

Anticipatory Shipment for Pickup Point Supply

Monica Rasini^{*a}, Niels Agatz^b, and Elena Tappia^a

^aDepartment of Management, Economics and Industrial Engineering, Politecnico di Milano, Via Lambruschini 4/B, 20156 Milano, Italy

^bRotterdam School of Management, Erasmus University, Burgemeester Oudlaan 50, 3062 PA Rotterdam The Netherlands

Please cite this document as: Rasini, M., Agatz, N., Tappia, E. (2020). Anticipatory shipment for pickup point supply. *Omega*, 93, 102089, <https://doi.org/10.1016/j.omega.2019.07.005>

Abstract

Many retailers allow customers to shop online and collect their orders at a pickup point nearby. In this paper, we study the anticipatory shipment of items to such pickup points in order to improve service and operational efficiency. We formulate a stochastic programming model to support the selection of products and associated quantities to ship to the pickup points in anticipation of customers' demand. The results of our numerical experiments suggest that anticipatory shipments can have substantial benefits both in terms of cost and lead-time. The benefits increase with the storage space at the pickup point. The anticipatory shipment strategy is especially beneficial in a setting which requires short delivery lead-times and when the e-fulfilment warehouse is further away.

Keywords: e-fulfilment, pickup point supply, anticipatory shipment, stochastic programming.

1 Introduction

Many online retailers allow customers to have their purchases shipped to a pickup point nearby. Especially in-store pickup services are becoming increasingly popular. This service lets customers take advantage of the convenience of online shopping, e.g., easy search and price comparison, without the shipping fees typically associated with home delivery. Moreover, in-store pickup involves less hassle than regular shopping as the items have already been picked and packed by store staff for pickup. A recent study by Accenture

^{*}Corresponding Author: monica.rasini@polimi.it

(2015) shows that 25% of consumers choose the pickup option when shopping online and this percentage is expected to grow further over the next years.

At the same time, online customers expect ever shorter lead-times including pickup or delivery within a few hours. Fast, same-day, service is attractive as it combines the advantages of online convenience with the immediacy of a traditional store. Several leading online retailers now offer a same-day pickup service as a cheaper alternative to home delivery. Amazon, for example, offers same-day and ‘instant’ pickup services in several cities in the U.S.(Dastin, 2017). Other examples include Target, Home Depot and Lowe’s (CourtneyReagan, 2017).

It may be possible to serve in-store pickup orders directly from the store’s inventory when the one store offers the same assortment as the physical store. However, this is not always the case. In the Netherlands, for example, general internet retailer Bol.com uses Albert Heijn supermarkets to serve as pickup points, offering both next-day and same-day pickup. This means that the retailer ships the pickup orders from a central (suburban) warehouse location to the (urban) store location. For next-day pickup, the volumes may justify a fixed delivery schedule between the warehouse and the pickup point. This allows the retailer to consolidate multiple pickup orders and thus create cost-efficient shipments. However, it is more challenging to efficiently organize an on-demand transportation system with short lead-times.

One way for the retailer to increase the efficiency of the fulfilment operations is to stock a selection of its assortment at the pickup location in anticipation of future customers’ demand. Amazon patented such an “anticipatory shipment” strategy that aims to predict when a customer will order an item and push it through the shipping process before the customer actually places an order (Bensinger, 2014). Products are moved to nearby hubs by shipping a package with a partial address. When the order arrives, the package can be delivered to the customer’s address far more quickly.

Using this strategy, the retailer must select the products and quantities to move to the pickup location in anticipation of demand. In this paper, we investigate under what circumstances ‘anticipatory shipment’ provides benefits for retailers and evaluate different shipment strategies.

The main contributions of this paper are as follows. First, we introduce and discuss a new anticipatory shipment strategy to supply pickup locations. Second, we present a stochastic programming model, as well as a simple heuristic procedure to allocate the limited space available at pickup point to different products. Finally, we conduct an extensive computational study to investigate under what circumstances anticipatory shipment is most beneficial. To evaluate the impact of the anticipatory shipment decisions, we use simple costs approximations and also simulate vehicle dispatches in more detail. Results show that cost saving and increased service levels are possible by applying the anticipatory

shipment strategy. The benefits increase with the storage space at the pickup point. The anticipatory shipment strategy is especially beneficial in settings where the distance between the warehouse and the pickup point is high and customers request short lead-times.

The remainder of the paper is organized as follows. The next section summarizes the relevant literature. In section 3, we describe the problem and the model. Section 4 presents our solution approach. In section 5, we describe our computational experiments and present the results in section 6. The final section summarizes our conclusions and presents directions for future research.

2 Related literature

Our work is related to the stream of literature on the fulfilment of goods ordered online or e-fulfillment. The use of (in-store) pickup points also links it to the area of multi-channel or omni-channel retailing (Agatz et al., 2008; Hübner et al., 2016; ?).

Most work on the so-called ‘last-mile’ of e-fulfillment focusses on the home delivery of goods ordered online in narrow delivery time windows, see, e.g., Boyer et al. (2009), Agatz et al. (2011), Ehmke and Campbell (2014), Yang et al. (2014) and ?. More recently, several researchers started to develop dispatching strategies for same-day delivery in which delivery requests need to be fulfilled within a short time period. The main challenge in this line of research is to determine the timing of the vehicle dispatches and the assignment of the delivery requests to the vehicle routes given the continuous arrival of new requests (???)

There is less research on the use of pickup points to organize the last-mile to the online customer. Mahar and Wright (2017) develop a model to determine the optimal subset of stores in which in-store pickup and return capabilities should be located. Mahar et al. (2012) present a policy to dynamically adjust the set of available pickup locations presented to the customers. Gao and Su (2016) show that offering an in-store pickup service may help retailers to expand their market coverage but that it is not suitable for all products. Gallino and Moreno (2014) empirically analyze the impact of such an in-store pickup service and conclude that it results in lower online sales but higher store sales. Kim et al. (2017) consider the customer perspective and study the factors that determine the intention to use the in-store pickup service. None of these papers study the specific challenges of same-day pickup point supply.

Conceptually, the selection of products to stock at a pickup point in anticipation of demand is similar to determining which subset of products to stock at a retail store. This area of research is known as assortment planning and typically aims to maximize sales given a certain customer choice model subject to various limitations on the number of products to stock, see Kök et al. (2008), Pentico (2008), and Hübner and Kuhn (2012) for reviews.

In contrast to most assortment planning problems in the literature, customer choice and substitution behavior is less relevant in our problem as we do not select the assortment for the customer. All customer orders are fulfilled, either directly from the pickup point inventory or via an on-demand shipment to the pickup point from the warehouse.

