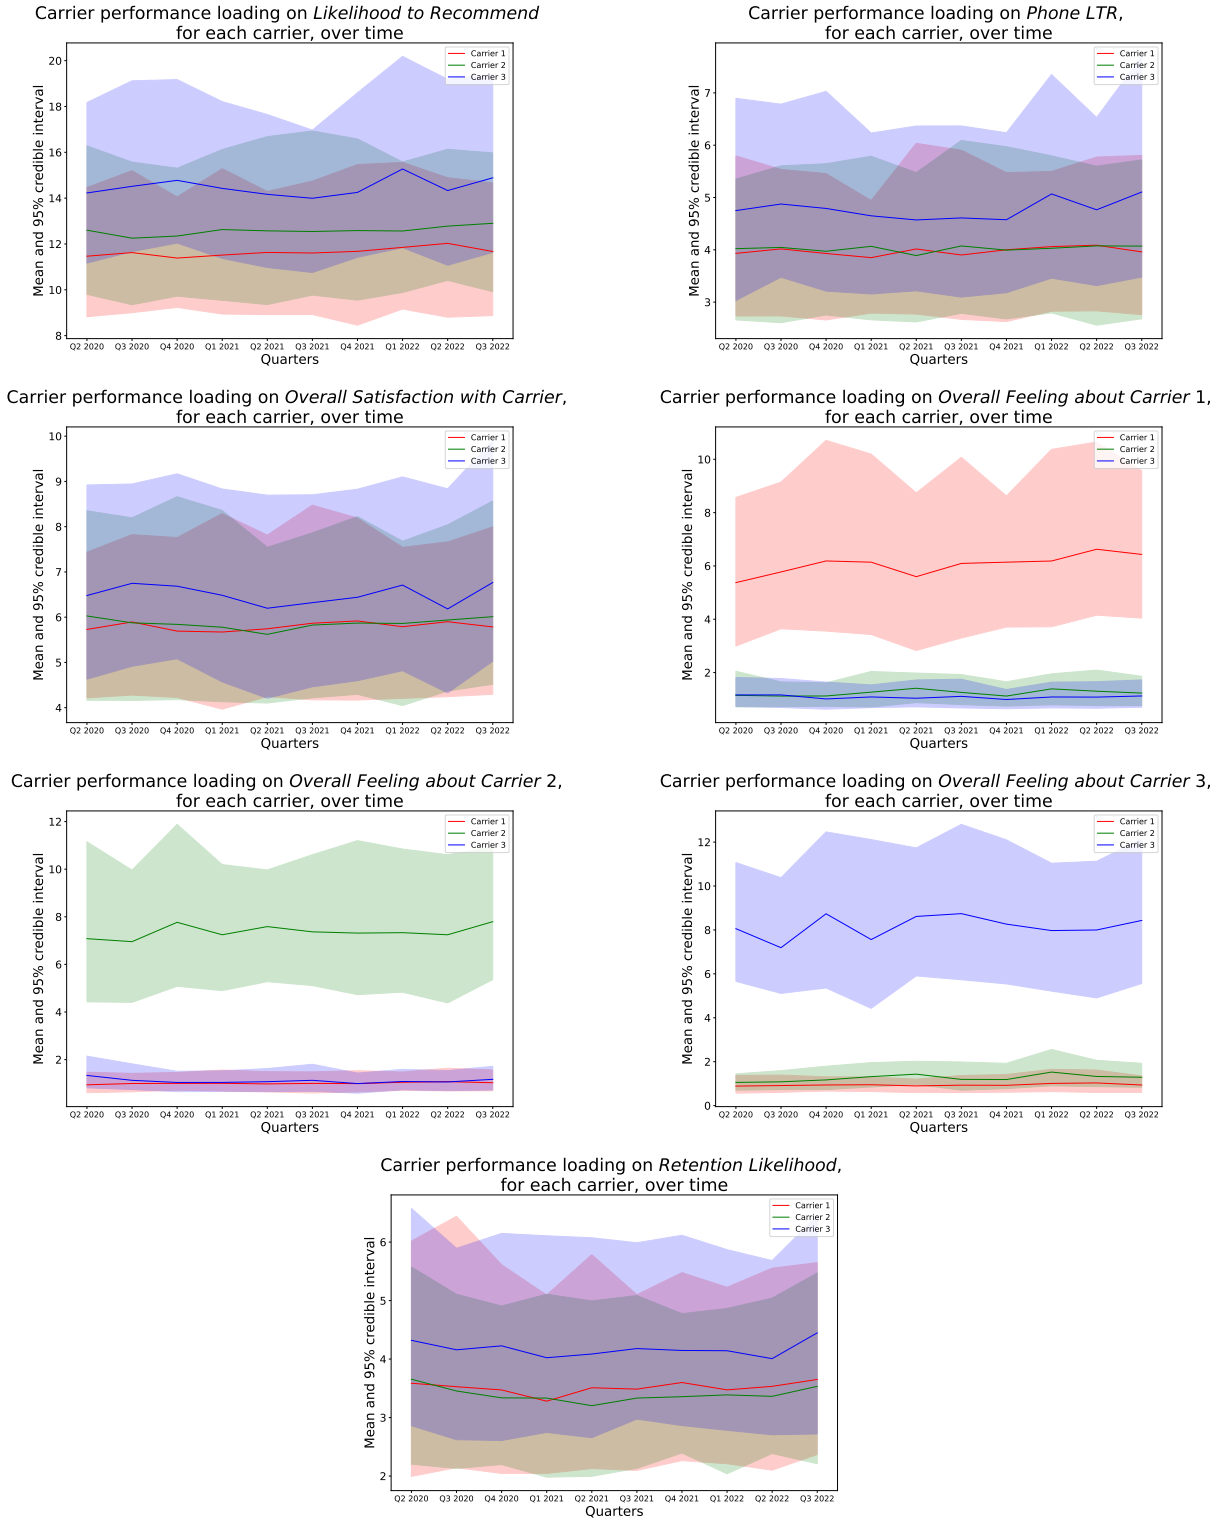
