

The Value of Online Interactions for Store Execution

Felipe Caro¹ • Victor Martínez-de-Albéniz² • Borja Apaolaza³

Abstract

Problem definition: Omnichannel retailers interact with customers both online and offline. So far, they have used the richer information available to optimize the sales process by designing the right channel and supply chain structures, and by personalizing offer, pricing, and promotions. We advance an additional dimension of omnichannel value: retailers can use online clickstreams to better understand customer needs, and optimize store layouts to maximize *webrooming conversion*, which we define as the ratio of sales to webrooming activity. **Methodology/results:** We develop a model in which in-store purchases depend on the customer’s shopping list, and the effort required to locate and reach the products within the store. Category location in the store thus drives the likelihood of a sale. We then apply our model to a large home improvement retailer and find that shoppers’ preferences are revealed by nearby online traffic, and hard-to-reach locations lead to lower webrooming conversion. Finally, we optimize category-location assignments using our demand model and find that putting higher-interest and higher-price items in the most effective locations can increase revenues by about 2-5% in comparison to models that ignore online clicks. **Managerial implications:** We show how using online clickstream information for optimizing offline operations can create significant value. More fundamentally, our results provide a word of caution that in some retailing segments like home improvement, longer in-store paths might not necessarily be better.

Keywords: Shopping lists, shopping baskets, clickstream, webrooming, demand estimation, layout

Submitted: November 18, 2021. Revised: September 8, 2023; August 20, 2024.

1. Introduction

In the last decade, omnichannel has become a dominating retail strategy in which retailers do not see online and offline as independent channels, but manage them jointly (Gallino and Moreno 2019, Caro et al. 2020). Omnichannel delivers value on multiple dimensions, because it allows customers to learn about the product in one channel, and fulfill the demand in another (Bell et al. 2014). This flexibility implies that retailers are no longer constrained to run a single-channel sales process, and have more freedom to optimize the funnel from need to purchase (Wiesel et al. 2011). The additional flexibility requires closer coordination of the operations and marketing functions (Bijmolt et al. 2021), but has the potential to groom more effective interactions with the customer, increasing their satisfaction and delivering higher profits to the retailer.

The literature has identified different ways to extract value from omnichannel. On the one hand, traditional marketing actions can be refined with more precise customer histories, such as targeted advertising or promotions (Goic and Olivares 2019). On the other hand, many decisions in the operational realm have been improved. The design of channels can be optimized by better

¹UCLA Anderson School of Management, Los Angeles, CA 90095, felipe.caro@anderson.ucla.edu

²IESE Business School, University of Navarra, Av. Pearson 21, 08034 Barcelona, Spain, valbeniz@iese.edu.

³The Wharton School, University of Pennsylvania, Philadelphia, PA 19104, apaolaza@upenn.edu.

understanding how offline and online affect each other (Bell et al. 2020, Kumar et al. 2019, Bar-Gill and Reichman 2020). Information provision can drive the channel choices of the consumers (Gallino and Moreno 2014). Information from online sources can also help predict demand better so inventory levels can be optimized (Huang and Van Mieghem 2014, Cui et al. 2018). Finally, fulfillment flexibility allows firms to better run their supply networks (Hübner et al. 2019).

Most of the strategies described above are effective for firms that have a large online sales channel, but it is less clear how offline-heavy retailers can take advantage of an omnichannel strategy. Moreover, despite the increase of online shopping, retailers keep investing in stores as these remain the primary instance to interact with consumers (Schaverien 2018, Dowsett 2019). According to Bell et al. (2014), the value can be obtained by providing better information online, in a webrooming model. In this paper, we intend to uncover one additional value creation strategy available to omnichannel retailers, which is especially important under webrooming: one can leverage online sessions to detect (true) customer needs at the category level and analyze the determinants of store effectiveness. Namely, we are interested in what we call *webrooming conversion*, measured as the number of sales in the category divided by the number of potential customers that showed interest online in the store vicinity. Note that this definition is different from the common metric used by retailers, who divide sales by footfall, which is problematic because footfall need not be made of shoppers with an interest in the product. In contrast, our definition accounts for such interest, while ignoring the actual number of store visitors.

The conversion process is complex, as visitors’ initial shopping intention must translate into store visit first, then into exposure to the wanted products (and others), then into consideration sets, and finally, into purchase. It involves time and effort from the consumers. As a result, a well-thought layout can help them access their desired products quicker, and they might end up buying with a higher probability (Underhill 2009). Indeed, convenience increases the chances that customers buy: more formally, time pressure and higher search costs decrease sales (Hui et al. 2009b, Brynjolfsson et al. 2011). This is the reason why impulse items such as chocolates are often located near the check-out line, and Amazon has patented the One-Click button to reduce cart abandonment (Wagner and Jeitschko 2017). Unfortunately, the understanding of the relationship between layouts and conversion is limited. While richer displays—i.e., displays with more products on the shelves—are known to increase conversion (Boada-Collado and Martínez-de-Albéniz 2020), there is a lack of empirical evidence linking product position in the store with sales. Causal evidence of this kind is hard to obtain, because retailers generally do not know the store visitor’s shopping list, and hence they only observe sales performance of a particular store area but not how effective it was in capturing potential purchase intentions.

A rare exception is Hui et al. (2013) who collected the planned shopping list of 275 shoppers across 99 categories, and subsequently showed that longer in-store paths increased unplanned spending. However, this type of effort is hard to scale. As Goic and Olivares (2019) put it, “In

contrast [to online channels], data regarding browsing behavior in retail stores have been, for the most part, nonexistent. Studies that seek to measure the effect of changes in the layout and display of a store have typically used aggregate store-level data to conduct causal analysis.” In this paper, we provide one scalable novel way to assess the effects of layout on conversion, which can be used for layout optimization. Our approach selects the subpopulations of consumers that showed interest in a product category in the vicinity of a store, and then connects it to sales for that same category in that store. Hence, it does not require individual-level data and should be applicable to many retailers, even those that cannot track consumer behavior individually.

For this purpose, we first build a theoretical model in which conversion is affected by the physical effort invested by the visitor to locate products in her shopping list. We then work with a large home improvement retailer for which we observe, during seven months, all offline and online activities. For 16 stores, we observe full transaction records, i.e., composition of individual receipts, category details, and precise location within each store. For the online channel, we observe full clickstreams, i.e., all the clicks with timestamps by distinct geolocated origins of internet traffic. For each store and category, we are thus able to count how many different potential customers might be interested in the category. This is a proxy for the number of store visitors genuinely interested in purchasing the category, and we show that it is indeed a strong predictor of category sales. We are then in a position to study how conversion is moderated by location in the store. After controlling for other shopping funnel factors including store and category fixed effects, we find that the distance from the store entrance is a critical determinant of conversion, and items easier to reach—i.e., closer to the store entrance—exhibit significantly higher conversion. In contrast, we find that spillovers from adjacent categories are not significant (recall that these are home improvement categories for which there is little impulse shopping), which suggests that using store visits to create cross-selling revenue may not always be possible or desirable, as also suggested by Gao and Su (2017a).

The empirical findings pave the way for optimizing store layout. We formulate this question as an assignment optimization problem, and show that revenues can be increased by about 2-5% when online information is used to decide category locations, in comparison to a model in which only offline information about sales is used. This involves a one-time layout change.

We illustrate the value of our framework with an example. Suppose that the manager at one of our data partner’s stores notices that a prime location of the store is underutilized because it is being used by a category with meager sales, such as the smart-home connectivity category, which consists of devices that connect remote sensors to monitor the status of a house. Then, the question is what else to place in that prime location. If the manager only has access to store sales, then she might consider placing a steady seller, such as the locks category. However, if she also had access to clickstream data from shoppers living nearby, she would be able to detect categories that are “punching below their weight”, such as laminated flooring in our case study. Optimizing the location of these categories allows them to achieve their full conversion potential.

Our work contributes to the growing literature on retail analytics in two ways. First, we theoretically and empirically disentangle the sales impact of a store location, measured by its distance from-entrance-to-exit, and the market potential of a category, measured by online clicks. Second, we formulate and solve a layout optimization problem that exploits the dependence of location on store revenues, which provides a blueprint for retailers on how to execute a data-driven strategy to maximize the value of their real estate. Our results show that having access to product preference lists—available in online interactions—as opposed to simply shopping baskets—typical in store transaction records—is very valuable. Our approach is thus a simpler alternative to in-store customer tracking (Hui et al. 2009a), and more importantly gives access to information about which categories were browsed online in the vicinity of each physical store (see Chen et al. 1999 for a similar idea applied to advertising). Besides establishing the connection between store layout and sales, we provide an integrative perspective where customer behavior is combined with layout design decisions, which goes beyond minimization of average travel distance (De Koster et al. 2007) or consideration of category adjacencies (Ozgormus and Smith 2020). Finally, note that our prescriptive results are applicable in retail settings where impulse purchases and cross-selling are small, such as home improvement or auto parts stores, to name a few. In other contexts, reducing in-store paths may have unintended consequences because it may reduce unplanned spending (Hui et al. 2013).

The rest of the paper is organized as follows. Section 2 reviews the relevant literature. Section 3 formulates the model of a shopping visit and formulates the layout optimization problem. We estimate the impact of layout on sales in Section 4. Section 5 includes a counterfactual analysis of alternative store layouts. Section 6 concludes the paper.

2. Literature Review

Our work is mainly related to three streams of literature. First, we build on the operations and marketing literature that has studied the shopping funnel. Second, we are connected to the works about offline-online channel interactions. Third, we contribute to the literature on prescriptive models for retail execution.

2.1 From shopping funnels to omnichannel

The concept of funnel is a natural approach to study the effects of different marketing strategies on the customer. The funnel applies to both physical channels, where store visits are transformed into units sold, and online channels, where leads become visits which in turn generate orders. Wiesel et al. (2011) provides a framework to integrate both channels. Lemon and Verhoef (2016) provide an excellent perspective on the customer journey that a potential shopper goes through, and highlight pre-purchase stages—webrooming in our case, defined as “search online, buy in store”—as

determinant of later purchasing decisions. They also review the rich marketing literature on these interactions. In terms of modelling, hierarchical models are a convenient way to capture that only a fraction of those showing an initial interest in shopping end up making a purchase (Arora et al. 1998, Martínez-de Albéniz et al. 2020). Detailed decision processes have been developed, such as the use of consideration sets (Wang and Sahin 2018) or product evaluation heuristics (Aouad et al. 2021). It is worth noting that one paper has theorized about the importance of order within the basket. Chen et al. (1999) highlights that some categories are more important than others as they are the reason behind a store visit. They develop the concept *marketing profits* to reflect that profit should be attributed to the category that brought the customer to the store.

When pre-purchase and in-purchase stages take place in different channels, the term “omnichannel” has frequently been employed. The phenomenon has been extensively studied in the last decade. Brynjolfsson et al. (2013) provides an early discussion of the potential of omnichannel for retailers. Gao and Su (2017a,b) develop analytical models for channel choice under omnichannel capabilities. As conceptualized in Bell et al. (2014), the benefits of omnichannel come from better category information, and from better fulfillment possibilities. In other words, there are advantages in showrooming, and in webrooming.

Physical interactions in the store allow retailers to engage with customers more effectively. Bell et al. (2020) show how the convenience and the store experience can help pure online players sell more. Kumar et al. (2019) identify the possibility of making in-store returns as another driver of sales increases.

Webrooming can also be valuable. Gallino and Moreno (2014) study the effect of Buy Online Pickup in Store (BOPS) on online and offline sales, and find that store traffic increases due to better information about in-store category availability. Interestingly, the quality of the online experience has an impact on offline sales and customers’ overall perception (Bar-Gill and Reichman 2020, Flavián et al. 2020).

Our study reveals a different value driver of omnichannel. It can be used to anticipate demand at the store level, and hence, study the impact of the store layout on the conversion process.