Moreover, unlike most assortment planning problems, we do not only focus on selecting which products to stock but also on how many items to stock. This links our work to the area of inventory management and production planning in which demand uncertainty is a key element (Mula et al., 2006). Most work in this area is based on stochastic optimization, see for example Eppen et al. (1989); Shapiro and Homem-de Mello (1998); Gupta and Maranas (2003); Rao et al. (2004). Most models in this area work based on expected demand or scenario analysis, e.g., see Escudero et al. (1993); Shapiro et al. (2014).

3 Problem description and model

We consider a retailer that sells P products online, indexed by $i \in \mathcal{P} = \{1, \dots, P\}$. Demand for the pickup of item i is a random variable d_i that follows a known probability distribution. The demand for different products is independent of each other. Online customers can collect their purchased goods at a pickup location of their choice within a few hours after placing the order. The selling season is composed of T sales periods, indexed by $t \in \mathcal{T} = \{1, \dots, T\}$, where a period could, for example, represent the opening time of the pickup point in a day.

The demand for a given item i can be supplied in two ways: (i) directly from the available stock at the pickup point or (ii) by an ‘on-demand’ emergency shipment from a dedicated e-fulfilment warehouse to the pickup point. To move stock to the pickup point, we incur an ‘anticipatory’ shipment cost of c_i^a per unit i . The unit shipment costs for an on-demand emergency shipment is c_i^d . The shipment costs include picking, packing, handling and transportation costs. We assume that the anticipatory shipment costs are typically lower than the on-demand shipment costs as there are more opportunities for consolidation, i.e., $c_i^a < c_i^d$. For ease of exposition, we use the same unit cost c_i^a for the anticipatory and the return shipments.

Let Q_i^t denote the inventory of product $i \in \mathcal{P}$ at the start of period $t \in \mathcal{T}$ and let I_i^t denote the inventory at the end of the period. That is, $I_i^t = (Q_i^t - d_i^t)^+$ for $i \in \mathcal{P}$. There is limited storage space K at the pickup location. Given k_i required storage capacity per unit of product $i \in \mathcal{P}$, this means that $\sum_{i \in \mathcal{P}} k_i Q_i^t \leq K$.

Our objective is to determine how many items Q_i of each product i to stock at the pickup location in each period as to minimize the total costs. We represent the state of the system at the end of each period by the vector $\mathbf{I}^t = (I_1^t, \dots, I_P^t)$. State \mathbf{I}^{t-1} and target level \mathbf{Q}^t together define the shipment quantities. When $Q_i^t - I_i^{t-1} > 0$, we ship $Q_i^t - I_i^{t-1}$ items of product i from the warehouse to the pickup point. When negative, $|Q_i^t - I_i^{t-1}|$ items of

product i are moved back to the warehouse. The pickup point supply is performed between period $t - 1$ and t so that the target level Q^t is available at the beginning of the period. The number of emergency shipments to the pickup point can be expressed as $(d_i^t - Q_i^t)^+$.

Hence, the problem is dynamic as we can update the pickup point inventories in each period and stochastic because we only have probabilistic information about the customers' demand. The t -stage costs for period t $F^t(\mathbf{I}^{t-1}, \mathbf{Q}^t, \mathbf{d}_i^t)$, including anticipatory, return and on-demand shipments, are:

$$F^t(\mathbf{I}^{t-1}, \mathbf{Q}^t, \mathbf{d}_i^t) = \sum_{i \in \mathcal{P}} [c_i^a |Q_i^t - I_i^{t-1}| + c_i^d (d_i^t - Q_i^t)^+]$$

The objective for this stochastic dynamic programming problem is to identify the optimal *decision policy* $\pi^* \in \Pi$ that minimizes the expected costs given an initial state \mathbf{I}^0 . A policy $\pi \in \Pi$ is a subsequence of decision rules $\pi = (Q^{0\pi}, Q^{1\pi}, \dots, Q^{T\pi})$, where each decision rule $Q^{t\pi}$ is a function that specifies the decision for state \mathbf{I}^{t-1} and follows policy π . An optimal policy π^* minimizes the total expected costs, conditional on initial state \mathbf{I}^0 :

$$\pi^* = \arg \min_{\pi \in \Pi} \mathbb{E} \left[\sum_{t=0}^T F(\mathbf{I}^t, \mathbf{Q}^{t+1\pi}, \mathbf{d}_i^{t+1}) | \mathbf{I}^0 \right] \quad (1)$$

Table 1 summarizes the main notation used in the paper.

<i>System sets</i>	
\mathcal{P}	Set of products (indexed by i)
\mathcal{T}	Set of periods (indexed by t)
\mathcal{S}	Set of scenarios (indexed by s)
<i>System parameters</i>	
K	Storage capacity at pickup point
k_i	Storage capacity required by one unit of product i
d_i^t	Demand of product i in period t
c^a	Anticipatory shipment cost per unit
c^d	On-demand shipment cost per unit
<i>Decision variables and derived quantities</i>	
Q_i^t	Inventory level of product i at the start of period t at the pickup point
I_i^t	Inventory level of product i at the end of period t
h_i^t	Number of units of product i transported by anticipatory shipment
w_i^t	Number of units of product i transported by on-demand shipment

Table 1: Main notation

4 Solution approach

Identifying a general optimal policy, e.g., by using backward dynamic programming, is computationally intractable due to the ‘‘curse of dimensionality’’ (Powell, 2007). Therefore, we simplify the model by limiting: (1) the look-ahead period that is taken into account at each decision point, and (2) the possible number of demand realizations at each stochastic transition. In particular, we consider a one-period look-ahead and we use the Sample Average Approximation (SAA) approach in which uncertainty is captured through a set S of equally likely scenarios (Shapiro and Philpott, 2007). This method is based on two steps: (i) sampling and (ii) deterministic optimization. First, a sample of N realizations of the random variables d_i^t is generated and consequently for each product the expected single-stage cost $\mathbb{E}[F_i^t(I_i^{t-1}, Q_i^t, d_i^t)]$ is approximated by the sample average function:

$$\mathbb{E}[F_i^t(I_i^{t-1}, Q_i^t, d_i^t)] \approx \frac{1}{N} \sum_{s \in \mathcal{S}} F_i^t(I_i^{t-1}, Q_i^t, d_i^{ts})$$

where d_i^{ts} denotes the demand of product i for period t in scenario s . Then, the SAA problem is modeled as a deterministic optimization problem.