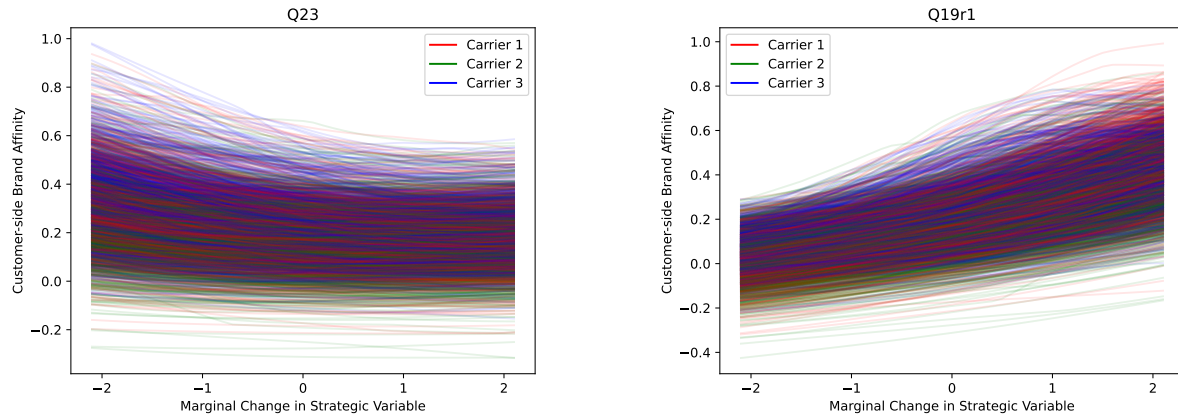


