Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Store-Wide Shelf-Space Allocation with Ripple Effects Driving Traffic

Tülay Flamand, Ahmed Ghoniem, Bacel Maddah

To cite this article:
Tülay Flamand, Ahmed Ghoniem, Bacel Maddah (2023) Store-Wide Shelf-Space Allocation with Ripple Effects Driving Traffic. Operations Research 71(4):1073-1092. https://doi.org/10.1287/opre.2023.2437

informs
https://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 71, No. 4, July–August 2023, pp. 1073–1092
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Contextual Areas

Store-Wide Shelf-Space Allocation with Ripple Effects Driving Traffic

Tülay Flamand,[a] Ahmed Ghoniem,[b,*] Bacel Maddah[c]

[a] Department of Economics and Business, Colorado School of Mines, Golden, Colorado 80401; [b] Department of Operations and Information Management, Isenberg School of Management, University of Massachusetts, Amherst, Massachusetts 01003; [c] Department of Industrial Engineering and Management, Maroun Semaan Faculty of Engineering and Architecture, American University of Beirut, Beirut, Lebanon

*Corresponding author

Contact: tflamand@mines.edu, https://orcid.org/0000-0003-3226-7087 (TF); aghoniem@isenberg.umass.edu, https://orcid.org/0000-0001-8326-1916 (AG); bacel.maddah@aub.edu.lb, https://orcid.org/0000-0001-7471-1655 (BM)

Received: January 13, 2021
Revised: June 2, 2022; September 28, 2022
Accepted: November 23, 2022
Published Online in Articles in Advance: March 16, 2023

Area of Review: Operations and Supply Chains

https://doi.org/10.1287/opre.2023.2437

Abstract. Given a store layout, product categories grouped into shelves, and historical sales data, we investigate how the allocation of product categories can be optimized in a fashion that guides in-store traffic and stimulates impulse buying. The latter constitutes an important shopping behavior that amounts to over 50% of the revenue in some retail settings. Considering a small-scale grocery store in Beirut, we analyze 40,000 customer receipts in order to relate in-store customer traffic to product shelf allocations and the store layout. This prompts the development of a predictive regression model that estimates traffic densities along a shelf as a function of the shelf-space allocation and the location of the shelf in the store. This traffic model captures a "ripple effect"—that is, the change in traffic throughout the store resulting from any change in product allocation to shelves. The customer traffic model is embedded within a mixed-integer nonlinear program that sets the shelf allocations across the entire store, thereby prescribing better location and shelf-space decisions for all product categories in a way that maximizes impulse buying. To overcome the computational challenges posed by the model, we develop a linear approximation for the traffic construct in the objective function, while keeping bilinear terms in the formulation, in order to derive lower and upper bounds on the optimal objective value. Our methodology produces a layout that yields a 65% improvement in the expected impulse profit for the grocery store in Beirut. Managerial insights into the structure of the proposed store configuration are also discussed. Specifically, the allocation of a fast-mover to a shelf directly drives traffic, not only through adjacent shelves, but also, indirectly, through more distant shelves that lead to it. This, in turn, creates advantageous locations for high-impulse products, which may not be in the immediate vicinity of fast-movers. Finally, the study suggests that, from an impulse profit perspective, the location of a product category in the store is more important than adjusting the amount of shelf space that it is allocated. This challenges the classical research approach whereby extensive effort is invested to determine the relative space of products along a single shelf, taken in isolation, without considering its location in the store.

Supplemental Material: The online companion is available at https://doi.org/10.1287/opre.2023.2437.

Keywords: shelf space allocation • impulse buying • in-store traffic • mixed-integer programming • retailing

1. Introduction

The central question that this paper investigates is twofold: (i) How do the store layout and the shelf-space allocation jointly affect in-store customer traffic; and (ii) how does the resulting traffic dictate product visibility and impulse buying? Impulse (or spontaneous, unplanned) buying constitutes a common shopping behavior and is believed to amount to 60%–70% of purchases in supermarkets (Underhill 2000), which are the primary focus of our study. For instance, stores like Walmart are increasingly viewed as attracting impulse buyers (Krasny 2012). Identified as a major source of revenue in retail, impulse purchases are actively stimulated via the visual display of product categories in the store and in-store marketing strategies (Bell et al. 2011). In fact, shelf-space allocation has long been recognized as a major determinant of a retailer's sales and operating costs (Corstjens and Doyle 1983). This is particularly important in self-service settings, where products are expected to visually stimulate demand in a spontaneous manner, without the encouragement of sales personnel (Anderson and Amato 1974).

Our study is predicated on the notion that in-store traffic depends on both the store layout and the allocation of product categories to the shelf space. For instance, areas that are in close proximity to the entrance and exit or the end of aisles are more visible to customers (Larson et al. 2005, Hui et al. 2009). Further, the allocation of fast-moving consumer goods (e.g., water, milk, and bread) to certain aisles directly boosts traffic in these aisles (Holmstrom 1997) and indirectly increases the customer footprint along paths in the store that lead to them. As a result, maximizing impulse profit requires a store-wide approach that can capture the complex interaction between shelf-space allocation and in-store traffic. This paper is a significant step in the direction of a holistic store-wide approach for managing impulse sales. Our work bridges exciting developments in two research areas: First, it develops predictive analytics for in-store traffic based on real data from a small-scale grocery store, at a time when tracking and sensory technologies are renewing interest in modeling in-store traffic. Second, it proposes prescriptive analytics for optimized store-wide shelf-space decisions using a mixed-integer nonlinear model that specifically targets determining the location and the amount of shelf space dedicated to each product category in the store (i.e., a set of substitutable products that serve similar needs, such as different brands of sugar or coffee). A common retail practice, which we tacitly assume in this paper, is that all the variants, also known as stock-keeping units (SKUs), in the category are assigned to adjacent shelf slots. Our approach embeds an empirically calibrated in-store traffic model, whereas, for simplicity, most recent studies consider the problem of optimizing shelf-space configurations assuming exogenous traffic densities along shelves.

By focusing on the effect of shelf-space allocation on store-wide impulse buying, the purpose of our study is not to ignore the substantial revenue stream that results from planned purchases. We implicitly assume that planned purchases tend to occur with certainty, regardless of the allocation decisions and the location of the shelves. That is, a customer who visits the store to buy such products as milk or coffee, say, will look for them and possibly seek assistance from employees to find them. Of course, the effect of allocation decisions of planned-purchases products on shopping convenience (i.e., effort to find products) is an interesting research question that is beyond the scope of the current work.

In what follows, we position the two main contributions of this paper in the context of the extant literature.

• Predictive Analytics for In-Store Traffic with "Ripple Effects." At the heart of this research is a predictive model for in-store traffic based on the store layout and the allocation of products. The model is calibrated using 40,000 customer receipts from a grocery store in Beirut. This departs from the recent literature, where the traffic density along shelves has been simplified to solely depend on the store layout (e.g., Botsali 2007; Flamand et al. 2016, 2018; Mowrey et al. 2018, 2019). Our traffic model is derived by a regression that predicts traffic at a store shelf based on its location and the sales volume of product categories allocated to it *and* to other shelves in the store. As such, our traffic model captures ripple effects (traffic variations) across the entire store due to change in product allocation, be it a minor swapping of two product-category locations or a major reassignment of shelf contents. For example, if soda is moved from the front to the back of the store, the ensuing ripples will be felt (i.e., traffic is varied) throughout the store. Although the grocery store in Beirut can be considered as "small," we believe that the developed methodology and insights extend to bigger stores. This store has 34 shelves, which provides enough room for impulse buying improving by swapping product groups across different shelves, which we numerically observe to be the main source of the improvement.
To avoid the demanding task of tracking customers throughout the store, we use an analytical method that estimates the dependent variable (traffic density) in the regression model based on receipts. Motivated by recent works on in-store customer traffic analysis (e.g., Hui et al. 2009), our approach assumes that customers shop around the store in a way that minimizes walking, in the spirit of the Traveling Salesman Problem (TSP).

• Prescriptive Analytics for Store-Wide Shelf-Space Allocation. The proposed in-store traffic model enables the formulation of a novel mixed-integer nonlinear program that seeks to improve product-category allocations in order to maximize impulse buying. The proposed model caters to several practical considerations: (i) product types (e.g., frozen, refrigerated, or ordinary products in supermarkets), which are associated with particular types of shelves; and (ii) product affinity, by requiring subsets of product categories that ought to be on the same shelf (e.g., pasta and pasta sauce) to be allocated to a same shelf in alternative, optimized configurations.

To the best of our knowledge, this is the first study that employs data and decision analytics for store-wide shelf-space management based on store layout and shelf-allocation considerations. This task is complex and numerically daunting due to several factors: (i) the nonlinearity of the optimization model due to the interplay between shelf-space allocation and traffic densities in different aisles, whereas classical shelf-space allocation models have either considered a few shelves in isolation or have adopted the simplifying assumption that in-store traffic is exogenous; and (ii) the complexity that governs the assignment of grouped product categories to shelves in the optimization model due to the embedded nonlinear traffic construct. Our solution methodology employs an approximation of the objective function by an upper-bounding piecewise linear function, which enables the computation of a valid optimality gap for the solution. We keep bilinear terms in the formulation, as these can be handled by recent versions of GUROBI.

Finally, we provide managerial insights for retailers by comparing the input store configurations, be it the current configuration or alternative configurations that are based on managerial directives, against optimized ones prescribed by our methodology and using scenario analysis over key model parameters. For example, we find that allocating fast-movers to the back of the store is highly beneficial, as it boosts traffic throughout the entire store. More generally, we find that the allocation of fast-movers to a shelf (e.g., at the back of the store) has important ripple effects in that it increases traffic through shelves (e.g., middle shelves) along the shopping paths that lead to it. This highlights the importance of optimizing shelf-space allocation across the entire store, and not limiting the decision scope to one or a few shelves, as is commonly done in the literature (as reviewed, for example, in Mou et al. 2018 and Ostermeier et al. 2021). A related interesting store-wide insight is that deciding on the location of a product category is significantly more important than adjusting the amount of shelf space that it is assigned. This departs from the classic literature that hinges on apportioning shelf space among competing products on one shelf. Finally, scenario analysis reveals that our optimization model is robust, in the sense that small changes in the input parameters do not have a big impact on the optimal impulse profit. This is useful, as some parameters may be difficult to estimate with great accuracy in practice (e.g., the impulse-purchase rates).

The remainder of this paper is organized as follows. Section 2 offers a literature background on the store-space planning process and reviews shelf-space allocation studies that are more closely related to our research, thereby highlighting our contribution. Section 3 provides a formal problem statement, along with the predictive analysis for in-store traffic using a regression model. Section 4 introduces the proposed mixed-integer nonlinear program and devises an upper-bounding linearization approach as a solution methodology. In Section 5, we calibrate our predictive model for in-store traffic using data from a grocery store in Beirut, for which we produce, in Section 6, improved shelf-space configurations using our methodology and conduct scenario analysis over key model parameters. Section 7 concludes the paper with a summary of our findings and directions for future research.

2. Background and Literature Review

In a recent paper, Ostermeier et al. (2021) present a general hierarchy for the space-planning process, which consists of store layout, category space assignment, and individual SKU allocation (see Section EC.1 of the online companion for more details). Our work spans aspects of the first two levels in this hierarchy, as we focus on category location (sequence) and space requirements. In the sequel, we briefly review related streams of the literature. In Section 2.1, we look at works on assigning individual SKUs to shelves and relate our work, especially our modeling assumptions, to this widely studied problem. In Sections 2.2–2.4, we discuss the literature pertaining to key concepts in our problem—namely, impulse buying, product visibility, and in-store traffic analysis. In Section 2.5, we position the contribution made by this paper in light of recent related works on shelf-space allocation.

2.1. Individual SKU Allocation

Consider a set of SKUs in a category of substitutable products (e.g., different coffee brands) with a well-defined space requirement and base demand. The problem here is to allocate the shelf space among the different SKUs and to form a planogram in both the horizontal and vertical dimensions. We present a brief account of this widely studied problem, along the lines of the comprehensive review by Bianchi-Aguiar et al. (2021).