Let h_i^t denote the number of units shipped in anticipation of demand to meet the target level Q^t at the beginning of period t for product i , and w_i^{ts} be the number of units shipped on-demand during this period in scenario s . This gives the following model.

$$\text{minimize} \quad \sum_{i \in \mathcal{P}} \left(c_i^a h_i^t + \frac{1}{N} \sum_{s \in \mathcal{S}} c_i^d w_i^{ts} \right) \quad (2)$$

$$\text{s.t.} \quad \sum_{i \in \mathcal{P}} (k_i Q_i^t) \leq K \quad (3)$$

$$h_i^t \geq Q_i^t - I_i^{t-1} \quad \forall i \in \mathcal{P} \quad (4)$$

$$h_i^t \geq I_i^{t-1} - Q_i^t \quad \forall i \in \mathcal{P} \quad (5)$$

$$w_i^{ts} \geq d_i^{ts} - Q_i^t \quad \forall i \in \mathcal{P} \quad \forall s \in \mathcal{S} \quad (6)$$

$$w_i^{ts} \geq 0 \quad \forall i \in \mathcal{P} \quad \forall s \in \mathcal{S} \quad (7)$$

$$Q_i^t \geq 0 \quad \forall i \in \mathcal{P} \quad (8)$$

$$Q_i^t \in \mathbb{Z} \quad \forall i \in \mathcal{P} \quad (9)$$

The objective function that aims to minimize the average total single-stage costs is given in (2). Constraint (3) makes sure that the pickup point capacity is respected. Constraints (4) and (5) ensure that for each product h_i^t is equal to $\max\{Q_i^t - I_i^{t-1}, I_i^{t-1} - Q_i^t\}$. Constraints (6) and (7) make sure that for each product and each scenario w_i^{ts} is equal to $(d_i^{ts} - Q_i^t)^+$. Finally, constraints (8) and (9) define non-negative and integer variables.

We solve the problem within a rolling horizon framework in which the identified policy Q^* is used to dynamically update the pickup point inventories at each decision point $t \in \mathcal{T}$.

Moreover, we extend the model to consider a two-period look-ahead to make the anticipatory shipment decisions. The two-period look-ahead solution (SAA 2) specifies the recommended shipping levels for the period t knowing decisions for all the previous periods and having only probabilistic information about the demand for periods t and $t + 1$. Decisions for period $t + 1$ are made at the next iteration of the rolling horizon.

$$\text{minimize } \sum_{i \in \mathcal{P}} \left[c_i^a h_i^t + \frac{1}{N} \sum_{s \in \mathcal{S}} \left(c_i^d w_i^{ts} + c_i^a h_i^{(t+1)s} + c_i^d w_i^{(t+1)s} \right) \right] \quad (10)$$

$$\text{s.t. } \sum_{i \in \mathcal{P}} (k_i Q_i^t) \leq K \quad (11)$$

$$\sum_{i \in \mathcal{P}} (k_i Q_i^{(t+1)s}) \leq K \quad \forall s \in \mathcal{S} \quad (12)$$

$$I_i^t \geq Q_i^t - I_i^{t-1} \quad \forall i \in \mathcal{P} \quad (13)$$

$$I_i^t \geq I_i^{t-1} - Q_i^t \quad \forall i \in \mathcal{P} \quad (14)$$

$$I_i^{(t+1)s} \geq Q_i^{(t+1)s} - I_i^{ts} \quad \forall i \in \mathcal{P} \quad \forall s \in \mathcal{S} \quad (15)$$

$$h_i^{(t+1)s} \geq I_i^{ts} - Q_i^{(t+1)s} \quad \forall i \in \mathcal{P} \quad \forall s \in \mathcal{S} \quad (16)$$

$$I_i^{ts} \geq d_i^{ts} - Q_i^t \quad \forall i \in \mathcal{P} \quad \forall s \in \mathcal{S} \quad (17)$$

$$I_i^{(t+1)s} \geq d_i^{(t+1)s} - Q_i^{(t+1)s} \quad \forall i \in \mathcal{P} \quad \forall s \in \mathcal{S} \quad (18)$$

$$I_i^{ts} \geq 0 \quad \forall i \in \mathcal{P} \quad \forall s \in \mathcal{S} \quad (19)$$

$$w_i^{(t+1)s} \geq 0 \quad \forall i \in \mathcal{P} \quad \forall s \in \mathcal{S} \quad (20)$$

$$I_i^t \geq 0 \quad \forall i \in \mathcal{P} \quad (21)$$

$$I_i^{(t+1)s} \geq 0 \quad \forall i \in \mathcal{P} \quad \forall s \in \mathcal{S} \quad (22)$$

$$Q_i^t \in \mathbb{Z} \quad \forall i \in \mathcal{P} \quad (23)$$

$$Q_i^{(t+1)s} \in \mathbb{Z} \quad \forall i \in \mathcal{P} \quad \forall s \in \mathcal{S} \quad (24)$$

Note that the solution time of this model grows quickly with the number of scenarios and products. Therefore, in our experiment, we simplify the model by considering only one $Q_i^{(t+1)s}$ across all scenarios. Experiments on small problems suggest that this does not deteriorate the solution quality.

5 Experimental design

In our computational experiments, we evaluate and compare the model as presented in Section 4 using a one period look-ahead based on the expected demand (ED) and with

multiple scenarios (SAA 1) and a two period look-ahead with multiple scenarios (SAA 2). In the scenario-based approaches, we use 20 scenarios as we believe this provides a good trade-off between run time and solution quality. Figure 1 shows the total cost in the setting with expected demand equal to the pickup point storage capacity ($K = 100\%$), applying SAA 1 and a different number of scenarios. The costs are normalized by setting the cost for SAA with one scenario equal to 100.




Figure 1: Total cost for SAA 1 and different number of scenarios in the setting with expected demand equal to the pickup point storage capacity; average results across five runs (each with ten instances and 20 periods)

Moreover, we provide the results of the following three benchmarks: (1) NOANT: solution when all customer pickup orders are dispatched from the warehouse on demand without any anticipatory shipments, (2) GRD: solution of greedy heuristic procedure that prioritises products with a higher expected demand-to-weight ratio and, for each selected product, stores the minimum of the expected demand and the remaining capacity at the pickup location and (3) HINDSIGHT: solution of the optimization model with “perfect information”, i.e., in which the demand in the look-ahead period is known when making decisions as described in Appendix A. This strategy provides a theoretical lower bound on the system-wide cost.

The proposed solution procedures are implemented in Java, with CPLEX V12.6.3 callable library routines for solving the mixed-integer linear programming problems. We generate twenty operating days based on actual demand data from a multi-channel retailer with pickup points in the Netherlands. We report the following statistics, where the averages are per period taken over ten random instances, each consisting of 20 periods: (1) the total cost (i.e., preparation and transportation costs) of satisfying all customers’ demand in a single period, (2) the number of items moved from and to the warehouse for anticipatory, return and on-demand shipments, (3) orders direct, i.e., the percentage of

customers’ orders served directly from the inventory at the pickup point, (4) Lead-time which represents the average time, in hours, between the order’s placement time and the earliest pickup time, (5) Capacity used, i.e., the percentage of pickup point capacity used at the beginning of a period, (6) the number of products stored as a percentage of the total assortment, and (7) the leftover inventory, i.e, the number of items at the pickup point at the end of a period.

