European Journal of Operational Research xxx (xxxx) xxx

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Invited Review

Can accessing much data reshape the theory? Inventory theory under the challenge of data-driven systems

Nesim Kohen Erkip

Industrial Engineering, Bilkent University, Ankara 06800, Turkey

A R T I C L E I N F O

Article history:
Received 25 November 2021
Accepted 18 August 2022
Available online xxx

Keywords:
Inventory
Data-driven inventory problems
Inventory problem with heterogenous customers
Review

A B S T R A C T

In this review, we discuss the data-driven systems and their effects on the implementation of the inventory theory. After overviewing the theory briefly, we group the data-driven approaches to simplify exposition. We consider the use of available data to estimate the parameters of more complex models, and propose developing the theory in that direction, as well. As a pedagogical example, an extension of the standard EOQ model with heterogenous customers is presented. The review proposes a research agenda for inventory problems and concludes with discussing challenges for the future.

1. Background and motivation

We consider environments with similar (if not the same) inventory/production decisions made repeatedly over time. Two characteristics of such environments are prevalent: A methodology for supporting the final decision (possibly leaving some room for the decision maker's own experience) and the existence of information on the current status of the operations, including some data. The methodology of making inventory decisions can include the findings of a related theory or be driven by the data. In other words, the theory and data-driven systems can attack the same decision problem.

In this overview, we discuss the developments in data acquisition, data processing and data-driven systems and the effects they are likely to have on implementing the theory. In specific, we consider decision-making environments where inventory theory may be implementable. The discussion is not restricted to results directly used in practice but includes the intuition gained from learning the theory.

The overview excludes discussing the statistical and econometrical methods or their mechanics. The growth of the research on the quality of estimators, structural estimation methods and likewise areas deserve a separate review article. Furthermore, the application of those methods to inventory-related problems with specific research questions is not covered in this review. Note that any decision problem will require knowledge of the statistical and econometrical methods and their mechanics at different levels of sophistication.

What we aim in this review is first to discuss theoretical and data-driven approaches. More specifically, for theoretical approaches, we summarize characteristics of the research work on the theory to support inventory decisions. For the data-driven strategies, we group the literature into main categories and discuss the elements of the suggested approaches. The second aim of the review is to point out the possibility of working with more sophisticated inventory-theoretic models. Note that these models would not have much practical relevance if there were no data for validation. The last objective of the review is to discuss the challenges and suggest possible research directions for inventory-related problems.

For the purposes outlined above, we define data acquisition and data-driven systems to possess at least one of the following characteristics:

1. Availability of direct and features data related to system inputs: In other words, relevant data is analogous to big data.
2. Availability of knowledge and techniques to process data in various ways and knowhow on meaningful interpretations: We presume that we can implement the statistical and econometrical methods.

Section 2 presents a historical sketch of inventory theory (starting with the 1950s), followed by a relatively shorter part that describes extensions of the classical theory under various data/information levels. These subsections do not intend to contain a thorough literature review. The section ends with a brief discussion on how the theory has been implemented (or presumed

E-mail address: nesim@bilkent.edu.tr

https://doi.org/10.1016/j.ejor.2022.08.024

for implementation) to solve real-life problems. In Section 3, we present data-driven approaches applicable for inventory-related decision-making. This part will be a kind of literature review without a claim on the exhaustiveness. The following section considers more general representations of the theoretical inventory models as another path to utilize the available, abundant data. As a pedagogical example, the derivation of the classical EOQ problem under backorders and heterogeneous customers is presented. Section 5 presents an overview of challenges for the future related to the development of data-driven approaches. The last section proposes a research agenda for inventory-related problems.

2. Inventory theory

2.1. A personal overview of the classical theory

Arrow (2002) gives a summary of the history of mathematical inventory theory. Of course, Arrow being one of the most important characters in the development of the theory, the document is more than a personal history; it is an account of the establishments (see Girlich & Chikan, 2001 for a review on inventory theory developed in earlier years). One can follow the original work in several different books and edited books published in the 1950s, and 60s, see Whitin (1953), Karlin, Scarf, & Arrow (1958), Arrow, Karlin, & Suppes (1960), Karlin, Scarf, & Arrow (1962), Hadley & Whitin (1963), Scarf, Daniel, Gilford, & Shelly (1963). Many studies followed the pioneering ones, as reported by two early review papers published around the same time, which consider the theory developed from two perspectives of decision making: Wagner (1980) takes a systems perspective and overviews past research work within the directions it specifies. The second review article is an account based on different environmental assumptions and targets an audience who will pursue research to develop models in possible known directions Silver (1981). It is interesting to note that their perspectives for future work have clear similarities with consequent research in inventory theory in the following years (there are reviews and OR/OM handbooks that summarize the state-of-the-art and are important guides for further use. See for example Balcik, Bozkir, & Kundakcioglu (2016), Basten & van Houtum (2014), Bijvank & Vis (2011), de Kok et al. (2018), Engebrethsen & Dauzère-Pérès (2019), as well as several chapters in Graves, Kan, & Zipkin (1993), Graves & de Kok (2003), Cochran, Cox, Keskinocak, Kharoufeh, & Smith (2010).

2.2. Broadening foundations with data/information

Of course, there have been many studies on inventory theory in the last few decades. In this subsection, we mention the ones different from the classical theory regarding data/information requirements.

Even in the early development days, researchers were interested in including additional features into the models to explain randomness better. However, those models did not constitute the mainstream at the time. We believe it is most likely because of their relative complexity and difficulty in implementation. Price-dependent demand as a natural extension of the neo-classical economic theory by Karlin & Carr (1962), backorder-dependent demand structure including empirical justification by Schwartz (1970) and inventory dependent demand considerations by Baker & Urban (1988), Urban (2005) are examples of specific features considered. Following the start of the e-commerce era in the 90s, researchers started to model inventory planning problems with additional feature information. Models used features in a variety of ways Hariharan & Zipkin (1995) use of advanced demand information, Tan, Güllü, & Erkip (2007) for advanced and uncertain demand information, Gavirneni, Kapuscinski, & Tayur (1999), and Savaşaneril & Erkip (2010) for the structural use of information made available by use of VMI and other similar approaches. Note that these studies presume that such data/information is either available or possible to collect.

A different line of research that does not require estimates of distributions or demand parameters is robust optimization, another approach to broaden the foundations. An early example of this type of work was named the min-max problem, a solution can be found in Scarf (1958) and two more recent robust optimization examples in Jackson, Muckstadt, & Li (2019), Shapiro & Xin (2020). For the remaining paper, we concentrate on models that explicitly require demand data to operate.

2.3. Theory developed and implementation

In this subsection, we characterize the theory developed and finally ended up with the assigned role of data for implementation. In the classical theory development, the connection of the theory developed and data is hierarchical; data is used to satisfy the input requirements of the models.

Problem environments are mostly “stylistic,” as researchers emphasize main tradeoffs rather than details. Of course, the other side of the coin in having stylistic models is the motivation to keep the environments as general as possible, rather than letting peculiarities drive the results. One can also call these models “fundamental building blocks.” The problems considered are the abstractions of decision situations that repeat over time. Usually, the modeling environment excludes part of the system and only finds the “significant” parts as observed or hypothesized. Researchers made some simplifying assumptions regarding parameter structures to ease the solution, such as linearity of costs. Additionally, they used the standard premise of representing demand as a random variable. Nevertheless, no specific distribution or specific parameters representing distribution and no restrictions on the parameter values are assumed.

The purpose of academic research has been to find structural results. Sometimes these structural results paved the way for policy structures such as (s, S), (R, Q), optimal in their respective environments. In some cases, solutions are found in closed form, or an algorithm is proved to converge to obtain the optimal solution, leading to engineering applications. Another line of research for more complex problems has emerged: By analyzing relatively more uncomplicated structures, researchers showed that some bounds exist on the worst possible performance of these complex problems. The bound can be either on the objective function or policy parameters. Finally, researchers working on the partial characterization of the optimal policy (inserting conditions where the structural properties hold) is a final resort for difficult situations. Note that the significant influence is generalization via portability of structures to more complex environments in all these findings.

Structural properties may not necessarily yield efficient computational algorithms. Research on computational algorithms complements the style mentioned above. The scientific work on the complement becomes especially crucial when one expands the scale of the problem to multiple items. Similarly, when computation becomes expensive, more efficiently computable heuristic methods are utilized, likely to use some of the structural properties mentioned above and any other knowledge or intuition available.

