

Supply and Demand Uncertainty in Multi-Echelon Supply Chains

Lawrence V. Snyder

Dept. of Industrial and Systems Engineering

Lehigh University

Bethlehem, PA, USA

`larry.snyder@lehigh.edu`

Zuo-Jun Max Shen

Department of Industrial Engineering and Operations Research

University of California–Berkeley

Berkeley, CA, USA

`shen@ieor.berkeley.edu`

September 2006

Abstract

This paper explores the differences between demand and supply uncertainty in multi-echelon supply chains. Our central premise is that the insights gained from the study of one type of uncertainty often do not apply to the other. In fact, the two types of uncertainty are in a sense mirror images of each other, in that the optimal strategy for coping with supply uncertainty may be exactly opposite to that for demand uncertainty. We present several studies, each involving a fundamental question of order frequency, inventory placement, or supply chain structure. In each study, we consider a simple multi-echelon supply chain, first under demand uncertainty and then under supply uncertainty in the form of disruptions. Using simulation, we demonstrate that the optimal strategy is different under the two types of uncertainty and discuss reasons for the differences.

We begin with a study demonstrating that the cost of failing to plan for disruptions is greater than that of failing to plan for demand uncertainty. We conclude with a study that demonstrates that resiliency to disruptions is often inexpensive, in that large improvements in service level can be achieved with only small increases in cost. The remaining studies identify a number of important properties of supply chains subject to disruptions that have not previously been discussed in the literature. For example, we introduce the “risk-diversification effect,” which describes the benefit of having multiple, decentralized stocking locations under supply uncertainty. This is the opposite of the classical risk-pooling effect, which describes the reverse tendency under demand uncertainty.

1. Introduction

At first glance, supply and demand uncertainty in supply chain models seem like two sides of the same coin. A toy store finding its supplier stocked out of a popular new toy sees supply uncertainty, while the toy manufacturer sees the demand uncertainty that caused the stockout in the first place. Moreover, similar strategies can be used to cope with both types of uncertainty: safety stock inventory, dual sourcing, improved forecasting, and so on.

Given this, it may seem that inventory and supply chain models that consider one type of uncertainty may be easily adapted to the other. This may partly explain why papers on demand uncertainty significantly outnumber those on supply uncertainty in the inventory and supply chain literature.

With the growing attention paid in recent years to supply chain disruptions and other types of supply uncertainty, it is worth examining in more detail the similarities and differences between the two types of uncertainty. The central thesis of this paper is that the two are quite different, and that the conventional wisdom gained from decades of research on demand uncertainty models does not always apply under supply uncertainty. In fact, the optimal strategy for dealing with supply uncertainty is, in many cases, the exact opposite from that for demand uncertainty.

In this paper, we present the results of a number of simulation studies that demonstrate these differences. Our simulations consider simple multi-echelon supply chains under both demand uncertainty (DU) and supply uncertainty (SU). Each study examines two different strategies (for example, one-for-one vs. batch ordering or centralization vs. decentralization). By examining the costs of each strategy, we can identify the optimal strategy under each type of uncertainty. We simulate the systems for a range of parameter values and, in many cases, find that the two types of uncertainty have different optimal strategies. These studies are bracketed by Study 1, which explores the cost of unreliability (i.e., the cost of not planning for uncertainty) and Study 8, which explores the cost of reliability (the increase in cost required to improve the system's resiliency).

In order to isolate the effect of each type of uncertainty, we consider them each separately. That is, when demand is uncertain, supply is deterministic, and vice-versa. One overarching theme that emerges from these studies is that SU and DU appear to

have a mirror-image relationship to one another, in the sense that the optimal strategy for one may be diametrically opposed to that for the other.

To our knowledge, this is the first paper to systematically examine the difference between these two types of uncertainty, and to highlight the different strategies that are appropriate. In addition, ours is one of the first papers to consider disruptions in multi-echelon models. We believe that it is important to examine disruptions in such settings because disruptions are not local; rather, they tend to propagate downstream through a supply chain.

We study these models using simulation because many of them are analytically intractable and because our objective is to gain insights using realistic models rather than to find optimal solutions to exact but vastly simplified models.

The remainder of this paper is structured as follows. In Section 2 we present a brief review of the literature on SU. We describe our simulation methodology and experimental design in detail in Section 3. In Sections 4–8, we present the results of our simulation studies, grouped according to the underlying question being addressed. We conclude in Section 9 with a summary of the insights gained from these studies.

2. Related Literature

In this paper, we consider supply uncertainty in the form of *supply disruptions*, during which a portion of the supply chain is completely inoperative. Disruptions tend to be infrequent and temporary but cause a significant change to the system when they occur. In contrast, *yield uncertainty* refers to a form of supply uncertainty in which the quantity produced or received differs from the quantity ordered by a random amount. The effects of yield uncertainty are more frequent (possibly affecting every order) but less severe than those of disruptions. See Yano and Lee (1995) for a review of yield uncertainty.

Like this paper, the paper by Chopra et al. (2005) considers the relationship between two types of uncertainty—in particular, supply uncertainty in the form of both disruptions and yield uncertainty. They focus on the importance of not “bundling” the two forms of supply uncertainty by treating the entire supply distribution as though it comes only from yield uncertainty. Schmitt and Snyder (2006) extend their results to a multiple-period setting.

A third form of supply uncertainty is *lead time uncertainty*, in which the lead time for each order is a random variable. Lead time uncertainty may actually be the result of a deterministic supply process; for example, a firm following a base-stock policy will produce stochastic lead times for its customers since it will occasionally experience stockouts, resulting in delays for the backordered units. See Çakanyildirim and Bookbinder (1999) for a review of stochastic lead time models.

The threat of major supply chain disruptions has recently attracted the attention of researchers and authors from a wide range of settings. For example, calls for increased study of supply chain vulnerabilities have appeared in business journals (Sheffi 2001, Simchi-Levi et al. 2002, Stauffer 2003), the popular press (Lynn 2005, Sheffi 2005), academic research centers (Lee 2004, Rice et al. 2003), and the national academies (Wein and Liu 2005). Empirical studies have reported the significant financial impact of supply chain disruptions (Hendricks and Singhal 2003, 2005a,b) and firms' lack of preparedness for them (Hillman and Sirkisoon 2006, Poirier and Quinn 2006).

The majority of the research on supply disruptions considers a single-echelon system with an unreliable supplier. The supplier's failure and repair times are random, usually exponentially or geometrically distributed. These models are often based on classical models such as the EOQ (Berk and Arreola-Risa 1994, Gürler and Parlar 1997, Parlar and Berkin 1991, Parlar and Perry 1995, 1996, Snyder 2006b), EPQ (Moinzadeh and Aggarwal 1997), (R, Q) (Gupta 1996, Parlar 1997), and (s, S) (Arreola-Risa and DeCroix 1998) models. More nuanced supply processes are considered by Dada et al. (2003), Li et al. (2004), and Song and Zipkin (1996). Disruption models are generally much less tractable than their deterministic-supply counterparts and require numerical optimization since closed-form solutions are rarely available. (An exception is the approximation method of Snyder (2006b).)

While the papers cited above generally deal with tactical or operational decisions, a few papers examine strategic questions. Tomlin (2006) considers strategies for mitigating disruptions, while Tomlin (2004) examines the supplier selection problem when the reliability of each supplier is unknown and is revealed over time. Tomlin and Wang (2004) examine the benefits of product mix flexibility and supply diversification in a multi-product setting. Tomlin and Snyder (2006) consider optimal mitigation strategies when the disruption risk evolves stochastically over time and the firm has some advanced warning of imminent disruptions. They consider finite-horizon (as well as

infinite-horizon) problems and examine the system behavior near the end of the time horizon. Lewis et al. (2005) consider the impact of transportation disruptions (e.g., border closures) on lead times and show that policies and costs are more dependent on the expected duration of a disruption than on its probability.

To date, very little research has considered disruptions in a multi-echelon setting. Hopp and Yin (2006) study the optimal placement and extent of inventory buffers and redundant capacity to protect a supply chain against disruptions. They consider acyclic assembly networks. They show that, under certain assumptions, it is optimal for at most one node on each path to the customer to hold additional inventory or backup capacity, and they observe that the more severe the upstream disruptions may be, the further upstream the inventory or capacity should be placed. Kim et al. (2005) consider yield uncertainty in a three-echelon supply chain and, unlike most SU models, consider risk-averse objective functions.

In addition, a growing body of literature considers disruptions in the context of facility location (Berman et al. 2004, 2005, Church and Scaparra 2005, Scaparra and Church 2005, Snyder and Daskin 2005, Snyder et al. 2006) or location-inventory models (Qi and Shen 2005, Qi et al. 2006). In general, these studies find that the optimal number of facilities increases when the disruption risk increases, a result that is similar to the tendency toward diversification discussed in Section 6.1. Whereas these studies aim to choose optimal facility locations or network configurations, ours examines optimal strategies for coping with disruptions within a network whose topology is already established.

Finally, the mirror-image relationship between supply and demand was also discussed by Hariharan and Zipkin (1995) in the context of demand lead times. Their study focuses on the relationship between supply and demand lead times rather than between supply and demand uncertainty.

