

Leveraging Loyalty Programs Using Competitor Based Targeting

Wayne Taylor

Brett Hollenbeck

Cox School of Business

Anderson School of Management

Southern Methodist University *

UCLA †

June 24, 2020

Abstract

Loyalty programs are widely used by firms but not well understood. These programs provide discounts and perks to loyal customers and are costly to administer, but produce uncertain changes in spending patterns. We use a large and detailed dataset on customer shopping behavior at one of the largest U.S. retailers before and after joining a loyalty program to evaluate how behavior changes. We combine this with detailed spatial data on customer and store locations, including the locations of competing firms. We find significant changes in behavior associated with joining the LP with a large amount of heterogeneity across customers. We find that location relative to competitors is the factor most strongly associated with increases in spending following joining the LP, suggesting that the LP's quantity discounts work primarily through business stealing and not through other demand expansion. We next estimate a set of predictive models to test what variables best predict how spending will change after joining the LP. We use high-dimensional data on spatial relationships between customers, the focal firm's stores, and competing stores as well as customers' historical spending patterns. These models are used to test whether past sales data reflecting customer's vertical value to the firm or spatial data reflecting customer's horizontal vulnerability are more predictive of post-LP spending increases. We show how LASSO regularization estimated on complex spatial relationships are more predictive than are models using past sales data or other spatial models including gravity models. Finally, we show how firms can use this type of model to leverage customer and competitor location data to substantially increase the performance of their LP through spatially driven segmentation rules.

Keywords: Loyalty programs, predictive analytics, spatial models, retail competition, LASSO estimation

*wjtaylor@mail.smu.edu

†brett.hollenbeck@anderson.ucla.edu

1 Introduction

Loyalty programs are now prevalent in many industries. The loyalty programs as marketers know them today originated in the airlines in the 1980s and have since permeated industries such as hotels, casinos, retailers, grocery stores, and restaurants, among others. The pervasiveness of loyalty programs is partly due to their flexibility in how they can be structured (e.g., earning rates and rewards) and managed for strategic decisions (e.g., targeting to increase customer engagement).

In spite of the popularity of loyalty programs with firms, their effectiveness at increasing profits has long been subject to debate. This debate has centered on the costs of giving discounts and perks to the most loyal customers, as well as the costs of administering the program itself, and whether these costs are justified by increases in spending by those customers. Loyalty programs have the potential to increase profits by increasing switching costs for existing customers, stealing business from rivals, or through second degree price discrimination. They may also indirectly increase profits by increasing customers psychological perceptions of the firm, by generating customer data that can be used for targeted promotions or CRM, or by exploiting agency issues such as flights booked by business travelers and paid for by their employers (Dreze and Nunes (2008), Roehm et al. (2002), Verhoef (2003), and Shugan (2005)).¹

Empirical studies of whether loyalty programs actually do increase profits have found mixed results. Verhoef (2003) finds that the effects are positive but very small, DeWulf et al. (2001) finds no support for positive effects of direct mail, Shugan (2005) finds that firms gain short term revenue at the expense of longer term reward payments, and Hartmann and Viard (2008) found no evidence that loyalty programs create switching costs.

In this paper, we use a large and detailed dataset on customer shopping behavior at one of the largest U.S. retailers before and after joining a loyalty program. This loyalty program takes a common form, tiered discounts after a certain level of spending, essentially operating as 2nd degree price discrimination or one particular form of behavior-based pricing (BBP). We therefore analyze the program through the lens of theoretical work in marketing on price discrimination and BBP, most notably Shin and Sudhir (2010). We focus on these research questions: when loyalty programs take this common form, how do consumer spending patterns change after they join? What factors predict who will increase spending after joining? How can firms use insight from readily available

¹For a more complete review, please see Bijmolt et al. (2011), Liu and Yang (2009), and McCall and Voorhees (2010)

spatial data on competitive structure to improve customer segmentation? Finally, we show how studying this contributes to the understanding of loyalty programs and price discrimination more broadly.

The practice of BBP in the form of rewarding one's best customers with discounted prices is controversial. Shin and Sudhir (2010) showed how it can be profitable to use this type of pricing strategy if two conditions are met. First, there must exist substantial heterogeneity in customer value. Second, customer preferences must be stochastic. To consider these conditions empirically in our retail spending data, we merge detailed customer location data with the data on spending patterns. We also collect data on the retailer's competitor's locations. We can therefore exploit variation across markets with different levels of competition and types of competitors at a granular level. We then observe how changes in spending associated with joining the LP varies at the customer level for customers facing different local competitive structures. We can observe, for instance, if customers in more isolated markets change their spending or if they merely start receiving discounts on their purchases and how this differs in highly competitive markets.

A crucial aspect of our data is that it results from merging credit card spending data with customer data. We therefore observe spending both before and after a customer joins the LP. Most prior work on LP's suffers from only observing spending data on customers in the program.² There is a clear selection effect of high spending customers into the LP that significantly complicates all efforts to measure the effect of LP's on behavior. Even with pre and post-join data there are still selection effects that would complicate an effort to isolate a purely causal effect. In particular, even when measuring the change in spending before and after within an individual customer, the decision to join the LP could coincide with a change in planned spending. Instead of trying to isolate a pure causal effect, therefore, we perform both descriptive analysis and predictive modeling to better understand the relationship between spending behavior, LPs, and complex spatial data and show how firms can use this to understand LP mechanics and improve segmentation of customers who join the LP.

We first show a broad set of descriptive results on customer behavior that take advantage of the unique nature of our data. We follow that with a predictive modeling exercise that measures how different data types can best predict which LP joiners are likely to increase spending, and therefore how the firm might use segmentation to identify profitable joiners and increase the profits of the

²See Meyer-Waarden and Benavent (2009) for an exception.

LP. In both these sets of results we use a common insight, that an individual’s spatial relationship with competitor stores is the main determinant of when spending increases and that the nature of this relationship can be complex but exploitable with freely available data. To our knowledge, this is the first paper to explicitly link the competitive structure of a market with the performance of a loyalty program.

Our descriptive analysis of the spending data finds the following results. First, behavior often changes significantly when a customer joins the loyalty program. As already noted, this change may or may not be caused by the LP itself, the causation could run in the opposite direction for some customers. However, across all customers the aggregate change in profits associated with customers joining the LP is small or even negative. Second, there is wide variation across customers in profitability, as many low-spenders greatly increase their spending when they join but many of retailer’s best customers simply gain discounts on their purchases. Certain types of LP joiners seem especially important, in particular we identify segments of customers who seem to consolidate their purchases at the focal retailer by increasing trip frequency and customers who maintain the same amount of products purchased but upgrade to higher priced products.³ Notably, the strongest predictor of whether a customer’s joining the LP will be associated with higher or lower profits is whether they are located near a competitor store or live in an isolated market with only the focal store present. In fact, all the profit gains associated with the LP are from customers in competitive markets. While this result is intuitive it has not been shown previously in a LP setting. This motivates the use of spatial data that accounts for the local competitive structure around each customer as the basis of a segmentation strategy to increase the LP’s effectiveness.

While simple descriptive results suggest spatial competitive structure is the crucial determinant of LP effectiveness, the full relationship is likely to be highly complex. To capture the complex spatial relationship between the customer, focal firm, and competition, we develop a new methodology for the treatment of spatial data in predicting customer outcomes. To do so we estimate the relationship between the change in spending (conditional a customer on joining the LP) and a very large number of variables on the complex spatial relationships with competitors at the individual level then select the metrics most predictive of the customer behavior of interest. This estimation is essentially a prediction problem with many potential predictive relationships and is therefore

³We emphasize that we do not observe spending at competitor stores, as is typical with firm-level data. We therefore do not literally observe consolidation and business stealing effects but infer them from how changes in behavior within customers vary, especially across customers with different spatial characteristics.

well suited to shrinkage type estimators like LASSO regularization. We find this method performs dramatically better at predictive tasks than standard treatments of spatial data. We also compare the results to specifications based on the theoretical insights from classic gravity models of Reilly (1931) and Huff (1964). We find the gravity models perform well compared to other simple models but our method does substantially better.

Next, we compare the performance of models using spatial data and past sales data as inputs. Following the insights of Shin and Sudhir (2010), we broadly define these data on spending patterns as capturing *vertical quality* or the value of a customer in terms of their overall demand, and define the data on local competitive structures as capturing *horizontal quality* or how likely it is to shift a customer's spending from a competitor's store to the focal store. This estimation serves two purposes. First, it validates the descriptive result that horizontal quality (i.e., spatial competitive structure) is a stronger predictor of LP effectiveness than vertical quality (i.e., past spending patterns or other RFM variables).

Second, we then use the output to show how firms can take advantage of this insight and leverage customer and competitor location data to increase the performance of their LP through spatially driven segmentation rules. A potential use for this is for firms to develop strategic targeting strategies to avoid promoting the LP to those whose behavior is unlikely to change and would be using the LP for discounts alone, based on their locations relative to the competition. Because location and travel costs form an important part of preferences over retailers, this can be thought of as a strategy for targeting price discounts on observable preference heterogeneity. Our results show that, consistent with the spirit of recent work in customer valuation (e.g., Ascarza (2018)), firms should leverage sources of observed heterogeneity and focus on the customers who are spatially vulnerable with easy access to competitors. We also show that our method for capturing the complex nature of spatial data substantially outperforms more standard treatments of this same data. We note that because the LP works through quantity discounts (second degree price discrimination) there is a complementary effect between spatial segmentation and this type of price discrimination. Because consumers self-select into the use of discounts based on desired spending level, segmenting them using vertical data in isolated markets can still be effective but in competitive markets horizontal data is substantially more effective.

In addition to contributing to the understanding of loyalty programs, we contribute to recent work that has explored the benefits of geotargeting, where promotion or other marketing activity

is a function of a customer’s real-time location using mobile data (Luo et al. (2014), Chen et al. (2016)). Fong et al. (2015) and Dube et al. (2017) have also considered “geoconquesting”, where the marketing activity focuses on instances when the customer is located near a competitor as opposed to near the focal firm. This literature is limited, but potential gains from geoconquesting are especially likely when the marketing activity is based in part on business stealing, as it is for loyalty programs.

While the previous literature therefore incorporates select spatial components, mostly it does not fully account for the complex customer-store-competitor spatial relationship. In prior work, competitor information is often integrated into models through simple customer-store or store-competitor distance metrics, thereby eliminating the possibility of complex spatial analyses. Beyond complex spatial analysis alone, the interplay between the competitive structure and loyalty program effectiveness has not yet been explored.

We also contribute to the study of competitive promotions and competitive price discrimination. Price discrimination strategies such as loyalty rewards should never lower profits by a monopolist but in oligopoly settings this is no longer true as firms may face a prisoner’s dilemma (Shaffer and Zhang (1995)). Chen et al. (2001) shows that when individual targeting is possible but imperfect, it can soften price competition among competing firms, but as targeting precision increases the prisoner’s dilemma reasserts itself. Ultimately then, it is an empirical question whether and when targeted price discrimination can increase profits. Previous work has shown that in practice the benefits of using targeted pricing can be quite high (Rossi et al. (1996), Besanko et al. (2003)). Li et al. (2018) specifically consider competitive price discrimination across markets and show that the profitability of tailoring prices to local markets depends on both the local market structure and competitive intensity across markets. Our analysis also considers local market structure but with price discounts in the form of a loyalty program. We show that in the context of quantity discounts, segmenting customers based on horizontal versus vertical characteristics are differentially profitable depending on the degree of local competition.

The remainder of the paper proceeds as follows. Section 2 describes the data on retail sales histories and competitor and customer locations. Section 3 provides model free evidence that the spatial competitive structure surrounding a customer is the key determinant of LP effectiveness. Section 4 describes a LASSO estimation of the complex interaction between competitive structure and LP effectiveness and then uses this result to evaluate the relative predictive power of horizontal

and vertical data and then use this to form a segmentation strategy. Section 5 discusses additional managerial implications and avenues for future research.

