Computers & Industrial Engineering 174 (2022) 108793

Contents lists available at ScienceDirect

Computers & Industrial Engineering

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

Simulation-optimization of inventory control of multiple products on a single machine with sequence-dependent setup times

Marco Aurélio Mesquita [a,b,*], João Vitor Tomotani [a,b]

[a] Department of Production Engineering, University of São Paulo, Brazil
[b] Av. Luciano Gualberto, 1380, Cidade Universitária, São Paulo, SP 05508-010, Brazil

ARTICLE INFO

Keywords:
Stochastic Economic Lot Scheduling Problem
Inventory Control
Production Scheduling
Simulation-Optimization

ABSTRACT

This work studies the problem of lot sizing and scheduling of multiple products on a single machine, with stochastic demand and sequence-dependent setup times, called Stochastic Economic Lot Scheduling Problem (SELSP). The present work differs from others in the literature by considering simple inventory control policies and using the simulation–optimization approach to calibrate their parameters. We consider two inventory control policies: (i) fixed cycling (First in Sequence - FIS) and (ii) dynamic scheduling based on inventory levels (Lowest Days of Supply - LDS), combined with an "order-up-to" lot sizing. The problem is solved using AnyLogic simulation software and the OptQuest search engine to minimize total inventory cost (ordering, holding and shortage costs). The experimental design included the following factors: number of items, coefficient of variation of demand, system workload, and degree of setup increment, allowing the comparison of the two inventory control policies in different scenarios. Experiments show that LDS outperforms FIS in all scenarios, achieving up to 4.6% cost savings for cases of more products, higher workload, and greater demand variance. The developed models proved to solve the problem, effectively generating reasonable solutions. Furthermore, as they are user-friendly, we believe they can be adapted, without great difficulties, to real-life scenarios of the process industry.

1. Introduction

Lot Sizing and Scheduling Problems (LSP) have received constant attention from researchers and practitioners in Production Planning and Control. Tomotani and Mesquita (2018) show that lot sizing and scheduling is a relevant problem, mainly in the process industry, directly impacting the service level and inventory costs.

A basic formulation of the LSP is to schedule multiple items on a single machine with limited capacity and significant setup times. This problem of multiple products and a single machine, continuous in time, where setup times, production times, and demands are known, is called the Economic Lot Scheduling Problem (ELSP). The objective of the ELSP is to obtain a cyclical pattern of production that will be repeated within a planning horizon, capable of meeting demand and capacity constraints, minimizing total holding and setup (Gascon et al., 1994).

The ELSP has a stochastic variant called Stochastic Economic Lot Scheduling Problem (SELSP), which considers the demand as a stationary stochastic process, and that production and setup times can also be random. Since demand is stochastic, to solve the SELSP, it is not sufficient to define a fixed production sequence with fixed lot sizes. It is necessary to define a policy capable of deciding when and how much to produce of each product dynamically, in response to actual demand, and considering the need for safety stocks (Glock et al., 2014).

Inventory control for SELSP can be classified according to scheduling and lot-sizing procedures. The scheduling may be fixed - the items are checked and produced in a predefined order (e.g., Federgruen and Katalan, 1998; Wagner and Smits, 2004; Jodlbauer and Reitner, 2012) or dynamic - the next item to be produced is chosen based on stock levels (e.g., Graves, 1980; Paternina-Arboleda and Das, 2005). The lot sizes can be fixed - production lots for each item have the same size (e.g., Graves, 1980; Vaughan, 2007; Jodlbauer and Reitner, 2012) or variable - lot sizes vary depending on stock levels (e.g., Smits et al., 2004; Wagner and Smits, 2004; Löhndorf and Minner, 2013). The present work compares fixed cycling and dynamic scheduling policies with variable lot sizes (order-up-to policy) in a single-machine and multiple-product environment with sequence-dependent setup times.

The contributions of the present work are twofold. First, from a theoretical standpoint, it aims to solve the SELSP using the

* Corresponding author at: Departamento de Engenharia de Produção, Escola Politécnica – USP, Av. Prof. Luciano Gualberto, 1380, São Paulo, SP 05508-010, Brazil.
E-mail address: marco.mesquita@poli.usp.br (M.A. Mesquita).

https://doi.org/10.1016/j.cie.2022.108793
Received 13 April 2022; Received in revised form 18 September 2022; Accepted 30 October 2022
Available online 3 November 2022
0360-8352/

simulation–optimization method, comparing two inventory control policies (fixed cycling and dynamic scheduling) in different scenarios (number of products, demand variation, system workload, and setup increment). Second, from a practical perspective, it proposes two simulation–optimization models, one for each inventory control policy, that allow calibration and application of them in real industrial environments. We highlight that the "single machine" mentioned above can be a complete automated line that produces only one product at a time (batch production) and need setup between lots, as is typical in the process industry (food, beverages, medicines etc.).

The article is structured as follows. The first section presents the research problem and objectives, and the second reviews the relevant literature. The third section presents the problem definition, while section four discusses the research method and simulation models. Next, section five presents the experimental design, and section six discusses the results. Finally, the last section concludes the article and proposes future research directions.

2. Literature review

The Economic Order Quantity (EOQ) model was one of the first applications of scientific method to industrial engineering and is considered the starting point of lot sizing and scheduling theory. The model calculates the optimal production quantity that minimizes total inventory cost, balancing holding and setup costs (Erlenkotter, 2014).

The limitation of the EOQ in considering only one product was addressed by Rogers (1958), who formulated the Economic Lot Scheduling Problem (ELSP). The original formulation of the ELSP assumes that: i) the machine produces only one item at a time, ii) production rates, set up times, and set up costs are known and independent of the production sequence, iii) the demand of each item is known and constant over an infinite planning horizon, iv) production capacity is finite, and all demand must be met without backlog. Due to these assumptions, the input data must be checked in advance to guarantee feasibility.

Holmbom and Segerstedt (2014) provide a historical overview from the EOQ to the ELSP. Their article details three main approaches to scheduling in ELSP: common cycle, basic period, and extended basic period. The authors also list practical challenges when implementing mathematical models, mainly because of the demand and processes uncertainties, which require the constant review of production plans. Beck and Glock (2019) also review the literature on ELSP and highlight that previous research on this topic has had a strong focus on mathematical modelling. According to them, several practical aspects directly linked to lot sizing and scheduling have not yet attracted much attention from researchers.

The ELSP has a stochastic variant called Stochastic Economic Lot Scheduling Problem (SELSP), which considers demand as a stationary stochastic process, while production and setup times can also be random. Vergin and Lee (1978) pointed out the practical importance and complexity of this problem and were the first to question the scarcity of academic literature on it. For a more detailed review on the SELSP, we recommend Sox et al. (1999) and Winands et al. (2011).

Goyal (1973) presented one of the first lot-sizing studies, where multiple items with stationary stochastic demand are produced on a single machine. The author proposed a model whose objective was to minimize the sum of holding, setup, and shortage costs. However, a significant limitation in this study, as pointed out by Vergin and Lee (1978), is that it considers the production lead time equal to zero, which is equivalent to unlimited capacity and, therefore, cannot be classified as SELSP.

Graves (1980) proposes the Multi-Product Production Cycling Problem (MPCP), equivalent to SELSP. The author presents a heuristic based on the case of a single product problem, which is simpler and easier to solve, and the notion of "composite product", a way of aggregating products in a family. Simulation experiments compare his composite product heuristic against four other heuristics derived from the reorder point (r, Q) and base stock level (s, S) policies. The composite product heuristic dominated three of these heuristics and had an equivalent performance against the fourth while being easier to parameterize.

Leachman and Gascon (1988) studied the SELSP (and were the first to name it stochastic ELSP, although they did not use the acronym SELSP) and proposed a model like Graves (1980), working with discrete time (time buckets) but allowing overtime per period. As in the work of Graves (1980), only one item can be produced in each time bucket. The model's objective is to define a production schedule that minimizes average inventory and setup costs, avoiding backorders. The authors used the concepts of runout times (estimated time until an item reaches its safety stock level) and slack times (time intervals in which production can remain idle without causing stockouts). When predicting a future "negative slack time", the heuristic seeks to reduce previous batches to eliminate this negative slack time. The heuristic starts with an initial solution, where lots are calculated considering a deterministic problem.