3. Model Assumptions and Experimental Design

3.1 Model Assumptions

Each of our simulation studies involves a simple multi-echelon supply chain under periodic review. Unless specified otherwise, all stages follow a base-stock policy. Each

stage consists of a processing function and an output buffer that stores finished goods inventory (though these “finished goods” may serve as raw materials for another stage). A stage may represent a physical location (factory, warehouse), a processing activity (production, assembly), or an SKU (raw material, component).

Unmet demands at each stage are backordered. Associated with each stage are the parameters h (the holding cost per unit of inventory per period), p (the stockout cost per unit of demand backordered, per period), and T (the processing time at the stage). Note that lead times take the form of processing times rather than transportation times; that is, they occur at the nodes rather than on the arcs of the network. However, the two notions are equivalent, and a system with one type of lead time can be transformed into a system with the other, whether or not disruptions are present. Other parameters will be introduced below as needed.

Under *demand uncertainty* (DU), the stage (or stages) at the furthest echelon downstream face iid random demands that are normally distributed with a mean of μ and a standard deviation of σ , denoted $N(\mu, \sigma^2)$. Under *supply uncertainty* (SU), the supply process at any stage that is subject to disruptions is governed by a two-state Markov chain in which the UP and DOWN states represents the non-disrupted and disrupted states, respectively. The process transitions from UP to DOWN with probability α (called the *failure probability*) and from DOWN to UP with probability β (the *repair probability*). Thus, failure and repair times are geometrically distributed with parameters α and β , respectively. One can easily show that the probability that the stage is UP or DOWN in a given period is given by $\beta/(\alpha + \beta)$ and $\alpha/(\alpha + \beta)$, respectively.

In each study, we test five pairs of values for (α, β) :

$$(\alpha, \beta) = (0.001, 0.1), (0.01, 0.3), (0.05, 0.5), (0.1, 0.7), (0.2, 0.9). \quad (1)$$

The first pair (0.001, 0.1) describes disruptions that are infrequent but long, while the last pair (0.2, 0.9) describes disruptions that are frequent but short. These characteristics make up what Tomlin and Snyder (2006) call the *disruption profile*.

During a disruption, no processing can take place at the disrupted stage; the disruption effectively pauses the production at that stage. Thus, if an item is in process at a stage whose processing time is T , and an m -period disruption occurs, then the item requires a total of $T + m$ periods to process. In addition, no inventory can be removed

from the output buffer of a disrupted stage. However, inventory is not destroyed by a disruption; it merely becomes unavailable until the disruption ends.

The simulations were performed using the software BaseStockSim v2.4 (Snyder 2006a). BaseStockSim is designed to simulate multi-echelon inventory systems and is available for download at www.lehigh.edu/~lvs2/software.html.


3.2 Experimental Design

In each study below, we first describe a “base” model that demonstrates the main point at issue in the study. We then simulate the model for a range of parameter values, generally varying the parameters one at a time from the base model. For each instance tested, we simulate the system for 10 trials, each of which consists of 10,000 periods. Each simulation is initialized by setting the failure state to UP (not disrupted) and the inventory level to S (the order-up-to level) at each stage. We use a 100-period warm-up interval (during which the costs and other performance measures are ignored). Given our initial conditions, the system tends to reach steady state rather quickly, which justifies the length of our warm-up interval.

BaseStockSim reports both the mean total cost across the 10 trials and the standard error of the mean (SEM)—i.e., the standard deviation of the mean cost across the trials—for each set of parameters tested. From this, we compute a 95% confidence interval as $\text{mean} \pm 1.96 \times \text{SEM}$, where $1.96 \approx |z_{0.05/2}|$. (See, e.g., Fishman (2001), Nelson (1995).) We say that one strategy is statistically better than another if its mean cost is smaller than that of the other strategy and the two confidence intervals are non-overlapping.

To evaluate a strategy accurately, the decision variables (e.g., base-stock levels) must be set nearly optimally. For example, it is unfair to claim that a decentralized inventory strategy outperforms a centralized one if the base-stock levels have not been set to near-optimal values in the two systems. In a few cases below, analytical methods are available for finding optimal levels. However, for most systems that we simulate, no such results are available in the literature. In these cases, we find the optimal base-stock levels approximately by performing a line search, simulating the system for a range of base-stock values (e.g., 100, 110, \dots , 200) and choosing the one with the smallest simulated total cost.

Figure 1: Two-stage supply chain (Study 1).



This approach, of course, has two main drawbacks. The first is that the true optimal value may not equal one of our chosen values exactly but may fall in between. The second is that a suboptimal solution may appear to be optimal due to simulation randomness. However, we feel that this methodology is sufficient since our aim in this paper is to derive qualitative insights rather than precise descriptions of optimal inventory levels. Moreover, the cost function is often reasonably flat around its optimum, suggesting some flexibility in the setting of decision variables.

In the sections that follow, we present seven studies, each of which addresses a particular strategic question related to ordering frequency, inventory placement, supply chain structure, and the cost of unreliability or reliability. Broadly speaking, these sections consider supply chains of gradually increasing complexity, as the questions we address become more complex.

4. The Cost of Unreliability

We begin with a study that evaluates the cost of *not* planning for uncertainty. We demonstrate that failing to plan for SU is a more costly proposition than failing to plan for DU.

4.1 Study 1: The Cost of Unreliability

4.1.1 Base Model

Consider a two-stage supply chain such as the one pictured in Figure 1. Stage 1 faces a demand with a mean of $\mu = 20$ per period and incurs a holding cost of $h = 1$ per unit of on-hand inventory and a stockout cost of $p = 20$ for each unmet demand. Stage 2 acts like an external supplier that cannot hold inventory. Stage 2 has zero processing time (when UP), while stage 1 has a 1-period processing time.

If either demand or supply is uncertain, a base-stock policy is optimal, and the optimal base-stock level S^* is greater than or equal to μ . (The optimality of a base-stock policy under SU is established by Tomlin (2006). That under DU is well known.)

Suppose, however, that the firm fails to plan for uncertainty; that is, it sets $S = \mu$, holding no safety stock. What is the cost of this policy?

To answer this question, we need a way of comparing DU and SU fairly. To this end, we define the *level of uncertainty* (LOU) for either type of uncertainty as the percentage of demands that are backordered (i.e., one minus the fill rate) when the base-stock level is set to $S = \mu$. We consider a DU process and an SU process to be *equivalent* if they have the same LOU.

Under DU, it is well known (see, e.g., Nahmias (2005)) that the fill rate in a single-period newsvendor-type system is given by $1 - \sigma\mathcal{L}(z)/\mu$, where $z = (S - \mu)/\sigma$ and $\mathcal{L}(z)$ is the standard normal loss function. Therefore, under DU the LOU is given by

$$LOU_D = \frac{\sigma\mathcal{L}(0)}{\mu} \approx 0.398942\frac{\sigma}{\mu}. \quad (2)$$

(To use (2) in a multi-period system with backorders, we are implicitly assuming that on-hand inventory is used to satisfy current-period demands before backorders. This is not the assumption made by our simulation, so the expression in (2) is only approximate.)

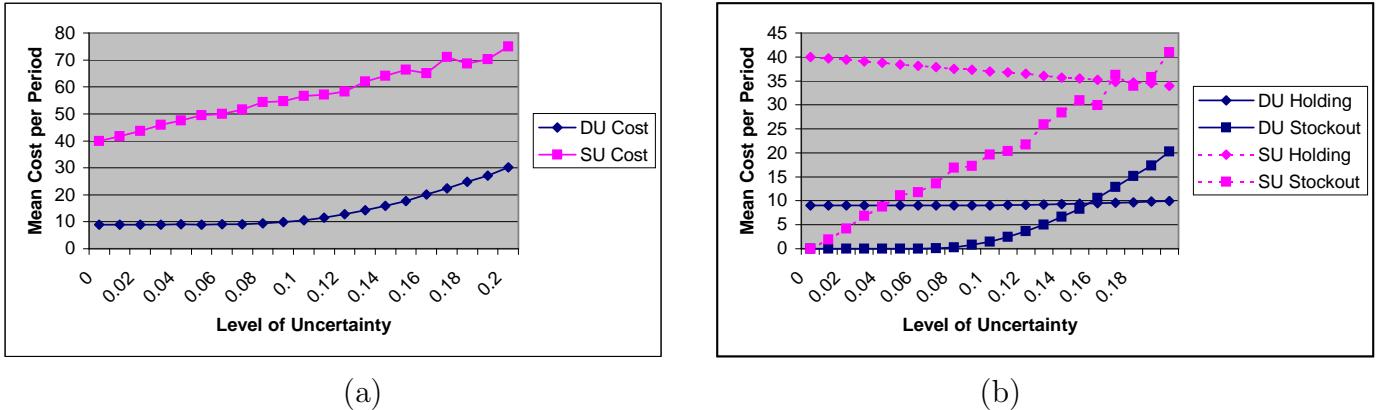
Under SU, if $S = \mu$, then the fill rate is exactly equal to the percentage of periods in which the supplier is UP, and therefore the LOU is given by

$$LOU_S = P(\text{supplier DOWN}) = \frac{\alpha}{\alpha + \beta}. \quad (3)$$

For the system in Figure 1, we varied the LOU from 0.0 to 0.2 in increments of 0.01 and found the values of σ and α that yield each LOU. (Under SU, we kept β fixed at 0.5.) We simulated the system with each value of σ under DU and α under SU. In all cases, we set $S = \mu$. The resulting mean total costs per period are displayed in Figure 2(a). Clearly, a given level of SU results in higher costs than an equivalent level of DU, and the mean total cost grows more quickly under SU than under DU. We can report anecdotally that similar patterns occur when S is set to a value larger than μ , although we did not test this finding rigorously.

