



# Embedding the production policy in location-allocation decisions

Emilio Carrizosa<sup>1</sup> · Alba V. Olivares-Nadal<sup>2</sup> · Pepa Ramírez-Cobo<sup>3</sup>

Received: 16 January 2019 / Revised: 20 July 2019 / Published online: 21 November 2019  
© Springer-Verlag GmbH Germany, part of Springer Nature 2019

## Abstract

This paper investigates how the production policy, as well as other factors, affect the facility location-allocation decisions. We focus on a  $p$ -median location problem in which one single perishable product is to be produced and shipped to a set of users. The time-correlated demands of the clients are generated by autoregressive processes, and they are forecasted from historical data. Empirically, we show that: (i) embedding the production policy in the location-allocation decision problem may lead to a facilities-clients assignment which does not necessarily correspond to the minimum cost allocation, but produces better profits, (ii) taking into account the autocorrelation of the demand can significantly improve the performance of the supply chain, and (iii) the variability of the demand strongly affects the performance of the supply chain, so a careful choice of production strategy is especially recommended in this case.

**Keywords** Facility layout · Allocation · Robust newsvendor · Autoregressive processes · Forecasting

**Mathematics Subject Classification** 90B80 · 90C10 · 90B30

---

This research is supported by Projects P11-FQM-7603 and FQM329 (Junta de Andalucía) and MTM2015-65915-R (Ministerio de Economía y Competitividad, Spain), all with ERD Funds. The authors are also supported by the project “*Cost-sensitive classification. A Mathematical Optimization approach*”, financed by BBVA Foundation.

---

✉ Alba V. Olivares-Nadal  
alba.nadal@chicagobooth.edu

<sup>1</sup> Institute of Mathematics of the University of Seville (IMUS), Sevilla, Spain

<sup>2</sup> Booth School of Business, University of Chicago, 5807 S Woodlawn Ave, Chicago, IL 60637, USA

<sup>3</sup> Department of Statistics and Operational Research, University of Cádiz, Cádiz, Spain

## 1 Introduction

Demand uncertainty affects decisions in supply chain management at two levels: strategic (location decisions) and tactical (production and inventory management), see for example Snyder (2006) and Melo et al. (2009). Although tactical decisions depend on the choice of the production policy, the research devoted to analyzing the impact of such a choice on strategic decisions is lacking. The work presented in this paper aims to alleviate the scarcity on this research topic by empirically analyzing whether or not production policies should affect the design of the supply chain. Although similar studies such as Güllü (1997) and Chen et al. (2000) have been carried out previously, they analyze the benefits of considering demand autocorrelation in a supply chain which is already set. Throughout our manuscript, we will compare the conclusions attained in our empirical work not only against these inventory studies, but also against the numerical results of various joint location-inventory works (Snyder et al. 2007; Shen and Qi 2007; Coullard et al. 2002; Miranda and Garrido 2004; Zhang and Unnikrishnan 2016).

Traditionally, demand uncertainty has been addressed in the location and inventory literature by considering a scenario analysis or by assuming that either the demands distribution functions (cdf) or their population means and variances are given. This literature is vast and a thorough review is not possible in this paper, but the interested reader is referred to Snyder (2006) in the location field, Levi et al. (2008), Shang and Song (2003), Wang et al. (2012) and Shu et al. (2005) in the inventory area. A special mention is deserved by the joint location-inventory models proposed in Snyder et al. (2007), Shen and Qi (2007), Coullard et al. (2002), Miranda and Garrido (2004) and Zhang and Unnikrishnan (2016), which allow for non-perishable products and thus inventory stocking. Despite dealing with settings that may account for risk pooling or multisourcing, these studies assume normally distributed demands. This assumption agrees with the main stream of work for inventory problems (Liao et al. 2011; Shen and Qi 2007; Shen et al. 2003), although sometimes the demand may also be modelled through a Poisson process (Ozsen et al. 2009; Wang et al. 2007). However, all these approaches are contrary to the recommendations given by Bandi and Bertsimas (2012), who show that assuming a known cdf of the demand may be prohibitive and inaccurate in practice. Dong and Lee (2003) also warn that assuming independent identically distributed demands may cause an underestimation of the system cost and inventory levels. As a consequence, alternative robust methods considering ambiguous distributions for the demand have proliferated; see, for instance, Gulpinar et al. (2013).

Modelling the demand through a cdf allows one to consider correlations amongst clients (Shahabi et al. 2014; Atamturk et al. 2012). However it also implies that the demand values along time are assumed independent, which may be unrealistic in practice (Graves 1999; Kahn 1986). Hence, inventory and location models with time-dependent demand have often been studied in the literature; see, for instance, Dong and Lee (2003), Lu et al. (2006), Wang et al. (2012), See and Sim (2010), Aviv (2003), Dogan and Goetschalckx (1999) and Snyder et al. (2007). Güllü (1997) studied the effect of demand forecasting on allocation using a Martingale Model of Forecast Evolution, whereas Chen et al. (2000) analyzed the effect of demand forecasting technique on the bullwhip effect. In order to perform our empirical study we generate

time-autocorrelated demands following autoregressive processes, a common choice in the inventory field (see Aviv 2002; Reyman 1989; Johnson and Thompson 1975). More specifically, as in Ganesh et al. (2014), Raghunathan (2003) and So and Zheng (2003), we will generate our demand from an  $AR(1)$  process, although as in Carrizosa et al. (2016) we will allow for heavy tailed shocks to ensure more realistic instances.

Despite this choice for the demand model allowing us to account for temporal autocorrelation and heavy tailed demands, it also constricts our study to a more simple setting than that of the joint location-inventory models of Snyder et al. (2007), Shen and Qi (2007), Coullard et al. (2002), Miranda and Garrido (2004) and Zhang and Unnikrishnan (2016). Additionally, since we restrict our test environment to a basic discrete  $p$ -facility location model in which the open facilities are to produce one single perishable product, our approach lacks the ability to deal with inventory stocking. We are also constrained to opening  $p$  facilities, where  $p$  is a parameter known beforehand. In contrast to previous works, we make no distributional assumptions and future demand is assumed unknown, but its historical data is available and used to determine the production levels at each plant. The problem we are addressing accomplishes three tasks: select the plants to be opened, allocate clients to such facilities, and decide the quantity of perishable product each plant must produce to serve the clients at every single period. The later task is performed by the chosen production policy, which usually does not take part in the former steps. Our aim is to embed the production strategy in the location-allocation decision problem and compare how different policies affect such decisions. In order to do so, we design a data-driven procedure which estimates the coefficients of the  $p$ -median problem using various production policies. In our numerical study we focus on the three newsvendor strategies compared in Carrizosa et al. (2016); namely, the classic newsvendor approach, an autoregressive production procedure assuming normally distributed errors, and the so called *ARUS*. The latter is a robust autoregressive newsvendor approach proposed in Carrizosa et al. (2016), which is distribution-free and has shown to behave well under heavy tailed demands.

The paper is structured as follows. Next section develops a data-driven approach for a location model, where the coefficients of the  $p$ -median problem are estimated through the chosen production policy. This approach is used to analyze the effects of such a choice on the location-allocation decisions. The design of the experiments is addressed in Sect. 3. Specifically, the data sets used to illustrate the approach are introduced in Sects. 3.1 and 3.2 describes the stochastic processes used to generate the demand time series, and Sect. 3.3 reviews the production policies to be considered in the numerical experiments. The remainder of the section specifies the choices on the profitability of the products and the different allocation criteria. Section 4 presents the numerical results under two perspectives. First, we analyze the effects of the different production policies when considering the data-driven location-allocation solutions. Second, we compare the performance of the data-driven allocation with the minimum cost allocation. Finally, last section presents some final remarks and recommendations for future work.

## 2 A data-driven approach for embedding the production policy in the location-allocation problem

In this section we describe a data-driven procedure that bonds the location-allocation decisions with the chosen production policy. Due to the large number parameters considered, a single-item perishable product is assumed in order to ease the conclusions arisen on this paper. In effect, no warehouses or stock levels are considered. Let  $I$  and  $J$  be the sets of potential locations for the facilities to be opened and the clients, respectively. Let  $p$  denote the total number of facilities the company wishes to open, where the problem to be solved encompasses locating  $p$  facilities amongst the candidate locations to serve all customers, whose demand values  $\{D_j(t)\}_{t=1}^T$  are available up to time  $T$ . For simplicity, such demand values are assumed to be independent amongst clients. Along with decisions related to the location of facilities and allocation of clients, the quantity of products  $Q_{ij}(t)$  to be produced at each location  $i$  to supply client  $j$  at each time instant  $t$  is to be determined. As the considered product is perishable, only one-period-ahead production levels need to be decided. The working mechanism of the considered supply chain is as follows. First, produce and ship  $Q_{ij}(t)$  from facility  $i$  to customer  $j$  at time instant  $t$ . Then, let the customer  $j$  buy only  $\min\{D_j(t), Q_{ij}(t)\}$ . For the sake of comprehension, we next present the notation used throughout this paper.

