Int. J. Production Economics 181 (2016) 107–112

Contents lists available at ScienceDirect

Int. J. Production Economics

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

Analytical insights into two-stage serial line supply chain safety stock

N. Grace Hua [a], Sean P. Willems [b,*]

[a] Microsoft Corporation, Redmond, WA 98052, USA
[b] University of Tennessee, Knoxville, TN 37996, USA

ARTICLE INFO

Article history:
Received 30 September 2014
Accepted 13 October 2015
Available online 23 October 2015

Keywords:
Multiechelon inventory system
Guaranteed service model of safety stock optimization
Inventory planning
Cost leadtime tradeoff
Real world applications

ABSTRACT

Effective inventory management is one of the most significant challenges facing today's global supply chains. Businesses are observing significant profitability gain by optimizing their inventory. This paper optimizes safety stock inventory in a two-stage serial line supply chain, inspired by real-life Cisco supply chains, under guaranteed-service safety stock model assumptions. We analytically show that the optimal safety stock levels depend on the cost and leadtime parameters of the supply chain. Intuitively, it is only worthwhile to hold safety stock inventory at the upstream stage when cost at the upstream stage is relatively low or its leadtime is relatively long. We also show that total supply chain safety stock cost can be reduced when cost allocated at the upstream stage is reduced or leadtime at the upstream stage is increased.

1. Cisco supply chain practice

Competition among companies has been shifted to competition among supply chains. In order to buffer demand variation, a large amount of capital is usually tied up in the form of inventories. Therefore, effective inventory management becomes one of the most significant challenges for global supply chains.

Many global companies like Cisco Systems Inc. (Cisco) are running extensively outsourced supply chain networks. Component suppliers, chip manufacturers, distributors, and other trading partners, although not owned by Cisco, are all key players in Cisco's supply chain. Cisco's key to success, in the absence of full ownership, is to be able to effectively plan and control the outcomes of each stage in the supply chain.

Fig. 1 presents an echelon-based view of Cisco's supply chain for it low end product space, including phones, simple routers and switches.

In this supply chain, vendors supply components for PCBA boards. These components arrive at an assembly site (referred to in the industry as a “Super Market”), where they are assembled into PCBA boards plus any further processing. Completed PCBA boards are then shipped to Direct Fulfillment (DF) sites where they are assembled into finished products. Finished products are then shipped to Strategic Logistics Center to meet external demands from customers.

Cisco owns the safety stock inventory decision-making at both Super Market and Direct Fulfillment (DF) sites. At Cisco, the DF sites always hold safety stock inventory of finished product. Whether Super Market sites hold safety stock of completed PCBA boards depends on the product's cost, the PCBA board building leadtime, as well as the number of DF sites supplied by the Super Market. For example, IP phone parts are lower in cost and shorter in leadtime so Super Markets usually do not hold PCBA safety stock for IP phones. In contrast, low-end routers have higher cost and more complicated parts that require longer assembly leadtime, and it is a common practice for their Super Market sites to hold PCBA safety stock inventory.

The Cisco example can be stylized to a two-stage serial-line supply chain, where the downstream stage (i.e., Direct Fulfillment) always needs to meet external demand from stock, and the upstream stage (i.e., Super Market) always has its components when needed. Cisco's objective is to minimize the total safety stock cost in this two-stage supply chain, with the primary inventory question being whether or not to maintain a safety stock at the upstream stage after its processing activity is complete.

Cisco's practices also seem to indicate that cost, leadtime and network structure jointly determine the optimal safety stock inventory decisions in their Super Market and DF sites. Under certain conditions, it makes sense to break the power of pooling and place inventory in both Super Market and DF sites. To gain deeper insight into these safety stock decisions and to prove or dispute their existing practices, this paper studies a two-stage

* Corresponding author at: Haslam Chair in Supply Chain Analytics, Professor of Analytics and Statistics, Haslam College of Business University of Tennessee, 237 Stokely Management Center, 916 Volunteer Boulevard Knoxville, TN 37996, USA. Tel.:0525 865-974-3991.

E-mail addresses: nih@microsoft.com (N. Grace Hua), swillems@utk.edu (S.P. Willems).

URL: http://willems.utk.edu/.

http://dx.doi.org/10.1016/j.ijpe.2015.10.010
0925-5273/

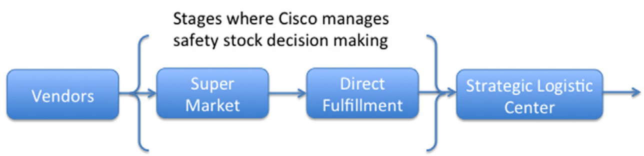

Fig. 1. Echelon-based view of Cisco Systems' supply chain for low end products.

serial line supply chain in order to analytically depict the joint impact of cost and leadtime on safety stock inventory decisions.

2. Literature review

The inventory management literature is sizable. We refer the readers to Federgruen (1993) and Porteus (2002) for a broad survey of the field. Simpson (1958) proposed an inventory model, later coined the guaranteed-service (GS) model by Graves and Willems (2003), that solves for the optimal safety stock levels and locations in serial line supply chains. In the GS framework, every stage faces a random but bounded demand from its downstream stage, and operates according to a base-stock policy. Each stage quotes an outgoing service time to its downstream stage, and the demand from the downstream stage has to be met within this service time. The outgoing service times are the decision variables in the model. Optimally setting the service times dictates the optimal safety stock placement in the supply chain.