Figure 6 Plot of the carrier k performance loading on each target variable j ($\sum_{i=1}^L \beta_{ji} \phi_{ikt}$) over time t .
 Notes: These values represent the contribution from KPI to target questions. These values are multiplied to brand affinity.



(a) "How likely are you to change anything about your current wireless service in the next 6 months?"

(b) How satisfied are you with your carrier's performance on the following aspect of your overall wireless service experience: Network Speed

Figure 7 Plotting variations in counterfactual customer-side brand affinities z_{ikt}^* in response to marginal changes in strategic variables, between -2 and 2. Each line represents an individual customer's response from a given carrier, in Q2 2020. These plots show evidence of individual-level nonlinearities in brand affinity responses to marginal changes in strategic variables.

rearrange current communication strategies automatically, to focus on the most critical aspects of customer satisfaction, at the individual level.

With a grid search on strategic variables for all customers in the test sample of the dataset, I decrease each customer's current value by 2 points for each strategic variable with at least five point scales, then gradually increase each current value by 0.1 increments, until it reaches +2 points. These changes are simultaneously implemented across all customers in 2020 Q2. Figure 7 plots individual-level responses in brand affinity after changing strategic variables from -2 to 2, using the posterior mean for z_{ikt} as a summary statistic. Carrier 2 seems to have a group of customers with lower brand affinity. Carriers 1 and 3 are remarkably close in terms of brand affinity level. The plots also offer strong evidence of customer-level nonlinearities for two strategic variables. A more positive response by customers to a change in wireless service within the next six months would induce a lower brand affinity, up to a certain point. However, a more positive response by customers to a satisfaction question on their providers' network speed would induce a rapid increase in brand affinity for most customers.

6. Optimizing Customer Satisfaction with Bayesian Optimization

The Digital Marketing Twins model, presented in Section 3, estimates a probabilistic generative model of customer-side brand affinity and firm-side performance factors. The analysis in Section 5.3 showed how posterior inference on the model can be used by managers to undertake campaign personalization, which fits into the broader category of discriminative, or segmentation, based analysis in marketing research. The assumption of partial equilibria in customers' preferences

and perception of brands and their offerings, conditional on firm-side (e.g., marketing) efforts, is fundamental to any applied discriminative marketing analysis; in the context of the proposed model’s inference for the U.S. wireless telecom industry, it is equivalent to assuming that the distribution across the customer base’s brand affinity is invariant, because of ongoing marketing strategies and realizations of capital expenditure.

This assumption is reflected in the *retroactive* and *tactical* nature of personalization campaigns, which, due to the current strengths and weaknesses across the dimensions of an individual’s brand affinity, resulting from customers’ experiences, usages, and interactions, should tailor forthcoming marketing communications to highlight strengths and ameliorate weaknesses. By leveraging the Digital Marketing Twins generative framework, this section highlights its potential as a *proactive* and *strategic* tool, grounded in the statistical perspective that counterfactual reasoning is a missing data problem, to rationalize and optimize a wireless telecom carrier’s marketing strategy. Specific marketing strategy dimensions lead to a more optimal distribution of a customer base’s brand affinity, by taking into consideration not only differences across customers, but also the competitive landscape.

From an optimization perspective, should the generative model be convex and smooth everywhere, then an optimal firm-level marketing policy – realized in this specific exercise as relative emphases across the 33 strategic dimensions – can be found through differential calculus on the model’s functional form. However, two computational challenges arise that require the use of Bayesian optimization: (1) the fitted model specification is almost surely nonconvex, and (2) a brute-force search of the policy grid creates an infeasible $O(n^d)$ time complexity.⁶

In dealing with these two computational challenges, Bayesian optimization is a powerful tool that can resolve black-box global optimization problems, particularly those with expensive function evaluations (Letham and Bakshy 2019). Bayesian optimization combines principles of exploration and exploitation, and uses acquisition functions to navigate the optimizable design space effectively and feasibly. I introduce a novel expected improvement (EI) acquisition likelihood that will enable marketers to identify *paths of least resistance* to improve overall and individual-specific brand affinity within a competitive landscape.

