UNIVERSIDAD DE CHILE
FACULTAD DE CIENCIAS FÍSICAS Y MATEMÁTICAS
DEPARTAMENTO DE INGENIERÍA INDUSTRIAL

MEASURING THE IMPACT OF SHELF POSITION ON A RETAILER'S SALES

MEMORIA PARA OPTAR AL TÍTULO DE INGENIERA CIVIL INDUSTRIAL

SOFÍA VICTORIA PONTIGO MARÍN

PROFESOR GUÍA:
ANDRÉS MUSALEM SAID

MIEMBROS DE LA COMISIÓN:
MARCELO OLIVARES ACUÑA
CARLOS NOTON NORAMBUENA

Este trabajo ha sido parcialmente financiado por:
CONICYT

SANTIAGO DE CHILE
2021

RESUMEN DE LA MEMORIA PARA OPTAR
AL TÍTULO DE INGENIERA CIVIL INDUSTRIAL
POR: **SOFÍA VICTORIA PONTIGO MARÍN**
FECHA: 2021
PROF. GUÍA: ANDRÉS MUSALEM SAID

MEASURING THE IMPACT OF SHELF POSITION ON A RETAILER'S SALES

Prior academic work in retailing suggests that a substantial fraction of purchasing decisions are made inside the store, where customers must choose from a large variety of products. In-store marketing tactics such as shelf space management then become particularly relevant as the number of products that could be displayed increases, while shelf space remains limited. Although extant research has shown a significant relationship between product location and sales, most studies are mainly focused on space elasticity and less is currently known about the effects of shelf position on sales. In addition, the potential reverse causality in the shelf position/sales relationship and the high costs of conducting controlled experiments have made it even more difficult to obtain managerial implications.

The purpose of this study is to contribute to this stream of research by examining the effects of shelf position on sales using a novel dataset about product position collected using a combination of robotics, internet of things and machine learning technologies. The data considers product position data for two stores of a major supermarket chain in Latin America and sales data provided by one of the retailer's vendors covering five product categories.

Using these data, we first perform a descriptive analysis to explore possible (non-causal) associations between variables. This analysis suggests that products located at the edges of an aisle sell between 37.0-69.6 % more than those located near the center, while products located closer to the checkout counters sell between 24.3-48 % more than those located farther away. The results also suggest that products located at the mid-height of the shelves sell 2.7 to 5.6 % higher sales than those placed on the bottom shelf.

We then propose a quasi-experimental approach to draw causal inferences about the impact of shelf position on sales. This requires us to rely on a semiautomatic method to detect changes in shelf position and hence identify suitable changes for a quasi-experimental approach, which we implement with a difference-in-differences analysis. In contrast with the results from the descriptive analysis, the quasi-experimental analysis finds no significant impact of the horizontal position of a product on its sales. However and consistent with the descriptive analysis, this analysis also suggests that by getting closer to the vertical center of a shelf, a product could increase its sales by 15.2-23.4 %, while increasing the height of a product can decrease its sales by 26.6 %.

RESUMEN DE LA MEMORIA PARA OPTAR
AL TÍTULO DE INGENIERA CIVIL INDUSTRIAL
POR: **SOFÍA VICTORIA PONTIGO MARÍN**
FECHA: 2021
PROF. GUÍA: ANDRÉS MUSALEM SAID

MEDICIÓN DEL IMPACTO DE LA POSICIÓN EN GÓNDOLA EN LAS VENTAS DE UN RETAILER

El trabajo académico previo en retail sugiere que una fracción sustancial de las decisiones de compra se toman dentro de la tienda, donde los clientes deben elegir entre una gran variedad de productos. Las tácticas de marketing dentro de la tienda, como la gestión del espacio en góndola, se vuelven particularmente relevantes a medida que aumenta la cantidad de productos que se pueden exhibir, mientras que el espacio en góndola sigue siendo limitado. Aunque la investigación existente ha demostrado una relación significativa entre la ubicación del producto y las ventas, la mayoría de los estudios se centran principalmente en la elasticidad espacial y actualmente se sabe menos sobre los efectos de la posición en las ventas. Además, la posible causalidad inversa en la relación posición/ventas y los altos costos de realizar experimentos controlados han hecho que sea aún más difícil obtener implicaciones gerenciales.

El propósito de este estudio es contribuir a esta corriente de investigación examinando los efectos de la posición en góndola en las ventas, utilizando un novedoso conjunto de datos sobre la posición del producto recopilado mediante una combinación de robótica, internet de las cosas y tecnologías de aprendizaje automático. Los datos consideran la posición de productos en dos tiendas de una importante cadena de supermercados en América Latina, y datos de ventas proporcionados por uno de los proveedores del retailer con productos pertenecientes a cinco categorías.

Utilizando estos datos, primero realizamos un análisis descriptivo para explorar posibles asociaciones (no causales) entre variables. Este análisis sugiere que los productos ubicados en los extremos de un pasillo venden entre un 37,0-69,6 % más que los ubicados cerca del centro, mientras que los productos ubicados más cerca de las cajas venden entre un 24,3-48 % más que los ubicados lejos. Los resultados también sugieren que los productos ubicados en la altura media de las góndolas venden entre un 2,7 y un 5,6 % más que los que se encuentran en nivel inferior.

Luego proponemos un enfoque cuasi-experimental para extraer inferencias causales sobre el impacto de la posición en las ventas. Este enfoque requiere de un método semiautomático para detectar cambios en la posición, y a partir de esto identificar los cambios adecuados para un enfoque cuasi-experimental, que implementamos con un análisis de diferencias-en-diferencias. En contraste con los resultados del análisis descriptivo, el análisis cuasi-experimental no encuentra un impacto significativo de la posición horizontal de un producto en sus ventas. Sin embargo, y en concordancia con el análisis descriptivo, este análisis también sugiere que al acercarse al centro vertical de una góndola un producto podría incrementar sus ventas en un 15,2-23,4 %, mientras que al aumentar la altura un producto puede disminuir sus ventas en un 26,6 %.

To my parents.

Acknowledgements

I wish to thank Zippedi for its contribution to this thesis. Zippedi not only provided the extensive shelf positioning database, but also managed the support of Unilever, who provided us the sales datasets. I'd also like to acknowledge partial funding support from the Fondecyt Regular 1181201 project.

I wish to express my sincere appreciation to my supervisor, Professor Andrés Musalem, for his unwavering support and guidance, and for the insightful suggestions through the development of this thesis.

Finally, I wish to acknowledge the support and great love of my parents, partner and female friends.

Contents

Tables

Ilustrations

Chapter 1

Introduction

1.1. Motivation

Extant research suggests that a substantial fraction of purchase decisions are made inside the store (e.g., Dagnoli (1987); Bell et al. (2010)), where customers must choose from a large variety of products. Strategies used to increase sales once the customer is at the store are known as in-store marketing. Within in-store marketing decisions, product allocation is an important topic particularly because as the number of products that could be displayed increases, shelf space becomes limited.

Product shelf allocation has been studied relying on two different approaches, depending on whether the effects of shelf space or shelf position on sales are measured. Regarding shelf space, prior research (e.g., Curhan (1972); Desmet and Renaudin (1998); Drèze et al. (1994)) analyzed space allocated to a product by estimating shelf-space elasticities using different functional forms. These researches reported an average space elasticity of approximately 0.2. Regarding shelf position, Drèze et al. (1994) analyzed shelf horizontal and vertical position effects on sales and concluded that by moving from the worst to the best horizontal and vertical positions, a brand could increase sales by 59 %.

Although previous studies have shown a positive relationship between product location and sales, evidence is still limited. The high costs of conducting controlled experiments in stores have made it even more difficult to address causality. Even in the case of the field experiments conducted by Drèze et al. (1994), the analysis relied on planograms, which not always reflect the actual position of products on the shelf. In particular, this thesis aims to study the effects of shelf position on sales. This question poses an important challenge for researchers because of a potential reverse causality in the relationship between shelf-position and sales: changes in shelf position may lead to changes in sales or vice versa. For instance, a manager anticipating increases in demand for highly seasonal products, may improve their shelf position just before the expected rise in sales. Similarly, a manager could allocate the best selling products to more visible aisles or locations that receive higher levels of traffic. In addition to reverse causality, both shelf space and shelf position are negotiated between retailer and manufacturers, who are willing to pay fees to increase facings or allocate products in more visible positions.

Another factor that limits the ability to conduct research on the impact of shelf position on sales is the availability of positioning data. In fact, some researchers have relied on multiple store visits to measure and record the position of products on the shelves (e.g., Frank and Massy (1970)). Nonetheless, new research possibilities emerge as the retail industry adopts new monitoring and management support technologies, such as robotics, internet of things and artificial intelligence. The store-execution data generated through these technologies are now available to complement previous studies. This thesis presents a novel data source provided by Zippedi that has promising implications for in-store marketing studies. Zipeddi uses robotics to collect daily positioning data for products in stores, obtaining the horizontal and vertical position of a product on the shelf. Using this data, we studied the effects of shelf position on sales in two stores from a major supermarket chain in Latin America.

We will first conduct a descriptive (i.e., non-causal) analysis to measure associations between shelf position and sales patterns for products in five categories, controlling for changes in prices and differences across stores, categories, brands and months. Interestingly, (Frank and Massy, 1970) also relied on an observational study, but used cross-sectional (as opposed to time-series) data and for a single product category. Moreover, (Frank and Massy, 1970) focused on the impact of the number of facings and vertical position on sales, without studying the role of horizontal position. Our descriptive analysis suggests that products located at the edges of an aisle sell between 37.0-69.6 % more that those located near the center. Additionally, products located closer to the checkout counters sell between 24.3-48 % more than those located farther away. Regarding the vertical position, we find that those located in central positions (between 0.75 to 1.25 meters from the bottom shelf) exhibit between 2.7 to 5.6 % higher sales than those placed on the bottom shelf. Products located between 1.50 and 1.75 meters sell 12.3 % to 18.1 % less than those located at the bottom of the shelves. However, these results have important limitations that prevent us from drawing conclusions about causality.

Given these limitations, we complement this analysis with the use of a quasi-experimental design. Under this approach we use changepoint algorithms to identify products that exhibited changes in their horizontal and/or vertical position at a given store. We then analyze the time series of sales for that product before and after the change and compare it to the pattern of sales for the same product in a control store where the position of the product on the shelf did not change. The comparison is performed using a differences-in-differences (DID) analysis. Although no evidence was found for causal relations between horizontal shelf position changes and sales, there were significant positive effects of changing the vertical shelf position of a product: by getting closer to the center of the shelve, a product could increase its sales in 15.2-23.4 %. Also, increasing the height of a product can decrease its sales by 26.6 %. Hence, the most detrimental vertical change is to get farther from the center of the shelves by moving a product upwards.

The remainder of the thesis is structured as follows. Chapter 2 reviews previous work on product shelf allocation studies and positions our contribution to the field in relation to the extant literature. In Chapter 3, we describe the dataset which includes the shelf position information provided by Zippedi in addition to sales data. These data will be used to estimate the associations and relationships between shelf position and sales. Chapter 4 then presents a descriptive analysis to measure the association between shelf position and sales. This section also discusses the limitations of this method. Chapter 5 then introduces the quasi-experimental approach. In particular, we implement a changepoint analysis to identify products that exhibited significant variations in their locations at a given store and among them we select those for which we can rely on the data from the other store to define a control group. We then evaluate the effects of such changes in sales through a DID estimator. Finally, we conclude with a summary of findings, a discussion of the limitations of this research and identifying avenues for future research.

1.2. Objectives

1.2.1. General objective

Estimate the effect of shelf position on the sales of a retailer's products.

1.2.2. Specific objectives

i) Determine whether certain horizontal or vertical positions are associated with greater sales levels.

ii) Identify changes in the horizontal and/or vertical shelf position of products in each store.

iii) Estimate the causal impact of these location changes on sales.

iv) Compare the causal impacts of the different location changes using regression methods.

1.3. Scope of the study

To estimate the effects of shelf position on sales, we implemented a quasi-experimental design and a difference-in-difference analysis using panel data. Our study covers only two stores from a supermarket chain located in the same geographical area of Santiago, Chile. Although we have positioning data for all products displayed, our sales data correspond to products manufactured by one company, hereafter referred to as the vendor. In addition, this thesis covers a one-year period from October 1, 2018 to September 30, 2019.

Chapter 2

Literature review

Shelf management studies are focused on developing evidence and insights to help managers increase sales through a better allocation of products to shelf positions. In general, these studies can be classified depending on the methodology used and their level of analysis. From a methodological perspective, some studies are based on historical observational data (Frank and Massy, 1970; Desmet and Renaudin, 1998), and others are based on field experiments (Chevalier, 1975; Curhan, 1972; Drèze et al., 1994). Within the latter class, most of the experimental studies analyze the total space allocated to a product category, and do not take into account special promotional locations inside the stores. On the contrary, Chevalier (1975) incorporated promotional displays into the analysis. In terms of the different levels of analysis, some studies carry out a product category level analysis, a stock keep unit (SKU) level analysis, or both. In this literature review, we describe studies which rely on the aforementioned approaches.

One of the first studies aiming to explain the effect of shelf management on supermarket sales was carried out by Frank and Massy (1970). The main purpose of this investigation was to estimate the effects of shelf position and space on sales for a frequently purchased, branded grocery product as well as the magnitude of this impact across different stores, brands, or package sizes. They collected data covering 30 different stores, including relevant information for 7 different product brands such as weekly sales, price, advertising, discounts, and package size. In terms of shelf-related information, the main variables of this study are the number of facings and the shelf level.[1] The authors found that within the range of 5-10 facings, adding an additional facing yields a significant increase in weekly sales in high, but not in low volume stores regardless of the package size. At the same time, the authors show that the effect of varying shelf level is modest, if any. Finally, the authors did not find a significant interaction between the effects of shelf level and the number of facings on sales.

The authors acknowledged the difficulty in addressing causality and attempted to lessen the bias produced by confounding variables. According to them, this difficulty arises because their study was based solely on observational data. In addition, the authors only analyzed a single product category.

Curhan (1972) revisited this issue adding a new modeling and methodological perspective.

[1] This variable measures whether or not the brand is allocated at 2, 3, 4 or more shelve levels from the floor.

In terms of the model, he proposed and tested the relationship between space elasticity and physical properties of the products. From a methodological perspective, Curhan (1972) filled several gaps left by Frank and Massy (1970). First, while Frank and Massy (1970) studied seven brands within a single grocery product category, Curhan (1972) included instead a large number of grocery products from both private and national brands. The importance of incorporating a greater variety of products is that the impact of position on sales may vary from one product to another. Second, Curhan (1972) conducted a controlled experiment over 28 supermarkets belonging to a regional chain, where 4 of them were used as a treatment group and 24 were used as a control group. In this experiment, the space allocated to different products was either increased or decreased at the treatment store. Third, while Frank and Massy incorporated a brand variable to mitigate selection biases, Curhan measured the influence of a brand on customer purchasing decisions through complementary variables such as brand type, market share and rate of sales. Furthermore, the researcher managed to incorporate supplementary variables, such as shelf capacity, merchandise variety, availability of substitutes and repurchase frequency. As a result, space elasticity averaged 0.212 for all items, showing a positive relationship between shelf space and unit sales for some products (Curhan, 1972). According to the author, for most products the impact of shelf space changes on unit sales is very small in comparison with the effects of other variables because there are unrecognized and uncontrolled interrelationships among item substitutes and complements that affect item unit sales. Despite not being able to fully explain the factors that influence space elasticity, Curhan's experiment contributed to a better understanding of the complexity of the purchasing process. In fact, Curhan (1972) was aware of the need to complete the study by incorporating "a good instrument to measure impulse buying" (p. 411) as well as more discriminating variables such as display location to better explain variations in space-elasticity of sales across stores.