In its most basic variant, this problem considers a single horizontal space dimension and assumes that the SKUs present a space-elastic demand, whereby the demand increases in the assigned space with diminishing returns. Earlier studies along these lines include Curhan (1972), Hansen and Heinsbroek (1979), Zufryden (1986), and Yang and Chen (1999). In the latter reference, the authors argue that with space constrained within a limited range, it is sufficient to assume that the SKU demand grows linearly in its assigned space. We adopt a similar linear approach in this paper for the demand of a product category as a function of its allocated space. Other studies (e.g., Corstjens and Doyle 1981, 1983; Irion et al. 2012) supplement the demand function with cross-space elasticities reflecting the fact that assigning a large space to an item may lower the demand for others. As noted in Irion et al. (2011), it is not likely that such cross-space effects carry over among different categories. Consequently, we ignore cross-space elasticity in our model.

Another stream of research on the SKU-allocation problem investigates demand substitution within a category in the event of an item stock-out (e.g., Borin et al. 1994 and Hübner and Schaal 2017a). Similar to cross-space elasticity, this effect is not likely present between different categories in the planning setting under investigation and is, therefore, not given consideration in our model. Moreover, the effect of a SKU relative location on demand has been examined in recent works. Some studies consider the effect of the vertical position of products, with the understanding that eye-level placements are usually favored (e.g., Hwang et al. 2005, 2009; Hariga et al. 2007). Others focus on the effect of a product horizontal position, whereby placements close to the end of aisles are typically favored (e.g., van Nierop et al. 2008 and Hansen et al. 2010). It is a common practice to assign SKUs from the same category along the entire height of a shelf segment, and, as such, we do not include the vertical location effect in our category-planning model. On the horizontal dimension, we assume that one shelf can accommodate multiple categories and apply a location-based density that favors closeness to the entrance or the exit of the store.

The literature has also given consideration to practical merchandising rules, such as displaying items in groups or families having the same brand, or similar colors and flavors, in close proximity to one another (e.g., Geismar et al. 2015, Bianchi-Aguiar et al. 2018, and Hübner et al. 2021b). One merchandising rule captured tacitly in our model is to require groups of categories to be allocated to the same shelf. For the case study in Beirut, these groups were based on the "current" configuration (at the time we visited the store).

Finally, noteworthy recent developments include demand stochasticity in studying the effect of stock-outs (e.g., Hübner and Schaal 2017b) and merchandising rules in three dimensions—namely, horizontally, vertically, and along the depth of the shelf (e.g., Düsterhöft et al. 2020 and Hübner et al. 2021a).

2.2. Impulse Buying

The Point-Of-Purchase Advertising International (POPAI 2014) categorizes purchases as *specifically planned*, *generally planned*, *substitute*, and *unplanned purchases*. Specifically, planned purchases include a choice of brand, prior to entering the store, whereas generally planned purchases are at the product-category level without committing to a brand. Although substitute purchases are not planned in advance, a customer might end up buying a different brand, rather than the planned one, upon seeing the store assortment. Unplanned purchases are not planned at all, neither at the category nor at the brand levels. Customers make such unplanned purchases as a result of in-store stimuli (Kollat and Willett 1967). For a more elaborate discussion of the definition of impulse purchases, the reader may refer to Piron (1991). Based on POPAI (2014), unplanned purchases in grocery stores have the largest share, amounting to 55% of the total purchases. Similarly, Phillips and Bradshaw (1993) report that, in grocery stores and other self-service sectors, at least half of the purchases are impulsive. In addition, Bellenger et al. (1978) find that 55% of purchases of bakery products, in a supermarket, and 51% of purchases of meals and snacks are impulsive.

The academic literature identifies product characteristics (Bellenger et al. 1978) and customer behavior (Kollat and Willett 1967, Bellenger et al. 1978, Phillips and Bradshaw 1993) as two major determinants for impulse buying. Inman et al. (2009) suggest that so-called *hedonic products*, which are associated with "pleasure" (e.g., chocolates or cakes), tend to have higher impulse-purchase rates. Likewise, Kacen et al. (2012) argue that hedonic products, low-priced products, and products on sale have higher impulse rates. Further, Abratt and Goodey (1990) elaborate on the rate of impulse purchases and emphasize the usefulness of in-store stimuli (e.g., price-off promotions, point-of-purchase displays, and coupons).

2.3. Product Visibility

Increasing impulse buying in a retail store is mostly related to improving the overall product visibility in a manner that encourages customers to explore products beyond their planned shopping lists. The visibility factor is studied with consideration to the relationship between sales and *space elasticity* (Cox 1970, Curhan 1972, Drèze et al. 1994, Desmet and Renaudin 1998), *cross-space elasticity* (Corstjens and Doyle 1981, Drèze et al. 1994, Chen et al. 2006), and *location effects* (Maddah and Bish 2009). Some studies suggest that sales of products that have higher impulse-purchase potential are more sensitive to changes in shelf space (Curhan 1972, Hübner and Schaal 2017b). Other studies also suggest that horizontal and vertical positioning of products have more important effects than space elasticity (Drèze et al. 1994, van Nierop et al. 2008) and that the location of product can significantly impact sales (Drèze et al. 1994). In addition, Drèze et al. (1994) suggest that vertical positioning of products has more impact on sales than horizontal positioning.

In a recent work, Mowrey et al. (2019) develop an analytical model to estimate the visibility of a rack by a customer walking through the main aisle of a retail store. This model is based on a thorough analysis of human vision and rack geometry, with a focus on rack orientation. They find that an obtuse shelf orientation offers more visibility, at the expense of space utilization. In another recent paper, Mowrey et al. (2018) employ this visibility concept in order to develop a mathematical program that optimizes the layout of a set of shelves (in terms of their orientation, length, and width) in a section of the store. The objective of their work is to maximize shelf visibility with constraints on space utilization. Guthrie and Parikh (2020) consider a similar problem with an added decision aspect on shelf curvature, allowing shelves to be tilted and curved. The authors also incorporate product allocation in their work. However, they focus on simple policies that allocate products based on straightforward heuristics once the shelf orientation is set.

2.4. In-Store Traffic Analysis

One limitation of the classical SKU-allocation models reviewed in Section 2.1 is studying shelves in isolation from their position in the store, which fails to capture the interplay between layout and product allocation.

Larson et al. (2005) and Hui et al. (2009) analyze customer travel behavior in grocery stores for a given layout using Radio Frequency Identification (RFID). They find that customer traffic density is not uniform throughout the store and that, for example, the middle-of-aisle locations are typically less visited than the end of aisles.

Customer traffic density also depends on the product-shelf allocation. Specifically, fast-moving products can draw consumer traffic into store aisles and improve customer footprint (Holmstrom 1997). Cummings et al. (1991) suggest that the majority of retailers sell tobacco products, regardless of their profitability, mainly as traffic drivers. Hui and Bradlow (2012) also suggest that less crowded areas of a store host less popular product categories.

Farley and Ring (1966) formulate a predictive model of supermarket traffic flow, using a regression model, which represents transition probabilities from one area to another in the store. Actual traffic-flow data are used, which were collected by following customers in the store during their shopping. However, besides the predictive model of traffic flow, no optimization model is proposed to inform better shelf-space allocation. We develop such a model in this paper that embeds an elaborate in-store traffic construct with ripple effects in an optimization framework.

2.5. Recent Contributions to Store-Wide Shelf-Space Planning

To the best of our knowledge, only a few recent studies consider a store-wide shelf-space allocation problem (e.g., Botsali 2007; Flamand et al. 2016; Ghoniem et al. 2016a, b; Ke 2023). Ke (2023) proposes a plan modification algorithm that examines the effect of a potential change in SKU locations on profitability. The algorithm evaluates submoves of a given plan and determines a combination of submoves that yields the best revenue improvement. Botsali (2007) investigates decisions pertaining to retail store layout configurations and shelf-space allocation of product categories. To estimate the visibility of products, shelves are discretized into segments having different traffic densities. The traffic density of each shelf segment solely depends on its location in the store; for example, segments that are closer to the entrance and exit have higher traffic. Given a customer shopping list, a simulated annealing heuristic generates a layout by configuring shelf units and heuristically assigns products to shelves. Flamand et al. (2016) propose an exact mathematical programming approach under a layout-based traffic density similar to that in Botsali (2007). To maintain a reasonable shelf allocation, and its familiarity to shoppers, product categories are grouped based on the original allocation and remain so as they get reallocated to different shelves. The mathematical program in Flamand et al. (2016) optimizes a one-to-one assignment of product-category groups to shelves, together with the locations of products in a group on the selected shelf. Flamand et al. (2018) examine a similar store-wide shelf-space allocation problem jointly with assortment planning with the objective of maximizing the overall store profit, again under the restrictive assumption that the traffic density is based on the store layout only.

In a recent work, Dorismond (2019) proposes two models for shelf-space allocation and store layout. The first model is applied in two phases and targets maximizing the visibility of impulse products. First, different blocks (departments) are allocated to the store floor plan, and then, product assignment to shelves is conducted at the block level. The second model considers the allocation of impulse products to end caps and other promotional areas of the store over multiple time periods (allowing the impulse products to be rotated over different selling seasons). The second model tacitly assumes that the store layout and the allocation of non-impulsive products are fixed. The visibility of the products in both models in Dorismond (2019) is generated from a taxing Monte Carlo simulation of customer profiles and shopping paths, which feeds into the optimization models. Another recent paper by Pak et al. (2020) considers an objective of allocating impulse products to end caps similar to Dorismond (2019). However, Pak et al. (2020) estimate the potential lift in sales of a SKU via regression analysis on historical data collected over multiple stores that exhibit similar customer behavior. The authors, however, do not capture how store layout and shelf allocation drive store traffic, product visibility, and impulse profit, in contrast to our current study.

Two related papers by Ostermeier et al. (2021) and Irion et al. (2011) follow a hierarchical approach for the last two phases of the space-planning process described in Section EC.1 of the online companion and present models that decide on the amount of space to allocate to categories in a store and the SKUs in each category. Both papers assume that the location of the categories in the store are predetermined. Irion et al. (2011) solve the SKU-allocation problem for each category multiple times with a different category capacity based on a linearization method introduced in a companion paper (Irion et al. 2012) that captures cross-space elasticity, substitution, and integrality in the number of facings. They then apply regression analysis for each category and link the category revenue to its assigned space via a monotone function with diminishing returns. This predictive construct is then fed into the category space-assignment problem in what Irion et al. (2011) refer to as the "store problem." Ostermeier et al. (2021) follow a similar approach to Irion et al. (2011), but utilize a well-structured preprocessing scheme instead of the power regression in the store problem. The approach in Ostermeier et al. (2021) seems better suited for handling larger problem instances (in terms of number of categories

Operations Research, 2023, vol. 71, no. 4, pp. 1073–1092, © 2023 INFORMS

per store) than that of Irion et al. (2011). Finally, along the same vein, Campo et al. (2000) investigate the optimal space to assign to each category in a set of stores. They adopt a demand function with space and cross-space elasticities typically used for allocating items within a category (similar to that in Corstjens and Doyle 1981) and analyze the effect of a store location on its layout.

Our work has several distinctive features. First, unlike Botsali (2007) and Flamand et al. (2016, 2018), we examine store-wide shelf-space allocation with the notion of in-store traffic, which depends on *both* the store layout and product allocation. Second, we utilize real data to calibrate a store-wide traffic-density model based on a regression analysis, which is then embedded within our product-allocation optimization model. This is arguably easier to apply and less computationally onerous than the Monte Carlo approach used to estimate product visibility in Dorismond (2019), for example. Third, we demonstrate the predictive ability of this regression model in answering "what-if" questions related to shelf-allocation adjustment in a real-world store. Fourth, for the same store, we demonstrate how mathematical programming can be used to obtain alternative shelf-space allocations that offer greater impulse-buying potential. Finally, we discuss useful managerial insights on the allocation of fast-movers, which can support and improve retailing practices. As such, assigning a product to a shelf should not be merely governed by the direct traffic it may attract to this space; rather, it should also account for indirect footprints that occur along the more distant shelves that lead to it throughout the entire network of walkways in the store. This effort is not tractable without predictive models for in-store traffic and optimization models for shelf-space allocation, as proposed in the current study. Whereas our managerial findings can provide general guidelines to practitioners, a detailed shelf-space allocation that takes into account store-wide interactions seems only attainable via such computational tools as the one proposed in this study.

