3033

REPRINT

Matias Schuster Puga, Stefan Minner and Jean-Sébastien Tancrez

Two-stage supply chain design with safety stock placement decisions

International Journal of Production Economics, 209, 183-193, 2019

CORE CENTER FOR OPERATIONS RESEARCH AND ECONOMETRICS

CORE
Voie du Roman Pays 34, L1.03.01
B-1348 Louvain-la-Neuve
Tel (32 10) 47 43 04
Email:lidam-library@uclouvain.be
https://uclouvain.be/en/research-institutes/lidam/core/reprints.html

International Journal of Production Economics 209 (2019) 183–193

Contents lists available at ScienceDirect

International Journal of Production Economics

journal homepage: www.elsevier.com/locate/ijpe

Two-stage supply chain design with safety stock placement decisions

Matías Schuster Puga [a,b], Stefan Minner [b], Jean-Sébastien Tancrez [a,*]

[a] Louvain School of Management, Center for Operations Research and Econometrics (CORE), Université catholique de Louvain, Belgium
[b] TUM School of Management, Technische Universität München, Germany

ARTICLE INFO

Keywords:
Location-inventory
Safety stock placement
Risk pooling
Service times
Customer classes

ABSTRACT

In this paper, we propose a supply chain design model that integrates facility location with safety stock placement and delivery strategy decisions, to reflect their interdependence and ultimately improve the resulting supply chain design. Safety stock placement decisions are modelled using the guaranteed-service approach. We also consider two customer classes that differ with respect to their delivery time preferences. The resulting non-linear model is formulated as a conic quadratic mixed-integer program that can be solved to optimality. We show that the model can be extended to include demand correlation and stochastic lead times. The computational experiments capture the trade-offs between location and safety stock placement decisions, between demand variability pooling and proximity to retailers, and between lead times and service times. We show that, when the distribution centres have short lead times to the retailers, distribution centres tend to hold safety stocks, and the retailers offer express deliveries. Conversely, when lead times are long, only the retailers hold safety stocks and they do not offer express deliveries.

1. Introduction

Supply chain management provides an effective way to increase competitive advantage by reducing operational costs and improving customer satisfaction. A critical decision that impacts the performance of a company in the long term is the strategic design of its supply chain. Facilities are a key driver of supply chain performance in terms of responsiveness and efficiency. On the one hand, companies can gain economies of scale when centralizing products in one location, which leads to an increase in efficiency. On the other hand, this efficiency increase comes at the expense of responsiveness, as many of the customers may be located far from the facilities (Chopra and Meindl (2016)). Moreover, in today's volatile economy, demand uncertainty is a major concern for companies, and it forces them to hold safety stocks. Therefore, the optimal placement of safety stocks is an important decision when designing the supply chain as it has a direct impact on the service quality. The decisions on where to locate facilities and where to place the safety stocks in these facilities are interdependent by nature. In this context, the integration of these decisions leads to a more efficient design of the supply chain.

In this paper, we propose a mathematical model that integrates safety stock placement decisions into the supply chain design problem. We consider two-stage supply chains composed of retailers, distribution centres (DCs) and one central plant. Our model simultaneously determines: (*i*) the number and location of opened DCs, (*ii*) the allocation of flows between DCs and retailers, (*iii*) for which retailers the opened DCs hold safety stocks, and (*iv*) the delivery time options at retailers. Moreover, we show that the model can be extended to include demand correlation and stochastic lead times. We use the guaranteed-service approach to model the safety stock placement decisions between DCs and retailers. The guaranteed-service model assumes that opened DCs serve each retailer with a deterministic service time, which is defined as the time elapsed between an order placed at the DC and the release of that order to the retailer. Several trade-offs are captured by integrating safety stock placement decisions into the strategic supply chain design problem. On the one hand, when safety stocks are only held at retailers, the lead time is assumed to be the delivery time between the central plant and the retailers. Consequently, the company can benefit from the lead time pooling effect in certain situations (e.g., for long lead times between the DCs and the retailers, and short order lead times between the central plant and the DCs). On the other hand, when safety stocks are held at DCs, the company can reduce safety stock costs by pooling demand variability at DCs (see Hua and Willems (2016)). Therefore, interesting managerial insights arise from the balance between lead time and

* Corresponding author. Université catholique de Louvain, Chaussée de Binche, 151, 7000 Mons, Belgium.
E-mail addresses: matias.schuster@uclouvain.be (M. Schuster Puga), stefan.minner@tum.de (S. Minner), js.tancrez@uclouvain.be (J.-S. Tancrez).

https://doi.org/10.1016/j.ijpe.2018.05.018
Received 18 October 2016; Received in revised form 9 April 2018; Accepted 16 May 2018
Available online 19 May 2018
0925-5273/

demand variability pooling strategies, which are influenced by the number of opened DCs, the safety stock placement and the delivery strategies.

Moreover, we differentiate customers at retailers regarding their delivery time preference. Customers like to be able to choose between different delivery options with different price points (Ray and Jewkes (2004), Teimoury et al., (2011)). The degree of flexibility at which a delivery system can meet particular customer needs can influence the decision of customers to place orders, and thus can be regarded as an important factor in attracting and retaining customers (Novich (1990), Gunasekaran et al., (2004)). In this paper, we include this differentiation and consider two classes of customers. Regular customers at retailers are served with a given service time. To improve its responsiveness, the company also offers express deliveries for those customers that request to be served immediately and accept to pay an extra fee. The income originated by the extra fee of express deliveries will balance the safety stock cost increase caused by the improvement of the delivery time. Interesting managerial insights are revealed when we consider different customer classes in a supply chain design model. The trade-offs arise from the selection of express and regular deliveries, which are dependant on the lead times, the express deliveries demand, the service times and the fees. In particular, when considering short lead times between DCs and the retailers, express deliveries are offered more often as safety stock costs are limited.

The main contributions of this paper are summarized as follows. Firstly, we propose a location-inventory model with demand uncertainty that integrates safety stock placement and delivery time decisions. Secondly, we extend the guaranteed-service model by considering different service times between DCs and the retailers, and different classes of customers. Thirdly, we reformulate our model as a conic quadratic mixed-integer program (CQMIP), which can be solved efficiently using standard optimization software packages such as CPLEX, as proposed in Atamtürk et al., (2012). Finally, to show the benefits of integrating facility location and safety stock placement decisions, we conduct a large set of computational experiments. We derive interesting managerial insights related to the impact of service times and lead times on the placement of safety stocks and on the delivery options selected.

The paper is organized as follows. The next section reviews the existing literature on supply chain design and strategic safety stock placement. In Section 3, we define our problem in detail and give the CQMIP formulation, extended with demand correlation and stochastic lead times. In Section 4, we show the results of the computational experiments and discuss managerial insights. We conclude in Section 5 and outline ideas for future work.

2. Literature review

Our research is grounded in the literature on the location-inventory and on the safety stock placement stream. In this section, we give a detailed review of each field, and on the papers that combine both streams, which are the closest to our research.

For a detailed review on the facility location stream, we refer the reader to Owen and Daskin (1998), ReVelle and Eiselt (2005), Klose and Drexl (2005), and Bravo and Vidal (2013). During the last 15 years, the basic location models have been generalized by integrating tactical and operational decisions. We refer the reader to Beamon (1998), Shen (2007), Melo et al. (2009), Mula et al. (2010), and Farahani et al. (2015) for a review on different integrated models such as location-inventory, location-routing or inventory-routing models.

We focus on the integrated location-inventory problem, and consider safety stocks for capturing demand uncertainty. The number of stocking locations has long been proven to directly affect the total amount of inventory held at these locations. Smykay (1973) notes that the safety stock held at a centralized location is equal to the aggregate safety stock held at multiple locations divided by the square root of the number of locations (Evers (1995)). Eppen (1979) compares a decentralized system where a separate inventory is held at each retailer with a centralized one where inventory is held at a DC. His model demonstrates that, due to the risk pooling effect, the consolidation of demand can reduce holding and penalty costs. In accordance with this feature, Shen et al. (2003) propose a location-inventory model with risk pooling. Their model determines the location of DCs, the allocation of flows, how often to reorder and which level of safety stock to maintain at DCs in order to minimize costs (i.e., opening, transportation and inventory). Computational results show that centralizing stocks at DCs leads to inventory reductions via risk pooling. Daskin et al. (2002) solve the same location-inventory problem by using Lagrangian relaxation, for instances with up to 150 customers. In a similar fashion, Miranda and Garrido (2004) combine a capacitated facility location problem with safety stocks. The resulting non-linear mixed-integer problem is solved through Lagrangian relaxation and sub-gradient methods. Kumar and Tiwari (2013) show the benefits of risk pooling in a three-layer supply chain and compare the savings from a centralized system with those from a decentralized one. Sourirajan et al. (2009) solve a location-inventory model using genetic algorithms and show the trade-off between lead times and inventory risk pooling. Park et al. (2010) formulate the location-inventory problem as a nonlinear integer programming model, for which a two-phase heuristic solution algorithm is derived based on the Lagrangian relaxation approach. Atamtürk et al. (2012) study the location-inventory problem for different cases such as capacitated facilities, correlated demands and multi-commodities. Their main contribution relies on the formulation of these models as conic quadratic mixed-integer programs (CQMIPs). Shahabi et al. (2013) also reformulate the location-inventory model as a compact CQMIP for a four-echelon supply chain where warehouses and hubs are located. Shahabi et al. (2014) extend the CQMIP formulation to include demand correlation in a three-echelon supply chain. Schuster Puga and Tancrez (2017) formulate the location-inventory problem as a non-linear continuous model that is solved using a heuristic algorithm, allowing the design of large supply chains. In their model, safety stocks are held at the DCs and at the retailers.