In contrast to the guaranteed service model, Graves and Willems (2003) characterize the stochastic service (SS) model as one that does not guarantee service times. In the SS model, each stage maintains sufficient safety stock to satisfy a certain service level target. When demand exceeds the available inventory, the stage can be thought of as a server in a queueing network where the stage has a random replenishment process. In practice, as pointed out by Graves and Willems (2003), the GS and SS models represent two extreme views of what happens in the presence of a stockout. When a stockout occurs, the GS model assumes countermeasures outside the model are taken in order to ensure the service time is guaranteed whereas the SS model assumes the existing stochastic replenishment process can continue without any special countermeasures to remedy the stockout situation. In reality, the solution is somewhere between these two extremes. However, the GS model has been extensively deployed in practice, as documented in papers like Willems (2008) and it matches the operating realities of Cisco. Furthermore, Gallego and Zipkin (1999) consider a serial line supply chain under the unbounded demand assumption, and show that the optimal policy in Simpson (1958) can be near optimal. As such, we focus our attention on the GS model in this paper.

The GS inventory model has been further developed along several paths: adapting the model to general acyclic network structures, developing efficient algorithms, and relaxing the original model assumptions. Humair and Willems (2006) extended the model to networks with clusters of commonality, and then the work is further developed to general acyclic networks in Humair and Willems (2011). The line of the literature on developing algorithms include Graves (1988), which solves the original GS model as a dynamic program; Graves and Lesnaia (2004), which report a branch and bound algorithm for solving the optimal safety stock placement problem in a general network structure; Lesnaia et al. (2005), which show that the Graves and Willems (2000) algorithm in a general network structure is NP-hard; and Magnanti et al. (2006), which propose a tight relaxation of the original GS algorithm, and gain faster computational performance on commercial solvers. Lastly, on the works that modify the original GS model assumption to incorporate other supply chain issues, Minner (2001) considers the GS model in a supply chain where materials can be returned. Graves and Willems (2008) consider the model under non-stationary demand, and Klosterhalfen et al. (2014) apply the model to static dual sourcing. Farasyn et al. (2011) documents a successful application of the GS model at Proctor and Gamble. Tian et al. (2011) employ the GS model to develop an iterative approach to jointly solve the problems of tactical safety stock placement and tactical production planning.

In this paper, we expand the understanding of the GS model for safety stock optimization by analytically characterizing the structure of the optimal inventory policy in a two-stage serial line. Among existing GS research, this is the first paper that analytically characterizes the properties associated with the optimal safety stock placement policy in this framework.

This paper is structured as follows. In Section 3, the GS model is briefly reviewed, and we characterize the structure of the optimal safety stock policy, which is the building block of the following sections. Section 4 investigates how cost and leadtime allocation affect the optimal safety stock policy, by fixing the total cost and leadtime of the supply chain. Section 5 concludes the paper and presents future research directions.

3. GS safety stock model

The following assumptions underlie the GS model for a two-stage serial line supply chain. The supply chain is modeled as a network where stages represent stocking points and arcs represent the precedence relationship between stages. The upstream stage is denoted as Stage 2 while the downstream stage is denoted as Stage 1, as shown in Fig. 2.

There is a cost and leadtime associated with each stage, denoted as C_i and T_i $(i=1,2)$. Cost includes the direct material and labor costs associated with transforming the stage's input to output and leadtime refers to the required time to accomplish this transformation assuming the stage's raw materials are available.

Each stage quotes a nonnegative guaranteed outgoing service time, S_i $(i=1,2)$, which is the amount of time between receiving an order from a downstream stage and satisfying the order. Service time is guaranteed in the sense that as long as the demand is less than a specified upper bound, the order is committed to be delivered within the outgoing service time. Consistent with Cisco's practice, we assume Stage 1 must provide immediate delivery for external demand and Stage 2 has immediate access to raw materials; this is equivalent to assuming $S_1=0$ and the incoming service time quoted to Stage 2 is zero. These service time assumptions for the edges of the network are common in GS papers, and can easily be relaxed, as discussed in Graves and Willems (2000).

We assume external demand occurs only at Stage 1 and one unit from Stage 2 is required to satisfy one unit of demand at Stage 1. Furthermore, we assume the demand process at Stage 1 is stationary with average per period demand μ and standard deviation σ. Demand at a stage is bounded by the function $D(\tau)$ which specifies an upper bound on demand for any interval of length τ. Such a demand bound does not mean that demand can never exceed the bound but it does mean that the safety stock in the system is only designed to satisfy the demand within the bound. Assuming a service factor of k, $D(\tau)$ can be defined as $\mu\tau+k\sigma\tau^{1/2}$.

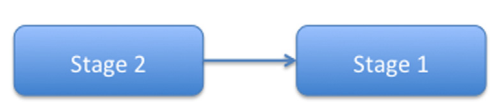

Fig. 2. In this two-stage serial line supply chain Stage 2 is the upstream stage supplying Stage 1. Stage 1 satisfies external demand.