Figure 2(b) breaks the total cost down into its two components, holding and stock-out costs. Under DU, the holding cost increases slightly and the stockout cost increases more sharply as the LOU increases. This is consistent with well known properties of base-stock systems under DU. Under SU, however, the stockout cost increases sharply with LOU while the holding cost *decreases*. The reason for the decrease in holding cost

Figure 2: Mean (a) total costs and (b) holding and stockout costs, for each LOU under DU and SU.



as LOU increases is that the on-hand inventory can fall below the safety stock level, $S - \mu$, during a disruption. The greater the LOU, the smaller the expected on-hand inventory. In contrast, under DU, the inventory level can fall either above or below the safety stock level.

To confirm that the cost under SU is greater than that under DU in general, we perform a more extensive computational test in the next section.

4.1.2 Extended Study

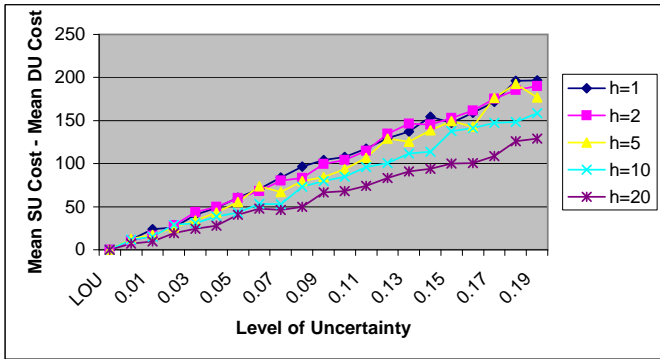
We generated 13 new instances of the system depicted in Figure 1 by varying the holding and stockout costs, as well as β for the SU tests. For each instance, we found the values of σ and α that resulted in LOUs equal to 0.00, 0.01, 0.02, ..., 0.20. As in Section 4.1.1, we set $S = \mu$ and simulated the system under each LOU.

The results are displayed in Figure 3, which plots the difference in cost (SU cost minus DU cost) for varying values of the holding cost (h), stockout cost (p), and recovery probability (β), respectively, as the LOU changes.¹ In nearly all cases, this difference is strictly positive and is statistically significant, confirming that the cost under a given level of SU is greater than that under an equivalent level of DU. In the few cases in which the SU cost is smaller than the DU cost, the magnitude of the difference is small and is likely a result of simulation randomness.

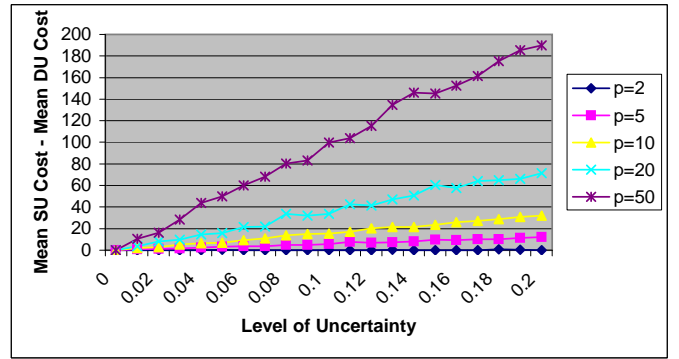
From Figure 3(a), it is evident that the magnitude of the difference between SU and

¹We omit detailed results of our computational experiments in table form due to space considerations. They are available from the authors upon request.

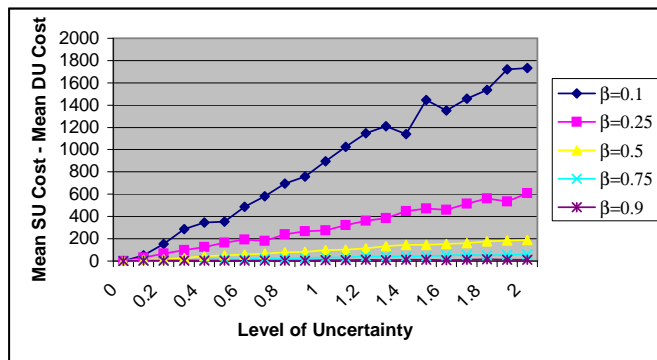
Figure 3: Mean total cost vs. LOU for varying values of (a) h , (b) p , and (c) β .



(a)



(b)



(c)

DU is greater when h is smaller. This is because, as discussed in Section 4.1.1, there is more inventory on-hand under DU than under SU; therefore, the DU cost increases as h increases and the difference between SU and DU decreases. From Figure 3(b), the cost difference increases as p increases, since backorders tend to linger longer under SU than under DU. Finally, from Figure 3(c), the cost difference increases as β decreases, since smaller β means longer disruptions.

Most firms already plan for DU, but this analysis suggests that it is even more critical to plan for SU. If SU is present, it is likely to represent a substantial portion of the firm's inventory-related costs (holding and stockout costs). Although the strategies described in subsequent sections of this paper are effective in reducing the cost under SU, they generally cannot eliminate it completely, or even reduce it to the levels of the DU cost. Additional research is warranted to explore alternate strategies for coping with SU, such as dual sourcing (Tomlin 2006) or product-mix flexibility (Tomlin and Wang 2004) that can reduce this cost further.

We summarize our findings from this study in the following remark.

Remark 1 *If the base-stock level is set to the mean demand (no safety stock), the cost of a system under a given level of SU is greater than that under an equivalent level of DU. The difference in cost is greater when the holding cost is smaller, the stockout cost is larger, or the recovery probability is smaller.*

5. Order Frequency

In the study that follows, we consider the following question: Would a firm prefer a large holding cost but zero fixed ordering cost or a small holding cost but a non-zero fixed cost? The firm follows a one-for-one replenishment policy under the first option and a batch-ordering policy under the second. Roughly speaking, the firm may be indifferent to the two options under DU. However, we show that batch ordering is preferable under SU since disruptions only have an impact during periods in which orders are placed.

5.1 Study 2: One-for-One vs. Batch Ordering

5.1.1 Base Model

Consider again the two-stage supply chain pictured in Figure 1. Now we assume that $p = 100$, and the 1-period processing time occurs at stage 2 rather than at stage 1.

First consider the case of DU. In this case, demand is normally distributed with a standard deviation of $\sigma = 5$. The question is, would the firm prefer a holding cost of $h = 0.1$ and a fixed order cost (per order placed to stage 2) of $K = 250$, or a holding cost of $h = 2.85$ and fixed cost of $K = 0$?

In the first option, the optimal policy at stage 1 is an (s, S) policy. We can find near-optimal values for s and S using the well known heuristic of first finding the optimal values of R and Q for an (R, Q) policy under the same parameters, and then setting $s = R$, $S = R + Q$ (Axsäter 2000). Using this approach, we get $R = 31$, $Q = 318$, and therefore $s^* = 31$, $S^* = 349$. The expected cost of this policy is 32.8 per period, calculated using equation (8.16) of Simchi-Levi et al. (2005).

Under the second option, the optimal policy is a base-stock policy with

$$S^* = \mu + \sigma \Phi^{-1} \left(\frac{p}{p+h} \right) \quad (4)$$

$$= 20 + 5\Phi^{-1}(0.9723) \approx 30, \quad (5)$$

where Φ^{-1} is the inverse of the standard normal cumulative distribution. Using the fact that the expected cost of this policy is given by

$$E[\text{optimal cost}] = (h+p)\sigma\phi \left(\Phi^{-1} \left(\frac{p}{p+h} \right) \right), \quad (6)$$

where ϕ is the standard normal pdf (Porteus 2002, Zipkin 2000), the expected cost of this policy is 32.8 per period. Thus, under DU, the firm is indifferent between these two cost structures: $h = 0.1, K = 250$ or $h = 2.85, K = 0$.

Now consider SU with disruptions possible at the supplier (stage 2). The failure and repair probabilities at stage 2 are $\alpha = 0.05$ and $\beta = 0.5$, respectively. If there is no fixed cost, then a base-stock policy is optimal, and the base-stock level for stage 1 can be set using the following formula, due to Tomlin (2006):

$$S^* = \mu \left[1 + F^{-1} \left(\frac{p}{p+h} \right) \right], \quad (7)$$

where F is the cumulative distribution of a random variable X defined as the number of consecutive DOWN periods (or 0 if the period is an UP period). Omitting the details of this calculation, we get $S^* = 60$ if $h = 2.85$. Simulating this system, we find a mean cost of 497.7 per period.