The strategic safety stock placement problem has been widely studied in the last 20 years. We refer the reader to Minner (2000) for a detailed presentation of the fundamental principles of safety stock planning. Graves and Willems (2003) give a review of the two general models for locating safety stocks: the stochastic service model and the guaranteed-service model. The two approaches differ in their demand treatment and service time characterization. In the stochastic model, the retailers may suffer stochastic delays when the DC does not have enough items available for delivery. Therefore, the service time is stochastic. In contrast, in the guaranteed-service model, the service time is a deterministic decision variable. The guaranteed-service approach can simplify the safety stock placement problem, because it can be formulated as a deterministic mathematical program. In a recent survey, Eruguz et al. (2016) give a detailed review on guaranteed-service models for multi-echelon inventory optimization.

In this paper, we use the guaranteed-service approach to model the safety stock placement decisions. The guaranteed-service model was first introduced by Simpson (1958), who models a serial chain production as a sequence of in-process inventories. The objective of the model is to determine the inventory service time at each production stage in order to satisfy the demand of its downstream stage at minimum inventory costs. Under this serial network structure, the paper shows that the minimum cost occurs when the inventory at each stage is either full or empty. Since then, the basic model has been extended in different ways. Inderfurth (1995) considers a multistage system where demands are correlated both between time and between products. The author shows that correlation between time has a greater impact on costs than product correlation, and that neglecting demand correlation may lead to incorrect levels of safety stocks and shortages. Inderfurth and Minner (1998) formulate the safety stock placement problem without internal delays and show that the optimal policy depends on the structure of the multi-stage system and the service measure used. Hua and Willems (2016) use the guaranteed-service model in a two-stage serial supply chain to give analytical insights on the optimal safety stock placement policy with regard to supply chain costs and lead times. Instead of optimizing the placement of safety stocks, Chen and Li (2015) use the guaranteed-service approach to optimize the (R,Q) inventory policy for a serial supply chain with Poisson demand and fixed order costs. The model is solved by an iterative procedure using dynamic programming. Inderfurth (1992) shows the benefits of lead time reduction in a multi-stage system with stochastic demand and stochastic lead times. Humair et al. (2013) also consider stochastic lead times and provide closed-form equations for the expected safety stock level at each stage of the supply chain. Osman and Demirli (2012) propose a decentralized and a centralized safety stock placement model considering uncertainty in demand and in lead times. The decentralized model is solved using a non-linear commercial solver, whereas a Benders decomposition technique is developed to solve the centralized model. Grahl et al. (2016) solve the safety stock placement problem with service time differentiation using metaheuristics in general acyclic networks. Li and Jiang (2012) also model the safety stock placement problem in general acyclic networks. The problem is solved by integrating constraint programming into a genetic algorithm. Magnanti et al. (2006) solve the safety stock placement problem using a commercial solver that generates flow cover cuts by considering redundant constraints.

Only few papers integrate safety stock placement decisions and supply chain design, as we aim to do in this paper. Graves and Willems (2005) incorporate multi-echelon stock placement into a network optimization problem for spanning-tree type supply chains with stationary demand. The supply chain design problem consists of the selection of the best option at each stage in terms of lead times and direct costs. They solve this problem using a dynamic programming heuristic algorithm. Following a similar approach, Funaki (2012) formulates a multi-echelon stock placement problem as a modification of the guaranteed-service model with due-date based and non-stationary demand. The optimization procedure is also based on dynamic programming. Moncayo-Martínez and Zhang (2013) formulate the problem of supply chain design and safety stock placement as a Max-Min ant system. Their solution approach is based on ant colony optimization to minimize a bi-objective function composed by the total cost (safety stock and production costs) and the lead times, to ensure product deliveries without delays. These papers relate supply chain design and safety stock placement but assume that the facilities already exist and are fixed. The combination of facility location and safety stock placement decisions is hardly ever included in the literature. You and Grossmann (2010) consider fixed opening costs of DCs, and use expected lead times to quantify the responsiveness of a three-echelon supply chain. The authors solve the model by using a spatial decomposition algorithm based on the integration of Lagrangian relaxation and piecewise linear approximation for cases up to 15 plants, 100 potential DC locations and 200 customers in around 10 h. This work is extended by Rodriguez et al. (2014) with the consideration of lost sales for redesigning the supply chain of spare part delivery. They solve the redesign problem for electric motor supply chains.

In the following, we differentiate our research in relation to these papers on supply chain network design, and in particular to Atamtürk et al. (2012). Unlike these papers, our paper integrates a multi-stage inventory system into the strategic supply chain network design, and considers the guaranteed-service approach for modeling safety stock placement decisions. Moreover, our model integrates features that are not included in the cited papers such as service times at retailers and customer classes depending on their delivery time preference. Both features have been integrated into the model and extensively analyzed using numerical experiments. We also differentiate our research from the papers that integrate safety stock placement decisions and supply chain design. In particular, we differentiate our contribution from You and Grossmann (2010) and from Rodriguez et al. (2014), which, to the best of our knowledge, are the only papers that explicitly integrate location decisions into the safety stock placement problem. Our paper combines features that are not included in these two papers, such as two customer classes with different delivery time preferences, demand correlation and stochastic lead times. Moreover, inspired by the all-or-nothing property of the guaranteed-service approach (proven in Simpson (1958)), we avoid the use of continuous service time variables between DCs and retailers for modeling the safety stock placement decisions. Instead, we include binary variables to explicitly decide for which retailers the opened DCs hold safety stocks, which is equivalent to provide service time differentiation between DCs and retailers. This is an important feature as the cited articles do not differentiate service times between DCs and retailers, and as it allows us to formulate our location-inventory model with safety stock placement decisions as a CQMIP. Compared to You and Grossmann (2010) and to Rodriguez et al. (2014), which provide a non-linear mixed-integer program that limits the computational efficiency, our CQMIP can be solved to optimality in significantly less computational time. Consequently, their numerical analysis is limited to few instances (the majority of them not solved to optimality), while we provide extensive computational results and multiple managerial insights.

3. Location-inventory model with safety stock placement decisions

3.1. Model presentation

In this section, we explain our model in detail. We consider two-stage supply chains with single sourcing and a single product. We assume that the vehicles have a given capacity, and consider full truckloads between DCs and retailers. The central plant has unlimited production capacity. The locations of the retailers are known and fixed. We suppose that the retailers' demands are uncorrelated and normally distributed, with a known mean and variance. The lead times between potential DCs and retailers are assumed to be known and dependent on the distance. At DCs, the order lead times to the central plant are also known.

In our model, two different delivery strategies can be applied by retailers, depending on whether express services are offered or not. When the first delivery strategy is followed, the retailers offer the possibility to have their product immediately (the product is available in stock) to customers. Therefore, a given percentage () of the customers at retailer r choose this option and is served immediately, with express deliveries, while the rest of the customers ($1 - P_r^{ED}$) is served with regular deliveries (with a given service time, S_r). The customers served with express deliveries accept to pay an additional fee (T_r) for the service. When the second delivery strategy is followed, the retailer r does not propose express deliveries, so all the customers are served with regular deliveries (with the given service time S_r).

Furthermore, our model includes cycle as well as safety inventory. On the one hand, cycle inventory is computed at the retailers and at the DCs. The cycle inventory at the retailers depends on the vehicles' capacity between DCs and retailers. At the DCs, the cycle inventory is determined by the size of the incoming order from the central plant, which is found by balancing ordering and inventory costs (transportation costs are supposed to be included in the ordering costs). Note that our model could be interpreted as locating factories (without a central plant) instead of locating DCs. In this case, the order cost would correspond to a production setup cost, and the order size to a batch size.

On the other hand, safety stocks are included to absorb demand uncertainty. The safety stock placement decisions are modelled using the guaranteed-service approach. As already mentioned in the previous section, Simpson (1958) shows that, under certain assumptions, safety stock inventories are either full or empty for two-stage supply chains. This result leads to two different safety stock placement strategies, depending on whether safety stocks are held at the DCs or not. It also allows us to avoid the use of continuous service time variables between the DCs and the retailers, as instead we can use binary variables that directly determine whether one strategy or the other is used for each retailer.

The net lead time, defined as the time during which safety stocks are required to cover demand variation, differs for the two strategies. The net lead times, at DCs and at retailers, are given for each strategy in Table 1, and are explained in the following. In safety stock placement strategy one, the DCs do not hold safety stocks and thus the net lead times at DCs are not relevant. The net lead time at a retailer r corresponds to the time from the central plant to the retailer, which is given by the order lead time from the central plant to the linked DC d, LT_d, plus the lead time between DC d and retailer r, LT_{dr}, and minus the service time at the retailer r, S_r. In safety stock placement strategy two, the DCs and the retailers hold safety stocks. At a DC d, the net lead time corresponds to the time from the central plant to DC d, LT_d. At a retailer r, the net lead time is given by the lead time between the linked DC d and the retailer r, LT_{dr}, minus the service time at the retailer r, S_r.

3.2. Notations

In the following, we introduce the notations used throughout this paper.

Indices and sets:

$r \in R = \{1, \ldots, n_r\}$ for retailers,
$d \in D = \{1, \ldots, n_d\}$ for potential DC locations,
$p \in P = \{1, 2\}$ for safety stock placement strategies (1, for safety stocks only at the retailer; 2, for safety stocks at the retailer and at the DC).