$p$	number of facilities to open
$I$	set of candidate locations for the facilities
$I_0$	set of opened facilities
$J$	set of clients
$D_j = \{D_j(t)\}_{t=1}^T \quad j \in J$	time series of demands of clients up to time $T$
$Q_{ij}(t) \quad i \in I, j \in J, t > 0$	the quantity of product to produce at location $i$ to supply client $j$ , at time $t$
$C^S(t) = (c_{ij}^S)(t)$	shipping costs matrix: $c_{ij}^S$ represents the cost of sending one unit of product from facility $i$ to client $j$ at time instant $t$
$C^P(t) = (c_i^P)(t)$	vector of production costs: $c_i^P$ represents the cost of producing one unit of product at plant $i$ at time instant $t$
$C(t) = (c_{ij})(t)$	total cost matrix: $c_{ij}(t) = c_{ij}^S(t) + c_i^P(t)$ at time instant $t$
$v_j(t)$	sale unit price of the perishable product for client $j$ at time instant $t$
$T + 1, \dots, T + r$	planning horizon

If  $b_{ij}(t, D_j(t))$  denotes the benefit of serving client  $j$  from facility  $i$  at time  $t$ , under the assumption of a known demand  $D_j(t)$ , then the benefit of serving client  $j$  from facility  $i$  at each period of the planning horizon  $T + 1, \dots, T + r$ , is given by

$$B_{ij}(T + 1, T + r, D_j) = \frac{1}{r} \sum_{t=T+1}^{T+r} b_{ij}(t, D_j(t)) \tag{1}$$

If the demand is known, we set  $Q_{ij}(t) = D_j(t)$  and hence coefficients  $b_{ij}(t, D_j(t))$  and the production levels  $Q_{ij}(t)$  are deterministic. If that were the case, the location-allocation problem to be solved would be formulated as the  $p$ -facility problem of maximizing total profit (1):

$$\begin{aligned} \max \quad & \sum_{j \in J} \sum_{i \in I} B_{ij}(T + 1, T + r)x_{ij} \\ \text{s.t.} \quad & \begin{cases} \sum_{i=1}^I x_{ij} = 1 & \forall j \in J \\ x_{ij} \leq y_i & \forall i \in I, j \in J \\ \sum_{i=1}^I y_i = p \\ x_{ij} \in \{0, 1\} & \forall i \in I, j \in J \\ y_i \in \{0, 1\} & \forall i \in I \end{cases} \end{aligned} \tag{2}$$

where the variables represent

$$\begin{aligned} y_i &= \begin{cases} 1 & \text{if facility } i \text{ is opened} \\ 0 & \text{otherwise} \end{cases} \\ x_{ij} &= \begin{cases} 1 & \text{if client } j \text{ is supplied by facility } i \\ 0 & \text{otherwise} \end{cases} \end{aligned}$$

However, since the values of  $\{D_j(t)\}_{t>T}$  are not available in practice, the benefit coefficients  $b_{ij}(t, D_j(t))$  need to be estimated. To this aim, we propose a data-driven distribution-free recursive approach, where no probabilistic model for the demand is assumed. To this aim, we propose a novel method taking into account the production policy to estimate the coefficients of the location-allocation problem. The procedure starts by dividing each client’s historical demand, given by a time series, into two sets: the training set (periods  $1, \dots, T_1$ ) and the testing set (periods  $T_1 + 1, \dots, T$ ). After choosing a specific production policy, the training set is used to estimate its parameters. Then, the test set is employed to estimate the benefit coefficients  $b_{ij}(t, D_j(t))$  (and therefore, the total benefit as in Eq. 1). This procedure is explained in more detail in Table 1.

Once the total benefit coefficients are estimated according to the procedure in Table 1, the location-allocation problem to be solved is as given in (2), where the benefit coefficients  $B_{ij}(T + 1, T + r)$  have been replaced by their estimates  $\hat{B}_{ij}(T + 1, T + r)$ . Note that (2) is a  $p$ -median problem in which the costs (or, equivalently, the benefits  $\hat{B}_{ij}$ ) do not only depend on the transportation and production costs ( $c_{ij}^S$  and  $c_i^P$ ), but also depend on  $Q_{ij}$ . Since the production levels are obtained through the production policy, erroneous decisions made by this policy may strongly perturb the coefficients  $\hat{B}_{ij}$ . In these cases, the location and allocation decisions obtained by (2) may significantly differ from the solution of the standard  $p$ -median problem. This phenomenon will be clearly illustrated in Sect. 4.

**Table 1** Recursive approach for estimating the parameters of the  $p$ -median problem

1. Split the demand time series in training  $\{D_j(t)\}_{t=1}^{T_1}$  and test  $\{D_j(t)\}_{t=T_1+1}^T$  sets for all clients  $j \in J$
2. For  $t_0 = T_1 + 1, \dots, T$  and for all  $j \in J$  repeat:
  - (a) Use the known values of the demand  $\{D_j(t)\}_{t=1}^{t_0-1}$  to fit the parameters of the selected production policy
  - (b) Use the production policy to calculate the quantity of product to send to each client at each instant of time from every plant  $i \in I$ ,  $Q_{ij}(t_0)$
  - (c) Using the real value for the demand,  $D_j(t_0)$ , calculate the benefit coefficients  $b_{ij}(t_0)$  as the obtained profit
 
$$b_{ij}(t_0) = v_j(t_0) \cdot \min\{D_j(t_0), Q_{ij}(t_0)\} - c_{ij}(t_0) \cdot Q_{ij}(t_0) \quad (3)$$
3. Estimate the benefit coefficients per period in the planning horizon as

$$\hat{B}_{ij}(T+1, T+r) = \frac{1}{T-T_1} \sum_{t=T_1+1}^T b_{ij}(t) \quad (4)$$

### 3 Description of the experiments

The aim of the empirical study we develop in this paper is two fold. First, we want to test how the choice of production strategy affects the design and management of the supply chain. The performance of the supply chain will not only depend on factors such as the demand's autocorrelation and variability, but also on the suitability of the chosen production policy to deal with the behaviour of the demand. Second, we compare the data-driven allocation against the classic minimum cost allocation because, as although it has been reported to be very restrictive (Weaver and Church 1986), the latter still constitutes the benchmark criterion. Despite the data-driven allocations providing better in-sample profit, we aim to check whether higher out-of-sample profit could be obtained if the clients were served by their closest open facility.

In order to accomplish our goals, we need to design an extensive empirical study. As a consequence, a large number of parameters, features and choices will be involved. We devote this section to describe in detail the design of the numerical experiments, whose results are presented in Sect. 4. Specifically, the data sets, stochastic models for the demand, production policies, product's profitability and allocation criteria are discussed here. Table 5, placed in the appendices, summarizes the numerical settings. In particular, the choices regarding the location problem (the number of facilities to open and the datasets) have been roughly made following Daskin and Owen (1999). The decisions more closely related to the production (generation of the demand, production policies, product profitability) have been based on the work presented by Carrizosa et al. (2016). We focus on time correlated demands because they are more realistic, and as shown by Helper et al. (2010), Güllü (1997) and Chen et al. (2000), taking into account the temporal autocorrelation can have a significant effect on the performance of the supply chain. Since it is well known that the variability of the demand is one of the key features that affect the performance of the supply chain (see, for instance, Lambert and Cooper (2000), Cachon (1999) and Germain et al. (2008)), we generate

these demands with different variabilities. In addition, as shown in Carrizosa et al. (2016), the profitability of the products may significantly affect the behaviour of the production policies, and thus the performance of the supply chain. Finally, we will also consider two different allocations: the data-driven allocation attained to maximize the in-sample profit (described in the previous section), and the classic minimum cost allocation.

### 3.1 Data sets

The data sets used in the experiments developed here closely follow those in Daskin and Owen (1999): the 55-node data set of Swain (1971) and the 88-node data set of Daskin (1995). The 88-node data set contains the coordinates of the fifty largest cities in the United States along with the state capitals, and its distance matrix was calculated using great circle distances. The 55-node data set of Swain represents 55 communities in the Washington D.C. area. This dataset has been used in ambulance problems (van den Berg 2016), and corresponds to a graph where the nodes are the users and the distances are the the closest paths between nodes. Average demands for the U.S. cities data set were given, and for the nodes of the Swain network were randomly generated following a uniform distribution as in Blanquero et al. (2014).