The work of Leachman and Gascon (1988) was revisited by Leachman et al. (1991), who propose improvements in the runout time heuristic, introducing new ways of calculating cycle times and slack times. In simulation experiments, the new heuristic showed a 3.0 % reduction in total costs compared to the original while maintaining similar inventory levels. This runout time heuristic was the basis for several other works, such as Brander et al. (2005) and Levén and Segerstedt (2007).

Vaughan (2007) studied two scheduling policies for SELSP: fixed cycling and dynamic scheduling. The production follows a predefined order in the first case, while the second defines the order according to all items' stock levels. The author sought to understand the conditions under which a fixed cycling would outperform the dynamic scheduling. The simulation results show that the fixed cycling is more suitable in scenarios with fewer items, longer setup times, and significant capacity limitations.

Löhndorf and Minner (2013) carried out a study in which they approach SELSP through simulation and optimization techniques, concluding, similarly to Vaughan (2007), that base-stock policies (dynamic scheduling) are suitable for the problem, but are overcome by fixed cycling in high workload scenarios. This work was extended in Löhndorf et al. (2014) when the authors included sequence-dependent setup times.

Gel et al. (2021) use queuing theory to estimate utilization and cycle times on a machine that produces multiple products with sequence-dependent setup times. The authors estimate them for four standard scheduling rules in industrial practice. Computer simulations conclude that the approximations provided by the proposed method are helpful for rough-cut capacity planning in real-world scenarios.

Tubilla and Gershwin (2021) study a version of SELSP in which the source of randomness are machine breaks. The authors test different heuristics for two problem sets, one where failures are frequent but repaired quickly, and another where breakdowns are rare but significantly impact the system. The authors propose a new policy that tightly controls the surplus of the highest-priority items using fixed base-stock levels and, for all other items, it determines their run lengths dynamically. Simulation experiments show that the new policy significantly outperforms the benchmarking approaches over a large set of operating conditions.

In previous works that applied simulation to study the SELSP, we identified two ways of modelling the demand process: a) as a compound Poisson process (e.g., Wagner and Smits, 2004; Vaughan, 2007; Löhndorf and Minner, 2013); b) as daily demands with independent random sampling (e.g., Kamath and Bhattacharya, 2007; Jodlbauer and Reitner, 2012; Rappold and Yoho, 2014). In the first case, the demand randomness results from the variance of order arrivals and order quantities. In the second, we can directly set the daily demand variance with different levels. As one of our goals is to evaluate the effect of variance on inventory control performance, we opted for the second approach, modelling daily demands with the lognormal distribution.

Another issue in inventory modelling is how to deal with stockouts. There are two basic approaches: i) loss of sales; ii) backorders (Hopp and Spearman, 2008). The first assumes that demand is partially met or lost when there is insufficient inventory and incurs a cost of lost sales. The second assumes that the demand will be met with an urgent production order, with a backorder cost proportional to the delay in customer service. It is still possible to find works dealing with mixed models in which customers are divided into classes, each with the option of backorder or loss of sales (Teunter and Haneveld, 2008). In the SELSP literature, we also identified these two approaches: a) lost sales (e.g., Karalli and Flowers, 2006; Liberopoulos et al., 2013; Löhndorf and Minner, 2013); b) backorder (e.g., Wagner and Smits, 2004; Jodlbauer and Reitner, 2012; Cunha Neto et al., 2015). In our work, we opted for the loss of sales premise, with partial or zero demand fulfilment, depending on the level of available inventory.

In this brief literature review, we have identified two main streams of research. The first one seeks to develop heuristics for solving the SELSP and usually evaluates them based on simulation experiments (e.g., Gascon et al., 1994; Vaughan, 2007; Cortés-Fibla et al., 2015). The second uses the simulation and optimization approach to calibrate fixed cycling and dynamic scheduling policies (Löhndorf and Minner, 2013; Löhndorf et al., 2014). The present work fits into the second one. It considers simple inventory control policies, using simulation and optimization models to calibrate their parameters and then running a complete factorial experiment, which allows analysing the effect of some operational factors on the performance of the models.

3. Problem formulation

This article compares two inventory control policies in a single machine and multiple items make-to-stock (MTS) environment. The machine produces batches of a single item at a time to replenish finished goods inventories consumed at the end of each day by daily demands. Each product has a deterministic production rate which, once the size of the production order is defined, allows the calculation of the production time. In addition to the production time, we consider sequence-dependent setup times. Production occurs continuously, seven days a week, 24 h a day.

The daily demand for each item is random and defined by a lognormal distribution. The inventory is consumed to meet each day's demand, considering the assumption of partial fulfilment in case the available stock is insufficient. There is no backorder, and the unmet demand is lost.

The inventory replenishment considers an "order-up-to" policy, with two parameters: minimum stock (s) and maximum stock (S). An item only goes into production if its inventory level drops below s, and the size of the production order will be the difference between S and its inventory position when scheduled. We consider two scheduling rules: i) First in Sequence (FIS); ii) Lowest Days of Supply (LDS). In the FIS rule, an ideal production sequence that minimizes setup times is defined a priori, and the next item in this sequence with an inventory below s is chosen. In LDS, out of all items below the minimum stock s, the one with the lowest inventory coverage is selected. Both rules allow the system to become idle if no item is below the minimum stock s after meeting daily demand.

Simulation-optimization models determine, for each of the two inventory control rules, which values of s and S minimize the Total Inventory Cost (TIC), which includes holding cost, setup cost and loss of sales/shortage cost. We considered s and S integers. Still, the models can be adapted to consider them and the lot sizes as multiples of some values (parameters) to cope with practical technological constraints of actual industrial processes.

4. Simulation models

Two discrete event simulation models were developed in AnyLogic software: i) First in Sequence (FIS) and ii) Lowest Days of Supply (LDS). The first follows the logic of fixed cycling, while the second follows a dynamic scheduling policy. Both simulation models have the same basic structure, differing only in the rule for defining the next item for production. Our SELSP conceptual model is based on Altiok and Melamed (2007) and considers three macro-processes: demand, production, and control. The production and demand processes interact with each other, determining the system dynamics. The control process collects data from the system to calculate the inventory performance metrics, used to calibrate parameters and compare stock policies.

In the demand process (Fig. 1), random daily demands are filled from the finished goods inventory. If there is insufficient stock, the demand is partially filled or lost if the stock is zero (stockout premise). The model does not consider backorders. If the machine is idle at the end of the day, the daily demand can trigger production if any item's stock drops below the minimum stock s.

The production process (Fig. 2) is a continuous cycle that switches between idle and busy states. The process consists of identifying the next item to produce or stopping if there is no item below the minimum stock.

There are two procedures for defining the next item, one for each model. The FIS model chooses the first item in the sequence below the minimum stock, while the LDS model selects the item with the lowest value of the "days of supply", a metric calculated by dividing the stock position by the average daily demand of the item. Algorithms 1 and 2 illustrate these two procedures.

Algorithm 1 First in Sequence – FIS

```
input k, inv, smin    // current item, stock levels, and minimum stock levels
output idle, j    // Boolean variable for machine status and next item to produce, if any
idle← True
j←k   // next item is the current item
for i←1 to n do:
   j←(j + 1) mod n   // tries next item in the sequence
   if (inv[j] ≤ smin[j]) then:
      idle←False
      break   // stop when you find the first item below minimum stock
return idle, j
```

Algorithm 2 Lowest Days of Supply – LDS

```
input inv, smin, d    // stock levels, minimum stock levels, and average demands
output idle, j    // Boolean variable for machine status and next item to produce, if any
idle← True
dosMin←infinity   // minimum days of supply
j←Null   // initializes j to null
for i←1 to n do:
   dos←inv[i]/d[i]   // calculates the days of supply of item i
   if (inv[i] ≤ smin[i]) and (dos ≤ dosMin) then:
      idle← False
      dosMin←dos   // sets the new minimum days of supply
      j←i   // sets i as the next item
return idle, j
```

Finally, the control process (Fig. 3) only collects daily data on the stock position of each item to calculate the final average inventory levels. The sales data, used to calculate the service level indicator and cost of lost sales, are collected directly in the demand process.