Parameters:

O_{dr}: unit transportation cost from DC d to retailer r, in €/item,
C_{dr}: capacity of vehicles from DC d to retailer r, in items/vehicle,
k_{dr}: fixed cost at retailer r of placing an order to the DC d, in €/order,
K_d: fixed cost at DC d of placing an order to the central plant, in €/order,
F_d: fixed cost of opening a DC d, in €/period,
h_r: unit inventory holding cost at retailer r, in €/(item · period),
H_d: unit inventory holding cost at DC d, in €/(item · period),
T_r: additional/expected fee per unit applied to the express customers of retailer r, in €/item,
z_α: standard normal deviate associated with service level α,
LT_d: order lead time between the central plant and the DC d, in periods,
LT_{dr}: lead time between the DC d and the retailer r, in periods,
S_r: Service time at retailer r, in periods,
μ_r: mean demand at retailer r, in items/period,
P_r^{ED}: percentage of customers at retailer r that request express deliveries ED, in %,
σ_r^{ED}: standard deviation of express deliveries' demand ED at retailer r, in items/period,
σ_r^{RD}: standard deviation of regular deliveries' demand RD at retailer r, in items/period,
σ_r: standard deviation of total demand at retailer r, in items/period,
CV_r: coefficient of variation of demand at retailer r.

Table 1
The net lead time at DCs and at retailers for each safety stock placement strategy.

Strategy	Net Lead Time	
	DC	Retailer
1	–	$LT_d + LT_{dr} - S_r$
2	LT_d	$LT_{dr} - S_r$

Decision variables:

y_d: equals one when DC d is opened, zero otherwise,
x_{dr}^p: equals one when DC d serves retailer r using safety stock placement strategy p, zero otherwise,
Q_d: size of order to the central plant at DC d, in items,
q_{dr}: shipment size from DC d to retailer r, in items/vehicle,
w_r^{ED}: equals one when retailer r proposes express deliveries, zero otherwise,
λ_d: total expected product flow through DC d, in items/period,
$\overline{\sigma}_d$: standard deviation of the product flow through DC d, in items/period.

3.3. Cost function

In this section, we describe the various costs included in our model. Our mathematical model minimizes transportation, facility opening, cycle inventory, ordering and safety stock costs. The **transportation cost** between a DC d and a retailer r is computed as the cost per unit, O_{dr}, times the retailer's demand, μ_r, and the allocation decision variable, x_{dr}^p. It leads to:

$$\sum_{d\in D, r\in R, p\in P} O_{dr} \cdot \mu_r \cdot x_{dr}^p \tag{1}$$

Next, the objective function includes the **fixed opening cost** of DCs. It can be written as:

$$\sum_{d\in D} F_d \cdot y_d \tag{2}$$

The **cycle inventory cost** at retailer r is computed as the unit inventory cost h_r, times the average inventory level. The latter equals half the shipment size, $\frac{q_{dr}}{2}$, times the allocation decision variable, x_{dr}^p. The cycle inventory cost at retailers is given by:

$$\sum_{d\in D, r\in R, p\in P} h_r \cdot \frac{q_{dr}}{2} \cdot x_{dr}^p \tag{3}$$

The shipment size, q_{dr}, can be computed using the EOQ model (prior to the optimization model), which leads to the following closed-form formula when the vehicle's capacity, C_{dr}, is considered:

$$q_{dr} = min\left(C_{dr} \quad ; \quad \sqrt{\frac{2 \cdot k_{dr} \cdot \mu_r}{h_r}}\right) \quad \forall\, r \in R, d \in D. \tag{4}$$

The **cycle inventory cost** at DC d is simply computed as the unit holding cost, H_d, times the average inventory level, $\frac{Q_d}{2}$. The **order cost** is given by the cost of placing an order, K_d, times the number of orders executed per period, $\frac{\sum_{r\in R, p\in P} \mu_r \cdot x_{dr}^p}{Q_d}$. Thus, the ordering and cycle inventory costs are:

$$\sum_{d\in D} H_d \cdot \frac{Q_d}{2} + \sum_{d\in D} K_d \cdot \frac{\sum_{r\in R, p\in P} \mu_r \cdot x_{dr}^p}{Q_d} \tag{5}$$

The order size variable can be computed as an EOQ, leading to $Q_d = \sqrt{\frac{2 \cdot K_d \cdot \sum_{r\in R, p\in P} \mu_r \cdot x_{dr}^p}{H_d}}$. The holding and order costs thus become:

$$\sum_{d\in D} \sqrt{2 \cdot H_d \cdot K_d \cdot \sum_{r\in R, p\in P} \mu_r \cdot x_{dr}^p} \tag{6}$$

The **safety stock costs at the retailers** are computed by multiplying the holding cost at retailer r, h_r, the normal standard deviate, z_α, where α represents the service level (i.e., the probability of satisfying all the demand during the net lead time through safety stocks), and the standard deviation of the demand during the net lead time. The standard deviation and the net lead time to account for depend on the safety stock placement and on the delivery strategies. When accounting for all the cases, we get Eq. (7) for the safety stock costs at retailers. The equation is given here below and the cases are explained in the following. When safety stocks are not held at DCs ($x_{dr}^1 = 1$), the order lead time LT_d is included in the net lead time. In the opposite case ($x_{dr}^2 = 1$), LT_d is not included. Regarding the delivery strategies, when express deliveries are offered ($w_r^{ED} = 1$), the service time S_r is only subtracted from the lead time for regular deliveries, as it equals zero for express deliveries. When express deliveries are not offered ($w_r^{ED} = 0$), regular and express customers are both served with regular deliveries, and the standard deviation is thus $\sigma_r = \sqrt{(\sigma_r^{ED})^2 + (\sigma_r^{RD})^2}$.

$$\sum_{r\in R} h_r \cdot z_\alpha \cdot \Bigg[w_r^{ED} \cdot \Bigg(\sum_{d\in D} x_{dr}^1 \cdot \left(\sqrt{(\sigma_r^{RD})^2 \cdot (LT_d + LT_{dr} - S_r) + (\sigma_r^{ED})^2 \cdot (LT_d + LT_{dr})} \right) + \sum_{d\in D} x_{dr}^2 \cdot \left(\sqrt{(\sigma_r^{RD})^2 \cdot (LT_{dr} - S_r) + (\sigma_r^{ED})^2 \cdot LT_{dr}} \right) \Bigg) + (1 - w_r^{ED}) \cdot \left(\sum_{d\in D} x_{dr}^1 \cdot \left(\sigma_r \cdot \sqrt{LT_d + LT_{dr} - S_r} \right) + \sum_{d\in D} x_{dr}^2 \cdot \left(\sigma_r \cdot \sqrt{LT_{dr} - S_r} \right) \right) \Bigg] \tag{7}$$

Eq. (7) is non-linear, which clearly hurts the tractability of the model. However, it can be approximated to remove this non-linearity. Eq. (7) is first reformulated as follows:

$$\sum_{r\in R} h_r \cdot z_\alpha \cdot \Bigg[w_r^{ED} \cdot \Bigg(\sum_{d\in D} x_{dr}^1 \cdot \left(\sqrt{(\sigma_r)^2 \cdot (LT_d + LT_{dr} - S_r) + (\sigma_r^{ED})^2 \cdot S_r} \right) + \sum_{d\in D} x_{dr}^2 \cdot \left(\sqrt{(\sigma_r)^2 \cdot (LT_{dr} - S_r) + (\sigma_r^{ED})^2 \cdot S_r} \right) \Bigg) + (1 - w_r^{ED}) \cdot \left(\sum_{d\in D} x_{dr}^1 \cdot \left(\sigma_r \cdot \sqrt{LT_d + LT_{dr} - S_r} \right) + \sum_{d\in D} x_{dr}^2 \cdot \left(\sigma_r \cdot \sqrt{LT_{dr} - S_r} \right) \right) \Bigg] \tag{8}$$

Looking at Eq. (8), clear similarities appear between the two first square roots and the two last ones (except for $(\sigma_r^{ED})^2 \cdot S_r$). By approximating the two first square roots by a sum of square roots of their two terms, Eq. (8) simplifies to the following Eq. (9), which is linear. By default, we will use this approximation (also an upper bound) for the safety stock costs at retailers, as it is more computationally efficient.[1]

$$\sum_{r\in R} h_r \cdot z_\alpha \cdot \left(\sigma_r \cdot \left[\sum_{d\in D} \left(x_{dr}^1 \cdot \sqrt{LT_d + LT_{dr} - S_r} + x_{dr}^2 \cdot \sqrt{LT_{dr} - S_r} \right) \right] + \sigma_r^{ED} \cdot w_r^{ED} \cdot \sqrt{S_r} \right) \tag{9}$$

When safety stock placement strategy two is selected, safety stocks are held at the retailers as well as at the DCs. The **safety stock costs at the DCs** are computed by multiplying the holding cost at the DCs, H_d, the normal standard deviate, z_α, and the standard deviation of the demand during the order lead time between the central plant and the DC d, $\sqrt{LT_d \cdot \sum_{r\in R} \sigma_r^2 \cdot x_{dr}^2}$. The safety stock cost at the DCs is thus given by:

$$\sum_{d\in D} H_d \cdot z_\alpha \cdot \sqrt{LT_d \cdot \sum_{r\in R} \sigma_r^2 \cdot x_{dr}^2} \tag{10}$$

Finally, when retailers offer express deliveries ($w_r^{ED} = 1$), an additional fee is charged to the customers, which results in an income for the company. It is computed by multiplying the extra fee per product, T_r, and the express deliveries demand, given by $P_r^{ED} \cdot \mu_r$. The total income originated by the extra fees is determined by:

$$\sum_{r\in R} P_r^{ED} \cdot \mu_r \cdot T_r \cdot w_r^{ED} \tag{11}$$

3.4. CQMIP formulation

In this subsection, we formulate our model as a conic quadratic mixed-integer program. We replace the non-linear terms in Eqs. (6) and (10) with auxiliary variables λ_d and $\overline{\sigma}_d$, and include the following quadratic constraints, using the fact that $x_{dr}^p = (x_{dr}^p)^2$:

$$\sum_{r\in R, p\in P} \mu_r \cdot (x_{dr}^p)^2 \le \lambda_d^2 \quad \forall\, d \in D, \tag{12}$$

$$\sum_{r\in R} \sigma_r^2 \cdot (x_{dr}^2)^2 \le \overline{\sigma}_d^2 \quad \forall\, d \in D. \tag{13}$$