In particular, display location is one of the many strategies that supermarkets rely on to promote products, and thus increase sales. Later, Chevalier (1975) studied the increase in sales due to in-store displays, which are generally set up in piles at the end of an aisle and are often price reduced. In accordance to Chevalier (1975), consumers tend to view them as special bargains and often buy a product which they had no previous intention of buying, therefore the impact of both impulsive buying and promotion are approached through these initiatives. Chevalier (1975) tested the impact of two levels of price cut on sales when a product is displayed: a threshold level and a deep level, with 6 % and 12 % reduction respectively. The experiment was conducted in four stores and includes 8 categories where 2 brands were tested. Through this experiment, Chevalier (1975) attempted to understand the impact of display for different types of products, by incorporating variables such as product lifecycle, competitive structure within a category and competitive position of a brand.

By means of his experiment results, Chevalier (1975) confirmed Curhan's intuition that promotions for substitutes or complements may affect the sales of a particular product. He found that when products have similar positions and no one has a clear market share advantage, the impact of promotional display is higher than in competitive structures where there is a clear leader. In addition, he also concluded that a product which sells twice as much as another when placed on a regular shelf layout should also sell twice as much when placed on promotional display (Chevalier, 1975).

All the studies discussed so far have recognized that factors such as brand reputation or advertising level influence the degree to which a product's shelf position or space can influence sales. At the same time, they have also recognized the difficulty in understanding how interactions between a product and its substitutes or complements can drive sales when a product is positioned in a particular place on the shelf.

Despite the importance of in-store and out-of-store factors, there are other forces that were not considered until the early nineties. Drèze et al. (1994) raised important issues for shelf management and space elasticity studies. Drèze et al. (1994) main critique to previous studies stated that although the elasticity estimated before might seem sizable to a manufacturer, it is not clear how significant it is to a retailer. Thus, these researchers realized that when it comes to the shelf space problem, there are two new opposing forces: the manufacturer's goal is tipically to maximize brand profits, while retailers seek to maximize profits of an entire product category.

In order to fully explain the effect of product position on sales, the authors designed two experiments with a wider set of brands and stores than previous studies and performing both a category and brand level analysis. In terms of the category level analysis, Drèze et al. (1994) implemented two different approaches for product display. Under one of them, product allocations were customized to meet historical sales, resulting in changes in sales and profits ranging from -2 % to 8 %. Under the second approach, planograms were rearranged either to facilitate cross-category merchandising or to increase/decrease ease of shopping, producing changes in sales of 5-6 %. In terms of brand level analysis, the authors tested a model in which, unlike previous studies, space and location elasticity were simultaneously incorporated. Horizontal and vertical position were modeled through second and third grade polynomials, respectively, while space elasticity was modeled using the Gompertz growth model (Drèze et al., 1994). Drèze et al. (1994) empirical results showed that whereas adding or removing a facing would not affect sales for most items, moving them from the worst to the best position could increase sales in 59 %, indicating that location parameters were more important than space.

Drèze et al. (1994) contributed greatly to reaching a deeper understanding of the effects of shelf location and space. However, the model proposed by the authors did not account for heterogeneity in brand and store level price elasticities, nor did it account for seasonality trends. More importantly, they rely on data collected from a set of planograms instead of directly from the store. Although using planograms could be cost effective compared to store audits, the information contained in them may differ from actual store execution. For example, when a product is out of stock, a retailer may choose to temporarily fill its shelf space with an available product. While this was an issue in the past due to the high costs associated with store audits, nowadays there are new technologies that can be used to monitor each aisle of the store. In this regard, our study draws on information collected by robots which travel through the aisles of a store collecting shelf execution information on a daily basis.

Subsequent research was useful to complement Dréze et al. (1994)'s findings, especially at the category level. One of this studies was carried out by Desmet and Renaudin (1998), who used data from a single retail chain to estimate space elasticities at the product category level. In addition, in order to test the relationship between impulse buying and space elasticity discussed before, they compared highly space-elastic product categories against product categories characterized by a high level of impulse buying (Desmet and Renaudin, 1998). These authors found an average elasticity of 0.21, along with a large variation from one product category to another, and a positive relationship between space elasticity and impulse buying (Desmet and Renaudin, 1998).

Besides complementing Dréze's category analysis, these results raise a question regarding brand level analysis: do the effects of shelf position on sales vary across different categories? This question is relevant both for manufacturers and retailers, which they usually bargain about the space and position of products on the shelf.

So far, we have presented studies which focus on explaining the effects of shelf position or total shelf space on sales, without necessary clarifying the underlying mechanisms that give rise to this relationship. Chandon and Hutchinson (2009) incorporated the consumer behavior perspective in order to understand the underlying process by which shelf management influences consumer choice, and thus, sales. The authors evaluated a series of in-store factors as drivers of attention and evaluation at the point of purchase, through an eye-tracking experiment in which participants viewed a planogram of two product categories on a screen under a particular shopping goal. Chandon and Hutchinson (2009) experiment measured attention (noting, reexamination and attention) as well as evaluation (recall, consideration, choice and evaluation) dependent variables. In-store factors were incorporated into this study as independent variables, which have been considered in the previous literature, such as number of facings, horizontal shelf position, vertical shelf position, and price.

The authors used a regression analysis to estimate the effects of in-store independent variables on attention and evaluation measures, and then a mediation analysis to examine whether the effects of in-store factors on evaluation are also entirely mediated by attention and therefore are effective even if they have no direct effect on evaluation (Chandon and Hutchinson, 2009). This study concluded that the number of shelf facings strongly influences visual attention and, through attention, brand evaluation. In addition, they found that the position of facings strongly influences attention, but that attention gains from shelf position do not always improve evaluation (Chandon and Hutchinson, 2009). In terms of shelf position, this research showed that locating a product on the left- or right-hand side of the shelf made no difference to either attention or evaluation, whereas positioning the brand on the top shelves had a positive influence on both measures. Although these results help us to understand the mechanisms by which shelf position affect attention and evaluation, Chandon et al's study (2009) was carried out just on two product categories (soap and pain relievers). In addition, this experiment was implemented in a laboratory (as opposed to in the field) and it was based on hypothetical purchasing decisions.

Another important factor for studying how the position of a product affects its sales is a product's location within the store. Some researchers have studied supermarket shopping paths to determine the store sections that are visited the most and the most common paths employed by customers when visiting a store (Hui et al., 2007; Larson et al., 2005; Pinna and Seiler, 2017). The findings in Larson et al. (2005) are particularly relevant to our study. The researchers used clustering techniques to study how shoppers travel through the store using

Radio Frequency Identification (RFID) tags placed on supermarkets shopping carts. After identifying three groups of shoppers depending on their travel times (from 2 to 10 minutes, from 10 to 17 minutes, and from 17 minutes to an hour), they applied a k-medoids clustering approach to identify different shopping paths within each group. The main finding of this study is that, independent of the group to which they belong, most customers follow a default path which begins at the entry of the store and continues through the racetrack around the aisles. The authors find that most shoppers tend to travel only through select aisles, and that shopping trips that display extensive aisle travel tend to do so via short excursions into and out of the aisles rather than traversing their entire length (Larson et al., 2005). These results are relevant for our study, because they can help us to identify locations within an aisle that may receive greater traffic and potentially lead to greater sales.

Among all of the studies reviewed so far, only two of them address the specific problem discussed in this thesis: the impact of product shelf position on sales (i.e., Frank and Massy (1970); Drèze et al. (1994)). The first is the observational study conducted by Frank and Massy (1970). As discussed earlier, one of the limitations of this study is that it only considers the effects of vertical position. Another limitation is that this study relies on cross-sectional data, which is unable to account for seasonality or exploit changes in product location for the same product. In contrast, our study analyzes the impact of both horizontal and vertical shelf position using one year of data (as opposed to a cross-section). The second study is the field experiment conducted by Drèze et al. (1994), in which space elasticity and shelf position are jointly analyzed. This research has the standard advantages of randomized experiments, however the integrity of the assigned treatment relies on planogram compliance. On the contrary, our analysis uses the actual daily position of products at the stores under study (see Table 2.1).

In general, the high costs of conducting controlled experiments and the low availability of positioning data for conducting observational studies have prevented the progress of research studying the impact of product shelf position on sales. In this research, we rely on a new dataset which tracks product shelf position on a daily basis. Our main contribution is that we rely on a semiautomatic method to detect changes in position and implement a quasi-experimental approach, i.e. without intervening the store. This allows us to use the observational data to perform a causal analysis using a DID approach. In summary, neither the data nor our methodological approach have been used in the extant research on this topic.

Table 2.1: Shelf management research.

Article	Research question	Research approach	Shelf positioning measurement approach	Dependent variable	Predictors of interest
Frank and Massy (1970)	Effects of shelf space and shelf position on sales.	Observational data.	Stores were monitored twice a week.	Normal sales.	Number of facings and vertical shelf position.
Curhan (1972)	Impact of shelf space changes on unit sales.	Field experiment.	Shelf management info system.	Space elasticity.	Product's physical properties and merchandising characteristics.
Drèze et al. (1994)	Effects of shelf space and shelf position on sales.	Field experiment.	Auditors monitored the integrity of the planograms bi-weekly.	Natural logarithm of sales.	Area occupied by a product, and the coordinates relative to the lower left corner of the shelves.
Desmet and Renaudin (1998)	Effects of shelf space on product category sales.	Observational data.	Pooled database.	Turnover share of the product category.	Share of allocated space to the product category.
Chandon and Hutchinson (2009)	Effects of the number and position of shelf facing on brand attention and evaluation.	Lab experiment.	Eye tracking.	Choice, consideration and recall.	Number of facings, and horizontal and vertical shelf position.
This study	Effects of shelf position on sales.	Observational data.	Shelf position measured daily using robotics.	Natural logarithm of unit sales.	Horizontal and vertical location on the shelf, and distance from the center of the aisle and from the middle point of the shelves.

Chapter 3

Empirical context and data description

The data for this study were obtained for two stores of a large supermarket chain in Latin America. Both stores are located in the same geographical area within the same city. They were selected to ensure that 1) customers of both stores present similar socioeconomic characteristics, and 2) their locations are far enough from each other so changes in one store are less likely to affect the sales of the other. Data used in this investigation comes from two different sources and covers a one-year period from October 1st, 2018 to September 30th, 2019. These data consider the position of products on the shelf and the sales for a subset of these products (i.e., those from one of the vendors of the supermarket chain). In this chapter, we describe both datasets and provide summary statistics for each of them.

3.1. Positioning data

Shelf location data was provided by Zippedi, which uses robotics and machine learning to collect and process store execution information with the aim of helping managers to identify stockouts and measure planogram compliance (Figure 3.1(a)). Every night at 11:00 pm, the Zippedi robot begins its journey through the store. The robot is programmed to navigate automatically across the aisles capturing the state of the supermarket at the end of the day (Figure 3.2). After 7 hours of operation, it has traveled about 4 km and scanned more than 25,000 products. This brand new technology creates a digital model of the store that helps managers to optimize in-store shelf execution. Since the robot detects stockouts, missing price tags and product mislocation, its operation has the potential to generate improvements in on-shelf availability (6-10 %), price integrity and planogram compliance.[1]

[1] https://www.zippedi.com.

(a) Photographed shelf. Green labels indicate products in stock, while red labels indicate products out of stock.

(b) Photographed price tag including SKU, EAN, product description and price.

Figure 3.1: Digital photograph and processed images.

The data collection process takes place as the robot photographs the shelves capturing the horizontal and vertical position of each price tag displayed. Regarding the horizontal position, every time the robot enters an aisle, it resets to zero and starts counting meters from left to right, so that when a price tag is found, it knows where it is located along the horizontal axis. Vertical position is measured starting from the first shelf level (from bottom to top). In addition, using artificial intelligence the robot is able to recognize relevant information from the digital images, such as the product's SKU, EAN code, price, and description (Figure 3.1(b)). Since the aisles are generally formed by two sets of shelves facing each other, each part of the aisle is labeled in terms of its cardinal orientation (N, S, E, O). Therefore, each shelf is identified by a code representing the aisle (i.e. '22B') and the orientation (i.e. 'E'). Finally, daily sessions of the robot are saved along with the date in which the measurements were performed.

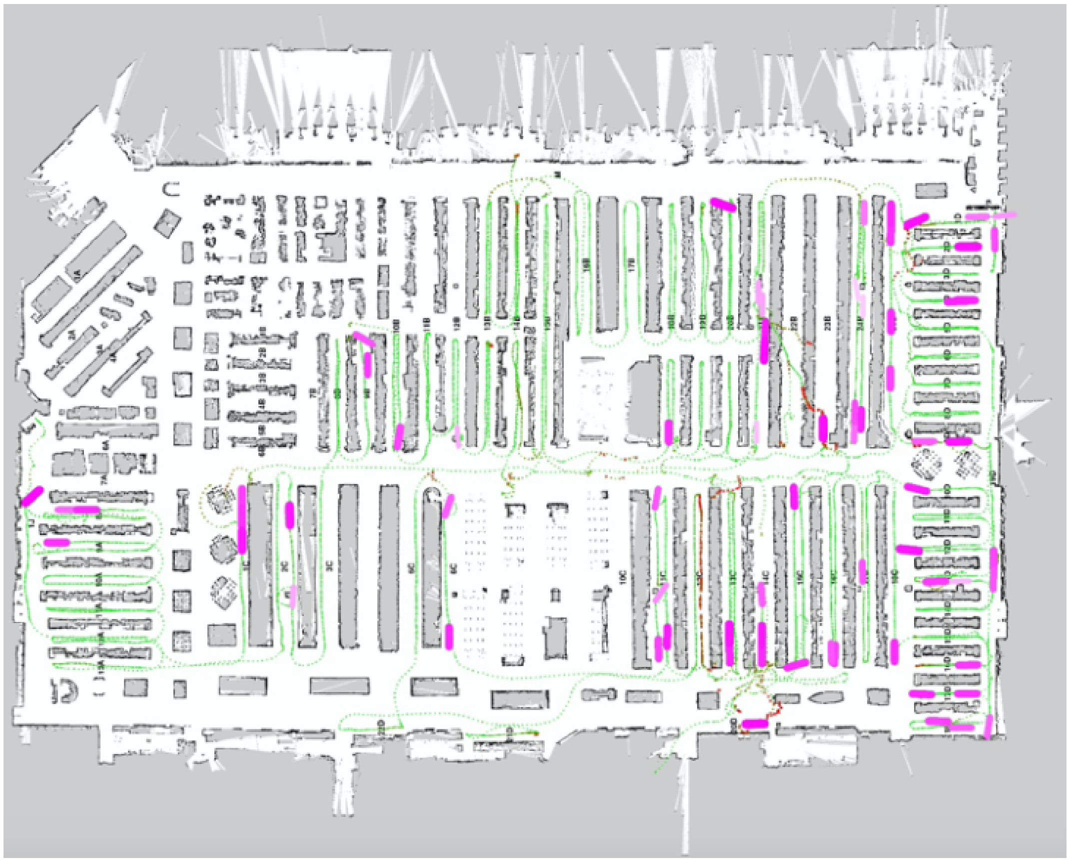

Figure 3.2: Example of a robot's path inside a store.

In summary, the positioning data contain the following information: SKU, EAN code, price, description, aisle, vertical and horizontal location, and date (see Table 3.1 for an example).

Table 3.1: Positioning data.

SKU	EAN	Price	Aisle	Horizontal location	Vertical location	Date
1647891	7805000182397	1699	22BE	5.13	1.67	2018-10-01
1647891	7805000182397	1699	22BE	5.15	1.33	2018-10-02
1647891	7805000182397	1699	22BE	5.12	1.68	2018-10-02
1647891	7805000182397	1699	22BE	5.17	1.33	2018-10-03
1647891	7805000182397	1699	22BE	5.23	1.34	2018-10-04
1647891	7805000182397	1699	22BE	5.18	1.69	2018-10-04

We now provide summary statistics for the shelf position variables. Between October 1, 2018 and October 15, 2019, the robot completed 259 sessions, corresponding to 259 days of operation. Over the entire period, it read 796 different SKUs belonging to the vendor brands in both stores, and on average, it read 608 different SKUs each month. Henceforth, we refer to stores as "store A" and "store B".