2 Data

In this section we describe the customer transaction data, competitive location data, and provide an overview of the spatial metrics used to characterize the competitive structure.

2.1 POS Transaction Data

The transaction data comes from a Fortune 500 specialty retailer. The retailer specializes in a range of product categories including lumber, electrical, and paint, among others. This point-of-sale (POS) data from the retailer is highly detailed: we observe the full basket of purchases at the SKU level from a random sample of 10,029 customers between March 2012 and March 2014 across a variety of product categories, regardless of transaction method (e.g., cash or credit card). These sum to over 2.4 million SKU level purchases across 897,819 store trips. On average, each customer has about 90 trips across nearly five different store locations through the two year observation period, spending about \$110 per store visit, \$684 per month (at the firm level, potentially across multiple store locations), and travels about 9.2 kilometers (5.7 miles). The large levels of monthly and per-trip spending are consistent with the nature of the categories sold at this retailer, which are relatively high-priced categories. In addition, many of the retailer's customers are professional customers whose purchases are related to their jobs. This includes many of the members of the loyalty program.⁴

⁴This is typical of loyalty programs in many settings, for instance an outsize share of members of airline and hotel loyalty programs are professionals who travel frequently for work.

Customers	10,029	
Date range	March 2012 to 2014	
	Mean	SD
Trips/customer	90	77
Spend/trip (net, after discounts)	\$110	\$102
Spend/month (net, after discounts)	\$684	\$714
Store locations visited	4.7	3.2
Distance by household (km)	9.2	6.2

Table 1: Summary Statistics

Crucially for this analysis, the firm providing the data uses a variety of methods to associate transactions with an individual customer (for example, matching addresses across credit cards used in multiple transactions). This allows the firm to observe all transaction activity regardless of loyalty program enrollment. Specifically, it allows us to observe changes in customer behavior at the firm upon joining the LP.⁵ We also observe all marketing activity for these customers through the firm’s email campaigns: over 900,000 emails were sent to about 39% of the customers and 11% of the customers received promotions specifically encouraging enrollment in the loyalty program.

Importantly, the data also contains the customer’s zip code and the latitude/longitude of each of the firm’s store locations, which provides us with a complete picture of all customer interactions with the firm across different store locations as well as the specific products purchased at each store location over time.

We further take advantage of a unique feature of this loyalty program in that the program discounts are earned and applied on only a single large category of the firm’s goods. We label this a *limited* loyalty program to emphasize the distinct structure. As requested by the firm providing the data, we cannot disclose the category of goods for which the discount applies. Instead, we refer to it as the *focal* category of interest.

The structure of the limited loyalty program is that of a second degree price discrimination or quantity discounts. Customers who reach spending thresholds in the focal category receive discounts

⁵For the average joiner, we observe about 7 months of transaction activity prior to enrollment and about 8.5 months of transactions after joining.

on future purchases within that category. The loyalty program consists of three such thresholds, with increasing discounts upon reaching each threshold. Customers who reach a given threshold retain their status until the end of the calendar year, at which point they begin the earning process again. The loyalty program is offered at the firm level, not the store level, so the thresholds can be attained from purchases across multiple locations.⁶

This unique feature of offering a discount only within a focal category provides an additional layer with which to study the interaction between competitive structure and LP effectiveness. This is because the firm competes with both generalists, who also sell across a large variety of categories, and specialists, who only sell the products in the specific category the LP applies. It also allows us to study whether and when the loyalty program purchases are associated with spillovers into other category purchases.

2.2 Competitor Data

We augment the focal firm's POS transaction data with location information (i.e., latitude and longitude) from four competitors. We recognize there may be other competitors in a given market however the firm providing the data explicitly stated these four competitors as their primary concern and others as inconsequential.⁷

The four competitors (see Table 2) vary in both size and product breadth. Competitor #1 is a big box store with a wide product variety (BB). The remaining three competitors are small-box stores. Competitor #2 and competitor #3 offer a wide assortment of products (SB1 and SB2), whereas competitor #4 specializes in the focal product category for which the limited loyalty program applies (SS). None of these competitors offer a loyalty program similar in structure to that of the focal firm.

⁶This loyalty program structure is consistent with that of a "customer-tier" LP, as opposed to a "frequency-based" LP, following Blattberg (2008).

⁷This industry is highly concentrated, with the focal firm and the primary big-box competitor capturing about 85% of market share, according to a 2019 industry report by IBISWorld.

Competitor	Reference	Footprint	Product Breadth
#1	BB	Big-Box	Wide
#2	SB1	Small-Box	Wide
#3	SB2	Small-Box	Wide
#4	SS	Small-Box	Specialized

Table 2: Competitor Types

The variety in the size and product breadth allows us to compare how the impact of the competitive structure might be competitor or type specific. More importantly, it allows us to gain insight into the extent of category-level versus store-level business stealing.

The latitude and longitude locations of each competitor are collected from the Google Maps API. We first pulled all competitors within a 50 kilometer radius of each of the focal stores, and then pulled all competitors within 50 kilometers of each customer located within a 50 kilometer mile radius of each focal store.⁸ This expanded footprint ensures that our definition of the competitive structure is customer centric rather than limited to the perspective of the focal store.

2.3 Quantifying the Competitive Structure

Our analysis measures and highlights the impact of the complex spatial relationship between the customer, the focal store, and the competitors on customer behavior and ultimately firm profits. Quantifying this relationship in such a way to accurately reflect the tradeoffs that an individual likely encounters when deciding which store to visit requires several complex considerations.

A common approach in quantifying competitive structure is to simply use the distance between a focal store and the customer along with the distance between the competitor and the focal store (or more commonly still, an indicator variable if they are both within, say, a 5 kilometer radius of the focal store). The drawback of this approach is that it does not jointly consider the customer-store-competitor location, resulting in a potential homogenization of very distinct competitive structures.

This limitation is illustrated in Figure 1: three competitors, C_1 , C_2 , and C_3 , and a customer, I ,

⁸We selected 50 kilometers because this was the maximum radius allowed by the Google Maps API at the time of pulling this information and exceeds the average distance traveled by household by a factor of about 7.

are positioned near the focal store S . Both C_1 and C_2 are nearly the same distance to the store, but their respective relationships to the customer and the focal store are considerably different. The customer has to pass by the focal store in order to visit the first competitor, C_1 , which suggests some spatial advantage for the focal store. However, the second competitor, C_2 is positioned right next to the customer, acting as a convenient alternative to the focal store. In a different case, C_2 and C_3 are both equidistant to the customer but one is between the customer and the store and the other is in the opposite direction. One goal of this analysis is to incorporate these complex spatial structures into the manager decision process, a strategy that has heretofore been ignored in the analyses of loyalty program effectiveness.

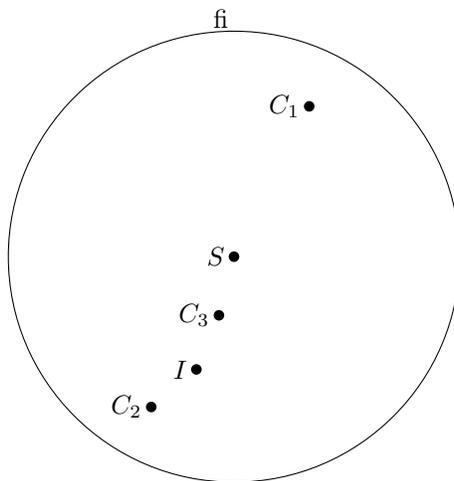


Figure 1: Limitation of Radii Approach

We recognize that it is unreasonable to expect a single metric to capture the complex nature of the spatial relationship between the customer, focal store, and competitor in its entirety. Instead, we approach the problem by proposing a large number of metrics, each of which captures at least some of the complex spatial relationship on its own, and then in our empirical analysis uncover which metrics or which interactions between metrics best predict customer behavior. By honing in on the right combination of distance metrics we can determine which *specific* features of the spatial relationship influence customer behavior the most.

We therefore consider standard distance metrics between each customer and competitor with the the focal firm in addition to the following:

- How much closer is the focal store to the customer, relative to the competitor?
- How *sparse* is the focal store and competitor, relative to the customer?
- Are the focal store and competitor *in the same direction* from the customer, and to what extent?

We also allow for interactions between population density and distance metrics to incorporate the difference in transportation costs between urban and rural areas. For brevity, the complete description of the distance metrics considered are contained in the appendix.

3 Spatial Competitive Structure and LP Performance

In this section we take advantage of the unique nature of our data and provide descriptive analysis of how customer behavior interacts with joining a loyalty program. Because we observe sales patterns before and after the customer joins the LP we can measure if this behavior changes and if so, how it changes. These results relate to the extant debates on how and whether LPs are effective at increasing profits.

We then provide model free evidence that the spatial competitive structure are key determinants of LP performance. Our analysis focuses on two metrics of LP effectiveness: the probability that a customer joins the program, and the change in average monthly spend for customers who decide to join. Here and throughout the paper, we use the change in spending net of the LP discounts. For customers who join the LP but do not change shopping patterns, this change is negative by default because of the discounts, and all else equal a positive change for this metric can therefore be taken as strictly beneficial for the retailer. In this section we look at how these metrics relate to the the distance between the customer and the focal store and the four competitor types.

One challenge in analyzing LP effectiveness is that the customer’s decision to join the loyalty program may be related to unobserved heterogeneity. While we observe purchases both before and after each customer joins the LP, which allows us to condition on individual-level time-invariant factors, it is still true that if customers join due to anticipated changes in their level of spending, the observed change in behavior can not all be attributed to the loyalty program. We therefore do not treat all changes in spending or other behavior associated with joining the LP as having been caused by the LP. Instead we provide descriptive results on how behavior changes and note when

those changes in behavior that coincide with joining the LP may be of direct interest even without a purely causal interpretation.

In addition, we treat the location of each customer as essentially exogenous prior to joining the LP, in which case the relative difference in outcomes across customers with different spatial characteristics can provide valid and useful comparisons. Finally, we replicate our results using only spending that does not qualify for the LP discounts and is thus less likely to result from selection on unobservables.

Descriptive Analysis:

Table 3 first presents the monthly sales at the firm level, split into qualified and non-qualified spend, for all customers and those that enroll in the loyalty program. In terms of overall spending and spending in the LP category, joiners and non-joiners are actually quite similar. As expected, customers who join the loyalty on average have slightly higher monthly spending on products that qualify for the LP discounts but not higher overall spending, and the differences are modest.

	All Customers	LP Joiners
Sales per Month	\$684	\$662
Qualified Sales per Month	\$31	\$47
Non-Qualified Sales per Month	\$652	\$615

Table 3: Overall Monthly Spending

We next show the difference between purchase behavior before and after a customer joins the LP. Table 4 presents the average change in sales, trip frequency, and basket level metrics for customers who join the loyalty program.⁹ On average, customers tend to increase monthly spending by about \$47 upon joining the loyalty program, with the majority of this change attributed to spend in categories that qualify for LP discounts. This is not surprising, but we also see a positive change in non-qualified spend, which could suggesting that the increase in spending associated with joining the LP spills over into other categories.

Joiners show a small increase in trip frequency, almost no change in number of categories shopped

⁹For all tables presented in the descriptive analysis, we exclude households that fall within the upper or lower .5% of monthly sales prior to joining the loyalty program to prevent the undue influence of outliers for the summary statistics.

in, and actually slightly decrease the number of unique products purchased. We also note there is substantial heterogeneity across customers.

In the second column, we limit attention to customers whose change in monthly spend was positive. Interestingly, for these customers, the vast majority of the increase is attributed to spend outside of the focal category and does not qualify for the LP discounts. Much of this is driven by a large increase in trip frequency, with more modest changes in basket diversity (distinct categories and SKUs), and basket size. We do not suggest that the LP alone caused all of these changes, but note that within the base of customers that do join the loyalty program, there is substantial heterogeneity in post-join behavior.