For the two databases used in our numerical experiments, we assume that the sets of clients and candidate locations are the same; that is to say,  $I = J$ . The number of facilities to be opened was chosen according to Daskin and Owen (1999), where they tested the two data sets for  $p = 3, \dots, 15$ . For the sake of brevity, we only considered the minimum, maximum and median of such set, therefore we tested the approach for three different values of  $p$ , namely  $p = 3, 9, 15$ .

### 3.2 Stochastic models for the demand

Since we aim to study how the autocorrelation and variability of the demand affects the location-allocation decisions, in our numerical results we restrict ourselves to the case where the demand follows an (uncertain) autoregressive process. Recall that the demand has often been modeled as an autoregressive process (Aviv 2002; Reyman 1989; Johnson and Thompson 1975). However, any other stochastic process considering correlated demands may also be tested by the proposed data-driven approach. The following constitutes a brief introduction to the autoregressive process, where a more detailed description is given by Box et al. (2008).

The demand time series of client  $j$ ,  $\{D_j(t)\}_{t>0}$ , follows an autoregressive process of order  $n \geq 1$  (noted  $AR(n)$ ) if it can be written as

$$D_j(t) = \alpha^j + \sum_{k=1}^n \theta_k^j D_j(t - k) + a_j(t), \quad t \geq n + 1 \tag{5}$$

where  $\alpha^j, \theta_1^j, \dots, \theta_n^j$  are coefficients and  $\{a_j(t)\}_{t>0}$  is the sequence of i.i.d model's error terms. In the real data sets presented in Sect. 3.1, the nodes' demands  $\{D_j(t)\}_{t=1}^{T+r}$  were unavailable (only the average demands were known) and therefore we needed to

simulate them artificially. In our experiments they were generated from autoregressive processes as in (5) with  $n = 1$ , in such a way that the expected value of the series matches the average known values. For simplicity, we choose  $\alpha^j = \alpha$  and  $\theta_k^j = \theta_k$  for all  $j \in J$ . The choice  $n = 1$  follows Carrizosa et al. (2016), from which we borrow the newsvendor policies under comparison in our numerical study. Within that paper, only results for  $n = 1$  were included, due to the performance of the production policies being virtually unaffected by changing the value of  $n$ . In our experimental study, two different distributions for the error terms are tested: the normal and the lognormal distributions. First we considered standard normally distributed errors, as we aimed to test the performance under the commonly considered  $AR(n)$  model. In addition, we found it of interest to check the effects of the production strategies under a high variability in the demand time series, as done in Huh et al. (2011) and Carrizosa et al. (2016). To such aim, the lognormal distribution  $LN(0, 3)$  was selected. Another aspect to be considered when simulating the data was the strength of the temporal dependence. In our experiments, two values of  $\theta_1 = \theta$  were set. Note that for  $n = 1$  the coefficient  $\theta_1$  represents the lag-1 autocorrelation coefficient, thus in order to test the methods on highly and moderately correlated time series,  $\theta_1 = 0.9$  and  $\theta_1 = 0.5$  were fixed. Figure 5, located in the Appendix, shows simulated demands time series generated under the two choices of error distributions and autocorrelation strengths.

In our experiments 25 instances were simulated for each data set, where series of length  $T + r = 288$  following an  $AR(1)$  process with i.i.d errors were generated for each node. The training set, which is used to fit the parameters of the production policy (see Table 1), is defined for time instants  $t = 1, \dots, 240$ . The test set, used to estimate the benefit coefficients of the location-allocation problem, considers the demand from  $t = T_1 + 1 = 241$  to  $t = T = 264$ . The validation set, used for calculating the out-of-sample profit, is composed of the rest of time series values  $\{D_j(t)\}_{t=T+1}^{T+r}$ , so the training set has a length equal to 240 and both the test and validation sets are of length 24. To manage the supply chain, each production policy is ran in the validation set, and the location ( $\{y_i\}_{i \in I}$ ) and allocation ( $\{x_{ij}\}_{i \in I, j \in J}$ ) solutions were obtained from the training and test sets. The total profit (1) is computed similarly as in Table 1 by using the validation set of each instance.

### 3.3 Single-period production policies

Traditionally, the production policy takes part once the set  $I_0$  of open facilities and the clients-facilities allocation have already been chosen. At each time period  $t$  the production level of each plant  $i \in I_0$  is determined in order to cover the demand of customer  $j \in J$ ,  $D_j(t)$ . In practice, the future demand  $D_j(t)$ , with  $t > T$  is unknown and only the historical data for  $t = 1, \dots, T$  is available. Although the purpose of the production strategy is usually restricted to determine such production levels, in our numerical experiments these strategies will also influence the location-allocation decisions, so a careful choice is needed. As we want to test both classic newsvendor approaches and strategies coping with time correlated demands, we consider the three production policies tested in Carrizosa et al. (2016). Also, we include here a benchmark approach in which the future demand is known.



### 3.3.1 Benchmark

In the benchmark approach we assume a perfect knowledge of the demand at each period. Hence, clients are supplied with their exact demand; i.e.:

$$Q_{ij}(t) = D_j(t).$$

As mentioned in Sect. 2, in this case the benefit coefficients (1) taking part in the  $p$ -median problem (2) are deterministic.

### 3.3.2 Static

In the classic version of the newsvendor problem, one makes a one-step decision on the quantity  $Q$  to be bought of one single product under the assumption that the demand is a random variable with known cdf  $F$ . If the decision maker buys each unit at cost  $c$  and sells it at price  $v$  then the expected profit is maximized by buying exactly  $Q = F^{-1}(1 - c/v)$  units (Khouja 1999). In our context, this implies that  $Q_{ij}(t)$ , which is the quantity of product to produce at plant  $i$  to deliver to client  $j$  at time instant  $t$ , is given by the expression:

$$Q_{ij}(t) = F^{-1}\left(1 - \frac{c_{ij}(t)}{v_j(t)}\right) \tag{6}$$

However, in practice the distribution of the demand is unknown and therefore in this approach an estimation of the distribution function must be employed. Usually, the empirical distribution function  $\hat{F}$ , which converges to the true cdf  $F$  for a large enough sample size, is considered. Note that under this approach the temporal dependence of the demand is ignored, thus this method will be called *static* hereafter.

### 3.3.3 The classic AR(n) approach

If the parameters in (5) are given (either known or estimated from sample data), then (5) can be used to forecast the demand. In particular, if the errors are assumed to follow a normal distribution, then the  $(1 - \tau)\%$  confidence prediction interval for a forecasted value at time  $T + 1$  is given by

$$\hat{D}_j(T + 1) \pm z_{1-\tau/2}\sigma_a \tag{7}$$

where  $\hat{D}_j(T + 1)$  is the estimated forecast, and  $z_{1-\tau/2}$  is the  $(1 - \tau/2)^{th}$  quantile of the standard normal distribution. Equivalently, under (5)  $D_j(T + 1)$  follows a normal distribution with mean  $\hat{D}_j(T + 1)$  and variance  $\sigma_a$ . Since  $\sigma_a$  is unknown in practice, it is replaced by its sample estimate  $\hat{\sigma}_a$ . From a similar argument as in (6), in the case of the classic  $AR(n)$  approach, the optimal solution to the newsvendor problem (if the expected profit is maximized) is

$$Q'_{ij}(t) = \Phi^{-1}_{(\hat{D}_j(T+1), \hat{\sigma}_a)} \left( 1 - \frac{c_{ij}(t)}{v_j(t)} \right), \tag{8}$$

where  $\Phi^{-1}$  denotes the inverse cdf of the normal distribution. Since  $Q'_{ij}(t)$  may be negative, the quantity  $Q_{ij}(t) = \max \{0, Q'_{ij}(t)\}$  will be considered instead.

If errors are not normally distributed, (7) may lead to inaccurate results (Carrizosa et al. 2016). This motivates the use of the robust inventory policy presented next.

### 3.3.4 The robust ARUS(n) approach

The *ARUS(n)* (Carrizosa et al. 2016) is a robust autoregressive newsvendor approach formulated in terms of a convex optimization problem. Robustness is addressed by *ARUS(n)* in two ways. First,  $Q_{ij}(t)$  is the solution of a maximin optimization problem. Second, the solution obtained by the *ARUS(n)* is distribution-free, unlike the classic *AR(n)* approach. By the use of uncertainty sets (Bandi and Bertsimas 2012) the method provides a solution for the newsvendor problem without making any distributional assumption.

