A survey of recent progress in the asymptotic analysis of inventory systems

David A. Goldberg
Department of Operations Research and Information Engineering, Cornell University, dag369@cornell.edu

Martin I. Reiman
Department of Industrial Engineering and Operations Research, Columbia University, mr3598@columbia.edu

Qiong Wang
Department of Industrial and Systems Engineering, University of Illinois at Urbana-Champaign, qwang04@illinois.edu

November 5, 2019

It has long been recognized that many inventory models most relevant to practice are inherently high-dimensional, and hence generally believed to become computationally intractable as certain problem parameters grow large (suffering from the "curse of dimensionality"). In the last decade, asymptotic analysis has shown that in many interesting settings such problems can actually be well-approximated by much simpler optimization problems, leading to new algorithms and insights. In this survey, we review the state-of-the-art as regards applying asymptotic analysis to such challenging inventory problems. In addition to surveying the literature, we present a detailed introduction to the relevant tools and methodologies through three in-depth case studies in which asymptotic analysis has recently led to major progress: lost sales models, dual-sourcing models, and assemble-to-order systems in the presence of large lead times.

1. Introduction

1.1. Overview of mathematical inventory research broadly.

The problem of deciding how much inventory a retailer should stock over time is no doubt an endeavor as ancient as commerce and civilization itself. It is generally accepted that the first modern treatment of the associated decision problems using formal mathematics dates back to Harris' seminal paper on the basic EOQ model (Harris (1913)). The field really took off in the 1950's when it was recognized that multi-stage inventory control problems could generally be framed as dynamic programs (Bellman et al. (1954)), and that depending on one's modeling assumptions these dynamic programs may have relatively simple optimal solutions, or may be truly high-dimensional and seemingly intractable (Arrow et al. (1951), Dvoretsky et al. (1952, 1953), Bellman et al. (1955), Arrow et al. (1958), Karlin and Scarf (1958)), where determining which of these scenarios holds for any particular set of modeling assumptions could be quite mathematically challenging. We note that the phenomenon of certain practically relevant dynamic programming formulations scaling poorly (computationally) with the ambient dimension came to be referred to as the "curse of dimensionality" (Bellman (1961)). In the following decades, researchers continued to define new application-dependent inventory models, and attempted to determine under which assumptions these models were tractable (often with simple and insightful optimal or near-optimal policies), where demonstrating tractability often involved non-trivial mathematical tools from dynamic programming and convex analysis. When these assumptions did not hold, the common wisdom (accumulated over time as certain basic models were studied by multiple researchers) was that these models were fundamentally intractable, typically due to the aforementioned curse of dimensionality, and instead heuristic methods were studied and validated numerically.

This general approach has by now led to a vast literature on mathematical inventory problems. For example, the following is a (non-comprehensive) list of surveys done on various aspects of mathematical inventory control in the past sixty years: Veinott (1966), Clark (1972), Aggarwal (1972), Silver (1981), Nahmias (1982), Aksoy et al. (1988), Porteus (1990), Raafat (1991), Kleijn et al. (1999), Petruzzi et al. (1999), Kennedy et al. (2002), Minner (2003), Elmaghraby et al. (2003), Chan et al. (2004), Urban (2005), Williams et al. (2008), Syntetos et al. (2009), Levi (2010), Moshe et al. (2011), Bijvank and Vis (2011), Winands et al. (2011), Fiestras-Janiero et al. (2011), Shi (2011), Krishnamoorthy et al. (2011), Bakker et al. (2012), Chen and Simchi-Levi (2012), Coelho et al. (2013), Brahimi et al. (2017), Yao and Minner (2017), Atan et al. (2017), Chen (2017b), Duong et al. (2018), along with at least one survey of surveys (Bushuev et al. (2015)), and several textbooks (Zipkin (2000), Porteus (2002), Simchi-Levi et al. (2005)).

1.2. Scope of this survey (and why this survey is needed).

1.2.1. Lack of surveys which go in-depth in a particular methodology. The majority of the aforementioned surveys each focus on a specific family of inventory models, and classify the work on those models without going into much mathematical depth of the associated techniques used in their analysis. This is problematic, as a survey which goes into considerable mathematical depth on a particular new methodology can be instrumental in moving an academic field forward, by providing a reference which can quickly introduce researchers to the key ideas and techniques needed to apply those ideas to other problems. Among the aforementioned surveys, it seems that some of the very few attempting such a presentation recently are Levi (2010) and Shi (2011), for the myopic balancing approximation algorithm methodology pioneered by Levi and co-authors; and Chen (2017b), for the use of $L^\natural$ convexity and related methodologies in inventory problems.

1.2.2. Lack of survey on asymptotic analysis. A different methodology which has recently led to substantial progress on several fundamental inventory models is **asymptotic analysis**. In asymptotic analysis, one proves that as certain parameters of the system grow large, various (probabilistic) phenomena manifest which lead to certain approximations performing well.

For example, it could be that for an inventory control problem which is naturally high-dimensional, such a phenomenon effectively (in some asymptotic sense) reduces the relevant optimization to one of a much lower dimension. Ideally, the parameters which "grow large" in the analysis are naturally large in some interesting setting, either for modeling reasons or e.g. in settings where other methods (such as dynamic programming) become impractical. Perhaps the most classical and fundamental example of asymptotic analysis is the celebrated central limit theorem, in which it is proven that under mild assumptions the distribution of the convolution of many independent and identically distributed random variables (properly renormalized) can be well-approximated by a Gaussian, with explicit bounds on the error of such an approximation coming from e.g. Stein's method (Ross (2011)).

There is by now a vast literature on such analyses in many academic disciplines, including probability theory, operations research, statistics, and computer science, and we make no attempt to survey the approach at such a level of generality here. Asymptotic analysis has been especially critical in solving some of the most challenging problems in the theory of queues, a discipline very close to inventory theory both in the problems studied and techniques used, and we note that a large number of surveys and books have appeared in the literature focusing on the application of asymptotic analysis specifically to the study of queues (Whitt (1974), Lemoine (1978), Glynn (1990), Harrison and Nguyen (1993), Schweitzer et al. (1993), Down and Meyn (1994), Williams (1998), Whitt (2002), Pang et al. (2007), Shanthikumar (2007), Ward (2012)). We note that in queueing theory, asymptotic analysis is often used to prove that a seemingly high-dimensional problem behaves asymptotically like a much lower-dimensional problem, with this phenomenon generally referred to as "state-space collapse" (Reiman (1984), Harrison and Van Mieghem (1997), Bramson (1998), Stolyar (2004), Dai and Tezcan (2011)).

Interestingly, it seems that no such analogous surveys exist explaining the application and importance of asymptotic analysis in inventory theory. This is in spite of the fact that: 1. at least one of the most practically important insights of inventory theory centers around such analyses, i.e. the implications of the central limit theorem for inventory pooling (Eppen (1979)); 2. there is by now a growing literature in inventory theory based on asymptotic analysis; and 3. the last decade has seen asymptotic analysis make substantial progress on several of the most fundamental and long-standing open problems in inventory theory, including the optimization of lost sales, dual-sourcing, and assemble-to-order (ATO) systems with large lead times.

1.2.3. Scope of this survey. For all of the above reasons, we believe the time is ripe for a survey on the application of asymptotic analysis to inventory theory, at a level which goes into some depth regarding the mathematical techniques and how they apply to different types of problems, using as case studies some of the aforementioned long-standing open problems in lost sales, dual-sourcing, and ATO systems. We have attempted exactly such an exposition in the present manuscript. Our survey will make no attempt to comprehensively review the state of mathematical inventory theory broadly (beyond a terse discussion of some of the more well-studied high-dimensional inventory control problems), nor the vast literature on more applied lines of research in inventory management. Nor will we attempt to review the even more vast literature on asymptotic analysis across other disciplines, instead focusing mostly on the methods most relevant to the particular inventory problems we consider in-depth.

1.3. Outline of survey.

Our survey will proceed as follows. In Section 2, we give a brief overview of some of the more well-studied inventory models which exhibit the “curse of dimensionality”. In Section 3, we first give a brief overview of the different methods researchers have used for attacking such problems. We then zoom in on the approach of primary interest in this survey, asymptotic analysis, and provide a broad overview of the methodology and its application to inventory theory. In the following four sections, we present a detailed introduction to the relevant tools and methodologies through three in-depth case studies in which asymptotic analysis has recently led to major progress. In Section 4, we provide a high-level overview of the three case studies, and place them in the broader context of asymptotic analysis. In Section 5, we consider the application of asymptotic analysis to lost sales models with large lead times. In Section 6, we consider the application of asymptotic analysis to dual-sourcing models with a large gap in the lead time between the fast and slow supplier. In Section 7, we consider the application of asymptotic analysis to ATO systems with large lead times. In Section 8, we provide some final thoughts, open problems, and directions for future research.

1.4. Preliminary notation and definitions.

Here we introduce some additional notation and definitions for use throughout the survey. Let r.v. denote random variable (or vector depending on context), and i.i.d. denote “independent and identically distributed”. For a real number x, let $x^+ \triangleq \max(x,0)$, and $x^- \triangleq \max(0,-x)$. For two r.v.s X,Y, let $X \sim Y$ denote equivalence in distribution. For two integers i,j, let $\delta_{i,j}$ equal 1 if $i=j$ and 0 otherwise. For a real-valued vector $\mathbf{x}$, let $x_{[k]} \triangleq (x_1,\ldots,x_k)$. Also, let the constraint $\mathbf{x} \geq \mathbf{0}$ denote the set of constraints $x_i \geq 0$ for all i. For a stochastic process $\{Z(t), t \geq 0\}$, and $t_1 \leq t_2$, let $Z(t_1,t_2) \triangleq Z(t_2) - Z(t_1)$. Recall that for two non-negative functions f,g of a common parameter n, we say that $f = O(g)$ if $\limsup_{n\to\infty} \frac{f(n)}{g(n)} < \infty$, $f = \Theta(g)$ if $0 < \liminf_{n\to\infty} \frac{f(n)}{g(n)} \leq \limsup_{n\to\infty} \frac{f(n)}{g(n)} < \infty$, and $f = o(g)$ if $\lim_{n\to\infty} \frac{f(n)}{g(n)} = 0$.

2. Overview of inventory models suffering from the curse of dimensionality

Although real-world supply chains are extremely complex, the academic inventory control literature has identified certain fundamental mathematical models which capture some of the basic trade-offs which manifest in more complex supply chains. Perhaps the best known of these models are the celebrated EOQ and newsvendor models (Harris (1913), Qin et al. (2011)), whose solutions are well-known. However, other such well-studied models, designed to capture additional features and trade-offs present in many real-world supply chains, remain poorly understood to this day, largely due to the associated dynamic programs exhibiting the curse of dimensionality. Here we briefly survey several such families of models, noting that our list of models (and associated references) are far from comprehensive, as they are mostly included for additional context and not the main focus of our survey. Some of these models, e.g. those corresponding to our case studies, will be described in greater depth later in the survey.

1. *Lost sales with positive lead times.* In lost sales models, demand which cannot be met in a given period is lost to the system (with an appropriate lost-demand penalty incurred), and cannot be met at a later time, resulting in the inventory updating in an inherently non-linear manner. This stands in contrast to models with demand backlogging, in which demand that cannot be met remains in the system (possibly resulting in a negative inventory). In the presence of lead times, i.e. a multi-period delay between when an order for more inventory is placed and when that order is received, models with lost sales become high-dimensional, as one must keep track of all orders already placed but not yet received (i.e. the pipeline vector). It is well-known that for models with backlogging, there exist simple optimal policies which only require knowledge of the net amount of inventory in the pipeline vector (i.e. the sum of all such orders), effectively reducing the problem to a 1-dimensional optimization (Karlin and Scarf (1958)). For lost sales models, it is known that no such simplification occurs, and there is a large body of work on alternative approaches to the resulting high-dimensional problems (Bijvank and Vis (2011), Xin and Goldberg (2016)).

2. *Multiple suppliers.* In models with multiple suppliers, one must decide not just how much inventory to order, but how much inventory to order from each of multiple suppliers. Different suppliers typically have different costs, and different lead times. A well-studied special case of such a model is the celebrated dual-sourcing problem, in which one must dynamically allocate orders for more inventory between a supplier which is cheaper but slower (i.e. has a larger lead time), and one which is more costly but faster. In such models, high-dimensionality can manifest in many ways. For example, in the presence of lead times, one must again keep track of all orders placed but not yet received, from all suppliers. Although a complete solution is known when the lead time gap between

any two suppliers is at most one, beyond this special case many formulations remain fundamentally high-dimensional, although alternative formulations (e.g. using queueing networks) may lead to reduced dimensionality (at the expense of having to make different modeling assumptions) (Minner (2003), Yao and Minner (2017), Song et al. (2016), Xin and Goldberg (2017)).

3. *ATO systems.* In ATO systems, one typically has an inventory for each type of component under consideration, and knowledge of a so-called bill of materials (BOM) which dictates how many of each component must be used to assemble a given product. Demand then arrives for products over time, and the inventory manager must decide when and how much to order of each component over time, as well as when and how to meet the demand for products over time (e.g. if two different products use the same component, one may prioritize one of those products over the other if there are only enough components on-hand to meet some of the demand). There may also be e.g. different costs and lead times associated with different components and products. Although the exact optimal policy can be computed when this BOM has a limited special structure (restricting the types of interactions between the components required by different products), in general the problem is high-dimensional, and many approaches have been considered in the literature (Atan et al. (2017)).

4. *Correlated demands.* When demand is not independent over time, e.g. when the demand in the current period depends in a complex way on past demands and/or some kind of high-dimensional state or feature vector, even simple models can become inherently high-dimensional, as the associated dynamic programs must keep track of all information needed to determine the conditional distribution of demand. We note that this phenomenon arises in many realistic models for demand, especially when forecasting methods are applied to the demand, but computational simplifications can occur if additional assumptions are made (e.g. that demand is linear in the features) (Baykal-Gursoy et al. (2010), Levi et al. (2007), Chen and Song (2001), Ban and Rudin (2018)).

5. *Perishable inventory.* When inventory is perishable, i.e. when inventory expires after some fixed amount of time, simple models can again become high-dimensional as one must in principle track the remaining life (or age) of each unit of inventory, and there is a vast literature on heuristics and approximations (Nahmias (1982), Duong et al. (2018), Chao et al. (2015, 2017), Zhang et al. (2016)).

6. *Multi-echelon.* In multi-echelon systems, an inventory manager makes ordering and routing decisions for multiple nodes in a supply chain, with possibly complex interactions between them. Unless this supply chain has a special (e.g. serial) structure, in which case classical results of Clark and Scarf (and follow-up work) demonstrate tractability, the problem remains high-dimensional (with dimension growing at least as fast as the number of nodes in the network) (Clark (1972), Federgruen (1993), Simchi-Levi and Zhao (2012)).

7. *Multiple items and/or retailers.* In models with multiple items and/or retailers, one must in general keep track of the inventory levels of all items and/or retailers, leading to high-dimensionality. Special cases here include the one-warehouse-multiple-retailer problem and joint-replenishment problem (Roundy (1985), Aksoy et al. (1988), Brahimi et al. (2017)). Many challenging lot-sizing problems also fall into this category (Bushuev et al. (2015)).

3. Overview of basic approaches to overcoming the curse of dimensionality

In total generality, the models discussed in Section 2 are intractable (especially if one desires a truly optimal solution). However, in some cases, various approaches have been able to overcome this difficulty (at least approximately). Here we survey these approaches.

3.1. Brief digression on "what is intractable".

Here we take a brief digression to discuss the notion of (in)-tractability. In theoretical computer science and the formal study of algorithms, these concepts have been formalized using e.g. Turing machines and complexity theory (Papadimitriou (2003)). Although the formal notion of (in)-tractability is less well-developed for inventory models, some formal hardness results (in the sense of complexity theory and classification as e.g. NP-Complete) are indeed known (Halman et al. (2009)). We note that the notion of tractable can also be subtle for inventory models, since (for example) a policy can have a simple (e.g. base-stock) structure, yet it could be that using backwards induction to compute the required base-stock levels is (at least in some sense) formally hard, e.g. in the same sense that computing the convolution of independent random variables is (in principle) computationally hard (in the worst-case). For further discussion along these lines, we refer the interested reader to Halman et al. (2009).

3.2. Overview of approaches.

Here we briefly survey several such approaches, again making no attempt to be comprehensive. We focus almost exclusively on approaches that provably help overcome the curse of dimensionality for seemingly high-dimensional inventory problems, making no attempt to survey solution methods for inventory control problems generally, nor the vast literature on heuristic methods for such problems.

3.2.1. Approaches not based on asymptotic analysis.

1. *Convexity.* Many of the classical insights into which inventory models are tractable stems from the observation that the relevant dynamic programming equations are convex (or some generalized notion thereof) for certain models. Such analyses have shown that (time-dependent) base-stock policies (in which one orders up to pre-determined levels), and more generally (s,S) policies (in

which one orders up to S if one's inventory is below s), are optimal (and can be computed using backwards induction) for certain models, including models with backlogged demand and positive lead times (Karlin and Scarf (1958), Veinott (1966)). Related concepts of integer (including $L^\natural$) convexity have also been used to prove interesting structural and algorithmic results for several models (Zipkin (2008), Chen (2017b)). Convexity (and related concepts such as submodularity) has also been a key tool used in many proofs appearing in the inventory literature.

2. *Dynamic programming techniques.* As inventory control problems can be formulated as dynamic programs, it is not surprising that techniques from dynamic programming have been able to shed additional insights into these problems. The canonical result here is that of Clark and Scarf, which showed that for serial multi-echelon systems the dynamic program has a certain "decomposition" structure which leads to tractability (Clark (1972)). Of course, as with convexity, in general dynamic programming has also been a key tool used in many proofs appearing in the inventory literature.

3. *Approximation algorithms.* For many problems, although it remains unknown how to efficiently find policies which are arbitrarily close to optimal, it is known how to efficiently compute policies which are guaranteed to be "close" to optimal. A typical guarantee for such an algorithm is that it incurs expected cost at most c times that incurred by the optimal policy, where c is (ideally) some constant independent of the problem parameters - such an algorithm is called a "c-approximation algorithm". Canonical results here include Roundy's "power-of-2" algorithm yielding a 1.02-approximation algorithm for the one-warehouse-multi-retailer problem (Roundy (1985)) and Levi's balancing-algorithms yielding constant-factor approximation algorithms for a wide range of inventory problems with correlated demand (Levi (2010)), as well as the many generalizations and extensions of these approaches (Chao et al. (2017)). A different literature for inventory control problems which derive approximation algorithms shows that by carefully analyzing the relevant dynamic programs and rounding/truncating/approximating relevant quantities, one can (with certain parameters held fixed) also derive efficient approximation methods (Halman et al. (2009)).

4. *Alternate modeling paradigms.* Some lines of work have attempted to overcome the intractability of certain dynamic programming formulations by modeling the basic problem differently, in some cases side-stepping the associated hardness by making different underlying assumptions. Prime examples here include framing models in the robust optimization paradigm (Bertsimas and Thiele (2006), Bertsimas et al. (2010), Gabrel et al. (2014)); structuring models to be more in line with results from the theory of queues and stochastic processing networks (Van Mieghem and Rudi (2002), Song et al. (2016)); or simply formulating as tractable linear programs (Dzielinski et al. (1965)).

3.2.2. Approaches based on asymptotic analysis. In the early inventory theory literature, asymptotic analysis was used almost exclusively through the application of the central limit theorem (CLT). The CLT was used to justify the normal approximation for the total demand over a lead time (Burgin (1975)), and to provide a simple to evaluate expression for the gain from inventory pooling (e.g. Eppen (1979)). However, such usages were relatively rare, let alone systematic. Following its prolific development in queueing systems and stochastic processing networks, asymptotic analysis has been applied to the study of production-inventory systems, which are equivalent or at least bear some close analogy to queueing systems (Bradley and Glynn (2002), Glasserman (1997), Glasserman and Liu (1997), Krichagina et al. (1993), Plambeck and Ward (2006), Wein (1992)), and are often analyzed under the "heavy traffic" condition (Krichagina et al. (1993), Plambeck and Ward (2006), Wein (1992)). More distinct from queuing systems are inventory models in which components are ordered instead of produced in-house. For these types of systems, asymptotic analysis took place in a more recent time and is becoming increasingly popular. The three case studies in sections 5, 6, 7 all belong to that category.

Probably the most important use of asymptotic analysis on inventory systems is to overcome the curse of dimensionality. Our survey will focus almost exclusively on this application. Before getting to that discussion, it is also worthwhile to briefly summarize some examples of other uses of the approach:

- **Inspiration for developing new types of policies:** Reiman (2004) defines a constant order policy for the lost sales model, and shows that it sometimes outperforms the best base-stock policy in systems with large lead times and large lost sales costs. Allon and van Mieghem (2010) develop a "tailored base-surge (TBS)" policy for a dual sourcing situation, and carry out related asymptotic analysis in a heavy traffic regime. Both studies are followed by many other substantial asymptotic analysis, which we will discuss in details in Sections 5 and 6 respectively.

Stolyar and Wang (2018) study a classical single-item backlog system with i.i.d. random lead times. Orders may cross in time, so it is not optimal to follow the commonly-used base stock policy, under which the average inventory cost scales with the demand rate r as $\Theta(\sqrt{r})$ $(r \to \infty)$. Based on asymptotic analysis, they develop a "Generalized Base Stock (GBS)" replenishment policy that can drive the average cost down to $o(\sqrt{r})$. The result is proven for systems with exponential lead times and tested by simulations for systems with other lead time distributions.

- **Qualitative insights of policies and heuristics:** Fu et al. (2019) consider a periodic-review inventory system of a perishable product that is either manufactured anew or remanufactured from returned units. They develop inventory policies for determining the replenishment amounts of both types of units in each period. They use asymptotic analysis to prove that with the increase of the total inventory level of aged units (i.e., those that will perish in the next period), the optimal replenishment amount of the new units converges to a constant.

