

Controlling for Retailer Synergies when Evaluating Coalition Loyalty Programs: A Bayesian Additive Regression Tree Approach

Wayne Taylor

Xiaojing Dong

February 17, 2022

Abstract

Spatial models in retailing allow for correlations among purchase decisions from consumers within predefined geographic areas. The purpose of these models is to control for unobserved demand side effects at the regional-level (e.g., a neighborhood), but they typically ignore synergies among individual retailers within a region. To capture the synergies on both the supply side (e.g., store density) and the demand side (e.g., socioeconomic differences across regions), we augment a traditional spatial model with a Bayesian Additive Regression Tree (BART). This allows us to account for unobserved regional differences and observed but potentially complex interactions among individual customers and retailers. We apply this model to a credit card coalition loyalty program (CLP). In our empirical setting, we are interested in analyzing the impact of the loyalty program earnings structure on monthly spend. We do this while controlling for the evolving coalition network, which contains hundreds of geographically dispersed partner retailers. Our data has two key features that permit this. First, the retail partner network evolves over time; this variation in retailer participation allows us to observe card spending patterns when individual retailers are both in and out of the coalition network. Second, the data contains a natural experiment where the loyalty program changed its earning structure, which allows us to estimate the impact of the rewards rate of the loyalty program on customer spend. Our findings show that failure to control for the dynamics of the coalition network results in severely biased estimates of CLP rewards effectiveness. We discuss the implications of BART in our empirical setting and highlight its potential in other marketing situations which contain numerous, interacting control variables.

1 Introduction

About 75% of consumers have a rewards program attached to their credit card, and three of the top four factors cited as most important when choosing a credit card relate to rewards.¹ The structure of a rewards program is flexible, but a recurring feature is that cardmembers earn points when purchasing at one or a few selected retailers, and the rewards can be redeemed during a certain period of time. At one extreme, rewards are earned and redeemed only at a single retailer (e.g., a co-branded credit card between the retailer and an issuing bank). At the other extreme are coalition loyalty programs, where rewards are earned and redeemed from a variety of retailers spanning multiple product and service categories. Estimating the effectiveness of a coalition loyalty program (CLP) is challenging when the network of participating retailers is large and changes over time. In this paper we estimate the impact of a change in a coalition loyalty program rewards structure on cardmember spend. To control for the complex and nonlinear interactions among coalition retail partners we augment our model with a Bayesian additive regression tree (BART) (Chipman et al., 2010). The primary goal of this paper is to measure the impact of the change in the loyalty program on cardmember spend. In doing so we highlight how BART can be integrated into our marketing model to solve a previously intractable, yet important consideration in the analysis of coalition loyalty problems. Specifically, we use BART to account for the complex and evolving network of partners, in order to evaluate the impact of a change in the rewards schedule of a large CLP. A secondary goal is that this paper serves as a motivating example on the versatility of BART and its potential to augment a variety of marketing models in similar situations facing a large number of interrelated control variables.

To estimate the impact of changing the coalition loyalty program, our modeling approach starts with a traditional econometric spatial model of monthly cardmember spend. Spatial models recognize that cardmembers within each predefined geographic region are similar, and their spending might be similar between neighboring regions. In our empirical application, this allows us to evaluate the influence of a change in the CLP rewards on spend while controlling for unobserved demand side factors such as differences in population density, socioeconomics, and level of competition associated with each region. However, a coalition loyalty program introduces many supply side issues that are not handled with such a spatial model setup. For instance, in a coalition loyalty program synergies may exist between retailers across regions. In addition, the sheer number of retailer partners associated with the coalition loyalty program may be too cumbersome to estimate with a traditional spatial model.

To account for supply side synergies among coalition retailers, we need to control for the structure of the coalition network in order to properly estimate the impact of a change in the rewards structure. This

¹From “The Payments Ecosystem”, 2021 report from Insider Intelligence.

introduces non-trivial challenges to the model; in our empirical application hundreds of retailers participate in the coalition and the network evolves over time. In addition to handling this high-dimensional data, synergies may exist among multiple partners, either within or across the predefined spatial regions. To address these modeling challenges, we augment our baseline spatial model with a Bayesian additive regression tree, which offers a flexible, non-parametric approach and accounts for the unobserved interactions among the retailers in the CLP. This allows us to control for both the unobserved demand side effects of each region and the observed supply side effects of partner retail locations. The augmented model reduces the potential for omitted variable bias, where supply side synergies would otherwise be absorbed into demand side regional-level spatial effects.

We estimate the proposed model using detailed, individual-level credit card data from a large European coalition loyalty program with hundreds of retailer partners spanning a broad geographic area. Of primary interest is the effect of a change in the rewards structure on monthly card usage. The spatial component captures the demand side effects from 26 regions of interest² while the BART component controls for supply side synergies among all of the coalition retailers. The supply side synergies we intend to capture can transcend the predefined spatial regions and include location advantages (e.g., a participating gas station may be more attractive if it is near a participating grocery store), perceptual advantages (e.g., the credit card may be more attractive if luxurious brands participate in the earnings network, even those not nearby), or economic advantages (e.g., complementary retailers). We use BART to capture these granular supply side synergies in tandem with the spatial component to control for demand side effects.

Two features of this data are key to this analysis. First, the coalition network evolves over time, allowing us to observe card spending patterns when individual retailers are both in and out of the coalition network. Second, the coalition changed the loyalty program rewards structure partway through the data. This shock allows us to separate out the impact of the loyalty program earnings structure on customer spending from the influence due to the dynamic network of the participating retailers.

Our empirical results show that failure to control for the evolving coalition network biases the estimated loyalty program effects, some of which become insignificant once the evolving network is accounted for. The analysis presents BART as one solution to account for CLP supply side synergies, which are infeasible to control for using a standard spatial model. Even though our empirical application focuses on augmenting a spatial model with BART, we hope marketers become inspired to use BART to enhance other marketing models in similar situations with numerous, potentially complex control variables. The modular nature of the BART model permits it to be included in almost all the traditional econometric models. The nonlinear

²The regions correspond to major city centers and were per-defined by the data provider, who is required to remain anonymous.

nature of the BART approach makes it suitable to capture sophisticated interactions of a large number of variables without predefined functional form. It offers great advantages when controlling for the effects of these input variables, while evaluating marginal impact of other variables of interest.

The remainder of the paper is organized as follows: Section 2 provides background and literature reviews related to loyalty programs and spatial models prior to introducing BART. Section 3 discusses the empirical setting. Section 4 presents the augmented spatial model with results in Section 5. Section 6 presents additional analyses that leverage BART and section 7 concludes.

2 Background and Literature Review

This paper draws primarily from two research streams. First, we review relevant work on loyalty programs, of which a coalition loyalty program is one type. Second, we discuss prior research on spatial models in marketing. Our aim is to motivate our empirical application and highlight the modeling challenges that prior work has not yet resolved, namely related to the inability of spatial models to capture granular supply side effects. We then introduce Bayesian additive regression trees as a solution to these challenges, both as it relates to our empirical application and its applicability in other marketing contexts.

2.1 Loyalty Programs

Loyalty programs have served as an important component for customer relationship management for a variety of firms. The marketing literature has seen a long list of studies dedicated to understanding their impacts across many industries, including airlines, financial services, hotel, gaming and retail (Ferguson and Hlavinka, 2007). The majority of these studies focused on understanding the overall impact of a loyalty program (LP) on changing the consumers' purchase behaviors before and after joining such a program (Lal and Bell (2003); Verhoef (2003); Taylor and Neslin (2005); Liu (2007); Zhang and Breugelmans (2012), to name a few). With a few exceptions, most of the existing studies examine loyalty programs offered by a single firm (see reviews provided by Breugelmans et al. (2015) and Liu and Yang (2009)). In contrast, a coalition loyalty program allows customers to buy and redeem rewards across multiple partner retailers, which have received only limited attention in the marketing literature (Dorotic et al., 2011; Stourm et al., 2017). Meyer-Waarden (2008) pointed out that with CLPs, it is even harder to influence consumers' purchase behavior than a loyalty program offered by a single firm. In fact, the literature has documented somewhat conflicting results of the impact of CLPs. Moore and Sekhon (2005) confirmed the lack of evidence that CLPs can improve market share. However, Lemon and Wangenheim (2009) finds a reinforcing mechanism in cross-buying across loyalty program partners using data from a European airline and its partners.

As summarized by Liu and Yang (2009), the design of a LP (and CLP) consists of many moving parts that likely influence its ultimate effects on changing consumers' purchase behaviors. These factors include, for example, the point structure, such as the point thresholds (O'Brien and Jones, 1995); and consumer characteristics. For example, Lal and Bell (2003) and Liu (2007) find that it is the light users who experienced the biggest increase in spending after joining the LP. In a coalition loyalty program, there is an added complexity due to the evolution of the participating retailers over time, which can influence the overall attractiveness of the program. Our paper contributes to the study of coalition loyalty programs, where we develop a model to evaluate the change in purchase behavior due to the adjustments in the points structure, while controlling for the potentially nonlinear and complex interactions from the evolving coalition retail network.

2.2 Spatial Models

The literature in marketing has an established history of incorporating spatial components into marketing models. The scope of the applications has been quite broad, including models of product adoption, satisfaction ratings, promotional strategies, and advertising allocation (Aravindakshan et al., 2012), among others (see Bradlow et al. (2005) for a more comprehensive review). Spatial models recognize that individuals (or, more generally, any unit of analysis) are located in space and that the decisions of individuals in the vicinity may be correlated. In spatial econometrics the spaces are typically represented on a map that is geographical in nature, but spatial closeness can be also defined using demographics (Yang and Allenby, 2003), latent constructs (Moon and Russell, 2004), or any other construct that determines the relative similarities of the individual units. A common assumption is that the behavior of one individual is conditionally independent of the behavior of another individual. In our empirical application, we assume that geographic proximity implies high correlation between individuals. That is, customers located near each other share similar unobserved traits.

The spatial model presented in this paper is designed to capture spatially correlated errors, the idea being that unobserved variables that drive purchase behavior can be inferred from proximity on a map (see Russell and Petersen (2000), Yang and Allenby (2003), and Bronnenberg and Sismeiro (2002) as examples). The spatial component of the model is formulated using a typical conditional auto-regressive (CAR) structure: given the random effects from all other areas, the distribution of the spatial effect only depends on its neighbors (Cressie, 1993). In short, each geographic area is similar to its neighbors.

Spatial modeling usually involves specifying a covariance matrix between regions, from which a correlation coefficient is estimated to measure influence between neighboring regions. Because of this matrix, empirical

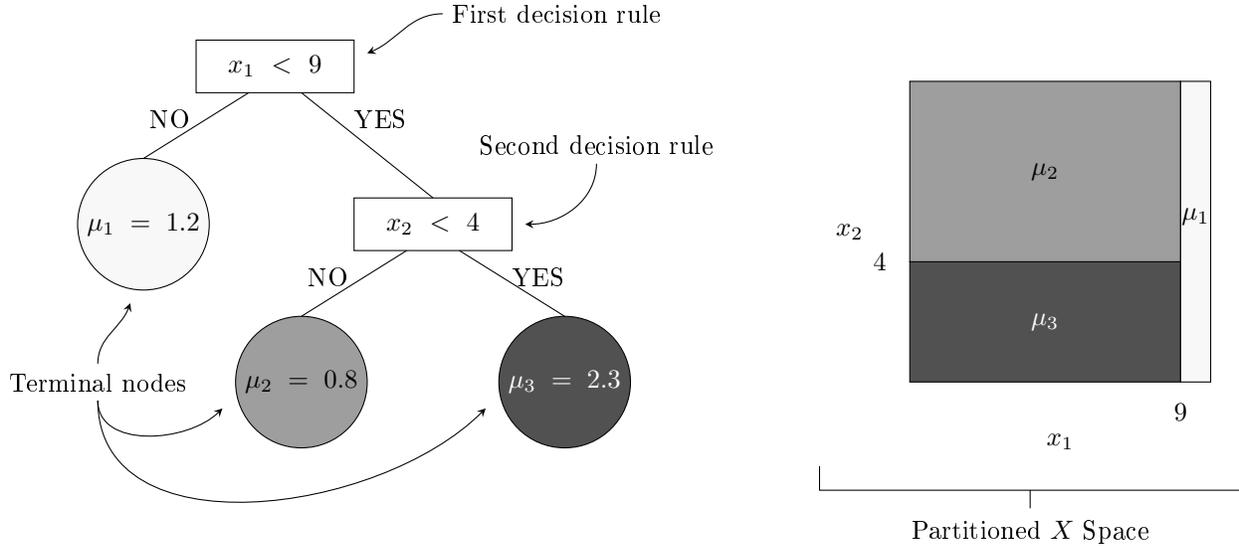
applications in prior work typically specify only a relatively small number of areas. However, in our empirical application, we observe hundreds of retailers. While we want to control for demand side effects at the regional level, it would be intractable to try and make the regions granular enough to capture supply side effects at the retailer level. Our solution is to retain a spatial modeling structure to capture demand side effects at the regional level, but offload the supply side effects from individual coalition retailers to a Bayesian additive regression tree.

2.3 Bayesian Additive Regression Trees (BART)

Marketing modeling has encountered many situations where the complexity of relationships among variables can be non-trivial. For example, optimizing store shelf layout (Van Nierop et al., 2008), identifying influential users in a social network (Trusov et al., 2010), and managing the marketing mix in customer relationship management (Rust and Verhoef, 2005), to name a few. In this paper we focus on the inherent complexity in controlling for the supply side effects of a large network of retailers whose locations could create synergies with each other (e.g., complementary stores located near each other) or could exert a highly non-linear relationships with respect to cardmember spend.

From the modeling perspective, conditioning on these cross-effects within a large coalition network is quite difficult. To account for the potentially high order interactions among multiple retailers and handle the high-dimensionality of the data we extend a spatial model with a Bayesian additive regression tree as proposed by Chipman et al. (2010). At its core, BART provides a Bayesian inference on a model of $y = f(X) + \varepsilon$ with $\varepsilon \sim N(0, \sigma^2)$ where f is a sum of regression trees, and therefore a non-parametric function of the input X variables.