More specifically, the value  $Q_{ij}(T + 1)$  will correspond to the worst case for the demand of client  $j$  at time  $T + 1$ ,  $D_j(T + 1)$ . Therefore,  $Q_{ij}(T + 1)$  is the minimizer of the so-called *ARUS(n)* problem, defined as

$$\min_{\alpha^j, \theta_1^j, \dots, \theta_n^j, a_j(1), \dots, a_j(T+1)} D_j(T + 1) \tag{9}$$

s.t

$$\left| \frac{1}{T - n} \sum_{t=n+1}^T a_j(t) \right| \leq \Gamma_1^j, \tag{10}$$

$$\frac{1}{T - n} \sum_{t=n+1}^T a_j^2(t) \leq \Gamma_2^j, \tag{11}$$

$$D_j(t) = \alpha^j + \sum_{k=1}^n \theta_k^j D_j(t - k) + a_j(t) \quad t = n + 1, \dots, T + 1, \tag{12}$$

$$|a_j(T + 1)| \leq \Delta^j \tag{13}$$

$$D_j(T + 1) \geq 0 \tag{14}$$

The rationale behind the constraints in the *ARUS(n)* problem is the following. Constraints (10) and (11) force both the mean and variance of the observed errors to be bounded, so as to bound the randomness of the variables. In (12) the consecutive values of the process are expressed in terms of the previous  $n$  values, according to the definition of the *AR(n)* process in (5). Finally, (13) implies that the absolute value of the random variable  $a_j(T + 1)$  (which represents the prediction error) is upper bounded by some constant  $\Delta^j$ .

By solving problem (9)–(14) a non-overly conservative yet robust solution for the production decision is chosen. However, there is a need for tuning the parameters

of problem (9)–(13) before solving it. As in Carrizosa et al. (2016), the choice of parameters is given by:

$$\hat{P}(a_j(t) \leq \Delta^j) = \left( \frac{c_{ij}(t)}{v_j(t)} \right)^2 \tag{15}$$

$$\Gamma_1^j = \frac{\Gamma_1^{j*} \sigma_a^j (1 + v/100) \sqrt{T - p}}{T - p} \tag{16}$$

$$\Gamma_2^j \approx (1 + \beta/100) \gamma_2^j \tag{17}$$

where  $\hat{P}$  is the empirical probability function of the errors,  $\sigma_a^j$  is estimated from the minimum variance of the errors imposing that the series follows an autoregressive process, and  $\gamma_2^j$  is obtained similarly. The full motivation for the choice of the unknown parameters can be found in Carrizosa et al. (2016).

### 3.4 Product’s profitability

Note that the quantity of product  $Q_{ij}(t)$  to produce at plant  $i$  to supply client  $j$  estimated via the production policies may depend on the total cost  $c_{ij}(t)$  and on the sale price  $v_j(t)$ . Hence it is expected that the supply chain performance is affected by the profitability of the products, as are the production policies themselves. Therefore, two types of profitability settings were considered for the perishable product: low- and high-profitability. For the sake of brevity we fix  $v_j(t) = v$  for all clients and time instants  $t$ , and  $c_{ij}(t) = c_{ij}$ . For low-profit products, the sale price was established as

$$v = 1.25 \max_{\substack{i=1,\dots,I \\ j=1,\dots,J}} \{c_{ij}^S, c_i^P\}, \tag{18}$$

where  $c_{ij}^S$  are the normalized transport costs, which are derived by normalizing the distance matrix, and  $c_i^P$  denotes the normalized production costs. Note that, in this case, the sale price is 25% higher than the largest cost. The high-profit products are defined as in (18) but with a higher value:

$$v = 1.75 \max_{\substack{i=1,\dots,I \\ j=1,\dots,J}} \{c_{ij}^S, c_i^P\}. \tag{19}$$

Hence, the sale price for high-profit products is 75% larger than the greatest cost.

### 3.5 Allocation criteria

In our numerical experiments, the different production policies outlined in Sect. 3.3 serve not only to determine the production levels of the plants, but they are also used to estimate the coefficients of the  $p$ -median problem (2). However, as noted in Shen et al. (2003) and Coullard et al. (2002), allocation of clients to their closest facility

may yield suboptimal solutions. Hence, although throughout our numerical results we always use the location provided by Problem (2), we decided to compare the profit obtained by each policy under two different allocation strategies:

- The data-driven maximum-profit allocation, i.e., the one given by the optimal values of variables  $x_{ij}$  when (2) is solved. We denote by  $r_t^{ps}$  the earnings obtained by this data-driven maximum-profit allocation at time instant  $t$ .
- The minimum-cost allocation, i.e., the one that allocates the clients to their minimum cost open facility ( $x_{ij} = 1$  if and only if  $i = \arg \min_{i \in I_0} c_{ij}$ ). The profit harvested by the minimum-cost allocation at time instant  $t$  is denoted by  $r_t^{mc}$ .

Although the data-driven maximum-profit allocation is designed to obtain the best in-sample profit, it may result in seemingly counterintuitive allocations. Hence, it is reasonable to wonder whether the traditional minimum cost allocation might provide better out-of-sample profit, and this analysis is given in Sect. 4.2. It will be shown that the allocation decisions under the proposed data-driven approach do not necessarily coincide with those of the minimum cost strategy, but usually yield better out-of-sample profit. Also, the number of clients optimally allocated with minimal cost depends on the production policy used to design the supply chain.

## 4 Numerical results

In this section, we empirically test the impact of the production policy, as well as other factors, over the location-allocation decisions. In Sect. 4.1 we analyze the influence of the production strategy, variability and autocorrelation in the demand, type of profitability and geographical position of the nodes under the data-driven maximum profit criterion. Section 4.2 will examine in detail the effect of considering instead the minimum cost allocation. All the optimization problems were solved using Xpress 7.2.

### 4.1 Results under the maximum profit criterion

Figures 1 and 2 show the box plots of the 25 total out-of-sample profits for all considered location-production approaches under the maximum profit criterion. The demand in all cases is generated from an  $AR(1)$  process with normal  $N(0, 1)$  and lognormal  $LN(0, 3)$  errors, respectively. Note that this implies that a low (respectively, high) dispersion in the realizations of the demand is expected in Fig. 1 (respectively, Fig. 2). Our findings showed few differences between the behavior of the location approach under different number of plants to be opened. Therefore, for brevity, only the case with  $p = 3$  is illustrated. The location-production strategies are named according to the associated production policy and the demand's autocorrelation, i.e., for instance *STATIC 0.9* denotes the classic newsvendor policy (6) for demand with a 0.9 lag-1 autocorrelation.

The results for  $N(0, 1)$  shocks are presented in Fig. 1. In this case, the profit samples concentrate in an interval of small length, especially for low autocorrelated demands. For the Swain database (bottom part of the figure) total profits are rather insensitive to the production policy for a given value of  $\theta$ . Obviously, the largest profits are obtained