Considering first store A, the horizontal position of products varies from 0.0 to 28.0 meters, while their vertical position varies from 0.0 to 2.0 meters. A map of vertical against horizontal position is shown in Figure 3.3. Most observations are concentrated in aisles with lengths no longer that 14 meters. Larger sections that are located in the store perimeter are not usually visited by the robot, leading to a lower density. In terms of the vertical position, this figure clearly shows greater density of observations around the different heights of the shelves.

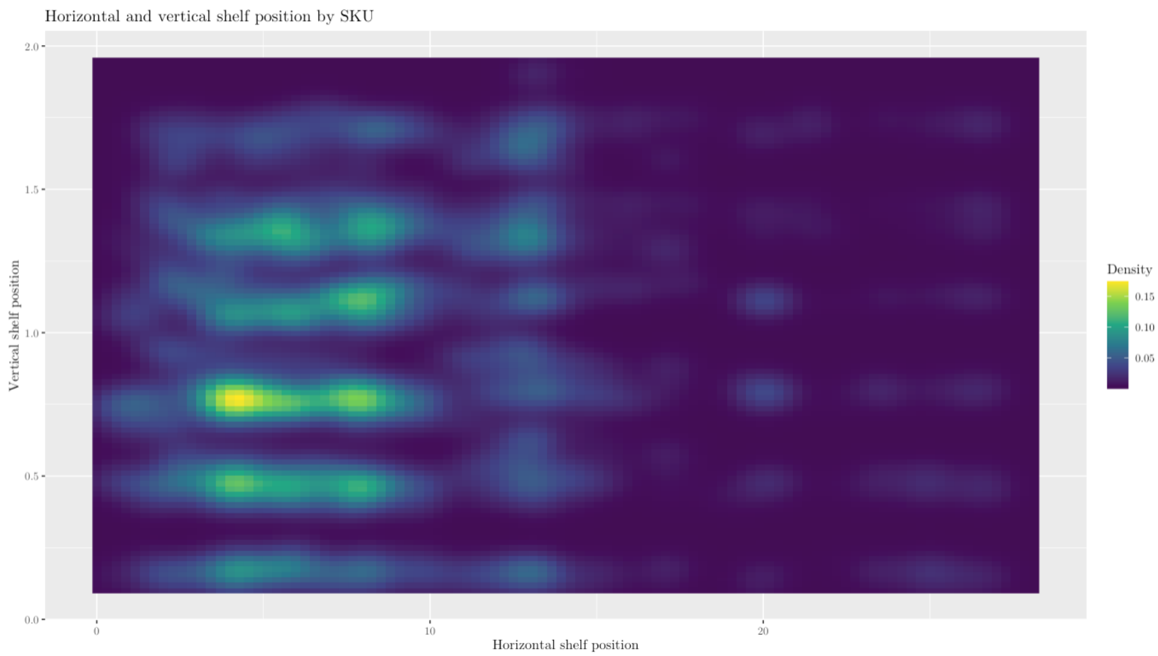

Figure 3.3: Map of shelf position in store A.

Considering next store B, the horizontal position varies from 0.0 to 22.0 meters, while their vertical position varies from 0.0 to 2.0 meters. Figure 3.4 shows the map of vertical against horizontal position of the store. Although observations in store B are also concentrated in smaller aisles, the robot frequently visited perimeter areas, so the density in these areas is larger than those in store A.

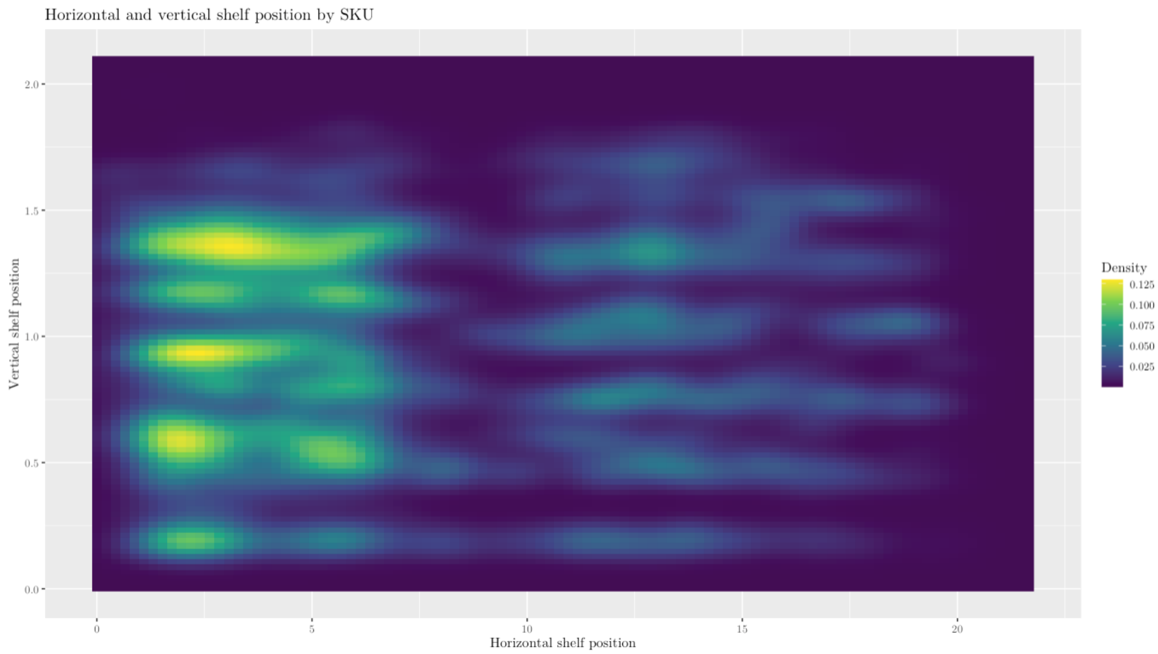

Figure 3.4: Map of shelf position in store B.

3.2. Sales data

The second data source includes sales information provided by the vendor. This dataset contains daily unit and dollar sales as well as brand, product category, market, submarket and group (Table 3.2). Note that product category, market, submarket and group are the way that the vendor classifies its products. This classification goes from the most general to the most specific level.

Table 3.2: Sales data.

Category	Brand	Market	Submarket	Group	Sales (Un)	Sales ($)
Cleaners	Quix	Dishwasher detergent	By-hand dishwasher	Liquid	11	9243.7
Cleaners	Quix	Dishwasher detergent	By-hand dishwasher	Liquid	6	5042.02
Cleaners	Quix	Dishwasher detergent	By-hand dishwasher	Liquid	6	5042.02
Cleaners	Quix	Dishwasher detergent	By-hand dishwasher	Liquid	8	6722.69
Cleaners	Quix	Dishwasher detergent	By-hand dishwasher	Liquid	4	3361.34
Cleaners	Quix	Dishwasher detergent	By-hand dishwasher	Liquid	4	3361.34

We now provide summary statistics regarding sales data. the vendor sells five product categories: Personal Care, Cleaners, Sauces and Dressings, Grocery and Dairy. In both stores, the first two offer the greatest variety of products and brands (Table 3.3), and sell the most (Figure B.1).

The average revenues of store A almost doubled those of store B: $203.417.066 and

Table 3.3: Category description.

Product category	Store	Brand	Market	SubMarket	Group	SKU
Grocery	A	7	6	10	16	99
	B	7	6	10	16	99
Sauces and Dressings	A	4	1	5	10	61
	B	4	1	5	10	62
Dairy	A	7	2	7	15	48
	B	7	2	7	15	48
Cleaners	A	10	6	14	23	164
	B	10	6	14	23	162
Personal Care	A	17	13	24	46	479
	B	17	13	24	46	479

$106.475.427 respectively. Figure 3.5 shows monthly revenues in both stores, indicating comparable increases in March and December and a considerable decline in February.

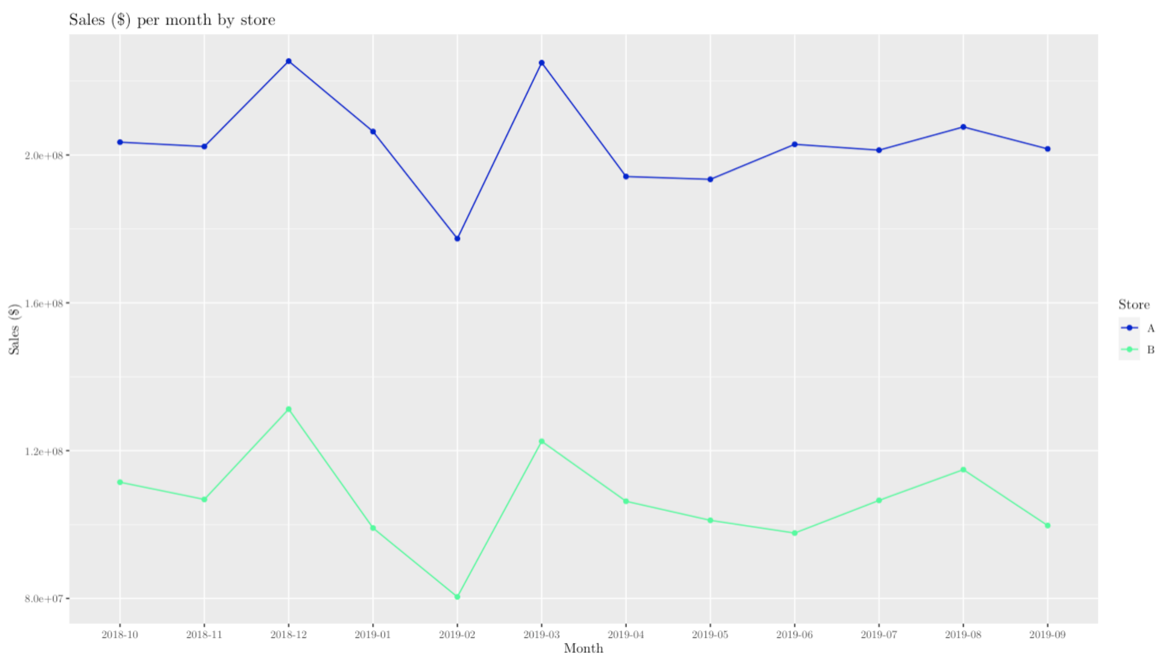

Figure 3.5: Monthly sales by store.

Figure 3.6 shows the sales against the horizontal position for each product category and store. Note that the horizontal position represents the fraction of the total length of the aisle (measured from its origin) where the price tag for a given SKU is read. These figures show where most products are located in each product category: Cleaners, Grocery, Personal Care and Sauces and Dressings product categories are displayed all over the shelves. Unlike these, dairy products are not placed across store A's or store B's whole aisles. In fact, in store A we observe a concentration of products at the beginning of the aisle and near the end, while a more uniform pattern is observed for store B.

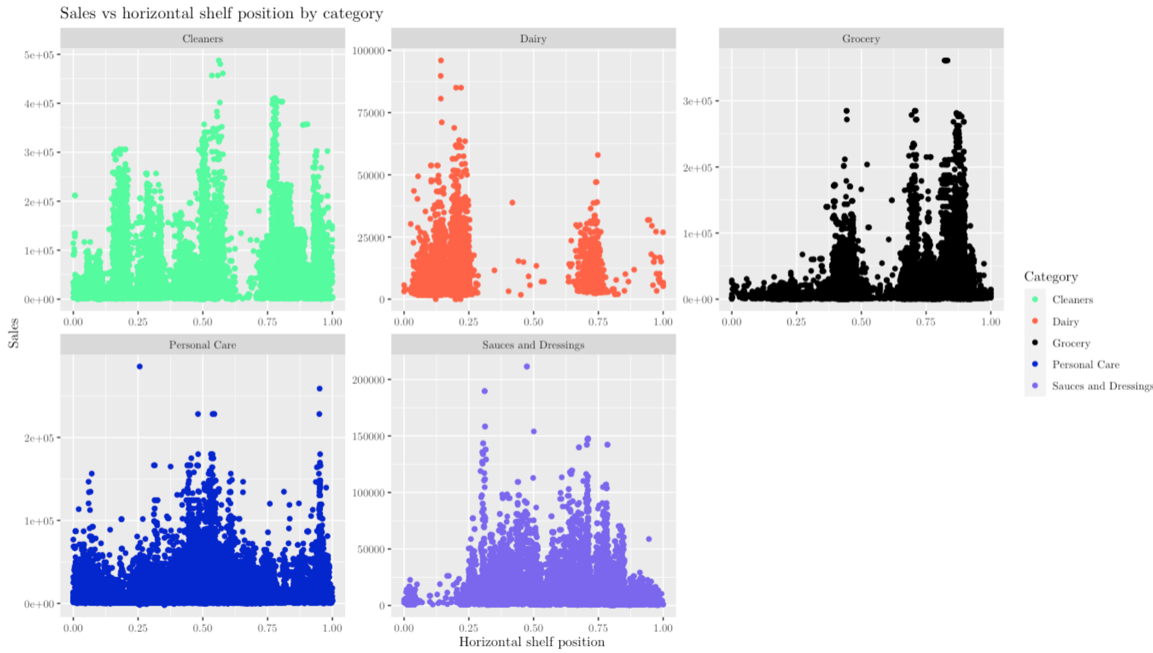

(a) Store A.

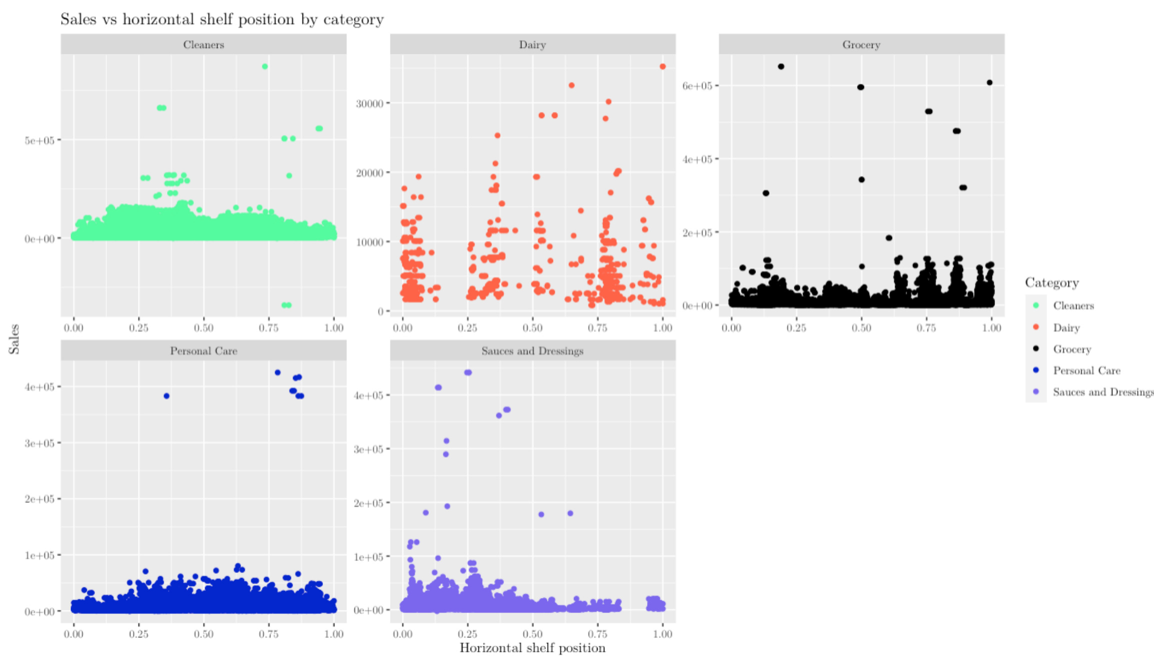

(b) Store B.

Figure 3.6: Sales against horizontal position by product category.

Figure 3.6 shows non-linear relationships between sales and horizontal shelf position in all categories. Nonetheless, this relation is not uniform across categories, and neither it is across stores. Regarding product categories, within each product category different curves may be noted for different sections of the aisles. This could mean that for a given section of the aisle, sales are smaller in the edges of said section. Although this result is more clear in store A, the same non-linearity is observed in store B. Regarding the comparison between stores, we observe that curves in store B are flatter that those in store A.