6 Results

This section presents a computational study to compare the performance of different shipment strategies (section 6.1), and analyze the impact of various parameters on the benefits of anticipatory shipments (section 6.2).

6.1 Shipment strategy comparison

The first set of experiments is aimed at assessing the performance of the different shipment strategies as presented in section 5. We consider an assortment \mathcal{P} of 100 products. The expect demand of each product $i \in \mathcal{P}$ is drawn from a truncated Beta distribution with $\alpha = 2$ and $\beta = 5$ between $[1, 10]$. Subsequently, we use a Poisson distribution to generate the demand realizations on each day based on these expected values. We draw the size of the products k_i from a truncated Exponential Distribution with $\lambda = 0.5$ between $[5, 100]$. We set the anticipatory shipment costs to 5 €/unit (c^a) and 25 €/unit for on-demand shipments. Table 2 summarizes the default parameter values used in the first set of experiments.

Note that in our experiments, we only consider the special case in which the probability distribution is the same for each period. In this case, we can simplify the problem as the decision problem is the same in each period. This means that we only have to determine one optimal ‘order-up-to’ policy Q^* that we can use in all periods and that there are no return shipments between periods.

We normalize the results by dividing by the cost of the NOANT strategy and then multiply by 100. Table 3 provides the results for three different pickup point capacities as a percentage of the total expected demand, i.e., $K = 50\%$, 100% and 150% . As expected, we see that using the anticipatory shipments result in cost savings and shorter lead-times to the customer. The cost savings increase with the available storage space at the pickup point and range between 49.8 percent and 75.6 percent. Looking at the theoretical hindsight benchmark, we see that further savings are possible if demand is known in advance. We can interpret the difference between HINDSIGHT and NOANT as the costs of the dynamic same-day pickup as compared to a next-day pickup. Our anticipation strategy can offset

<i>System-related parameters</i>	
Number of periods [days]	$T = 20$
<i>Assortment-related parameters</i>	
Number of products	$P = 100$
Size of product [units of capacity]	$k_i \sim \epsilon(\lambda = 0.5), k_i \in [5; 100]$
<i>Demand-related parameters</i>	
Mean demand [units]	$\mu_i \sim \beta(\alpha = 2; \beta = 5), \mu_i \in [1; 10]$
Delivery lead-time [hours]	$L = 2$
<i>Cost-related parameters</i>	
Anticipatory and return shipment cost [€/unit]	$c^a = 5$
Shipment cost for on-demand transportation system [€/unit]	$c^d = 25$

Table 2: Default parameter values

some of the additional costs, especially if there is sufficient storage capacity available at the pickup point.

Looking at the lead-times, we see that the customers receive faster service in the anticipatory shipment strategy. Figure 2 shows the distribution of the time between the order’s placement time and the time it can be collected by the customer in the over-capacity setting. We see that SAA 1 satisfies 91 percent of customers’ orders immediately, while this takes one hour and twenty minutes on average in the NOANT setting.

Overall, the stochastic policies (SAA 1 and SAA 2) perform better than the deterministic policies (ED and GRD). The cost difference is between 3 and 12 percent points. To better understand the differences in performance, we can look at the solutions in terms of the used capacity and the number of stored products. In the under-capacity setting, we see that while all policies fully use all capacity, the SAA policies store a wider range of products at the pickup point. With more available capacity, all methods store the full range of products but the SAA policies store more items of each product. In both cases, this results in more orders that can be served directly from the pickup point and less on-demand dispatches. This comes at the costs of producing more leftovers at the end of each period but this does not lead to additional costs as these items can be used in the next period.

For the deterministic policies (ED and GRD), there are no benefits in increasing the look-ahead period, as well as in considering the expected demand for all the remaining planning horizon; while for the stochastic policy (SAA) using a two-period look-ahead instead of a one-period look-ahead provides slightly better results. Especially in the setting with abundant capacity, we observe that a longer look-ahead period is beneficial, i.e., a 3.3 percent point decrease in the total cost. With a longer look-ahead period, we observe an

Strategy	Total cost	Units moved	Orders direct	Lead-time	Used capacity	Products stored	Leftover inventory
<i>Under-capacity setting ($K = 50\%$)</i>							
NOANT	100.0	356	0	1.22	0	0	0
GRD	53.1	359	59	0.49	100	67	50
ED	51.9	359	60	0.48	100	77	54
SAA 1	50.9	359	62	0.46	100	84	59
SAA 2	50.2	360	63	0.45	100	80	79
HINDSIGHT	39.6	356	76	0.30	100	71	0
<i>Expected demand equal to capacity ($K = 100\%$)</i>							
NOANT	100.0	356	0	1.22	0	0	0
GRD	37.0	360	79	0.25	95	98	71
ED	37.0	360	79	0.25	95	100	72
SAA 1	33.4	362	84	0.19	100	100	111
SAA 2	32.6	362	85	0.18	100	99	141
HINDSIGHT	20.8	356	99	0.01	94	93	0
<i>Over-capacity setting ($K = 150\%$)</i>							
NOANT	100.0	356	0	1.22	0	0	0
GRD	36.6	360	80	0.24	64	100	72
ED	36.6	360	80	0.24	64	100	72
SAA 1	27.7	364	91	0.11	88	100	161
SAA 2	24.4	365	95	0.06	99	100	245
HINDSIGHT	20.0	356	100	0.00	64	94	0

Table 3: Base results, with 100 products and $c^a/c^d = 0.2$; average results across ten instances and 20 periods

increase in the number of items moved to the pickup point in anticipation of customers' demand as it is less likely that the stored products are not requested by the customers if we consider more periods.




Figure 2: Lead-time distribution in the over-capacity setting, with 100 products and $c^a/c^d = 0.2$; results of ten instances and 20 periods

6.2 Impact of system parameters affecting anticipatory shipments

In this section, we investigate the impact of different parameters affecting the anticipatory shipment decisions on the performance of the system. In particular, we vary the number of products offered online and anticipatory shipment cost compared to the on-demand transportation system. In these experiments, we compare the NOANT strategy and the anticipatory shipment strategy based on SAA method with two-period look-ahead.

6.2.1 Impact of assortment size

First, we study the impact of varying the number of products offered P . In addition to the base case of 100, we consider 50 and 200 products. To isolate the effect of the number of products, we assume that the expected demand remains the same but is split among a different number of products. Figure 3 illustrates the main results of these experiments.

We generally see that the cost savings decrease with the number of products P . Considering the over-capacity setting, the cost is 32 percent of the cost for NOANT strategy in the case of 200 products, leading to a cost savings of up to 68 percent. The saving increases up to 77.5 percent for the case with only 50 products. This is because with less products, more orders can be fulfilled directly from the pickup point stock.