My framework is novel in that it combines a generative model to use with Bayesian optimization for policy search (e.g., Athey et al. (2017)). The Digital Marketing Twins generative model can be viewed as a form of offline simulators (Letham and Bakshy 2019), that elicits the effects of changes to a system more efficiently, which in this case is the brand affinity distribution of the customer

⁶ For example, a grid of 1.0, by 0.1 grid size, would translate into 10^{33} calculations to be evaluated. State-of-art GPUs have 20,000 cores, each of which can (optimistically) evaluate the model to calculate the individual-level brand affinity of all customers in one millisecond. This search thus would still take 1.58 quintillion years to complete.

base. Formally, as an offline simulator in a policy search framework, the Digital Marketing Twins framework allows (1) exploration of brand affinity distributions under different marketing policy configurations (2) computation of both individual- and aggregate- level “rewards” (i.e., increase in brand affinity relative to other brands), and (3) a counterfactual estimate of alternative policies on the same customer or segment (Bottou et al. 2013, Dudík et al. 2014).

The modeling of the counterfactual brand affinity response surface uses the popular Bayesian optimization implementation through a Gaussian process (GP) prior, denoted as:

$$z_{ikt}^* \sim \mathcal{GP}(m(\cdot), K(\cdot, \cdot)) \quad (12)$$

I denote brand affinity evaluations from the GP as z_{ikt}^* to indicate that they are counterfactual outcomes that do not necessarily map onto z_{ikt} inferred from the data used to train the Digital Marketing Twins generative model. For this reason, the GP prior is seen as a distribution over function spaces, and analogously, the goal of Bayesian optimization is to learn a global policy response set and identify feasible optima. The prior is characterized by a mean function $m(\mathbf{x}^{(str)}) = \mathbb{E}[f(\mathbf{x}^{(str)})]$ and a covariance function $K(\mathbf{x}^{(str)}, \mathbf{x}'^{(str)}) = \text{cov}[f(\mathbf{x}^{(str)}), f(\mathbf{x}'^{(str)})]$. The covariance function specifies the covariance between any two points in the design space, which as noted, consists of thirty-three strategic variables. Specifically, I use the ARD-RBF kernel, hyperparameterized by τ (amplitude) and l_j (policy-dimension specific lengthscale). The advantages of the ARD-RBF kernel are that it undertakes variable selection, and it is infinitely differentiable, which enables me to capture complex nonlinear interactions through the policy space:

$$K(\mathbf{x}^{(str)}, \mathbf{x}'^{(str)}) = \tau^2 \exp\left(-\frac{1}{2} \sum_{j=1}^J \left(\frac{\mathbf{x}_j^{(str)} - \mathbf{x}'_j^{(str)}}{l_j}\right)^2\right) \quad (13)$$

In the current study context, the nonparametric GP prior is especially useful as a policy response surface model for Bayesian optimization. First, it provides uncertainty estimates for unobserved points, which are crucial for applying an explore/exploit algorithm. Second, the mean and variance predictions are available in closed form, enabling fast gradient optimization when identifying the next optimal point to test. Third, the smoothness assumption allows efficient exploration of the complex, nonlinear relationships between strategic and target variables from the surveys, which were learned using neural networks during training.

Finally, to choose policies for future evaluation, I use the Expected Improvement (EI) acquisition function that is integrated over the posterior distribution in brand affinity, and rationalized over explore/exploit dynamics in this context. The optimization problem I address aims to maximize an objective $\mathbb{E}[z^*]$, with the constraints $c_j(\mathbf{x}_j^{(str)}) \geq 0$ for $j = 1, \dots, J$; such that $J = 33$, representing the strategic variables to be optimized within feasible constraints – which can be implemented as

+/- 1.0 of the current values. Accordingly, the individual-level improvement at any x^* over the current feasible policy ($\mathbf{x}^{(str)}$) can be expressed as:

$$I_i(\mathbf{x}^{*(str)}, \mathbf{x}^{(str)}) = \max \left\{ 0, \frac{z_{ikt}^* - z_{ikt}}{K(\mathbf{x}^{*(str)}, \mathbf{x}^{(str)})^{-1}} \right\} \mathbb{I} \{ c(\mathbf{x}^{*(str)}) > 0 \} \quad (14)$$