When implementing the theory in practical decision-making situations, two distinct ways can be observed: The first and the obvious one is implementing the theoretical results in a practical case with minor adjustments. When minor adjustments are inadequate, one can still use the accumulated knowledge, but using the intuition gained becomes essential. The second one is statistical modeling, assuming sufficient data is available. Data processing in its classical meaning requires empirical modeling and statistical analysis research. For applying inventory theory, the work needed includes the estimation of parameters and demand distributions (first family and then family parameters), all performed using available data and information.

Hence, one can think of the approach for solving inventory-related decisions in two stages: the first stage is collecting and using the available data to estimate various input parameters. The second stage uses the estimated data structures in the first stage to optimize for the decision variables of the model. In short, we name this approach "estimate then optimize" (with the abbreviation EtO), indicating that we have two different stages.

3. Inventory problems: Data-driven methods

21st century marks many things for different scientific fields. One common feature is the advance of information technologies. There are numerous references - a few review-type related articles are Arunachalam, Kumar, & Kawalek (2018), Choi, Wallace, & Wang (2018), Hazen, Skipper, Boone, & Hill (2018), Mortenson, Doherty & Robinson (2015), Nguyen, Li, Spiegler, Ieromonachou, & Lin (2018) and Wang, Gunasekaran, Ngai, & Papadopoulos (2016).

In OR, we tend to name these developments with different phases of analytics: Descriptive analytics, the use of the technology for monitoring and storing data, and visualizing data. In predictive analytics, we can further assess the effects of the current decision structure in the future follow demand changes, follow state changes, and observe anomalies. Statistics, data mining, and machine learning are among the essential methods to enable such predictions. Prescriptive analytics is after understanding what will happen under "optimized" decisions and working on what should have been the decision. Prescriptive analytics, in this respect, is more dynamic and interested in the evolution of decisions. Therefore, traditionally known as complex problem structures with multiple objectives, multiple decision-makers fall into this category. There are several advancements in this analytics category, even if limitations on the frequency of decision-making and partial availability of relevant information make real-time implementation challenging.

These developments have increased the use of methodologies that implement the so-called data-driven approaches. The studies overviewed in this section use data as an interface. There are various definitions of "data-driven": Here, we use it as the defining part of a context where decisions are based predominantly on the available data. In other words, data is the primary interface used to arrive at decisions. However, as we see in this section, there are various conceptualizations either in nonparametric models or through structural assumptions driving the use of data.

To review data-driven approaches, we categorize the studies on inventory problems into five groups. We use the term "inventory" casually, as some of the work cited may contain work on different decisions. There is no claim concerning the non-jointness of the groups. However, we think the central assertion representing the approaches is reasonably distinct in different groups. The groups are defined and studied further below:

Application of Artificial Intelligence approaches with no policy restriction

Artificial Intelligence (AI) plays an important role in transforming business practices (See Brynjolfsson & Mcafee, 2017 and Brynjolfsson, Hui, & Liu, 2019 as examples). AI, in this review, represents a set of technologies and tools that can process and analyze available information and "optimize" actions to achieve prespecified goals and perform tasks. AI has been around for some time. However, only recently have the advances in machine learning become a tool, fueled by the developments in computing power and the acquisition of digitized data.

Here, we consider using AI techniques to solve non-recurring inventory-related problems. In other words, we delay the discussion of learning to a later group.

These methods do not explicitly estimate demand distribution or use any structural property of the optimal solution (or optimal solution structure for a more straightforward problem). They utilize the available data (which may include feature data); hence, the formulation becomes similar to robust optimization with the data input.

Bertsimas & Thiele (2006a) is one of the essential early examples of this group of approaches. They use robust optimization ideas to build an equivalent model without any uncertainty. Under some mild conditions, structural properties are investigated regarding the policies obtained. Bertsimas & Thiele (2006b) report more results for environments other than the newsvendor. Ban & Rudin (2019) (also mentioned in the next group, "Nonparametric approaches under a policy") consider the newsvendor problem with features data. Utilizing empirical risk minimization algorithms and kernel optimization method, they show that they obtain the same objective function for the newsvendor problem. The approach resembles Liyanage & Shanthikumar (2005) which will also be mentioned in the next group, although it is more general and uses practically no information on the structure of demand.

Nonparametric approaches under a policy

These methods use at least one property of the "optimal" solution obtained under some assumptions. A good example would be the newsvendor problem with the optimal percentile of the cumulative demand distribution. The aim is not to explicitly estimate this distribution but to use the data to develop the decision variable values sought.

A subgroup of such studies is related to applying more traditional mathematical programming models with the knowledge of a policy that drives the decision-making. Iyer & Schrage (1992) is an early example of such data-driven approaches. It is well-known that a time-dependent (s, S) policy is optimal for the deterministic multi-period dynamic inventory problem. Using this structural property as the starting point for data-driven optimization, one can construct a mathematical model to find time-dependent policy parameters. Beutel & Minner (2012) propose a linear programming formulation where newsvendor problem with features data is solved. The implication is that the standard percentile solution for the newsvendor problem is now a function of the features and hence dependent on features. Features do not need to be identified, as they will likely be problem-dependent. They consider service level benchmarks to compare models that solve the same problem. The LP problem solves the demand model's order quantity decision and regression coefficients as a function of features (independent) data. van der Laan, Teunter, & Romeijnders (2019) further consider a service level approach with chance-constrained formulation (part of robust optimization) to prevent over-fitting with a smaller-sized data sample. Turgut, Taube, & Minner (2018) consider a very detailed retail inventory problem with multiple items, joint ordering, backroom considerations, and case pack-size limitations. The well-known can-order point policies with the case pack size limitations are the implemented policy. A mixed-integer program is formulated to minimize total cost over known demand incidences to determine the policy parameters.

A second subgroup of the nonparametric approaches uses statistical tools to solve for the policy parameters. We start with studies on the newsvendor problem. For the newsvendor problem, Bertsimas & Thiele (2005) proposes a ranking procedure called empirical quantile to find the historical demand value that satisfies the desired percentile. Levi, Roundy, & Shmoys (2007) generalize the approach by suggesting the use of the "sample average approximation" (SAA) method to estimate the demand value corresponding to the percentile of the empirical demand distribution. They also provide some bounds on the quality of the solution. Levi, Perakis, & Uichanco (2015) present improved bounds on the same problem. Ban & Rudin (2019) consider the newsvendor problem with features data. Liu, Letchford, & Svetunkov (2022) extends Ban & Rudin (2019)to non-linear objective function forms. Under censored demand data, there are few studies available; please refer to Huh, Levi, Rusmevichientong, & Orlin (2011) that uses the Kaplan-Meier estimator. Neghab, Khayyati, & Karaesmen (2022) extends Ban & Rudin (2019) when features are not observable. There is a limited number of studies for environments other than the newsvendor. Levi et al. (2007) propose a dynamic programming procedure to deal with multi-period inventory problems, this time using the knowledge of optimal base stock policies. They also show some bounds on the quality of the solution obtained. Ban (2020) shows the existence of objective function bounds for limited data (finite sample size) for the general multi-period inventory problems, using (s, S) policy structure as the basis. Huh & Rusmevichientong (2009) considers a multi-period lost sales problem. They propose using a stochastic gradient method under the policy structure that utilizes newsvendor-like percentiles.

We end this subsection by stating the significance of using features data in addition to demand or sales data, as is the case in Beutel & Minner (2012). All the above methods can be extended to utilize feature data. Note that a combination of statistical and classical mathematical programming approaches is used simultaneously to solve the formulated problem.

Parametric approaches: Focus on the final operational objective function

The approach focuses on defining a measure and operating on that measure. The measure corresponds to a function of the deviation from the optimization solution. The idea was presented by Liyanage & Shanthikumar (2005) proposes using a function that combines the estimation and optimization tasks and optimizes this function to find a decision rule. Of course, one needs to make assumptions on the statistic's functional form to use and proceed; in this work, they confine themselves to statistical estimators of the "optimal" solution. In their study, they implemented the approach for the newsvendor problem with exponentially distributed demand. Finally, they consider two approaches: Estimating the parameter of the exponential distribution from data and obtaining the empirical distribution of the data and using SAA. Then they compare these two with the approach using operational statistics and show that the a priori expected profit is better than the other two approaches.

Chu, Shanthikumar, & Shen (2008) propose a more general optimization function so that the restrictive character of the statistical estimation function is removed. However, knowledge of the shape parameter is still required. To follow the same line of work, Ramamurthy, Shanthikumar, & Shen (2012) first suggests a heuristic based on the previous work and corrects the results to consider the shape parameter. Besbes, Phillips, & Zeevi (2010) use a demand function derived from customers' perspective in a newsvendor setting and approach finding optimal decision (purchasing quantity) statistically using the individual consumer's utility function, as well as the available data. A more recent work, Siegel & Wagner (2021) uses an asymptotically unbiased estimate of the expected profit function. The approach requires the use of an adjustment term derived from the distribution. In general, an approximate adjustment term is obtained by the Taylor series analysis. In a particular case of demand being exponentially distributed, an exact adjustment term can be specified.