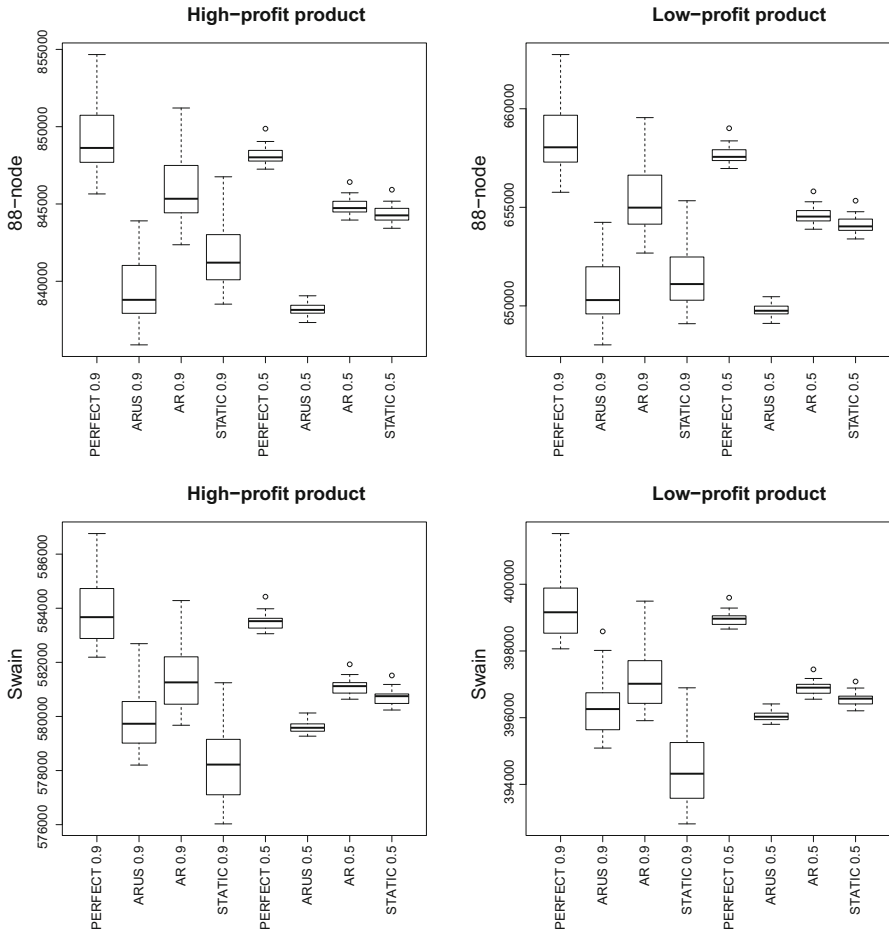


Fig. 1 Out-of-sample profits for autoregressive demands with  $N(0, 1)$  shocks

under the benchmark approach as perfect knowledge of the demand is assumed. Concerning the static approach, a better performance is achieved for  $\theta = 0.5$ . This is unsurprising since this strategy disregards time autocorrelation. Slightly better values, close to those of the benchmark approach, are obtained when the classic *AR* model is considered. This is also a reasonable result since the *AR* process assumes normal errors. The conclusions are similar for the 88-node data set (top part of Fig. 1), though much worse results are obtained for *ARUS*. In conclusion, when the demand variability is low the results are highly dependent on both the spatial distribution of the nodes (since the obtained results are different for the two data sets) and the autocorrelation in the demand. However, the effect of the production strategy and the type of profitability seems to be weaker.

Now we analyze the boxplots for Lognormal shocks (Fig. 2). As expected, the heavy-tailed behavior in the demand implies a higher variability in the profits. In

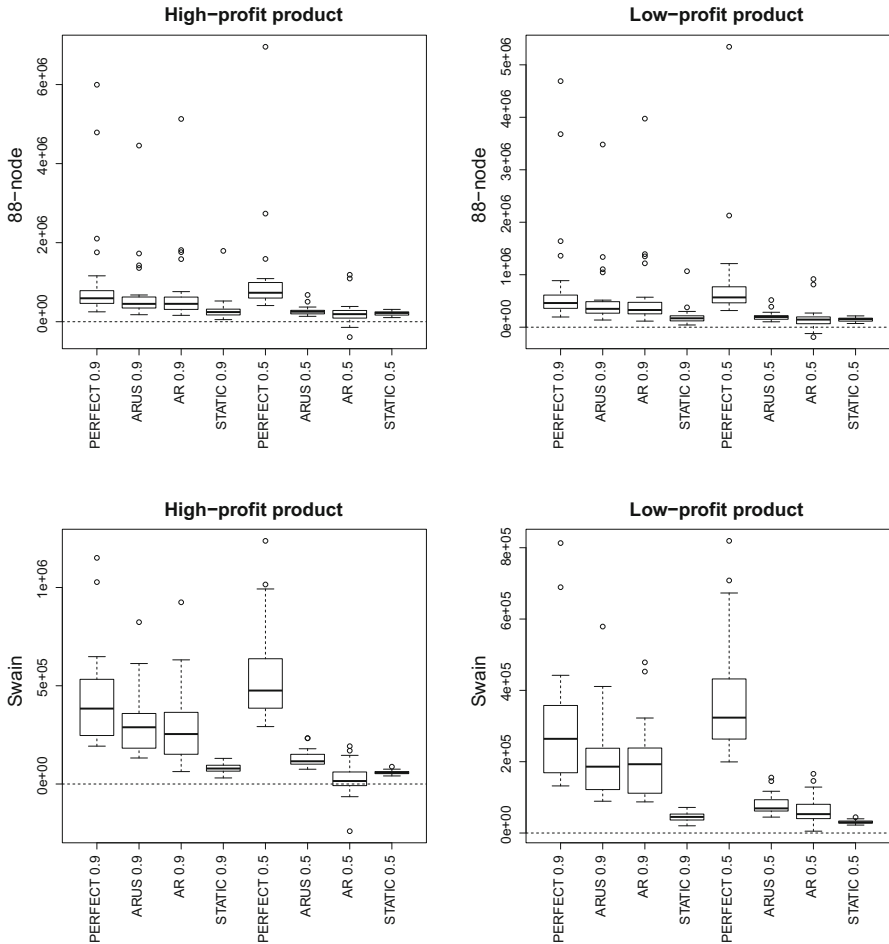


Fig. 2 Out-of-sample profits for autoregressive demands with  $LN(0, 3)$  shocks

particular, the results for the 88-node data set are more dispersed than those for the Swain network, especially for high-profit products. This variability increases with the autocorrelation parameter  $\theta$ . Consider now the effect of the different production policies on the profits. Once again, better profits are obtained under the benchmark approach, this outperformance being more evident for low correlated demands. The performance of the static strategy is similar in all cases. Although the robust approach never obtains losses, the classic *AR* produces negative total profits when  $\theta = 0.5$  for high-profit products, and also for low-profit products on the 88-node database. The effect of the profitability seems to be negligible for the 88-node data sets. However, for the Swain data the type of profitability affects the results under the *AR* strategy. Therefore, when extreme values for the demand are considered the performance of the supply chain is not only affected by the spatial distribution of the nodes and the autocorrelation of the demand, but also by the production policy.

**Table 2** Average out-of-sample improvement by using the minimum cost allocation for the **Swain** database

Profitability	Production	$\theta = 0.9$			$\theta = 0.5$		
		$p = 3$	$p = 9$	$p = 15$	$p = 3$	$p = 9$	$p = 15$
High	STATIC	2.62	1.91	1.24	6.18	8.06	8.94
	AR	-18.53	-18.62	-17.61	-8002.84	-1059.04	-938.30
	ARUS	0.98	1.30	1.27	2.15	2.30	2.51
Low	STATIC	3.11	4.11	4.12	6.79	8.42	8.88
	AR	-11.55	-17.61	-17.15	-374.73	-406.14	-380.50
	ARUS	1.62	0.94	0.75	2.84	2.83	3.11

These results are consistent not only with Carrizosa et al. (2016), who showed that the classic newsvendor production policy is not appropriate under highly auto-correlated demands, but also with Aviv (2003), who highlighted the benefits of using previous demand signals to estimate inventory levels under time autocorrelated demands. Unlike Raghunathan (2003), who concluded that increasing the autocorrelation would diminish variance in the demand forecasts and production levels, it seems that the variability of the data-driven profits increases with correlation. A different conclusion was reported by Lu et al. (2006), who observed that lowering the correlation enhanced the performance of the inventory strategy. We can agree with such affirmation if the performance is measured in terms of variance, since the median profit seems to be unchanged, but the higher the autocorrelation, the higher the dispersion of the data-driven profits.

### 4.2 The effect of the allocation

As mentioned in Sect. 3.5, we are interested on comparing the effect of the production strategy on the allocation between facilities and clients. The results shown in the previous section were obtained under the data-driven allocation, that yields the maximum in-sample profit according to the production policy. In this section, we also assign clients to the minimum cost open facility. Note that the minimum cost allocation may produce better out-of-sample profits than the so-called maximum profit allocation, because the later is optimal in-sample. Hence, for comparison purposes we now manage the supply chain under both allocations.

Tables 2 and 3 report the average percentage of improvement by using the minimum cost allocation instead of the maximum profit allocation for the Swain and 88-node datasets, respectively. Recall that, in Sect. 3.5 we explained the difference between both criteria and denoted the minimum cost and the maximum profit allocations at time instant  $t$  by  $r_t^{mc}$  and  $r_t^{ps}$ , respectively. Hence, the average percentage of improvement is given by

$$\overline{\Delta r} = \sum_{t=265}^{288} \frac{(r_t^{mc} - r_t^{ps}) / r_t^{ps}}{24} \cdot 100 \tag{20}$$

**Table 3** Average out-of-sample improvement by using the minimum cost allocation for the **88-node** database

Profitability	Production	$\theta = 0.9$			$\theta = 0.5$		
		$p = 3$	$p = 9$	$p = 15$	$p = 3$	$p = 9$	$p = 15$
High	STATIC	-13.45	-17.81	-18.31	1.29	1.81	2.42
	AR	-10.40	-8.94	-8.17	-123.60	-133.66	-178.07
	ARUS	0.10	0.08	0.12	0.08	0.07	0.08
Low	STATIC	-10.66	-11.89	-13.93	1.31	1.55	1.95
	AR	-11.71	-11.14	-10.63	-201.53	-146.40	-131.90
	ARUS	0.11	0.10	0.14	0.09	0.08	0.10

In order to test the effect of the number of plants to be opened, three values of  $p$ ,  $p = 3, 9, 15$ , were considered in the experiment. Since the the results for normally distributed errors were not strongly dependent on the production policy (see Sect. 4.1), we only report here the results for lognormally distributed shocks.

