



Interfaces

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Incorporating Stochastic Lead Times Into the Guaranteed Service Model of Safety Stock Optimization

Salal Humair, John D. Ruark, Brian Tomlin, Sean P. Willems

To cite this article:

Salal Humair, John D. Ruark, Brian Tomlin, Sean P. Willems (2013) Incorporating Stochastic Lead Times Into the Guaranteed Service Model of Safety Stock Optimization. *Interfaces* 43(5):421-434. <http://dx.doi.org/10.1287/inte.2013.0699>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2013, INFORMS

Please scroll down for article—it is on subsequent pages



INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Incorporating Stochastic Lead Times Into the Guaranteed Service Model of Safety Stock Optimization

Salal Humair

Harvard School of Public Health, Boston, Massachusetts 02115; and School of Science and Engineering, Lahore University of Management Sciences, D.H.A. Lahore 54792, Pakistan, shumair@hsph.harvard.edu

John D. Ruark

Logility, Inc., Burlington, Massachusetts 01803, jruark@logility.com

Brian Tomlin

Tuck School of Business, Dartmouth College, Hanover, New Hampshire 03755, brian.tomlin@tuck.dartmouth.edu

Sean P. Willems

School of Management, Boston University, Boston, Massachusetts 02215, willems@bu.edu

Effective end-to-end supply chain management and network inventory optimization must account for service levels, demand volatility, lead times, and lead-time variability. Most inventory models incorporate demand variability, but far fewer rigorously account for lead-time variability, particularly in multiechelon supply chain networks. Our research extends the guaranteed service model of safety stock placement to allow random lead times. The main methodological contribution is the creation of closed-form equations for the expected safety stock in the system; this includes a derivation for the early-arrival stock in the system. The main applied contributions are the demonstration of real stochastic lead times in practice and a discussion of how our approach outperforms more traditional heuristics that either ignore lead-time variability or consider the maximum lead time at each stage.

Key words: base-stock policy; multiechelon inventory optimization; stochastic lead times; supply chain application.

In recent years, supply chain executives have had to grapple with the increased globalization of their internal supply chains and those of their customers and suppliers, while also navigating the heightened volatility in supply and demand brought about by significant economic uncertainty. The effective use of inventory, both strategically and tactically, is crucial for success in this global and volatile world. The companies that can best manage the trade-off between customer service and inventory are the ones best placed to compete and win. Not surprisingly, inventory optimization is a top priority for executives as evidenced by a 2012 survey of 153 senior operations executives conducted by Supply Chain Management World. This report (SCMWorld 2012, p. 11) states:

We asked the supply chain community to identify the specific areas which they will be targeting for improvement this coming year and, with the statistics on

demand and supply volatility in mind, it should come as little shock that inventory optimisation stands tall at the very top with a substantial 64 percent selection rate. Not only is this ambition number one overall, it is also well up from last year. As a general remedy to not only volatile demand and supply, but also increasing complexity, inventory optimisation promises lean operations with high levels of customer service.

This finding echoes that of other surveys, including a 2011 report by Chief Supply Chain Officer (CSCO) Insights (CSCO 2011) in which 76 percent of respondents indicated that inventory management excellence was either their top supply chain priority or a highly important one. To truly manage the inventory and service trade-off in globally dispersed supply chains, companies must take an end-to-end perspective when determining where to locate their inventories and how much inventory to hold. Interestingly,

this CSCO report found, on page 10, that the largest barrier to “more effective cross network inventory management” was that firms “can’t optimize their network holistically” (i.e., take an end-to-end perspective), with 83.0 percent of respondents reporting this inability as a medium or high barrier. In addition to technological and organizational challenges, 73.5 percent and 69.5 percent of respondents also identified demand and supply volatility, respectively, as important barriers.

Although network inventory optimization has been the subject of academic research for more than 50 years (see Simpson 1958, Clark and Scarf 1960), only in the past decade have concerted efforts been made to transfer this multiechelon inventory research to widespread real-world practice. With provable optimality a common objective in academic research, simplification of reality is a hallmark of academic models. Unfortunately, this simplification can be a roadblock to industry adoption. Companies are not typically concerned with whether a network inventory solution gives a provably optimal policy in a setting they perceive to be an oversimplification of the reality they face on a daily basis. They are concerned with whether the solution can effectively accommodate enough of their key pressure points so that it reflects their reality with sufficient fidelity. Only then does a company adopt the solution and implement its recommendations. In our experience with implementing network inventory solutions at numerous companies, lead-time variability is one such pressure point.

All supply chains face lead-time variability. However, it manifests itself differently in different parts of the supply chain. For raw material locations, lead-time variability occurs in big chunks of time. It is not that lead time is 30 days plus or minus five days. It is that lead time is 30 days 60 percent of the time, 60 days 30 percent of the time, and 120 days 10 percent of the time. For finished goods distribution, the lead-time variability might be two days plus or minus a day; however, at this point in the supply chain, the product is at its most expensive and buffering the variability in time is extremely costly.

In this paper, we extend the guaranteed service (GS) model for safety stock optimization to incorporate random lead times in multiechelon networks. This

work on lead-time variability is critical to the widespread adoption of the GS approach—an overall approach that has led to the following documented benefits: savings of \$100 million by Proctor & Gamble, \$50 million by HP, and \$20 million by Kraft Foods; a 26 percent inventory reduction by Boston Scientific; and a 25 percent finished goods inventory reduction by Black & Decker.

The rest of the paper is organized as follows: We first summarize the GS model for safety stock placement in the case of deterministic lead times. We then present our modeling approach for stochastic lead times, and then a numerical analysis where we compare our approach to known heuristics. We also discuss real-world information systems constraints and model implementation issues. We relegate mathematical expressions and derivations to the appendix.

Modeling Framework for Deterministic Lead Times (GS-DET)

As Graves and Willems (2003) define, approaches to optimize safety stock in networks can be classified into one of two frameworks depending on how they model service between stages in the supply chain. Following their nomenclature, we chose to adopt and extend the GS modeling framework. This framework is a pragmatic modeling approach that has been successfully deployed across a wide range of industries to optimize safety stock quantities and locations in complex real-world networks under the assumption that lead times are deterministic, as Willems (2008), Humair and Willems (2011), and Farasyn et al. (2011) discuss. In this project, we generalize the GS framework to allow for random lead times in the network.

Before describing the GS model with random lead times, it is helpful to review the most salient features of the GS model with deterministic lead times (hereafter GS-DET). The reader can also find a more detailed overview of GS-DET in Graves and Willems (2000). In the GS-DET model, a supply chain is modeled as a network of stages, where each stage represents an item at a processing function (e.g., procurement, transportation, production, assembly) and a potential inventory stocking location for the item processed at the stage. All stages are assumed to use

a periodic-review base-stock policy with an identical review period. This common review period serves as the underlying time unit in the model. At the start of a period, each stage observes its demand and places an order to replenish that demand. By convention, if demand is zero in a period, then an order of zero units is placed. All stages are assumed to have deterministic lead times. As an example, imagine that a stage represents an assembly step in which it assembles two components—each from a different upstream stage—and that the assembly stage has a lead time of L periods. Then, a processing order placed by the assembly stage is completed (i.e., reaches the stocking location for that stage) L periods after the order has been placed, assuming that the two upstream stages have adequate inventory so that the assembly stage can begin its processing immediately. Note that the GS framework assumes that a stage single sources any particular item; therefore, if multiple upstream stages feed a stage, these stages supply different components.