The proposed acquisition function differs from extant improvement functions used in Bayesian optimization through the division of the inverse kernel evaluation, $K(x^*, x)^{-1}$. Intuitively, the denominator term penalizes any policies in the strategic variables that deviate too far from the current policy. Instead, the improvement function rewards policies that give *the paths of least resistance* to be executed by the marketing organization. Taken together, the final EI acquisition function is given by the following Monte Carlo integration (over both customers and the posterior distribution in brand affinity):

$$\alpha_{EI}(\mathbf{x}^{(str)}) = \sum_{i=1}^N E_{p(z)} [I_i(\mathbf{x}^{*(str)}, \mathbf{x}^{(str)})] \quad (15)$$

Through this approach, the Digital Marketing Twins framework can serve as a policy simulator, providing a mechanism to draw from the posterior predictive distribution of a customer profile across time and under different firms, in terms of their target satisfaction variables. This approach offers a proactive and strategic tool for rationalizing and optimizing a carrier’s marketing strategies in a realistic and relevant way, which ultimately should improve customer satisfaction.

7. Conclusion

The novel Digital Marketing Twins methodology promises to address the challenges of analyzing large-scale customer surveys, as demonstrated in the context of the U.S. wireless telecommunications retail market. This study thus addresses two major issues: (1) the theoretical difficulty of integrating customer surveys into a prescriptive framework and (2) the practical problem posed by repeated cross-sectional surveys, such that it contributes to the literature on digital twins, machine learning methods for competitive environments, and customer satisfaction.

The proposed methodology provides counterfactual responses under different scenarios, which can serve as a powerful tool in the realm of customer analytics. The technique also addresses the missing data problem that is typical of repeated cross-sectional surveys, thereby presenting a comprehensive approach to understanding and leveraging customer survey data at scale.

The implementation of the methodology involves the development of a deep generative and probabilistic latent factor model that captures customer-side brand affinity at the individual level, for each brand and each time period, while controlling for observed heterogeneity and firm-side factors. The methodology leverages Bayesian optimization to maximize individual-level, latent, customer-side brand affinity, thereby leading to a “path of least resistance” at the individual level.

The findings have implications for marketers who seek to improve customer satisfaction by understanding the causes of satisfaction from surveys. Furthermore, because the methodology appears generalizable to other sectors and contexts, and therefore, it suggests new avenues for research and applications in the field of marketing and customer analytics.

References

- Ascarza E (2018) Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research* 55(1):80–98.
- Ascarza E, Neslin SA, Netzer O, Anderson Z, Fader PS, Gupta S, Hardie BG, Lemmens A, Libai B, Neal D, et al. (2018) In pursuit of enhanced customer retention management: Review, key issues, and future directions. *Customer Needs and Solutions* 5:65–81.
- Athey S, Wager S, et al. (2017) Efficient policy learning. Technical report.
- Betancourt M (2020) Ordinal regression case study, section 2.2. URL https://betanalpha.github.io/assets/case_studies/ordinal_regression.html, [Online; accessed 30-June-2023].
- Blei DM, Kucukelbir A, McAuliffe JD (2017) Variational inference: A review for statisticians. *Journal of the American statistical Association* 112(518):859–877.
- Bottou L, Peters J, Quiñonero-Candela J, Charles DX, Chikering DM, Portugaly E, Ray D, Simard P, Snelson E (2013) Counterfactual reasoning and learning systems: The example of computational advertising. *Journal of Machine Learning Research* 14(11).
- Braun M, Schweidel DA (2011) Modeling customer lifetimes with multiple causes of churn. *Marketing Science* 30(5):881–902.
- Clinton J, Jackman S, Rivers D (2004) The statistical analysis of roll call data. *American Political Science Review* 98(2):355–370.
- Dudík M, Erhan D, Langford J, Li L (2014) Doubly robust policy evaluation and optimization .
- Fader PS, Hardie BG, Lee KL (2005) RFM and CLV: Using iso-value curves for customer base analysis. *Journal of Marketing Research* 42(4):415–430.
- Gómez-Bombarelli R, Wei JN, Duvenaud D, Hernández-Lobato JM, Sánchez-Lengeling B, Sheberla D, Aguilera-Iparraguirre J, Hirzel TD, Adams RP, Aspuru-Guzik A (2018) Automatic chemical design using a data-driven continuous representation of molecules. *ACS central science* 4(2):268–276.
- Griffiths RR, Hernández-Lobato JM (2020) Constrained Bayesian optimization for automatic chemical design using variational autoencoders. *Chemical science* 11(2):577–586.
- Gustafsson A, Johnson MD, Roos I (2005) The effects of customer satisfaction, relationship commitment dimensions, and triggers on customer retention. *Journal of Marketing* 69(4):210–218.
- Jackman S (2001) Multidimensional analysis of roll call data via Bayesian simulation: Identification, estimation, inference, and model checking. *Political Analysis* 9(3):227–241.
- Kannan P, Wright GP (1991) Modeling and testing structured markets: A nested logit approach. *Marketing Science* 10(1):58–82.
- Kapteyn MG, Pretorius JV, Willcox KE (2021) A probabilistic graphical model foundation for enabling predictive digital twins at scale. *Nature Computational Science* 1(5):337–347.