The regression tree serves as the basic building block to a BART model. Figure 1 illustrates a simple regression tree based on our empirical application. In this example, only two variables (x_1 and x_2) are considered, which represent the distances between a customer and two partner retailers. The regression tree partitions the X space into subgroups. Observations are assigned to each subgroup by working down the tree’s decision rules. If the distance to the first retailer, x_1 , is less than 9km then the distance of the second retailer, x_2 , is considered and the predicted monthly spend on the credit card is either μ_2 when x_2 is less than 4km or μ_3 otherwise. The values at these terminal nodes (μ) are typically the average of the observations in each subgroup (e.g., the average monthly spend for observations that meet all the decision rules on the branches leading to the terminal node).



On the left is a regression tree with two decision rules and three terminal nodes. The decision rules are processed from top to bottom until a terminal node is reached, where the terminal node contains the predicted outcome value. The predicted value is typically the average of the outcome variable y associated with all X that satisfy the decision rules in the data used to fit the model. On the right is a visual representation of the total X space partitioned, according to the regression tree presented on the left.

Figure 1: Example Regression Tree Reproduced from Hill et al. (2020)

BART extends this single-tree setup by combining multiple such trees, where a regularization prior controls the tree growing process to keep each tree small to explain only a modest portion of the data. Models with sums of regression trees, such as BART, have greater ability and flexibility than single trees to capture interactions and non-linearities among the input variables (Hill, 2011). Relative to a traditional tree model, BART offers two additional key advantages in practice, as outlined in Allenby et al. (2014). First, the default priors have been shown to consistently provide reasonable estimates, so there is less parameter tuning using cross-validation. Second, the depth of the tree (i.e., its complexity) is inferred naturally through the MCMC process so that the level of interactions is inferred statistically. The non-parametric nature of BART lends itself unique advantages in capturing sophisticated interactions among a large number of input predictors and their possible nonlinear effects. Furthermore, the Bayes posterior distributions obtained through MCMC provides the usual representation of uncertainty in the parameter estimates.

Although there are many modeling incidences that are well suited to a flexible approach such as BART, most of the current applications have been mostly focused on evaluating heterogeneous treatment effects (for example, Hill (2011), Green and Kern (2012), Carnegie et al. (2019), Heckman et al. (2014), Hahn et al. (2020), and Athey and Imbens (2017)). This paper applies the approach in a distinct context.

A significant feature of BART is that it can be easily incorporated into a variety of established models, as its modularity allows it to be simply inserted into the MCMC sampling process when estimating the model.

This presents BART as an attractive option to control for potentially complex, auxiliary variables while still adhering to a traditional functional form on the key econometric variables of interest. As demonstrated in Chipman et al. (2010), the BART model can be applied in a probit model, with only minor modifications of the methods. Kindo et al. (2016) further expanded it to a multinomial probit model. In addition, BART has also been applied in the context of survival analysis. Bonato et al. (2011) incorporated BART into parametric survival models, including Cox Proportional Hazard model, Weibull and Accelerated Failure time models; and Sparapani et al. (2016) applied BART in non-parametric survival analysis using a non-proportional hazard approach. In this application we augment a spatial model with BART, similar to Zhang et al. (2007) who use the approach for merging datasets.

The implementation of BART in the marketing literature is surprisingly limited. In our empirical application, we use BART to control for the complex supply side effects among many coalition retailer partners in order to address our managerial question of interest: what is the impact of changing the coalition loyalty program on customer spending patterns? We find that in doing so, we reduce the bias of the coefficients of interest related to the effects of the loyalty program. To our knowledge this paper is one of the first to examine the spatially driven impacts of a coalition loyalty program and one of the early attempts to incorporate the BART model into research in marketing.

3 Empirical Setting: Coalition Loyalty Program

In this section we present the empirical context and describe the data used for model estimation. We highlight the unique aspects of the data, namely the availability of a natural experiment that allows us to estimate the impact of changes in the LP rewards structure on credit card spend while controlling for the evolving network of participating retailers.

The dataset comes from a large European coalition loyalty program. Cardmembers sign up for the loyalty program combined with a credit card. The credit card is used for transactions as any other typical credit card, including purchases and cash withdrawals. Purchases can be made (and points accumulated) anywhere credit cards are accepted, however points are earned at a faster rate with purchases at in-network retailers (i.e., “partner” stores) relative to out-of-network stores. Earned points are redeemed for cash vouchers, which can be used to purchase additional goods and services at selected partner retailers. In addition to managing the points tracking and voucher generation and redemption, the firm sends outbound marketing campaigns to encourage sign-up, purchases, and voucher usage.

The network of partners is considerably large and diverse, both in terms of partner type (e.g., gas stations, clothing retailer, etc.) and regional footprint – in an average month there are 1,956 unique partners.

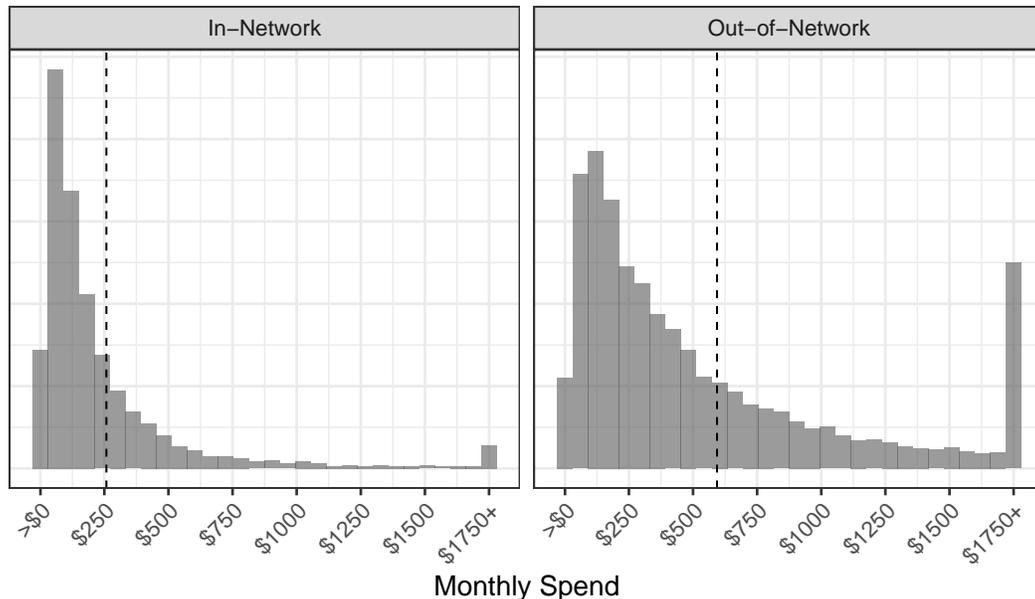
Typically, partners are independent retailers looking to benefit from the network effects of the coalition, but there is a substantial number of chain retailers included as well. Over the observed time frame of four years (November 2007 - November 2011) the network size ranges from 1,891 partners to 2,017 partners. While the overall network size remains relatively stable over time, in each month there are an average of nine partners joining the network and nine partners leaving the network.

Total Cardmembers	619
Avg. Monthly In-Network Spend	\$258
Avg. Monthly Out-of-Network Spend	\$594
Avg. In-Network Distance Traveled (km)	12
Avg. Unique Branches Visited	8
Avg. # of Stores within 10km	46

Table 1: Cardmember Summary Statistics

We observe monthly spend from 619 randomly selected customers. About 71% of the months observed have zero in-network spend and 44% of the months observed have zero out-of-network spend. The data summary is presented in Table 1. As shown in this Table, the average monthly in-network spend is \$258, about half of the the out-of-network monthly spend at \$594 (for months with nonzero spend), but there is a long tail to this distribution. Figure 2 shows the distribution of non-zero monthly spend.

Non-Zero Monthly Card Usage

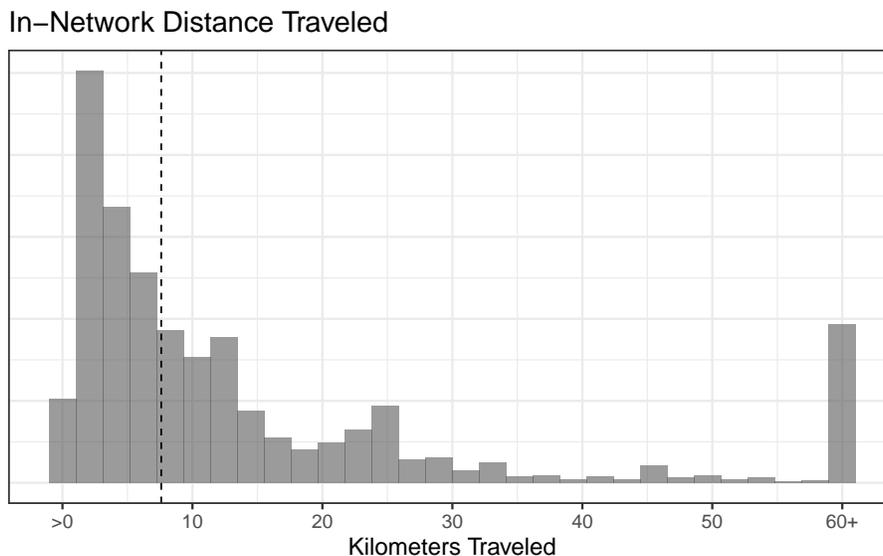


The average cardmember monthly spend is \$258 in-network and \$594 out-of-network in months with nonzero spend.

Figure 2: Monthly Card Spend

The geographic footprint of the retail partners in the coalition network is rather large – the average

distance between each cardmember and all retail partners is about 100km. However, about half of all transactions occur within 8km. Figure 3 shows the distribution of distance traveled for all in-network transactions, with 58% occurring within 10km and 87% within 30km.

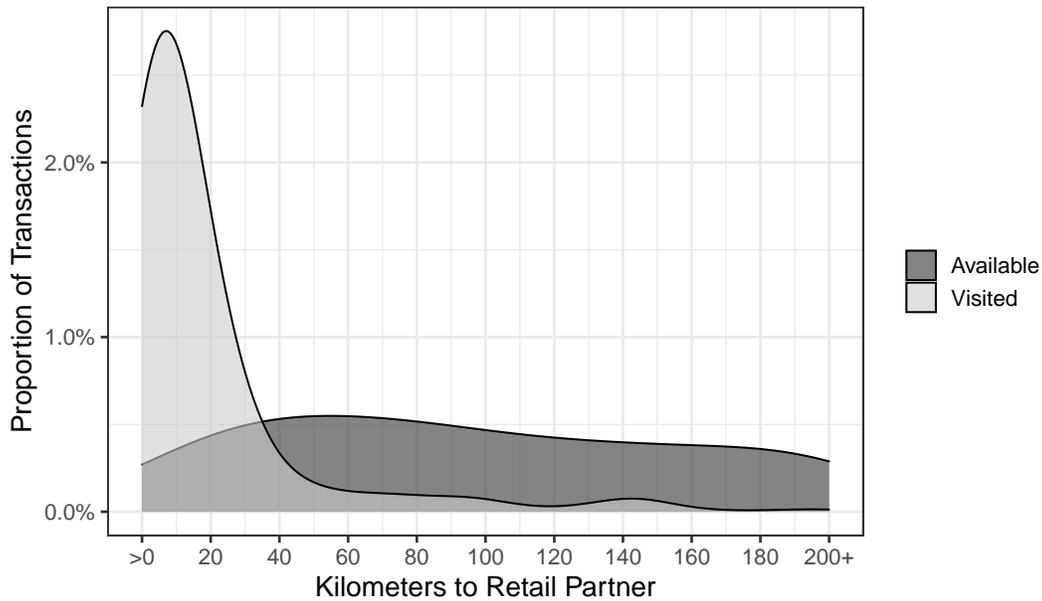


The dashed line represents the median of the distance traveled by each cardmember for in-network transactions, about 8km.

Figure 3: In-Network Distance Traveled

This suggests that the geographic distribution of retail partners will impact cardmembers' spending differently, as most cardmembers shop primarily at stores close to them. Figure 4 compares the distribution of distance traveled to retailers "available" to each cardmember, depending on their location. We see the vast majority of cardmembers stick with a local network of partners, and do not engage with the coalition network further away. The implication of this is that a partner entry or exit from the coalition would have heterogeneous influences on cardmembers' behaviors. The impact would be much larger to those located nearby, but almost negligible for those far away. Considering the many entry and exit among the large number of partners over the four years of the data period, controlling for the retailer partner changes is critical.

Coalition Partner Distribution

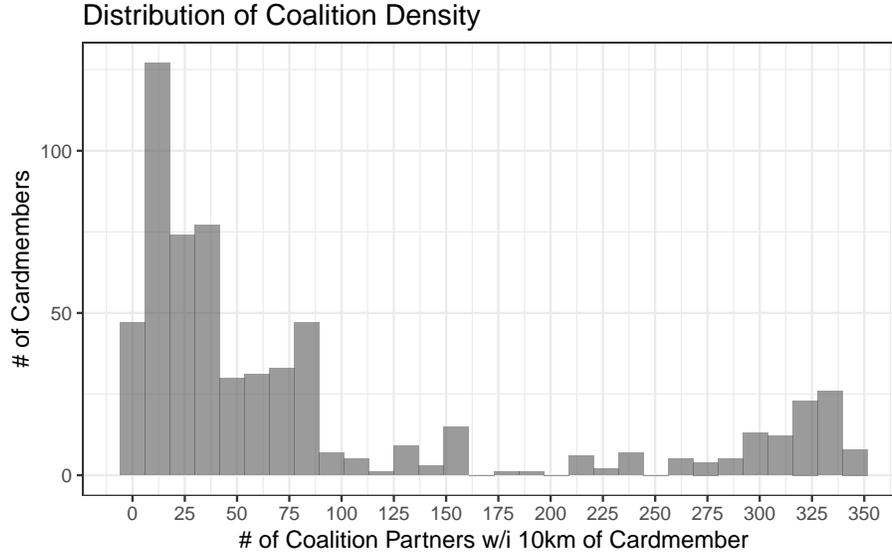


The distribution of coalition partner retailers available to each cardmember is relatively uniform in terms of distance. However, cardmembers tend to frequent retailers located close by.

