

Building Supply Chain Resilience through Virtual Stockpile Pooling

Fang Liu

50 Nanyang Ave. Singapore, 639798, liu_fang@ntu.edu.sg, Phone: +65-6790-6318, Fax: +65-6792-2313

Jing-Sheng Song

100 Fuqua Drive. Durham, NC, 27708, USA, jssong@duke.edu, Phone: +1-919-660-7828, Fax:
+1-919-681-6245

Jordan D Tong

975 University Ave. Madison, WI, 53706, USA, jordan.tong@wisc.edu, Phone: +1-608-265-8220, Fax:
+1-608-263-3142

Abstract: Stockpiling inventory is an essential strategy for building supply chain resilience. It enables firms to continue operating while finding a solution to an unexpected event that causes a supply disruption or demand surge. While extremely valuable when actually deployed, stockpiles incur large holding costs and usually provide no benefits until such a time. To help to reduce this cost, this paper presents a new approach for managing stockpiles. We show that if leveraged intelligently, stockpiles can also help an organization better meet its own regular demand by enabling a type of virtual pooling we call virtual stockpile pooling (VSP). The idea of VSP is to first integrate the stockpile into several locations' regular inventory buffers and then dynamically reallocate the stockpile among these locations in reaction to the demand realizations to achieve a kind of virtual transshipment. To study how to execute VSP and determine when it can provide the most value, we formulate a stylized multi-location stochastic inventory model and solve for the optimal stockpile allocation and inventory order policies. We show that VSP can provide significant cost savings: in some cases nearly the full holding cost of the stockpile (i.e., VSP effectively maintains the stockpile for free), in other cases nearly the savings of traditional physical inventory pooling. Last, our results prescribe implementing VSP with many locations for large stockpiles, but only a few locations for small stockpiles.

Keywords: Supply chain disruption risk management, demand surge, multi-location inventory model, inventory pooling, transshipment

Received: August, 2015; accepted April, 2016 by Albert Ha after two revisions.

1 Introduction

Creating redundancy is a key component to most firms' strategies for increasing a supply chain's resilience. Redundancy provides firms the capability to continue operating while finding a solution to an unexpected supply chain event that causes a supply disruption or demand surge (Sheffi 2005). One of the most common forms of redundancy is inventory stockpiling (Tomlin and Wang 2011). Hospitals and pharmaceutical companies stockpile medical supplies and drugs for disaster-response (Sheffi 2005, Shen et al. 2010, Adida et al. 2011); manufacturers stockpile key components to hedge against supply disruptions (Maddah et al. 2014); financial organizations stockpile cash to hedge against bank runs ("Reserve Requirements" 2014); and energy companies are required to maintain oil reserves to hedge against supply disruptions ("IEA" 2007). Although highly valuable to provide resilience in case of rare events, maintaining these stockpiles can be extremely costly on a normal basis. For example, the US government invests over \$500 million annually for the strategic stockpiling of vaccinations and medications ("HHS Budget in Brief" 2015). This paper presents an idea - *virtual stockpile pooling* (VSP) - to help solve this dilemma by significantly reducing the cost of stockpiling inventory.

To explain the idea of VSP, it is helpful to first discuss existing strategies for maintaining stockpiles in increasing order of sophistication. The most straightforward approach is a *dedicated stockpile*: simply set aside a constant amount of inventory in a dedicated warehouse or storage area. Such an approach is also common, perhaps because the individuals responsible for risk management often belong to a different team (such as a crisis management team), separated from the operational managers who are responsible for managing inventory for routine random demand fluctuations. This separation is also evident in conventional inventory models, which typically focus on the latter type of demand uncertainty isolated from other types of risk.

Often a more cost-efficient approach, however, is an *integrated stockpile*: commingle the stockpile into the inventory buffers already used in the supply chain. Several companies have begun to use this approach. For example, Cisco establishes the sizes and locations of inventory buffers of key products as part of their supply chain risk management program (Anupindi 2012). Similarly, Johnson and Johnson (J&J) holds stockpiles of inventory for the US government in the form of additional inventory buffers which are integrated with their regular inventory (Sheffi 2005). The primary benefit of integrating the stockpile is that it is cycled due to regular demand, thereby reducing expiration costs. Consider how J&J executes an integrated stockpile, discussed in Sheffi (2005). As part of its disaster preparedness program, the US government contracts with J&J,




Figure 1: Depiction of static stockpiles and virtual stockpile pooling

requiring them to have enough inventory be able to fill a large order of a certain size at any time. In turn, J&J establishes a “red-line” at each of its various plants and distribution centers and follows a Sell-One-Stock-One (SOSO) policy for this red-line inventory at that location. If the total inventory level ever falls to this red-line, the location must report stock-out to its regular customers. By setting the red-lines at each location so that the sum is large enough, J&J ensures they can fulfill the government’s contract. Moreover, the stockpile is continually being rotated with new inventory so that the stockpiled goods do not expire.

Although under J&J’s practices the stockpile is rotated, it still “cannot be used to compensate for day-to-day variations [...] J&J’s everyday processes have to operate as if such inventory does not exist” (Sheffi 2005 pp. 173-174). The underlying reason for this inability is that although the stockpile is integrated, it is still a *static stockpile* (SS). The focus of this paper is to present a new, *dynamic* strategy we call VSP, which seeks to improve upon J&J’s red-line practice so that the integration of the stockpile not only reduces the cost of expiration for the stockpile, but also helps J&J to cope with its regular demand variations.

The key idea behind VSP is to dynamic allocate the integrated stockpile amongst multiple locations. It is based on a simple observation, depicted in Figure 1: one location can go below their red-line by a certain amount if another location can “cover” for them by increasing their red-line by an equal amount. Consequently, the stockpile can be leveraged to execute a type of virtual transshipment of inventory from one location to another. In the figure, location 1 can “transship” inventory to location 2 by increasing its red-line (thereby having less inventory available to sell to regular demand) so that location 2 can decrease its red-line (thereby having more inventory available to sell to its regular demand.) In this way, if J&J intelligently changes its red-line practice from being static to dynamic, such virtual transshipments may enable J&J to achieve an inventory

pooling effect.

Now, a natural question is, how much cost savings can VSP provide over SS? Also, is there a relationship between the cost savings from VSP and traditional inventory pooling? These questions raise another: what is the optimal way to execute VSP? To address these questions, we develop and analyze two multi-location stochastic inventory models with periodic review - one for SS and one for VSP. Under SS, the allocation of the red-lines at each location are fixed and each location follows a SOSO policy so as to never violate the red-line. The resulting optimal inventory order policy and inventory cost can be obtained by adapting standard inventory theory. We establish these results in Section 4, after reviewing related literature and describing the model setting in Sections 2 and 3.

In contrast, the optimal way to execute VSP is much more challenging. Under VSP, a reallocation of red-lines can occur after demands are realized each period but before they must be satisfied. This *virtual* reallocation problem presents analytical obstacles similar to those faced by the *physical* inventory transshipment literature. For this reason, several of our modeling assumptions follow a typical model setup in that literature (e.g., Robinson 1990, Archibald and Thomas 1997). We are able to show that the optimal stockpile (re)allocation policy in each period is myopic. It turns out that the single-period allocation problem is a convex program. By analyzing its dual problem we develop a greedy algorithm to solve it. Interestingly, the optimal reallocation prioritizes allocating more stockpile to locations based on their holding plus backorder cost parameters - the holding cost is not weighted more. Furthermore, we show it can be executed by locations forming a chain based on their holding and backorder costs and, every period, each location needs only to provide one neighbor at most two numbers. We are then able to characterize the optimal inventory replenishment policy, showing it is a stationary base-stock policy. The optimal base-stock levels are not independent across locations. Nevertheless, we can derive intuitively-attractive bounds and other qualitative properties and develop an algorithm to compute them. These findings are presented in Section 5.

Lastly, by comparing the lowest possible cost under SS to the optimal cost under VSP, we establish a lower-bound on the potential value of VSP. Through limiting behaviors and numerical examples, we characterize the factors that determine the value of VSP. We show that VSP can provide significant cost savings over SS and we identify when these savings are greatest. In particular, if the stockpile is small enough or the number of locations is large enough, then VSP provides cost savings equal to the full holding cost of the stockpile. In other words, VSP effectively maintains the stockpile for free under these conditions. Remarkably, the majority of these cost savings can often be achieved even with a reasonably small number of locations. On the other hand, when the stockpile is large enough relative to the number of locations, VSP can achieve the cost savings of traditional

physical inventory pooling. That is, VSP enables the full benefits of physical inventory pooling without incurring physical transportation costs under these conditions. Interestingly, we find that the majority of these cost savings can often be achieved even with a reasonably small stockpile. Furthermore, to achieve the most cost savings per location, we find that one should implement VSP with many locations for a large stockpile, but only a few locations for a small stockpile. Finally, like traditional physical inventory pooling, we prove that VSP savings are decreasing in demand correlation. These findings are presented in Section 6. We conclude in Section 7. The Appendix contains selected proofs. The online Supplementary Appendix contains discussions on various extensions.

2 Related Literature

Building supply chain resilience through operational levers has received tremendous attention in the operations management literature in recent years. A dominant theme of these levers is to invest in flexibility. See, for example, Van Mieghem (2003, 2011), Tomlin and Wang (2005, 2011), and the references therein. In the domain of inventory management, such flexibility is characterized by dual (or multiple) sourcing. Many researchers have studied how to effectively deploy the alternative supply sources in inventory replenishment for regular demand at a *single* location; recent works include Tomlin (2006), Veeraraghavan and Scheller-Wolf (2008), Song and Zipkin (2009), Allon and Van Mieghem (2010), Sheopuri et al. (2010), Hua et al. (2015), and Janakiraman et al. (2014). Different from the focus of these studies, the current paper concerns another lever – inventory stockpiling. While the management of a stockpile at a single location can follow standard inventory theory, as shown in Section 4, we are interested reducing the stockpile cost by dynamically managing stockpiles across multiple locations. To our knowledge, there is no prior work in the literature that studied such a problem.

The “red-line” concept for maintaining stockpile in any given location is very similar to the rationing policy for managing inventory with multiple demand classes studied by many scholars. See, e.g., Topkis (1968), Ha (1997), de Vèricourt et al. (2002), Gayon et al. (2009), and references therein. A rationing policy specifies inventory thresholds called the rationing levels, one for each demand class. A unit of inventory can be used to satisfy a demand from a particular class only if the inventory level is above the rationing level for that demand class. The rationing levels are in descending order of the importance (i.e., priority) of the demand class. In the integrated stockpile system, there are two demand classes – the regular demand and the disruption demand – with the latter being more important. The red-line thus corresponds to the rationing level for disruption

demand. What is different in the present paper is that the disruption demand rarely occurs, so it is very expensive to always reserve this stock. The VSP idea is precisely to increase the utility of these reserves by virtually pooling them across locations.

Our model is closely related to the literature on inventory pooling, which represents a large stream of research within operations management beginning with Eppen’s (1979) seminal paper. That work derives the value of consolidating inventory in a multi-location newsvendor problem, also showing that the value increases as locations become less correlated. While Eppen leverages normal demand distributions for tractability, following works have been able to generalize inventory pooling results. For example, Eppen and Schrage (1981) extend pooling results to include positive lead times. Federgruen and Zipkin (1984) study more general settings including other types of demand distributions and non-identical retailers, also with positive lead times. Corbett and Rajaram (2006) provide a generalization of Eppen (1979)’s results to non-normal demand. Other representative works include Gerchak and Mossman (1992), Anupindi and Bassok (1999), Alfaro & Corbett (2003), Benjaafar et al. (2005), Yang and Schrage (2009), Berman et al. (2011). Alptekindglu et al. (2013) studies a pooled inventory system where a supplier must meet different minimum service levels (in-stock probabilities) for different customers. Our paper differs in that there is a system-wide constraint that can be satisfied using inventory from any location.