The literature provides no guidance on setting s and S under SU if the fixed cost is non-zero, though we do know that both should be multiples of the demand, μ . Therefore, in the case in which $h = 0.1$ and $K = 250$, we find approximate values for s and S by setting $S = 340$ (obtained from the DU case) and testing $s = 0, 20, 40, 60$. Simulating the system for each value of s , we find that the minimum cost occurs when $s = 40$, with a mean cost of 391.1. Therefore, batch ordering is preferred to one-for-one ordering in the case of SU. The mean cost for all four cases (batch vs. one-for-one ordering under SU and DU) and the SEM of the costs for the SU cases are summarized in Table 1. (The DU costs were computed analytically, rather than by simulation, so SEMs are not required.) Under SU, the difference in cost is significant at the 95% level, since the confidence intervals are non-overlapping ($391.1 + 1.96 \times 29.6 < 497.7 - 1.96 \times 23.4$).

Table 1: Summary of inventory parameters and mean and SEM of costs (Study 2).

Type of Uncertainty	Measure	Batch Ordering	One-for-One Ordering
DU	s	31	—
	S	349	30
	Mean Cost ^a	32.8	32.8
SU	s	40	—
	S	340	60
	Mean Cost ^b	391.1	497.7
	SEM of Cost ^c	29.6	23.4

^aMean total cost per period, obtained analytically.

^bMean total cost per period, across all trials and all periods period in each trial, obtained by simulation.

^cSD (across all trials) of average total cost per period (across periods), obtained by simulation.


The reason that batch ordering is preferred under SU but not under DU is that, the less frequent the orders are, the less likely it is that a disruption will have an effect. If $s = 40$, $S = 340$, then orders are placed roughly every 15 periods. Many of stage 2's disruptions will start after an order and end before the next order, and stage 1 will suffer no ill effects from these disruptions. In contrast, in the base-stock policy, stage 1 places an order in every period, so it suffers from every one of stage 2's disrupted periods.

Moreover, if stage 2 is disrupted when stage 1 places an order, stage 1 incurs μ stockouts per period once the safety stock runs out until the disruption ends, whether it uses batch or one-for-one ordering. The optimal safety stock levels— $S - \mu$ under the one-for-one policy, $s - \mu$ under the batch policy—tend to be roughly equal to each other. Once the disruption ends, all outstanding orders are filled, and the backlog clears. Therefore, the *severity* of a disruption's impact is the same under both policies, but the *likelihood* of stage 1 feeling a given disruption is smaller under the batch policy. Thus, the batch policy is preferred.

5.1.2 Extended Study

We now consider a range of values for the cost and disruption parameters to demonstrate that the effect discussed in Section 5.1.1 holds more broadly than for just the specific instance given. We generated 100 instances for the DU case by setting $K = 250$, $p = 10, 20, \dots, 100$ and $h = 0.01, 0.02, \dots, 0.1$. For each combination of these values, we first calculated a near-optimal (s, S) policy using the method described in Section 5.1.1. We then set $K = 0$ and found the value of h (call it \hat{h}) that results in the

Figure 4: Mean costs for (s, S) and base-stock policies under (a) DU and (b) SU.




same expected cost when the base-stock level is set optimally using (4). Therefore, for each of the 100 instances, we have an (s, S) policy and a corresponding base-stock policy with the same expected cost under DU.

We simulated each of these instances under both policies (even though analytical expressions for the expected cost are available for DU). The resulting mean costs are plotted in Figure 4(a). The x -axis plots the cost of the base-stock policy while the y -axis plots that of the (s, S) policy. Notice that the points are clustered tightly along the line $y = x$, confirming that the two policies produce approximately the same expected cost for each set of the parameter values.

Next, we set (α, β) to the five pairs of values in (1). For each of the resulting $100 \times 5 = 500$ instances, we simulated a batch policy and a one-for-one ordering policy. In the one-for-one policy, we used holding costs \hat{h} and set the base-stock level set optimally using (7). In the batch policy, we used holding costs h and $K = 250$; we set S equal to the S from the DU case and chose the value of $s = 0, 20, 40, 60$ that resulted in the lowest mean cost per period.

The resulting mean costs for the SU simulations are plotted in Figure 4(b). It is apparent from the figure that the batch policy tends to be preferable to the one-for-one policy. Indeed, the batch policy has a smaller cost for 286 out of 500 instance (57.2%). The cost difference is only statistically significant for 40 of the 500 instances (8%). The one-for-one policy tends to be preferred when the ratio of h to p is large (Figure 5); in these cases, the high holding cost incurred in the batch policy outweighs the benefit, since disruptions and the resulting stockouts are not particularly costly.

Figure 5: Percentage difference in cost between (s, S) and base-stock policies vs. h/p .



We summarize our findings from this study in the following remark.

Remark 2 *Under DU, a firm is indifferent between a batch and a one-for-one ordering policy, provided that the parameters are chosen such that the expected costs of the two policies are the same. However, under SU, the batch ordering policy is preferable, since the severity of a disruption’s impact is roughly the same under both policies, but the likelihood of the disruption affecting a given order is smaller under the batch policy.*

6. Inventory Placement

In this section we present two studies that examine fundamental questions of inventory placement in multi-echelon supply chains. In Study 3, we consider the question of centralization in a one-warehouse, multi-retailer (OWMR) system. Under DU, we know that centralization is optimal, but we show that decentralization is preferable under SU. Study 4 examines whether it is preferable to hold inventory upstream or downstream in a serial system. Again, we show that the conventional wisdom from DU does not hold under SU.

6.1 Study 3: Centralization vs. Decentralization

6.1.1 Base Model

Figure 6 depicts a one-warehouse, multi-retailer (OWMR) system with three retailers. Eppen (1979) proves that, if the demand is iid random at each retailer and the holding costs are equal at the two echelons, then it is optimal to store inventory at the warehouse (stage 4) rather than at the individual retailers (stages 1–3). This is the classical *risk-pooling effect*: The centralized system (in which all inventory is held at the warehouse)

Figure 6: OWMR system (Study 3).




Table 2: Summary of base-stock levels and mean, SEM, and SD of costs (Study 3).

Type of Uncertainty	Measure	Decentralized System	Centralized System
DU	S	30	77
	Mean Cost ^a	51.29	29.61
SU	S	40	180
	Mean Cost ^b	785.30	793.10
	SEM of Cost ^c	38.08	45.28
	SD of Cost ^d	2549.21	4031.98

^aMean total cost per period, obtained analytically.

^bMean total cost per period, across all trials and all periods period in each trial, obtained by simulation.

^cSD (across all trials) of average total cost per period (across periods), obtained by simulation.

^dSD of total cost per period, across all trials and all periods in each trial, obtained by simulation.

is cheaper than the decentralized one (in which inventory is held only at the retailers) since centralization causes a pooling of the demand variance across retailers.

Suppose that $h = 1.5$ at all stages and $p = 50$ at the retailers. Demands at the retailers are iid $N(20, 5^2)$. The warehouse has a processing time of $T = 1$ while the processing time at the retailers equals 0. A base-stock policy is optimal in both the centralized and decentralized systems. The optimal base-stock level can be calculated using (4), with $\mu = 20$, $\sigma = 5$ in the decentralized system and $\mu = 60$, $\sigma = \sqrt{3 \times 5^2} = 8.66$ in the centralized system. The costs of these solutions, obtained using (6), are 51.29 and 29.61, respectively, confirming that centralization is optimal for this instance. Table 2 summarizes these findings.

Now consider the case of SU. We assume that $\sigma = 0$ and that disruptions are possible at the stage(s) that hold inventory (that is, at the warehouse in the centralized system and at the retailers in the decentralized system), with failure and repair probabilities $\alpha = 0.05$ and $\beta = 0.5$, respectively. We assume that all stages follow a base-stock policy. Unfortunately, the form of the optimal policy is not known, nor are the optimal base-stock levels. (Equation (7) no longer applies because it assumes that disruptions occur upstream from the site holding inventory, while we assume that the disruption

occurs at the site.) One can show, however, that the optimal base-stock level is a multiple of $\mu = 20$.

Therefore, we simulated this system using $S = 0, 20, 40, 60$ in the decentralized system and $S = 60, 120, 180, 240$ in the centralized system. The optimal values of S and the corresponding mean, SEM, and standard deviation of the cost per period are listed in Table 2. The mean costs are quite close; since the confidence intervals are overlapping, no conclusions can be drawn about which system has the lower mean cost. On the other hand, the SD of the cost per period is significantly smaller in the decentralized system.

These findings can be explained as follows. The retailers experience disruptions in the same percentage of periods in the two systems, though the disruptions originate from the warehouse in the centralized system. Therefore, the mean costs are approximately equal in the two systems. In the centralized system, disruptions are less frequent (because there is one site to be disrupted rather than three) but more severe (since each disruption affects all three retailers), hence the centralized system produces greater variance in the cost.


From this preliminary example, we conclude that *centralization is preferred under DU* due to the risk-pooling effect, but that *under SU, decentralization produces roughly the same mean cost but less variability in cost than centralization*. In the following section, we perform a more extensive computational study to verify that these results hold under a broader range of parameters.

6.1.2 Extended Study

We created a number of new data sets based on the base case described in Section 6.1.1 by varying h_1, \dots, h_4 and p_1, \dots, p_3 . For the SU tests, we also varied the failure and repair probabilities (α, β) using the values in (1). This resulted in 21 instances for DU and 105 for SU, since each of the parameter settings is tested under five values of (α, β) . In all tests, $h_4 \leq h_1 = h_2 = h_3$ (warehouse holding costs are no greater than retailer ones).