Fig. 4 presents a screenshot of the general simulation model for $n = 5$ products in AnyLogic software. The two versions of the model (FIS and LDS) are used in simulation–optimization experiments with the OptQuest engine to compare the two inventory control policies and analyse the effect of some operational factors on their performance.

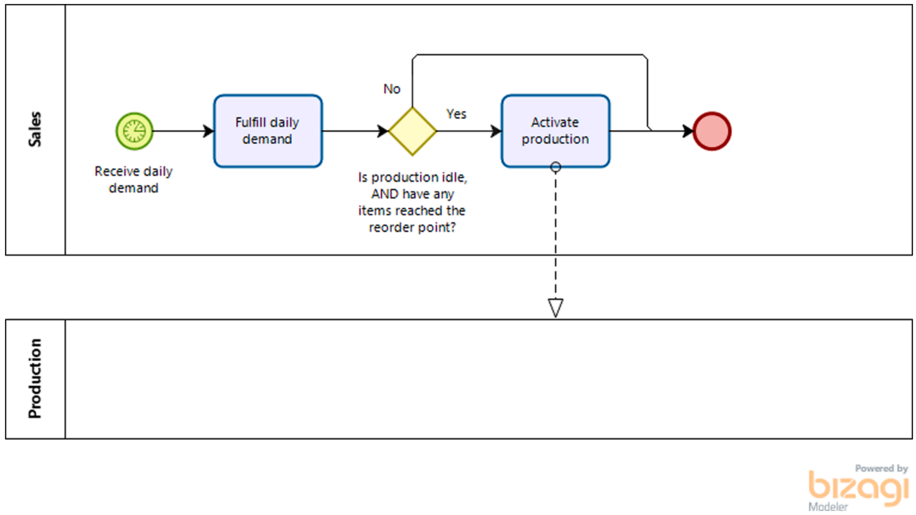

Fig. 1. Caption: Demand Process. Alt Text: The figure describes the process of meeting daily demands.

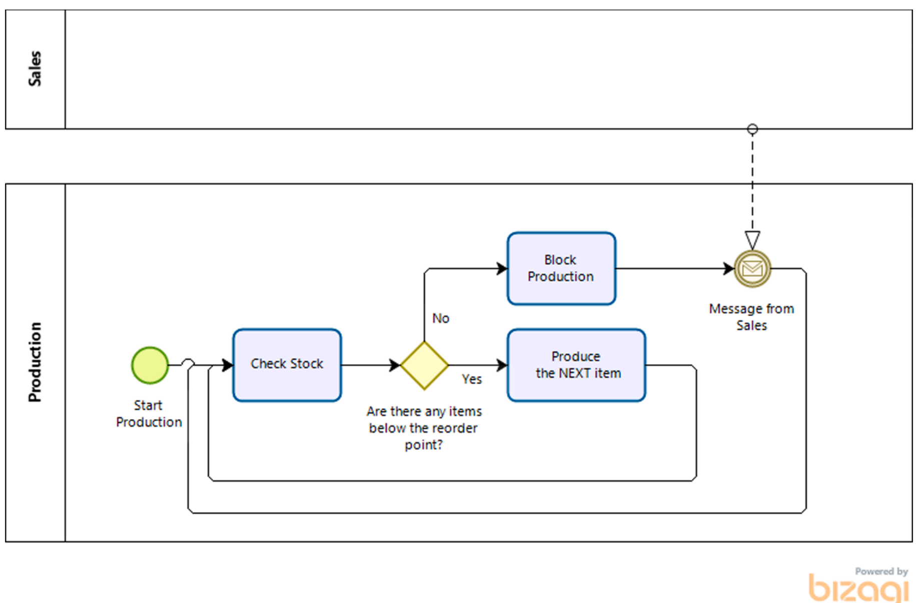

Fig. 2. Caption: Production Process. Alt Text: The figure describes the production scheduling decision process.

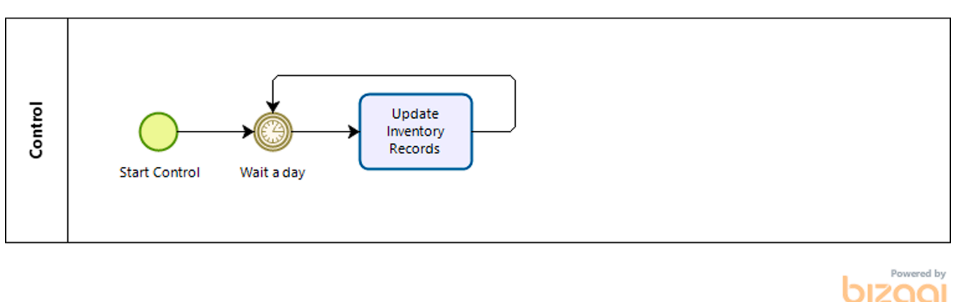

Fig. 3. Caption: Control Process. Alt Text: The figure describes the control process.

5. Experimental design

This section presents the design of the experiments carried out to compare the two inventory control policies in SELSP. We used a full factorial design, which considers the following factors:

– inventory model (mod): FIS, LDS
– number of items (n): 5, 10
– coefficient of variation of demand (cv): 20 %, 50 %, 100 %
– utilization rate (ρ): 70 %, 80 %, 90 %
– setup increment (α): 50 %, 100 %, 200 %

