Basket data-driven approach for omnichannel demand forecasting

Haytham Omar[1,2] Walid Klibi[1,*] M. Zied Babai[1] Yves Ducq[2]

[1]The Centre of Excellence in Supply Chain (CESIT), Kedge Business School, Bordeaux, France

[2]University of Bordeaux, CNRS UMR 5218, IMS, Talence, France

**Corresponding author: walid.klibi@kedgebs.com*

Abstract

Omnichannel retailing has changed the purchasing behavior of customers in recent years, especially in online shopping, which has led to higher complexity in supply chain demand forecasting. Nowadays customers buy a variety of products in baskets that do not share similar characteristics and across various channels. In this article, we propose a new approach to forecasting demand, driven by data on customers shopping baskets. Drawing on network graph theory and market basket analysis, we identify three attributes for a product to promote the connection with other products sold together in a basket: *degree*, *strength*, and *support*. These attributes are used as predictor variables with an autoregressive integrated moving average model. We conduct an empirical investigation using sales and basket data related to an assortment of 24,000 products of a major cosmetics retailer in France selling through online and physical retail channels. We provide empirical evidence that using the shopping basket data improves the forecasting accuracy in omnichannel retailing. We also show that there is a considerable benefit from a joint forecasting of the online and store channels and a shared inventory between both channels.

Keywords: *Omnichannel retailing, demand forecasting, shopping basket, network analysis, inventory*

1 Introduction

1.1 General introduction and motivation

With the rise of omnichannel retailing in recent years, the shopping experience has changed drastically. Omnichannel retailing uses a variety of channels to interact with customers and fulfill their orders—a seamless shopping experience, enabling customers to order anytime from anywhere, in person or through digital devices, and to have their purchase be delivered at their preferred time and location (Strang 2013, Bell et al. 2014, Chopra 2018). As discussed by Brynjolfsson and Smith (2000), advanced technologies are blurring the distinctions between store and online retailing, creating an omnichannel environment. The shift to such an environment was accelerated by the COVID-19 pandemic (Knowles et al. 2020). From an operations management perspective, omnichannel retailing implies the integration of different demand streams, which raises several key issues, related to pricing optimization, assortment planning, demand forecasting, and inventory management. These issues are far from being resolved today.

Lapide (2016) and Byrne (2016) highlighted that omnichannel retailing adds complexity and that part of the change retailers would need to make involves their forecasting and planning processes, which work well for shelf sales but not necessarily for online order fulfillment. The common practice among multichannel retailers has been a siloed approach to planning and forecasting the different channels (Byrne 2016). Rooderkerk and Kök (2019) have argued that inventory replenishment traditionally relies on a separate forecasting demand of channels, whereas the demand cannot be truly captured because the retailer is unaware of the customer's journey that has led to such demand. For example, a customer could see a product in the showroom and order the product online. This often leads to an increase of the demand uncertainty, which makes the forecasting a more challenging task. In the same way, Byrne (2016) argued that among the most common trends in omnichannel distribution is the recognition of the need for a shared view of inventory across all channels.

Furthermore, the shopping basket notion has gained an increasing importance in omnichannel retailing. It has become a crucial vehicle to capture sales through continuous promotional periods, fast accessibility of a wide range of products, an abundance of social media advertising, and various bundling offers. For instance, the data panel of a major cosmetics retailer used in this work reveals that, out of 2.2 million online orders over the year, 50% of them occurred in baskets with two distinct products or more, and there were 2.4 ordered products on average per basket. Previously, Kumar and Rao (2006) argued that marketing departments should use basket composition data to target advertisement to consumers and find purchasing patterns in baskets. This raises the question of how data from multiple product baskets could be considered by operations managers to improve the accuracy of demand forecasting in omnichannel retailing. Surprisingly, little research has been devoted to demand forecasting within an omnichannel context, and to the best of our knowledge, no research has incorporated basket characteristics and behavior in demand forecasting. We endeavor to fill this gap in this article.

To analyze the purchasing patterns in shopping baskets, a mining approach, referred to as market basket analysis, is used. This approach enables to characterize the associations between products that are frequently sold together in a basket based on measures, such as the support, the lift, and the confidence (Berry and Linoff 2004, Ghoshal and Sarkar 2014). For instance, the support measures the percentage of sold baskets that contain a set of products sold together. Market basket analysis has been used by many retailers, such as Walmart, Amazon, Flipkart, to boost their marketing capabilities. Furthermore, graph theory is another field that is used to model the connections between items within a network and analyze the magnitude of their connectivity. It enables to calculate attributes such as the degree of connectivity, frequency, strength etc. Such approaches should be considered as a basis for operations managers to forecast demand given the data on shopping baskets, which constitutes a challenge that we address in this work.

1.2 Business context

Our research work is part of a project in collaboration with a large retailer in the cosmetics industry that sells a broad catalog of products through store and online channels. The project aims at the integration of in-store and online channels from demand forecasting and fulfillment perspectives. Figure 1 illustrates the retailer's downstream supply chain schema composed of a set of customers and their locations, a selected store, a retail warehouse, and a fulfillment center.

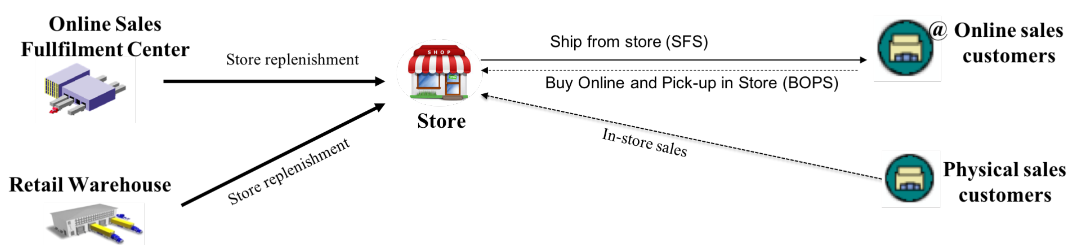

Figure 1: Omnichannel fulfillment

The retailer has been operating its in-store sales for a long time, relying on the deployment of a brick-and-mortar store network, replenished periodically from a central warehouse (denoted BMS). With the advent of e-commerce, the retailer started offering online sales and moved recently to the implementation of "buy online pick-up at store" (BOPS) and "ship-from store" (SFS) policies through the online channel to ensure a competitive response time that is supported via a dedicated fulfillment center. SFS takes advantage of the existing physical network by turning certain store locations to ship-from points for the online sales. Figure 1 depicts the replenishment-storage-shipment flows at the store under the omnichannel context, with dual product flow replenishment to the store (bold arrows), ship-to customer flows for SFS (a regular arrow), and customer moves to the store for BMS and for BOPS options (dotted arrows). Hence, a key question tackled in this work is, how can omnichannel demand forecasting be used to improve the inventory management and fulfillment performance in retailing?

1.3 Contributions

Our paper contributes to the operations management literature in three ways:

1. Building on market basket analysis and graph theory, we propose a novel forecasting approach for online and store sales that is driven by data on customers shopping baskets. To the best of our knowledge, our article is the first that considers the shopping basket in demand forecasting.
2. Through assessing a dataset of an assortment of 24,029 products in the online and store channels of a major cosmetics retailer, we empirically show the outperformance of the proposed forecasting approach compared to other benchmark forecasting methods commonly used in the retail context. We also empirically show the benefit

of joint forecasting of the online and store sales.

3. We provide empirical evidence that using joint forecasting and shared inventory in an omnichannel context leads to a reduction of inventory shortages. This is shown by comparing the inventory performance of three fulfillment scenarios where the sales of both channels are forecasted separately or jointly and the inventory is either dedicated to each channel or shared by both channels.

The remainder of the article is organized as follows. Section 2 is dedicated to a review of the literature by presenting the research background on omnichannel retailing and demand forecasting in the retail context. In Section 3, we present our basket data–driven forecasting approach. We empirically assess the performance of the different forecasting approaches in Section 4. In Section 5, we analyze the impact of omnichannel forecasting on the inventory and fulfillment performance. Finally, in Section 6, we present the conclusions of the article in addition to suggesting some avenues for further research.