• **Models with other operational decisions**: Xin et al. (2019) consider a discrete-time lost sale model, the control of which is exercised by both replenishment and pricing decisions. They show that a constant order policy, coupled with dynamic pricing decisions that are based on the current inventory level, is asymptotically optimal as the lead time increases.

• **Characterizing policy parameters:** Ang et al. (2017) consider the (r, q) policy in a continuous-review inventory system with i. i. d. lead times. They prove the asymptotic independence of the inventory position and inventory on-order in high-demand rate regime and develop close-form expressions of the relationship between policy parameters and the inventory cost.

• **Demand learning**: As a related development, several researchers combine learning algorithms with an inventory policies to address situations where the demand distribution is not known a priori. The effectiveness of such an approach is normally measured by the regret, i.e., the difference between the resulting inventory cost and the optimal cost under full knowledge of the demand distribution. There has been a stream of work on the convergence rate of regret as the length of the time horizon grows (Huh and Rusmevichientong (2009), Zhang et al. (2018b), Shi et al. (2016), Chen et al. (2019), Agrawal and Jia (2019), Levi et al. (2015), Ban and Rudin (2018)).

For studies that address high dimensionality of inventory systems, there are multiple dimensions along which one could potentially classify the relevant asymptotic analyses, including: specific model studied; asymptotic regime applied; and methods of analysis utilized. Here we focus on the first two (i.e. model studied and asymptotic regime applied), as for any given model different papers may apply quite different asymptotic techniques. Such techniques often include: generalizations of the CLT, stochastic comparison arguments, relaxation methods (linear programming and otherwise), decomposition techniques, and combinations thereof (also with tools from convexity and dynamic programming). Indeed, we will see examples of all these techniques in our three case studies. We note that in general, similar techniques have appeared in the asymptotic analysis of queueing systems and stochastic processing networks, and that several researchers have worked on both the asymptotic analysis of queues and inventories.

The asymptotic results of interest here will in general be of the following form:

"For intractable high-dimensional inventory model ..., if one lets model parameter ... grow large (all else fixed), then "simple" policy ... is asymptotically optimal (in an appropriate sense) as said parameter grows to infinity".

We note that even for the same model, there are often works letting different parameters scale (often yielding different asymptotically optimal simple policies). We also note that "simple" may have different meanings in different settings, with some models and asymptotic regimes allowing for very simple and intuitive policies (sometimes agreeing with methods already adopted by practitioners), with others requiring more complex policies (albeit far simpler than those dictated by the optimality equations of the original high-dimensional dynamic program).

Tables surveying and classifying relevant asymptotic analyses. Here we survey and classify those works which have used asymptotic methods to overcome the curse of dimensionality in inventory control problems. We note only the model studied and asymptotic regime applied, and point the interested reader to the relevant literature for further details (e.g. nature of the asymptotically optimal policy, modes of analysis applied, etc.). Below, there is one table for each model studied, and the row represents the parameter which is scaled asymptotically. We have opted to focus on readability, at the expense of not providing too many details into the relevant asymptotic regime beyond the high-level quantity which is scaled up in the analysis. For example, a work may be classified as one in which demand is scaled up, or costs are scaled up. Of course, for any particular model there may be multiple ways to do such a scaling, and we refer the reader to the relevant papers for further details. The category "Other" is used as a catch-all for non-standard asymptotic regimes, in which quantities such as time, service level constraints, or other policy parameters grow large. Also, some works appear in multiple boxes, as they study multiple asymptotic regimes in the same paper. For those works relevant to our three case studies, we will go into much greater detail later in the review.

Demand	Plambeck and Ward (2006), Plambeck (2008), Wan and Wang (2015), Lu et al. (2015)
Cost	-
Lead time	Reiman and Wang (2015), Reiman et al. (2018)
Network size	-
Other	-

Table 1 Work applying asymptotic analysis to high-dimensional ATO systems.

Demand	Kaminsky et al. (2008)
Cost	Glasserman (1996)
Lead time	-
Network size	-
Other	Glasserman (1996)

Table 2 Work applying asymptotic analysis to high-dimensional make-to-stock / order systems.

Demand	-
Cost	Huh et al. (2009), Bu et al. (2017), Xin (2019)
Lead time	Goldberg et al. (2016), Xin and Goldberg (2016), Bu et al. (2017), Xin (2019)
Network size	-
Other	Wei et al. (2018), Xin et al. (2019)

Table 3 Work applying asymptotic analysis to high-dimensional lost sales systems.

Demand	Jasin et al. (2015)
Cost	-
Lead time	Xin and Goldberg (2017)
Network size	-
Other	-

Table 4 Work applying asymptotic analysis to high-dimensional multiple source / fulfillment center systems.

Demand	-
Cost	-
Lead time	-
Network size	-
Other	Federgruen et al. (1994)

Table 5 Work applying asymptotic analysis to high-dimensional joint replenishment problems.

Demand	Shi et al. (2018)
Cost	-
Lead time	-
Network size	Romeijn and Morales (2003), Ahuja et al. (2007), Simchi-Levi and Wei (2012)
Other	-

Table 6 Work applying asymptotic analysis to high-dimensional inventory network design problems.

Demand	Cudina and Ramanan (2011)
Cost	Huh et al. (2016), Rong et al. (2017)
Lead time	Chen et al. (2017), Rong et al. (2017)
Network size	Rong et al. (2017)
Other	-

Table 7 Work applying asymptotic analysis to high-dimensional multi-echelon systems.

Demand	Zhang et al. (2018), Bu et al. (2019)
Cost	Bu et al. (2019)
Lead time	-
Network size	-
Other	Bu et al. (2019)

Table 8 Work applying asymptotic analysis to high-dimensional perishable inventory systems.

4. Overview of case studies and asymptotic analysis

In the following sections, we explore three case studies in which asymptotic analysis has recently led to fundamental progress on several long-standing problems in inventory control. Here we provide an overview of these studies, as well as some additional context within the landscape of asymptotic analysis.

- **Case Study I: Lost sales models**. Here two lines of work are analyzed. Huh et al. (2009) prove that when the lost sales penalty grows large in lost sales inventory models, a base-stock policy is asymptotically optimal. Xin and Goldberg (2016) prove that when the lead time grows large in lost sales inventory models, a constant-order policy (in which the same constant is ordered in every period independent of the state of the system) is asymptotically optimal. Several key concepts from the asymptotic analysis of stochastic systems manifest here.

— *Stochastic comparison and coupling can be used to bound complex dynamics.* This concept is used to sandwich the cost incurred between that incurred in two "simpler" inventory systems, which have different dynamics but which see the same sequence of demands (i.e. are coupled).

— *Asymptotically negligible quantities can be ignored.* The concept that many "complex features" of a model contribute only a vanishingly small (as the system scales) amount to the associated costs is used to show that certain lower and upper bounds agree in the limit.

— *Complex Markov chains simplify in steady-state.* This concept is used to simplify the analysis of the Markov chain induced by an "optimal policy", which is implicitly defined and may be very complex, by exploiting certain symmetries which must hold in the long-run.

— *Convexity can be used to bound complex dynamics.* This concept is also used to simplify the analysis of the Markov chain induced by an "optimal policy", as the inventory size is a convex function of the previous quantities ordered. Jensen's inequality then bounds relevant quantities by those in which the "random" order quantities are replaced by their means.

— *Asymptotic independence simplifies analysis.* The concept that random quantities which are "far apart" in some distance (e.g. time) can be treated as if they were independent allows one to decouple the state of the system "one lead-time away" from the state of the system "now" as the lead-time grows large.

— *Vanishingly small perturbations introduce regularity.* This concept is used to reduce the associated Markov chains to chains which are more regular (e.g. having a nice regenerative structure).

- **Case Study II: dual-sourcing models with lead time gap**. Xin and Goldberg (2017) prove that when the lead time of the slow supplier grows large, a so-called tailored base-surge policy (in which a constant order is placed at the slow supplier while a base-stock policy is used at the express supplier) is asymptotically optimal in dual-sourcing models. This case study uses many of the same key concepts as Case Study I (e.g. simplicity of the steady-state, convexity, asymptotic independence, vanishingly small perturbations). It also introduces an additional key concept from the asymptotic analysis of stochastic systems.

— *Long-run-average problems can be approximated by discounted problems.* This concept is used to reduce questions about the long-run-average behavior of certain Markov chains to analogous questions when the relevant costs are discounted with a discount factor approaching one, which induces additional regularity and stationarity properties.

- **Case Study III: ATO Systems with large lead times**. Dogru et al. (2010), Reiman and Wang (2012, 2015), Reiman et al. (2018) prove that when the lead times of the components grow large, a "target-tracking" control strategy (in which one attempts to keep the state of the system at certain ideal levels computed by re-solving stochastic programs) is asymptotically optimal. This case study uses the key concepts of stochastic comparison (to derive a stochastic-programming lower bound for the cost incurred in a given period, by considering a modified system in which several key features are relaxed), the fact that asymptotically negligible quantities can be ignored, as well as the technique of introducing vanishingly small perturbations, from Case Study I. It also uses multiple additional key concepts from the asymptotic analysis of stochastic systems.

— *Functional CLT reduces complex processes to simple ones in the limit.* This concept is used to show that certain processes associated with the inventory system behave like "nice" Gaussian processes as the lead time grows large.

— *Continuity of optimization problems implies key quantities change slowly over time in the limit.* This concept is used to show that the policy one gets by repeatedly re-solving certain stochastic programs behaves continuously, which is used to show that one can keep the system "near" the desired ideal levels.

— *Maximal inequalities show processes concentrate around their means.* This concept is used to show that certain stochastic processes indeed track their ideal levels, by showing that the (normalized) deviations from these ideal levels (with high probability) never deviate far from zero (uniformly).

5. Case Study I: Lost sales models with large lead times

5.1. Introduction to lost sales models with lead times.

5.1.1. High-level introduction. In many settings of practical relevance, if demand arrives for a product and there is not enough inventory on-hand to meet that demand, then that demand is lost to the system (and cannot be met at a later date). The canonical example here is the retail industry, where a potential customer may simply go to a different supplier in case of an out-of-stock event (Bijvank and Vis (2011)). Another important feature in many inventory systems is a lead time, i.e. a delay between when an order for new inventory is placed, and the time at which the supplier receives that inventory, due e.g. to the need to transport, manufacture, etc. additional product (Minner (2003)). In inventory models with a positive lead time, the state-space of the underlying dynamic program becomes high-dimensional (with the dimension growing linearly in the lead time), since one must account for all outstanding orders which have already been placed but not yet received. We note that this makes the large lead time regime interesting both due to the fact that in many applications it is at least moderately large, and also because it is exactly the setting of large lead times in which methods such as dynamic programming become intractable.

5.1.2. Literature review prior to recent progress using asymptotic analysis. In the same paper that essentially "solved" inventory models with backlogging and positive lead times by showing that the high-dimensionality of the underlying dynamic programs collapses and the problem becomes (effectively) 1-dimensional, Karlin and Scarf (1958) observe that no such simplification occured in the setting of lost sales and positive lead times. Indeed, much of the early literature on such lost sales models with positive lead times was of such a negative nature (Yaspan (1961), Morton (1969)). As such, the natural formulation of the problem as a dynamic program becomes high-dimensional as the lead time grows large, rendering known computational methods intractable due to the curse of dimensionality. Until the mid 2000's, much of the subsequent literature thus focused on analyzing a vast array of heuristics for this problem, typically without any performance guarantees. Some work also proved limited structural properties of heuristics, and in some cases the optimal policy itself. For a comprehensive overview of such work, and also other related variants of lost sales models, we refer the interested reader to the excellent survey by Bijvank and Vis (2011). A particularly relevant example of such a result here (as we will later see) is the work of Reiman (2004), who proves some asymptotic comparisons between the constant-order policy and certain base-stock policies. Here we recall that a constant-order policy is one in which the same constant amount is ordered in every period (or depending on the problem at some fixed intervals), independent of the state of the system. Starting in the mid-2000's, a sequence of results substantially improved our understanding of these models. First, Zipkin (2008) proves that (under a certain transformation) the value function associated with the natural dynamic program for this problem has a property called $L^{\natural}$-convexity"(analogous to convexity) (Zipkin (2008a)). Next, Levi et al. (2008) prove that a certain tractable balancing heuristic yielded a 2-approximation for the problem (i.e. incurs at most twice the expected cost of an optimal policy). Substantial progress was also made regarding algorithms which could efficiently solve the relevant dynamic programs when the lead time is small (Chen, Dawande and Janakiraman (2014), Halman et al. (2009)), and additional progress was made using various methods in approximate dynamic programming and stochastic control (Brown and Smith (2014)). For a numerical comparison of several of these algorithms and heuristics, we refer the interested to Zipkin (2008). **However, the fundamental question of how to compute nearly optimal policies for such models when the lead time is large remained a fundamentally open problem**.

5.1.3. Overview of recent progress using asymptotic analysis. In the last decade, asymptotic analysis has transformed our understanding of the optimal policy for such systems when the lead time is large. There were two main strands of work along these lines. Both assume the simplest non-trivial setup and cost structure, in which time is discrete, there is a single product, demand is i.i.d. from a known distribution, order quantities and demands can be fractional, and there are two types of costs incurred. First, there is a linear per-unit holding cost incurred in every period for each unit of on-hand inventory, where we note that no cost is incurred for inventory in the pipeline vector. Second, there is a linear per-unit lost-sales penalty incurred in every period for each unit of inventory which is lost. We note that such a simple setup with linear costs is the "canonical model" in this context, going back to the original work of Karlin and Scarf (1958), and that many other cost structures are equivalent to this one under an appropriate transformation. In all cases, we restrict our discussion to the corresponding stationary/infinite-horizon problem in which the objective is to minimize the long-run average cost, although several of the results also hold in the finite horizon setting (Goldberg et al. (2016)).

Work showing that a base-stock policy is asymptotically optimal as the lost sales penalty grows large. Huh et al. (2009) show that if one fixes the holding cost, lead time, and demand distributions, and lets the lost sales penalty grow large, then a base-stock policy is asymptotically optimal, where we note that (as discussed in Zipkin (2008)) such a heuristic had been proposed multiple times in the literature (although never rigorously justified). This is intuitive, as when the lost sales penalty is very large, any policy that performs well (including the optimal policy) must keep the inventory level sufficiently high such that almost no losses incur (or the expected cost would be too high). But in such a regime, the fact that unfulfilled demand is lost (as opposed to backlogged) hardly matters, since essentially no demand goes unfulfilled. Thus the policy which is optimal in the setting of backlogged demand, which is (as previously noted) a simple base-stock policy, will similarly be (at least approximately) optimal in the lost sales case when this penalty is large. We note that as the optimal base-stock level can be computed efficiently by solving a convex program, this leads to efficient approximately optimal policies. In this setting, the lead time can be any fixed (possibly large) integer, which is held fixed as the lost sales penalty grows large. We also note that the setting of large lost sales penalty (relative to e.g. holding costs) arises in several applications, and point the reader to Huh et al. (2009)

for further details. The main proof technique involves: 1. proving several formal comparisons (in the sense of stochastic comparison (Muller and Stoyan (2002)) between lost sales models and simpler backlogging models under related problem parameters; and 2. proving that as the lost sales penalty grows large, certain features of the problem become asymptotically negligible due to the fact that any optimal policy will ensure that very little demand is lost.

Work showing that a constant-order policy is asymptotically optimal as the lead time grows large. Xin and Goldberg (2016) show that if one fixes the holding cost, lost sales penalty, and demand distribution, and lets the lead time grow large, then a constant-order policy is asymptotically optimal, where we note that the use of constant-order policies for such problems was introduced in Reiman (2004), and further supported numerically in Zipkin (2008). This is again intuitive, as when the lead time is very large, such a significant amount of randomness is injected into the system between when an order for more inventory is placed and when the order is received, that "being smart" algorithmically provides almost no benefit, and thus one might as well use an open-loop policy. We note that Xin and Goldberg (2016) actually explicitly bound the optimality gap, with the optimality gap decaying to zero at least as fast as the exponential rate of convergence of the expected waiting time in a related single-server queue to its steady-state value. We also note that here the "best" constant to use in the constant-order policy does not depend on the lead time (as when the same constant is ordered in every period the lead time becomes irrelevant in steady-state), and that this optimal constant can be computed by solving a convex program related to the waiting time in a single-server queue (Xin and Goldberg (2016)). The main proof technique involves: 1. showing that over a single lead time the on-hand inventory behaves like a certain queueing system; and 2. applying convexity-type arguments (including Jensen's inequality) and several additional bounds to show that a constant-order policy performs nearly optimally, with the constant (intuitively) coinciding with the expected value of each component of the random pipeline vector induced by an optimal (stationary) policy. We note that previously Goldberg et al. (2016) have derived a similar result (in the finite-horizon setting, and under different assumptions) using a more complicated coupling argument; and that closely related convexity arguments appeared previously in the work of Hajek (1983).

5.2. Formal model and problem statement.

In this section, we formally define the lost-sales inventory model and optimization problem. The inventory system consists of a single item, with integer lead time L, to be managed in discrete time. We assume demand is i.i.d., and let $\{D_t, t \geq 1\}$ be an associated sequence of i.i.d. demand realizations, distributed as the non-negative r.v. D with distribution $\mathcal{D}$, which we assume to have finite mean, and (to rule out certain trivial degenerate cases) to have strictly positive variance. At each time t, the system evolves as follows.

- A new amount of inventory (ordered L periods ago and now at the head of the pipeline vector) is delivered and added to the on-hand inventory.
- A new order is placed (to be received in L periods, initially joins the end of the pipeline vector).
- The demand in period t is realized.
- The on-hand inventory and pipeline vector are updated.
- Costs for period t are incurred.

5.2.1. Formal definition of inventory control policy. The goal is to select an inventory control policy π, which consists of a mapping from states (in principle also entire histories) of the inventory system to order quantities. Namely, a feasible policy π will be equated with the valid definition of a discrete-time ordering process $\{o^{\pi}_t, t \geq 1\}$. We now describe the set of admissible policies $\pi = \{o^{\pi}_t, t \geq 1\}$. First it will be helpful to formalize two additional processes implicitly defined by a policy π (along with initial conditions).

- There is an L-dimensional pipeline vector process $\{\mathbf{x}^{\pi,t}, t \geq 1\}$, with $x^{\pi,t}_i$ equal to $o^{\pi}_{t-(L-i+1)}$ for $t \geq L+1$ and $i \in [1, L]$. Note that $x^{\pi,t+1}_i = x^{\pi,t}_{i+1}$ for $i \in [1, L-1]$, while $x^{\pi,t+1}_L = o^{\pi}_t$.
- There is an inventory process $\{\mathcal{I}^{\pi}_t, t \geq 1\}$, defined through the dynamical inventory equation $\mathcal{I}^{\pi}_{t+1} = \max\big(0, \mathcal{I}^{\pi}_t + x^{\pi,t}_1 - D_t\big)$.

Then we say that a policy $\pi = \{o^{\pi}_t, t \geq 1\}$ is admissible if it satisfies the following properties.

- **Adaptivity**. Letting $\mathcal{F}$ denote the filtration generated by the demand process, we require that π_t is adapted to $\mathcal{F}_{t-1}$ for all $t \geq 1$.
- **Non-negativity**. We require that π is non-negative.

We let Π denote the set of all admissible policies.

5.2.2. Formal problem statement We assume a per-unit linear inventory holding cost of h per unit of time; and per-unit linear lost sales penalty of p per unit of time. Thus under a given admissible policy π, at each time t an expected inventory cost $C^{\pi}(t) \triangleq h\mathbb{E}\big[\big(\mathcal{I}^{\pi}_t + x^{\pi,t}_1 - D_t\big)^+\big] + p\mathbb{E}\big[\big(\mathcal{I}^{\pi}_t + x^{\pi,t}_1 - D_t\big)^-\big]$ is incurred. We then define the long-run average cost of a policy π as $\mathcal{C}^{\pi} \triangleq \limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \mathcal{C}^{\pi}(t)$. In that case, the formal problem of interest is $\inf_{\pi\in\Pi} C^{\pi}$.

5.3. Detailed overview of Huh et al. (2009).

We now review the results and approach of Huh et al. (2009). We begin by making several additional definitions. Let $\mathrm{OPT}(L, h, \mathcal{D}, p) \triangleq \inf_{\pi\in\Pi} C^{\pi}$, i.e. the optimal value with the dependence on the underlying parameters $L, h, \mathcal{D}, p$ made explicit. Let $\mathrm{ORDERUPL}\big(L, h, \mathcal{D}, p, S\big)$ denote the expected long-run performance of the following policy: if the inventory position is less than S, order an amount of inventory which brings it up to S; else order nothing. Here we recall that the inventory position is the sum of all orders in the pipeline and the on-hand inventory. Namely, this coincides with the performance of a base-stock policy for the associated lost sales system, where we recall that such a policy (with appropriate S) is actually optimal for the corresponding system in which demand is backlogged (instead of lost). Finally, let $\mathrm{ORDERUPB}\big(L, h, \mathcal{D}, b, S\big)$ be the expected long-run performance of the same exact policy, but implemented in a system in which demand is backlogged instead of lost, where this system with backlogging incurs a linear holding cost of h and a linear backlogging penalty of b. As systems with backlogging will not be the primary object of study, we will not formally define this system here, instead referring the reader to Huh et al. (2009) for further details.

5.3.1. Stochastic comparison between lost sales and backlogging models. The proof of Huh et al. (2009) relies heavily on two stochastic comparison results, which compare the costs incurred by various policies (optimal or otherwise) in backlogging and lost sales models with related (but different) parameters. We begin by stating the two stochastic comparison results, which were proven earlier in Janakiraman et al. (2007).