Optimal base-stock levels for both the centralized and decentralized systems are available analytically for DU (see Section 5.1.1). For SU, the optimal base-stock level is a multiple of the demand per period; therefore, we simulated the decentralized system using $S_1, \dots, S_3 = 0, 20, 40, 60$ and the centralized system using $S_4 = 60, 120, 180, 240$.

Figure 7: Mean optimal costs for decentralized and centralized systems under (a) DU and (b) SU.



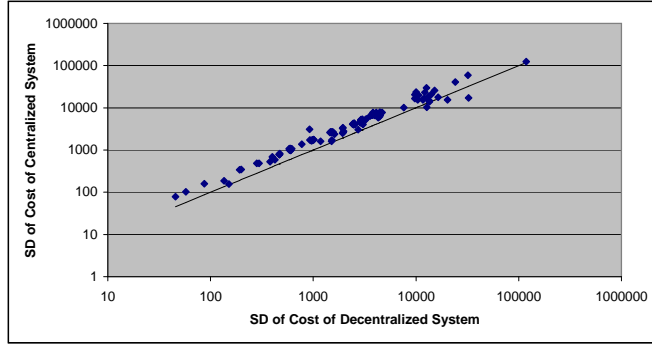
The optimal costs for the DU tests are displayed in Figure 7(a). Each point in the figure represents one of the 21 tests. The horizontal axis represents the optimal cost of the decentralized system while the vertical axis represents that of the centralized system with the same input parameters. These costs are generated via simulation for consistency with other studies, even though analytical formulas are available. (The simulated costs differed from the analytical costs by an average of 0.27%.)

As expected, the centralized system always outperforms the decentralized one under DU. This is evident from the fact that all points in Figure 7(a) lie below the line $y = x$. On average, the centralized system is 58.1% less expensive than the decentralized one, and all confidence intervals are non-overlapping (so each of the differences is statistically significant). For instances in which upstream holding costs are strictly less than downstream ones, the savings come from both the smaller inventory requirements and the lower holding cost per unit under centralization.

The SU results are displayed in Figure 7(b). (Note that the axes are drawn using a logarithmic scale.) The tight clustering of the points along the line $y = x$ suggests that the mean cost under the centralized and decentralized strategies are roughly equal, as in the example in Section 6.1.1. Indeed, of the 105 instances tested, only 12 (11.5%) resulted in non-overlapping confidence intervals for the mean cost. In all of these instances, the upstream holding costs are strictly less than downstream ones, which tends to bias the costs in favor of centralization because of holding costs rather than because of uncertainty.

Figure 8 displays the standard deviation of the cost per period (across all trials and

Figure 8: Standard deviation of costs for decentralized and centralized systems under SU.



periods per trial) under the decentralized and centralized systems. The SD is greater for the centralized system for 102 out of 105 instances (97.1%), confirming our suggestion from Section 5.1.1 that decentralization results in less variability in cost. The three instances for which the SD is smaller under centralization all have $\alpha = 0.001$, $\beta = 0.1$, which results in the least frequent (but longest) disruptions, the greatest amounts of simulation variability overall, and the least statistically significant results.

We summarize our findings from these experiments in the following remarks.

Remark 3 *Under DU, centralization is always the preferred strategy in an OWMR system. This is due to the classical **risk-pooling effect** and to smaller (or equal) holding costs at the upstream stage.*

Remark 4 *Under SU, decentralization is the preferred strategy in an OWMR system with equal holding costs at upstream and downstream stages. By “preferred,” we mean that decentralization results in an approximately equal mean cost but less cost variability than centralization does. This is due to what we call the **risk-diversification effect**, which says that disruptions are equally frequent in either a centralized or decentralized system but less severe in a decentralized one. If the holding costs are strictly less upstream than downstream, the preferred strategy depends on the balance between the difference in holding costs and the magnitude of the risk-diversification effect.*

6.2 Study 4: Inventory Placement

6.2.1 Base Model

Consider the three-stage serial supply chain pictured in Figure 9. Stage 3 represents a factory, stage 1 represents a retailer, and stage 2 represents an intermediate facility that

Figure 9: Serial system (Study 4).




Table 3: Summary of base-stock levels and mean and SEM of costs (Study 4).

Type of Uncertainty	Measure	Upstream	Downstream
DU	S	31	29
	Mean Cost ^a	12.24	21.57
	SEM of Cost ^b	0.17	0.15
SU	S	20	60
	Mean Cost ^a	188.62	126.56
	SEM of Cost ^b	17.94	8.6

^aMean total cost per period, across all trials and all periods period in each trial, obtained by simulation.

^bSD (across all trials) of average total cost per period (across periods), obtained by simulation.

performs some operation (like packaging or cross-docking) but cannot store inventory. The question we address in this section is whether it is preferable to hold inventory at the upstream stage (3) or the downstream stage (1).

Under DU, the conventional wisdom is that, lead times aside, it is preferable to hold inventory upstream since holding costs generally increase as one moves downstream. On the other hand, suppose that disruptions are possible at stage 2—say, because it is located in an area with poor weather, or because its workforce is prone to strikes. In this case, it may be preferable to hold inventory downstream since such inventory can protect against disruptions at stage 2, while upstream inventory is ineffective in that regard.

In our base instance, we set the processing time (T) at stages 1 and 2 to 0 and at stage 3 to 1. Holding costs are given by $h_1 = 2$, $h_3 = 1$, and the stockout cost is $p = 50$. Consider the DU case first. Demands are normally distributed with mean $\mu = 20$ and SD $\sigma = 5$. It is well known that a base-stock policy is optimal whether we hold inventory upstream or downstream (Clark and Scarf 1960), and since only stage 3 has a non-zero processing time, we can set the base-stock level at either location using (4). The optimal base-stock levels and costs are listed in Table 3. As expected, it is optimal to hold inventory upstream at stage 3, and the cost difference is statistically significant.

Under SU, we set the disruption and repair probabilities at stage 2 to $(\alpha, \beta) = (0.05, 0.5)$ and find the corresponding optimal base-stock levels for both the upstream

and downstream strategies by simulation. The results are also reported in Table 3. In this case, the cost to hold inventory downstream is (statistically significantly) smaller than that to hold inventory upstream, since downstream inventory can help to buffer against disruptions. That also explains why the base-stock level is 60 downstream (versus only 20 upstream)—the additional 40 units serve as safety stock inventory to use during disruptions, whereas there is no reason to hold additional inventory under the upstream strategy.

In the next section, we conduct similar experiments on a broader set of instances.

6.2.2 Extended Study


We created several additional tests sets based on the instance in Section 6.2.1 by varying h_1 and p_1 . For the SU tests, we tested the five values of α and β from (1) at stage 2, resulting in 9 instances for DU and 45 for SU. Optimal base-stock levels under DU were set using (4), while those for SU were set by testing $S = 20, 40, 60, 80$. (The optimal base-stock level under SU is a multiple of the demand per period.)

The results of the DU experiment are displayed in Figure 10(a), in which the horizontal axis represents the cost under the upstream strategy and the vertical axis represents that under the downstream strategy. (The stacked points at $x \approx 10.5$ represent five instances in which $h_3 = 1$ and $p = 20$, each of which yields the same cost for the upstream strategy.) From the figure it is evident that the upstream strategy dominates the downstream one. Except for the one instance in which $h_1 = h_3 = 1$ (in which the two strategies have equal costs), the upstream strategy is statistically better than the downstream one for every instance.

The results for SU are given in Figure 10(b). The downstream strategy dominates in this case, in the sense that no instance has a lower upstream cost than downstream. In 10 of the 45 cases (22.2%), the upstream strategy is strictly cheaper, and in 8 of those 10 cases, the difference is statistically significant.

However, the two strategies produce equal costs for the majority of instances. At first this seems puzzling, since the holding costs may be different at stages 1 and 3. To explain these results, we first note that the optimal base-stock level in the upstream strategy is always 20, i.e., the demand per period, since additional inventory will not help during a disruption. Therefore, under the upstream strategy, stage 3 ends every period with an inventory level of at most 0, so it incurs no holding costs, only stockout

Figure 10: Mean optimal costs for upstream and downstream strategies under (a) DU and (b) SU.




costs. Under the downstream strategy, if the optimal base-stock level at stage 1 also happens to be 20, then it, too, will incur no holding costs. In this case, neither the upstream nor the downstream strategy holds safety stock inventory—both strategies simply accept the disruption risk without protecting against it. Indeed, this is the case for all of the instances that lie on the line $y = x$ in Figure 10(b).

The optimal base-stock level under the downstream strategy tends to increase as α increases. That is, as disruptions become more frequent, it becomes preferable to hold safety stock to protect against them, rather than simply accepting the disruption risk. This trend is evident from Figure 11, which plots the percentage savings from holding inventory downstream on the vertical axis and α_2 on the horizontal one. As α_2 increases, we find more instances for which the optimal base-stock level downstream is greater than 20, and consequently more instances for which the downstream strategy is strictly better than the upstream one. In addition, the magnitude of the savings tends to increase as α_2 increases.