2.2 Path studies

An interesting variation of the study of shopping funnels is possible when interactions between customer and firm occur multiple times, requiring us to consider the sequence in which they occur. Sequential decisions have been considered in the marketing and economics literature across visits (Chintagunta et al. 2012) or within a single visit (Larson et al. 2005, Hui et al. 2009a,b, 2013, Ruiz et al. 2020). Hui et al. (2009a) provides a review of marketing research that considers paths of consumers in store settings. Ruiz et al. (2020) include memory effects as well as one-step forward considerations. This literature suggests that consumers rarely shop the entire store, and quite

possibly come with a plan in mind on how to traverse the store.

Closer to our context is Hui et al. (2009b), who model the path between supermarket categories, using a conditional model in which transitions from category to category are driven by destination characteristics and path history, which they validate with path data in one store. The data is obtained from individuals moving through the store, tracked by RFID tags, and the model is calibrated using the actual transitions that took place. Interestingly, store layout is included in the decision process through the distance between store regions. We use a similar approach, with some differences. First, we do not have individual path data, but have access to category-level aggregates. Second, we focus on conversion from needs to purchases, so our consideration of cross-category interactions is operationalized in conversion spill-overs, between needs for one category and sales for another, which is moderated by the location of categories in the store. Third, we see variation of category locations across different stores, which allows us to control for category characteristics and separately measure the impact of category location.

Also very close to our approach is Hui et al. (2013), who document that in-store path length strongly increases unplanned purchases, in a grocery retail setting. In this case, RFID tracking was complemented by an entry survey that elicited the customer’s shopping list, which allowed the authors to instrument the endogenous in-store path length. In addition to this study, a field experiment using in-store promotions was conducted to validate that longer paths increased unplanned spending. In comparison, we do not find evidence of any strong cross-category spill-over effects, which may be due that our context is home improvement in which impulse shopping is less salient. This suggests that our prescriptions, which seek to reduce in-store distance to increase conversion, may need to be revisited in other retail contexts in which impulse shopping is important.

Finally, while there is a lack empirical evidence linking store layouts with sales, extensive work has been done to connect the position on a shelf with sales. Chandon et al. (2009) is a classic study on how in-store marketing decisions shape consumer behavior. More recently, Jalali et al. (2023) partner with a convenience store chain to conduct a field experiment that allows them to empirically estimate the effect of vertical product location on sales. They find that, on average, the eye-level position generates higher sales relative to the bottom- and top-shelves, with the magnitude of this effect depending on product characteristics and how other products are laid out vertically. Based on these insights, they write and solve a planogram optimization problem for their retail partner, in the spirit of Corstjens and Doyle (1981). Hübner et al. (2019) tackle the joint challenges of limited shelf space and in-store replenishment constraints, combining shelf space allocation, assortment and replenishment decisions. The shelf space allocation problem also has implications on the supply chain and supplier pricing, as shown by Martínez-de Albéniz and Roels (2011) or Heese and Martínez-de Albéniz (2018).

2.3 Optimization for retail execution

Our work is also connected to works that develop models to improve retail execution. Interventions have focused on different dimensions, which we briefly review. Perdikaki et al. (2012), and Mani et al. (2015) measure the impact of staffing levels on sales, which Chuang et al. (2016) use to develop a labor planning methodology. Caro and Gallien (2010) study inventory distribution across stores and combine demand forecasting and inventory allocation optimization to improve sales at Zara; Gallien et al. (2015) apply a similar approach to new product distribution. Inventory inaccuracy is another cause of suboptimal retail performance. DeHoratius et al. (2008) measure the extent of inaccuracies and DeHoratius and Raman (2008) use inventory replenishment and audits to mitigate their effects. Montoya and Gonzalez (2019) develop a hidden Markov chain model to predict phantom stock-outs based on sales time-series. The effect of store congestion has also been explored: Lu et al. (2013) measure how queues reduce sales conversion.

We discuss here an understudied aspect of retail execution. Indeed, we are not aware of existing work that studies the role of store layout in generating sales. In particular, Larson et al. (2005), Hui et al. (2009b), and Hui et al. (2013) do not investigate store design because they only have data on one store and hence cannot disentangle the effect of category location from the category itself. In contrast, we study layout decisions. The design of a store layout resembles that of designing a warehouse. There exists a broad literature on warehouse layout optimization, see De Koster et al. (2007) for an excellent review. Usually, the design problem is formulated as a large integer program that is solved with heuristic techniques. The methods have also been applied to store layout design, e.g., Mowrey et al. (2018). In these models, customer behavior is integrated through simplified customer behavior assumptions such as considering penalties for categories that are not adjacent (Ozgormus and Smith 2020). In contrast, we use the moderating effect of location on conversion to propose improved layouts.

3. Model

3.1 The Shopping Process

In the same vein as the shopping funnel discussed in Section 2, we make the following assumptions for the shoppers in our model:

1. Store choice: consumers prefer buying at a store that is closest to where they live. Hence, each store has a “natural catchment area” that consists of all the households within a certain radius.
2. Shopping lists: a significant fraction of consumers start their purchasing process with a prioritized list of items in mind that they would like to buy or are considering buying. A list can

have multiple categories—e.g., paint, indoor lamps, and curtains—or just a single one, e.g., exhaust fans. Categories that are more important to the consumer are held higher in the list.

3. Webrooming: a significant fraction of consumers does research online about the products of interest on the retailer’s website, and then follow through by visiting the store to purchase (some of) the items they researched online. A consumer’s (mental) shopping list dictates the order in which they search the items on the retailer’s website. The first item on the shopping list can be understood as the “lead category” for that given consumer (Chen et al. 1999).
4. Store sales moderated by effort: once at the store, consumers try to purchase all the items on their shopping list but might give up on some if they run out of time or if they are not willing to exert the necessary effort to find and fetch the item.

This sequential funnel makes some assumptions regarding customer behavior. First, it requires, implicitly, that consumers highly value their time, so they make their store choice based on proximity, and limit their willingness to shop to fill functional needs, thereby disregarding potential impulse purchases that would require extra effort for a small additional utility. This assumption is reasonable in many retail settings, including home improvement, in which all stores are alike and that carry most of the categories offered online.

Second, we ignore competing stores. Note that we are not assuming that consumers are captive to a particular store, but rather that households are representative of the demand faced by the neighboring stores, even if they do not necessarily shop there.

Third, the shopping list assumption can be justified in retail settings where choices are made before entering the store. Hence, there is prior choice set that is mostly unaffected by the layout. This assumption is consistent with choice models where each customer has a preference list. In the literature, these categories are substitutes and the customer ends up buying a single, preferred category out of the available ones. In our context, we extend this view to consider a preference list of complementary categories, so this can be interpreted as a shopping list.

Fourth, the webrooming assumption is based on a common pattern observed in omnichannel retailing. In fact, industry reports show that the percentage of shoppers doing online research prior to visiting the store can range from 69% to 88% (Accenture 2013, Harris 2013, Deloitte 2017).

Finally, the moderation effect that effort has on sales is justified by the value of time premise. This assumption is consistent with behavioral models in which consumers have a time budget for in-store purchases (such as groceries as shown in Hui et al. 2009b), and is more amenable to functional categories such as home improvement, for which the time spent enjoying the store experience is not a major driver of conversion. In addition, and relatedly, we assume that shoppers know where items are located in the store and walk the shortest path. Based on the assumption that consumers value their time, they will try to find the shortest path to fetch the items they want, and in that spirit we use the shortest path as a proxy for the distance they traverse in the stores.

Note that in our context, basket sizes are actually smaller during week-ends, when time budgets should be more generous, thus consistent with consumers highly valuing their time. This pattern is also consistent with professional home-builders making high-ticket purchases during the week and amateur DIY-ers making smaller purchases on weekends. This may not be the case in other retail contexts where impulse spending is high, such as fashion retailing. This provides the boundary conditions for our model and empirical findings: we expect them to carry over to retail contexts in which the priority is facilitating the conversion from explicit potential needs into purchases, regardless of the geographical location (our study uses data from Chile but the results should be applicable to the US or any other country). This includes home improvement of course, but also electronics and white goods. On the contrary, it probably does not hold for retailing with a higher impulse content such as fashion or groceries.

3.2 Empirical Approach

Our empirical approach is based on the assumptions presented in the previous section. Conceptually at a high level, it has the following form:

$$sales_{ist} = \alpha_{m_i s} + \alpha_t + f(online_visits_{ist}, effort_{is}) + e_{ist}, \quad (1)$$

where m_i denotes the macro category that contains category i . In our context, the macro category m_i corresponds to level 0 in the product hierarchy—to be introduced in Section 4.1—which is typically a product family such as lighting, and it includes multiple (level-1) categories like indoor lamps, outdoor lamps, light bulbs, etc., that are indexed by i .

The dependent variable $sales_{ist}$ should be considered in log form, so as to justify an additive structure of independent drivers and to better fit the empirical distribution which is Bell-shaped (this is not the case without applying the log-transformation). The terms $\alpha_{m_i s}$ and α_t correspond to macro category-store ($m_i s$) and time (t) fixed effects, which represent the baseline demand. Note that the interaction of macro category and store will capture any specificity that the local market may have with respect to a product family. The next term amplifies demand as a function of webrooming moderated by effort, through a generic function $f(\cdot)$ that increases with online visits and decreases with effort. Here, $online_visits_{ist}$ represents a vector of relevant metrics that characterize online traffic, and $effort_{is}$ should capture the time (disutility) involved in finding category i at store s , based on the store layout. Note that the latter excludes the fixed time/cost it takes to arrive to the store, which would be captured by the macro category-store fixed effect. Finally, e_{ist} is the usual error term.

A few more remarks are noteworthy. We consider two amplification components in Equation (1). Namely, (i) primary demand: people that came to the store with the intention of buying, and exerted the effort to find the category; and (ii) secondary demand: people that came to the store

searching for something else, but got exposed to the category and ended up buying (spillover in path, spillover nearby, spillover within aisle). Both components should be captured by $online_visits_{ist}$.

We can observe that there are no substitution effects included in Equation (1). This formulation is appropriate when different categories are solutions to non-overlapping functional needs. Our empirical analysis is performed at the category level, with 165 different ones. At this level, substitution effects across categories should be negligible.

The effect of store execution is captured mainly by the macro category-store fixed effects. This includes the impact of assortments, which are quite stable over time, staff intervention, which are limited because these are large stores with emphasis on self-service, and display—also stable because there are no changes in layout within a store. In particular, all stores may have different assortment breadths, but service level is extremely high, meaning that there is always at least one product with available stock within each category, i.e., there is large variety every day during the entire period. Moreover, all the results are unchanged when restricting our attention to categories in which the service level is 100%. In other settings, such as apparel or groceries, it may be necessary to include inventory levels in Equation (1) to control for potential demand censoring when there are stockouts (Boada-Collado and Martínez-de-Albéniz 2020), but in our case it is not needed.

Finally, promotional activities may be an important control to include, but these tend to be the same in all stores, and hence, the effect of promotions is absorbed by the fixed effects and cannot be identified.

4. Application to Home Improvement Retailing

4.1 Context

We collaborated with a South American chain of home improvement stores, a leader in this industry, which operated 60 stores across Chile and an online channel at the time of the collaboration. We obtained a comprehensive proprietary dataset providing information about stores, categories and customer interactions, which we describe below.

The retailer sells a variety of home improvement categories, such as tools or materials. For the sake of illustration, the items in the assortment belong to categories such as paint, sawn timber, gardening tools, roofing, electric extension cords, or interior car accessories, to name a few. The same assortment is sold in stores and online. In the categories available in the data, we list 82,178 SKUs that are categorized in different hierarchical levels in the following manner: 5 level-D clusters, 21 level-0 clusters (our macro categories denoted m_i), 168 level-1 clusters (our category focus indexed by i), 787 level-2 clusters, and finally other more fine-grained clusters. We select for our analysis 165 level-1 categories, after removing 3 legacy categories with zero sales.

At this retailer, the weight of the online channel is small, as it is responsible for only 2.63% and

6.20% of total receipts and revenues, respectively. At the same time, in this industry webrooming is known to be an important factor affecting the shopping process; for instance, Home Depot states that it influences about 60% of store purchases even though the online channel only contributes to 6% of sales (Digital Commerce 360 2017). Because categories are functional and product research is typically done in advance, this seems to be the ideal setting to assume that customers build a shopping list before entering the store, and to empirically connect online browsing to purchases.