To understand this, we need to consider that at the beginning of the period, before a realization of the demand is known, the retailer must take two decisions: (i) which products to store at pickup point and (ii) how much inventory to carry at pickup point for each product selected. With a smaller number of products the first decision is easier. The capacity allocation to a larger assortment leads to a smaller percentage of products of the overall assortment stored at pickup point and, thus, there is a higher risk that products not selected to be stocked at the pickup point are requested by customers. This creates more leftovers and more emergency on-demand shipments and thus lower benefits from the anticipatory shipment strategy.




Figure 3: Results varying the number of products, $n = 50; 100; 200$; average results across ten instances and 20 periods

The detailed results of the different shipment strategies are reported in Table 5 (see Appendix B). Results are aligned with comments provided in Section 6.1 in terms of comparison of the different strategies.

6.2.2 Impact of the anticipatory shipment cost

Next, we investigate the impact of the ratio between the anticipatory shipment costs (c^a) and the on-demand shipment costs (c^d). In particular, we fix c^d and change c^a , leading to c^a/c^d ratios of 0.2, 0.4 and 0.8. Figure 4 presents the main results of these experiments.

As expected, the use of anticipatory shipments decreases as c^a increases relative to c^d . With c^a/c^d increased from 0.2 to 0.8, the percentage of the pickup point capacity used to store products in anticipation of customers' demand is greatly reduced, e.g., from 99 percent to 62 percent in the over-capacity setting. This means more emergency on-demand shipments and, thus, lower costs savings. When the costs of the emergency shipments are similar to the costs of the anticipatory shipments, there is no reason to stock products at pickup point in anticipation of customers' demand. Conversely, the retailer waits for the customer order and sends to the pickup point exactly the required products and quantities.

The detailed results of the different shipment strategies are reported in Table 6 (see

Appendix B). As before, SAA2 provides slightly better results compared to SAA1. However, stochastic policies do not provide benefits with a ratio c^a/c^d of 0.8 coherently with the reduction in the use of anticipatory shipments when increasing the related costs.

The costs of the emergency on-demand shipment relate to various factors, including the distance between the warehouse and the pickup point and the available time to consolidate shipments. We explicitly model these factors in section 6.3.




Figure 4: Results varying the anticipatory shipment cost with $c^a/c^d = 0.2; 0.4; 0.8$; average results across ten instances and 20 periods

6.3 Impact of system parameters affecting on-demand dispatches

Up to now, we did not explicitly model the operational level vehicle dispatch decisions associated with the emergency on-demand shipments. In this section, we introduce a simple vehicle dispatching policy to study the consolidation of orders for joint delivery at an operational level. We particularly investigate the impact of customer lead-time, distance between warehouse and pickup point and dispatch flexibility on the benefits of the anticipatory shipment strategy. We use the values reported in Table 2 for the system-, assortment- and demand-related parameters, while the cost-related parameters are explained in the following sections. The detailed results of the different shipment strategies are reported in Appendix C.

6.3.1 Vehicle dispatching policy

At an operational level, the retailer continuously receives customers' pickup requests. After a customer places an order j with delivery lead-time L , the retailer checks the inventory availability at the associated pickup point. If the item is not available at the pickup point, it is shipped there from the warehouse, within the delivery lead-time L .

To manage the on-demand transportation system, we apply the following vehicle dispatching policy. Let \mathcal{J} be the set of orders to be served from the warehouse. Each order $j \in \mathcal{J}$ has a dispatch time window $[e_j; l_j]$. The earliest dispatch time $e_j = \tau_j + t^p$ indicates

at what time the product is ready to be shipped and the latest dispatch time $l_j = \tau_j + L - t^t$ corresponds to the time before which it has to be shipped to arrive in time. Let B be the set of dispatch times within the order period. To ensure that we can serve all orders in time, the time between the dispatches should be less than or equal to $L - t^t$. All orders $j \in \mathcal{J}$ that have a latest dispatch time l_j between $t \in B$ and $t + 1 \in B$ are dispatched at time b .

In applying this vehicle dispatching policy, we assume a sufficient number of vehicles available to ensure all orders are dispatched in time. We introduce the vehicle operating cost for a dispatch, fixed to 200 €/dispatch. We consider eight dispatch moments within a period and a vehicle with 10,000 units of capacity, half an hour as the travel time between the warehouse and the pickup location, 15 minutes as the preparation time for all the orders and 5 €/unit as preparation costs for each unit moved regardless of the shipping mode. Furthermore, we consider a two-hour delivery lead-time. Finally, we assume that the fulfillment location for the online orders also supplies the traditional stores and hence the retailer can combine the store and pickup point replenishments in one route, without the need for additional vehicles dedicated to anticipatory and returns shipments.

6.3.2 Vehicle dispatches and cost evaluation

Our anticipatory shipment model uses a fixed unit transshipment costs to model the costs of the emergency on-demand dispatches. We now evaluate whether this simplified cost structure can be used to capture the actual costs of managing an on-demand delivery system. We simulate the costs by recording the number of vehicles dispatches, the distance-based operating cost, and the preparation costs.

Table 4 compares the “real” total cost, and the cost estimated through the proportional cost model considered in the previous sections, referred to as “model” total cost. All the costs are normalized by setting the “model” total cost for NOANT equal to 100. Results show that the model with proportional costs provides a fairly accurate approximation of the real costs. The cost difference is explained by the order consolidation decisions: the model assumes ten units per vehicle on average, while different values can occur when simulating the vehicle dispatches. Lower is the percentage of orders served from the inventory available at the pickup point, higher is the retailer’s capability of aggregating orders and shipping more units per vehicle.

In any case, the anticipatory shipment strategy allows cost savings in all the settings. The anticipatory shipment strategy, as opposed to the on-demand shipment strategy (NOANT), results in a significant reduction of the number of vehicles dispatches, and therefore of the transportation costs, due to an increased number of orders fulfilled from the pickup point during the order period (from 19.5 vehicles for NOANT to 8.3 for SAA 1

in the over-capacity setting).

Strategy	K	“Model” total cost	“Real” total cost	Vehicles dispatches	Orders direct
NOANT	-	100.0	63.8	19.5	0
SAA 1	50%	50.9	55.6	15.8	62
SAA 1	100%	33.4	45.1	11.0	84
SAA 1	150%	27.7	39.0	8.3	91

Table 4: Base results, with 100 products, $c^a/c^d = 0.2$, eight dispatches moments, customer lead-time of two hours and travel time of half an hour; average results across ten instances and 20 periods

6.3.3 Impact of the customer lead-time

First, we investigate the impact of changing the delivery lead-time L offered by the system. We consider two additional scenarios: one with a shorter delivery lead-time of one hour, and one with a longer delivery lead-time of four hours. Figure 5 shows the changes in the total cost compared to the base case of two hours.

