

ERASMUS UNIVERSITY ROTTERDAM
Erasmus School of Economics

Factoring health into the equation: Promoting healthy purchasing
decisions through shelf-space optimization

Abstract

The optimal allocation of scarce shelf space is an everlasting challenge in retailing. Nowadays, this problem can be tackled by numerous optimization-based approaches to shelf-space planning. Many, if not all, of these methods attempt to promote impulse purchases, which are often unhealthy, hedonic products such as ice cream and cookies. In this paper, we extend one of these methods such that it promotes healthy purchasing decisions, as opposed to impulse purchases of unhealthy products.

For this purpose, a visibility penalty is introduced to the objective function of the model. This penalty penalizes the placement of unhealthy products on prominent shelf segments. The results indicate that by forgoing just 0.2% of profit, one might already attain a 4% gain in both store-wide availability of healthy products and store healthiness when we consider product visibility. Even larger gains can be attained by forgoing a larger share of profit: a 5% decrease in profit, yields gains in availability and visibility of healthy products of around 18% and 26%, respectively. These results indicate that store healthiness can be substantially improved by introducing a visibility penalty. Most notably, even stores that are not willing to forgo large shares of their profit could gain substantially in terms of store healthiness by introducing this penalty.

The visibility penalty is finally combined with a healthy-left, unhealthy-right ordering. The results suggest that this integrated approach is able to attain a visible product ordering with near-equal profit and store healthiness, when compared to a store that is optimized using the visibility penalty only.

Bachelor Thesis

BSc² in Econometrics and Economics

Name: Stefan van Berkum

Student ID: 467315

Supervisor: Olga Kuryatnikova, Ph.D.

Second assessor: Luuk van Maasackers, M.Sc.

Date final version: July 1, 2021

The views stated in this thesis are those of the author and not necessarily those of the supervisor, second assessor, Erasmus School of Economics, or Erasmus University Rotterdam.

Contents

1	Introduction	2
2	Related Work	4
3	Data	5
3.1	Replication	5
3.2	Extensions	8
4	Methodology	11
4.1	APSA	11
4.2	Optimization-based heuristic approach	13
4.3	Visibility penalty	14
4.4	Healthy-left, unhealthy-right	15
4.5	Integrated approach	15
4.6	Evaluation	16
5	Results	17
5.1	Replication	17
5.1.1	Heuristic performance	17
5.1.2	Neighborhood size	18
5.1.3	Optimality gap	19
5.1.4	Affinity relationships	19
5.2	Visibility penalty	20
5.3	Healthy-left, unhealthy-right	22
6	Conclusion	23
6.1	Replication	23
6.2	Extensions	24
6.3	Application	24
Appendix A	Code description	25
Appendix A.1	Source code	25
Appendix A.2	Output	27
Appendix A.3	Graphics	28

Factoring health into the equation: Promoting healthy purchasing decisions through shelf-space optimization

Stefan van Berkum^a

^a*Erasmus University Rotterdam, Burgemeester Oudlaan 50, 3062 PA, Rotterdam, The Netherlands*

Abstract

The optimal allocation of scarce shelf space is an everlasting challenge in retailing. Nowadays, this problem can be tackled by numerous optimization-based approaches to shelf-space planning. Many, if not all, of these methods attempt to promote impulse purchases, which are often unhealthy, hedonic products such as ice cream and cookies. In this paper, we extend one of these methods such that it promotes healthy purchasing decisions, as opposed to impulse purchases of unhealthy products.

For this purpose, a visibility penalty is introduced to the objective function of the model. This penalty penalizes the placement of unhealthy products on prominent shelf segments. The results indicate that by forgoing just 0.2% of profit, one might already attain a 4% gain in both store-wide availability of healthy products and store healthiness when we consider product visibility. Even larger gains can be attained by forgoing a larger share of profit: a 5% decrease in profit, yields gains in availability and visibility of healthy products of around 18% and 26%, respectively. These results indicate that store healthiness can be substantially improved by introducing a visibility penalty. Most notably, even stores that are not willing to forgo large shares of their profit could gain substantially in terms of store healthiness by introducing this penalty.

The visibility penalty is finally combined with a healthy-left, unhealthy-right ordering. The results suggest that this integrated approach is able to attain a visible product ordering with near-equal profit and store healthiness, when compared to a store that is optimized using the visibility penalty only.

Keywords: Shelf-space optimization, healthy choices, nudging, retail

1. Introduction

With an ever-expanding range of products, retailers have to decide which products to include in their limited shelf space, and where to put them. This problem of shelf-space planning has been addressed in academic literature since

Email address: stefanvanberkum@gmail.com (Stefan van Berkum)

as early as the 1960s, with an increasing focus on optimization-based approaches (Bianchi-Aguiar, Hübner, Carravilla & Oliveira, 2021).

Despite this increasing attention for optimization-based approaches, many decision support systems that are implemented in practice still use simplistic rules of thumb (Hübner & Kuhn, 2012). This points out the need for simple, yet effective approaches that can be easily implemented in decision support systems.

Most, if not all, optimization-based approaches aim to maximize profits and hence they often promote impulse purchases (e.g., Flamand, Ghoniem, Haouari & Maddah (2018)), which are often unhealthy hedonic products such as ice cream and cookies (Inman, Winer & Ferraro, 2009; Kacen, Hess & Walker, 2012). This stands in stark contrast to the recent shift in society towards healthier diet choices (Shan et al., 2019) and an apparent willingness of retailers to promote this healthy purchasing behavior (Martinez, Rodriguez, Mercurio, Bragg & Elbel, 2018).

There are many ways to promote healthy purchasing behavior, but in the context of shelf-space planning there are at least two promising approaches: enhancing visibility (i.e., prominent and convenient display of healthy products) and availability (i.e., a higher ratio of healthy to unhealthy products).

While there is no general consensus on the effectiveness of visibility nudges, the majority of studies have shown positive results (Shaw, Ntani, Baird & Vogel, 2020; Vecchio & Cavallo, 2019), most notably: placement of healthy products near the checkout counter (e.g., Kroese, Marchiori & De Ridder (2016); Van Gestel, Kroese & De Ridder (2018)), placing healthy products at eye-level (e.g., Adam, Jensen, Sommer & Hansen (2017); Foster et al. (2014); Sigurdsson, Saevarsson & Foxall (2009)), and even placing healthy products to the left of unhealthy products (e.g., Romero & Biswas (2016)).

Regarding availability enhancements, one study found positive effects when the assortment consisted of 75% percent healthy snacks as opposed to 25% (Van Kleef, Otten & Van Trijp, 2012). This indicates that increasing the sheer ratio of healthy to unhealthy products in the assortment could already prove effective in promoting healthy purchasing behavior.

This paper investigates several possible adaptations of a state-of-the-art optimization-based approach to shelf-space planning (Flamand et al., 2018). For this purpose, we first attempt to replicate the results of the original paper, to investigate whether it indeed performs as well as expected. The model is then adapted such that it promotes healthy purchasing behavior, as opposed to impulse purchases of unhealthy products. Possible approaches include incorporating relative attractiveness of different shelf heights, penalizing prominent placement of unhealthy products in the shelf-space optimization, and promoting healthy-left and unhealthy-right shelf-space planning. For the evaluation of each of the considered approaches, we use an adapted version of the data in Flamand et al. (2018) such that it incorporates a measure of relative nutritional value for each product category and relative attractiveness of different shelf heights.

The rest of this paper is structured as follows: Section 2 briefly outlines the academic literature that is related to this research. Section 3 discusses the data that is used in this paper and how it is obtained. Section 4 discusses the original

model, each of the considered adaptations, and how their performance will be evaluated. Section 5 discusses the obtained results and their interpretation. Finally, Section 6 briefly summarizes the results and discusses some suggestions for future work.

2. Related Work

Shelf-space planning can be split into two distinct parts: how much of each product to include in the assortment (assortment planning) and where to place them in the store (shelf-space allocation). Note that in other literature, sometimes these definitions are used in a different context (Bianchi-Aguiar et al., 2021), but for this paper we will stick to the aforementioned distinction in line with Flamand et al. (2018).

The first part of shelf-space planning is assortment planning. This is the process of determining which items to include in the assortment, how many stock keeping units (SKUs) to keep if a product is selected, and the inventory level of each SKU (Gürhan Kök, Fisher & Vaidyanathan, 2008). This requires careful consideration of the trade-off between different assortments, which is largely independent of the physical characteristics of the store (i.e., layout and individual shelf capacity). The other dimension of shelf-space planning is shelf-space allocation. This is the process of allocating products to shelves in the store, given a set of products and their corresponding SKUs. This process deals with the optimal allocation of scarce shelf space in a store, which is by definition based on its physical characteristics (as opposed to assortment planning). For a detailed overview of both components of shelf-space planning, interested readers can refer to Bianchi-Aguiar et al. (2021); Hübner & Kuhn (2012); Karampatsa, Grigoroudis & Matsatsinis (2017).

While much of the academic literature considers either one of these aspects of shelf-space planning separately (either assortment planning or shelf-space allocation), some papers consider both these aspects in tandem (e.g., Chen & Lin (2007); Flamand et al. (2018); Hübner & Schaal (2017)). In this paper, we will extend one of such models: a state-of-the-art integrated assortment planning and shelf-space allocation approach that is proposed in Flamand et al. (2018). This optimization-based heuristic approach yields near-optimal solutions within manageable computation times.

Visibility enhancements are mostly based on ideas that originate in behavioral sciences and economics. Placing healthy products in more prominent positions in the store are so called nudges (affecting the choice architecture in a non-limiting way (Thaler & Sunstein, 2009)), and they have shown positive results in promoting healthy choices (Vecchio & Cavallo, 2019). Another approach that one might consider is a healthy-left, unhealthy-right placement of products, which appears to positively influence healthy purchasing behavior (Romero & Biswas, 2016).