The notion of quoted service times is central to the GS framework. Instead of a stage guaranteeing that it will provide off-the-shelf service to its direct customers (i.e., have adequate inventory to fill an order immediately), a stage quotes an outbound service time $S \geq 0$ and guarantees that it will fill an order in exactly S periods. Thus, two service times are associated with each stage. The outbound service time for a stage, described previously, is the delivery time that a stage guarantees to its customers (i.e., those stages that are immediately downstream and adjacent to the stage). The inbound service time at a stage is the maximum delivery time a stage is quoted by its upstream and adjacent supplying stages. The GS approach assumes that if a stage stocks out, it will take extraordinary measures (e.g., expediting) to meet its guaranteed outbound service time. Determining the percentage of time a stage should have before it resorts to extraordinary measures is a managerial decision. This is captured by a user-defined service level, and safety stocks are set to ensure the outbound service time is met from inventory with a probability equal to that service level. An alternative interpretation of the guarantee is that managers specify a demand bound and agree that demand within the bound is met from inventory, but demand that

exceeds the bound is met by extraordinary measures. These two interpretations are somewhat analogous; however, we adopt the service level interpretation, because we find that practitioners are more comfortable with the service-level framing.

The GS-DET model determines the amount of safety stock required to achieve the quoted service time guarantee. If end-customer demands are Normally distributed and independent across periods (which we assume for ease of exposition in this paper), then for a given set of service times in the network, the GS-DET model sets the safety stock at a stage using Equation (1) in the appendix. This expression is reminiscent of the classical square root expression commonly used in practice; however, this classical expression is modified to account for the inbound and outbound service times. Instead of the safety stock being driven by the lead time, it is driven by the net-replenishment lead time, which is lead time plus inbound service time less outbound service time; the intuition is that safety stock needs only to protect against demand uncertainty over the net-replenishment lead time and not the actual lead time. We refer the reader to the appendix for details.

Users can deploy the GS-DET model to determine the required safety stocks for a set of user-defined service times. They can also allow the model to minimize the total network inventory cost by having the model optimize the service times in the network. When using the cost-optimization feature, a user must also specify per-unit inventory holding costs at each stage.

Inventory Optimization Under Random Lead Times (GS-RAN)

Although companies have applied the GS-DET framework in many settings, a key barrier to wider adoption was its assumption of deterministic lead times. Many users were hesitant to deploy a modeling paradigm that did not capture a key reality—that they experienced variable lead times at certain processing stages. To facilitate wider adoption, we had to generalize the GS-DET framework to allow for random lead times at any stage in the user-defined network. Next, we describe and develop the GS model with random lead times (hereafter GS-RAN).

The GS-RAN model adopts the same modeling paradigm as we describe previously for the GS-DET model; however, it allows any stage to have a random lead time. This is a very significant extension to the underlying framework and required two key advancements: (1) the development of an appropriate safety stock expression for the case of random lead times in a network, and (2) the ability to optimize service times to minimize inventory cost across a network with random lead times.

Safety Stock

We now develop the safety stock expression for any arbitrary stage in the network for any given set of network service times. We assume that replenishment orders do not cross in time. That is, earlier orders are always received before later orders. Later in this paper, we comment on how that assumption might be relaxed.

Consider the inventory level at some stage at some arbitrary period. The inventory level equals the starting inventory that is given by the base stock level, plus the cumulative replenishments received up to and including that period, less the cumulative demand filled up to and including that period. We define the shortfall as the difference between what the stage has shipped out and what it has replenished. Unless extraordinary measures are taken, a stockout will occur if the shortfall exceeds the base stock level. Because demands and lead times are random, the shortfall itself is a random variable. Intuitively, safety stock, which inflates the base stock level, is a buffer that accommodates this randomness in the shortfall. As long as the randomness does not result in an excessively large shortfall, the safety stock can fully absorb the shock and prevent a stock out. Characterizing the shortfall expression and its associated statistics is central to developing a safety stock expression.

Let us first consider the case where the inbound and outbound service times are both zero. As we show in the appendix, the shortfall is the cumulative demand over the most recently realized lead time, and standard approaches can be used to derive the shortfall statistics. If one models the shortfall as having an approximately normal distribution, then the safety stock at a stage is given by Equation (2) in the appendix. The reader might recognize this as the

commonly prescribed single-stage safety stock expression when lead times are random. Many classic operations management textbooks, such as Nahmias (1997), Silver et al. (1998), and Hopp and Spearman (2008), include this expression or close variants of it.

In the case of a multistage network, we cannot simply adopt the commonly used single-stage safety stock prescription, because it implicitly assumes that a stage quotes a zero outbound service time and experiences a zero inbound service time. In a network, a stage may quote a positive outbound service time and (or) experience a positive inbound service time. Therefore, we have to develop a general safety stock equation for random lead times when service times can be positive. The appendix shows the mathematical development of the shortfall expression, its associated statistics, and an appropriate safety stock expression. The result is the safety stock expression given by Equation (5). This generalized safety stock expression in Equation (5) is structurally similar to the well-known single-stage expression given in Equation (2). The only difference is that the mean and variance of the lead time are replaced with what are essentially the mean and variance of the net-replenishment lead-time random variable, conditional on its being positive. That is, the single-stage expression is modified by replacing the lead-time statistics with the equivalent statistics for the positive part of the net-replenishment lead-time random variable.

When the inbound and outbound service times are both zero, the general safety stock expression given in Equation (5) is identical to that in Equation (2). That is, we obtain the widely used single-stage safety stock expression as a special case of our general expression. This is beneficial from a user-adoption perspective because (1) we can relate the general expression to an expression familiar to practitioners, and (2) the generalization has an intuitive explanation: the lead-time statistics in the familiar expression are simply replaced by the equivalent statistics for the positive part of the net-replenishment lead time. For the deterministic lead-time model, GS-DET, we see earlier that the general safety stock expression is identical to the classic single-stage expression, but with the net-replenishment lead time replacing the lead time. By analogy, for the case of random lead times, one might be tempted to generalize the

single-stage safety stock expression in Equation (2) by replacing the mean and variance of the lead time with the mean and variance of the net-replenishment lead time. That approach is flawed, however, because the shortfall is positive only if the net-replenishment lead time is positive. Using the mean and variance for the net-replenishment lead time can lead to significant errors in setting the safety stock, especially if the probability of a negative net-replenishment lead time is nontrivial.

We remind the reader that the aforementioned development is predicated on the assumption that replenishment orders do not cross in time. That is, an earlier replenishment order never arrives after a later replenishment order. Ruling out order crossing by either assumption (e.g., Kaplan 1970) or by construction (e.g., Song and Zipkin 1996) is quite common in the operations management literature. However, certain papers have explicitly considered order crossing (e.g., Zalkind 1978, Hayya et al. 1995, Bradley and Robinson 2005). If order crossing is of concern, then an alternative safety stock expression to Equation (5) could be developed by adapting the single-stage approach in Bradley and Robinson (2005) to the case of positive inbound and outbound service times.