This functional form is well documented to be the method the GS model implements $D(\tau)$ in practice. Specific examples are provided in Graves and Willems (2008) and Bossert and Willems (2007), while Willems (2008) presents 38 supply chains from 29 companies that all employ this functional form for specifying their demand.

The safety stock optimization problem **P** for this two-stage GS model can be formulated as

$$\textbf{P} \quad \min \alpha C_2 k\sigma\sqrt{T_2-S_2}+\alpha(C_1+C_2)k\sigma\sqrt{S_2+T_1}$$
$$\text{s.t.} \quad 0 \le S_2 \le T_2$$

where

C_i=Cost at stage i, i=1, 2
T_i=Leadtime at stage i, i=1, 2
S_2=Outgoing service time at stage 2. This is the decision variable in this two-stage network
α=Safety stock holding cost rate
k=Service factor chosen
σ=Standard deviation of demand per period

The first and second terms in the objective function represent the safety stock holding cost at Stage 2 and Stage 1, respectively. The only constraint requires the outgoing service time at Stage 2 to be nonnegative and to not exceed Stage 2's leadtime. In our two-stage serial-line supply chain, there are only two possible optimal safety stock policies. Either Stage 2 does not hold safety stock (i.e., $S_2=T_2$) which causes Stage 1 to hold sufficient safety stock for demand over the combined leadtimes from both stages or Stage 2 does hold safety stock (i.e, $S_2=0$) and each stage holds safety stock to cover demand over its own leadtime.

The safety stock required at a stage can also be determined by defining net replenishment time at a stage. The net replenishment time at Stage 2, calculated as T_2-S_2, represents the amount of time Stage 2 is responsible for covering with safety stock. When Stage 2's net replenishment time is zero (i.e., $S_2=T_2$), Stage 2 does not need to hold safety stock which necessitates Stage 1 carrying a safety stock of $k\sigma\sqrt{T_2+T_1}$. When Stage 2's net replenishment time equals Stage 2's leadtime, (i.e., $S_2=0$ so $T_2-S_2=T_2$), Stage 2 is responsible for meeting the demand over its leadtime which requires safety stock in the amount of $k\sigma\sqrt{T_2}$ at Stage 2 and $k\sigma\sqrt{T_1}$ at Stage 1.

To gain deeper insight into the two-stage problem, we employ the concept of cost allocation and leadtime allocation between the two stages. If the total cost for a unit is $C=C_2+C_1$, we let ω denote the percentage of the total cost allocated to Stage 2; i.e., $\omega=C_2/(C_2+C_1)$. Similarly, if the total leadtime in the supply chain is $T=T_1+T_2$, we let γ denotes the percentage of the total leadtime allocated to Stage 2; i.e., $\gamma=T_2/(T_2+T_1)$. Furthermore, we let ρ denotes the net replenishment time allocated to Stage 2, where $\rho=(T_2-S_2)/(T_2+T_1)$.

With these new variables defined, problem **P** can be re-written to problem **P'** as follows:

$$\textbf{P'} \quad \min \alpha\omega Ck\sigma\sqrt{\rho T}+\alpha Ck\sigma\sqrt{(1-\rho)T}$$
$$\text{s.t.} \quad 0 \le \rho \le \gamma$$

where

C=Total per unit cost, i.e. $C=C_1+C_2$
ω=Percentage of total per unit cost allocated at Stage 2; $\omega=C_2/C$
T=Total supply chain leadtime, i.e. $T=T_1+T_2$
γ=Percentage of total supply chain leadtime allocated at Stage 2; $\gamma=T_2/T$
ρ=Decision variable, percentage of total net replenishment time allocated at Stage 2; $\rho=(T_2-S_2)/T$
α=Safety stock holding cost rate
k=Service factor chosen
σ=Standard deviation of demand per period

Since **P'** is simply a rewriting of **P**, there remain two possible optimal safety stock policies: $\rho=0$ or $\rho=\gamma$. When $\rho=\gamma$, safety stock is held at both stages, forcing the system as a whole to hold more units of safety stock versus when $\rho=0$. Intuitively, to break the power of pooling at Stage 1, the cost at Stage 2 must be low enough that decoupling Stage 2's leadtime results in a lower total safety stock cost.

Lemma 1 quantitatively characterizes the optimal safety stock policy.

Lemma 1. *When* $1-\omega\sqrt{\gamma}-\sqrt{1-\gamma}\ge 0$, $\rho^*=\gamma$; *and when* $1-\omega\sqrt{\gamma}-\sqrt{1-\gamma}\le 0$, $\rho^*=0$, *where* ρ^* *is the optimal solution to* **P'**.

All proofs appear in Appendix. This lemma shows that the cost and leadtime allocation across the stages jointly determine the optimal safety stock inventory policy. This point will be further elaborated in the following sections.

4. Impact of system parameters on the safety stock placement

In this section, we fix the total unit cost and total leadtime for the two-stage network, and investigate the impact of cost allocation ω and leadtime allocation γ on the optimal safety stock policy and corresponding optimal total safety stock cost. Subsection 4.1 investigates the impact of cost allocation and Subsection 4.2 studies the impact of leadtime allocation.