Table 1 positions our work with respect to the recent literature classified based on the following:

- Decisions on (i) department location, (ii) department size (of shelf space allocated), (iii) category location, (iv) category size, SKU size, and (v) promotional end caps allocation;
- Traffic model whether (i) based on the location of the shelf in the store, (ii) based on the allocation of products to shelves, or (iii) exogenously specified;
- Methodology, whether it is based on (i) simulation or (ii) mathematical programming.

Table 1 reveals that our paper is the first to consider store-wide decisions related to category location and size of assigned shelf space via a tractable mathematical programming formulation, while utilizing a traffic model that captures both the effect of shelf location and allocation of products across the store. As discussed in Section 1, the main contribution of this paper vis-à-vis Flamand et al. (2016, 2018) is to integrate an allocation-based component of the traffic in the decision model. This has proven to be challenging to (1) quantify, as we resort to an elaborate preprocessing involving analyzing individual baskets and nonlinear regression; and (2) analyze, as the end model includes complex nonlinear terms that we work hard on simplifying. This work makes an important step forward with regard to including such a traffic component, which captures ripple effects from product reassignments across the store.

3. Problem Statement and Predictive Analytics for In-Store Traffic

We first describe our problem in Section 3.1. We then develop a predictive model for in-store traffic as a function of the store layout and shelf-space allocation using regression in Sections 3.2 and 3.3. This predictive model plays a central role in our study. Throughout this paper,

Table 1. Positioning of This Paper with Respect to the Recent Literature

	Decisions						Traffic			Method	
Work	Dept. location	Dept. size	Category location	Category size	SKU size	End caps	Location	Allocation	Exog.	Sim	MP
Botsali (2007)	–	–	✓	✓	–	–	✓	–	–	✓	–
Campo et al. (2000)	–	–	–	✓	–	–	–	–	✓	–	✓
Dorismond (2019) (1)	✓	✓	✓	✓	–	–	–	–	–	✓	✓
Dorismond (2019) (2)	–	–	–	–	–	✓	–	–	–	✓	✓
Flamand et al. (2016, 2018)	–	–	✓	✓	–	–	✓	–	–	–	✓
Ghoniem et al. (2016a, 2016b)	–	–	✓	✓	–	–	–	–	✓	–	✓
Irion et al. (2011)	–	–	–	✓	✓		–	–	✓	–	✓
Ke (2023)	–	–	✓	✓	–	–	✓	–	–	✓	–
Ostermeier et al. (2021)	–	–	–	✓	✓	–	–	–	✓	–	✓
Pak et al. (2020)	–	–	–	–	–	✓	–	–	–	–	✓
This paper	–	–	✓	✓	–		✓	✓	–	–	✓

Notes. Dorismond (1) and (2) refer to specific chapters in Dorismond's (2019) thesis. Dept., department; Exog., exogenous; MP, Mathematical Programming; Sim, simulation.

we consider the allocation of an entire product category, at an aggregate level, without deciding on the location of specific brands or SKUs within the space allocated to this category. For example, we look at the shelf location and the total shelf-space allocation for the coffee category, without further specifying the relative space assigned therein for individual coffee brands, flavors, and package sizes. The latter aspect is beyond the focus of our study and requires, at a more operational level, ascertaining vertical and horizontal shelf-space allocations to build planograms for a specific product category based on the store-wide solution that our model produces, as discussed in Section 2.1. Henceforth, we shall use the terms "product category" and "product" interchangeably.

3.1. Problem Statement

In this section, we formally introduce our notation and the store-wide shelf-space management problem with impulse-buying maximization. Without loss of generality, and for illustrative purposes, our discussion is placed in the context of a grocery store and can be specialized for other retail businesses. We address a setting where a retailer seeks to optimize shelf-space allocation in order to guide in-store traffic in a manner that maximizes the expected impulse buying under a given grocery store layout. Grocery stores commonly adopt grid layouts that comprise shelves (internal and peripheral), end caps, and a network of walkways, as depicted in Figure 1. To suit a broad spectrum of perishable and nonperishable product categories, shelves may include *refrigerated*, *frozen*, and *ordinary* shelves. Refrigerators serve the purpose of displaying perishable product categories, such as dairy products, juices, or eggs, whereas freezers are used for such product categories as frozen vegetables, pizza, ice cream, and precooked meals, to name a few. Ordinary shelves are dedicated to all other product categories that can be stored at ambient temperature—for example, rice, pasta, sauce, and oil. In order to abide by practical product-affinity requirements and to reduce the computational overhead, we assume that products are grouped, and each group can be assigned to exactly one shelf. In generating such groups, one may examine each shelf individually and determine all product groups that can be feasibly allocated to it. Specifically, one can utilize minimum shelf-space requirements for products, in concert with affinity considerations between products, in order to generate a large number of groups that can be feasibly allocated to a shelf. However, in our experience, this approach of building groups tends to be computationally onerous for the resulting optimization model, with little benefit. Instead, we form groups in our case study based on the following more pragmatic approach: (i) maintaining the group structure in the current configuration (so that the model would swap shelf contents), (ii) allocating fast-movers to the back of the store, and (iii) spreading fast-movers throughout the store. We refer to the set of groups resulting from each setting as an *input configuration*. We also find it numerically pertinent not to mix

Figure 1. (Color online) Two Customers' Shopping Paths Obtained by Solving a TSP for Each

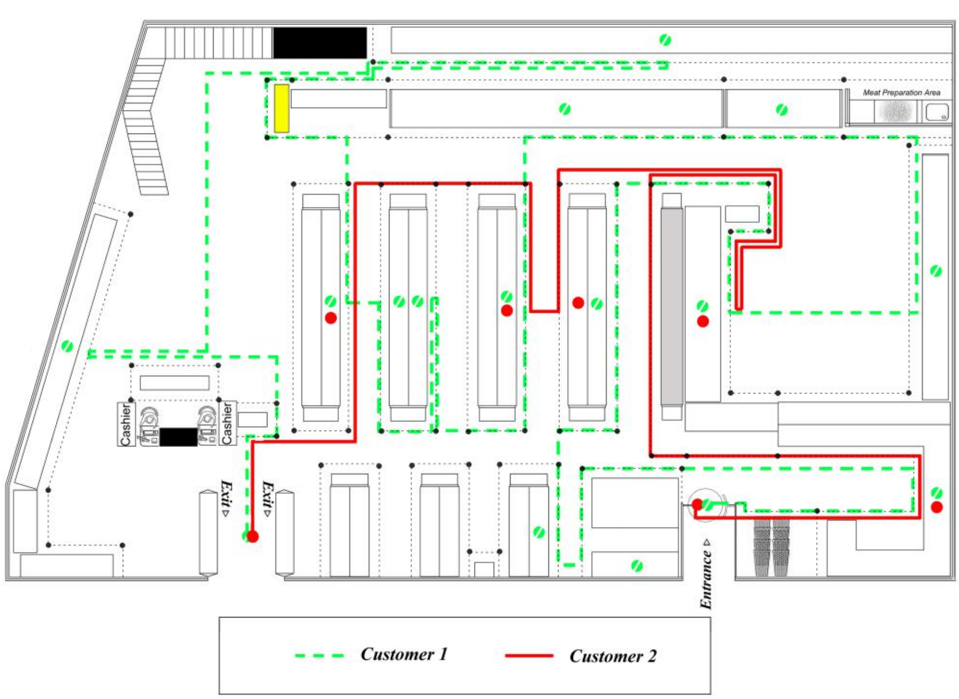

the input configurations in building the optimization model, in order to mitigate the associated computational burden. Consequently, the problem amounts to assigning grouped product categories to shelves throughout the entire store and adjusting the shelf space allocated to products on each shelf within prespecified lower and upper bounds on space requirements.

In practice, retailers seem to be applying such a shelf-content-swapping approach, but in a more localized manner. For example, one observes sometimes that the Soda and Water shelf, formerly in Aisle 3, has been moved to Aisle 7, whereby it was swapped with the Pasta and Pasta Sauce shelf. Our model considers all such swaps concurrently and selects the ones that maximize impulse buying. We introduce our model notation in Table 2.

Product demand is a resultant of both planned and impulse purchases. Products having high demand, henceforth referred to as *fast-movers* (e.g., milk, bread, or eggs), are commonly purchased in a planned fashion as part of a customer's shopping list. Such categories typically constitute a large fraction of the store demand volume, display stable demand patterns, and attract customers to their allocated shelves in a sustained fashion. The allocation of such categories is therefore critical in directing traffic throughout the store. Products with low demand, referred to as *slow-movers* (e.g., cakes or cosmetics), tend to have higher impulse rates (Kollat and Willett 1967). In this regard, Cox (1970) and Curhan (1972) found that sales of slow-movers are more sensitive to changes in shelf-space allocation than sales of

Table 2. Model Notation

Notation	Definition
Sets	
$\mathcal{T}$	Set of distinct shelf types. For example, refrigerated, frozen, and ordinary shelves.
$\mathcal{B}_t$	Set of the shelves of type t, $\forall t \in \mathcal{T}$.
$\mathcal{B} \equiv \cup_{t\in\mathcal{T}} \mathcal{B}_t$	Set of all store shelves.
$\mathcal{P}_t$	Set of product categories to be allocated to shelves of type t, $\forall t \in \mathcal{T}$.
$\mathcal{P} \equiv \cup_{t\in\mathcal{T}} \mathcal{P}_t$	Set of all product categories.
$\mathcal{J}$	Set of all preformed groups of product categories.
$\mathcal{Z}_g$	Set of product categories in group g, $\forall g \in \mathcal{J}$.
$\mathcal{G}_b$	Subset of product groups in an input layout that can be feasibly allocated to shelf b, $\forall b \in \mathcal{B}$.
$\mathcal{R}_b$	Set of shelves that are adjacent to shelf b, $\forall b \in \mathcal{B}$.
Parameters	
C_b	Length (capacity) of shelf b, $\forall b \in \mathcal{B}$.
λ_p	Basket sales volume of product p, $\forall p \in \mathcal{P}$.
λ_w	Basket sales volume of a subset of products $w \subseteq \mathcal{P}$.
ℓ_p and u_p	Minimum and maximum shelf-space requirement for product p, $\forall p \in \mathcal{P}$.
ρ_p	Unit profit margin for product p, $\forall p \in \mathcal{P}$, estimated as the mean profit margins of SKUs in product category p.
$i_p \in [0,1]$	Impulse-purchase rate of product p, $\forall p \in \mathcal{P}$.
$k_b \in [0,1]$	The layout-based customer traffic density along shelf b, $\forall b \in \mathcal{B}$.
d_{0b}	The shortest rectilinear distance between shelf b and the entrance, $\forall b \in \mathcal{B}$.
d_{bn}	The shortest rectilinear distance between shelf b and the exit, $\forall b \in \mathcal{B}$.
d_{qr}	The shortest walking distance between shelf q and shelf r.
m_g	The total number of people shopping for the product categories in group g, $\forall g \in \mathcal{J}$.
$\Delta_{qbr} \in \{0,1\}$	$\Delta_{qbr} = 1$ if the shortest path between shelves q and r includes shelf b, $\forall b \in \mathcal{B}, q \in \mathcal{R}_b, r \in \mathcal{B} : r \neq b, q$.
α	Intercept of the beta regression model for the traffic density.
β_1	Coefficient of the layout-based component in the beta regression model for the traffic density.
β_2	Coefficient of the allocation-based component in the beta regression model for the traffic density.
h	Short-hand notation of β_2 ($h = \beta_2$).
a_b	Intercept of the beta regression model when written in function of the allocation based traffic only, $\theta_b(z_b)$.
$\tilde{z}_b$	Coordinate along the x axis of the inflection point of $\theta_b(z_b)$.
z_1	Coordinate along the x axis of a point on $\theta_b(z_b)$ to the right of the inflection point, $z_1 > \tilde{z}_b$.
θ_{b0} and I_b	Intercepts of the first two components of the piece-wise linear approximation of $\theta_b(z_b)$, $\forall b \in \mathcal{B}$.
b_{1Lb} and b_{2Lb}	Slopes of the first two components of the piece-wise linear approximation $\theta_b(z_b)$, $\forall b \in \mathcal{B}$.
$\tilde{\theta}'_b$ and θ'_1	Slopes of $\theta_b(z_b)$ at $z_b = \tilde{z}_b$ and $z_b = z_1$, respectively, $\forall b \in \mathcal{B}$.
Decision variables	
$x_{bg} \in \{0,1\}$	$x_{bg} = 1$ if and only if group g is assigned to shelf b, $\forall b \in \mathcal{B}, g \in \mathcal{G}_b$.
s_{pb}	The shelf space allocated to product p along shelf b, $\forall p \in \mathcal{P}, b \in \mathcal{B}$.
Φ_b	The allocation-based customer traffic density along shelf b, $\forall b \in \mathcal{B}$.
$\Phi_b^{\text{norm}} \in [0,1]$	Normalized value of Φ_b.
z_b	Short-hand notation of Φ_b^{norm} ($z_b = \Phi_b^{\text{norm}}$).
f_{qr}	The force (amount of attraction) exerted on shelf q by shelf r, $\forall q, r \in \mathcal{B} \vert q \neq r$.
F_{qb}	The force on shelf b exerted through any adjacent shelf q, $\forall b \in \mathcal{B}, q \in \mathcal{R}_b$.
θ_b	The traffic density along shelf b that is the likelihood of shelf b to be visited by a customer, $\forall b \in \mathcal{B}$.
y_{1b}, y_{2b}, y_{3b}	Binary variables indicating which component of the piece-wise linear approximation of $\theta_b(z_b)$ applies.

fast-movers. Hence, such product categories require careful display to catch customer attention and stimulate impulse purchases. Because the profit generated from planned or fast-moving products could be viewed as a "constant," regardless of shelf-space-allocation decisions, we focus in the sequel on the profit that results from impulse buying only. The store-wide shelf-space management with impulse-buying maximization problem (IBMP) consists of allocating shelf space to the assortment of products offered by the store so that the capacity of shelves is not exceeded, lower and upper bounds on the space allocated to products are enforced, and the resulting store-wide traffic yields a most profitable impulse-buying potential.