Availability enhancements are based on findings that increased shelf space for healthy products can already promote healthy purchasing behavior (e.g.,

Van Kleef et al. (2012)), which is supported by findings that increased shelf space can yield additional sales for a product (Eisend, 2014).

Approaches that combine visibility and availability enhancements appear to have the greatest potential for inducing healthy purchasing behavior (Shaw et al., 2020). Therefore, it makes sense to also investigate possible combinations of visibility enhancements (placing healthy products in more prominent positions in the store) and availability enhancements (increasing the amount of healthy products in the assortment).

To place healthy products in more prominent positions in the store, one needs to determine what these ‘prominent’ positions are exactly. One measure of shelf attractiveness used in this paper is traffic density. For the purpose of this paper, this parameter is simply simulated where shelves are divided over relative attractiveness categories, as in Flamand et al. (2018). For the application of this model in practice, one might consider modeling the traffic in a more sophisticated way such as in Tsai & Huang (2015).

Another measure that can influence shelf attractiveness is shelf height. Research into the relative attractiveness of different shelf heights suggests that there is a range roughly between eye and knee level that is the most attractive (Hübner, Düsterhöft & Ostermeier, 2021). This finding is confirmed by eye-tracking studies, where it appears that consumers focus on the center of a shelf (Drexler & Souček, 2016; Huddleston, Behe, Driesener & Minahan, 2018). Moreover, in general, it appears that the top shelf is slightly more attractive than the bottom shelf (Chandon, Hutchinson, Bradlow & Young, 2009; Drèze, Hoch & Purk, 1994).

3. Data

In this paper, we use a simulated dataset similar to the one used in Flamand et al. (2018). The authors simulate five store instances ranging from 30 shelves and 240 product categories to 100 shelves and 800 product categories. For the purpose of replication, we simulate the smallest, middle, and largest store instances (30/240, 50/400, and 100/800). For the purpose of evaluating the extensions, a single store with 30 shelves and 240 product categories is sufficient as the considered methods can simply be applied to larger stores as well.

3.1. Replication

Before discussing the data simulation, we define some notation that is used throughout the rest of this paper. The notation that is used in this paper consists of the notation used in the original model that is relevant to our replication and extension, extended with the health score and three-dimensional adaptation. Interested readers can refer to Flamand et al. (2018) for the original model with corresponding notation. For the purpose of this paper, the following notation is utilized:

- Sets and indices:

- $\mathcal{N} \equiv \{1, \dots, n\}$: The set of product categories, indexed by j .
- \mathcal{L} : The set of product category pairs $(j, j') \in \mathcal{N}^2$ that have allocation disaffinity (should not be placed on the same shelf).
- \mathcal{H}_1 : The set of product category pairs $(j, j') \in \mathcal{N}^2$ that have symmetric assortment affinity (should both be selected and placed on the same shelf, or neither should be selected).
- \mathcal{H}_2 : The set of product category pairs $(j, j') \in \mathcal{N}^2$ that have asymmetric assortment affinity (if j is selected, j' must be selected as well, and placed on the same shelf).
- \mathcal{H}_3 : The set of product category pairs $(j, j') \in \mathcal{N}^2$ that have allocation affinity (should be placed on the same shelf).
- $\mathcal{B} \equiv \{1, \dots, m\}$: The set of shelves, indexed by i .
- $\mathcal{K}_i \equiv \{1, 2, \dots, n_h, n_h + 1, \dots, n_v n_h\}$: The set of consecutive shelf segments along shelf i ($i \in \mathcal{B}$), indexed by k . Here, n_h and n_v denote the number of horizontal and vertical positions for each shelf, respectively.
- $\mathcal{K} \equiv \cup_{i \in \mathcal{B}} \mathcal{K}_i$: The set of all shelf segments in the store.
- $\mathcal{R} \equiv \{(k_1, k_2, j) \in \mathcal{K} \times \mathcal{K} \times \mathcal{N} : k_1 < k_2 \text{ and } \sum_{h=k_1+1}^{k_2-1} c_h > u_j - 2\varphi_j\}$: The set of triplets (k_1, k_2, j) , such that product j cannot cover all intermediate shelf segments between k_1 and k_2 .
- $\mathcal{A}_k \equiv \begin{cases} \emptyset, & \text{if } k \bmod n_h = 0 \\ \{k + 1, \dots, k + (n_h - k \bmod n_h)\}, & \text{otherwise} \end{cases}$: The set of shelf segments to the right of shelf segment k , where n_h denotes the number of horizontal positions for each shelf, $k \in \mathcal{K}$.

• Parameters:

- $f_k \in (0, 1]$: The attractiveness of segment k , $k \in \mathcal{K}$.
- Φ_j : The maximum attainable profit for product category j on a single segment, which would be obtained when $f_k = 1$, $j \in \mathcal{N}$. Note that in the original model, this measure consists of the profit margin, demand volume, and impulse purchase potential. For the purpose of this paper, however, we omit these elements and simply define Φ_j as before. In this paper, we directly simulate Φ_j , so its determination in practice can be done in any manner.
- ℓ_j/u_j : The minimum/maximum amount of shelf space allocated to product j (if it is selected in the assortment), $j \in \mathcal{N}$.
- φ_j : The minimum space that needs to be allocated to product j on a segment (if it is assigned to that segment), $j \in \mathcal{N}$.
- β_i : The largest index among the segments on shelf i , $i \in \mathcal{B}$.
- c_k : The capacity of segment k , $k \in \mathcal{K}$.

- h_j : The health score of each product category j , an integer on the interval $[1, 100]$, where 1 is the least healthy and 100 is the most healthy, $j \in \mathcal{N}$.
- Decision variables:
 - $x_{ij} \in \{0, 1\}$: $x_{ij} = 1$ if and only if product category j is allocated to shelf i , $\forall i \in \mathcal{B}, j \in \mathcal{N}$.
 - $y_{kj} \in \{0, 1\}$: $y_{kj} = 1$ if and only if product category j is allocated to shelf segment k , $\forall k \in \mathcal{K}, j \in \mathcal{N}$.
 - $s_{kj} \in \mathbb{R}^+$: The amount of space allocated to product category j on shelf segment k , $\forall k \in \mathcal{K}, j \in \mathcal{N}$.
 - $z_{jj'} \in \{0, 1\}$: $z_{jj'} = 1$ if and only if product category j and j' are both selected into the assortment, $\forall j, j' \in \mathcal{N}$.
 - $q_{kj} \in \{0, 1\}$: $q_{kj} = 1$ if and only if product category j is allocated to both segments k and $k + 1$, $\forall k \in \mathcal{K} \setminus \{\beta_i : i \in \mathcal{B}\}, j \in \mathcal{N}$.

For the purpose of replication of the original paper, we utilize following data generation scheme:

- All shelves have a capacity of 18 units, and each shelf is split into three horizontal positions with one vertical level (i.e., $n_h = 3$ and $n_v = 1$), of equal capacity (i.e., $c_k = 6$, $\forall k \in \mathcal{K}, j \in \mathcal{N}$).
- We simulate three store instances that consist of 30, 50, and 100 shelves and 240, 400, and 800 products.
- The minimum space ℓ_j for each product category j is randomly generated using a uniform distribution over the interval $[1, 3]$, rounded to the nearest integer.
- The maximum space u_j for each product category j is randomly generated using a uniform distribution over the interval $[\ell_j + 1, 6]$, rounded to the nearest integer. Note that this step is slightly different from the one used in [Flamand et al. \(2018\)](#). This adaptation avoids the case where $\ell_j = u_j$, which would effectively result in the exclusion of product j .
- The largest possible profit Φ_j for each product category j is randomly generated using a uniform distribution over the interval $[1, 25]$, rounded to two decimal places.
- The base attractiveness t of each shelf is randomly generated such that the shelves are evenly divided over the following five categories (i.e., 20% each): $t \in \{0.05, 0.25, 0.45, 0.65, 0.85\}$. The attractiveness f_k of each segment is generated by increasing the base attractiveness t of its shelf by a randomly generated amount. The center shelf segments are assumed to be less attractive than the right- and leftmost shelf segments. Consequently,

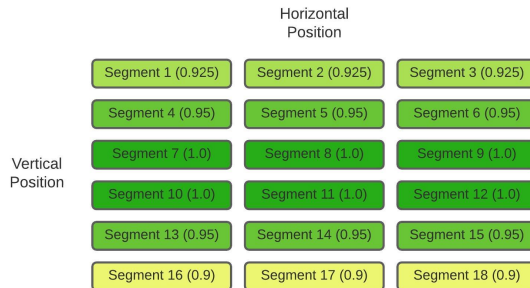


Figure 1: Example of a 6x3 shelf with vertical attractiveness score in parentheses.

the attractiveness of the center shelf segments is generated using a uniform distribution over the interval $[t, t + 0.05]$. The attractiveness of left- and rightmost shelf segments is generated using a uniform distribution over the interval $[t + 0.06, t + 0.1]$.

- The minimum allocated space is set to 0.1 units for each product category.

3.2. Extensions

As mentioned before, for the extensions we simulate a single store instance with 30 shelves and 240 product categories. For easy comparison, we use a similar data generation procedure to the one used in [Flamand et al. \(2018\)](#), but we adapt it such that it includes a health score between 1-100 for each product.

Moreover, shelf height needs to be incorporated into the model if we are to place healthy products at an attractive vertical position. This requires us to extend the model into a three-dimensional space. Conceptually, one would split each shelf into a given number of levels (e.g., six levels for each shelf) and weight the attractiveness score f_k of segment k based on its shelf height. For example, if each shelf has six levels, one might weight the attractiveness score of each segment by the following vertical attractiveness scores f_k^v : 0.925 (top shelf), 0.95 (fifth level), 1.0 (fourth level), 1.0 (third level), 0.95 (second level), 0.9 (bottom shelf). These particular scores are based on the relative attractiveness of shelves at the center and the relative attractiveness of top shelves as opposed to bottom shelves.