Note that one can generalize the approach, and some of the work mentioned in the previous subsections may also fit the framework mentioned here. Beutel & Minner (2012) is an example that combines statistical estimation under the original problem's objective function. Nevertheless, the approach to defining an "operational objective" function is more general, although all the work described here uses the newsvendor model as the basis. Further work is needed to extend the results to other environments.

Approaches with learning

Traditionally, Bayesian learning was the main tool for learning the unknown demand distribution. The procedure starts with a prior distribution and improves the prior with incoming new data. These methods are parametric, as at least one parameter is learned and used for implementation. Azoury (1985) is one of the early works that implements the Bayesian approach for a general inventory problem. More recent studies are Chen & Chao (2019) and Ding, Puterman, & Bisi (2002) for the censored data case.

On the other hand, the new trend is to use search mechanisms with updates (updating corresponds to machine learning). When demand distribution is unknown, a learning algorithm provides policies that depend on the observed information and current time. Hence, one may call this "policy learning" compared to parameter learning. The policy is then updated with the previous period's information. The objective function used for such learning will vary, but it is often represented as a function of regret. There are usually two main streams of techniques: Reinforcement Learning and Statistical Learning. Here we focus on the implementation of these approaches for inventory models.

Bertsimas & Kallus (2020) applies machine learning methods for a general optimization problem. The data may include feature data over the standard demand data. The newsvendor problem is used as an example in the study. Chen, Chao, & Ahn (2019) propose a nonparametric learning algorithm that considers ordering and pricing decisions in a multi-period environment. Gijsbrechts, Boute, Van Mieghem, & Zhang (2022) implement a deep reinforcement learning algorithm for the dual-sourcing and dual-mode problems, which are known to have high dimensionality issues while solving with classical dynamic programming.

Combination of approaches

There are several studies where combinations of the techniques described above are implemented. The basic idea is to use a beneficial approach under a circumstance to improve another approach's performance. Here we report a few of those studies:

Saghafian & Tomlin (2016) study a problem where some information on the moments or tail of the demand distribution is available. Under these conditions, an enhanced algorithm with Bayesian learning works well for the newsvendor problem. This study is particular as it may pave the way for most of the earlier approaches to be improved by such information. Huber, Müller, Fleischmann, & Stuckenschmidt (2019) propose novel approaches to solve the data-driven newsvendor problem based on machine learning and quantile regression. Oroojlooyjadid, Snyder, & Takáč (2020) implement a deep learning algorithm for a newsvendor problem when features data for the demand is available. Hence, forecasting demand using features and inventory decisions are embedded into a single problem, which is then tackled by a deep learning algorithm with the knowledge that the newsvendor problem's optimal solution is a percentile of the demand distribution. De Moor, Gijsbrechts, & Boute (2022) consider a deep learning approach for perishable inventory problems. However, they "reinforce" the algorithm by using existing well-performing heuristics. Yuan, Luo, & Shi (2021) consider a periodic review single product inventory system with fixed cost under censored demand. Note that the existence of fixed cost results in loss of joint convexity of the expected cost function over the decision parameters, in this case, (s, S) policy parameters. A stochastic gradient descent algorithm is employed in an algorithm that minimizes regret. Huber et al. (2019) combines AI-type learning with optimization models and implements it to the newsvendor problem with features data available.

The premise for implementation: When it comes to implementation, the methods outlined in this section may vary in terms of the details. The following characteristics are common in almost all the approaches:

1. Data processing and optimization are performed either simultaneously or iteratively. This characteristic is what some authors name as the "integrated" approach, though environments and what is meant by integration varies (see Goltsos, Syntetos, Glock, & Ioannou, 2022; Huber et al., 2019; Liu et al., 2022; Neghab et al., 2022; Syntetos, Babai, Boylan, Kolassa, & Nikolopoulos, 2016 as examples).
2. The objective function selected for the optimization may not necessarily be the objective function that will be utilized by the corresponding inventory problem as described in the classical theory.

These studies, when compared to the "estimate then optimize" (EtO) approach mentioned before, are different as they attempt to handle the two stages (estimation and optimization) simultaneously. We name this approach "estimate & optimize" (with the abbreviation EO). EO is almost similar to the "predict-and-optimize" idea in Elmachtoub & Grigas (2022). The idea is the same, though we prefer to use EO as inventory problems may have other parameters to estimate, and the outcome of the process may not always be used for prediction. Note that this similarity in approaches is a good indication of a possible cross-fertilization of available knowledge.

4. Inventory problems: Better representation of uncertainty

Traditionally, demand and possible uncertain quantities have been modeled in the simplest possible fashion: an outcome set with an accompanying probability distribution. However, in another area of operations management forecasting - we analyze uncertain factors using several techniques which investigate the effects of possible features selected to be related to the uncertain phenomenon we are studying. The critical assumption is that all these features are exogenous to the decision-making environment and are represented in the demand distribution in an aggregate way. In the previous section, data-driven methods made use of these features, and here we suggest establishing analytical models which include these features. The models are expected to be more complex and harder to solve. Nevertheless, with the available data, validation of such models may be possible, a situation that did not exist at the early stages of the theory development.

This section considers an essential subset of such models: demand to be a function of the decisions, rather than exogenous (see Huang, Leng, & Parlar, 2013 for a general discussion of demand functions in decision modeling). We confine ourselves to single-location inventory models in the following subsection. Note that the literature on competitive inventory problems also models the exogeneity of the demand. Those models reflect a natural consequence of the competition rather than a result of the customer choice. Hence research work on competition is not considered here. We conclude the section by presenting a simple model to demonstrate the idea.

4.1. Endogenous demand inventory problems

Hotelling (1929) is one of the earlier works which considers a single additional feature "distance" affecting the decision of an individual consumer. In this way, "demand" is represented as a function of distance and a random component. Of course, adding such features is attractive conceptually but complex in implementation as one requires to find relevant data on the feature and a representing metric of the feature to support the decision-making process. Another way of observing the same phenomenon is to consider a subset of these features as additional decision variables for the same situation. The exogenous demand concept is then replaced by endogenous demand decisions directly affect the demand. A well-known example is a relation between price and demand. Dana Jr. & Petruzzi (2001) and Wang (2022) consider endogenous demand as affected by pricing. In recent years, the revenue management area has grown as it has been possible to store and utilize vast data to realize the requirements of such data-hungry models.

These models can also be considered under the title of consumer choice models. The assumption that demand is exogenous is removed in such models, and decisions affect the demand in a defined (regardless of being deterministic or stochastic) way. There are many such examples, especially in service industries. For some recent work on the general consumer choice modeling, we refer the reader to Gallego & Topaloglu (2019) and Feng, Li, & Wang (2017). The inventory literature with consumer choice is more limited, but there are some recent work; see Farahat & Lee (2018), Transchel, Buisman, & Haijema (2022) and Martínez-de Albéniz & Kunnumkal (2021) as examples.

Implementing such models is comparable to the classical inventory models, with an additional load for the estimation work and the increasing difficulties in obtaining practical structural characterizations. More recently, a lot of researchers considered new approaches for the estimation problem (see, Akşin, Ata, Emadi, & Su, 2013; Gallino, Karacaoglu, & Moreno, 2022; Hathaway, Emadi, & Deshpande, 2022; Hu, Allon, & Bassamboo, 2022; Jagabathula, Rusmevichientong, Venkataraman, & Zhao, 2022 and Musalem, Olivares, & Schilkrut, 2021). Note that all are very recent examples that show the richness of the research in the area. Nevertheless, these efforts may be worthwhile as it is likely to learn more with less stylization.

The models specified above represent situations with data used to feed the model, as described with EtO in Section 2.

A model with simple assumptions is presented in the next subsection to illustrate the idea. This example shows the possible use of consumer choice models for the classical EOQ problem environment. We model this case and obtain a set of meaningful structural results. We name the model pedagogical since it is based on the EOQ model, the first model in inventory taught in related courses. It is interesting to note that the properties obtained for the classical EOQ model will be a part of the structural results.