From Tables 2 and 3 we observe that the worst results are obtained under the classic autoregressive production policy. For the *AR*, the minimum cost criterion always worsens the total profit of the supply chain, especially for the Swain database. Note that in this case, when  $p = 3$ ,  $\theta = 0.5$  and for high profitability, the obtained loss is greater than 8000%. This disagrees with Zhang and Unnikrishnan (2016) results, which observed that a larger number of facilities would increase the system costs. Note also that the average percentage of improvement gets worse as the autocorrelation of the demand decreases. The *ARUS* is the least affected production strategy since its average percentage of improvement varies between 0.07 and 3.11%. On the other hand the results under the static newsvendor policy are more variable: for the Swain network it is the policy that best responds to the change of allocation, but for the 88-node database and a high autocorrelation this change results in a worsening between 10.66 and 18.31%. These results contradict Coullard et al. (2002), whose numerical experiments showed that reallocating clients to their least cost facility would not much impact the profits.

In conclusion, although a minimum cost allocation may be thought a sensible criterion, the previous findings show that it can lead to worse out-of-sample profits than those obtained under the maximum profit strategy. This conclusion agrees with Shen et al. (2003) and Coullard et al. (2002), who claim suboptimality of closest allocation when inventory costs take part. Moreover, substantial changes in the results are observed for the different data sets, production policies and autocorrelation values. However, the behavior under different values of  $p$  does not seem to be affected. Once again, this conclusion differs from the affirmation in Zhang and Unnikrishnan (2016), who found out that the higher the number of the facilities, the larger the system costs.

In order to better understand the reason why the minimum cost allocation criterion impacts differently the performance of supply chains designed by different production policies, we report in Table 4 the average percentage of clients (for  $p = 3, 9, 15$ ) that were assigned with minimum cost when the maximum profit criterion was assumed.



**Table 4** Average percentage of clients allocated to the minimum cost open facility

Profitability	Production	Swain		88-node	
		$\theta = 0.9$	$\theta = 0.5$	$\theta = 0.9$	$\theta = 0.5$
High	STATIC	66.57	81.96	71.95	81.52
	AR	64.80	26.74	89.59	67.80
	ARUS	93.19	88.82	96.83	96.82
Low	STATIC	68.78	84.51	71.32	83.12
	AR	56.19	18.62	86.89	59.92
	ARUS	92.02	87.10	96.74	96.50

It is expected that the higher this percentage is, the less impact the minimum cost allocation criterion has.

Frequently, for the 88-node database more clients are allocated to the minimum cost facility than for the Swain network. Also the robust approach *ARUS* is the one which assigns more clients with minimum cost. It is interesting to note that, unlike the autoregressive approaches *AR* and *ARUS*, the performance of the static forecasting technique improves when the autocorrelation decreases. For the Swain network, *AR* never assigns more than 65% of clients to their minimum cost facility. Moreover, for low-correlated demand this percentage strongly decreases: for low-profit products only 18.62% of clients are assigned with minimal cost. This case corresponds to the worst average percentage of improvement derived from applying minimum cost allocation instead of the original assignment of clients in Table 2.

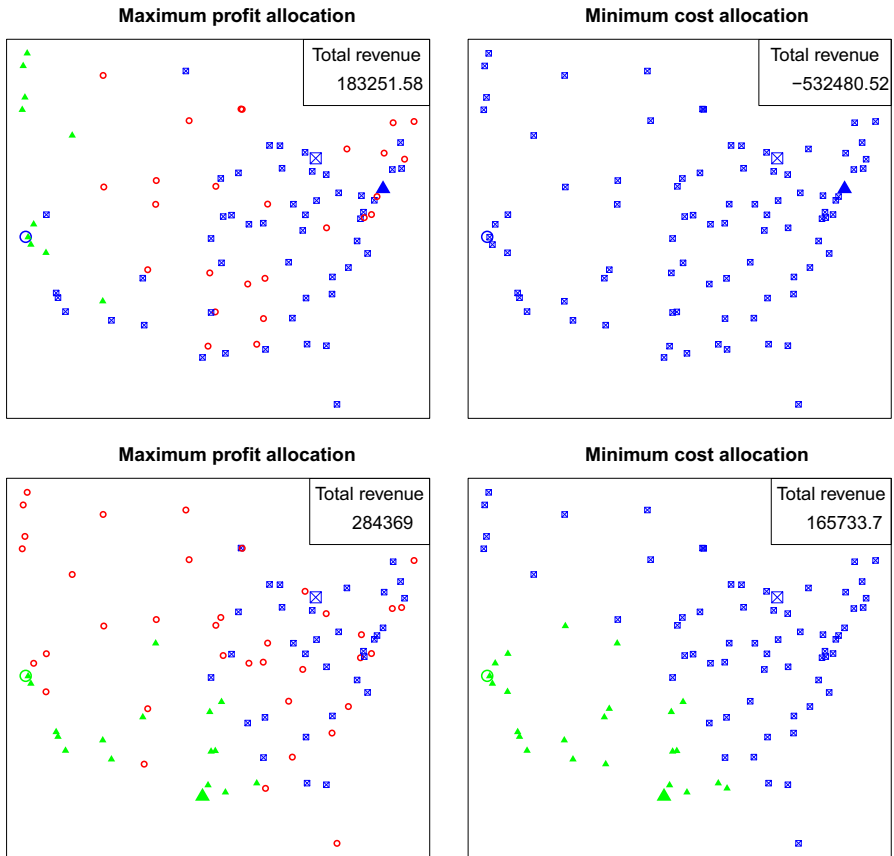
These results may be a consequence of the behavior observed in Carrizosa et al. (2016): for demands with high variability, the *AR* approach performs better for low-profit products than for high-profit products. Hence, the *AR* production strategy diminishes the profitability by opening facilities with high production costs and supplying clients with high serving costs. This will usually lead to a better joint performance of the location-production solution.

As an illustration of this phenomenon, we depict in Figs. 3 and 4 the locations of the opened plants for the case  $p = 3$  and the assignment of clients in the 88-node data set. Two instances have been randomly selected (top and bottom panels) and we depict both allocation criteria (left and right panels). In all cases the demand is low-correlated and heavy-tailed, and final total profit is displayed for each case. Figures 3 and 4 show the results for high and low profit product, respectively.

The three open facilities are represented within the figures by either a large solid green triangle, an empty red circle or a crossed blue square. Their assigned clients have been represented by smaller but analogous symbols with the same color. It may be the case that an open facility is not served by itself but by another open plant. In such a case, only the color (not the symbol) changes to denote which plant will supply it.

Figures 3 and 4 show a number of interesting issues. In particular:

1. First, the minimum cost allocation is usually suboptimal (agreeing, for instance, with Shen et al. (2003) and Coullard et al. (2002)). Second, in some cases the