Lemma 1.

1. $OPT(L, h, \mathcal{D}, p) \geq \inf_S ORDERUPB\big(L, h, \mathcal{D}, \frac{p}{L+1}, S\big)$.
2. $\inf_S ORDERUPL\big(L, h, \mathcal{D}, p, S\big) \leq \inf_S ORDERUPB\big(L, h, \mathcal{D}, p + L \times h, S\big)$.

Note that the first part of Lemma 1 states that the optimal cost in a lost sales system with one set of parameters is at least the optimal cost in a related inventory system with backlogging cost a fraction of the lost sales penalty in the original lost sales system, which follows from the fact that a base-stock policy is optimal for the associated system with backlogging (and thus the optimal value is achieved by minimizing over S). The second part of the lemma states that the cost incurred by a (suboptimal) base-stock policy in a lost sales system with one set of parameters is at most the optimal cost in a related inventory system with backlogging cost a multiple of the lead time larger than the lost sales penalty in the original lost sales system.

Proof sketch of Lemma 1: For simplicity, suppose $\mathcal{D}$ has support on the positive integers (this is not needed, but will simplify our exposition here), and that $\mathrm{OPT}(L, h, \mathcal{D}, p)$ is attained by some policy π_L. We proceed by showing that (under an appropriate coupling between the systems and policies) the total cost incurred by π_L in the first t periods (under the empty initialization) exceeds that incurred in a different inventory system in which demand is backlogged (instead of lost) and in which decisions are made by a carefully constructed suboptimal (i.e. not base-stock) policy, where we note that the desired long-run-average result then follows from a limiting argument (the details of which we omit here). Thus let inventory system $\mathcal{S}_L$ be the lost sales system with parameters $L, h, \mathcal{D}, p$ (under the empty initialization), operating under the optimal policy π_L for that system (note that π_L is in general very complicated). Alternatively, let inventory system $\mathcal{S}_B$ be the backlogging system with parameters $L, h, \mathcal{D}, \frac{p}{L+1}$ (under the empty initialization), operating under the following non-standard inventory control π_B, where we note that both systems ($\mathcal{S}_L$ and $\mathcal{S}_B$) will see the same (coupled) sequence of demands. The policy π_B is non-standard in that it maintains both a non-negative on-hand inventory, as well as a non-negative queue holding all backlogged units of demand (allowing both of these to be positive at the same time, incurring both a per-unit holding cost for every unit in the on-hand inventory and a per-unit backlogging cost for every unit in the backlogging queue, in contrast to a standard policy which would always meet as much backlogged demand as possible at each time), and also maintains two pipeline vectors (a regular pipeline and a backlogged pipeline). It is also non-standard in that in each period, it places two types of orders. First, in each period, it places a "regular" order which enters the back of the regular pipeline (and arrives to the system L periods later, joining the on-hand inventory at that time). The regular order it places is equivalent to the amount of inventory π_L orders in that same period (recall that demands are coupled in the two systems). Second, in each period it also places a "backlogged" order equal to the number of units of demand that joined the backlogged queue in the previous period, which enters the back of the backlogged pipeline and also arrives L periods later. Here (in a given period) an arriving unit of demand joins the backlogged queue if there is no on-hand inventory in that period to meet it. The policy π_B also allocates units of inventory to units of demand in a non-standard manner. In particular, when a unit of demand first arrives, π_B attempts to meet that demand using its on-hand inventory (derived entirely from regular orders which have gone through the regular pipeline). Any units of demand that arrive and cannot be satisfied in this way join the backlogged queue, to eventually be satisfied exclusively by the backlogged order made in the following period (which by construction will be for exactly the "right" amount to meet this backlogged demand, i.e. every unit of inventory which arrives from the backlogged pipeline is immediately used to meet items in the backlogged queue, never joining the on-hand inventory). We suppose that the decision of precisely which "units" of demand are to be met if e.g. there is enough on-hand inventory to satisfy some but not all of the new units of demand in a given period is made under some arbitrary indexing of the units of demand. Then the following may be proven by natural induction arguments, where we omit the details here.

1. The pipeline vector in $\mathcal{S}_L$ and the regular pipeline vector in $\mathcal{S}_B$ are equal at all times.

2. The on-hand inventory levels in $\mathcal{S}_L$ and $\mathcal{S}_B$ are equal at all times, and hence the holding costs incurred in both systems are identical.

3. The amount of inventory in the backlogged pipeline in $\mathcal{S}_B$ is at most the number of units of demand in the backlogged queue (in $\mathcal{S}_B$) at all times.

4. The number of units of demand that join the backlogged queue in $\mathcal{S}_B$ equals the number of units of demand that are lost in $\mathcal{S}_L$, and each unit of demand that joins the backlogged queue in $\mathcal{S}_B$ spends exactly L + 1 time units in the backlogged queue. As the unit backlogging cost in $\mathcal{S}_B$ is exactly a $\frac{1}{L+1}$ fraction of the lost sales penalty in $\mathcal{S}_L$, it follows that the lost sales penalties incurred in $\mathcal{S}_L$ are identical to the backlogging penalities incurred in $\mathcal{S}_B$ (under an appropriate accounting scheme).

It can also be shown that

- The total inventory and backlogging cost incurred by non-standard service mechanism π_B in $\mathcal{S}_B$ is larger than the corresponding costs incurred in the same backlogging system under an optimal base-stock policy, i.e. $\inf_S \text{ORDERUPB}\big(L, h, \mathcal{D}, \frac{p}{L+1}, S\big)$, as the non-standard service mechanism is suboptimal.

Combining the above completes the proof of the first comparison result. The second comparison result again follows from arguments bounding the average cost of the relevant lost sales system by that of the associated system with backlogging, and we omit the details.

5.3.2. Formal statement of asymptotic optimality and additional proof details. With Lemma 1 in hand, we can now complete the proof of the main asymptotic result of Huh et al. (2009). First, we formally introduce a certain technical condition from Huh et al. (2009), required for their results to hold.

Definition 1 (Asymptotically sublinear mean residual life (ASMRL) Assumption). Let us say that the demand distribution $\mathcal{D}$ has asymptotically sublinear mean residual life if either $\mathcal{D}$ has bounded support, or $\lim_{x\to\infty} \frac{E[D-x|D>x]}{x} = 0$.

Then the formal asymptotic optimality result proven in Huh et al. (2009) is the following.

Theorem 1. *If $\mathcal{D}$ satisfies the ASMRL Assumption, then* $\lim_{p\to\infty} \frac{\inf_S ORDERUPL\left(L,h,\mathcal{D},p,S\right)}{OPT(L,h,\mathcal{D},p)} = 1$. *Namely, order-up policies are asymptotically optimal as the lost-sales penalty grows to ∞.*

Proof sketch of Theorem 1: It follows from Lemma 1 that

$$\frac{\inf_S \text{ORDERUPL}\big(L,h,\mathcal{D},p,S\big)}{\text{OPT}(L,h,\mathcal{D},p)} \le \frac{\inf_S \text{ORDERUPB}\big(L,h,\mathcal{D},p+L\times h,S\big)}{\inf_S \text{ORDERUPB}\big(L,h,\mathcal{D},\frac{p}{L+1},S\big)}.$$

That

$$\lim_{p\to\infty} \frac{\inf_S \text{ORDERUPB}\big(L,h,\mathcal{D},p+L\times h,S\big)}{\inf_S \text{ORDERUPB}\big(L,h,\mathcal{D},\frac{p}{L+1},S\big)} = 1$$

then follows from several additional stochastic comparison arguments, combined with the basic properties of the 1-period newsvendor problem, some straightforward algebra, and the ASMRL Assumption, and we omit the details. We note that the intuition behind this final limit is that so long as $p\to\infty$, whether the backlogging penalty is $p+L\times h$ or $\frac{p}{L+1}$ does not really matter, as in both cases the penalty for incurring any backlogging is so high that (up to lower order considerations) neither system incurs an asymptotically meaningful amount of backlogging, where we note that the ASMRL assumption intuitively ensures that paying a very high price on very little backlogged demand still leads to an asymptotically negligible cost.

5.4. Detailed overview of Xin and Goldberg (2016).

We now review the results and approach of of Xin and Goldberg (2016). We begin by making several additional definitions. Let us say that a policy $\pi\in\Pi$ is stationary if there exists a measurable function $f:\mathcal{R}^{L+1}\to\mathcal{R}^{+}$ such that w.p.1, $o^{\pi}_t = f\big(\mathbf{x}^{\pi,t},\mathcal{I}^{\pi}_t\big)$ for all $t\ge 1$. Due to the i.i.d. nature of the demand process, it follows from the basic theory of MDP that there exists an optimal policy $\pi\in\Pi$ which is stationary, and here we assume π_L to be some fixed such stationary policy.

5.4.1. Formalizing the steady-state of π_L and its properties. Here we formalize the notion of the steady-state pipeline and inventory induced by π_L, and review some important properties of the associated r.v.s. Note that π_L induces an $L+1$-dimensional Markov chain (L dimensions for the pipeline vector, 1 for the inventory level) with continuous state-space and discontinuous dynamics, which are not easy to explicitly characterize or control (as the policy is in general defined abstractly through certain dynamic programming and fixed-point equations). As discussed in Xin and Goldberg (2016), it seems impossible to rule out e.g. the pathological situation that the Markov chain induced by π_L is (for example) periodic, or more generally does not induce a “nice” stationary distribution. Fortunately, it can be shown through standard perturbative and limiting arguments (i.e. perturb π_L slightly to force regenerative behavior, and then let the perturbation size go to zero) that there must exist a probability measure jointly on the pipeline vector and inventory level with several intuitive and desirable properties, coinciding with how one would “expect” the steady-state of such a stationary optimal policy to behave. In particular, the following is proven in Xin and Goldberg (2016).

THEOREM 2. *One may construct an L-dimensional random vector $\boldsymbol{\chi}^*$ (intuitively the steady-state pipeline vector), r.v. $\mathcal{I}^*$ (intuitively the steady-state inventory level), and $\{D_i, i \geq 1\}$ on a common probability space such that the following are true.*

1. $(\boldsymbol{\chi}^*, \mathcal{I}^*)$ *is w.p.1 non-negative, has finite mean, and is independent of* $\{D_i, i \geq 1\}$.
2. $\mathbb{E}[\mathcal{I}^*] = \mathbb{E}\left[\max_{j=0,\ldots,L}\left(\sum_{i=1}^{j}(\chi^*_{L+1-i} - D_{L+1-i}) + \delta_{j,L}\mathcal{I}^*\right)\right]$.
3. $\mathbb{E}[\chi^*_i] = \mathbb{E}[\chi^*_1] \leq \mathbb{E}[D_1]$ *for all* i.
4. $h\mathbb{E}[\mathcal{I}^*] + p\mathbb{E}[D_1] - p\mathbb{E}[\chi^*_1] = OPT(L, h, \mathcal{D}, p)$.

Here we note that Theorem 2.(2) follows from the max-plus style updates dictated by the dynamic inventory equations, and the celebrated Lindley recursion which shows that such a process has a supremum-type representation (Lindley (1952)). We also note that Theorem 2.(3) follows from stationarity and the manner in which the pipeline vector updates over time.

5.4.2. Applying convexity over a lead time. Next, we observe that since the demands over a lead time are necessarily independent from the orders received over that lead time and the inventory level at the start of that lead time, the right-hand side of Theorem 2.(2) is a jointly convex function of χ^* and $\mathcal{I}^*$ (for each fixed realization of $\{D_i, i = 1, \ldots, L\}$), which will allow us to apply the (conditional) multi-variate Jensen's inequality, to conclude the following.

LEMMA 2.

$$\mathbb{E}[\mathcal{I}^*] \geq \mathbb{E}\left[\max_{j=0,\ldots,L}\left(j\mathbb{E}[\chi^*_1] - \sum_{i=1}^{j} D_i + \delta_{j,L}\mathbb{E}[\mathcal{I}^*]\right)\right] \geq \mathbb{E}\left[\max_{j=0,\ldots,L}\left(j\mathbb{E}[\chi^*_1] - \sum_{i=1}^{j} D_i\right)\right].$$

Combining Theorem 2.(4) with Lemma 2, we conclude the following.

COROLLARY 1.

$$OPT(L, h, \mathcal{D}, p) \geq h\mathbb{E}\left[\max_{j=0,\ldots,L}\left(j\mathbb{E}[\chi^*_1] - \sum_{i=1}^{j} D_i\right)\right] + p\mathbb{E}[D_1] - p\mathbb{E}[\chi^*_1].$$

5.4.3. Connecting to the constant-order policy. As we will now see, the bound of Corollary 1 looks very similar to the performance of the constant-order policy which always orders $\mathbb{E}[\chi^*_1]$ in every period. Indeed, we now review in greater depth the constant-order policy, and characterize the best constant-order policy. For any $r \in [0, \mathbb{E}[D_1])$, the constant-order policy π_r is the policy that places the constant order r in every period. It is well-known that for any such $r \in [0, \mathbb{E}[D_1])$, the sequence of on-hand inventory levels in the associated lost sales system (controlled by π_r and initially empty) converges in distribution, and in expectation, to a limiting r.v. with finite mean, which we denote by I^r_∞. I^r_∞ has the same distribution as the steady-state waiting time in the corresponding $GI/GI/1$ queue with interarrival distribution $\mathcal{D}$ and processing time distribution the constant r. It is well-known (e.g. Asmussen (2003), again due to Lindley's recursion) that $I^r_\infty \sim$

$\sup_{j\geq 0}\left(jr-\sum_{i=1}^{j}D_i\right)$. It follows that for any $r\in[0,\mathbb{E}[D_1])$, the long-run average cost incurred by the corresponding constant-order policy equals

$$\mathrm{CONST}\big(L,h,\mathcal{D},p,r\big)=\ h\mathbb{E}\left[\sup_{j\geq 0}\left(jr-\sum_{i=1}^{j}D_i\right)\right]+p\mathbb{E}[D_1]-pr, \tag{1}$$

independent of the lead time L.

Note that by combining Corollary 1 and (1), we immediately derive the following bound on the performance of the constant-order policy which orders $\mathbb{E}[\chi_1^*]$ in every period. Here we recall that $\mathbb{E}[\chi_1^*]$ really depends on L (as well as $h,\mathcal{D},p$) as it is derived from the behavior of the optimal policy π_L, and here we make this dependence explicit through the notation $\chi_1^*(L,h,\mathcal{D},p)$. Then we deduce the following bound.

Corollary 2. *$CONST\big(L,h,\mathcal{D},p,\mathbb{E}\big[\chi_1^*(L,h,\mathcal{D},p)\big]\big)-OPT(L,h,\mathcal{D},p)$ is at most*

$$h\left(\mathbb{E}\left[\sup_{j=0,\dots,\infty}\left(j\mathbb{E}\big[\chi_1^*(L,h,\mathcal{D},p)\big]-\sum_{i=1}^{j}D_i\right)\right]-\mathbb{E}\left[\max_{j=0,\dots,L}\left(j\mathbb{E}\big[\chi_1^*(L,h,\mathcal{D},p)\big]-\sum_{i=1}^{j}D_i\right)\right]\right).$$

Although we will later see that such differences between suprema can in some sense be nicely controlled, to do so we would have to have some concrete understanding of $\mathbb{E}[\chi_1^*]$, and especially the gap $\mathbb{E}[D_1]-\mathbb{E}[\chi_1^*]$. Although some relevant (yet crude) bounds are derived in Goldberg et al. (2016), here we will take a more refined approach which allows us to sidestep this lack of understanding of $\mathbb{E}\big[\chi_1^*(L,h,\mathcal{D},p)\big]$.

Indeed, as it is easily verified from (1) that $\mathrm{CONST}\big(L,h,\mathcal{D},p,r\big)$ is a convex function of r on $[0,\mathbb{E}[D_1])$, to find the best possible constant-order policy, it suffices to select the r minimizing this one-dimensional convex function over the compact set $[0,\mathbb{E}[D_1]]$. Let $r_\infty\in\arg\min_{0\leq r\leq\mathbb{E}[D_1]}\mathrm{CONST}\big(L,h,\mathcal{D},p,r\big)$ denote the infimum of this set of optimal constant-order levels, in which case the best constant-order policy will refer to π_{r_∞}. For $r\in\big[0,\mathbb{E}[D_1]\big]$ and $L\geq 1$, let $C_L(r)\triangleq h\mathbb{E}\big[\max_{j=0,\dots,L}\left(jr-\sum_{i=1}^{j}D_i\right)\big]+p\mathbb{E}[D_1]-pr$, and let r_L denote the infimum of $\arg\min_{r\in[0,\mathbb{E}[D_1]]}C_L(r)$. We also define $C_\infty(r)$ to be the associated function with $L=\infty$, which we note coincides with $\mathrm{CONST}\big(L,h,\mathcal{D},p,r\big)$, and thus the previously defined quantity r_∞ is consistent also with this new definition. Indeed, note that $C_L(r)$ is a certain "truncation" of the expression describing the performance of the constant-order policy, in which the relevant supremum is only taken over L periods. Then the following more refined optimality gap is proven in Xin and Goldberg (2016), which side-steps the problem of needing to know $\mathbb{E}[\chi_1^*]$ (as in Corollary 2).

Lemma 3. *$CONST\big(L,h,\mathcal{D},p,r_\infty\big) - OPT(L,h,\mathcal{D},p)$ is at most*

$$
\begin{aligned}
&h\Bigg(\mathbb{E}\big[\sup_{j\geq 0}(jr_\infty - \sum_{i=1}^{j} D_i)\big] - \mathbb{E}\big[\max_{0\leq j\leq L}(jr_\infty - \sum_{i=1}^{j} D_i)\big]\Bigg)\\
&+\quad h\Bigg(\mathbb{E}\big[\max_{0\leq j\leq L}(jr_\infty - \sum_{i=1}^{j} D_i)\big] - \mathbb{E}\big[\max_{0\leq j\leq L}(jr_L - \sum_{i=1}^{j} D_i)\big]\Bigg)\\
&-\quad p(r_L - r_\infty).
\end{aligned}
$$

Proof sketch of Lemma 3: Corollary 1 implies that

$$
\begin{aligned}
\mathrm{OPT}(L,h,\mathcal{D},p) &\geq h\mathbb{E}\left[\max_{j=0,\ldots,L}\left(j\mathbb{E}[\chi_1^*] - \sum_{i=1}^{j} D_i\right)\right] + p\mathbb{E}[D_1] - p\mathbb{E}[\chi_1^*]\\
&= C_L(\mathbb{E}[\chi_1^*]) \quad\geq\quad C_L(r_L)\\
&= h\mathbb{E}\big[\max_{j=0,\ldots,L}\left(jr_L - \sum_{i=1}^{j} D_i\right)\big] + p\mathbb{E}[D_1] - pr_L.
\end{aligned}
$$

Combining with (1) and some straightforward algebra then completes the proof.

5.4.4. Formal statement of asymptotic optimality and additional proof details. For $\theta \geq 0$, let us define

$$
\phi(\theta) \triangleq \exp(\theta r_\infty)\mathbb{E}[\exp(-\theta D_1)] \quad , \quad \gamma \triangleq \inf_{\theta\geq 0}\phi(\theta),
$$

and $\vartheta \in \arg\min_{\theta\geq 0}\phi(\theta)$ denote the supremum of the set of minimizers of $\phi(\theta)$, where we define ϑ to equal ∞ if the above infimum is not actually attained. It follows from definitions and standard arguments with generating functions (see Xin and Goldberg (2016)) that $r_\infty < \mathbb{E}[D_1], \vartheta > 0$, and $\gamma \in [0,1)$. It follows from the celebrated Cramér's Theorem (Deuschel and Strook (1989)), and more generally the theory of large deviations, that up to exponential order $\mathbb{P}(kr_\infty \geq \sum_{i=1}^{k} D_i)$ decays like γ^k as $k\to\infty$. Let $g(h,\mathcal{D},p) \triangleq \inf_{x\in\mathcal{R}}\mathbb{E}\big[h(x-D)^+ + p(x-D)^-\big] > 0$, and Q denote the $\frac{p}{p+h}$ quantile of the demand distribution, where we note that Q is the optimal order quantity in the single-stage newsvendor problem and $\mathrm{OPT}(L,h,\mathcal{D},p) \geq g(h,\mathcal{D},p)$, which (when there is no ambiguity) we will simply denote by g.

Then the formal asymptotic optimality result proven in Xin and Goldberg (2016) is the following.

Theorem 3. *For all $L \geq 1$,*

$$
\frac{CONST\big(L,h,\mathcal{D},p,r_\infty\big)}{OPT(L,h,\mathcal{D},p)} \leq 1 + h\big((1-\gamma)g\big)^{-1}\left(\mathbb{E}[D_1] - r_\infty + \big(e\vartheta(L+1)\big)^{-1}\right)\gamma^{L+1}. \tag{2}
$$

Namely, constant-order policies are asymptotically optimal as the lead time grows to ∞, with the (relative) optimality gap decaying exponentially at rate γ.

The formal proof proceeds by bounding each of the terms appearing in Lemma 3. A key insight is that expectations of maxima of random walks (such as those appearing in Lemma 3) have an elegant representation through the celebrated Spitzer's identity (Spitzer (1956)), which can be used to bound several of the relevant terms (combined with closely related arguments of Kingman (1962)). The remaining terms are bounded by using the fact that $r_L (r_\infty)$ is a minimizer of the convex function $C_L(r)\big(C_\infty(r)\big)$, and applying the relevant first-order optimality conditions (along with the fact that $r_\infty \le r_L$, also proven in Xin and Goldberg (2016), and a straightforward application of Chernoff's bound).