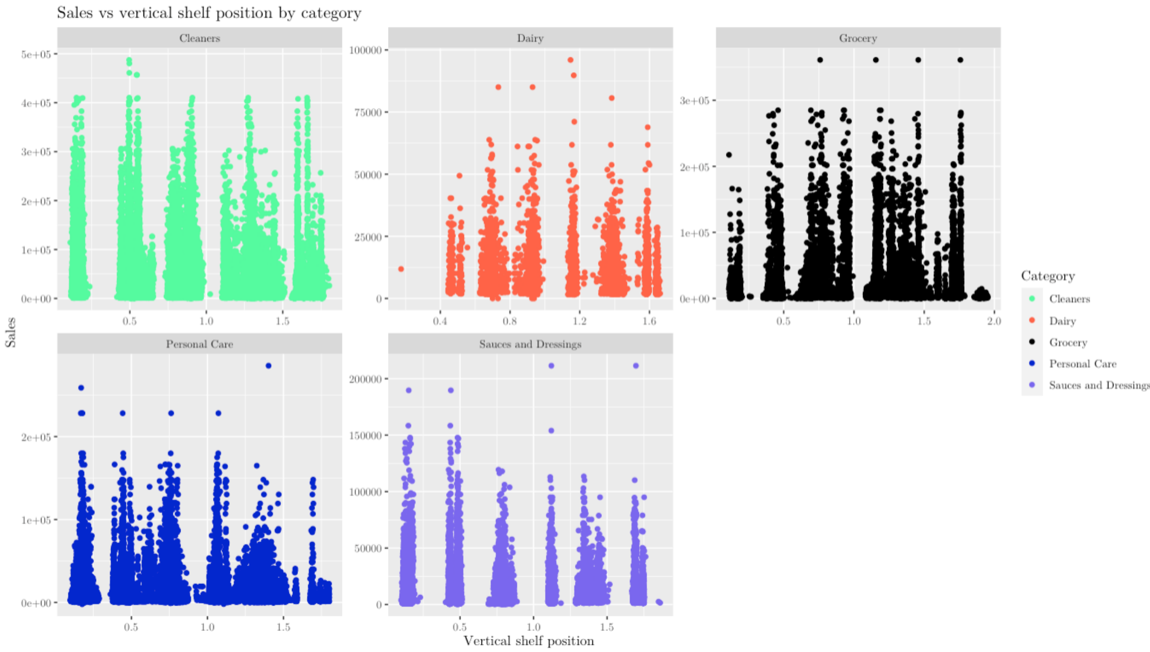

(a) Store A.

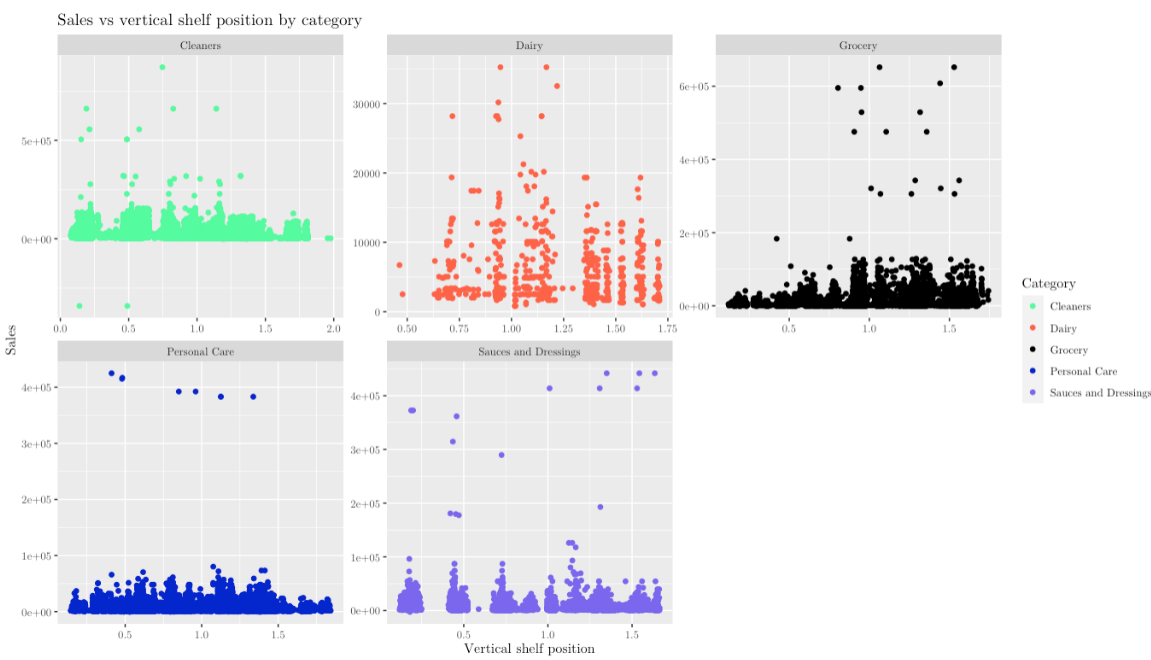

(b) Store B.

Figure 3.7: Sales against vertical position by product category.

Sales against vertical position are shown in Figure 3.7. As for the horizontal location, a non-linear relation between shelf position and sales is observed. For most categories it is difficult to infer differences in sales due to product allocation, since the curves are similar across different shelf levels. However, in store A's Sauces and Dressings and Personal Care products, slightly decreases in sales are observed as the height increases. In store B's Grocery products, a mild increase in sales is observed as the height increases.

In general, the non-linear relation between sales and shelf position is consistent across stores and product categories. In addition, Figure 3.6 and Figure 3.7 reveal the distribution of products within the horizontal and vertical axes of the aisles. Along the horizontal axis, we observe that the vendor products are located across the whole length of the aisle for most product categories, while only in some segments for the remaining product categories (e.g., dairy). Regarding vertical axis, products are located in all levels.

Chapter 4

Descriptive analysis

This section presents a descriptive analysis of the relationship between sales and shelf position. We first formulate the statistical specification for modeling sales and then present the results with their corresponding limitations.

4.1. Modeling sales

As a first step towards measuring the relationship between sales and shelf position, a model of sales is formulated as follows:

$$\begin{aligned}\ln(1 + SALES_{ist}) = \beta_0 + \beta_c CENTRAL_{ist} + \sum_{k=1}^{8} \beta_{sk} SHELFLEVEL_{kist} + \\ \beta_{ac} AISLEPART_{ist} \cdot CENTRAL_{ist} + \\ \beta_{ap} AISLEPART_{ist} \cdot PERIMETER_{ist} + \\ \beta_{amc} AISLEMIDDLE^2_{ist} \cdot CENTRAL_{ist} + \\ \beta_{amp} AISLEMIDDLE^2_{ist} \cdot PERIMETER_{ist} + \\ \beta_p PRICE_{ist} + \beta_m MONTH_t + \beta_s STORE_s + \beta_b BRAND_i + \\ \beta_c CATEGORIZATION_i + \epsilon_{ist}\end{aligned} \tag{4.1}$$

where $SALES_{ist}$ denotes the unit sales of product i at store s on date t. Unit sales are chosen as the dependent variable instead of revenues, because of the correlation between the latter and price, which is one of the control variables in this descriptive analysis. In addition, due to the significant differences between the two stores under analysis, we implement a logarithmic transformation of unitary sales.

The main covariates of interest are those related to product location on the shelf. Product location considers three components: aisle location, horizontal shelf position and vertical shelf position. Regarding aisle location, the aisles visited by the robot can be classified into two groups: those along the perimeter of the store, and those in the center (Figure 3.2). This classification is analogous to the one used by Larson et al. (2005) in their study of shopping paths over the outer ring and the central aisle zones of a store. This distinction is useful since the findings in Larson et al. (2005) suggest that customers might spend more time in zones closer to the outer ring. Thus, let $CENTRAL_{ist} = 1$ denote whether the aisle of store

s in which product i was detected in period t is located in the center of the store, otherwise $CENTRAL_{ist} = 0$. We also define, $PERIMETER_{ist} = 1 - CENTRAL_{ist}$, which indicates whether the aisle is located in the store perimeter.

The horizontal location of product i at store s in period t is incorporated via the $AISLEPART_{ist}$ covariate, which represents the distance from the beginning of the aisle.[1] Accordingly, this variable is defined as the fraction of the total length of the aisle (measured from its origin) where the price tag for a given SKU is read. The interactions between $AISLEPART_{ist}$ and the indicators C_{ist} and P_{ist} allow us to compare the effect of horizontal position across the two aisle locations. In addition, since places at the end of the aisle could be more attractive than those in the center [2], we also include the $AISLEMIDDLE_{ist}$ variable, which represents the squared distance from the center of the aisle. This variable is also interacted with the indicators C_{ist} and P_{ist}.

In terms of the vertical position, we use a set of eight dummy variables to flexibly estimate the effect of vertical position on sales. Denoting by z_{ist} the vertical position of product i at store s in period t, these dummy variables are defined in 0.25 meter increments as follows: $SHELFLEVEL_{1ist} = 1$ if $z_{ist} \in [0.0, 0.25)$ and equal to 0, otherwise; $SHELFLEVEL_{2ist} = 1$ if $z_{ist} \in [0.25, 0.5)$ and equal to 0, otherwise;...; $SHELFLEVEL_{8ist} = 1$ if $z_{ist} \in [1.75, 2.00]$, and equal to 0 otherwise.

In addition to product location covariates, several control variables are included to control for changes in prices and differences across months, stores and brands. In the descriptive analysis, we consider four different specifications, each of them with different sets of fixed effects to control for differences among product categorizations. These four sets correspond to those used by the vendor: $CATEGORY$ (Model 1), $MARKET$ (Model 2), $SUBMARKET$ (Model 3) and $GROUP$ (Model 4); where $CATEGORY$ corresponds to the broadest categorization, while $GROUP$ is the most specific one (See Table 3.2).

After this first analysis, a second model to measure the relationship between sales and shelf position among different product types is formulated as follows:

[1] For center-of-the-store aisles, the origin is defined at the section of the aisle located farthest from the checkout counters. For perimeter aisles, the origin is defined as the section of the aisle located to the right of the main entrance when standing inside the store facing the checkout counters.

[2] See the results of Larson et al. (2005) in Chapter 2.

$$
\begin{aligned}
\ln(1 + SALES_{ist}) = \beta_0 + \beta_c CENTRAL_{ist} + \beta_h HEDONIC_{ist} + \\
\sum_{k=1}^{8} \beta_{sk} SHELFLEVEL_{kist} + \\
\beta_a AISLEPART_{ist} + \\
\beta_{am} AISLEMIDDLE^2_{ist} + \\
\sum_{k=1}^{8} \beta_{shk} SHELFLEVEL_{kist} \cdot HEDONIC_{ist} + \\
\beta_{ah} AISLEPART_{ist} \cdot HEDONIC_{ist} + \\
\beta_{amh} AISLEMIDDLE^2_{ist} \cdot HEDONIC_{ist} + \\
\beta_p PRICE_{ist} + \beta_m MONTH_t + \beta_s STORE_s + \beta_b BRAND_i + \epsilon_{ist}
\end{aligned}
\tag{4.2}
$$

In order to test these associations, we classified products into two groups: utilitarian products, which are meant to be practical and useful, and are part of regular purchases; and hedonic products, which are associated with satisfaction and pleasure. The product type is incorporated via the $UTILITARIAN_{ist}$ and $HEDONIC_{ist}$ indicators. Table C.1 reports the classification of each product category. Similar to the first model, we added controls for changes in prices and differences across months, stores and brands. In the next section, we present the results of both models.

4.2. Results

Estimation results for Model 4.1 are presented in Table 4.1. The discussion begins with the results related to the role of aisle location, which are then followed by those associated with the horizontal and vertical position of products on the shelf.

First, in terms of aisle orientation, note that the coefficient for $CENTRAL$ is significant ($p < 0.001$) under all four models. Both the direction and magnitude of the effects are consistent across models. These results indicate that aisles located in the store's central area sell about 30 % less than those located at the outer ring of the store.

Considering the horizontal location of a product within an aisle with a central location (i.e., $AISLEPART$ when $CENTRAL = 1$), we find that it is significantly positive ($p < 0.001$) for all models. The magnitude of the effects varies across models, decreasing as they are adjusted using more specific fixed effects. The $CATEGORY$-adjusted model implies a 48 % greater sales for products far from the beginning of the aisle, while the corresponding increase for the GROUP-adjusted model is 24.3 %. Overall, these results implies that products located closer to the checkout counters are associated with higher levels of sales.

Considering next perimeter aisles (i.e., aisles located at the outer ring of the store), under the models (2)-(4), horizontal location is either significant or marginally significantly negative. This implies that locations at the left hand of the main entrance are associated with somewhat lower sales.

Recall that the horizontal location is measured not only in terms of the distance to the origin, but also the distance from the center of the aisle. In particular, $AISLEMIDDLE \times CENTRAL$ is significant ($p < 0.001$) and positive across all models. This covariate presents large effects on sales, and varies from 37.0 % for the GROUP-adjusted model to 69.6 % for the $MARKET$-adjusted model. Thus, locations closer to the extremes are associated with greater sales for products in aisles located in the central area of the store. These conclusions, along with those observed on the $CENTRAL$ variable, are consistent with those of Larson et al. (2005). In their study, they find that shoppers are prone to perform short excursions in and out of the aisles rather than traversing the entire length of each aisle, so products placed at the center of aisles will receive much less "face time" than those placed toward the ends (Larson et al., 2005). In contrast, the distance from the center for perimeter aisles has inconsistent results across models in terms of the direction, magnitude and significance.

The effects of vertical location are presented in reference to the first shelf, from bottom to top. The difference between being located at the first versus the second shelf is not significant for all but one model. In contrast, products located at the next three levels present larger sales compared to the first level. Although the magnitude of these increases are small in comparison to those observed for the horizontal position, locations at the fourth and fifth shelves, i.e. at 0.75 and 1.00 meters high, are associated with 2.7-5.6 % increases in sales.

These higher sales levels are not observed, however, for shelves placed at 1.25 meters and above, as they exhibit negative effects on sales compared to the bottom shelf. The results for the sixth and seventh positions are significant ($p < 0.001$) and consistent across models. Furthermore, the magnitude of the decrease in sales is considerably worse when being located at the seventh shelf. Results for the last shelf level are also negative and significant for three of the four specifications.

Figure 4.1 shows graphically the regression estimates for each specification and its corresponding confidence interval. Curves of shelf level coefficients are similar across models, showing small effects for levels between 0.50 and 1.00 meters. The figure also illustrates the difference between the effects of horizontal shelf position on sales for aisles located at the outer ring and in the central area of the store.

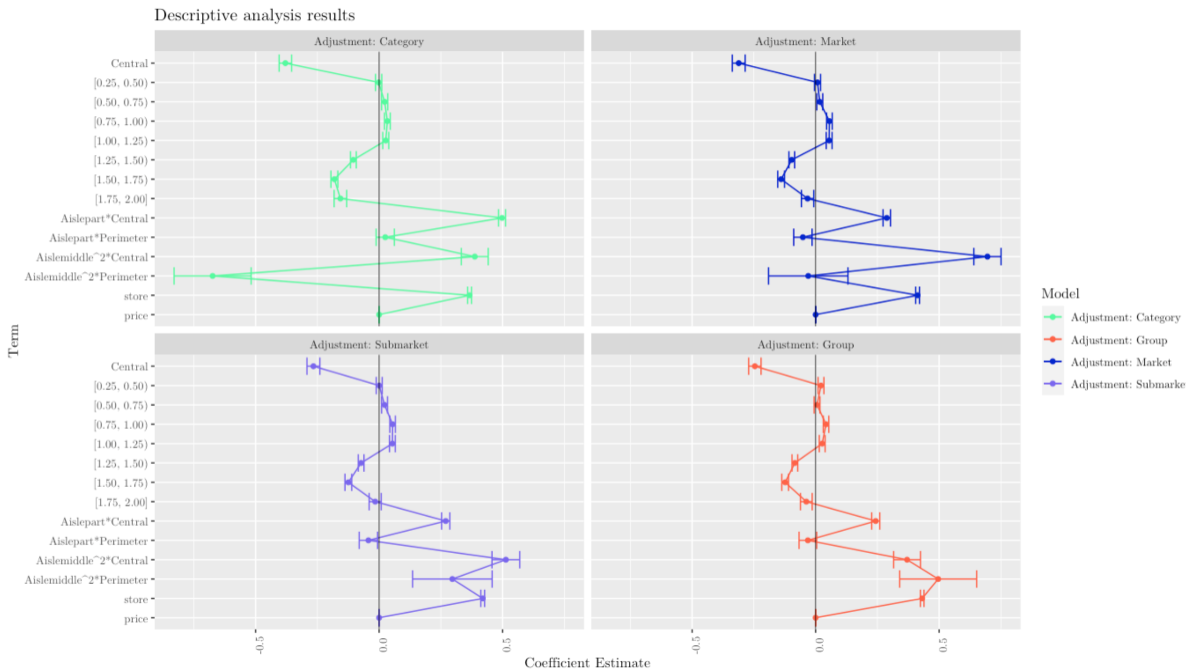

Figure 4.1: Regression results.