4.2. A simple inventory model with backorders and heterogeneous customers

There is fixed ordering cost K and inventory holding cost h/unit/unit time. Backorders are allowed with a positive backorder cost p/unit/unit time. Lead-time for procurement is assumed to be zero. The total available demand rate is defined as Λ. Let Q denote the ordering quantity and b denote the maximum number of backorders we allow. In this standard problem average total cost, denoted by $\mathrm{ATC}(Q, b)$, is minimized. The optimal solution for this problem is given by two equations:

$$Q^* = \sqrt{\frac{2K\Lambda}{h}}\sqrt{\frac{h+p}{p}} \tag{1}$$

$$b^* = \frac{h}{h+p}Q^* \tag{2}$$

Furthermore, the optimal average total cost is given by

$$\mathrm{ATC}(Q^*, b^*) = \sqrt{2K\Lambda h}\sqrt{\frac{p}{p+h}} \tag{3}$$

Of course these solutions are only valid with the below implicit assumptions additional to the ones present in the standard EOQ model:

Assumption 1: Total profit margin received from sales is larger than the cost of operation. Hence if r is defined as the revenue per unit sold and c as the unit cost, the following condition holds:

$$\Lambda(r-c) \geq \sqrt{2K\Lambda h}\sqrt{\frac{p}{p+h}} \tag{4}$$

Of course, we expect this to hold in general, so at this point, it seems like a technicality. However, also notice that given all parameters are constants, there is a minimum demand rate below which is no longer economical to operate. Call this min demand as $\Lambda_{\min}$. Hence the following equality can be found for this minimum demand.

$$\Lambda_{\min} = \frac{2Kh\,(p/(p+h))}{(r-c)^2} \tag{5}$$

Furthermore, we can also specify the $p_{\min}$ as the minimum value of backorder cost per unit per unit time, to ensure the profitability of satisfying all demand, as:

$$p_{\min} = \frac{h}{\frac{2Kh}{\Lambda(r-c)^2} - 1} \tag{6}$$

Assumption 2: Presume that if there is a lost sale, the profit margin per unit $(r-c)$ is lost. Under these circumstances, we would not consider any lost sales in this system as selling means profit and hence it is not economical to have lost sales. This follows the result in Assumption 1 and Eq. (4) presented above.

Assumption 3: All the customers have the same sensitivity (or insensitivity) towards backorders. In other words, they are homogeneous. Furthermore, they accept any delay for their order to be fulfilled as the maximum waiting time is dictated by the seller. Let tf be defined as the maximum delay dictated by the seller for demand fulfillment. According to the described model this term is given by the following equation:

$$tf = \frac{b^*}{\Lambda} = \sqrt{\frac{2Kh}{\Lambda(h+p)p}} \tag{7}$$

Now, we are ready to remove Assumption 3 and present a model for heterogeneous customers. Before proceeding, we note that when we have heterogeneous customers, Assumption 2 is still valid for the seller, but some customers may prefer not to buy the item. Hence there may be lost sales. Similarly, Assumption 1 will be valid as long as rate of satisfied demand is not below $\Lambda_{\min}$.

We define a utility structure for each customer. We use the following utility function:

$$\mathrm{U}(t,d) = u - r + d - f(t) - \eta \tag{8}$$

where:

$\mathrm{U}(t,d)$ is the utility that a customer will get from buying this item, as a function of two parameters, as we presume all the other parameters of the environment are already discussed and set: t, is the maximum backorder duration that a customer will observe, and d is the discount that the seller is willing to give to the customer to ensure that she is willing to buy even with a delay. We assume without loss of generality that the utility of not buying the product is zero.

$f(t)$ is any non-decreasing function of time t, representing the disutility of any customer with the delay.

ν is the valuation of the customer for the product received; we assume this term to be the same for all customers. However, we differentiate customers by a random taste parameter, η.

η is the taste parameter, defined as a uniformly distributed random variable representing the disutility of being a buyer in this monopolistic environment and is a reaction that can be attributed partly to the possible delay in satisfying the demand. We consider $\eta \sim \mathrm{U}(0,\mathrm{B})$.

Without loss of generality, we assume that $\nu - r$ is always positive. Hence using the distribution of η, one may come up with the probability of any customer demanding the item, as described below.

$$\text{Prob}\{\text{No demand}\} = 1, \qquad \text{if } \nu - r + d - f(t) \leq 0 \tag{9}$$

$$*\text{Prob}\{\text{No demand}\} = 1 - \frac{\nu - r + d - f(t)}{B}, \quad \text{if } 0 \leq \nu - r + d - f(t) \leq B \tag{10}$$

$$*\text{Prob}\{\text{Demand}\} = \frac{\nu - r + d - f(t)}{B}, \quad \text{if } 0 \leq \nu - r + d - f(t) \leq B$$
$$\text{Prob}\{\text{Demand}\} = 1, \qquad \text{if } \nu - r + d - f(t) \geq B \tag{11}$$

Define t_1 to be the maximum value for the maximum delay such that for all $t \leq t_1$ condition defining Eq. (11) holds. Similarly, define t_2 to be the minimum value for the maximum delay such that for all $t \geq t_2$ condition defining Eq. (9) holds. More specifically, we can write the following identities:

$$f(t) \geq \nu - r + d - B, \qquad \text{if } t \leq t_1 \tag{12}$$

$$f(t) \leq \nu - r + d, \qquad \text{if } t \geq t_2 \tag{13}$$

As a result, we determine the behavior of any heterogeneous customer as follows:

- If the maximum delay in satisfying demand is greater than or equal to t_1, demand is 0.
- If the maximum delay in satisfying demand is less than or equal to t_2, demand is 1.
- For any maximum delay in satisfying demand is in the range $[t_1, t_2]$, we have positive probabilities for both to have demand and no demand.

We employ a standard equilibrium analysis to find possible solutions where the customers and the seller would agree.

The final step of the derivation is to bring two sides of the transaction together: On one hand, we determined the operational policy of the seller. On the other hand, we determined the behavior of each customer, of course dependent on the random taste parameter with the distribution assumed. Now to initiate the analysis we make the following assumptions:

Assumption 4: Customers are independent, and do not collude. Hence total demand for a given maximum delay to satisfy demand and discount can simply be defined as the total available customers multiplied by the expected demand of a customer.

Assumption 5: We assume that even if some part of the demand is lost, the remaining demand still has a constant rate. This means that the rate becomes slower but still constant.

Assumption 6: We assume that Assumption 1 still holds. We further assume that the minimum demand defined in Eq. (5) is never observed as t_1 prevents those cases. Of course, this assumption is just to make things simpler and avoid a trivial solution.

Now we can further analyze the joint model: Define $\mathrm{ATP}_\lambda(Q,b)$ as the average total profit function for a given demand rate λ after introducing unit revenue and unit cost values. Using Eq. (3) we obtain

$$\mathrm{ATC}_\lambda(Q^*,b^*) = (r-c)\lambda - \sqrt{2K\lambda h}\sqrt{\frac{p}{p+h}} \tag{14}$$

Similarly, we can write the following equation representing the maximum backorder duration dictated by the seller for demand fulfillment for a given demand rate λ, tf_λ following Eq. (7).

$$tf_\lambda = \sqrt{\frac{2Kh}{\lambda(h+p)p}} \tag{15}$$

$ATP_\Lambda(Q^*, 0)$ defines the average total profit function of EOQ with the original demand rate. We then find t_3, which is the maximum value for the delay seller will permit as otherwise average total profit of operating an EOQ is greater. To find t_3 we use Eq. (15), Eq. (14), and $ATP_\Lambda(Q^*, 0)$. A quadratic function of t_3 is obtained and can be solved for a unique positive t_3, as it has one positive and one negative root.

Now, we can determine the behavior of seller as follows:

- For a given demand rate λ, if the associated tf_λ value given by Eq. (15) is more than t_3 seller prefers not to backorder and use EOQ as it yields a better average total profit value.
- For a given demand rate, if the associated tf_λ value given by Eq. (15) is less than t_3 seller prefers to use the backorder model as it yields a better average total profit value.

Bringing all the above together and assuming that $f(t)$ defined in the utility function is an arbitrary function, one can show that the equilibrium value for the maximum delay to satisfy customer demand, t^E is one of the below:

- $t^E = \min\{t_1, t_3\}-$ this is the case where the seller decides not to backorder at all. None of the demand is lost.
- $t^E = t_2-$ this is the case where the seller operates with backorders and still none of the demand is lost.
- $t^E \in (t_2, \min\{t_1, t_3\})-$ There might be multiple equilibrium points depending on the functional form of $f(t)$. In this case among the equilibrium points the one which yields the maximum average total profit for the seller will be enforced. Note that any of these solutions will yield some lost customers.

As a special case, if $f(t)$ is defined as a quadratic function of t, then one can show that there is a unique equilibrium.