The manner in which VSP achieves inventory-pooling benefits is related to physical inventory transshipment (see, e.g., Robinson 1990; Archibald and Thomas 1997; Rudi et al. 2001; Axsater 2003; Dong and Rudi 2004; Zhang 2005; Herer and Tzur 2006; Susic 2006; Hu et al. 2008; Zhao 2008; Huang and Susic 2010). Under inventory transshipment, pooling benefits are achieved by sometimes serving demand at one location using inventory at another (see Paterson et al. 2011 for a review). These papers typically consider a cost for physical transshipment, which limits the number of physical transshipments that are desirable to make such that full inventory-pooling benefits are not reached. In a similar vein, VSP does not allow unlimited “virtual transshipments” between locations because in order to conduct “virtual transshipment” one location must have enough stockpile available to deplete and the other location must have enough inventory to increase her stockpile. However, our paper differs from the transshipment literature in that we do not consider the physical transshipment of any real goods, but rather the virtual transfer of stockpile, which have no associated physical transshipment costs and affect the inventory dynamics differently from physical inventory transshipments.

Finally, we note that the practice of stockpiling is an important component of disaster preparedness, and has therefore been a topic of interest in humanitarian operations. Perhaps most closely




Figure 2: Order of events in a period.

related is in this literature is Shen et al. (2010), who examines the management of perishable inventory in a single-location deterministic Economic Manufacturing Quantity model with a minimum volume constraint, taking motivation from contracts between pharmaceutical companies and the US government. We refer the reader to Ergun et al. (2010), Van Wassenhove (2005), and Starr and Van Wassenhove (2014) for more general reviews of humanitarian operations and disaster supply chain management. While VSP may be of interest to those studying humanitarian supply chains, our model primarily takes the viewpoint of the firm and does not take into account any specific features of emergency response.

3 Model Setting & Review

First, we describe the model without stockpiles and review some standard results. We consider a single product that has a regular demand at multiple locations. There are N locations for the item, identified by superscript i , $i = 1, 2, \dots, N$. Inventory at each location is managed using a periodic-review system over an infinite horizon. A “location” can be thought of generally. For example, it can be at a manufacturing plant, a warehouse, or a distribution center.

The events in each period are depicted in Figure 2, which are consistent with commonly-made assumptions in the transshipment literature. Inventory is ordered at the beginning of each period from an outside supplier with ample supply. Demand is collected from customers throughout the period. Then the replenishment order arrives. At the end of the period, the current-period demand and backorders are fulfilled as much as possible from the on-hand inventory; the remaining demand is backlogged. Location i ’s demands across periods are i.i.d. random variables D^i with a finite mean μ^i and distribution F^i . Location i incurs a unit holding cost rate h^i and a unit backorder cost rate b^i . We assume $0 < h^i < b^i$ for all i . The objective is to minimize the long-run average holding and backorder costs.

For any two real numbers p and q , define $p \vee q = \max\{p, q\}$, $p \wedge q = \min\{p, q\}$, $p^+ = \max\{p, 0\}$ and $q^- = \max\{-q, 0\}$. Throughout the paper bold letters denote $N \times 1$ column vectors. For example, $\mathbf{h} = (h^i)_{i=1}^N$ and \mathbf{h}^T denotes the transpose. The vector \mathbf{e} denotes the vector of ones and \mathbf{e}^i denotes the vector of zeros with a one in the i th location. For any two vectors $\mathbf{p} = (p_i)$ and $\mathbf{q} = (q_i)$, $\mathbf{p} \cdot \mathbf{q} = \sum_{i=1}^N p_i q_i$, the inner product of the two vectors. All other vector operations are component-wise operations. For example, $\mathbf{p} \vee \mathbf{q} = (p_i \vee q_i)$.

3.1 No Stockpile

Before considering stockpiles, it will be convenient to review some well-known inventory results when there is no stockpiling. For any given period, let y be the net inventory after orders arrive. The expected one-period holding and backorder cost at location i is

$$C^i(y) = E[\widehat{C}^i(y, D^i)], \quad \text{where} \quad (1)$$

$$\widehat{C}^i(y, d) = h^i(y - d)^+ + b^i(y - d)^-. \quad (2)$$

Under these model assumptions it is well known that a stationary myopic basestock policy at each location is optimal (see, for example, Zipkin 2000). In particular, the optimal basestock level for location i , denoted as $\bar{y}^i = \arg \min C^i(y) \geq 0$, can be determined according to

$$P(D^i < \bar{y}^i) = \frac{b^i}{h^i + b^i}.$$

The optimal long-run average cost for location i is $J^{\mathcal{B}, i} = C^i(\bar{y}^i)$. The optimal cost over all of the locations is

$$J^{\mathcal{B}} = \sum_{i=1}^N J^{\mathcal{B}, i} = C(\bar{\mathbf{y}}), \quad (3)$$

where $C(\mathbf{y}) = \sum_{i=1}^N C^i(y^i)$.

The following special cases will also help us make clear and more intuitively appealing comparisons. Locations are *normal* if their demand distributions are normally distributed $D^i \sim \text{Normal}(\mu^i, \sigma^i)$. They are *symmetric* if they have identical costs and marginal demand distribution functions $h^i = h$, $b^i = b$, $D^i = D$. They are *independent* if their demands are independent. We will use subscripts n , s , and i under equalities to denote these special conditions. For example, $\stackrel{n,s}{=}$ denotes equality under the special case that locations are normal and symmetric (but not necessarily independent). Let ϕ and Φ denote the standard normal density and cumulative distribution functions, respectively. It is well known that the optimal basestock level for each location i under

the normal and symmetric assumption is

$$\bar{y}_{n,s}^{\mathcal{B},i} = \mu + \sigma z^*, \quad \forall i, \quad (4)$$

where $z^* = \Phi^{-1}(\frac{b}{h+b})$. The associated optimal total long-run average cost over all of the locations is

$$J_{n,s}^{\mathcal{B}} = N\sigma(h+b)\phi(z^*). \quad (5)$$

3.2 Physical inventory pooling

Furthermore, although locations are *not* physically pooled in this model setting, it will be convenient to imagine the performance of the system *as if it were* physically pooled. Under the well-known physical inventory pooling scenario (Eppen 1979), demands from all locations are satisfied using one consolidated pool of inventory. Therefore, there is one demand $D^{\mathcal{T}} = \sum_i^N D^i$ and one basestock level Υ . The system's expected one period inventory cost would be

$$C^{\mathcal{T}}(\Upsilon|m) = hE\left[\Upsilon - \sum_{i=1}^N D^i\right]^+ + bE\left[\Upsilon - \sum_{i=1}^N D^i\right]^-,$$

where h is the holding cost rate and b is the backorder rate of the central warehouse. Thus, we can find the optimal basestock level at the central warehouse excluding reserves through equation $\text{Prob}\left(\Upsilon^{\mathcal{T}} \geq \sum_{i=1}^N D^i\right) = \frac{h}{h+b}$. Furthermore, if D^i is normal and independent for all i , then the optimal basestock level $\Upsilon^{\mathcal{T}}$ for the physically pooled demands is

$$\Upsilon_{n,s,i}^{\mathcal{T}} = N\mu + \sqrt{N}\sigma z^* \quad (6)$$

and the associated optimal system-wide long-run average inventory cost is

$$J_{n,s,i}^{\mathcal{T}} = \sqrt{N}\sigma(h+b)\phi(z^*). \quad (7)$$

Comparing (5) and (7), we have $\frac{J^{\mathcal{T}}}{J^{\mathcal{B}}} = \frac{1}{\sqrt{N}}$. Thus, if the inventory could be pooled in this model setting, the total inventory cost would reduce by a factor of $O(\sqrt{N})$.

4 Static Stockpiles

Now we introduce stockpiles into the model. We first model static stockpiling (SS), which aligns with the red-line practice at J&J discussed in the introduction. Under static stockpiles, each location j is assigned a static red-line m^j and the sum of these must meet a total minimum requirement, $\sum_j m_t^j = m$. In turn, each location follows a SOSO policy for this red-line inventory so that it can never fall below the red-line. We formally list the sequence of events in period t below:

1. Observe the state variables at each location - inventory on hand, x_t^i , and backorders, u_t^i . (Note that it is possible that both x_t^i and u_t^i are both strictly positive, in which case $x_t^i = m^i$.)
2. Place an order z_t^i for each location, bringing net inventory after ordering to $y_t^i = x_t^i - u_t^i + z_t^i \geq m^i$.
3. The order z_t^i arrives and random demand D_t^i is realized.
4. Each location serves demand (as much as possible while still honoring their m^i) and assesses holding and backorder costs.
5. The state variables update to $x_{t+1}^i = m^i + [y_t^i - m^i - D_t^i]^+$ and $u_{t+1}^i = [y_t^i - m^i - D_t^i]^-$.

In what follows, we confirm that the optimal way to manage inventory with static stockpiles is to use an augmented base-stock policy that effectively ignores the stockpile. Given the allocation \mathbf{m} , for any period t , $y_t^i = y^i$ and demand realization $D_t^i = d$, the one-period inventory cost at location i is

$$\begin{aligned} \widehat{C}^{\mathcal{S},i}(y^i, d | m^i) &= h^i[y^i - m^i - d]^+ + b^i[y^i - m^i - d]^- + h^i m^i \\ &= \widehat{C}^i(y^i - m^i, d) + h^i m^i. \end{aligned} \quad (8)$$

The second expression rewrites the one-period cost in a convenient way: it shows that the costs can be interpreted as simply setting aside m^i units, and then evaluating the standard inventory cost on the remaining inventory. Taking expectation over D^i , we obtain the expected one-period inventory cost $C^{\mathcal{S},i}(y^i | m^i) = C^i(y^i - m^i) + h^i m^i$, which depends on the inventory decision only through the standard inventory cost function C . Consequently, by following the standard inventory theory reviewed in the last section, the location's optimal inventory policy for a single period is a basestock policy that ignores the stockpiled inventory.

The lowest cost possible among all possible allocations of m can be established as follows. Define $\mathcal{M} = \{\mathbf{m} | \mathbf{m} \cdot \mathbf{e} = m, \mathbf{m} \geq 0\}$ the set of all possible stockpiles allocations. Also, let $\mathcal{A}_t = \{(\mathbf{y}_t, \mathbf{m}_t) | \mathbf{y}_t \geq (\mathbf{x}_t - \mathbf{u}_t) \vee \mathbf{m}_t, \mathbf{m}_t \in \mathcal{M}\}$ be the set of all feasible actions in period t and $\Pi^{\mathcal{S}} = \{\pi | (\mathbf{y}_t, \mathbf{m}_t)_\pi \in \mathcal{A}_t, t = 1, 2, \dots\}$ be the set of all feasible stockpiling policies. The system's optimal long run average cost is

$$J^{\mathcal{S}}(m) = \min_{\pi \in \Pi^{\mathcal{S}}} \lim_{T \rightarrow \infty} \frac{\sum_{t=1}^T E_\pi \left[\sum_{i=1}^N \widehat{C}^{\mathcal{S},i}(y_t^i, D_t^i | m^i) \right]}{T}.$$

Suppose h^i 's are all different, and let $i^* = \arg \min_i \{h^i\}$. The following proposition formalizes the above findings.