3.2. The Effect of Store Layout and Shelf-Space Allocation on Traffic Density

An important element of our problem is the notion of customer traffic density. Each shelf is characterized by its customer traffic density, ranging between zero and one, that reflects the likelihood of a shelf to be visited by a customer. This traffic density involves two components: (i) a static layout-based component, k_b, which depends on the location of the shelf within the store (e.g., near entrance or near exit); and (ii) a variable component, Φ_b, that depends on the products allocated to the shelf. Both components can be estimated using store data. This intertwining between traffic densities, store layout and the location of products, and their allocation enables a more realistic, albeit more complex and computationally challenging, representation of in-store dynamics. As illustrated in Figure 1, we subdivide the network of walkways into a set of areas that span shelves in the entire store. Note that, although we assume that an area represents a shelf, the granularity of the analysis and modeling can be adjusted for an area to be viewed as a shelf segment or a set of shelves. We formulate the layout-based customer traffic density of shelf b based on the proximity of shelf b to the entrance and exit—that is, the closer the shelf is to the entrance and exit, the higher its layout-based traffic density is. Specifically,

$$k_b = \frac{1}{\min(d_{0b}, d_{bn})}, \quad \forall b \in \mathcal{B}, \tag{1}$$

where d_{0b} and d_{bn} are the distances from shelf b to the entrance and the exit, respectively. In general, the dependency of traffic on location may be more complex than what (1) advocates. However, it can be argued that the regression model in Section 3.3 will assign to k_b in (1) the right coefficient(s) to better estimate the more involved reality. For example, the density k_b can be scaled by a positive constant without loss of generality, as this will simply lead to a different, but equivalent, calibration of the regression model. That is, the traffic expression in (1) provides a good starting point for the regression analysis to capture the effect of a shelf location on traffic. Our empirical analysis in Section 5 attests to the predictive ability of (1). With the scalability of k_b in (1), we assume that d_{0b} and d_{bn} are considered in terms of a small unit (e.g., one foot) in order to ensure that $k_b \le 1$, which is required for the beta regression in Section 3.3.

In a grocery store, a customer follows a shopping path, from one shelf to another, until purchasing all items in the shopping list. While traveling from one shelf to another, if there is no direct path between the two, customers necessarily pass by intermediate shelves to reach their target shelf. Hence, the allocation-based traffic density Φ_b is derived from the traffic passing by shelf b in order to reach another destination. To calculate Φ_b, we follow a two-step procedure similar to Farley and Ring (1966): (i) the calculation of the force (amount of attraction) exerted on shelf q by shelf r for any pair of shelves in the store, f_{qr}; and (ii) the calculation of the force on shelf b exerted through any adjacent shelf q, F_{qb}, by aggregating forces f_{qr}, where the shortest path of a shopper walking from shelf q to shelf r passes by shelf b.

The force f_{qr} is defined as follows:

$$f_{qr} = \frac{\sum_{g\in\mathcal{G}_r} m_g x_{rg}}{d_{qr}}, \quad \forall q, r \in \mathcal{B} | q \neq r, \tag{2}$$

where $\mathcal{G}_r$ is the set of groups that can be allocated to shelf r, $x_{rg} = 1$ if and only if group g is assigned to shelf r, and m_g is the total number of shopper for the products in group g,

$$m_g = \sum_{p\in\mathcal{Z}_g}\left(\lambda_p - \sum_{\substack{q\in\mathcal{Z}_g:\\ q<p}}\lambda_{pq} + \sum_{\substack{q\in\mathcal{Z}_g:\\ q<p}}\sum_{\substack{r\in\mathcal{Z}_g:\\ r<q<p}}\lambda_{pqr} - \cdots + (-1)^{|\mathcal{P}|-1} \sum_{\substack{q\in\mathcal{Z}_g:\\ q<p}}\cdots\sum_{\substack{s\in\mathcal{Z}_g:\\ s<\cdots<q<p}}\lambda_{pq\ldots s}\right), \quad \forall g \in \mathcal{J}, \tag{3}$$

where $\mathcal{Z}_g$ is the set of products in group g and λ_w is the number of baskets containing all products in the set w. The subset w could be of any cardinality, involving pairs, triplets, quadruplets, or quintuplets of products, etc. For example, λ_p is the number of customer baskets (receipts) having product p over a selling horizon, λ_{pq} is the number of baskets that contain product pairs (p,q) together, λ_{pqr} is the number of baskets that contain product triplets (p,q,r) together, and so on. The calculation of m_g is similar to the "inclusion-exclusion principle" for calculating the cardinality of the union of sets (e.g., Ross 2014).

Let $\mathcal{R}_b$ is a set of shelves that are adjacent to shelf b—that is, it is possible to travel from shelf b to each of these shelves without passing by an intermediate shelf.

Figure 2. Illustration of the Attraction Model for Traffic

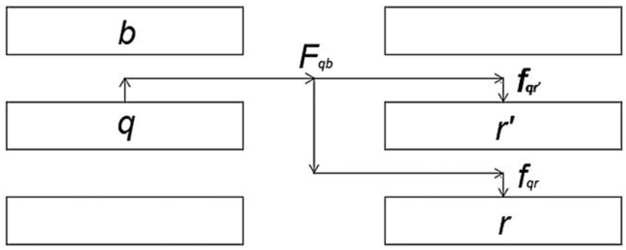

The force F_{qb} is calculated as follows:

$$F_{qb} = \sum_{r \in \mathcal{B}: r \neq b, q} \Delta_{qbr} f_{qr}, \quad \forall b \in \mathcal{B}, q \in \mathcal{R}_b, \tag{4}$$

where $\Delta_{qbr} = 1$, if the shortest path between shelves q and r includes shelf b, and $\Delta_{qbr} = 0$ otherwise. Figure 2 illustrates the estimation of F_{qb}. Finally, to calculate the allocation-based traffic density of shelf b, we aggregate the forces on all of its adjacent shelves exerted through it:

$$\Phi_b = \sum_{q \in \mathcal{R}_b} F_{qb}, \quad \forall b \in \mathcal{B}. \tag{5}$$

Equations (2)–(5) indicate that the allocation-based traffic density, Φ_b, is affected by the allocation at a wide range of other shelves, as the network of walkways usually allows easy access between any two pairs of shelves. As consequence, the model captures ripple effects in traffic, whereby altering the allocation of any shelf affects traffic in various areas of the store.

3.3. A Predictive Model for In-Store Traffic

The predictive model for the traffic density along shelf b, θ_b, is stated as a function of a layout component, k_b in (1), and an allocation component, Φ_b in (5). We use regression analysis to express θ_b as a function of k_b and Φ_b. Because $\theta_b \in [0, 1]$, it is not suitable to use a linear regression. Instead, we use a logit transformation of the dependent variable, θ_b, which yields,

$$\theta_b = \frac{e^{(\alpha + \beta_1 k_b + \beta_2 \Phi_b^{\text{norm}})}}{1 + e^{(\alpha + \beta_1 k_b + \beta_2 \Phi_b^{\text{norm}})}}, \quad \forall b \in \mathcal{B}, \tag{6}$$

where $\Phi_b^{\text{norm}} = \frac{\Phi_b}{\sum_{q \in \mathcal{B}} \Phi_q}$ is the normalized value of Φ_b. In performing the regression analysis in (6), we find it useful to follow the recommendation of Ferrari and Cribari-Neto (2004) to utilize a beta regression, where the dependent variable is assumed to have a beta distribution.

The objective function can now be estimated based on the probability-tree model for customer impulse buying depicted in Figure 3. This probability tree can help identify the likelihood that a customer impulsively purchases product p from its shelf b, where it has been assigned a shelf space s_{pb} such that $\ell_p \leq s_{pb} \leq u_p$. To illustrate, a customer visits shelf b with probability θ_b and then notices product p with a probability that is proportional to s_{pb}, its allocated shelf space (similar to other works in the literature, e.g., Yang and Chen 1999, Botsali 2007, and Flamand et al. 2016). A customer who notices p would purchase it impulsively with a probability i_p, which reflects the likelihood that the product, when made visible to a consumer, would stimulate an impulse purchase. This may be estimated via customer surveys.

4. Optimizing Store-Wide Shelf-Space Allocation

We propose an optimization model that involves the predictive traffic-density construct developed in Section 3. The problem of splitting the space of a given shelf among products in a given group is addressed and solved to optimality with a simple algorithm, in Section 4.1. Then, in Section 4.2, the model is cast as a nonlinear mixed-integer program that assigns groups of products to shelves, with product-space amount determined by the algorithm in Section 4.1. For this mathematical program, we propose a solution methodology that linearizes constraints and approximates nonlinear terms in the objective function.

Figure 3. Probability Tree Model of Customer Impulse-Purchase Behavior for a Product p Assigned to Shelf b

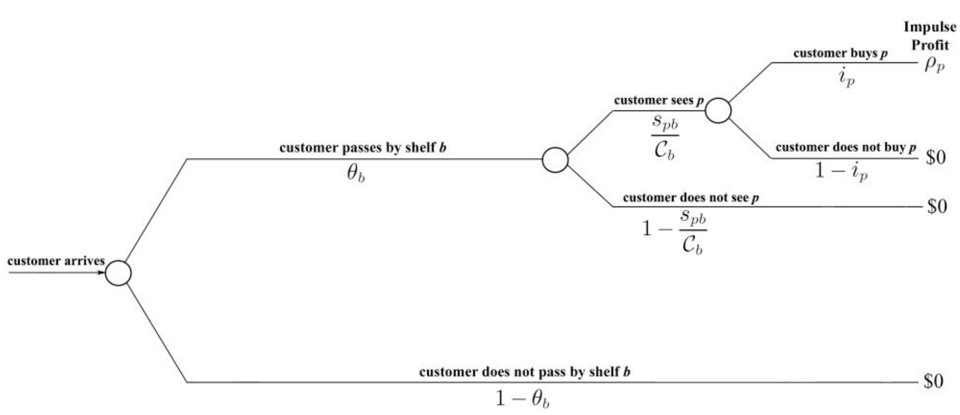

4.1. Structure of Group-Shelf-Space Allocation

The first step of our proposed methodology is to determine the optimal shelf space that should be allocated to products in a group g on any shelf b, considering any feasible group-shelf assignments. This requires solving a continuous knapsack problem for any such (g, b) pair. The following greedy algorithm, denoted by Algorithm 1, provides an optimal solution to this problem, which is facilitated by the existence of lower and upper bounds (ℓ_p and u_p, $\forall p \in \mathcal{P}$) on the space that can be assigned to a product. These bounds can be determined by the management (Yang and Chen 1999) or dictated by historical demand levels. In particular, lower bounds may be enforced on shelf space for one or several of the following reasons: (i) to project a specific "store image," irrespective of immediate profitability (Corstjens and Doyle 1981); (ii) to bring attention to new products that are added to the assortment (Corstjens and Doyle 1981, Yang 2001); (iii) to advertise and give greater visibility to a specific product that is carried in the store (Borin et al. 1994, Amrouche and Zaccour 2007); or (iv) to avoid stock-outs (Drèze et al. 1994). Upper bounds, in contrast, may be set for products at a later stage of the life cycle to keep the store up-to-date (Corstjens and Doyle 1981) and to leave more space for refreshing the product mix in the store (Yang 2001).