Figure 4: Coalition Usage

There is also substantial heterogeneity in the size of the coalition network around each cardmember. Figure 5 shows the distribution of the number of stores within 10km of each cardmember.³ The average cardmember has access to 88 retail partners within 10km, but we see this distribution is bi-modal: some cardmembers are in high areas of saturation (e.g., a city center with many partners) while others are located in areas with only a dozen or so partners nearby (e.g., an area of low interest to potential coalition partners). Recognizing both the limited travel distances and the variation in nearby network composition, it is important for the model to consider the variations in accessing the different partners among cardmembers. This motivates the need to account for the unique distances between each cardmember and all partner retail locations.

³Since the coalition evolves, this distribution represents the average number of stores across time.



Within a 10km travel range (slightly above the median travel distance of 8km), there is substantial variation in the number of coalition partners available.

Figure 5: Stores within 10km of Cardmembers

3.1 Change in the Loyalty Program Structure

A key feature of this data is that it contains a natural experiment where the firm decided to change the LP earnings structure. The change in the LP design occurred in November 2009 which was intended to make the program more competitive, and is therefore exogenous with respect to individual cardmember locations or spend behavior. The four year observation period spans the two years on either side of this change, which serves as the focal point of our analysis: to what extent did the change in the LP structure influence cardmember behavior? Since we do not have access to a control group (that is, individuals who did not experience a change in the coalition loyalty program), evaluating the causal impact of the program design change on all the cardmembers is not feasible. Instead, we focus our attention to comparing the differential impacts across two customer segments: high spenders and low spenders, defined by a median split in their total spend prior to the change.⁴ This allows us to determine if the change was equally attractive to each segment, or if one segment was more sensitive to the change.

The change introduced by the firm involved two primary components: 1) The relative attractiveness of using the card out-of-network and 2) the exchange rate (i.e., the monetary voucher value of each point earned). Table 2 outlines the structure of the LP both before and after the change occurred. The in-network earning rate remained constant at .5 or 1 point per dollar spent, depending on the partner. However the

⁴While there are countless ways to segment customers, we chose a median split since it maximizes the number of cardmembers in the fewest segments.

out-of-network earnings rate doubled from .1 to .2 points per dollar, thereby making out-of-network spend relatively more attractive. In addition, the points were devalued by a factor of three in that after the change it requires three times as many points to earn an equivalent voucher value. This change played out in a predictable pattern in the data: overall spending declined (perhaps from point devaluation) and the proportion of spend in-network declined (perhaps due to the relative gain in attractiveness of out-of-network card usage).

	Pre LP Change	Post LP Change
Points Earned In-Network	.5 or 1*	.5 or 1*
Points Earned Out-Of-Network	.1	.2
Exchange Rate	$33\frac{1}{3}$	100
Avg. In-Network Spend	\$90	\$61
Avg. Out-Of-Network Spend	\$330	\$334
Avg. Total Spend	\$420	\$395

*Depending on the retail partner.

Table 2: Coalition Program Schedule

Interestingly, we find that the change in the behavior varies drastically between “High” and “Low” card-member segments, which we defined using a median split of total spend prior to the loyalty program change. Figure 6 shows the monthly spend by segment before and after the LP change. After the LP changed, high spenders reduced their in-network and out-of-network spend and low spenders reduced in-network spend but increased out-of-network spend. These patterns are reproduced with a regression discontinuity model shown in Table 3. The challenge with this simple regression is that we have not conditioned on the evolving coalition network of retailers. This makes it difficult to disentangle the effect of the change in the LP structure from the influences of the retailers’ participation and their accessibility to the cardmembers. Furthermore, the changes in the participation of the retailers in the coalition network may be correlated with the spending behaviors of the customers. For example, the decrease in in-network spend from the low segment may reflect coalition partners leaving the network who are attractive to this group, and not have any relation to the change in the loyalty program.

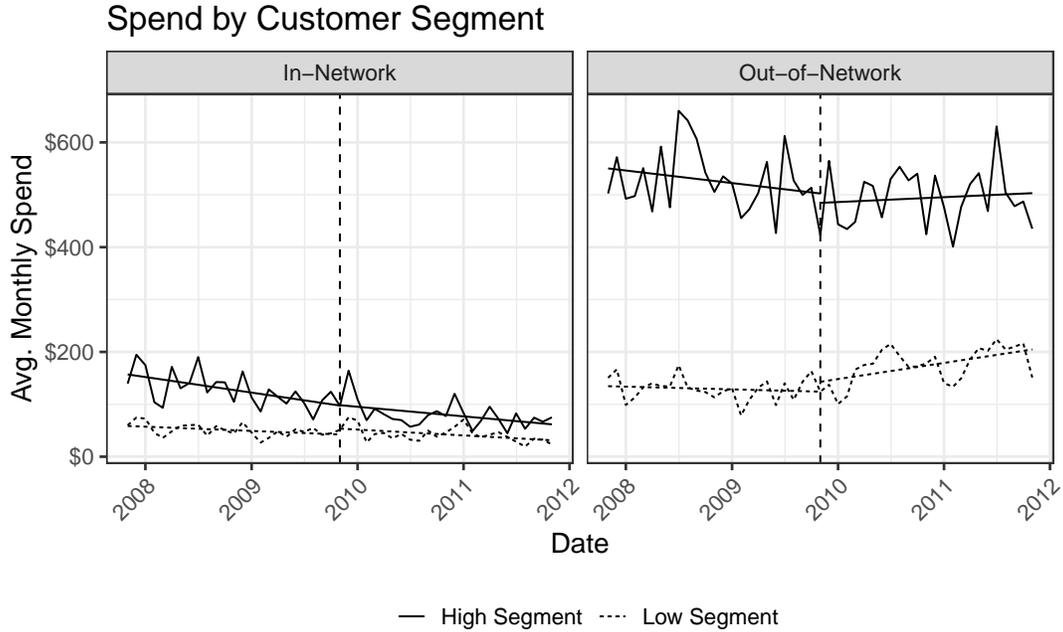


Figure 6: Monthly Spend by Segment

	<i>Monthly Spend DV:</i>	
	In-Network	Out-of-Network
	(1)	(2)
Intercept	50.53*** (2.93)	129.62*** (6.64)
High Spend Segment	78.06*** (4.15)	401.10*** (9.39)
Post LP Change	-8.08** (4.11)	44.21*** (9.29)
Interaction	-40.44*** (5.81)	-81.13*** (13.14)
Observations	30,280	30,280
R ²	0.02	0.09

Note: * p<0.1; ** p<0.05; *** p<0.01

The results from this regression discontinuity suggest that the change in the loyalty program negatively impacted all spend except for out-of-network spending in the low segment.

Table 3: Regression Discontinuity

To isolate the influence of the change in the LP structure, we need to control for the state of the coalition

network. This is a non-trivial task for a variety of reasons. The first is that the network is very large, involving around two-thousand retailers, so including individual partner level controls is simply not feasible using a traditional functional form model. The second concern is that synergies may exist among retailers in such a large network. In other words, it is unlikely that a single partner will make a material difference with its entry or exit but rather it is the holistic network composition that plays a role in the attractiveness of the network. For example, a handful of complementary stores, as a group, makes the card attractive and only if all these stores exit card usage may drop significantly. Capturing these higher order interactions is especially challenging, and in the next section we present a model to handle both issues.

4 Model Specification

In this section we first introduce our conditional auto-regressive (CAR) spatial model, which is then augmented with BART. The primary purpose of this model is to estimate the effect of the change in the loyalty program, conditional on unobserved regional level effects and the influence from the evolving network structure. We begin with the base model specification with spatially correlated errors and discuss the modifications to account for the correlations between our two dependent variables and their censored distributions. Then, we augment this model with the BART component to control for the complex coalition network structure.

4.1 A Conditional Auto-Regressive (CAR) Spatial Model

In this sub-section we outline our base model for each dependent variable, which takes the form of a CAR spatial model. We observe in-network (y_{it}^I) and out-of-network (y_{it}^O) credit card spend for cardmembers $i = 1 \dots N$ in month t . Since the independent variables are identical in both models, we use superscripts to differentiate between in-network (I) and out-of-network (O) coefficients and for clarity we drop the superscripts when feasible. The base models are specified as follows:

$$y_{it}^I = \alpha_i^I + \Gamma^I \mathbb{S}_{[it, \cdot]} + \kappa^I G_i + \zeta^I Z_{it} + \omega^I G_i Z_{it} + \theta_a^I + \varepsilon_{it}^I \quad (1)$$

$$y_{it}^O = \alpha_i^O + \Gamma^O \mathbb{S}_{[it, \cdot]} + \kappa^O G_i + \zeta^O Z_{it} + \omega^O G_i Z_{it} + \theta_a^O + \varepsilon_{it}^O \quad (2)$$

Both dependent variables are log transformed to accommodate their skewed distributions. In months with zero spend we treat the outcome as a censored variable, where the realized $y = \tilde{y}$ if the latent variable $\tilde{y} > 0$ and 0 otherwise, as in a standard Tobit Type II specification.

Among the coefficients, α_i captures individual-level fixed effect for each cardmember i . Γ controls for

the month-specific fixed effect, with December as the reference. \mathbb{S} corresponds with the (1×11) seasonality matrix associated with cardmember i at month t . Of primary interest are the coefficients associated with the spend group of the customers, the time period after the LP change, and their interaction: κ , ζ , and ω . G_i is an indicator which equals one if customer i is in the high spend group, determined by a median split of total spend prior to the LP change. Z_{it} is an indicator that equals one for all observations after the LP changed. Lastly, the coefficient ω for the interaction term captures the differential effect on behavior changes between the two spending groups, due to the change in the coalition loyalty program.

We control for spatially correlated errors and estimate random effects θ_a for areas $a = 1, \dots, A$ using a conditional auto-regressive prior (Cressie, 1993). The key idea of the CAR component is that each area is similar to its neighboring areas. The density of each regional effect θ_a is defined by the set of conditional distributions:

$$p(\theta_a | \theta_{-a}) \sim N \left(\frac{\rho}{h_a} \sum_{k \neq a} \mathbb{C}_{ak} \theta_k, \frac{\delta^2}{h_a} \right), \quad a = 1 \dots A \quad (3)$$

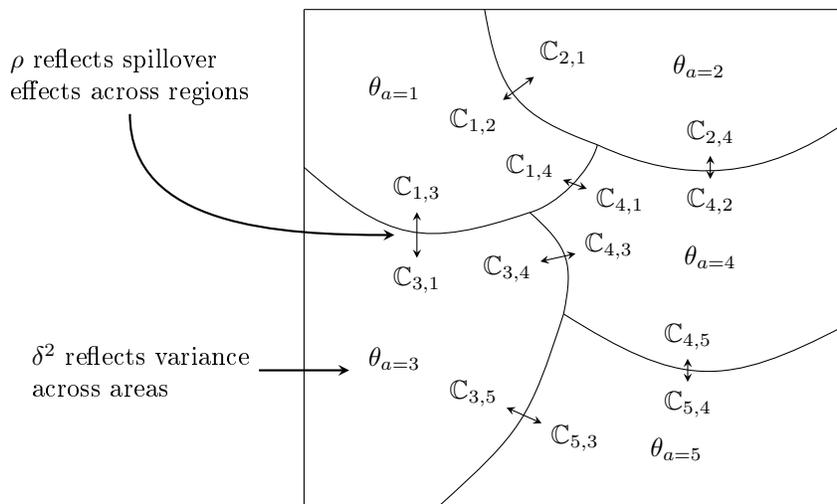
The distribution of each area effect θ_a follows a normal distribution where the mean and variance of the distribution are influenced by the regional effects of the neighbors and the spillover effects across the areas. In particular, the mean is a function of the following three components: 1) an adjacency matrix \mathbb{C} identifying the neighbor areas of a , where $\mathbb{C}_{ak} = 1$ if area k is a neighbor to a and $\mathbb{C}_{ak} = 0$ otherwise (including $\mathbb{C}_{aa} = 0$), 2) the total number of adjacent areas to a , h_a , which is defined as $h_a = \sum_{k=1}^A \mathbb{C}_{ak}$ and where A refers to the total number of areas, and 3) the estimated spillover coefficient, ρ . Positive (negative) values of ρ indicate positive (negative) values of correlation between adjacent regions. If ρ is not significantly different from zero, we conclude that there are no spillover effects across regions. When ρ approaches 1 the conditional mean of θ_a equals the average of its neighbors. The variance of this distribution is a function of the estimated cross-area variance δ^2 , which is scaled by the number of neighboring regions h_a . In our empirical application each cardmember is assigned to one of 26 regional areas (as predefined by the data provider) corresponding to the major city centers.

A primary difference between the CAR model and an alternative simultaneous auto regressive (SAR) model (for example, as used by Yang and Allenby (2003)) is that the SAR model allows for a non-symmetric spatial matrix θ . However, as Cressie (1993) points out, a space-time model is probably more appropriate if the modeler desires non-symmetric spatial dependence matrix. Since the parameters are more naturally interpreted from a conditional-expectation structure, it is argued that the CAR model be used from the outset.

The joint distribution of $p(\theta)$ implied from equation 3 is

$$p(\theta|\rho, \delta^2) = N\left(0, \delta^2 (\mathbb{H} - \rho\mathbb{C})^{-1}\right) \quad (4)$$

where \mathbb{H} is an $A \times A$ diagonal matrix with h_a as the diagonal elements. We employ a uniform prior on ρ over a specified range. The parameter ρ must lie in the interval $[\lambda_{\min}^{-1}, \lambda_{\max}^{-1}]$ where λ_{\min} and λ_{\max} denote the minimum and maximum eigenvalues of \mathbb{C} for which the matrix $(\mathbb{H} - \rho\mathbb{C})$ can be inverted (Sun et al., 1999). Figure 7 provides a simple illustration of the CAR structure in a hypothetical region with five areas.



$$\mathbb{C} = \begin{bmatrix} 0 & 1 & 1 & 1 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 1 & 1 \\ 1 & 1 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 \end{bmatrix}$$

$$\mathbb{H} = \begin{bmatrix} 3 & 0 & 0 & 0 & 0 \\ 0 & 2 & 0 & 0 & 0 \\ 0 & 0 & 3 & 0 & 0 \\ 0 & 0 & 0 & 4 & 0 \\ 0 & 0 & 0 & 0 & 2 \end{bmatrix}$$

A CAR structure in a hypothetical five area region. ρ controls spillover among areas, and δ^2 captures the variance in the area fixed effects – both of these parameters are estimated. The corresponding \mathbb{C} and \mathbb{H} matrices are fixed and set by the analyst depending on where the area boundaries lie.