We acknowledge that the uniform dominance of the downstream strategy over the upstream one under SU is a result of our assumptions that disruptions only occur at stage 2 and that inventory cannot be used by a stage while it is disrupted. If stage 3 could experience disruptions and could use its inventory while disrupted—in other words, if safety stock at stage 3 had some value—then the upstream strategy would sometimes be optimal, depending on the disruption characteristics and on the relative holding costs between stages 1 and 3.

We summarize our findings from these experiments in the following remarks.

Figure 11: Relative savings from the downstream strategy under SU vs. failure probability at stage 2.



Remark 5 *In a serial supply chain under DU, it is optimal to hold inventory upstream rather than downstream, assuming that there is zero lead time between the upstream inventory point and the customer. If this assumption does not hold, the upstream strategy will tend to be less attractive as the lead time between the upstream inventory point and the customer increases and to be more attractive as the ratio between the upstream and downstream holding costs decreases.*


Remark 6 *In a serial supply chain under SU, it is optimal or equivalent to hold inventory downstream rather than upstream, regardless of the holding costs at each site, assuming that upstream inventory has no value in protecting against disruptions. If this assumption does not hold, the downstream strategy will tend to be less attractive as the ratio between the upstream and downstream holding costs decreases and to be more attractive as the lead time between the upstream inventory point and the customer increases.*

7. Supply Chain Structure

The three studies that follow each consider a question of supply chain structure. Study 5 compares “hub-and-spoke” and “point-to-point” networks (i.e., networks with an intermediate warehouse echelon and those in which plants ship directly to retailers). It examines the preferred network under each type of uncertainty and relates the results to risk pooling and risk diversification (Section 6.1).

Studies 5 and 6 examine the benefit of multiple sourcing. Study 6 considers a single retailer with one or more suppliers, examining the value of “backup” suppliers

Figure 12: (a) Hub-and-spoke (H-S) and (b) point-to-point (P-P) networks (Study 5).



under both DU and SU. Study 7 considers multiple retailers and an equal number of suppliers, varying the number of suppliers that may serve each retailer. Borrowing from the notion of process flexibility (Jordan and Graves 1995), we evaluate the change in cost as additional links are added to the warehouse–retailer network.

7.1 Study 5: Hub-and-Spoke vs. Point-to-Point Networks

7.1.1 Base Model

In this section we consider a question of supply chain design—in particular, the appropriate number of echelons for a distribution system. Figure 12 depicts two potential designs, which we refer to as “hub and spoke” (H-S; Figure 12(a)) and “point to point” (P-P; Figure 12(b)). In the H-S design, the factory (stage 9) distributes product to two warehouses (stages 7 and 8), which hold inventory and, in turn, distribute product to the retailers (stages 1–6). In the P-P design, the factory bypasses the intermediate warehouse level, distributing product directly to the retailers, which hold the inventory. The factory does not hold inventory in either system.

The H-S approach is a popular one for a number of reasons, primarily related to economies of scale and reduced holding costs that arise from holding inventory at warehouses rather than at individual retailers. Therefore, we expect that this design is optimal under DU. On the other hand, we have seen in Study 3 that decentralization may be optimal under SU due to the risk-diversification effect, which leads us to suspect that the P-P design may be optimal under SU.

Our initial test of these hypotheses involves the following parameters. In the H-S network, the processing time at the factory is $T_9 = 4$ and that at the warehouses is

Table 4: Summary of base-stock levels and mean and SEM of costs (Study 5).

Type of Uncertainty	Measure	H-S Network	P-P Network
DU	S	320	110
	Mean Cost ^a	116.90	135.83
	SEM of Cost ^b	0.80	1.81
SU	S	540	150
	Mean Cost ^a	778.07	437.81
	SEM of Cost ^b	24.38	13.4

^aMean total cost per period, across all trials and all periods period in each trial. Values for DU were obtained analytically; values for SU were obtained by simulation.

^bSD (across all trials) of average total cost per period (across periods), obtained by simulation.

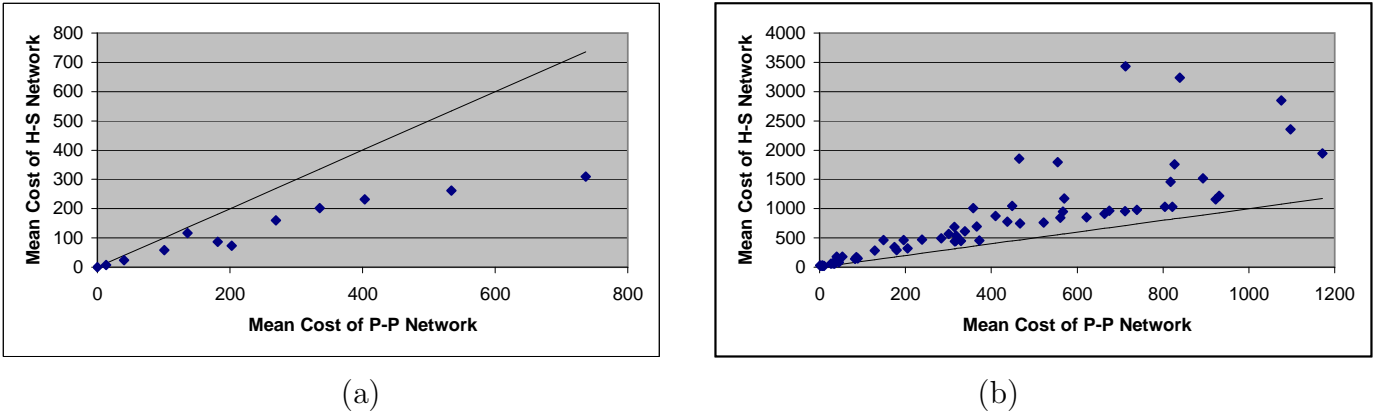
$T_7 = T_8 = 1$. In the P-P network, the factory takes on the processing function that the warehouses perform in the H-S system, so the processing time at the factory is $T_9 = 5$. The holding cost at the warehouses (in the H-S network) and at the retailers (in the P-P network) is $h = 1$, and the stockout cost is $p = 10$. Each stage follows a base-stock policy. Demand per period is $N(20, 5^2)$ at each retailer under DU (and equals 20 under SU). Disruptions can occur at both the plant and the warehouses (if they exist) under SU, with $\alpha = 0.05$ and $\beta = 0.5$.

The optimal base-stock levels and the mean and SEM of the total cost for each network design and for both types of uncertainty are listed in Table 4. As expected, the H-S design is optimal under DU, while the P-P design is optimal under SU. Both differences in cost are statistically significant. As in Study 3, under DU the savings from centralization (holding inventory at two warehouses rather than six retailers) comes from the risk-pooling effect, while the savings from decentralization under SU comes from the risk-diversification effect. We perform a more extensive numerical experiment in the next section to confirm these results.

7.1.2 Extended Study

We modified the base instance described in Section 7.1.1 to create a number of additional test instances by varying the stockout cost at the retailers and the holding cost at the warehouses in the H-S network and at the retailers in the P-P network. In all instances, stages of a given type (warehouse, retailer) have identical parameters. For the SU tests, we tested the five values of the failure and repair probabilities given in (1) at the warehouses (stages 7 and 8) in the H-S network, and at the plant (stage 9) in both networks. This resulted in 12 tests for DU and 60 for SU. Under DU, we

Figure 13: Mean optimal costs for P-P and H-S networks under (a) DU and (b) SU.



tested base-stock levels 240, 260, \dots , 360 at the warehouses in the H-S network and 50, 70, \dots , 170 at the retailers in the P-P network. Under SU, we tested base-stock levels 460, 500, \dots , 700 in the H-S network and 70, 90, \dots , 190 in the P-P network.


The results for DU are displayed in Figure 13(a). As is evident from the figure, the mean cost for the H-S network is less than that for the P-P network for every instance, and this difference is statistically significant for every instance. The opposite result holds in the SU case, displayed in Figure 13(b), in which the P-P network is preferred for every instance. The cost difference is statistically significant for 47 out of 60 (78.3%) of the instances. The percent difference in cost between the two networks exhibited no clear trend with respect to h , p , h/p , or the disruption parameters.

As suggested in Section 7.1.1, these results can be viewed as an extension of those in Study 3: centralization of inventory is optimal under DU, while decentralization is optimal under SU. Note that, unlike in Section 7.1.1, in this section we assumed that the holding cost is equal at either echelon (warehouse or retailer). We summarize the results from this section in the following remarks.

Remark 7 *In a distribution network under DU, a hub-and-spoke network (in which plants ship to warehouses, which hold inventory and ship to retailers) is preferred to a point-to-point network (in which plants bypass the warehouses and ship directly to retailers, which hold the inventory). This is because of the risk-pooling effect, which encourages consolidation of inventory stocking points.*

Remark 8 *In a distribution network under SU, a point-to-point network is preferred to a hub-and-spoke network because of the risk-diversification effect, which encourages*

Figure 14: One retailer with one or more suppliers (Study 6).



decentralization of inventory stocking points.

7.2 Study 6: Supplier Redundancy

7.2.1 Base Model


In this section, we investigate the benefit that arises from multiple sourcing. Consider a single retailer with one or more suppliers (Figure 14). The suppliers are identical in terms of cost and other parameters. The value of having multiple suppliers arises from the fact that each supplier has a fixed processing capacity (defined as the maximum number of units that may *begin* production in a single period), so additional suppliers can fill in when the primary supplier is maxed out or disrupted.