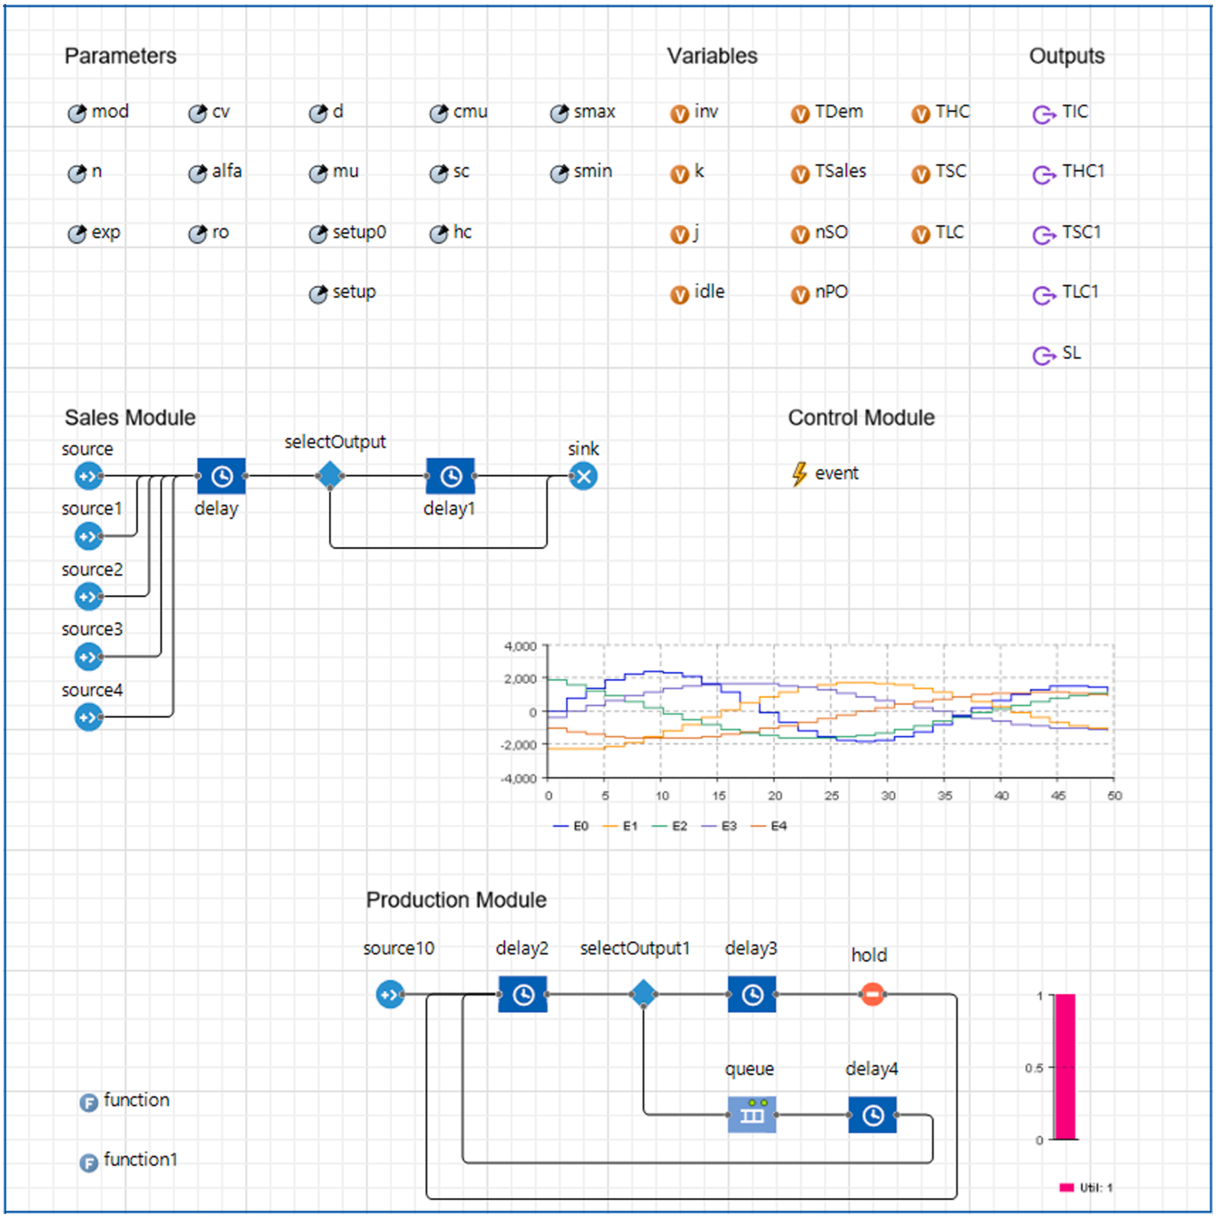

Fig. 4. Caption: Screenshot of the general simulation model for n = 5 products in AnyLogic. Alt Text: The figure shows a screenshot of the simulation model for n = 5 products in the AnyLogic software.

In the design of the experiments, we assumed that all products have the same production and demand rates, holding and setup costs, and loss of sales costs. Due to this homogeneity, all products will have the same parameters of inventory control (values of minimum and maximum stock). These assumptions enable the factorial analysis of variance, presented in the following section. However, the simulation–optimization models developed accept different demand and cost parameters for specific case evaluations.

To directly compare the results from different instances, we sought to scale the demand so that the total contribution margin was potentially the same for all instances. We set the total contribution margin target at Y=$1.0 million per year, a contribution margin per unit of $40, and 250 working days per year. Then, we calculated the average daily demand (d) for each product, which resulted in 10 and 20 units/day for cases with 10 and 5 products, respectively. The daily demand of each product is a lognormal random variable, with variance given by the coefficient of variation factor.

Once the daily demand for each product is defined, we establish the production rate, which corresponds to the total daily demand divided by the utilization rate of each experiment ($\mu = n \bullet d/\rho$). We consider deterministic production and setup times; only demand is stochastic.

To analyse the effect of sequence-dependent setup times, we established a basic setup (su0), inversely proportional to the number of items, and a linear setup growth rate (α). With su0 and α, we determine the instances setup matrix. To define the basic setup, we considered a target of 300 h/year (5 % of total working hours per year) and 24 setups per product per year (each item being produced twice a month). Thus, we got the values of 2.5 and 1.25 h of setup times for the cases of 5 and 10 products, respectively. The α factor corresponds to the largest increment in setup, which occurs when we jump from one item j to n-1 items further in the production cycle. Table 1 exemplifies the setup matrix for the hypothetical case of n = 5, su0 = 1 h, and α = 100 %.

The simulation–optimization models seek to determine the inventory control parameters that minimize the Total Inventory Cost (TIC), given by the sum of Total Holding Cost (THC), Total Setup Cost (TSC) and Total Loss of Sales Cost (TLC). We defined the setup and holding costs

Table 1
Setup matrix for n = 5, su0 = 1 h, and α = 100 %.

		To				
		1	2	3	4	5
From	1	0.0	1.000	1.333	1.667	2.000
	2	2.000	0.0	1.000	1.333	1.667
	3	1.667	2.000	0.0	1.000	1.333
	4	1.333	1.667	2.000	0.0	1.000
	5	1.000	1.333	1.667	2.000	0.0

using the economic order quantity concept as an approximation. We set the holding cost to $25 per unit per year (or 25 % of the item cost) and the setup cost to $250 (only the setup times are sequence-dependent, not the setup costs). For the shortage cost, we considered the cost of a lost sale equal to the item's contribution margin. In addition to TIC, we also evaluate the Service Level (SL) indicator, which corresponds to the percentage of demand met throughout the simulation.

The simulations start with all stocks full and last for 250 working days. The combination of factors provided 108 instances, each being simulated five times, totalling 540 experiments. Using the OptQuest engine, the optimization process was carried out in three runs. First, we consider a range for s and S values from 10 to 2000, with step 10 and the additional constraint of $s \leq S - 1$. Thus, we obtained five pairs of (s, S) for each instance and the corresponding values of the total inventory cost (TIC). Then, we repeated the 540 experiments twice by reducing the step to 5 and 1 and the search space as follows. Let S_1 and S_2 be the lower and the upper S reached in the previous run. We set the new S range from S_1 minus k to S_2 plus k, with k being the step in the previous run, and the new s range from 10 to S_2, including the constraint $s \leq S - 1$. After the third run, for those instances with a higher variation of S, we did extra few runs to meet convergence. The results of the last runs were considered the final outputs of the optimization process. Table 2 summarizes the parameters and variables of the experimental design.

6. Results and discussion

This section describes the results of 540 experiments (108 instances with five replicates each). The design of the experiment considers 5 variable factors: i) control model (mod), ii) number of items (n), iii) coefficient of variation of demand (cv), iv) utilization rate (ρ) and v) setup increment (α). From n and ρ, we determine the three semi-variable factors: i) average daily demand (d), ii) production rate (μ), iii) basic setup (su0). Finally, we have three fixed factors: i) holding cost (hc), ii) setup cost (sc), iii) loss of sales cost (lc). For each experiment, we calculated: total inventory cost (TIC), minimum stock (s), and maximum stock (S).

Table 2
Summary of parameters and variables.

Class	Parameter	Values
Variable factor	Inventory model	FIS, LDS
Variable factor	Number of items (n)	5, 10
Variable factor	Demand coefficient of variation (cv)	20 %, 50 %, 100 %
Variable factor	Utilization rate (ρ)	70 %, 80 %, 90 %
Variable factor	Setup increment (α)	50 %, 100 %, 200 %
Semi variable factor	Daily demand (d, unit/day)	$1.000.000/n/cmu/L
Semi variable factor	Production rate (μ, unit/day)	n*d/ρ
Semi variable factor	Basic setup (su0, h)	300 h/(24*n)
Fixed factor	Holding cost (hc, $/year)	$25
Fixed factor	Setup cost (sc, $)	$250
Fixed factor	Loss of sales cost (lc, $/unit)	$40
Fixed factor	Number of simulated days (L)	250
Response Variable	Total Inventory Cost (TIC, $)	
Response Variable	Minimum stock (s)	
Response Variable	Maximum stock (S)	
Response Variable	Service level (sl)	
Experiment Parameter	Replications per instance (r)	5
Experiment Parameter	Total number of experiments (N)	540

We ran the experiments on a personal computer with a microprocessor Intel Core i7 3.6-GHz PC and 16 GB of RAM. Table 3 presents the means and standard deviations of the computational times. These times depend on the size of the search space and the step considered in each run (10, 5 and 1, for the first, second and third runs). Table 4 presents the first and last lines of the experiment results table.