	All Joiners	> 0 Δ Sales
Δ Monthly Sales	\$47	\$821
Δ Monthly Qualified Sales	\$37	\$133
Δ Non-Qualified Sales	\$10	\$688
Δ Trip Frequency	0.07	3.95
Δ Unique Categories	-0.01	0.14
Δ Unique SKUs	-0.30	0.41
Δ # Items	-2.44	1.22

Table 4: Change in Behavior

We next consider how a rich but often overlooked source of variation across customers, the competitive structure, can help in attributing the variation in behavior for customers who join the loyalty program. To start, we first label customers as "isolated" if they are located more than eight kilometers (five miles) away from any of the focal store's competitors and "competitive" otherwise, which indicates there is a competitor relatively close to the customer. We selected this cutoff based on average distance driven by joining customers to surrounding focal stores. Table 5 summarizes a few metrics of interest across these types of customers, namely the probability of joining and the average changes in monthly qualified and non-qualified spend. We also include the average customer sales per month to alleviate concerns that our region labels are associated with substantially different customer types, at least in terms baseline spending patterns.

Approximately one quarter of the customers are located in "competitive" regions and are relatively close to one of the four competitors of interest, as identified by the firm providing the data. Relative to customers in more isolated regions, these customers have essentially the same probability of joining the loyalty program but exhibit a substantially larger change in monthly spend. As expected, the change in qualified spend is positive across both groups. However, customers located near the competition also exhibit an equally high change in non-qualified spend. For these customers, joining the loyalty program may present an opportunity to consolidate purchase activity at the focal store. For customers in relatively isolated areas, they are likely devoting most of their budget to the focal store prior to joining the LP and enrollment is unlikely to generate additional spend from competition.

	Competitive	Isolated
Share of Customers	25%	75%
Pr(Join)	4.4%	4.5%
Δ Sales Join	\$109	\$28
Δ Qualified Sales Join	\$52	\$33
Δ Non-Qualified Sales Join	\$57	-\$5
Sales per Month	\$707	\$647

Table 5: Effect of Competition

In Table 6 we further split the change in monthly sales based on the median qualified spending amounts prior to joining the loyalty program. In both cases, the larger change in spend comes from below-median spenders pre-LP. This is somewhat counterintuitive since the program targets discounts at high spenders. But most notably, this change is more than twice as large in competitive regions, suggesting that customers with relatively low spend prior to joining are likely splitting their purchase behavior across competitors. For customers with relatively high pre-join spend, the change is not as large, but still higher in the competitive regions. High spenders in isolated markets actually decrease spending somewhat, once discounts are taken into account. This suggests that the degree of competition near a customer can have a substantial impact on customer engagement with the loyalty program.

Pre LP Qualified Spend:	Competitive		Isolated	
	Low	High	Low	High
Δ Sales Join	\$197	\$33	\$82	-\$23

Table 6: Interaction Between Competition and Pre-Join Spending

To further analyze what is driving these findings we break down the results by competitor type. Figure 2 shows the average change in monthly spending based on the relative isolation of the customer.¹⁰ The findings are relatively consistent with the aggregate results, with the exception of a notably larger increase from customers near the big-box generalist, suggesting this type of competitor presents the most opportunities from which to steal business from, and a notably smaller increase from the small-box specialist, who presents the fewest opportunities from which to steal business from.

To better understand this mechanism, recall that the LP only applies to a specific category, not all products, and that one of the competitors (SS) specializes in selling only that focal category. We thus note the portion of change in spend that occurs in the LP category of interest. Figure 2 highlights this by distinguishing the change in spend as either qualifying for the LP discount (within the focal category) or not, where the combined change is noted by a black square. As expected, a large portion of the change in spend is driven by purchases that qualify for LP discounts. However, there is substantial positive changes outside of the focal category. If there was no spillover effects of the loyalty program, we would expect zero change in purchases that will not impact the LP rewards. Upon joining the LP, customers may decide to consolidate purchase behavior with one store rather than cherry picking rewards from the focal store (in the qualifying category) and continuing with the same purchase patterns at other stores in the non-qualifying categories. For customers located in relatively isolated regions, we see that the gains from focal category purchases tend to be offset by declines in other purchase categories, regardless of competitor type.

Our descriptive evidence suggests that a competitor’s location can have a strong influence on how spending changes when customers join the LP. The analysis motivates the need for a more sophisticated model that captures these complex spatial effects that we present in a later section.

¹⁰Note that not all customers may have access to all competitors within a market. Because of these, some customers will not be associated with certain competitors.

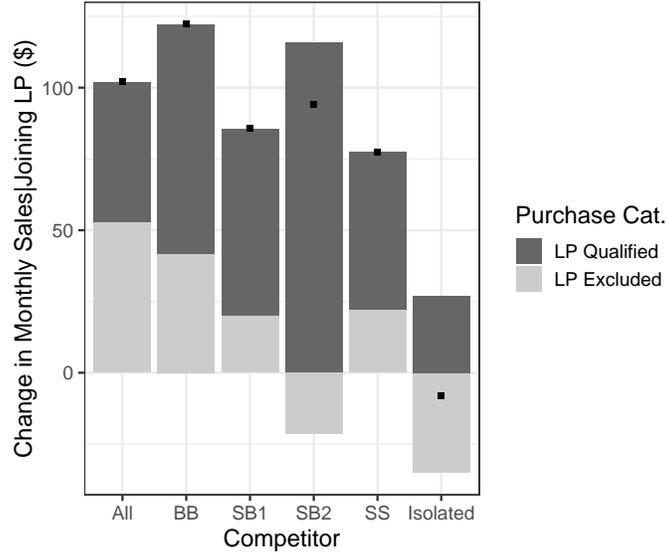


Figure 2: Change in Monthly Spend by Competitor Type and Purchase Category

Change in Spend Descriptive Analysis:

Before specifying our models of join and spend activity, we first provide a deeper analysis of the change in spend based on pre-join customers characteristics. Ideally, the focal firm would be able to use this information to determine which customers are likely to exhibit the largest change in spend after joining the LP and then segment these customers accordingly.

First, Table 7 displays a variety of change metrics based on a median split of five metrics of interest: qualified spending levels prior to joining the LP, number of stores visited, and the three distance metrics of interest: customer-store, customer-competition, and store-competition. We are interested, for instance, in how the firm might predict ex ante what customer characteristics are associated with post-LP increases in spending.

Overall, we see the change in sales is higher for customers 1) with relatively low qualified spend prior to joining the LP, 2) who visit more than two stores pre-join, 3) who are relatively far from the focal store, 4) who are relatively close to the competition, and 5) who shop at stores relatively close to a competitor.

Trip frequency increases substantially more for customers with relatively low qualified sales prior to joining the LP, and for customers who are relatively further from the focal store.

The basket composition becomes more diverse (change in unique categories and SKUs) for

customers with relatively low qualified sales (pre LP), who visit relatively fewer focal stores, and who are relatively close to the competition. Together, these findings continue to suggest that business stealing is more likely for customers with room to grow with the focal firm but also have access to competitors nearby.

Finally, we redo this analysis but focusing on customers who increase spending after joining the loyalty program. The patterns remain relatively consistent but with the magnitudes more pronounced. As in the section with all joiners, the change in spending (in all categories: overall, qualified, and non-qualified) is substantially higher in situations where the customer is close to focal store and the focal store is close to the competition. A similar pattern holds for changes in trip frequency. Interestingly, when the customer is relatively close to the focal store, the changes are lower than those who are relatively far from the focal store.

These initial results begin to illustrate how the competitive structure, defined as the joint relationship between the customer, focal store, and competitor locations, are strongly related with a customer's engagement with the LP. More importantly, we note that the difference between the high and low categories is substantially smaller when the median split is based on pre-join qualified sales or number of stores visits, suggesting that in some cases spatial characteristics appear to be more informative in identifying change in spend versus traditional behavior based metrics.

All Joiners										
	Pre-LP Qual. Sales		# Stores		Cust. Store		Cust. Comp.		Store Comp.	
Median:	\$7.57		2		8.44 km		14.0 km		13.6 km	
	low	high	low	high	low	high	low	high	low	high
Δ Sales	\$109	-\$14	\$28	\$62	\$12	\$82	\$127	-\$32	\$95	\$0
Δ Qual. Sales	\$53	\$22	\$41	\$35	\$24	\$51	\$64	\$12	\$46	\$29
Δ Non-Qual. Sales	\$56	-\$36	-\$13	\$27	-\$12	\$32	\$63	-\$43	\$49	-\$29
Δ Trip Freq.	0.37	-0.23	0.11	0.04	-0.13	0.26	0.13	0.01	0.02	0.11
Δ Unique Cat.	0.05	-0.06	0.00	-0.01	0.05	-0.06	0.06	-0.08	-0.01	-0.01
Δ Unique SKU	-0.02	-0.57	-0.14	-0.42	-0.19	-0.41	0.09	-0.68	-0.27	-0.32
Δ # Items	-3.53	-1.36	-1.55	-3.11	-0.43	-4.45	-1.58	-3.30	-0.40	-4.47

Δ Sales Join > 0										
	Pre-LP Qual. Sales		# Stores		Cust. Store		Cust. Comp.		Store Comp.	
Median:	\$7.75		2		8.24 km		14.4 km		13.7 km	
	low	high	low	high	low	high	low	high	low	high
Δ Sales	\$911	\$731	\$813	\$826	\$734	\$907	\$960	\$683	\$1,073	\$571
Δ Qual. Sales	\$100	\$166	\$116	\$145	\$65	\$200	\$147	\$120	\$157	\$109
Δ Non-Qual. Sales	\$812	\$565	\$697	\$681	\$668	\$707	\$813	\$563	\$915	\$462
Δ Trip Freq.	3.72	4.18	3.61	4.19	3.81	4.08	4.44	3.46	4.62	3.28
Δ Unique Cat.	0.19	0.08	0.29	0.03	0.21	0.06	0.18	0.09	0.15	0.12
Δ Unique SKU	0.71	0.12	0.92	0.05	0.76	0.07	0.49	0.33	0.46	0.37
Δ # Items	1.84	0.61	4.48	-1.07	5.33	-2.85	4.53	-2.05	-1.57	3.98

Table 7: Change in Behavior Based on Pre-Join Characteristics

Patterns Among Profitable Joiners: Within the customers who increase spending upon joining the LP, we investigate other changes in behavior to determine if there are patterns as to what is driving the positive change in spend. We focus on using the non-spatial variables to see if we might identify segments to provide insight into specifically how engagement with the LP changes upon joining.

In Table 8 we provide a cross-tabulation of selected measures of shopping behavior for these customers. The cells indicate the proportion of customers that fall into each quadrant. We construct the quadrants based on whether there is a substantial increase (which we define as at least a 15% increase in the metric) or if the behavior is flat or even decreasing. For example, of the customers with a positive change in spend, 13.9% of them showed no appreciable increase in trip frequency paired with a significant increase basket size. Within each cell we also note the average increase in spending upon joining to highlight the heterogeneity across customers and relative value of each segment.

As expected, most consumers show an increase in the given variables, since we are selecting only customers who increase spending overall. Of greater interest are the customers who land in the off-diagonal groups - that is, customers who exhibit a positive change in one metric at the expense of a negative change in another. In the complementary Table 9 we provide additional information about the customers within these off-diagonal cells, namely their pre-join behavior and simple spatial metrics.

The first cross tabulation shows that more than one third of those who increase spending do so via increased trip frequency but with no increase in per-trip basket size. This behavior is consistent with consolidating purchases from multiple competing stores to only the focal store after joining the LP and we refer to this group as *consolidators*. The average change in spend of this group at \$749 is nearly twice that of the complementary group, who increase basket size but not trip frequency. This is despite the fact that in the complementary table these customers (in the "lower left" cell) have similar pre-join spending levels as those in the opposite cell ("top right"). The spatial information provides some indication of why these customers may have proven to be so valuable. The more valuable group is actually further away from the focal stores, at a distance of more than 10km versus about 8km, and on average but essentially the same distance to the competition. That is, it appears that after joining the customers shift trips that occur elsewhere to the focal store.