Under DU, the second supplier (stage 3 in Figure 14) is used only when the demand exceeds the capacity of the primary supplier, and the third supplier (stage 4) is used only when the demand exceeds the sum of the capacities of stages 2 and 3. The savings that results from the backup suppliers should therefore decrease as the capacity increases. Moreover, assuming that the capacity is strictly greater than the mean demand, we should expect to see a moderate savings from the first backup supplier but virtually no savings from the second backup supplier, since the probability that the demand is more than twice the mean is essentially 0.

Under SU, the backup suppliers play two roles. The first is to step in when the primary supplier fails. The second is to provide additional capacity to help recover after a *total* disruption in which all suppliers fail simultaneously. In our tests, we assume that each supplier's capacity (when UP) is sufficient to meet the demand per period. Therefore, the backup suppliers' ability to perform their first role is largely independent of the supply capacity, while it is critically dependent for the second role.

In our base case, the mean demand per period is 20, with a standard deviation of 5 under DU. The costs at the retailer are given by $h = 1$ and $p = 20$. We simulated the DU systems with one, two, and three suppliers, with the base-stock level at stage

Figure 15: Mean optimal costs for retailer with one, two, and three suppliers as a function of capacity under (a) DU and (b) SU.



1 set to 20, 24, 28, and 32, then chose the base-stock level that resulted in the lowest mean cost per period. The results are displayed in Figure 15(a), which plots the mean optimal cost per period as a function of the capacity per supplier. As expected, the cost for the single-supplier system decreases as the capacity increases. However, the two- and three-supplier costs remain flat as the capacity increases, since those systems have ample capacity even at the lowest capacity level. In addition, the benefit from the second supplier is substantial when the capacity equals 25 but is much smaller for greater capacity values, since a capacity of 30 or greater is sufficient to meet the entire demand with 98% probability ($30 = \mu + 2\sigma$). Finally, the third supplier provides essentially no benefit since a two-supplier system is more than adequate to meet demand. (The difference in cost between the two- and three-supplier systems is not statistically significant, nor is that between the one- and two-supplier systems with capacity 30 or greater.)

In contrast, the same tests performed under SU, with the failure and repair probabilities at each supplier set to $(\alpha, \beta) = (0.05, 0.5)$, indicate that the second supplier continues to provide some benefit even as the capacity continues to increase (Figure 15(b)). This is due to the fact that the second supplier is used whenever there is a disruption, and the disruption risk is independent of the supplier capacity. (In these trials, we tested values of 20, 30, 40, 50, and 60 for the base-stock level at the retailer. Note that the optimal base-stock level is not necessary a multiple of the demand under SU because of the finite supplier capacity.)

Similarly, the single-supplier cost continues to decrease as the capacity increases

because the additional capacity helps the system recover quickly after a disruption. The cost when the capacity equals 35 is 13.3% smaller than that when capacity equals 30, and the cost when capacity equals 40 is 13.2% smaller again than that when capacity equals 35. Under DU, on the other hand, the corresponding cost differences are 1.4% and 0.6%, respectively. (The differences are not statistically significant under either DU or SU.)

As in the DU case, the two- and three-supplier curves in Figure 15(b) are flat. This is because the benefit of the backup suppliers comes less from the additional capacity they provide and more from their simple presence during disruptions. In other words, as long as a backup supplier's capacity is sufficient to meet the demand (in this case, 20), it can provide adequate supply during a disruption to the primary supplier.

Unlike under DU, the third supplier does provide some benefit over and above that from the second supplier. The cost of the three-supplier system is (averaging over the four capacity levels) 21.8% less expensive than that of the 2-supplier system, and the cost difference is statistically significant for each of the four capacity levels. In contrast, the corresponding differences for the DU case are negligible.


Evidently, at least under these settings, redundant supply is a much more important tool under DU than under SU. We investigate this finding for a broader range of instances in the next section.

7.2.2 Extended Study

We simulated the system described in Section 7.2.1 with varying values of the holding cost h at stage 1. For each value of h , we tested systems with one, two, and three suppliers, and for each we tested values of the processing capacity equal to 25, 30, 35, and 40. For the SU tests, we tested the five pairs of values of (α, β) from (1). (Each supplier has the same failure and repair probabilities.) This resulted in 60 instances for DU and 180 for SU. For each problem instance, we tested base-stock levels at stage 1 equal to 20, 24, 28, 32 under DU and 20, 30, 40, 50, 60 under SU.

The results from the DU tests are summarized in Figure 16, which reports the average percent savings due to the second supplier and the average additional savings due to the third supplier for each capacity level. (Note that, whereas Figure 15 plots the cost, Figure 16 plots the savings, that is, the difference in cost.) The figure indicates that the second supplier provides a modest benefit (7.3%) for tight capacity and little

Figure 16: Average percent savings due to second and third suppliers under DU.



benefit otherwise. The third supplier provides essentially no benefit in either case, for reasons discussed in the previous section. (The values in Figure 16 with slight negative savings are due to simulation randomness.)

Figures 17(a) and (b) report the average percent savings under SU for the second and third suppliers, respectively, as a function of the supplier capacity. Each curve in the figures represents a given value of (α, β) , or the overall average. From Figure 17(a) it is evident that the second supplier provides a substantial benefit, with an overall average savings of 75.3% over the single-supplier cost. The savings generally decreases as the capacity increases (since a single supplier with large capacity can recover quickly after disruptions). The exception is the curve for $(\alpha, \beta) = (0.001, 0.1)$, which is quite flat. The reason for this is that, since the probability of two suppliers failing simultaneously is negligible, the backup supplier’s role is only to provide supply during disruptions, and never to help recover after them (see Section 7.2.1). Since the capacity is sufficient to meet the demand per period, the size of this benefit is independent of the capacity. The savings also tends to decrease as the failure probability increases, since the backup suppliers become less beneficial as they become less reliable.²

The differences in cost are statistically significant for all of the instances with $\alpha \geq 0.01$ but for only 10% of the instances with $(\alpha, \beta) = (0.001, 0.1)$. Although the percentage savings are greatest when $(\alpha, \beta) = (0.001, 0.1)$, those savings are the least statistically significant because failures are quite rare, and rarer disruptions result in larger SEMs in the simulation.

²Strictly speaking, a supplier’s reliability is dictated not by its failure probability alone but by the probability that it is in the DOWN state. For our two-state failure process, this probability is given by $\alpha/(\alpha + \beta)$, which equals 0.01, 0.03, 0.09, 0.13, and 0.18 for $(\alpha, \beta) = (0.001, 0.1), (0.01, 0.3), (0.05, 0.5), (0.1, 0.7), (0.2, 0.9)$, respectively. Therefore, as α increases in our instances, so does the DOWN probability.

Figure 17: Average percent savings due to (a) second and (b) third supplier under SU.





Figure 17(b) indicates that the third supplier, too, is beneficial, providing an average additional savings of 14.7%. As with the second supplier, the savings tends to decrease as the capacity or the failure probability increase. In addition, the statistical significance of the savings increases as α increases: For $(\alpha, \beta) = (0.001, 0.1)$, the cost difference is not statistically significant (and is in fact negative for a few instances due to simulation randomness), but it is significant for 90% of instances with $(\alpha, \beta) = (0.2, 0.9)$.

We summarize our findings from this section in the following remarks.

Remark 9 *Under DU, a second supplier provides moderate savings if each supplier's capacity is tight (less than 2 SDs above the mean demand) and little benefit otherwise. A third supplier provides essentially no benefit, regardless of capacity, assuming that the combined capacity of the first two suppliers is sufficient to meet the demand with high probability.*

Remark 10 *Under SU, multiple sourcing provides a substantial benefit since it provides both emergency capacity during and ramp-up capacity after disruptions. The savings tends to decrease as each supplier's capacity increases and as the suppliers' reliability decreases. A third supplier still provides some benefit, though less so than the second supplier.*

Figure 18: Four retailers with four suppliers, partially linked (Study 7).