Proposition 1. *Under SS and given an allocation \mathbf{m} , the optimal order policy at each location is*

a base stock policy with base stock level

$$y^{S,i} = \bar{y}^i + m^i.$$

The lowest optimal long-run average cost among all allocations in \mathcal{M} is

$$\begin{aligned} J^S(m) &= C(\bar{\mathbf{y}}) + mh^{i^*} \\ &= J^B + mh^{i^*}. \end{aligned} \tag{9}$$

In particular, the smallest possible additional cost of maintaining the stockpile is $J^S(m) - J^B = mh^{i^*}$.

Depending on the context, it may not be possible to hold all of the inventory at the cheapest holding cost location. For example, the system may require that at least a certain amount of the stockpile be held at some preferred locations. Or, a contract may require that at least a certain amount of the stockpile be spread out to multiple locations. These kinds of additional considerations can be included as constraints, and the problem can be formulated and solved as a mixed integer and linear program (for both static stockpiles in this section and virtual stockpile pooling in the next section); see Section 1 in the online Supplementary Appendix. In general, adding these additional constraints increase the costs to hold the stockpile and preclude both the derivation of a closed-form optimal policy and tractable analytical comparison of static stockpile versus virtual stockpile pooling systems. Therefore, we choose to ignore context-dependent constraints to facilitate a clear comparison between static stockpiles and virtual stockpile pooling.

5 Virtual Stockpile Pooling

We now propose and analyze an alternative strategy that also meets the objective of guaranteeing m units at all times but can increase cost efficiency: Instead of maintaining separate isolated stockpiles at each location, one can strategically reallocate the stockpiles among the locations after observing demand realizations each period. We call this strategy *virtual stockpile pooling* (VSP). Under VSP, the sequence of events in each period is as follows:

1. Observe the state variables at each location - inventory on hand, x_t^i , and backorders, u_t^i .
2. Place an order z_t^i for each location such that $\sum_{i=1}^N y_t^i \geq m$, bringing the net inventory after ordering to $y_t^i = x_t^i - u_t^i + z_t^i$.
3. The order z_t^i arrives and random demand D_t^i realizes.

4. Allocate the stockpile among the locations, \mathbf{m}_t .
5. Serve demand at each location (as much as possible while still honoring m^i) and assesses holding and backorder costs.
6. The state variables update to $x_{t+1}^i = m_t^i + [y_t^i - m_t^i - D_t^i]^+$ and $u_{t+1}^i = [y_t^i - m_t^i - D_t^i]^-$.

In contrast to static stockpiles, the key difference is the ability to reallocate the stockpile to different locations every period in event 4.

Now, the inventory decision \mathbf{z}_t depends on the state of the system only through the net inventory $\mathbf{x}_t - \mathbf{u}_t$. Also, although the allocation decision \mathbf{m}_t depends on \mathbf{z}_t and $\mathbf{x}_t - \mathbf{u}_t$ through \mathbf{y}_t , it can be shown that it does not affect the future net inventory $\mathbf{x}_{t+1} - \mathbf{u}_{t+1}$ and hence will not affect the future inventory decision \mathbf{z}_{t+1} nor the future allocation decision \mathbf{m}_{t+1} . Therefore, we have:

Lemma 1. *Under VSP, the optimal allocation decision \mathbf{m} is myopic.*

Consequently, we can solve the VSP problem in a backwards fashion. We first find the optimal myopic allocation policy, i.e., the allocation of m units that minimizes that period's inventory costs. Then, we solve for the optimal order policy.

5.1 Optimal Allocation Policy

The *allocation problem* is to decide how to distribute the stockpile (or “red-lines”) among the locations to minimize inventory costs after observing demand realizations. Without loss of generality, we index the locations in increasing order of $h^i + b^i$ and define

$$a^i = h^i + b^i, \quad a^0 = 0, \quad \Delta a^i = a^i - a^{i-1} \geq 0, \quad i = 1, \dots, N.$$

Now, given basestock level \mathbf{y} , allocation \mathbf{m} , and demand realization \mathbf{d} , the one-period inventory cost across all locations is

$$\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{d} | m) = \widehat{C}^{\mathcal{S}}(\mathbf{y}, \mathbf{d} | \mathbf{m}) = \sum_{i=1}^N \widehat{C}^{\mathcal{S}, i}(y^i, d | m^i) \quad (10)$$

$$= \mathbf{a} \cdot [\mathbf{y} - \mathbf{m} - \mathbf{d}]^- + \mathbf{h} \cdot (\mathbf{y} - \mathbf{d}). \quad (11)$$

In the last expression, we have rewritten the costs by leveraging the fact that $z = z^+ - z^-$. The allocation problem is then

$$\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{d} | m) = \min_{\mathbf{m} \in \mathcal{M}(\mathbf{y})} \left\{ \widehat{C}^{\mathcal{S}}(\mathbf{y}, \mathbf{d} | \mathbf{m}) \right\}, \quad (12)$$

where $\mathcal{M}(\mathbf{y}) = \{\mathbf{m} | \mathbf{m} \in \mathbb{R}_+^N, \mathbf{m} \leq \mathbf{y}, \mathbf{m} \cdot \mathbf{e} = m\}$ is the set of all possible allocations given inventory level \mathbf{y} .

Because the last term of (11) is independent of \mathbf{m} , the allocation problem is equivalent to

$$\text{Problem A: } \min_{\mathbf{m} \in \mathcal{M}(\mathbf{y})} \mathbf{a} \cdot [\mathbf{m} - (\mathbf{y} - \mathbf{d})]^+.$$

The objective function of Problem A is separable and convex. Its feasible region is convex and the equality constraint is linear. Therefore, its dual problem exists and there is no duality gap. We can obtain the optimal allocation policy by solving the dual problem. Let γ , $\boldsymbol{\varpi}$ and $\boldsymbol{\nu}$ be the Lagrange multipliers for constraints $\mathbf{m} \in \mathbb{R}_+^N$, $\mathbf{m} \leq \mathbf{y}_t$, $\mathbf{m} \cdot \mathbf{e} = m$, respectively. The dual problem can be written as

$$\begin{aligned} \text{(Dual)} \quad \max \quad & q(\boldsymbol{\varpi}, \boldsymbol{\nu}, \gamma) = [(\mathbf{y} - \mathbf{d}) \cdot \mathbf{e} - m]\gamma - \mathbf{d} \cdot \boldsymbol{\varpi} - (\mathbf{y} - \mathbf{d}) \cdot \boldsymbol{\nu} \\ \text{s.t.} \quad & 0 \leq \boldsymbol{\nu} - \boldsymbol{\varpi} - \gamma \mathbf{e} \leq \mathbf{a}, \quad \boldsymbol{\varpi}, \boldsymbol{\nu} \geq 0. \end{aligned}$$

Fixing γ , the problem is separable across i and is equivalent to solving the following N linear programs separately

$$\begin{aligned} \text{(Problem-}i\text{)} \quad \min \quad & q^i(\mu^i, \nu^i) = d^i \mu^i + (y^i - d^i) \nu^i \\ \text{s.t.} \quad & \nu^i - \mu^i \geq \gamma \\ & \nu^i - \mu^i \leq \gamma + a^i \\ & \nu^i, \mu^i \geq 0. \end{aligned}$$

Here, $i = 1, \dots, N$. By solving each of these problems as a function of γ and then solving for optimal γ , we can show that the optimal allocation can be obtained by a greedy allocation – it tries to allocate the stockpiles to locations with smaller a^i values. More specifically, for any given \mathbf{y} and \mathbf{d} , define

$$\begin{aligned} \theta^j(\mathbf{y}, \mathbf{d}) &= \sum_{i=1}^{j-1} y^i + \sum_{i=j}^N (y^i - d^i)^+, \quad j = 1, \dots, N, \\ \widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m) &= \sum_{j=1}^N \Delta a^j [m - \theta^j(\mathbf{y}, \mathbf{d})]^+, \end{aligned}$$

where $\sum_{i=1}^0 y^i = 0$. Here, $\theta^j(\mathbf{y}, \mathbf{d})$ is the amount of inventory available for stockpiles if we backlog all current-period demands at locations 1 to $j-1$; $\widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$ is the sum of additional inventory-related costs if the total required m exceeds $\theta^j(\mathbf{y}, \mathbf{d})$ over all j . We have:

Proposition 2. *Under VSP, given \mathbf{y} and \mathbf{d} , the following greedy allocation $\mathbf{m}^{\mathcal{P}}(\mathbf{y}, \mathbf{d})$ is optimal*

$$\mathbf{m}^{\mathcal{P},i}(\mathbf{y}, \mathbf{d}) = y^i \wedge \left[[y^i - d^i]^+ + \left(m - \sum_{j=1}^{i-1} y^j - [y^i - d^i]^+ \right)^+ \right], \quad i = 1, \dots, N. \quad (13)$$

The resulting one-period cost is $\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m) = \widehat{C}(\mathbf{y}, \mathbf{d}) + \widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$.

Proposition 2 implies that the one-period cost of pooling under optimal allocation can be decomposed into two parts: the inventory cost without stockpiles and the additional holding and backorder cost caused by the required stockpiles.

To better illustrate the optimal allocation policy in Proposition 2, the following procedure provides a more detailed implementation algorithm. We then illustrate how one can execute this algorithm in an example. For any given period t , assume a location's order decision $\mathbf{y}_t = \mathbf{y}$ and a demand realization $\mathbf{d}_t = \mathbf{d}$.

Algorithm 1. (*VSP Allocation*)

Step 1. locations calculate their resulting on hand inventory if they were to satisfy as much demand as possible, $\hat{x}^i = [y^i - d^i]^+$, $i=1, \dots, N$.

Step 2. If $m \leq \sum_i \hat{x}^i$, allocate $m^1 = \hat{x}^1$, $m^2 = \hat{x}^2 \wedge (m - m^1)^+$, $m^3 = \hat{x}^3 \wedge (m - m^1 - m^2)^+$, etc., If at any time all units are allocated, stop. If at any time $m > \sum_i \hat{x}^i$, allocate $m^i = \hat{x}^i$ for each i . Set $i = 1$ and go to Step 3.

Step 3. Allocate $m^i = \max \left\{ m - \sum_{j=1}^{i-1} \hat{x}^j, y^i - \hat{x}^i \right\}$ to location i and update $\hat{x}^i = m^i$. Go to Step 4.

Step 4. If all m units have been allocated, stop; the resulting \mathbf{m} is the final allocation. Otherwise, set $i = i + 1$ and return to Step 3.

In short, the optimal allocation is to first allocate as much stockpiles to the locations who would have positive on-hand inventory after satisfying all his demand anyway (without causing backorders). If there is still stockpiles to be allocated, then one starts to allocate to locations (causing backorders) in the ascending order of their indices. Note that the optimal allocation is unique if every $h^i + b^i$ is unique and $m > \sum_{i=1}^N [y^i - d^i]^+$. When these conditions do not hold, there can be multiple optimal allocations. For example, if $m \leq \sum_{i=1}^N [y^i - d^i]^+$, then any allocation that satisfies $\mathbf{m} \in \mathcal{M}(\mathbf{y})$ and $m^i \in [0, [y^i - d^i]^+]$ is optimal.