The middle cross-tabulation compares changes in number of items purchased and the average price per item. The largest group increases spending by buying more items but without a substantial increase in average price. But about 31% of customers show no increase in basket size but a substantial increase in the price paid per item, suggesting that these customers may be responding to the LP discounts either by upgrading to more expensive items after joining the LP or shifting higher purchases away from competitor stores. We refer to this group as *upgraders*. These customers

also increase their total spending by more than twice as much as those who simply buy more items.

We see a similar pattern as the first section: the customers in the more valuable off-diagonal cell (in terms of change in spend after joining) are relatively further away from the focal store. Again, customers that join but are close to the focal store are likely devoting most of their budget to the focal store as is, so joining is unlikely to have a pronounced impact on spend (along with other metrics) unless there is a competitor from which to take business from. The upgraders are also higher spenders pre-join, suggesting they are more likely to receive the LP discounts.

Finally, the bottom cross-tabulation shows again that close to a third of customers have no increase in qualified sales but a substantial increase in non-qualified sales. In terms of change in spend upon joining, those in the top right cell are more than five times more valuable than those in the lower left cell, even though the pre-join sales per month are actually lower. This behavior is again consistent with being a consolidator, shifting non-qualified spending in response to joining the LP. As with before, the spatial metrics provide additional insight into why we might see a marked increase in non-qualified sales: they are more than 2km closer to a competitor, on average, and the focal store is about 1km closer to a competitor, on average.

Our deep analysis of the joining customers continues to provide insight on how behavior changes upon joining the LP and illustrate that spatial metrics provide incremental value in identifying valuable customer groups. In these tables, many of the behavior-based metrics are nearly identical, so any effort to pinpoint the valuable customers based on behavior based metrics alone may be difficult. However, we have provided preliminary evidence that spatial metrics appear to indicate which customers are valuable, both in aggregate analyses as shown earlier and in very detailed analyses shown in this section.

In the next section, we integrate spatial metrics into predictive models of customer behavior. In addition to the relatively simple spatial metrics already presented in the descriptive analysis, we consider a variety of other subtle spatial metrics which may also influence probabilities of joining the loyalty program and change in spend behavior.

		Δ Basket Size (# Items)	
		Flat/Decrease	Sig. Increase
Δ Trip Frequency	Flat/Decrease	14.4%	13.9%
		\$814	\$312
	Sig. Increase	35.8%	35.8%
		\$749	\$1,094
		Δ Price/Item	
		Flat/Decrease	Sig. Increase
Δ Basket Size (# Items)	Flat/Decrease	19.4%	30.8%
		\$233	\$1,103
	Sig. Increase	34.3%	15.4%
		\$591	\$1,506
		Δ Non-Qualified Sales	
		Flat/Decrease	Sig. Increase
Δ Qualified Sales	Flat/Decrease	4.0%	27.9%
		\$59	\$636
	Sig. Increase	12.4%	55.7%
		\$115	\$1,125

Table 8: Positive Δ Sales|Join Proportions and Δ Sales

Pre LP Metrics	1: Trip Freq v. # Items		2: # Items v. Price/Item		3: Qual v. Non-Qual Sales	
	Top Right	Lower Left	Top Right	Lower Left	Top Right	Lower Left
Sales per Month	\$506	\$464	\$604	\$414	\$448	\$581
Non-Qual Sales/Month	\$471	\$442	\$577	\$388	\$410	\$543
Qual Sales/Month	\$35	\$22	\$27	\$26	\$38	\$38
Trips/Month	6.66	4.36	5.33	4.97	5.32	6.07
Unique Categories	1.73	2.06	2.08	1.86	2.08	1.90
Unique SKUs	3.12	4.10	4.17	3.49	4.02	3.80
Basket Size	6.42	18.32	18.50	7.32	9.48	11.23
Price/Item	\$17.86	\$10.72	\$9.88	\$19.82	\$11.34	\$13.36
Stores Visited	1.96	2.29	2.21	2.06	2.04	2.44
Store-Customer km	8.24	10.06	9.46	8.40	9.27	9.98
Cust.-Competitor km	15.71	16.13	15.79	13.75	14.34	16.44
Store-Competitor km	14.59	14.18	14.02	13.51	14.20	15.14

Table 9: Quadrants of Interest for Positive Δ Sales|Join

4 LASSO Regularization and Competitive Segmentation

In this section we propose and estimate a predictive model of how customer behavior changes after joining the LP. This model takes as inputs detailed data on pre-join spending patterns as well as complex representations of the spatial relationships between customers and their local competitive structures. We broadly define these data on spending patterns as capturing *vertical quality* or the value of a customer in terms of their overall demand, and define the data on local competitive structures as capturing *horizontal quality* or how likely it is to shift a customers spending from a competitor’s store to the focal store. Because spending at competitors is unobserved, spatial data must fill this role. Estimating a model with these two inputs serves two purposes. First, we can analyze the estimation results directly to compare the relative predictive power of our complex competitive structure (horizontal) variables and traditional predictors of LP effectiveness like sales history (vertical) variables, as well as simple rule-of-thumb variables or simple measures of customer

location. This relative predictive power is informative on how LPs manage to increase spending, i.e., via demand expansion or business stealing.

Second, the estimation suggests a segmentation strategy for a firm seeking to improve LP effectiveness by providing predictions of which customers are the most likely to increase spending based on their past spending histories and unique competitive structures. The wide variation across customers in profitability suggests the gains can be quite large and spatial data on customer location may be more accessible than historical sales data in addition to having strong predictive power.

We estimate the impact of the competitive structure on customer behavior with two models. We first model the change in spending conditional on participating in the LP and the competitive structure. Next, we model the probability that a customer joins the firm’s loyalty program, conditional on the competitive structure. We estimate separately to allow these variables to influence the changes in spend and join probabilities differently.

The goal of the estimation is to predict which customers are likely to increase net spending as members of the LP based on their past sales and spatial characteristics. Because the goal is prediction and the data are high-dimensional, the problem is well-suited to statistical methods built around dimension reduction or “regularization.” We ultimately use a LASSO approach because this will allow us to test inclusion of a large number of possible spatial measures and let the model select the most predictive variables. The output also provides a clearer interpretation than other methods and one goal of the analysis is to assess which individual factors best predict LP effectiveness. Finally, the results can be used to assess whether the vertical measures representing spending-based customer value or the horizontal variables representing location-based customer vulnerability to competitors are more effective at predicting behavior. In our discussion of the results we show how this comparison is informative of how and when quantity discounts or behavior-based pricing schemes are likely to succeed.

4.1 Change in Spend

We first estimate the change in monthly store-level spend, conditional on joining the loyalty program. Let y_{is} be the *change* in monthly sales for customer i at store s . We model this relationship as:

$$y_{is} = \beta' X_{is} + \delta' Z_{is} + \varepsilon_{is} \tag{1}$$

where ε_{is} is an unobserved, normally distributed random error centered at zero. Here X_{is} contains the competitive structure parameters between customer i at focal store location s and Z contains firm-level marketing and store-specific customer purchase behavior. These metrics are intended to jointly capture the complex spatial relationships between the customer, focal store, and competitors. Since there can be more than one focal store in the vicinity of an individual, these metrics vary across each store s for a given individual i . Likewise, since each customer’s location is unique, the spatial metrics also vary across each individual i from all observations for store s .

β reflects the impact of the competitive structure on an individual’s likelihood of joining the program. We avoid specific interpretations within this vector of coefficients (e.g. representing the cost of transportation) and instead interpret the coefficients as holistically capturing the many spatial factors that a customer considers before committing to the firm’s loyalty program (e.g., convenience of the focal store relative to the competition, isolation of the focal store, distance to the store versus other competitors, etc.).¹¹

Z_{is} includes traditional recency, frequency, and monetary value (RFM) variables along with basket specific metrics: monthly trip frequency, overall sales, distinct number of items, number of product categories, basket size, discounts received, sales in the category of interest, and sales in the category of interest for those products that are included in the limited LP. We also include the proportion of sales in the category of interest, an indicator for whether the customer received marketing activity specifically encouraging joining the LP, and an indicator for whether the customer received any marketing promotion.

δ captures the effects of the marketing and past purchase behavior on the change in monthly spend. For instance, customers with relatively low trip frequencies may be associated with greater changes in monthly spend due to trip consolidation.

¹¹We recognize that self-selection could potentially lead to biased estimates of change-in-spend parameters since customers choose to enroll into the loyalty program. While this does not effect the predictive power of the model or our ability to compare the relative importance of different data types in contributing to this predictive power, it may effect the direct interpretation of coefficients. We mainly refrain from interpreting these directly, but in addition, to mitigate these concerns, we estimated a Heckman-style two-stage correction method with a flexible control function in an attempt to correct any remaining selection bias. The results suggested that self-selection is inconsequential and the coefficients estimates of interest were nearly identical to what we present in this paper. The full results of this are available in the Appendix. See Heckman (1979) for an introduction to the two-stage estimator and Ahn and Powell (1993) for extensions to more flexible selection correction functions.

4.2 Join Rate

We next estimate the probability that a given customer joins the LP. Having customers enroll in the LP may be of direct interest to the firm regardless of changes in spending that occur due to joining if, for instance, this allows the firm to target future marketing at the customer or use their data for other purposes. Let the indirect utility of joining the firm’s loyalty program for customer i at focal store s be

$$u_{is} = \gamma' X_{is} + \kappa' Z_{is} + \eta_{is} \quad (2)$$

Here X_{is} and Z_{is} are defined in the same way as in the change in spend model. The error term η_{is} captures the idiosyncratic variation in utility across customers and stores. Assuming the $\{\eta_{is}\}$ are independent and identically distributed Type I extreme value, we can derive the join probabilities as follows:

$$\Pr(j_i = 1 | X_{is}, Z_{is}, \gamma, \kappa) = \frac{\exp(\gamma' X_{is} + \kappa' Z_{is})}{1 + \exp(\gamma' X_{is} + \kappa' Z_{is})} \quad (3)$$

where $j_i = 1$ if customer i joins the limited loyalty program at some point during the observed data.

4.3 LASSO Estimation

Our proposed model contains a large number of spatial variables defined to capture the competitive structure. On the one hand, capturing the spatial relationship between the customer, focal store, and competitors is complex and may require many variables. However, all else being equal a concise model is preferable from a managerial standpoint. In addition, many of the spatial metrics we devised may be redundant or highly correlated with each other. To systematically remove variables that are either unnecessary or redundant we employ the Least Absolute Shrinkage and Selection Operator (LASSO) estimator, introduced by Frank and Friedman (1993) and Tibshirani (1996).

In general, for a model with a k -dimensional parameter vector θ the LASSO method performs the following:

$$(\theta^*) = \arg \min \left\{ -\log L(\theta) + \lambda \sum_k |\theta_k| \right\} \quad (4)$$

Where L is the likelihood of data given the model specified, and λ is a tuning parameter that represents the penalty incurred if we choose a nonzero value for any parameter θ_k . This approach regulates the trade-off between an accurate model with more predictive power and a concise model that is more readily interpretable by managers.¹² We select the tuning parameter λ through ten fold cross-validation, which is perhaps the simplest and most widely used method for this task.¹³

4.4 Estimation Results

The results from the LASSO estimation are presented in Table 10. For model validation purposes, we split the original data into a 75% training set and 25% test set. The results presented below were estimated using the training set alone, while the test set is used to evaluate predictive performance. For brevity, the table only displays variables that are significant in at least one of the two models. First we discuss how the competitive structure influences the join probabilities before discussing the estimated impact on the change in monthly spending, as both are influenced differently. From a statistical standpoint, it is relatively easy to interpret the marginal effects of each the variables. However, we must emphasize that many of the spatial impacts are codependent and interpreting the marginal effect of one coefficient while holding another spatial variable constant is sometimes impossible. Later, we present stylized visual representations to better understand the impact of the spatial relationships on change in spend.