4.1. Impact of cost allocation

Fig. 3 shows total safety stock cost as a function of Stage 2's net replenishment time allocation. The four lines correspond to four ways to allocate cost between the stages (ω_i=0.1, 0.4, 0.7 and 0.9). In this example, the holding cost rate is 45%, the average daily demand is 100 per day and the standard deviation of demand is 80. The safety factor k is 3. Total supply chain cost and total supply chain leadtime are chosen to be 100.

Fig. 3 shows that when $\gamma=0.6$, $\rho^*=\gamma$ for ω_1=0.1 and ω_2=0.4 and $\rho^*=0$ for ω_3=0.7 and ω_4=0.9. It can be intuited that there seem to exist a threshold of ω, lying between 0.4 and 0.7, which separates the optimal policies of Stage 2 holding and not holding safety stock inventory. Moreover, Fig. 2 also seems to suggest that, fixing all parameters other than ω, as ω increases the minimum safety stock cost steadily increases even though the optimal safety

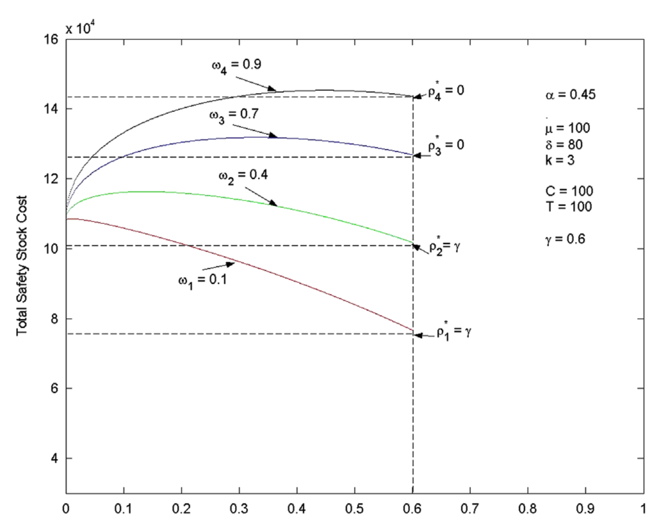

Fig. 3. Change of optimal safety stock policy with changing ω.

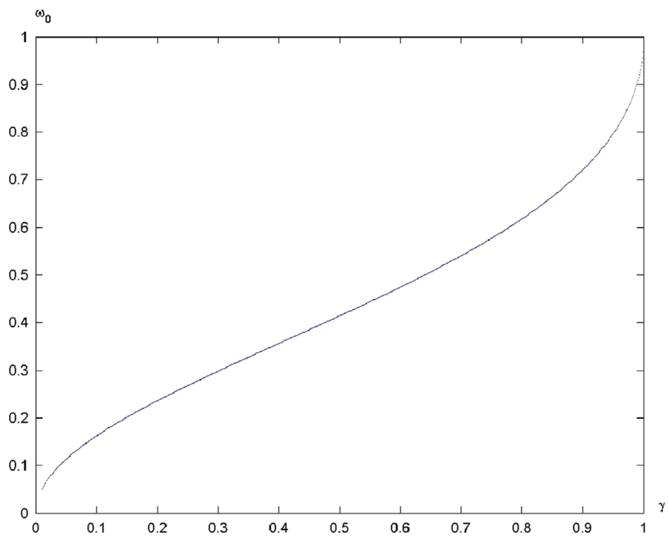

Fig. 4. ω_0 as a function of γ.

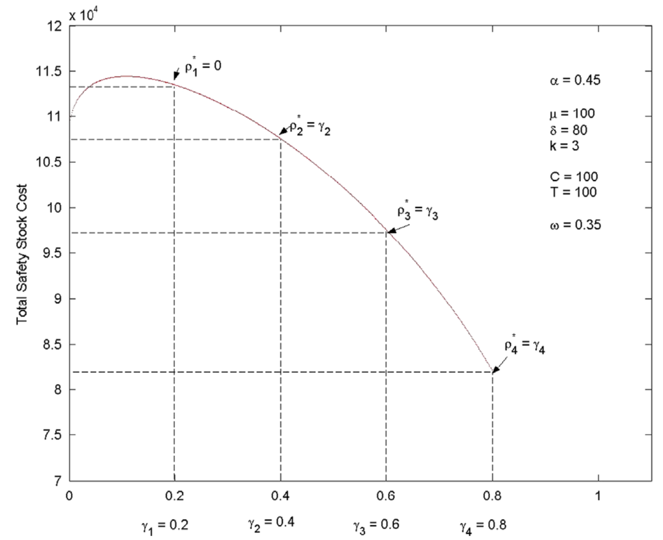

Fig. 5. Change of optimal safety stock policy with changing γ.

stock policy may change. Theorems 1 and 2 analytically prove those intuitions.

Theorem 1. *Fixing all other parameters and changing only the cost allocation ω, there exists $\omega_o = (1-\sqrt{1-\gamma})/\sqrt{\gamma}$, such that for $\omega \in [0, \omega_0]$, $\rho^* = \gamma$, and for $\omega \in [\omega_0, 1]$, $\rho^* = 0$.*

Theorems 1 shows that when the cost allocated to Stage 2 is less than a threshold based on the leadtime allocation between stages, it makes economic sense to break the power of pooling and place a decoupling stock at Stage 2.

Fig. 4 intuitively shows that ω_0 is an increasing function of γ, and Corollary 1 analytically proves it.