6. Case Study II: Dual-sourcing models with large lead time gap

6.1. Introduction to dual-sourcing models with lead time gap.

6.1.1. High-level introduction. Dual-sourcing is a common practice in supply chain management, in which a company usually purchases its materials from a regular supplier at a lower cost, but is also able to obtain materials from an expedited supplier at a higher cost when needed. For example, in the summer of 2003, Amazon used FedEx to deliver the new Harry Potter more promptly and maintained regular shipping via UPS (Kelleher (2003), Veeraraghavan and Scheller-Wolf (2008)). Similarly, Allon and van Mieghem (2010) describe an example of a U.S. company that has two suppliers, one in Mexico and one in China. The Mexican supplier has shorter lead time but higher ordering cost; the Chinese supplier has longer lead time and lower ordering cost. The company takes advantage of the dual-sourcing strategy to meet their demand more efficiently. Although dual-sourcing is practically attractive, optimizing a dual-sourcing inventory system is notoriously challenging, as the state-space of the underlying dynamic program becomes high-dimensional (with the dimension growing linearly in the gap between the regular and expedited lead times).

6.1.2. Literature review prior to recent progress using asymptotic analysis. Although dual-sourcing systems have been studied since the 1960's, the structure of the optimal policy remains poorly understood, with the exception of when the system is consecutive, i.e., the lead time difference between the two sources is exactly one. Early studies of periodic review dual-sourcing inventory models include Barankin (1961), Daniel (1963), Neuts (1964), which showed that base-stock policies are optimal when the lead times of the two sources are zero and one respectively. Fukuda (1964) extends the result to general lead time settings as long as the lead time difference remains one. Whittmore and Saunders (1977) show that the optimal policy is no longer a simple base-stock policy when the lead time difference is larger than one. In the subsequent decades, a vast literature has arisen for such models, and we refer the interested reader to the survey of Minner (2003), as well as e.g. the more recent works of Feng et al. (2006), Fox et al. (2006), Chen, Xue and Yang (2013), Huggins and Olsen (2010), Angelus and Özer

(2015), Boute and Van Mieghem (2015), Gong, Chao and Zheng (2014), Song and Zipkin (2009), and the references therein. As an exact solution seems out of reach, researchers have instead investigated certain structural properties of the optimal policy (Hua et al. (2014)), and analyzed various heuristic policies including single index (SI) policies (Scheller-Wolf, Veeraraghavan and van Houtum (2008)), dual index (DI) policies (Veeraraghavan and Scheller-Wolf (2008)), and additional policies (Sheopuri, Janakiraman and Seshadri (2010), Boute and Van Mieghem (2015)).

A simple policy that is implemented in practice, and which will be quite relevant to our subsequent analysis, is the so-called *Tailored Base-Surge* (TBS) policy. It was first proposed and analyzed in Allon and van Mieghem (2010), where we note that closely related standing order policies had been studied previously (Rosenshine and Obee (1976), Janssen and De Kok (1999)). Under such a TBS policy, a constant order is placed at the regular source in each period to meet a *base* level of demand, while the orders placed at the express source follow an order-up-to rule to manage demand *surges*. We refer to Mini-Case 6 in Van Mieghem (2008) for more about the motivation and background of TBS policies. Note that dual-sourcing inventory systems in which a constant-order policy is implemented for the regular source are essentially equivalent to single-sourcing inventory systems with constant returns, which have been investigated in the literature (Fleischmann and Kuik (2003), DeCroix, Song and Zipkin (2005)). Allon and van Mieghem (2010) analyze TBS policies in a continuous review model, and their focus was to find the best TBS policy. Numerical results in Klosterhalfen, Kiesmüller and Minner (2011) and Rossi, Rijpkema and van der Vorst (2012) show that TBS policies are comparable to DI policies, and outperform DI policies for some problem instances. Allon and van Mieghem (2010) conjecture that the TBS policy performs more effectively as the lead time difference between the two sources grows. Janakiraman et al. (2015) analyze a periodic review model and studied the performance of TBS policies. They provide an explicit bound on the performance of TBS policies compared to the optimal one when the demand had a specific structure, and provide numerical experiments suggesting that the performance of the TBS policy improves as the lead time difference grows large. **However, the fundamental question of why TBS policies seem to perform so well, and more generally the question of how to compute nearly optimal policies for such models when the lead time gap is large, remained fundamentally open problems**.

6.1.3. Overview of recent progress using asymptotic analysis. Recently, substantial progress was made towards answering these fundamental questions using asymptotic analysis. In particular, it was shown in Xin and Goldberg (2017) that when the lead time of the express source is held fixed (along with the relevant costs and demand distribution), a simple TBS policy is asymptotically optimal as the lead time of the regular source grows large. Furthermore, as the "best" TBS policy can be computed by solving a convex program that does not depend on the lead time of the regular source (Janakiraman et al. (2015)), these results lead directly to very efficient algorithms (with complexity independent of the lead time of the regular source) with asymptotically optimal performance guarantees. Xin and Goldberg (2017) also provide explicit bounds on the optimality gap, showing it decays inverse-polynomially in the lead time of the regular source. Perhaps most importantly, since some companies are already implementing such TBS policies (Allon and van Mieghem (2010)), these results provide theoretical support for the use of TBS policies in practice. The main proof technique can be viewed as a much more sophisticated implementation of the technique used by Xin and Goldberg (2016) in Case Study I above. In particular, the proof involves: 1. applying convexity-type arguments (specifically the conditional Jensen's inequality) to show that over a single lead time the dual-sourcing optimization problem can be "lower-bounded" by a much simpler inventory optimization problem in which the same constant (intuitively the expected amount ordered by an optimal stationary policy from the regular source) is always ordered from the regular source; 2. recognizing that a stationary base-stock policy is asymptotically optimal for this simpler problem; and 3. showing that this lower bound is asymptotically matched by the TBS policy which (intuitively) orders the same constant from the regular source and applies the same base-stock policy at the express source. To formalize these arguments, additional methodologies from the MDP literature such as the vanishing discount factor approach (in which the long-run-average problem is approximated by a sequence of stationary discounted problems) are applied as well.

6.2. Formal model and problem statement.

In this section, we formally define the dual-sourcing inventory model and optimization problem. For simplicity, we will restrict ourselves to the setting in which the express source has lead time zero, which captures essentially all of the complexity of the problem, and refer the reader to Xin and Goldberg (2017) for the more general results which hold for any fixed lead time at the express source, and which are very similar in nature to the results for the zero express lead time case. The inventory system consists of a single item, with lead time 0 from the express source and lead time L from the regular source, to be managed in discrete time. We assume demand is i.i.d., and let $\{D_t, t \geq 1\}$ be an associated sequence of i.i.d. demand realizations, distributed as the non-negative r.v. D with distribution $\mathcal{D}$, which we assume to have finite mean, and (to rule out certain trivial degenerate cases) to have strictly positive variance. We also suppose demand is backlogged, which in the dual-sourcing setting is consistent with the literature and in which the curse of dimensionality manifests (recall that in the single-source setting the backlogged case is easy). At each time t, the system evolves as follows.

- A new amount of inventory (ordered L periods ago and now at the head of the pipeline vector) is delivered and added to the on-hand inventory.
- New orders are placed from the regular source (to be received in L periods, initially joins the end of the pipeline vector) and from the express source (received immediately).
- The demand in period t is realized.
- The on-hand inventory and pipeline vector are updated.
- Costs for period t are incurred.

6.2.1. Formal definition of inventory control policy. The goal is to select an inventory control policy π, which consists of a mapping from states (in principle also entire histories) of the inventory system to order quantities. Namely, a feasible policy π will be equated with the valid definition of two discrete-time ordering processes: $\{o_t^{\pi,R}, t \geq 1\}$ (the ordering process from the regular source) and $\{o_t^{\pi,E}, t \geq 1\}$ (the ordering process from the express source). We now describe the set of admissible (i.e. feasible) policies $\pi = \{(o_t^{\pi,R}, o_t^{\pi,E}), t \geq 1\}$. First it will be helpful to formalize two additional processes implicitly defined by a policy π (along with initial conditions).

- There is an (L-1)-dimensional *truncated* regular pipeline vector process $\{\mathbf{x}^{\pi,t}, t \geq 1\}$, with $x_i^{\pi,t}$ equal to $o_{t-L+i}^{\pi,R}$ for $t \geq L$ and $i \in [1, L-1]$. Note that $x_i^{\pi,t+1} = x_{i+1}^{\pi,t}$ for $i \in [1, L-2]$, and $x_{L-1}^{\pi,t+1} = o_t^{\pi,R}$. In contrast to the lost sales setting, here we use an "accounting trick" by letting the "on-hand inventory" at time t represent said inventory *after* receiving the regular order in that period. Noting that the regular order to be received in period $t+1$ is that which, from the perspective of period t, was ordered (only) $L-1$ periods in the past (as opposed to L periods in the past), it turns out that under such an accounting trick one can express all relevant updates of the system without ever explicitly tracking the regular order placed L periods before time t, and may instead track only the $L-1$ most recent regular orders. This leads to the above notion of truncated regular pipeline, allowing us to reduce the dimension of the regular pipeline vector. Although not a major conceptual difference, such non-standard accounting turns out to simplify several manipulations relevant to the dual-sourcing problem (Sheopuri, Janakiraman and Seshadri (2010), Xin and Goldberg (2017)), and e.g. reduces the relevant notion of lead time to the *gap* between the the regular and expedited lead times in the setting of a general expedited lead time. We note that in the lost sales setting (i.e. Case Study I), such reductions were not used, and as such (for simplicity) we never fully specified certain analogous details (instead leaving them implicit).
- There is an inventory process $\{\mathcal{I}_t^\pi, t \geq 1\}$, defined through the dynamical inventory equation $\mathcal{I}_{t+1}^\pi = \mathcal{I}_t^\pi + x_1^{\pi,t} + o_t^{\pi,E} - D_t$. As noted above, $\mathcal{I}_t^\pi$ represents the inventory level after receiving the regular order in that period.

Then we say that a policy $\pi = \{(o_t^{\pi,R}, o_t^{\pi,E}), t \geq 1\}$ is admissible if it satisfies the following properties.

- **Adaptivity**. Letting $\mathcal{F}$ denote the filtration generated by the demand process, we require that π_t is adapted to $\mathcal{F}_{t-1}$ for all $t \geq 1$.
- **Non-negativity**. We require that π is non-negative.

We let Π denote the set of all admissible policies.

6.2.2. Formal problem statement We assume per-unit linear inventory holding cost of h per unit of time; and per-unit linear backlogging penalty of b per unit of time. It is known that one can reduce the problem to that in which the regular source has zero cost, and we hence assume that here, supposing a per-unit linear ordering cost of c from the express source. We also let $G(y)$ be the sum of the holding and backlogging costs when the inventory level equals y in the end of a time period, i.e. $G(y) \triangleq hy^+ + by^-$. Then under a given admissible policy π, at each time t an expected inventory cost $C^\pi(t) \triangleq \mathbb{E}\big[c \times o_t^{\pi,E} + G\big(\mathcal{I}_t^\pi + o_t^{\pi,R} + o_t^{\pi,E} - D_t\big)\big]$ is incurred. We then define the long-run average cost of a policy π as $\mathcal{C}^\pi \triangleq \limsup_{T\to\infty} \frac{1}{T}\sum_{t=1}^{T} \mathcal{C}^\pi(t)$. In that case, the formal problem of interest is $\inf_{\pi\in\Pi} C^\pi$.

6.3. Detailed overview of Xin and Goldberg (2017).

We now review the results and approach of of Xin and Goldberg (2017). We let $\text{OPT}(L)$ denote $\inf_{\pi\in\Pi} C^\pi$, with the given regular lead time specified.

6.3.1. Formalizing the steady-state induced by an optimal policy and its properties. As in Xin and Goldberg (2016), first we will formalize the notion of a stationary optimal policy here, and the steady-state induced by the associated Markov chain. In contrast to Xin and Goldberg (2016), here the relevant questions of existence, convergence, etc. are even more subtle (in part due to the lack of boundedness from below, as well as the lack of "natural" regeneration points), as explained in Xin and Goldberg (2017). They again overcome this by proving the existence of certain random vectors with the relevant properties such a stationary distribution would naturally have, through a more sophisticated perturbation and limiting argument Xin and Goldberg (2017). In particular, the following result is proven in Xin and Goldberg (2017), in analogy with Theorem 2 from Case Study I.

THEOREM 4. *For all $L \geq 1$, one may construct an $(L-1)$-dimensional random vector $\boldsymbol{\chi}^{*,L}$ (intuitively the steady-state truncated regular pipeline vector), an L-dimensional random vector*

$\mathbf{q}^{*,L}$ *(intuitively the steady-state vector of expedited orders to be placed over the* next L *periods including the current period), and a random variable* $\mathcal{I}^{*,L}$ *(intuiutively the net on-hand inventory after receiving the regular order in the current period), as well as* $\{D_i, i \geq 1\}$*, on a common probability space such that the following are true. Here* $q_1^{*,L}$ *is intuitively the expedited order placed in the "current" period (i.e. the period in which the on-hand inventory after receiving the regular order is* $\mathcal{I}^{*,L}$*), while* $q_L^{*,L}$ *is intuitively the expedited order placed* L *periods into the future (in which regular order* $\chi_{L-1}^{*,L}$ *is received).*

1. *W.p.1* $(\boldsymbol{\chi}^{*,L}, \mathbf{q}^{*,L})$ *is non-negative. Also,* $(\boldsymbol{\chi}^{*,L}, \mathcal{I}^{*,L})$ *is independent of* $\{D_i, i \geq 1\}$*, and* $q_i^{*,L}$ *is independent of* $\{D_j, j \geq i\}$ *for* $i \in [1, L]$*.*
2. $\chi_i^{*,L} \sim \chi_1^{*,L}$ *for* $i \in [1, L-1]$*, and* $q_i^{*,L} \sim q_1^{*,L}$ *for* $i \in [1, L]$*.*
3. *For all* $k \in [1, L]$*,*

$$\mathcal{I}^{*,L} + \sum_{i=1}^{k-1} (q_i^{*,L} + \chi_i^{*,L} - D_i) + q_k^{*,L} - D_k \sim \mathcal{I}^{*,L} + q_1^{*,L} - D_1.$$

4. $(\boldsymbol{\chi}^{*,L}, \mathbf{q}^{*,L}, \mathcal{I}^{*,L})$ *has finite mean.*
5. $\mathbb{E}[\chi_1^{*,L}] + \mathbb{E}[q_1^{*,L}] = \mathbb{E}[D_1]$*.*
6.

$$OPT(L) \geq c\left(\mathbb{E}[D_1] - \mathbb{E}[\chi_1^{*,L}]\right) + \mathbb{E}\left[G\left(\mathcal{I}^{*,L} + q_1^{*,L} - D_1\right)\right].$$

6.3.2. Vanishing discount factor approach. Although Theorem 4.(6) relates OPT(L) to a certain expectation, it will be convenient to further simplify Theorem 4 by rewriting in terms of a certain infinite-horizon discounted problem, which will later help us formally connect to a TBS policy since such discounted problems have stationary optimal policies (as the TBS policy is stationary in nature). We will then "close the loop" by using the fact that by general MDP theory, as the discount factor approaches one the discounted problem should (in an appropriate sense) "converge" to the original long-run-average (non-discounted) problem. More precisely, we introduce a discount factor α to implement the so-called "vanishing discount factor" approach to analyzing infinite-horizon MDP (Huh, Janakiraman and Nagarajan (2011)), which will allow for a simpler analysis when we pass to the limit as $L \to \infty$ since the "simpler" inventory optimization problem arising from our lower bound (after applying the conditional Jensen's inequality) will then be stationary. Formally, Theorem 4 (combined with some straightforward algebra) implies the following corollary. Let $r_L \triangleq \mathbb{E}[\chi_1^{*,L}]$.

Corollary 3. *For all $L \geq 2$ and $\alpha \in (0,1)$,*

$$\begin{aligned} OPT(L) &\geq c\left(\mathbb{E}[D_1] - r_L\right) + \frac{1-\alpha}{1-\alpha^L} \sum_{k=1}^{L} \alpha^{k-1} \mathbb{E}\left[G\left(\mathcal{I}^{*,L} + q_1^{*,L} - D_1\right)\right] \\ &\geq c\left(\mathbb{E}[D_1] - r_L\right) + (1-\alpha) \sum_{k=1}^{L} \alpha^{k-1} \mathbb{E}\left[G\left(\mathcal{I}^{*,L} + \sum_{i=1}^{k-1} (q_i^{*,L} + \chi_i^{*,L} - D_i) + q_k^{*,L} - D_k\right)\right]. \end{aligned}$$

6.3.3. Applying the conditional Jensen's inequality and relating to a single-source inventory model. We now apply the conditional Jensen's inequality to Corollary 3, which will allow us to lower-bound OPT(L) by the optimal value of a certain finite-horizon discounted single-source inventory model with backlogged demand. We will then relate this finite-horizon problem to an associated infinite-horizon discounted problem, which has an optimal stationary policy. Furthermore, we will connect the behavior of such an optimal stationary policy to the performance of an associated TBS policy, ultimately allowing us to prove our main results. In particular, it follows from Theorem 4 and the independence structure of the relevant r.v.s that for $k \in [1, L]$,

$$\mathbb{E}\left[\mathcal{I}^{*,L} + \sum_{i=1}^{k-1} (q_i^{*,L} + \chi_i^{*,L} - D_i) + q_k^{*,L} - D_k \middle| D_{[k]}\right] =$$

$$\mathbb{E}[\mathcal{I}^{*,L}] + \sum_{i=1}^{k-1} (\mathbb{E}[q_i^{*,L} | D_{[i-1]}] + r_L - D_i) + \mathbb{E}[q_k^{*,L} | D_{[k-1]}] - D_k.$$

Further combining with Corollary 3, the convexity of G, and Jensen's inequality for conditional expectations, we obtain the following result.

Proposition 1. *For any $\alpha \in (0,1)$ and $L \geq 2$, $OPT(L) - c\left(\mathbb{E}[D_1] - r_L\right) \geq$*

$$\begin{aligned} (1-\alpha) \sum_{k=1}^{L} \alpha^{k-1} \mathbb{E}\Bigg[G\Bigg(\mathbb{E}[\mathcal{I}^{*,L}] - r_L + \sum_{i=1}^{k-1} \left(\mathbb{E}[q_i^{*,L} | D_{[i-1]}] - (D_i - r_L)\right) \\ + \mathbb{E}[q_k^{*,L} | D_{[k-1]}] - (D_k - r_L)\Bigg)\Bigg]. \end{aligned} \tag{3}$$

We note that (3) is the discounted cost incurred (during periods $1, \ldots, L$) by the policy ordering $\mathbb{E}[q_i^{*,L} | D_{[i-1]}]$ in period i, of a single-sourcing L-period backlog inventory problem with unit holding cost h, backlogging cost b, zero ordering cost, discount factor α, i.i.d. demand distributed as $D - r_L$ (which we note can be positive *or negative*), zero lead time, and initial inventory $\mathbb{E}[\mathcal{I}^{*,L}] - r_L$, multiplied by $(1-\alpha)$ (Karlin and Scarf (1958)). Such models, and their optimal policies, have been studied in-depth in the literature (Karlin and Scarf (1958), Zipkin (2000), Fleischmann and Kuik (2003)). Let $\overline{\Pi}$ denote the family of all feasible non-anticipative policies for the aforementioned inventory problem (as it is typically defined (Zipkin (2000))). For $\pi \in \overline{\Pi}$ and $i \geq 1$, let $C_i^{\pi}(r_L)$ denote the cost incurred by policy π in the aforementioned inventory problem in period i, if the initial inventory is $-\infty$ (which can at no cost be raised by a policy to any desired level). Let us define

$$V_\alpha^n(r_L) \triangleq \inf_{\pi\in\overline{\Pi}} \mathbb{E}\left[\sum_{i=1}^n \alpha^{i-1} C_i^\pi(r_L)\right] \quad ; \quad V_\alpha^\infty(r_L) \triangleq \inf_{\pi\in\overline{\Pi}} \mathbb{E}\left[\sum_{i=1}^\infty \alpha^{i-1} C_i^\pi(r_L)\right].$$

Then combining the above, we derive the following lower bound for OPT(L).

Lemma 4. *For all $L \geq 2$, and $\alpha \in (0,1)$,*

$$OPT(L) \geq c(\mathbb{E}[D_1] - r_L) + (1-\alpha) V_\alpha^L(r_L). \tag{4}$$

6.3.4. Formal statement of asymptotic optimality and additional proof details. Although explicit bounds on the optimality gap are proven in Xin and Goldberg (2017), here for simplicity we simply state the main asymptotic result, referring the reader to Xin and Goldberg (2017) for the bound on the optimality gap.

TBS Policies. Before stating the main result, let us formally review the defintion of a TBS policy and the "best" TBS policy. A TBS policy $\pi_{r,S}$ with parameters (r,S) is defined as the policy that places a constant order r from the regular supplier in every period, and follows an order-up-to rule from the express supplier which in each period raises the on-hand inventory to S (if it is below S after receiving the regular order in that period), and otherwise orders nothing. More formally, it holds that $o_t^{\pi_{r,S},R} = r$, and $o_t^{\pi_{r,S},E} = \max(0, S - \mathcal{I}_t^{\pi_{r,S}})$, for all t. Furthermore, letting $I_\infty^r \triangleq \sup_{j\geq 0}\left(jr - \sum_{i=1}^j D_i\right)$, and D_1' an independent realization of $\mathcal{D}$, it is shown in Janakiraman et al. (2015) that

$$C^{\pi_{r,S}} = \; c(\mathbb{E}[D_1] - r) + \mathbb{E}\left[G\left(I_\infty^r + S - D_1'\right)\right]. \tag{5}$$

Furthermore, it is shown in Janakiraman et al. (2015) that an optimal choice of r, S, which minimizes (5), can be found efficiently (independent of the lead time L) using convex optimization, and we let r^*, S^* denote some such optimal choice.