Algorithm 1 first allocates the minimum shelf-space requirement, ℓ_p, for any product category p. Thereafter, it sorts the product categories in the group by their nonincreasing value (contribution to the impulse buying) per unit of space. Starting from the product category having the highest value, the algorithm assigns shelf space to product categories in turn, as much as possible, up to their maximum requirements (i.e., u_p) and until the entire shelf space has been filled.

Algorithm 1 (Space-Allocation Algorithm for (g, b))

Sort product categories in group g ($\forall p \in \mathcal{Z}_g$) by their decreasing value per unit of space (i.e., $\frac{\rho_p i_p}{\mathcal{C}_b}$) and assign them to a sorted set $\mathcal{Z}'_g$
Initialize the total occupied space: $S = 0$
for all $p \in \mathcal{Z}_g$ **do**
  $s_{pb} \leftarrow \ell_p$
  $S \leftarrow S + s_{pb}$
end for
for all $p \in \mathcal{Z}'_g$ **do**
  if ($S = \min(\mathcal{C}_b, \sum_{p\in\mathcal{Z}_g} u_p)$) **then**
    break
  end if
  $s_{pb} \leftarrow \ell_p + \min(u_p - \ell_p, \mathcal{C}_b - S)$
  $S \leftarrow \sum_{p\in\mathcal{Z}_g} s_{pb}$
end for
Report the s_{pb} values for each $p \in \mathcal{Z}_g$

The optimality of Algorithm 1 is established in Lemma 1.

Lemma 1. *Algorithm 1 yields an optimal solution to this continuous knapsack problem.*

Proof. See Section EC.3 of the online companion. □

Remark 1. The following nonlinear knapsack model can be employed to account for shelf-space elasticity in the allocation of product-category space for group g along shelf b, in lieu of Algorithm 1 (where linearity is assumed): {Maximize $\sum_{p\in Z_g} \rho_p i_p \frac{s_{pb}^{\gamma}}{C_b}$: $\sum_{p\in\mathcal{Z}_g} s_{pb} \le \mathcal{C}_b$, $\ell_p \le s_{pb} \le u_p$, $\forall p \in Z_g$}, where γ is a parameter that can be set between 0.1 and 0.2, based on empirical evidence in the academic literature (e.g., Irion et al. 2011 and Eisend 2014). However, our results for the supermarket in Beirut indicate that the optimized product space allocations obtained via Algorithm 1 are optimal in most cases for the nonlinear knapsack model and occasionally near-optimal with a deviation below 0.2% from the optimal nonlinear objective value.

4.2. Base Optimization Model and Simplification

We first propose a nonlinear mixed-integer optimization model in Section 4.2.1, which is then simplified in Section 4.2.2 using linearization and approximation schemes.

4.2.1. Base Model. We develop an optimization model that allocates individual product categories to capacitated shelves in a store. The impulse-buying maximization problem under a given store layout is modeled as the following mixed-integer nonlinear program, where s_{pb} values in the objective function are predetermined using Algorithm 1:

$$\textbf{IBMP}: \text{Maximize} \sum_{b\in\mathcal{B}} \sum_{g\in\mathcal{G}_b} \sum_{p\in\mathcal{Z}_g} \rho_p i_p \frac{s_{pb}}{\mathcal{C}_b} x_{bg} \theta_b, \tag{7a}$$

$$\text{subject to} \sum_{b\in\mathcal{B}|g\in\mathcal{G}_b} x_{bg} = 1, \quad \forall g \in \mathcal{J}, \tag{7b}$$

$$\sum_{g\in\mathcal{G}_b} x_{bg} = 1, \quad \forall b \in \mathcal{B}, \tag{7c}$$

$$f_{qr} = \frac{\sum_{g\in\mathcal{G}_r} m_g x_{rg}}{d_{qr}}, \quad \forall q, r \in \mathcal{B} | q \ne r, \tag{7d}$$

$$F_{qb} = \sum_{r\in\mathcal{B}: r\ne b, q} \Delta_{qbr} f_{qr}, \quad \forall b \in \mathcal{B}, q \in \mathcal{R}_b, \tag{7e}$$

$$\Phi_b = \sum_{q\in\mathcal{R}_b} F_{qb}, \quad \forall b \in \mathcal{B}, \tag{7f}$$

$$\Phi_b^{\text{norm}} = \frac{\Phi_b}{\sum_{q\in\mathcal{B}} \Phi_q}, \quad \forall b \in \mathcal{B}, \tag{7g}$$

$$\theta_b = \frac{e^{(\alpha+\beta_1 k_b+\beta_2 \Phi_b^{\text{norm}})}}{1 + e^{(\alpha+\beta_1 k_b+\beta_2 \Phi_b^{\text{norm}})}}, \quad \forall b \in \mathcal{B}, \tag{7h}$$

$$x \text{ binary}, f, F, \Phi, \Phi^{\text{norm}}, \theta_b \ge 0. \tag{7i}$$

Figure 4. (Color online) Outer Approximations for Different k_b and z_1 Values

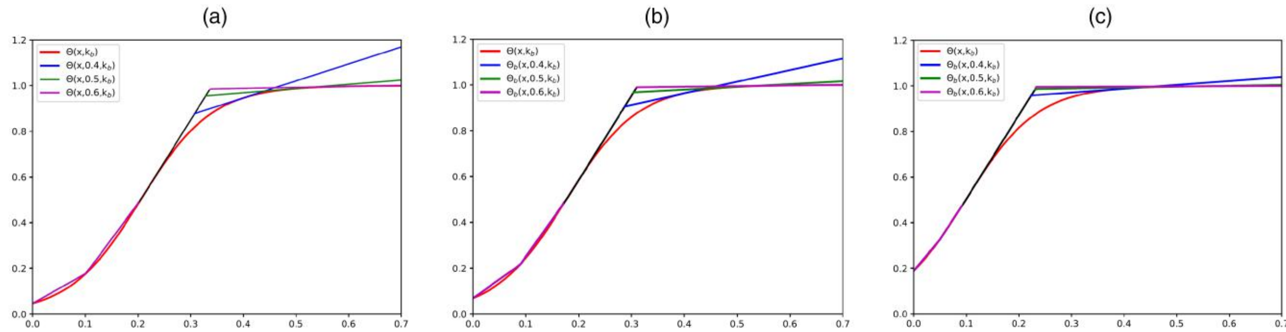

Notes. (a) $k_b = 0.048$, $z_1 = 0.4, 0.5, 0.6$. (b) $k_b = 0.124$, $z_1 = 0.4, 0.5, 0.6$. (c) $k_b = 0.333$, $z_1 = 0.4, 0.5, 0.6$.

The Objective Function (7a) maximizes the average impulse-buying profit per customer, which is estimated following the customer behavior depicted in Figure 4. Constraint (7b) guarantees that each group of products is assigned to a single shelf, whereas Constraint (7c) ensures that each shelf accommodates exactly one group. Constraints (7d–g) inform the calculation of the allocation-based traffic density, as discussed in Section 3.2. Constraint (7h) is a predictive construct that estimates the traffic density of each shelf, as discussed in Section 3.3. Constraint (7i) enforces logical binary and nonnegativity restrictions on the decision variables.

4.2.2. Approximated Model. The formulation developed in Section 4.2.1 poses a computational challenge to modern-day solvers, due to the combinatorial nature of the problem and its nonlinearity. To this end, we propose in this section an approximation that yields an upper bound for the traffic-density variable θ_b. We otherwise keep bilinear terms that arise in the formulation, which can be handled by more recent versions of the commercial solver GUROBI. To build the approximation, the density θ_b in (6) is rewritten as

$$\theta_b(z_b) = \frac{a_b e^{hz_b}}{1 + a_b e^{hz_b}}, \quad 0 \le z_b \le 1, \tag{8}$$

where $a_b = e^{(\alpha+\beta_1 k_b)}$, $z_b = \Phi_b^{\text{norm}}$, $h = \beta_2$.

The following lemma avails of a convexity analysis on $\theta_b(z_b)$ presented in Section EC.2 of the online companion in order to develop useful upper bounds. Let $\tilde{z_b} = -ln(a_b)/h$ be the x-axis coordinate for the inflection point of $\theta_b(z_b)$, as defined in the online companion. Furthermore, define b_{1Lb} as the slope of the line connecting the two points $(0, \theta_b(0))$ and $\left(\frac{\tilde{z_b}}{2}, \theta_b\left(\frac{\tilde{z_b}}{2}\right)\right)$, and define b_{2Lb} and I_b as the slope and intercept of the line connecting the two points $\left(\frac{\tilde{z_b}}{2}, \theta_b\left(\frac{\tilde{z_b}}{2}\right)\right)$ and $(\tilde{z_b}, \theta_b(\tilde{z_b}))$. In addition, define $t_b(z_0)$ as the equation of a line tangent to $\theta_b(z_b)$ at a point with an x-axis coordinate equals to z_0.

Lemma 2. *Suppose* $0 < \tilde{z_b} < 1$. *For any* $z_1 \in (\tilde{z_b}, 1]$, *an upper-bound approximation of* $\theta_b(z_b)$ *is*:

$$\bar{\theta_b}(z_b) = \begin{cases} \theta_b(0) + b_{1Lb} z_b, & z_b \in \left[0, \frac{\tilde{z_b}}{2}\right] \\ I_b + b_{2Lb} z_b, & z_b \in \left(\frac{\tilde{z_b}}{2}, \tilde{z_b}\right] \\ \min[t_b(\tilde{z_b}), t_b(z_1)], & z_b \in (\tilde{z_b}, 1] \end{cases}. \tag{9}$$

Proof. See Section EC.3 of the online companion. □

One issue that needed to be resolved before applying Lemma 2 to obtain upper-bound approximations is the selection of the tangential point $z_1 \in (\tilde{z_b}, 1]$. One can try out several values of z_1 for different a_b. In our case, we found that the approximation works well for $z_1 = 0.5$, as illustrated in Figure 4.

Using the upper-bound approximation in Lemma 2, θ_b in Constraint (7h) of Model IBMP is replaced by the linear constraints given in Section EC.4 of the online companion.

The Objective Function (7a) in the Model IBMP involves bilinear terms $x_{bg}\theta_b$, which can be handled by recent versions of the commercial solver GUROBI. Similarly, Constraints (7d–g) can be restated as follows (see Section EC.4 of the online companion):

$$\sum_{b_2 \in \mathcal{B}} \sum_{q \in \mathcal{R}_{b_2}} \sum_{r \in \mathcal{B} | r \neq q \& b_2} \Delta_{qb_2r} \frac{\sum_{g \in \mathcal{G}_r} m_g x_{rg} \Phi_b^{\text{norm}}}{d_{qr}} = \sum_{q \in \mathcal{R}_b} \sum_{r \in \mathcal{B} | r \neq b \& q} \Delta_{qbr} \frac{\sum_{g \in \mathcal{G}_r} m_g x_{rg}}{d_{qr}}, \quad \forall b \in \mathcal{B}, \tag{10}$$

where the bilinear terms $x_{rg}\Phi_b^{\text{norm}}$ can also be handled by an adequate solver, such as GUROBI.

The overall approximated model IBMP with bilinear terms, denoted by IBMPa, has the same objective function as IBMP, (7a), and is subject to (7b), (7c), (10), and other linear approximation constraints, as shown in Section EC.3 of the online companion.

In the absence of a suitable solver, the bilinear terms in IBMPa can be linearized using auxiliary variables and related constraints (see, for example, Ghoniem and Maddah 2015). In our experience, however, the direct use of GUROBI to solve Model IBMPa, as opposed to a fully linearized model, yields a better optimality gap (5% versus 21%) within the same time limit (of 10 CPU hour). Solving Model IBMPa with a time limit not only provides a valid optimality gap on its solution, but also yields a valid optimality gap on the exact (nonlinear) objective value (obtained by recalculating the approximated model's solution objective value using the nonlinear objective function). This is achieved due to the fact that the approximated objective value is an upper bound on the exact nonlinear one (Lemma 2).