4.4.1 Join Probabilities

The LASSO estimation results are shown in Table 10, and the full list of spatial variable definitions is provided in Table 14. In this join model many of the traditional distance metrics drop out of the model (e.g., direct distance between the focal store and the customer) in favor of secondary spatial metrics.

For instance, as the angle between the big box competitor (BB) increases, the probability of joining the LP at the focal store also increases. This indicates that the probability of joining

¹²An alternative approach would be a ridge regression or similar method. In this case the penalty is applied more smoothly, shrinking coefficients on highly correlated variables towards each other. We prefer the LASSO approach because the penalty structure removes coefficients entirely, resulting in clearer model interpretation. Specifically, we can see that many non-spatial variables will drop out altogether. This result is a notable output of interest. A ridge regression would lack this clean interpretation but would result in similar quantitative predictions and can be provided upon request.

¹³See James et al. (2013) for more information.

increases if the competitor is in the opposite direction relative to the customer, rather than in the same direction.

The last two sets of spatial components (*ccbar* and *icbar*) are designed to account for differences in markets with the presence of multiple competitors of the same type. The first of these variables, *ccbar*, is the average distance between each competitor and its competitive center (defined as the average latitude and longitude across competitors of the same type). A relatively high value indicates that the competitors are relatively spread out in a market. The negative coefficient on the big box generalist (*ccbar_1*) indicates that the more dispersed these stores are in a market, the less likely the customer is to join the LP. The second variable, *icbar*, measures the average distance between each customer and each of the competitors. While a customer may be close to the competitive center, a high value of *icbar* indicates that each individual competitor (of a given type) is still relatively far away. This coefficient is positive for the two small-box generalists (SB1 and SB2), suggesting that a relatively sparse distribution of these competitors may be beneficial to the focal firm.

Finally, we see a strong positive effect on the proportion of sales in the category of interest: customers who dedicate a greater portion of store spend to the LP category are more likely to benefit from the LP discounts. There is also a positive impact of whether or not they received a marketing promotion, for both general emails and those specifically designed to enroll them in the LP.

4.4.2 Change in Monthly Spend

In this section we now review the importance of competitive structure as a predictor of the estimated LP effectiveness, as measured by change in monthly spending at the focal store, conditional on joining the firm's loyalty program. Recall the dependent variable in this estimation is the change in monthly spend after joining the loyalty program, net of any LP discounts. For this table, we only included customers who had at least some spending both before and after joining. Since we are interested in how spending *changes* after joining, rather than *whether* spending occurs or not, this seems to be a reasonable filter.

First, there is a positive impact on the squared distance between the customer and the focal store, suggesting that the greater gains come from those located further away from the focal store. This aligns with our previous intuition: customers located near the focal store are more likely to

already dedicate the majority of their spend with the focal store, so the potential for increases in spend are limited, relative to those who are further away and thus more likely to be sharing spend with the competition.

Importantly, the distance between the customer and the competition also influences changes in spend at the focal store. Customers who exhibit the largest changes in spend are those who are relatively far from the the big-box generalist (ic1sq) and the second small-box generalist (ic3sq) and close to from the other competitors, holding all other variables constant.

Many of the radius band coefficients drop out of this model. However, we see that the intercept angle is retained. The difference in the estimated change in spend for a customer where a big-box (BB) competitor is in the same line of travel as the focal store (intercept angle of zero) versus in the opposite direction (intercept angle of 180 degrees) is about \$124, holding all other variables constant. This is a non-trivial amount that cannot be captured using traditional spatial measures. However, it is important to again point out that many of these distance metrics cannot vary independently of others. The visual results presented later will illustrate the degree to which the intercept angle influences the change in spend while simultaneously accounting for other changes.

Finally, population density along with its interactions with the direct distance measures has a significant impact on the change in monthly spend. On its own, the effect is positive: more densely populated areas are associated with a greater change in spend. This is intuitive as more densely populated areas tend to have a greater level of competition and higher incomes. There are also numerous interactions with our simple distance metrics. For instance, increases in the distance between the store and the customer influence customers in dense areas more negatively, relative to customers in less dense areas. This appears to be capturing the challenges of traveling in more densely populated areas (e.g., traveling a given distance in a city versus a rural setting).

As an important robustness check, we also include a column where the estimation proceeds as before but only using the change in spending on products that do not qualify for the LP discounts. Comparing this column and the column for all spending shows a very high correlation in the LASSO coefficients and perfect correlation in which variables are retained.

4.4.3 Summary

The LASSO results suggest that the customer behavior appears to be strongly influenced by the competitive structure of the local market. The regularization tends to keep quite a few of the

spatial metrics, suggesting that relatively simple distance metrics are, on their own, unable to fully characterize the predicted behavior. This is not very surprising: there are many intuitive reasons why a customer may change their spending patterns after joining a loyalty program with respect to their location.

4.5 Visual Representation of Results

Even after variable reduction via LASSO, it is still difficult to interpret how the competitive structure influences estimated changes in spend due to the codependency between variables. We therefore visually illustrate the LASSO's results on how the competitive structure influences LP effectiveness. For brevity, we limit our discussion to two heatmaps that show the predicted change in spending for hypothetical customers. Additional heatmaps are provided in the Appendix. These maps show how there is large heterogeneity across different customers with different competitive structures in the potential gains from having them join the LP. However, since these maps are mostly stylized representations of the actual data we focus on comparing the magnitudes across maps rather than specific prediction levels.

Figure 3 plots all four competitors, each set an equal distance from the focal store. The color on the heatmap reflects the estimated change in spend from a customer in that position if they were to join the focal store's loyalty program, with shades closer to white representing greater increases. This map shows that most of the gains are from customers near the big-box generalist, and less so from the small-box specialist. This map highlights that the value of a potential customer is heavily dependent on the extent to which the focal store can steal business from the competition. This also represents visually how a targeting strategy might be employed using complex representations of spatial relationships constructed from readily available data, even in the absence of past sales data on customers.

Table 10: Lasso Pr(Join) and Change in Sales|Join

	Coefficient	Pr(Join)	Change in Sales Join	Change in Non-Qualified Sales Join
1	(Intercept)	-3.5018	-27.9689	-46.4571
2	sisq		0.0194	0.0280
3	ic2	0.0010		
4	ic3	0.0054		
5	ic1sq		0.2158	0.2088
6	ic2sq		-0.1387	-0.1415
7	ic3sq		0.1429	0.1315
8	ic4sq	-0.0001	-0.1381	-0.1043
9	sc1	0.0026		
10	sc1sq	0.0001	-0.0308	-0.0322
11	sc2sq		-0.0615	-0.0375
12	sc3sq		-0.0950	-0.1016
13	sc4sq	-0.0001	0.1128	0.0874
14	s4	-0.3020		
15	nAll_1	0.0225		
16	n1_1	0.6830		
17	n3_1	0.1946		
18	n5_1	0.1593		
19	n10_1	-0.0474		
20	n15_1	0.0514		
21	nAll_2	-0.0014	0.4183	0.3750
22	n1_2	-0.2532		
23	n10_2	0.0332		
24	n15_2	-0.0158		
25	nAll_3	-0.0141		
26	n1_3	0.2901		
27	n3_3	-0.0033		
28	n5_3	-0.0690		
29	n10_3	0.0579		
30	nAll_4	-0.0077		
31	n3_4	0.1371		
32	n15_4	-0.0044		
33	cInt1: BB intercept angle	0.0013	0.6900	0.7909
34	cInt2: SB1 intercept angle		-0.0945	-0.1817
35	cInt3: SB2 intercept angle	0.0005	0.5190	0.5555
36	ccbar_1 (avg. BB-BB comp. center km)	-0.0126		
37	ccbar_2 (avg. SB1-SB1 comp. center km)	0.0097	0.4599	0.5875
38	icbar_2 (avg. customer-SB1 km)	0.0036		
39	icbar_3 (avg. customer-SB2 km)	0.0039		
40	icbar_4 (avg. customer-SS km)	-0.0001		
41	Prop. of sales in LP category (pre LP)	5.7842		
42	Received promotional email	0.0762		
43	Received promotional email for LP	1.3227		
44	Trips/month (pre LP)	0.0106	-24.5945	-24.1253
45	Total sales (pre LP)	0.0002	0.0037	0.0002
46	Total sales in LP category (pre LP)		0.1177	0.1106
47	Total qualified sales in LP category (pre LP)		-0.1678	-0.1058
48	Number of focal stores w/i 60km	0.0035	1.0196	1.0950
49	Population Density (in sq miles)		0.0290	0.0262
50	Population Density x si		-0.0016	-0.0016
51	Population Density x ic1		-0.0023	-0.0021
52	Population Density x ic2		-0.0007	-0.0011
53	Population Density x ic3		-0.0003	-0.0002
54	Population Density x ic4		0.0043	0.0043
55	Population Density x sc1		0.0001	
56	Population Density x sc2		0.0014	0.0012
57	Population Density x sc3		-0.0007	
58	Population Density x sc4		-0.0026	-0.0028
59	distinctSKU	-0.0055		
60	numItems	0.0019	-4.1178	-3.6378
61	distinctCategories	-0.0989		

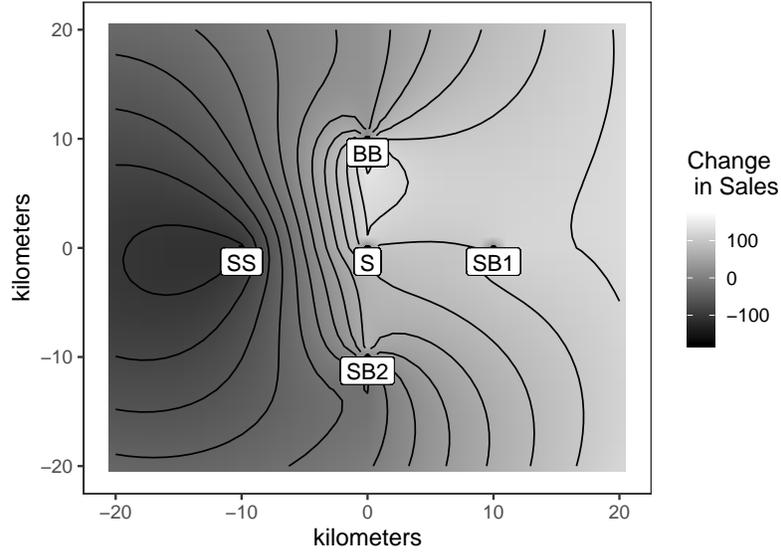


Figure 3: All Competitors

Figure 4 illustrates the predicted effects from an actual competitive structure randomly drawn from the data. Customers located on the far side of the second small-box specialist (SB2), and the big-box generalist (BB), have the lowest predicted change in spend. These customers are relatively far away from the focal store with a competitor directly in their path of travel, resulting in limited predicted gains. On the other hand, the other two competitors (SB1 and SS) are relatively close and the line of travel is not as critical with these competitors. The customers that show the most promise are those that are relatively close to the focal store but in the direction of the second small box generalist (SB2) and the big-box generalist (BB).

We also included two customer locations (I1 and I2) to illustrate more specifically the impacts of the competitive structure. The first customer, I1, is roughly halfway between the focal store and the second small box competitor (SB2). This customer has relatively high predicted change in spend in part, we suspect, of the potential to consolidate purchases. In contrast, the second customer, I2, is located right near two other competitors. Because this customer is so close to the competition, relative to the focal store, the expected change in sales is slightly dampened.