The location-inventory model with safety stock placement decisions and delivery strategies is formulated by:

$$\min \sum_{d\in D, r\in R, p\in P} O_{dr} \cdot \mu_r \cdot x_{dr}^p \tag{14}$$

$$+ \sum_{d\in D} F_d \cdot y_d \tag{15}$$

$$+ \sum_{d\in D, r\in R, p\in P} h_r \cdot \frac{q_{dr}}{2} \cdot x_{dr}^p \tag{16}$$

$$+ \sum_{d\in D} \sqrt{2 \cdot H_d \cdot K_d} \cdot \lambda_d \tag{17}$$

$$+ \sum_{r\in R} h_r \cdot z_\alpha \cdot \left(\sigma_r \cdot \left[\sum_{d\in D} \left(x_{dr}^1 \cdot \sqrt{LT_d + LT_{dr} - S_r} + x_{dr}^2 \cdot \sqrt{LT_{dr} - S_r} \right) \right] + \sigma_r^{ED} \times \cdot w_r^{ED} \cdot \sqrt{S_r} \right) \tag{18}$$

$$+ \sum_{d\in D} H_d \cdot z_\alpha \cdot \sqrt{LT_d} \cdot \overline{\sigma}_d \tag{19}$$

$$- \sum_{r\in R} P_r^{ED} \cdot \mu_r \cdot T_r \cdot w_r^{ED} \tag{20}$$

$$\text{s.t} \sum_{r\in R, p\in P} \mu_r \cdot (x_{dr}^p)^2 \le \lambda_d^2 \quad \forall\, d \in D, \tag{21}$$

$$\sum_{r\in R} \sigma_r^2 \cdot (x_{dr}^2)^2 \le \overline{\sigma}_d^2 \quad \forall\, d \in D, \tag{22}$$

$$\sum_{d\in D, p\in P} x_{dr}^p = 1 \quad \forall\, r \in R, \tag{23}$$

$$x_{dr}^p \le y_d \quad \forall\, r \in R, d \in D, p \in P, \tag{24}$$

$$\lambda_d, \overline{\sigma}_d \ge 0 \quad \forall\, d \in D, \tag{25}$$

$$w_r^{ED}, x_{dr}^p, y_d \in \{0, 1\} \quad \forall\, r \in R, d \in D, p \in P. \tag{26}$$

The objective function includes the transportation cost (14); the facility opening cost (15); the inventory cost at retailers (16); the order cost

[1] We tested the accuracy of the approximation by running 144 experiments with the parameter values shown in Table 2, except for $LT_{dr}^{Fixed} = (1, 5, 10, 20)$, for $CV_r = 0.8$, for $P_r^{ED} = (10, 50, 90)$, and for $T_r = (0.05, 0.25)$, where we use subsets of the values. The average computational time for solving the model using Eq. (7) equals 7415 s, whereas the model using Eq. (9) only takes 953 s on average. The safety stock costs at retailers are increased by 1.15% on average using Eq. (9).

at DCs (17); the safety stock cost (18) and (19) and the income from additional fees (20). Constraints (21) and (22) are the quadratic constraints defining the auxiliary variables. Constraints (23) ensure that each retailer is assigned to one DC using one safety stock placement strategy. Constraints (24) guarantee that the retailers are linked to opened DCs. Finally, constraints (25) give the non-negativity constraints of the auxiliary variables and (26) define the domain of the binary decision variables.

The main advantages of the CQMIP approach is that it is direct, efficient and flexible, as it can be solved using standard optimization software packages without the need of specialized algorithms (Atamtürk et al. (2012)). Standard optimization software packages such as CPLEX offer the possibility of solving CQMIPs based on interior-point methods and branch and bound algorithms. However, the performance of the branch and bound algorithm can be improved by strengthening the formulations with cutting planes. Therefore, we have reformulated the quadratic constraints (21) and (22) with polymatroid inequalities to strengthen the convex relaxation of the CQMIP. The separation problem for the extended polymatroid inequalities is solved by the greedy algorithm of Edmonds (1970). We refer the reader to Atamtürk and Narayanan (2008) and to Atamtürk et al. (2012) for a detailed explanation of the separation problem and the greedy algorithm.

3.5. Extension: demand correlation across retailers and stochastic lead times

In this section, we extend our model by incorporating demand correlation and stochastic lead times. First, demand correlation impacts the safety stocks at the DCs, and its inclusion can better support the safety stock placement decision within the location-inventory problem. Demand correlation is considered across every pair of retailers, defined by the demand correlation coefficient (ρ_{rl}) between retailer r and retailer l. The demand at a DC d is assumed to follow a normal distribution with mean $\sum_{r\in R, p\in P} \mu_r \cdot x_{dr}^p$ and variance $\sum_{r\in R, l\in R, p\in P} \rho_{rl}\cdot\sigma_r\cdot\sigma_l\cdot x_{dr}^p\cdot x_{dl}^p$. Second, stochastic lead times affect safety stocks at the retailers and at the DCs, increasing the total supply chain costs. We suppose that the order lead times at the DCs, and the lead times between DCs and retailers are normally distributed random variables with means LT_d and LT_{dr}, and variances $\sigma_{LT_d}^2$ and $\sigma_{LT_{dr}}^2$, respectively. Assuming that orders do not cross (which is consistent when using single sourcing) and that successive lead times are independent, we can apply the same expression as in Nahmias and Olsen (2015). This approach can be used in our model due to the fact that when DCs hold safety stocks, they quote a zero service time to the retailers (see Humair et al. (2013)).

Consequently, we reformulate the safety stock costs at the retailers and at the DCs. The safety stock cost term at the retailers (18) is replaced by:

$$\sum_{r\in R} h_r\cdot z_\alpha\cdot\left[\sum_{d\in D}\left(x_{dr}^1\cdot\sqrt{(LT_d+LT_{dr})\cdot\sigma_r^2+\left(\sigma_{LT_d}^2+\sigma_{LT_{dr}}^2\right)\cdot\mu_r^2-S_r\cdot\sigma_r^2}\right.\right.$$
$$\left.\left.+x_{dr}^2\cdot\sqrt{LT_{dr}\cdot\sigma_r^2+\sigma_{LT_{dr}}^2\cdot\mu_r^2-S_r\cdot\sigma_r^2}\right)+\sigma_r^{ED}\cdot w_r^{ED}\cdot\sqrt{S_r}\right] \tag{27}$$

For the safety stocks at the DCs (19), constraint (22) is replaced by the following quadratic constraint:

$$LT_d\cdot\left(\sum_{r\in R, l\in R}\rho_{rl}\cdot\sigma_r\cdot\sigma_l\cdot x_{dr}^2\cdot x_{dl}^2\right)+\sigma_{LT_d}^2\cdot\left(\sum_{r\in R}\mu_r\cdot x_{dr}^2\right)^2\le\overline{\sigma}_d^2\quad\forall\, d\in D, \tag{28}$$

4. Numerical experiments

In this section, we present computational experiments to illustrate the application of our approach and obtain managerial insights.

Table 2
Parameter values used for the numerical experiments.

Parameters	Values
Indexes	
$n_r = n_d$	88
Network parameters	
LT_d	(1, 5) days
LT_{dr}^{Fixed}	(0, 1, 5, 10, 20) days
Speed	60 km/h
$S_r^{\%}$	(20, 60, 100) %
Demand parameters	
CV_r	(0.4, 0.8)
α	99%
z_α	2.33
P_r^{ED}	(10, 30, 50, 70, 90) %
Cost parameters	
F_d	1000 €/day
K_d	1000 €/order
$H_d = h_r$	0.05 €/item · day
T_r	(0.05, 0.25, 0.50) €/item

4.1. Computational results

We use the data from the 1990 U.S Census given in Daskin (1995). We employ the 88-node data set, which reports data of the lower U.S state capitals plus Washington DC and the 50 largest U.S cities. In Table 2, we summarize the values used for the different parameters. The potential DC locations are chosen to be the retailers' locations ($n_d = n_r$). The retailers' demand, μ_r, is determined by dividing the population of each city by 1000 (as in Atamtürk et al. (2012)). The standard deviation of demand is computed using the coefficient of variation CV_r. Transportation unit costs, O_{dr}, are proportional to the distances between the DCs and the retailers, with a unit distance cost of 0.001€/km. The lead times between DCs and the retailers, LT_{dr}, are computed based on a fixed time (LT_{dr}^{Fixed}) that varies between 0 and 20 days, plus a variable time that depends on the distance (assuming an average vehicle speed of 60 km/h). We assume that the order lead times between the central plant and the DCs, LT_d, are the same for all DCs, and restricted to being either 1 or 5 days. For the stochastic lead times, σ_{LT_d} and $\sigma_{LT_{dr}}$ are computed using the coefficient of variation (equal to 0.8). For the sake of simplicity, in our experiments, we consider that the shipment sizes equal the vehicle's capacity ($q_{dr} = C_{dr}$).[2] We compute the vehicle's capacity, C_{dr}, as 10% of the largest retailer's demand. The percentage of express deliveries demand, P_r^{ED}, is the same at all retailers and varies between 10 and 90%. The service time, S_r, is considered the same at all retailers. This assumption is grounded on the fact that, in many real situations, companies offer identical service time, regardless of where their customers are located. It is computed as a percentage (using $S_r^{\%}$) of LT_{dr}^{Fixed}, the fixed lead time between the DCs and the retailers. For a better understanding of the results and the different trade-offs, we consider uncorrelated demands and deterministic lead times for the majority of the experiments, and study their effect separately. In Section 4.2, we include a detailed analysis of their impact on the safety stock placement (see Fig. 4) and on the overall supply chain cost (see Fig. 3). The parameter values lead to $2\cdot5\cdot3\cdot2\cdot5\cdot3 = 900$ parameter combinations, and we use them to carry out 2800 experiments.