Figure 3 illustrates how one can execute the optimal allocation algorithm such that each location need only communicate with only one other location, and there is at most $2N - 2$ numbers that need to be exchanged. In the example, there are three locations and the demand realizations indicate a near-worst-case situation in which 4 numbers need to be exchanged. First, each location calculates how much inventory is available with and without backordering. Location 3, who has the highest holding plus backorder cost, begins by raising her red-line as high as possible without causing backorders (in this case, 1 unit). She then reports the remaining stockpile needed to location 2 (9 units). Location 2 continues by agreeing to accommodate as high of a red-line as possible without

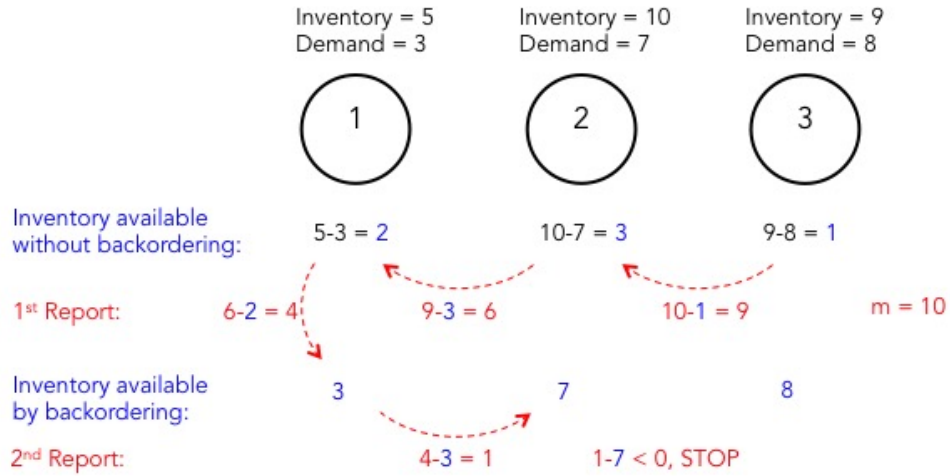


Figure 3: Illustration of reporting to execute the optimal VSP allocation algorithm.

causing backorders (in this case, 3 units) and reporting the remaining amount needed to location 1 ($9-3 = 6$ units). Then location 1 accepts as many units as she can both without *and* with causing backorders (i.e., all her available inventory, 5 units), before reporting back up the chain the remaining need ($6-5 = 1$ unit). Finally, location 2 must incur one more backorder to finish covering the total stockpile requirement.

Recall that under SS, the lowest possible cost occurs when all of the stockpile is held at the location with the smallest h^i . In contrast, under VSP, locations are prioritized to hold the stockpile according to the sum of their holding and backorder costs, $a^i = h^i + b^i$. Why is the rank based on the holding plus the backorder cost parameter? Why is the holding cost not given more weight? The VSP allocation algorithm reveals insight into the reason why this ranking occurs. When the system has the ability to reallocate stockpile after the demand realization, then the additional costs of stockpiling occur only when the system must start to allocate stockpiles that also induces backorders (Step 2, or the second report in the example). Before this point, there is no additional cost to maintaining the stockpile. After this point, however, each stockpiled unit creates *both* an additional unit to hold *and* an additional backordered demand. Therefore, it is optimal to allocate to locations with the lowest $h^i + b^i$.

5.2 Optimal Order Policy

To find the optimal order policy for the VSP problem, one can follow standard arguments in dynamic programming to prove the following lemma:

Lemma 2. *Under VSP, the optimal ordering decision \mathbf{y} is myopic.*

Next we characterize the optimal order policy by solving the single period problem. Define $C^{\mathcal{P}}(\mathbf{y}|m) = E[\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{D}|m)]$, $A(\mathbf{y}|m) = E[\widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{D}|m)]$. The second part of Proposition 2 yields an expression of the expected one-period cost in terms of the standard one-period cost function:

$$C^{\mathcal{P}}(\mathbf{y}|m) = C(\mathbf{y}) + A^{\mathcal{P}}(\mathbf{y}|m).$$

For example, under symmetric locations with identical basestock levels $\mathbf{y} = y\mathbf{e}$ we have

$$A^{\mathcal{P}}(\mathbf{y}|m) \underset{s}{=} (h+b)E\left[m - \sum_{i=1}^N [y - D^i]^+\right]^+,$$

so the one period cost simplifies to

$$C^{\mathcal{P}}(\mathbf{y}|m) \underset{s}{=} (h+b)E\left\{\left[m - \sum_{i=1}^N [y - D^i]^+\right]^+ + \sum_{i=1}^N (y - D^i)^-\right\} + Nhy - Nh\mu.$$

In general, we can show the following properties, which will be useful when characterizing the optimal order policy and the optimal cost.

Lemma 3. *For any given \mathbf{y} and \mathbf{d} :*

- (i) $\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$, $\widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$, $A^{\mathcal{P}}(\mathbf{y}|m)$ and $C^{\mathcal{P}}(\mathbf{y}|m)$ are positive, increasing and convex in m .
- (ii) $\widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$ and $A^{\mathcal{P}}(\mathbf{y}|m)$ are decreasing in y^i for all i .
- (iii) $\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$ and $C^{\mathcal{P}}(\mathbf{y}|m)$ are convex and supermodular in \mathbf{y} .
- (iv) $\partial C^{\mathcal{P}} / \partial y^i \partial m \leq 0$ for all i .
- (v) $\partial A^{\mathcal{P}} / \partial y^i$ increases when adding a new location to the system. Under symmetric locations, $\partial A^{\mathcal{P}} / \partial y^i$ is increasing in the number of locations N .

From Lemma 3, the optimal order policy can be characterized as the solution to the following first order conditions.

Proposition 3. *Under VSP, the optimal order policy is a stationary basestock policy with basestock levels $\mathbf{y}^{\mathcal{P}} = \operatorname{argmin} C^{\mathcal{P}}(\mathbf{y}|m)$ satisfying*

$$P(D^i + \mathbf{m}^{\mathcal{P},i}(\mathbf{y}^{\mathcal{P}}, \mathbf{D}) < y^{\mathcal{P},i}) = \frac{b^i}{h^i + b^i} \quad i = 1, \dots, N. \quad (14)$$

The optimal basestock levels $\mathbf{y}^{\mathcal{P}}$ are decreasing (coordinate-wise) in N , increasing (coordinate-wise) in m , and bounded by $\bar{\mathbf{y}} \leq \mathbf{y}^{\mathcal{P}} \leq \bar{\mathbf{y}} + m\mathbf{e}$. The optimal long run average cost is $J^{\mathcal{P}}(m) = C^{\mathcal{P}}(\mathbf{y}^{\mathcal{P}}|m)$.

Thus, even though the stockpile allocation can change every period, the optimal ordering policy is still a stationary basestock policy. Unlike under SS, however, the optimal basestock levels under VSP must be determined jointly because the minimum requirement can be met using inventory

at any location. Proposition 3 provides some natural bounds: the basestock level at each location should be no smaller than the case when there is no stockpile and no larger than the case when each location take the entire stockpile. These basestock levels also have some intuitively attractive properties. When more locations participate, each location should (weakly) decrease their basestock level. And, in order to accommodate a larger m , each location should (weakly) increase their basestock levels.

Importantly, even though the allocation policy takes place after the demand realization, Proposition 3 also implies that if one follows the optimal order policy and the optimal allocation policy, she will achieve an in-stock probability equal to $b^i/(b^i + h^i)$. Therefore, one can achieve a target in-stock probability of ISP^i at each location i by simply choosing a backorder cost that satisfies $b^i/(b^i + h^i)$ and then solving our model: the optimal solution to our model will achieve the desired target in-stock probability at each location.

It is also interesting to note that based on Proposition 3, there is a one-to-one correspondance between the basestock levels and the average time until a unit rotates. Under a base stock policy, we know that there always exists a fixed number inventory units in the system that have not yet been consumed by demand. Therefore, from Little’s Law, under a base stock policy y^i , we know that the expected time until rotation is $y^i/E[D]$, where $E[D]$ is the expected demand per period. From this observation, we can make some immediate observations about the expected time until rotation under the SS and VSP. In particular, the bounds in Proposition3 imply that moving from SS to VSP reduces the expected time until rotation at the “worst” location (i.e., the longest expected time until rotation is worse under SS than under VSP). One can also derive further metrics about the time until rotation under the optimal policy, such as the percent of units that rotate before a given age or the percent of time a location has at least one expired unit; see Section 2 in the outline Supplementary Appendix. However, directly incorporating age into the optimization problem is not straightforward and is outside the scope of this paper.

5.2.1 Basestock Level Computation

To obtain an exact computation of the optimal basestock policy $\mathbf{y}^{\mathcal{P}}$, we develop an efficient algorithm. Let \mathbf{y}^{-i} denote the sub-vector of \mathbf{y} with its i^{th} component removed. Denote $y^{\mathcal{P},i}(\mathbf{y}^{-i})$ to be the optimal basestock level under VSP for retailer i given that the other retailers order according to a stationary basestock policy with basestock levels \mathbf{y}^{-i} . We have the following intuitively appealing property.

Lemma 4. $y^{\mathcal{P},i}(\mathbf{y}^{-i})$ is decreasing in y^j for $j \neq i$, $i, j = 1, \dots, N$.

Lemma 4 states that the optimal basestock level for a given location will decrease if another location increases its basestock level. The property holds because of Lemma 3 part (iii). Based on this property, we can employ the following algorithm to compute the optimal basestock levels, which is an adaptation of a standard procedure developed for calculating the set of Nash equilibrium in supermodular games.

Algorithm. (*VSP Basestock Levels*)

Step 0: Set all location's lower bound to $\bar{\mathbf{y}}$ and upper bound to $\bar{\mathbf{y}} + m \cdot \mathbf{e}$.

Step i , $i=1, \dots, N$: Reset location i 's upper (lower) bound to the optimal basestock level for location i assuming other location's basestock levels are fixed at their lower (upper) bounds.

Repeat Steps 1 to N until the bounds cannot be improved for any location. The resulting region between each upper and lower bounds is the optimal basestock level sets.

Note that in each iteration the algorithm searches for the optimal basestock level in one dimension given the other basestock levels fixed at their upper or lower bound values. Thus, within each iteration the complexity of the algorithm is $O(N)$ line searches, implying that the optimal basestock levels can be computed in polynomial time. If the optimal basestock levels are unique, then the upper and lower bounds converge to a single vector - the optimal basestock vector $\mathbf{y}^{\mathcal{P}}$. The proof for the algorithm follows directly from Theorem 12.8 of Fudenberg and Tirole (1991).

6 Value of Virtual Stockpile Pooling

We now derive limiting behaviors and numerical examples to better understand under how VSP creates value and when it is expected to provide the most value. To accomplish this objective, we investigate the behavior of the following measure, which we call the value of VSP:

$$J^{\Delta}(m) = J^{\mathcal{S}}(m) - J^{\mathcal{P}}(m) \tag{15}$$

$$= C(\bar{\mathbf{y}}) + mh^{i^*} - C(\mathbf{y}^{\mathcal{P}}|m). \tag{16}$$

Recall that $J^{\mathcal{S}}(m)$ is the minimum possible cost under any allocation of a static stockpile. In this sense, $J^{\Delta}(m)$ is a lower bound for the cost savings that optimally-managed VSP can provide. However, note that adding other restrictions to stockpile management such as those discussed at the end of section 4 may affect $J^{\mathcal{P}}(m)$ and $J^{\mathcal{S}}(m)$ differently, in which case $J^{\Delta}(m)$ can be considered an approximation. To better understand the value of VSP, we will also make comparisons to no-stockpile case and traditional physical inventory pooling case reviewed in Section 3.




Figure 4: Value of VSP as a function of m for $N = 2, 3, 4, 5$ measured by absolute cost difference (SS cost - VSP cost; left) and in percent cost savings ($[\text{SS cost} - \text{VSP cost}] / \text{SS cost}$; right) .