<u>What do we learn from this exercise?</u>

The classical EOQ model gives two critical messages: The first is that it simply outlines the trade-off in an inventory system. The second one regards the robustness of the solution concerning average total cost: any reasonable mistake that one makes in estimating one of the model's parameters results in a relatively more minor effect on the average total cost term. Nevertheless, the intuition gained is critical. On the other hand, all the above statements are likely applicable to the model derived in this subsection. In the developed model, the waiting time preferences of the customers dictate the possible equilibrium solutions. Nevertheless, the model requires additional effort in estimating parameters (traditionally challenging to estimate). New techniques have been developed to increase the chances of having reasonable and consistent estimates for those parameters.

One significant finding is that the EOQ model of describing a seller is still valid in the new environment, as the seller's structural properties will still be the same for a given demand level. Of course, the model presented is more complicated, and the demand is the outcome of an equilibrium imposed by the customers' sensitivity to the waiting time. Another observation is that a few of the individual results for the EOQ come together to describe the given equilibrium - in this case, the EOQ model with no backorders allowed, EOQ with backorders and finally, EOQ with partial backorders and lost sales.

There are challenges in the presented model. We presumed a discount given to customers to ensure that some will be willing to buy even with a delay. All the results are for a given discount value assumed to be set initially. One may wish to analyze further to find the "optimal" discount to offer. However, this issue might be tricky as the discount is considered part of the backorder cost, and there is a minimum backorder cost stated by equation (6) that allows the system with backorders to be economically viable for an arbitrary demand level. Hence optimization over possible discount values which are economical and, at the same time, feasible may be a complicated issue.

5. Challenges for the future

Most scientific fields are faced with the challenges brought by data-driven systems. Many studies are available to discuss various aspects of these challenges. We mention some of these aspects below.

Operations Management Field

There are several recent overviews regarding future expectations of the OM field. As inventory theory is a part of the field, we believe these overviews are essential sources of future outlook. Cachon, Girotra, & Netessine (2020) emphasize the importance of impactful OM and interestingly advocate generality of results with more actively engaging with diverse audiences. One may interpret this expectation as favorable for data-driven and theoretical works, as the expectations are not for a method but the span of impact. Nevertheless, state-of-the-art for data-driven methods is still far from being sufficiently general. Olsen & Tomlin (2020) review Industry 4:0 idea, which combines the physical with digital worlds. Hopefully, the data obtained from the physical world will be exploited by data-driven methods and theoretical models with more features. Mišić & Perakis (2020) presents a summary of applications of data-driven systems and pinpoints several future directions. Most of those directions are covered in this paper, except for the idea of interpretability. Interpretability is the possibility of a human quickly seeing and understanding how the data-driven model maps an observation to a prediction. Hence, an interpretable model is a model where the user can explicitly learn its operation logic. Interpretability is desirable as it becomes part of the knowledge and hence is portable to other problems. Finally, Song, van Houtum, & Van Mieghem (2020) clearly states that human decision-making should still be in the center and can be augmented by data-driven decision models in various settings.

Analytics Age and Businesses

Mortenson, Doherty & Robinson (2015) presents a historical sketch of the developments in the OR field, bringing to the current analytics age. A research agenda is then outlined to emphasize that the OR community should internalize these developments. Furthermore, OR community is expected to develop research directions that will consider areas like significant data volumes, new data architectures, incorporating unstructured data in decision making, visualizing data for decision making and using real-time analytics. Hindle, Kunc, Mortensen, Oztekin, & Vidgen (2020) in their introduction for a special issue emphasizes the critical role of business analytics for the OR community. They present a review of the intensity of these topics within OR-related fields and other fields. These two reviews strongly advocate that researchers "jump into the wagon" of data-driven models. In support of the ideas discussed in the above studies, the following two papers give examples using cases. Zhan & Tan (2020) use a case study to show the role of building an infrastructure to harvest big data to enhance performance. Kraus, Feuerriegel, & Oztekin (2020) approach to a case using deep neural networks to solve operational problems and show that they can enhance the performance. They conclude that data-driven models still need much customization, implying that a full generic model is still far away.

Several researchers emphasize the role of data-driven approaches within business models. The advance of companies that center on data as the focal point of their business models brings new research problems to the field. Sorescu (2017) present some of the aspects of data-driven business models and state that far more effective research is needed to relate data-driven issues with the structural aspects of businesses.

Effects on Decision-makers

Most of the paper is devoted to overviewing approaches related to decision-making. However, one would expect that the type of support would influence the decision-makers. Brynjolfsson & McElheran (2016) devise a survey responded to by about 40 000 US manufacturers. One of the main findings is the reluctance of a more significant proportion of companies not fully utilize the data-driven approaches. Haller & Satell (2020), on the other hand, recommends decision-makers to be cautious concerning data. A decision-maker must ask four questions on various issues before employing such techniques: the source of the data, the approach used in analyzing the data, the aspects not told by the data, and finally, is it possible to use data to redesign products and business models. Of course, this might be why most decision-makers still believe in the concept of intuition rather than passive obedience to data.

6. What is the takeaway? Proposing a research agenda

It is essential to understand what the new possibilities bring for research. On one side, the more classical EtO is still valid, as researchers thrive on approaching complex problems with some rigor, where the structure of the problem is the focus. On the other hand, the approaches we see in Section 3 are robust when the problem is more complex than the ones which could be structurally analyzed. Most methods combine these two steps and handle them simultaneously, using the EO approach. Of course, depending on the method, what they optimize is not usually what is optimized for the corresponding problem in the EtO approach. Hence, it leaves a vital area for research. Both approaches are valid and can be utilized in practice. Nevertheless, analysts need to approach each style cautiously.

To wrap up the overview, we address the following research agenda:

1. Research to compare data-driven methods is needed. We can give three such comparisons. Lim, Shanthikumar, & Shen (2006) uses structural properties for comparison. Ban & Rudin (2019), on the other hand, uses both structural properties and numerical data sets for comparison. Finally, Meller, Taigel, & Pibernik (2018) compares the performance of methods under a controlled simulation environment. We believe all these approaches are valuable and can be used simultaneously.
2. Another dimension is comparing data-driven approaches with conventional ones. Feldman, Zhang, Liu, & Zhang (2022) is an exciting study comparing the models described in Section 4 with those in Section 3. The comparison is complicated and requires severe preparation if it is performed with the actual data.
3. Classical theory is limited to less complex inventory problems. Hence, new approaches should be implementable for more complex situations. Several such studies have been published recently. For example, for non-stationary newsvendor situations, see Meller et al. (2018), Keskin, Min, & Song (2021b) and Huber et al. (2019). Perishable inventory problems constitute another such challenge for the classical approach. Three recent studies, Li, Tang, Zhou, & Fan (2021), Keskin, Li, & Song (2021a) and De Moor et al. (2022) show that it is possible to use the new data-driven paradigm for reasonable solutions. Gijsbrechts et al. (2022) considers structurally difficult problems and proposes deep reinforcement learning to improve performance. Of course, other challenging problems are still open for investigation. Another fruitful line of research is exemplified by Boute, Gijsbrechts, van Jaarsveld, & Vanvuchelen (2021). Their article presents a road map for implementing deep reinforcement learning for inventory control problems.
4. Feature data is becoming an integral part of most OM models. Models that specify the type of feature data to be used (rather than keep it general) together with a data-driven approach are gaining momentum in the literature. Choi (2018), Hauser, Flath, & Thiesse (2021), Huang & Van Mieghem (2014) and Weißhuhn & Hoberg (2021) are some of the recent examples. Research at this detail level looks very promising. Nevertheless, one needs to make sure t hat relevant feature data is available. One research direction would be to analyze big data and recommend additional feature data to be collected for the specific type of inventory problem (see Ikegwu, Nweke, Anikwe, Alo, & Okonkwo, 2022 for a more general perspective; Boone, Ganeshan, Jain, & Sanders, 2019, and See-To & Ngai, 2018 for forecasting). Finally, for the approaches described in Section 4, the analytical models require behavioral data to be built and validated. Choi (2018), Gallino et al. (2022), See-To & Ngai (2018) and Weißhuhn & Hoberg (2021) may be consider as studies in that direction. Nevertheless, this line of research seems to be very promising.
5. Another challenging issue is the existence of censored data. If decisions affect the consequent uncertainty observed, we have censored data. Researchers have some experience in lost-sales situations when it becomes essential to differentiate sales and demand. However, more general approaches will be instrumental with the availability of feature data and other complications. One such study is proposed by Lee, Homem-de Mello, & Kleywegt (2012).
6. The overview excluded the empirical methods used for estimation and prediction and the application of those methods to inventory-related problems with specific research questions. Some recent related references were supplied in Section 4. More research to develop empirical methods in Management Science, as well as research more specific to the needs of inventory problems, is needed for the data-driven research to expand.
7. We understand that what we call “uncertainty” can be further analyzed using the classical approach. Such an example is introduced in Section 4.2, bringing together the well-known EOQ problem environment with the notion of heterogeneous customers. We are encouraged to use such models as we observe the development of the behavioral operations management field (please see Donohue, Özer, & Zheng, 2020 for a recent review on the topic). Further, as an example, several studies complement the model stated in Section 4.2. As a departure from the classical EOQ model, the model presented in Section 4.2 considers backorder time as a factor for demand. Allon, Federgruen, & Pierson (2011), and Chen, Kumar, Singhal, & Singhal (2021) consider waiting-time costs (both structurally and empirically), indicating that it is possible to have such a representation as given in the model. Backorder cost is also another issue Schwartz (1970) uses empirical data to verify a model, which then can be used to estimate backlogging cost (please see Argon, Güllü, & Erkip, 2001, and Liberopoulos, Tsikis, & Delikouras, 2010 for details). Further research in this direction will help enrich the classical theory and increase the chances of real-life implementations.
8. The models mentioned in Sections 2 and 4 use the EtO approach. In practice, plans are implemented on a rolling basis; hence, we expect this sequence to be implemented over time, though the cycle of renewing estimate may not be the same as optimized. The studies detailed in Section 3, on the other