Figure 7: CAR Structure

Lastly, in our model $\varepsilon_{it} \sim N(0, \sigma^2)$ reflects unobservable factors modeled as random error that is independent over time. We allow for correlated errors between monthly in- and out-of-network spend by framing the two equations as a seemingly unrelated regression (SUR) (Zellner, 1962). Estimation of our censored SUR follows the approach of Huang (2001), where the only modification from a typical SUR

estimation is that we substitute augmented data \tilde{y} in the estimation rather than the censored data y . During estimation this is generated from a truncated normal distribution.⁵

4.2 Augmenting the Spatial Model with BART

In section 3.1, we pointed out the challenges when leveraging a simple regression discontinuity approach to evaluate the impact of the change in the loyalty program. Even in the presence of a natural experiment, the naive regression discontinuity fails to control for potentially important demand side and supply side effects, which may bias the results and managerial interpretations. The unobserved demand side effects are captured with the spatial CAR specification discussed in the prior section. That leaves the observed supply side effects to be accounted for, which considers the individual retailer partners’ locations and their entry/exit behaviors. We augment our CAR spatial model with BART to control for these supply side variations. The addition of BART is intended to control for the state of the network in each month, so that we have an unbiased estimate of our effects of interest: the impact of the change in the LP. Our augmented model is specified as:

$$y_{it} = \underbrace{\alpha_i}_{\text{Individual FE}} + \underbrace{\Gamma S_{[it, \cdot]}}_{\text{Seasonality}} + \underbrace{\kappa G_{it} + \zeta Z_{it} + \omega G_{it} Z_{it}}_{\text{Segment and LP Change}} + \underbrace{\theta_a}_{\text{Area RE}} + \underbrace{f(\mathbb{X}_{[it, \cdot]})}_{\text{BART}} + \varepsilon_{it} \quad (5)$$

where $\mathbb{X}_{[it, \cdot]}$ is a $(1 \times 2,329)$ matrix which represents the state of the coalition network with respect to cardmember i at month t for all 2,329 coalition partners. Each element in this matrix represents the geographic influence measured by $1/d_{is}^*$, where d_{is}^* is the distance in kilometers between cardmember i and store s . If the coalition partner is not in the network at time t , the element takes the value of zero.⁶ For example, $\mathbb{X}_{[12,5,760]}$ denotes the geographic influence of store 760 to cardmember 12 in month 5 of the data.⁷

The building block of the BART model is a regression tree. Let \mathcal{T} denote a single binary tree consisting of a set of decision rules and terminal nodes. The collection of terminal nodes associated with this tree is denoted by $\mathcal{M} = \{\mu_1, \mu_2, \dots, \mu_{\mathcal{N}}\}$ for the \mathcal{N} terminal nodes associated with this tree. In the example provided in section 2.3, we had $\mathcal{M} = \{1.2, 0.8, 2.3\}$ to represent the three expected values of y given a particular partition of X as determined by the decision tree. We let the function $g(\mathbb{X}; \mathcal{T}, \mathcal{M})$ represent a

⁵More details on the estimation procedure are available in the appendix.

⁶Similar to all the other tree-based approach, the BART specification is somewhat agnostic to the exact functional form of the geographic influence and its form is mostly dictated to properly denote a retailer’s out-of-network status with a zero. This is done so that the decision tree can allocate a prediction associated with an out-of-network status when needed. In other words, unlike a standard regression model, in a decision tree a value of zero does not imply that the variable will “drop out” of the prediction.

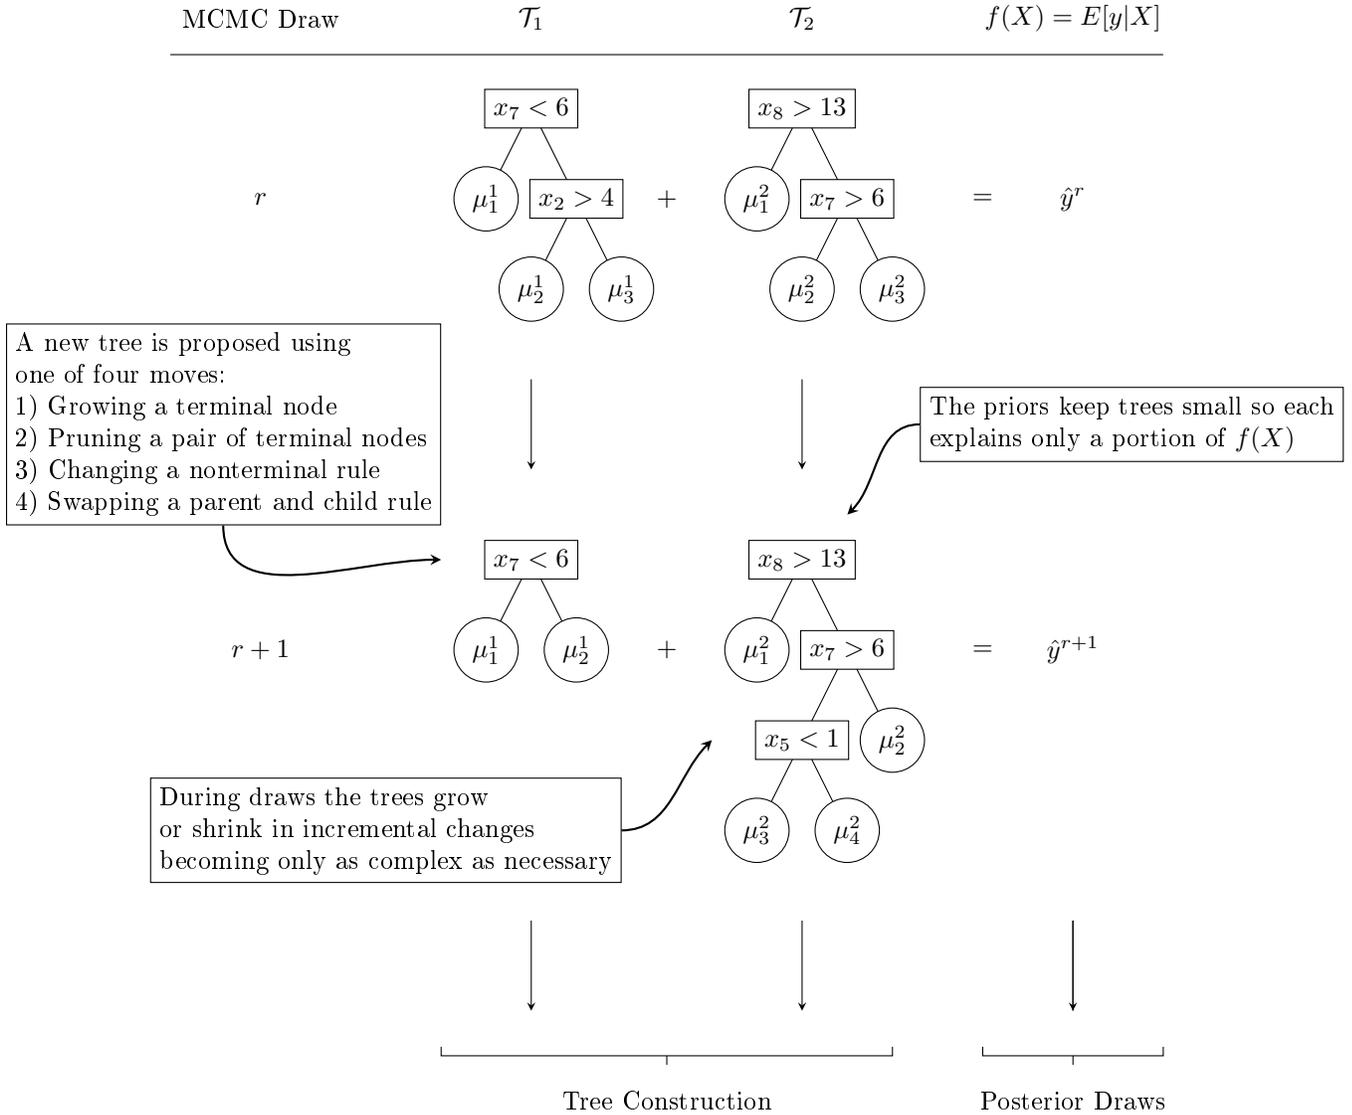
⁷In some CLPs, it is possible that the decision of a retailer to join the coalition is endogenous with respect to cardmember spending patterns. Controlling for this endogeneity would introduce much sophistication to the model, especially considering a CLP with a large number of stores. However, in the web appendix we provide evidence suggesting that in our empirical application the decision of an individual retailer is exogenous with respect to overall coalition spend and therefore do not consider endogeneity a concern in our model.

non-parametric function that assigns the terminal node $\mu_n \in \mathcal{M}$ to \mathbb{X} , based on the decision rules followed for each record in the data. With this notation, the sum-of-trees model is defined as follows across m trees:

$$f(\mathbb{X}) \equiv \sum_{j=1}^m g(\mathbb{X}; \mathcal{T}_j, \mathcal{M}_j) \quad (6)$$

Through combining such regression trees, BART allows any potential interactions and non-linearities among the input \mathbb{X} variables. To estimate this sum-of-trees model, the number of trees is pre-specified. The rest of the parameters, including the depth and the structure of each tree \mathcal{T}_j and its terminal nodes values \mathcal{M}_j are obtained through a variant of Bayesian backfitting methods outlined by Hastie and Tibshirani (2000). It iteratively constructs and fits residuals across MCMC draws, an approach similar to the gradient boosting approach of Friedman (2001). To avoid overfitting, each tree is intended to be small and explain only a portion of $f(\mathbb{X})$, by limiting its depth and each parameter μ_k being shrunk towards 0, which is achieved through a regularization prior.

The sum-of-trees, backfitting MCMC algorithm, and regularization prior form the essential features of BART. We review the prior specifications in Appendix 2, and the full marginal distributions in the MCMC estimation are detailed in Appendix 1. Additional technical details on the backfitting MCMC algorithm can be found in Chipman et al. (2010). Essentially, the backfitting MCMC algorithm draws the parameters associated with each tree, including \mathcal{T}_j and \mathcal{M}_j , from their joint conditional distribution given the parameters from all the other trees and the rest of the parameters associated with the model as specified in equation 5. Figure 8 presents the basic features of a Bayesian additive regression tree: many individual trees are simultaneously fit over the MCMC sampling process and then combined based on the posterior probabilities to obtain an estimate.



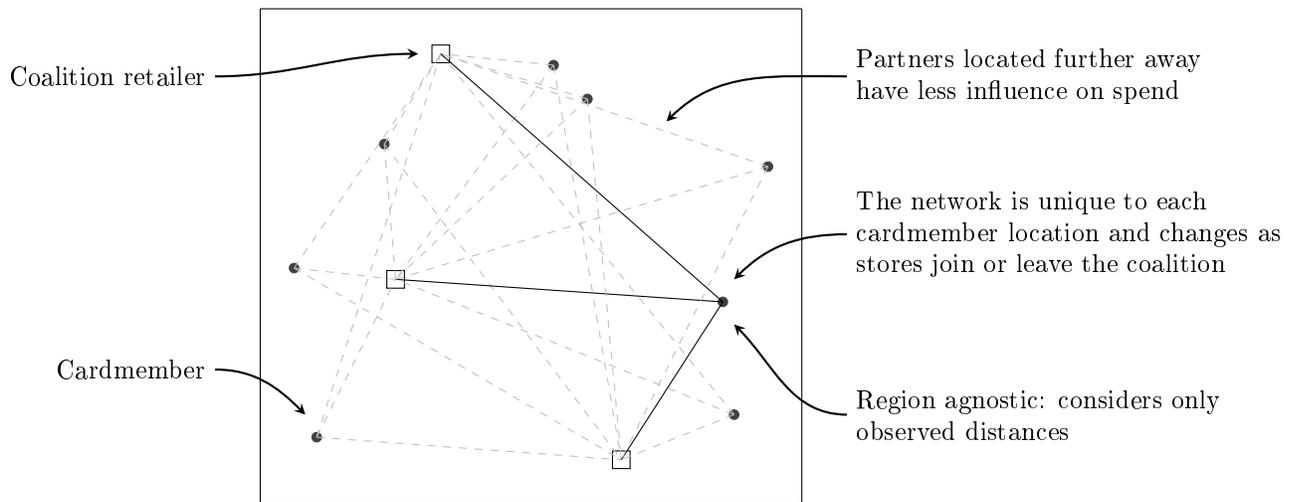
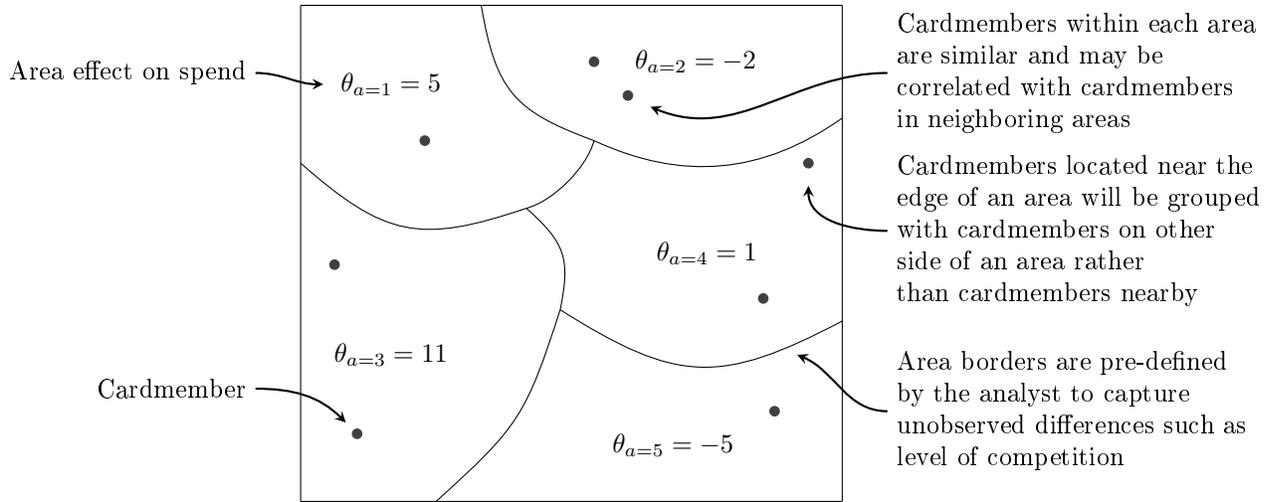
Here we present the basic building blocks of a two tree BART model. The first tree, \mathcal{T}_1 , in the left column ends up getting “trimmed” during the MCMC sampling process moving from draw r to draw $r + 1$ while the second tree, \mathcal{T}_2 , “grows” a terminal node. Each individual tree explains a relatively small portion of the data, but by adding the trees together the predictions can accommodate a highly nonlinear mapping from X to y .