Number of participating locations

First, we study the impact of increasing the number of participating locations into the system. We have:

Proposition 4. *The value of VSP, $J^\Delta(m)$, is increasing in the number of locations N . As the number of locations grow large, the optimal basestock levels approach the optimal basestock levels with no stockpile, $\lim_{N \rightarrow \infty} \mathbf{y}^P = \bar{\mathbf{y}}$, and the value of VSP is the full cost of the stockpile under SS $\lim_{N \rightarrow \infty} J^\Delta(m) = J^S(m) - J^B = mh^{i^*}$.*

In other words, if there are enough locations, then locations can effectively ignore the stockpile when deciding their basestock levels. The stockpile requirement can be met without any “extra” inventory at no additional cost - the cost savings equals the cost of holding the stockpile under SS. The key idea is that VSP allows the flexibility to allocate stockpiles to any location. Thus, if there are enough locations, there will always be enough locations that have extra inventory after meeting their own demand.

Does VSP only provide value when N is large, or does it provide significant value even if there are only a few participating locations? Interestingly, the value of VSP can be quite large even for only a few participating locations. Figure 4 shows the value of VSP for symmetric and independent locations with $h = 1$, $b = 9$, and exponential demand with mean 100. Before noting the shape of the graph, note that these cost savings are large from an absolute standpoint. As a point of comparison, the total inventory costs under the base case is only about 228 per location, and Figure 4 (right) shows the percent cost savings over the total cost under SS. Now, under SS the minimum additional




Figure 5: The optimal cost under VSP (left) and the percent cost increase (over no-stockpile case) to hold stockpile under VSP (right).

cost to hold m units is mh . Therefore, the figure shows that as long as m is not too large, the value of VSP is close to its maximum value, mh . And, as m grows large, the value is large, but approaches a limit. We study the impact of m next.

Amount of stockpiles

The numerical example suggests that the value of VSP is increasing in m . The following result confirms that this is the case, and provides an insightful results about its limiting behavior. (Here, to make a fair comparison, we assume the holding and backorder costs to be equal at all locations, including the consolidated warehouse in the physical inventory pooling case.)

Proposition 5. *The value of VSP, $J^\Delta(m)$, is increasing in m . For small m , $\lim_{m \rightarrow 0} \mathbf{y}^P = \bar{\mathbf{y}}$, and the value of VSP is the full cost of holding the stockpile under SS, $\lim_{m \rightarrow 0} J^\Delta(m) = J^S - J^B = mh^{i^*}$. As m grows large, the sum of the optimal basestock levels approach the optimal basestock level under physical inventory pooling, $\lim_{m \rightarrow \infty} \left(\sum_{i=1}^N y^{P,i} - m \right) = \Upsilon^T$, and the value of VSP approaches the value of physical pooling, $\lim_{m \rightarrow \infty} J^\Delta(m) = J^B - J^T = \left(N - \sqrt{N} \right) \sigma(h+b)\phi(z^*)$.*

Thus, if m is small enough, VSP effectively provides the stockpile for free and without increasing the basestock levels above the base case. Figure 5 shows the optimal cost under VSP and the percent cost increase over the no-stockpile case for various values of m , under the same parameters as Figure 4. In this example, as long as m is not too large (say, less than 200) and there are sufficient numbers of locations (say, 4) then it can be accommodated with only a small increase in cost ($<1\%$).

On the other hand, if m is large, then Proposition 5 states that the optimal basestock levels share a special relationship to the physical inventory pooling case: the total safety stock held equals the safety stock required under physical inventory pooling. In other words, locations can reduce their safety stock as much as they would be able to do under physical inventory pooling. The value of VSP approaches the benefits of physical inventory pooling. The intuition is that if m is large enough there will always be enough stockpile to facilitate a “virtual transshipment” - simultaneously increasing the red-line at one location and decreasing the red-line at another - to avoid a stockout if there is available inventory anywhere in the system. Therefore, in general, VSP can be thought of as physical inventory pooling with a limit on the number of transshipments based on m . Of course, it is also important to remember that VSP does not execute transshipments in the same manner as physical inventory transshipments - under VSP no real inventory actually moves between locations - however, it is useful to draw this conceptual comparison in order to help to understand when VSP will be most valuable. For instance, recall that under certain natural conditions, traditional inventory pooling allows the system to reduce its total inventory while achieving the same service level. Therefore, in these situations, Proposition 5 also suggests that VSP can help reduce the overall average time until rotation in comparison to SS (refer to our earlier discussion in subsection 5.2 regarding the relationship between base-stock levels and expected time until rotation.)

Does the stockpile need to be large to facilitate enough “virtual transshipments” to provide savings close to traditional inventory pooling? Interestingly, the value can be nearly as large even if m is only moderate in magnitude. Again, refer to Figure 4. In this example, the value of physical inventory pooling is about 143 when $n = 2$ and the value of VSP reaches nearly 90% of this value when $m = 200$. As a point of comparison, the optimal basestock levels in the base case are about 230 units per location in this example.

Cost savings per location

The previous analysis suggests that, in a sense, m and N are complements: VSP enables one to maintain m for free so long as N is large enough, and for a fixed N VSP enables one to gain all the benefits of physical inventory pooling so long as m is large enough. The following result shows that, indeed, the most cost savings occur for large m and N .

Proposition 6. *The value of pooling $J^\Delta(m)$ can be arbitrarily large with $\limsup_{m,N \rightarrow \infty} J^\Delta(m) = \infty$.*

Thus, VSP can provide an arbitrary large amount of savings as long as there are enough




Figure 6: The cost savings of VSP per location.

participating locations and a large enough stockpile. Furthermore, this result suggests that the value of adding more participating locations is greater when there is a larger requirement, and vice versa. Figure 6 illustrates the complementary nature of m and N by plotting the cost savings per location. Notice that for small m , a small number of locations provides the maximum cost savings. However, as m increases, so does the number of participating locations that maximizes the cost savings per location. One may be interested in the cost savings per location if, for example, there is a fixed cost to coordinate more locations under VSP. Thus, our results suggest that one should pool many locations for a large amount of stockpiles, but only a few if there is only a small amount of stockpiles.

Demand Correlation

We examine the effect of demand correlation on the value of VSP. We first need to leverage the following definition of correlation for demand vectors:

Definition 1. (Shaked and Shanthikumar 2007, p. 395) Let \mathbf{X} and \mathbf{Y} be two n -dimensional random vectors such that

$$E[\varphi(\mathbf{X})] \leq E[\varphi(\mathbf{Y})]$$

for all supermodular functions $\varphi : R^n \rightarrow R$, provided the expectations exist. Then \mathbf{X} is said to be smaller than \mathbf{Y} in the supermodular order (denoted by $\mathbf{X} \leq_{sm} \mathbf{Y}$).

Supermodular order characterizes the inter-correlation of the random vectors. $\mathbf{X} \leq_{sm} \mathbf{Y}$ implies that the random variables in \mathbf{Y} are more positively correlated than in \mathbf{X} . The following results

show the effect of the demand's inter-dependence on the one period cost saving. Denote $J_{\mathbf{D}}^{\Delta}(m)$ as the value of pooling under demand \mathbf{D} . We have:

Proposition 7. *Let \mathbf{D}_1 and \mathbf{D}_2 be two demand vectors. (i) If $\mathbf{D}_1 \leq_{sm} \mathbf{D}_2$, then the value of VSP is larger under \mathbf{D}_1 , $J_{\mathbf{D}_1}^{\Delta}(m) \geq J_{\mathbf{D}_2}^{\Delta}(m)$. (ii) If \mathbf{D}_1 and \mathbf{D}_2 have the same marginal distributions and the demands under \mathbf{D}_1 are weakly positively associated while the demands under \mathbf{D}_2 are weakly negatively associated, then $J_{\mathbf{D}_1}^{\Delta}(m) \geq J_{\mathbf{D}_2}^{\Delta}(m)$.*

(For a detailed definition of negatively associated, see Shaked and Shanthikumar 2007 p. 401.)

Given that we have already established a close relationship between VSP and traditional physical inventory pooling, one would expect to find that the demand characteristics that lead to more cost savings under traditional physical inventory pooling also lead to the most savings under VSP. Proposition 7 confirms that indeed, as it is under physical inventory pooling, the value of VSP is greater under more negatively correlated demands.

7 Concluding Remarks

In this paper, we identified and analyzed a new way to increase supply chain resilience in the form of stockpiled inventory at lower cost than current practices. By integrating stockpiles into existing inventory buffers and then dynamically conducting virtual transshipments of stockpiled inventory among locations, a firm can achieve risk-pooling benefits through what we call virtual stockpile pooling (VSP). We characterized the optimal stockpile allocation and replenishment policies under VSP, and provided simple algorithms to calculate and execute these policies. Using limiting behaviors and numerical examples, we showed that VSP can provide significant cost savings. This analysis also provided insight into when VSP is most valuable by drawing relationships between VSP, static stockpiling, no stockpiling, and traditional physical inventory pooling. For instance, in certain situations, VSP effectively maintains the stockpile for free. In other situations, it can achieve nearly the savings of traditional physical inventory pooling.

We conclude with a discussion of opportunities for future research and some of the limitations of our model. First, the model is stylized in order to focus on the main idea of VSP and to study its performance analytically. Thus, a specific organization may need to relax some assumptions in order to generate exact policy parameters. For example, our one-period leadtime assumption (which must also be assumed for much of the transshipment literature) allows us to derive the optimal replenishment policies in order to gain insights into the value of VSP and how those benefits

behave and are achieved. While the insights of the benefits of VSP should hold for multi-period leadtimes, future research may examine the implementation issues that arise from the increasingly complicated optimal inventory policies. The present model also does not capture factors that might make storing stockpiles at one location more desirable than storing inventory at another location. However, these factors tend to be very context specific (and are in fact negligible in the case of J&J). Last, we have assumed an entirely centralized system in which each location operates to minimize the entire system's cost. Such an assumption is reasonable for multiple locations within the same firm, as is the case with the J&J example. Nevertheless, if one can extend VSP to also apply to decentralized settings, then even further savings are possible. However, the question of exactly how one can optimally execute VSP with decentralized locations requires significant further research. Our results serve to provide a helpful starting point for these directions.

Acknowledgement. We would like to thank the Senior Editor and the referee for helpful comments to improve the paper. We would also like to thank Professor Yossi Sheffi for helpful discussions. This research was partially supported by Grant No. 71390331 from the Chinese Natural Science Foundation.

References

- Adida, E., P.C. DeLaurentis, M. Lawley. 2011. Hospital stockpiling for disaster planning. *IIE Transactions* **43**(5) 348-362.
- Alfaro, J.A., C.J. Corbett. 2003. The value of SKU rationalization in practice (the pooling effect under suboptimal inventory policies and nonnormal demand). *Production and Operations Management* **12**(1) 12-29.
- Allon, G., J.A. Van Mieghem. 2010. Global dual sourcing: Tailored basesurge allocation to near- and offshore production. *Management Science* **56**(1) 110-124.
- Alptekindglu, A., A. Banerjee, A. Paul, N. Jain. 2013. Inventory pooling to deliver differentiated service. *Manufacturing & Service Operations Management* **15**(1) 33-44.
- Anupindi, R. 2012. Supply chain risk management at Cisco: Response to H1N1. Michigan Ross School of Business Case 1-428-881.
- Anupindi, R, Y. Bassok. 1999. Centralization of stocks: suppliers vs. manufacturer. *Management Science* **45**(2) 178-191.
- Archibald, T.W., L.C. Thomas. 1997. An optimal policy for a two depot inventory problem with stock transfer. *Management Science* **43**(2) 173-183.