2 Literature review

2.1 Omnichannel retailing

Omnichannel retailing is a nascent research area. In a recent review Melacini et al. (2018) underlined that key topics, such as the evolution of retail distribution networks, assortment planning over multiple channels, and the logistics role played by stores in the delivery process, are still under-represented in omnichannel retailing. Jasin et al. (2019) reported that the term omnichannel was introduced by Strang (2013) as a boundaryless retail experience to customers where a customer can research a product online or through a catalog, mobile applications, showrooms, or physical stores and then decide to buy though one of these channels. This is congruent with most of the definitions referring to omnichannel retailing as the use of a variety of channels to interact with customers and fulfill their orders (Bell et al. 2014, Chopra 2018). Chopra (2018) presented a framework of relative costs of the four omnichannel alternatives introduced in Bell et al. (2014). He highlighted, for instance, the contrast between the high inventory costs incurred by traditional retail and the high transportation costs for retailers using an online channel. Gallino and Moreno (2014) are among the first to cover the integration of online and offline channels in retail. They discussed important effects of the implementation of a BOPS channel, especially the cross-selling effect and the channel-shift effect. From a consumer point of view, Ailawadi and Farris (2017) underlined that there are clear benefits to an omnichannel distributional structure and discussed the heterogeneity in consumer reasons for buying online and the unknown path to purchasing across channels, which clearly is challenging from a forecasting perspective. This is congruent with recent discussions in Bell et al. (2018) and Bijmolt et al. (2019), who introduced the notion of the omnichannel customer journey around showrooming and web-rooming processes, and underlined the interplay between marketing and operations. The study of Ishfaq et al. (2016), relying on interviews with supply chain executives, underlined how omnichannel retailing could create separate demand streams and complicate the inventory positioning and allocation decisions for retailers. Goic and Olivares (2019) reported that in a data-driven omnichannel framework, the boundaries between online and offline data are disappearing and the methodologies to analyze these data are converging. Since about 2015, the number of studies on omnichannel retailing has risen significantly. First, a number of studies in the literature rely on empirical studies to explore this novel topic. Hübner et al. (2016) provided an exploratory study with 33 retailers that highlighted the advantages and challenges of centralized versus decentralized and of integrated versus dedicated distribution schema. The authors reported that optimizing the cross-channel processes in distribution centers and stores, and inventory integration and allocation, are among the most important areas for fulfilling distribution requirements. However, as underlined by the authors, only a few studies and quantitative models have modeled delivery schema when online and offline channels are integrated. Furthermore, the qualitative review proposed by Ishfaq et al. (2016) underlined the role of store-based retailing on the omnichannel fulfillment strategy, which would be helped by forward placement of inventory. In the same way, a study of German retailers conducted by Mena et al. (2016) to investigate specifically the transition to omnichannel logistics provides a framework on the level of integration at the warehouse with regards to the inventory, picking, and assortment processes. Gallino et al. (2016) empirically studied how the deployment of ship-to-store policy at a major US retailer generated sales dispersion and affected inventory-ordering models with an increase of cycle and safety inventories. In a similar way, Gallino and Moreno (2014) studied the importance of sharing reliable available inventory information with customers when channels are integrated. All these empirical investigations highlight the need to integrate store and online channel operations and to investigate cost-service trade-offs when channels are planned and operated jointly.

Furthermore, some quantitative models were proposed in the omnichannel distribution. More specifically, contributions on network design (Arslan et al. 2020, Guerrero-Lorente et al. 2020), dynamic pricing (Harsha et al. 2019), and inventory and fulfillment (Gao and Su 2016) were recently proposed. Recent work on designing omnichannel distribution networks underlined the challenge to minimize replenishment, delivery, and fulfillment costs when channels are integrated, and the benefit of using SFS channel (Arslan et al. 2020). Using a newsvendor setting with inventory decisions, Gao and Su (2016) studied store inventory optimization when a BOPS channel is available, and they concluded that BOPS tends to increase traffic in the store by providing inventory information and increased convenience to customers.The work of Harsha et al. (2019) developed a dynamic price optimization problem in the presence of cross-channel interactions and proposed two pricing policies that partitioned the inventory between channels. Gallino et al. (2019) and Jasin et al. (2019) discussed how relying on a warehouse to fulfill customer online orders can reduce operational costs and argued that adapting an SFS model can increase the service level perceived by customers.

Recently, Rooderkerk and Kök (2019) introduced the challenge of assortment planning, which is a strategic problem faced by retailers when opening a new channel or coordinating between channels in an omnichannel context. These authors discuss the notions of asymmetric integration in which part of the offering is common, and the other part is specific to each channel, and finally there is symmetric integration, where what is sold in the stores is exactly what is sold online. The authors, rightfully, appealed to revoking the classical predictive and prescriptive modeling approaches in an omnichannel context that all work by a separate demand stream and by location (stores or fulfillment centers). The authors reported that academic research has been largely silent on the challenges of forecasting the demand streams of the various types of omnichannel flows. To the best of our knowledge, there is no research work in the literature that considers forecasting and inventory planning to operate an omnichannel distribution network under online and store demand.

2.2 Demand forecasting in the retail context

A plethora of forecasting methods have been studied since the 1980s (Syntetos et al. 2016, Petropoulos et al. 2020). Such methods include simple extrapolative methods, such as exponential smoothing and moving averages (Gardner Jr 1985, Svetunkov and Petropoulos 2018), autoregressive integrated moving average (ARIMA)-type models (Gilbert 2005, Babai et al. 2013), machine learning methods (Zhang and Qi 2005, Punia et al. 2020), and judgmental methods (Petropoulos et al. 2016, 2018). A considerable amount of research work has been dedicated to analyzing and comparing the performance of such methods through empirical investigations and international forecasting competitions using supply chain data (Petropoulos and Makridakis 2020). The last forecasting competition (referred to as the M5 competition) was built on the case of a Walmart retail supply chain with more than 30,000 products (Makridakis et al. 2020). It showed the outperformance of machine learning methods compared to standard statistical forecasting methods. The M5 competition has also shown the ability to improve the accuracy of demand forecasting by considering exogenous/explanatory variables. The inclusion of exogeneous variables in the forecasting approach will be part of this research.

In the omnichannel retail context, Armstrong (2016) argued that forecasting strongly depends on the retailer's omnichannel strategy, noting that both forecasts of the online orders and sales at stores are important for fulfillment decisions. This issue can be addressed through hierarchical forecasting. Fildes et al. (2019) presented an overview on the practice and research of retail forecasting. They reported that, at the product level, the time horizon of forecasting changes across the supply chain from quarterly to monthly to daily depending on the location (e.g., monthly forecasting at the distribution center level and quarterly forecasting at the factory level). They highlighted the complexity of product forecasting due to different physical attributes such as color, size, and packaging. They also highlighted the importance of product mix across product categories and they recommended that category management starts with forecasting the category level with a judgmental approach, taking into account inter-category purchasing behavior or the product mix. Ma and Fildes (2021) presented a meta-learning framework based on deep convolutional neural networks to forecast retail sales. Based on weekly data of a grocery and drug chain related to a sample of 50 stores and 30 product categories, they showed the superior forecasting performance of the proposed meta-learner. They recommended, for forecasting retail sales, building a pool of base forecasters using both individual and pooled forecasting methods, to target finding the best combination of forecasts instead of the best individual method. Shang et al. (2020) showed that online retailers, especially in fashion, face a much bigger forecasting issue compared to physical retailers, which relates to product returns.

Furthermore, Fildes et al. (2019) argued that retail sales at the product and daily levels are characterized by a high degree of intermittence, that is, frequent zero sales. This demand characteristic attracted a considerable amount of research, and several forecasting methods were proposed to deal with demand intermittence. Syntetos and Boylan (2005) developed a commonly used framework to categorize intermittent demand into four categories—namely, smooth, lumpy, erratic, and intermittent. For a recent overview on the research dealing with intermittent demand, readers are referred to Nikolopoulos (2020). Croston (1972) established a benchmark method in theory and practice for intermittent demand forecasting. It is the only method that is available in major ERP-type solutions such as SAP and specialized forecasting software packages. The M5 competition is the first forecasting competition that has included the Croston method to deal with intermittent demand. However, it is worth noting that for a high degree of intermittence and lumpiness, the Croston method often leads to poor forecasting accuracy, and simple methods, such as single exponential smoothing (SES), may lead to a much better performance (Syntetos et al. 2015). The Croston method is included in this research work as a benchmark for the performance of the proposed approach.