Early-Arrival Stock

An interesting phenomenon can occur in a network with random lead times. If a stage's outbound service time exceeds its inbound service time, then some lead-time realizations may result in a replenishment order arriving at the stocking location before the associated customer order has shipped. That is, the stage has replenished more demand than it has shipped out. This can only occur if the outbound service time exceeds its inbound service time, and even then only if the realized lead time is lower than the difference between outbound and inbound service times; stock arrived earlier than needed to fill customer demand.

For practical reasons, this early-arrival stock cannot be simply passed on to a stage's downstream customer(s). Recall that a downstream stage operates under the assumption that items from an upstream stage will be available exactly S periods after they were requested, where S is the outbound service time quoted by the upstream stage. For reasons including storage constraints or coordination requirements,

the downstream stage might not be willing to take ownership of this early-arrival stock until the agreed-upon time. Even if the downstream stage is willing to take the items earlier than planned, it may not be able to start processing them until the originally planned time. In either case, the early-arrival stock will exist somewhere in the system in practice. In the appendix, we develop an expression for the average early-arrival stock at a stage; see Equation (6).

Network Inventory Optimization

In the GS model, inventory optimization implies finding an optimal combination of service times in the network, because inventory levels, and hence costs (the objective function), are driven by service times. The problem is usually formulated as a nonlinear optimization problem with linear constraints. Frequently, the optimization must be carried out in the presence of constraints on the amount of safety stock that can be carried at a stage. Any constraints on carrying inventory are incorporated as constraints on the net-replenishment lead times allowed at stages.

For the GS-DET model, the objective function is concave and the feasible region is a polytope. This leads to the well-known property that the optimal solution is an extreme point; that is, each stage sets its outbound service time either to zero or to its maximum allowable service time (i.e., the sum of its inbound service time and its lead time). To solve GS-RAN, we developed a new optimization algorithm to handle cost functions that employ stochastic lead times in general network topologies. Humair and Willems (2011) report the details of the algorithm. In this paper, we focus only on the reasons that necessitated the development of a new algorithm.

Optimizing supply chains with stochastic lead times entails several complications. First, the objective function is no longer constrained to be concave; therefore, the optimal solution may lie in the interior of the feasible region. Second, the objective function may not be smooth (e.g., the lead times may have a discrete probability mass function (PMF) that cannot be approximated by a normal distribution). Hence, standard nonlinear optimization approaches that rely on differentiability to characterize optimal solutions (e.g., Bazaraa et al. 1993) cannot be used. Third, as previously described, random lead times

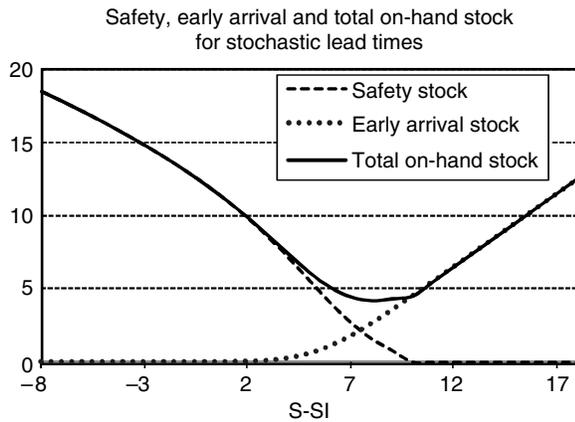


Figure 1: The function that describes expected total inventory at a stage can be nonconcave.

can lead to early-arrival stock in the network, a feature of the optimal solution that does not arise in the GS-DET model.

For a given pair of inbound and outbound service times, the average inventory at a stage is the sum of the safety stock and early-arrival stock, given by Equations (5) and (6), respectively. With increasing outbound service times and fixed inbound service time, safety stock decreases and early-arrival stock increases. With fixed outbound service time and increasing inbound service time, the safety stock increases but early-arrival stock decreases. This behavior leads to the nonconcavity of the objective function (see Figure 1). In addition, if the lead time has an arbitrary discrete distribution with spikes in probability at different times, the objective function can become nondifferentiable (see Figure 2).

Quantifying this objective function for GS-RAN, and developing an algorithm to optimize it for general networks is a key contribution to GS modeling literature, and to the adoption of the GS model among practitioners.

Numerical Analysis of Stochastic Lead Times

To demonstrate the value of our approach, we consider 12 real-world supply chains documented in Willems (2008). Willems (2008) presents 38 supply chains from 22 companies. The company's lead

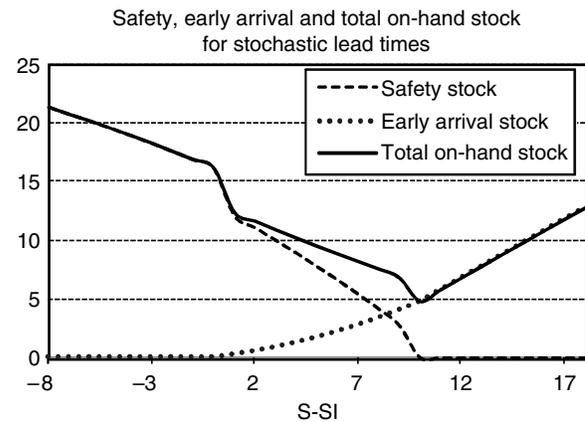


Figure 2: The function that describes expected total inventory at a stage can be nondifferentiable when the lead time has a discrete distribution.

modeler, as representative of that company's business, identified each of these supply chains. Of the 38 chains, 26 use stochastic lead times that are solved using the approach we present in this paper. The supply chains represent industries as varied as pharmaceuticals and farming equipment. We consider the first 12 supply chains from Willems (2008) that employ stochastic lead times (see Table 1).

We begin with the smallest network example from Willems (2008). Chain 01 is a chemical industry supply chain that consists of eight stages and 10 links. Furthermore, the analyst that created this model chose to model only stochastic lead times at a single raw material stage, which represents the key ingredient to the process. This simple example demonstrates how our stochastic lead-times approach produces inventory levels that differ from the alternative approximations sometimes used in practice. Some of the approximations, although they may appear conservative at face value (e.g., fixing the lead time to its maximum possible value), can give misleading results.

Figure 3 presents Chain 01's supply chain map. All three retail stages have a lead time of zero days and both manufacturing stages have a lead time of 10 days. Supply stages Part_0002 and Part_0003 have lead times of 15 and 10 days, respectively. Part_0001, however, can have a lead time of 20 days or 25 days, each with probability 40 percent, or a lead time of 50 days with probability 20 percent. Table 2 shows the remainder of the network data.

Chain name	SIC code	SIC description	Average chain length	Maximum chain length	No. of normal stage times	No. of discrete stage times	Total stages	Total arcs
1	2869	Industrial Organic Chemicals, Not Elsewhere Classified	38	38	—	1	8	10
3	3577	Computer Peripheral Equipment, Not Elsewhere Classified	77.04	79.8	—	8	17	18
5	2099	Food Preparations, Not Elsewhere Classified	45.28	47.35	7	9	27	31
6	3421	Cutlery	88.35	96	16	—	28	28
7	3531	Construction Machinery and Equipment	81.14	85	38	—	38	78
8	3845	Electromedical and Electrotherapeutic Apparatus	91.04	91.04	—	23	40	48
9	2043	Cereal Breakfast Foods	44.32	47.38	9	2	49	52
10	5064	Electrical Appliances, Television, and Radio Sets	116.97	162	21	—	58	176
11	3531	Construction Machinery and Equipment	55.89	60	39	6	68	108
12	2043	Cereal Breakfast Foods	105.03	108.6	16	12	88	107
14	4731	Arrangement of Transportation of Freight and Cargo	90.12	131.63	36	—	116	119
15	2841	Soap and Other Detergents, Except Specialty Cleaners	21.72	26	77	—	133	164

Table 1: The table shows summary data for the first 12 chains from Willems (2008) that employ stochastic lead times.