Corollary 1. *$\omega_0 = (1-\sqrt{1-\gamma})/\sqrt{\gamma}$ is an increasing function of γ.*

Corollary 1 analytically shows that when γ, the leadtime allocated to Stage 2, decreases, the cost threshold, ω_0, also decreases. Combining this result with Theorem 1, we can see that, in general, the benefit of decoupling Stage 2's leadtime is less possible to be realized when γ decreases, unless the cost allocated to Stage 2 is small enough. In other words, when leadtime allocated to Stage 2 decreases, unless the cost allocation at Stage 2 is sufficiently low, holding safety stock inventory at Stage 2 will result in a higher minimum total safety stock cost than holding no safety stock at Stage 2 and having Stage 1 maintain sufficient safety stock to cover the leadtime at both stages.

Theorem 2 investigates the impact of ω on total safety stock cost in the supply chain.

Theorem 2. *Keeping all parameters other than ω constant, the minimum total safety stock cost is an increasing function of ω.*

As ω increases, the optimal safety stock policy changes. When ω is between 0 and ω_0, Stage 2 holds safety stock inventory and the total safety stock cost increases as ω increases; when ω is between ω_0 and 1, Stage 2 does not hold safety stock inventory, and the total safety stock cost is a constant with regard to ω. As is shown in the proof, the total safety stock cost is greater when $\omega \in [\omega_0, 1]$ versus when $\omega \in [0, \omega_0]$. Thus the safety stock cost is nondecreasing as ω increases.

In summary, Theorem 2 shows that increasing the allocation of cost to Stage 2, while keeping all other parameters the same, increases the safety stock cost. From a managerial perspective, Theorem 2 shows that for two supply chain configurations with identical parameters other than the cost allocation, the one allocating less cost at Stage 2 always incurs a smaller total safety stock cost.

4.2. Impact of leadtime allocation

Keeping the cost allocation constant, we next investigate the impact of leadtime allocation on the optimal safety stock policy and total safety stock cost. Fig. 5 shows four allocations of leadtime between Stages 1 and 2 equal to 0.2, 0.4, 0.6 and 0.8. Cost allocation ω is chosen to be 0.35. All other parameters remain the same as in Fig. 3.

Based on Fig. 5, we can intuit similar results to those found in the previous subsection. Namely, a leadtime allocation threshold exists. However, the total safety stock cost seems to decrease as the leadtime allocated to Stage 2 increases, and this is the opposite effect compared to the impact of increasing cost allocation to Stage 2. Theorems 3 and 4 prove these observations analytically.

Theorem 3. *Fixing all other parameters in the supply chain and considering only changes to the leadtime allocation γ, then for any cost allocation ω there exists $\gamma_0 = 4\omega^2/(\omega^2+1)^2 \in [0, 1]$, such that for $\gamma \in [0, \gamma_0]$, $\rho^* = 0$, and for $\gamma \in [\gamma_0, 1]$, $\rho^* = \gamma$.*

If the leadtime allocated to Stage 2 is small enough, then it is optimal not to hold safety stock at Stage 2. But if Stage 2's leadtime exceeds the threshold, then it is optimal to hold safety stock at Stage 2. This result is opposite the result of Theorem 1 in the sense that a greater cost allocation indicates not holding safety stock while a greater leadtime allocation indicates holding safety stock at Stage 2.

Taken together, Theorems 1 and 3 demonstrate that to break the power of pooling generated by holding safety stock inventory at only Stage 1, the cost at Stage 2 has to either be relatively cheap or its leadtime has to be relatively long.

Fig. 6 intuitively shows the leadtime threshold γ_0 as an increasing function of cost allocation ω, and Corollary 2 analytically proves it.

Corollary 2. *The threshold γ_0 is an increasing function of ω.*

In the same spirit of Theorem 1, we can say that $1-\gamma_0$ measures the size of the leadtime allocation set where it is optimal for Stage 2 to hold safety stock inventory. Therefore, Corollary 2 analytically shows that when the cost allocated at Stage 2 increases, in general, the benefit of decoupling the total supply chain

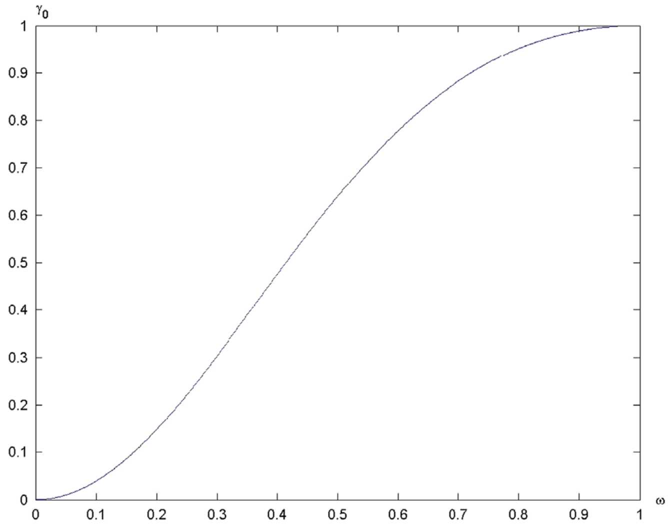

Fig. 6. Leadtime threshold γ_0 as a function of cost allocation ω.

leadtime is less possible to be realized, unless the leadtime allocated to Stage 2 is also fairly large.