Figure 8: Fundamental BART Components

We now have a fully specified model that controls for both 1) demand side effects through the regional fixed effects and spatially correlated error structure, and 2) supply side retailer level effects through BART. As discussed in Chipman et al. (2010), the BART component of this model can simply be added into the MCMC sampling process, allowing us to easily incorporate it into our spatial model. The primary advantage of this specification over our baseline model is that it explicitly controls for the retailer level effects on the supply side and their dynamics. In the data, variations both over time and across cardmembers will assist

the BART component with estimation. The variation over time comes from partners entering and leaving the network each month. The variation across cardmembers stems from the geographic distribution of the network with respect to each cardmember. As previously discussed, due to the limited distance traveled for in-network transactions, it is as if each cardmember essentially has their own local network with which they interact.

Figure 9 shows the primary differences between the CAR and BART components and highlights why both are advantageous in our model specification. The CAR spatial specification allows us to control for unobserved, regional level differences and focuses on the demand side influences among cardmembers. In contrast, the BART component accounts for the effects from the observed, but complex, retailer locations, which represents the supply side effects and their variations over time as retailers join or exit the coalition. Excluding either of these components could bias our main effects of interest related to the change in the loyalty program. Furthermore, including CAR but excluding BART could confound the estimated spatial effects since all coalition partner entry and exit behaviors would be absorbed (incorrectly) into the regional effects. Our augmented model contrasts with most prior spatial work in marketing, which generally clusters firm locations into regional areas thus potentially confounds location effects specific to the firm with regional fixed effects.



In both the top and bottom maps 8 cardmembers (the black dots) are located in the same spots. In the CAR prior (on top), no consideration is given to individual locations within an area: only differences across areas matter. This means that two cardmembers located near each other will be separated (and have different regional effects) if a border crosses between them. The analyst needs to be careful that the specified borders reflect natural groupings that would reasonably capture unobserved differences specific to an area. In the BART model (on bottom), no consideration is given to unobserved area level differences. In our specification, only the distance between each cardmember and each retailer (the white squares) matters. In essence, each cardmember experiences whichever “local” network surrounds them, which may evolve over time. The local network for one of the cardmembers is highlighted by the solid black lines. This variation across cardmembers allows BART, combined with the CAR prior to evaluate the value of each partner that participates in the coalition.

Figure 9: CAR (top) and BART (bottom) Comparison with 8 Cardmembers and 3 Coalition Retailers

5 Results

In this section we discuss the estimation results. We briefly review the SUR, spatial, and monthly results before moving onto the primary coefficients of interest related to the change in the LP structure.

5.1 SUR, Spatial, and Monthly Estimates

Table 4 shows the correlation estimate between the in-network and out-of-network spend from the SUR specification. A positive correlation is not surprising, since a high spending in-network cardmember is likely to be a high out-of-network spender, but the low magnitude of .09 suggests that in-network spend is only loosely indicative of out-of-network spend once the coalition network has been accounted for. This may be a reflection of the large heterogeneity in coalition partner access near each cardmember.

	95% HPD		
	Median	Lower	Upper
$\mathbb{E}[\varepsilon^I \varepsilon^O \cdot]$.088	.073	.104

Table 4: Estimated Correlation Between Spend and Redemption

Figure 10 presents the coefficients from the eleven month dummy variables (with December as the reference). For each of the two dependent variables (in-network and out-of-network monthly spend) we estimate the coefficients both excluding and including BART to control for the coalition network structure. The purpose of doing this is to determine to what extent the presence of BART changes the other coefficients in the model. We see that most of the monthly coefficients remain relatively stable after adding BART – this is somewhat expected since the structure of the network is unlikely to have a large influence on monthly seasonality effects.

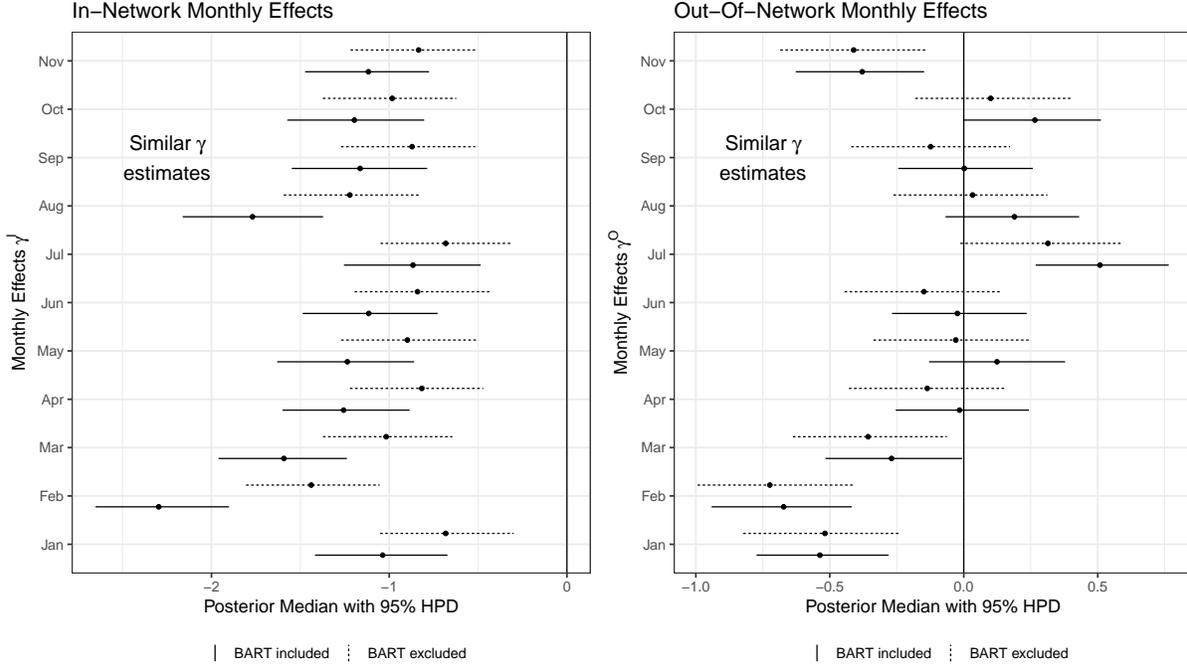


Figure 10: Monthly Coefficients

Table 5 presents the estimated model standard deviations (SD) σ and the parameters related with the CAR spatial component, including the regional standard deviations δ and the regional spillover ρ . To demonstrate the influence of including BART in the modeling, the results are presented for both models with and without the BART component. As shown in the table, by including BART, the model variance σ decreases slightly, suggesting that the BART component helps to explain some of the model variations. More interestingly, adding BART leads to a dramatic decrease in the cross-regional variance (captured through δ^2) for both dependent variables. Further, the regional spillover coefficient, ρ , changes from positive and statistically significant to insignificant in the model with BART. Summarizing the above results, we can see that by controlling for retailer locations using BART, information that otherwise would be attributed to unobserved demand side regional variations and spillover effects has been properly allocated into observed supply side effects. This applies to both in-network and out-of-network spending analysis.

Coef.	Description	y^I : In-Network Spend		y^O : Out-Of-Network Spend	
		BART Excluded	BART Included	BART Excluded	BART Included
σ	Model SD	6.00 (5.89, 6.11)	5.32 (5.23, 5.42)	4.99 (4.93, 5.05)	4.00 (3.96, 4.05)
δ	Regional SD	11.76 (8.79, 15.65)	2.84 (1.76, 4.11)	6.99 (5.16, 9.28)	2.80 (1.78, 4.08)
ρ	Reg. spillover	0.16 (0.05, 0.19)	-0.07 (-0.33, 0.17)	0.16 (0.04, 0.19)	-0.06 (-0.32, 0.18)

Table 5: Model Variance and Spatial Results. Posterior Median and 95% HPD

To further illustrate the effect of augmenting a traditional spatial model with BART, Figure 11 plots the posterior median of each of the 26 region effects, along with the 95% highest posterior density. Once again, we include the estimates both with and without the BART component. The regional effects were reduced dramatically once we control for retailer specific information via BART. Without BART, the analyst would mistakenly believe that the variation in the data is a result of the unobserved demand side effects unique to each region, when in fact the variation appears primarily related to observed retailer locations in the coalition.

Recall that regional effects, and similarly regional spillovers, are intended to capture unobserved effects unrelated to individual retailers. Our spatial BART specification allows us to separate the influence of the retailers' contribution at each location against regional level effects and better portray the uncertainty associated with regional level estimates.

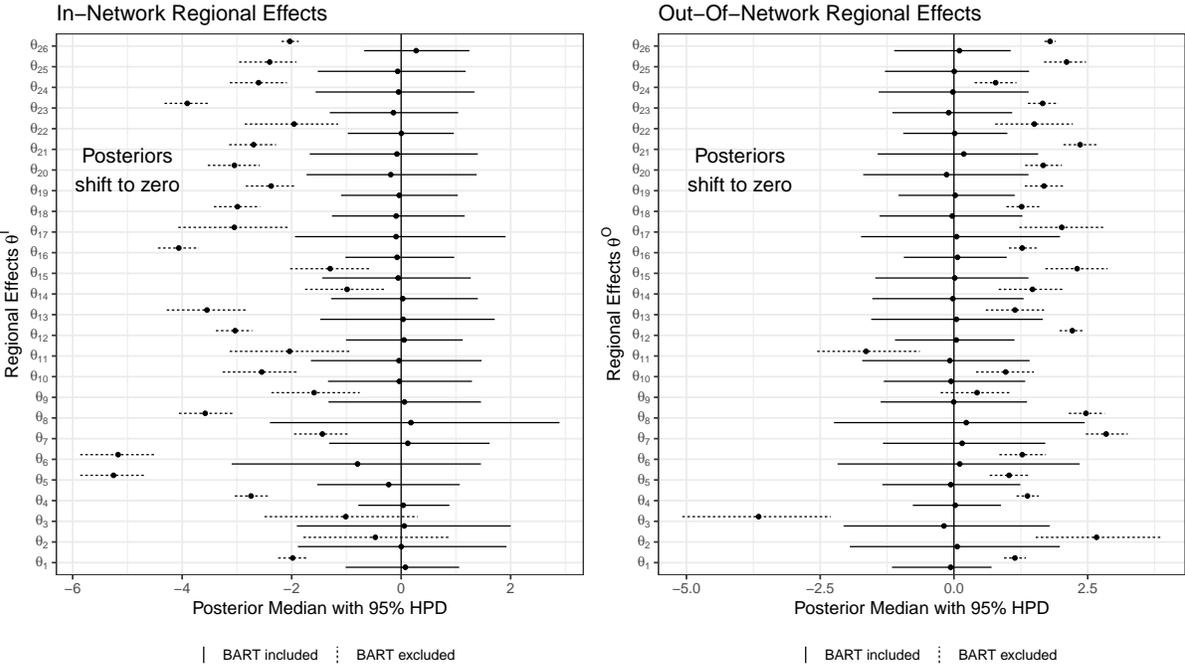


Figure 11: Regional Coefficients

5.2 Impact of Change in CLP Structure

In this section we focus on our coefficients of interest, the influence of the LP change and how it varies by customer segment. Our premise is that to properly measure the loyalty program rewards effect, even in the presence of a natural experiment, we must control for the evolving coalition network. Otherwise, any changes in the coalition network over time may be confounded with the estimated impact of the change in the LP. As previously discussed, we decided to employ BART to control for these complex relationships among retail

partners. While BART will not explain the mechanism by which the store relationships influence the LP estimates, controlling for these auxiliary variables enhances the model through reducing omitted variable bias in our coefficients of interest.

To examine the importance of including the CAR and BART, the results are presented from three different models. The first model applies the regression discontinuity approach while ignoring both CAR and BART; the second model adds only the CAR component, and the third model includes both.

First, we consider a model similar to the regression discontinuity presented in Table 3 from section 3.1. We include this table because after transforming the dependent variables, the regression discontinuity coefficients from before are not directly comparable to any coefficients presented in this section. Specifically, we take the full model specification but “turn off” both the CAR and BART components, but still retain individual cardmember fixed effects, seasonality, and our LP and segment coefficients of interest. With these components turned off, this model fails to control for any demand side effects for each region through the CAR specification, nor does it control for supply side effects from the coalition retail partners through BART. The results for both the in-network and out-of-network spend are presented in Table 6. The estimate of parameter κ shows that relative to the low spending group, the high spending group does not experience a statistically significant change in their in-network spending, but an increase only in the out-of-network spending. The estimate for ζ implies that after the LP changed its rewards structure, the in-network spending for the low segment is reduced (ζ is negative and statistically significant for y^I) while the out-of-network spending is increased (ζ is positive and statistically significant for y^O). This is expected, as shown in Table 2; the updated rewards structure makes the out-of-network purchase slightly more attractive than before. In addition, the spending for both in- and out-of-network are decreased more for the high spending group, as evaluated by ω , and the results are statistically significant.

Coef.	Description	y^I : In-Network Spend	y^O : Out-Of-Network Spend
κ	High Segment	1.44 (-0.18, 2.98)	3.72 (2.18, 5.29)
ζ	Post LP	-0.28 (-0.46, -0.10)	0.38 (0.17, 0.58)
ω	HS x Post LP	-0.53 (-0.77, -0.27)	-0.94 (-1.23, -0.66)

These are results are from the model specified but with the CAR prior and BART components removed. The coefficients have a similar interpretation to the regression discontinuity shown earlier in Table 3.