As expected, we see that the costs increase with a shorter lead-time. Furthermore, we observe that the impact of decreasing the delivery lead-time is much stronger than the impact of increasing it. With the reduction of the delivery lead-time from two hours to one hour, the total cost has more than tripled, e.g., from 100 to 374 with NOANT strategy. With delivery lead-time of four hours, the variation is less pronounced, e.g., for NOANT the total cost decreases with only 18 percent points.

Moreover, we observe that when the delivery lead-time is short the anticipatory shipment strategy produces higher benefits, as well as the improvement obtained with perfect information (HINDSIGHT) becomes less important. With enough capacity at the pickup point, the use of the anticipatory shipment strategy allows offering a one hour delivery service at the same costs as a two hour service without anticipation of customers’ demand (NOANT).

6.3.4 Impact of the distance between warehouse and pickup point

Next, we study the impact of the distance between the warehouse and the pickup location t^t . For a given customer lead-time, a larger travel distance means that there is less time to combine orders for the on-demand shipments.

We analyze two additional scenarios: one with a travel time of 15 minutes, and one with travel time of one hour. Figure 6 shows the changes in the total cost compared to the




Figure 5: Results varying the customer lead-time, $L = 1; 2; 4$ hours; average results across ten instances and 20 periods

base case of half an hour. We observe that reducing the travel time to 15 minutes does not have a significant impact on the total cost. However, with travel time of one hour, we see a substantial increase in the total cost occurs with all the analyzed strategies and under all the settings. The cost increase reaches up to 374 with NOANT. The longer transportation time provides the retailer less time to consolidate orders for joint shipment, leading to an increase in vehicles dispatches and associated costs. As expected, we see that the benefits of the anticipatory shipments increases with the distance to the pickup point.




Figure 6: Results varying the travel time, $t^t = 0.25; 0.5; 1.0$; average results across ten instances and 20 periods

6.3.5 Impact of the dispatch flexibility

Finally, to evaluate the potential of the anticipatory shipment strategy, we consider another policy that the retailer may use to efficiently handle the pickup points supply: increasing

the number of potential dispatch moments within the order period. A more flexible dispatch may potentially allow more efficient use of the vehicles and reduce costs.

We analyze two scenarios: (i) 16 dispatch moments, i.e., vehicles dispatches take place each thirty-minutes; (ii) 32 dispatch moments, i.e., vehicles dispatches take place each fifteen-minutes. We observe that when the number of dispatch moments doubles, there are no significant changes in terms of number of vehicles dispatches and thus transportation costs (Figure 7). With 32 dispatch moments, the retailer can further postpone the dispatching decision and consolidating more orders for joint shipment. Therefore, there is a reduction of the vehicles dispatches for all the analyzed strategies. However, differences are very small, e.g., one vehicle per period for NOANT, which suggest that the costs do not depend too much on the specific operational strategy introduced for vehicle dispatching. The delivery lead-time seems the constraining factor. For fast deliveries (e.g., within two hours), there is little room to consolidate multiple shipments. Other strategies should be used in order to provide fast and cost-effective deliveries. In this regard, we observe that the anticipatory shipment strategy brings significant benefits in all the analyzed settings.




Figure 7: Results varying the number of dispatch moments, $B = 8; 16; 32$; average results across ten instances and 20 periods

7 Conclusions

This paper studies a new anticipatory shipment strategy to supply pickup locations. The problem entails the decision of which products and quantities to stock at the pickup point in anticipation of future customers' orders so to minimize the expected total cost incurred in moving products from the warehouse to the pickup points.

We formulate the problem as a stochastic programming model using sample average approximation to capture the demand uncertainty. To evaluate the effects of different anticipatory strategies on the operational costs, we introduce a simple vehicle dispatching

policy.

Our experiments suggest that significant benefits are possible by applying the anticipatory shipment strategy as compared to an on-demand transportation system that makes frequent small dispatches. In particular, the anticipatory shipment strategy leads to cost saving and increased service level. The benefits increase with the availability storage space at the pickup location and the limited assortment offered to online customers. We also demonstrate that for fast deliveries the retailer can not focus on the drop size to efficiently manage the transportation system. Whereas, the benefits of the anticipatory shipment strategy increase with shorter customer lead-times and longer distances between the warehouse and the pickup location.

As this is a first attempt to address the anticipatory shipment strategy for pickup point supply, we see many potential areas for further research. First of all, the model can easily incorporate changing demand characteristics between periods. Future work can introduce more detailed models (e.g., two-dimensional packing problems, see Lodi et al. (2002)) and supply chain structures (e.g., multiple pickup points, limited stock at central warehouse, consolidating routing for different locations and transshipment). Other promising directions of future research are the incorporation of multi-item orders and the analysis of demand dependency.

Appendix A HINDSIGHT MODEL FORMULATION

In the HINDSIGHT model the demand in the look-ahead period is known when making decisions. Let \hat{d}_i^t denote the real demand of product i in period t , the model can be formulated as follows.

$$\text{minimize } \sum_{i \in \mathcal{P}} \left(c_i^a h_i^t + c_i^d w_i^t \right) \quad (25)$$

$$\text{s.t. } \sum_{i \in \mathcal{P}} (k_i Q_i^t) \leq K \quad (26)$$

$$h_i^t \geq Q_i^t - I_i^{t-1} \quad \forall i \in \mathcal{P} \quad (27)$$

$$h_i^t \geq I_i^{t-1} - Q_i^t \quad \forall i \in \mathcal{P} \quad (28)$$

$$w_i^t \geq \hat{d}_i^t - Q_i^t \quad \forall i \in \mathcal{P} \quad (29)$$

$$w_i^t \geq 0 \quad \forall i \in \mathcal{P} \quad (30)$$

$$Q_i^t \geq 0 \quad \forall i \in \mathcal{P} \quad (31)$$

$$Q_i^t \in \mathbb{Z} \quad \forall i \in \mathcal{P} \quad (32)$$

Appendix B Impact of system parameters: detailed results for the different shipment strategies

Table 5 and Table 6 illustrate the results of the different shipment strategies varying individual system parameters affecting the anticipatory shipment decisions: number of products offered online (N) and anticipatory shipment cost compared to the on-demand transportation system (c^a/c^d ratio), respectively. For each experiment, the costs are normalized by dividing all the costs by the cost of the NOANT strategy and multiplied by 100.