Table 5 presents the Average Total Inventory Cost (Av. TIC) for each level of the five factors. The "Diff." columns show the percentage difference from the best result. In bold are the highest values of Av. TIC and Diff. Similarly, we have the Average Maximum Stocks (Av. S).

We emphasize that the differences in the previous table depend on the arbitrated values for the holding, setup, and loss of sales costs. We consider that the relationship between inventory cost parameters in this paper is realistic for stable market economies.

Analysing the first factor (n), we found that the scenarios with n = 10 present inferior performance. We stipulated that the demand would be inversely proportional to the number of products, so that the instances would have the same potential total margin of contribution. We defined the production rates according to the total daily demand and the utilization rate of the instance. Despite this, the number of items still was the factor with the most statistically significant effect, as seen below.

The effects of demand uncertainty and capacity constraint are as expected. The higher the variance of demand, the higher the TIC. Likewise, the higher the workload, the higher the TIC. Interestingly, the increase in demand variance and workload are offset by the rise in maximum stock level (S), which reflects on average stock levels; the optimization model seeks to compensate for uncertainty in demand and capacity constraints with the increasing stock.

The next factor is the setup increment (α), which defines the sequence-dependent setup times. This factor had a more negligible impact on TIC, despite the wide range of values tested (up to a 200 % increase in setup). One possible interpretation is that, despite the high penalty, the loss of productivity due to higher setup times is offset by the reduction in loss of sales by producing items closer to stockout.

Finally, in the last three columns of Table 5, we compare the two control models - FIS and LDS. The LDS model surpassed the FIS not only in the general average but also in all combinations of the other four control factors. Table 6 presents the average total inventory cost for the different combinations of n, cv, and ρ, disregarding factor α which had the lowest impact on the cost. The most significant differences are seen in the instances with higher demand uncertainty and workloads.

Figs. 5 and 6 present the results of the experiments with different cv, ρ, and mod combinations in a boxplot format for n = 5 and 10, respectively. On the left half, we have the plots for the FIS model and, on the right, the ones for the LDS model. The first three boxplots on the left correspond to ρ = 0.7, the next three to ρ = 0.8, and the next three to ρ = 0.9. This sequence repeats on the right half. Finally, the first three boxplots on the left correspond to the values of cv = 0.2, 0.5, and 1.0. This sequence is repeated for the other plots on the right. These boxplots show the apparent effect of the cv and ρ factors on TIC. In both cases (n = 5 and 10), one can see the superiority of the results with the LSD model on the right compared to the FIS model on the left.

Table 3
Mean and standard deviation of computational times per run and model.

	1st Run		2nd Run		3rd Run	
	Mean (s)	S.D. (s)	Mean (s)	S.D. (s)	Mean (s)	S.D. (s)
FIS05	12.9	1.3	5.6	2.2	10.9	0.9
FIS10	25.6	2.0	5.4	1.4	19.3	3.1
LDS05	12.5	1.2	5.3	2.6	10.7	1.9
LDS10	25.0	1.3	5.0	1.4	21.2	1.8

Table 4
Head and tail of the experiment results table, with the respective input parameters.

#	mod	n	cv	α	ρ	d	μ	su0	hc	sc	lc	TIC	s	S	time
001	FIS	5	0.2	0.5	0.7	20.0	142.86	2.50	0.1	250.0	40.0	46093.9	84	363	10.1
002	FIS	5	0.2	0.5	0.7	20.0	142.86	2.50	0.1	250.0	40.0	46670.8	83	355	9.9
003	FIS	5	0.2	0.5	0.7	20.0	142.86	2.50	0.1	250.0	40.0	46475.9	84	360	11.3
004	FIS	5	0.2	0.5	0.7	20.0	142.86	2.50	0.1	250.0	40.0	46459.6	89	369	11.6
005	FIS	5	0.2	0.5	0.7	20.0	142.86	2.50	0.1	250.0	40.0	46400.2	90	364	11.1
006	FIS	5	0.2	0.5	0.8	20.0	125.00	2.50	0.1	250.0	40.0	48072.8	105	416	10.3
007	FIS	5	0.2	0.5	0.8	20.0	125.00	2.50	0.1	250.0	40.0	47848.3	95	410	11.3
008	FIS	5	0.2	0.5	0.8	20.0	125.00	2.50	0.1	250.0	40.0	47943.4	101	423	12.3
009	FIS	5	0.2	0.5	0.8	20.0	125.00	2.50	0.1	250.0	40.0	47540.8	102	417	10.2
010	FIS	5	0.2	0.5	0.8	20.0	125.00	2.50	0.1	250.0	40.0	47989.2	109	415	10.5
…	…	…	…	…	…	…	…	…	…	…	…	…	…	…	
536	LDS	10	1.0	2.0	0.9	10.0	111.11	1.25	0.1	250.0	40.0	74786.7	101	287	19.0
537	LDS	10	1.0	2.0	0.9	10.0	111.11	1.25	0.1	250.0	40.0	74974.2	101	278	21.1
538	LDS	10	1.0	2.0	0.9	10.0	111.11	1.25	0.1	250.0	40.0	75500.0	103	281	21.6
539	LDS	10	1.0	2.0	0.9	10.0	111.11	1.25	0.1	250.0	40.0	74860.5	102	292	23.8
540	LDS	10	1.0	2.0	0.9	10.0	111.11	1.25	0.1	250.0	40.0	75085.6	106	276	22.6

Table 5
Average Total Inventory Cost (TIC) and Average Maximum Stock (S) by factor.

n	Av. TIC	Diff.	cv	Av. TIC	Diff.	ρ	Av. TIC	Diff.	α	Av. TIC	Diff.	mod	Av. TIC	Diff.
5	**53527.8**	0 %	0.2	**55978.0**	0 %	0.7	**59736.3**	0 %	0.5	**61163.2**	0 %	FIS	61966.9	2.3 %
10	69004.8	28.9 %	0.5	60016.3	7.2 %	0.8	61272.0	2.6 %	1.0	61257.0	0.15 %	LDS	**60565.7**	0 %
			1.0	67804.6	21.1 %	0.9	62790.6	5.1 %	2.0	61378.7	0.35 %			
n	**Av. S**	**Diff.**	**cv**	**Av. S**	**Diff.**	**ρ**	**Av. S**	**Diff.**	**α**	**Av. S**	**Diff.**	**mod**	**Av. S**	**Diff.**
5	423.0	63.0 %	0.2	**314.2**	0 %	0.7	**331.5**	0 %	0.5	**338.5**	0 %	FIS	346.1	2.9 %
10	**259.5**	0 %	0.5	333.8	6.2 %	0.8	345.6	4.3 %	1.0	342.7	1.2 %	LDS	**336.4**	0 %
			1.0	375.7	19.6 %	0.9	346.6	4.5 %	2.0	342.5	1.2 %			

Table 6
Average Total Inventory Cost (Av. TIC) by factor.