We now present the estimation results for Model 4.2, which are reported in Table D.1. Recall that this model allows us to compare the relationship between shelf position and sales of two different types of products: utilitarian and hedonic.

First we note from the results that, on average, hedonic products sell more than utilitarian products. Considering the horizontal shelf position of utilitarian products, we find that the extremes of an aisle are associated with greater sales, and that this is particularly more pronounced for the products located at the end of the aisle.[3] Note that similar results were found for Model 4.1. However, the corresponding implications for hedonic products are different. Considering the fraction of the aisle in which the product is located, although being located far from the beginning of the aisle is also associated with greater sales, the magnitude of this effect is less than that for utilitarian products (18.3 %). In terms of the distance from the center of the aisle, hedonic products located closer to the extremes are associated with 50.8 % smaller sales.

Considering the vertical shelf position of utilitarian products, the best locations are the third and fourth levels, which are associated with 8.3 % and 10.6 % increases in sales, respectively. Similarly to the results of the first model, shelves located at 1.25 meters and higher are the worst positions. In particular, utilitarian products located at the seventh shelf, i.e. at 1.50 meters from the floor, are associated with 17.7 % decreases in sales. On the contrary, the best vertical shelf position of hedonic products is the bottom level, while the worst locations are the third and fourth levels, which are associated with 17.7 % and 22.0 % decreases in sales.

[3] Recall that for center-of-the-store aisles, the origin is defined at the section of the aisle located farthest from the checkout counters. For perimeter aisles, the origin is defined as the section of the aisle located to the right of the main entrance when standing inside the store facing the checkout counters.

Table 4.1: Descriptive analysis results.

	Dependent variable:			
	logsales			
	(1)	(2)	(3)	(4)
Central	−0.380***	−0.311***	−0.266***	−0.246***
	(0.013)	(0.013)	(0.013)	(0.013)
[0.25, 0.50)	−0.002	0.008	0.0003	0.022***
	(0.006)	(0.006)	(0.006)	(0.006)
[0.50, 0.75)	0.022***	0.017***	0.022***	0.006
	(0.006)	(0.006)	(0.006)	(0.006)
[0.75, 1.00)	0.033***	0.056***	0.055***	0.042***
	(0.006)	(0.006)	(0.006)	(0.006)
[1.00, 1.25)	0.027***	0.055***	0.054***	0.027***
	(0.006)	(0.006)	(0.006)	(0.006)
[1.25, 1.50)	−0.104***	−0.097***	−0.073***	−0.084***
	(0.006)	(0.006)	(0.006)	(0.006)
[1.50, 1.75)	−0.181***	−0.140***	−0.125***	−0.123***
	(0.007)	(0.007)	(0.007)	(0.007)
[1.75, 2.00]	−0.157***	−0.033**	−0.016	−0.038***
	(0.013)	(0.013)	(0.013)	(0.012)
Aislepart · Central	0.498***	0.288***	0.270***	0.243***
	(0.007)	(0.008)	(0.009)	(0.008)
Aislepart · Perimeter	0.025	−0.052***	−0.043**	−0.031*
	(0.019)	(0.019)	(0.019)	(0.018)
(aislemiddle^2) · Central	0.387***	0.696***	0.513***	0.370***
	(0.028)	(0.028)	(0.029)	(0.028)
(aislemiddle^2) · Perimeter	−0.674***	−0.030	0.297***	0.496***
	(0.080)	(0.082)	(0.082)	(0.079)
Store	0.366***	0.413***	0.419***	0.431***
	(0.004)	(0.004)	(0.004)	(0.004)
Price	0.00001***	0.00001***	0.00000***	−0.00002***
	(0.00000)	(0.00000)	(0.00000)	(0.00000)
Constant	0.830***	0.472***	1.311***	0.395***
	(0.050)	(0.067)	(0.063)	(0.102)
Observations	266,949	266,949	266,949	266,949
R^2	0.268	0.315	0.346	0.401
Adjusted R^2	0.268	0.315	0.346	0.400
Residual Std. Error	0.758 (df = 266880)	0.734 (df = 266862)	0.717 (df = 266835)	0.686 (df = 266792)
F Statistic	1,436.960*** (df = 68; 266880)	1,425.846*** (df = 86; 266862)	1,249.323*** (df = 113; 266835)	1,143.764*** (df = 156; 266792)

Note: The four model specifications differ in the product categorization used: category, market, submarket and group, respectively.

*p<0.1; **p<0.05; ***p<0.01

Finally, it is important to note that the analysis in this section corresponds to a descriptive (i.e., non-causal) approach to study the relationship between shelf position and sales. Nonetheless, the fundamental challenge is to deal with omitted variables that could influence both sales and shelf location. This applies both to cross-sectional comparisons of different products and it is also relevant for time series analysis of the performance of the same product placed at different locations. In terms of the former, locating the best selling products at certain locations (e.g., at eye level or closer to the ends of the aisles) may lead to a reverse causality problem when estimating the relationship between shelf position and sales. In addition, changing a product's shelf position in anticipation of an expected increase/decrease in demand may also yield a simultaneity problem between shelf position and sales. An approach to address this omitted variable issues and provide causal inferences is described and implemented in the next Chapter.

Chapter 5

Estimation of the effects of shelf position on sales using a quasi-experimental approach

This chapter presents the estimation of the effects of shelf position on sales using a quasi-experimental approach. We first present the quasi-experimental design and then introduce a method to detect changes in shelf position that allows us to identify treatment and control units. Finally, we conduct a DID analysis and present the results with the corresponding discussion.

5.1. Quasi-experimental design

Randomized controlled experiments are usually the preferred method for measuring causal effects. However, conducting controlled experiments with random assignment may sometimes be highly costly or impractical. In contrast, quasi-experimental designs are less expensive, time consuming and pragmatic and have more realism with regard to how the program or treatment is administered (Schweizer et al., 2016). Accordingly, a difference-in-difference (DID) quasi-experimental approach is used to estimate the effects of shelf position on sales, drawing on the observational data presented in previous chapters. We begin by introducing the definitions for quasi-experimental treatment and control units and then discuss two necessary assumptions of the DID estimator.

The design treatment was defined as a product exhibiting a sharp change in shelf position. We note that a precise definition of what constitutes a sharp change in our setting is provided in the next section. The treatment is then represented by the following expression:

$$SP_{itsa} = \begin{cases} sp_0, \; t \; < \; t^* \\ sp_1, \; t \; \geq \; t^* \end{cases} \tag{5.1}$$

where SP_{itsa} is the shelf position of SKU i, at time t, in store s, along axis A (i.e., horizontal or vertical axis). If the change occurred at time t^*, we require constant shelf positions before and after t^*: sp_0 is the shelf position at the pre-treatment period, while sp_1 is the shelf position at the post-treatment period.

Finding such changes in a non experimental setting is challenging because product placement modifications in supermarket stores might be relatively infrequent. However, having a sufficiently long time series makes it more likely to find these instances. With this definition of the treatment, one approach to measure the causal impact of shelf position is to compare the sales patterns of the product before and after the intervention. A limitation of this approach is that there could be omitted variables that may explain the differences between the sales patterns observed before and after the treatment. Hence, it is useful to identify and rely on a control group to obtain a more robust estimate. In this research, we rely on the sales pattern of the treated product at a different store to construct a control group.

Defining these control groups might be challenging because stores belonging to the same supermarket chain may have the same in-store marketing policies, so it might be common to observe similar product locations when visiting them. Shelf position policies, however, are not always implemented simultaneously, so these time windows can be used to define the treatment and the control group. The data for the treatment an control group can then be analyzed using a DID approach.

Two necessary assumptions for the DID estimator are consistency and parallel trends. The consistency assumption of the DID estimator is expressed as follows (Lechner, 2011):

$$Y(t) = (1 - T) \cdot Y^0(t) + T \cdot Y^1(t) \tag{5.2}$$

where $Y(t)$ is the outcome at time t and T is the treatment status. $Y^0(t)$ represents the sales of a product that was assigned to the control group and $Y^1(t)$ represents the sales of a product assigned to the treatment group. Note that when a unit is assigned to the treatment, i.e., $T = 1$, only $Y^1(t)$ is observed, while the potential outcome without the treatment is not. On the contrary, if a unit was not assigned to the treatment, i.e., $T = 0$, only Y^0 is observed, while the potential outcome with the treatment is not. In other words, the consistency assumption demands that treated and untreated status remains constant: the treatment group is treated in the post treatment period only, and the control group is never treated. This assumption is also called Stable Unit Treatment Value assumption (SUTVA). In particular, the sharp change defined in Equation 5.1 and the selection method to find them ensure that our data satisfy the consistency assumption.

Another assumption of the DID estimator is that the error term is not correlated with the other variables of the model, in particular: T, t and $T \cdot t$. This assumption is relevant to obtain an unbiased estimate of the treatment effect. This parallel trends are part of this assumption, and it is satisfied when $corr(\epsilon, T \cdot t) = 0$. This key assumption of the difference-in-differences design means that in the absence of treatment, the average change in the outcome of interest for the treated units would have been equal to the average change in the outcome for the untreated ones (Bundell and Dias, 2005). Since we will consider a relatively short period before and after the treatment, it is less likely that some events might change the sales trend only for one store but not for the other within our period of analysis.

5.2. Identification of treatment and control units

In this section we describe the procedure used to identify products with sharp changes in their shelf position in one store and exhibiting no changes in the other. This will allow us to identify experimental and control units. The main stages of this process are the following:

i) Select a store $s_i \in \{A, B\}$

ii) Detect products (SKUs) with sharp changes as defined in Equation 5.1 in store s_i.

iii) Check the shelf position of the same product in the other store, s_j.

iv) If no changes are found in s_j, define s_i as the treatment group and s_{-i} as the control group.

v) Ensure that treated units are thoroughly following the treatment assigned.

vi) Verify that other possible explanations for changes in sales are mitigated.

A description of these steps is provided next. First, a changepoint analysis was implemented in order to detect significant changes in the horizontal or vertical position of a product. Changepoints are defined as those points in a data sequence where a change in a statistical property is observed (mean, variance or distribution). Examples of multiple changepoint search algorithms that have been proposed so far are Binary Segmentation (Scott and Knott, 1974), Segment Neighborhood (Auger and Lawrence, 1989), and Pruned Exact Linear Time, PELT (Killick et al., 2012). The latter was chosen for this study because it is an exact algorithm, and its computational cost is smaller.

In practice, changepoint detection begins with a visual inspection of the time series, followed by the implementation of the search algorithm selected. Thus, a visual inspection of the horizontal and vertical shelf position of each SKU was performed. The main objective of this first step was to verify if there are any changes within the data. We aimed to find sharp changes in position, as illustrated in Figure 5.1. This kind of change reveals a transition that can be used as a treatment experiment, and ensures that consistency assumption is satisfied.

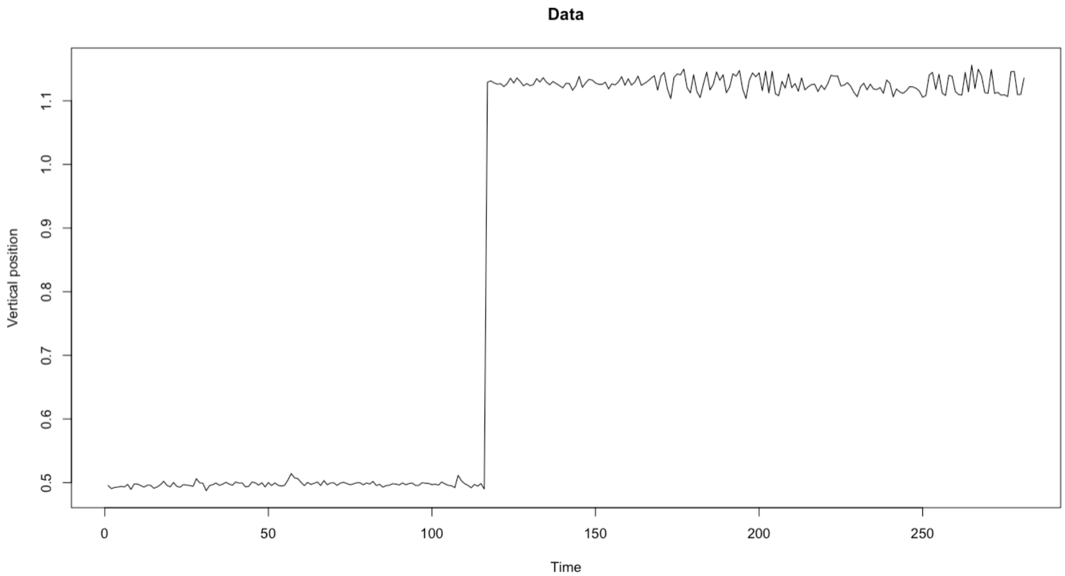

Figure 5.1: Sharp change in vertical shelf position.

Thereafter, the PELT algorithm was applied to all SKUs that had a sharp change either in its horizontal or its vertical shelf position. Multiple changepoint detection requires the specification of a penalty to guard against under/over-fitting. PELT was first applied with standard penalties such as BIC, AIC and Asymptotic with theoretical type I error of 0.1 and 0.01. Since the three penalties under-fitted the data (Figure E.1), PELT was then applied along with the CROPS[1] method (Haynes et al., 2014). As a result, the exact time at which the change in position occurred was detected. Thereafter, the remaining store was inspected to verify that no changes in shelf position occurred.

Since the robot is not always able to enter an aisle, some units have incomplete positioning data.[2] To ensure that the treated product followed the treatment with greater confidence, we impose a data frequency condition such that we exclude periods with more than three consecutive days without shelf-position data. Finally, the first day in which a change in position was observed is removed from the DID analysis, due to the uncertainty about the exact moment when shelf position change occurred.

Considering the possibility of alternative explanations for sales changes, a set of additional conditions was imposed. Firstly, the product location inside the store, that is, the aisle in which it is displayed, must remain constant. Since occasional display of a product in other aisles could explain changes in sales, treatment periods which had such events were excluded from the analysis. Secondly, the price of the product was required to exhibit no more than a 1 % change.

Figure 5.2 illustrates the conditions imposed to select a treatment and control unit. Vertical

[1] Changepoint for a Range of Penalties, CROPS.

[2] When the robot meets an obstacle in its path, such as supermarket carts or pallet jacks, it does not enter the aisle.

position in store A had a sharp change as required in Equation 5.1, while it remained constant at store B. In particular, a pre- and post-treatment period of 7 days was available for this unit. Note that there were no changes in price.