Another important body of literature in retail forecasting deals with demand aggregation. In fact, demand in a retail supply chain can be aggregated at several levels, such as the product level (e.g., stock keeping unit, family, etc.), geographic location (e.g., store, region, etc.), or time (e.g., day, week, month, quarter, etc.). These different aggregation strategies can be either hierarchical and/or temporal and necessitate different forecasting methods (Rostami-Tabar et al. 2013). For a review on the advances related to supply chain forecasting by aggregation, interested readers are referred to Syntetos et al. (2016). It is worth pointing out that despite the richness of the forecasting literature in the retail context and the integration of different exogenous variables, the attributes of basket shopping practices have never been considered in terms of demand forecasting. This gap in the literature motivates our research to consider market basket analysis and linkages between sold products when forecasting sales in the omnichannel retailing context.

3 Basket data-driven forecasting approach

Based on the data of the shopping baskets, the proposed forecasting approach starts by developing the network of the sold products to identify their attributes, which are then used within an ARIMAX regression model (Box et al. 1994). This two-step forecasting approach is detailed in the following subsections.

3.1 Network development and attributes extraction

Graph theory is widely applied to conceptualize and analyze complex networks in supply chains (Gross and Yellen 2005, Kim et al. 2015, Dooley et al. 2019). In our context, we use graph theory to capture the products' relationships based on their sales in shopping baskets. More specifically, we consider a weighted graph to model the relationship between sold products. The nodes correspond to the products in the considered assortment, and when there is a pair of products sold together within at least one basket (we refer to this as a connection between the two products), the edge linking the two products is associated with the number of baskets in which the two products are sold together. This number is referred to as the frequency of this connection. We note that in our graph representation we assume that the relational attributes of products are the only associations and there is no causal relationship; thus, the network is undirected. With this in mind, we identify the three attributes of the product considered in the proposed forecasting approach: (1) *Degree* of the product, the number of connections (or arrows) with the other products of the assortment; (2) *Strength* of the product, the total frequency with all connected products; and (3) *Support* of the product, the strength of the product divided by the total number of sold baskets. In order to illustrate the graph and the characterization of the attributes, we consider an example of four products (A, B, C, D). In this example, we consider a total of 20 sold baskets. The product A is sold nine times. It is sold in two baskets with product B, in two baskets with product C, and in five baskets with product D. Figure 2a presents the graph of this example and Figure 2b shows the calculation of the three attributes for the four products. For example, the degree of product A is 3 because this product is sold with three products, whereas the strength is 9, which is the sum of the frequency of the arrows. The support is $9/20$, which is the strength of product A divided by the total number of baskets.

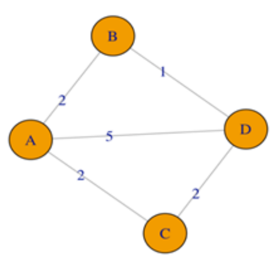

(a) Graph associated with the example

Product	Degree	Strength	Support
A	3	9	9/20
B	2	3	3/20
C	2	4	4/20
D	2	4	8/20

(b) Attributes calculation

Figure 2: Illustrative example for the calculation of the attributes

From graph-theoretic perspective, we can conceptualize the network of connections between products in baskets as follows. Let $t \in T$ be the set of historical sales periods (days) and $b \in B$ be the set of sold baskets, where B^t is the subset of baskets sold in period $t \in T$. Let $p \in P$ be the set of products, where P_b is the subset of products included in sold basket b. We define G as a undirected graph, denoted by $G = (N, E)$, composed by N nodes, representing the number of products ($N = P$) and E edges representing pairs of products, where $E : (p, p'), \forall p, p' \in P$. Accordingly, the graph G^t is defined for a given period t, and is composed by a set of subgraphs G^t_b per basket $b \in B^t$.

The network is produced on a daily basis with a rolling time, and the attributes are determined for each product and for each forecasting period t. Let $h(t)$ be the historic horizon used to build the network related to the estimation of period t, composed by $[t-1, ..., t-h]$, for the collection of baskets $B^{h(t)}$. Consequently, $G^{h(t)}$ defines the network graph associated to the set of periods in $h(t)$ such that $G^{h(t)} = \cup_{[t-1,...,t-h]} G^t$. For a given period t, based on the historic horizon $h(t)$ data, the three attributes are computed as follows.

We introduce $\lambda^t_{b(p,p')}$ as a counter that takes 1 if edge $(p,p') \in G^t_b$ (i.e., when both products are in the same basket), and 0 otherwise. It is then used to compute $\lambda^t_{(p,p')} = \sum_{t \in h(t)} \sum_{b \in B^{h(t)}} \lambda^t_{b(p,p')}, \forall p,p' \in P$, which assesses the number of occurrences of a given edge (p,p') (i.e. a pair of products), in the set of baskets, over the historic horizon $h(t)$. The *Strength* attribute in period t is based on the number of incident edges for product p in basket b, which is computed as $\mathit{Strength}_{pt} = \sum_{p' \in P} \lambda^t_{(p,p')}$. The *Support* attribute in period t of product p is given by $\mathit{Support}_{pt} = \mathit{Strength}_{pt}/|B^{h(t)}|$, where $|B^{h(t)}|$ is the cardinality of $B^{h(t)}$. The *Degree* attribute in period t of product p, is expressed as $\mathit{Degree}_{pt} = \sum_{p' \in P} \min(1, \lambda^t_{(p,p')})$. Note that the product quantity information is captured by the historical sales of the product and it is not represented in the graph.

3.2 Forecasting approach

The proposed forecasting approach can be summarized into two steps and is illustrated in Figure 3 :

- **Step 1.** Shopping baskets data during the historic horizon $h(t)$ are used to build the network, which is then considered to calculate the three attributes: *Strength*, *Support*, and *Degree*.
- **Step 2.** The three attributes are used as regressors along with historical sales data within an ARIMAX model.

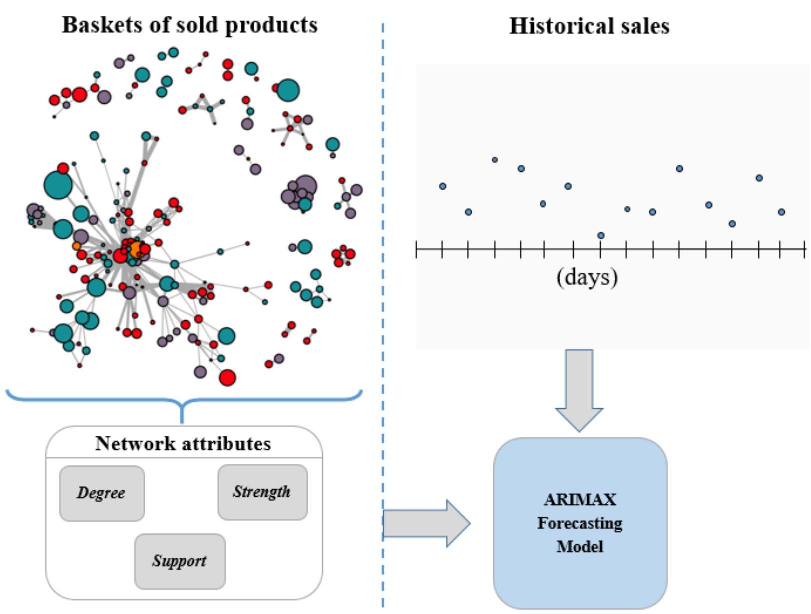

Figure 3: Illustration of the forecasting approach

Hence, at any time period t and for each product p, the forecast, denoted by y_{pt}, is calculated as:

$$y_{pt} = a\mathit{Strength}_{pt} + b\mathit{Support}_{pt} + c\mathit{Degree}_{pt} + n_{pt} \tag{1}$$

where $Strength_{pt}$, $Support_{pt}$, $Degree_{pt}$ are the *strength*, *support*, and *degree* regressors, respectively (calculated for a network historic horizon $h(t)$), and a, b, c are their respective coefficients. We also assume that n_{pt} is given by an ARIMA model.

4 Empirical investigation

4.1 Supply chain and data description

We empirically analyze the performance of the proposed forecasting approach by considering real data of a major French cosmetics retailer. We use an assortment of 24,029 products that are divided into six product families: perfume, care, makeup, bath, hair, and accessories. The retailer is globally deployed in more than 25 countries and is specifically well established in France with hundreds of stores in many cities and an online sales platform. As illustrated in Figure 1, the retailer operated a retail warehouse and an online fulfillment platform, both located in France. The retailer data used for the purpose of this research relates to sales in 2018. The data contain the product reference, the description of the product, the sold quantity, the selling date, the delivery date, and the invoice (order) number. Based on this information, two datasets are built. A first dataset, referred here to as Panel A, contains the online sales in all France for the year 2018 with about 2.2 million orders (single- and multiple-item baskets). Panel A is used to make the network analysis and to show the accuracy of the proposed forecasting approach. The second dataset, referred to here as Panel B, represents the orders of the online and store channels in one of the largest region in France in 2018. Panel B is used to assess the omnichannel forecasting accuracy and inventory fulfillment performance.