When all service times are zero, some common approximations are to assume Part_0001’s lead time is fixed at its mean value, or at its maximum possible value. In either case, the resulting supply chain has deterministic lead times and we can use GS-DET to optimize its safety stock. These two heuristic solution procedures, however, significantly underestimate and overestimate inventory levels at Part_0001. Table 3 shows that fixing the lead time to its mean underestimates the safety stock at Part_0001 by approximately 7,400 units, which is very large compared to the average demand seen at that stage (418 units). Fixing the lead time to its maximum value still

Stage name	Stage cost	Demand average	Demand std. dev.	Max service time	Service level (%)	Lead-time average
Manuf_0001	39					10
Manuf_0002	36					10
Part_0001	12					28
Part_0002	5					15
Part_0003	9					10
Retail_0001	0	253	36.62	0	95	0
Retail_0002	0	45	1	0	95	0
Retail_0003	0	75	2	0	95	0

Table 2: The table shows the data for supply chain 01, that Figure 3 presents graphically.

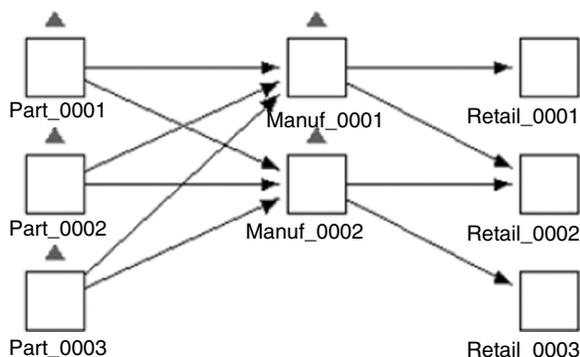


Figure 3: Supply chain 01 from Willems (2008) is an industrial chemical product for which the analyst modeled only stage-time variability for a single raw material Part_0001.

underestimates the safety stock by approximately 7,300 units, but overestimates the pipeline stock by about 9,200 units, overestimating the total inventory at Part_0001 by about 1,900 units.

These overestimates and underestimates produce an –18 percent and +4 percent difference in total inventory investment in the supply chain (i.e., over the value obtained under our stochastic lead-times approach).

Table 4 summarizes the results for all 12 chains. We focus our attention on Chain 12, an end-to-end supply chain for cereal that consists of 88 stages and 107 arcs. Of the 88 stages, 16 employ Normally distributed lead times and 12 employ discrete stochastic lead times. Figure 4 depicts the supply chain.

Stock level	GS-RAN	GS-DET with mean LT	GS-DET with max LT
Safety	7,724	319	427
Pipeline	11,704	11,704	20,900
Total	19,428	12,023	21,327

Table 3: This table reports the resulting quantities of each inventory type for three solution approaches—our approach (GS-RAN), and two heuristic approaches that convert the problem to a deterministic formulation that allow GS-DET to be solved. These heuristic approaches significantly underestimate the expected safety stock for Part_0001. LT = lead times.

We again compare our stochastic lead-time approach to the heuristics of replacing a stage's lead-time variability with the average lead-time value or its maximum lead-time value.

As in the case of Chain 01, the average lead-time heuristic underestimates total inventory required and the maximum lead-time heuristic overestimates total inventory required by –3.2 percent and 2.4 percent, respectively. Of even greater significance, the optimal inventory locations change for the heuristic solutions

Chain name	Solution approach	Pipeline stock	Safety stock	Early-arrival stock	Total stock	No. of stages holding safety stock	No. of stages with stochastic lead times	Total stages
01	GS-RAN	26,334	8,351		34,685	5	1	8
01	GS-DET using mean LT	26,334	946		27,280	5	1	8
01	GS-DET using max LT	35,530	1,054		36,584	5	1	8
03	GS-RAN	51,723	15,744		67,467	9	8	17
03	GS-DET using mean LT	51,723	5,004		56,728	7	8	17
03	GS-DET using max LT	65,157	5,546		70,703	7	8	17
05	GS-RAN	4,200,254	1,043,257		5,243,511	25	16	27
05	GS-DET using mean LT	4,200,254	630,507		4,830,762	22	16	27
05	GS-DET using max LT	4,835,046	678,482		5,513,529	22	16	27
06	GS-RAN	1,461,409	50,062	17,838	1,529,310	16	16	28
06	GS-DET using mean LT	1,461,409	805		1,462,215	22	16	28
06	GS-DET using max LT	1,539,907	864		1,540,772	22	16	28
07	GS-RAN	2,058	363	18	2,441	34	38	38
07	GS-DET using mean LT	2,058	265		2,324	19	38	38
07	GS-DET using max LT	2,268	279		2,548	19	38	38
08	GS-RAN	775,578	46,030	60,691	882,300	5	23	40
08	GS-DET using mean LT	805,525	41,180		846,706	4	23	40
08	GS-DET using max LT	849,073	40,050		889,123	3	23	40
09	GS-RAN	3,348,802	385,497		3,734,300	35	11	49
09	GS-DET using mean LT	3,348,802	106,308		3,455,110	24	11	49
09	GS-DET using max LT	3,439,581	113,757		3,553,338	24	11	49
10	GS-RAN	31,698	4,890	17	36,606	55	21	58
10	GS-DET using mean LT	31,698	3,939		35,637	52	21	58
10	GS-DET using max LT	33,219	4,024		37,244	52	21	58
11	GS-RAN	2,393	490	7	2,891	48	45	68
11	GS-DET using mean LT	2,393	402		2,795	36	45	68
11	GS-DET using max LT	2,784	442		3,227	36	45	68
12	GS-RAN	9,921,038	1,511,516	101,743	11,534,298	80	28	88
12	GS-DET using mean LT	9,921,038	1,255,293		11,176,331	75	28	88
12	GS-DET using max LT	10,572,177	1,264,372		11,836,550	75	28	88
14	GS-RAN	642,693	6,002	1,576	650,272	84	36	116
14	GS-DET using mean LT	642,693	2,926		645,620	55	36	116
14	GS-DET using max LT	648,425	3,237		651,662	64	36	116
15	GS-RAN	1,684,500	775,010		2,459,510	77	77	133
15	GS-DET using mean LT	1,684,500	511,227		2,195,727	77	77	133
15	GS-DET using max LT	2,151,861	580,678		2,732,540	77	77	133

Table 4: The table shows the resulting quantities of each inventory type for three solution approaches—our approach (GS-RAN), and two heuristic approaches that convert the problem to a deterministic formulation that allows GS-DET to be solved. LT = lead times.