Table 6: Regression Discontinuity LP x Segment Results. Posterior Median and 95% HPD

We next move onto the two primary models we’ve discussed: the CAR model without BART, and then the full model which includes the BART component. The coefficients for the high spend segment and post LP change indicators, along with their interaction, are shown in Table 7.

The estimates for κ from the three models show that, with the CAR prior added and BART excluded, the second model has statistically significant estimate for the high spend group in both in- and out-of-network

spend, which was statistically insignificant in the first model for in-network spend. After adding BART, the estimate become more positive and statistically significant for both in- and out-of-network spend. The estimates for ζ are similar between the first two models, with or without the CAR prior (Tables 7 and 6). That is, we find that change in the LP appears to have a significant negative effect on the low segment in-network spend. However, after controlling for supply side effects with BART the coefficient is no longer statistically significant. This suggests that the change in the LP rewards structure is biased downward if the coalition structure is not accounted for. In other words, it is likely that key coalition partners exited after the coalition changed the LP rewards structure. The firm managing the coalition would fail to capture this insight unless they conditioned on coalition network structure. The qualitative result also applies to the out-of-network spend model. In particular, adding BART shifts the estimated impact of LP change to be more positive from 0.32 to 0.60. Considering the importance of evaluating the LP change on in-network spend, and the impact of BART, next we are going to focus on our discussion regarding this results.

Coef.	Description	y^I : In-Network Spend		y^O : Out-Of-Network Spend	
		BART Excluded	BART Included	BART Excluded	BART Included
κ	High Segment	1.90 (0.75, 3.07)	2.79 (1.46, 4.23)	1.89 (0.70, 2.97)	4.20 (2.44, 5.86)
ζ	Post LP	-0.32 (-0.56, -0.10)	-0.07 (-0.37, 0.27)	0.32 (0.14, 0.49)	0.60 (0.36, 0.86)
ω	HS x Post LP	-0.64 (-0.96, -0.33)	-0.68 (-0.99, -0.35)	-0.83 (-1.06, -0.58)	-0.91 (-1.12, -0.71)

Table 7: Results from models with LP x Segment adding CAR with and without BART. Posterior Median and 95% HPD.

The biases from excluding the BART are illustrated in Figure 12. Here we present the posterior distributions of post LP effect by spend segment for each of the two dependent variables. In creating this plot, the effects on the low segment refer to the estimates of ζ , and those related to the high segment are calculated as the posterior sum of $\zeta + \omega$. By controlling for the evolving coalition network structure, we find that the change in the loyalty program shifts all spend predictions up, so much so that for the low spend group the impact of the LP change on in-network spend becomes insignificant. In short, changes in the network structure that lead to decreases in spend were being absorbed into the LP coefficients. After the evolving network was controlled for using BART, the downward biases were reduced.

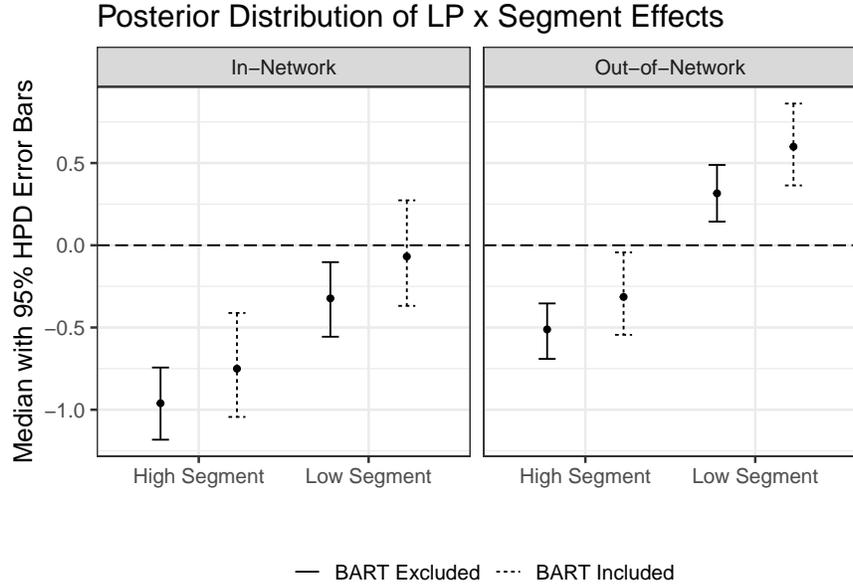


Figure 12: LP Effects

The findings from this section highlight the importance of implementing (or at the very least considering) methods like BART in marketing models. Even in the presence of a natural experiment, there may be numerous, complex ancillary variables that need to be conditioned on and accounted for. In our empirical setting, the coalition loyalty program would fail to see that the change in the LP had no negative impact on the in-network spending from the low segment.

6 Supplementary Analyses

In this section we present two supplementary analyses that highlight the value of incorporating BART into the model in our empirical context. In the first analysis, we estimate the total value of the coalition relative to the impact of changing the rewards structure, and in the second analysis we explore the relationship between partner store distance and cardmember spend.

6.1 Relative Impact of Network to LP Change

A natural extension to the prior analysis is to determine which is more impactful: a change in the LP design or a change in the retail network structure? In other words, what type of coalition network reconfiguration would have the same impact as changing the design of the loyalty program? For a coalition loyalty program, there is much more strategic flexibility in selecting individual partners to add or remove from the coalition (which can be tailored to individual cardmember locations) than changing the LP design, which affects

all cardmembers equally. By using BART, we are able to estimate spend patterns under various network structures using coefficients that would have been intractable to obtain using a traditional econometric approach alone.

We present the two simplest comparisons: 1) change in expected monthly spend with no retail partners in the coalition versus all partners in the coalition (assuming a post-LP design); and 2) change in expected monthly spend pre-LP and post-LP (assuming the average geographic influence across all cardmembers over time – which can be interpreted as the average network structure). Since our model is highly non-linear, we cannot isolate the impacts of the coalition network without considering the design of the loyalty program, and vice versa. For instance, a change in the coalition structure (i.e., which retailers are participating at a given point in time) depends on the current LP structure (pre or post change).

The firm level spend estimates are shown in Figure 13. We find that going from an empty to a full network (i.e., no retail partners to all possible retail partners) increases estimated monthly spend by about \$24,930 (about \$40/cardmember per month), but the change in the LP decreases estimated monthly spend by about \$33,460 (about \$52/cardmember per month). In short, the change in the LP structure was more detrimental than the entire benefits of the coalition network. This highlights the sensitivity of these cardmembers to the rewards structure – their engagement with the card appears highly dependent on the rewards structure. Thus the negative modification to the LP (from the perspective of a customer earning rewards) is met with a substantial decline in use. We note that this is total spend across segments, so even if there was no impact to the in-network spend for low segment cardmembers, the negative effects from the high segment prevailed.

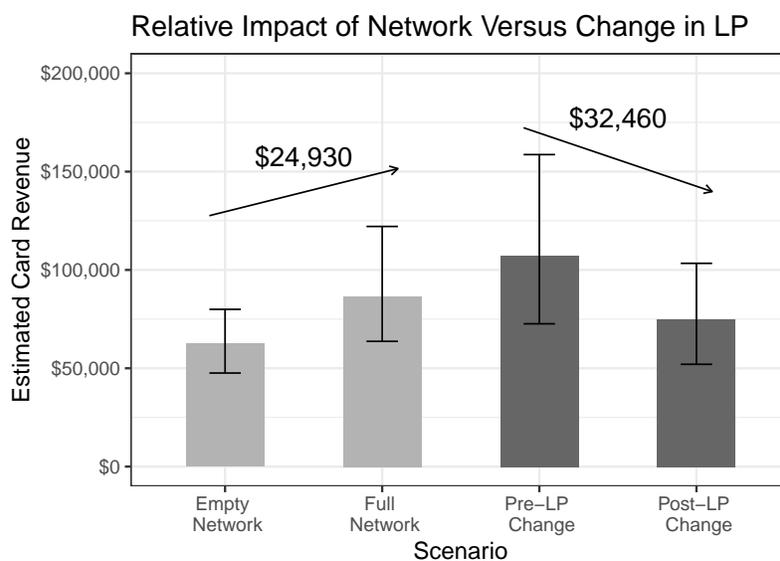


Figure 13: Impact of Network Structure Relative to LP Change

While this simplistic example gives a sense of the relative impact of the change in the LP structure relative to potential changes in the network structure, it gets us no further in understanding which of the partners may be contributing most to the coalition network. While BART does not provide individual coefficients on each of the partner stores, we can easily characterize the marginal effects by using a partial dependent function approach. Specifically, we evaluated the difference in estimated spend between when a coalition partner is in-network (at its current distance from each cardmember) and when the partner is out-of-network (and thus the geographic influence is 0). We do this for all 2,329 partners in the data, and find individual marginal effects are overwhelmingly negligible. In fact, over 80% of the locations have no effect on monthly spend predictions and drop out of the BART predictions entirely. This may seem unusual given that we previously showed a substantial change in spend when all stores are in the network versus removed, but it reflects a key advantage of the BART model and its ability to manage higher order interactions between individual retail locations. In other words, the effects among stores are not necessarily additive but instead can be synergistic among partners. Once estimation is complete, identifying these key partners is non-trivial and outside the scope of this paper, as it involves an extensive search routine.

6.2 Marginal Effect of Distance

In this section we further investigate the relationship between customers spend and their distance to retail partners. We first consider how distance to each partner influences spend at an aggregate level, focusing on the tradeoff between in-network and out-of-network spend. To do this, we set the average customer-store distance as a baseline and then deviate from the average distance to each store at fixed intervals. Figure 14 presents the results when we grow or contract the average distance by increments of 10%. For example, 1 is the current network at the average distance, whereas 1.1^5 is the distance after five 10% increases, or a 61% increase relative to the baseline ($1.1^5 = 1.61$). Since the model contains other components (such as customer level fixed effects), in this graph we simply plot the changes in the BART component $f(X)$. We see that as distance increases, the log of in-network spend decreases, everything else is held constant. However, some of this decline appears to be offset by increases in out-of-network spending in more sparse (i.e., larger) networks. For relatively concentrated networks, in-network spend tends to increase at the expense of out-of-network spending.

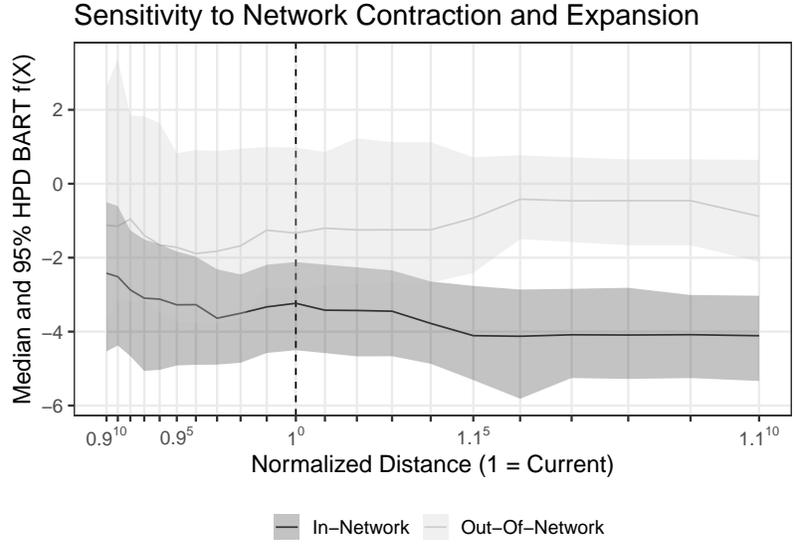


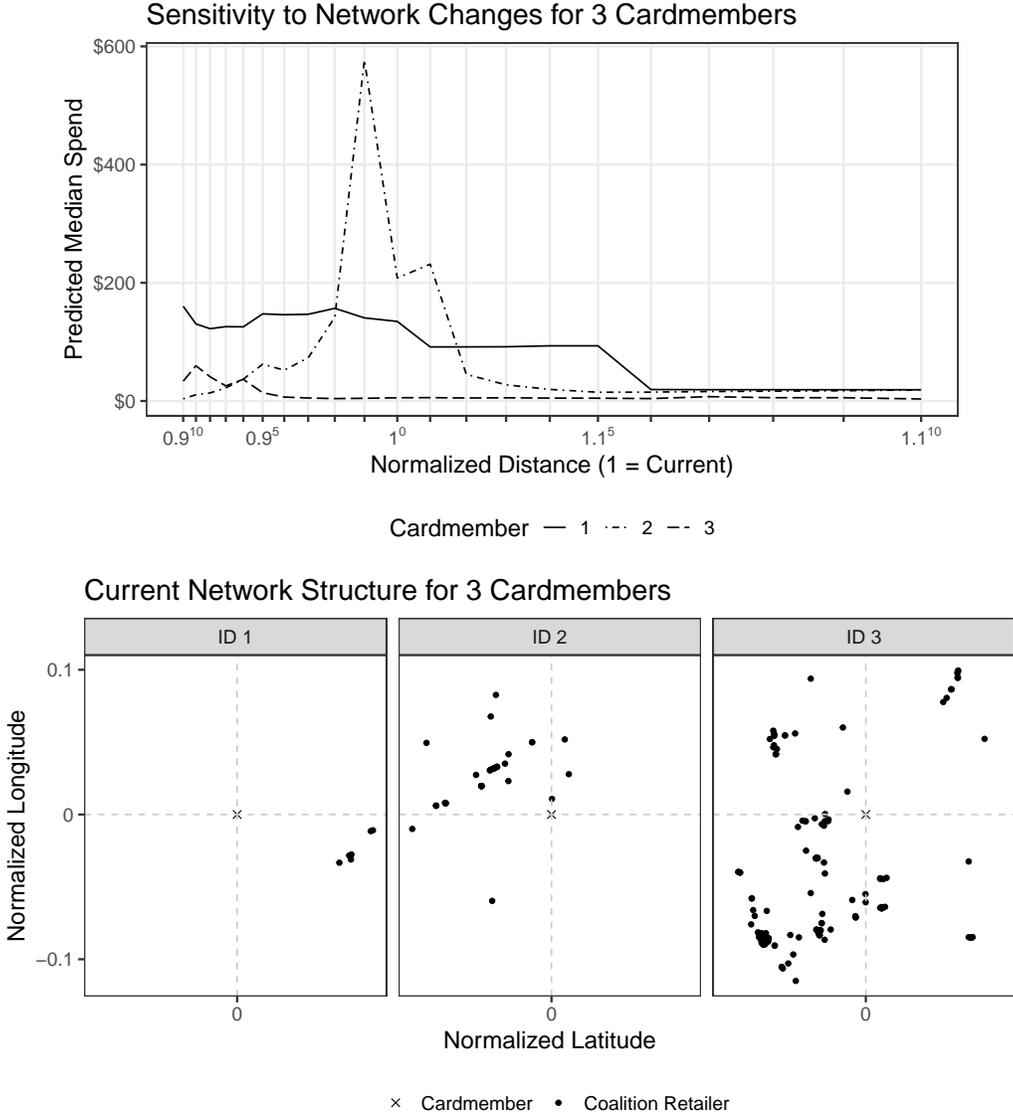
Figure 14: Aggregate Distance Sensitivity

To investigate this phenomenon at the individual customer level, we select three cardmembers from the dataset, each with slightly different sensitivities to distance. Each of these cardmembers will have their own fixed effect, which will account for different baseline spending patterns, and their own regional effect, which will control for unobserved demand side effects. The predicted spend patterns and their local network structure (the coalition network within 10km, since most transactions do not occur much further away), are presented in Figure 15.