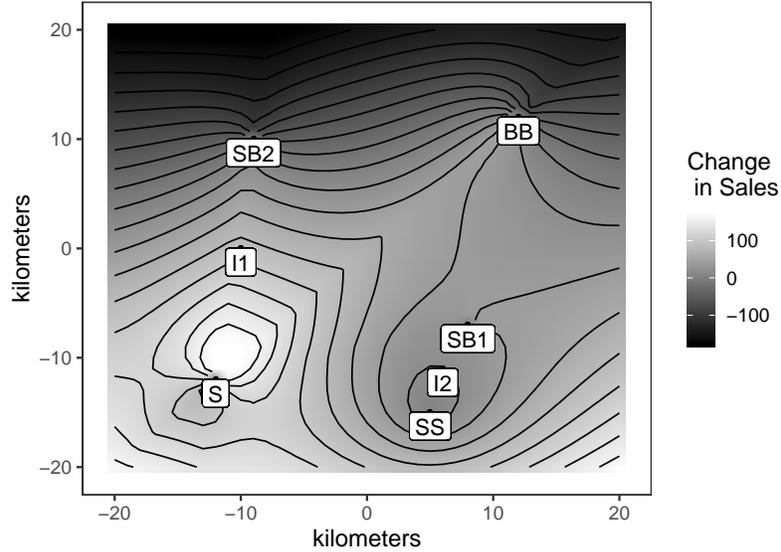


Figure 4: Randomly Selected Competitive Structure

The heatmaps presented convey two key advantages of this analysis. First, the geographic influence of a competitor is relatively complex. It makes clear how simple radii surrounding either the competitor or the focal store would not sufficiently capture the influence of the competitive structure on customer behavior. Second, these complex relationships are not fixed and can vary by the type of competitor in an area. This is an important consideration for the firm in the formation of targeting strategies: the combined relationship between the location of each competitor, customer, and focal store strongly influences the customer’s interaction with the focal firm.

4.6 Competitive Segmentation: Out of Sample Validation

In this section we use the estimated model to demonstrate the value of augmenting traditional sales data with readily available spatial data on competitors. We use the LASSO results to form a segmentation strategy designed around predicting how customer spending changes upon joining the LP. The estimation in the previous sections allows us to predict change in spending for each potential customer as a flexible function of their proximity to each store type. Our aim is to motivate firms to augment their detailed customer transaction data with information on competitor locations to identify customers who are likely to be most profitable upon joining a LP. By identifying which

segments are more likely to increase spending upon joining the LP the firm can focus its limited marketing budget on promoting the LP to those customers and avoid promoting it to those who are predicted to not alter spending patterns and simply take discounted prices. More generally, our results suggest that firms should consider segmenting on both horizontal quality and vertical quality to identify their most valuable customers.

To evaluate the potential increase in profits from this type of targeting, we use a validation or holdout sample of customers who were not used in estimation. For simplicity, we segment customers at the store level, rather than the firm level, since the competitive structure varies across each customer-store interaction. Our primary dependent variable of interest is the expected change in monthly spending.

Importantly, we only predict the estimated change in spend from customers who actually joined the LP. We then measure the success of each segmentation approach by calculating the actual change in spending from those customers whose predicted incremental profits from joining the LP is positive. In this way, we can evaluate each segmentation strategy by focusing on different subsets of customers who actually joined, acknowledging the fact that there is self-selection into the LP that may be motivated by planned changes in spending not caused by the LP itself. In recognition of this selection we do not use the model to forecast how spending would change among customers we do not observe joining the LP. Rather, the goal is to illustrate how spatial information helps to segment customers upon joining the LP to customer value predictions. A hypothetical targeting strategy that might follow this approach would therefore work by trying to prompt the most profitable customers to join the LP and simultaneously not advertise the LP at all to customers outside this group. We also note that a closely related strategy using spatial competitive structure variables to target LP promotion could be conducted at the store level rather than the customer level. This is not possible with RFM based targeting and could easily be implemented with in-store promotions or even store-level eligibility restrictions for LP discounts.

Rule-of-Thumb Segmentation Strategies:

The predictive performance of the spatial LASSO strategy is compared to numerous other potential segmentation rules and the results are shown in Table 11. We first describe a series of simple segmentation criteria that are not based on the full competitive structure but instead rely on simpler spatial information or more traditional RFM metrics. In each case we compare the predictive performance of the rules by evaluating the total change in spending of all customers who

would be identified by the rule.

First, a mass marketing strategy would be to simply avoid segmentation and treat the customers as a homogeneous group. In aggregate, customers show a slightly positive change in monthly spending after joining. This is consistent with the descriptive result that spending only mildly increases on average after customers join the LP.

We also consider a simple strategy based on vertical quality, consisting of focusing on only those that exhibit relatively high spending amounts (prior to joining the LP) at a few reasonable cutoff points (based on the empirical distribution of monthly spend at a single store location). We consider three levels as a cutoff: \$80 (the median), \$100, and \$200 (about the 75th quantile). In each of these the change in monthly spend actually declines substantially. Segmentation in this manner would therefore be a highly unprofitable strategy.

We also consider a naive location-based strategy that selects customers within a few kilometers of the focal store. If we limit attention to only those customers within 5 kilometers of the target store we get a small positive boost, which is only slightly stronger than a mass marketing approach.

A second type of naive location-based segmentation would be to emphasize horizontal variables and try to take advantage of the business stealing aspect of loyalty programs by focusing on all customers located near a competitor. We test this strategy and find that customers within 5 km of a competitor exhibit a substantial decrease in profits.

Gravity Model-Based Segmentation Strategies:

We next consider segmentation strategies based on formally estimated models of customer change in spending. To provide a comparison, we begin with three spatial models in the spirit of classic gravity models. These models date back to Reilly (1931) who coined the “law of retail gravitation” and are extended to a more flexible model by Huff (1964). These gravity models are designed to predict store choice while we have richer data on spending, but only at the focal store. We attempt to preserve the theoretical insight of these models but adjust the specification to instead predict the probability of a positive change in spend with $p_{k=1} = \frac{w_1/d_1^\alpha}{\sum_{k=1}^5 w_k/d_k^\alpha}$, where d represents the distance between the customer and each store type k (the focal store and the four competitors) and w are the weights or “attractiveness” of each store.

We consider several Huff-style model specifications. In the first, we set the store attractiveness or mass to be the square footage of each competitor type and in the second specification we treat each store attractiveness as a parameter to estimate. In both of these cases, the decay parameter α

is estimated using maximum likelihood. In our third specification we allow α_k to differ over store type. For segmentation, we take the models predictions and group the customers based on the percentile of positive spend in the training data.¹⁴

We find that the gravity models perform quite well overall at identifying profitable customer groups. They outperform the mass marketing approach and the rule of thumb strategies. They also do better out-of-sample than the naive models using spatial and RFM data, suggesting that those models are prone to overfitting in a way the more parsimonious gravity models are not. In the next subsection we show that our preferred specification performs about three times as well as the Huff-style gravity models, however.

LASSO Model-Based Segmentation Strategies:

Next, we compare several series of models distinguished by which data are used as inputs as well as how they are estimated. Each model is estimated on the training observations and validated on the holdout set. Note that we expect a higher degree of noise between the train and test set than in other applications, particularly LASSO applications, because we are not simply estimating change in spend but estimating a predictive model of change in spend and then grouping all customers with a predicted spend greater than zero. This comparison is more suited to the actual firm application of segmentation compared to a simple forecast of change in spend.

First we test a naive spatial model that only considers the standard basic spatial data, i.e. distances between the focal store and customer, and each of the four competitors. We see that the out-of-sample performance is actually worse than a mass marketing approach. In other words, while horizontal data is valuable in principle a simplistic approach to incorporating it that misses the nuances of competitive structure is not necessarily effective.

The second model considers only vertical quality or historical sales information: proportion of sales in the category of interest, overall sales, category level sales, and category level sales that specifically qualify for the LP. This model also under performs a mass segmentation approach and results in negative profits, suggesting that these data alone have poor predictive power and are prone to overfitting.

The third model uses the LASSO approach but ignores competitor spatial information. We estimate the model using traditional RFM metrics alone. This model outperforms all others thus

¹⁴We attempted numerous variations on the Huff model given that it is not a direct translation of the original method, and kept only those that performed best at predicting change in spending. More details are available upon request.

far on nearly every metric both in-sample and out-of-sample, attesting to the strength of using the LASSO approach as a baseline for improving predictive power.

Next, we show results for our preferred spatial LASSO approach. Here we augment the LASSO model with the spatial information of the competition. It outperforms all other segmentation strategies in the validation sample, with the actual change in sales over six times higher than that of a mass marketing targeting scheme. The gains compared to the RFM LASSO model are also substantial: about a 30% increase in the out-of-sample change in sales and change in non-qualified sales. In short, the spatial LASSO does the best job of predicting which LP joiners end up being most valuable. In addition to its superior out-of-sample performance, the spatial LASSO shows greater stability between in-sample and out-of-sample predictions, relative to most other approaches.

Finally, we compare the performance of an RFM approach to a spatial approach when estimated and validated separately in competitive and isolated markets.¹⁵ The spatial model outperforms the RFM only model in all comparisons where the market type (isolated or competitive) or data subset (train or test) is held constant.

In the training set, the RFM model performs slightly better at predicting overall sales in isolated markets and slightly better at non-qualified sales in competitive markets. The spatial model, by contrast, performs substantially better in competitive markets than in isolated markets. This shows how the data on horizontal relationships primarily add predictive value in settings where customers are vulnerable to competition. In isolated markets, there is close to zero value added from horizontal metrics compared to vertical metrics.

These results are broadly informative on how 2nd degree price discrimination (quantity discounts or behavior-based pricing) and 3rd degree price discrimination (targeted price discounts) interact in practice. Properly designed quantity discounts should be profitable by increasing spending by high value customers. Therefore, in isolated markets it is effective to target these discounts using a measure of vertical quality constructed from past sales data. The quantity discounts work by letting customers self-select into the discounts based on their level of spending and marginal demand. In markets with close competitors, targeting can be done based on the horizontal vulnerability dimension and within the targeted group customers will self-select into the discounts based on vertical characteristics. It is better then to target customers who might shift their purchases away

¹⁵These market types are defined in section 3.

from competitors than to target the highest value customers alone, since they may not change their spending at all in response to the LP discounts.

One argument might be that since marketing is done via email and is relatively inexpensive, why not simply send marketing campaigns to everyone and not bother with segmenting the customers? The key here is that in some cases it might be preferable to limit LP enrollment to certain individuals. While the firm cannot prevent interested customers from joining its loyalty program, it can choose which customers to invest in after they join (e.g., special offers). Identifying the right type of customer for this type of promotion can substantially increase its profitability. While this has generally been known, we show two distinct contributions. First, location data adds substantially more predictive power to a segmentation strategy than historical sales patterns. Second, simply incorporating traditional spatial metrics like distance between customer and store does not capture the complexity of competitive structure. Instead, we show how firms could take advantage of rich spatial data on both customers and competitors locations using a more sophisticated but still straightforward approach.

Targeting Strategy	Avg. Δ Sales		Avg. Δ Non-Qual. Sales	
	Train	Test	Train	Test
Mass Marketing	-\$18.69	\$6.91	-\$24.32	\$2.59
Top spender (>\$80)	-\$83.07	-\$38.10	-\$88.18	-\$33.29
Top spender (>\$100)	-\$95.38	-\$45.16	-\$99.92	-\$40.19
Top spender (>\$200)	-\$181.13	-\$103.67	-\$178.76	-\$95.73
Within 5 kilometers of focal store	\$46.60	\$12.60	\$38.86	\$13.23
Within 5 kilometers of any competitor	-\$47.57	-\$54.54	-\$53.00	-\$41.93
Gravity fixed w ($\hat{\alpha} = .227, .232$)	\$13.76	\$20.16	\$13.11	\$22.34
Gravity estimated w ($\hat{\alpha} = .039, .041$)	\$11.05	\$24.90	-\$4.28	\$30.47
Gravity α_k	\$17.93	\$19.43	\$17.80	\$13.90
Naive spatial	\$40.72	\$3.97	\$55.13	\$11.21
Historical sales only	\$29.06	-\$23.06	\$28.96	\$15.58
LASSO RFM only	\$58.75	\$36.43	\$54.92	\$24.56
LASSO w/Spatial	\$78.34	\$47.09	\$70.06	\$31.64
LASSO RFM only (Competitive)	\$58.16	\$31.60	\$56.67	\$31.76
LASSO RFM only (Isolated)	\$57.89	\$37.94	\$53.25	\$22.31
LASSO w/Spatial (Competitive)	\$86.65	\$72.39	\$81.63	\$67.72
LASSO w/Spatial (Isolated)	\$75.85	\$38.08	\$66.60	\$18.78

Table 11: Targeting Validation Performance

Comparison of Sales Predictions by Customer Type:

To better understand the contribution of horizontal variables in prediction exercise, we show how the different segmentation strategies would select customers who differ on vertical and horizontal metrics.