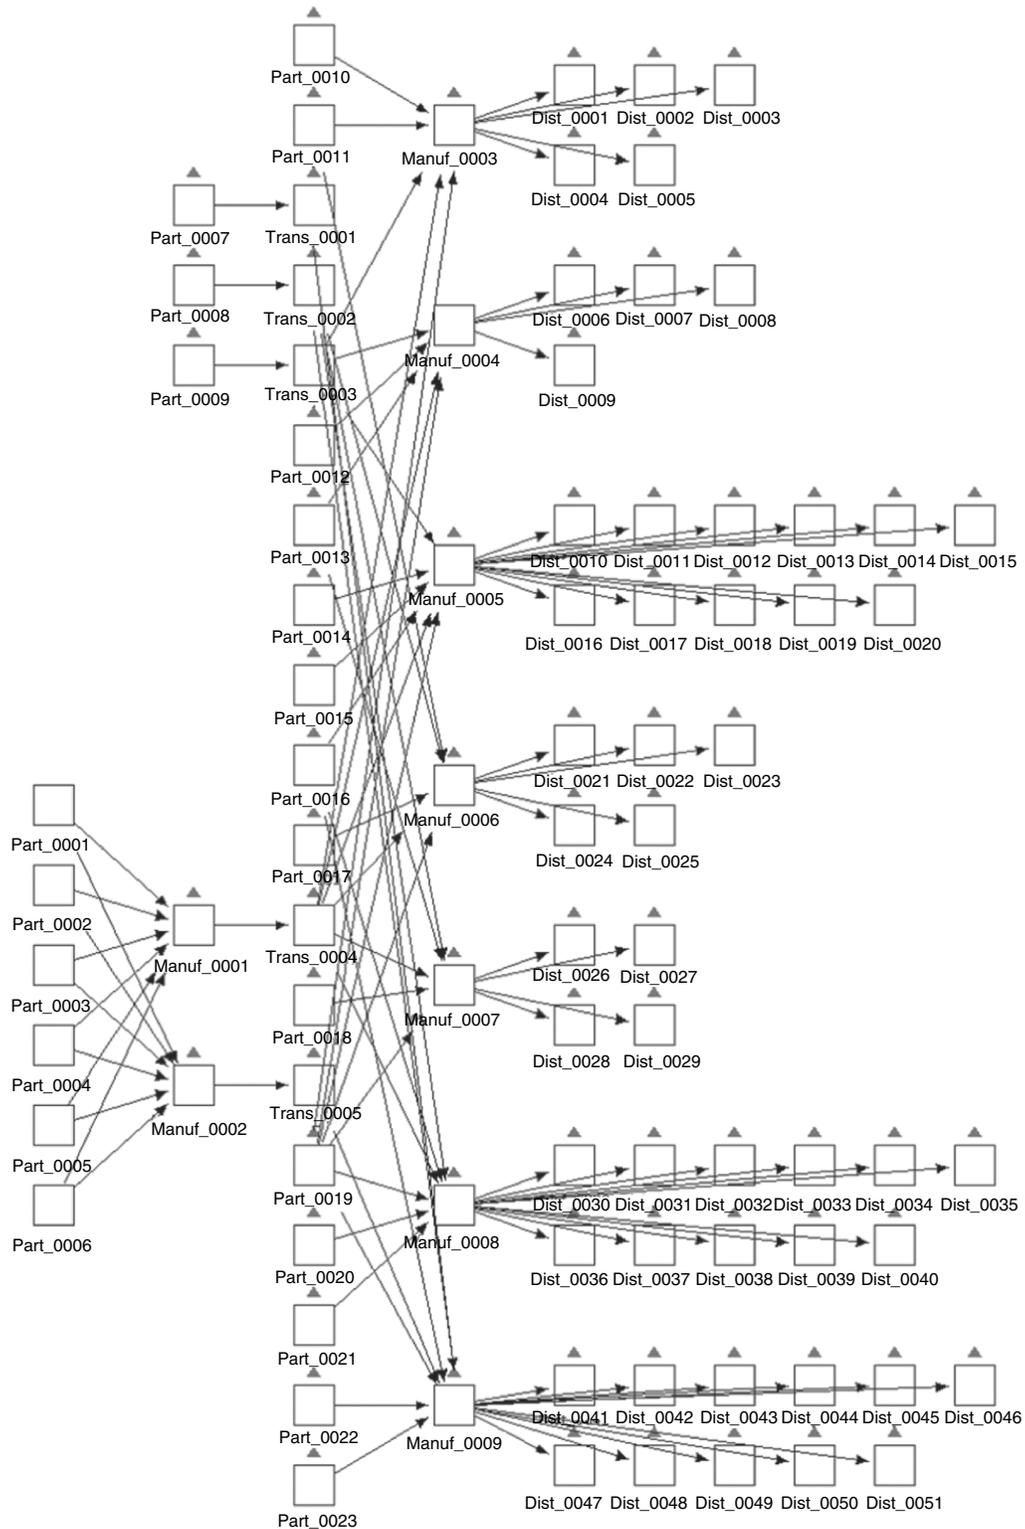


Figure 4: Chain 12 from Willems (2008) is a cereal supply chain comprised of 88 stages and 107 arcs. Of the 88 stages, 28 employ stochastic lead times.

Downloaded from informs.org by [168.122.65.189] on 12 December 2013, at 21:01 . For personal use only, all rights reserved.

versus the stochastic lead-time solution. For the heuristic solutions, 75 stages hold safety stock; however, in the stochastic lead-time solution, 80 stages hold safety stock. Intuitively, chains that properly model stochastic lead times have more stages that hold inventory to avoid early-arrival stock. By assuming deterministic lead times, the heuristics assume away this problem, which allows fewer stages to hold inventory. In effect, stochastic lead times can break the power of pooling over stages, because the early-arrival stock ramifications make this type of pooling uneconomical.

Real-World Considerations

Companies that implement a sufficiently large, automated, repeatable GS model into their planning process typically do not already have the lead-time input parameters required for the GS-RAN model. In particular, although analysts have often calculated lead-time variability for strategic inventory projects via tools such as spreadsheets and desktop database systems, very few companies have these data calculated and stored in the centralized enterprise resource planning (ERP) or data warehouse systems that normally provide the other data required for large-scale GS model implementations.

Next, we detail some challenges encountered in measuring lead-time variability.

How to Determine Lead-Time Variability

Lead-time variability, or a lead-time PMF, is most readily calculated from transactional data available in the ERP system. For some item locations, shipments occur often enough to a stocking location that taking the most recent history of shipments, either over a fixed time frame or a count of shipments, is sufficient to provide a suitable sample for estimating an average and standard deviation or for generating a histogram that can serve as a PMF.

For other item locations, usually those with long review cycles or with intermittent demand, the individual shipment history does not provide enough data points for a valid sample. In these cases, items can be partitioned by a common characteristic into buckets, and the lead-time distribution is calculated on shipment history for all item locations within each bucket. An example of a partitioning scheme is to

bucket all items at each stocking location by supplier. This might be appropriate when few suppliers provide many items to the location, and each supplier's shipment characteristics are independent of which items are being shipped. For example, this can apply to transportation within the company's own distribution network from central to regional distribution centers, where goods are normally batched together into heterogeneous truckloads.

Another technique to account for insufficient data is to carry forward the lead-time parameters used previously for an item location or its bucket. Typically, if the number of available shipment transactions for an item location is less than a specified threshold, then the shipment history of that item location's bucket will be used. If the history is also less than the threshold, then the prior lead-time parameters for the item location will be used. If a set of prior parameters exists, then the prior specified coefficient of variation (COV) can be applied. The prior COV can be calculated before implementation based on historical behavior of similar items for the business.