The CQMIP is solved in CPLEX on a 2.70 GHz computer with 16 GB of RAM. Table 3 shows, for the different versions of our model, the averages

[2] We tested this assumption, running 900 experiments (all the parameter combinations, $k_{dr} = 100$) with shipment size decisions, computed using Eq. (4). We obtain a slight change in the results: the total costs are reduced by 0.65% on average, 78% of the vehicles are loaded to full capacity, and the average shipment size equals 89%.

Table 3
Computational performance of the CQMIP.

	# Instances	Time (Sec.)			# Cuts			# DCs			Gap (%)
		Min	Avg	Max	Min	Avg	Max	Min	Avg	Max	
Base Case	900	546	919	3103	9	27	70	4.00	6.44	7.00	0.00
Stochastic LT	270	203	864	8072	9	10	11	6.00	6.83	7.00	0.02
Correlated Demands	1050	128	926	5799	9	10	13	6.00	6.86	7.00	0.01

over all our experiments, for the computational time, the number of polymatroid cuts applied, the number of opened DCs and the optimality gap. The CQMIP finds the optimal solution for the base cases (using the parameter values in Table 2) in an average time of around 15 min. Similar results are obtained for the instances with stochastic lead times and correlated demands (using the parameter values described in Figs. 3 and 4, respectively). The addition of these features does not hurt significantly the computational performance of our model. To complete this discussion of the computational performance of our model, we compare the CQMIP with and without the reformulation of the quadratic constraints with polymatroid inequalities. The formulation without polymatroid cuts gets an optimality gap of 0.3% in an average computational time of 35 min (maximum of 2 h), whereas our model with polymatroid cuts finds the optimal solution in all instances in 18 min on average.[3] All in all, our approach is therefore efficient in solving the location-inventory problem with safety stock placement decisions and delivery strategies.

4.2. Managerial insights

In this section, we derive managerial insights from the computational results. We start by analysing the trade-offs that influence the selection of safety stock placement strategies. In our model, risk pooling has two combined effects that can be inferred from Eqs. (9) and (10): lead time pooling and demand variability pooling. From Eq. (10), we see that demand variability pooling benefits arise for proportionally long order lead times between the central plant and the DCs (LT_d), which will result in the selection of safety stock placement strategy two (x^2_{dr}). Conversely, when the order lead times between the central plant and the DCs (LT_d) are proportionally short, and the lead times between DCs and the retailers (LT_{dr}) proportionally long, the demand variability pooling becomes less valuable, which leads decision makers towards choosing strategy one (x^1_{dr}) to benefit from lead time pooling (first term of Eq. (9)). These tendencies can be observed in Table 4, where we show the number of retailers served with safety stock placement strategies one and two, respectively. Strategy two is more broadly chosen for long order lead times between the central plant and the DCs (LT_d), while strategy one is more frequently selected for long lead times between DCs and retailers (LT^{Fixed}_{dr}). We also see in Table 4 that safety stock placement strategy one tends to be selected slightly more often when the coefficient of variation (CV_r) increases. Although this may seem contradictory to intuition, it is due to the fact that the number and location of DCs varies when increasing the coefficient of variation. By investigating our results more deeply, we find that, when the coefficient of variation increases, DCs are located closer to large cities to mitigate the safety stock increase in these cities. Then, and this is what we observe in our cases, it may happen that a DC located in a new city serves less retailers, which decreases the value of demand variability pooling and leads to a change in the selected safety stock placement strategy (from two to one). This also illustrates the complex and sometimes unintuitive interdependency of the decisions (location, safety stock sizing and placement) and thus the value of an integrated model.

[3] This comparison is conducted on 24 experiments, with the parameter values $LT_d = (1, 5)$, $LT_{dr} = (1, 10, 20)$, $CV_r = (0.4, 0.8)$, $S^{\%}_r = 60\%$, $P^{ED}_r = (30, 70\%)$, and $T_r = 0.25$.

Table 4
Number of retailers served with safety stock placement strategy one (left|) and safety stock placement strategy two (|right), for different coefficients of variation (CV_r), service times ($S^{\%}_r$), order lead times (LT_d) and lead times between DCs and retailers (LT^{Fixed}_{dr}). The values in the table are averages over 15 instances (for each value of P^{ED}_r and T_r).

CV_r	LT^{Fixed}_{dr}	$S^{\%}_r = 20$		$S^{\%}_r = 60$		$S^{\%}_r = 100$	
		$LT_d = 1$	$LT_d = 5$	$LT_d = 1$	$LT_d = 5$	$LT_d = 1$	$LT_d = 5$
0.4	0	0|88	0|88	0|88	0|88	0|88	0|88
	1	15|73	0|88	12|76	0|88	0|88	0|88
	5	88|0	9|79	88|0	7|81	0|88	0|88
	10	88|0	22|66	88|0	9|79	0|88	0|88
	20	88|0	88|0	88|0	22|66	0|88	0|88
0.8	0	0|88	0|88	0|88	0|88	0|88	0|88
	1	16|72	0|88	13|75	0|88	0|88	0|88
	5	88|0	9|79	88|0	7|81	0|88	0|88
	10	88|0	29|59	88|0	9|79	0|88	0|88
	20	88|0	88|0	88|0	29|59	0|88	0|88

Another interesting insight arises from the impact of service times on the selected safety stock placement strategies. As seen in Table 4, strategy two tends to be selected when the service time ($S^{\%}_r$) increases, especially if the order lead time between the central plant and the DCs (LT_d) is long. An increase of the service time is equivalent to reducing the lead time between DCs and retailers, increasing the importance of the order lead time at DCs, and therefore motivating the use of demand variability pooling strategies. This trend is further captured in Fig. 1, where we show the safety stock placement strategy when service times at retailers differ from each other. In particular, we divide the retailers into three groups according to their demand size (lower than 200, between 200 and 450, and higher than 450 products), and assign a distinct service time to the retailers of each group. The coefficient of variation (CV_r) is fixed to 0.8, which implies that large retailers also have large variances. Interestingly, we observe in Fig. 1 that the effect of the service times on the safety stock placement strategies selected depends on the size of the retailer. The large retailers have a greater impact on the safety stock placement decision than the smaller retailers. In Fig. 1 a, only the small retailers with long service times (decreasing the effective retailer-DC lead time) and small distances to the DCs benefit from demand variability pooling strategies and hold safety stocks at the DCs. In Fig. 1b, large retailers provide long service times and are thus encouraged to pool demand variability at the DCs (the effective lead time is short), holding safety stock at the DCs. Consequently, the small and medium size retailers can take advantage and pool their smaller variability with the large retailers in common DCs by holding safety stocks in those DCs. Looking at Fig. 1a and 1b, it appears that the service time at the retailers clearly affects the optimal safety stock placement strategy, and that the dynamics among retailers are quite complex, as changing the service time for a retailer may impact the safety stock placement of another retailer.

Fig. 2 details the supply chain design in four instances, with the number of DCs opened, the allocation of flows between DCs and retailers, and the selected safety stock placement strategies. Several trends can be observed in Fig. 2. To begin with, the retailers assigned to a given DC tend to select the same safety stock placement strategy, except for those retailers with higher demands. For these retailers, strategy one tends to

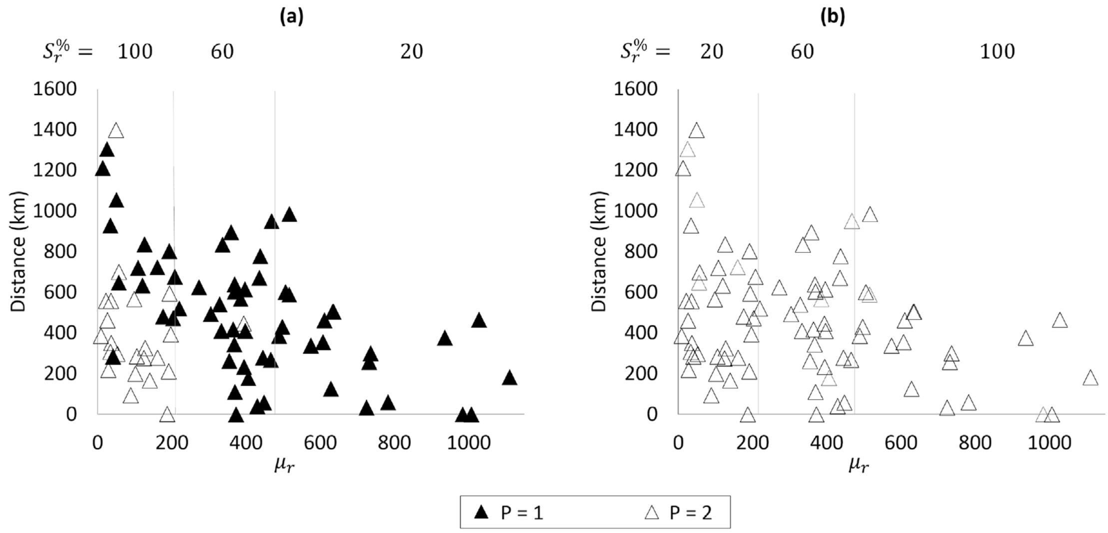

Fig. 1. Safety stock placement strategy selected depending on the retailer's demand and on the distance between the retailer and the linked DC. The two graphs are for different service times at retailers ($S_r^{\%}$), with the service times changing according to the demand size of the retailer: in (a) small retailers have longer service time, while in (b) small retailers have shorter service time. We use $LT_d = 1$, $LT_{dr}^{Fixed} = 5$, $CV_r = 0.8$, $P_r^{ED} = 30\%$ and $T_r = 0.25$.