n	cv	ρ	mod	Av. TIC	Diff.	n	cv	ρ	mod	Av. TIC	Diff.
5	0.2	0.7	FIS	46564.0	0 %	10	0.2	0.7	FIS	63990.7	0 %
5	0.2	0.7	LDS	**46485.9**	−0.2 %	10	0.2	0.7	LDS	**63519.0**	−0.7 %
5	0.2	0.8	FIS	47899.3	0 %	10	0.2	0.8	FIS	64832.7	0 %
5	0.2	0.8	LDS	**47774.1**	−0.3 %	10	0.2	0.8	LDS	**64288.7**	−0.8 %
5	0.2	0.9	FIS	48042.6	0 %	10	0.2	0.9	FIS	65689.5	0 %
5	0.2	0.9	LDS	**47967.0**	−0.2 %	10	0.2	0.9	LDS	**64682.3**	−1.5 %
5	0.5	0.7	FIS	51173.8	0 %	10	0.5	0.7	FIS	67023.9	0 %
5	0.5	0.7	LDS	**50272.3**	−1.8 %	10	0.5	0.7	LDS	**65848.3**	−1.8 %
5	0.5	0.8	FIS	52857.0	0 %	10	0.5	0.8	FIS	68725.6	0 %
5	0.5	0.8	LDS	**51641.3**	−2.3 %	10	0.5	0.8	LDS	**67025.0**	−2.5 %
5	0.5	0.9	FIS	54451.1	0 %	10	0.5	0.9	FIS	70625.8	0 %
5	0.5	0.9	LDS	**52547.5**	−3.5 %	10	0.5	0.9	LDS	**68003.5**	−3.7 %
5	1.0	0.7	FIS	58968.6	0 %	10	1.0	0.7	FIS	73413.8	0 %
5	1.0	0.7	LDS	**57873.6**	−1.9 %	10	1.0	0.7	LDS	**71701.8**	−2.3 %
5	1.0	0.8	FIS	61457.9	0 %	10	1.0	0.8	FIS	75739.7	0 %
5	1.0	0.8	LDS	**59794.5**	−2.7 %	10	1.0	0.8	LDS	**73227.8**	−3.3 %
5	1.0	0.9	FIS	65253.4	0 %	10	1.0	0.9	FIS	78695.1	0 %
5	1.0	0.9	LDS	**62476.0**	**−4.3 %**	10	1.0	0.9	LDS	**75053.5**	**−4.6 %**

Table 7 presents the factorial analysis of variance applied to the results of the experiments. The control factors n, cv, and ρ have the most significant effect on TIC. Next, we have the mod factor (FIS and LDS models). The α factor, the fourth control factor in the table, presented a lower significance level than the others, although it was quite significant (p-value = 1.09E-07).

Exploring the mod factor's effect further, we compared each of the 54 FIS instances with their counterpart LDS instances using t-tests (instance 1 × instance 55, 2 × 56, and so on). LDS instances outperformed FIS in 53 of 54 tests, 46 of which were statistically significant at a 5 % significance level. Fig. 7 shows the relative differences between LDS and FIS ($\%Diff = (TIC_{LDS} - TIC_{FIS})/TIC_{FIS}$). The relative differences range from −4.96 % (instance 24) to + 0.03 % (instance 6), with a mean relative difference of −2.13 % and a standard deviation of 1.39 percentage points.

Thus, the results presented here show the superiority of LDS over FIS for the cases considered in our experimental design. We can explain the dominance of LDS over FIS because the former is more flexible than the latter, allowing us to choose items independently of their position in the sequence. Essential to note that FIS is not, in fact, a fixed sequence but fixed cycling, which only means that the items are checked in a given sequence. During simulation, we observe "jumps" in the FIS model as well in LDS, and the difference is that in LDS, those jumps are to the "lowest day of supplies" items and not to the next item below the minimum stock (s).

As this is an important outcome of this research, we performed a sensitivity analysis to verify if the previous result holds for other costs. In Table 8, we present an example comparing instances 54 and 108 (n =

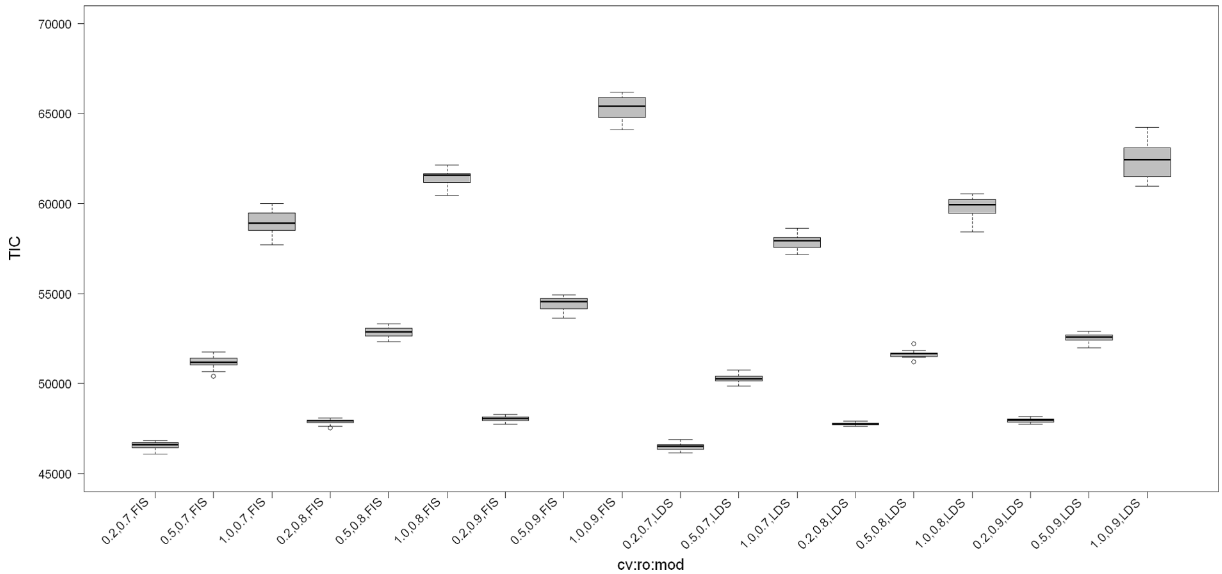

Fig. 5. Caption: Boxplots of samples with n = 5 and different combinations of cv, ρ, and mod factor levels. Alt Text: The figure shows boxplots of total inventory cost (TIC) for n = 5 and different combinations of cv, ρ, and mod factor levels.

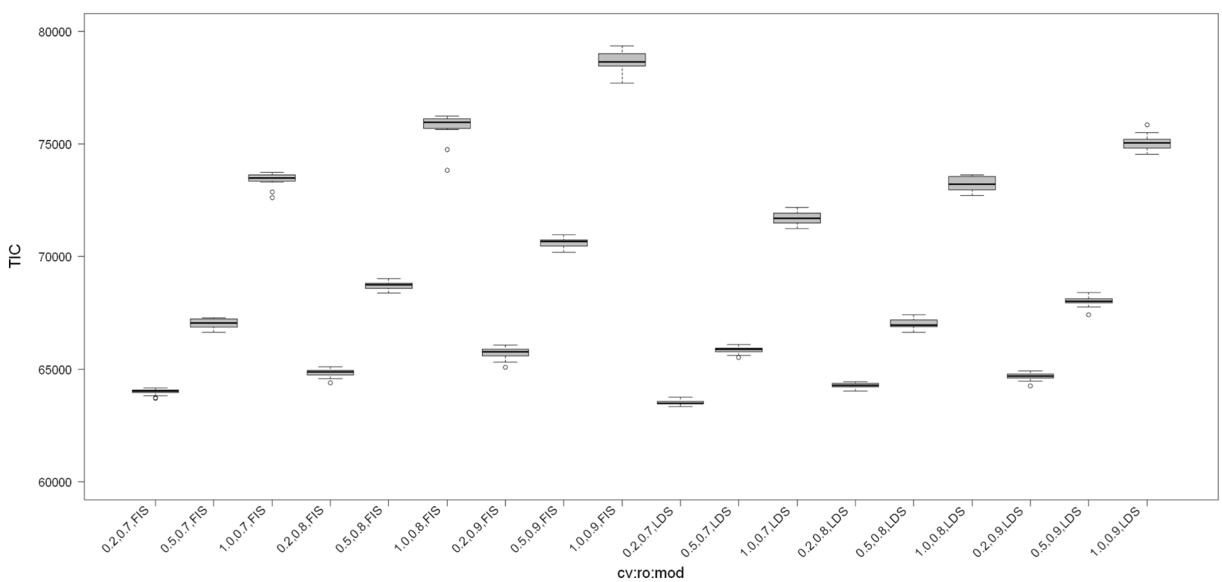

Fig. 6. Caption: Boxplots of samples with n = 10 and different combinations of cv, ρ, and mod factor levels. Alt Text: The figure shows the equivalent of Fig. 5 for n = 10.