- Axsater, S. 2003. A new decision rule for lateral transshipments in inventory systems. *Management Science* **49**(9) 1168-1179.
- Benjaafar, S., W. Cooper, J.-S. Kim. 2005. On the benefits of pooling in production-inventory systems. *Management Science* **51**(4) 548-565.
- Berman, O., D. Krass, M. M. Tajbakhsh. 2011. On the benefit of pooling in inventory management. *Production and Operations Management* **20**(1) 57-71.
- Corbett, C., K. Rajaram. 2006. A generalization of the inventory pooling effect to nonnormal dependent demand. *Manufacturing & Service Operations Management* **8**(4) 351-358.
- Dong, L., N. Rudi. 2004. Who benefits from transshipment? Exogenous vs. endogenous wholesale prices. *Management Science* **50**(5) 645-657.
- Eppen, G. 1979. Effects of centralization on expected costs in a multi-location newsboy problem. *Management Science* **25**(5) 498-501.
- Eppen, G., L. Schrage. 1981. Centralized ordering policies in a multi-warehouse system with lead times and random demand. *TIMS Studies in the Management Sciences* **16** 51-67.
- Ergun, O., G. Karakus, P. Keskinocak, J. Swann, M. Villarreal. 2010. Operations research to improve disaster supply chain management. *Wiley Encyclopedia of Operations Research and Management Science*.
- Federgruen, A., P. Zipkin. 1984. Approximations of dynamic, multilocation production and inventory problems. *Management Science* **30**(1) 69-84.
- Fudenberg, D., J. Tirole. 1991. *Game Theory*. The MIT Press.
- Gayon, J-P., F. de Véricourt, F. Karaesmen. 2009. Stock rationing in an M/Er/1 multi-class make-to-stock queue with backorders. *IIE Transactions* **41**(12): 1096-1109.
- Gerchak, Y., D. Mossman. 1992. On the effect of demand randomness on inventories and costs. *Operations Research* **40**(4) 804-807.
- Ha, A. 1997. Inventory rationing in a make-to-stock production system with several demand classes and lost sales. *Management Science* **43**(8) 1093-1103.
- Herer, Y., M. Tzur. 2006. The multilocation transshipment problem. *IIE Transactions* **38**(3) 185-200.
- HHS Budget in Brief. 2015. HHS FY2016 Budget in Brief. <http://www.hhs.gov/about/budget/budget-in-brief/cdc/index.html>. Retrieved August 10, 2015.
- Hu, X., I. Duenyas, R. Kapuscinski. 2008. Optimal joint inventory and transshipment control under uncertain capacity. *Operations Research*. **56**(4) 881-897.

- Hua, Z., Y. Yu, W. Zhang, X. Xu. 2015. Structure properties of the optimal policy for dual-sourcing systems with general lead times. *IIE Transactions* **47**(8) 841-850.
- Huang, X., G. Sošic. 2010. Transshipment of inventories: Dual allocation vs. transshipment prices. *Manufacturing & Service Operations Management* **12**(2) 299-318.
- IEA. 2007. Oil supply security. Retrieved July 28, 2014, from http://www.iea.org/publications/freepublications/publication/oil_security.pdf.
- Janakiraman, G., S. Seshadri, A. Sheopuri. 2014. Analysis of tailored base-surge policies in dual sourcing inventory systems. *Management Science* **61**(7) 1547-1561.
- Maddah, B., A. Yassine, M. K. Salameh, L. Chatila. 2014. Reserve stock models: Deterioration and preventive replenishment. *European Journal of Operational Research* **232**(1) 64-71.
- Paterson, C., G. Kiesmuller, R. Teunter, K. Glazebrook. 2011. Inventory models with lateral transshipments: A review. *European Journal of Operational Research* **210**(2) 125-136.
- Reserve Requirements. <http://www.federalreserve.gov/monetarypolicy/reservereq.htm>. Retrieved August 15, 2014.
- Robinson, L. W. 1990. Optimal and approximate policies in multiperiod, multilocation inventory models with transshipments. *Operations Research* **38**(2) 278-295.
- Rudi, N., S. Kapus, D. F. Pyke. 2001. A two-location inventory model with transshipment and local decision making. *Management Science* **47**(12) 1668-1680.
- Sheffi, Y. 2005. *The resilient enterprise: overcoming vulnerability for competitive advantage*. MIT Press Books.
- Shaked, M., G. Shanthikumar. 2007. *Stochastic Orders*. Springer Science+Business Media, LLC.
- Shen, Z., M. Dessouky, F. Ordonez. 2010. Perishable inventory management system with a minimum volume constraint. *Journal of Operations Research Society* **62**(1) 2063-2082.
- Sheopuri, A., G. Janakiraman, S. Seshadri. 2010. New policies for the stochastic inventory control problem with two supply sources. *Operations Research* **58**(3) 734-745.
- Sošic, G. 2006. Transshipment of inventories among the retailers: Myopic vs. farsighted stability. *Management Science* **52**(10) 1493-1508.
- Song, J., P. Zipkin. 2009. Inventories with multiple supply sources and networks of queues with overflow bypasses. *Management Science* **55**(3) 362-372.
- Starr, M., L. Van Wassenhove. 2014. Introduction to the Special Issue on Humanitarian Operations and Crisis Management. *Production and Operations Management* **23**(6) 925-937.
- Tomlin, B. 2006. On the value of mitigation and contingency strategies for managing supply-chain disruption risks. *Management Science* **52**(5) 639-657.

- Tomlin, B., Y. Wang. 2005. On the value of mix flexibility and dual sourcing in unreliable news vendor networks. *Manufacturing Service Operations Management* **7**(1) 37-57.
- Tomlin, B., Y. Wang. 2011. Operational strategies for managing supply chain disruption risk. Kouvelis, P et al. eds., *The Handbook of Integrated Risk Management in Global Supply Chains*. John Wiley & Sons.79-101.
- Topkis, D. M. 1968. Optimal ordering and rationing policies in a nonstationary dynamic inventory model with n demand classes. *Management Science* **15**(1) 160-176.
- Topkis, D. M. 1998. *Supermodularity and complementarity*. Princeton University Press.
- Van Mieghem, J. A. 2003. Capacity management, investment, and hedging: review and recent developments. *Manufacturing Service Operations Management* **5**(3) 269-302.
- Van Mieghem, J. A. 2011. Risk management and operational hedging: An overview. Kouvelis, P et al. eds., *The Handbook of Integrated Risk Management in Global Supply Chains*. John Wiley & Sons Inc, Hoboken, NJ.
- Van Wassenhove, L. N. 2005. Humanitarian aid logistics: supply chain management in high gear. *Journal of the Operational Research Society* **57**(5) 475-489.
- Veeraraghavan, S., A. Scheller-Wolf. 2008. Now or later: A simple policy for effective dual sourcing in capacitated systems. *Operations Research* **56**(4) 850-864.
- de Véricourt, F., F. Karaesmen, Y. Dallery. 2002. Optimal stock allocation for a capacitated supply system. *Management Science* **48**(11) 1486-1501.
- Yang, H., L. Schrage. 2009. Conditions that cause risk pooling to increase inventory. *European Journal of Operational Research* **192**(3) 837-851.
- Zhang, J. A. 2005. Transshipment and its impact on supply chain members' performance. *Management Science* **51**(10) 1534-1539.
- Zhao, H. 2008. Optimal dynamic production and inventory transshipment policies for multi-location make-to-stock systems. *Operations Research* **56**(2) 400-410.
- Zipkin, P. 2000. *Foundations of Inventory Management*. McGraw Hill.

Appendix: Selected Proofs

Proof of Proposition 2: We first find \mathbf{m}^* and then derive the expression for $\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$. Finding \mathbf{m}^* is equivalent to solving Problem- i . We discuss the solution to Problem- i by considering the following two cases:

Case 1. $y^i \geq d^i$. If $\gamma + a^i \geq 0$, then the optimal solution is $\iota^{i*} = 0$ and $\nu^{i*} = \gamma^+$. If $\gamma + a^i < 0$,

then $\iota^{i*} = (\gamma + a^i)^-$ and $\nu^{i*} = (\gamma + a^i)^+ = 0 = \gamma^+$. We have

$$\iota^{i*} = (\gamma + a^i)^-, \quad \nu^{i*} = \gamma^+; \quad q^i(\iota^{i*}, \nu^{i*}) = d^i(\gamma + a^i)^- + (y^i - d^i)\gamma^+.$$

Case 2. $0 \leq y^i < d^i$. If $\gamma + a^i \geq 0$, then $\iota^{i*} = 0$ and $\nu^{i*} = \gamma + a^i$. If $\gamma + a^i \leq 0$, then $\iota^{i*} = (\gamma + a^i)^-$ and $\nu^{i*} = 0 = (\gamma + a^i)^+$. We have

$$\iota^{i*} = (\gamma + a^i)^-, \quad \nu^{i*} = (\gamma + a^i)^+; \quad q^i(\iota^{i*}, \nu^{i*}) = d^i(\gamma + a^i)^- + (y^i - d^i)(\gamma + a^i)^+.$$

Combining cases 1 and 2, we can write

$$\begin{aligned} q(\boldsymbol{\iota}^*, \boldsymbol{\nu}^*, \gamma) &= \gamma \left[\sum_{i=1}^N (y^i - d^i) - m \right] - \sum_{i=1}^N q^i(\iota^{i*}, \nu^{i*}) \\ &= -\gamma m + \sum_{i=1}^N (y^i - d^i)\gamma - \sum_{i=1}^N d^i(\gamma + a^i)^- - \sum_{i=1}^N (y^i - d^i)^+ \gamma^+ + \sum_{i=1}^N (y^i - d^i)^- (\gamma + a^i)^+ \\ &= -\gamma m - \sum_{i=1}^N (y^i - d^i)^+ \gamma^- - \sum_{i=1}^N d^i(\gamma + a^i)^- + \sum_{i=1}^N (y^i - d^i)^- [(\gamma + a^i)^+ - \gamma]. \end{aligned}$$

Recall that $0 = a^0 \leq a^1 \leq a^2 \leq \dots \leq a^N$. Thus, q can be written as a piecewise linear function of γ

$$q(\boldsymbol{\iota}^*, \boldsymbol{\nu}^*, \gamma) = \begin{cases} -\gamma m + \sum_{j=1}^N a^j [y^j - d^j]^-, & \gamma > 0 \\ -\gamma m + \gamma \left(\sum_{j=1}^{i-1} y^j + \sum_{j=i}^N [y^j - d^j]^+ \right) \\ \quad + \sum_{j=1}^{i-1} a^j d^j + \sum_{j=i}^N a^j [y^j - d^j]^-, & -a^i \leq \gamma < -a^{i-1}, i = 1, \dots, N. \end{cases}$$

For fixed \mathbf{y} and \mathbf{d} , where $\mathbf{y} \cdot \mathbf{e} \geq m$, define

$$\kappa^i(\mathbf{y}, \mathbf{d}) = \sum_{j=1}^{i-1} a^j d^j + \sum_{j=i}^N a^j [y^j - d^j]^-, \quad i = 1, \dots, N+1.$$

Because both \mathbf{y} and \mathbf{d} are non-negative, we have

$$\Delta \kappa^i = \kappa^{i+1} - \kappa^i = a^i d^i - a^i [y^i - d^i]^- \geq a^i (d^i \wedge y^i) \geq 0, \quad i = 1, \dots, N.$$

Both $\theta^i(\mathbf{y}, \mathbf{d})$ and κ^i are increasing in i . Moreover, $\theta^{N+1}(\mathbf{y}, \mathbf{d}) = \sum_{j=1}^N y^j \geq m$. Using these notations and making a change of variable from γ to $-\gamma$ yields

$$q(\boldsymbol{\iota}^*, \boldsymbol{\nu}^*, \gamma) = \begin{cases} m\gamma + \kappa^1, & \gamma < 0 \\ (m - \theta^i(\mathbf{y}, \mathbf{d}))\gamma + \kappa^i, & -a^{i-1} < \gamma \leq -a^i, i = 1, \dots, N. \end{cases}$$