- Kekre S, Krishnan MS, Srinivasan K (1995) Drivers of customer satisfaction for software products: implications for design and service support. *Management Science* 41(9):1456–1470.
- Kingma DP, Ba J (2014) Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* .
- Kingma DP, Welling M (2013) Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* .
- Lee TY, Bradlow ET (2011) Automated marketing research using online customer reviews. *Journal of Marketing Research* 48(5):881–894.
- Letham B, Bakshy E (2019) Bayesian optimization for policy search via online-offline experimentation. *Journal of Machine Learning Research* 20:145–1.
- Malshe A, Colicev A, Mittal V (2020) How main street drives wall street: Customer (dis) satisfaction, short sellers, and abnormal returns. *Journal of Marketing Research* 57(6):1055–1075.
- Netzer O, Feldman R, Goldenberg J, Fresko M (2012) Mine your own business: Market-structure surveillance through text mining. *Marketing Science* 31(3):521–543.
- Phan D, Pradhan N, Jankowiak M (2019) Composable effects for flexible and accelerated probabilistic programming in numpyro. *arXiv preprint arXiv:1912.11554* .
- Reichheld FF (2003) The one number you need to grow. *Harvard Business Review* 81(12):46–55.
- Ringel DM (2023) Multimarket membership mapping. *Journal of Marketing Research* 60(2):237–262.
- Tirunillai S, Tellis GJ (2014) Mining marketing meaning from online chatter: Strategic brand analysis of big data using latent Dirichlet allocation. *Journal of Marketing Research* 51(4):463–479.
- Tsialiamanis G, Wagg DJ, Dervilis N, Worden K (2021) On generative models as the basis for digital twins. *Data-Centric Engineering* 2:e11.
- Venkatesan R, Kumar V (2004) A customer lifetime value framework for customer selection and resource allocation strategy. *Journal of Marketing* 68(4):106–125.
- Yu J, Song Y, Tang D, Dai J (2021) A digital twin approach based on nonparametric Bayesian network for complex system health monitoring. *Journal of Manufacturing Systems* 58:293–304.
- Zhang C, Bütepage J, Kjellström H, Mandt S (2018) Advances in variational inference. *IEEE transactions on pattern analysis and machine intelligence* 41(8):2008–2026.