Table 12 presents the simplest spatial metrics (distances between the customer, focal store, and each competitor) and monthly spending of the customers targeted in each strategy. Since there is considerable overlap between the customers across these strategies, we parse out the observations that were predicted to increase sales in the spatial LASSO but not by the alternative segmentation strategies. In the table we label these as “missed” opportunities.

We see that customers predicted to increase spending by the model using traditional RFM

metrics alone have lower pre-join sales per month relative to customers selected by the spatial LASSO and are positioned significantly further from the focal store. This finding is made evident when looking at the missed opportunities: these customers are much closer to the focal store and the sales per month is much higher. The other distance metrics show that these customers and store locations also tend to be closer to the competition. This emphasizes one of the key insights of the spatial model: on average, regardless of location, customers with higher than average sales per month are unlikely to show large increases in spend after joining the LP (likely because they are already dedicating most spend to the focal store). However, we posit that customers located relatively closer to both the focal store and the competition are more likely to consolidate purchases away from the competition and will exhibit higher than average changes in spend upon joining.

We next evaluate the customers from the rule of thumb segmentation strategy that focuses on customers close to the focal store. This illustrates a similar issue as the previous strategy. On average, customers close to the focal store are likely to have access to a competitor as well, so at first glance these customers appear to be promising prospects for targeting. However, they may be *too* close to the focal store, in the sense that they already dedicate a substantial amount of spend to the focal store (as evidenced by the relative high sales per month). The spatial LASSO recognizes that customers relatively further away may prove to be promising, so long as they have relatively similar access to competitors.

Finally, we consider the second Huff-style gravity model with estimated weights w , since, of the gravity models, this variation best segmented customers. Under this model, valuable customers are excluded simply because they are relatively far away from the focal store. However, for these customers all four competitors are closer to the focal store, compared with customers identified by the Huff model. We again find that the most valuable customers, in terms of expected change in spend, are those with the right balance distance between them and the focal store locations and the competition. This highlights a potential weakness of a traditional gravity model, which does not explicitly consider distances between the store locations, only the distances to the customer.

This analysis shows how the spatial LASSO model can capture more subtle drivers of customer behavior that are typically missed when looking at transaction data alone (e.g., RFM variables) or relatively simple spatial components. We find that there are substantial gains in segmentation actions by accounting for more complex spatial relationships among the customer, focal store, and relevant competitors.

	Out-of-Sample Performance						
	LASSO (RFM only)			w/i 5km of Focal Store		Gravity (est. w)	
	LASSO	Targeted	Missed	Targeted	Missed	Targeted	Missed
Δ Monthly Sales	\$47.09	\$36.43	\$32.78	\$12.60	\$58.15	\$24.90	\$70.08
Δ Qualified Sales	\$15.46	\$11.87	\$18.45	-\$0.63	\$20.88	\$0.15	\$23.99
Δ Non-Qualified Sales	\$31.64	\$24.56	\$14.33	\$13.23	\$37.27	\$24.75	\$46.09
Sales/Month (pre LP)	\$161.05	\$138.07	\$377.26	\$424.25	\$111.97	\$385.13	\$86.40
Sales/Month (pre LP) SD	\$260.13	\$241.78	\$351.43	\$534.57	\$187.42	\$482.54	\$126.30
Sales/Month (pre LP) 25th	\$21.43	\$22.41	\$146.88	\$107.78	\$15.76	\$91.43	\$15.53
Sales/Month (pre LP) 75th	\$150.33	\$123.19	\$464.56	\$507.74	\$111.06	\$398.32	\$167.25
Sales/Month (post LP)	\$209.67	\$177.76	\$410.03	\$430.11	\$171.02	\$398.32	\$167.25
si: Store-Customer (km)	9.48	13.78	5.81	3.00	12.52	5.20	13.86
ic1: Customer-BB (km)	14.19	14.61	14.20	13.00	15.05	14.28	13.76
ic2: Customer-SB1 (km)	13.90	16.83	12.16	15.55	14.41	14.44	15.59
ic3: Customer-SB2 (km)	16.83	17.46	17.79	14.98	18.44	17.14	16.92
ic4: Customer-SS (km)	14.96	16.17	13.23	13.62	16.02	14.69	15.46
sc1: Store-BB (km)	14.31	14.72	14.90	12.78	15.18	15.01	12.65
sc2: Store-SB2 (km)	13.96	16.58	13.02	15.47	14.42	14.34	15.05
sc3: Store-SB1 (km)	16.20	17.61	19.09	14.92	17.49	18.02	14.50
sc4: Store-SS (km)	14.24	14.87	14.07	13.55	14.83	15.15	12.95

Table 12: Targeted Customers

4.7 Spatial Variable Contribution

In the previous section, we emphasized the value of integrating spatial or horizontal information into the segmentation process. In this section we provide an additional measure of the contribution of spatial variables to the model fit relative to other types of variables. Specifically, we highlight that our spatial metrics contribute substantially more towards R^2 relative to traditional vertical variables.

Table 13 presents the model fit using our two dependent variables of interest, where each row represents a single subset of the available variables. The first column displays the pseudo- R^2 from

a logistic regression and the second column the standard R^2 from a linear regression. The variables are separated into one of four labels: our proposed spatial metrics, RFM (trip frequency and sales information), basket composition (number of items per basket, distinct items, distinct categories, etc.), and finally the marketing indicators.

The spatial metrics explain considerably more than any of the alternative variable types. The sum of the R^2 values across the partial models is similar to the fit from the aggregate model in the final row. This suggests that the variable types tend to be relatively orthogonal to each other. In other words, the spatial information appears to be adding a non-trivial amount of unique information to the model that would otherwise be absorbed into the residuals.

A natural rebuttal to this table is that the results are driven simply by the sheer number of spatial variables relative to the other variable types. This is only partially true: recall that all of the spatial metrics are derived solely from the latitude and longitude of the focal store, competition, and the customer. From a managerial perspective it is very easy to recreate the diverse set of spatial metrics once these few location points are obtained. In this sense, it is more the variety among the spatial variables we designed, rather than simply the number of variables, that contributes a substantial amount of information to each model.

Variable Type	Pr(Join)	Δ Sales Join
Spatial	3.4%	14.9%
RFM	0.6%	9.8%
Basket composition	1.3%	3.8%
Marketing	0.1%	0.0%
All variable types	5.2%	25.5%

Table 13: Pseudo- R^2 and R^2 by variable type

5 Discussion and Conclusion

Both academics and managers have been interested in understanding the factors contributing to a successful loyalty program. We use a large new dataset on customer spending behavior at a major U.S. retailer to develop a better understanding of the role of spatial data and competitive structure in LP performance and in the effectiveness of price discrimination strategies more broadly. We

argue that a significant but largely overlooked driver of a loyalty program’s success is the joint spatial relationship among the focal firm’s stores, the competition, and the customer, or what we collectively label the competitive structure.

We first document large changes in behavior after joining the LP, with a wide degree of heterogeneity across customers. Notable among this customer heterogeneity are customer segments we describe as consolidators or as upgraders. This heterogeneity suggests that targeting the LP with promotions could be effective in dealing with the well-known issue that in many cases firms believe that LPs fail to increase profits.

Second, we find strong model-free evidence that access to competitor stores is the strongest predictor of whether a customer will increase spending upon joining the LP. This suggests that LPs using the common form of quantity discounts work via business stealing rather than overall demand expansion. Third, in informal comparisons we show that past spending behavior is not very predictive of how spending will change when a customer joins the LP. These data are less predictive than spatial data and in particular the more complex representations of spatial relationships.

We then develop formal predictive models of join rates and changes in spend post-join. This provides several important results. First, segmenting using quantity discounts (or performing second degree price discrimination) naively to the highest spenders would be very unprofitable. Naively segmenting customers based on simple spatial rules would not be very effective either, performing roughly the same as not treating all customers the same. Next, using LASSO to select among many possible relationships in the estimation data helps predict changes in spending substantially better than simple models. Then, comparing performance of LASSO models relying on vertical variables as inputs to using horizontal variables shows a substantial improvement associated with the readily available spatial data.

These results suggest a segmentation strategy for firms interested in identifying their most profitable customers. They are also informative regarding how and when second degree price discrimination or one type of behavior-based pricing are likely to be successful. As Shin and Sudhir (2010) note, it can be profitable to use this type of pricing strategy if there exists substantial heterogeneity in customer value and customer preferences are stochastic. Our results confirm this insight and operationalize it by defining a customer’s vertical quality using past spending data and their horizontal quality using spatial relationships to competitors and the focal store. Quantity discounts based on vertical quality can be effective in relatively isolated markets but in competitive

settings targeting on horizontal variables is much more effective.

Still, our analysis is not without limitations. Due to the nature of the data we cannot make very strong statements on the extent of business stealing that may be occurring. In the absence of data from multiple competitors or direct responses from the customers themselves, we can only infer the degree of business stealing based on changes in shopping behavior as a function of competitive structure. We hope that our research provides the foundation for future work that integrates the complex spatial information into models of customer behavior, especially in their relationship to making more strategic firm decisions.

References

- Ahn, H. and Powell, J. L. (1993). Semiparametric estimation of censored selection models with a nonparametric selection mechanism. *Journal of Econometrics*, 58(1-2):3–29.
- Ascarza, E. (2018). Retention futility: Targeting high-risk customers might be ineffective. *Journal of Marketing Research*, 55(1):80–98.
- Belloni, A., Chernozhukov, V., and Hansen, C. (2014). Inference on treatment effects after selection among high-dimensional controls. *The Review of Economic Studies*, 81(2):608–650.
- Besanko, D., Dube, J.-P., and Gupta, S. (2003). Competitive price discrimination strategies in a vertical channel using aggregate retail data. *Marketing Science*, 49(9):1121–1138.
- Bijmolt, T. H., Dorotic, M., Verhoef, P. C., et al. (2011). Loyalty programs: Generalizations on their adoption, effectiveness and design. *Foundations and Trends® in Marketing*, 5(4):197–258.
- Blattberg, R. C. (2008). Byung-do kim, and scott a. neslin (2008), database marketing: Analyzing and managing customers. *International Series in Quantitative Marketing*.
- Chen, Y., Li, X., and Sun, M. (2016). Competitive mobile geo targeting. *Marketing Science*, 36(5):645–812.
- Chen, Y., Narasimhan, C., and Zhang, Z. J. (2001). Individual marketing with imperfect targetability. *Marketing Science*, 20(1):1–98.
- DeWulf, K., Odekerken-Schröder, G., and Iacobucci, D. (2001). Investments in consumer relationships: A cross-country and cross-industry exploration. *Journal of Marketing*, 65(4):33–50.
- Dreze, X. and Nunes, J. C. (2008). Feeling superior: The impact of loyalty program structure on consumers’ perceptions of status. *Journal of Consumer Research*, 35(6):890–905.
- Dube, J.-P., Fang, Z., Fong, N., and Luo, X. (2017). Competitive price targeting with smartphone coupons. *Marketing Science*, 36(6):813–1017.
- Ellickson, P. and Misra, S. (2012). Enriching Interactions: Incorporating Revenue and Cost Data into Static Discrete Games. *Quantitative Marketing and Economics*, 10:1–26.
- Fong, N. M., Fang, Z., and Luo, X. (2015). Geo-conquesting: Competitive locational targeting of mobile promotions. *Journal of Marketing Research*, 52(5):726–735.
- Hartmann, W. and Viard, B. (2008). Do frequency reward programs create switching costs? a dynamic structural analysis of demand in a reward program. *Quantitative Marketing and Economics*, 6(2):109–137.
- Heckman, J. J. (1979). Sample selection bias as a specification error. *Econometrica*, 47(1):153–161.
- Huff, D. L. (1964). Defining and estimating a trading area. *The Journal of Marketing*, pages 34–38.
- James, G., Witten, D., Hastie, T., and Tibshirani, R. (2013). *An introduction to statistical learning*, volume 112. Springer.