5. Application to a Supermarket in Beirut

Our study is based on real data that we obtained from the store management of what we refer to, for confidentiality, as the "B Supermarket" in Beirut. B Supermarket currently offers 64 product categories that are classified into 21 major departments, as described in Section EC.5 of the online companion. In terms of the shelf configuration, B Supermarket has a grid layout and involves 34 shelves (four refrigerator shelves, two freezers, and 28 ordinary shelves, as illustrated in Figure 1. We also describe the estimation of the input parameters on product impulse-purchase rates, i_p and profit margins, ρ_p in Section EC.5 of the online companion. Our main data set on product demand includes 40,000 customer receipts that have been collected over two months. We present descriptive analytics and visualization of this data in Section EC.5 of the online companion. The main takeaway from this analysis is that the number of baskets with a large number of categories is limited. As such, when estimating the shelf traffic resulting from product allocation, m_b in (3), we limit the analysis to quintuplets.

Next, we analyze in-store traffic of B Supermarket and develop a regression predictive model, as discussed in Section 3.3. To estimate the independent variable related to the allocation-based traffic density of each shelf b, Φ_b, Equations (2)–(5) in Section 3.2 are used. The other independent variable related to layout-based traffic, k_b, has been estimated based on the proximity of the shelf to the entrance and exit (measured in feet) in Equation (1) of Section 3.2. To estimate the dependent variable, θ_b, we conduct a traveling salesman problem type of analysis based on our basket data and the current allocation. The TSP analysis is based on the following assumptions:

- Each customer starts shopping from the entrance of the store and finishes it at the exit.
- Each customer stops at the cashier before going to the exit.
- Customers shop the store in a way that minimizes their walking distance.

Based on these assumptions, a shopping path is created for each basket (receipt) in our data set by solving a classical TSP formulation using CPLEX. A total of $n = 40{,}000$ TSPs have been solved. To illustrate, Figure 1 shows, in two distinct line patterns, the shopping paths of two customers based on their respective TSP solutions. Based on the receipt, the shelves where a customer shops are marked by the circles with the matching pattern. Each TSP resulted in a shopping path that minimizes the customer's travel distance, such that all marked shelves are visited, starting from the entrance and ending at the exit. In our experience, the optimization effort associated with solving such TSPs with 20–30 categories per receipt (as in our data set) is below two CPU seconds using modern-day solvers. However, for scalability purposes, it can be pertinent to substitute the use of a solver with an efficient TSP heuristic, such as that by Lin and Kernighan (1973).

Repeating the analysis in Figure 1 for each basket, based on the TSP-generated shopping paths, all visited shelves are determined. Visited shelves include the shelves where the customer goes to buy a product and the ones she passes by to reach another shelf. For example, in Figure 1, Customer 2 passes by the shelf that is highlighted in gray without purchasing anything from it. By analyzing all baskets (receipts) in our data in this manner, we obtain the total number of customers that visited shelf b, n_b, $\forall b \in \mathcal{B}$. The customer-traffic density of shelf b is then estimated as $\theta_b = \frac{n_b}{n}$. The resulting traffic densities are shown in Figure 5, where low, medium, high, and very high categories have the traffic-density ranges of (0, 0.07), [0.07, 0.10), [0.10, 0.20), and [0.2, 1), respectively. In Figure 5 and subsequent figures, fast-moving products are marked by circles, and high-impulse products are marked by squares. High-impulse products are determined by sorting products based on their impulse-purchase rates and selecting the top 10. For determining fast-movers, we first calculate the planned demand of product $p \in \mathcal{Z}_g$ by isolating the impulse portion of its sales—that is, $\lambda_p - \sum_{b\in\mathcal{B}}\lambda_p i_p \theta_b \frac{s_{pb}}{C_b} x_{bg}$. Then, the products are sorted based on their planned demand volume, and the top 10 products are considered. Note that a few products appear on both lists, and we keep them only on the list in which it appears at a higher ranking. After eliminating the products that overlap, we finally have eight high-impulse products and nine fast-movers.

Using our TSP methodology is highly advantageous for the retailers to delineate the approximate shopping paths of customers and estimate traffic in different areas of the store, without utilizing expensive tools such as RFID technology. Note that figuring the density of traffic across the store has several practical usages for in-store operations and promotions, which extend beyond the impulse-buying scope of this paper. Our TSP framework offers a low-cost, analytical approach for measuring store traffic.

Figure 5. (Color online) Current Shelf-Space Allocation at B Supermarket and Estimated In-Store Traffic

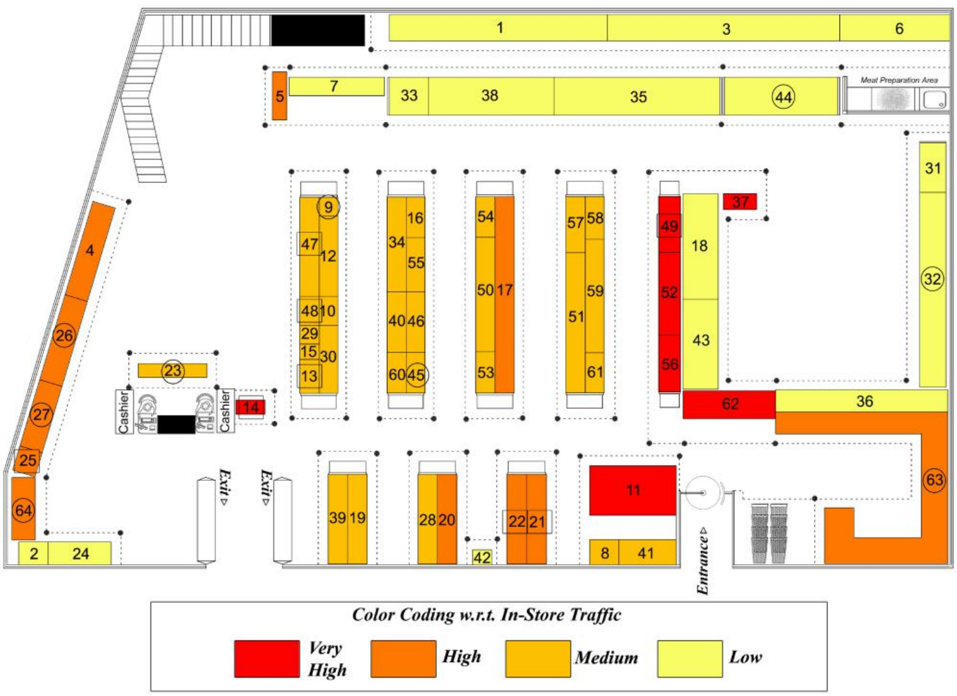

Once we estimate the dependent variable, θ_b, and the independent variables related to the allocation (Φ_b^{norm} in (5)) and the layout (k_b in (1)) traffic components, we perform the beta regression described in Section 3.3, which yields the results in Table 3. Because the associated p-values for all regression coefficients are less than 0.05, we consider that these coefficients are statistically significant.

Predictive models such as that in Table 3 can be very useful in practice for understanding the effect of product allocation on traffic. To demonstrate the predictive power of the model in Table 3, we modify the current allocation of Figure 5 and estimate the new store traffic. We have examined two alternative input configurations: (i) Most of the fast-movers are allocated to the back of the store, referred to as "Fast Movers Back"; and (ii) fast-movers are allocated to different aisles with the purpose of spreading them throughout the store, referred to as "Fast Movers Spread." These alternative input configurations are depicted in Figures EC.5 and EC.6 in Section EC.6 of the online companion, respectively. Among the current allocation, the Fast Movers Back, and the Fast Movers Spread input layouts, Fast Movers Back provides the highest value of impulse profit, as shown in the second column of Table 4, with an estimated impulse profit of $1.38 per customer. This is calculated utilizing the exact nonlinear objective function in IBMP.

Table 3. Results of the Beta-Regression Analysis

Parameter	Coefficient	p-value
$\boldsymbol{\Phi}_b^{\text{norm}}$	$\beta_1 = 14.751$	0.000
$\boldsymbol{k_b}$	$\beta_2 = 5.491$	0.004
Constant	$\alpha = -3.285$	0.000

6. Findings and Managerial Insights

With the necessary input parameters estimated and validated, as discussed in Section 5, we now apply our store-wide shelf-space-optimization model of Section 4.2 to B Supermarket.

6.1. Optimized Solution for B Supermarket and Managerial Insights

The three input configurations discussed in Section 5 provide examples of rule-of-thumb managerial policies, or directives, that could potentially increase the traffic density in the store. We utilize these input configurations as instances for Model IBMPa described in Section EC.4 of the online companion, which we ran with a time limit of one CPU hour. The results are summarized in Table 4, where the Input Objective column reports the objective

Table 4. Comparison of the Initial and Optimal Objective Values of Different Settings

Input configuration	Input objective ($)	Optimized objective ($)
Current Allocation	1.14	1.59
Fast Movers Back	1.38	1.89
Fast Movers Spread	1.15	1.59

values of the aforementioned three input configurations estimated using the original nonlinear objective function, and the Optimized Objective column in Table 4 displays the objective values where the allocation of the groups in each setting is optimized using our proposed approximation methodology in Section 4.2.2, and the reported objective values are postcalculated using the original nonlinear objective function.

Our findings in the second column of Table 4 demonstrate that *ad hoc* or manual alterations of the Current Allocation can bring about relative improvements in the impulse-buying profit. However, more importantly, the results in the third column indicate that it is only through optimization that the impulse-purchase potential of any input configuration is fully attained. Given the combinatorial nature of the problem, it is practically impossible for managers to track in an *ad hoc* manner the direct and indirect impacts of swapping shelf allocations on in-store traffic.

Next, we analyze in greater detail the Current Allocation, which has a baseline customer impulse profit margin of $1.14, against the optimized Fast Movers Back configuration, which yields the greatest impulse-purchase profit in our study of $1.89, amounting to a 65% improvement. The corresponding solution is depicted in Figure 6, where the symbols (+), (−), and (0) represent whether the traffic density of a shelf increased, decreased, or remained unchanged with respect to the Current Allocation. The following observations are in order:

• In the optimized allocation, *fast-movers are assigned to shelves with relatively low traffic density when compared with the Current Allocation*. For example, Juice (#26) and Soda (#27) are moved from the leftmost shelf in the current allocation in Figure 5 to a shelf having a lower traffic density in the back of the store in the optimized allocation in Figure 6. Similarly, Bread (#9) is reallocated to a shelf having a lower traffic density in the optimized allocation (left of the entrance at the lower right corner in Figure 6), compared with the current allocation, where it is at the right of the cashier in Figure 5. This has the desirable effect of increasing the traffic density along the shelf where Chocolate Chips (#22), a high-impulse category, is allocated (opposite to Bread (#9) in Figure 6).

• *Secondary (ripple) effects*: *Allocating fast-movers to the back of the store strengthens traffic along middle shelves.* For example, moving Juice (#26) and Soda (#27) from the lower left corner in Figure 5 to the back of the store (upper right corner) in Figure 6 clearly drives traffic through the middle shelves, especially to the right of the cashier. This illustrates the importance of the traffic model and the manner in which it tracks direct and indirect allocation effects throughout the entire network of walkways in the store.