be selected as a large retailer tends to take proportionally less benefit from demand variability pooling at a DC (it represents a large share of the pool itself). Furthermore, strategy two tends to be selected when DCs have a large number of retailers assigned, to benefit from demand variability pooling. In Fig. 2a, we observe that relatively long order lead times (LT_d) result in DCs holding safety stocks and a limited number of DCs being opened to benefit from demand variability pooling. When the order lead time is reduced in Fig. 2b, the DCs do not hold safety stocks and a larger number of DCs is opened (as DCs are not pooling safety stocks). When comparing Fig. 2a and c, we see that an increase of the lead times between DCs and retailers (LT_{dr}^{Fixed}) motivates a broader use of strategy one, and a movement further to the east of the south eastern DC. In Fig. 2d, we show how DCs are centralized when the facility opening costs are increased, leading to four opened DCs. For this case, safety stocks are held at those DCs that serve a larger number of retailers and benefit largely from demand variability pooling.

The previous results show the interdependency between location decisions, lead times, safety stock placement decisions, demand variability pooling and lead time pooling. In this context, Figs. 3 and 4 aim at quantifying the importance of integrating safety stock placement decisions when designing the supply chain. In Fig. 3, we compare the total costs found by our model, where we decide whether to hold safety stocks at the DCs or not, with those found by the model that does not integrate the decision and forces DCs to hold safety stocks. Moreover, we compare these two models for three different situations: when all the lead times are deterministic, when the order lead times are deterministic but the lead times between DCs and retailers are stochastic, and when all the lead times are stochastic. For supply chains with deterministic lead times, the option to have safety stocks only at the retailers reduces the overall cost by up to 2.6%. The cost reduction clearly increases when the lead time between DCs and retailers is stochastic (lead time pooling prevails), and even more when the order lead times at the DCs are also stochastic. When stochastic lead times are applied, integrating safety stock placement decisions reduces the total cost by up to 22.3%. Looking at Eq. (27), we see that lead time variances are pooled when safety stocks are only held at the retailers, encouraging safety stock placement strategy one, and justifying the larger cost difference compared to deterministic lead times. When lead times between DCs and retailers are increasing, the cost difference goes up and then down. The cost difference (in absolute value) keeps increasing with long lead times between DCs and retailers, but it represents a lower percentage due to the significant total cost increase. In Fig. 3, we also see that the number of retailers served with safety stock placement strategy one increases when stochastic lead times are considered. When stochastic lead times are applied, the increase of the lead times (at the DCs, and between DCs and retailers) motivates the use of lead time pooling strategies to mitigate the increase of safety stock costs. We note that the average number of DCs opened increases from 5.8 to 7, when changing from deterministic lead times to stochastic lead times. The increase on the number of DCs reduces the lead times between DCs and retailers, and consequently, reduces the impact of lead time uncertainty. Overall, we see that selecting the right safety stock strategy may affect the location decisions and allows to significantly reduce the supply chain costs in all cases, and even more when the lead time variability increases.

Furthermore, Fig. 4 shows the importance of integrating safety stock placement decisions when positive and negative demand correlation is considered. For the computation of the negative correlation matrix, we fix a negative correlation (value given in Fig. 4) between each retailer and the two closest retailers, and the rest of the correlation matrix is completed using the CVX package (see Grant and Boyd (2014)), which provides the minimum correlation coefficients that still result in a positive-semidefinite matrix. Therefore, in the negative correlation case, the average correlation coefficient is very close to zero. For the positive correlation matrix, we fix the same correlation coefficients for every pair of retailers (value given in Fig. 4, except for $\rho_{rr} = 1$). In Fig. 4a, we show the cost of demand correlation, i.e. the cost difference when it is considered or not, for two distinct models: our model, where holding safety stocks at the DCs is a decision, and the model where DCs are forced to hold safety stocks. As safety stocks can be reduced at DCs, total supply chain costs are slightly reduced with negative correlation, compared to the case with no demand correlation (see Eq. (28). However, when increasing the correlation, the total supply chain cost increases by up to 10.4% when DCs are forced to hold safety stocks, and only by up to 5.6% when we allow safety stocks to be held only at the retailers. The inclusion of safety stock placement decisions can thus reduce the cost by up to 5% when demand correlation is considered. Optimal safety stock placement enables a mitigation of the positive demand correlation impact. In Fig. 4b, which shows the safety stock costs at retailers and at DCs, we observe that indeed, when the demand correlation increases, the safety stock costs increase and strategy one, with no safety stock at DCs, is used by more and more retailers. Lead time pooling strategies mitigate the increase of safety stock costs, while risk pooling is less beneficial when

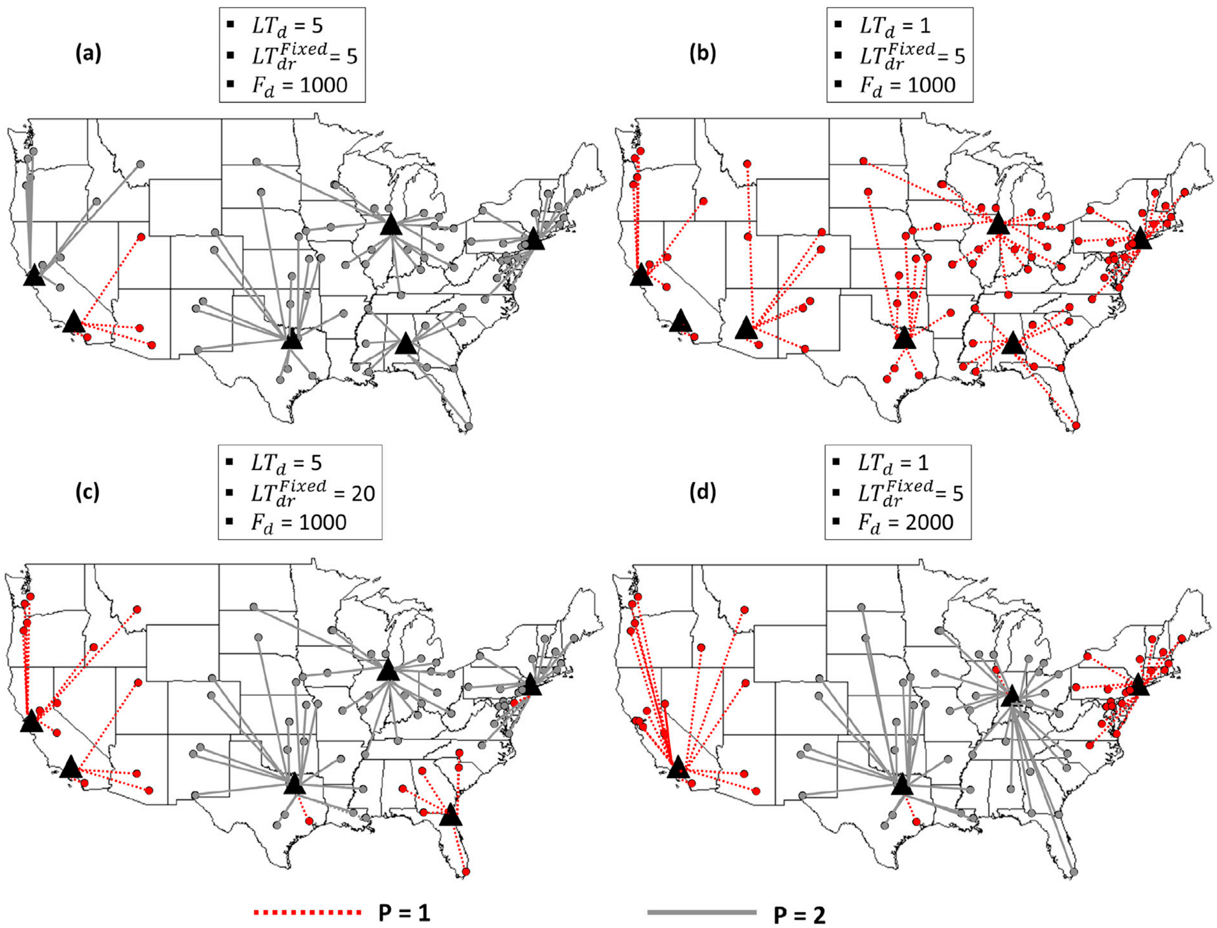

Fig. 2. Safety stock placement strategy selected for different values of order lead time (LT_d), lead time between DCs and retailers (LT_{dr}^{Fixed}) and facility opening cost (F_d), with $CV_r = 0.8$, $S_r^{\%} = 60\%$, $P_r^{ED} = 70\%$ and $T_r = 0.5$.