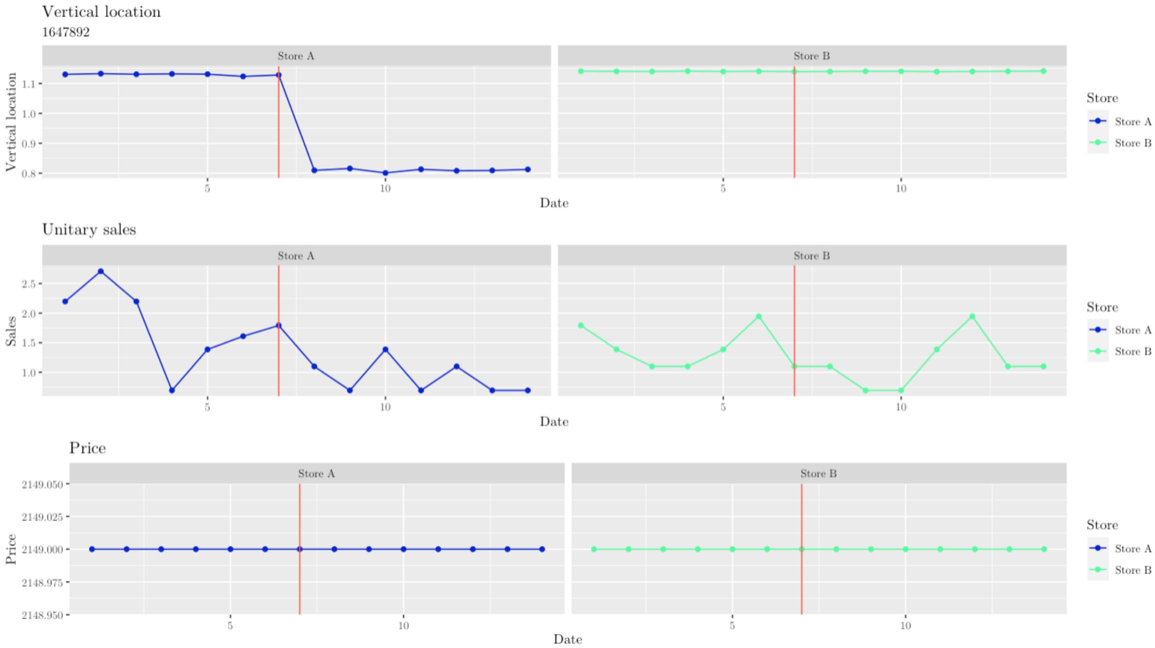

Figure 5.2: Vertical position, unitary sales and price of a selected unit.

As a result of the selection method, a dataset containing treated and non-treated units was generated. In store A, 16 SKUs successfully passed the conditions, with prices ranging between $499 and $4890 Chilean pesos and daily unitary sales of 1 - 33 units. In store B, 22 SKUs passed the conditions, with prices ranging between $499 and $4929 Chilean pesos and daily unitary sales of 1 - 68 units. SKUs of both stores belonged to the following categories: Sauces & Dressing, Grocery, Personal Care, and Cleaning Products. No sharp changes were found in the Dairy product category.

In practice, the final dataset was composed of: 16 treated SKUs from store A, with their corresponding control group in store B; and 22 treated SKUs from store B, with their corresponding control group in A. Most of the treated units belong to the Sauces and Dressings and Deodorants markets, with 16 and 12 units respectively. The remaining group belong to the following markets: Hair care, Laundry detergents, Dishwashing detergents, Soaps, Red class, Jam, honey and sweet, and Herbal infusions.

Most of the SKUs experienced changes in its vertical shelf position only: 8 units decreased their position, while 13 increased it. 9 SKUs changed its horizontal shelf position only: 7 came closer to the center of the aisle, while 2 came farther. Finally, 8 SKUs experienced changes both in its vertical and horizontal shelf position. A summary is reported in Table 5.1.

The next section describes the DID approach used to measure the effect of shelf position on sales.

Table 5.1: Type of changes summary.

Type of change	Vertical change	Horizontal change	Total units
Vertical	Decrease	No change	8
	Increase	No change	13
Horizontal/Vertical	Decrease	Closer	4
	Increase	Closer	1
	Increase	Farther	3
Horizontal	No change	Closer	7
	No change	Farther	2

5.3. Difference-in-Difference analysis

This section aims to evaluate the impact of a treatment, T, over an outcome, Y, where the treatment is a sharp change in position as defined in Equation 5.1, and the outcome is the variable $\ln(1 + SALES)$ as defined in Chapter 4. For subsequent analysis, standard treatment status and the pre/post treatment periods are defined:

$$T = \begin{cases} 1, & \text{if the unit received the treatment.} \\ 0, & \text{otherwise.} \end{cases} \tag{5.3}$$

$$post_t = \begin{cases} 1, & \text{if the period t corresponds to a time after the treatment group received the treatment.} \\ 0, & \text{otherwise.} \end{cases} \tag{5.4}$$

The variable T determined the treatment group ($T = 1$) and the control group ($T = 0$), while variable *post* indicates the pre-treatment period ($post = 0$), and the post-treatment period ($post = 1$). Note that if a product received the treatment, $T = 1, \forall t$. Besides the standard variables for the DID analysis, a set of covariates regarding shelf position were added. Before performing the difference in difference analysis, a naive regression was used to examine the 38 selected units separately. Both the model and the results are reported in Appendix F.

We now present the difference-in-difference analysis conducted jointly considering all treatment units. To measure the effects of shelf position on sales, we constructed 8 different metrics of changes in position. We now present this metrics and then explain the proposed models and report its results.

A set of 4 binary variables were created to capture the type of the change. The first two indicates whether the unit changed its horizontal or its vertical position:

$$CHANGE.H = \begin{cases} 1, & \text{if the unit changed its horizontal shelf position.} \\ 0, & \text{otherwise.} \end{cases} \tag{5.5}$$

$$CHANGE.V = \begin{cases} 1, & \text{if the unit changed its vertical shelf position.} \\ 0, & \text{otherwise.} \end{cases} \tag{5.6}$$

The remaining two binary variables were created to capture the direction of the change, based on the results of the descriptive analysis performed in Chapter 4. Recall that we found a positive association between sales and being located at the edges of an aisles, and a positive relation between sales and being located at the middle heights of the shelves. The variables are the following:

$$CLOSER.H = \begin{cases} 1, & \text{if the unit came closer to the center of the aisle.} \\ 0, & \text{otherwise.} \end{cases} \tag{5.7}$$

$$CLOSER.V = \begin{cases} 1, & \text{if the unit came closer to the center of the shelves (1.0 meters).} \\ 0, & \text{otherwise.} \end{cases} \tag{5.8}$$

Additionally, a set of 4 variables to measure the direction and the magnitude of the change were constructed. First, two variables to measure whether or not it moved away from the beginning of the aisle and whether or not the product increased its height in the shelves:

$$CHANGE.AISLEPART = AISLEPART_1 - AISLEPART_0 \tag{5.9}$$

$$CHANGE.VERTICAL = VERTICAL_1 - VERTICAL_0 \tag{5.10}$$

Where 1 and 0 represents the post and the pre-treatment period respectively. Note that Equation 5.9 is constructed based on the $AISLEPART$ variable defined in Chapter 4, which represent the fraction of the aisle in which the price tag was read. Finally, two variables to measure how much a unit came closer or farther from the center of the aisle and from the center of the shelve, respectively:

$$CHANGE.AISLEMIDDLE = AISLEMIDDLE_1 - AISLEMIDDLE_0 \tag{5.11}$$

$$CHANGE.VERTICALMIDDLE = VERTICALMIDDLE_1 - VERTICALMIDDLE_0 \tag{5.12}$$

Where 1 and 0 represents the post and the pre-treatment period respectively. The change in the horizontal shelf position was calculated with the $AISLEMIDDLE$ variable described in Chapter 4 and therefore, this change indicates how much the unit came closer or farther from the center of the aisle. The variable $VERTICALMIDDLE$ indicates the distance from the center of the shelve, which correspond to 1.0 meter. Note that these eight metrics take values different from zero if and only if a unit belongs to the treatment group and $t = 1$.

Finally, we adjust for the store and the day of the week, and added fixed effects for each unit. The following formulation summarizes the model:

$$\begin{aligned} \ln(1 + SALES_{ist}) = \beta_0 + \beta_1 post_t + \beta_2 T_{is} + \beta_h \cdot HORIZONTAL_{metrics} + \beta_v \cdot VERTICAL_{metrics} + \\ \beta_3 STORE_{is} + \beta_4 DAY_t + \beta_5 SKU_i + \epsilon_{ist} \end{aligned} \tag{5.13}$$

We estimated 5 models. All of them incorporated the variables corresponding to the treatment period, *post*, the treatment status, T, and the adjustments for store, day of the week and the fixed effects for each unit. The difference lies in the combination of metrics used, that is, $HORIZONTAL_{metrics}$ and $VERTICAL_{metrics}$ matrix. Table 5.2 shows the combination of metrics used in each model tested.

Table 5.2: DID metrics.

Variable	Model 1	Model 2	Model 3	Model 4	Model 5
Change.h	✓	✓			✓
Change.v	✓	✓			✓
Closer.h		✓			✓
Closer.v		✓			✓
Change.aislepart			✓		✓
Change.vertical			✓		✓
Change.aislemiddle				✓	
Change.verticalmiddle				✓	

Table 5.3 shows the results for different subsets of explanatory variables. Under all estimated models, the treatment variable, T, was significant ($p < 0.001$), and had negative effects on sales of about 20 %. The coefficient estimated for this variable indicates that the treated units had on average lower sales than the control units. Also note that *post* is not significant under any model, implying no significant differences between the pre- and post-treatment periods for the control units.

In terms of the first model, Model 1 considers whether or not there were changes in the vertical or horizontal position, and neither of these effects is significant. We also note that the magnitude of the coefficients for these variables is small and it is not consistent across models. Model 2 shows that getting closer or farther from the vertical center of the shelves is significant ($p < 0.05$), while getting closer or farther from the center of the aisle is not. This results indicates that on average the increase in sales is about 15 %. Both the significance and the direction of this effects are consistent across models. Model 3 considers the magnitude of the vertical and horizontal changes. Neither of these coefficients are significant. Model 4 shows a significant negative effect of getting farther from the vertical center of the aisle, a result that is consistent with Model 2. Model 5 reaffirms the findings of the previous models: no significant results for horizontal changes, and significant results for vertical changes. Combining the effects of the *closer.v* and *change.vertical* variables, the model indicates that moving away from the center of the shelves upwards has a negative impact on sales.

All five models had a similar fit, with R^2 going from 0.503 to 0.508. Model five reported the highest R^2 and adjusted R^2. Compared with the descriptive analysis performed in Chapter 4, the results of the vertical changes are consistent: moving away from the vertical center of the shelves upwards results in a decline in sales. However, in the descriptive analysis it was also observed that the products located closer to the ends of the aisles and closer to the checkout counters reported higher sales than those located near the center. The difference-in-difference analysis shows no significant effects. This could reveal a selection bias: products with the highest sales were located near the checkout counters, hence when a control store is added to the analysis, the effect is not significant.

In addition to this analysis, we perform a second DID analysis incorporating the product type variables defined in Chapter 4 to compare the effects of the treatment of this two types of products. As in the case of the first model, we adjusted for the store and the day of the week as well as added fixed effects for each unit. Analogously, we estimated 5 models with the same combination of metrics. Estimation results are reported in Table 5.4. Since results for the T, *post* and $STORE$ variables are similar to the results of the first model, we did not incorporated them in Table 5.4. We now discuss the effects of the treatment by product type.

First of all, the $UTILITARIAN$ variable is significant (p <0.001) and negative under all four models, with magnitudes varying from 70.6 % to 78.8 %. This means that, on average, products classified as utilitarian sell approximately 75.3 % less than hedonic products.

Model (1) and (3) show a large negative impact on sales for utilitarian products that changed its horizontal position farther from the beginning of the aisle. According to model (3), an utilitarian product that moves closer to the checkout counters, on average decrease their sales by 97.6 %. By contrast, hedonic products increases its sales by 41.5 % by getting closer to the checkout counters. In terms of vertical changes in position, hedonic products experienced a decrease in its sales when they increase their height on the shelves, by about 28.2 %.

Model (2) and (4) show negative effects on sales for hedonic product that move away from the center of the shelves, i.e. the least favorable position change for a hedonic product is to move away from the center of the shelf by increasing its vertical position. In terms of horizontal changes, moving closer to the extremes of the aisle has large negative effects on utilitarian product's sales. In fact, on average, they decrease its sales by about 97.0 %.

These results, in summary, show us that the treatment has different effects depending on the type of product in question. In general, if utilitarian and hedonic products experience the same change in position, opposite effects on their sales can be expected.

Table 5.3: DID analysis results.

	Dependent variable:				
	logsales				
	(1)	(2)	(3)	(4)	(5)
change.h	0.010	−0.001			0.041
	(0.056)	(0.081)			(0.106)
change.v	0.071	0.001			−0.0001
	(0.052)	(0.063)			(0.063)
closer.h		−0.013			−0.080
		(0.097)			(0.106)
closer.v		0.152**			0.234***
		(0.077)			(0.083)
change.aislepart			−0.023		−0.147
			(0.201)		(0.273)
change.vertical			−0.102		−0.266***
			(0.089)		(0.102)
change.aislemiddle				−0.261	
				(0.276)	
change.verticalmiddle				−0.282**	
				(0.135)	
T	−0.212***	−0.210***	−0.175***	−0.181***	−0.206***
	(0.037)	(0.037)	(0.030)	(0.028)	(0.037)
Post	−0.039	−0.035	−0.003	−0.006	−0.031
	(0.037)	(0.037)	(0.029)	(0.028)	(0.037)
Store	0.560***	0.552***	0.564***	0.553***	0.552***
Constant	1.846***	1.830***	1.849***	1.830***	1.862***
	(0.081)	(0.082)	(0.084)	(0.082)	(0.086)
Observations	1,216	1,216	1,216	1,216	1,216
R^2	0.503	0.505	0.503	0.504	0.508
Adjusted R^2	0.483	0.484	0.482	0.484	0.486
Residual Std. Error	0.481 (df = 1167)	0.481 (df = 1165)	0.481 (df = 1167)	0.480 (df = 1167)	0.479 (df = 1163)
F Statistic	24.617*** (df = 48; 1167)	23.768*** (df = 50; 1165)	24.598*** (df = 48; 1167)	24.732*** (df = 48; 1167)	23.113*** (df = 52; 1163)

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

Table 5.4: DID analysis results by product type

	Dependent variable: logsales				
	(1)	(2)	(3)	(4)	(5)
Utilitarian	−0.786***	−0.742***	−0.788***	−0.743***	−0.743***
	(0.129)	(0.130)	(0.134)	(0.130)	(0.134)
change.h · Utilitarian	−0.190*				−0.235
	(0.099)				(0.240)
change.h · Hedonic	0.077				0.403*
	(0.074)				(0.210)
change.v · Utilitarian	0.029				−0.073
	(0.068)				(0.091)
change.v · Hedonic	0.035				0.019
	(0.074)				(0.082)
closer.h · Utilitarian		0.151			0.373
		(0.229)			(0.270)
closer.h · Hedonic		−0.050			−0.513***
		(0.071)			(0.182)
closer.v · Utilitarian		0.073			0.129
		(0.079)			(0.118)
closer.v · Hedonic		0.259***			0.108
		(0.088)			(0.176)
change.aislepart · Utilitarian			−0.976**		−0.247
			(0.389)		(0.905)
change.aislepart · Hedonic			0.415*		0.011
			(0.245)		(0.301)
change.vertical · Utilitarian			−0.080		−0.097
			(0.130)		(0.158)
change.vertical · Hedonic			−0.282**		−0.517***
			(0.130)		(0.152)
change.aislemiddle · Utilitarian				−0.970**	
				(0.379)	
change.aislemiddle · Hedonic				0.530	
				(0.418)	
change.verticalmiddle · Utilitarian				−0.143	
				(0.209)	
change.verticalmiddle · Hedonic				−0.308*	
				(0.178)	
Constant	1.826***	1.795***	1.807***	1.773***	1.781***
	(0.082)	(0.083)	(0.086)	(0.085)	(0.088)
Observations	1,216	1,216	1,216	1,216	1,216
R^2	0.506	0.506	0.507	0.508	0.517
Adjusted R^2	0.484	0.485	0.486	0.487	0.493
Residual Std. Error	0.480 (df = 1165)	0.480 (df = 1165)	0.479 (df = 1165)	0.479 (df = 1165)	0.476 (df = 1157)
F Statistic	23.834*** (df = 50; 1165)	23.906*** (df = 50; 1165)	23.991*** (df = 50; 1165)	24.024*** (df = 50; 1165)	21.383*** (df = 58; 1157)

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

Chapter 6

Conclusion

The aim of this study was to examine the effects of shelf position on sales. To do so, we first performed a regression analysis to study possible associations between variables. The regression analysis revealed that there is a positive relationship between sales and being located at the edges of the aisle or near to the checkout counters. The results also suggested a positive relationship between sales and being located at a height of 0.75 to 1.00 meters on the shelves. Thereafter, we conducted a DID analysis on a set of products that had changes in their horizontal or vertical position. Such changes were detected through a semiautomatic method. One of the more significant findings emerging from this analysis is that a product can increase its sales by moving closer to the vertical center of the shelf, and decrease its sales by moving from the center upwards. Changing its horizontal shelf position made no significant difference on sales. Taken together, these findings suggest 1) that there is a causal relationship between the vertical position of a product and its sales, and 2) the possibility that retailers allocate the best selling products at the edges of the aisle and near the checkout counters.