10, cv = 1.0, ρ = 0.9, and α = 2.0) under six combinations of setup cost (sc=$250, $500 and $1000) and annual holding cost (hc=$25 and $50 per year). The results show that the relative difference between LDS and FIS did not change significantly despite changing the total cost (TIC).

Lastly, in addition to analysing the effect of operational factors on system performance, the models and methods presented in this paper can be applied to real-world scenarios, with minor adaptations. Items don't need to have the same parameters values, which will result in different control parameters for each item. The users can run the models, find the best results for each model, and decide how to run their systems. Fig. 8 compares FIS and LDS for a specific experimental condition.

7. Conclusion

This work studied the problem of scheduling and lot sizing of multiple products in a single machine, with stochastic demand and sequence-dependent setup times, called Stochastic Economic Lot Scheduling Problem (SELSP). We applied the simulation–optimization approach with the AnyLogic simulation software and the OptQuest search engine to calibrate the inventory control models, which determine the sequence and the sizes of the lots, to minimize the total inventory cost, which includes holding costs, setup costs and loss of sales/shortage costs.

In the study, we considered two inventory control policies: (i) fixed cycling (First in Sequence - FIS) and (ii) dynamic scheduling based on stock levels (Lowest Days of Supply - LDS), combined with an "order-up-to" lot sizing. The experimental design included the following factors: number of items, coefficient of variation of demand, system workload, and degree of setup increment, allowing for comparing the two control policies (FIS and LDS) under comprehensive experimental conditions. The LDS model outperformed the FIS model in all instances, the difference being higher in the more restrictive conditions, with higher demand variation and greater workload.

The present work can be extended in two directions: inventory

Table 7
Analysis of Variance for Total Inventory Cost (TIC).

Factor	Df	Sum Sq	Mean Sq	F value	Pr(>F)	Sig.
n	1	3.23E + 10	3.23E + 10	**256000.0**	<2.0E-16	***
cv	2	1.30E + 10	6.51E + 09	**51500.0**	<2.0E-16	***
ro	2	8.40E + 08	4.20E + 08	**3324.0**	<2.0E-16	***
alfa	2	4.20E + 06	2.10E + 06	16.64	1.09E-07	***
mod	1	2.65E + 08	2.65E + 08	**2099.0**	<2.0E-16	***
n:cv	2	2.61E + 08	1.30E + 08	1031.0	<2.0E-16	***
n:ro	2	3.26E + 06	1.63E + 06	12.89	3.66E-06	***
cv:ro	4	1.88E + 08	4.71E + 07	372.9	<2.0E-16	***
n:alfa	2	7.33E + 05	3.66E + 05	2.900	0.0561	.
cv:alfa	4	9.61E + 05	2.40E + 05	1.902	0.1092	
ro:alfa	4	5.96E + 05	1.49E + 05	1.180	0.3188	
n:mod	1	1.28E + 07	1.28E + 07	101.6	<2.0E-16	***
cv:mod	2	7.93E + 07	3.97E + 07	314.0	<2.0E-16	***
ro:mod	2	2.80E + 07	1.40E + 07	110.7	<2.0E-16	***
alfa:mod	2	6.14E + 05	3.07E + 05	2.432	0.0891	.
n:cv:ro	4	9.69E + 06	2.42E + 06	19.17	1.57E-14	***
n:cv:alfa	4	6.24E + 05	1.56E + 05	1.234	0.2956	
n:ro:alfa	4	2.33E + 05	5.83E + 04	0.461	0.7640	
cv:ro:alfa	8	1.35E + 06	1.69E + 05	1.340	0.2217	
n:cv:mod	2	4.75E + 05	2.38E + 05	1.880	0.1538	
n:ro:mod	2	9.63E + 05	4.82E + 05	3.813	0.0228	*
cv:ro:mod	4	9.19E + 06	2.30E + 06	18.18	8.14E-14	***
n:alfa:mod	2	1.36E + 05	6.82E + 04	0.540	0.5832	
cv:alfa:mod	4	5.04E + 05	1.26E + 05	0.997	0.4090	
ro:alfa:mod	4	3.56E + 05	8.91E + 04	0.705	0.5888	
n:cv:ro:alfa	8	1.38E + 06	1.72E + 05	1.365	0.2098	
n:cv:ro:mod	4	2.42E + 05	6.06E + 04	0.480	0.7506	
n:cv:alfa:mod	4	8.63E + 05	2.16E + 05	1.708	0.1471	
n:ro:alfa:mod	4	3.33E + 05	8.32E + 04	0.659	0.6208	
cv:ro:alfa:mod	8	3.99E + 05	4.98E + 04	0.394	0.9235	
n:cv:ro:alfa:mod	8	9.00E + 05	1.13E + 05	0.891	0.5242	

Signif. codes: '***' – 0.001, '**' – 0.01, '*' – 0.05, '.' – 0.1.

control policies and resolution methods. We considered an "order-up-to" model with freely variable lot sizes in this work. An extension would be to consider fixed or multiple lot sizes, which can be an operational constraint adherent to the reality of many factories in the process industry. Another possibility would be to consider the backorder premise rather than the loss of sales, where unmet demands generate production orders with an associated backorder cost, also a real-life problem faced by many factories. Finally, an interesting operational issue for modelling would be related to the weekly work shift. While some process companies operate 24/7, it is more common to see production stop for a day or two a week, depending on demand and inventory levels. Lot sizing in this scenario of weekly stops would also be interesting to study.

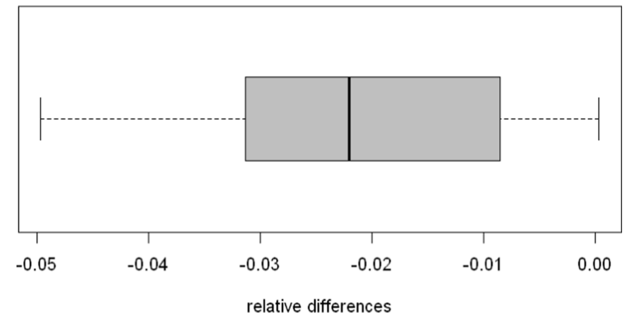

Fig. 7. Caption: Boxplots of FIS and LDS relative differences. Alt Text: The figure shows the boxplots of the relative differences between the results of the FIS and LDS models.

Table 8
Comparison of instances 54 and 108 for different setup and holding costs.

Setup Cost ($)	Holding Cost ($/yr)	Av. TIC – FIS	Av. TIC – LDS	Diff.	t	p-value
250	25	78822.0	75191.3	−4.6 %	−10.42	2.7E-05
500	25	102870.4	99262.6	−3.5 %	−13.05	1.8E-06
1000	25	137353.6	133645.8	−2.7 %	−12.59	1.5E-06
250	50	115564.4	110624.1	−4.3 %	−16.96	3.68E-05
500	50	148381.6	143463.8	−3.3 %	−14.79	8.3E-06
1000	50	194967.7	189564.2	−2.8 %	−13.58	7.7E-05

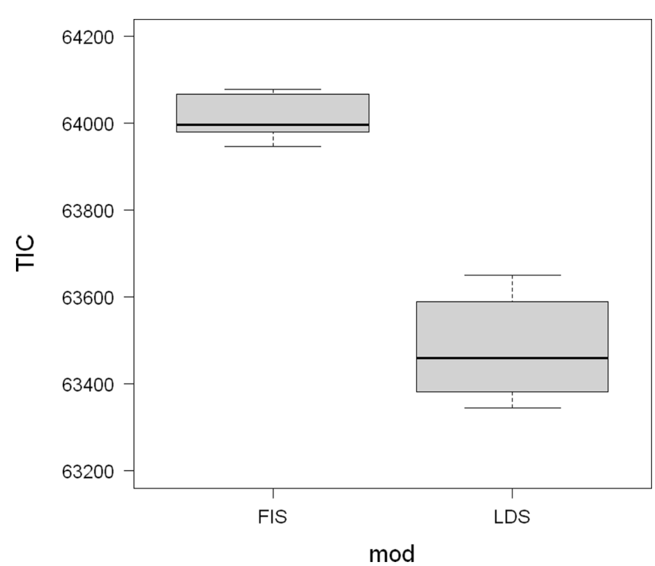

Fig. 8. Caption: Boxplots of FIS and LDS results for n = 10, cv = 0.2, ρ = 0.7, and α = 1.0. Alt Text: The figure shows the boxplots of FIS and LDS results for n = 10, cv = 0.2, ρ = 0.7, and α = 1.0.