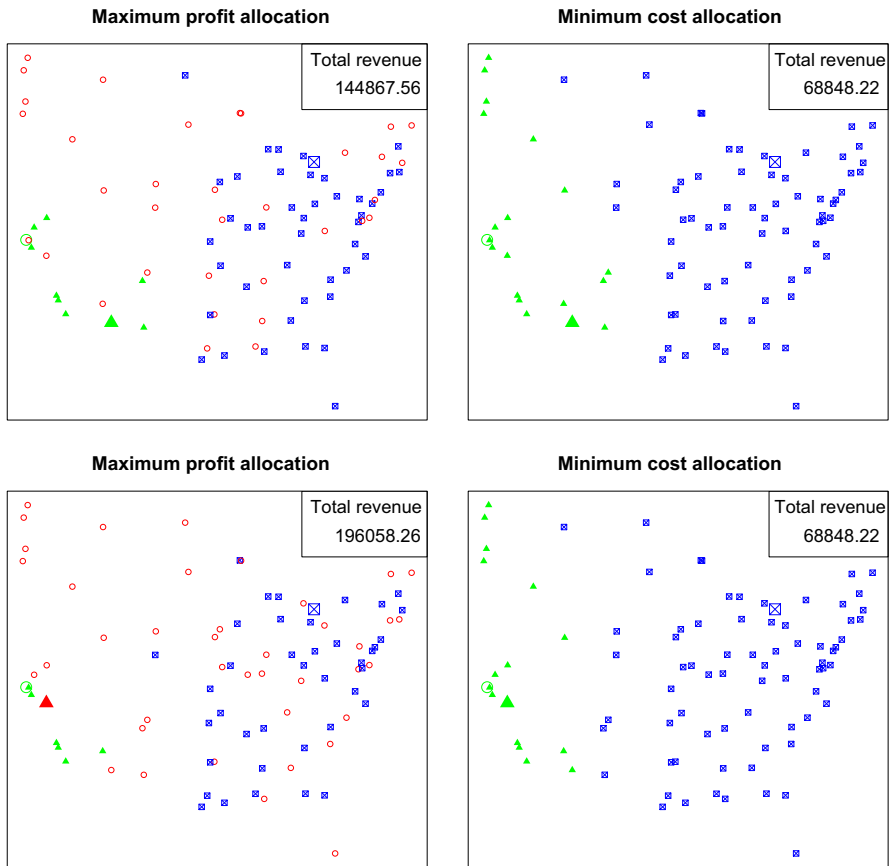


**Fig. 3** Location-allocation solutions in two instances of the 88-node dataset for *AR* production policy with high-profit products and autocorrelation  $\theta = 0.5$

clients are allocated to their furthest open facility, and geographically close clients may be optimally allocated to facilities far from each other. Moreover, we have observed that the facilities opened may have high production costs.

2. In some cases the plants are not served by themselves (see the plant denoted by a circle in the left panel of both figures). This is due to extremely high production costs: it is more profitable to be served by a plant with higher transportation costs but lower overall expenses. This behaviour was already observed in Coullard et al. (2002).
3. Since the location of the plants is decided by optimizing the in-sample empirical profit, some of the  $p$  open plants may be useless under the minimum cost allocation. Therefore, some open plants may not serve any client (see the plant denoted by a circle in the right panels of both figures).

Summarizing, we can conclude that embedding the production policy into the location-allocation problem may significantly improve the out-of-sample profit, leading to different location-allocation decisions that may not follow a minimum cost



**Fig. 4** Location-allocation solutions in two instances of the 88-node dataset for *AR* production policy with low-profit products and autocorrelation  $\theta = 0.5$

allocation strategy. In particular, the worse the production policy performs under highly variable demand, the lower the percentage of clients assigned with minimal cost. In this case, also plants with higher production and transportation costs might be opened. As a consequence, although seemingly senseless, the data-driven allocations may provide a better out-of-sample performance of the supply chain.

### 5 Conclusions, recommendations and extensions

In this paper we have empirically analyzed the impact of numerous factors over the location-allocation decisions, with special emphasis on the effects of the production policy. From the numerical study carried out on Sect. 4, we can conclude that embedding the production policy into the location-allocation problem may significantly improve the out-of-sample profit of the supply chain. These location-allocation

decisions may not follow a minimum cost strategy, specially for those production policies that are not able to cope with highly variable demand.

The dispersion of the demand strongly affects the performance of the joint location-production model. As a consequence, a careful choice of the production policy is specially needed in the presence of high variances. If the chosen strategy is not able to deal with heavy-tailed demands, the supply chain may attain losses. The effect caused by the production policy chosen is not as pronounced if the demand has lower variance. In this case, the performance of the supply chain relies more on both the spatial distribution of the nodes and the temporal autocorrelation of the demand. Hence, also taking into account the autocorrelation of the demand when managing the supply chain can significantly improve the final profit.

Gathering the conclusions from the empirical study carried out in this paper, we venture to give some tips on how to improve the design and performance of a supply chain:

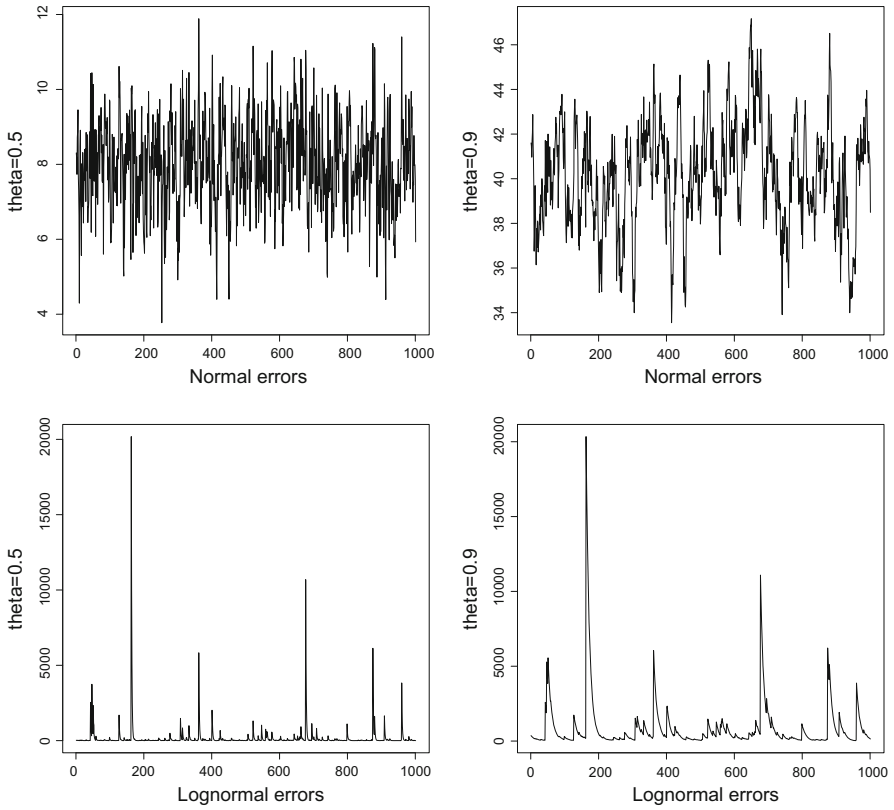
1. We recommend to embed the production policy into the location-allocation problem. To do so, it is necessary to use the historical data to choose the production policy before designing the supply chain. A careful choice must be made, taking into account the characteristics of the demand.
2. The most relevant feature is the variability of the demand. If the demand doesn't show high variations, it might be advisable the use of the classic autoregressive newsvendor procedure, which assumes normality. For heavy tailed demands, we recommend a robust production policy that is able to deal with no distributional assumptions and strong variability. In our numerical experiments, the *ARUS* approach (Carrizosa et al. 2016) performs well.
3. If the demand shows negligible autocorrelation, the classic newsvendor policy might be appropriate to design the supply chain. Nevertheless, the higher the autocorrelation of the demand, the worse the supply chain will perform.
4. We do not advise the use of minimum cost allocation; even though the data-driven clients-facilities allocations might seem unreasonable, they will probably lead to a better out-of-sample performance of the supply chain.

These recommendations arise naturally from the conclusions of our empirical study, where it is important to take into account that our numerical experiments are constrained to our test environment. First, we have considered only one single facility model ( $p$  facilities, with  $p$  fixed) and second, we have made strong assumptions over demand distribution (demands of different clients are considered as independent, and thus the spatial correlation is ignored). Furthermore, the simplest production model has been used in order to not add more degrees of freedom to the problem. Extensions to more realistic models, including multi-item inventory or non-perishable storable products deserve further study. Also, dynamic allocation variables might be considered in order to adjust clients-facilities assignments according to changes of the demand over time.