One possible way to extend the model to three-dimensional space is to represent a shelf as one consecutive set of segments. An example of such a shelf is depicted in Figure 1. In this example, the segments are ordered based on their vertical position and on their horizontal position within each level. In this way, the original computation based on the two-dimensional model can be used. In essence, this representation simply squeezes the long 18-segment shelves in [Flamand et al. \(2018\)](#) into a three-dimensional 6x3 shelf.

This three-dimensional representation of the original flat 18-segment shelf requires a small adjustment to one of the assumptions in [Flamand et al. \(2018\)](#). The authors assume that shelf segments in the middle of the store are less




Figure 2: Example a double aisle that consists of eight shelves, with three horizontal shelf segment positions each, viewed from above. Dashed lines depict possible other shelves.

attractive than end-of-aisle segment. In the context of an 18-segment aisle, this makes sense as end-of-aisle segments are assumed to have a higher traffic density. However, when we squeeze these shelves into a 6x3 shelf, the shelves do not span an entire aisle anymore. Therefore, this assumption that left- and rightmost segments are relatively attractive does not necessarily hold in this case. To the contrary, it can be argued that the center segments are relatively attractive as consumers tend to focus on the center of a shelf (Drexler & Souček, 2016; Huddleston et al., 2018). In line with this reasoning, we assume that center shelves are more attractive than the left- and rightmost segments, contrary to the assumption in Flamand et al. (2018).

The previously described assumptions still allow us to model the relative attractiveness of end-caps, albeit through a different mechanism. One might imagine that a store that consists of 30 shelves, could position them in aisles as depicted in Figure 2. This would, for example, allow the store to form three of these double aisles, and the remaining six shelves can be positioned on the outer perimeter of the store. Contrary to the original two-dimensional model, the end-caps are now defined as standalone 6x3 shelves, rather than the outer segments of a flat 18-segment shelf. The relative attractiveness of end-caps is then simply incorporated into the base attractiveness of the end-of-aisle shelf. Similarly to Flamand et al. (2018), in this paper the segment attractiveness is mainly based on the base attractiveness of its shelf, and only slightly affected by the horizontal and vertical position of a segment *within* a shelf.

For the purpose of incorporating the relative attractiveness of different shelf heights, we can split the relative attractiveness of each segment f_k into a horizontal component $0 < f_k^h \leq 1$ and a vertical component $0 < f_k^v \leq 1$, such that:

$$f_k = f_k^h f_k^v, \quad \forall i \in \mathcal{B}, k \in \mathcal{K}, \quad (1)$$

provided that:

$$f_a^h = f_b^h, \quad \forall i \in \mathcal{B}, a, b \in \mathcal{K}_i | (a < b \wedge (b - a) \bmod h = 0), \quad (2)$$

where h denotes the number of segments per level (horizontal positions).

Equation 2 ensures that the horizontal segment attractiveness component is equal for any two segments that are on the same horizontal position along the same shelf.

These adaptations result in the following data generation scheme:

- All shelves have a capacity of 18 units, and each shelf is split into three horizontal positions with six vertical levels (i.e., $n_h = 3$ and $n_v = 6$), of equal capacity (i.e., $c_k = 1, \forall k \in \mathcal{K}, j \in \mathcal{N}$).
- We simulate a single store that consists of 30 shelves and 240 products.
- The minimum space ℓ_j for each product category j is randomly generated using a uniform distribution over the interval $[1, 3]$, rounded to the nearest integer.
- The maximum space u_j for each product category j is randomly generated using a uniform distribution over the interval $[\ell_j + 1, 6]$, rounded to the nearest integer. Note that this step is slightly different from the one used in [Flamand et al. \(2018\)](#). This adaptation avoids the case where $\ell_j = u_j$, which would effectively result in the exclusion of product j .
- The largest possible profit Φ_j for each product category j is randomly generated using a uniform distribution over the interval $[1, 25]$, rounded to two decimal places.
- The health score h_j for each product category j is randomly generated using a uniform distribution over the interval $[1, 100]$, rounded to the nearest integer. Here, a score of 1 is the least healthy and a score of 100 is the most healthy.
- The base attractiveness t of each shelf is randomly generated such that the shelves are evenly divided over the following five categories (i.e., 20% each): $t \in \{0.05, 0.25, 0.45, 0.65, 0.85\}$. The horizontal attractiveness f_k^h of each segment is generated for each shelf column (i.e., all six levels) by increasing the base attractiveness t of its shelf by a randomly generated amount. The left- and rightmost shelf segments are assumed to be less attractive than the center shelf segments. Consequently, the attractiveness of the left- and rightmost shelf segments is generated using a uniform distribution over the interval $[t, t + 0.05]$. The attractiveness of center shelf segments is generated using a uniform distribution over the interval $[t + 0.06, t + 0.1]$. All shelf segments on the same horizontal position along a certain shelf are required to have equal horizontal attractiveness in accordance with Equation 2.
- The overall attractiveness f_k of each segment k is computed by multiplying its horizontal attractiveness score f_k^h by its corresponding vertical attractiveness score $f_k^v \in \{0.925, 0.95, 1.0, 1.0, 0.95, 0.9\}$ (as in Figure 1).
- The minimum allocated space is set to 0.1 units for each product category.

- All affinity relationships (\mathcal{L} , \mathcal{H}_1 , \mathcal{H}_2 , and \mathcal{H}_3), are simulated by randomly selecting 1% of the product categories and randomly assigning another product category to each of them for each relation.

4. Methodology

In [Flamand et al. \(2018\)](#), an integrated Assortment Planning and Shelf-space Allocation (APSA) model is proposed. This model is outlined in Subsection 4.1. Moreover, the authors of [Flamand et al. \(2018\)](#) propose an optimization-based heuristic approach to solve their model within manageable computation times. This approach is briefly discussed in Subsection 4.2. In this paper, we consider two possible extensions of this model: a visibility penalty for unhealthy products and healthy-left, unhealthy-right shelf positioning, discussed in Subsections 4.3 and 4.4, respectively. Subsection 4.5 outlines an integrated approach which combines these two extensions. The evaluation of the considered approaches is discussed in Subsection 4.6.

4.1. APSA

The MIP problem **APSA** is a model for shelf-space optimization that was introduced in [Flamand et al. \(2018\)](#). The model is defined as follows:

$$\text{APSA : Maximize } \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{N}} \Phi_j \frac{f_k s_{kj}}{c_k} \quad (3a)$$

$$\text{subject to } \sum_{i \in \mathcal{B}} x_{ij} \leq 1, \quad \forall j \in \mathcal{N} \quad (3b)$$

$$\sum_{j \in \mathcal{N}} s_{kj} \leq c_k, \quad \forall j \in \mathcal{N} \quad (3c)$$

$$\ell_j \sum_{i \in \mathcal{B}} x_{ij} \leq \sum_{k \in \mathcal{K}} s_{kj} \leq u_j \sum_{i \in \mathcal{B}} x_{ij}, \quad \forall j \in \mathcal{N} \quad (3d)$$

$$\varphi_j y_{kj} \leq s_{kj} \leq \min\{c_k, u_j\} y_{kj}, \quad \forall j \in \mathcal{N}, k \in \mathcal{K} \quad (3e)$$

$$s_{k_2, j} \geq c_{k_2} (y_{k_1, j} + y_{k_3, j} - 1), \quad \forall i \in \mathcal{B}, j \in \mathcal{N}, \\ k_1, k_2, k_3 \in \mathcal{K}_i | k_1 < k_2 < k_3, (k_1, k_3, j) \notin \mathcal{R} \quad (3f)$$

$$y_{kj} \leq x_{ij}, \quad \forall i \in \mathcal{B}, j \in \mathcal{N}, k \in \mathcal{K}_i \quad (3g)$$

$$x_{ij} \leq \sum_{k \in \mathcal{K}_i} y_{kj}, \quad \forall i \in \mathcal{B}, j \in \mathcal{N} \quad (3h)$$

$$q_{kj} \geq y_{kj} + y_{k+1, j} - 1, \quad \forall i \in \mathcal{B}, j \in \mathcal{N}, k \in \mathcal{K}_i \setminus \{\beta_i\} \quad (3i)$$

$$\sum_{j \in \mathcal{N}} q_{kj} \leq 1, \quad \forall i \in \mathcal{B}, k \in \mathcal{K}_i \setminus \{\beta_i\} \quad (3j)$$

$$x_{ij} + x_{ij'} \leq 1, \quad \forall (j, j') \in \mathcal{L}, i \in \mathcal{B} \quad (3k)$$

$$x_{ij} - x_{ij'} = 0, \quad \forall (j, j') \in \mathcal{H}_1, i \in \mathcal{B} \quad (3l)$$

$$x_{ij} \leq x_{ij'}, \quad \forall (j, j') \in \mathcal{H}_2, i \in \mathcal{B} \quad (3m)$$

$$x_{ij} - x_{ij'} \leq 1 - z_{jj'}, \quad \forall (j, j') \in \mathcal{H}_3, i \in \mathcal{B} \quad (3n)$$

$$x_{ij} - x_{ij'} \geq -1 + z_{jj'}, \quad \forall (j, j') \in \mathcal{H}_3, i \in \mathcal{B} \quad (3o)$$

$$z_{jj'} \leq \sum_{i \in \mathcal{B}} x_{ij}, \quad \forall (j, j') \in \mathcal{H}_3 \quad (3p)$$

$$z_{jj'} \leq \sum_{i \in \mathcal{B}} x_{ij'}, \quad \forall (j, j') \in \mathcal{H}_3 \quad (3q)$$

$$z_{jj'} \geq \sum_{i \in \mathcal{B}} x_{ij} + \sum_{i \in \mathcal{B}} x_{ij'} - 1, \quad \forall (j, j') \in \mathcal{H}_3 \quad (3r)$$

$$y_{k_1, j} + y_{k_2, j} \leq 1, \quad \forall (k_1, k_2, j) \in \mathcal{R} \quad (3s)$$

$$\mathbf{x}, \mathbf{y}, \mathbf{z}, \mathbf{q} \in \{0, 1\}, \mathbf{s} \geq 0. \quad (3t)$$