This is a concave, piecewise linear function of γ , so the optimal γ that maximizes this function is $\gamma^* = a^i$, if $\theta^i(\mathbf{y}, \mathbf{d}) \leq m < \theta^{i+1}(\mathbf{y}, \mathbf{d})$, and the corresponding m^j s are:

if $\gamma^* > 0$, then

$$m^{j*} = \begin{cases} y^j & j < i \\ [y^i - d^i]^+ & j > i; \\ m - \sum_{k \neq i} m^k & j = i \end{cases}$$

if $\gamma^* = 0$, then

$$m^{j*} = \begin{cases} [y^j - d^j]^+ & m \geq \sum_{k=1}^j [y^k - d^k]^+ \\ 0 & m \leq \sum_{k=1}^{j-1} [y^k - d^k]^+ \\ m - \sum_{k \neq i} m^k & \sum_{k=1}^j [y^k - d^k]^+ > m > \sum_{k=1}^{j-1} [y^k - d^k]^+ \end{cases}.$$

Substituting γ^* into q , we obtain $q^* = \sum_{i=1}^N [a^i(m - \theta^i) + \kappa^i] 1_{\{\theta^i \leq m < \theta^{i+1}\}}$, where 1_U is the indicator of event U . Note that $\kappa^i = \sum_{j=1}^i a^j (\theta^{j+1}(\mathbf{y}, \mathbf{d}) - \theta^j(\mathbf{y}, \mathbf{d}))$. Therefore,

$$q^* = \sum_{i=1}^N \left[a^i(m - \theta^i(\mathbf{y}, \mathbf{d})) + \sum_{j=1}^i a^j (\theta^{j+1}(\mathbf{y}, \mathbf{d}) - \theta^j(\mathbf{y}, \mathbf{d})) \right] 1_{\{\theta^i(\mathbf{y}, \mathbf{d}) \leq m < \theta^{i+1}(\mathbf{y}, \mathbf{d})\}},$$

and hence

$$q^* = \sum_{i=1}^N \left[\sum_{j=1}^i \Delta a^i(m - \theta^i(\mathbf{y}, \mathbf{d})) \right] 1_{\{\theta^i(\mathbf{y}, \mathbf{d}) \leq m < \theta^{i+1}(\mathbf{y}, \mathbf{d})\}},$$

That is, $q^* = \sum_{j=1}^N \Delta a^i(m - \theta^i(\mathbf{y}, \mathbf{d}))^+$. Finally, adding back $[\mathbf{y} - \mathbf{d}] \cdot \mathbf{h}$, we obtain $\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m) = \widehat{C}(\mathbf{y}, \mathbf{d}) + \widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$. Rewriting m^i , we have equation (13). ■

Proof of Lemma 3: We will leverage the following additional lemma (see Topkis 1998; Lemma 2.6.4).

Lemma 5. *Let $f(x) : R \rightarrow R$ and $g(\mathbf{z}) : R^n \rightarrow R$, $f \circ g : R^n \rightarrow R$. If f is increasing and convex and g is supermodular and coordinate-wise decreasing (increasing), then $f \circ g$ is supermodular.*

(i) Note that function $[m - y]^+$ is increasing and convex in m and $\widehat{A}^{\mathcal{P}}$ is a positive sum of these functions, so $\widehat{A}^{\mathcal{P}}$ is positive, increasing and convex in m . Because $\widehat{C}(\mathbf{y}, \mathbf{d})$ is independent of m , we have $\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$ is increasing and convex in m .

(ii) Similar to (i), the proof follows because each term of $\widehat{A}^{\mathcal{P}}$ is decreasing and convex in y^i , and $\widehat{A}^{\mathcal{P}}$ is a positive sum of these terms.

(iii) We first show $\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m) = \min_{\mathbf{m} \in \mathcal{M}} \{ \widehat{C}^{\mathcal{S}}(\mathbf{y}, \mathbf{d}|\mathbf{m}) \}$ is convex in \mathbf{y} . Choose any two vectors \mathbf{y}_i , $i = 1, 2$, let \mathbf{m}_i be the minimizer of $\widehat{C}^{\mathcal{P}}(\mathbf{y}_i, \mathbf{d}|m)$. Note that for any $0 \leq \lambda \leq 1$, $\lambda \mathbf{m}_1 + (1 - \lambda) \mathbf{m}_2$ is a feasible point in $\mathcal{M}(\mathbf{y})$ for $\mathbf{y} = \lambda \mathbf{y}_1 + (1 - \lambda) \mathbf{y}_2$. Because $\widehat{C}^{\mathcal{S}}(\mathbf{y}, \mathbf{d}|\mathbf{m})$ is

convex in \mathbf{y} , we have

$$\begin{aligned} \lambda \widehat{C}^{\mathcal{P}}(\mathbf{y}_1, \mathbf{d}|m) + (1 - \lambda) \widehat{C}^{\mathcal{P}}(\mathbf{y}_2, \mathbf{d}|m) &= \lambda \widehat{C}^{\mathcal{S}}(\mathbf{y}_1, \mathbf{d}|\mathbf{m}_1) + (1 - \lambda) \widehat{C}^{\mathcal{S}}(\mathbf{y}_2, \mathbf{d}|\mathbf{m}_2) \\ &\geq \widehat{C}^{\mathcal{S}}(\lambda \mathbf{y}_1 + (1 - \lambda) \mathbf{y}_2, \mathbf{d}|\lambda \mathbf{m}_1 + (1 - \lambda) \mathbf{m}_2) \geq \widehat{C}^{\mathcal{P}}(\lambda \mathbf{y}_1 + (1 - \lambda) \mathbf{y}_2, \mathbf{d}|m) \end{aligned}$$

Therefore, $\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$ is convex in \mathbf{y} . Next we show $\widehat{C}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m) = \widehat{C}(\mathbf{y}, \mathbf{d}) + \widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$ is supermodular in \mathbf{y} . Note that $\widehat{C}(\mathbf{y}, \mathbf{d})$ is separable convex in \mathbf{y} , and hence supermodular in \mathbf{y} . It is sufficient to show $\widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{d})$ is supermodular in \mathbf{y} . Moreover, since $\widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$ is a positive sum of N terms, we only need to show that each of these terms is supermodular. Take term j in $\widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$, let $f(x) = [x]^+$ and $g(\mathbf{y}) = m - \theta^j(\mathbf{y}, \mathbf{d})$. Then, f is increasing and convex and g is coordinate-wise decreasing and separable hence supermodular. Therefore from Lemma 5, we know $f \circ g$ is supermodular in \mathbf{y} .

(iv) and (v) Fixing $y^j, j \neq i$, we can show that each term in $\widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{d}|m)$ has decreasing differences (y^i, m) , which leads to a negative cross partial. Therefore, $C^{\mathcal{P}}(\mathbf{y}|m) = E[\widehat{C}(\mathbf{y}, \mathbf{D}) + \widehat{A}^{\mathcal{P}}(\mathbf{y}, \mathbf{D}|m)]$ has decreasing differences in (y^i, m) and hence $\partial C^{\mathcal{P}}(\mathbf{y})/\partial y^i \partial m \leq 0$ is true. In order to prove $\partial C^{\mathcal{P}}(\mathbf{y}|m)/\partial y^i$ is increasing in N , we only need to show $\partial A^{\mathcal{P}}(\mathbf{y}|m)/\partial y^i$ is increasing in N because $E[\widehat{C}(\mathbf{y}, \mathbf{D})]$ is independent of N . Note that

$$\partial A^{\mathcal{P}}(\mathbf{y}|m)/\partial y^i = - \sum_{j=1}^i \Delta a^j P(\theta^j(\mathbf{y}, \mathbf{D}) \leq m, D^i < y^i) - \sum_{j=i+1}^N \Delta a^j P(\theta_N^j(\mathbf{y}, \mathbf{D}) \leq m). \quad (17)$$

Now, fix the basestock level and the marginal demand distribution for all of the suppliers, then add another supplier to the system and let h^e and b^e be the holding cost rate and the backorder cost rate of the new supplier. Suppose $a^l \leq a^e \leq a^{l+1}$ for some l . Because $\theta^j(\mathbf{y}, \mathbf{D})$ is increasing in N , each probability term in (17) is decreasing in N . Note that $\Delta a^{l+1} = a^{l+1} - a^l = a^{l+1} - a^e + a^e - a^l = \Delta a(l+1, e) + \Delta a(e, l)$. Then $\partial A^{\mathcal{P}}(\mathbf{y}|m)/\partial y^i$ with and without e can be written as:

$$\begin{aligned} &- \sum_{j=1, j \neq l+1}^i \Delta a^j [P(\theta^j(\mathbf{y}, \mathbf{D}) \leq m, D^i < y^i) - P(\theta^j(\mathbf{y}, \mathbf{D}) \leq m, D^i < y^i)] \\ &- \sum_{j=i+1, j \neq l+1} \Delta a^j [P(\theta^j(\mathbf{y}, \mathbf{D}) \leq m) - P(\theta^j(\mathbf{y}, \mathbf{D}) \leq m)] \\ &- \Delta a(l+1, e) \left[P(\theta^l(\mathbf{y}, \mathbf{D}) + (y^e - D^e)^+ \leq m, [D^i < y^i] \mathbf{1}_{\{i>e\}}) - P(\theta^l(\mathbf{y}, \mathbf{D}) \leq m, [D^i < y^i] \mathbf{1}_{\{i>e\}}) \right] \\ &- \Delta a(e, l) \left[P(\theta^l(\mathbf{y}, \mathbf{D}) + y^e \leq m, [D^i < y^i] \mathbf{1}_{\{i>e\}}) - P(\theta^l(\mathbf{y}, \mathbf{D}) \leq m, [D^i < y^i] \mathbf{1}_{\{i>e\}}) \right] \geq 0. \end{aligned}$$

Therefore, part (v) is true. ■

Proof of Proposition 3: First, observe that the myopic ordering policy for the relaxed VSP problem is optimal. Thus, $\mathbf{y}^{\mathcal{P}} = \operatorname{argmin} C^{\mathcal{P}}(\mathbf{y}|m)$ is the optimal basestock level. Because $C^{\mathcal{P}}(\mathbf{y}|m)$ is convex in \mathbf{y} , $\mathbf{y}^{\mathcal{P}}$ satisfies the first order condition. By applying the Envelope theorem on

$\partial C^{\mathcal{P}}(\mathbf{y}^{\mathcal{P}}|m)/\partial y^i$, we have equation (14). That $\mathbf{y}^{\mathcal{P}}$ coordinate-wise decreases if we add an additional location and increases in m follows from Lemma 3 part (v). That it is increasing in m follows directly from equation (14) and the fact that $m^{\mathcal{P},i} \geq 0$, *a.e.* ■

Proof of Proposition 4: Note that for any given \mathbf{d} and $\mathbf{y} \geq \bar{\mathbf{y}}$, we have the following

$$[m - \theta^j(\mathbf{y}, \mathbf{d})]^+ \leq [m - \theta^1(\mathbf{y}, \mathbf{d})]^+ \leq [m - \theta^1(\bar{\mathbf{y}}, \mathbf{d})]^+$$

where the right most term converges to zero as N increases. Because there exists $\epsilon, \delta > 0$ such that $Pr(D^i > \epsilon) > \delta$ and D^i has continuous support for all i , we have $\lim_{N \rightarrow \infty} A^{\mathcal{P}}(\mathbf{y}|m) = 0$. Thus, $\lim_{N \rightarrow \infty} \mathbf{y}^{\mathcal{P}} = \bar{\mathbf{y}}$, and hence the value of VSP is the full cost of the stockpiles under separate stockpiles $\lim_{N \rightarrow \infty} J^{\Delta}(m) = J^{\mathcal{S}}(m) - J^{\mathcal{B}} = mh^{i^*}$. ■