Corollaries 1 and 2, taken together, show that leadtime and cost jointly determine the optimal safety stock policy. When cost allocated at Stage 2 increases, it is possible but harder for leadtime to break the power of pooling created by holding safety stock at only Stage 1; at the same time, when leadtime allocated at Stage 2 increases, it is possible but harder for cost allocation to realize any inventory cost savings from decoupling the supply chain safety stock created by holding safety stock at Stage 2.

Corollaries 1 and 2 provide analytical justification for Cisco's supply chain practices as described in Section 1. For phone supply chains, where the cost allocated at the Super Market (corresponding to Stage 2 in our model) is low enough there is no need to hold safety stock at the super market. While for low-end routers, which have higher Stage 2 cost and a longer leadtime, it makes sense to hold safety stock at the Super Market.

Theorem 4. *Keeping all the parameters other than γ the same, the total safety stock cost is a decreasing function of γ.*

Proof of Theorem 4 demonstrates that although the inventory placement in the supply chain varies as γ increases from 0 to 1, the optimal total safety stock cost is decreasing. In practice, Theorem 4 shows that for two supply chains with identical parameters other than the leadtime allocation, the one with more leadtime incurred at Stage 2 always has lower total safety stock cost.

Taken together, Theorems 2 and 4 provide several insights. First, fixing all system parameters other than leadtime and cost allocation, the configuration that minimizes the total safety stock cost assigns as little cost as possible and as much leadtime as possible to Stage 2. Second, with total cost and total leadtime fixed, there would never be a benefit to shift the cost allocated to Stage 2 from Stage 1 in exchange for decreasing Stage 2's leadtime allocation if one wants to reduce the safety stock cost.

Linking the insights back to the Cisco example, Theorems 2 and 4 provides some potential directions for safety stock cost reduction. To reduce total safety stock cost, it would be beneficial for Cisco to further examine options that assign as little cost as possible and as much leadtime as possible to Stage 2. Also, decreasing Stage 2's cost allocation while increasing its leadtime allocation will never result in a supply chain with lower safety stock cost. This second result was a real learning for Cisco's procurement team.

5. Conclusions

In this paper, we consider a two-stage serial line supply chain and employ the GS safety stock inventory model to determine the optimal safety stock placement policy and safety stock cost. We then analytically characterize the supply chain safety stock decision with regard to supply chain cost and leadtime changes.

We start by varying the cost and leadtime allocation when fixing the total cost and leadtime in the supply chain. We show that only when the cost at Stage 2 is relatively low or its leadtime is relatively long, is it worthwhile to decouple the total supply chain leadtime and hold safety stock inventory at Stage 2. We also show that when we fix all supply chain parameters other than cost and leadtime allocation, the way to reduce the total safety stock cost is to increase the leadtime allocation and decrease the cost allocation at Stage 2 as much as possible.

This research fills a gap in the GS literature by analytically characterizing the GS model's optimal safety stock policy with regard to supply chain cost and leadtime. It still has great potential for further development. First, it would be interesting to analytically consider other network structures such as simple distribution networks and assembly networks. Second, it would be interesting to try to relax several key model assumptions, such as the bounded demand assumption and guaranteed service assumption, and consider whether the derived insights still hold.

Appendix: Proofs

Proof of Lemma 1. It is easy to show that for $0 \le \rho \le \gamma$, $\frac{d^2 f(\rho)}{d\rho^2} = -(\frac{1}{4})\omega\rho^{-\frac{3}{2}} - (\frac{1}{4})(1-\rho)^{-3/2} < 0$. Therefore, $f(\rho) = \omega\sqrt{\rho} + \sqrt{1-\rho}$ is a concave function of ρ.Thus $\alpha\omega Ck\sigma\sqrt{\gamma T} + \alpha Ck\sigma\sqrt{(1-\gamma)T} \le \alpha Ck\sigma\sqrt{T}$, i.e. $\omega\sqrt{\gamma} + \sqrt{1-\gamma} \le 1$, we have $\rho^* = \gamma$. Otherwise, $\rho^* = 0$. □

Proof of Theorem 1. For any fixed γ, choose $\omega_0 \in [0, 1]$ such that $1 - \omega_0\sqrt{\gamma} - \sqrt{1-\gamma} = 0$.

Therefore, for $\omega \in [0, \omega_0]$, we have $1 - \omega\sqrt{\gamma} - \sqrt{1-\gamma} \ge 1 - \omega_0\sqrt{\gamma} - \sqrt{1-\gamma} = 0$. According to Lemma 1, $\rho* = \gamma$.

Similarly, when $\omega \in [\omega_0, 1]$, we have $1 - \omega\sqrt{\gamma} - \sqrt{1-\gamma} \le 1 - \omega_0\sqrt{\gamma} - \sqrt{1-\gamma} = 0$, by Lemma 1, $\rho* = 0$. □

Proof of Corollary 1. By definition, $\omega_0 = (1 - \sqrt{1-\gamma})/\sqrt{\gamma}$, and when $\gamma \in (0, 1)$, $\frac{\partial\omega_0}{\partial\gamma} = \frac{1-\sqrt{1-\gamma}}{\gamma\sqrt{\gamma}\sqrt{1-\gamma}} > 0$, so ω_0 is an increasing function of $\gamma \in (0, 1)$. □

Proof of Theorem 2. To prove this result, we discuss $\omega \in [\omega_0, 1]$, $\omega \in [0, \omega_0]$ and compare when $\omega_1 \in [0, \omega_0]$ and $\omega_2 \in [\omega_0, 1]$.