Formal statement of asymptotic optimality. The formal asymptotic optimality result proven in Xin and Goldberg (2017) is the following, where again we refer the reader to Xin and Goldberg (2017) for precise analytical bounds on the optimality gap.

Theorem 5. $\lim_{L\to\infty} \frac{C^{\pi_{r^*,S^*}}}{OPT(L)} = 1$. *Namely, TBS policies are asymptotically optimal as the lead time grows to ∞.*

Overview of remainder of the proof by connecting Lemma 4 to a TBS policy. The remainder of the proof follows from the fact that the optimization problem defined by $V_\alpha^\infty(r_L)$ has an optimal stationary base-stock policy, with order-up level (say) $\hat{S}_{\alpha,L}$, and that $V_\alpha^\infty(r_L)$ is "close" to $V_\alpha^L(r_L)$ for appropriate sequencese of α and L with $\alpha \uparrow 1$ and $L \to \infty$ simultaneously. It follows that the lower bound from Lemma 4, i.e.

$$c(\mathbb{E}[D_1] - r_L) + (1-\alpha)V_\alpha^L(r_L), \tag{6}$$

is "almost" $c(\mathbb{E}[D_1] - r_L) + (1-\alpha)V_\alpha^\infty(r_L)$, itself "almost" the cost incurred by $\pi_{r_L,\hat{S}_{\alpha,L}}$ in the corresponding discounted infinite-horizon version of the original dual-sourcing problem (normalized by $1-\alpha$), which as $\alpha \uparrow 1$ is "almost" the long-run-average cost incurred by $\pi_{r_L,\hat{S}_\alpha}$ for the original long-run-average non-discounted problem (by the vanishing discount factor methodology). Here the "almost" arises from subtleties such as the fact that: 1. in (6) the first term $c(\mathbb{E}[D_1] - r_L)$ conceptually comes from the long-run-average (not discounted) setting; 2. in the problem defined by $V_\alpha^\infty(r_L)$ a policy can choose its initial conditions while in the long-run-average problem the initial conditions are implicit from the policy itself; and 3. one must show appropriate boundedness and convergence (e.g. of $\hat{S}_{\alpha,L}$) as $\alpha \uparrow 1$ and $L \to \infty$. Xin and Goldberg (2017) use techniques from MDP, queueing, and the theory of random walks (including again results of Spitzer and Kingman) to show that all these technical issues can be overcome, with all relevant convergences holding in a suitably uniform manner.

7. Case Study III: ATO Systems with large lead times

7.1. Introduction to ATO Systems.

7.1.1. High-level introduction. ATO systems are those in which a manufacturer manages an inventory of raw materials (i.e. components), which are converted into products (each of which requires some amount of each component) over time as demand arrives for different products. To make the idea concrete, let us consider some well-studied special cases, depicted in Figure 1. On the left is the so-called "W system" that has two products to be assembled from three components. The system embodies a commonality strategy: instead of using completely separate sets of components, the two products are assembled from one unit of a common component (component 0) and another unit of a product-specific component (1 or 2). The common component is used to exploit the pooling effect to mitigate uncertainty in demands for individual products. In the middle is the "M system", which has three products assembled from two components. This system provides customers with the flexibility of choosing a product with a single functionality embedded in either component, or both of them. On the right is the "N system", which is a special case of both the W and M systems.

A fundamental assumption of the ATO model is that assembly of products takes a negligible amount of time. Obviously, in this case it is most efficient to keep inventories at the component level and assemble products only for existing demands. This Assemble-to-Order (ATO) strategy has been widely deployed, especially by PC manufacturers such as Dell. Its popularity provides strong motivations for doing research on these systems, which generally splits along the lines of whether one is studying so-called ATO production-inventory systems, or ATO inventory systems. In ATO production-inventory systems, components are produced within the system, and inventory problems arise due to costs of and/or constraints on production capacities. In ATO inventory systems, components are ordered from outside the system, and inventory problems are caused by delays between ordering components and receiving them (lead times), which requires component ordering decisions to be made before the exact demand is known. Our survey focuses on the latter model, which, for simplicity, we refer to as ATO systems when there is no ambiguity.

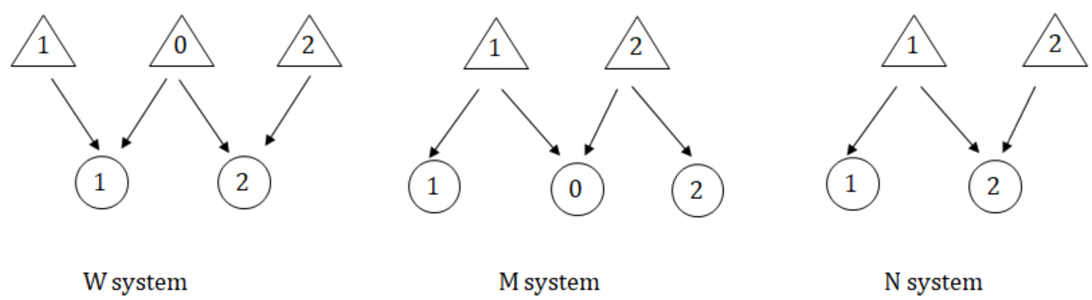

Figure 1 **Instances of ATO Systems**

Like many canonical inventory problems, inventory control in ATO systems needs to address a fundamental tradeoff between ordering more components to reduce backlog costs and ordering less to reduce inventory holding costs. When products are assembled from multiple components, ordering decisions need to be coordinated to minimize idle inventories of some components because of the lack of others, especially those with longer lead times. When different products share a common component, decisions on ordering components also need to take into account the way they are allocated. For instance, in the W system, if one product has higher priority for using the common component, then the product-specific component used by the other product should have its inventory position set at a lower level. On the other hand, decisions on allocating components also need to consider their replenishment process. For instance, in the M system, a manager may need to decide whether to use a unit of component 1 to serve demand for product 1 or reserve it to serve product 0 (if it is more valuable) after component 2 becomes available. The decision depends on the availability of both components, which is a result of replenishment decisions. Addressing these issues to optimize inventory control in ATO systems is often formulated using models of minimizing inventory costs. Among them, the backlog model with a general Bill of Materials (BOM) and deterministic lead times is a difficult and important problem. In such models, the state-space of the underlying control problem becomes high-dimensional, with the dimension growing both in the lead time and the size of the BOM. As before, this makes the large lead time regime interesting not just as in many applications it is at least moderately large, but also because it is exactly the setting of large lead times in which methods such as dynamic programming become intractable.

7.1.2. Literature review prior to recent progress using asymptotic analysis. Optimizing ATO systems for a single period gives rise to a two-stage stochastic linear program that can, in principle, be solved numerically (Song and Zipkin (2003)). However, developing an optimal control policy over an infinite time horizon is a daunting challenge. At the time when the survey of the field, Song and Zipkin (2003), was published, the optimal control policy was known only for systems with a single product (Rosling (1989)). The survey concludes that:

"As indicated above, little is known about the forms of optimal policies for multi-period models. The research to date mostly assumes particular policy types. It would be valuable to learn more about truly optimal policies. Even partial characterizations would be interesting. Also, better heuristic policy forms would be useful."

Many heuristic policies have been developed within the confines of specific assumptions, most notably, policies that assume the use of independent base stock policies for replenishment and FIFO for allocation (Hausman et al. (1998), Lu et al. (2003), Lu and Song (2005), Zhang (1997)). This stream of research is majorized in a sense by Lu and Song (2005), which formulates an explicit computational procedure for determining the optimal base stock levels. (It is worth noting that Jaarsveld and Scheller-Wolf (2015) point out the poor scaling of this computation as the number of components and products grow.) However, if a policy class itself is inherently inefficient, then finding the best parameters has limited effect on achieving optimality. There have been efforts to develop better types of policies. For instance, the "no-hold back" allocation policy, which outperforms FIFO in some cases (Lu et al. (2010)), but is also proven to be inadequate in other cases (Wan and Wang (2015)). Some studies on periodic-review systems (Aksay and Xu (2004), Deshpande (2005)) deviate from the assumption of FIFO allocation, but only for demands that arrive in the same period. These policies revert to FIFO in continuous-review systems where the length of the period diminishes to zero.

On the other hand, finding the exact optimal policies remains a formidable challenge. Beyond the single-product systems (Rosling (1989)), recent studies have only developed optimal policies for a few rare cases, in which the cost parameters take particular values that make the allocation decisions trivial (Dogru et al. (2010), Reiman and Wang (2012)). The proof of optimality depends on the latter factor, so the result cannot be extended to more general settings. This gives rise to the following question: **between exact optimality, which is beyond the reach of current methods, and heuristics that restrict attention to suboptimal policy forms, can we develop a new approach that, while not exactly optimal nonetheless performs well for many ATO systems?** An asymptotic analysis of ATO systems with large lead times provides such an approach.

It should be pointed out that, prior to the work using asymptotic analysis on ATO inventory systems that is described here, Plambeck and Ward (2006) provide an asymptotic analysis of ATO production-inventory systems. Their asymptotic analysis is carried out in the "high-volume" region where demand arrival rates are scaled up. The production capacities (along with prices) are determined in a one-time optimization that yields a system in heavy traffic. They prove asymptotic optimality of an allocation policy that involves solving a particular linear program periodically, where this period is chosen in a manner similar to that in Harrison's "Big-Step" method (Harrison (1996)).

7.1.3. Overview of recent progress using asymptotic analysis. Recently, substantial progress was made in the analysis of ATO systems using asymptotic analysis. The results in Dogru et al. (2010), Reiman and Wang (2012, 2015), Reiman et al. (2018), taken together show that if one fixes the BOM and cost vectors, and lets the lead times grow large, then a "target-tracking" control strategy is asymptotically optimal. The approach can be summarized as the following four-step procedure:

1. *Step 1:* Prove that one can lower-bound the cost rate in terms of a certain stochastic program which is a relaxation of the true problem.
2. *Step 2:* Develop methodologies to solve the stochastic program.
3. *Step 3:* Develop an inventory policy that makes replenishment and allocation decisions to follow targets for inventory positions and backlog levels that are set by solving the lower-bound stochastic program. (As the system evolves, these targets typically need to be updated by re-solving the stochastic program.)
4. *Step 4:* Prove that the policy can asymptotically (as the lead time grows large) keep the system "close" to this lower bound, by using arguments centered on weak convergence (by CLT-like arguments), continuity, and the "nice behavior" of the relevant limiting processes.

Here we note that related stochastic programs have been considered for certain one-period models going back at least to Gerchak and Henig (1986), and we refer the interested reader to Song and Zipkin (2003) for additional references. In contrast to the algorithms provided in Case studies I and II, here the algorithm is more sophisticated, and conceptually can be thought of as providing a "completely new algorithm" as opposed to "justifying" intuitive heuristics.

7.2. Formal model and problem statement.

In this section, we formally define the ATO inventory optimization problem. The inventory system has m products that are assembled from n components according to a BOM, with corresponding $n \times m$ matrix representation A, with each component a_{ji} equals to the amount of component j i needed to assemble a single unit of product i ($1 \leq i \leq m$, $1 \leq j \leq n$). Here we assume that all components have a common lead time L, but will later generalize to the setting in which different components may have different lead times. The problem can be formulated as either a periodic-review or a continuous-review model, and here we focus on the continuous-review model. We assume the following model for demand. Let $\{\mathbf{d}^k, k \geq 1\}$ be an i.i.d. sequence of m-dimensional vectors. Intuitively, d_i^k will represent the demand for product i which arrives during the kth "demand arrival event". Furthermore, demand arrival events occur at times dictated by a Poisson process with rate λ. We let $\{\mathcal{D}(t), t \geq -L\}$ denote the vector of demands that have arrived on $[-L, t]$. At each time t, the system evolves as follows.

- New demands arrive and components that were ordered exactly one lead time ago are received.
- Products are assembled from available components to serve demands (allocation) according to one's chosen control policy.
- New orders for components are placed (replenishment) according to one's chosen control policy.
- Inventory on-hand of components and backlog levels of products are updated.
- The cost rate is updated according to the new inventory on-hand and backlog levels.

Here we note that as the model is in continuous time, the above sequence of events must be interpreted through that lens, where e.g. the set of times at which demands or components actually arrive and/or are actually ordered will typically be measure zero. In addition, costs are incurred continuously through time, not just at event epochs. The cost rate stays constant between event epochs.

7.2.1. Formal definition of inventory control policy.

The goal is to select an inventory control policy π, which consists of both a replenishment and allocation policy. A feasible policy π will be equated with the valid definition of two processes, as follows.

- There is an n-dimensional replenishment process $\{\mathcal{R}^\pi(t), t \geq -L\}$, with $\mathcal{R}_j^\pi(t)$ (intuitively) equal to the number of units of the ith component ordered in $[-L, t]$.

- There is an m-dimensional allocation process $\{\mathcal{Z}^{\pi}(t), t \geq 0\}$, with $\mathcal{Z}_i^{\pi}(t)$ (intuitively) the number of units of demand of the ith product type met in $[0, t]$.

We now describe the set of admissible such policies $\pi = \left(\mathcal{R}^{\pi}, \mathcal{Z}^{\pi}\right)$. First it will be helpful to formalize two additional processes implicitly defined by a policy π, as follows.

- There is an n-dimensional on-hand inventory process $\{\mathbf{I}^{\pi}(t), t \geq 0\}$, with $I_i^{\pi}(t)$ (intuitively) the on-hand inventory level of the ith component at time t. For a given policy π and fixed initial conditions, the process $\mathbf{I}^{\pi}$ is uniquely defined through the following dynamical component inventory equation, which requires that for all $0 \leq t_1 \leq t_2$, $\mathbf{I}^{\pi}(t_2) = \mathbf{I}^{\pi}(t_1) + \mathbf{R}^{\pi}(t_1 - L, t_2 - L) - A \cdot \mathbf{Z}^{\pi}(t_1, t_2)$.
- There is an m-dimensional product backlog process $\{\mathbf{B}^{\pi}(t), t \geq 0\}$, with $B_i^{\pi}(t)$ (intuitively) the backlog for product i at time t. For a given policy π and fixed initial conditions, the process $\mathbf{B}^{\pi}$ is uniquely defined through the following dynamical product backlog equation, which requires that for all $0 \leq t_1 \leq t_2$, $\mathbf{B}^{\pi}(t_2) = \mathbf{B}^{\pi}(t_1) + \mathbf{D}(t_1, t_2) - \mathbf{Z}^{\pi}(t_1, t_2)$.

Then we say that a policy $\pi = \left(\mathcal{R}^{\pi}, \mathcal{Z}^{\pi}\right)$ is admissible if it satisfies the following properties.

- **Adaptivity**. Letting $\mathcal{F}$ denote the filtration generated by the demand process, we require that π is adapted to $\mathcal{F}$.
- **Non-negativity**. The intuitive definitions of all four processes require that $\mathcal{R}^{\pi}, \mathcal{Z}^{\pi}, \mathbf{I}^{\pi}, \mathbf{B}^{\pi}$ are all non-negative.
- **Monotonicity**. As over time the total number of components ordered and demand met is non-decreasing, we require that $\mathcal{R}^{\pi}$ and $\mathcal{Z}^{\pi}$ are non-decreasing.
- **Appropriate measurability, existence of limits, and continuity**. We also require that π is appropriately measurable, and that all relevant processes are r.c.l.l., although for clarity of exposition we do not dwell on these technical matters here.

We let Π denote the set of all admissible policies.

7.2.2. Formal problem statement. We assume per-unit linear inventory holding costs of $\mathbf{h}$ per unit of time (with h_i the holding cost of component i); and per-unit linear inventory backlogging costs of $\mathbf{b}$ per unit of time (with b_i the backlogging cost of product i). Thus under a given admissible policy π, at each time t,the inventory cost is incurred at the expected rate $C^{\pi}(t) \stackrel{\Delta}{=} \mathbf{h} \cdot \mathbb{E}[\mathbf{I}^{\pi}(t)] + \mathbf{b} \cdot \mathbb{E}[\mathbf{B}^{\pi}(t)]$. We then define the long-run-average cost of a policy π as $\mathcal{C}^{\pi} \stackrel{\Delta}{=} \limsup_{T \to \infty} \frac{1}{T} \int_0^T C^{\pi}(t) dt$, where we note that the appropriate measurability and continuity requirement of all admissible policies ensures that all relevant integrals are well-defined. In that case, the formal problem of interest is $\inf_{\pi \in \Pi} \mathcal{C}^{\pi}$.

7.3. Detailed overview of Dogru et al. (2010), Reiman and Wang (2012, 2015), Reiman et al. (2018).

We now review the results and approach of Dogru et al. (2010), Reiman and Wang (2012, 2015), Reiman et al. (2018).

7.3.1. Stochastic program lower bound. We begin by describing the relevant lower bound on the cost rate at a given time t, i.e. $C^{\pi}(t)$, from the perspective of time $t-L$. Here we derive the lower bound intuitively, without formal mathematical arguments, referring the reader to Dogru et al. (2010), Reiman and Wang (2012, 2015), Reiman et al. (2018) for related formal arguments. Since we only require a lower bound, it will be fine (at least intuitively) if we make the problem "easier" by "relaxing" certain features of the problem. The major relaxation we make is that we allow the inventory manager to defer all allocation and assembly in $[t-L,t]$ until time t, as assembling components into products to clear some desired backlog of demand (for various products) can be done instantaneously at that time, and at that time one then has the maximum amount of information and flexibility (as all demand to arrive in $[t-L,t]$, and components in the pipeline at time t, will have arrived). Of course, the system may incur heavy inventory holding and backlog costs before time t as no demand is served during the period $(t-L,t)$. However these costs do not affect the objective here, which is to minimize the cost rate at time t only. Under such a relaxation, just before the allocation at time t:

- The n-dimensional on-hand component inventory equals $\mathbf{I}^{\pi}(t-L)+\mathbf{R}^{\pi}(t-2L,t-L)$.
- The m-dimensional product backlog equals $\mathbf{B}^{\pi}(t-L)+\mathbf{D}(t-L,t)$.

Now, subject to the above, how small could we make $C^{\pi}(t)$ by "optimally" allocating the on-hand component inventory to the product backlog? Let us denote by $\mathbf{z}$ the m-dimensional vector such that z_i equals the amount of backlogged demand for product i met in the allocation at time t. Note that for any such $\mathbf{z}$; and fixed realizations of $\mathbf{I}^{\pi}(t-L),\mathbf{R}^{\pi}(t-2L,t-L),\mathbf{B}^{\pi}(t-L),\mathbf{D}(t-L,t)$; the rate at which the inventory cost is incurred at time t equals $\mathbf{h}\cdot\big(\mathbf{I}^{\pi}(t-L)+\mathbf{R}^{\pi}(t-2L,t-L)-A\mathbf{z}\big)+\mathbf{b}\cdot\big(\mathbf{B}^{\pi}(t-L)+\mathbf{D}(t-L,t)-\mathbf{z}\big)$. Denoting $\mathbf{I}^{\pi}(t-L)+\mathbf{R}^{\pi}(t-2L,t-L)$ by $\mathbf{y}$, and $\mathbf{B}^{\pi}(t-L)$ by $\boldsymbol{\alpha}$, we thus derive a lower bound for $C^{\pi}(t)$ of

$$\mathbb{E}\Big[\mathbf{h}\cdot\mathbf{y}-\sum_{j=1}^{n}h_j\sum_{i=1}^{m}a_{ji}z_i+\mathbf{b}\cdot\boldsymbol{\alpha}+\mathbf{b}\cdot\mathbf{D}(t-L,t)-\sum_{i=1}^{m}b_iz_i\Big].$$

Letting $\mathbf{c}$ denote the vector such that $c_i=\sum_{j=1}^{n}h_ja_{ji}+b_i$, we may rewrite this lower bound as

$$\mathbb{E}\Big[\mathbf{h}\cdot\mathbf{y}+\mathbf{b}\cdot\boldsymbol{\alpha}+\mathbf{b}\cdot\mathbf{D}(t-L,t)-\mathbf{c}\cdot\mathbf{z}\Big].$$

Now, we have already asserted that we get to observe $\mathbf{D}(t-L,t)$ before choosing $\mathbf{z}$, i.e. $\mathbf{z}$ can be a function of $\mathbf{D}(t-L,t)$. Let us now think on what constraints $\mathbf{z}$ must satisfy. First, it must hold that

$\mathbf{z} \leq \boldsymbol{\alpha} + \mathbf{D}(t-L,t)$, as one cannot clear non-existent backlog. Second, it must hold that $A\mathbf{z} \leq \mathbf{y}$, as one cannot use components one does not have. Imposing these constraints (and non-negativity of $\mathbf{z}$), for any fixed values of $\boldsymbol{\alpha}$ and $\mathbf{y}$ we derive the following lower bound for $C^\pi(t)$:

$$\mathbf{h} \cdot \mathbf{y} + \mathbf{b} \cdot \big(\boldsymbol{\alpha} + \mathbb{E}\big[\mathbf{D}(t-L,t)\big]\big) - \mathbb{E}\Big[\sup_{\mathbf{z} \geq \mathbf{0}} \big\{\mathbf{c} \cdot \mathbf{z} \ \mid \ \mathbf{z} \leq \boldsymbol{\alpha} + \mathbf{D}(t-L,t), A\mathbf{z} \leq \mathbf{y}\big\}\Big]. \tag{7}$$

We complete the lower bound by noting that we may "optimize" this lower bound over all non-negative choices for $\mathbf{y}$ and $\boldsymbol{\alpha}$, although note that we will do so in such a way that $\mathbf{y}$ and $\boldsymbol{\alpha}$ must be selected without knowledge of the realized value of $\mathbf{D}(t-L,t)$. Combining all of the above, and making rigorous with formal stochastic comparison arguments, the following is proven in Dogru et al. (2010). Let $\mathbf{D}$ be a random vector distributed as the lead time demand $\mathbf{D}(t-L,t)$.