Outlier Detection

Once all the data have been collected, investigating whether all the data should be used to calculate lead-time variability or whether some outliers should be removed is a valuable exercise. Safety stock requirements increase significantly as a function of lead-time variability. If specific outliers are true once-off or exception events, then pruning them from the data set is often advisable.

Calculating Lead-Time Variability at the Right Level of the Supply Chain Hierarchy

Inventory targets are calculated at the stock-keeping unit (SKU) location level; therefore, lead-time variability must ultimately be specified at the SKU location level. However, first calculating lead-time variability at a higher level in the product hierarchy and then populating the SKU location level can be useful. If a single shipment is delayed, then only the SKUs on that shipment will have lead-time variability. In this case, aggregating all the lead-time information for the entire product family and then using the resulting lead-time data for each SKU in the product family are better actions to take.

Frequency of Updating Lead Times

For the purposes of setting safety stock levels, updating lead-time variability with the same frequency with which inventory targets are updated does no harm; however, it also provides little benefit. Whereas inventory targets are often updated weekly or monthly, updating lead-time variability quarterly or semi-annually seems to balance the effort with the reward.

Conclusion

In this paper, we extend the GS modeling framework to incorporate stochastic lead times. The primary modeling challenge involves properly characterizing inventory requirements at a stage when it faces positive inbound and (or) outbound service times. When placed in a multiechelon supply chain, this requires properly characterizing both safety stock and early-arrival stock. We compare our approach against two reasonable heuristics and find that our solution (1) more accurately calculates total inventory levels, and (2) produces a different inventory stocking strategy; stochastic lead times make it optimal to have more stages hold safety stock in the supply chain. Because most ERP systems do not include variability information, determining the inputs for stochastic lead times presents some real-world challenges; however, these challenges are solvable.

The solution procedure we describe in this paper represents a pragmatic solution approach to a hard supply chain problem, namely, optimizing safety stock levels and locations in a multiechelon supply chain that faces both demand and supply variability.

Appendix

Deterministic Lead-Time Safety Stock

We assume that all stages have deterministic lead times; we denote stage k 's lead time as L_k , where stage k is any arbitrary stage in the network. If end-customer demands are normally distributed and independent across periods (which we assume for ease of exposition in this paper), then for a given set of service times in the network, the GS-DET model sets the safety stock at stage k as

$$SFTY_k = z_k \sigma_k^d \sqrt{SI_k + L_k - S_k}, \quad (1)$$

where z_k is a safety factor that reflects the service level at stage k ; that is, $z_k = \text{normsinv}(\text{service}_k)$, σ_k^d is the standard

deviation of demand at stage k , S_k is the outbound service time quoted by stage k , and SI_k is the inbound service time experienced by stage k (i.e., the maximum of the service times quoted by stage k 's direct supplier stages). In Equation (1), we note that $SI_k + L_k - S_k \geq 0$, because it is never necessary in GS-DET to quote an outbound service time that exceeds the sum of the inbound service time and the lead time. The service level and demand standard deviation for internal stages (i.e., stages that do not directly serve end-customer demand) can be imputed from the service levels and demand characteristics of the end-customer stages.

Defining the net-replenishment time for a stage as $NRLT_k = SI_k + L_k - S_k$ is instructive. The safety stock expression in Equation (1) can then be written as $SFTY_k = z_k \sigma_k^d \sqrt{NRLT_k}$. When the inbound (SI_k) and outbound (S_k) service times are both zero, then $NRLT_k = L_k$ and the safety stock expression takes on the exact form of the widely-used single-stage expression (i.e., $SFTY_k = z_k \sigma_k^d \sqrt{L_k}$). When either or both of the service times can be positive, then the net-replenishment lead time simply replaces the lead time in this widely-used expression.

Random Lead-Time Safety Stock

As we describe in the main body of this paper, safety stock is a buffer that absorbs randomness in the shortfall, where the shortfall at time t is the difference between what a stage has shipped out and what it has replenished. For any arbitrary stage k , we define $REPL_k(t)$ as the total replenishment orders received by stage k from period 1 up to and including period t , and we define $SHIP_k(t)$ as the cumulative demand filled by stage k from period 1 up to and including period t . It then follows that $SHORT_k(t) = SHIP_k(t) - REPL_k(t)$. Characterizing the shortfall expression is central to developing the safety stock expression. To promote clarity, we first consider the case of instantaneous inbound and outbound service times, and then move on the case of general service times.

Inbound and Outbound Service Times (Both Zero). Because stage k quotes an outbound service time of $S_k = 0$, at time t , it will have shipped out all the demands it received in periods $1, \dots, t$ (i.e., $SHIP_k(t) = d_k^1 + d_k^2 + \dots + d_k^t$). Let $p_k(t)$ denote the period in which stage k placed its most recently replenished order. Then, because replenishment orders do not cross, stage k has replenished all the demands it received in periods $1, \dots, p_k(t)$, that is, $REPL_k(t) = d_k^1 + d_k^2 + \dots + d_k^{p_k(t)}$. Because stage k experiences an inbound service time of $SI_k = 0$, it can start processing a replenishment order immediately and then take the random lead time \tilde{L}_k to process the replenishment order. Therefore, the period in which the most recently replenished order was placed is given by $p_k(t) = t - l_k(t)$, where $l_k(t)$ is the realized lead time of the most recently replenished order. Because $SHORT_k(t) = SHIP_k(t) - REPL_k(t)$, it then follows that

$$SHORT_k(t) = \sum_{i=0}^{l_k(t)-1} d_k^{t-i}.$$

That is, the shortfall at time t is the sum of the most recent $l_k(t)$ periods of demand. Applying results from statistics for the random sum of random variables (e.g., Drake 1967, p. 112), the standard deviation of the shortfall at stage k is

$$\sigma_{s_k} = \sqrt{\mu_k^l (\sigma_k^d)^2 + (\mu_k^d)^2 (\sigma_k^l)^2},$$

where μ_k^d and σ_k^d are the mean and standard deviation of the demand at stage k and μ_k^l and σ_k^l are the mean and standard deviation of the lead time at stage k . If one models the shortfall as having an approximately normal distribution, then the safety stock at stage k should be set as

$$SFTY_k = z_k \sqrt{\mu_k^l (\sigma_k^d)^2 + (\mu_k^d)^2 (\sigma_k^l)^2}, \quad (2)$$

where z_k is a safety factor reflecting the service level at stage k (i.e., $z_k = \text{normsinv}(\text{service}_k)$).

General Inbound and Outbound Service Times.

Because stage k quotes a service time of $S_k \geq 0$, at time t , it will have shipped out all the demands it received in periods $1, \dots, t - S_k$ (i.e., $SHIP_k(t) = d_k^1 + d_k^2 + \dots + d_k^{t-S_k}$). Let $p_k(t)$ denote the period in which stage k placed its most recently replenished order. Then, because replenishment orders do not cross, stage k has replenished all the demands it received in periods $1, \dots, p_k(t)$ (i.e., $REPL_k(t) = d_k^1 + d_k^2 + \dots + d_k^{p_k(t)}$). Now, stage k cannot start processing any replenishment order until all input items are available, which occurs exactly $SI_k \geq 0$ periods after the replenishment order is placed, and it then takes the random lead time \tilde{L}_k to process the replenishment order. Therefore, $p_k(t) = t - (SI_k + l_k(t))$, where $l_k(t)$ is the realized lead time of the most recently replenished order.