We see that for Cardmember 1, expected monthly spend is around \$150 but slowly tapers off as the stores move further away. This first cardmember has relatively strong spend considering the network is relative sparse. This could be a reflection of limited outside spend opportunities. For Cardmember 2, expected revenue is slightly higher at \$200 for the current network, and then increases as the distances to stores contract slightly but then declines if the distances contract too much. This could reflect a synergy between store locations which is only maximized when a “sweet spot” is reached in terms of relative distances to each other. Finally, Cardmember 3 has very low projected spend, even though their network is most dense. Cardmember 3 may be in a region that has a high level of competition, or the composition of stores may not be attractive. Alternatively, this may reflect what we discussed in section 6.1: that most coalition partners simply do not have any influence on monthly spending patterns.

These predictions demonstrate the flexibility of incorporating BART into the spatial model. Not only can we capture the granular spatial relationship between each retail partner and cardmember, we are also able to estimate very flexible (i.e., non-linear) relationships between the distance to a particular retailer and

its influence on customer spend. Recall that this type of analysis is infeasible without a method such as BART to deal with a large number of potentially interacting parameters associated with the retail partners' distances from the customer. While our primary interest is evaluating the effect of the change in the LP, the coalition can leverage supplementary analyses like this one to strategically identify retail partners that are most influential on cardmember spend.



The top graph shows the changes in the expected median total spend as the network around three cardmember from the data expands or contracts. The bottom graph shows the actual network of stores the cardmember has access to, where the graph area is limited to a 10km radius and the cardmember placed in the center. The latitude and longitude has been normalized to ensure anonymity for the firm. The cardmembers each have very different spend patterns, which do not necessarily correlate strictly with the number of stores nearby. BART identifies which retail partner locations add value to the coalition.

Figure 15: Network Sensitivity for Three Cardmembers

7 Conclusion

In this paper we evaluate the impact of a change in a coalition loyalty program on cardmember spend. We argue that to properly estimate this effect it is important to control for both the unobserved demand side effects at the regional level, and the observed supply side effects specific to the coalition network structure. In our model, the demand side effects are captured via a CAR framework, as is typically done with extant spatial models. To control for the supply side effects we augmented our baseline spatial model with a BART component.

Our empirical results show that both a naive regression discontinuity and a traditional spatial model that accounts for unobserved effects at the regional level lead to the conclusion that the change in the LP had a negative effect on spend within the firm’s low spend segment. However, this effect is reduced dramatically when the model was augmented with the BART component to control for the evolving coalition network. The shift in the results suggests that the decrease in spend after the loyalty program change was likely influenced by changes in the coalition network (i.e., partner composition). Our empirical analysis offers one feasible solution to the common situation in marketing where a modeler is faced with a large number of potentially complex and interacting variables that are too cumbersome to specify using a traditional functional form. By augmenting our core model with BART, we were able to deal with these control variables while still retaining a specified functional form on our coefficients of interest. In other words, we did not need to send the whole analysis through a nonlinear machine learning approach – BART was applied only to the control variables that we suspected would influence the outcome variable in a complex and nonlinear manner.

While our primary goal was to estimate the impact of the change in the LP on spend, we also hoped to illustrate the potential of augmenting familiar marketing models with BART so that marketers can accommodate evermore complex modeling situations. Leveraging our empirical context, we demonstrate the implications derived from including the supply side effects using BART, via two supplementary analyses. In the first, by comparing the impact of the coalition structure with the change in the loyalty program, we find that the change in the LP point structure, although appears to be inconsequential, is more detrimental than removing all the retail partners in influencing members’ purchases. The second analysis illustrates how purchase reduces as the distance to stores is increased, through both aggregate level and individual level visualizations. It shows the change in travel distance to the store partners do have important impact on both in- and out-of-network spend, but the impact varies across cardmembers.

In the field of machine learning there has been a common debate between linear and nonlinear models. The advantages of linear models, such as the regression discontinuity model applied in our analysis as the baseline model, lie primarily in its tractability in estimation and interpretability in modeling results.

However, these models are generally restrictive and need pre-defined functional form of the input features (X variables). Nonlinear models, such as BART, offer more flexibility in capturing the possible impacts of a large number of input variables and their synergies. By enriching the classical spatial model with BART, our approach takes the advantages of both linear and nonlinear models in order to meet the objective of our analysis: estimating the impact of changing the loyalty program on cardmember spend, while controlling for the interactions and dynamics among the participating retailers. We believe this approach will find more applications in analyzing marginal impacts of firm actions while controlling for many other related input variables, especially when these control variables are interrelated in complex ways.

References

- Allenby, G. M., Bradlow, E. T., George, E. I., Liechty, J., and McCulloch, R. E. (2014). Perspectives on bayesian methods and big data. *Customer Needs and Solutions*, 1(3):169–175.
- Aravindakshan, A., Peters, K., and Naik, P. A. (2012). Spatiotemporal allocation of advertising budgets. *Journal of Marketing Research*, 49(1):1–14.
- Athey, S. and Imbens, G. W. (2017). The state of applied econometrics: Causality and policy evaluation. *Journal of Economic Perspectives*, 31(2):3–32.
- Bonato, V., Baladandayuthapani, V., Broom, B. M., Sulman, E. P., Aldape, K. D., and Do, K.-A. (2011). Bayesian ensemble methods for survival prediction in gene expression data. *Bioinformatics*, 27(3):359–367.
- Bradlow, E. T., Bronnenberg, B., Russell, G. J., Arora, N., Bell, D. R., Duvvuri, S. D., Ter Hofstede, F., Sismeiro, C., Thomadsen, R., and Yang, S. (2005). Spatial models in marketing. *Marketing Letters*, 16(3-4):267–278.
- Breugelmans, E., Bijmolt, T. H., Zhang, J., Basso, L. J., Dorotic, M., Kopalle, P., Minnema, A., Mijnlief, W. J., and Wunderlich, N. V. (2015). Advancing research on loyalty programs: a future research agenda. *Marketing Letters*, 26(2):127–139.
- Bronnenberg, B. J. and Sismeiro, C. (2002). Using multimarket data to predict brand performance in markets for which no or poor data exist. *Journal of Marketing Research*, 39(1):1–17.
- Carnegie, N., Dorie, V., and Hill, J. L. (2019). Examining treatment effect heterogeneity using bart. *Observational Studies*, 5(2):52–70.
- Chipman, H. A., George, E. I., McCulloch, R. E., et al. (2010). Bart: Bayesian additive regression trees. *The Annals of Applied Statistics*, 4(1):266–298.
- Cressie, N. A. C. (1993). *Statistics for spatial data*. Wiley series in probability and mathematical statistics. Applied probability and statistics. Wiley, New York, rev. ed.. edition.
- Dorotic, M., Fok, D., Verhoef, P., and Bijmolt, T. (2011). Do vendors benefit from promotions in a multi-vendor loyalty program? *Marketing Letters*, 22(4):341–356.
- Ferguson, R. and Hlavinka, K. (2007). The colloquy loyalty marketing census: sizing up the us loyalty marketing industry. *Journal of Consumer Marketing*.

- Friedman, J. H. (2001). Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232.
- Gelfand, A. E. and Smith, A. F. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American statistical association*, 85(410):398–409.
- Green, D. P. and Kern, H. L. (2012). Modeling heterogeneous treatment effects in survey experiments with bayesian additive regression trees. *Public opinion quarterly*, 76(3):491–511.
- Hahn, P. R., Murray, J. S., Carvalho, C. M., et al. (2020). Bayesian regression tree models for causal inference: Regularization, confounding, and heterogeneous effects (with discussion). *Bayesian Analysis*, 15(3):965–1056.
- Hastie, T. and Tibshirani, R. (2000). Bayesian backfitting (with comments and a rejoinder by the authors. *Statistical Science*, 15(3):196–223.
- Heckman, J. J., Lopes, H. F., and Piatek, R. (2014). Treatment effects: A bayesian perspective. *Econometric reviews*, 33(1-4):36–67.
- Hill, J., Linero, A., and Murray, J. (2020). Bayesian additive regression trees: a review and look forward. *Annual Review of Statistics and Its Application*, 7:251–278.
- Hill, J. L. (2011). Bayesian nonparametric modeling for causal inference. *Journal of Computational and Graphical Statistics*, 20(1):217–240.
- Huang, H.-C. (2001). Bayesian analysis of the sur tobit model. *Applied Economics Letters*, 8(9):617–622.
- Kindo, B. P., Wang, H., and Peña, E. A. (2016). Multinomial probit bayesian additive regression trees. *Stat*, 5(1):119–131.
- Lal, R. and Bell, D. (2003). The impact of frequent shopper programs in grocery retailing. *Quantitative Marketing and Economics*, 1(2):179–202.
- Lemon, K. N. and Wangenheim, F. V. (2009). The reinforcing effects of loyalty program partnerships and core service usage: a longitudinal analysis. *Journal of Service Research*, 11(4):357–370.
- Liu, Y. (2007). The long-term impact of loyalty programs on consumer purchase behavior and loyalty. *Journal of marketing*, 71(4):19–35.
- Liu, Y. and Yang, R. (2009). Competing loyalty programs: Impact of market saturation, market share, and category expandability. *Journal of Marketing*, 73(1):93–108.

- Meyer-Waarden, L. (2008). The influence of loyalty programme membership on customer purchase behaviour. *European Journal of marketing*.
- Moon, S. and Russell, G. J. (2004). A spatial choice model for product recommendations. *Marketing Science Institute Report*, (04-120).
- Moore, G. and Sekhon, H. (2005). Multi-brand loyalty cards: a good idea. *Journal of Marketing Management*, 21(5-6):625–640.
- O’Brien, L. and Jones, C. (1995). Do rewards really create loyalty? *Harvard Business Review*, 73(3):75–82.
- Russell, G. J. and Petersen, A. (2000). Analysis of cross category dependence in market basket selection. *Journal of Retailing*, 76(3):367–392.
- Rust, R. T. and Verhoef, P. C. (2005). Optimizing the marketing interventions mix in intermediate-term crm. *Marketing science*, 24(3):477–489.
- Sparapani, R. A., Logan, B. R., McCulloch, R. E., and Laud, P. W. (2016). Nonparametric survival analysis using bayesian additive regression trees (bart). *Statistics in medicine*, 35(16):2741–2753.
- Stourm, V., Bradlow, E., and Fader, P. (2017). Market positioning using cross-reward effects in a coalition loyalty program. *HEC Paris Research Paper No. MKG-2017-1242*.
- Sun, D., Tsutakawa, R. K., and Speckman, P. L. (1999). Posterior distribution of hierarchical models using car (1) distributions. *Biometrika*, 86(2):341–350.
- Taylor, G. and Neslin, S. (2005). The current and future sales impact of a retail frequency reward program. *Journal of Retailing*, 81(4):293–305.
- Trusov, M., Bodapati, A. V., and Bucklin, R. E. (2010). Determining influential users in internet social networks. *Journal of marketing research*, 47(4):643–658.
- Van Nierop, E., Fok, D., and Franses, P. H. (2008). Interaction between shelf layout and marketing effectiveness and its impact on optimizing shelf arrangements. *Marketing Science*, 27(6):1065–1082.
- Verhoef, P. (2003). Understanding the effect of customer relationship management efforts on customer retention and customer share development. *Journal of Marketing*, 67(4):30–45.
- Yang, S. and Allenby, G. M. (2003). Modeling interdependent consumer preferences. *Journal of Marketing Research*, 40(3):282–294.

- Zellner, A. (1962). An efficient method of estimating seemingly unrelated regressions and tests for aggregation bias. *Journal of the American statistical Association*, 57(298):348–368.
- Zhang, J. and Breugelmans, E. (2012). The impact of an item-based loyalty program on consumer purchase behavior. *Journal of Marketing research*, 49(1):50–65.
- Zhang, S., Shih, Y.-C. T., Müller, P., et al. (2007). A spatially-adjusted bayesian additive regression tree model to merge two datasets. *Bayesian Analysis*, 2(3):611–633.

Appendix 1: Markov Chain Monte Carlo Estimation

We carried out estimation through sequentially generating draws from the following full conditional distributions. With the exception of the SUR components, the drawing process between in-network spend (y^I) and out-of-network spend (y^O) spend is nearly identical. The outline below refers to the in-network spend draws.