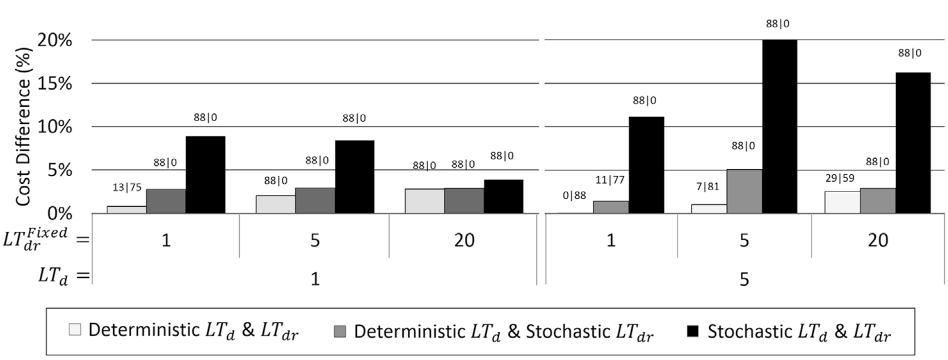

Fig. 3. Total cost difference (in percentage) between our model and the model forcing safety stocks to be held at DCs, for three different combinations of deterministic and stochastic lead times, and for different lead times between DCs and retailers (LT_{dr}^{Fixed}) and different order lead times (LT_d). Above each bar, we give the average number of retailers served with safety stock placement strategy one (left|) and safety stock placement strategy two (|right). The values in the figure are averages over 15 instances (for each value of P_r^{ED} and T_r), with $CV_r = 0.8$ and $S_r^{\%} = 60\%$.

demand correlation increases. Overall, Figs. 3 and 4 show that the integration of safety stock placement decisions can lead to significant cost savings, and particularly with stochastic lead times and demand correlation.

Another interesting insight arises from the selection of the delivery strategies at the retailers. In all the computational experiments, we found that, for each specific instance, all the retailers apply the same delivery strategy. Table 5 shows in which situations express deliveries (ED) or only regular deliveries (RD) are offered. Several trends are captured in Table 5. Firstly, with longer service times ($S_r^{\%}$), safety stocks are reduced for regular deliveries, and retailers thus tend to select this type of delivery. Secondly, we observe that increasing the demand uncertainty (CV_r) and the lead time between the DCs and the retailers (LT_{dr}^{Fixed}) leads to the selection of only regular deliveries in order to mitigate the increase of the safety stock cost. Finally, by increasing the fee (T_r), the retailers have incentives to improve the customer service, and consequently, they tend to offer express deliveries.

A classic trade-off in the supply chain design problem is whether to centralize facilities in order to gain economies of scale, or to decentralize them to become more responsive by being closer to retailers. One interesting result from our model is that the service times may impact this trade-off, along with the location and the number of DCs opened. This is illustrated in Fig. 5 for two instances. We see that safety stock placement strategy one is selected in both cases, so that DCs do not hold safety stocks. Regarding the delivery strategies, the increase in service time leads to drop off express deliveries as they become comparatively more

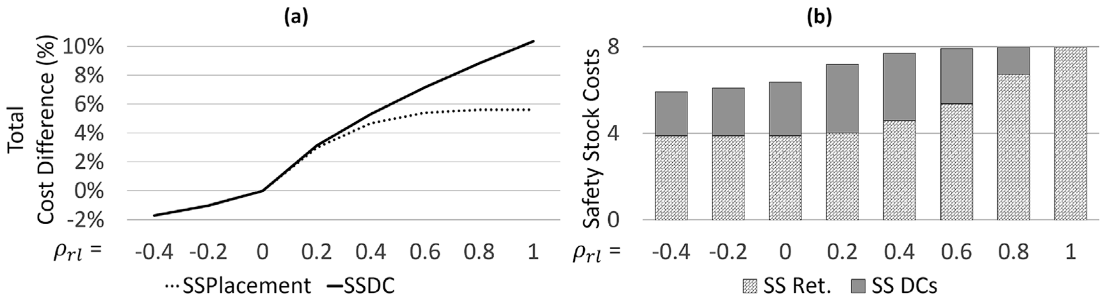

Fig. 4. (a) Total cost difference (in percentage) when demand correlation is included, for our model with safety stock placement decisions (SSPlacement) and the model forcing safety stocks to be held at DCs (SSDC), for different correlation coefficients (ρ_{rl}). (b) Safety stock costs (in thousands) at retailers (SS Ret.) and at DCs (SS DCs) for different correlation coefficients (ρ_{rl}), applying our model. The values in the figures are averages over 150 instances (for each value of LT_d, LT_{dr}^{Fixed}, P_r^{ED}, T_r, and with $S_r^{\%} = 100\%$ and $CV_r = 0.8$).

Table 5
Express deliveries (ED) and regular deliveries (RD) applied for different coefficients of variation (CV_r), fee (T_r), service times ($S_r^{\%}$) and lead times between DCs and retailers (LT_{dr}^{Fixed}). The values in the table are found over 10 instances (we get the same delivery strategy for each value of LT_d and P_r^{ED}).

T_r	LT_{dr}^{Fixed}	$CV_r = 0.4$			$CV_r = 0.8$		
		$S_r^{\%} =$ 20	$S_r^{\%} =$ 60	$S_r^{\%} =$ 100	$S_r^{\%} =$ 20	$S_r^{\%} =$ 60	$S_r^{\%} =$ 100
0.05	1	ED	ED	ED	ED	RD	RD
	5	ED	RD	RD	RD	RD	RD
	10	RD	RD	RD	RD	RD	RD
	20	RD	RD	RD	RD	RD	RD
0.25	1	ED	ED	ED	ED	ED	ED
	5	ED	ED	ED	ED	ED	ED
	10	ED	ED	ED	ED	ED	RD
	20	ED	ED	ED	ED	RD	RD

expensive. Finally, and most interestingly, we observe in Fig. 5 that, when DCs do not hold safety stocks, longer service times may lead to an increase in the number of DCs opened, which reduces the lead times between DCs and retailers. Due to the concavity of the square root function (see the first term of Eq. (9)), a decrease of these lead times (LT_{dr}) has a larger impact when the net lead times are short. Longer service times may thus encourage a reduction of lead times and an increase in the number of DCs. We can conclude from Fig. 5 that the integration of delivery strategies in the supply chain design problem is justified as it may impact the number and location of DCs.

5. Conclusions

This paper presents an approach for integrating safety stock placement and delivery strategy decisions into the strategic supply chain design. We propose a mathematical model that considers differentiated service times and two different classes of customers regarding their delivery time preferences. It can be extended to include demand correlation and stochastic lead times. Our model minimizes the costs of transportation, facility opening, cycle inventory, ordering and safety stocks. The resulting non-linear model can be formulated as a conic quadratic mixed-integer program, as proposed in Atamtürk et al. (2012), by including binary variables that explicitly decide for which connected retailers the opened DCs hold safety stocks.

Managerial insights are then inferred from extensive computational experiments. We discuss the trade-off between demand variability pooling and lead time pooling. In particular, proportionally short lead times between DCs and retailers result in a supply chain where DCs and retailers tend to hold safety stocks in order to benefit from pooling demand variability at the DCs. In such a supply chain, the retailers tend to offer express deliveries to their customers, except if the service times are long or if the fees are very small. When proportionally long lead times between DCs and retailers are considered, the supply chain is configured to benefit from lead time pooling so that DCs do not hold safety stocks. In this case, the retailers only propose regular deliveries, except if the service times are short or if the fees are large. We also show the impact of demand uncertainty on safety stock placement decisions. The retailers with high demands tend to hold safety stocks and to refrain from pooling demand variability at the DCs. Moreover, the DCs serving a large number of retailers tend to hold safety stocks to benefit from demand variability

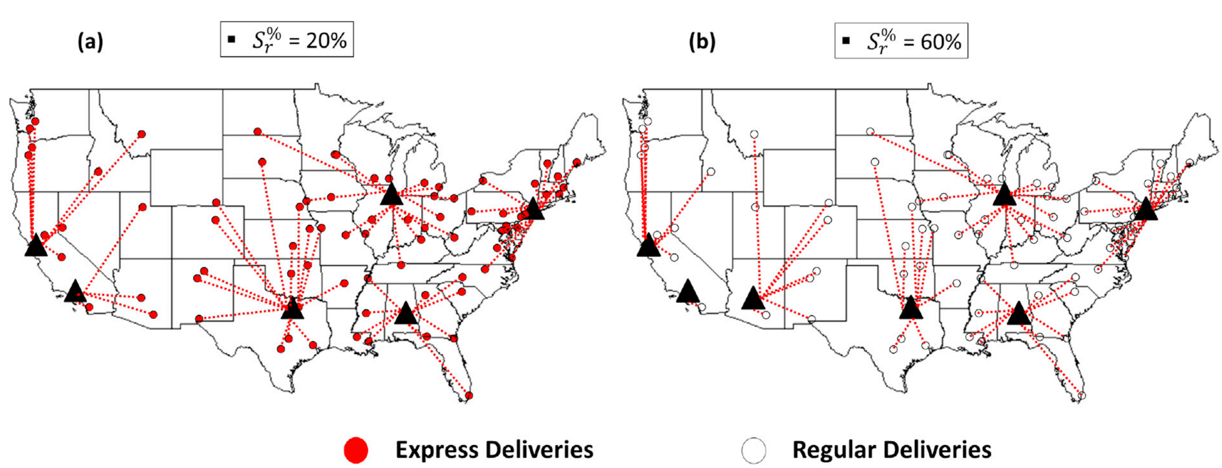

Fig. 5. Safety stock placement ($p = 1$) and delivery strategy selected for different service times ($S_r^{\%}$), with $LT_d = 1$, $LT_{dr}^{Fixed} = 5$, $CV_r = 0.4$, $P_r^{ED} = 30\%$ and $T_r = 0.05$.

pooling at the DCs. Regarding the impact of the service time at the retailers, we observe that increasing the service time is equivalent to decreasing the lead time between DCs and retailers, which encourages demand variability pooling and holding safety stocks at the DCs. Moreover, we show that the dynamics among retailers are quite complex, as changing the service time for a retailer may impact the safety stock placement strategy of another retailer. Finally, we show how optimal safety stock placement decisions help mitigating the cost impact of demand correlation and stochastic lead times, by moving from strategies in which safety stocks are held at the DCs to strategies in which safety stocks are held at the retailers. The integration of safety stock placement and delivery decisions into the strategic design of supply chains is justified as we have shown that these decisions and the location decisions are interrelated and as it allows to balance all the aforementioned trade-offs and reduce the overall supply chain cost.