Regarding the simulation–optimization method, we used proprietary software in this research, which may be a limitation of its practical application. A natural development of the present work would be implementing the simulation models in free software and coding a specific numerical optimization method to calibrate the inventory control parameters.

Finally, we emphasize that the models developed in this work were applied in hypothetical scenarios. Still, we believe that, because they are user-friendly models, they can be adapted and used, without great difficulties, in real systems of the process industry, thus characterizing themselves in a theoretical contribution to production planning and control in pull production environments.

Data availability statement

The data that support the findings of this study are openly available in:

Mesquita, Marco Aurélio de; Tomotani, João Vitor (2022), "SimOpt SELSP", Mendeley Data, V1, https://doi.org/10.17632/zzxc77h3gg.2

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

Data will be made available on request.

References

Altiok, T., & Melamed, B. (2007). *Simulation modeling and analysis with Arena*. Academic Press.

Beck, F. G., & Glock, C. (2019). The economic lot scheduling problem: A content analysis. *International Journal of Production Research, 58*(3), 1–18.

Brander, P., Levén, E., & Segerstedt, A. (2005). Lot sizes in a capacity constrained facility: A simulation study of stationary stochastic demand. *International Journal of Production Economics, 93–94*(2), 375–386.

Cortés-Fibla, R., P.I. Vidal-Carreras, and J.P. García-Sabater. 2015. "Considering the effect of demand diversity on the performance of different production strategies for the economic lot scheduling problem." *IESM Conference.*

Cunha Neto, E., Ferreira Filho, V. J. N., & Arruda, E. F. (2015). Stochastic economic lot sizing and scheduling problem with pitch interval, reorder points and flexible sequence. *International Journal of Production Research, 53*(19), 5948–5961.

Erlenkotter, D. (2014). Ford Whitman Harris's economical lot size model. *International Journal of Production Economics, 155*, 12–15.

Federgruen, A., & Katalan, Z. (1998). Determining production schedules under base-stock policies in single facility multi-item production systems. *Operations Research, 46*(6), 883–898.

Gascon, A., Leachman, R. C., & Lefrançois, P. (1994). Multi-item, single-machine scheduling problem with stochastic demands: A comparison of heuristics. *International Journal of Production Research, 32*(3), 583–596.

Gel, E. S., Fowler, J. W., & Khowala, K. (2021). Queuing approximations for capacity planning under common setup rules. *IISE Transactions, 53*(11), 1177–1195.

Glock, C. H., Grosse, E. H., & Ries, J. M. (2014). The lot sizing problem: A tertiary study. *International Journal of Production Economics, 155*, 39–51.

Goyal, S. K. (1973). Lot Size Scheduling on a single machine for stochastic demand. *Management Science, 19*(11), 1322–1325.

Graves, S. C. (1980). The Multi-Product Production Cycling Problem. *IIE Transactions, 12*(3), 233–240.

Holmbom, M., & Segerstedt, A. (2014). Economic Order Quantities in production: From Harris to Economic Lot Scheduling Problems. *International Journal of Production Economics, 155*, 82–90.

Hopp, W. J., & Spearman, M. L. (2008). *Factory Physics: Foundations of Manufacturing Management* (3rd ed.). Waveland Press.

Jodlbauer, H., & Reitner, S. (2012). Optimizing service-level and relevant cost for a stochastic multi-item cyclic production system. *International Journal of Production Economics, 136*(2), 306–317.

Kamath, B. N., & Bhattacharya, S. (2007). Lead time minimization of a multi-product, single-processor system: A comparison of cyclic policies. *International Journal of Production Economics, 106*(1), 28–40.

Karalli, S. M., & Flowers, A. D. (2006). The Multiple-Family ELSP with Safety Stocks. *Operations Research, 54*(3), 523–531.

Leachman, R. C., & Gascon, A. (1988). A heuristic scheduling policy for multi-item, single-machine production systems with time-varying, Stochastic Demands. *Management Science, 34*(3), 377–390.

Leachman, R. C., Xiong, Z. K., Gascon, A., & Park, K. (1991). Note: An improvement to the dynamic cycle lengths heuristic for scheduling the multi-item, single-machine. *Management Science, 37*(9), 1201–1205.

Levén, E., & Segerstedt, A. (2007). A scheduling policy for adjusting economic lot quantities to a feasible solution. *European Journal of Operational Research, 179*(2), 414–423.

Liberopoulos, G., Pandelis, D. G., & Hatzikonstantinou, O. (2013). The stochastic economic lot sizing problem for non-stop multi-grade production with sequence-restricted setup changeovers. *Annals of Operations Research, 209*, 179–205.

Löhndorf, N., & Minner, S. (2013). Simulation optimization for the Stochastic Economic Lot Scheduling Problem. *IIE Transactions, 45*(7), 796–810.

Löhndorf, N., Riel, M., & Minner, S. (2014). Simulation optimization for the stochastic economic lot scheduling problem with sequence-dependent setup times. *International Journal of Production Economics, 157*, 170–176.

Paternina-Arboleda, C. D., & Das, T. K. (2005). A multi-agent reinforcement learning approach to obtaining dynamic control policies for stochastic lot scheduling problem. *Simulation Modelling Practice and Theory, 13*(5), 389–406.

Rappold, J. A., & Yoho, K. D. (2017). Setting safety stocks for stable rotation cycle schedules. *International Journal of Production Economics, 156*, 146–158.

Rogers, J. (1958). A computational approach to the Economic Lot Scheduling Problem. *Management Science, 4*(3), 264–291.

Smits, S., Wagner, M., & de Kok, T. (2004). Determination of an order-up-to policy in the stochastic economic lot scheduling model. *International Journal of Production Economics, 90*(3), 170–176.

Sox, C. R., Jackson, P. L., Bowman, A., & Muckstadt, J. A. (1999). A review of the stochastic lot scheduling problem. *International Journal of Production Economics, 62*(3), 181–200.

Teunter, R. H., & Haneveld, W. K. K. (2008). Dynamic inventory rationing strategies for inventory systems with two demand classes, Poisson demand and backordering. *European Journal of Operational Research, 190*(1), 156–178.

Tomotani, J. V., & de Mesquita, M. A. (2018). Lot sizing and scheduling: A survey of practices in Brazilian companies. *Production Planning and Control, 29*(3), 236–246.

Tubilla, F., & Gershwin, S. B. (2021). Dynamic scheduling in make-to-stock production systems with setup times and random breakdowns: Performance analysis and improved policies. *International Journal of Production Research, 59*(1), 1–19.

Vaughan, T. (2007). Cyclical schedules vs. dynamic sequencing: Replenishment dynamics and inventory efficiency. *International Journal of Production Economics, 107*(2), 518–527.

Vergin, R. C., & Lee, T. (1978). Scheduling rules for the multiple product single machine system with stochastic demand. *Information Systems and Operational Research., 16*(1), 64–73.

Wagner, M., & Smits, S. R. (2004). A local search algorithm for the optimization of the stochastic economic lot scheduling problem. *International Journal of Production Economics, 90*(3), 391–402.

Winands, E. M. M., Adan, I. J. B. F., & van Houtum, G. J. (2011). The stochastic economic lot scheduling problem: A survey. *European Journal of Operational Research, 210*(1), 1–9.