4.2 Network and basket analysis

We start by presenting the empirical results of the network analysis and the exploratory analysis of the sold products and baskets.

4.2.1 Characteristics of the network

We empirically analyze the obtained networks through some characteristics, namely, the density, the assortativity, and the average path length. The density of the network is the ratio between the number of connections in the network and all the possible connections. This characteristic enables us to show if the products in the network have a high connectivity (i.e., degree). The empirical results show that the density of the network is equal to 0.3%, which proves that the network is characterized by a low density. However, from our analysis, we find that a few products have a very high degree of connectivity, whereas most of the products have a low degree. This shows that the network is heavily skewed in terms of degree distribution to some products with high connectivity. The assortativity of the network enables testing whether the baskets are composed of products from the same product family. A network is highly assortative if its correlation coefficient is close to 1, non-assortative if it is 0, and disassortative if it has a coefficient less than 0 (Newman 2002, Noldus and Van Mieghem 2015). The empirical results show that the overall network assortativity is 0.23, which means that 77 % of the retail baskets are cross-family orders. Finally, we investigate if the omnichannel retail network follows a small-world network. To do so, we calculate the average path length of the network. Note that small-world networks are often characterized by a short average path length. The empirical results show a high average path length of 2.85 compared with the case of randomized networks in which the average path length is 2.8 (for the latter, we used 50 randomly generated networks with the same number of nodes and density). This empirical result shows that this retail network is not considered a small-world network.

4.2.2 Exploratory analysis of products and sold baskets

We start by analyzing the sales patterns and the degree of intermittence. To do so, we analyze the data by using the demand classification proposed by Syntetos and Boylan (2005), referred to hereafter as SBC. Recall that in the SBC scheme, the demand is classified based on the average demand interval (ADI) and squared coefficient of variation of demand sizes (CVZ^2) with the cut-off values of ADI = 1.32 and $CVZ^2 = 0.49$. Four categories are identified: smooth (ADI < 1.32 and $CVZ^2 < 0.49$), lumpy (ADI > 1.32 and $CVZ^2 > 0.49$), erratic (ADI <1.32 and CVZ $^2 > 0.49$), and intermittent (ADI > 1.32 and CVZ $^2 < 0.49$). We report in Figure 4 the percentage of products within each category.

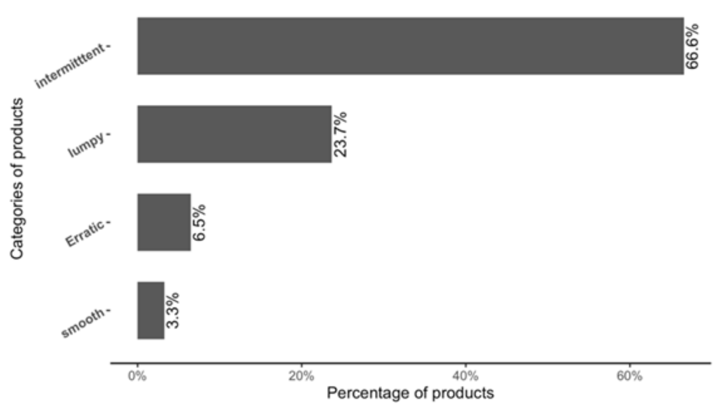

Figure 4: Percentage of products in each category in the online channel

The results in Figure 4 show that the majority of products (66.6%) are characterized by an intermittent demand pattern. This confirms that the online sales considered in our case are in line with the case presented by Fildes et al. (2019) in the physical retail space, where 52.2% of the sold products had an intermittent demand pattern. Figure 4 shows as well that products with lumpy and erratic demand patterns represent more than 30%, whereas those with a smooth demand constitute only 3.3% of the assortment. A focus on the average demand intervals of the products in each category shows that a product with an intermittent demand is sold on average once every 14 days, whereas a product with a smooth demand is sold on average every 1.1 days.

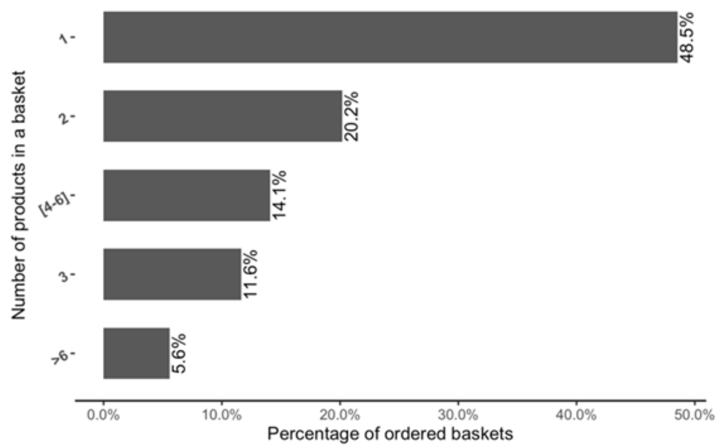

Figure 5: Percentage of singe-product baskets and multiple-product Baskets

We now analyze the composition of orders made for the considered assortment of products. We present in Figure 5 the percentage of orders with the number of products in these orders. The results in Figure 5 show that 48.5% of the total baskets have only one product, whereas 51.5% of the baskets have two products or more. More precisely, the average number of products in a basket is 2.4; a third quartile equals 3, and a maximum equals 90. This concurs with our initial inference that there could be singularity of buying for certain products, and effectively these products will have low degree, strength of frequency flow, and strength of support, whereas many products will have higher degree, which means higher connectivity to other products as well as a higher strength of frequency flow.

We now check whether a product is usually sold in a single-product basket or in a multiple-product basket. Our initial intuition is that the new forecasting approach would work well for a product that is sold at least once with another product in the training horizon. Obviously, a product that is sold all the time as a single product in a basket would be equally well forecasted using the proposed ARIMAX or a basic ARIMA approach. To conduct this analysis, we report in Figure 6 the percentage of products and the percentage of times (across all orders) each product is sold in multiple-product baskets.

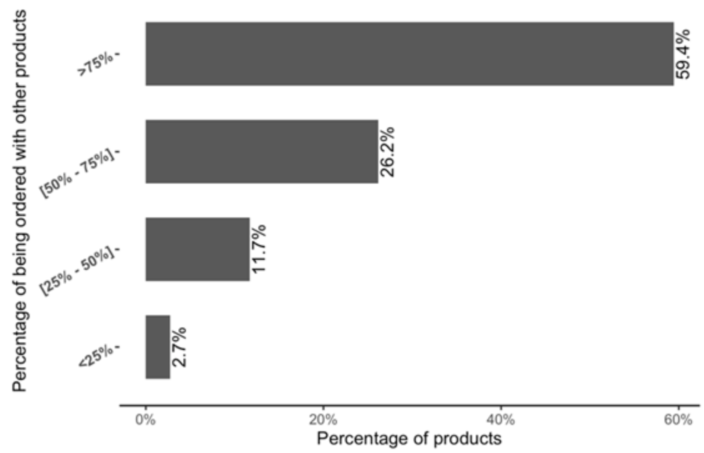

Figure 6: The number of times (percentage wise) each product was sold with other products

The results in Figure 6 show that only 2.7% of the products are sold individually, whereas 59.4% of the products are sold more than 75% of the time with other products. This confirms that the majority of products are sold with other products in a basket more than once, which also shows the importance of degree and strength attributes when forecasting demand of these products. Finally, we report in Figure 7 the degree distribution of the four categories of products. The results show that the majority of products in the intermittent category have a daily degree between 0 and 2. This asserts that the products characterized by intermittent demand have low connectivity compared with products in the erratic, lumpy, and smooth categories. Moreover, because strength is concerned with the frequency of common ordering of connected products, we report the correlation between the degree and strength of products in the four categories.

The results in Figure 8 show that the correlation between degree and strength is 0.97, indicating that the frequency of ordering intermittent demand products with other products is low compared to lumpy, erratic, and smooth demands.