Three types of data are available to us, which reflect customer behavior in online and offline channels:

- *Transaction data.* It describes the subset of the assortment’s categories that are purchased together. Each product bought belongs to a receipt, which is assigned to a physical store and a date. We refer to the purchase data as shopping-cart or shopping-basket data hereafter. From the raw information, we compute how many receipts issued by a certain store in a certain date included products of each category.
- *Clickstream data.* It describes the online journey that potential customers navigate when visiting the retailer’s website. It consists of time-stamped observations of category-level visits, with an IP address identifier (totalling 3,691,442 different identifiers). We refer to the clickstream data as shopping-list data or webrooming data hereafter. To process clickstream data, we first define a session as the web journey that a potential customer (given by an IP identifier) navigates in one day, i.e., at any time within a given date. One session is formed by a list of ranked categories, represented in a ranked vector. IP identifiers are geolocated, so we are able to associate each session with stores nearby. Specifically, the catchment area of a given store is a 5km radius for stores in the Santiago Metropolitan area and 20km elsewhere. There are two main reasons why we use different catchment areas. First, because Santiago is more densely populated than other regions in Chile, one could expect the concentration of home improvement stores to be higher there, what makes it unlikely that someone in Santiago will not find a home-improvement store within a few kilometers and will need to travel longer distances. In contrast, home improvement stores will not be as numerous in more sparsely populated areas of Chile, so potential consumers will need to travel further to their nearest store. Second, the cost of travel is higher in densely populated areas like Santiago as opposed to other areas in the country, what will make the people that live in Santiago less willing to travel a longer distance to a home improvement store.

There is one special IP identifier that is worth mentioning: it corresponds to a gateway assigned to all wireless connections from mobile networks. Despite this point being geolocated in Santiago, it comprises all the mobile connections that originate in Chile. For this reason, mobile traffic cannot reliably be assigned to a nearby store, so it is left out of our analysis. This traffic represents only 0.158% of all clicks (we replicate the analysis allocating mobile visits to

the stores proportionally, and results remain unchanged). Furthermore, we do not consider web visits that are thought to be generated by bots through web-scraping. To remove those visits, we filter the visits by those IP identifiers that either visit one category more than 40 times or visit more than 300 products in a given day (we also replicate the analysis including these visits, and results again remain unchanged).

- *Store layout data.* It describes the layout of the store, i.e., it details each level-1 category’s location in each store. The layouts of 16 brick-and-mortar stores are available in pdf files. We process these files automatically and we obtain the locations of the category labels within the layouts. These locations are described in (x, y) coordinates, and measured in pixels, but for each file, the scale conversion is available, through the width of checkout corridors which measures 1.65 meters. Hence, we can compute the distance in meters that a potential customer has to walk in the retailer’s store, so as to buy from a category. From this map, we can thus compute the distance between categories and between a category, the store entrance and the checkout lanes. We use Manhattan distances in meters, so as to reflect the true walking distance given the existence of horizontal and vertical aisles in the stores.

Given the information about layouts, we focus our study on 16 of the retailer’s brick-and-mortar stores (26.6% of the total), and its online channel. From these stores, nine are located in the Santiago Metropolitan area, while the remaining stores belong to other regions. We use daily data from December 1st, 2018 to June 30th, 2019, with the exception of 19 days that were removed from the analysis due to missing values. The total study period is thus 30 weeks long (we replicate the analysis excluding holiday dates, and again results remain unchanged).

Tables 1 and 2 compare stores and category metrics for these 16 stores compared to the entire network. We observe that stores included in our subsample are slightly larger in scale (number of receipts, size, assortment) but similar to the rest regarding basket composition (price, basket size), hence suggesting that no bias is introduced by focusing on our chosen store subset. Thus, the 16 in-sample stores are representative of the average store of our retail partner.

4.2 Descriptive Statistics

In this section we operationalize the variables from our conceptual Model (see Section 3). The data is aggregated weekly to avoid within-week fluctuations: each observation corresponds to a week t , a level-1 category i (which we call category for simplicity), and a store s . Hence, we define the following variables of interest:

- N_{ist} : Number of receipts that include category i issued at store s during week t .
- N_{st} : Total number of receipts issued at store s during week t . Note that $N_{st} \leq \sum_i N_{ist}$ because a receipt may include multiple categories.

Table 1: Store descriptive statistics, 30 week sample.

Statistic	<i>In-sample stores</i>						<i>Out-of-sample stores</i>					
	N	Mean	St. Dev.	Min	Median	Max	N	Mean	St. Dev.	Min	Median	Max
General												
Total revenue (mn. Chilean Pesos)	16	17,581	6,230	6,133	17,861	31,099	44	12,863	5,405	5,034	11,474	29,020
Total number of receipts	16	437,020	131,293	181,833	456,678	653,842	44	319,913	107,947	101,343	315,076	559,126
Total number of online sessions (mn.)	16	1.115	3.604	0.048	0.235	14.620	44	0.945	3.106	0.009	0.178	14.837
Average basket value (Chilean Pesos)	16	39,977	6,007	26,674	38,988	49,203	44	39,971	7,550	26,058	39,173	62,756
Average basket size	16	3.25	0.34	2.76	3.25	4.14	44	3.18	0.25	2.68	3.21	3.61
Layout												
Store size (square meters)	16	11,825	3,012	7,040	11,796	18,461	44	9,792	2,506	5,000	9,363	15,080
Number of aisles	16	79.6	19.0	50	81.5	123	-	-	-	-	-	-
Number of checkout lanes	16	17.3	4.7	8	17	27	-	-	-	-	-	-
Average distance on avenue (meters)	16	29.75	8.99	15.08	31.65	52.55	-	-	-	-	-	-
Average distance on aisle (meters)	16	18.63	4.57	11.32	19.46	25.65	-	-	-	-	-	-
Inventory												
Number of distinct SKUs carried	16	31,258	4,236	23,892	32,273	36,024	44	25,864	5,012	13,306	24,830	38,093
Number of distinct SKUs sold	16	27,860	4,232	20,437	28,838	33,400	44	22,755	4,676	11,238	21,807	33,640
Number of distinct categories carried	16	163.6	1.0	162	164	165	44	161.5	6.3	122	162.5	165
Number of distinct categories sold	16	161.5	1.3	159	162	165	44	161.8	5.9	125	163	165
Prices												
Average SKU price (Chilean Pesos)	16	4,534	1,180	2,154	4,668	6,765	44	4,653	1,472	2,667	4,619	12,899

Table 2: Category descriptive statistics, 30 week sample.

Statistic	<i>In-sample stores</i>						<i>Out-of-sample stores</i>					
	N	Mean	St. Dev.	Min	Median	Max	N	Mean	St. Dev.	Min	Median	Max
General												
Average revenue, per store (Chilean Pesos)	165	953,587	2,119,414	5,110	454,995	24,727,340	165	890,978	2,487,860	8,772	367,424	30,012,816
Average number of receipts, per store	165	5,721	6,889	1.3	3,128	43,929	165	4,163	5,097	3.2	2,206	32,031
Average number of online visits, per store	165	561.24	623.62	0.03	367.66	4,337.30	165	456.91	506.75	0.03	300.63	3,520.34
Inventory												
Average number of distinct SKUs carried, per store	165	187.28	223.92	1.36	107.63	1,295.5	165	155.37	183.52	1.486	93.16	1,078.5
Average number of distinct SKUs sold, per store	165	165.76	209.37	0.25	89.69	1,418.81	165	135.50	171.61	1.00	75.25	1,223.39
Prices												
Average SKU price (Chilean Pesos)	165	17,621	28,389	296	7,783	203,485	165	17,089	26,773	291	7,485	151,007

Note: In some categories, there are SKUs that are sold without presenting a positive stock in the inventory data.

- $V_{1,ist}$: Number of online sessions in which category i is viewed as the first item, within the catchment area of store s during week t .
- $V_{2-4,ist}$: Number of online sessions in which category i is viewed as the second, third, or fourth item, within the catchment area of store s during week t .
- $V_{5+,ist}$: Number of online sessions in which category i is viewed as the fifth item or further, within the catchment area of store s during week t .
- V_{ist} : Number of online sessions in which category i is viewed in any order, within the catchment area of store s during week t . It follows that $V_{ist} = V_{1,ist} + V_{2-4,ist} + V_{5+,ist}$.
- V_{st} : Total number of online sessions within the catchment area of store s during week t . It follows that $V_{st} = \sum_i V_{i,ist}$. Note that $V_{st} \leq \sum_i V_{ist}$ because a session may include multiple categories.
- V_{it} : Total number of online sessions in which category i is viewed in any order during week t . It follows that $V_{it} = \sum_s V_{ist}$.
- D_{is} : Distance to pick item i in store s measured in meters, i.e., the distance between the store entrance and category i plus the distance between category i and the checkout lanes. Note that there are no changes in layout during the time window of study, so this variable is independent of time.

In our study, we use all variables except distance in log form for ease of interpretation of the coefficients and to remove skewness, i.e., we transform variable X into $x := \log(1+X)$ (we add one to avoid problems with zero values of X). With this notation, the variable n_{ist} is our proxy for $sales_{ist}$, and D_{is} is our proxy for $effort_{is}$. Our proxy for $online_visits_{ist}$ includes $v_{ist}, v_{1,ist}, v_{2-4,ist}, v_{5+,ist}$, and might also include these same variables for other categories j whose traffic is relevant to the sales of category i .

Table 3 contains the descriptive statistics of the logged variables, and Table 4 their correlations. One can observe that the amount of generic online traffic v_{st} has a small correlation with sales indicators n_{st} or n_{ist} . However, category-specific online traffic $v_{it}, v_{ist}, v_{1,ist}, v_{2-4,ist}$ and $v_{5+,ist}$ has a high positive correlation with category-specific sales n_{ist} . This indicates that indeed online activity can be used as a key input for store demand forecasting, and this insight is a promising starting point to develop a more sophisticated model as discussed in Section 3.

To further illustrate the available data, Figure 1 shows the joint evolution of V_{ist} and N_{ist} for two stores and two categories. We can see that both series tend to move together, although their relative values (i.e., their ratio) changes across stores and categories, which is natural given that some categories may require relatively more browsing to achieve a certain level of sales, and the customers

Table 3: Descriptive statistics of the main model variables.

Statistic	N	Mean	St. Dev.	Min	Pctl(25)	Median	Pctl(75)	Max
n_{ist}	78,750	4.13	1.85	0.00	3.18	4.43	5.51	8.71
n_{st}	78,750	9.28	1.55	0.00	9.23	9.52	9.82	10.45
$v_{1,ist}$	78,750	2.96	1.41	0.00	1.95	3.00	3.95	7.77
$v_{2-4,ist}$	78,750	2.76	1.34	0.00	1.79	2.77	3.71	7.59
$v_{5+,ist}$	78,750	1.85	1.25	0.00	0.69	1.79	2.71	6.63
v_{ist}	78,750	3.73	1.41	0.00	2.83	3.78	4.72	8.42
v_{st}	78,750	8.66	0.84	6.74	8.01	8.78	9.31	10.24
v_{it}	78,750	5.88	1.37	0.00	5.17	5.98	6.78	9.79
D_{is}	78,750	119.09	51.80	28.55	77.77	114.13	151.40	377.37

Table 4: Correlation matrix between the variables of interest.

	n_{ist}	n_{st}	$v_{1,ist}$	$v_{2-4,ist}$	$v_{5+,ist}$	v_{ist}	v_{st}	v_{it}
n_{ist}	1							
n_{st}	0.4018 ***	1						
$v_{1,ist}$	0.3999 ***	-0.0386 ***	1					
$v_{2-4,ist}$	0.3015 ***	-0.0590 ***	0.8798 ***	1				
$v_{5+,ist}$	0.2209 ***	-0.0672 ***	0.7434 ***	0.8793 ***	1			
v_{ist}	0.3534 ***	-0.0527 ***	0.9499 ***	0.9669 ***	0.8700 ***	1		
v_{st}	-0.0074 **	-0.0618 ***	0.5400 ***	0.6277 ***	0.6227 ***	0.6178 ***	1	
v_{it}	0.4902 ***	-0.0234 ***	0.8165 ***	0.7036 ***	0.5616 ***	0.7746 ***	0.0773 ***	1

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

around some stores may have higher natural conversion between browsing and purchasing, compared to others. These structural, static differences will be captured by the macro category-store fixed effects in our model.