Strategy	N	K	Total cost	Units moved	Orders direct	Lead-time	Used capacity	Products stored	Leftover inventory
GRD	50	50%	51.1	357	61	0.46	100	64	34
ED	50	50%	49.6	357	63	0.44	100	74	37
SAA 1	50	50%	48.8	357	64	0.43	100	82	38
SAA 2	50	50%	48.3	357	65	0.42	100	79	55
HINDSIGHT	50	50%	41.1	355	74	0.33	100	74	0
GRD	50	100%	32.9	357	84	0.18	97	97	50
ED	50	100%	32.8	357	84	0.18	97	99	50
SAA 1	50	100%	30.4	359	87	0.15	100	100	74
SAA 2	50	100%	29.7	359	88	0.14	100	100	98
HINDSIGHT	50	100%	21.6	355	98	0.03	96	97	0
GRD	50	150%	32.0	357	85	0.17	66	100	51
ED	50	150%	32.0	357	85	0.17	66	100	51
SAA 1	50	150%	25.6	360	93	0.07	83	100	111
SAA 2	50	150%	22.5	361	97	0.03	99	100	202
HINDSIGHT	50	150%	20.0	355	100	0.00	67	99	0
GRD	200	50%	60.9	411	49	0.62	99	59	66
ED	200	50%	59.7	412	51	0.60	100	69	72
SAA 1	200	50%	58.9	412	52	0.58	100	74	79
SAA 2	200	50%	58.3	412	52	0.57	100	72	99
HINDSIGHT	200	50%	45.1	408	69	0.38	100	57	0
GRD	200	100%	45.1	413	69	0.37	100	89	96
ED	200	100%	44.8	413	69	0.37	100	94	99
SAA 1	200	100%	42.9	415	72	0.34	100	95	131
SAA 2	200	100%	42.0	415	73	0.32	100	93	163
HINDSIGHT	200	100%	24.8	408	94	0.07	100	78	0
GRD	200	150%	41.4	413	74	0.31	77	100	106
ED	200	150%	41.4	413	74	0.31	77	100	106
SAA 1	200	150%	33.0	418	84	0.18	100	100	208
SAA 2	200	150%	31.7	419	86	0.17	100	100	246
HINDSIGHT	200	150%	20.0	408	94	0.07	100	78	0

Table 5: Results varying the assortment size; average results across ten instances and 20 periods

Strategy	c^a/c^d	K	Total cost	Units moved	Orders direct	Lead-time	Used capacity	Products stored	Leftover inventory
GRD	0.4	50%	65.0	359	59	0.49	100	67	50
ED	0.4	50%	64.1	359	60	0.48	100	77	54
SAA 1	0.4	50%	64.6	358	59	0.48	100	87	41
SAA 2	0.4	50%	63.0	359	62	0.45	100	81	68
HINDSIGHT	0.4	50%	54.7	356	76	0.30	100	71	0
GRD	0.4	100%	53.0	360	79	0.25	95	98	71
ED	0.4	100%	53.0	360	79	0.25	95	100	72
SAA 1	0.4	100%	52.5	361	80	0.23	97	100	82
SAA 2	0.4	100%	50.0	361	84	0.19	100	99	125
HINDSIGHT	0.4	100%	40.6	356	99	0.01	94	93	0
GRD	0.4	150%	52.7	360	80	0.24	64	100	72
ED	0.4	150%	52.7	360	80	0.24	64	100	72
SAA 1	0.4	150%	52.1	361	81	0.23	67	100	85
SAA 2	0.4	150%	44.5	361	93	0.08	96	100	208
HINDSIGHT	0.4	150%	40.0	356	100	0.00	64	94	0
GRD	0.8	50%	88.8	359	59	0.49	100	67	50
ED	0.8	50%	88.5	359	60	0.48	100	77	54
SAA 1	0.8	50%	90.6	357	48	0.63	96	89	15
SAA 2	0.8	50%	88.7	357	57	0.51	100	85	38
HINDSIGHT	0.8	50%	84.9	356	76	0.30	100	71	0
GRD	0.8	100%	85.0	360	79	0.25	95	98	71
ED	0.8	100%	85.0	360	79	0.25	95	100	72
SAA 1	0.8	100%	90.5	357	48	0.62	50	90	15
SAA 2	0.8	100%	85.2	357	75	0.29	91	97	73
HINDSIGHT	0.8	100%	80.2	356	99	0.01	94	93	0
GRD	0.8	150%	84.9	360	80	0.24	64	100	72
ED	0.8	150%	84.9	360	80	0.24	64	100	72
SAA 1	0.8	150%	90.5	357	48	0.62	33	90	15
SAA 2	0.8	150%	85.1	357	76	0.29	62	97	75
HINDSIGHT	0.8	150%	80.0	356	100	0.00	64	94	0

Table 6: Results varying the anticipatory shipment cost; average results across ten instances and 20 periods

Appendix C Impact of dispatching strategy: detailed results for the different shipment strategies

Table 7 reports the results of the different shipment strategies varying individual parameters describing the operational strategy used for vehicle dispatching: lead-time (L), travel time (t^t) and dispatch moments (B). All the costs are normalized by setting the “model” total cost for NOANT in the base case (i.e., $L = 2$, $t^t = 0.5$ and $B = 8$) equal to 100.

Strategy	L	t^t	B	$K = 50\%$		$K = 100\%$		$K = 150\%$	
				“Real” total cost	Vehicles dispatches	“Real” total cost	Vehicles dispatches	“Real” total cost	Vehicles dispatches
GRD	2	0.5	8	87.7	15.9	76.5	12.7	75.3	12.4
ED	2	0.5	8	87.4	15.9	76.2	12.7	75.3	12.4
SAA 1	2	0.5	8	87.2	15.8	70.6	11.0	61.2	8.3
SAA 2	2	0.5	8	85.5	15.3	70.8	11.0	53.2	5.9
HINDSIGHT	2	0.5	8	70.5	11.1	37.2	1.7	31.4	0.0
GRD	1	0.5	8	197.6	47.2	130.1	28.0	127.2	27.2
ED	1	0.5	8	193.8	46.1	129.7	27.9	127.2	27.2
SAA 1	1	0.5	8	189.9	45.0	111.5	22.6	82.9	14.4
SAA 2	1	0.5	8	185.3	43.7	108.8	21.9	64.5	9.2
HINDSIGHT	1	0.5	8	138.4	30.4	39.6	2.3	31.4	0.0
GRD	4	0.5	8	70.1	11.0	63.2	9.0	62.4	8.7
ED	4	0.5	8	69.8	10.9	63.0	8.9	62.4	8.7
SAA 1	4	0.5	8	69.8	10.8	58.1	7.5	51.6	5.6
SAA 2	4	0.5	8	68.0	10.3	57.5	7.3	46.2	4.0
HINDSIGHT	4	0.5	8	52.9	6.1	34.6	0.9	31.4	0.0
GRD	2	0.25	8	87.6	15.9	75.3	12.4	74.2	12.1
ED	2	0.25	8	87.1	15.8	75.0	12.3	74.2	12.1
SAA 1	2	0.25	8	86.6	15.6	69.3	10.6	59.7	7.9
SAA 2	2	0.25	8	85.1	15.2	69.2	10.6	51.9	5.6
HINDSIGHT	2	0.25	8	70.1	11.0	36.7	1.5	31.4	0.0
GRD	2	1.0	8	197.6	47.2	130.1	28.0	127.2	27.2
ED	2	1.0	8	193.8	46.1	129.7	27.9	127.2	27.2
SAA 1	2	1.0	8	185.3	43.7	111.5	22.6	82.9	14.4
SAA 2	2	1.0	8	185.3	43.7	108.8	21.9	64.5	9.2
HINDSIGHT	2	1.0	8	138.4	30.4	39.6	2.3	31.4	0.0
GRD	2	0.5	16	87.2	15.8	75.5	12.5	74.5	12.2
ED	2	0.5	16	86.7	15.7	75.3	12.4	74.5	12.2
SAA 1	2	0.5	16	86.4	15.6	69.6	10.7	60.1	8.0
SAA 2	2	0.5	16	84.8	15.1	69.5	10.7	52.3	5.7
HINDSIGHT	2	0.5	16	69.0	10.7	36.7	1.5	31.4	0.0
GRD	2	0.5	32	83.6	14.8	73.0	11.7	72.0	11.5
ED	2	0.5	32	83.0	14.6	72.7	11.7	72.0	11.5
SAA 1	2	0.5	32	82.5	14.5	67.3	10.1	58.5	7.5
SAA 2	2	0.5	32	81.0	14.0	67.0	10.0	51.2	5.4
HINDSIGHT	2	0.5	32	65.6	9.7	36.4	1.4	31.4	0.0