In the MIP problem **APSA**, the objective value **3a** maximizes a measure of store profit. Constraint **3b** ensures that each product is assigned to at most one shelf. Constraint **3c** ensures that the allocated space on any given shelf does not exceed its capacity. Constraint **3d** ensures that the allocated space for each product is between its minimum and maximum space requirement. Constraint **3e** ensures that the allocated space for any given product on any given shelf segment is between the minimum allocated space for this product, and the minimum of the capacity of this segment and the maximum space requirement of this product. Constraint **3f** is the second valid inequality in [Flamand et al. \(2018\)](#), it ensures that any product that is allocated to a pair of segments k_1 and k_3 ($k_1 < k_2 < k_3$) on any given shelf, is also allocated to any segment in between (k_2). This constraint is included as recommended by the authors. Constraint **3g** ensures that any product is only allocated to a shelf segment if it is assigned to the corresponding shelf. Constraint **3h** ensures that any product is placed on at least one shelf segment if it is allocated to the corresponding shelf. Constraint **3i** ensures that for any shelf, q_{kj} is equal to one if and only if product j is assigned to both segments k and $k + 1$, and zero otherwise. Constraint **3j** ensures that only one product category runs over any two consecutive shelf segments. Constraint **3k** ensures that any two products that have allocation disaffinity, are not placed on the same shelf. Constraint **3l** ensures that any two products that have symmetric assortment affinity, are selected together, and placed on the same shelf, or neither is selected. Constraint **3m** ensures that for

any two products that have asymmetric assortment affinity: if the first product is selected, the second product is selected as well, and placed on the same shelf. Constraints 3n-3r ensure that any two products that have allocation affinity are placed on the same shelf. Constraint 3s is the first valid inequality in Flamand et al. (2018), it ensures that any product that cannot cover all intermediate sections k_2 between segments k_1 and k_3 ($k_1 < k_2 < k_3$) on any given shelf, is not allocated to both segments k_1 and k_3 . This constraint is included, as recommended by the authors. Finally, Constraint 3t ensures that the decision variables x_{ij} , y_{kj} , $z_{jj'}$, and q_{kj} are binary, and that decision variable s_{kj} is a continuous variable larger or equal to zero.

4.2. Optimization-based heuristic approach

As briefly mentioned before, this paper builds upon the optimization-based heuristic approach that was proposed in Flamand et al. (2018). This method consists of two steps: an initialization procedure and a subsequent MIP-based re-optimization procedure.

In the first step of the method, the problem is initialized by optimizing each shelf separately. For each shelf, only the products that have not yet been selected are considered in its optimization. The shelves are traversed from most attractive to least attractive, such that the best products are placed on the most promising shelves.

In the second step of the method, the shelves are split into τ levels according to their current objective value contribution. From each level, a shelf is selected and these τ shelves are jointly re-optimized starting from the initial solution provided by the initialization procedure. After re-optimization, the shelves are taken out of consideration and the next set of τ shelves is re-optimized, until the number of remaining shelves falls below τ . This procedure repeated until any of the stopping conditions are met. In this paper, we utilize the following stopping conditions:

- The gap ϵ between the incumbent solution and the upper bound provided by the continuous relaxation of the problem falls below 0.5%.
- All shelves are traversed ten times with a change in the objective value that is smaller than 0.01. That is, no change for *every* set of τ shelves that is re-optimized, in each of the ten loops.
- The re-optimization procedure is run for more than five hours.

After each shelf optimization in the initialization procedure, the original algorithm described in Flamand et al. (2018) checks all products j that are placed on this shelf for allocation affinity relationships (\mathcal{H}_3) with products j' that are not selected on this shelf, and removes these products j' from consideration in subsequent shelf optimizations.

While the aforementioned affinity relationship check is indeed important, it is not sufficient to ensure solution feasibility with respect to the affinity relationships. Again, consider the case that product j is placed on a shelf that is

not considered. Then, one must also consider any set of products (j', j) that have asymmetric assortment affinity (\mathcal{H}_2). Consider for example cake icing and cake: if cake icing is selected in the assortment, then cake must be selected as well. The converse need not hold (i.e., cake can be selected without selecting cake icing). Therefore, the situation might arise that cake is selected on a shelf i , and cake icing is not. Then, in any subsequent shelf optimization cake icing must not be considered as cake has already been placed on shelf i . In other words, for any set $(j', j) \in \mathcal{H}_2$, one must remove j' from consideration when product j is placed on shelf that is not considered.

Moreover, when a product j' is removed from consideration in either of the aforementioned checks, this can affect other affinity relationships that concern product j' . Firstly, one must remove any product j'' from consideration when $(j', j'') \vee (j'', j') \in \mathcal{H}_1$. This is because both products must be selected or neither of them should be included in the assortment. As j' is removed from consideration, j'' must be removed from consideration as well. Secondly, one must again remove any product j'' from consideration when $(j'', j') \in \mathcal{H}_2$. As before, if cake is removed from consideration, cake icing must also be removed. These two additional check must then be recursively applied to product j'' if it is removed from consideration.

Finally, this procedure is not only applicable to the initialization procedure, but also to the subsequent re-optimization. This is because when we consider subsets of shelves, the same problems arise as when we optimize individual shelves (products can be placed on a shelf that is not considered). For this reason, each of the aforementioned checks must also be applied after the re-optimization of each set of τ shelves. In summary:

1. When a product j is placed on a shelf that is not considered.
 - \mathcal{L} : No action required (always satisfied).
 - \mathcal{H}_1 : No action required (j' would have been placed as well).
 - \mathcal{H}_2 : Remove j' for any pair $(j', j) \in \mathcal{H}_2$.
 - \mathcal{H}_3 : Remove j' for any pair $(j, j') \vee (j', j) \in \mathcal{H}_3$.
2. For any product j' that is removed from consideration.
 - \mathcal{L} : No action required (always satisfied).
 - \mathcal{H}_1 : Remove j'' for any pair $(j', j'') \vee (j'', j') \in \mathcal{H}_1$.
 - \mathcal{H}_2 : Remove j'' for any pair $(j'', j') \in \mathcal{H}_2$.
 - \mathcal{H}_3 : No action required (always satisfied).
3. Repeat Step 2 for any product j'' that is removed from consideration in the previous iteration, until no more products are removed.

4.3. Visibility penalty

One possible way to promote healthy purchase decisions is to place healthy products in more prominent positions in the store. Conversely, we can penalize

the prominent positioning of unhealthy products. We can do the latter by adapting the objective function of the **APSA** problem as follows:

$$\mathbf{APSA} : \text{Maximize } \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{N}} \Phi_j \frac{f_k s_{kj}}{c_k} - \frac{\gamma}{h_j} \text{func}\left(\frac{f_k s_{kj}}{c_k}\right), \quad (4)$$

where $\frac{f_k s_{kj}}{c_k}$ represents the visibility of product category j on shelf segment k ; γ is a penalty parameter for the visibility unhealthy food; $\text{func}(x)$ denotes any function of x .

For the purpose of this paper, we simply utilize a linear function of the visibility: $\text{func}(x) = x$. Note that this can be extended to any other function, if that were to be preferred.

Note that this visibility penalty also inherently penalizes the inclusion of unhealthy products in the assortment ($\frac{s_{kj}}{c_k}$), unweighted by the segment attractiveness score f_k . Because of this, the visibility penalty is particularly interesting if one wants to combine visibility and availability enhancements.

4.4. Healthy-left, unhealthy-right

A healthy-left, unhealthy-right shelf-space allocation can be implemented by penalizing unhealthy products being placed to the left of healthy products. If we include the previously mentioned extension to three-dimensional space, this can be done by defining a penalty function:

$$\psi(k_1) = \sum_{k_2 \in \mathcal{A}_{k_1}} \sum_{j_2 \in \mathcal{N}} (s_{k_2 j_2} h_{j_2}) - |\mathcal{A}_{k_1}| \sum_{j_1 \in \mathcal{N}} (s_{k_1 j_1} h_{j_1}). \quad (5)$$

We can use this in the objective function as follows:

$$\mathbf{APSA} : \text{Maximize } \sum_{k_1 \in \mathcal{K}} \left(\sum_{j \in \mathcal{N}} \left(\Phi_j \frac{f_{k_1} s_{k_1 j}}{c_{k_1}} \right) - \theta \psi(k_1) \right), \quad (6)$$

where θ is a parameter for the importance of healthy-left, unhealthy-right placement. This procedure rewards healthy products placed to the left of unhealthy products and penalizes the cases where unhealthy products are placed to the left of healthy products.

Note that this function is deliberately written out in different sums as this allows us to model the healthy-left, unhealthy-right approach without introducing quadratic terms to the objective function.

4.5. Integrated approach

These aforementioned approaches can be integrated to obtain a Health-adjusted Assortment Planning and Shelf-space Allocation problem **HAPSA**, with an objective function as follows:

$$\mathbf{HAPSA} : \text{Maximize } \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{N}} \Phi_j \frac{f_k s_{kj}}{c_k} - \frac{\gamma}{h_j} \frac{f_k s_{kj}}{c_k} - \theta \psi(k), \quad (7)$$

where $\psi(k)$ is the penalty function defined in Equation 5.