4.2.3 Attributes modeling, testing, and validation

We tested the significance of the three attributes to predict sales using the generalized linear Lasso regression model (Tibshirani 1996). To do so, we tested the significance of the different combinations of the attributes (i.e., one, two, or three attributes) to find the best combination. The test is conducted by using 10 random samples of Panel A. Recall that the Lasso regression model penalizes unimportant regressors by shrinking some of the coefficients of the regression variables to zero if they are insignificant. We report in Table 1 the results of the two best linear regression models; the first model shows the degree and strength attributes and the second one shows the three attributes: degree, strength, and support. Table 1 reports the coefficients of the two models, the root mean squared error (RMSE), the Akaike information criterion (AIC), and the Bayesian information criterion (BIC). The results of the tests show that both models are significant. Note that all the coefficients of

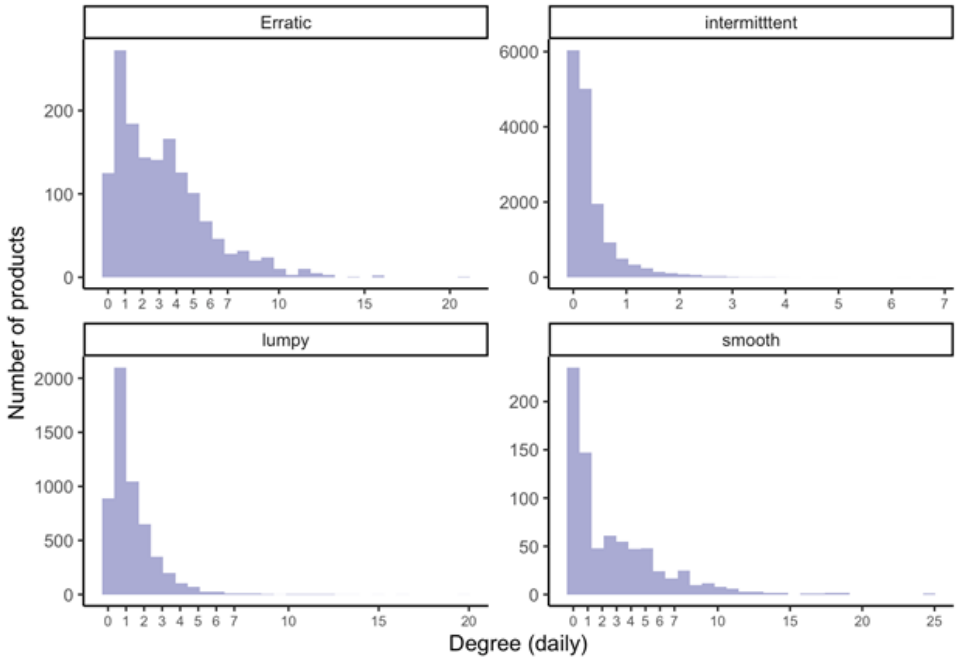

Figure 7: Degree distribution of products per category of products

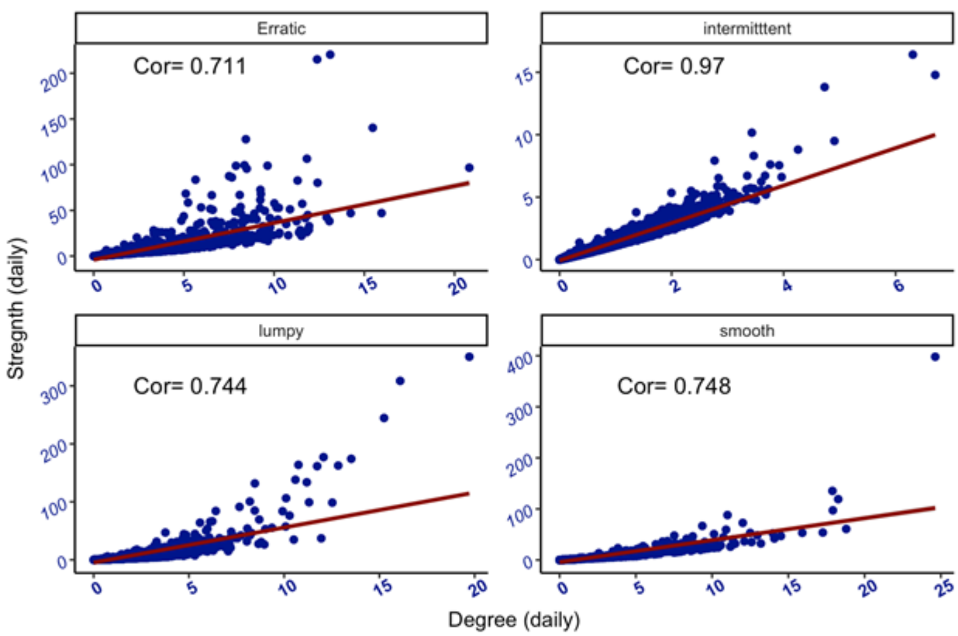

Figure 8: Pearson correlation test between degree and strength

Model	Adj. R^2	Intercept	Degree	Strength	Support	RMSE	AIC	BIC
Degree & Strength	0.235	0.071	0.030	0.019		6.674	100029610	100029668
Three attributes	0.235	0.077	0.034	0.017	148.2	6.670	100021142	100021214

Table 1: Two best linear regression models using the attributes

the two regression models have a p-value less than 0.05. However, Table 1 shows that using the three attributes of degree, strength, and support, provides the best fit. Hence, these three attributes are considered to be regressors in the proposed ARIMAX regression model. Because the proposed forecasting approach requires a network horizon to determine the exogenous variables, we tested several network horizons, that is, a network horizon with 7 days, 21 days, 28 days, and 56 days. The preliminary results show that the 7 days horizon overall provides the best performance, and will be further used in the paper.

4.3 Forecasting accuracy analysis

The forecasting accuracy of the proposed approach is first evaluated using the Panel A data. We use three forecasting methods as a benchmark to show the performance of the proposed forecasting methods, namely, the ARIMA model, SES, and the Croston method. As argued in the literature review, these methods are commonly used in the retail context when dealing with products with intermittent demand patterns. The performance of the four methods is evaluated using one-step-ahead forecasts. We consider a within sample from January 1 to October 19, 2018, and we use the out-of sample from October 20 to December 31, 2018, to evaluate the performance of the forecasting methods. The smoothing coefficients of SES and Croston are selected such that they optimize the mean squared error (MSE) over the within sample. We evaluate the accuracy of the forecasting methods by means of four measures that are commonly in the intermittent demand context: the mean absolute error (MAE), the root mean square error (RMSE), the mean absolute scaled error (MASE), and the symmetric mean absolute percentage error (sMAPE). The MAE measure is used to show the magnitude of the error without having the issue of a bias sign. RMSE is an indication of the standard deviation of forecast errors. Because MAE and RMSE do suffer from a scale dependence issue, we use MASE and sMAPE as scale-free measures of the forecasting accuracy. Both measures represent a scaled version of the absolute errors. In the former, the scaling of the errors is based on the in-sample MAE of the naïve forecasting method (Hyndman and Koehler 2006, Steinker et al. 2017, Babai et al. 2020). In the latter, the scaling of the error in each period in the out-of sample is based on an average between the actual demand and its forecast.

4.3.1 Forecasting accuracy results of the online channel

Table 2 shows the empirical results using the data of Panel A. For each forecasting method and accuracy measure, we report two figures: the upper one represents the average error (across all products) and the one below represents the percentage of products where the forecasting method is the best for the accuracy measure. The lower forecasting errors are highlighted by a bold font.