• In the optimized allocation, s*ome high-impulse products may actually be less visible*. For example, Iced Tea (#25) is moved from the lower left corner in Figure 5 to the upper right corner in Figure 6, and Popcorn (#49) is moved from rightmost middle shelf in Figure 5 to two shelves to the left in Figure 6. This is due to the aforementioned secondary ripple effects. Moving Iced Tea and Soda in this manner diminishes their own impulse

Figure 6. (Color online) Optimized Solution for the "Fast Movers Back" Input Configuration for B Supermarket

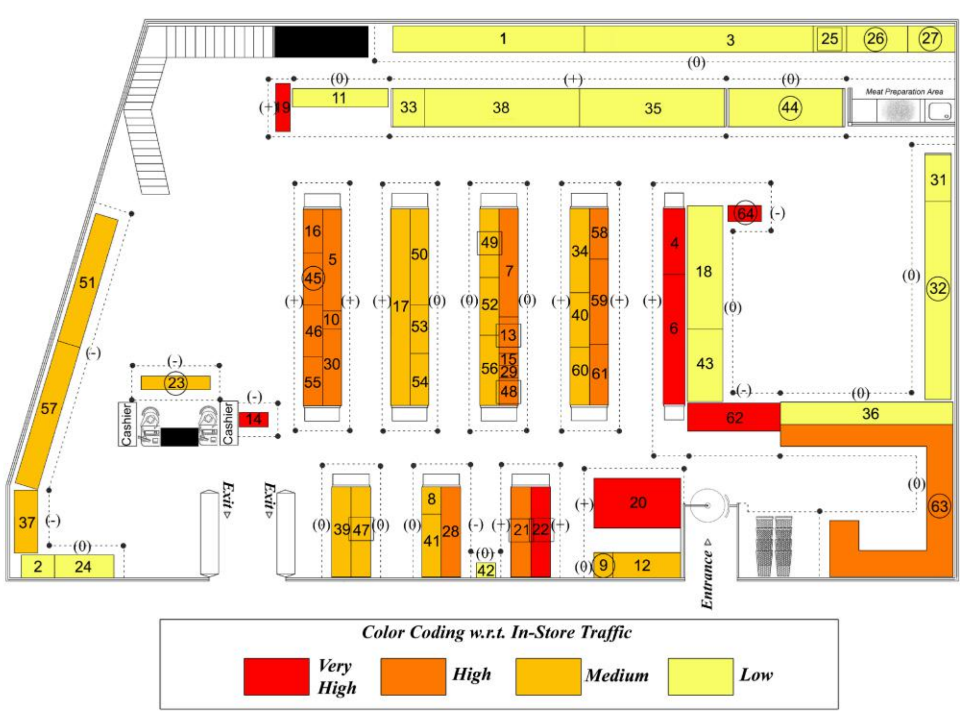

Table 5. The Effect of Static Location vs. Variable Location on the Solution

Configuration	Location	Allocation	Objective value ($)	Improvement (%)
Input allocation	Static	$(\ell_p = u_p)$	1.38	0.0
Configuration 1	Variable	$(\ell_p = u_p)$	1.82	31.9
Configuration 2	Static	$(\ell_p < u_p)$, 10%	1.43	3.6
Optimized allocation	Variable	$(\ell_p < u_p)$, 10%	1.89	37.0

sales, while creating overall better traffic and impulse sales elsewhere in the store.

To further study the value of our approach, we compare our methodology herein against that in Flamand et al. (2016) using B Supermarket. In Flamand et al. (2016), a simplifying assumption is made, whereby the product visibility to the consumer depends solely on the *exogenous* attractiveness of the shelf (which depends on its location in the store layout) and the amount of shelf space it is allocated along the shelf. In contrast, the proposed approach in the current study captures in a more elaborate manner how the allocation to one shelf directly drives traffic to it, but also indirectly through ripples boosting sales at other intermediate shelves. In our experience, the richer model herein outperforms that in Flamand et al. (2016) with a 19% improvement in the prescribed solution for the B Supermarket.

6.2. Sensitivity Analysis

We also perform an analysis under alternative parameter configurations, always starting with the Fast Movers Back input layout. Tables 5–8 and Table EC.3 in Section EC.7 of the online companion report the results of our analysis, by solving the approximated linear model by GUROBI. The first row of Table 5 reports the objective value of the Fast Movers Back input layout without any optimization. Configuration 1 investigates the problem, where the size of each product is fixed, $\ell_p = u_p$, $\forall p \in \mathcal{P}$, and we only decide, via the optimization model in Section 4.2, the assignment of product categories to shelves (variable location). Configuration 2 examines the problem where we only decide on the shelf-space amount between ℓ_p and u_p (at $\pm 10\%$ of the current allocation) for each product p, assuming that p stays on its currently assigned shelf (static location).

Table 5 reveals the following: (i) In configuration 1, when the problem is solved for the fixed shelf-space requirements ($\ell_p = u_p$), the shelf-product assignments yield a 32% improvement over the input allocation (which is close to that of the optimized allocation reported in the last row of Table 5); (ii) in configuration 2, in the reverse case, where the location of each product is fixed, deciding only on shelf-space amount of each product p between ℓ_p and u_p yields a modest 3.6% improvement over the input allocation. In conclusion, *the main contribution to the improvement of the optimized allocation is due to the product location, rather than the amount of shelf space assigned to the product*. This is interesting, as it challenges most of the findings in the open literature, whereby the location of products is essential fixed (same shelf), and optimization is done on the amount of shelf space.

In addition, we test configurations 3 and 4 in Table 6, where the ℓ_p and u_p values are further varied by $\pm 15\%$ and $\pm 20\%$, respectively, from the current allocation. Table 6 confirms the insight that location matters much more than shelf space, as modest additional improvements are achieved when the shelf-space constraints are loosened.

Next, we test the effect of impulse-purchase rates and profit margins on the solution quality. Tables 7 and 8 report our findings. In Table 7, we consider configurations 5–8. To generate these configurations, products are ordered according to the nondecreasing order of their impulse rates, i_p. Impulse rates of the first half (high-impulse products) are increased by 5% and 10%, respectively, whereas those of the second half (low-impulse products) are decreased by 5% and 10%, respectively, for configurations 5 and 6. For configurations 7 and 8, this process is reversed; impulse rates of the first half are decreased by 5% and 10%, respectively, whereas those of the second half are increased by 5% and 10%, respectively.

Table 7 indicates that the solution is not sensitive to the impulse-purchase rates of low-impulse products. As such, *the impulse-rate-estimation effort could focus mainly on high-impulse products.* This is useful, given the difficulty in estimating impulse rates.

In Table 8, we compare configurations 9–12. To generate these instances, products are ordered according to the nondecreasing order of their impulse rates, i_p. Profit margins, ρ_p, of the first half are increased by 5% and

Table 6. The Effect of the Flexibility in the Minimum and Maximum Space Requirements on the Solution

Configuration	Location	Allocation	Objective value ($)	Improvement (%)
Configuration 3	Variable	$(\ell_p < u_p)$, 15%	1.93	39.9
Configuration 4	Variable	$(\ell_p < u_p)$, 20%	1.96	42.0

Table 7. The Effect of Impulse-Purchase Rates on the Solution

Configuration	High-impulse products (%)	Low-impulse products (%)	Objective value ($)	Improvement (%)
Configuration 5	+5	−5	1.94	40.6
Configuration 6	+10	−10	1.99	44.2
Configuration 7	−5	+5	1.84	33.3
Configuration 8	−10	+10	1.78	29.0

10%, respectively, whereas those of the second half are decreased by 5% and 10%, respectively, for configurations 9 and 10. For the other two configurations, this process is reversed; profit margins of the first half are decreased by 5% and 10%, respectively, whereas those of the second half are increased by 5% and 10%, respectively. Table 8 reveals that the solution is not too sensitive to the profit margins. All in all, the analysis in Tables 5–8 indicate that the model is robust. *Small changes of the order of 5%–10% in the key input parameters do not lead to significant loss of profitability.* This is comforting in practical applications, where a highly accurate estimation of some parameters, like the impulse-purchase rates, may not be possible. Further sensitivity analysis on the demand volume for product groups, m_g, confirms these observations (see the online companion, Section EC.7).

7. Conclusions and Directions for Future Research

The central question investigated in this paper is how to manage store-wide shelf space in a manner that guides in-store traffic and boosts impulse buying. To this end, we analyze data pertaining to product categories, their allocations, and sales using 40,000 receipts from a grocery store in Beirut with the reasonable assumption that customers shop the store in a way that minimizes walking, as in the classical traveling salesman problem. As a result, we develop a predictive nonlinear regression model that relates in-store traffic to the store layout and product shelf-space allocations. This predictive model is of practical value on its own, as it allows quantifying the effect of product allocation on traffic, and, consequently, impulse profit. Moreover, the predictive model serves as a cornerstone for a mixed-integer nonlinear optimization model that pursues optimized shelf-space allocations. Because of the nonlinearity and computational challenges that the proposed optimization model poses, we linearize it by employing a piece-wise linear approximation of the regression construct, while keeping bilinear terms in the model, which can be directly handled by recent versions of GUROBI.

The approximated optimization model is useful in obtaining a dual bound, allowing for the estimation of an optimality gap. We demonstrate the usefulness of our methodology by implementing it on the grocery store in Beirut. The enhanced shelf-space allocation potentially yields up to 65% improvement in the impulse-buying profit per customer over the current allocation of the store. Based on our prescriptive analysis, from a qualitative viewpoint, the optimized allocation advocates two main insights: (i) assigning fast-movers to locations of the stores that typically have low traffic; and (ii) assigning products in a way as to maximize overall store traffic, where often ripple effects, such as driving traffic to neighboring shelves, are important.

We also conduct a detailed sensitivity analysis over the base values of the key input parameters estimated from our Beirut store case study. The main insight we obtain from the analysis is that the improvement in impulse-buying profitability that we observe is due mostly to optimizing the location of different products in the store, rather than adjusting the shelf space of each product—that is, it is (almost) all about location. The analysis also reveals that the model is robust in the sense of not being too sensitive to small errors in the estimation of input parameters, which is an important feature in practice, where highly accurate estimates may be hard to achieve.

The key novelty in this work stems from the in-store traffic model, which incorporates the ripple effects of any change in allocation decisions at one shelf throughout

Table 8. The Effect of Profit Margins on the Solution

Configuration	High-impulse products (%)	Low-impulse products (%)	Objective value ($)	Improvement (%)
Configuration 9	+5	−5	1.95	41.3
Configuration 10	+10	−10	2.00	44.9
Configuration 11	−5	+5	1.84	33.3
Configuration 12	−10	+10	1.78	29.0

the entire store, a major departure from the extant literature. A common, but simplifying, assumption in the literature views the product visibility to the consumer as a function of the exogenous attractiveness of the shelf (based on its location in the store layout) and the amount of shelf space it is given (e.g., in Botsali 2007 and Flamand et al. 2016). This, however, ignores store-wide network effects and how the traffic directed to one shelf influences traffic in neighboring shelves, in particular, and even more distant shelves in the store, in general. The proposed in-store traffic construct therefore achieves improved results, in comparison with our previous work in Flamand et al. (2016), due to its ability to track in-store traffic globally, with the aforementioned store-wide ripple effects of allocation decisions. As such, the proposed methodology can be valuable for practitioners, in that it takes as an input current or candidate store configurations, with pregrouped product categories that reflect product affinity, cross-selling potential, and managerial imperatives, and optimizes the store-wide allocation of grouped products to shelves, guided by the underlying mixed-integer nonlinear in-store traffic model. The exponential number of solutions render this effort, of course, beyond reach for manual decision making or total enumeration, thereby highlighting the value of the predictive (in-store traffic modeling) and prescriptive (store-wide allocation) analytics presented in this work.

The proposed methodology relies on three key data inputs: (i) the store layout; (ii) customer receipts as a practical, readily available alternative to the more technologically involved and onerous sensory devices and tracking robot; and (iii) a survey of the managers and customers to estimate the impulse rates of products. These three key input parameters inform the proposed mathematical models and make this approach adequate for a variety of grocery stores, including in Western nations. One can, in fact, argue that larger data sets (of the type needed for our methodology) may be readily available at major Western stores, based on sensors and RFID tags mounted on carts, hand scanners used by shoppers in some stores, and mobile apps.

We recommend for future research an analysis of end caps, which we did not address in this paper. In particular, it would be interesting to investigate how end caps could be used for promotions without negatively affecting impulse buying. (The downside of end caps is that they could induce customers to skip an aisle altogether, as a result of the "deal" at the end.) A related interesting direction of future research is that of cross-selling, along the same shelf or across aisle, whereby placing two or more complementary categories in close proximity to one another (utilizing end caps or regular shelves) could further boost sales (e.g., Ghoniem et al. 2016c and Maddah and Bish 2009). In particular, as one type of sales (e.g., impulse) could lead to another (e.g., cross-selling), it would be interesting to analyze the "compounding" effect of cross-selling and impulse buying on demand. In addition, it is tacitly assumed in this paper that customers do not mind the extra walking that may be needed to search for (the typically fast-moving) products on their shopping list. In reality, such search may come at a cost. Adding such a search cost to our model is an interesting area for future research.

The estimation of some of our model parameters may be alternatively examined in future research given some recent works—for example, the product-visibility estimation in Mowrey et al. (2019). Finally, the current paper lays important foundations for the optimization of retail physical store layouts, in terms of shelf location, size, and orientation.