4.6. Evaluation

We can define three measures to evaluate the performance of each approach. The first measure is simply the Profit (P):

$$P = \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{N}} \Phi_j \frac{f_k s_{kj}}{c_k}. \quad (8)$$

A second measure, which measures the store-wide availability of healthy products, is the Store-average Health Score (SHS):

$$SHS = \frac{1}{\sum_{k \in \mathcal{K}} c_k} \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{N}} h_j s_{kj}. \quad (9)$$

The SHS measures the average health score for one capacity unit.

A third measure that we consider, which measures the store-wide visibility healthy products, is the Store-average Visibility-weighted Health Score (SVHS):

$$SVHS = \frac{1}{\sum_{k \in \mathcal{K}} c_k} \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{N}} h_j f_k s_{kj}. \quad (10)$$

The SVHS measures the average health score for one capacity unit, weighted by the segment attractiveness. Note that this measure also partially includes the availability of healthy products (as in the SHS).

These three measures can be compared against a benchmark model (without the extensions) for each approach, to examine the trade-off between profit (which is expected to decline), visibility of healthy products, and availability of healthy products.

For the purpose of this paper, we first focus on the visibility penalty approach, as it integrates both visibility and availability enhancements, and is therefore the most promising approach. We vary the parameter γ and by doing so we investigate the effects of various parameter settings. Then, with an adequate parameter setting for γ , the healthy-left, unhealthy-right approach is introduced and the effect of this approach is visually examined to see how large θ must be to achieve an acceptable ordering.

All algorithms are coded in Java, using CPLEX 20.1.0. The CPLEX parameters are tuned using the CPLEX tuning tool for the three-dimensional version of a store with 30 shelves and 240 products, for each of the three subtasks (initialization, continuous relaxation, and re-optimization). We set $\tau = 4$ shelves, and the re-optimization of each set of four shelves is run using a relative MIP gap tolerance of 0.001 and time limit of 120 seconds. All source code can be found on GitHub¹. A brief description of the files in this repository can be found in [Appendix A](#). Everything is run on a Lenovo Ideapad 510S-14IKB laptop with an Intel(R) Core(TM) i5-7200U CPU 2.50GHz processor and 8 GB RAM.

¹<https://github.com/stefanvanberkum/HAPSA>

5. Results

In this section, we will discuss the results of the replication and extensions that we consider in this paper. The replication results are outlined in 5.1. As to the extensions, we investigate the trade-off between profit and the previously defined scores that measure store healthiness (SHS and SVHS). Subsection 5.2 discusses the results of the visibility penalty. Subsection 5.3 discusses the results of the healthy-left, unhealthy-right approach.

5.1. Replication

For the purpose of checking whether the original results replicate, we run three instances of the smallest, middle, and largest stores that are considered in Flamand et al. (2018). We measure the total running time in seconds, as CPLEX has no built-in method to measure CPU time in Java. Moreover, we use the CPLEX parameters that are tuned for our three-dimensional adaptation of APSA, for a 30-shelf, 240-product store. We use these parameters as they significantly improve performance of CPLEX in our three-dimensional adaptation of APSA. This can be expected to improve performance of the two-dimensional model as well, since this is the same problem, except for the fact that the shelves are split into less segments. For each run in this section, the time limit stopping condition is set to one hour (wall-clock time).

Note that differences in performance are difficult to interpret, as we do not have access to the original implementation code. With respect to the stopping conditions, we are not aware of the number of loops required for termination in the original runs. Moreover, we have no information on the utilized parameter settings with respect to the relative MIP gap, time limit, and other tunable parameters.

In Subsection 5.1.1 we discuss the general performance of the heuristic approach, as opposed to the full model. In Subsection 5.1.2 we discuss the effect of different values of the neighborhood size τ (the number of shelves that is selected for re-optimization in each iteration of the re-optimization procedure). In Subsection 5.1.3 we discuss the effect of different values of the target optimality gap ϵ . In Subsection 5.1.4 we discuss the effect of adding the different affinity relationships (\mathcal{L} , \mathcal{H}_1 , \mathcal{H}_2 , and \mathcal{H}_3).

5.1.1. Heuristic performance

In this subsection, the optimization-based heuristic approach is compared against running the full model APSA as a whole. The result of this comparison is depicted in Table 1. The results are consistent with the findings in Flamand et al. (2018), in the sense that the heuristic substantially outperforms the regular APSA model. We observe that the optimality gap for regular APSA is larger, when compared to the original findings. This can be due to many reasons. Possibly, it is due to the fact that we measure wall-clock time in seconds, as opposed to CPU time. Another possible explanation is that the original runs were done using a substantially more powerful computer, using a different programming language (C++). We can also note that the heuristic is faster in our

Table 1: The performance of the optimization-based heuristic approach, compared to the full model APSA.

$(\mathcal{B} , \mathcal{N})$	Inst.	APSA		Heuristic $\tau = 4$ $\epsilon = 0.5\%$	
		T(s)	G(%)	T(s)	G(%)
(30, 240)	1	3603	2.03	63	0.50
	2	3601	2.20	43	0.49
	3	3601	1.80	32	0.45
(50, 400)	1	3605	3.05	173	0.45
	2	3602	1.85	109	0.50
	3	3602	2.75	92	0.48
(100, 800)	1	3609	6.88	623	0.48
	2	3618	7.72	617	0.49
	3	3611	7.51	465	0.46

Note. T(s): Time in seconds, G(%): Optimality gap.

Table 2: The effect of different neighborhood sizes τ on the performance of the optimization-based heuristic.

$(\mathcal{B} , \mathcal{N})$	Inst.	APSA		Heuristic					
		T(s)	G(%)	$\tau = 2$		$\tau = 3$		$\tau = 4$	
				$\epsilon = 0.5\%$		$\epsilon = 0.5\%$		$\epsilon = 0.5\%$	
T(s)	G(%)	T(s)	G(%)	T(s)	G(%)	T(s)	G(%)		
(30, 240)	1	3603	2.03	101	0.89	68	0.50	63	0.50
	2	3601	2.20	51	0.89	44	0.47	43	0.49
	3	3601	1.80	85	0.52	29	0.47	32	0.45
(50, 400)	1	3605	3.05	215	0.69	86	0.48	173	0.45
	2	3602	1.85	180	0.77	103	0.49	109	0.50
	3	3602	2.75	318	0.50	59	0.49	92	0.48
(100, 800)	1	3609	6.88	908	0.59	330	0.48	623	0.48
	2	3618	7.72	1504	0.63	494	0.46	617	0.49
	3	3611	7.51	1686	0.61	375	0.48	465	0.46

Note. T(s): Time in seconds, G(%): Optimality gap.

runs, especially for larger store instances. Possibly, the increased performance can be explained by our relative MIP gap tolerance and time limit for each re-optimization iteration. The improved scalability of the heuristic can be due to the parameter tuning, which is done for a three-dimensional version of a 30-shelf, 240-product store. As this problem is quite large, it can be expected that the heuristic with these parameters performs relatively well for large instances.

5.1.2. Neighborhood size

In this subsection, we investigate the effect of different neighborhood sizes τ on the performance of the optimization-based heuristic approach. The result of this comparison is depicted in Table 2. The results are not entirely consistent with the results obtained in the original paper, in the sense that that $\tau = 3$ yields results superior to $\tau = 4$, especially for large instances. This could be due to our CPLEX parameter settings (relative MIP gap tolerance, time limit, and other tuned parameters). It might be the case that these yield even better results when $\tau = 3$.

Table 3: The effect of different values for the target optimality gap ϵ on the performance of the optimization-based heuristic.

$(\mathcal{B} , \mathcal{N})$	Inst.	APSA		Heuristic					
				$\tau = 4$		$\tau = 4$		$\tau = 4$	
		T(s)	G(%)	$\epsilon = 1.5\%$		$\epsilon = 1.0\%$		$\epsilon = 0.5\%$	
		T(s)	G(%)	T(s)	G(%)	T(s)	G(%)	T(s)	G(%)
(30, 240)	1	3603	2.03	37	1.07	36	0.96	63	0.50
	2	3601	2.20	32	0.97	25	0.87	43	0.49
	3	3601	1.80	17	0.92	20	0.77	32	0.45
(50, 400)	1	3605	3.05	46	0.98	67	0.92	173	0.45
	2	3602	1.85	70	0.89	99	0.97	109	0.50
	3	3602	2.75	63	0.88	55	0.90	92	0.48
(100, 800)	1	3609	6.88	355	0.81	460	0.85	623	0.48
	2	3618	7.72	354	1.12	266	0.99	617	0.49
	3	3611	7.51	276	0.91	275	0.93	465	0.46

Note. T(s): Time in seconds, G(%): Optimality gap.

Table 4: The effect of adding affinity relationships \mathcal{L} and \mathcal{H}_1 on the performance of the optimization-based heuristic.

$(\mathcal{B} , \mathcal{N})$	Inst.	No affinity		\mathcal{L}		\mathcal{H}_1	
		T(s)	G(%)	T(s)	G(%)	T(s)	G(%)
(30, 240)	1	63	0.50	74	0.49	64	0.48
	2	43	0.49	38	0.49	169	0.66
	3	32	0.45	32	0.49	56	0.49
(50, 400)	1	173	0.45	131	0.45	360	0.49
	2	109	0.50	110	0.47	229	0.50
	3	92	0.48	106	0.43	106	0.48
(100, 800)	1	623	0.48	313	0.46	917	0.48
	2	617	0.49	540	0.46	688	0.50
	3	465	0.46	662	0.49	545	0.46

Note. T(s): Time in seconds, G(%): Optimality gap.