By definition, the shortfall at time t is $SHORT_k(t) = SHIP_k(t) - REPL_k(t)$, which is positive when stage k has filled some demands that it has not yet replenished. That is, the shortfall is positive when the demand associated with the most recently replenished order has not yet been filled (i.e., when $p_k(t) < t - S_k$). Because $p_k(t) = t - (SI_k + l_k(t))$, the condition that $p_k(t) < t - S_k$ is equivalent to $l_k(t) > S_k - SI_k$. Now, if $l_k(t) = S_k - SI_k$, then $p_k(t) = t - S_k$, stage k has replenished and filled exactly the same set of demands, and $SHORT_k(t) = 0$. However, if $l_k(t) < S_k - SI_k$, then $p_k(t) > t - S_k$ and stage k has replenished all the demands that it has filled and also some demands it has yet to fill. The shortfall is negative in this case, showing that the stage has had more replenishments come in than orders it has shipped out.

Observe that when inbound and outbound service times are both zero (i.e., $S_k = SI_k = 0$), then we can never have a negative shortfall because lead times are nonnegative, which rules out the scenario $l_k(t) < 0$. With general service times, however, it may be possible for a lead-time realization to be low enough such that $l_k(t) < S_k - SI_k$. Defining the realized net-replenishment time for the most recently received order as $nrlt_k(t) = SI_k + l_k(t) - S_k$, we then see that

the shortfall is positive if $nrlt_k(t) > 0$, the shortfall is zero if $nrlt_k(t) = 0$, and the shortfall is negative if $nrlt_k(t) < 0$.

Define $SHORT_k^+(t)$ as the positive part of $SHORT_k(t)$ and $SHORT_k^-(t)$ as the negative part of $SHORT_k(t)$; that is, $SHORT_k^+(t) = \max\{SHORT_k(t), 0\}$ and $SHORT_k^-(t) = \min\{SHORT_k(t), 0\}$. From above, $SHORT_k(t) \geq 0$ if and only if $nrlt_k(t) \geq 0$, and $SHORT_k(t) \leq 0$ if and only if $nrlt_k(t) \leq 0$. Applying $SHORT_k(t) = SHIP_k(t) - REPL_k(t)$, it then follows that

$$SHORT_k^+(t) = \begin{cases} \sum_{i=0}^{nrlt_k(t)-1} d_k^{t-S_k-i}, & nrlt_k(t) > 0; \\ 0, & nrlt_k(t) \leq 0; \end{cases} \quad (3)$$

$$SHORT_k^-(t) = \begin{cases} -\sum_{i=1}^{-nrlt_k(t)} d_k^{t-S_k+i}, & nrlt_k(t) < 0; \\ 0, & nrlt_k(t) \geq 0. \end{cases} \quad (4)$$

That is, the positive shortfall at time t is the sum of the $nrlt_k(t)$ periods of demand counting back from $t - S_k$ (assuming that $nrlt_k(t)$ is positive) and the negative shortfall at time t is the negative sum of the $-nrlt_k(t)$ periods of demand counting forward from $t - S_k$ (assuming that $nrlt_k(t)$ is negative). For completeness, we note that $SHORT_k(t) = SHORT_k^+(t) + SHORT_k^-(t)$.

A stockout occurs only if the shortfall exceeds the base stock level. As such, a stockout cannot occur at time t if the shortfall at time t is negative. Therefore, we are only interested in the positive shortfall when developing a safety stock expression. In particular, we are interested in the statistics of the positive shortfall. We develop the shortfall statistics making no assumption on the lead-time distribution and define the following expressions based on the lead-time distribution for stage k :

$$H_k^1(T) = \sum_{l=0}^T P[L_k = l],$$

$$H_k^2(T) = \sum_{l=T}^{\infty} l P[L_k = l],$$

$$H_k^3(T) = \sum_{l=T}^{\infty} l^2 P[L_k = l],$$

where $P[L_k = l]$ is the probability that the realized lead time for stage k equals l . If instead of specifying a general discrete distribution, the user wishes to specify a continuous distribution, then these expressions become $H_k^1(T) = \int_{l=0}^T f_{L_k}(l) dl$, $H_k^2(T) = \int_{l=T}^{\infty} l f_{L_k}(l) dl$, and $H_k^3(T) = \int_{l=T}^{\infty} l^2 f_{L_k}(l) dl$, where $f_{L_k}(l)$ is the lead-time density function. For certain distributions (e.g., normal), closed-form expressions exist for these H_k functions.

For any given pair of inbound and outbound service times (SI_k and S_k , respectively), the realized positive shortfall at time t , $SHORT_k^+(t)$, is given by Equation (3). Using Equation (3) and recalling that μ_k^d is the mean demand

for stage k , and $nrlt_k(t) = SI_k + I_k(t) - S_k$ is the realized net-replenishment time, then the mean positive shortfall is $nrlt_k(t)\mu_k^d$ if the realized net-replenishment lead time is non-negative, but the mean positive shortfall is 0 if the realized net-replenishment lead time is negative. Using the conditioning identity that $E[X] = E[E[X | Y]]$, where X and Y are random variables and $E[\cdot]$ denotes expectation, we can express (after some algebraic manipulation) the mean positive shortfall as

$$\mu_{s_k^+}(SI_k, S_k) = \mu_k^d Q(\max\{S_k - SI_k, 0\}),$$

where

$$Q(T) = H_k^2(T) - T(1 - H_k^1(T))$$

is the expected value of the positive part of the net-replenishment lead time when $T = \max\{S_k - SI_k, 0\}$. Using Equation (3) and recalling that σ_k^d is the standard deviation of demand for stage k , then the variance of the positive shortfall is $nrlt_k(t)(\sigma_k^d)^2$ if the realized net-replenishment lead time is nonnegative, but the variance of the positive shortfall is 0 if the realized net-replenishment lead time is negative. Using the conditioning identity that $V[X] = E[V[X | Y]] + V[E[X | Y]]$ where $V[\cdot]$ denotes variance, we can express (after some algebraic manipulation) the standard deviation of the positive shortfall as:

$$\begin{aligned} \sigma_{s_k^+}(SI_k, S_k) &= \sqrt{Q(\max\{S_k - SI_k, 0\})(\sigma_k^d)^2 + (\mu_k^d)^2 R(\max\{S_k - SI_k, 0\})}, \end{aligned}$$

where $Q(T)$ is given above, and where

$$\begin{aligned} R(T) &= T^2 H_k^1(T)(1 - H_k^1(T)) - 2TH_k^1(T)H_k^2(T) \\ &\quad + H_k^3(T) - (H_k^2(T))^2, \end{aligned}$$

is the variance of the positive part of the net-replenishment lead time when $T = \max\{S_k - SI_k, 0\}$. We note that $Q(0) = \mu_k^l$ and $R(0) = (\sigma_k^l)^2$; therefore, in the special case in which $SI = S = 0$, we obtain $\sigma_{s_k^+}(0, 0) = \sqrt{\mu_k^l(\sigma_k^d)^2 + (\mu_k^d)^2(\sigma_k^l)^2}$.