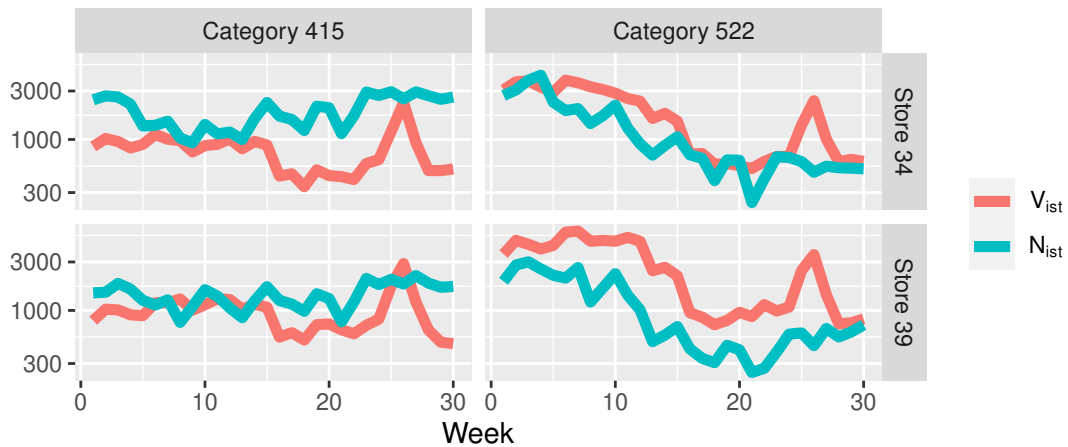


Figure 1: Evolution of clickstream and sales figures, for two categories and two stores.

As shown in Table 4, n_{ist} and v_{ist} are highly correlated. Hence, clickstream activity seems a useful lead indicator of category i 's performance. Taking into account this relationship, we can further study the impact of the store layout on conversion. Figure 2 plots conversion for a given store, measured as $n_{ist} - v_{ist} = \log((1 + N_{ist})/(1 + V_{ist})) \approx \log(N_{ist}/V_{ist})$, averaged over 30 weeks. We observe that, while conversion fluctuates, we see a clear trend showing that the conversion of distant categories is lower than those near the entrance or center. This model-free evidence suggests that a category's location in the store strongly affects the conversion from category interest to actual sales.




Figure 2: Model-free representation, with sales (left) and webrooming conversion (right) in quintiles. The arrows correspond to store entrances and exit. The layout on the left also shows the current and proposed location for the laminated flooring category that is discussed in Section 5.2.

4.3 Model Specification

Based on the variables of interest defined above, we write Equation (1) as the following regression model:

$$n_{ist} = \alpha_{m_{is}} + \alpha_t + \beta v_{ist} + \gamma D_{is} + e_{ist}. \quad (2)$$

Since the data does not contain any changes in store layouts, our main analysis aggregates over time to make it clear that identification comes from cross-sectional variation (i.e., variation across stores). In Section 4.6 we report the time-dependent analysis and show that the results remain the same. Hence, we aggregate Equation (2) across time and obtain

$$\bar{n}_{is} = \alpha_{m_{is}} + \beta \bar{v}_{is} + \gamma D_{is} + \epsilon_{is}, \quad (3)$$

where $\bar{n}_{is} = \frac{1}{T} \sum_{t=1}^T n_{ist}$, and $\bar{v}_{is} = \frac{1}{T} \sum_{t=1}^T v_{ist}$. This equation preserves the logic of Equation (2) and represents our main empirical specification. We also consider variations of this specification, in

which we incorporate quality-segregated online visits, via $\bar{v}_{1,ist}$, $\bar{v}_{2-4,ist}$, $\bar{v}_{5+,ist}$ instead of \bar{v}_{ist} , where these metrics are computed as averages of the corresponding time-varying variable.

Furthermore, this equation captures the direct influence where sales are simply the consequence of a true demand need existing prior to a store visit. Beyond this direct influence, the literature has identified other indirect influences, namely spill-over effects between categories. In other words, if there is a flow of shoppers interested in buying a certain category, these visitors will be exposed to other categories on their way to their primary shopping objective. For instance, Hui et al. (2013) documents a significant positive spill-over in a grocery context.

To study these cross-category interactions, we consider three potential drivers of sales arising from spill-over effects. First, for a certain category i , we consider the primary demand associated with categories j that require the shopper to walk by i in their path to j . For this purpose, we define the binary variable $INPATH_{ijs}$, which is equal to one if the shortest path from entrance to j and then to exit coincides with either the shortest path from entrance to i to j to exit, or from entrance to j to i to exit. Otherwise, it equals zero. For the paths to coincide, they must have the same Manhattan distance with up to a 10% deviation. We then define $PATH_{ist}$ as the number of online sessions within the catchment area of the store that include any category $j \neq i$ such that $INPATH_{ijs} = 1$:

$$PATH_{ist} = \sum_{j \neq i} INPATH_{ijs} \times V_{jst}, \quad (4)$$

and let $path_{ist} = \log(1 + PATH_{ist})$, and $\overline{path}_{is} = \frac{1}{T} \sum_{t=1}^T path_{ist}$. This variable should thus capture spill-overs into items that are in central locations within the store, that see a high amount of traffic for primary items that are further inside the store. It results in the following specification:

$$\bar{n}_{is} = \alpha_{m_{is}} + \beta \bar{v}_{is} + \gamma D_{is} + \delta \overline{path}_{is} + \epsilon_{is}. \quad (5)$$

One may note that visitor flows in the store might not necessarily coincide with the shortest path, since it is common that visitors prioritize the main ‘avenues’ within the store to move faster. Moreover, as visitors move faster, it is possible that they will pay less attention to other products, in contrast to when they move slower in the smaller aisles closer to the product that they are looking for. Hence, we should focus on more localized traffic. For this reason, we consider a second potential spill-over effect, with two specific forms. We consider the primary demand of categories j in the vicinity of i , defined through the binary variable $NEARBY_{ijs}$ which is equal to one if the distance between i and j is less than 20 meters. We then define $NEAR_{ist}$ as the number of online sessions within the catchment area of the store that include any category $j \neq i$ such that $NEARBY_{ijs} = 1$:

$$NEAR_{ist} = \sum_{j \neq i} NEARBY_{ijs} \times V_{jst}, \quad (6)$$

and let $near_{ist} = \log(1 + NEAR_{ist})$, and $\overline{near}_{is} = \frac{1}{T} \sum_{t=1}^T near_{ist}$.

Alternatively, we define the binary variable $SAMEAISLE_{ijs}$ which is equal to one if categories i and j are located in the same aisle in store s . We then let $AISLE_{ist}$ as the number of online sessions within the catchment area of the store that include any category $j \neq i$ such that $SAMEAISLE_{ijs} = 1$:

$$AISLE_{ist} = \sum_{j \neq i} SAMEAISLE_{ijs} \times V_{jst}, \quad (7)$$

and let $aisle_{ist} = \log(1 + AISLE_{ist})$, and $\overline{aisle}_{is} = \frac{1}{T} \sum_{t=1}^T aisle_{ist}$.

Both \overline{near}_{is} and \overline{aisle}_{is} variables capture spill-overs related to proximity to store hot spots, and lead to model specifications similar to Equation (5) in which the \overline{path}_{is} variable is replaced by either \overline{near}_{is} or \overline{aisle}_{is} .

4.4 Estimation

The distance variable D_{is} exhibits cross-sectional variation because stores have different sizes and shapes. We could estimate Equation (3) using ordinary least squares (OLS), but one may be concerned about the endogeneity of the distance variable D_{is} . Namely, it is possible that the store manager, knowing that category i is popular in the region of store s , decides to place the category in a more accessible store location. If this is the case, then the number of receipts \bar{n}_{is} might be negatively associated with D_{is} , though this would not be a causal relationship but rather a link due to the latent sales expectation. Because of the potential endogeneity problem, we proceed using a standard two-stage least squares (2SLS) approach with a Hausman-type instrumental variable.

The Hausman instrument can be used to address the endogeneity issue under the (standard) assumption that, once we control for macro category-store, local demand becomes independent across stores. This is needed to satisfy the exogeneity condition. Following Cachon et al. (2019), we select stores that are further than 150 km away from the focal store in order to make the local demand independence assumption more likely to be satisfied.

In addition, in-store logistics to replenish items is a costly activity that can add up to somewhere between 38% and 48% of retail logistics costs (Hübner and Kuhn 2012). For that reason, a proper store layout must take replenishment into account. For instance, Reiner et al. (2013) argue that “store design should not only consider customer needs but also logistics requirements so as to enable the trouble-free replenishment of products”. Our retail partner is aware of the implications of store layout on replenishment cost, and therefore, provides guidelines that store managers must adhere to, e.g., bulky and tall items should be placed in locations that can be easily accessed by a forklift, so they tend to be closer to the unloading dock. These guidelines are shared across stores, implying that the Hausman instrument satisfies the relevance condition once we control for local demand by means of the fixed effects.

Based on this reasoning, we instrument the distance of a category in a given store with the average of the distance measures of that category in the two closest stores that are beyond 150 km from the focal store. Specifically, in the first stage of the 2SLS procedure, we instrument D_{is} with $D_{is}^H := (\sum_{s' \in \mathcal{S}_s} D_{is'}) / |\mathcal{S}_s|$, where \mathcal{S}_s is the set that contains the two closest stores to s that are at least 150 km apart from s . Then, we regress D_{is} with D_{is}^H as well as all other covariates in Equation (3), and we use \widehat{D}_{is} to denote the predicted value. We expect D_{is}^H to be positively correlated with D_{is} due to the replenishment cost considerations that are similar across stores, which is confirmed by the first-stage results reported in Section 4.5. Furthermore, because we control for macro category-store fixed effects and only take stores that are far apart, we argue that the predicted value \widehat{D}_{is} is exogenous. In the second stage, the distance D_{is} in Equation (3) is replaced by the first-stage predicted value \widehat{D}_{is} . We thus obtain an unbiased estimate of the distance coefficient γ . In our results, we report both the OLS and the 2SLS estimates, and we discuss below their differences.

4.5 Results

We first study the use of nearby online interactions as determinants of store sales. For that purpose, we set $\gamma = 0$ in Equation (5). The results are reported in Table 5. Model (1) presents a benchmark model that only incorporates macro category-store fixed effects. As we can see, fixed effects alone lead to an R^2 of 0.37, which suggests that cross-category and cross-store heterogeneity are high in our context. Model (2) incorporates the total online traffic for each store, in the same way Gallino and Moreno (2014) used online interactions as a driver of store sales. In comparison to them, we find that general online traffic is not significant and does not help predict category-level sales, suggesting that accounting for store variation through fixed effects is sufficient and store-level online traffic simply contains redundant information. In contrast, when we consider category-level online interactions in Models (3) and (4), the model fit improves steeply, to $R^2 = 0.66 - 0.67$. This implies that category-level clicks provide a strong signal about sales. Moreover, the coefficient in Model (3) is equal to 1.1857 and is highly significant, which suggests that the relationship between (logged) clicks and sales is approximately proportional. In other words, if online clicks increase by 10%, sales also increase by a similar amount. Model (4) breaks down clicks into different ‘quality grades’, by considering separately clicks in which the focal category was the first one in the session (the sequence of categories viewed by the consumer; the first one should be the most important to the consumer), and the clicks in which the focal category was in positions 2 to 4, or 5+. We can observe that indeed clicks in the first position have the highest coefficient 1.0142, whereas later clicks had lower coefficients and lower statistical significance. This supports our interpretation that online interactions are a proxy for true consumer interest, and it is revealed especially when it appears early in the online search sequence of the consumer. Namely, a 10% increase in first-

position of clicks leads to an increase of about 10% of store sales, while a 10% increase in clicks in the fifth position increases sales by about 3%.