5.1.3. Optimality gap

In this subsection, we investigate the effect of different values for the target optimality gap ϵ . The result of this comparison is depicted in Table 3. The results are roughly consistent with the findings of the original paper. As expected, increasing the target optimality gap decreases the running times but yields larger optimality gaps.

5.1.4. Affinity relationships

In this subsection, we investigate the effect of adding affinity relationships (\mathcal{L} , \mathcal{H}_1 , \mathcal{H}_2 , and \mathcal{H}_3) to the model. The effect of adding five of each affinity relationship separately is depicted in Table 4 (\mathcal{L} and \mathcal{H}_1) and Table 5 (\mathcal{H}_2 and \mathcal{H}_3). The effect of adding five of each affinity relationship together is depicted in Table 6. These results are roughly consistent with the findings of the original paper. The separate affinity relationships \mathcal{L} , \mathcal{H}_2 , and \mathcal{H}_3 do not appear to have a substantial impact on the performance of the heuristic. The addition of the affinity relationship \mathcal{H}_1 does have a substantial impact on the performance of the heuristic. Running time and convergence are negatively impacted by this relationship, and even more so by adding all relationships together (\mathcal{L} , \mathcal{H}_1 , \mathcal{H}_2 , and \mathcal{H}_3).

Table 5: The effect of adding affinity relationships \mathcal{H}_2 and \mathcal{H}_3 on the performance of the optimization-based heuristic.

$(\mathcal{B} , \mathcal{N})$	Inst.	No affinity		\mathcal{H}_2		\mathcal{H}_3	
		T(s)	G(%)	T(s)	G(%)	T(s)	G(%)
(30, 240)	1	63	0.50	106	0.45	79	0.50
	2	43	0.49	42	0.50	43	0.48
	3	32	0.45	39	0.47	34	0.46
(50, 400)	1	173	0.45	110	0.50	113	0.49
	2	109	0.50	107	0.47	140	0.47
	3	92	0.48	103	0.49	77	0.48
(100, 800)	1	623	0.48	478	0.48	392	0.50
	2	617	0.49	650	0.49	625	0.48
	3	465	0.46	611	0.44	608	0.46

Note. T(s): Time in seconds, G(%): Optimality gap.

Table 6: The effect of adding all affinity relationships on the performance of the optimization-based heuristic.

$(\mathcal{B} , \mathcal{N})$	Inst.	No affinity		$(\mathcal{L}, \mathcal{H}_1, \mathcal{H}_2, \mathcal{H}_3)$	
		T(s)	G(%)	T(s)	G(%)
(30, 240)	1	63	0.50	282	0.65
	2	43	0.49	151	0.96
	3	32	0.45	53	0.49
(50, 400)	1	173	0.45	414	0.51
	2	109	0.50	164	0.46
	3	92	0.48	137	0.49
(100, 800)	1	623	0.48	798	0.49
	2	617	0.49	812	0.48
	3	465	0.46	615	0.50

Note. T(s): Time in seconds, G(%): Optimality gap.

5.2. Visibility penalty

As to the extension, we run the APSA model with a visibility penalty 40 times, with the visibility penalty parameter γ between 5-200 (i.e., increments of five). The percentage change in each of the three measures (profit, SHS, and SVHS) for different values of γ is displayed in Figure 3. The plot includes curves fitted by means of LOESS, with the corresponding 95% confidence intervals. Here, we can clearly discern a pattern for each of the three measures. The profit initially falls only by a minor amount, and appears to fall more steeply as γ increases. As to the other scores, both SHS and SVHS increase very steeply at first, and appear to flatten out as γ increases. The results imply that by forgoing just 0.2% of profit, one might attain an approximate gain of 4% in both the store-wide availability of healthy products (SHS) and store healthiness when we consider product visibility (SVHS). Moreover, by forgoing approximately 1% of profits, one might attain an approximate gain of 7% and 10% in SHS and SVHS, respectively. Stores that are willing to surrender approximately 5% of profits might attain gains around 18% and 26% in SHS and SVHS, respectively.

Another way we might analyze these results is by plotting the change in SHS and SVHS relative to the absolute change in profit. This is displayed for SHS and SVHS in Figures 4a and 4b, respectively. These plots also include curves fitted by means of LOESS, with the corresponding 95% confidence intervals. Again, we can clearly see that the initial gain in each of these scores is relatively high compared to the loss in profit. These ratios initially fall sharply for both the SHS




Figure 3: Percentage change in profit, SHS, and SVHS using different values of γ .




Figure 4: Ratio between change SHS/SVHS and change in profit using different values of γ .

and SVHS. As γ increases, the ratios for SHS and SVHS appear to respectively converge to approximately 3.5% and 5.1% percent for each percentage point in forgone profits.

We can also consider the effect of the visibility penalty parameter γ on the average shelf. Figure 5 displays a front view of the average shelf for different values of γ . For each segment, this front view displays the average health score for that particular segment. What is clear from this figure is that by increasing γ , the average shelf gets healthier. However, we cannot discern a clear pattern for the placement of healthy products. One might expect these healthy products to get pulled towards the center of the shelf, as these segments are both horizontally and vertically more attractive than the others. The fact that this does not seem to happen might be due to the relatively large effect of the base attractiveness of each shelf, compared to the horizontal and vertical difference in segment attractiveness within a shelf. Moreover, the model assumes that if a product is placed on multiple shelf segments, these segments must be adjacent (horizontally). For this reason, the average health scores of horizontally adjacent shelf segments are not independent. Another factor worth considering

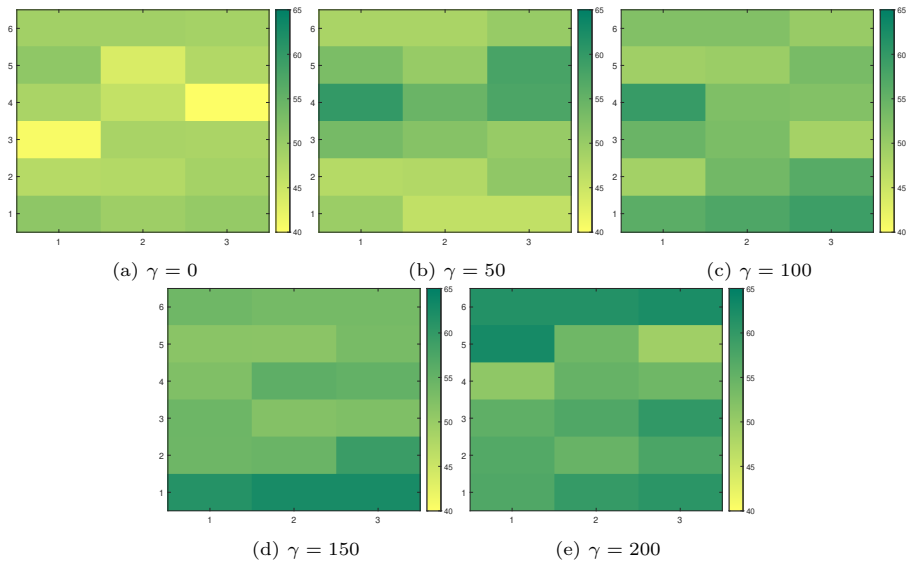


Figure 5: Segment-average health score for different values of the visibility penalty parameter γ , as viewed from the front of a shelf.

is that the optimization algorithm still needs to balance profits and health scores. While these aspects are simulated independently, it might coincidentally be the case that some unhealthy products are very profitable. The combination of these two elements might cause the seemingly random patterns we see for increasing values of γ . If the aim is to achieve a more visible product ordering, one might consider adjusting the effect of horizontal and vertical positions on segment attractiveness.

5.3. Healthy-left, unhealthy-right

For the purpose of this paper, we focus on the integrated approach HAPSA (visibility penalty combined with healthy-left, unhealthy-right ordering) as the aim is to make the store healthier. If one were to consider the healthy-left, unhealthy-right ordering separately, then this would promote the placement of very unhealthy products at the rightmost side of the shelf. The visibility penalty parameter γ parameter is frozen at $\gamma = 100$. Note, however, that this can be any value and is simply for the purpose of demonstration. We proceed by increasing the healthy-left, unhealthy-right ordering parameter θ until a visual ordering appears for the average shelf. This gives an indication of how much profit must be forgone to obtain an acceptable ordering.

The aforementioned procedure yields an acceptable ordering for $\theta = 0.001$, which is visualized in Figure 6. Here we clearly see that healthy products are on average placed on the left side of the shelf, while unhealthy products are on average placed on the right side of the shelf. The effect of this ordering on profits, SHS, SVHS is outlined in Table 7. These results indicate that the profit

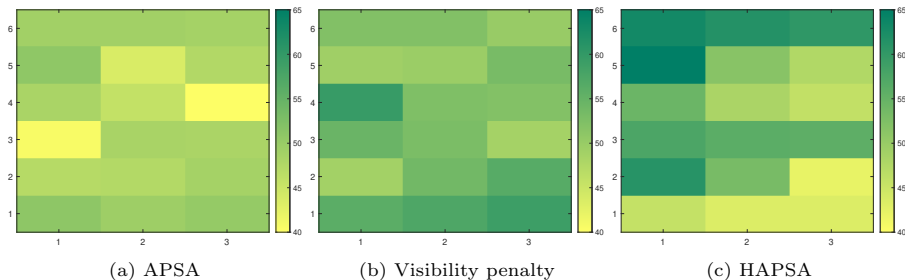


Figure 6: Segment-average health score for regular APSA, APSA with a visibility penalty ($\gamma = 100$), and HAPSA ($\gamma = 100$, $\theta = 0.001$), as viewed from the front of a shelf.

that is attained for HAPSA is practically equal to the profit that is obtained by APSA with a visibility penalty. Moreover, the SHS and SVHS scores are also nearly equal to those for APSA with a visibility penalty.