Firstly, when $\omega \in [\omega_0, 1]$, $\rho* = 0$, so the total safety stock cost can be expressed as $\alpha Ck\sigma\sqrt{T}$, and it is not a function of ω.

Secondly, when $\omega \in [0, \omega_0]$, $\rho* = \gamma$, so the total safety stock cost can be expressed as $\alpha\omega Ck\sigma\sqrt{\gamma T} + \alpha Ck\sigma\sqrt{(1-\gamma)T}$, and it is obviously an increasing function of ω.

Thirdly, when $\omega_1 \in [0, \omega_0]$ and $\omega_2 \in [\omega_0, 1]$, we have

$$\alpha\omega_1 Ck\sigma\sqrt{\gamma T} + \alpha Ck\sigma\sqrt{(1-\gamma)T} \le \alpha Ck\sigma\sqrt{T}.$$

The inequality holds because $\rho* = \gamma$ when $\omega_1 \in [0, \omega_0]$.

Therefore, the conclusion of this theorem holds. □

Proof of Theorem 3. Denote $f(\rho) = \alpha\omega Ck\sigma\sqrt{\rho T} + \alpha Ck\sigma\sqrt{(1-\rho)T}$, and define

$$g(\gamma) = f(\gamma) - f(0) = \alpha\omega Ck\sigma\sqrt{\gamma T} + \alpha Ck\sigma\sqrt{(1-\gamma)T} - \alpha Ck\sigma\sqrt{T}.$$

Therefore $g(0)=0$.

And we have $\frac{\partial^2 g(\gamma)}{\partial\gamma^2} = -\alpha\omega Ck\sigma\frac{\sqrt{T}}{4\sqrt{\gamma^3}} - \alpha Ck\sigma\frac{\sqrt{T}}{4\sqrt{(1-\gamma)^3}} \le 0$.

Thus $g(\gamma)$ is a concave function of $\gamma \in [0, 1]$.

Choose $\gamma_0 \in [0, 1]$, such that $f(\gamma_0) = f(0)$, i.e., $g(\gamma_0) = g(0) = 0$.

According to the concavity of $g(\gamma)$, when $\gamma \in [0, \gamma_0]$, $g(\gamma) \geq 0$, i.e., $f(\gamma) - f(0) \geq 0$, therefore $\rho* = 0$; and when $\gamma \in [\gamma_0, 1]$, $g(\gamma) \leq 0$, therefore $\rho* = \gamma$. □

Proof of Corollary 2. By definition, γ_0 satisfies $1 - \omega\sqrt{\gamma_0} - \sqrt{1-\gamma_0} = 0$, therefore, $\gamma_0 = 4\omega^2/(\omega^2+1)^2$. This is a monotinic increasing function of $\omega \in [0, 1]$ because when $\omega \in [0, 1]$, $\frac{\partial \gamma_0}{\partial \omega} = 8\frac{\omega - \omega^3}{(\omega^2+1)^3} \geq 0$. □

Proof of Theorem 4. We will discuss when $\gamma \in [0, \gamma_0]$, when $\gamma \in [\gamma_0, 1]$ and compare when $\gamma_1 \in [0, \gamma_0]$ $\gamma_2 \in [\gamma_0, 1]$.

Firstly, when $\gamma \in [0, \gamma_0]$, $\rho* = 0$, . So the safety stock cost is $\alpha Ck\sigma\sqrt{T}$, and it is not a function of γ.

Second, when $\gamma \in [\gamma_0, 1]$, $\rho* = \gamma$, so the total safety stock cost is $\alpha\omega Ck\sigma\sqrt{\gamma T} + \alpha Ck\sigma\sqrt{(1-\gamma)T}$ we need to show that this is a decreasing function of γ for $\gamma \in [\gamma_0, 1]$.

To see this, first we have that

$$\frac{\partial^2(\alpha\omega Ck\sigma\sqrt{\gamma T} + \alpha Ck\sigma\sqrt{(1-\gamma)T})}{\partial \gamma^2} = -\alpha\omega Ck\sigma\frac{\sqrt{T}}{4\sqrt{\gamma^3}} - \alpha\omega Ck\sigma\frac{\sqrt{T}}{4\sqrt{(1-\gamma)^3}} \leq 0,$$

so this is a concave function of $\gamma \in [\gamma_0, 1]$, and its maximum is reached when

$$\frac{\partial(\alpha\omega Ck\sigma\sqrt{\gamma T} + \alpha Ck\sigma\sqrt{(1-\gamma)T})}{\partial \gamma} = \alpha\omega Ck\sigma\frac{\sqrt{T}}{2\sqrt{\gamma}} - \alpha Ck\sigma\frac{\sqrt{T}}{2\sqrt{1-\gamma}} = 0,$$

i.e., $\omega\frac{1}{\sqrt{\gamma}} - \frac{1}{\sqrt{1-\gamma}} = 0$, i.e., when $\hat{\gamma} = \frac{\omega^2}{1+\omega^2}$. So the total safety stock cost is a decreasing function for all $\gamma \in [\hat{\gamma}, 1]$. Notice that $\gamma_0 = \frac{4\omega^2}{(1+\omega^2)^2} > \frac{\omega^2}{1+\omega^2} = \hat{\gamma}$, so the total safety stock cost $\alpha\omega Ck\sigma\sqrt{\gamma T} + \alpha Ck\sigma\sqrt{(1-\gamma)T}$ is a decreasing function when $\gamma \in [\gamma_0, 1]$.