This work contributes to existing knowledge by providing a causal link between shelf position and sales. The method used in this study to detect changes in position permits to design a quasi-experiment relying on observational data without intervening the stores, but ensuring that units follow the treatment. In addition, it may be applied to other stores to find multiple controls.

An issue that was not addressed in this study was whether the state of the other products of the same aisle remained constant, that is, if they changed their position, their prices, etc. Also, this study did not include information on stockouts and missing price tag alerts generated by the robot. In addition, although the descriptive regression analysis was tested on a large sample size, our difference-in-difference analysis was limited to only 38 units. It was not possible to study the effects of changes in position to the whole set of products sold by the retailer, because of the absence of sales data for products belonging the other vendors. This limitation also affected the variety of products studied.

Further research might explore a methodological improvement to complement our findings. A natural progression of this work is to repeat the study using multiple controls. Zippedi is currently operating in other stores belonging to the same chain. Stores that are in the same geographical area, but far enough away to ensure changes in one store are less likely to affect the sales of the other, could be used as multiple control groups to strengthen our findings. If these remain similar, our DID analysis results would be more robust. Further work could also be conducted using other methodologies to estimate the causal effects of the treatment, such as the Synthetic Control Methods introduce by Abadie and Gardeazabal (2003) and Abadie et al. (2010) that construct synthetic control units driven by the data. This method could be useful to choose the control units in a non-arbitrary way. Finally, using a broader range of products could be useful not only to increase the sample size of the descriptive analysis, but also of the quasi-experiment. Additionally, this could add variation regarding the products that are being studied.

In sum, we hope that this study might contribute to improve our understanding of the impact of in-store marketing actions on sales and profits and that might stimulate further research aimed at taking advantage of the new data sources that become available with the use of new technologies such as robotics, internet of things and machine learning.

Bibliography

Abadie, A., Diamond, A., and Hainmueller, J. (2010). Synthetic Control Methods for Comparative Case Studies: Estimating the Effect of California's Tobacco Control Program. *Journal of the American Statistical Association*, 105(409):493–505.

Abadie, A. and Gardeazabal, J. (2003). The Economic Costs of Conflict: A Case Study of the Basque Country. *American Economic Review*, 93(1):113–132.

Auger, I. E. and Lawrence, C. E. (1989). Algorithms for the optimal identification of segment neighborhoods. *Bulletin of Mathematical Biology*, 51:39–54.

Bell, D. R., Corsten, D., and Knox, G. (2010). Unplanned Buying on Shopping Trips. *Marketing Science Institute Working Paper Series*, Report No. 10-109.

Bundell, R. and Dias, M. C. (2005). Evaluation Methods for Non- Experimental Data. *Fiscal Studies*, 21(4):427–468.

Chandon, P. and Hutchinson, J. W. (2009). Does In-Store marketing works? Effects of the Number and Position of Shelf Facings on Brand Attention and Evaluation at the Point of Purchase. *Journal of Marketing*, 73:1–17.

Chevalier, M. (1975). Increase in Sales Due to In-Store Display. *Journal of Marketing Research*, 12(4):426–431.

Curhan, R. C. (1972). The Relationship between Shelf Space and Unit Sales in Supermarkets. *Journal of Marketing Research*, 9(4):406–412.

Dagnoli, J. (1987). Impulse Governs Shoppers. *Advertising Age*, October 5th:93.

Desmet, P. and Renaudin, V. (1998). Estimation of product category sales responsiveness to allocated shelf space. *International Journal of Reseach in Marketing*, 15(5):443–457.

Drèze, X., Hoch, S. J., and Purk, M. E. (1994). Shelf Management and Space Elasticity. *Journal of Retailing*, 70(4):301–326.

Frank, R. E. and Massy, W. F. (1970). Shelf Position and Space Effects on Sales. *Journal of Marketing Research*, 7(1):59–66.

Haynes, K., Eckley, I., and Fearnhead, P. (2014). Efficient penalty search for multiple changepoint problems.

Hui, S. K., Bradlow, E. T., and Fader, P. S. (2007). An Integrated Model of Grocery Store Shopping Path and Purchase Behavior. *Journal of Consumer Research*, 36(3):478–478.

Killick, R. and Eckley, I. A. (2014). changepoint: An R Package for Changepoint Analysis. *Journal of Statistical Software*, 58(Issue 3).

Killick, R., Fearnhead, P., and Eckley, I. A. (2012). Optimal Detection of Changepoints

With a Linear Computational Cost. *Journal of the American Statistical Association*, 107(500):1590–1598.

Larson, J. S., Bradlow, E. T., and Fader, P. S. (2005). An Exploratory Look at Supermarket Shopping Paths. *International Journal of Reseach in Marketing*, 22(4):395–414.

Lavielle, M. (2005). Using penalized contrasts for the change-point problem. *Signal Processing*, 85(8):1501–1510.

Lechner, M. (2011). The estimation of causal effects by difference-in-difference methods. *Foundations and Trends® in Econometrics*, 4(3):165–224.

Pinna, F. and Seiler, S. (2017). Estimating Search Benefits from Path-tracking Data: Measurement and Determinants. *Marketing Science*, 36(4).

Schweizer, M. L., Braun, B. I., and Milstone, A. M. (2016). Research Methods in Healthcare Epidemiology and Antimicrobial Stewardship—Quasi-Experimental Designs. *Infect Control Hosp Epidemiol*, 37(10):1135–1140.

Scott, A. J. and Knott, M. (1974). A Cluster Analysis Method for Grouping Means in the Analysis of Variance. *Biometrics*, 30(3):507–512.

Appendix A

Preprocessing and data cleaning

Negative values in horizontal and vertical position variables were eliminated from the dataset. In addition, SKUs with less than 6 observations per month were not considered in the analysis.

Outliers detected in vertical position, corresponding to values greater than 2.1 meters, were also eliminated. Outliers in sales were observed in laundry detergents, so a filter consisting in removing 10 % of each tail was applied.

Noise in price was cleaned through a process consisting in the following steps: firstly, price standard deviation and mean of each SKU were calculated without considering values belonging to the first five percentiles and the last five percentiles; secondly, observations with prices exceeding ± 3 standard deviations were eliminated; third, price standard deviation and mean of each market were calculated, with the same aforementioned conditions, and observations with prices exceeding ± 3 standard deviations were eliminated. The objective of applying the first filter at the SKU level is to eliminate anomalous price values considering historical observations of each product. However, if a product has plenty of anomalous values, the filter at the SKU level is not enough. Thus, the second filter was applied to compare how large these remaining values are in comparison to the values presented by the market to which it belongs. After the cleaning process, 40 observations that still presented anomalous values were eliminated. Missing values in price were linearly interpolated.

Appendix B

Data description

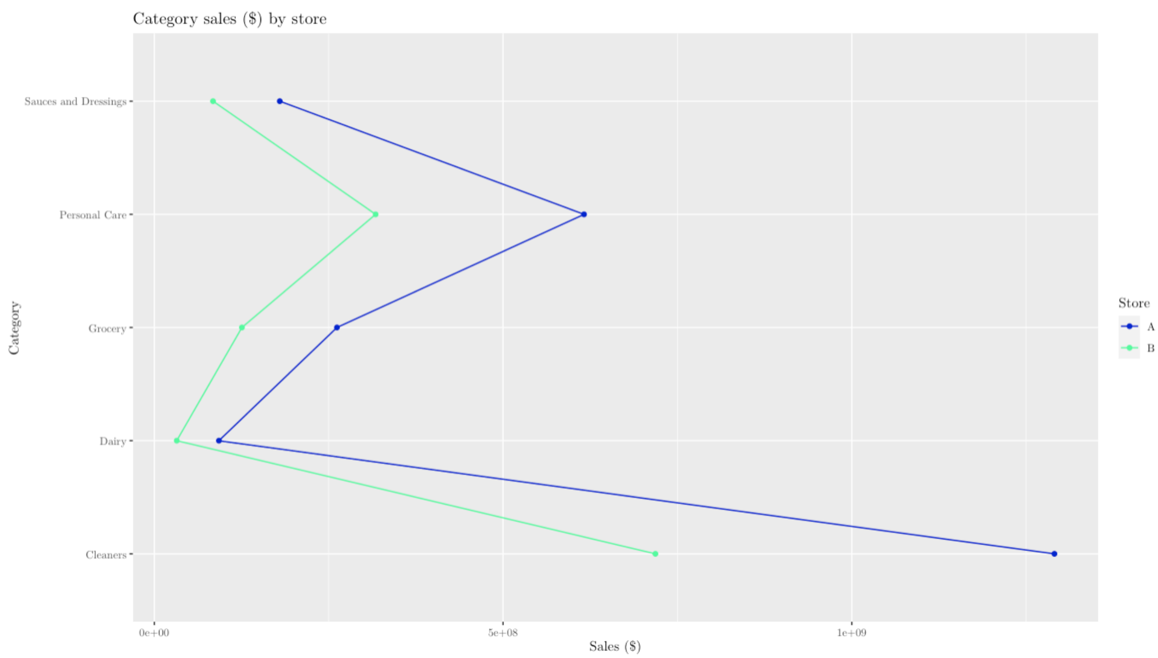

Figure B.1: Category sales by store.

Appendix C

Vendor's market description

Table C.1: Market description.

Market	SKUs	Brands	General description	Type of product
Deodorant	118	5	Stick and spray deodorants	Utilitarian
Hair care	111	7	Shampoo, conditioner and styling cream	Utilitarian
Soaps	66	6	Liquid soap, bar and shower gel	Utilitarian
Sauces and dressings	62	4	Mayonnaise, mustard, ketchup, chili sauces	Hedonic
Laundry detergent	55	4	Laundry detergent	Utilitarian
Herbal infusion	49	3	Tea and infusions	Utilitarian
Oral hygiene	41	2	Toothbrushes and toothpastes	Utilitarian
Personal care	37	6	Face creamd and exfoliating scrubs	Utilitarian
Ice cream	36	5	Ice cream and desserts	Hedonic
Household cleaner	36	5	Glass cleaner, grease remover, and multi surface cleaners	Utilitarian
Dishwashing detergent	33	4	Dishwashing detergent	Utilitarian
Softener	20	2	Softeners and laundry stain removers	Utilitarian
Jam, honey and sweet	16	1	Jams	Utilitarian
Red class	15	3	Tomato sauces	Utilitarian
Gift kits	4	14	-	Hedonic
Textile	2	5	Cleaning cloths and wet wipes	Utilitarian
Margarine	2	5	Margarine	Utilitarian
Noodles	1	8	Noodles	Utilitarian
Foot and manicure products	1	5	Powder and spray foot deodorants	Utilitarian

Appendix D

Descriptive analysis

Table D.1: Descriptive analysis results: product type comparison.

	Dependent variable:
	logsales
Central	-0.047^{***}
	(0.007)
[0.25, 0.50)	0.020^{***}
	(0.007)
[0.50, 0.75)	0.083^{***}
	(0.007)
[0.75, 1.00)	0.106^{***}
	(0.007)
[1.00, 1.25)	0.057^{***}
	(0.007)
[1.25, 1.50)	-0.078^{***}
	(0.007)
[1.50, 1.75)	-0.177^{***}
	(0.008)
[1.75, 2.00]	-0.123^{***}
	(0.013)
Aislepart	0.463^{***}
	(0.007)
(Aislemiddleˆ2)	0.432^{***}
	(0.030)
Hedonic	0.081^{***}
	(0.021)
[0.25, 0.50) · Hedonic	-0.082^{***}
	(0.015)
[0.50, 0.75) · Hedonic	-0.260^{***}
	(0.017)
[0.75, 1.00) · Hedonic	-0.326^{***}
	(0.015)
[1.00, 1.25) · Hedonic	-0.095^{***}
	(0.017)
[1.25, 1.50) · Hedonic	-0.057^{***}
	(0.017)
[1.50, 1.75) · Hedonic	0.047^{***}
	(0.018)
[1.75, 2.00] · Hedonic	0.035
	(0.118)
Aislepart · Hedonic	-0.280^{***}
	(0.018)
(Aislemiddleˆ2) · Hedonic	-0.940^{***}
	(0.065)
Store	0.389^{***}
	(0.004)
Price	0.00001^{***}
	(0.00000)
Constant	0.704^{***}
	(0.051)
Observations	266,073
R^2	0.270
Adjusted R^2	0.270
Residual Std. Error	0.758 (df = 265997)
F Statistic	$1{,}314.274^{***}$ (df = 75; 265997)
Note:	*p<0.1; **p<0.05; ***p<0.01

Appendix E

Changepoint detection method: PELT and CROPS algorithms

In many industries, a large amount of data is recorded and stored every day thanks to the incorporation of new technologies. In particular, the retail stores considered in this study have incorporated a robot that daily records the shelf position of each product, making these sequences available for analysis. With more than 600 products read each month, a method to detect changes in position within these sequences is necessary. The method used is called changepoint analysis and is described in this appendix.

Changepoint detection is the name given to the problem of estimating the point at which the statistical properties of a sequence of observations change. In an ordered sequence of data, $y_{1:n} = (y_1, \ldots, y_n)$, a changepoint is said to occur when there exist a time $\tau \in (1, \ldots, n-1)$ such that the statistical properties of $(y_1, \ldots, y_\tau)$ and $(y_{\tau+1}, \ldots, y_n)$ are different in some way (Killick et al., 2012). When multiple changes occur, said m, the m changepoints split the data into $m+1$ segments. To estimate the number and position of multiple changepoints, the PELT search algorithm proposed and tested by Killick et al. (2012) minimizes:

$$\sum_{i=1}^{m+1} C(y_{(\tau_{i-1}+1):\tau_i}) + \beta f(m) \tag{E.1}$$

Where $\beta f(m)$ is a penalty to guard against over fitting, and C, a segment specific cost function, is twice the negative log likelihood:

$$C(y_{(t_{i+1}):s}) = 2(-\max_{\theta} \sum_{i=t+1} s \log f(y_i|\theta)) \tag{E.2}$$

Regarding the penalty, PELT algorithm is designed for penalty functions that are linear in the number of changepoints: $\beta f(m) = \beta m$. In practice, the common approach to select the penalty constant β, is to plot the segmentation estimated by the PELT algorithm for a given value, and then evaluate if the chosen penalty fits the data well. The underlying problem is that the choice of the penalty has a impact on the accuracy of the segment estimate (Haynes et al., 2014).

The R package for Changepoint Analysis, changepoint (Killick and Eckley, 2014), provides some standard penalties. If p denotes the number of additional parameters introduced by adding a constant, PELT could be applied with the following penalty constants: $\beta = 2p$ (AIC); $\beta = p \log n$ (SIC); $\beta = 2p2 \log n$ (Hannan-Quinn); and the theoretical type I error (Asymptotic). We tested all these penalty constants in all markets inspected, but all of them under-fitted the data. The tested penalties tend to find fewer segments than those shown in the time series, and, in some cases, even the sharp changes were not detected (Figure E.1). In other words, the standard penalties are too small for our application. For this reason, we implement the PELT algorithm along with CROPS. CROPS (Changepoints for a Range of Penalties), is a method to solve the same penalized minimisation problem than PELT, but using different values of the penalty, $\beta \in [\beta_{min}, \beta_{max}]$ (Haynes et al., 2014). Both algorithms can be implemented together in the changepoint package. The procedure to find the changes in position for each SKU sequence is described next.