Lemma 5. *For all $\pi \in \Pi$ and $t \geq L$,*

$$C^\pi(t) \geq \inf_{\boldsymbol{\alpha},\mathbf{y} \geq \mathbf{0}} \Big(\mathbf{h} \cdot \mathbf{y} + \mathbf{b} \cdot \big(\boldsymbol{\alpha} + \mathbb{E}\big[\mathbf{D}\big]\big) - \mathbb{E}\Big[\sup_{\mathbf{z} \geq \mathbf{0}} \big\{\mathbf{c} \cdot \mathbf{z} \ \mid \ \mathbf{z} \leq \boldsymbol{\alpha} + \mathbf{D}, A\mathbf{z} \leq \mathbf{y}\big\}\Big]\Big).$$

Furthermore, in Reiman and Wang (2015), it is shown that Lemma 5 can be stated in a more convenient form as follows, through some elementary substitutions and manipulations of the relevant optimizations.

Lemma 6. *For all $\pi \in \Pi$ and $t \geq L$,*

$$C^\pi(t) \ \geq \ \underline{C}^L \triangleq \inf_{\mathbf{y}} \Big(\mathbf{h} \cdot \mathbf{y} + (\mathbf{b} - \mathbf{c}) \cdot \mathbb{E}\big[\mathbf{D}\big] + \mathbb{E}\Big[\inf_{\mathbb{B} \geq \mathbf{0}} \big\{\mathbf{c} \cdot \mathbb{B} \ \mid \ A\mathbb{B} \geq A\mathbf{D} - \mathbf{y}\big\}\Big]\Big). \tag{8}$$

Note that the minimization in (8) is a two-stage stochastic linear program. Explicit computational procedures for solving this stochastic program as a discrete optimization problem have been developed for special cases, such as the W system (Dogru et al. (2010)) and the M system (Dogru et al. (2017)). Discrete convexity results, which can be used to develop efficient algorithms, have been proven for systems with particular BOM structures (see, e.g., Dogru et al. (2017), Zipkin (2016)). Moreover, DeValve et al. (2018) provides a thorough analysis on computational and integrality questions, showing that the optimal solutions for general cases (also with integrality constraints) can be efficiently computed by applying a so-called primal-dual methodology.

7.3.2. Development of asymptotically optimal target-based inventory policies. The average cost of an ATO system can reach the lower bound $\underline{C}^L$ if replenishment and allocation outcomes meet the optimal solutions of (8). More formally, let $\mathbf{y}^*$ denote an optimal choice of $\mathbf{y}$ in (8), which (intuitively) represents an "idealized" component inventory position vector. It is easy to keep the inventory position at $\mathbf{y}^*$: simply use a base stock policy with base stock levels $\mathbf{y}^*$. Let us also define the random vector $\mathbf{Q}^L(t) \triangleq A\mathbf{D}(t-L,t) - \mathbf{y}^*$, where (intuitively) $Q^L_j(t)$ represents the "shortage" of component j with regards to what would be needed to clear all demand which arrives in $(t-L,t]$, given access to the components on-hand and in the pipeline. Then the final term of (8) becomes

$$\mathbb{E}\big[\inf_{\mathbb{B}\geq\mathbf{0}}\big\{\mathbf{c}\cdot\mathbb{B}\ \mid\ A\mathbb{B}\geq\mathbf{Q}^L(t)\big\}\big] \tag{9}$$

Intuitively, (9) is an optimization over all backlog vectors $\mathbb{B}$ "consistent" with the realized component shortages at time t, $\mathbf{Q}^L(t)$, subject to the inventory position at time $t-L$ adhering to the optimal level $\mathbf{y}^*$. Let $\mathbb{B}^*$ denote the function that maps realizations of $\mathbf{Q}^L(t)$ to optimal solutions of the relevant linear program, i.e.

$$\mathbb{B}^*\big(\mathbf{Q}^L(t)\big)\in\operatorname*{arg\,min}_{\mathbb{B}\geq\mathbf{0}}\big\{\mathbf{c}\cdot\mathbb{B}\ \mid\ A\mathbb{B}\geq\mathbf{Q}^L(t)\big\}.$$

Existence and uniqueness of such a solution is addressed in Reiman and Wang (2015), Reiman et al. (2018), and for simplicity of discussion here let us simply assume existence and uniqueness. For a given policy π, let us define $\mathbf{y}^\pi(t)\triangleq\mathbf{I}^\pi(t-L)+\mathbf{R}^\pi(t-2L,t-L)-A\mathbf{B}^\pi(t-L)$. Then if it were possible for a policy π to satisfy $\mathbf{y}^\pi(t)=\mathbf{y}^*$, $\mathbf{B}^\pi(t)=\mathbb{B}^*\big(\mathbf{Q}^L(t)\big)$ w.p.1 (or even in expectation due to linearity) for all $t\geq L$, it would hold (as proven in Reiman and Wang (2015)) that π would be optimal, as it would match the lower bound $\underline{C}^L$. Unfortunately, in general such an exact match will be impossible, since an admissible policy cannot re-select the backlog levels at each time regardless of the state of the system.

To overcome this dilemma, the authors use a "target tracking" strategy. They reason as follows. Note that $\mathbf{Q}^L(t)$ is "observable" at time t. Thus at time t we may "plug in" this observed realized value for $\mathbf{Q}^L(t)$ into the linear program and compute the "idealized" backlog vector $\mathbb{B}^*\big(\mathbf{Q}^L(t)\big)$, i.e., the target. Supposing that: 1. $\mathbf{Q}^L(t)$ changes "sufficiently slowly" over time; 2. $\mathbb{B}^*\big(\mathbf{Q}^L(t)\big)$ is a "sufficiently continuous" function; and 3. we had been able to keep the actual backlog process $\mathbf{B}^\pi$ "sufficiently close" to the idealized process $\mathbb{B}^*\big(\mathbf{Q}^L(s)\big)$ for all s in an interval "sufficiently close" to t, the tracking strategy would suggest to allocate components in the interval "near" t in such a way that $\mathbf{B}^\pi$ continues to be "sufficiently close" to the backlog target, and that this should be achievable by "pushing the process" in an appropriate direction over a short time period. This logic is then repeated as $\mathbf{Q}^L$ evolves, ensuring that $\mathbf{B}^\pi$ stays "sufficiently close" to $\mathbb{B}^*\big(\mathbf{Q}^L(t)\big)$ for all $t\geq 0$. Note that, as indicated, the target $\mathbb{B}^*\big(\mathbf{Q}^L(t)\big)$ changes over time, and this target is obtained by re-solving a linear program when the right hand side ($\mathbf{Q}^L(t)$) changes. Although the context is different, there have been several successful uses of re-solving in the literature (Reiman and Wang (2008), Jasin et al. (2012), Bumpensanti and Wang (2018), Vera et al. (2019)).

This plan is carried out in Reiman and Wang (2015) to derive an asymptotically optimal policy. More precisely, Reiman and Wang (2015) formulates the following "allocation principle", which dictates how to "push the process" in the "right direction" based on the solution $\mathbb{B}^*\big(\mathbf{Q}^L(t)\big)$ to the re-solved linear program. This allocation principle can be roughly described as follows: **at each time t, do not serve a product if its current backlog level is not above the target (even if components are available), and serve all other products to maximally eliminate their excess over backlog targets (myopically), subject to component availability.** We refer the interested reader to Reiman and Wang (2015) for a complete specification of the sequence of asymptotically optimal policies $\{\pi_{*,L}, L > 0\}$. (There may be several allocation policies that satisfy the allocation principle.) We also note that the same allocation principle applies to ATO production-inventory systems, where components are produced instead of ordered (Plambeck and Ward (2006)), leading to a policy that is asymptotically optimal in the high demand volume regime, while also more efficient in all other cases (Wan and Wang (2015)).

7.3.3. Formal statement of asymptotic optimality and additional proof details. The formal asymptotic optimality result proven in Reiman and Wang (2015) is the following.

Theorem 6. $\lim_{L\to\infty} \frac{C^{\pi_{*},L}}{\underline{C}^L} = 1$, *with both* $C^{\pi_{*},L}$ *and* $\underline{C}^L$ *scaling as* $\Theta(\sqrt{L})$ *as* $L \to \infty$. *Thus, the sequence of policices* $\{\pi_{*,L}, L > 0\}$ *is asymptotically optimal as* $L \to \infty$.

The basic intuition behind the formal proof is related to the time and space scaling of the functional CLT that arises in this setting, where we note that such simultaneous time and space scalings are standard in the analysis of queueing systems in heavy traffic (cf. Iglehart and Whitt (1970), Whitt (2002)). Indeed, although for technical reasons Reiman and Wang (2015) do not formally use such a CLT (instead using various concentration / martingale / maximal inequalities), here we explain the relevant intuition through the lens of such CLT. In particular, let $\hat{\mathbf{B}}^{\pi_{*},L}(t) \triangleq \frac{\mathbf{B}^{\pi_{*},L}(Lt)}{\sqrt{L}}$, and $\hat{\mathbf{Q}}^L(t) \triangleq \frac{\mathbf{Q}^L(Lt)}{\sqrt{L}}$. Then the sequence of processes $\{\hat{\mathbf{Q}}^L(t), L > 0\}$ satisfies a functional CLT as $L \to \infty$. Thus this sequence of processes converges (as $L \to \infty$) to a "nice" limiting process $\hat{\mathbf{Q}}(t)$, namely Brownian motion, which has continuous sample paths. Classical results in the continuity of linear program (e.g. Hoffman's lemma) then imply that $\mathbb{B}^*\big(\hat{\mathbf{Q}}^L(t)\big)$ also has continuous sample paths (for large L). The intuition behind the proof that the actual backlog level is able to stay close to its target value is as follows. By the allocation principle, whenever the backlog of a product falls below its ideal level dictated by $\mathbb{B}^*\big(\mathbf{Q}^L(t)\big)$ none of that product is serviced, which induces a strong drift "back up" towards this ideal level. This corrects the shortfall so quickly that it is not asymptotically significant, leading to the resulting continuity on the correct asymptotic scale. Since a "conservation of mass" argument from Reiman and Wang (2015) implies that as long as no backlogs are substantially below their ideal levels, it must be that no backlogs are substantially above their ideal levels, we may conclude that the allocation principle (combined with the asymptotic balancing of supply and demand by basic limit theorems of probability) essentially keeps all backlog levels at their ideal levels at all times. More precisely, it is proven that

$$\lim_{L\to\infty}\sup_{t\geq 1}\left|\mathbb{E}\big[\hat{\mathbf{B}}^{\pi_*,L}(t)\big]-\mathbb{E}\big[\hat{\mathbb{B}}^*\big(\mathbf{Q}^L(t)\big)\big]\right|=\mathbf{0},$$

which by linearity of expectation, Lemma 6, and some additional arguments and bounds (the details of which we omit) is sufficient to ensure the desired asymptotic optimality and prove Theorem 6.

7.3.4. Systems with non-identical lead times. Obviously, inventory management becomes more complex for ATO systems when components can have different lead times. It turns out that in comparison with systems with identical lead times, the real difference is in component replenishment. Except for components with the longest lead times, it is not efficient to order them according to a base stock policy. While ideal inventory positions can be prescribed by an analogous SP that also sets a lower bound on the average cost, attaining these ideal inventory positions is typically infeasible. The good news is that this new obstacle does not make asymptotic optimality unreachable, it just takes a more complex policy to get there (Reiman et al. (2018)).

For simplicity of discussion, we will only consider systems with two distinct lead times, L_1 and L_2, where $L_2 > L_1$, and there are $n_1(n_2)$ components with lead time $L_1(L_2)$. Let $A^1(A^2)$ denote the BOM for those components with lead time $L_1(L_2)$. Also, let $\mathbf{D}^1 \triangleq \mathbf{D}(t-L_1,t)$, and $\mathbf{D}^2 \triangleq \mathbf{D}(t-L_2,t-L_1)$, where we note that $\mathbf{D}^1+\mathbf{D}^2=\mathbf{D}(t-L_2,t)$. Similarly, let $\mathbf{h}^k$ be the vector of component holding costs for those components with lead time L_k, $k\in\{1,2\}$, and let $\mathbf{c}$ denote the vector such that $c_i=\sum_{k=1}^{2}\sum_{j=1}^{n_k}h_j^k a_{ji}^k+b_i$. It will be evident from the discussion that not much effort is required to extend the analysis and results to systems with any finite number of distinct lead times. Instead of formally re-defining the associated problem, we simply note that the formal problem definition is essentially the same, with the caveat that now one associates a given policy π with two on-hand component inventory processes, one each for those components with lead time $L_1(L_2)$ respectively, with the associated defining dynamical inventory equations reflecting these different lead times (see Reiman et al. (2018) for details).

Similar to the development that leads to Lemma 6, Reiman and Wang (2012), Reiman et al. (2018) establish a lower bound on the expected cost at time t, which is also a lower bound on the average inventory cost of the ATO system.

LEMMA 7. *For all* $\pi \in \Pi$ *and* $t \geq L_2$,

$$
\begin{aligned}
C^{\pi}(t) \geq \quad & \underline{C}^{L_1,L_2} \triangleq \inf_{\mathbf{y}^2}\{\mathbf{h}^2 \cdot \mathbf{y}^2 + \mathbb{E}[\varphi^1(\mathbf{y}^2,\mathbf{D}^2)]\} + \mathbf{b}\cdot\mathbb{E}[\mathbf{D}^1+\mathbf{D}^2], \\
& \textit{with} \quad \varphi^1(\mathbf{y}^2,\mathbf{D}^2) \;=\; \inf_{\mathbf{y}^1}\{\mathbf{h}^1\cdot\mathbf{y}^1 + \mathbb{E}[\varphi^0(\mathbf{y}^1,\mathbf{y}^2,\mathbf{D}^1+\mathbf{D}^2)\mid \mathbf{D}^2]\}, \\
& \textit{and} \quad \varphi^0(\mathbf{y}^1,\mathbf{y}^2,\mathbf{D}^1+\mathbf{D}^2) \;= \\
& \quad \inf_{\mathbb{B}\geq\mathbf{0}}\{\mathbf{c}\cdot\mathbb{B} \mid A^2\mathbb{B} \geq A^2(\mathbf{D}^1+\mathbf{D}^2)-\mathbf{y}^2, A^1\mathbb{B} \geq A^1(\mathbf{D}^1+\mathbf{D}^2)-\mathbf{y}^1\} - \mathbf{c}\cdot\mathbb{E}[\mathbf{D}^1+\mathbf{D}^2].
\end{aligned}
$$

As in the identical lead time case, the average cost of an ATO system can reach the lower bound $\underline{C}^{L_1,L_2}$ if replenishment and allocation outcomes meet the optimal solutions of the optimization problem in Lemma 7. Also as in the identical lead time case, Reiman et al. (2018) again reformulate the relevant optimization problem in terms of certain observable shortages, and again develops policies to keep the inventory positions and backlog levels near their ideal target levels. The targets are updated by applying new demand inputs and re-solving the stochastic program of Lemma 7. A major additional complexity in the non-identical lead time setting is that even certain of the "ideal" component inventory positions are now random processess that vary over time. For example, in the three-stage stochastic program given in Lemma 7, although the optimal choice for $\mathbf{y}^2$ is simply a fixed vector, the optimal choice for $\mathbf{y}^1$ is now a complex function of past demands. In Reiman et al. (2018) a policy π_{*,L_1,L_2} is derived that keeps: 1. the inventory positions for components of lead time L_1 "sufficiently close" to the (state-dependent) targets of $\mathbf{y}^1$ given by the optimal solutions to the stochastic program, and are updated by re-solving the problem with new inputs as the system evolves; 2. the product backlogs "sufficiently close" to their ideal levels given by the same stochastic program through essentially the same allocation principle used in the identical lead time case; and 3. the component inventories for components of lead time L_2 equal to the fixed optimal value for $\mathbf{y}^2$ in the stochastic program by applying a base-stock strategy as in the identical lead time case. For systems with a single product, the policy reduces to an algorithm that is equivalent to that of Rosling (1989), and thus is exactly optimal. We also note that for systems where the different lead times are very close to one-another, the same level of asymptotic optimality can be achieved by following an independent base stock policy (Reiman et al. (2016)).

For a formal statement of asymptotic optimality in the non-identical lead time setting, we refer the reader to Reiman et al. (2018). Due to the additional complexity introduced by non-identical lead times, Reiman et al. (2018) introduce several additional methodological innovations. First, to show that all relevant quantities are "asymptotically close" to the idealized processes given by the re-solved stochastic programs, the authors formulate a general "stochastic tracking model" within which one can seamlessly demonstrate all relevant convergences (as opposed to requiring separate notations, tools, and analyses for each required convergence). Second, proving continuity of the relevant maps (especially those involved with the ideal levels for components with shorter lead times) is much more involved, as the dimension of the relevant linear programs grows with the lead time, where the authors overcome this by more explicitly using the special structure of the relevant programs. Also, a perturbation analysis is used to ensure uniqueness of the relevant optimal solutions. Here we also note that in contrast to the stochastic programs arising in the identical lead time case, in the setting of non-identical lead times these programs have an inherently nested structure. As such, the question of the computational complexity of (approximately) solving these programs remains poorly understood, and we refer the reader to Reiman et al. (2018) for related discussion.

8. Conclusion and directions for future research

In this survey, we: 1. reviewed several fundamental high-dimensional inventory models which exhibit the curse of dimensionality; 2. reviewed different methodologies for approaching these problems; and 3. provided a detailed exposition of how asymptotic analysis has significantly advanced our understanding of these models. We also provided an in-depth introduction to the relevant tools and methodologies through three case studies in which asymptotic analysis has recently led to major progress: lost sales models, dual-sourcing models, and assemble-to-order systems in the presence of large lead times. Our survey suggests several interesting directions for future research.

- **Tackle other models**. Section 2 contains many challenging inventory control models which have resisted solution to date, and which make natural candidates for trying to apply asymptotic analysis. Indeed, the very recent work Bu et al. (2019) on perishable inventory models, Xin (2019) on lost sales inventory models (in which a capped base-stock policy is proven to exhibit uniformly good performance in a certain asymptotic regime), and Stolyar and Wang (2018) on inventory models with random lead times, suggest that asymptotic analysis has much more to say here. This also includes understanding at what level of model generality, e.g. as regards allowing for more general cost functions, joint distribution of demand, number of suppliers and products, lead time distribution, time-dependent behavior, more realistic constraints on ordering and production, etc. such asymptotic analysis can be applied. Here we note that developing a better understanding of the general tracking model defined in Reiman et al. (2018), and more broadly defining other general frameworks and sets of assumptions under which such asymptotic analyses hold, may be a good first step along these lines.

- **Address computational issues and complexity**. As noted in Section 7.3.1, even in some cases where we know particular policies to be asymptotically optimal, we do not fully understand the computational complexity of implementing the relevant policies. This relates to the fact that a formal theory of computational complexity for the structured stochastic dynamic programs which arise in inventory control remains incomplete (as discussed in Section 3.1), although we refer the interested reader to Halman et al. (2009, 2014), Halman and Nannicini (2019), and more generally Dyer and Stougie (2006), Shmoys and Swamy (2006), Papadimitriou et al. (1987), Sidford et al. (2018) for relevant work on the complexity of stochastic dynamic programming. For the setting of ATO systems, DeValve et al. (2018) have also made recent progress along these lines, where we note that relevant questions such as how often one must "re-solve" certain approximating optimization problems in online optimization is a question currently of high interest across multiple academic communities (see e.g. Vera et al. (2019), Bumpensanti and Wang (2018)).

- **Turn asymptotic results into efficient algorithms for all parameters**. Related to the above computational issues is the fact that even in settings where certain policies are known to be asymptotically optimal, it remains poorly understood how to efficiently "bridge" such asymptotic results with policies which are effective outside of those regimes, yielding implementable algorithms which are practically efficient across a broad range of parameters. This may also involve developing a finer understanding of the error in asymptotic approximations, and how one could develop more sophisticated algorithms with stronger optimality guarantees (see Xin and Goldberg (2016) for further relevant discussion). For recent progress along such lines, see Xin (2019), DeValve et al. (2018).

- **Unify asymptotic analysis with robust optimization and machine learning**. There has been considerable recent effort within the academic inventory control community to account for model uncertainty, through robust optimization, machine learning, and statistical approaches (see e.g. Xin et al. (2015), Zhang et al. (2018b), Agrawal and Jia (2019), Sun and Van Mieghem (2019), Ban and Rudin (2018), Gijbrechts et al. (2018)), yet a unified understanding of how asymptotic analysis in inventory control interacts with such methodologies remains an interesting open question. Although some results exist along these lines (see e.g. Huh et al. (2009), and to some extent Zhang et al. (2018b), Chen et al. (2019), Agrawal and Jia (2019) and related work), our understanding of the extent to which insights from asymptotic analysis can be modified and extended to handle model uncertainty and learning remains incomplete.

- **Make it relevant to practice**. The gap between the practice of inventory control in industry and the theoretical study of inventory models in academia remains large, where we note that this gap has existed for much of the history of the field of inventory theory (Silver (1981), Kumar et al. (2013)). Although it remains a long-standing challenge (across many academic disciplines) to convert theoretical insights into practical results, the applied nature of the operations research discipline is certainly amenable to such "tech transfer" and we point to Xin et al. (2017) for a recent effort to convert insights from asymptotic analysis in inventory theory into practical results. Such an endeavor will also likely involve the creation of new models (possibly in which inventory control is integrated with other features) arising from novel applications, and we point to Govindarajan et al. (2017) for a recent application of asymptotic methods to a new model for inventory control in omnichannel retailing, and to Banerjee et al. (2017) for a recent application to a model of pricing and inventory control in ride-sharing.

Acknowledgements

D. A. Goldberg gratefully acknowledges support from the National Science Foundation [Grant1757394]. M. I. Reiman and Q. Wang also would like to thank support from the National Science Foundation [Grant1363314]. All authors thank the many students, collaborators, and colleagues with whom they have collaborated on inventory-related research over the years.