If one models the positive shortfall as having an approximately normal distribution, then the safety stock at stage k should be set as

$$\begin{aligned} SFTY_k(SI_k, S_k) &= z_k \sqrt{Q([S_k - SI_k]^+)(\sigma_k^d)^2 + (\mu_k^d)^2 R([S_k - SI_k]^+)}, \end{aligned} \quad (5)$$

where $[S_k - SI_k]^+ = \max\{S_k - SI_k, 0\}$, and where z_k is a safety factor reflecting the service level at stage k ; that is, $z_k = \text{normsinv}(\text{service}_k)$. Adopting a normal approximation for the (positive) shortfall may not be reasonable in all circumstances; for example, if the lead-time distribution is very bimodal. A more robust general safety stock expression can be developed by adapting the single-stage convolution approach in Eppen and Martin (1988) to the general service time case by recognizing that the positive part of the net-replenishment lead time takes on the role of the lead time.

Early-Arrival Stock

As we discuss in the main body of this paper, the early-arrival stock at stage k at time t is the positive part of the difference between the cumulative replenishments and the cumulative shipments. Equivalently, the early-arrival stock at time t is zero minus the negative part of the shortfall at time t . For any given pair of inbound and outbound service times (SI_k and S_k , respectively), the realized negative shortfall at time t , $SHORT_k^-(t)$, is given by Equation (4). Using Equation (4) and recalling that μ_k^d is the mean demand for stage k and $nrlt_k(t) = SI_k + I_k(t) - S_k$ is the realized net-replenishment time, then the mean negative shortfall is $-nrlt_k(t)\mu_k^d$ if the realized net-replenishment lead time is nonpositive, but the mean negative shortfall is 0 if the realized net-replenishment lead time is positive. Using the conditioning identity that $E[X] = E[E[X | Y]]$, where X and Y are random variables and $E[\cdot]$ denotes expectation, we can express (after some algebraic manipulation) the mean negative shortfall as

$$\mu_{s_k^-}(SI_k, S_k) = \mu_k^d (\mu_k^l - \max\{S_k - SI_k, 0\} - Q(\max\{S_k - SI_k, 0\})),$$

where $Q(T)$ is given above. Because early-arrival stock is zero minus the negative part of the shortfall, we can write the average early-arrival stock as

$$\begin{aligned} EARLY_k(SI_k, S_k) &= \mu_k^d (Q([S_k - SI_k]^+) - \mu_k^l + [S_k - SI_k]^+), \end{aligned} \quad (6)$$

where $[S_k - SI_k]^+ = \max\{S_k - SI_k, 0\}$. We note that $Q(0) = \mu_k^l$ and therefore $EARLY_k(SI_k, S_k) = 0$ if $S_k \leq SI_k$; that is, we have no early-arrival stock unless the outbound service time exceeds the inbound service time. Because single-stage models are associated with zero inbound and outbound service times, it follows that early-arrival stock only becomes relevant in multistage networks.

References

- Bazaraa MS, Sherali HD, Shetty CM (1993) *Nonlinear Programming: Theory and Algorithms*, 2nd ed. (John Wiley & Sons, Hoboken, NJ).
- Bradley JR, Robinson LW (2005) Improved base-stock approximations for independent stochastic lead times with order crossover. *Manufacturing Service Oper. Management* 7(4):319–329.
- Clark AJ, Scarf H (1960) Optimal policies for a multi-echelon inventory problem. *Management Sci.* 6(4):475–490.
- CSCO (2011) Five strategies for improving inventory management across complex supply chain networks. Accessed June 4, 2013, http://www.scdigest.com/assets/repes/exec_brief_network_inventories.pdf.
- Drake AW (1967) *Fundamentals of Applied Probability Theory*, 3rd ed. (McGraw-Hill, New York).
- Eppen GD, Martin RK (1988) Determining safety stock in the presence of stochastic lead time and demand. *Management Sci.* 34(11):1380–1390.

- Farasyn I, Humair S, Kahn JI, Neale JJ, Rosen O, Ruark J, Tarlton W, Van de Velde W, Wegryn G, Willems SP (2011) Inventory optimization at Procter & Gamble: Achieving real benefits through user adoption of inventory tools. *Interfaces* 41(1): 66–78.
- Graves SC, Willems SP (2000) Optimizing strategic safety stock placement in supply chains. *Manufacturing Service Oper. Management* 2(1):68–83.
- Graves SC, Willems SP (2003) Supply chain design: Safety stock placement and supply chain configuration. de Kok AG, Graves SC, eds. *Supply Chain Management: Design, Coordination, and Operation* (Elsevier, Amsterdam), 95–132.
- Hayya JC, Xu SH, Ramasesh RV, He XX (1995) Order crossover in inventory systems. *Stochastic Models* 11(2):279–309.
- Hopp WJ, Spearman ML (2008) *Factory Physics*, 3rd ed. (Waveland Press, Inc., Long Grove, IL).
- Humair S, Willems SP (2011) Technical note: Optimizing strategic safety stock placement in general acyclic networks. *Oper. Res.* 59(3):781–787.
- Kaplan RS (1970) A dynamic inventory model with stochastic lead times. *Management Sci.* 16(7):491–507.
- Nahmias S (1997) *Production and Operations Analysis*, 3rd ed. (Irwin, Burr Ridge, IL).
- SCMWorld (2012) 2012 supply chain predictions. Webinar, SCMWorld, London.
- Silver EA, Pyke DF, Peterson R (1998) *Inventory Management and Production Planning and Scheduling*, 3rd ed. (John Wiley & Sons, Hoboken, NJ).
- Simpson KF Jr (1958) In-process inventories. *Oper. Res.* 6(6):863–873.
- Song J-S, Zipkin PH (1996) Inventory control with information about supply conditions. *Management Sci.* 42(10):1409–1419.
- Willems SP (2008) Real-world multiechelon supply chains used for inventory optimization. *Manufacturing Service Oper. Management* 10(1):19–23.
- Zalkind D (1978) Order-level inventory systems with independent stochastic leadtimes. *Management Sci.* 24(13):1384–1392.

Salal Humair is a research scientist at the Harvard School of Public Health, and an associate professor at the School of Science and Engineering, Lahore University of Management Sciences in Pakistan. He received his doctorate in operations research at the Massachusetts Institute of Technology. His research interests are in applying Operations Research to large-scale systems such as healthcare and water systems.

John D. Ruark is Chief Technology Officer, Inventory Optimization Solutions, at Logility, Inc. and was CTO and cofounder of Optiant, Inc., an award-winning provider of multiechelon inventory optimization applications. He received a PhD in operations research from the Massachusetts Institute of Technology.

Brian Tomlin is an associate professor of business administration at the Tuck School of Business at Dartmouth. His research focuses on operations strategy and supply chain management. Brian received his PhD from MIT's Sloan School of Management, where he was awarded the Zannetos PhD Dissertation Prize. His undergraduate degree is from University College Dublin in Ireland. Prior to becoming an academic, Brian worked full time for a number of companies, including General Electric and the Boston Consulting Group.

Sean P. Willems is associate professor of operations and technology management at Boston University's School of Management. His research focuses on supply chain design and optimization problems. Sean received his BSE in decision sciences at The Wharton School, University of Pennsylvania, and his master's in operations research and PhD in operations management from MIT.