Firstly, we applied the PELT algorithm along with the CROPS algorithm to the horizontal and vertical time series of each SKU. The horizontal position range of values is greater than the vertical one (20 and 2 meters, respectively), thus, in order to obtain a good fit, the interval $[\beta_{min}, \beta_{max}]$ was different for each analysis: we used $\beta_{min} = 1.5$ for the horizontal position sequences, and $\beta_{min} = 0.001$ for the vertical position sequences; for both sequences, we used $\beta_{max} = 50$. The output for each product is a matrix containing m optimal segmentations and the corresponding β values.

Secondly, we selected products that had a sharp change in its position as shown in Figure E.2(a), and choose the optimal β for those products. As we mentioned above, the process of selecting a penalty is an important part of the changepoint analysis. Nonetheless, is often assessed by visual inspection and there is still no standard method. For our analysis, we used the approach proposed by Lavielle (2005) and implement in Haynes et al. (2014), consisting in plotting a diagnostic as shown in Figure E.2(b). The main idea is that if a true changepoint is added to the model then the improvement in fit will be large. Then, we looked for the slope break in the diagnostic plot, which indicates the optimal number of changepoints.

Thirdly, we applied the PELT algorithm with the optimal β to each SKU's sequence. Thereafter, the exact time at which the change in position occurred as well as the segments before and after the change were detected. Figure E.2(c) illustrates the result of this process.

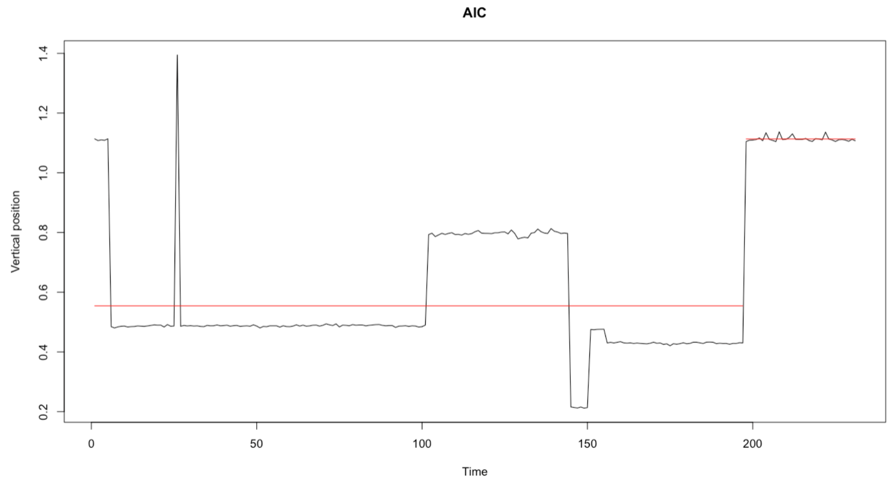

(a) PELT changepoints with AIC penalty.

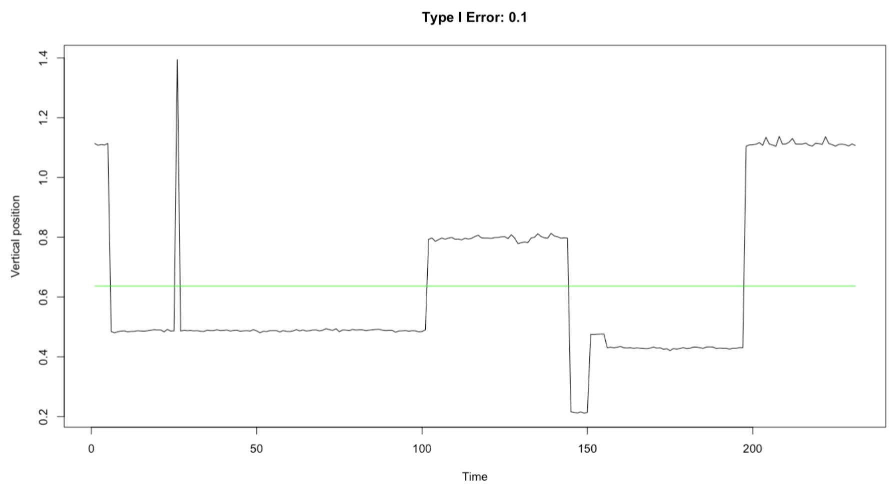

(b) PELT changepoints with Asymptotic penalty (Type I Error = 0.1).

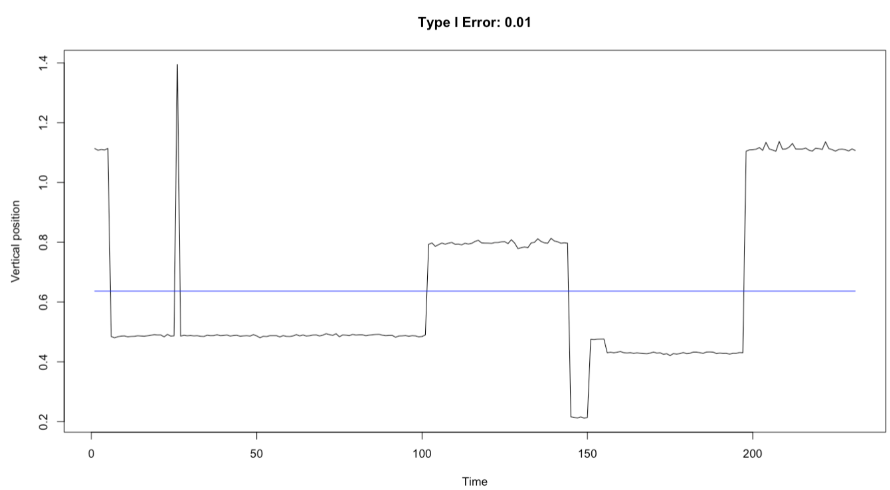

(c) PELT changepoints with Asymptotic penalty (Type I Error = 0.01).

Figure E.1: Tested penalties.

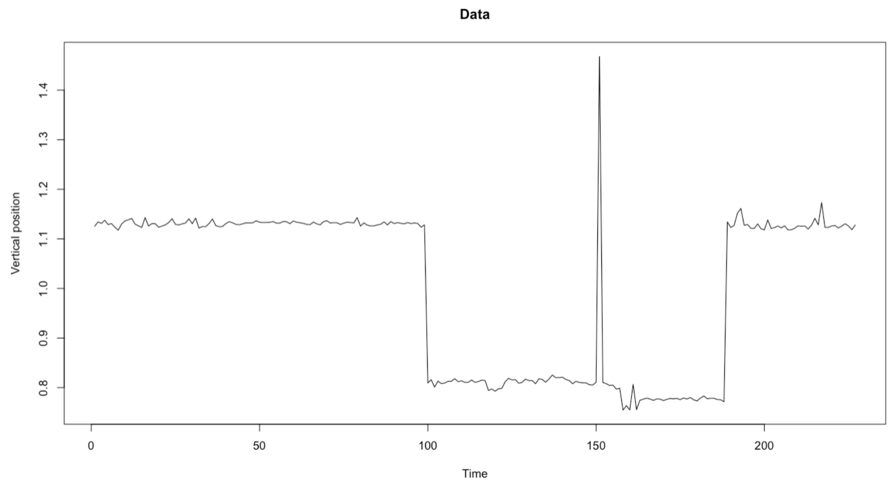

(a) Vertical shelf position sequence of a product.

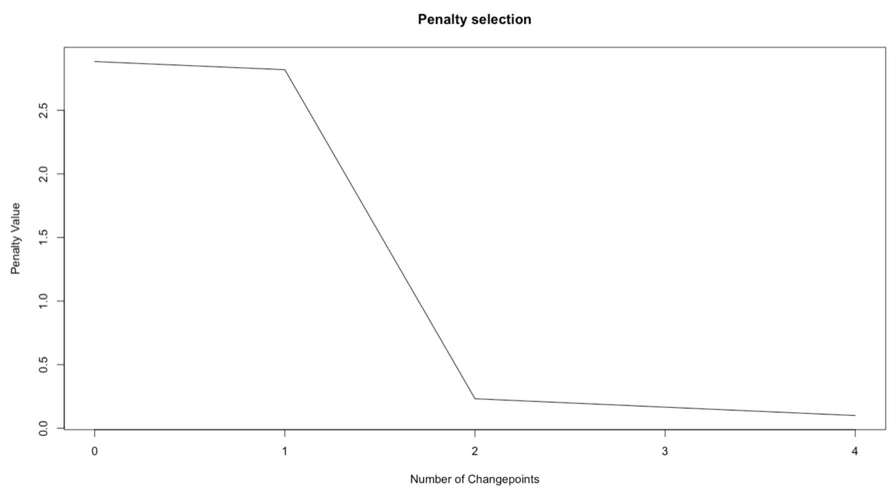

(b) CROPS diagnostic plot.

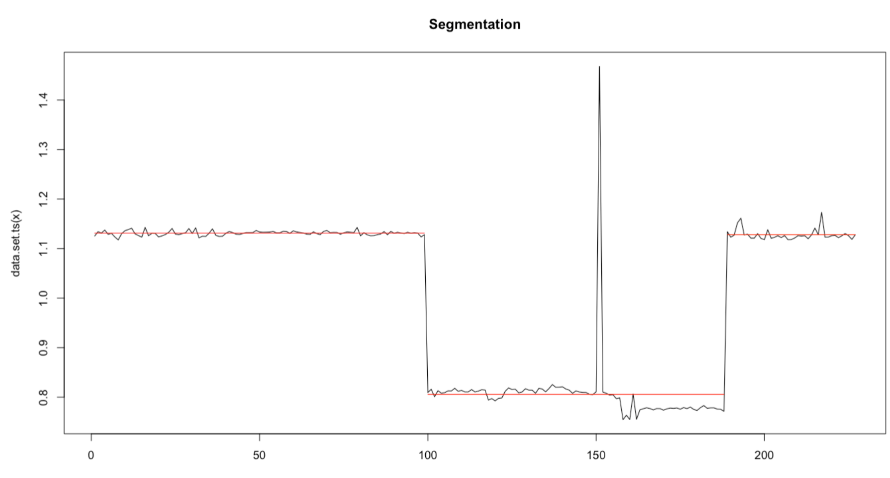

(c) Changepoints detected.

Figure E.2: PELT and CROPS methodology.

Appendix F

Difference in difference individual analysis

Before performing the difference in difference analysis, a naive regression was used to examine the 38 selected units separately:

$$\ln(1 + SALES) = \beta_0 + \beta_1 STORE + \beta_2 T \cdot t + \epsilon \tag{F.1}$$

Results are reported in Table F.1. The treatment was significant only in 9 cases: 2 at the 0.01 level and 7 at the 0.05 level. In 5 cases the treatment had a negative impact on sales, with magnitudes ranging between 0.299 and 0.657; the impact of the treatment on sales of the remaining 4 units was positive, with magnitudes going from 0.375 to 0.844. Cases with negative estimated coefficients are: a unit that came closer to the center of the aisle; 2 units that increased its vertical position; and 2 to units that decreased it. Cases with positive estimated coefficients are: 2 units with changes in its vertical position (increase and decrease); a unit that came closer to the center of the aisle: and a unit with a decrease in its vertical position and that came closer to the center of the aisle.

From all cases with significant treatment effects, 5 of them belong to the Sauces and Dressings market. Recall that this market is composed by products like mayonnaise, mustard, ketchup and chili sauces (Figure C.1). There are no clear patterns for this market either. Although the results of these cases were not conclusive, the small number of cases with significant treatment effects is a relevant insight.

Table F.1: DID previous results.

	Market	Type of change	Vertical change	Horizontal change	Vertical direction	Horizontal direction	T coefficient	p-value	Period length
1	Herbal infusion	Vertical	-0.31	0.00	Decrease	No change	-0.49	0.002**	10.00
2	Sauces and Dressings	Vertical	0.88	0.00	Increase	No change	-0.62	0.002**	14.00
3	Sauces and Dressings	Horizontal	0.00	-0.12	No change	Closer	-0.66	0.012*	5.00
4	Deodorant	Vertical	0.32	0.00	Increase	No change	-0.41	0.018*	15.00
5	Sauces and Dressings	Horizontal	0.00	-0.07	No change	Closer	0.84	0.032*	3.00
6	Sauces and Dressings	Horizontal/Vertical	-0.29	-0.14	Decrease	Closer	0.38	0.029*	5.00
7	Sauces and Dressings	Vertical	-0.30	0.00	Decrease	No change	-0.30	0.047*	15.00
8	Hair Care	Vertical	-0.27	0.00	Decrease	No change	0.69	0.029*	5.00
9	Deodorant	Vertical	0.28	0.00	Increase	No change	0.84	0.023*	4.00
10	Sauces and Dressings	Horizontal/Vertical	0.63	0.25	Increase	Farther	-0.14	0.512	12.00
11	Sauces and Dressings	Horizontal/Vertical	-0.31	-0.09	Decrease	Closer	0.01	0.967	8.00
12	Deodorant	Vertical	-0.30	0.00	Decrease	No change	0.16	0.484	9.00
13	Deodorant	Vertical	0.28	0.00	Increase	No change	-0.10	0.633	7.00
14	Deodorant	Vertical	0.28	0.00	Increase	No change	0.09	0.629	9.00
15	Dishwashing detergent	Vertical	-0.32	0.00	Decrease	No change	-0.40	0.12	5.00
16	Hair Care	Vertical	1.26	0.00	Increase	No change	-0.42	0.054	4.00
17	Deodorant	Vertical	0.30	0.00	Increase	No change	0.50	0.228	4.00
18	Deodorant	Vertical	0.30	0.00	Increase	No change	0.57	0.117	6.00
19	Sauces and Dressings	Horizontal	0.00	-0.23	No change	Closer	-0.03	0.821	9.00
20	Soaps	Horizontal/Vertical	-0.85	-0.41	Decrease	Closer	-0.42	0.079	3.00
21	Deodorant	Vertical	0.33	0.00	Increase	No change	0.21	0.363	10.00
22	Jam, honey and sweet	Vertical	0.23	0.00	Increase	No change	0.14	0.669	5.00
23	Sauces and Dressings	Horizontal/Vertical	0.29	0.04	Increase	Farther	0.27	0.128	14.00
24	Sauces and Dressings	Horizontal/Vertical	0.29	-0.12	Increase	Closer	0.06	0.666	14.00
25	Sauces and Dressings	Horizontal/Vertical	0.13	0.06	Increase	Farther	0.28	0.302	6.00
26	Sauces and Dressings	Horizontal/Vertical	-0.27	-0.06	Decrease	Closer	0.01	0.933	14.00
27	Sauces and Dressings	Horizontal	0.00	-0.07	No change	Closer	-0.25	0.311	9.00
28	Deodorant	Vertical	0.28	0.00	Increase	No change	-0.40	0.16	4.00
29	Red Class	Horizontal	0.00	-0.09	No change	Closer	-0.34	0.344	3.00
30	Deodorant	Vertical	-0.21	0.00	Decrease	No change	-0.47	0.171	4.00
31	Sauces and Dressings	Vertical	0.30	0.00	Increase	No change	0.07	0.602	15.00
32	Sauces and Dressings	Horizontal	0.00	-0.18	No change	Closer	0.14	0.57	5.00
33	Sauces and Dressings	Horizontal	0.00	-0.11	No change	Closer	0.66	0.389	4.00
34	Laundry detergent	Horizontal	0.00	0.21	No change	Farther	0.11	0.483	20.00
35	Laundry detergent	Horizontal	0.00	0.27	No change	Farther	-0.31	0.138	12.00
36	Hair Care	Vertical	0.59	0.00	Increase	No change	0.65	0.068	4.00
37	Deodorant	Vertical	-0.42	0.00	Decrease	No change	0.24	0.48	4.00
38	Deodorant	Vertical	-0.26	0.00	Decrease	No change	0.07	0.775	5.00