References

Aggarwal, Sumer C. 1974. A review of current inventory theory and its applications. *International Journal of Production Research*, 12(4), 443-482.

Agrawal, Narendra, and Morris A. Cohen. 2001. Optimal material control in an assembly system with component commonality. *Naval Research Logistics (NRL)*, 48 (5), 409-429.

Agrawal, Shipra, and Randy Jia. 2019. Learning in structured MDPs with convex cost functions: Improved regret bounds for inventory management. arXiv preprint arXiv:1905.04337.

Ahuja, Ravindra K., Wei Huang, H. Edwin Romeijn, and Dolores Romero Morales. 2007. A heuristic approach to the multi-period single-sourcing problem with production and inventory capacities and perishability constraints. *INFORMS Journal on Computing*, 19(1), 14-26.

Akcay, Yalcin, and Susan H. Xu. 2004. Joint inventory replenishment and component allocation optimization in an Assemble-to-Order system. *Management Science*, 50(1), 99-116.

Aksoy, Yasemin, and S. Selcuk Erenguc. 1988. Multi-item inventory models with co-ordinated replenishments: a survey. *International Journal of Operations and Production Management*, 8(1), 63-73.

Allon, Gad, and Jan A. Van Mieghem. 2010. Global dual sourcing: Tailored base-surge allocation to near-and offshore production. *Management Science*, 56(1), 110-124.

Ang, Marcus, Karl Sigman, Jing-Sheng Song, and Hanqin Zhang. 2017. Closed-form approximations for optimal (r, q) and (S, T) policies in a parallel processing environment. *Operations Research*, 65(5), 1414-1428.

Angelus, A., Ö. Özer, 2015. Knowledge you can act on: optimal policies for assembly systems with expediting and advance demand information. *Operations Research*, 64(6), 1338-1371.

Anily, Shoshana, and Julien Bramel. 2004. An asymptotic 98.5%-effective lower bound on fixed partition policies for the inventory-routing problem. *Discrete Applied Mathematics*, 145(1), 22-39.

Arapostathis, A., V. S. Borkar, E. Fernández-Gaucherand, M. K. Ghosh, S. I. Marcus. 1993. Discrete time controlled Markov processes with average cost criterion: a survey. *SIAM Journal on Control and Optimization*, 31(2), 282-344.

Arrow, Kenneth J., Theodore Harris, and Jacob Marschak. 1951. Optimal inventory policy. *Econometrica: Journal of the Econometric Society*, 250-272.

Arrow, Kenneth J., Samuel Karlin, and Herbert Scarf. 1958. The nature and structure of inventory problems. Studies in the mathematical theory of inventory and production. Stanford University Press, Stanford, Ca (1958).

Asmussen, S. 2003. *Applied probability and queues*. Second ed., Springer.

Atan, Zumbul, Taher Ahmadi, Clara Stegehuis, Ton de Kok, and Ivo Adan. 2017. Assemble-to-Order systems: A review. *European Journal of Operational Research*, 261(3), 866-879.

Bakker, Monique, Jan Riezebos, and Ruud H. Teunter. 2012. Review of inventory systems with deterioration since 2001. *European Journal of Operational Research*, 221(2), 275-284.

Ban, Gah-Yi, and Cynthia Rudin. 2018. The big data newsvendor: Practical insights from machine learning. *Operations Research*, 67(1), 90-108.

Banerjee, Siddhartha, Daniel Freund, and Thodoris Lykouris. 2017. Pricing and Optimization in Shared Vehicle Systems: An Approximation Framework. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, 517-517.

Barankin, E. W. 1961. A delivery-lag inventory model with an emergency provision. *Naval Research Logistics Quarterly*, 8(3), 285-311.

Baykal-Gursoy, Melike, and Nesim K. Erkip. 2010. Forecasting for inventory planning under correlated demand. *Wiley Encyclopedia of Operations Research and Management Science.*

Bellman, Richard. 1954. The theory of dynamic programming. *Bulletin of the American Mathematical Society*, 60(6), 503-515.

Bellman, Richard, Irving Glicksberg, and Oliver Gross. 1955. On the optimal inventory equation. *Management Science*, 2(1), 83-104.

Bellman R.E. 1961. *Adaptive Control Processes*. Princeton University Press, Princeton, NJ.

Bertsimas, Dimitris, and Aurlie Thiele. 2006. A robust optimization approach to inventory theory. *Operations Research*, 54(1), 150-168.

Bertsimas, Dimitris, Dan A. Iancu, and Pablo A. Parrilo. 2010. Optimality of affine policies in multistage robust optimization. *Mathematics of Operations Research*, 35(2), 363-394.

Bijvank, Marco, and Iris FA Vis. 2011. Lost-sales inventory theory: A review. *European Journal of Operational Research*, 215(1),1-13.

Billingsley, P. 1999. *Convergence of Probability Measures*. Wiley, New York.

Boute, R.N., J.A. Van Mieghem. 2015. Global dual sourcing and order smoothing: The impact of capacity and lead times. *Management Science*, 61(9), 2080-2099.

Bradley, James R., and Peter W. Glynn. 2002. Managing capacity and inventory jointly in manufacturing systems. *Management Science*, 48(2), 273-288.

Bradley, James R., and Lawrence W. Robinson. 2005. Improved base-stock approximations for independent stochastic lead times with order crossover. *Manufacturing and Service Operations Management*, 7(4), 319-329.

Brahimi, Nadjib, Nabil Absi, Stphane Dauzre-Prs, and Atle Nordli. 2017. Single-item dynamic lot-sizing problems: An updated survey. *European Journal of Operational Research*, 263(3), 838-863.

Bramson, M. J. 1962. The Variable lead time Problem in Inventory Control: A Survey of the Literature-Part 1. *Journal of the Operational Research Society*, 13(1), 41-53.

Bramson, Maury. 1998. State space collapse with application to heavy traffic limits for multiclass queueing networks. *Queueing Systems*, 30(1-2), 89-140.

Brown, David B., and James E. Smith. 2014 Information relaxations, duality, and convex stochastic dynamic programs. *Operations Research*, 62(6), 1394-1415.

Bu, Jinzhi, Xiting Gong, and Dacheng Yao. 2017. Constant-order Policies for Lost-sales Inventory Models with Random Supply Functions: Asymptotics and Heuristic. Available at SSRN 3063730.

Bu, Jinzhi, Xiting Gong, and Xiuli Chao. 2019. Asymptotic Optimality of Base-stock Policies for Perishable Inventory Systems. Preprint.

Bumpensanti, Pornpawee, and He Wang. 2018. A Re-solving Heuristic with Uniformly Bounded Loss for Network Revenue Management. arXiv preprint arXiv:1802.06192.

Burgin, T. A. "The gamma distribution and inventory control." Journal of the Operational Research Society 26, no. 3 (1975): 507-525.

Bushuev, Maxim A., Alfred Guiffrida, M. Y. Jaber, and Mehmood Khan. 2015. A review of inventory lot sizing review papers. *Management Research Review*, 38(3), 283-298.

Chan, Lap Mui Ann, Awi Federgruen, and David Simchi-Levi. 1998. Probabilistic analyses and practical algorithms for inventory-routing models. *Operations Research*, 46(1), 96-106.

Chan, Lap MA, ZJ Max Shen, David Simchi-Levi, and Julie L. Swann. 2004. Coordination of pricing and inventory decisions: A survey and classification. In *Handbook of Quantitative Supply Chain Analysis*, 335-392. Springer, Boston, MA, 2004.

Chao, Xiuli, Xiting Gong, Cong Shi, and Huanan Zhang. 2015. Approximation algorithms for perishable inventory systems. *Operations Research*, 63(3), 585-601.

Chao, Xiuli, Xiting Gong, Cong Shi, Chaolin Yang, Huanan Zhang, and Sean X. Zhou. 2017. Approximation algorithms for capacitated perishable inventory systems with positive lead times. *Management Science*, 64(11), 5038-5061.

Chen, Fangruo, and Jing-Sheng Song. 2001. Optimal policies for multiechelon inventory problems with Markov-modulated demand. *Operations Research*, 49(2), 226-234.

Chen, Xin, and David Simchi-Levi. 2012. Pricing and inventory management. *The Oxford handbook of pricing management*, 784-824.

Chen, W., M. Dawande, G. Janakiraman. 2014. Fixed-Dimensional Stochastic Dynamic Programs: An Approximation Scheme and an Inventory Application. *Operations Research.* 62(1), 81-103.

Chen, Y., W. Xue, J. Yang. 2013. Optimal inventory policy in the presence of a long-term supplier and a spot market. *Operations Research.* 61(1), 88-97.

Chen, Li, Jing-Sheng Song, and Yue Zhang. 2017. Serial inventory systems with markov-modulated demand: Derivative bounds, asymptotic analysis, and insights. *Operations Research*, 65(5), 1231-1249.

Chen, Xin. 2017. $L^\natural$-convexity and its applications in operations. *Frontiers of Engineering Management.*

Chen, Xin, Alexander Stolyar, and Linwei Xin. 2019. Asymptotic optimality of constant-order policies in joint pricing and inventory control models. Available at SSRN.

Chen, Beryl and Cong Shi. 2019. Tailored Base-Surge Policies in Dual-Sourcing Inventory Systems with Demand Learning. Preprint.

Clark, Andrew J. 1972. An informal survey of multiechelon inventory theory. *Naval Research Logistics Quarterly*, 19(4), 621-650.

Coelho, Leandro C., Jean-Franois Cordeau, and Gilbert Laporte. 2013. Thirty years of inventory routing. *Transportation Science*, 48(1), 1-19.

Cudina, Milica, and Kavita Ramanan. 2011. Asymptotically optimal controls for time-inhomogeneous networks. *SIAM Journal on Control and Optimization*, 49(2), 611-645.

de Kok, Ton, Christopher Grob, Marco Laumanns, Stefan Minner, Jrg Rambau, and Konrad Schade. 2018. A typology and literature review on stochastic multi-echelon inventory models. *European Journal of Operational Research*, 269(3), 955-983.

2011. Dai, Jim, and Tolga Tezcan. State space collapse in many-server diffusion limits of parallel server systems. *Mathematics of Operations Research*, 36(2), 271-320.

Daniel, K. H. 1963. A delivery-lag inventory model with emergency order. H. Scarf, D. Gilford, M. Shelly, eds. *Multistage Inventory Models and Techniques*, chapter 2. Stanford University Press, Stanford, CA.

DeCroix, G., J.S. Song, P. Zipkin. 2005. A series system with returns: Stationary analysis. *Operations Research*, 53(2), 350-362.

Deshpande, Vinayak, Morris A. Cohen, and Karen Donohue.2005. A threshold inventory rationing policy for service-differentiated demand classes. *Management Science*, 49(6), 683-703.

Deuschel, Jean-Dominique, and Daniel Stroock. 1989. Large deviations. Academic Press, Boston

DeValve, Levi, Aleksandar Pekec, and Yehua Wei. 2018. A primal-dual approach to analyzing ATO systems. Available at SSRN 299143.

Dogru, Mustafa K., Martin I. Reiman, and Qiong Wang. 2010. A stochastic programming based inventory policy for Assemble-to-Order systems with application to the W model. *Operations research*, 58(4-part-1), 849-864.

Dogru, Mustafa K., Martin I. Reiman, and Qiong Wang. 2017. AssembletoOrder Inventory Management via Stochastic Programming: Chained BOMs and the MSystem. *Production and Operations Management*, 26(3), 446-468.

Down, D., and S. Meyn. 1994. A survey of Markovian methods for stability of networks. In *11th International Conference on Analysis and Optimization of Systems Discrete Event Systems*, 490-504, Springer, Berlin, Heidelberg.

Downs, B., R. Metters, S. Semple. 2001. Managing inventory with multiple products, lags in delivery, resource constraints and lost sales: A mathematical programming approach. *Management Science*, 47(3), 464479.

Duong, L., Lincoln Wood, and W. Wang. 2018. A review and reflection on inventory management of perishable products in a single-echelon model. *International Journal of Operational Research*, 31(3), 313-329.

Dvoretzky, Aryeh, Jack Kiefer, and Jacob Wolfowitz. 1952. The inventory problem: I. Case of known distributions of demand. *Econometrica (pre-1986)*, 20(2), 187.

Dvoretzky, Aryeh, Jack Kiefer, and Jacob Wolfowitz. 1953. On the optimal character of the (s, S) policy in inventory theory. *Econometrica: Journal of the Econometric Society*, 21(4), 586-596.

Dyer, Martin, and Leen Stougie. 2006. Computational complexity of stochastic programming problems. *Mathematical Programming*, 106(3), 423-432.

Dzielinski, Bernard P., and Ralph E. Gomory. 1965. Optimal programming of lot sizes, inventory and labor allocations. *Management Science*, 11(9), 874-890.

Elmaghraby, Wedad, and Pinar Keskinocak. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science*, 49(10), 1287-1309.

Eppen, Gary D. 1979. Note : effects of centralization on expected costs in a multi-location newsboy problem. *Management Science*, 25(5), 498-501.

Erdos, P., M. Kac. 1946. On certain limit theorems of the theory of probability. *Bull. Amer. Math. Soc*, 52(4), 292-302.

Federgruen, Awi. 1993. Centralized planning models for multi-echelon inventory systems under uncertainty. *Handbooks in Operations Research and Management Science 4*, 133-173.

Federgruen, Awi, and Michal Tzur. 1994. The joint replenishment problem with time-varying costs and demands: Efficient, asymptotic and ϵ-optimal solutions. *Operations Research*, 42(6), 1067-1086.

Feinberg, E.A. 2011. Total expected discounted reward MDPs: existence of optimal policies. *Wiley Encyclopedia of Operations Research and Management Science*.

Feng, Q., S. Sethi, H. Yan, H. Zhang. 2006. Are base-stock policies optimal in inventory problems with multiple delivery modes? *Operations Research*, 54(4), 801-807.

Fiestras-Janeiro, M. Gloria, Ignacio Garca-Jurado, Ana Meca, and Manuel A. Mosquera. 2011. Cooperative game theory and inventory management. *European Journal of Operational Research*, 210(3), 459-466.

Filar, J. 2007. *Controlled Markov chains, graphs and Hamiltonicity*. Now Publishers Inc.

Folland, G. B. 1999. *Real Analysis: Modern Techniques and Their Applications*. Second Edition, Wiley.

Fox, E.J., R. Metters, J. Semple. 2006. Optimal inventory policy with two suppliers. *Operations Research*, 54(2), 389-393.

Fleischmann, M., R. Kuik. 2003. An optimal inventory control with independent stochastic item returns. *European Journal of Operational Research*, 151, 25-37.

Fu, Ke, Xiting Gong, and Guitian Liang. 2019. Managing Perishable Inventory Systems with Product Returns and Remanufacturing. *Production and Operations Management*, 28(6), 1366-1386.

Fukuda, Y. 1964. Optimal policies for the inventory problem with negotiable lead time. *Management Science*. 10(4), 690-708.

Gabrel, Virginie, Ccile Murat, and Aurlie Thiele. 2014. Recent advances in robust optimization: An overview. *European Journal of Operational Research*, 235(3), 471-483.

Gerchak, Yigal, and Mordechai Henig. 1986, An inventory model with component commonality. *Operations Research Letters*, 5(3), 157-160.

Gijsbrechts, Joren, Robert N. Boute, Jan A. Van Mieghem, and Dennis Zhang. 2018. Can deep reinforcement learning improve inventory management? performance and implementation of dual sourcing-mode problems. *Performance and Implementation of Dual Sourcing-Mode Problems*, December 17, 2018.

Glasserman, Paul. 1996. Allocating production capacity among multiple products. *Operations Research*, 44(5), 724-734.

Glasserman, Paul. 1997. Bounds and asymptotics for planning critical safety stocks. *Operations Research*, 45(2), 244-257.

Glasserman, Paul, and Tai-Wen Liu. Corrected diffusion approximations for a multistage production-inventory system. *Mathematics of Operations Research*, 22(1), 186-201.

Glynn, Peter W. 1990. Diffusion approximations. *Handbooks in Operations Research and Management Science 2*, 145-198.

Goldberg, David A., Dmitriy A. Katz-Rogozhnikov, Yingdong Lu, Mayank Sharma, and Mark S. Squillante. 2016. Asymptotic optimality of constant-order policies for lost sales inventory models with large lead times. *Mathematics of Operations Research*, 41(3), 898-913.

Gong, X., X. Chao, S. Zheng. 2014. Dynamic pricing and inventory management with dual suppliers of different lead times and disruption risks. *Production and Operations Management*. forthcoming.

Govindarajan, Aravind, Amitabh Sinha, and Joline Uichanco. 2017. Inventory optimization for fulfillment integration in omnichannel retailing. SSRN Electron. J.

Hajek, B. 1983. The proof of a folk theorem on queueing delay with applications to routing in networks. *Journal of the ACM*, 30(4), 834-851.

Halman, Nir, Diego Klabjan, Mohamed Mostagir, Jim Orlin, and David Simchi-Levi. 2009. A fully polynomial-time approximation scheme for single-item stochastic inventory control with discrete demand. *Mathematics of Operations Research*, 34(3), 674-685.

Halman, N., Diego Klabjan, Chung-Lun Li, Jim Orlin and David Simchi-Levi. 2014. Fully Polynomial Time Approximation Schemes for Stochastic Dynamic Programming, *SIAM J. Discrete Math.*, 28, 1725-1796.

Halman, N., and Giacomo Nannicini. 2019. Toward Breaking The Curse of Dimensionality: An FPTAS for Stochastic Dynamic Programs With Multidimensional Action and Scalar State, *SIAM Journal on Optimization*, 29, 1131-1163.

Harris, F. 1913. How many parts to make at once. *Factory, The Magazine of Management*, 10(2), 135-136.

Harrison, J. Michael. 1988. Brownian models of queueing networks with heterogeneous customer populations. In *Stochastic Differential Systems, Stochastic Control Theory and Applications*, 147-186. Springer, New York, NY.

Harrison, J. Michael, and Lawrence M. Wein. 1990. Scheduling networks of queues: Heavy traffic analysis of a two-station closed network. *Operations research*, 38(6), 1052-1064.

Harrison, J. Michael, and Vien Nguyen. 1993. Brownian models of multiclass queueing networks: Current status and open problems. *Queueing Systems*, 13(1-3), 5-40.

Harrison, J. Michael. 1996. The BIGSTEP approach to flow management in stochastic processing networks. *Stochastic Networks: Theory and Applications*, 4, 147-186.

Harrison, J. Michael, and Jan A. Van Mieghem. 1997. Dynamic control of Brownian networks: state space collapse and equivalent workload formulations. *The Annals of Applied Probability*, 7(3), 747-771.

Harrison, J. Michael, and Marcel J. Lopez. 1999. Heavy traffic resource pooling in parallelserver systems. *Queueing Systems*, 33(4), 339-368.

Harrison, J. Michael, and Jan A. Van Mieghem. 1999. Multi-resource investment strategies: Operational hedging under demand uncertainty. *European Journal of Operational Research*, 113(1), 17-29.

Hausman, W. H., H. L. Lee, and A. X. Zhang. 1998. Order response time reliability in a multi-item inventory system. *European Journal of Operational Research*, 109(3), 646-659.

Heyman, D., M. Sobel. 1984. *Stochastic Models in Operations Research*, Vol. II. McGraw-Hill, New York.

Hordijk, A., H. Tijms. 1974. Convergence Results and Approximations for Optimal (s, S) Policies. *Management Science*, 20(11), 1432-1438.

Hordijk, Arie, and Henk Tijms. 1975. On a conjecture of Iglehart. *Management Science*, 21 (11), 1342-1345.

Hua, Z., Y. Yu, W. Zhang, X. Xu. 2014. Structural properties of the optimal policy for dual-sourcing systems with general lead times. *IIE Transactions*. forthcoming.

Huang, Kai, and Ton de Kok. 2015. Optimal FCFS allocation rules for periodicreview assembletoorder systems. *Naval Research Logistics (NRL)*, 62(2), 158-169.

Huggins, E. L., T. L. Olsen. 2010. Inventory control with generalized expediting. *Operations Research*, 58(5), 1414-1426.

Huh, Woonghee Tim, Ganesh Janakiraman, John A. Muckstadt, and Paat Rusmevichientong. 2009. Asymptotic optimality of order-up-to policies in lost sales inventory systems. *Management Science*, 55(3), 404-420.

Huh, Woonghee Tim, and Paat Rusmevichientong. 2009. A nonparametric asymptotic analysis of inventory planning with censored demand. *Mathematics of Operations Research*, 34(1), 103-123.

Huh, W.T., G. Janakiraman. 2010. On the optimal policy structure in serial inventory systems with lost sales. *Operations Research*, 58(2), 486-491.

Huh, W.T., G. Janakiraman, M. Nagarajan. 2011. Average cost single-stage inventory models: an analysis using a vanishing discount approach. *Operations Research*, 59(1), 143-155.

Huh, Woonghee Tim, Ganesh Janakiraman, and Mahesh Nagarajan. 2016. Capacitated multiechelon inventory systems: Policies and bounds. *Manufacturing and Service Operations Management*, 18(4), 570-584.

Humblet, P. A. 1982. Determinism minimizes waiting time in queues. Technical report. *MIT Laboratory for Information and Decision Systems*. LIDS-P-1207.

Iglehart, D. L. 1963. Optimality of (s, S) policies in the infinite horizon dynamic inventory problem. *Management Science*, 9(2), 259-267.

Iglehart, Donald L., and Ward Whitt. 1970. Multiple channel queues in heavy traffic. I. *Advances in Applied Probability*, 2(1), 150-177.

van Jaarsveld, Willem, and Alan Scheller-Wolf. 2015. Optimization of industrial-scale assemble-to-order systems. *INFORMS Journal on Computing*, 27(3), 544-560.