We next incorporate the three spill-over variables from Section 4.3 in Models (5)-(8). We can see that primary demand \bar{v}_{is} remains significant and with a coefficient similar to that in Model (3). In contrast, \overline{near}_{is} is insignificant, whereas \overline{path}_{is} and \overline{aisle}_{is} are positive and significant but very small in magnitude. The coefficient of \overline{aisle}_{is} is in fact the largest among the spill-over variables considered, indicating that an increase in interest (online clicks) for neighboring items within an aisle slightly increases sales. This suggests that spill-over effects are positive but of second-order importance, which is understandable given the functional, non-impulse nature of the categories sold in our home improvement context. Another possible interpretation of this result is that webrooming informs a more focused consumer that will spend less time roaming at the store, and therefore, opportunities for cross-selling are diminished. Other authors have discussed similar effects of webrooming, see for instance Gao and Su (2017a).

The previous models establish that online clicks are a valuable determinant of store sales. We can now study the impact of category location on sales, corresponding to Equation (3) with $\gamma \neq 0$. As described earlier, we operationalize ease of access to the category in the store via the distance from entrance to the category and then to exit. Table 6 shows the result of the estimation. The table first provides the results without instrumenting the distance variable, in a standard OLS estimation, in Models (9) and (10), without and with online clicks respectively.

We first observe in Models (9) and (10), that distance is significant but only marginally improves the result of Models (1) and (3). The coefficient for \bar{v}_{is} in Model (10) compared to Model (3) remains almost the same, which means that the role of distance, driver of conversion, seems orthogonal to that of online interactions, a proxy for true consumer needs.

Similar to the models in Table 5, we observe that adding online visits improves the accuracy of the model significantly, as shown by the R^2 which increases from 0.37 to 0.66. This implies that a model that ignores online visits will confound the attribution of sales potential to categories, measured exclusively by their fixed effect $\alpha_{m_{is}}$ in Model (9). Specifically, Models (9) and (10) differ in how they attribute sales potential to categories after accounting for the effect of distance. Namely, Model (9) without online visits cannot distinguish between actual product potential and being at a good location: higher online visits might be explaining a higher performance, which leads to a confounded estimate of sales potential. Similarly, some categories with a high level of online visits but low conversion will be attributed a sales potential that is lower than their real potential, which can represent a missed opportunity. In other words, the more refined Model (10) ‘cleans up’ errors in attribution of sales potential to categories. As a result, optimizing the store layout without a proxy for product interest (or demand) will likely identify the wrong product-location assignments, as we show later in Section 5.2 where we discuss the case of the laminated flooring category highlighted in Figure 2.

Table 5: Models using online interactions.

	<i>Dependent variable:</i>							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
\bar{v}_s		-0.6418 (0.4284)						
\bar{v}_{is}			1.1857*** (0.0409)		1.1863*** (0.0409)	1.1883*** (0.0409)	1.1491*** (0.0402)	1.1503*** (0.0402)
$\bar{v}_{1,is}$				1.0142*** (0.0666)				
$\bar{v}_{2-4,is}$				-0.1279 (0.1588)				
$\bar{v}_{5+,is}$				0.3083* (0.1341)				
\overline{path}_{is}					0.0006* (0.0004)			0.0006* (0.0004)
\overline{near}_{is}						0.0028** (0.0014)		0.0028* (0.0014)
\overline{aisle}_{is}							0.0533*** (0.0111)	0.0562*** (0.0111)
Fixed effects	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store
Observations	2,625	2,625	2,625	2,625	2,625	2,625	2,625	2,625
R ²	0.3651	0.3651	0.6568	0.6596	0.6574	0.6575	0.6612	0.6626
Adjusted R ²	0.2725	0.2725	0.6066	0.6209	0.6070	0.6072	0.6115	0.6128
Residual Std. Error	1.4570 (df = 2290)	1.4570 (df = 2290)	1.0714 (df = 2289)	1.0518 (df = 2287)	1.0708 (df = 2288)	1.0706 (df = 2288)	1.0648 (df = 2288)	1.0630 (df = 2286)
F Statistic	3.9426*** (df = 334; 2290)	3.9426*** (df = 334; 2290)	13.0776*** (df = 335; 2289)	13.7514*** (df = 337; 2287)	13.0636*** (df = 336; 2288)	13.0715*** (df = 336; 2288)	13.2911*** (df = 336; 2288)	13.2842*** (df = 338; 2286)

Note: Robust standard errors in parentheses.

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

Table 6 then reports the results of the 2SLS procedure, which controls for potential endogeneity of in-store distance, see Section 4.4. First, we observe that the first stage in Models (11) and (12) is working correctly: in-store distance is indeed significantly correlated with the instrument (relevance criterion). Second, the results show a strong effect of distance, with a coefficient of -0.0176 in Model (12). The coefficient is about seven times higher than the estimate coming out of OLS in Models (9) and (10), suggesting that the distance D_{is} exhibits some level of endogeneity. Surprisingly, the retailer is not necessarily placing higher selling items in locations with easier access. In fact, the opposite is happening, e.g., higher selling categories are placed at the back of the store. For example in the right image of Figure 2, we see that some high-conversion categories (in orange) are placed in locations far from the entrance. This was perhaps done with the hope of generating cross-selling—despite our analysis not uncovering strong cross-selling patterns. As a result, OLS underestimates the effect of distance, and it becomes necessary to pursue a 2SLS analysis. Moreover, the coefficient γ has a relatively high value: distance within the store roughly varies between 50 and 250 meters, which implies that the difference in sales between the closest and furthest categories is about $-0.0176 \times (250 - 50) = -3.52$, a 97% decrease (since $e^{-3.52} = 0.029$).

4.6 Robustness

While our main models in Tables 5 and 6 are kept simple to focus on the direct impact of online visits and effort, we run several robustness checks to confirm that our empirical results hold across a variety of empirical specifications and instrumental variable configurations. We discuss below the findings and include details in the Appendix.

First, as we have explained in Section 4.4, we build the instrument for our the endogenous distance variable D_{is} by averaging the distance of category i in the two closest stores that are at least 150 km apart from the focal store s . Tables 7 and 8 in the Appendix replicate the models in Table 6 and show the corresponding estimates using different instrumental variables. Table 7 reports the robustness for the models without online visits, where Model (9) corresponds to the OLS estimate without instrumentation, and the rest correspond to 2SLS estimates using different Hausman instruments: Model (11a) selects the two furthest stores; Model (11b) selects the two closest stores beyond 150 km, which is our Model (11) in Table 6; Model (11c) selects two random stores beyond 150 km; Model (11d) selects the two stores immediately to the South and the two immediately to the North of the focal store; and Model (11e) also selects four stores—two immediately to the South and two immediately to the North of the focal store—but they must be beyond 150 km. Table 8 follows the same logic and pertains to the models with online visits, where Model (10) corresponds to the OLS estimate without instrumentation, and Models (12a)-(12e) use the same instrument as their counterpart in Table 7. As one can see, the coefficient of the distance variable D_{is} is very stable across different variations of the Hausman instrument. It ranges between

Table 6: Models using online interactions and category location, time-aggregated data.

<i>Dependent variable:</i>				
	\bar{n}_{is}			
	<i>OLS</i>		<i>2SLS</i>	
	(9)	(10)	(11)	(12)
\bar{v}_{is}		1.1861*** (0.0409)		1.1888*** (0.0414)
D_{is}	-0.0020 (0.0012)	-0.0024** (0.0010)	-0.0145* (0.0088)	-0.0176*** (0.0070)
First stage, $IV(D_{is})$ $R^2 = 0.8695$			0.2248*** (0.0465)	0.2247*** (0.0465)
Fixed effects	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store
Observations	2,625	2,625	2,625	2,625
R^2	0.3656	0.6575	0.3465	0.6290
Adjusted R^2	0.2728	0.6072	0.2509	0.5746
Residual Std. Error	1.4567 (df = 2289)	1.0706 (df = 2288)	1.4785 (df = 2289)	1.1142 (df = 2288)
F Statistic	3.94*** (df = 335; 2289)	13.07*** (df = 336; 2288)		
Wald test			3.83*** (df = 335; 2289)	12.08*** (df = 336; 2288)

Note: Robust standard errors in parentheses.

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

-0.0145 and -0.0212 in Models (11a) through (11e), and between -0.0155 and -0.0215 in Models (12a) through (12e). In all cases the magnitude and statistical significance of the 2SLS estimates is larger and stronger than the OLS estimates.

Second, even though we do have cross-store layout variation, one may think that, if stores are far apart, then customers may be intrinsically different so in reality separate estimations should be conducted for each of the stores. To remedy this, we focus on the nine stores located in the Santiago Metropolitan area, which serve a common pool of customers living in the same city who arguably have more homogeneous tastes than across cities. We replicate the estimation of Models (9) through (12) with the data from these nine stores, see Table 9 in the Appendix. Again, we observe that the main findings are preserved.

Third, although we have used time-aggregated variables (across the 30 weeks that the data spans) in our main empirical specification—namely \bar{n}_{is} and \bar{v}_{is} —the original transaction and click-stream data are timestamped, which allows for a more granular analysis that leverages the panel

structure of the original data, in the vein of Equation (2). Tables 10 and 11 in the Appendix present the time-disaggregated version of the models in Tables 5 and 6 respectively. We note that the key insights are preserved and the estimates for visits and in-store distance remain stable.

Finally, the granular data also allows us to explore a possible lag between webrooming and store visits. We thus expand the model in Equation (2) by including $v_{is,t-1}$, i.e., the amount of online visits in the prior week. Since v_{ist} and $v_{is,t-1}$ are highly correlated, we may introduce colinearity in this new model, thereby introducing noise in the coefficients of these variables. We find that indeed the coefficient of v_{ist} drops from 1.0386 down to 0.5377 when lagged online visits are introduced—see Models (28) and (31) in Tables 11 and 12 respectively—while the lagged variable itself $v_{is,t-1}$ has a coefficient of 0.5517. Most importantly, the coefficient of distance remains at -0.0175 , hence nearly identical to that of Model (28) without lagged online visits.

5. Store Layout Optimization

5.1 A Category-Position Assignment Problem

Our model assumes and empirically demonstrates that a category’s location within the store has a significant impact on the conversion it generates. In this section, we are interested in prescribing improved layouts that increase total sales, taking consumer true needs as fixed captured via their online interactions.

Product location optimization is a relatively well-studied area of research, mainly in warehouse settings, see De Koster et al. (2007) for a review. In these contexts, one usually minimizes picking costs, which results in placing high-rotation items in easily accessible locations, while slow-movers are sent to more remote locations. In a store, the costs to bring items to the shelf are relatively small and insensitive to location within the store. As a consequence, we focus on the main driver of profits coming from the impact of category location on sales conversion.

We can formulate the layout design problem as the following assignment problem. Let x_{isp} be a binary variable that equals one when category $i \in \mathcal{I}$ is located in position $p \in \mathcal{P}$, in store s . One category can go into one position, and one position can only take one category.