The results in Table 2 show that the proposed ARIMAX forecasting method overall leads to the highest accuracy when compared to the three forecasting benchmarks. The results also show that the ARIMAX method is the best-performing method for the highest number of products within the assortment. The second-best-performing method is the ARIMA method, whereas the method that is associated with the lowest forecasting

		ARIMAX	ARIMA	SES	Croston
MAE	Average error	**0.474**	0.484	0.587	1.351
	% of products	54.66%	37.26%	3.56%	4.53%
RMSE	Average error	**0.875**	0.885	1.010	1.762
	% of products	46.43%	30.59%	8.84%	14.14%
MASE	Average error	3.304	3.320	**2.502**	4.437
	% of products	54.62%	37.29%	3.56%	4.53%
sMAPE	Average error	**178.028**	178.204	178.664	179.357
	% of products	31.84%	12.22%	17.10%	38.84%

Table 2: Forecasting accuracy results

Intermittent					**Erratic**				
	ARIMAX	ARIMA	SES	Croston		ARIMAX	ARIMA	SES	Croston
MAE	**0.194**	0.206	0.236	0.297	MAE	**2.467**	2.502	2.941	8.171
RMSE	0.342	**0.340**	0.358	0.413	RMSE	**4.783**	4.923	5.500	10.812
MASE	**1.135**	1.187	1.218	1.822	MASE	27.027	27.673	**17.656**	32.693
sMAPE	189.986	**189.038**	189.563	189.892	sMAPE	**115.645**	123.746	122.716	128.725
Lumpy					**Smooth**				
	ARIMAX	ARIMA	SES	Croston		ARIMAX	ARIMA	SES	Croston
MAE	**0.733**	0.734	0.917	2.169	MAE	1.088	**1.075**	1.138	5.013
RMSE	**1.404**	1.420	1.645	2.852	RMSE	1.481	**1.456**	1.558	5.388
MASE	4.021	4.037	**2.453**	3.928	MASE	**1.224**	1.236	1.418	11.427
sMAPE	**162.838**	167.287	165.836	166.451	sMAPE	**137.151**	144.420	142.361	146.587

Table 3: Forecasting accuracy results per category of products

accuracy is Croston method. Note that the outperformance of ARIMAX compared to ARIMA is relatively small (ranging from 0.1% under sMAPE and 2% under MAE), whereas it can be much higher, going up to 65%, when compared to the Croston method. Note that under the MASE measure, SES shows a lower average error but there is a low number of where SES is the best (3.56%). This means that this outperformance is due to some products, whereas the MASE values are much higher than those of ARIMAX. To better understand the performance of ARIMAX and the relatively lower accuracy of SES and Croston, despite their forecast performance evidence in the context of intermittent demand, we analyze the forecasting accuracy results per category according the SBC classification scheme. The results for the four categories are reported in Table 3 and the lower forecast errors are highlighted by a bold font.

The results per category show that the proposed ARIMAX method is overall still associated with a high forecasting accuracy compared with the other benchmarks for the lumpy, erratic, and smooth categories. However, the relative error reduction by using ARIMAX is much higher under the smooth, erratic, and lumpy categories, going up to 6% when compared to ARIMA and 78% when compared to Croston. The outperformance of the proposed approach under these three categories can be explained by the fact that under non-intermittent demands, demand occurrences are higher and the approach better captures the connection between the sold products, which improves its forecasting accuracy. This in line with the findings of Section 4.2, where we have shown that the connectivity of products is higher under the smooth, erratic, and lumpy categories. Moreover, as expected, the results reveal that for the intermittent demand category, the relative performance of SES and Croston improves. These two benchmarks, especially the Croston method, have been developed for the specific case of intermittent demand with stationary demand intervals and demand sizes, which means that when the

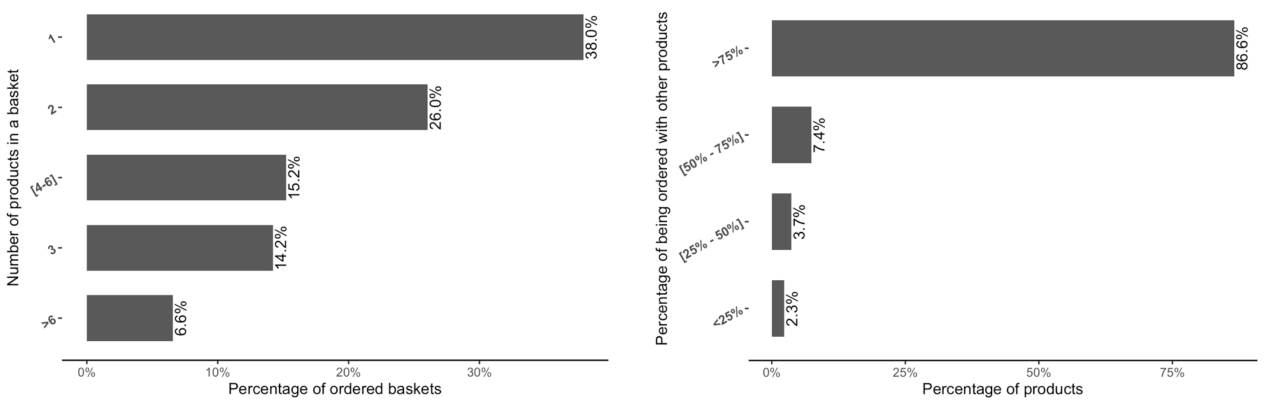

Figure 9: Baskets characteristics for the omnichannel case

degree of lumpiness or erraticness increases, the method loses its performance. We also note that SES may lead in some cases of non-smooth demands to a good forecasting accuracy. This is expected because SES is known to deal better with the non-stationarity of the data in the intermittent demand context as shown in the literature by Babai et al. (2014).

4.3.2 Omnichannel forecasting results

For the purpose of omnichannel forecasting, we consider the data of Panel B, which relate to the online and store channels. We first analyze the characteristics of the baskets in the omnichannel case. The results are shown in Figure 9. In the left we report the percentage of singe-product baskets and multiple-product Baskets and in the right we report the number of times (percentage wise) each product was sold with other products. The results in Figure 9 show that 62% of the total baskets have two products or more. This exceeds the percentage in the online case (reported in Figure 5), which further endorses the importance of the shopping basket behavior in the omnichannel case. Moreover the results in Figure 9 show that 86.6% of the products are sold more than 75% of the time with other products. This percentage is higher than the online channel case (reported in Figure 6). This further confirms that the majority of products are sold with other products in a basket more than once, which is more pronounced in the omnichannel case.

Next, a joint sales forecast is made after aggregating the daily online and store sales. We analyze the forecasting accuracy of the four forecasting methods using the online sales channel, the sales of the store channel, and the omnichannel aggregated sales. We report in Table 4 the average sMAPE results for the three cases. Results are shown using only sMAPE because this is a relative error measure that enables a fair comparison among the three cases. The results in Table 4 show that by using the omnichannel case, the sMAPE error decreases, which means that the forecasting accuracy improves by forecasting the sales based on an aggregation of the two channels. It is important to note that the improvement of the performance with the joint forecasting of both sales channels is realized under the four forecasting methods. Table 4 also shows that although our proposed forecasting method is not the most accurate for the store sales channel, it remains the best in the omnichannel sales case.

	Forecast accuracy (sMAPE)			
	ARIMAX	ARIMA	SES	Croston
Online sales	197.883	198.106	198.641	198.704
Store sales	190.244	189.753	190.655	189.870
Omnichannel sales	**188.611**	**189.479**	**188.968**	**188.921**

Table 4: Omnichannel forecasting accuracy results

4.4 Omnichannel inventory and fulfillment performance

An important question is whether the proposed forecasting approach can be used to improve the inventory and fulfillment performance in the omnichannel network. To answer this question, we analyze the fulfillment performance at the store with a separate versus a joint forecasting method and a dedicated versus a shared inventory. We do so by contrasting the performances of the three alternative scenarios on Panel B data. Recall that, as illustrated in Figure 1, the business context considers that the store has a dual source replenishment (a central warehouse and a fulfillment center) and that inventories kept at store are for online and store sales. Hence, the three scenarios are as follows:

- **Scenario 1**: Dedicated inventory and separate forecasting scenario (referred to as the DISF scenario). In this baseline scenario, the store inventory is split by sales channel and the forecast is separated by channel.
- **Scenario 2**: Shared inventory and separate forecast scenario (referred to as the SISF scenario). In this scenario, the two sales channels share the same inventory at store. However, the replenishment system depends separately on the forecast of demand of the store channel and the online channel.
- **Scenario 3**: Shared inventory and joint forecasting scenario (referred to as the SIJF scenario). In this scenario, the store and online channels share the same inventory at the store. However, the forecast is made jointly using the aggregated data from both channels.

To evaluate the inventory performance for the three scenarios, we first conduct a goodness-of-fit analysis to test the fit of different distributions with the demand in the different channels. We consider five distributions: normal, Poisson, gamma, negative binomial distribution (NBD), and stuttering Poisson (StPoisson). These distributions are commonly considered in modeling intermittent and non-intermittent demand data (Syntetos et al. 2013, Turrini and Meissner 2019). We show in Table 5 the results obtained with a Kolmogorov-Smirnov test at a 5% significance level.