1. Update BART

- (a) Create $y_{it}^* = y_{it} - \theta_a - W_i \Delta - \frac{\sigma_{OI}}{\sigma_{OO}} e_i^O$ where e^O is the residual between the predicted out-of-network spend and augmented data y^O and W_i is the matrix specific to cardmember i containing all the fixed effects, seasonality variables S , spend group G , post-LP indicator Z , and the interaction GZ .
- (b) Estimate function BART $f()$ in $y^* = f(X)$

2. Update missing data z

- (a) Create $y_{it}^* = f(X_{it}) + \theta_a + W_i \Delta - \frac{\sigma_{OI}}{\sigma_{OO}} e_i^O$
- (b) For observations with missing data in y_{it} (e.g., zero in-network spend), draw $f(z|\cdot) \sim \text{Truncated Normal}(y^*, -\infty, 0)$
- (c) Replace the missing data in y_{it} with the draws from z

3. Update Σ

- (a) Update $e_{it}^O = \hat{y}_{it}^O - y_{it}^O$ and $e_{it}^I = \hat{y}_{it}^I - y_{it}^I$ and join the columns into E
- (b) Draw $p(\Sigma|\cdot) \sim \text{IW}\left(\nu_\Sigma + n, \left(E' E + V\right)^{-1}\right)$
- (c) $\nu_\Sigma = 3$ and V is a 2×2 identity matrix multiplied by ν_Σ

4. Update regional effects θ

- (a) Create $y_{it}^* = y_{it} - f(X_{it}) - W_i \Delta$
- (b) Set $\theta_\Sigma = \left[\frac{U' U}{\sigma_{ss}} + \frac{(H - \rho C)}{\delta^2}\right]^{-1}$ where U is an indicator design matrix of size $I \times A$ (number of cardmembers by number of regions)
- (c) Set $\theta_\mu = \theta_\sigma \left(\frac{U' y^*}{\sigma_{ss}}\right)$
- (d) Draw $p(\theta|\cdot) \sim \text{MVN}(\theta_\mu, \theta_\Sigma)$

5. Update regional variance δ^2

- (a) $p(\delta^2|\cdot) \sim \text{IG}(a, b)$

- (b) $a = a_2 + A/2$ with $a_2 = .001$ and $A = 26$ regions
- (c) $b = b_2 + \left(\theta' B' \theta B\right)$ where $B = (H - \rho C)$ and $b_2 = .001$
6. Update regional correlation ρ using Metropolis-Hastings algorithm with a random walk, similar to Yang and Allenby (2003)
- (a) Let the proposed draw be given by $\rho^{(n)} = \rho^{(p)} + \epsilon$ where $\rho^{(p)}$ is the previous draw and ϵ is a draw from the normal density with mean 0 and variance .0025
- (b) Accept the proposed draw with the following probability $\alpha = \min \left\{ \frac{|B(\rho^{(n)})| \exp \left[-.5(1/\delta^2) \theta' B(\rho^{(n)})' B(\rho^{(n)}) \theta \right]}{|B(\rho^{(p)})| \exp \left[-.5(1/\delta^2) \theta' B(\rho^{(p)})' B(\rho^{(p)}) \theta \right]}, 1 \right\}$ where $B(\rho) = H - \rho C$ and the probability of acceptance is zero if $\rho^{(n)}$ is outside the range of $[\lambda_{\min}^{-1}, \lambda_{\max}^{-1}]$
7. Update all other coefficients Δ
- (a) Create $y_{it}^* = y_{it} - f(X_{it}) - \theta_a$
- (b) Estimate Δ from model of $y_{it}^* = W_i \Delta + \epsilon$ where $\epsilon \sim N(0, \sigma^2)$ where σ^2 is already drawn in Σ
- (c) Draw $p(\Delta|\cdot) \sim N(\bar{\Delta}, V_{\Delta})$ where V_{Δ} equals the inverse of an identity matrix times .01, with number of columns equal to that of W

Appendix 2: Prior Specification

We estimate the conditional spatial BART model using Markov Chain Monte Carlo. The method of estimation requires specification of the prior distributions for the model parameters and derivation of the full conditional distributions. We set the priors to be diffuse and use conjugate priors when possible. First, we complete the Bayesian hierarchical model with priors $p(f)$, $p(\theta)$, $p(\Delta)$ for f , θ , and Δ , respectively. We save our discussion of $p(\sigma^2)$ for later when we introduce the SUR structure. We assume *a priori* independence.

The BART model is indexed by $\{(T_j, M_j), j = 1, \dots, m\}$. We then have:

$$p(f) = \prod_{j=1}^m p(T_j, M_j) = \prod_{j=1}^m \{p(T_j) \cdot p(M_j|T_j)\} \quad (7)$$

As in Chipman et al. (2010), we define $p(T_j)$ by three factors, corresponding to a node being non-terminal, the selection of the splitting variable for the a non-terminal node, and the choice of the splitting value conditional on a chosen splitting variable. The probability that a node at depth d is non-terminal is assumed to be

$$\alpha (1 + d)^{-\gamma} \quad (8)$$

where $\alpha \in (0, 1)$ and $\gamma \in [0, \infty)$ are two hyperparameters reflecting the prior belief about the tree. If the depth of the tree is believed to be small, a large value is assigned for γ , so that the probability decays fast with d . We use the proposed $\alpha = 0.95$ and $\gamma = 2$ as default values, which implies that with prior probability 0.05, 0.55, 0.38, 0.09 and 0.03 the tree as 1, 2, 3, 4, and 5 or more terminal nodes, respectively.

A natural choice for the selection of the splitting variable, conditional on a node being non-terminal, is a uniform prior over all available variances. A default choice for the prior distribution of the splitting value is a uniform choice of available splitting values.

We also define a prior for M_j . Let M_{jk} be the k th element of M_j . Conditional on T_j , we assume i.i.d. normal priors for M_{jk} . The mean and variance of the normal prior are specified in such a way that each tree is a weak learner, and each individual tree plays a small role in the overall fit. See section 3.2 of Chipman et al. (2010) for more details.

For all other coefficients Δ we use a multivariate normal prior:

$$\Delta \sim MVN(\bar{\Delta}, V_{\Delta}) \quad (9)$$

We complete the prior distribution specifications with those for the hyperparameters ρ , δ^2 , $\bar{\Delta}$, and V_{Δ} .

We assume $p(\delta^2)$ to be an inverse Gamma distribution, denoted by $IG(a, b)$, with density function

$$p(\delta^2) \propto \frac{1}{(\delta^2)^{a+1}} \exp\left(-\frac{b}{\delta^2}\right) \quad (10)$$

where a and b are fixed hyperparameters. Finally, the hyperparameters on the demographic effects $\bar{\Delta}$ and V_{Δ} are fixed.

The Markov Chain proceeds by generating draws from the set of conditional posterior distributions of the parameters (Gelfand and Smith, 1990). After subtracting out the spatial random effects θ_a and demographic effects Δ , from y_{it} the estimation of f is a standard BART model. The conditional distributions of the remaining model parameters, given f , are of standard form.

Appendix 3: Full Model Results

The first four columns adhere to a traditional spatial model and exclude BART. Models one through three incrementally add the customer and group variables of interest to the baseline spatial model. The fifth column augments the full spatial model with BART.

Coefficient	In-Network Spend: y^I					Description
	Model 1	Model 2	Model 3	BART Excluded	BART Included	
σ	6.00 (5.90, 6.09)	6.00 (5.90, 6.10)	6.00 (5.90, 6.10)	6.00 (5.89, 6.11)	5.32 (5.23, 5.42)	model variance
γ_1	-0.62 (-0.97, -0.23)	-0.69 (-1.07, -0.31)	-0.67 (-1.05, -0.30)	-0.68 (-1.05, -0.30)	-1.04 (-1.42, -0.67)	January
γ_2	-1.39 (-1.77, -1.02)	-1.44 (-1.84, -1.07)	-1.43 (-1.80, -1.06)	-1.44 (-1.81, -1.05)	-2.30 (-2.65, -1.90)	February
γ_3	-0.97 (-1.33, -0.60)	-1.02 (-1.39, -0.63)	-1.01 (-1.38, -0.63)	-1.02 (-1.37, -0.64)	-1.59 (-1.96, -1.24)	March
γ_4	-0.77 (-1.15, -0.41)	-0.83 (-1.20, -0.43)	-0.81 (-1.18, -0.41)	-0.82 (-1.22, -0.47)	-1.26 (-1.60, -0.89)	April
γ_5	-0.84 (-1.22, -0.48)	-0.90 (-1.28, -0.54)	-0.88 (-1.27, -0.51)	-0.90 (-1.27, -0.51)	-1.24 (-1.63, -0.86)	May
γ_6	-0.79 (-1.15, -0.42)	-0.85 (-1.23, -0.46)	-0.84 (-1.19, -0.44)	-0.84 (-1.19, -0.43)	-1.12 (-1.49, -0.73)	June
γ_7	-0.63 (-1.00, -0.24)	-0.69 (-1.05, -0.30)	-0.67 (-1.05, -0.27)	-0.68 (-1.05, -0.31)	-0.87 (-1.25, -0.49)	July
γ_8	-1.17 (-1.57, -0.80)	-1.24 (-1.61, -0.87)	-1.21 (-1.61, -0.85)	-1.22 (-1.59, -0.83)	-1.77 (-2.16, -1.37)	August
γ_9	-0.82 (-1.19, -0.46)	-0.88 (-1.26, -0.51)	-0.86 (-1.25, -0.48)	-0.87 (-1.27, -0.52)	-1.16 (-1.55, -0.79)	September
γ_{10}	-0.93 (-1.31, -0.56)	-0.99 (-1.39, -0.63)	-0.97 (-1.36, -0.59)	-0.98 (-1.37, -0.62)	-1.20 (-1.57, -0.80)	October
γ_{11}	-0.84 (-1.21, -0.50)	-0.90 (-1.26, -0.55)	-0.83 (-1.18, -0.48)	-0.83 (-1.22, -0.50)	-1.12 (-1.47, -0.78)	November
κ		1.44 (0.28, 2.55)	1.75 (0.60, 2.92)	1.90 (0.75, 3.07)	2.79 (1.46, 4.23)	Customer Group
ζ			-0.64 (-0.80, -0.49)	-0.32 (-0.56, -0.10)	-0.07 (-0.37, 0.27)	Post LP
ω				-0.64 (-0.96, -0.33)	-0.68 (-0.99, -0.35)	Group x LP
δ	11.75 (8.92, 15.74)	11.72 (8.60, 15.44)	11.65 (8.51, 15.29)	11.76 (8.79, 15.65)	2.84 (1.76, 4.11)	regional SD
ρ	0.16 (0.05, 0.19)	0.16 (0.06, 0.19)	0.16 (0.06, 0.19)	0.16 (0.05, 0.19)	-0.07 (-0.33, 0.17)	reg. spillover

Table 8: In-Network Partial Model Results

Out-Of-Network Spend: y^O						
Coefficient	Model 1	Model 2	Model 3	BART Excluded	BART Included	Description
σ	4.99 (4.93, 5.05)	4.99 (4.92, 5.04)	4.99 (4.93, 5.05)	4.99 (4.93, 5.05)	4.00 (3.96, 4.05)	model variance
γ_1	-0.47 (-0.75, -0.18)	-0.51 (-0.79, -0.21)	-0.51 (-0.81, -0.23)	-0.52 (-0.82, -0.24)	-0.54 (-0.77, -0.28)	January
γ_2	-0.67 (-0.96, -0.38)	-0.71 (-1.00, -0.43)	-0.72 (-1.02, -0.43)	-0.72 (-0.99, -0.41)	-0.67 (-0.94, -0.42)	February
γ_3	-0.30 (-0.61, -0.04)	-0.35 (-0.64, -0.07)	-0.35 (-0.63, -0.05)	-0.36 (-0.64, -0.06)	-0.27 (-0.52, -0.01)	March
γ_4	-0.09 (-0.38, 0.21)	-0.13 (-0.42, 0.16)	-0.14 (-0.41, 0.17)	-0.14 (-0.43, 0.16)	-0.02 (-0.25, 0.24)	April
γ_5	0.02 (-0.26, 0.32)	-0.02 (-0.30, 0.28)	-0.03 (-0.32, 0.27)	-0.03 (-0.34, 0.25)	0.12 (-0.13, 0.38)	May
γ_6	-0.09 (-0.39, 0.20)	-0.14 (-0.44, 0.14)	-0.15 (-0.42, 0.16)	-0.15 (-0.44, 0.14)	-0.02 (-0.27, 0.23)	June
γ_7	0.36 (0.10, 0.67)	0.32 (0.03, 0.61)	0.32 (0.02, 0.61)	0.31 (-0.01, 0.59)	0.51 (0.27, 0.77)	July
γ_8	0.08 (-0.20, 0.37)	0.04 (-0.25, 0.33)	0.03 (-0.23, 0.34)	0.03 (-0.26, 0.31)	0.19 (-0.07, 0.43)	August
γ_9	-0.07 (-0.36, 0.23)	-0.12 (-0.41, 0.18)	-0.12 (-0.42, 0.17)	-0.12 (-0.42, 0.17)	0.00 (-0.24, 0.26)	September
γ_{10}	0.15 (-0.12, 0.46)	0.11 (-0.19, 0.39)	0.10 (-0.18, 0.41)	0.10 (-0.18, 0.41)	0.27 (0.00, 0.51)	October
γ_{11}	-0.37 (-0.65, -0.10)	-0.41 (-0.70, -0.15)	-0.41 (-0.66, -0.12)	-0.41 (-0.68, -0.13)	-0.38 (-0.63, -0.15)	November
κ		1.61 (0.49, 2.76)	1.67 (0.56, 2.81)	1.89 (0.70, 2.97)	4.20 (2.44, 5.86)	Customer Group
ζ			-0.10 (-0.21, 0.03)	0.32 (0.14, 0.49)	0.60 (0.36, 0.86)	Post LP
ω				-0.83 (-1.06, -0.58)	-0.91 (-1.12, -0.71)	Group x LP
δ	6.95 (5.10, 9.22)	6.97 (5.12, 9.19)	6.97 (5.15, 9.31)	6.99 (5.16, 9.28)	2.80 (1.78, 4.08)	regional SD
ρ	0.15 (0.05, 0.19)	0.15 (0.05, 0.19)	0.16 (0.05, 0.19)	0.16 (0.04, 0.19)	-0.06 (-0.32, 0.18)	reg. spillover

Table 9: Out-Of-Network Partial Model Results