**Conflict of interest** The authors declare that they have no conflict of interest.

## Appendix

See Fig. 5 and Table 5



**Fig. 5** Examples of high and low correlated autoregressive demands generated with errors following  $N(0, 1)$  and  $LN(0, 3)$

**Table 5** Summary of the numerical setting

Name	Description	Section
<i>Number of facilities to open</i>		
Following Daskin and Owen (1999)		
$p = 3$	Minimum number of facilities tested in Daskin and Owen (1999)	3.1
$p = 9$	Median number of facilities tested in Daskin and Owen (1999)	3.1
$p = 15$	Maximum number of facilities tested in Daskin and Owen (1999)	3.1
<i>Data sets</i>		
Tested in Daskin and Owen (1999)		
Swain	55-node data set of Swain (1971)	3.1
88-node	88-node data set of Daskin (1995)	3.1
<i>Generation of the demand</i>		
Following Carrizosa et al. (2016)		
Model	AR(1) Autoregressive process or order one (Eq. 5)	3.2
Temporal autocorrelation	Moderate Lag-1 autocorrelation of 0.5	3.2
	High Lag-1 autocorrelation of 0.9	3.2
Variability	Low Normal errors ( $N(0, 1)$ )	3.2
	High Lognormal errors ( $LN(0, 3)$ )	3.2
<i>Production policies</i>		
Tested in Carrizosa et al. (2016)		
Benchmark	Future demand is known (i.e., $Q_{ij}(t) = D_j(t)$ )	3.3.1
Static	Classic newsvendor policy (Eq. (7))	3.3.2
AR	Classic newsvendor policy assuming normally distributed errors (Eq. 8)	3.3.3
ARUS	Robust newsvendor problem assuming uncertain AR demand (Pb. (9)–(14))	3.3.4
<i>Products' profitability</i>		
Similar to Carrizosa et al. (2016)		
Low	Sale price is 25% larger than the greatest cost (Eq. 18)	3.4
High	Sale price is 75% larger than the greatest cost (Eq. 19)	3.4
<i>Allocation criteria</i>		
Data-driven maximum profit	Maximum in-sample profit with chosen production policy (Pb. (2) and Table 1)	3.5
Minimum cost	Clients allocated to the minimum cost open facility	3.5

## References

- Atamturk A, Berenguer G, Shen Z (2012) A conic integer programming approach to stochastic joint location-inventory problems. *Oper Res* 60:366–381
- Aviv Y (2002) Gaining benefits from joint forecasting and replenishment processes: the case of auto-correlated demand. *Manuf Serv Oper Manag* 4(1):55–55
- Aviv Y (2003) A time-series framework for supply-chain inventory management. *Oper Res* 51:210–227
- Bandi C, Bertsimas D (2012) Tractable stochastic analysis in high dimensions via robust optimization. *Math Program* 134:23–70
- Blanquero R, Carrizosa E, Nogales-Gómez A, Plastria F (2014) Single-facility huff location problems on networks. *Ann Oper Res* 222(1):175–195
- Box E, Jenkins G, Reinsel G (2008) Time series analysis: forecasting and control. Wiley series in probability and statistics. Wiley, Hoboken, New Jersey
- Cachon GP (1999) Managing supply chain demand variability with scheduled ordering policies. *Manag Sci* 45(6):843–856
- Carrizosa E, Olivares-Nadal AV, Ramírez-Cobo P (2016) Robust newsvendor problem with autoregressive demand. *Comput Oper Res* 68:123–133
- Chen F, Drezner Z, Ryan JK, Simchi-Levi D (2000) Quantifying the bullwhip effect in a simple supply chain: the impact of forecasting, lead times, and information. *Manag Sci* 46(3):436–443
- Coullard C, Daskin M, Shen Z (2002) An inventory-location model: formulation, solution algorithm and computational results. *Ann Oper Res* 100:86–106
- Daskin M (1995) Network and discrete location: models, algorithms and applications. Wiley, New York
- Daskin MS, Owen SH (1999) Two new location covering problems: the partial p-center problem and the partial set covering problem. *Geogr Anal* 31(3):217–235
- Dogan K, Goetschalckx M (1999) A primal decomposition method for the integrated design of multi-period production-distribution systems. *IIE Trans* 31:1027–1036
- Dong L, Lee H (2003) Optimal policies and approximations for a serial multiechelon inventory system with time-correlated demand. *Oper Res* 51:969–980
- Ganesh M, Raghunathan S, Rajendran C (2014) The value of information sharing in a multi-product, multi-level supply chain: impact of product substitution, demand correlation, and partial information sharing. *Decis Support Syst* 58:79–94
- Germain R, Claycomb C, Dröge C (2008) Supply chain variability, organizational structure, and performance: the moderating effect of demand unpredictability. *J Oper Manag* 26(5):557–570
- Graves SC (1999) A single-item inventory model for a nonstationary demand process. *Manuf Serv Oper Manag* 1(1):50–61
- Güllü R (1997) A two-echelon allocation model and the value of information under correlated forecasts and demands. *Eur J Oper Res* 99(2):386–400
- Gulpinar N, Pachamanova D, Canakoglu E (2013) Robust strategies for facility location under uncertainty. *Eur J Oper Res* 225:21–35
- Helper CM, Davis LB, Wei W (2010) Impact of demand correlation and information sharing in a capacity constrained supply chain with multiple-retailers. *Comput Ind Eng* 59(4):552–560
- Huh WT, Levi R, Rusmevichientong P, Orlin JB (2011) Adaptive data-driven inventory control with censored demand based on kaplan-meier estimator. *Oper Res* 59(4):929–941
- Johnson G, Thompson H (1975) Optimality of myopic inventory policies for certain dependent demand processes. *Manag Sci* 21:1303–1307
- Kahn JA (1986) Inventories and the volatility of production. Ph.D. thesis, Massachusetts Institute of Technology, Department of Economics
- Khouja M (1999) The single-period (news-vendor) problem: literature review and suggestions for future research. *Omega* 27(5):537–553
- Lambert DM, Cooper MC (2000) Issues in supply chain management. *Ind Mark Manag* 29(1):65–83
- Levi R, Roundy R, Shmoys D, Pal M (2008) Approximation algorithms for capacitated stochastic inventory control models. *Oper Res* 56(5):1184–1199
- Liao S, Hsieh C, Lai P (2011) An evolutionary approach for multi-objective optimization of the integrated location-inventory distribution network problem in vendor-managed inventory. *Expert Syst Appl* 38:6768–6776
- Lu X, Song J, Regan A (2006) Inventory planning with forecast updates: approximate solutions and cost error bounds. *Oper Res* 54:1079–1097

- Melo M, Nickel S, Saldanha da Gama F (2009) Facility location and supply chain management: a review. *Eur J Oper Res* 196:401–412
- Miranda PA, Garrido RA (2004) Incorporating inventory control decisions into a strategic distribution network design model with stochastic demand. *Transp Res Part E Logist Transp Rev* 40(3):183–207
- Ozsen L, Daskin M, Coullard C (2009) Facility location modeling and inventory management with multi-sourcing. *Transp Sci* 43:455–472
- Raghuathan S (2003) Impact of demand correlation on the value of and incentives for information sharing in a supply chain. *Eur J Oper Res* 146(3):634–649
- Reyman G (1989) State reduction in a dependent demand inventory model given by a time series. *Eur J Oper Res* 41:174–180
- See C, Sim M (2010) Robust approximation to multiperiod inventory management. *Oper Res* 58:583–594
- Shahabi M, Unnikrishnan A, Jafari-Shirazi E, Boyles SD (2014) A three level location-inventory problem with correlated demand. *Transp Res Part B Methodol* 69:1–18
- Shang K, Song J (2003) Newsvendor bounds and heuristic for optimal policies in serial supply chains. *Manag Sci* 49:618–638
- Shen Z, Qi L (2007) Incorporating inventory and routing costs in strategic location models. *Eur J Oper Res* 179:372–389
- Shen Z, Coullard C, Daskin M (2003) A joint location-inventory model. *Transp Sci* 37:40–55
- Shu J, Teo C, Shen Z (2005) Stochastic transportation-inventory network design problem. *Oper Res* 53:48–60
- Snyder L, Daskin M, Teo C (2007) The stochastic location model with risk pooling. *Eur J Oper Res* 179:1221–1238
- Snyder LV (2006) Facility location under uncertainty: a review. *IIE Trans* 38(7):547–564
- So KC, Zheng X (2003) Impact of supplier's lead time and forecast demand updating on retailer's order quantity variability in a two-level supply chain. *Int J Prod Econ* 86(2):169–179
- Swain R (1971) A decomposition algorithm for a class of facility location algorithms. Ph.D. thesis, Cornell University, Ithaca, N.Y
- van den Berg PL, van Essen JT, Harderwijk EJ (2016) Comparison of static ambulance location models. In: 2016 3rd International conference on logistics operations management (GOL), pp 1–10. IEEE
- Wang T, Atasu A, Kurtulus M (2012) A multiordering newsvendor model with dynamic forecast evolution. *Manuf Serv Oper Manag* 14:472–484
- Wang Z, Yao DQ, Huang P (2007) A new location-inventory policy with reverse logistics applied to B2C e-markets of China. *Int J Prod Econ* 107(2):350–363
- Weaver JR, Church RL (1986) A location model based on multiple metrics and multiple facility assignment. *Transp Res Part B Methodol* 20(4):283–296
- Zhang ZH, Unnikrishnan A (2016) A coordinated location-inventory problem in closed-loop supply chain. *Transp Res Part B Methodol* 89:127–148

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.