		Poisson	NBD	Normal	Gamma	StPoisson
Online sales	Fit	99.23%	100.00%	94.04%	99.65%	99.98%
	No fit	0.77%	0.00%	5.96%	0.35%	0.02%
Store sales	Fit	98.54%	99.80%	93.49%	97.77%	99.78%
	No fit	1.30%	0.04%	6.35%	2.07%	0.07%
Omnichannel sales	Fit	97.83%	99.87%	90.97%	96.84%	99.81%
	No fit	2.17%	0.13%	9.03%	3.16%	0.19%

Table 5: Empirical goodness-of-fit results

The results show the strong empirical fit of the NBD followed by the StPoisson distribution, with normal being the distribution associated with the lowest fit. These results are expected knowing that the demand data are characterized by a high degree of intermittence. Note though that the fit of normal increases when the joint sales are considered. Therefore, for the purpose of the inventory performance investigation, NBD is selected to model the lead-time demand.

We then assess the inventory performance of the three scenarios by measuring their resulting stock on hand and the inventory backordering. To do so, we consider a periodic order-up-to-level inventory control policy, where the order-up-to level is calculated to satisfy a target cycle service level (CSL, the fraction of replenishment periods in which all of the demand can be met from stock). Recall that under this policy, each day, the inventory position is reviewed and an order is triggered if it is found to be below the order-up-to level to raise it up to the order-up-to level. The order arrives after a lead time, and any demand that is not satisfied from stock on hand is backordered. The lead time is fixed to 2 days. For the purpose of the analysis, we fix three target CSLs—90%, 95%, and 99%—and for target CSL and each scenario, we measure the average inventory holding volumes and the average backordering volumes (averages calculated over the evaluation period and across all products). Note that the last 73 days (from October 20, 2018, to December 31, 2018) are used to evaluate the performance. In order to have a fair comparison of the three scenarios, we plot efficiency curves of inventory holding versus backordering volumes. The efficiency curves are reported in Figure 10. In these efficiency curves the method that has its curve closer to the x-axis for a certain inventory holding volume implies a lower backordering volume and thus a higher efficiency. We do not report the achieved CSLs because for the DISF scenario the CSL cannot be reported and compared to that of the two other scenarios.

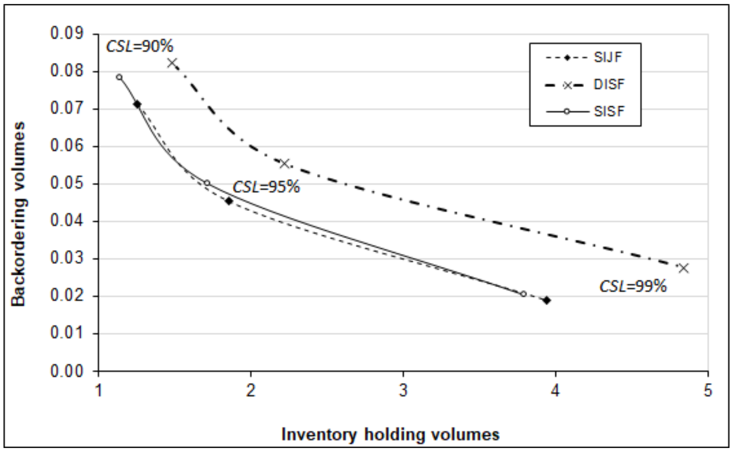

Figure 10: Efficiency curves: holding Versus backordering volumes of the fulfillment scenarios

The results in Figure 10 show that the shared inventory and joint forecast scenario leads to the lowest inventory backordering for fixed inventory holding volumes, which indicates that this scenario is associated with the highest inventory efficiency. Hence, this empirical investigation of the inventory performance reveals the benefit of using the joint demand to make forecasts and to use a shared inventory for both the online and retail channels.

5 Conclusion and Future Work

There is an agreement in the omnichannel retail literature that forecasting omnichannel sales is a challenging task, and research devoted to deal with this issue is lacking. This study has proposed a new approach to forecasting demand in an omnichannel retail context using data on customers shopping baskets. The forecasting approach builds on market basket analysis and graph theory to identify attributes of the products, which can be used with a time-series forecasting model. We have conducted an empirical analysis of the proposed forecasting approach and other benchmark forecasting methods commonly used in the retail context. A dataset of an assortment of more than 24,000 products in the online and store channels from a global leading cosmetics retailer was used for this purpose. Our study is the first to make an empirical analysis of both online and store sales and the first to use information on shopping baskets to forecast demand.

We have characterized the empirical behavior of more than 2 million online orders. Our investigation has shown that more than 95% of the sold products are characterized by an intermittent demand pattern, with 30% of them having a high lumpiness. By analyzing the composition of orders, we find that more than 50% of products are sold in baskets with more than two products per basket. Our findings reveal as well that the network of omnichannel retail sales is characterized by low density and a high average path length, and thus cannot be considered to be a small-world network. However, the connectivity (i.e., degree) distribution is heavily skewed to a few products in the network. Through a regression analysis, our study also reveals the importance of considering the degree, strength, and support of the network as good regressors for forecasting purposes.

The empirical assessment of the performance of the forecasting methods shows that one of the most popular methods for intermittent demand forecasting (the Croston method) leads to a poor forecasting accuracy in the omnichannel context. This underperformance is accentuated in the product category with lumpy demand patterns. Our proposed forecasting approach, which uses the degree, strength, and support attributes within an ARIMAX model, tackles this issue by improving the forecasting accuracy for both intermittent and lumpy demand patterns. These empirical findings highlight the value of considering attributes of the linkages between sold products in baskets for omnichannel forecasting purposes. The empirical study also shows that a joint forecasting based on the sales of both online and store sales leads to a higher forecasting accuracy and inventory performance. Such an omnichannel forecasting approach is recommended to consolidate inventories at stores, likely leading to a considerable reduction of inventory shortages. These findings enable an omnichannel network designer to gain valuable insights on how to deploy inventories in a set of fulfillment centers and on how stores could play a major role in an efficient urban fulfillment. The findings also provide interesting insights on omnichannel assortment planning based on the data-driven network analysis of the baskets.

An interesting avenue for further research would be to deepen the analysis using market basket analysis and graph theory to identify other attributes than those that can be considered for forecasting purposes in the omnichannel context. Such attributes can also be considered within machine learning approaches, which have been recommended in the literature for their superior performance when using data of exogenous variables. An extension of this work would be to investigate alternative inventory rationing and fulfillment scenarios that integrate joint forecasting. Finally, another avenue for future research could be to study the impact of product assortments at stores on the customer channel choice using the linkages between sold product baskets.

References

Ailawadi, K. L. and Farris, P. W. (2017). Managing multi-and omni-channel distribution: metrics and research directions. *Journal of retailing*, 93(1):120–135.

Armstrong, C. (2016). Omnichannel retailing and demand planning. *The Journal of Business Forecasting*, 35(4):10.

Arslan, A. N., Klibi, W., and Montreuil, B. (2020). Distribution network deployment for omnichannel retailing. *European Journal of Operational Research*.

Babai, M., Tsadiras, A., and Papadopoulos, C. (2020). On the empirical performance of some new neural network methods for forecasting intermittent demand. *IMA Journal of Management Mathematics*, 31(3):281–305.

Babai, M. Z., Ali, M. M., Boylan, J. E., and Syntetos, A. A. (2013). Forecasting and inventory performance in a two-stage supply chain with arima (0, 1, 1) demand: Theory and empirical analysis. *International Journal of Production Economics*, 143(2):463–471.

Babai, M. Z., Syntetos, A., and Teunter, R. (2014). Intermittent demand forecasting: An empirical study on accuracy and the risk of obsolescence. *International Journal of Production Economics*, 157:212–219.

Bell, D. R., Gallino, S., and Moreno, A. (2014). How to win in an omnichannel world. *MIT Sloan Management Review*, 56(1):45.

Bell, D. R., Gallino, S., and Moreno, A. (2018). Offline showrooms in omnichannel retail: Demand and operational benefits. *Management Science*, 64(4):1629–1651.

Berry, M. J. and Linoff, G. S. (2004). *Data mining techniques: for marketing, sales, and customer relationship management*. John Wiley & Sons.

Bijmolt, T. H., Broekhuis, M., De Leeuw, S., Hirche, C., Rooderkerk, R. P., Sousa, R., and Zhu, S. X. (2019). Challenges at the marketing–operations interface in omni-channel retail environments. *Journal of Business Research*.