Proof of Proposition 5: We first show that $J^{\Delta}(m)$ is increasing in m . From the definition of $J^{\Delta}(m)$ we know that it is sufficient to show $C(\mathbf{y}_m^{\mathcal{P}*}|m)$'s increasing rate in m is always no more than h^{i^*} , where $\mathbf{y}_m^{\mathcal{P}*}$ is the optimal order up to levels when the minimum requirement level is m . Consider a small increment ϵ in m . Then, construct a hybrid inventory reserve policy such that the first m units of stockpile is pooled after demand realization while the last ϵ units is stockpiled at location i^* using static stockpiling policy. Then, we can show similar to Proposition 1 that the optimal cost under this hybrid policy is $C(\mathbf{y}_m^{\mathcal{P}*}|m) + \epsilon h^{i^*}$, the optimal inventory policy is to order up to $\mathbf{y}_m^{\mathcal{P},j*}$ for all locations $j \neq i$, and $\mathbf{y}_m^{\mathcal{P},i*} + \epsilon$ for location i , and the optimal allocation of the stockpile after demand realization is $\mathbf{m}^{\mathcal{P},j}(\mathbf{y}_m^{\mathcal{P}*}, \mathbf{D})$ for all locations $j \neq i$, and $\mathbf{m}^{\mathcal{P},i}(\mathbf{y}_m^{\mathcal{P}*}, \mathbf{D}) + \epsilon$ for location i . Because this hybrid policy is a feasible policy under stockpile pooling, we must have $C(\mathbf{y}_{m+\epsilon}^{\mathcal{P}*}|m+\epsilon) \leq C(\mathbf{y}_m^{\mathcal{P}*}|m) + \epsilon h^{i^*}$, therefore, the cost increment from m to $m+\epsilon$ is less than ϵh^{i^*} implying an increment rate smaller than h^{i^*} . This proves that $J^{\Delta}(m)$ is increasing in m . Next, from Proposition 3 we know that for small amounts of stockpiles, the optimal basestock levels approach the optimal basestock levels in the base case $\lim_{m \rightarrow 0} \mathbf{y}^{\mathcal{P}} = \bar{\mathbf{y}}$, and the value of VSP is the full cost of holding the stockpiles under separate stockpiles $\lim_{m \rightarrow 0} J^{\Delta}(m) = J^{\mathcal{S}} - J^{\mathcal{B}} = mh^{i^*}$. To compare with traditional pooling, without loss of generality, we assume $h + b = 1$. Note that the optimal policy, $\mathbf{y}^{\mathcal{P}}$ has identical entries and we write each entry as $y^{\mathcal{P},i} = y^{\mathcal{P}*}(m) + \frac{m}{N}$ where the first term is correlated with m . If we select the basestock level under stockpiles pooling to be $y^i = \frac{y^*}{N} + \frac{m}{N}$, the one period cost difference between stockpiles pooling and traditional pooling is larger compared to under the optimal basestock levels, $C^{\mathcal{P}}(\mathbf{y}^{\mathcal{P}}|m) - C(y^*|m) \leq C^{\mathcal{P}}(\mathbf{y}|m) - C(y^*|m)$. Recall that the cost under stockpiling can be written in two parts, similarly we can write the cost under traditional

pooling in two parts:

$$\begin{aligned} C(y|m) &= hE \left[y - \sum_{i=1}^N D^i \right]^+ + bE \left[y - \sum_{i=1}^N D^i \right]^- + hm \\ &= hE \left[y + m - \sum_{i=1}^N D^i \right]^+ + bE \left[y + m - \sum_{i=1}^N D^i \right]^- + E \left(m - \left(y + m - \sum_{i=1}^N D^i \right)^+ \right)^+ \end{aligned}$$

thus the one period cost difference under $y^i = \frac{y^*}{N} + \frac{m}{N}$ is

$$C^{\mathcal{P}}(\mathbf{y}|m) - C(y^*|m) = E \left[\left[m - \sum_{i=1}^N \left(\frac{y^*}{N} + \frac{m}{N} - D^i \right)^+ \right]^+ - \left(m - \left(y^* + m - \sum_{i=1}^N D^i \right)^+ \right)^+ \right].$$

Then since $\sum_{i=1}^N \left(\frac{y^*}{N} + \frac{m}{N} - d^i \right)^+ \geq \sum_{i=1}^N \left(\frac{y^*}{N} + \frac{m}{N} - d^i \right) = y^* + m - \sum_{i=1}^N d^i$, we have the above cost difference is always negative.

Similarly, if we substitute $y = Ny^{\mathcal{P}*}(m)$ into the cost difference, then we have $C^{\mathcal{P}}(\mathbf{y}^{\mathcal{P}}|m) - C(y|m) \leq C^{\mathcal{P}}(\mathbf{y}^{\mathcal{P}}|m) - C(y^*|m)$, and by similar argument that $\sum_{i=1}^N (y^{\mathcal{P}*}(m) + \frac{m}{N} - d^i)^+ \geq \sum_{i=1}^N (y^{\mathcal{P}*}(m) + \frac{m}{N} - d^i) = y + m - \sum_{i=1}^N d^i$, we have

$$C^{\mathcal{P}}(\mathbf{y}^{\mathcal{P}}|m) - C(y|m) = E \left[\left[m - \sum_{i=1}^N \left(y^{\mathcal{P}*}(m) + \frac{m}{N} - D^i \right)^+ \right]^+ - \left(m - \left(y + m - \sum_{i=1}^N D^i \right)^+ \right)^+ \right] \leq 0.$$

Thus,

$$0 \geq C^{\mathcal{P}}(\mathbf{y}^{\mathcal{P}}|m) - C(y^*|m) \geq E \left[\left[m - \sum_{i=1}^N \left(\frac{y^*}{N} + \frac{m}{N} - D^i \right)^+ \right]^+ - \left(m - \left(y^* + m - \sum_{i=1}^N D^i \right)^+ \right)^+ \right].$$

Note that when $m \rightarrow \infty$, because $E[D^i] < \infty$, we have

$$E \left[\left(m - \sum_{i=1}^N \left(\frac{y^*}{N} + \frac{m}{N} - D^i \right)^+ \right)^+ \right] \rightarrow E \left[\left(m - \sum_{i=1}^N \left(\frac{y^*}{N} + \frac{m}{N} - D^i \right) \right)^+ \right] = E \left[\left(\sum_{i=1}^N D^i - y^* \right)^+ \right]$$

and $E \left(m - \left(y^* + m - \sum_{i=1}^N D^i \right)^+ \right)^+ \rightarrow E \left(\sum_{i=1}^N D^i - y^* \right)^+$. Thus,

$$0 \geq \lim_{m \rightarrow \infty} [C^{\mathcal{P}}(\mathbf{y}^{\mathcal{P}}|m) - C(y^*|m)] \geq \lim_{m \rightarrow \infty} E \left[\left(\sum_{i=1}^N D^i - y^* \right)^+ \right] - E \left(\sum_{i=1}^N D^i - y^* \right)^+ = 0.$$

and we have that $\lim_{m \rightarrow \infty} [C^{\mathcal{P}}(\mathbf{y}^{\mathcal{P}}|m) - C(y^*|m)] = 0$. The cost under stockpiles pooling converges

to the cost under traditional inventory pooling. When each location is symmetric,

$$\begin{aligned} C^{\mathcal{P}}(\mathbf{y}^{\mathcal{P}}|m) &= (h+b)E \left\{ \left[m - \sum_{i=1}^N [y^{\mathcal{P}} - D^i]^+ \right]^+ + \sum_{i=1}^N (y^{\mathcal{P}} - D^i)^- \right\} + Nhy^{\mathcal{P}} - Nh\mu \\ &\rightarrow (h+b)E \left\{ \left[\sum_{i=1}^N D^i - (y^{\mathcal{P}} - m) \right]^+ \right\} + h(Ny^{\mathcal{P}} - m) - Nh\mu. \end{aligned}$$

Thus, solving for $\mathbf{y}^{\mathcal{P}} - m$ in this case is similar to finding the inventory level for physical inventory pooling. ■

Proof of Proposition 6: We prove this proposition by finding a case for which $J^{\Delta}(m)$ can be arbitrarily large. From Proposition 5, we know that the cost under pooling for a network with symmetric holding and backorder cost rates converges to the traditional pooling $\lim_{m \rightarrow \infty} [C^{\mathcal{P}}(\mathbf{y}^{\mathcal{P}}|m) - C(y^*|m)] = 0$. Alternatively, if the system follows the stockpiling policy, the cost difference under the optimal policies is $C^{\mathcal{S}}(\mathbf{y}^{\mathcal{S}}|m) - C(y^*|m)$, where $y^{\mathcal{S},i} = \bar{y} + \frac{m}{N}$ is identical among all locations. The cost difference $C^{\mathcal{S}}(\mathbf{y}^{\mathcal{S}}|m) - C(y^*|m)$ equals

$$\begin{aligned} & hE \left(\sum_{i=1}^N \left[\bar{y} + \frac{m}{N} - D^i \right]^+ - \left[y^* + m - \sum_{i=1}^N D^i \right]^+ \right) + bE \left(\sum_{i=1}^N \left[\bar{y} + \frac{m}{N} - D^i \right]^- - \left[y^* + m - \sum_{i=1}^N D^i \right]^- \right) \\ & + E \left[\sum_{j=1}^N \left(\frac{m}{N} - \left(y^{\mathcal{P}*}(m) + \frac{m}{N} - D^j \right)^+ \right)^+ - \left(m - \left(y^* + m - \sum_{i=1}^N D^i \right)^+ \right)^+ \right]. \end{aligned}$$

When $m \rightarrow \infty$, because $E[D^i] < \infty$, then the first term goes to $N\bar{y} - y^*$, the second term goes to 0, and the third term goes to $E \left[\sum_{i=1}^N (\bar{y} - D^i)^- - \left(y^* - \sum_{i=1}^N D^i \right)^- \right]$. Note that when D^i are not perfectly positively correlated, i.e. $D^i = D^j$ for all $i, j = 1, \dots, N$, then $y^* < N\bar{y}$, and hence

$$\begin{aligned} E \left[\sum_{i=1}^N (\bar{y} - D^i)^- - \left(y^* - \sum_{i=1}^N D^i \right)^- \right] & \geq E \left[\sum_{i=1}^N (\bar{y} - D^i)^- - \left(N\bar{y} - \sum_{i=1}^N D^i \right)^- \right] \\ & \geq E \left[\sum_{i=1}^N (\bar{y} - D^i)^- - \sum_{i=1}^N (\bar{y} - D^i)^- \right] = 0. \end{aligned}$$

Therefore, the cost difference when $m \rightarrow \infty$ satisfies $\lim_{m \rightarrow \infty} C^{\mathcal{S}}(\mathbf{y}^{\mathcal{S}}|m) - C(y^*|m) \geq N\bar{y} - y^* > 0$. Note that the lower bound of the difference is increasing in N , thus as $N \rightarrow \infty$ the difference goes to ∞ . This completes the proof of Proposition 6. ■

Proof of Proposition 7: Note that $J^{\Delta}(m) = J^{\mathcal{S}}(m) - J^{\mathcal{P}}(m)$. (i) follows from Definition 1. (ii) follows from Shaked and Shanthikumar (2007) Theorem 9.A.23. ■