Thirdly, we consider the scenario when we have $\gamma_1 \in [0, \gamma_0]$, $\rho* = 0$ and $\gamma_2 \in [\gamma_0, 1]$ $\rho* = \gamma$. For $\gamma_1 \in [0, \gamma_0]$, the total safety stock cost is $\alpha Ck\sigma\sqrt{T}$, but for $\gamma_2 \in [\gamma_0, 1]$, the total safety stock cost is $\alpha\omega Ck\sigma\sqrt{\gamma_2 T} + \alpha Ck\sigma\sqrt{(1-\gamma_2)T}$, and it is obvious that $\alpha\omega Ck\sigma\sqrt{\gamma_2 T} + \alpha Ck\sigma\sqrt{(1-\gamma_2)T} \leq \alpha Ck\sigma\sqrt{T}$ because $\gamma_2 \in [\gamma_0, 1]$.

Therefore, in summary, we have proved that the total safety stock cost is a decreasing function of γ.

References

Bossert, J.M., Willems, S.P., 2007. A periodic review modeling approach for guaranteed service supply chains. Interfaces 37 (5), 420–435.

Farasyn, I., Humair, S., Jahn, J.I., Neale, J.J., Rosen, O., Ruark, J., Tarlton, W., de Velde, W.V., Wegryn, G., Willems, S.P., 2011. Inventory optimization at procter & gamble: achieving real benefits through user adoption of inventory tools. Interfaces 41 (1), 66–78.

Federgruen, A., 1993. Centralized planning models for multi-echelon inventory systems under uncertainty. In: Graves, S.C., Rinnooy Kan, A.H., Zipkin, P.H. (Eds.), Handbooks in Operations Research and Management Science, Vol 4. North-Holland, Amsterdam, Elsevier Science, B.V..

Gallego, G., Zipkin, P., 1999. Stock positioning and performance estimation in serial production-transportation systems. Manuf. Serv. Oper. Manag. 1, 77–88.

Graves, S.C., 1988. Safety stocks in manufacturing systems. J. Manuf. Oper. Manag. 1, 67–101.

Graves, S.C., Lesnaia, E., 2004. Optimizing safety stock placement in general network supply chains. In: Proceedings of Singapore-MIT Alliance Annual Symposium Conference. Singapore.

Graves, S.C., Willems, S.P., 2000. Optimizing strategic safety stock placement in supply chains. Manuf. Serv. Oper. Manag. 2 (1), 68–83.

Graves, S.C., Willems, S.P., 2003. Supply chain design: safety stock placement and supply chain configuration. In: de Kok, A.G., Graves, S.C. (Eds.), Handbooks in Operations Research and Management Science, Vol. 11. North-Holland Publishing Company, Amsterdam, The Netherlands, Supply Chain Management: Design, Coordination and Operation.

Graves, S.C., Willems, S.P., 2008. Strategic inventory placement in supply chains: nonstationary demand. Manuf. Serv. Oper. Manag. 10 (2), 278–287.

Humair, S., Willems, S.P., 2006. Optimizing strategic safety stock placement in supply chain with clusters of commonality. Oper. Res. 54 (4), 725–742.

Humair, S., Willems, S.P., 2011. Technical note: optimizing strategic safety stock placement in general acyclic networks. Oper. Res. 59 (3), 781–787.

Klosterhalfen, S.T., Minner, S., Willems, S.P., 2014. Strategic safety stock placement in assembly networks with dual sourcing. Manuf. Serv. Opera. Manag. 16 (2), 204–219.

Lesnaia, E., Vasilescu, I., Graves, S.C., 2005. The complexity of safety stock placement in general-network supply chains. In: Proceedings of Singapore-MIT Alliance Annual Symposium Conference Singapore.

Magnanti, T.L., Shen, Z.M., Shu, J., Simchi-Levi, D., Teo, C., 2006. Inventory placement in acyclic supply chain networks. Oper. Res. Lett. 34, 228–238.

Minner, S., 2001. Strategic safety stocks in reverse logistics supply chains. Int. J. Prod. Econ. 71, 417–428.

Porteus, E.L., 2002. Foundations of stochastic inventory theory. Stanford University Press, Stanford, CA, USA.

Simpson, K.F., 1958. In-process inventories. Oper. Res. 6, 863–873.

Tian, F., Willems, S.P., Kempf, K.G., 2011. An iterative approach to item-level tactical production and inventory planning. Int. J. Prod. Econ. 133 (1), 439–450.

Willems, S.P., 2008. Data set: real-world multi-echelon supply chains used for inventory optimization. Manuf. Serv. Oper. Manag. 10 (1), 19–23.