Box, G. E., Jenkins, G. M., and Reinsel, G. C. (1994). Time series analysis, forecasting and control. englewood clifs.

Brynjolfsson, E. and Smith, M. D. (2000). Frictionless commerce? A comparison of Internet and conventional retailers. *Management Science*, 46(4):563–585.

Byrne, T. M. M. (2016). Omnichannel: How will it impact retail forecasting and planning processes? *The Journal of Business Forecasting*, 35(4):4.

Chopra, S. (2018). The evolution of omni-channel retailing and its impact on supply chains. *Transportation Research Procedia*, 30:4–13.

Croston, J. D. (1972). Forecasting and stock control for intermittent demands. *Journal of the Operational Research Society*, 23(3):289–303.

Dooley, K. J., Pathak, S. D., Kull, T. J., Wu, Z., Johnson, J., and Rabinovich, E. (2019). Process network modularity, commonality, and greenhouse gas emissions. *Journal of Operations Management*, 65(2):93–113.

Fildes, R., Ma, S., and Kolassa, S. (2019). Retail forecasting: Research and practice. *International Journal of Forecasting*.

Gallino, S. and Moreno, A. (2014). Integration of online and offline channels in retail: The impact of sharing reliable inventory availability information. *Management Science*, 60(6):1434–1451.

Gallino, S., Moreno, A., and Rooderkerk, R. P. (2019). Omnichannel fulfillment dilemmas: Customer preferences and manager perceptions. *Available at SSRN 3399664*.

Gallino, S., Moreno, A., and Stamatopoulos, I. (2016). Channel integration, sales dispersion, and inventory management. *Management Science*, 63(9):2813–2831.

Gao, F. and Su, X. (2016). Omnichannel retail operations with buy-online-and-pick-up-in-store. *Management Science*, 63(8):2478–2492.

Gardner Jr, E. S. (1985). Exponential smoothing: The state of the art. *Journal of forecasting*, 4(1):1–28.

Ghoshal, A. and Sarkar, S. (2014). Association rules for recommendations with multiple items. *INFORMS Journal on Computing*, 26(3):433–448.

Gilbert, K. (2005). An arima supply chain model. *Management Science*, 51(2):305–310.

Goic, M. and Olivares, M. (2019). Omnichannel analytics. In *Operations in an Omnichannel World*, pages 115–150. Springer.

Gross, J. L. and Yellen, J. (2005). *Graph theory and its applications*. CRC press.

Guerrero-Lorente, J., Gabor, A. F., and Ponce-Cueto, E. (2020). Omnichannel logistics network design with integrated customer preference for deliveries and returns. *Computers & Industrial Engineering*, 144:106433.

Harsha, P., Subramanian, S., and Uichanco, J. (2019). Dynamic pricing of omnichannel inventories. *Manufacturing & Service Operations Management*, 21(1):47–65.

Hübner, A., Holzapfel, A., and Kuhn, H. (2016). Distribution systems in omni-channel retailing. *Business Research*, 9(2):255–296.

Hyndman, R. J. and Koehler, A. B. (2006). Another look at measures of forecast accuracy. *International journal of forecasting*, 22(4):679–688.

Ishfaq, R., Defee, C. C., Gibson, B. J., and Raja, U. (2016). Realignment of the physical distribution process in omni-channel fulfillment. *International Journal of Physical Distribution & Logistics Management*.

Jasin, S., Sinha, A., and Uichanco, J. (2019). Omnichannel operations: Challenges, opportunities, and models. In *Operations in an Omnichannel World*, pages 15–34. Springer.

Kim, Y., Chen, Y.-S., and Linderman, K. (2015). Supply network disruption and resilience: A network structural perspective. *Journal of operations Management*, 33:43–59.

Knowles, J., Ettenson, R., Lynch, P., and Dollens, J. (2020). Growth opportunities for brands during the covid-19 crisis. *MIT Sloan Management Review*, 61(4):2–6.

Kumar, N. and Rao, R. (2006). Research note—using basket composition data for intelligent supermarket pricing. *Marketing Science*, 25(2):188–199.

Lapide, L. (2016). Retail omnichannel needs better forecasting & planning. *The Journal of Business Forecasting*, 35(3):12.

Ma, S. and Fildes, R. (2021). Retail sales forecasting with meta-learning. *European Journal of Operational Research*, 288(1):111–128.

Makridakis, S., Spiliotis, E., and Assimakopoulos, V. (2020). The m5 accuracy competition: Results, findings and conclusions. *Int J Forecast*.

Melacini, M., Perotti, S., Rasini, M., and Tappia, E. (2018). E-fulfilment and distribution in omni-channel retailing: a systematic literature review. *International Journal of Physical Distribution & Logistics Management*, 48(4):391–414.

Mena, C., Bourlakis, M., Hübner, A., Wollenburg, J., and Holzapfel, A. (2016). Retail logistics in the transition from multi-channel to omni-channel. *International Journal of Physical Distribution & Logistics Management*.

Newman, M. E. (2002). Assortative mixing in networks. *Physical review letters*, 89(20):208701.

Nikolopoulos, K. (2020). We need to talk about intermittent demand forecasting. *European Journal of Operational Research*.

Noldus, R. and Van Mieghem, P. (2015). Assortativity in complex networks. *Journal of Complex Networks*, 3(4):507–542.

Petropoulos, F., Apiletti, D., Assimakopoulos, V., et al. (2020). Forecasting: theory and practice. *arXiv preprint arXiv:2012.03854*.

Petropoulos, F., Fildes, R., and Goodwin, P. (2016). Do 'big losses' in judgmental adjustments to statistical forecasts affect experts' behaviour? *European Journal of Operational Research*, 249(3):842–852.

Petropoulos, F., Kourentzes, N., Nikolopoulos, K., and Siemsen, E. (2018). Judgmental selection of forecasting models. *Journal of Operations Management*, 60:34–46.

Petropoulos, F. and Makridakis, S. (2020). The m4 competition: Bigger. stronger. better. *International Journal of Forecasting*, 36(1):3–6.

Punia, S., Nikolopoulos, K., Singh, S. P., Madaan, J. K., and Litsiou, K. (2020). Deep learning with long short-term memory networks and random forests for demand forecasting in multi-channel retail. *International journal of production research*, 58(16):4964–4979.

Rooderkerk, R. P. and Kök, A. G. (2019). Omnichannel assortment planning. In *Operations in an Omnichannel World*, pages 51–86. Springer.

Rostami-Tabar, B., Babai, M. Z., Syntetos, A., and Ducq, Y. (2013). Demand forecasting by temporal aggregation. *Naval Research Logistics (NRL)*, 60(6):479–498.

Shang, G., McKie, E. C., Ferguson, M. E., and Galbreth, M. R. (2020). Using transactions data to improve consumer returns forecasting. *Journal of Operations Management*, 66(3):326–348.

Steinker, S., Hoberg, K., and Thonemann, U. W. (2017). The value of weather information for e-commerce operations. *Production and Operations Management*, 26(10):1854–1874.

Strang, R. (2013). Retail without boundaries. *Supply Chain Management Review*, 17(6).

Svetunkov, I. and Petropoulos, F. (2018). Old dog, new tricks: a modelling view of simple moving averages. *International Journal of Production Research*, 56(18):6034–6047.

Syntetos, A., Lengu, D., and Babai, M. Z. (2013). A note on the demand distributions of spare parts. *International Journal of Production Research*, 51(21):6356–6358.

Syntetos, A. A., Babai, M. Z., and Gardner Jr, E. S. (2015). Forecasting intermittent inventory demands: simple parametric methods vs. bootstrapping. *Journal of Business Research*, 68(8):1746–1752.

Syntetos, A. A., Babai, Z., Boylan, J. E., Kolassa, S., and Nikolopoulos, K. (2016). Supply chain forecasting: Theory, practice, their gap and the future. *European Journal of Operational Research*, 252(1):1–26.

Syntetos, A. A. and Boylan, J. E. (2005). The accuracy of intermittent demand estimates. *International Journal of forecasting*, 21(2):303–314.

Tibshirani, R. (1996). Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288.

Turrini, L. and Meissner, J. (2019). Spare parts inventory management: New evidence from distribution fitting. *European Journal of Operational Research*, 273(1):118–130.

Zhang, G. P. and Qi, M. (2005). Neural network forecasting for seasonal and trend time series. *European journal of operational research*, 160(2):501–514.