Table 7: Results varying customer lead-time, travel time and number of dispatch moments; average results across ten instances and 20 periods

References

- Niels Agatz, Ann Campbell, Moritz Fleischmann, and Martin Savelsbergh. Time slot management in attended home delivery. *Transportation Science*, 45(3):435–449, 2011.
- Niels AH Agatz, Moritz Fleischmann, and Jo AEE Van Nunen. E-fulfillment and multi-channel distribution—a review. *European Journal of Operational Research*, 187(2):339–356, 2008.
- Greg Bensinger. Amazon wants to ship your package before you buy it. <http://blogs.wsj.com/digits/2014/01/17/amazon-wants-to-ship-your-package-before-you-buy-it>, 2014.
- Kenneth K Boyer, Andrea M Prud’homme, and Wenming Chung. The last mile challenge: evaluating the effects of customer density and delivery window patterns. *Journal of Business Logistics*, 30(1):185–201, 2009.
- Courtney Reagan. Putting buy online, pick up in store to the test: How the retailers stacked up. <https://www.cnbc.com/2017/11/22/putting-buy-online-pick-up-in-store-to-the-test-how-the-retailers-stacked-up.html>, 2017.
- Jeffrey Dastin. Amazon adds ‘instant pickup’ in u.s. brick-and-mortar push. <https://www.reuters.com/article/us-amazon-com-pickup/amazon-adds-instant-pickup-in-u-s-brick-and-mortar-push-idUSKCN1AV1EF>, 2017.
- Jan Fabian Ehmke and Ann Melissa Campbell. Customer acceptance mechanisms for home deliveries in metropolitan areas. *European Journal of Operational Research*, 233(1):193–207, 2014.
- Gary D Eppen, R Kipp Martin, and Linus Schrage. Or practice—a scenario approach to capacity planning. *Operations Research*, 37(4):517–527, 1989.
- Laureano F Escudero, Pasumarti V Kamesam, Alan J King, and Roger JB Wets. Production planning via scenario modelling. *Annals of Operations Research*, 43(6):309–335, 1993.
- Santiago Gallino and Antonio Moreno. Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information. *Management Science*, 60(6):1434–1451, 2014.
- Fei Gao and Xuanming Su. Omnichannel retail operations with buy-online-and-pick-up-in-store. *Management Science*, 2016.

- Anshuman Gupta and Costas D Maranas. Managing demand uncertainty in supply chain planning. *Computers & Chemical Engineering*, 27(8):1219–1227, 2003.
- Alexander Hübner, Andreas Holzapfel, and Heinrich Kuhn. Distribution systems in omnichannel retailing. *Business Research*, 9(2):255–296, 2016.
- Alexander H. Hübner and Heinrich Kuhn. Retail category management: State-of-the-art review of quantitative research and software applications in assortment and shelf space management. *Omega*, 40(2):199 – 209, 2012. ISSN 0305-0483.
- Eunhye Kim, Myeong-Cheol Park, and Jongtae Lee. Determinants of the intention to use buy-online, pickup in-store (bops): The moderating effects of situational factors and product type. *Telematics and Informatics*, 2017.
- A Gürhan Kök, Marshall L Fisher, and Ramnath Vaidyanathan. Assortment planning: Review of literature and industry practice. In *Retail supply chain management*, pages 99–153. Springer, 2008.
- Andrea Lodi, Silvano Martello, and Michele Monaci. Two-dimensional packing problems: A survey. *European journal of operational research*, 141(2):241–252, 2002.
- Stephen Mahar and P Daniel Wright. In-store pickup and returns for a dual channel retailer. *IEEE Transactions on Engineering Management*, 2017.
- Stephen Mahar, Peter A Salzarulo, and P Daniel Wright. Using online pickup site inclusion policies to manage demand in retail/e-tail organizations. *Computers & Operations Research*, 39(5):991–999, 2012.
- Josefa Mula, Raul Poler, JP Garcia-Sabater, and Francisco Cruz Lario. Models for production planning under uncertainty: A review. *International journal of production economics*, 103(1):271–285, 2006.
- David W Pentico. The assortment problem: A survey. *European Journal of Operational Research*, 190(2):295–309, 2008.
- Warren B Powell. *Approximate Dynamic Programming: Solving the curses of dimensionality*, volume 703. John Wiley & Sons, 2007.
- Uday S Rao, Jayashankar M Swaminathan, and Jun Zhang. Multi-product inventory planning with downward substitution, stochastic demand and setup costs. *IIE Transactions*, 36(1):59–71, 2004.
- Alexander Shapiro and Tito Homem-de Mello. A simulation-based approach to two-stage stochastic programming with recourse. *Mathematical Programming*, 81(3):301–325, 1998.

Alexander Shapiro and Andy Philpott. A tutorial on stochastic programming. *Manuscript*. Available at www2.isye.gatech.edu/ashapiro/publications.html, 17, 2007.

Alexander Shapiro, Darinka Dentcheva, et al. *Lectures on stochastic programming: modeling and theory*, volume 16. SIAM, 2014.

Xinan Yang, Arne K Strauss, Christine SM Currie, and Richard Eglese. Choice-based demand management and vehicle routing in e-fulfillment. *Transportation science*, 50(2): 473–488, 2014.