- Li, Y., Gordon, B., and Netzer, O. (2018). An empirical study of national vs. local pricing by chain stores under competition. *Marketing Science*, 37(5):812–837.
- 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.
- Lockhart, R., Taylor, J., Tibshirani, R. J., and Tibshirani, R. (2014). A significance test for the lasso. *Annals of statistics*, 42(2):413.
- Luo, X., Andrews, M., Fang, Z., and Phang, C. W. (2014). Mobile targeting. 60(7):1738–1756.
- McCall, M. and Voorhees, C. (2010). The drivers of loyalty program success: An organizing framework and research agenda. *Cornell Hospitality Quarterly*, 51(1):35–52.
- Meyer-Waarden, L. and Benavent, C. (2009). Grocery retail loyalty program effects: self-selection or purchase behavior change? *Journal of the Academy of Marketing Science*, 37(3):345–358.
- Reilly, W. (1931). The law of retail gravitation. *wj reilly. Inc., New York*.
- Roehm, M. L., Pullins, E. B., and Roehm Jr, H. A. (2002). Designing loyalty-building programs for packaged goods brands. *Journal of Marketing Research*, 39(2):202–213.
- Rossi, P., McCulloch, R., and Allenby, G. (1996). The value of purchase history data in target marketing. *Marketing Science*, 15(4):301–394.
- Shaffer, G. and Zhang, Z. J. (1995). Competitive coupon targeting. 14(4):343–466.
- Shin, J. and Sudhir, K. (2010). A customer management dilemma: When is it profitable to reward one’s own customers? *Marketing Science*, 29(4):671–689.
- Shugan, S. (2005). Brand loyalty programs: Are they shams? *Marketing Science*, 24(2):185–193.
- 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.

Appendix 1: Distance Metrics

In the table below, x denotes the competitor type. In our empirical application this is either the big box generalist (BB), the small box generalists (SB1 or SB2), or the small box specialist (SS).

Distance Metric	Description
si	Distance between the focal store s and individual i (in km)
icx	Distance between the individual and competitor x (in km)
scx	Distance between the focal store and competitor x (in km)
sx	Sparsity metric for competitor x , defined as $scx/(icx + si)$
$nAll_x$	Total number of x competitors within the region
nd_x	Total number of x competitors within d km of the individual
$cIntx$	Angle formed between focal store and competitor with individual at apex (0 means s and x are in same direction relative to c 180 means s in opposite direction of x)
$icbar_x$	Average km between the individual and each individual competitor x (if only one competitor of type x in an area, this equals icx)
$ccbar_x$	Average km distance between each competitor and the competitive center (if only one competitor of type x in an area, this equals 0)

Table 14: Distance Metric Descriptions

Appendix 2: Selection Correction

We also test a specification of our LASSO model that includes a Heckman-style two-stage correction method in an attempt to offset for or correct any selection bias in the change in spending. To do so we combine a two-step estimation of the decision to join the LP and the post-join change in spending, with exclusion restrictions identifying the selection effect. First we estimate the joint probabilities using the model specification above. Then, we take the estimated joint probabilities and include them in the second stage model of change in spend as a control function. Our approach is to place a flexible function ϕ over \hat{p} using high-order polynomials to control for unobserved heterogeneity. See Heckman (1979) for an introduction to the two-stage estimator, Ahn and Powell (1993) for extensions to more flexible selection correction functions, and Belloni et al. (2014) for a discussion of estimating selection models with high-dimensional controls and Lockhart et al. (2014) for significance testing in LASSO. For a recent application, see Ellickson and Misra (2012).

The change in spend model used in the second stage model is modified as follows:

$$y_{is} = \beta' X_{is} + \delta' Z_{is}^* + \phi(\hat{p}_{is}) + \varepsilon_{is} \quad (5)$$

Where Z_{is}^* is the same as Z_{is} but with three variables removed that satisfy the exclusion restriction: the proportion of sales in the category of interest, an indicator for whether the customer received marketing activity specifically encouraging joining the LP, and an indicator for whether the customer received any marketing promotion. The variables arguably influence the probability of joining the loyalty program but are unlikely to influence change in spend behavior once joined.

Table 15 shows the coefficient estimates when employing a two-stage correction method. The coefficients on the polynomials shrink to zero in the LASSO model, suggesting that self-selection is unlikely to be a concern (analogous to a coefficient of zero on an inverse Mills ratio in a traditional selection model). This leads to coefficient estimates, and subsequent targeting predictions, that are similar to what is already presented in the main findings. In the interest of conciseness, we decided to keep the simpler modeling framework.

Table 15: Lasso Pr(Join) and Change in Sales|Join

	Coefficient	Pr(Join)	Change in Sales Join	Change in Non-Qualified Sales Join
1	(Intercept)	-3.5018	-19.3914	-44.8786
2	sisq		0.0042	0.0149
3	ic2	0.001		
4	ic3	0.0054		
5	ic1sq		0.2113	0.2076
6	ic2sq		-0.1402	-0.1407
7	ic3sq		0.146	0.135
8	ic4sq	-1e-04	-0.1257	-0.0991
9	sc1	0.0026		
10	sc1sq	1e-04	-0.0207	-0.0237
11	sc2sq		-0.0453	-0.0284
12	sc3sq		-0.1096	-0.1168
13	sc4sq	-1e-04	0.1089	0.0917
14	s4	-0.302		
15	nAll_1	0.0225		
16	n1_1	0.683		
17	n3_1	0.1946		
18	n5_1	0.1593		
19	n10_1	-0.0474		
20	n15_1	0.0514		
21	nAll_2	-0.0014	0.4925	0.4141
22	n1_2	-0.2532		
23	n10_2	0.0332		
24	n15_2	-0.0158		
25	nAll_3	-0.0141		
26	n1_3	0.2901		
27	n3_3	-0.0033		
28	n5_3	-0.069		
29	n10_3	0.0579		
30	nAll_4	-0.0077		
31	n3_4	0.1371		
32	n15_4	-0.0044		
33	cInt1: BB intercept angle	0.0013	0.692	0.752
34	cInt2: SB1 intercept angle		-0.2691	-0.2939
35	cInt3: SB2 intercept angle	5e-04	0.53	0.5497
36	ccbar_1 (avg. BB-BB comp. center km)	-0.0126		
37	ccbar_2 (avg. SB1-SB1 comp. center km)	0.0097	1.2723	0.8758
38	icbar_2 (avg. customer-SB1 km)	0.0036		
39	icbar_3 (avg. customer-SB2 km)	0.0039		
40	icbar_4 (avg. customer-SS km)	-1e-04		
41	Prop. of sales in LP category (pre LP)	5.7842		
42	Received promotional email	0.0762		
43	Received promotional email for LP	1.3227		
44	Trips/month	0.0106	-27.3642	-25.6736
45	Total sales (pre LP)	2e-04	0.0066	0.0017
46	Total sales in LP category (pre LP)		0.0353	0.0136
47	Total qualified sales in LP category (pre LP)		-0.1075	-0.0428
48	Number of focal stores w/i 60km	0.0035	1.059	1.0569
49	Population Density (in sq miles)		0.0279	0.0251
50	Population Density x si		-0.0019	-0.0017
51	Population Density x ic1		-0.0019	-0.0018
52	Population Density x ic2		-6e-04	-9e-04
53	Population Density x ic4		0.0036	0.0036
54	Population Density x sc1		-2e-04	-3e-04
55	Population Density x sc2		0.0014	0.001
56	Population Density x sc3		-7e-04	
57	Population Density x sc4		-0.002	-0.0021
58	distinctSKU	-0.0055	-0.8616	
59	numItems	0.0019	-4.3895	-3.9053
60	distinctCategories	-0.0989		
61	poly(phat,1)		-617.5849	-382.4511
62	poly(phat,2)		4871.7517	4215.9458
63	poly(phat,3)		-4705.8528	-4272.9904
64	poly(phat,4)		-3849.4443	-3371.5503
65	poly(phat,5)		432.9157	467.4562
66	poly(phat,6)		4049.8262	3559.0346

Appendix 3: Additional Heatmaps

Figure 5 shows additional heatmaps of interest.

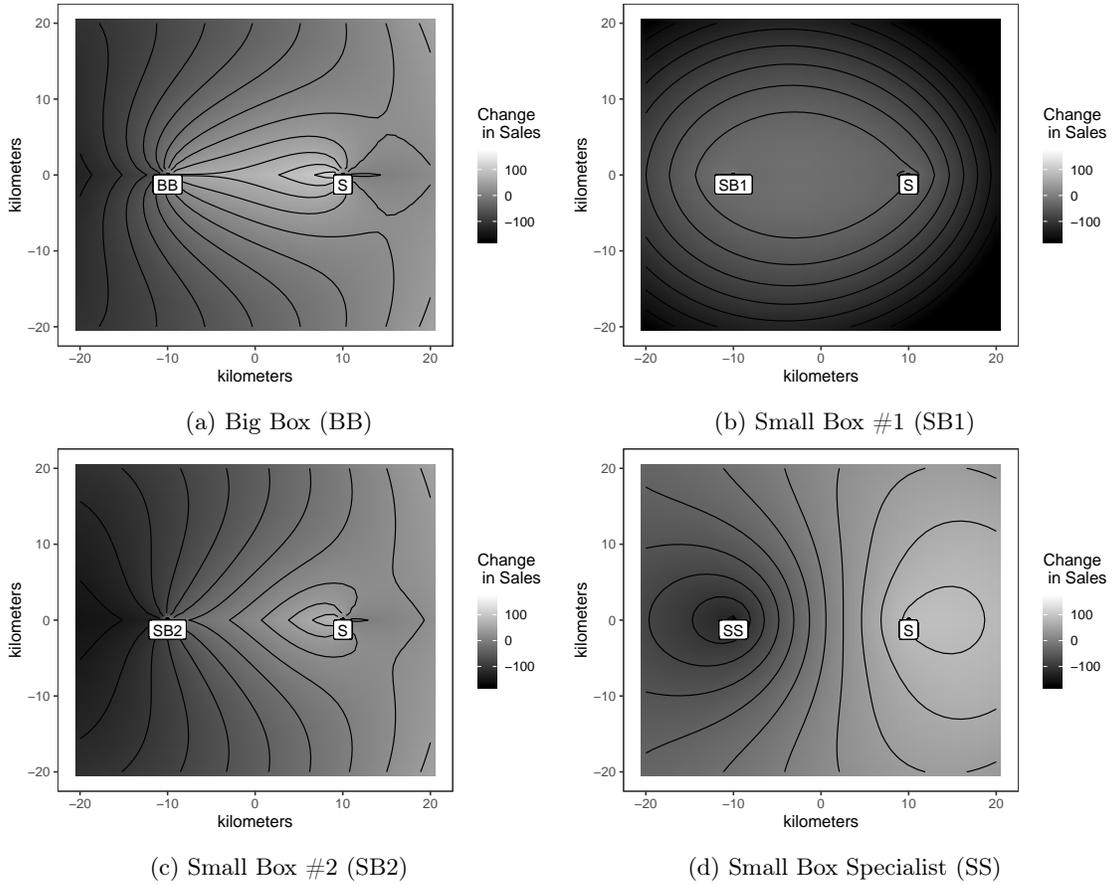


Figure 5: Change in Monthly Spend by Competitive Structure