Acknowledgments

The authors thank Brenda Allen-Toon, Insights Consulting at Walmart, for useful discussions and input. The authors also acknowledge the valuable assistance of David Abou Chacra, Dany Abou Jaoude, and Elias Mkanna in collecting the data set utilized in this research.

References

Abratt R, Goodey SD (1990) Unplanned buying and in-store stimuli in supermarkets. *Managerial Decision Econom.* 11(2):111–121.

Amrouche N, Zaccour G (2007) Shelf-space allocation of national and private brands. *Eur. J. Oper. Res.* 180(2):648–663.

Anderson EE, Amato HN (1974) A mathematical model for simultaneously determining the optimal brand-collection and display-area allocation. *Oper. Res.* 22(1):13–21.

Bell DR, Corsten D, Knox G (2011) From point of purchase to path to purchase: How preshopping factors drive unplanned buying. *J. Marketing* 75(1):31–45.

Bellenger DN, Robertson DH, Hirschman EC (1978) Impulse buying varies by product. *J. Advertising Res.* 18(6):15–18.

Bianchi-Aguiar T, Hübner A, Carravilla MA, Oliveira JF (2021) Retail shelf space planning problems: A comprehensive review and classification framework. *Eur. J. Oper. Res.* 289(1):1–16.

Bianchi-Aguiar T, Silva E, Guimarães L, Carravilla MA, Oliveira JF (2018) Allocating products on shelves under merchandising rules: Multi-level product families with display directions. *Omega* 76:47–62.

Borin N, Farris PW, Freeland JR (1994) A model for determining retail product category assortment and shelf space allocation. *Decision Sci.* 25(3):359–384.

Botsali AR (2007) Retail facility layout design. PhD dissertation, Texas A&M University.

Campo K, Gijsbrechts E, Goossens T, Verhetsel A (2000) The impact of location factors on the attractiveness and optimal space shares of product categories. *Internat. J. Res. Marketing* 17(4):225–279.

Chen YL, Chen JM, Tung CW (2006) A data mining approach for retail knowledge discovery with consideration of the effect of shelf-space adjacency on sales. *Decision Support Systems* 42(3):1503–1520.

Corstjens M, Doyle P (1981) A model for optimizing retail space allocations. *Management Sci.* 27(7):822–833.

Corstjens M, Doyle P (1983) A dynamic model for strategically allocating retail space. *J. Oper. Res. Soc.* 34(10):943–951.

Cox KK (1970) The effect of shelf space upon sales of branded products. *J. Marketing Res.* 7(1):55–58.

Cummings KM, Sciandra R, Lawrence J (1991) Tobacco advertising in retail stores. *Public Health Rep.* 106(5):570–575.

Curhan RC (1972) The relationship between shelf space and unit sales in supermarkets. *J. Marketing Res.* 9(4):406–412.

Desmet P, Renaudin V (1998) Estimation of product category sales responsiveness to allocated shelf space. *Internat. J. Res. Marketing* 15:443–457.

Dorismond J (2019) Data-driven models for promoting impulse items in supermarkets. PhD dissertation, State University of New York at Buffalo.

Drèze X, Hoch SJ, Purk ME (1994) Shelf management and space elasticity. *J. Retailing* 70(4):301–326.

Düsterhöft T, Hübner A, Schaal K (2020) A practical approach to the shelf-space allocation and replenishment problem with heterogeneously sized shelves. *Eur. J. Oper. Res.* 282(1):252–266.

Eisend M (2014) Shelf space elasticity: A meta-analysis. *J. Retailing* 90(2):168–181.

Farley JU, Ring LW (1966) A stochastic model of supermarket traffic flow. *Oper. Res.* 14(4):555–567.

Ferrari SLP, Cribari-Neto F (2004) Beta regression for modelling rates and proportions. *J. Appl. Statist.* 31(7):799–815.

Flamand T, Ghoniem A, Maddah B (2016) Promoting impulse buying by allocating retail shelf space to grouped products. *J. Oper. Res. Soc.* 67(7):953–969.

Flamand T, Ghoniem A, Haouari M, Maddah B (2018) Integrated assortment planning and store-wide shelf space allocation: An optimization-based approach. *Omega* 81:134–149.

Geismar HN, Dawande M, Murthi B, Sriskandarajah C (2015) Maximizing revenue through two-dimensional shelf-space allocation. *Production Oper. Management* 24(7):1148–1163.

Ghoniem A, Maddah B (2015) Integrated retail decisions with multiple selling periods and customer segments: Optimization and insights. *Omega* 55:38–52.

Ghoniem A, Flamand T, Haouari M (2016a) Optimization-based very large-scale neighborhood search for generalized assignment problems with location/allocation considerations. *INFORMS J. Comput.* 28(3):575–588.

Ghoniem A, Flamand T, Haouari M (2016b) Exact solution methods for a generalized assignment problem with location/allocation considerations. *INFORMS J. Comput.* 28(3):589–602.

Ghoniem A, Maddah B, Ibrahim A (2016c) Optimizing assortment and pricing of multiple retail categories with cross-selling. *J. Global Optim.* 66(2):291–309.

Guthrie B, Parikh PJ (2020) The rack orientation and curvature problem for retailers. *IISE Trans.* 52(10):1081–1097.

Hansen JM, Raut S, Swami S (2010) Retail shelf allocation: A comparative analysis of heuristic and meta-heuristic approaches. *J. Retailing* 86(1):94–105.

Hansen P, Heinsbroek H (1979) Product selection and space allocation in supermarkets. *Eur. J. Oper. Res.* 3(6):474–484.

Hariga MA, Al-Ahmari A, Mohamed AA (2007) A joint optimisation model for inventory replenishment, product assortment, shelf space and display area allocation decisions. *Eur. J. Oper. Res.* 181(1):239–251.

Holmstrom J (1997) Product range management: A case study of supply chain operations in the European grocery industry. *Supply Chain Management* 2(3):107–115.

Hübner A, Schaal K (2017a) Effect of replenishment and backroom on retail shelf-space planning. *Bus. Res.* 10:123–156.

Hübner A, Schaal K (2017b) A shelf-space optimization model when demand is stochastic and space-elastic. *Omega* 68:139–154.

Hübner A, Düsterhöft T, Ostermeier M (2021a) Shelf space dimensioning and product allocation in retail stores. *Eur. J. Oper. Res.* 292(1):155–171.

Hübner A, Schäfer F, Schaal K (2021b) Maximizing profit via assortment and shelf-space optimization for two-dimensional shelves. *Production Oper. Management* 29(3):547–570.

Hui SK, Bradlow ET (2012) Bayesian multi-resolution spatial analysis with applications to marketing. *Quant. Marketing Econom.* 10(4):419–452.

Hui SK, Fader PS, Bradlow ET (2009) The traveling salesman goes shopping: The systematic deviations of grocery paths from TSP optimality. *Marketing Sci.* 28(3):566–572.

Hwang H, Choi B, Lee MJ (2005) A model for shelf space allocation and inventory control considering location and inventory level effects on demand. *Internat. J. Production Econom.* 97(2):185–195.

Hwang H, Choi B, Lee MJ (2009) A genetic algorithm approach to an integrated problem of shelf space design and item allocation. *Comput. Indust. Engrg.* 56(3):809–820.

Inman JJ, Winer RS, Ferraro R (2009) The interplay among category characteristics, customer characteristics, and customer activities on in-store decision making. *J. Marketing* 73(5):19–29.

Irion J, Lu J-C, Al-Khayyal F, Tsao YC (2011) A hierarchical decomposition approach to retail shelf space management and assortment decisions. *J. Oper. Res. Soc.* 62(10):1861–1870.

Irion J, Lu J-C, Al-Khayyal F, Tsao YC (2012) A piecewise linearization framework for retail shelf space management models. *Eur. J. Oper. Res.* 222(1):122–136.

Kacen JJ, Hess JD, Walker D (2012) Spontaneous selection: The influence of product and retailing factors on consumer impulse purchases. *J. Retailing Consumer Services* 19(6):578–588.

Ke W (2023) Merchandise placement optimization. Ghoniem A, Maddah B, eds. *Retail Space Analytics* (Springer, New York).

Kollat DT, Willett RP (1967) Customer impulse purchasing behaviour. *J. Marketing Res.* 4(1):21–31.

Krasny J (2012) How Walmart's one-stop shop concept gets us to impulse buy. Accessed February 20, 2013, https://www.businessinsider.com/walmarts-one-stop-shop-makes-us-spend-more-2012-5.

Larson JS, Bradlow ET, Fader PS (2005) An exploratory look at supermarket shopping paths. *Internat. J. Res. Marketing* 22(4):395–414.

Lin S, Kernighan BW (1973) An effective heuristic algorithm for the traveling-salesman problem. *Oper. Res.* 21(2):498–516.

Maddah B, Bish EK (2009) Locational tying of complementary retail items. *Naval Res. Logist.* 56(5):421–438.

Mou S, Robb DJ, DeHoratius N (2018) Retail store operations: Literature review and research directions. *Eur. J. Oper. Res.* 265(2):399–422.

Mowrey CH, Parikh PJ, Gue KR (2018) A model to optimize rack layout in a retail store. *Eur. J. Oper. Res.* 271(3):1100–1112.

Mowrey CH, Parikh PJ, Gue KR (2019) The impact of rack layout on visual experience in a retail store. *INFOR Inform. Systems Oper. Res.* 57(1):75–98.

van Nierop E, Fok D, Franses PH (2008) Interaction between shelf layout and marketing effectiveness and its impact on optimizing shelf arrangements. *Marketing Sci.* 27(6):1065–1082.

Ostermeier M, Düsterhöft T, Hübner A (2021) A model and solution approach for store-wide shelf space allocation. *Omega* 102: 102425.

Pak O, Ferguson M, Perdikaki O, Wu SM (2020) Optimizing stock keeping unit selection for promotional display space at grocery retailers. *J. Oper. Management* 66(5):501–533.

Phillips H, Bradshaw R (1993) How customers actually shop: Customer interaction with the point of sale. *J. Marketing Res. Soc.* 35(1):51–62.

Piron F (1991) Defining impulse purchasing. Holman RH, Solomon MR, eds. *Advances in Consumer Research*, vol. 18 (Association for Consumer Research, Provo, UT), 509–514.

POPAI (2014) POPAI Mass Merchant Shopper Engagement Study. Accessed June 28, 2016, https://cdn2.hubspot.net/hub/73834/file-1640923392-pdf/docs/popai_-_2014_mass_merchant_shopper_engagement_study.pdf.

Ross SM (2014) *Introduction to Probability Models*, 11th ed. (Academic Press, New York).

Underhill P (2000) *Why We Buy: The Science of Shopping* (Simon & Schuster, New York).

Yang MH (2001) An efficient algorithm to allocate shelf space. *Eur. J. Oper. Res.* 131(1):107–118.

Yang MH, Chen WC (1999) A study on shelf space allocation and management. *Internat. J. Production Econom.* 60–61(1):309–317.

Zufryden FS (1986) A dynamic programming approach for product selection and supermarket shelf-space allocation. *J. Oper. Res. Soc.* 37(4):413–422.

Tulay Flamand is an assistant professor in the Department of Economics and Business, Colorado School of Mines. Her research focuses on data analytics and optimization, applied to retail analytics, sports, mining, and energy. It also includes developing effective solution methodologies for solving challenging optimization problems. Her work has been published in journals such as *INFORMS Journal on Computing*, *IISE Transactions, Journal of the Operational Research Society* and *Operations Research*.

Ahmed Ghoniem is a professor of operations management at the Isenberg School of Management, University of Massachusetts Amherst. His work focuses on developing novel optimization models and effective solution methodologies for challenging problems that arise in retail analytics, transportation and logistics, and supply chain management in general. He is an associate editor for the *Journal of the Operational Research Society*. His research has been published in such journals as *European Journal of Operational Research*, *INFORMS Journal on Computing*, and *Operations Research*.

Bacel Maddah is a professor and the inaugural chairperson in the Industrial Engineering and Management Department at AUB, Lebanon. His research interests are in retailing, supply chain management, logistics, and stochastic processes. His research is published in leading journals such as *Operations Research*, *Management Science*, *IISE Transactions*, *European Journal of Operational Research*, and *Omega*. He is an associate editor for *International Transactions in Operational Research*.