Janakiraman, G., R. O. Roundy. 2004. Lost-sales problems with stochastic lead times: Convexity results for base-stock policies. *Operations Research*, 52(5), 795803.

Janakiraman, Ganesh, Sridhar Seshadri, and J. George Shanthikumar. 2007 A comparison of the optimal costs of two canonical inventory systems. *Operations Research*, 55(5), 866-875.

Janakiraman, G., S. Seshadri, A. Sheopuri. 2015. Analysis of Tailored Base-Surge Policies in Dual Sourcing Inventory Systems. *Management Science*, 61(7), 1547-1561.

Janssen, F., T. De Kok. 1999. A two-supplier inventory model. *International Journal of Production Economics*, 59, 395-403.

Jasin, Stefanus, and Sunil Kumar. 2012. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research*, 37(2), 313-345.

Jasin, Stefanus, and Amitabh Sinha. 2015. An LP-based correlated rounding scheme for multi-item ecommerce order fulfillment. Operations Research 63, no. 6 (2015): 1336-1351.

Kaminsky, Philip, and Onur Kaya. 2008. Scheduling and duedate quotation in a maketoorder supply chain. *Naval Research Logistics (NRL)*, 55(5), 444-458.

Karlin, S. and H. Scarf. 1958. Inventory models of the Arrow-Harris-Marschak type with time lag. *Studies in the Mathematical Theory of Inventory and Production.*

Karush, W. 1957. A queuing model for an inventory problem. Operation Research, 5(5), 693703.

Kelleher, K. 2003. Why FedEx is gaining ground. *Business 2.0* (October) 56-57.

Kennedy, W. J., J. Wayne Patterson, and Lawrence D. Fredendall. 2002. An overview of recent literature on spare parts inventories. *International Journal of Production Economics*, 76(2), 201-215.

Kingman, J.F.C. 1962. Some inequalities for the queue GI/G/1. *Biometrika*, 49, 315-324.

Kleijn, Marcel J., and Rommert Dekker. 1999. An overview of inventory systems with several demand classes. In *New Trends in Distribution Logistics*, 253-265. Springer, Berlin, Heidelberg.

Klosterhalfen, S., G. Kiesmüller. S. Minner. 2011. A comparison of the constant-order and dual-index policy for dual sourcing. *International Journal of Production Economics*, 133(1) 302-311.

Korolev, V. Yu, and I. G. Shevtsova. 2010. On the upper bound for the absolute constant in the Berry-Esseen inequality. *Theory of Probability and Its Applications*, 54(4) 638-658.

Krichagina, Elena V., Sheldon XC Lou, Suresh P. Sethi, and Michael I. Taksar. 1993. Production control in a failure-prone manufacturing system: diffusion approximation and asymptotic optimality. *The Annals of Applied Probability*, 3(2), 421-453.

Krishnamoorthy, A., B. Lakshmy, and R. Manikandan. 2011. A survey on inventory models with positive service time. *Operations Research*, 48(2), 153-169.

Kumar, Navneet, Abhaya Mishra, and Sajeev Abraham George. 2013. Gaps between inventory management theory and practice: a critical examination of emerging trends from the FMCG industry. *International Journal of Logistics Economics and Globalisation*, 5(1), 1-14.

Lemoine, Austin J. 1978. State of the ArtNetworks of Queues: A Survey of Weak Convergence Results. *Management Science*, 24(11), 1175-1193.

Levi, Retsef, Martin Pl, Robin O. Roundy, and David B. Shmoys. 2007. Approximation algorithms for stochastic inventory control models. *Mathematics of Operations Research*, 32(2), 284-302.

Levi, Retsef, Ganesh Janakiraman, and Mahesh Nagarajan. 2008. A 2-approximation algorithm for stochastic inventory control models with lost sales. *Mathematics of Operations Research*, 33(2), 351-374.

Levi, Retsef. 2010. Provably near-optimal approximation algorithms for operations management models. In *Risk and Optimization in an Uncertain World*, 179-192.

Levi, Retsef, Georgia Perakis, and Joline Uichanco. 2015. The data-driven newsvendor problem: new bounds and insights. *Operations Research*, 63(6), 1294-1306.

Lindley, David V. 1952. The theory of queues with a single server. In *Mathematical Proceedings of the Cambridge Philosophical Society*, 48(2), 277-289, Cambridge University Press.

Lu, Yingdong, Jing-Sheng Song, and David D. Yao. 2003. Order fill rate, leadtime variability, and advance demand information in an assemble-to-order system. *Operations Research*, 51(2), 292-308.

Lu, Yingdong, and Jing-Sheng Song. 2005. Order-based cost optimization in assemble-to-order systems. *Operations Research*, 53(1), 151-169.

Lu, Yingdong, Jing-Sheng Song, and Yao Zhao. 2010. No-holdback allocation rules for continuous-time Assemble-to-Order systems. *Operations Research*, 58(3), 691-705.

Lu, Lijian, JingSheng Song, and Hanqin Zhang. 2015. Optimal and asymptotically optimal policies for assembletoorder nand Wsystems. *Naval Research Logistics (NRL)*, 62(8), 617-645.

Mangasarian, Olvi L., and T-H. Shiau. 1987. Lipschitz continuity of solutions of linear inequalities, programs and complementarity problems. *SIAM Journal on Control and Optimization*, 25(3), 583-595.

Minner, Stefan. 2003. Multiple-supplier inventory models in supply chain management: A review. *International Journal of Production Economics*, 81, 265-279.

Morton, T. E. 1969. Bounds on the solution of the lagged optimal inventory equation with no demand backlogging and proportional costs. *SIAM Review*, 11(4), 572-596.

Dror, Moshe, and Bruce C. Hartman. 2011. Survey of cooperative inventory games and extensions. *Journal of the Operational Research Society*, 62(4), 565-580.

Muller, Alfred, and Dietrich Stoyan. 2002. *Comparison Methods for Stochastic Models and Risks*, vol. 389, New York: Wiley.

Nadar, Emre, Mustafa Akan, and Alan Scheller-Wolf. 2014. Optimal structural results for Assemble-to-Order generalized M-systems. *Operations Research*, 62(3), 571-579.

Nahmias, Steven. 1982. Perishable inventory theory: A review. *Operations research*, 30(4), 680-708.

Neuts, F. 1964. An inventory model with optimal time lag. *SIAM Journal on Applied Mathematics*, 12, 179-185.

Pang, Guodong, Rishi Talreja, and Ward Whitt. 2007. Martingale proofs of many-server heavy-traffic limits for Markovian queues. *Probability Surveys*, 4, 193-267.

Papadimitriou, Christos H. 2003. *Computational Complexity*. John Wiley and Sons Ltd.

Papadimitriou, Christos H., and John N. Tsitsiklis. 1987. The complexity of Markov decision processes. *Mathematics of Operations Research*, 12(3), 441-450.

Petruzzi, Nicholas C., and Maqbool Dada. 1999. Pricing and the newsvendor problem: A review with extensions. *Operations research*, 47(2), 183-194.

Plambeck, Erica L., and Amy R. Ward. 2006. Optimal control of a high-volume assemble-to-order system. *Mathematics of Operations Research*, 31(3), 453-477.

Plambeck, Erica L. 2008. Asymptotically optimal control for an assemble-to-order system with capacitated component production and fixed transport costs. *Operations Research*, 56(5), 1158-1171.

Porteus, Evan L. 1990. Stochastic inventory theory. *Handbooks in Operations Research and Management Science 2*, 605-652.

Porteus, Evan L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press.

Qin, Yan, Ruoxuan Wang, Asoo J. Vakharia, Yuwen Chen, and Michelle MH Seref. 2011. The newsvendor problem: Review and directions for future research. *European Journal of Operational Research*, 213(2), 361-374.

Raafat, Fred. 1991. Survey of literature on continuously deteriorating inventory models. *Journal of the Operational Research Society*, 42(1), 27-37.

Rao, U., A. Scheller-Wolf, S. Tayur. 2000. Development of a rapidresponse supply chain at Caterpillar. *Operations Research*, 48(2), 189-204.

Reiman, Martin I. 1984. Some diffusion approximations with state space collapse. In *Modelling and Performance Evaluation Methodology*, 207-240, Springer, Berlin, Heidelberg.

Reiman, Martin I. 2004. A new and simple policy for the continuous review lost sales inventory model. *Unpublished manuscript.*

Reiman, M. I., Q. Wang. 2008. An Asymptotically Optimal Policy for a Quantity-Based Network Revenue Management Problem. *Mathematics of Operations Research*, 33(2), 257-282.

Reiman, Martin I., and Qiong Wang. 2012. A stochastic program based lower bound for Assemble-to-Order inventory systems. *Operations Research Letters*, 40(2), 89-95.

Reiman, Martin I., and Qiong Wang. 2015. Asymptotically optimal inventory control for Assemble-to-Order systems with identical lead times. *Operations Research*, 63(3), 716-732.

Reiman, Martin I., Haohua Wan, and Qiong Wang. 2016. On the use of independent base-stock policies in Assemble-to-Order inventory systems with nonidentical lead times. *Operations Research Letters*, 44(4), 436-442.

Reiman, Martin I., Haohua Wan, and Qiong Wang. 2018. Asymptotically Optimal Inventory Control for Assemble-to-Order Systems. arXiv preprint arXiv:1809.08271.

Roberts, Donald Merle. 1959. Approximations to optimal policies in a dynamic inventory model. PhD dissertation, Department of Statistics, Stanford University.

Romeijn, H. Edwin, and Dolores Romero Morales. 2003. An asymptotically optimal greedy heuristic for the multiperiod singlesourcing problem: The cyclic case. *Naval Research Logistics (NRL)*, 50(5), 412-437.

Rong, Ying, Zmbl Atan, and Lawrence V. Snyder. 2017. Heuristics for BaseStock Levels in MultiEchelon Distribution Networks. *Production and Operations Management*, 26(9), 1760-1777.

Rosenshine, M., D. Obee. 1976. Analysis of a standing order inventory system with emergency orders. *Operations Research*, 24(6), 1143-1155.

Rosling, Kaj. 1989. Optimal inventory policies for assembly systems under random demands. *Operations Research*, 37(4), 565-579.

Ross, Nathan. "Fundamentals of Steins method." Probability Surveys 8 (2011): 210-293.

Rossi, R., W. A. Rijpkema, J. G. A. J. van der Vorst. 2012. The impact of dual sourcing on food supply chain networks: the case of Egyptian strawberries. In *Proc. of the 11th Wageningen International Conference on Chain and Network Management*, Wageningen.

Roundy, Robin. 1985. 98%-effective integer-ratio lot-sizing for one-warehouse multi-retailer systems. *Management Science*, 31(11), 1416-1430.

Scarf, H. 1960. The Optimality of (s, S) Policies in the Dynamic Inventory Problem. *Mathematical Methods in the Social Sciences*. Stanford University Press. 196-202.

Schäl, M. 1993. Average optimality in dynamic programming with general state space. *Mathematics of Operations Research*, 18(1), 163-172.

Scheller-Wolf, A., S. Veeraraghavan, G.-J. van Houtum. 2008. Effective Dual Sourcing with a Single Index Policy. Working paper.

Schrijver, Alexander. 1998. *Theory of Linear and Integer Programming*. John Wiley and Sons.

Schweitzer, Paul J., Giuseppe Serazzi, and Marco Broglia. 1993. A survey of bottleneck analysis in closed networks of queues. In *Performance Evaluation of Computer and Communication Systems*, 491-508. Springer, Berlin, Heidelberg.

Sennott, L.I., 1989. Average cost optimal stationary policies in infinite state MDP with unbounded costs. *Operations Research*, 37(4), 626-633.

Sethi, S.,M. Taksar, and Q. Zhang. 1992. Capacity and production decisions in stochastic manufacturing systems: An asymptotic optimal hierarchical approach. *Production and Operations Management*, 1(4), 367-392.

Shanthikumar, J. George, Shengwei Ding, and Mike Tao Zhang. 2007. Queueing theory for semiconductor manufacturing systems: a survey and open problems. *IEEE Transactions on Automation Science and Engineering*, 4(4), 513-522.

Shapiro, Alexander, Darinka Dentcheva, and Andrzej Ruszczyski. 2009. Lectures on stochastic programming: modeling and theory. *Society for Industrial and Applied Mathematics.*

Shen, Z. 2007. Integrated supply chain design models: a survey and future research directions. *Journal of Industrial and Management Optimization*, 3(1), 1-27.

Zheng, Yu-Sheng. 1992. On properties of stochastic inventory systems. *Management Science*, 38(1), 87-103.

Sheopuri, A., G. Janakiraman, S. Seshadri. 2010. New policies for the stochastic inventory control problem with two supply sources. *Operations Research*, 58(3), 734-745.

Shi, Cong. 2011. Approximation Algorithms for Stochastic Optimization Problems in Operations Management. *Wiley Encyclopedia of Operations Research and Management Science*, 1-20.

Shi, Cong, Weidong Chen, and Izak Duenyas. 2016. Nonparametric data-driven algorithms for multiproduct inventory systems with censored demand. *Operations Research*, 64(2), 362-370.

Shi, Cong, Yehua Wei, and Yuan Zhong. 2018. Process flexibility for multi-period production systems. Available at SSRN 2655790.

Shmoys, David B., and Chaitanya Swamy. 2006. An approximation scheme for stochastic linear programming and its application to stochastic integer programs. *Journal of the ACM (JACM)*, 53(6), 978-1012.

A. Sidford, M. Wang, C. Wu, L. Yang, Y. Ye. 2018. Near-Optimal Time and Sample Complexities for Solving Markov Decision Processes with a Generative Model. NeurIPS.

Silver, Edward A. 1981. Operations research in inventory management: A review and critique. *Operations Research*, 29(4), 628-645.

Simchi-Levi, David, and Yehua Wei. 2012. Understanding the performance of the long chain and sparse designs in process flexibility. *Operations research*, 60(5), 1125-1141.

Simchi-Levi, David, and Yao Zhao. 2012. Performance evaluation of stochastic multi-echelon inventory systems: A survey. *Advances in Operations Research.*

Simchi-Levi, David, Xin Chen, and Julien Bramel. 2005. The logic of Logistics: Algorithms, and Applications for Logistics and Supply Chain Management Theory.

Song, Jing-Sheng, and Paul Zipkin. 2003. Supply chain operations: Assemble-to-Order systems. *Handbooks in Operations Research and Management Science 11*, 561-596.

Song, J., P. Zipkin. 2009. Inventories with Multiple Supply Sources and Networks of Queues with Overflow Bypasses. *Management Science*, 55(3), 362-372.

Song, Jing-Sheng, Li Xiao, Hanqin Zhang, and Paul Zipkin. 2016. Optimal policies for a dual-sourcing inventory problem with endogenous stochastic lead times. *Operations Research*, 65(2), 379-395.

Spitzer, Frank. 1956. A combinatorial lemma and its application to probability theory. *Trans. Amer. Math. Soc.*, 82, 323-339.

Stolyar, Alexander L. 2004. Maxweight scheduling in a generalized switch: State space collapse and workload minimization in heavy traffic. *The Annals of Applied Probability*, 14(1), 1-53.

Stolyar, Alexander, and Qiong Wang. 2018. Exploiting random lead times for significant inventory cost savings. arXiv preprint arXiv:1801.02646.

Sun, Jiankun, and Jan A. Van Mieghem. 2019. Robust dual sourcing inventory management: optimality of capped dual index policies and smoothing. *Manufacturing and Service Operations Management.*

Syntetos, Aris A., John E. Boylan, and Stephen M. Disney. 2009. Forecasting for inventory planning: a 50-year review. *Journal of the Operational Research Society*, 60 (sup1), S149-S160.

Talluri, Kalyan, and Garrett Van Ryzin. 1998. An analysis of bid-price controls for network revenue management. *Management Science*, 44(11-part-1), 1577-1593.

Thorisson, H. 1992. Construction of a stationary regenerative process. *Stochastic Processes and Their Applications*, 42(2), 237-253.

Urban, Timothy L. 2005. Inventory models with inventory-level-dependent demand: A comprehensive review and unifying theory. *European Journal of Operational Research*, 162(3), 792-804.

Mieghem, Jan A. Van, and Nils Rudi. 2002. Newsvendor networks: Inventory management and capacity investment with discretionary activities. *Manufacturing and Service Operations Management*, 4(4), 313-335.

Van Mieghem, J. A. 2008. *Operations Strategy: Principles and Practice*. Dynamic Ideas, Belmont, MA.

Veeraraghavan, S., A. Scheller-Wolf. 2008. Now or later: Dual index policies for capacitated dual sourcing systems. *Operations Research*, 56(4), 850-864.

Veinott, Jr, Arthur F. 1966. On the Opimality of (s,S) Inventory Policies: New Conditions and a New Proof." *SIAM Journal on Applied Mathematics*, 14(5), 1067-1083.

Veinott Jr, Arthur F. 1966. The status of mathematical inventory theory. *Management Science*, 12(11), 745-777.

Vera, Alberto, and Siddhartha Banerjee. 2019. The bayesian prophet: A low-regret framework for online decision making." In Abstracts of *The 2019 SIGMETRICS/Performance Joint International Conference on Measurement and Modeling of Computer Systems*, 81-82.

Wan, Haohua, and Qiong Wang. 2015. Asymptotically-optimal component allocation for Assemble-to-Order productioninventory systems." *Operations Research Letters* 43(3), 304-310.

Ward, Amy R. Asymptotic analysis of queueing systems with reneging: A survey of results for FIFO, single class models." *Surveys in Operations Research and Management Science*, 17, no. 1 (2012): 1-14.

Wei, Lai, Stefanus Jasin, and Linwei Xin. Asymptotic Optimality of Order-Up-To Control for Stochastic Inventory Systems with Sequential Probabilistic Service Level Constraints." Available at SSRN 3114348 (2018).

Wein, Lawrence M. Dynamic Scheduling of a Multiclass Make-to-Stock Queue.” *Operations Research*, no. 4 (1992):724-735.

Whitin, Thomson M. Inventory control research: A survey.” *Management Science* 1, no. 1 (1954): 32-40.

Whitt, Ward. Heavy traffic limit theorems for queues: a survey. In Mathematical Methods in Queueing Theory, pp. 307-350. Springer, Berlin, Heidelberg, 1974.

Whitt, Ward. Stochastic-process limits: an introduction to stochastic-process limits and their application to queues. Springer Science and Business Media, 2002.

Whittmore, A. S., S. C. Saunders. 1977. Optimal inventory under stochastic demand with two supply options. *SIAM Journal on Applied Mathematics*. **32** 293-305.

Williams, R. J. Some recent developments for queueing networks. In Probability Towards 2000, pp. 340-356. Springer, New York, NY, 1998.

Williams, Brent D., and Travis Tokar. A review of inventory management research in major logistics journals: Themes and future directions. The International Journal of Logistics Management 19, no. 2 (2008): 212-232.

Winands, Erik MM, Ivo JBF Adan, and Geert-Jan van Houtum. The stochastic economic lot scheduling problem: A survey. European Journal of Operational Research 210, no. 1 (2011): 1-9.

Xin, Linwei, and David A. Goldberg. Distributionally robust inventory control when demand is a martingale. arXiv preprint arXiv:1511.09437 (2015).

Xin, Linwei, and David A. Goldberg. Optimality gap of constant-order policies decays exponentially in the lead time for lost sales models. Operations Research 64, no. 6 (2016): 1556-1565.

Xin, Linwei, and David A. Goldberg. Asymptotic optimality of tailored base-surge policies in dual-sourcing inventory systems. Management Science 64, no. 1 (2017): 437-452.

Xin, Linwei, Long He, Jagtej Bewli, John Bowman, Huijun Feng, and Zhiwei Qin. On the performance of Tailored Base-Surge policies: theory and application at Walmart. com.” Com (December 18, 2017) (2017).

Xin, Linwei. Understanding the Performance of Capped Base-Stock Policies in Lost-Sales Inventory Models. Available at SSRN (2019).

Xin, Linwei. A 1.79-approximation algorithm for a continuous review lost-sales inventory model. Preprint.

Yao, Man, and Stefan Minner. 2017. Review of multi-supplier inventory models in supply chain management: An update.

Yaspan, A. 1961. An inclusive solution to the inventory problem. *Operations Research*, 9, 371-382.

Zhang, Alex X. 1997. Demand Fulfillment Rates in an Assembel-to-Order System with Multiple Products and Dependent Demands. *Production and Operations Management*, 6(3), 309-324.

Zhang, Can, Turgay Ayer, and Chelsea C. White III. 2-Approximation Algorithms for Perishable Inventory Control When FIFO Is an Optimal Issuing Policy." arXiv preprint arXiv:1605.01442 (2016).

Zhang, Hailun, Jiheng Zhang, and Rachel Zhang. 2018. Managing Perishable Inventory Systems as Non-Perishable Ones.

Zhang, Huanan, Xiuli Chao, and Cong Shi. 2018. Closing the gap: A learning algorithm for the lost-sales inventory system with lead times.

Zhou, S.X., Z. Tao, X. Chao. 2011. Optimal control of inventory systems with multiple types of remanufacturable products. *Manufacturing and Service Operations Management*, 13(1), 20-34.

Zipkin, Paul H. 1986. Stochastic lead times in continuoustime inventory models. *Naval Research Logistics Quarterly*, 33(4), 763-774.

Zipkin, Paul H. 2000. *Foundations of inventory management*

Zipkin, Paul H. 2008. Old and new methods for lost-sales inventory systems. *Operations Research*, 56(5), 1256-1263.

Zipkin, Paul H. 2008. On the structure of lost-sales inventory models. *Operations Research*, 56(4), 937-944.

Zipkin, Paul H. Some specially structured Assemble-to-Order systems.2017. *Operations Research Letters*, 44(1), 136-142.