hand, use the EO approach, simultaneously dealing with the two stages of problems. The studies on predict-and-optimize in the literature (see, Elmachtoub & Grigas, 2022 as an example) use a similar approach to EO. Note that this similarity in approaches is a good indication of a possible cross-fertilization of available knowledge, presenting the researchers in different fields to contribute to the methodologies which can be applied and extended for inventory problems.

9. The final item we mention for the research agenda follows naturally: Why not utilize different approaches to model a decision situation and select the model that gives a more accurate representation using the available data. For further analysis, this model (approach) can then be implemented for decision-making. This idea looked futuristic when mentioned by Geoffrion (2008) in his foreword for a book advocating the use of intuition for decision making. Some authors utilize the idea; see den Boer & Sierag (2021) as an example. Note that this opens a new research direction requiring convincing comparison methodologies besides other technicalities.

Epilogue

Discussion on utilizing data-driven systems can be found in other fields; see Zappone, Di Renzo, & Debbah (2019) for the control of wireless networks and Saltelli (2019) for mathematical model validation in general. It is interesting to note that although Saltelli (2019) reports experiences from another field, the resemblance of the approach selects a model that gives a more accurate representation using the available data.

One can find many studies that pinpoint dangers in employing data-driven systems in general. Matzner (2019), and Rainie, Anderson, & Page (2017) present two examples, the former from a human standpoint and the latter from a governance standpoint. It is interesting to note that some researchers are cautiously advocating the implementation of data-driven algorithms for social good (see Shi, Wang, & Fang (2020) as an example). We will see much discussion on the pros and cons of data-driven systems from all perspectives in the future.

We want to conclude the discussion for our field. We believe that data-driven systems have definite merit, and we should prepare ourselves and our community to become knowledgeable (if not experts) on those approaches. Boutilier & Chan (2021) report courses they designed to include the teaching of these techniques together with optimization ideas. However, we will advocate teaching theoretical models more than ever, as it is hard to keep track of the changes with data-driven systems without keeping the intuition alive. We refer to Chhajed & Lowe (2008), an edited book with chapters all related to the intuition we obtain using theoretical models for different areas of OM and Geoffrion (1976), an article, under-cited, in my opinion, which talks about how one can obtain intuition from a mathematical model. The exposition in Section 4.2 is an example of gaining insight into inventory theory.

Acknowledgment

I thank the editor and the anonymous referees for their constructive comments that have helped improve the paper. I presented a keynote on the topic at the 14th ISIR Summer School in Leuven and developed the work further. I want to thank the organizers, Prof. Robert Boute and Prof. Maxi Udenio, for asking me to have the keynote and all the summer school participants for their inspiring comments and questions. Finally, I would like to thank Emin Ozyoruk for his help in converting the manuscript to LATEX format throughout the preparation process.

References

Akşin, Z., Ata, B., Emadi, S. M., & Su, C. L. (2013). Structural estimation of callers delay sensitivity in call centers. *Management Science, 59*(12), 2727–2746.
Allon, G., Federgruen, A., & Pierson, M. (2011). How much is a reduction of your customers' wait worth? An empirical study of the fast-food drive-thru industry based on structural estimation methods. *Manufacturing & Service Operations Management, 13*(4), 489–507.
Argon, N. T., Güllü, R., & Erkip, N. (2001). Analysis of an inventory system under backorder correlated deterministic demand and geometric supply process. *International Journal of Production Economics, 71*(1–3), 247–254.
Arrow, K. J. (2002). The genesis of "optimal inventory policy". *Operations Research, 50*(1), 1–2.
(1960). *Mathematical methods in the social sciences, proceedings of stanford symposium on mathematical methods in the social sciences*. In K. J. Arrow, S. Karlin, & P. Suppes (Eds.). Stanford University Press.
Arunachalam, D., Kumar, N., & Kawalek, J. P. (2018). Understanding big data analytics capabilities in supply chain management: Unravelling the issues, challenges and implications for practice. *Transportation Research Part E: Logistics and Transportation Review, 114*, 416–436.
Azoury, K. S. (1985). Bayes solution to dynamic inventory models under unknown demand distribution. *Management Science, 31*(9), 1150–1160.
Baker, R. A., & Urban, T. L. (1988). A deterministic inventory system with an inventory-level-dependent demand rate. *Journal of the Operational Research Society, 39*(9), 823–831.
Balcik, B., Bozkir, C. D. C., & Kundakcioglu, O. E. (2016). A literature review on inventory management in humanitarian supply chains. *Surveys in Operations Research and Management Science, 21*(2), 101–116.
Ban, G. Y. (2020). Confidence intervals for data-driven inventory policies with demand censoring. *Operations Research, 68*(2), 309–326.
Ban, G. Y., & Rudin, C. (2019). The big data newsvendor: Practical insights from machine learning. *Operations Research, 67*(1), 90–108.
Basten, R. J. I., & van Houtum, G. J. (2014). System-oriented inventory models for spare parts. *Surveys in Operations Research and Management Science, 19*(1), 34–55.
Bertsimas, D., & Kallus, N. (2020). From predictive to prescriptive analytics. *Management Science, 66*(3), 1025–1044.
Bertsimas, D., & Thiele, A. (2005). A data-driven approach to newsvendor problems. *Working Paper*. Massachusetts Institute of Technology.
Bertsimas, D., & Thiele, A. (2006a). A robust optimization approach to inventory theory. *Operations Research, 54*(1), 150–168.
Bertsimas, D., & Thiele, A. (2006b). Robust and data-driven optimization: Modern decision making under uncertainty. In *Models, methods, and applications for innovative decision making* (pp. 95–122). INFORMS.
Besbes, O., Phillips, R., & Zeevi, A. (2010). Testing the validity of a demand model: An operations perspective. *Manufacturing & Service Operations Management, 12*(1), 162–183.
Beutel, A. L., & Minner, S. (2012). Safety stock planning under causal demand forecasting. *International Journal of Production Economics, 140*(2), 637–645.
Bijvank, M., & Vis, I. F. (2011). Lost-sales inventory theory: A review. *European Journal of Operational Research, 215*(1), 1–13.
Boone, T., Ganeshan, R., Jain, A., & Sanders, N. R. (2019). Forecasting sales in the supply chain: Consumer analytics in the big data era. *International Journal of Forecasting, 35*(1), 170–180.
Boute, R. N., Gijsbrechts, J., van Jaarsveld, W., & Vanvuchelen, N. (2021). Deep reinforcement learning for inventory control: A roadmap. *European Journal of Operational Research, 298*(2), 401–412.
Boutilier, J. J., & Chan, T. C. (2021). Introducing and integrating machine learning in an operations research curriculum: An application-driven course. In *INFORMS transactions on education*.
Brynjolfsson, E., Hui, X., & Liu, M. (2019). Does machine translation affect international trade? Evidence from a large digital platform. *Management Science, 65*(12), 5449–5460.
Brynjolfsson, E., & Mcafee, A. (2017). Artificial intelligence, for real. *Harvard Business Review, 1*, 1–31.
Brynjolfsson, E., & McElheran, K. (2016). The rapid adoption of data-driven decision–making. *American Economic Review, 106*(5), 133–139.
Cachon, G. P., Girotra, K., & Netessine, S. (2020). Interesting, important, and impactful operations management. *Manufacturing & Service Operations Management, 22*(1), 214–222.
Chen, B., & Chao, X. (2019). Parametric demand learning with limited price explorations in a backlog stochastic inventory system. *IISE Transactions, 51*(6), 605–613.
Chen, B., Chao, X., & Ahn, H. S. (2019). Coordinating pricing and inventory replenishment with nonparametric demand learning. *Operations Research, 67*(4), 1035–1052.
Chen, R. R., Kumar, S., Singhal, J., & Singhal, K. (2021). Note: The value and cost of the customer's waiting time. *Manufacturing & Service Operations Management, 23*(6), 1539–1542.
Chhajed, D., & Lowe, T. J. (2008). *Building intuition: Insights from basic operations management models and principles*: Vol. 115. Springer Science & Business Media.
Choi, T. M. (2018). Incorporating social media observations and bounded rationality into fashion quick response supply chains in the big data era. *Transportation Research Part E: Logistics and Transportation Review, 114*, 386–397.