Table 7: Effects of APSA, APSA with a visibility penalty ($\gamma = 100$), and HAPSA ($\gamma = 100$, $\theta = 0.001$) on profit, SHS, and SVHS.

	APSA	Visibility penalty	HAPSA
Profit	5305.18 (-)	5156.90 (-2.8%)	5160.05 (-2.7%)
SHS	47.6 (-)	53.3 (+12.0%)	53.2 (+11.8%)
SVHS	21.7 (-)	25.6 (+18.1%)	25.5 (+17.6%)

Note. Change compared to regular APSA is in parentheses.

6. Conclusion

In this section, we will briefly summarize all results and provide some final remarks. Subsection 6.1 briefly outlines the results obtained from the replication of the original research. Subsection 6.2 discusses the results of the extensions to the three-dimensional adaptation of the model. Subsection 6.3 provides some final remarks on further application of the investigated methods.

6.1. Replication

In short, we were able to replicate most of the findings of the original paper. Their proposed optimization-based heuristic approach was found to substantially outperform the regular model. A finding worth noting is that our implementation of the heuristic appeared to be even faster than in the original paper. Moreover, the scalability of the heuristic was substantially improved. This may be due to the tuning of CPLEX parameters, which were tuned for a large-scale model. Moreover, the relative MIP gap tolerance and the time limit for each re-optimization iteration might have affected the performance. These parameters might have also caused another finding of interest: the neighborhood size $\tau = 3$ appeared to perform better than $\tau = 4$. This could be further investigated in future research.

6.2. Extensions

In this paper, we investigate possible approaches to promote healthy purchasing behavior through shelf-space optimization. For this purpose, we adapt and extend a state-of-the-art optimization-based approach to shelf space planning. The method is adapted to model a three-dimensional store. Moreover, the method is complemented by a visibility penalty and a healthy-left, unhealthy-right ordering to promote healthy purchasing behavior.

The proposed visibility penalty substantially increases both the store-wide availability of healthy products (SHS) and their visibility (SVHS). The results suggest that by forgoing 1% of profits, one might attain gains in SHS and SVHS of approximately 7% and 10%, respectively. For stores that are willing to forgo 5% of their profits, these gains in SHS and SVHS rise to around 18% and 26%, respectively. It is worth noting that the gains in SHS and SVHS are relatively large compared to the forgone profits for small values of the visibility penalty parameter γ . This implies that stores that are willing to forgo minimal amounts of profit might still benefit from this visibility enhancement. For example, a store that is willing to surrender just 0.2% of profits may gain approximately 4% in both SHS and SVHS. This loss in profit might easily be covered by marketing gains that are based around a healthier store image.

Secondly, the visibility penalty combined with a healthy-left, unhealthy-right ordering is able to attain a visible product ordering with near-equal profit and store healthiness, when compared to a store that is optimized using the visibility penalty only. Moreover, this ordering can be used to sort on any numerical product aspect (e.g., environmental impact).

Another finding worth noting is that the running time significantly increases due to the three-dimensional adaptation. This makes sense, as we use six times as many segments as in the original model. Moreover, we observe longer running times and worse convergence when the visibility penalty parameter γ and healthy-left, unhealthy-right ordering parameter θ are increased. Future research might pinpoint why this is the case, and possibly mitigate this behavior. Perhaps, the performance of the algorithm might be improved by randomly removing products from shelves (much like a dropout technique used in the context of deep neural networks). This is also left as a question for future research.

In this paper, we incorporate a visibility penalty and healthy-left, unhealthy-right ordering into a single state-of-the-art optimization-based approach to shelf-space planning. Note, however, that both of the proposed methods can be incorporated into almost any optimization-based approach as long as it includes some kind of measure of segment attractiveness.

6.3. Application

The application of the investigated methods in a real-world setting would require some additional research. Firstly, using real-world products would require an accurate health score for each considered product, which could be determined for each product by using new or existing health scores such as the Healthy Eating Index (Krebs-Smith et al., 2018).

Secondly, the horizontal and vertical attractiveness scores used in this paper are determined solely for the sake of demonstration. In a real-world setting, one might consider extensive studies that combine traffic measurement such as Tsai & Huang (2015) with eye-tracking approaches such as Drexler & Souček (2016); Huddleston et al. (2018). Preferably, one would tailor the attractiveness for each segment to the specific store that is considered in the optimization. For example, end-cap shelves could be given a relatively high base attractiveness, as well as shelves near the checkout.

As to the extensions, in practice one might optimize a given store both without the extensions and with a visibility penalty, based on the store owner’s preference for store healthiness. If a health score between 1-100 is used, one would expect the effects to be roughly similar to the results found in this paper. By also optimizing the store without the extensions, one can confirm whether the results are satisfactory. If not, the γ parameter can be adjusted accordingly. Moreover, any product ordering might be attained, as long as it is based on a numerical product characteristic (e.g., price). However, the required value of θ might differ between stores. Therefore, one might consider a similar procedure as used in this paper. By increasing the parameter value until an acceptable ordering appears, one ensures that the products are ordered at minimal cost to any other measures.

For HAPSA with both a visibility penalty and product ordering, note that these penalties need not serve the same purpose, and that they might conflict with each other. For example, the healthy-left, unhealthy-right ordering on its own would try to place very unhealthy products at the rightmost side of the shelf. In practice, this causes very stringent orderings to actually decrease store healthiness. This must be taken into account when the parameter settings are determined, as per what is deemed the most important by the store owner.

Appendix A. Code description

As mentioned before, all source code used in this paper can be found on GitHub². This appendix briefly summarizes the files that can be found within this repository. Appendix A.1 discusses the main Java source code, required to run all methods. Appendix A.2 outlines the automatically generated output files. Appendix A.3 discusses the R and MATLAB code used to obtain the graphs and colormaps, respectively.

Appendix A.1. Source code

Path: *main* → *HAPSA* → *src* → *hapsa* → ...

The following different classes are utilized:

- *AsymmetricPair.java*: The *AsymmetricPair* class defines a basic framework for pairs of products that have an asymmetric affinity relationship (\mathcal{H}_2).

²<https://github.com/stefanvanberkum/HAPSA>

- *HAPSA.java*: The *HAPSA* class provides a framework for the models used in the MIP-based re-optimization procedure of the optimization-based heuristic approach to assortment planning and shelf space optimization. It extends the *Model* class.
- *Main.java*: The *Main* class provides the main execution environment, and contains methods for the optimization-based heuristic approach.
- *Model.java*: The *Model* class is an abstract class that defines the basic framework for a model used in the optimization-based heuristic approach to assortment planning and shelf space optimization (*SSP* and *HAPSA*). It extends the CPLEX *IloCplex* class.
- *ParameterTuner.java*: The *ParameterTuner* class provides methods to tune the CPLEX parameters for the optimization-based heuristic approach to shelf space planning. The optimal parameters are determined for the first iteration of the initialization procedure, the continuous relaxation, and the first iteration of the re-optimization procedure. The results are saved to a file and automatically loaded into the *Main* class.
- *Product.java*: The *Product* class provides definitions and methods for a product within the context of a store.
- *Segment.java*: The *Segment* class provides definitions and methods for a shelf segment.
- *Shelf.java*: The *Shelf* class provides definitions and methods for a store shelf.
- *Solution.java*: The *Solution* class provides a framework and methods for any store planning solution.
- *SSP.java*: The *SSP* class provides a framework for the Single Shelf Problem (SSP) model used in the initialization procedure of the optimization-based heuristic approach to assortment planning and shelf space optimization. It extends the *Model* class.
- *Store.java*: The *Store* class provides definitions and methods for a store.
- *StoreSimulator.java*: The *StoreSimulator* class provides methods for the simulation of a store.
- *SymmetricPair.java*: The *SymmetricPair* class defines a basic framework for pairs of products that have a symmetric affinity relationship (\mathcal{L} , \mathcal{H}_1 , and \mathcal{H}_3).
- *Utils.java*: The *Utils* class provides some general utility methods for the HAPSA method.

Appendix A.2. Output

Path: *main* → *HAPSA* → ...

We can distinguish two general types of output, parameter files and results:

- *Parameters* → ...

The parameter files are generated by the *ParameterTuner* class, and written to a separate directory for each subproblem:

- *CONT*: This directory contains CPLEX parameter files for the continuous relaxation.
- *HAPSA*: This directory contains CPLEX parameter files for the MIP-based re-optimization procedure.
- *SSP*: This directory contains CPLEX parameter files for the initialization procedure.

- *Results* → ...

The result files are generated by the *Main* class, and written to a separate directory for each problem type:

- *APSA*: The regular, three-dimensional APSA.
- *APSA_2D*: The regular, two-dimensional APSA.
- *AVA*: Three-dimensional APSA with an availability penalty (not used in this paper).
- *HAPSA*: The integrated approach HAPSA.
- *HLUR*: Three-dimensional APSA with healthy-left, unhealthy-right ordering (not used in this paper).
- *VIS*: Three-dimensional APSA with a visibility penalty.

Each of these directories contains three subdirectories:

- *CSV*: This directory contains CSV files of the scores (profit, SHS, and SVHS), and front-view shelf segment-average health scores. These files are solely written for the purpose of creating graphs and colormaps, as discussed in [Appendix A.3](#). These files are omitted from the GitHub repository.
- *Summary*: This directory contains summary files for each run. These files contain the profit, SHS, SVHS, front-view shelf segment-average health scores, and some other statistics.
- *Variables*: This directory contains CSV files with the final values for the decision variables of interest (s_{kj} , x_{ij} , and y_{kj}). These files are omitted from the GitHub repository.

Appendix A.3. Graphics

Path: *main* → ...