Appendix A

Variable Name	Description
Q14R1	[pipe:hCurrentProvider]'s Network - Which of the following directly contributed to you giving [pipe:hCurrentProvider] a rating of ?
Q14R2	[pipe:hCurrentProvider]'s Price / Value - Which of the following directly contributed to you giving [pipe:hCurrentProvider] a rating of ?
Q14R8	[pipe:hCurrentProvider]'s Billing process - Which of the following directly contributed to you giving [pipe:hCurrentProvider] a rating of ?
Q14R3	[pipe:hCurrentProvider]'s Customer Service - Which of the following directly contributed to you giving [pipe:hCurrentProvider] a rating of ?
Q14R4	General feelings about [pipe:hCurrentProvider] - Which of the following directly contributed to you giving [pipe:hCurrentProvider] a rating of ?
Q14R5	[pipe:hCurrentProvider]'s Plans - Which of the following directly contributed to you giving [pipe:hCurrentProvider] a rating of ?
Q14R9	[pipe:hCurrentProvider]'s Rewards and benefits - Which of the following directly contributed to you giving [pipe:hCurrentProvider] a rating of ?
Q14R6	[pipe:hCurrentProvider]'s Devices - Which of the following directly contributed to you giving [pipe:hCurrentProvider] a rating of ?
Q14R7	Other (please specify) - Which of the following directly contributed to you giving [pipe:hCurrentProvider] a rating of ?
Q19R1	Network speed - How satisfied are you with [pipe:hCurrentProvider]'s performance on the following aspects of your overall wireless service experience?
Q19R2	Network reliability - How satisfied are you with [pipe:hCurrentProvider]'s performance on the following aspects of your overall wireless service experience?
Q19R3	Data plans that meet my needs - How satisfied are you with [pipe:hCurrentProvider]'s performance on the following aspects of your overall wireless service experience?
Q19R4	Value for the price paid - How satisfied are you with [pipe:hCurrentProvider]'s performance on the following aspects of your overall wireless service experience?
Q19R5	Accuracy of billing - How satisfied are you with [pipe:hCurrentProvider]'s performance on the following aspects of your overall wireless service experience?
Q19R6	Rewards and recognition - How satisfied are you with [pipe:hCurrentProvider]'s performance on the following aspects of your overall wireless service experience?
Q19R7	Easy to do business with - How satisfied are you with [pipe:hCurrentProvider]'s performance on the following aspects of your overall wireless service experience?
Q19R8	Solves problems the first time you contact them - How satisfied are you with [pipe:hCurrentProvider]'s performance on the following aspects of your overall wireless service experience?
Q19R9	Is a brand for me - How satisfied are you with [pipe:hCurrentProvider]'s performance on the following aspects of your overall wireless service experience?
Q19R10	Total cost of wireless service - How satisfied are you with [pipe:hCurrentProvider]'s performance on the following aspects of your overall wireless service experience?
Q19R11	Device Selection - How satisfied are you with [pipe:hCurrentProvider]'s performance on the following aspects of your overall wireless service experience?
Q19R1aux (resp. Q19R2- Q19R11)	Aware of [pipe:hCurrentProvider]'s performance on the aspects mentioned in Q19R1? (resp. Q19R2-Q19R11)
Q23	How likely are you to change anything (plan, provider, device) about your current wireless service in the next 6 months?
Q35	On a scale of 1 to 5, how well do you feel you understand the details of your wireless plan with [pipe:hCurrentProvider]?

Table 4 List of Strategic Variables

Variable Name	Description
Q12	Thinking about your overall experience with your wireless service provider, on a scale of 0 to 10, how likely are you to recommend [pipe:hCurrentProvider] to a friend or family member?
Q18	Q18: Overall, how satisfied are you with [pipe:hCurrentProvider]?
Q20	How likely are you to switch wireless service providers within the next 12 months?
Q20AR1	Carrier 1 - What best describes your overall feeling about each wireless service provider?
Q20AR2	Carrier 2 - What best describes your overall feeling about each wireless service provider?
Q20AR3	Carrier 3 - What best describes your overall feeling about each wireless service provider?
Q27	How likely are you to recommend your [pipe:Q26] phone to a friend or a colleague?