In the future, this research could be extended by integrating other features such as three stage supply chains, multiple modes of transportation, multiple products, direct shipments from the production sites to the retailers, capacity constraints at the retailers and at the DCs, or the consideration of more than two customer classes. One could also consider a finer customer behaviour modeling approach with respect to the relation between the fee and the delivery times, including a penalty when express deliveries' demands are not proposed.

Acknowledgments

This research was supported by the Belgian Fund for Scientific Research (FNRS) through the FRESH grant. It was partly conducted while the first author was a visiting Ph.D. student at the Technical University of Munich.

References

Atamtürk, A., Berenguer, G., Shen, Z.-J.M., 2012. A conic integer programming approach to stochastic joint location-inventory problems. Oper. Res. 60 (2), 366–381.
Atamtürk, A., Narayanan, V., 2008. Polymatroids and mean-risk minimization in discrete optimization. Oper. Res. Lett. 36 (5), 618–622.
Beamon, B.M., 1998. Supply chain design and analysis: models and methods. Int. J. Prod. Econ. 55 (3), 281–294.
Bravo, J.J., Vidal, C.J., 2013. Freight transportation function in supply chain optimization models: a critical review of recent trends. Expert Syst. Appl. 40 (17), 6742–6757.
Chen, H., Li, P., 2015. Optimization of (R, Q) policies for serial inventory systems using the guaranteed service approach. Comput. Ind. Eng. 80, 261–273.
Chopra, S., Meindl, P., 2016. Supply Chain Management: Strategy, Planning, and Operation, sixth ed.
Daskin, M., Coullard, C., Shen, Z.-J., 2002. An inventory-location model: formulation, solution algorithm and computational results. Ann. Oper. Res. 110 (1–4), 83–106.
Daskin, M.S., 1995. Network and Discrete Location: Models, Algorithms, and Applications. Wiley, New York.
Edmonds, J., 1970. Submodular functions, matroids, and certain polyhedra. Combin. Struct. Appl. 69–87.
Eppen, G.D., 1979. Effects of centralization on expected costs in a multi-location newsboy problem. Manag. Sci. 25 (5), 498–501.
Eruguz, A.S., Sahin, E., Jemai, Z., Dallery, Y., 2016. A comprehensive survey of guaranteed-service models for multi-echelon inventory optimization. Int. J. Prod. Econ. 172, 110–125.
Evers, P.T., 1995. Expanding the square root law: an analysis of both safety and cycle stocks. Transport. Res. E Logist. Transport. Rev. 31 (1), 1–20.
Farahani, R.Z., Rashidi Bajgan, H., Fahimnia, B., Kaviani, M., 2015. Location-inventory problem in supply chains: a modelling review. Int. J. Prod. Res. 53 (12), 3769–3788.
Funaki, K., 2012. Strategic safety stock placement in supply chain design with due-date based demand. Int. J. Prod. Econ. 135 (1), 4–13.
Grahl, J., Minner, S., Dittmar, D., 2016. Meta-heuristics for placing strategic safety stock in multi-echelon inventory with differentiated service times. Ann. Oper. Res. 2 (242), 489–504.
Grant, M., Boyd, S., 2014. CVX: Matlab Software for Disciplined Convex Programming, Version 2.1. http://cvxr.com/cvx.
Graves, S.C., Willems, S.P., 2003. Supply chain Design: Safety Stock Placement and Supply Chain Configuration, vol. 11. Elsevier, pp. 95–132.
Graves, S.C., Willems, S.P., 2005. Optimizing the supply chain configuration for new products. Manag. Sci. 51 (8), 1165–1180.
Gunasekaran, A., Patel, C., McGaughey, R.E., 2004. A framework for supply chain performance measurement. Int. J. Prod. Econ. 87 (3), 333–347.
Hua, N.G., Willems, S.P., 2016. Analytical insights into two-stage serial line supply chain safety stock. Int. J. Prod. Econ. 181, 107–112.
Humair, S., Ruark, J.D., Tomlin, B., Willems, S.P., 2013. Incorporating stochastic lead times into the guaranteed service model of safety stock optimization. Interfaces 43 (5), 421–434.
Inderfurth, K., 1992. Valuation of Leadtime Reduction in Multi-stage Production Systems. Operations Research in Production Planning and Inventory Control, pp. 413–427.
Inderfurth, K., 1995. Multistage safety stock planning with item demands correlated across products and through time. Prod. Oper. Manag. 4 (2), 127–144.
Inderfurth, K., Minner, S., 1998. Safety stocks in multi-stage inventory systems under different service measures. Eur. J. Oper. Res. 106 (1), 57–73.
Klose, A., Drexl, A., 2005. Facility location models for distribution system design. Eur. J. Oper. Res. 162 (1), 4–29.
Kumar, S., Tiwari, M., 2013. Supply chain system design integrated with risk pooling. Comput. Ind. Eng. 64 (2), 580–588.
Li, H., Jiang, D., 2012. New model and heuristics for safety stock placement in general acyclic supply chain networks. Comput. Oper. Res. 39 (7), 1333–1344.
Magnanti, T.L., Shen, Z.-J.M., Shu, J., Simchi-Levi, D., Teo, C.-P., 2006. Inventory placement in acyclic supply chain networks. Oper. Res. Lett. 34 (2), 228–238.
Melo, M.T., Nickel, S., Saldanha-Da-Gama, F., 2009. Facility location and supply chain management–a review. Eur. J. Oper. Res. 196 (2), 401–412.
Minner, S., 2000. Strategic Safety Stocks in Supply Chains. Springer.
Miranda, P.A., Garrido, R.A., 2004. Incorporating inventory control decisions into a strategic distribution network design model with stochastic demand. Transport. Res. E Logist. Transport. Rev. 40 (3), 183–207.
Moncayo-Martínez, L.A., Zhang, D.Z., 2013. Optimising safety stock placement and lead time in an assembly supply chain using bi-objective max–min ant system. Int. J. Prod. Econ. 145 (1), 18–28.
Mula, J., Peidro, D., Díaz-Madroñero, M., Vicens, E., 2010. Mathematical programming models for supply chain production and transport planning. Eur. J. Oper. Res. 204 (3), 377–390.
Nahmias, S., Olsen, T.L., 2015. Production and Operations Analysis, seventh ed. Waveland Press.
Novich, N., 1990. Distribution strategy: are you thinking small enough? Sloan Manag. Rev. 32 (1), 71–77.
Osman, H., Demirli, K., 2012. Integrated safety stock optimization for multiple sourced stockpoints facing variable demand and lead time. Int. J. Prod. Econ. 135 (1), 299–307.
Owen, S.H., Daskin, M.S., 1998. Strategic facility location: a review. Eur. J. Oper. Res. 111 (3), 423–447.
Park, S., Lee, T.-E., Sung, C., 2010. A three-level supply chain network design model with risk-pooling and lead times. Transport. Res. E Logist. Transport. Rev. 46 (5), 563–581.
Ray, S., Jewkes, E.M., 2004. Customer lead time management when both demand and price are lead time sensitive. Eur. J. Oper. Res. 153 (3), 769–781.
ReVelle, C.S., Eiselt, H.A., 2005. Location analysis: a synthesis and survey. Eur. J. Oper. Res. 165 (1), 1–19.
Rodriguez, M.A., Vecchietti, A.R., Harjunkoski, I., Grossmann, I.E., 2014. Optimal supply chain design and management over a multi-period horizon under demand uncertainty. Part I: MINLP and MILP models. Comput. Chem. Eng. 62, 194–210.
Schuster Puga, M., Tancrez, J.-S., 2017. A heuristic algorithm for solving large location–inventory problems with demand uncertainty. Eur. J. Oper. Res. 259 (2), 413–423.
Shahabi, M., Akbarinasaji, S., Unnikrishnan, A., James, R., 2013. Integrated inventory control and facility location decisions in a multi-echelon supply chain network with hubs. Network. Spatial Econ. 13 (4), 497–514.
Shahabi, M., Unnikrishnan, A., Jafari-Shirazi, E., Boyles, S.D., 2014. A three level location-inventory problem with correlated demand. Transp. Res. Part B Methodol. 69, 1–18.
Shen, Z.-J.M., 2007. Integrated supply chain design models: a survey and future research directions. J. Ind. Manag. Optim. 3 (1), 1–27.
Shen, Z.-J.M., Coullard, C.R., Daskin, M.S., 2003. A joint location-inventory model. Transport. Sci. 37 (1), 40–55.
Simpson, K.F., 1958. In-process inventories. Oper. Res. 6 (6), 863–873.
Smykay, E.W., 1973. Physical Distribution Management. Macmillan.
Sourirajan, K., Ozsen, L., Uzsoy, R., 2009. A genetic algorithm for a single product network design model with lead time and safety stock considerations. Eur. J. Oper. Res. 197 (2), 599–608.
Teimoury, E., Modarres, M., Monfared, A.K., Fathi, M., 2011. Price, delivery time, and capacity decisions in an m/m/1 make-to-order/service system with segmented market. Int. J. Adv. Manuf. Technol. 57 (1–4), 235–244.
You, F., Grossmann, I.E., 2010. Integrated multi-echelon supply chain design with inventories under uncertainty: MINLP models, computational strategies. AIChE J. 56 (2), 419–440.