Choi, T. M., Wallace, S. W., & Wang, Y. (2018). Big data analytics in operations management. *Production and Operations Management, 27*(10), 1868–1883.
Chu, L. Y., Shanthikumar, J. G., & Shen, Z. J. M. (2008). Solving operational statistics via a Bayesian analysis. *Operations Research Letters, 36*(1), 110–116.
Cochran, J. J., Cox, L. A., Jr., Keskinocak, P., Kharoufeh, J. P., & Smith, J. C. (2010). *The wiley encyclopedia of operations research and management science*. Wiley.
Dana, J. D., Jr., & Petruzzi, N. C. (2001). Note: The newsvendor model with endogenous demand. *Management Science, 47*(11), 1488–1497.
de Kok, T., Grob, C., Laumanns, M., Minner, S., Rambau, J., & Schade, K. (2018). A typology and literature review on stochastic multi-echelon inventory models. *European Journal of Operational Research, 269*(3), 955–983.
De Moor, B. J., Gijsbrechts, J., & Boute, R. N. (2022). Reward shaping to improve the performance of deep reinforcement learning in perishable inventory management. *European Journal of Operational Research, 301*(2), 535–545.
den Boer, A. V., & Sierag, D. D. (2021). Decision-based model selection. *European Journal of Operational Research, 290*(2), 671–686.
Ding, X., Puterman, M. L., & Bisi, A. (2002). The censored newsvendor and the optimal acquisition of information. *Operations Research, 50*(3), 517–527.
Donohue, K., Özer, O., & Zheng, Y. (2020). Behavioral operations: Past, present, and future. *Manufacturing & Service Operations Management, 22*(1), 191–202.
Elmachtoub, A. N., & Grigas, P. (2022). Smart “Predict, then optimize”. *Management Science, 68*(1), 9–26.
Engebrethsen, E., & Dauzère-Pérès, S. (2019). Transportation mode selection in inventory models: A literature review. *European Journal of Operational Research, 279*(1), 1–25.
Farahat, A., & Lee, J. (2018). The multiproduct newsvendor problem with customer choice. *Operations Research, 66*(1), 123–136.
Feldman, J., Zhang, D. J., Liu, X., & Zhang, N. (2022). Customer choice models vs. machine learning: Finding optimal product displays on Alibaba. *Operations Research, 70*(1), 309–328.
Feng, G., Li, X., & Wang, Z. (2017). On the relation between several discrete choice models. *Operations Research, 65*(6), 1516–1525.
Gallego, G., & Topaloglu, H. (2019). Introduction to choice modeling. In *Revenue management and pricing analytics* (pp. 109–128). New York, NY: Springer.
Gallino, S., Karacaoglu, N., & Moreno, A. (2022). Need for speed: The impact of in--process delays on customer behavior in online retail. *Operations Research*.
Gavirneni, S., Kapuscinski, R., & Tayur, S. (1999). Value of information in capacitated supply chains. *Management Science, 45*(1), 16–24.
Geoffrion, A. (2008). *Foreword for building intuition: Insights from basic operations management models and principles*: 115. Springer Science and Business Media.
Geoffrion, A. M. (1976). The purpose of mathematical programming is insight, not numbers. *Interfaces, 7*(1), 81–92.
Gijsbrechts, J., Boute, R. N., Van Mieghem, J. A., & Zhang, D. J. (2022). Can deep reinforcement learning improve inventory management? Performance on lost sales, dual-sourcing, and multi-echelon problems. *Manufacturing & Service Operations Management*.
Girlich, H. J., & Chikan, A. (2001). The origins of dynamic inventory modelling under uncertainty: (The men, their work and connection with the stanford studies). *International Journal of Production Economics, 71*(1–3), 351–363.
Goltsos, T. E., Syntetos, A. A., Glock, C. H., & Ioannou, G. (2022). Inventory–forecasting: Mind the gap. *European Journal of Operational Research, 299*(2), 397–419.
Graves, S. C., Kan, A. H. G. R., & Zipkin, P. H. (1993). Logistics of production and inventory. In *Handbooks in operations research and management science: vol. 4*. Amsterdam: NorthHolland.
Graves, S. C., & de Kok, A. G. (2003). Supply chain management: Design, coordination and operation. In *Handbooks in operations research and management, volume 4*. Elsevier.
Hadley, G. F., & Whitin, T. M. (1963). *Analysis of inventory systems*. Englewood Cliffs, NJ.
Haller, E., & Satell, G. (2020). Data-driven decisions start with these 4 questions. *Harvard Business Review*. https://hbr.org/2020/02/data-driven-decisions-start-with-these-4-questions
Hariharan, R., & Zipkin, P. (1995). Customer-order information, leadtimes, and inventories. *Management Science, 41*(10), 1599–1607.
Hathaway, B. A., Emadi, S. M., & Deshpande, V. (2022). Personalized priority policies in call centers using past customer interaction information. *Management Science, 68*(4), 2806–2823.
Hauser, M., Flath, C. M., & Thiesse, F. (2021). Catch me if you scan: Data-driven prescriptive modeling for smart store environments. *European Journal of Operational Research, 294*(3), 860–873.
Hazen, B. T., Skipper, J. B., Boone, C. A., & Hill, R. R. (2018). Back in business: Operations research in support of big data analytics for operations and supply chain management. *Annals of Operations Research, 270*(1), 201–211.
Hindle, G., Kunc, M., Mortensen, M., Oztekin, A., & Vidgen, R. (2020). Business analytics: Defining the field and identifying a research agenda. *European Journal of Operational Research, 281*(3), 483–490.
Hotelling, H. (1929). Stability in competition. *Economic Journal, 39*(153), 41–57.
Hu, K., Allon, G., & Bassamboo, A. (2022). Understanding customer retrials in call centers: Preferences for service quality and service speed. *Manufacturing & Service Operations Management, 24*(2), 1002–1020.
Huang, J., Leng, M., & Parlar, M. (2013). Demand functions in decision modeling: A comprehensive survey and research directions. *Decision Sciences, 44*(3), 557–609.
Huang, T., & Van Mieghem, J. A. (2014). Clickstream data and inventory management: Model and empirical analysis. *Production and Operations Management, 23*(3), 333–347.
Huber, J., Müller, S., Fleischmann, M., & Stuckenschmidt, H. (2019). A data-driven newsvendor problem: From data to decision. *European Journal of Operational Research, 278*(3), 904–915.
Huh, W. T., Levi, R., Rusmevichientong, P., & Orlin, J. B. (2011). Adaptive data-driven inventory control with censored demand based on Kaplan-Meier estimator. *Operations Research, 59*(4), 929–941.
Huh, W. T., & Rusmevichientong, P. (2009). A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research, 34*(1), 103–123.
Ikegwu, A. C., Nweke, H. F., Anikwe, C. V., Alo, U. R., & Okonkwo, O. R. (2022). Big data analytics for data-driven industry: A review of data sources, tools, challenges, solutions, and research directions. *Cluster Computing*, 1–45.
Iyer, A. V., & Schrage, L. E. (1992). Analysis of the deterministic (s, s) inventory problem. *Management Science, 38*(9), 1299–1313.
Jackson, P. L., Muckstadt, J. A., & Li, Y. (2019). Multiperiod stock allocation via robust optimization. *Management Science, 65*(2), 794–818.
Jagabathula, S., Rusmevichientong, P., Venkataraman, A., & Zhao, X. (2022). Estimating large-scale tree logit models. *Operations Research, 70*(2), 641–655.
Karlin, S., & Carr, C. R. (1962). Prices and optimal inventory policy. In K. J. Arrow, S. Karlin, & H. Scarf (Eds.), *Studies in applied probability and management science* (pp. 159–172). Stanford University Press.
Karlin, S., Scarf, H., & Arrow, K. J. (1958). *Studies in the mathematical theory of inventory and production*. Stanford University Press.
Karlin, S., Scarf, H., & Arrow, K. J. (1962). *Studies in applied probability and management science*. Stanford University Press.
Keskin, N. B., Li, Y., & Song, J. S. J. (2021a). Data-driven dynamic pricing and ordering with perishable inventory in a changing environment. *Management Science*. Forthcoming
Keskin, N. B., Min, X., & Song, J. S. J. (2021b). The nonstationary newsvendor: Data-driven nonparametric learning. Available at SSRN 3866171.
Kraus, M., Feuerriegel, S., & Oztekin, A. (2020). Deep learning in business analytics and operations research: Models, applications and managerial implications. *European Journal of Operational Research, 281*(3), 628–641.
Lee, S., Homem-de Mello, T., & Kleywegt, A. J. (2012). Newsvendor-type models with decision-dependent uncertainty. *Mathematical Methods of Operations Research, 76*(2), 189–221.
Levi, R., Perakis, G., & Uichanco, J. (2015). The data-driven newsvendor problem: New bounds and insights. *Operations Research, 63*(6), 1294–1306.
Levi, R., Roundy, R. O., & Shmoys, D. B. (2007). Provably near-optimal sampling-based policies for stochastic inventory control models. *Mathematics of Operations Research, 32*(4), 821–839.
Li, L., Tang, O., Zhou, W., & Fan, T. (2021). Backroom effect on perishable inventory management with IoT information. *International Journal of Production Research*, 1–23.
Liberopoulos, G., Tsikis, I., & Delikouras, S. (2010). Backorder penalty cost coefficient “b”: What could it be? *International Journal of Production Economics, 123*(1), 166–178.
Lim, A. E. B., Shanthikumar, G., & Shen, Z. J. M. (2006). Model uncertainty, robust optimization, and learning. In *Models, methods, and applications for innovative decision making, INFORMS tutorials in operations research* (pp. 66–94).
Liu, C., Letchford, A. N., & Svetunkov, I. (2022). Newsvendor problems: An integrated method for estimation and optimisation. *European Journal of Operational Research, 300*(2), 590–601.
Liyanage, L. H., & Shanthikumar, J. G. (2005). A practical inventory control policy using operational statistics. *Operations Research Letters, 33*(4), 341–348.
Martínez-de Albéniz, V., & Kunnumkal, S. (2021). A model for integrated inventory and assortment planning. *Management Science*.
Matzner, T. (2019). The human is dead–long live the algorithm! human-algorithmic ensembles and liberal subjectivity. *Theory, Culture & Society, 36*(2), 123–144.
Meller, J., Taigel, F., & Pibernik, R. (2018). Prescriptive analytics for inventory management: A comparison of new approaches. Available at SSRN: 3229105.
Mišić, V. V., & Perakis, G. (2020). Data analytics in operations management: A review. *Manufacturing & Service Operations Management, 22*(1), 158–169.
Mortenson, M. J., Doherty, N. F., & Robinson, S. (2015). Operational research from Taylorism to Terabytes: A research agenda for the analytics age. *European Journal of Operational Research, 241*(3), 583–595.
Musalem, A., Olivares, M., & Schilkrut, A. (2021). Retail in high definition: Monitoring customer assistance through video analytics. *Manufacturing & Service Operations Management, 23*(5), 1025–1042.
Neghab, D. P., Khayyati, S., & Karaesmen, F. (2022). An integrated data-driven method using deep learning for a newsvendor problem with unobservable features. *European Journal of Operational Research*.
Nguyen, T., Li, Z. H. O. U., Spiegler, V., Ieromonachou, P., & Lin, Y. (2018). Big data analytics in supply chain management: A state-of-the-art literature review. *Computers & Operations Research, 98*, 254–264.
Olsen, T. L., & Tomlin, B. (2020). Industry 4.0: Opportunities and challenges for operations management. *Manufacturing & Service Operations Management, 22*(1), 113–122.
Oroojlooyjadid, A., Snyder, L. V., & Takáč, M. (2020). Applying deep learning to the newsvendor problem. *IISE Transactions, 52*(4), 444–463.
Rainie, L., Anderson, J., & Page, D. (2017). Code-dependent: Pros and cons of the algorithm age. *Pew Research Center, 8*.
Ramamurthy, V., Shanthikumar, J. G., & Shen, Z. J. M. (2012). Inventory policy with parametric demand: Operational statistics, linear correction, and regression. *Production and Operations Management, 21*(2), 291–308.