Let d_p be the distance a consumer must travel from the entrance when a category is located in position p (at a given store s , subindex removed for simplicity). Then, the location-dependent demand of category i can be written as $r_{isp} = r_i \times d_{isp}$, where r_i is the average revenue from the category per receipt in which the category is present, and $d_{isp} = \exp(\alpha_{m_{is}} + \beta \bar{v}_{is} + \gamma d_p)$, as

predicted by Model (12). We then formulate the layout design problem as:

$$J_s := \max_x \sum_{i \in \mathcal{I}} \sum_{p \in \mathcal{P}} r_{isp} x_{isp} \quad (8)$$

$$s.t. \sum_{i \in \mathcal{I}} x_{isp} \leq 1 \quad \forall p \in \mathcal{P} \quad (9)$$

$$\sum_{p \in \mathcal{P}} x_{isp} \leq 1 \quad \forall i \in \mathcal{I} \quad (10)$$

$$x_{isp} \in \{0, 1\}. \quad (11)$$

The formulation J_s only includes constraints pertaining to the impossibility of placing two categories in the same location, or one category being sent to two locations. It is easy to incorporate additional linear constraints reflecting business conditions for the category in the store. For example, if a category can only be located in a particular part of the store, then we can set $x_{isp} = 0$ for infeasible locations. If categories i and j must be adjacent, then we can set $x_{isp} \leq \sum_{p'} A_{pp'} x_{jsp'}$ with $A_{pp'} = 1$ if p and p' are adjacent and zero otherwise; in other words, if $x_{isp} = 1$, then one adjacent p' (with $A_{pp'} = 1$) is such that $x_{jsp'} = 1$. In the absence of additional constraints, Equations (9)-(10) make a Totally-Unimodular Matrix (TUM), and hence constraint $x_{isp} \in \{0, 1\}$ can be replaced with $0 \leq x_{isp} \leq 1$ without changing the optimal solution of (8). In other words, J_s can be obtained by solving a linear program. Otherwise, we solve an integer program.

Note that we can write $r_{isp} = \bar{r}_{is} u_p$, with $u_p = \exp(\gamma d_p)$ and $\bar{r}_{is} = r_i \times \exp(\alpha_{m_i s} + \beta \bar{v}_{is})$, which will allow us to find the optimal assignment in closed form. Indeed, we can write the objective as $\sum_{i \in \mathcal{I}} \sum_{p \in \mathcal{P}} \bar{r}_{is} u_p x_{isp}$. This is maximized by assigning the location p with the largest u_p to the category i with the largest \bar{r}_{is} : assign the best in-store position (highest u_p) to the best-selling category (highest \bar{r}_{is}).

Finally, observe that the objective function in problem (8) has a separable structure. In other words, the sales of category i are independent of the location of other categories. This formulation is thus applicable to settings in which cross-selling is small. It requires a more complex, non-linear objective function when there are cross-category interactions, such as unplanned spending effects (Hui et al. 2013).

5.2 Improving on Existing Layouts

We can now apply the method of Section 5.1 to reengineer the actual layouts observed in our data. We first provide an in-depth analysis for one store and then provide results for the complete set of stores.

We define positions p in the same way as categories, i.e., we let $\mathcal{P} = \mathcal{I}$ and $p \in \mathcal{P}$ denotes the (current) location of category p . We compute r_i to be equal to the average spending per receipt that contains category i over the season of 30 weeks. To limit the number of changes, we formulate

the following decision problem:

$$J_s(z) := \max_x \sum_{i \in \mathcal{I}} \sum_{p \in \mathcal{P}} r_{isp} x_{isp} \quad (12)$$

$$s.t. \sum_{i \in \mathcal{I}} x_{isp} \leq 1 \quad \forall p \in \mathcal{P} \quad (13)$$

$$\sum_{p \in \mathcal{P}} x_{isp} \leq 1 \quad \forall i \in \mathcal{I} \quad (14)$$

$$\sum_{p \in \mathcal{P}} x_{psp} \geq |\mathcal{P}| - z \quad (15)$$

$$x_{isp} \in \{0, 1\} \quad (16)$$

In contrast to J_s , the problem formulation $J_s(z)$ includes the additional parametric constraint (15), where z is an integer variable. This constraint limits the number of actual category assignment changes to be z at the most. For example, if $z = 0$, the only feasible solution is to set $x_{psp} = 1$ for all $p \in \mathcal{P}$. If $z = |\mathcal{P}|$, then the constraint is innocuous. When z takes intermediate values, it provides us with interventions with varying degrees of complexity. Note, however, that constraint (15) breaks the TUM structure of the constraint matrix, and thus requires us to solve a set of integer programs. In addition, we consider two versions of the decision set \mathcal{P} : one that includes all categories, and another one that excludes construction categories that are typically bulkier and located at the side of the store, and hence, are difficult to place in any other store position. Our formulation ignores all other business constraints, e.g., adjacencies, space limitations, etc. but still our results are useful to understand the potential of layout optimization as suggested by our empirical findings.

Consider store 51, depicted in Figure 2. In this store, we have 165 different categories assigned to 165 positions shown in the map. As z increases, $J_s(z)$ increases from $J_s(0)$ (current layout). Of course, this improvement is due to the ability to optimize product locations compared to the status quo. This can be achieved with our Model (12) or with simpler models that also identify distance as a driver of sales, such as Model (11). To capture the value of our empirical findings, it is thus more appropriate to compare the incremental gain of using online information and a more sophisticated and accurate model, vs. that of the simpler model without such online information.

As discussed in Section 4.5, Models (11) and (12) differ in how they attribute sales potential to categories. Given this difference in demand estimation, the optimization program $J_s(z)$ may propose different product swaps in the layout to maximize impact. The difference is driven by \bar{r}_{is} , since the estimates of γ (and hence u_p) are almost identical for Models (11) and (12). If the ranking of \bar{r}_{is} is the same for Models (11) and (12), then the recommended layout will be the same. In contrast, if the estimates \bar{r}_{is} vary, the recommended swaps might be different and the more sophisticated model will lead to higher performance. We thus solve $J_s(z)$ with r_{isp} determined by Models (11) and (12) separately, and then evaluate the performance obtained by

the recommended layout using Model (12), which is taken as the ground truth because of its better accuracy. Let $REV_s^{full_info}(z) = J_s(z)$ be the performance achieved with r_{isp} from Model (12), and let $REV_s^{no_online_info}(z)$ be the objective value of the optimal layout under Model (11), but evaluated with r_{isp} coming from Model (12), implying $REV_s^{no_online_info}(z) \leq J_s(z)$.

Figure 3 shows the incremental value of using Model (12) vs. (11) in store 51, measured by $100 \times REV_s^{full_info}(z)/REV_s^{no_online_info}(z)$ and as a function of the number of changes allowed z . Interestingly, the revenue lift originating from using the full data when re-designing store layouts (as opposed to not accounting for online interactions) is greatest when the number of changes z is small. This means that the discrepancy in how the models with and without online interactions assign sales potential to the categories, and the economic implications that come with it, are most salient when the number of allowed changes are small. As we budget for more changes (e.g., after about 10 changes allowed if construction categories are included, and 50 if they are not), we see that the model with full information generates about 5% higher revenues. At the other extreme, when the number of changes is unlimited ($z = 165$), the lift in revenue comes from errors in attribution leading to changes in the ranking, which is 2-5%, depending on whether we exclude construction items from the layout optimization program. Note that a lift of 2% is very significant for a home improvement retailer, where margins are thin and increasing the top line typically has a very strong effect on net margins.

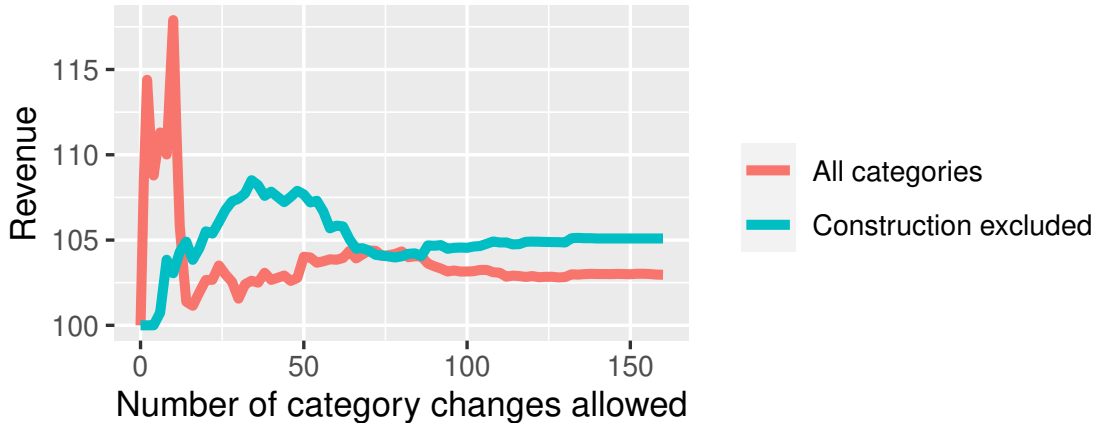


Figure 3: Improvement in store 51 due to the layout optimization using our model with online clicks, compared to the layout optimized using the model without online information, assuming that the true model is given by the former. Optimizing with full information suggests swaps that have a higher impact and results in a 2-5% revenue lift when the number of changes is unlimited.

When examining in more detail the type of changes recommended in this scenario, we see that top-selling categories located towards the back of the store are moved to the front, and some minor ones in front positions are relegated to the back of the store. As a more concrete example, consider

the laminated flooring category highlighted in Figure 2. This category is a high seller (as shown in the left image of Figure 2), but exhibits moderate conversion with respect to its online visits (as seen in the right image). Its current location p has a distance $d_p = 104.4$ meters, which is just slightly above the average distance in the store (98.7 meters). We solve $J_s(z)$ allowing for 10 changes (or 5 swaps), which is a realistic number of changes that could be implemented at a store, and excluding the construction categories (so there are 150 categories in total).

We first use Model (12) to estimate the sales potentials \bar{r}_{is} . Under this model, the laminated flooring category ranks 6th in terms of sales potential, but the top two categories—paint and ceramic tiles—are already located in positions with a short distance, so they do not need to be moved. Hence, the laminated flooring category becomes a candidate to be relocated. The optimized layout suggests putting it in the location currently used by the smart-home connectivity category, which has a sales potential that ranks 138th out of 150, and a distance of 46.2 meters (see the left image in Figure 2). This change means a 178% sales uplift for the laminated flooring category, since $e^{-0.0176 \times (46.4 - 104.4)} = 2.78$. The optimized layout using Model (11) also identifies the smart-home connectivity category as a prime location that could be better used. However, Model (11) misses the opportunity of the laminated flooring category because it is unaware of its high level of online visits and moderate conversion. Indeed, under Model (11), the sales potential of the laminated flooring category ranks 21st. Instead, under Model (11) the smart-home connectivity location is used to place the locks category, which is not the best choice because the locks category already sells quite well relative to its online visits (i.e., it has high webrooming conversion). In fact, the sales potential of the the locks category under Model (12) ranks 15th, so it would only be relocated if at least 30 changes (or 15 swaps) were allowed at store 51.

Finally, we can extend the optimization to our subsample of 16 stores for which we can reengineer the layout. Figure 4 shows the distribution of the revenue improvements achieved with an unconstrained layout change and one limited excluding construction categories, again comparing the performance of the optimized layout using Model (12) vs. Model (11). As we can see, the revenue lift of using the more sophisticated Model (12) can be significant, with some stores achieving improvements of more than 8%.

6. Conclusion

In this paper, we have provided a new perspective on how omnichannel, via webrooming customer interactions, can help retailers manage their physical stores better. Specifically, we have posited that, when sales are preceded by a need that crystallizes into a shopping list and pre-purchase category search, then store sales are driven by both the amount of nearby online visits and the effort that it takes to fetch the category in the store. We validate our conceptual model with data from a home improvement chain, over multiple categories and locations. The data provides

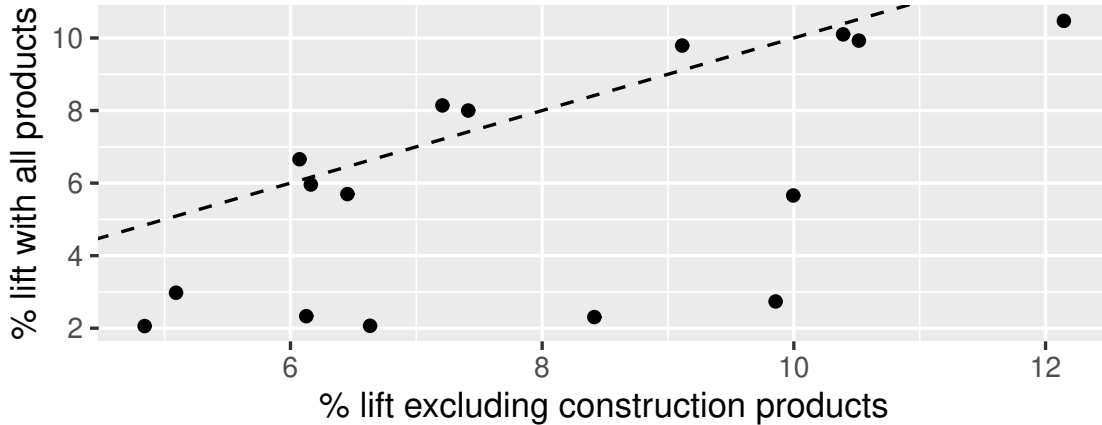


Figure 4: Distribution of $100 \times (REV_s^{full.info}(\infty)/REV_s^{no.online.info}(\infty) - 1)$ across 16 stores. Each point corresponds to a store.

variation of category interest and in-store location, and allow us to identify the effect of online visits and effort on sales. We find that sales grow proportionally with online visits, and that easy-to-reach store positions lead to significantly higher conversion. In addition, we show that there is a small, second-order cross-selling effect in this context, although this is not the focus of our research.