Two scripts are utilized for the purpose of graphical output representation:

- *plotResults.R*: Creates all sorts of different plots from the user-specified results. Implemented in R (run in RStudio).
- *plotShelf.m*: Creates colormaps from the user-specified results. Implemented in MATLAB R2020a.

References

- Adam, A., Jensen, J. D., Sommer, I., & Hansen, G. L. (2017). Does shelf space management intervention have an effect on calorie turnover at supermarkets? *Journal of Retailing and Consumer Services*, *34*, 311–318. doi:[10.1016/j.jretconser.2016.07.007](https://doi.org/10.1016/j.jretconser.2016.07.007).
- Bianchi-Aguiar, T., Hübner, A., Carravilla, M. A., & Oliveira, J. F. (2021). Retail shelf space planning problems: A comprehensive review and classification framework. *European Journal of Operational Research*, *289*, 1–16. doi:[10.1016/j.ejor.2020.06.018](https://doi.org/10.1016/j.ejor.2020.06.018).
- Chandon, P., Hutchinson, J. W., Bradlow, E. T., & Young, S. H. (2009). Does in-store marketing work? Effects of the number and position of shelf facings on brand attention and evaluation at the point of purchase. *Journal of Marketing*, *73*, 1–17. doi:[10.1509/jmkg.73.6.1](https://doi.org/10.1509/jmkg.73.6.1).
- Chen, M. C., & Lin, C. P. (2007). A data mining approach to product assortment and shelf space allocation. *Expert Systems with Applications*, *32*, 976–986. doi:[10.1016/j.eswa.2006.02.001](https://doi.org/10.1016/j.eswa.2006.02.001).
- Drexler, D., & Souček, M. (2016). The influence of sweet positioning on shelves on consumer perception. *Food Packaging and Shelf Life*, *10*, 34–45. doi:[10.1016/j.fpsl.2016.09.001](https://doi.org/10.1016/j.fpsl.2016.09.001).
- Drèze, X., Hoch, S. J., & Purk, M. E. (1994). Shelf management and space elasticity. *Journal of Retailing*, *70*, 301–326. doi:[10.1016/0022-4359\(94\)90002-7](https://doi.org/10.1016/0022-4359(94)90002-7).
- Eisend, M. (2014). Shelf space elasticity: A meta-analysis. *Journal of Retailing*, *90*, 168–181. doi:[10.1016/j.jretai.2013.03.003](https://doi.org/10.1016/j.jretai.2013.03.003).
- 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. doi:[10.1016/j.omega.2017.10.006](https://doi.org/10.1016/j.omega.2017.10.006).

- Foster, G. D., Karpyn, A., Wojtanowski, A. C., Davis, E., Weiss, S., Brensinger, C., Tierney, A., Guo, W., Brown, J., Spross, C., Leuchten, D., Burns, P. J., & Glanz, K. (2014). Placement and promotion strategies to increase sales of healthier products in supermarkets in low-income, ethnically diverse neighborhoods: A randomized controlled trial. *The American Journal of Clinical Nutrition*, *99*, 1359–1368. doi:[10.3945/ajcn.113.075572](https://doi.org/10.3945/ajcn.113.075572).
- Gürhan Kök, A., Fisher, M. L., & Vaidyanathan, R. (2008). Assortment planning: Review of literature and industry practice. In N. Agrawal, & S. A. Smith (Eds.), *Retail Supply Chain Management* (pp. 99–153). Boston, MA: Springer volume 122 of *International Series in Operations Research & Management Science*. doi:[10.1007/978-0-387-78902-6_6](https://doi.org/10.1007/978-0-387-78902-6_6).
- Hübner, A., Düsterhöft, T., & Ostermeier, M. (2021). Shelf space dimensioning and product allocation in retail stores. *European Journal of Operational Research*, *292*, 155–171. doi:[10.1016/j.ejor.2020.10.030](https://doi.org/10.1016/j.ejor.2020.10.030).
- Hübner, A., & Schaal, K. (2017). An integrated assortment and shelf-space optimization model with demand substitution and space-elasticity effects. *European Journal of Operational Research*, *261*, 302–316. doi:[10.1016/j.ejor.2017.01.039](https://doi.org/10.1016/j.ejor.2017.01.039).
- Hübner, A. H., & Kuhn, H. (2012). Retail category management: State-of-the-art review of quantitative research and software applications in assortment and shelf space management. *Omega*, *40*, 199–209. doi:[10.1016/j.omega.2011.05.008](https://doi.org/10.1016/j.omega.2011.05.008).
- Huddleston, P. T., Behe, B. K., Driesener, C., & Minahan, S. (2018). Inside-outside: Using eye-tracking to investigate search-choice processes in the retail environment. *Journal of Retailing and Consumer Services*, *43*, 85–93. doi:[10.1016/j.jretconser.2018.03.006](https://doi.org/10.1016/j.jretconser.2018.03.006).
- Inman, J. J., Winer, R. S., & Ferraro, R. (2009). The interplay among category characteristics, customer characteristics, and customer activities on in-store decision making. *Journal of Marketing*, *73*, 19–29. doi:[10.1509/jmkg.73.5.19](https://doi.org/10.1509/jmkg.73.5.19).
- Kacen, J. J., Hess, J. D., & Walker, D. (2012). Spontaneous selection: The influence of product and retailing factors on consumer impulse purchases. *Journal of Retailing and Consumer Services*, *19*, 578–588. doi:[10.1016/j.jretconser.2012.07.003](https://doi.org/10.1016/j.jretconser.2012.07.003).
- Karampatsa, M., Grigoroudis, E., & Matsatsinis, N. F. (2017). Retail category management: A review on assortment and shelf-space planning models. In E. Grigoroudis, & M. Doumpos (Eds.), *Operational Research in Business and Economics* Springer Proceedings in Business and Economics (pp. 35–67). Springer, Cham. (1st ed.). doi:[10.1007/978-3-319-33003-7_3](https://doi.org/10.1007/978-3-319-33003-7_3).

- Krebs-Smith, S. M., Pannucci, T. R. E., Subar, A. F., Kirkpatrick, S. I., Lerman, J. L., Tooze, J. A., Wilson, M. M., & Reedy, J. (2018). Update of the Healthy Eating Index: HEI-2015. *Journal of the Academy of Nutrition and Dietetics*, *118*, 1591–1602. doi:[10.1016/j.jand.2018.05.021](https://doi.org/10.1016/j.jand.2018.05.021).
- Kroese, F. M., Marchiori, D. R., & De Ridder, D. T. (2016). Nudging healthy food choices: A field experiment at the train station. *Journal of Public Health*, *38*, e133–e137. doi:[10.1093/pubmed/fdv096](https://doi.org/10.1093/pubmed/fdv096).
- Martinez, O., Rodriguez, N., Mercurio, A., Bragg, M., & Elbel, B. (2018). Supermarket retailers' perspectives on healthy food retail strategies: In-depth interviews. *BMC Public Health*, *18*, 1019. doi:[10.1186/s12889-018-5917-4](https://doi.org/10.1186/s12889-018-5917-4).
- Romero, M., & Biswas, D. (2016). Healthy-left, unhealthy-right: Can displaying healthy items to the left (versus right) of unhealthy items nudge healthier choices? *Journal of Consumer Research*, *43*, 103–112. doi:[10.1093/jcr/ucw008](https://doi.org/10.1093/jcr/ucw008).
- Shan, Z., Rehm, C. D., Rogers, G., Ruan, M., Wang, D. D., Hu, F. B., Mozaffarian, D., Zhang, F. F., & Bhupathiraju, S. N. (2019). Trends in dietary carbohydrate, protein, and fat intake and diet quality among US adults, 1999-2016. *JAMA*, *322*, 1178–1187. doi:[10.1001/jama.2019.13771](https://doi.org/10.1001/jama.2019.13771).
- Shaw, S. C., Ntani, G., Baird, J., & Vogel, C. A. (2020). A systematic review of the influences of food store product placement on dietary-related outcomes. *Nutrition Reviews*, *78*, 1030–1045. doi:[10.1093/nutrit/nuaa024](https://doi.org/10.1093/nutrit/nuaa024).
- Sigurdsson, V., Saevarsson, H., & Foxall, G. (2009). Brand placement and consumer choice: An in-store experiment. *Journal of Applied Behavior Analysis*, *42*, 741–745. doi:[10.1901/jaba.2009.42-741](https://doi.org/10.1901/jaba.2009.42-741).
- Thaler, R. H., & Sunstein, C. R. (2009). *Nudge: Improving decisions about health, wealth, and happiness*. Penguin Books.
- Tsai, C.-Y., & Huang, S. (2015). A data mining approach to optimise shelf space allocation in consideration of customer purchase and moving behaviours. *International Journal of Production Research*, *53*, 850–866. doi:[10.1080/00207543.2014.937011](https://doi.org/10.1080/00207543.2014.937011).
- Van Gestel, L. C., Kroese, F. M., & De Ridder, D. T. (2018). Nudging at the checkout counter: A longitudinal study of the effect of a food repositioning nudge on healthy food choice. *Psychology & Health*, *33*, 800–809. doi:[10.1080/08870446.2017.1416116](https://doi.org/10.1080/08870446.2017.1416116).
- Van Kleef, E., Otten, K., & Van Trijp, H. C. (2012). Healthy snacks at the checkout counter: A lab and field study on the impact of shelf arrangement and assortment structure on consumer choices. *BMC Public Health*, *12*, 1072. doi:[10.1186/1471-2458-12-1072](https://doi.org/10.1186/1471-2458-12-1072).

Vecchio, R., & Cavallo, C. (2019). Increasing healthy food choices through nudges: A systematic review. *Food Quality and Preference*, 78, 103714. doi:[10.1016/j.foodqual.2019.05.014](https://doi.org/10.1016/j.foodqual.2019.05.014).