Saghafian, S., & Tomlin, B. (2016). The newsvendor under demand ambiguity: Combining data with moment and tail information. *Operations Research, 64*(1), 167–185.
Saltelli, A. (2019). A short comment on statistical versus mathematical modelling. *Nature Communications, 10*(1), 1–3.
Savaşaneril, S., & Erkip, N. (2010). An analysis of manufacturer benefits under vendor-managed systems. *IIE Transactions, 42*(7), 455–477.
Scarf, H. (1958). *A min-max solution of an inventory problem in studies in the mathematical theory of inventory and production*. Stanford University Press.
Scarf, H., Daniel, K. H., Gilford, D. M., & Shelly, M. W. (1963). *Multistage inventory models and techniques*. Stanford University Press.
Schwartz, B. L. (1970). Optimal inventory policies in perturbed demand models. *Management Science, 16*(8), B–509.
See-To, E. W., & Ngai, E. W. (2018). Customer reviews for demand distribution and sales nowcasting: A big data approach. *Annals of Operations Research, 270*(1), 415–431.
Shapiro, A., & Xin, L. (2020). Time inconsistency of optimal policies of distributionally robust inventory models. *Operations Research, 68*(5), 1576–1584.
Shi, Z. R., Wang, C., & Fang, F. (2020). Artificial intelligence for social good: A survey. arXiv preprint arXiv:2001.01818.
Siegel, A. F., & Wagner, M. R. (2021). Profit estimation error in the newsvendor model under a parametric demand distribution. *Management Science, 67*(8), 4863–4879.
Silver, E. A. (1981). Operations research in inventory management: A review and critique. *Operations Research, 29*(4), 628–645.
Song, J. S., van Houtum, G. J., & Van Mieghem, J. A. (2020). Capacity and inventory management: Review, trends, and projections. *Manufacturing & Service Operations Management, 22*(1), 36–46.
Sorescu, A. (2017). Data-driven business model innovation. *Journal of Product Innovation Management, 34*(5), 691–696.
Syntetos, A. A., Babai, Z., Boylan, J. E., Kolassa, S., & Nikolopoulos, K. (2016). Supply chain forecasting: Theory, practice, their gap and the future. *European Journal of Operational Research, 252*(1), 1–26.
Tan, T., Güllü, R., & Erkip, N. (2007). Modelling imperfect advance demand information and analysis of optimal inventory policies. *European Journal of Operational Research, 177*(2), 897–923.
Transchel, S., Buisman, M. E., & Haijema, R. (2022). Joint assortment and inventory optimization for vertically differentiated products under consumer-driven substitution. *European Journal of Operational Research, 301*(1), 163–179.
Turgut, O., Taube, F., & Minner, S. (2018). Data-driven retail inventory management with backroom effect. *OR Spectrum, 40*(4), 945–968.
Urban, T. L. (2005). Inventory models with inventory-level-dependent demand: A comprehensive review and unifying theory. *European Journal of Operational Research, 162*(3), 792–804.
van der Laan, N., Teunter, R. H., & Romeijnders, W. (2019). The data-driven newsvendor problem: Achieving on-target service levels. *Working Paper*.
Wagner, H. M. (1980). Research portfolio for inventory management and production planning systems. *Operations Research, 28*(3-part-i), 445–475.
Wang, G., Gunasekaran, A., Ngai, E. W., & Papadopoulos, T. (2016). Big data analytics in logistics and supply chain management: Certain investigations for research and applications. *International Journal of Production Economics, 176*, 98–110.
Wang, R. (2022). The threshold effects on consumer choice and pricing decisions. *Manufacturing & Service Operations Management, 24*(1), 448–466.
Weißhuhn, S., & Hoberg, K. (2021). Designing smart replenishment systems: Internet-of-things technology for vendor-managed inventory at end consumers. *European Journal of Operational Research, 295*(3), 949–964.
Whitin, T. M. (1953). *The theory of inventory management*. Princeton University Press.
Yuan, H., Luo, Q., & Shi, C. (2021). Marrying stochastic gradient descent with bandits: Learning algorithms for inventory systems with fixed costs. *Management Science*.
Zappone, A., Di Renzo, M., & Debbah, M. (2019). Wireless networks design in the era of deep learning: Model-based, AI-based, or both? *IEEE Transactions on Communications, 67*(10), 7331–7376.
Zhan, Y., & Tan, K. H. (2020). An analytic infrastructure for harvesting big data to enhance supply chain performance. *European Journal of Operational Research, 281*(3), 559–574.