Our results have important implications for the management of physical stores. First, they suggest that layout reengineering using the information provided by online visits can provide a tempting lift in revenues, of around 2-5% compared to that obtained when online information is ignored. Second, they imply that the efforts to generate store visits, in the hope that they will generate unplanned purchases, may not be fruitful. In other words, it may be better that stores do not accept new roles as delivery points (Faithfull 2018, Jones 2019), if the categories on sale are related to a functional need that requires previous research. Third, our results identify the effort to find categories in the store as a hindrance to conversion. In other words, actions to make in-store category search simpler may lead to increased sales. One such action could be to provide category ‘addresses’ to consumers when they prepare their shopping lists, as Target does, see Figure 5. Finally, our results have been demonstrated for home improvement retailing. When impulse purchases are important, it may be advisable to pursue a different strategy with longer in-store paths, to sustain unplanned spending (Hui et al. 2013).

This study highlights the importance of better understanding the role of store design on customer experiences. This is a promising direction for future research. Indeed, the adoption of Internet Of Things technologies in stores provides new data sources for a more granular understanding of the trajectories of customers over time (the funnel view) and space (transitions between home, work and shopping destinations). This requires the full digitalization of the store conditions, and precise category locations, a piece of information that to date is rarely available, with the exception

of supermarket planograms, common in grocery retailing, or RFID sources, installed by Walmart or Zara among others. It can potentially reveal the causal impact of different interventions such as category viewing, staff advice or fitting (Musalem et al. 2021), environmental stimuli such as music or temperature (Martínez-de Albéniz and Belkaid 2021), as well as category information provision—to differentiate the effect of reducing cognitive load from search vs. that of physical movement to reach an item. Furthermore, if both online and offline activities could be connected at the customer level, one may be able to separate primary demand—from those that searched online before visiting the store—versus secondary demand made of spill-overs from entering an aisle. Indeed, combining on-premise data with online interactions is particularly interesting, so that conceptual frameworks such as Bell et al. (2014) can be operationalized and translated into prescriptive advice for retailers.




Figure 5: Standard category page on target.com, where the address of the category in the store of your choice is indicated.

Acknowledgements

Victor Martínez-de-Albéniz acknowledges the financial support provided by the Agencia Estatal de Investigación (AEI) from the Spanish Ministry of Science and Innovation, with project reference PID2020-116135GB-I00 MCIN/ AEI / 10.13039/501100011033.

References

- Accenture 2013. Accenture Study Shows U.S. Consumers Want a Seamless Shopping Experience Across Store, Online and Mobile. Technical report, Accenture plc, Available: <https://accntu.re/3oFRJo2>. Last accessed Oct 7, 2021.
- Aouad, A., V. Farias, and R. Levi. 2021. Assortment optimization under consider-then-choose choice models. *Management Science* 67 (6): 3368–3386.
- Arora, N., G. M. Allenby, and J. L. Ginter. 1998. A hierarchical Bayes model of primary and secondary demand. *Marketing Science* 17 (1): 29–44.
- Bar-Gill, S., and S. Reichman. 2020. Stuck Online: When Online Engagement Gets in the Way of Offline Sales. *MIS Quarterly* Forthcoming:NA.
- Bell, D. R., S. Gallino, and A. Moreno. 2014. How to win in an omnichannel world. *MIT Sloan Management Review* 56 (1): 45.
- Bell, D. R., S. Gallino, and A. Moreno. 2020. Customer supercharging in experience-centric channels. *Management Science* 66 (9): 4096–4107.
- Bijmolt, T. H., M. Broekhuis, S. De Leeuw, C. Hirche, R. P. Rooderkerk, R. Sousa, and S. X. Zhu. 2021. Challenges at the marketing–operations interface in omni-channel retail environments. *Journal of Business Research* 122:864–874.
- Boada-Collado, P., and V. Martínez-de-Albéniz. 2020. Estimating and optimizing the impact of inventory on consumer choices in a fashion retail setting. *Manufacturing & Service Operations Management* 22 (3): 582–597.
- Brynjolfsson, E., Y. Hu, and D. Simester. 2011. Goodbye pareto principle, hello long tail: The effect of search costs on the concentration of product sales. *Management Science* 57 (8): 1373–1386.
- Brynjolfsson, E., Y. J. Hu, and M. S. Rahman. 2013. Competing in the Age of Omnichannel Retailing. *MIT Sloan Management Review* 54 (4): 23.
- Cachon, G. P., S. Gallino, and M. Olivares. 2019. Does adding inventory increase sales? Evidence of a scarcity effect in US automobile dealerships. *Management Science* 65 (4): 1469–1485.
- Caro, F., and J. Gallien. 2010. Inventory Management of a Fast-Fashion Retail Network. *Operations Research* 58 (2): 257–273.
- Caro, F., A. G. Kök, and V. Martínez-de Albéniz. 2020. The future of retail operations. *Manufacturing & Service Operations Management* 22 (1): 47–58.
- Chandon, P., J. W. Hutchinson, E. T. Bradlow, and S. H. Young. 2009. Does in-store marketing work? Effects of the number and position of shelf facings on brand attention and evaluation at the point of purchase. *Journal of marketing* 73 (6): 1–17.
- Chen, Y., J. D. Hess, R. T. Wilcox, and Z. J. Zhang. 1999. Accounting profits versus marketing profits: A relevant metric for category management. *Marketing science* 18 (3): 208–229.
- Chintagunta, P. K., J. Chu, and J. Cebollada. 2012. Quantifying transaction costs in online/off-line grocery channel choice. *Marketing Science* 31 (1): 96–114.
- Chuang, H. H.-C., R. Oliva, and O. Perdikaki. 2016. Traffic-based labor planning in retail stores. *Production and Operations Management* 25 (1): 96–113.

- Corstjens, M., and P. Doyle. 1981. A model for optimizing retail space allocations. *Management Science* 27 (7): 822–833.
- Cui, R., S. Gallino, A. Moreno, and D. J. Zhang. 2018. The operational value of social media information. *Production and Operations Management* 27 (10): 1749–1769.
- De Koster, R., T. Le-Duc, and K. J. Roodbergen. 2007. Design and control of warehouse order picking: A literature review. *European journal of operational research* 182 (2): 481–501.
- DeHoratius, N., A. J. Mersereau, and L. Schrage. 2008. Retail inventory management when records are inaccurate. *Manufacturing & Service Operations Management* 10 (2): 257–277.
- DeHoratius, N., and A. Raman. 2008. Inventory record inaccuracy: An empirical analysis. *Management science* 54 (4): 627–641.
- Deloitte 2017. 2017 pre-Thanksgiving pulse survey. Technical report, Deloitte Touche Tohmatsu Limited, Available: <https://bit.ly/2IOKEg6>. Last accessed Oct 7, 2021.
- Digital Commerce 360 2017. Home Depot plans to spend \$5.4 billion to sharpen its omnichannel strategy. *Digital Commerce 360* December 8:online.
- Dowsett, S. 2019. Less is more? Inditex cuts stores but boosts space in home market Spain. *Reuters* June 20:online.
- Faithfull, M. 2018. Innovation: do lockers create increased footfall? *Beyond Retail Industry* June 14:online.
- Flavián, C., R. Gurrea, and C. Orús. 2020. Combining channels to make smart purchases: The role of webrooming and showrooming. *Journal of Retailing and Consumer Services* 52:101923.
- Gallien, J., A. J. Mersereau, A. Garro, A. D. Mora, and M. N. Vidal. 2015. Initial shipment decisions for new products at Zara. *Operations Research* 63 (2): 269–286.
- Gallino, S., and A. Moreno. 2014. Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information. *Management Science* 60 (6): 1434–1451.
- Gallino, S., and A. Moreno. 2019. *Operations in an omnichannel world*. Springer.
- Gao, F., and X. Su. 2017a. Omnichannel retail operations with buy-online-and-pick-up-in-store. *Management Science* 63 (8): 2478–2492.
- Gao, F., and X. Su. 2017b. Online and offline information for omnichannel retailing. *Manufacturing & Service Operations Management* 19 (1): 84–98.
- Goic, M., and M. Olivares. 2019. Omnichannel Analytics. In *Operations in an Omnichannel World*, 115–150. Springer.
- Harris 2013. Showrooming and Webrooming: A Tale of Two Trends. Technical report, Harris Interactive, Available: <https://bit.ly/3AfNbqB>. Last accessed Oct 7, 2021.
- Heese, H. S., and V. Martínez-de Albéniz. 2018. Effects of assortment breadth announcements on manufacturer competition. *Manufacturing & Service Operations Management* 20 (2): 302–316.
- Huang, T., and J. A. Van Mieghem. 2014. Clickstream data and inventory management: Model and empirical analysis. *Production and Operations Management* 23 (3): 333–347.
- Hübner, A., A. Holzapfel, H. Kuhn, and E. Obermair. 2019. Distribution in Omnichannel Grocery Retailing: An Analysis of Concepts Realized. In *Operations in an Omnichannel World*, 283–310. Springer.

- Hübner, A. H., and H. Kuhn. 2012. Retail category management: State-of-the-art review of quantitative research and software applications in assortment and shelf space management. *Omega* 40 (2): 199–209.
- Hui, S. K., P. S. Fader, and E. T. Bradlow. 2009a. Path data in marketing: An integrative framework and prospectus for model building. *Marketing Science* 28 (2): 320–335.
- Hui, S. K., P. S. Fader, and E. T. Bradlow. 2009b. Testing behavioral hypotheses using an integrated model of grocery store shopping path and purchase behavior. *Journal of consumer research* 36 (3): 478–493.
- Hui, S. K., J. J. Inman, Y. Huang, and J. Suher. 2013. The effect of in-store travel distance on unplanned spending: Applications to mobile promotion strategies. *Journal of Marketing* 77 (2): 1–16.
- Jalali, Z., M. C. Cohen, N. Ertekin, and M. Gumus. 2023. Vertical Product Location Effect on Sales: A Field Experiment in Convenience Stores. *Available at SSRN 4476326*.
- Jones, C. 2019. Can't wait for that delivery? Amazon, Rite Aid team up to make it easier to get packages. *USA Today* June 27:online.
- Kumar, A., A. Mehra, and S. Kumar. 2019. Why do stores drive online sales? Evidence of underlying mechanisms from a multichannel retailer. *Information Systems Research* 30 (1): 319–338.
- Larson, J. S., E. T. Bradlow, and P. S. Fader. 2005. An exploratory look at supermarket shopping paths. *International Journal of research in Marketing* 22 (4): 395–414.
- Lemon, K. N., and P. C. Verhoef. 2016. Understanding customer experience throughout the customer journey. *Journal of marketing* 80 (6): 69–96.
- Lu, Y., A. Musalem, M. Olivares, and A. Schilkrut. 2013. Measuring the effect of queues on customer purchases. *Management Science* 59 (8): 1743–1763.
- Mani, V., S. Kesavan, and J. M. Swaminathan. 2015. Estimating the impact of understaffing on sales and profitability in retail stores. *Production and Operations Management* 24 (2): 201–218.
- Martínez-de Albéniz, V., and A. Belkaid. 2021. Here comes the sun: Fashion goods retailing under weather fluctuations. *European Journal of Operational Research* 294 (3): 820–830.
- Martínez-de Albéniz, V., A. Planas, and S. Nasini. 2020. Using clickstream data to improve flash sales effectiveness. *Production and Operations Management* 29 (11): 2508–2531.
- Martínez-de Albéniz, V., and G. Roels. 2011. Competing for shelf space. *Production and Operations Management* 20 (1): 32–46.
- Montoya, R., and C. Gonzalez. 2019. A hidden Markov model to detect on-shelf out-of-stocks using point-of-sale data. *Manufacturing & Service Operations Management* 21 (4): 932–948.
- Mowrey, C. H., P. J. Parikh, and K. R. Gue. 2018. A model to optimize rack layout in a retail store. *European Journal of Operational Research* 271 (3): 1100–1112.
- Musalem, A., M. Olivares, and A. Schilkrut. 2021. Retail in high definition: Monitoring customer assistance through video analytics. *Manufacturing & Service Operations Management* 23 (5): 1025–1042.
- Ozgormus, E., and A. E. Smith. 2020. A data-driven approach to grocery store block layout. *Computers & Industrial Engineering* 139:105562.
- Perdikaki, O., S. Kesavan, and J. M. Swaminathan. 2012. Effect of traffic on sales and conversion rates of retail stores. *Manufacturing & Service Operations Management* 14 (1): 145–162.