Table 5 List of Target Variables

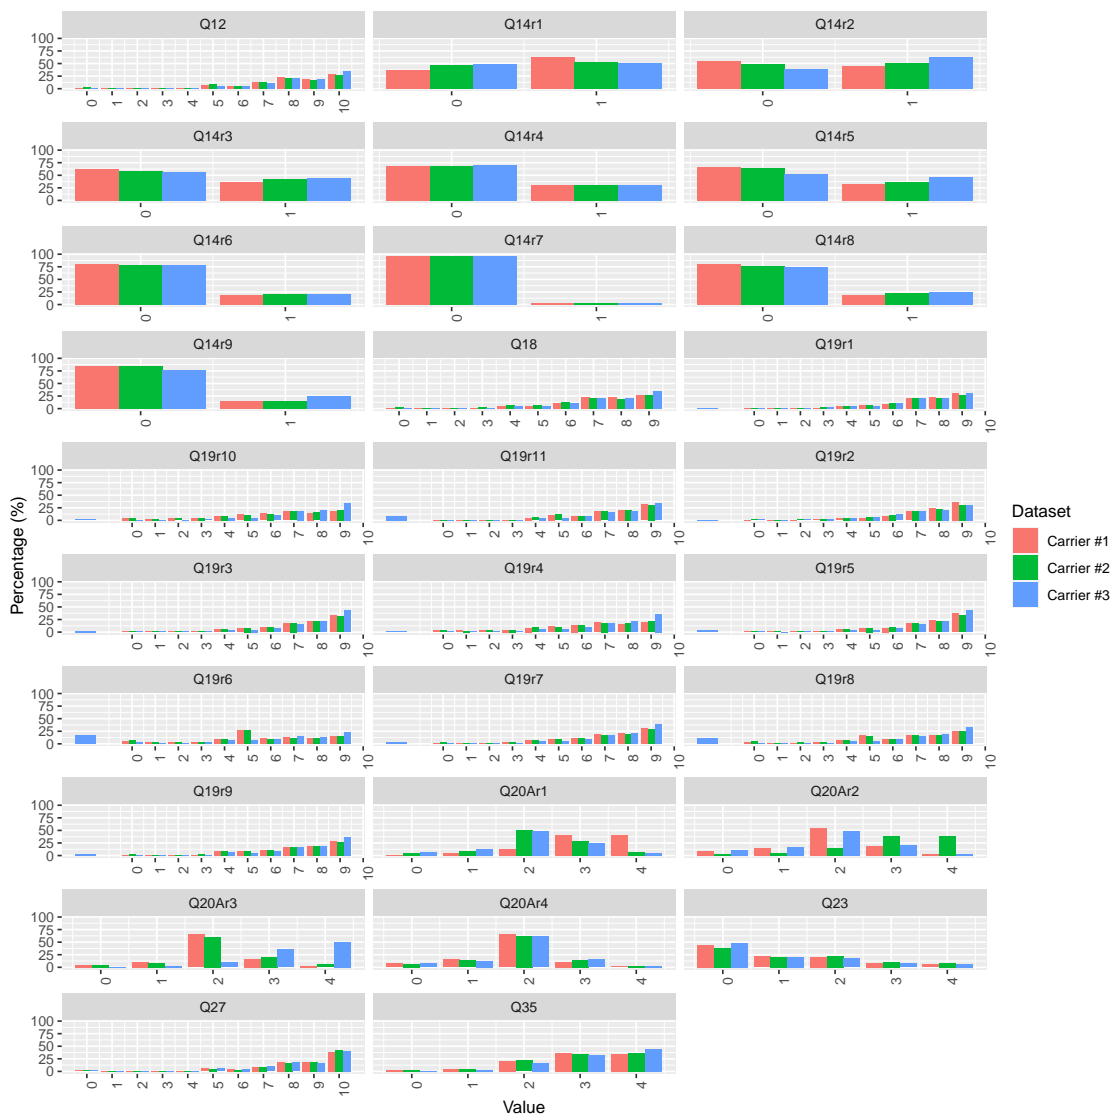


Figure AP.1 Summary Statistics for Target Variables and Strategic Variables, Per Carrier