- Reiner, G., C. Teller, and H. Kotzab. 2013. Analyzing the efficient execution of in-store logistics processes in grocery retailing—The case of dairy products. *Production and Operations Management* 22 (4): 924–939.
- Ruiz, F. J., S. Athey, D. M. Blei et al. 2020. Shopper: A probabilistic model of consumer choice with substitutes and complements. *Annals of Applied Statistics* 14 (1): 1–27.
- Schaverien, A. 2018. Five Reasons Why Amazon Is Moving Into Bricks-And-Mortar Retail. *Forbes* December 2018:online.
- Underhill, P. 2009. *Why we buy: The science of shopping—updated and revised for the internet, the global consumer, and beyond*. Simon and Schuster.
- Wagner, R. P., and T. Jeitschko. 2017. Why Amazon’s 1-Click Ordering Was A Game Changer. Knowledge@Wharton.
- Wang, R., and O. Sahin. 2018. The impact of consumer search cost on assortment planning and pricing. *Management Science* 64 (8): 3649–3666.
- Wiesel, T., K. Pauwels, and J. Arts. 2011. Practice Prize Paper-Marketing’s Profit Impact: Quantifying Online and Off-line Funnel Progression. *Marketing Science* 30 (4): 604–611.

Appendix: Supporting Tables for Robustness

Table 7: Models (9) and (11), with multiple instrumental variable configurations.

	<i>Dependent variable:</i>					
	\bar{n}_{is}					
	<i>OLS</i>				<i>2SLS</i>	
	(9)	(11a)	(11b)	(11c)	(11d)	(11e)
D_{is}	-0.0020 (0.0012)	-0.0170*** (0.0054)	-0.0145* (0.0088)	-0.0174** (0.0068)	-0.0156** (0.0074)	-0.0212** (0.0070)
First stage, $IV(D_{is})$		0.3373*** (0.0550)	0.2248*** (0.0465)	0.2713*** (0.0507)	0.2876*** (0.0615)	0.2940*** (0.0526)
First stage, R^2		0.8749	0.8695	0.8714	0.8698	0.8713
Fixed effects	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store
Observations	2,625	2,625	2,625	2,625	2,625	2,625
R^2	0.3656	0.3379	0.3465	0.3366	0.3428	0.3200
Adjusted R^2	0.2728	0.2410	0.2509	0.2395	0.2466	0.2204
Residual Std. Error	1.4567 (df = 2289)	1.4882 (df = 2289)	1.4785 (df = 2289)	1.4897 (df = 2289)	1.4827 (df = 2289)	1.5082 (df = 2289)
F Statistic	3.94*** (df = 335; 2289)					
Wald test		3.79*** (df = 335; 2289)	3.83*** (df = 335; 2289)	3.78*** (df = 335; 2289)	3.81*** (df = 335; 2289)	3.69*** (df = 335; 2289)

Note: Robust standard errors in parentheses.

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

Table 8: Models (10) and (12), with multiple instrumental variable configurations.

	<i>Dependent variable:</i>					
	\bar{n}_{is}					
	<i>OLS</i>				<i>2SLS</i>	
	(10)	(12a)	(12b)	(12c)	(12d)	(12e)
\bar{v}_{is}	1.1861*** (0.0409)	1.1884*** (0.0413)	1.1888*** (0.0414)	1.1889*** (0.0415)	1.1895*** (0.0417)	1.1893*** (0.0416)
D_{is}	-0.0024** (0.0010)	-0.0155*** (0.0041)	-0.0176*** (0.0070)	-0.0185*** (0.0060)	-0.0215*** (0.0072)	-0.0203*** (0.0061)
First stage, $IV(D_{is})$		0.3374*** (0.0551)	0.2247*** (0.0465)	0.2712*** (0.0507)	0.2875*** (0.0612)	0.2940*** (0.0526)
First stage, R^2		0.8749	0.8695	0.8714	0.8698	0.8713
Fixed effects	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store
Observations	2,625	2,625	2,625	2,625	2,625	2,625
R^2	0.6575	0.6364	0.6290	0.6254	0.6125	0.6177
Adjusted R^2	0.6072	0.5830	0.5746	0.5704	0.5556	0.5616
Residual Std. Error	1.0706 (df = 2288)	1.1031 (df = 2288)	1.1142 (df = 2288)	1.1196 (df = 2288)	1.1388 (df = 2288)	1.1311 (df = 2288)
F Statistic	13.07*** (df = 336; 2288)					
Wald test		12.34*** (df = 336; 2288)	12.08*** (df = 336; 2288)	11.97*** (df = 336; 2288)	11.57*** (df = 336; 2288)	11.73*** (df = 336; 2288)

Note: Robust standard errors in parentheses.

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

Table 9: Models (9) through (12) estimated with the subset of stores within the Santiago Metropolitan area.

<i>Dependent variable:</i>				
\bar{v}_{is}				
	<i>OLS</i>		<i>2SLS</i>	
	(13)	(14)	(15)	(16)
\bar{v}_{is}		1.1763*** (0.0549)		1.1841*** (0.0560)
D_{is}	-0.0018 (0.0018)	-0.0027* (0.0016)	-0.0153 (0.0098)	-0.0217** (0.0092)
First stage, $IV(D_{is})$			0.2539*** (0.0342)	0.2535*** (0.0342)
$R^2 = 0.8894$				
Fixed effects	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store
Observations	1,470	1,470	1,470	1,470
R^2	0.4072	0.6762	0.3890	0.6402
Adjusted R^2	0.3202	0.6284	0.2993	0.5871
Residual Std. Error	1.4334 (df = 1281)	1.0599 (df = 1280)	1.4553 (df = 1281)	1.1172 (df = 1280)
F Statistic	4.68*** (df = 188; 1281)	14.14*** (df = 189; 1280)		
Wald test			4.55*** (df = 188; 1281)	12.75*** (df = 189; 1280)

Note: Robust standard errors in parentheses.

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

Table 10: Models in (1) through (8) from Table 5, granular data (without time-aggregation).

		<i>Dependent variable:</i>							
		<i>n_{ist}</i>							
		(17)	(18)	(19)	(20)	(21)	(22)	(23)	(24)
<i>v_{st}</i>			-0.0762 (0.0726)						
<i>v_{ist}</i>				1.0363*** (0.0061)		1.0367*** (0.0061)	1.0372*** (0.0061)	1.0019*** (0.0061)	1.0017*** (0.0061)
<i>v_{1,ist}</i>					0.7492*** (0.0079)				
<i>v_{2-4,ist}</i>					0.2372*** (0.0108)				
<i>v_{5+,ist}</i>					0.0658*** (0.0089)				
<i>path_{ist}</i>						0.0007*** (0.0001)			0.0007*** (0.0001)
<i>near_{ist}</i>							0.0023*** (0.0002)		0.0022*** (0.0002)
<i>aisle_{ist}</i>								0.0613*** (0.0021)	0.0640*** (0.0021)
Fixed effects	Week	Week	Week	Week	Week	Week	Week	Week	Week
	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store	MacroCat-Store
Observations	78,750	78,750	78,750	78,750	78,750	78,750	78,750	78,750	78,750
R ²	0.3303	0.3303	0.5690	0.5809	0.5695	0.5694	0.5740	0.5752	
Adjusted R ²	0.3272	0.3272	0.5670	0.5789	0.5675	0.5674	0.5720	0.5732	
Residual Std. Error	1.5208 (df = 78386)	1.5208 (df = 78385)	1.2201 (df = 78385)	1.2032 (df = 78383)	1.2193 (df = 78384)	1.2195 (df = 78384)	1.2195 (df = 78384)	1.2129 (df = 78384)	1.2112 (df = 78382)
F Statistic	106.5128*** (df = 363; 78386)	106.2232*** (df = 364; 78385)	284.2384*** (df = 364; 78385)	296.7827*** (df = 366; 78383)	284.0855*** (df = 365; 78384)	283.9372*** (df = 365; 78384)	289.3962*** (df = 365; 78384)	289.2072*** (df = 367; 78382)	

Note: Robust standard errors in parentheses.

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

Table 11: Models (9) through (12) from Table 6, granular data (without time-aggregation).

<i>Dependent variable:</i>				
n_{ist}				
	<i>OLS</i>		<i>2SLS</i>	
	(25)	(26)	(27)	(28)
v_{ist}		1.0366*** (0.0061)		1.0386*** (0.0062)
D_{is}	-0.0020*** (0.0002)	-0.0023*** (0.0002)	-0.0145*** (0.0014)	-0.0172*** (0.0013)
First stage, $IV(D_{ist})$			0.2248*** (0.0078)	0.2247*** (0.0077)
$R^2 = 0.8695$				
Fixed effects	Week MacroCat-Store	Week MacroCat-Store	Week MacroCat-Store	Week MacroCat-Store
Observations	78,750	78,750	78,750	78,750
R^2	0.3308	0.5695	0.3146	0.5465
Adjusted R^2	0.3277	0.5675	0.3114	0.5443
Residual Std. Error	1.5203 (df = 78385)	1.2193 (df = 78384)	1.5386 (df = 78385)	1.2515 (df = 78384)
F Statistic	106.4*** (df = 364; 78385)	284.1*** (df = 365; 78384)		
Wald test			104.0*** (df = 364; 78385)	269.8*** (df = 365; 78384)

Note: Robust standard errors in parentheses.

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

Table 12: Models (19), (26) and (28) including one-week lagged online visits.

	<i>Dependent variable:</i>		
	<i>n_{ist}</i>		
	<i>OLS</i>		<i>2SLS</i>
	(29)	(30)	(31)
<i>v_{ist}</i>	0.5358*** (0.0134)	0.5360*** (0.0134)	0.5377*** (0.0135)
<i>v_{is,t-1}</i>	0.5510*** (0.0136)	0.5511*** (0.0135)	0.5517*** (0.0137)
<i>D_{is}</i>		-0.0024*** (0.0002)	-0.0175*** (0.0013)
First stage, <i>IV(D_{ist})</i>			0.2247*** (0.0046)
<i>R</i> ² = 0.8695			
Fixed effects	Week	Week	Week
	MacroCat-Store	MacroCat-Store	MacroCat-Store
Observations	76,125	76,125	76,125
<i>R</i> ²	0.5802	0.5808	0.5565
Adjusted <i>R</i> ²	0.5782	0.5788	0.5543
Residual Std. Error	1.1997 (df = 75760)	1.1988 (df = 75759)	1.2331 (df = 75759)
F Statistic	287.7*** (df = 364; 75760)	287.6*** (df = 365; 75759)	
Wald test			272.0*** (df = 365; 75759)

Note: Robust standard errors in parentheses.

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