7.3 Study 7: Supplier Flexibility

7.3.1 Base Model

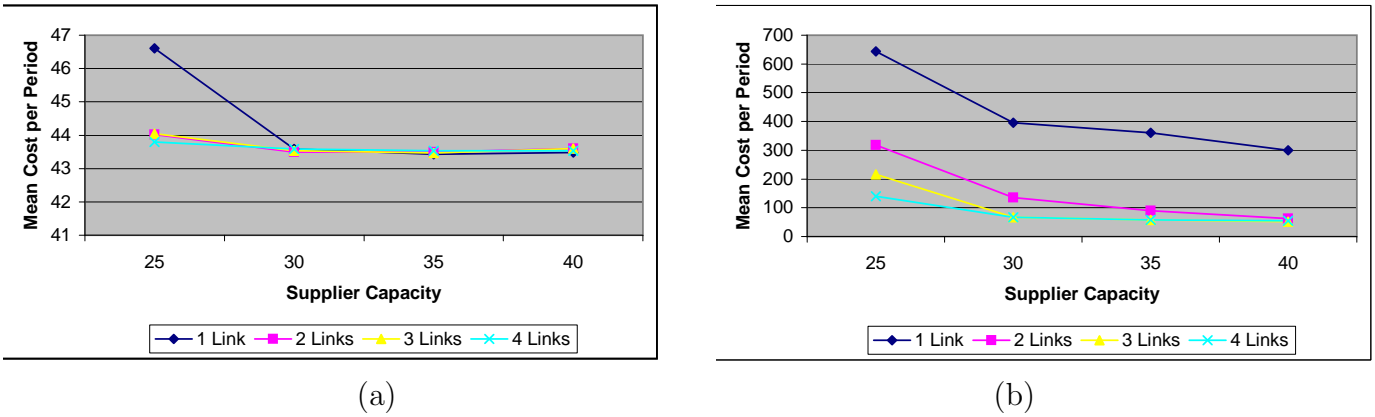
Consider a supply chain with four retailers and four suppliers, as in Figure 18. As in Study 6 (Section 7.2.1), the suppliers are identical in terms of cost and reliability, and we are interested in the question of how many suppliers should serve each retailer. Unlike in Study 6, however, in this study the retailers compete for the suppliers' limited capacity, reducing the benefit that each additional supplier provides. In Figure 18, each supplier serves two retailers, but we will also consider the cases in which each supplier serves one, three, or all four retailers. In essence, we are asking how dense the supplier–retailer graph should be. Our analysis is motivated in part by the work on process flexibility by Jordan and Graves (1995), in which the question is how many tasks each worker should be trained in—i.e., how dense the graph linking workers and tasks should be.

As in Study 6, supplier flexibility provides a retailer with additional capacity to meet excess demand under DU, while under SU the additional suppliers play two roles: to provide product when the primary supplier has failed, and to help recover more quickly after a disruption in which all suppliers that serve a given retailer have failed.

We simulated the system under DU and SU with one through four arcs (i.e., suppliers) per retailer. For the SU tests, we used $(\alpha, \beta) = (0.01, 0.3)$.³ The other parameter values are the same as those described in Section 7.2.1, as are the ranges of base-stock levels tested. Figure 19 plots the mean cost as a function of the supplier capacity for DU (Figure 7.2.1(a)) and SU (Figure 7.2.1(b)). As in Study 6, we find that the second link provides a modest savings for the tightest capacity value (25), but essentially none for looser capacity levels, and no benefit from the third or fourth links. In addition,

³We used $(\alpha, \beta) = (0.01, 0.3)$ rather than $(0.05, 0.5)$ as in other studies because the larger α values resulted in an unstable system for capacity level 25. See Section 7.3.2 for more details.

Figure 19: Mean optimal costs for supply chain with one through four suppliers per retailer, as a function of capacity under (a) DU and (b) SU.



the costs generally remain flat as the capacity increases, since a total capacity of 30 or more is sufficient to meet all of the demand in a period with 98% probability. Except under capacity 25, the differences in cost between the one- and two-link systems are not statistically significant, nor are those between the two- and three-link or three- and four-link systems under any capacity level.

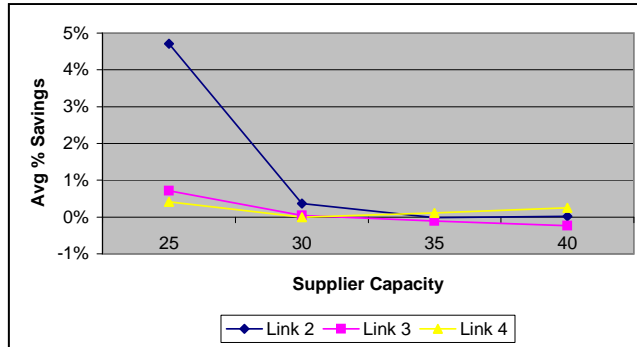
In contrast, the second link provides a substantial benefit, and the third link provides a modest benefit, at all capacity levels under SU. (The cost differences are statistically significant for the second link, and for two of the four capacity levels for the third link.) These savings result from the critical role that backup suppliers can play during and immediately after disruptions. Moreover, each of the four cost curves continue to decline slightly as the capacity increases. This is a difference from Study 6, in which the costs remained flat for capacity levels of 30 or greater, even under SU. The reason for this difference is that here, the retailers compete for the suppliers' capacity, and therefore additional capacity continues to provide additional savings.

These findings suggest that supply flexibility provides only minimal benefit under DU but provides substantial benefit under SU. This mirrors the findings from Study 6, in which we studied supplier redundancy rather than supplier flexibility. We repeat these experiments for a broader range of instances in the next section.

7.3.2 Extended Study

We simulated the system described in Section 7.3.1 with varying values of the holding cost at the retailers, keeping the stockout cost fixed at $p = 20$. For each value of

Figure 20: Average percent savings due to second, third, and fourth links under DU.



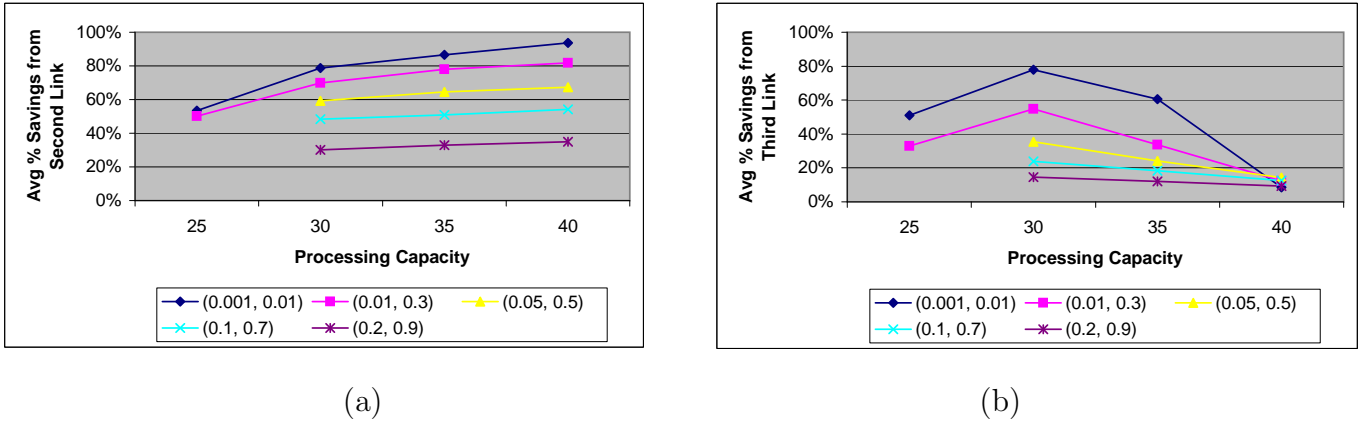
h , we simulated the system with one through four links per retailer, and for each we tested capacity values of 25, 30, 35, and 40. For the SU tests, we used the values of (α, β) from (1). This resulted in 80 instances for DU and 400 for SU. However, for capacity = 25 and $\alpha \geq 0.05$, the total expected system capacity (after accounting for disruptions) was less than 20, resulting in an unstable system. Therefore, we omitted these 60 instances (3 values of $\alpha \times 5 h \times 4$ graph densities) from our results.

Figure 20 summarizes the average percent savings from the second, third, and fourth links under DU, for each capacity level. The results confirm our findings from the base case in Section 7.3.1 that the second link provides only a modest benefit (4.7%, which is statistically significant) if the capacity equals 25, but no benefit otherwise; and that the third and fourth links provide essentially no benefit.

Figures 21(a) and (b) display the average percent savings due to the second and third links (resp.) under SU. Clearly, the second link provides a substantial benefit for all disruption profiles and capacity levels, with an overall average savings of 60.9%. Unlike in the corresponding graph in Study 6 (Figure 17(a)), in Figure 21(a) the relative savings *increases* as the capacity increases. Additional units of capacity provide increasing marginal savings (for the capacity values tested) when the retailers compete for supplier capacity but decreasing marginal savings when there is only a single retailer (as in Study 6). As in Study 6, the savings decrease as the failure probability (and the DOWN percentage) increases.

The savings from the second link is statistically significant for all instances in which $\alpha > 0.001$, and for 8 of the 20 instances (40%) for which $\alpha = 0.001$. As in Study 6, although the mean savings are greatest when $\alpha = 0.001$, so is the variability, making

Figure 21: Average percent savings due to (a) second and (b) third link under SU.



the differences less statistically significant.

Figure 21(b) suggests that, under SU, the third link continues to provide a benefit, with an average savings of 29.2% over and above that provided by the second link. The savings initially increases with capacity but begins to decrease when the capacity equals 35. The statistical significance of the savings decreases with α : the savings are significant for all instances in which $(\alpha, \beta) = (0.1, 0.7)$ or $(0.2, 0.9)$, for 13 out of 15 (87%) of instances with $(\alpha, \beta) = (0.05, 0.5)$, for 8 out of 20 (40%) of instances with $(\alpha, \beta) = (0.01, 0.3)$, and for 1 out of 20 (5%) of instances with $(\alpha, \beta) = (0.001, 0.1)$.

We omit detailed results for the fourth link. In general, the fourth link provides a modest benefit at the smallest capacity level (25), and essentially no benefit otherwise.

Comparing Figures 17 (Study 6) and 21 (Study 7), one can see that the average savings from each additional supplier is smaller in the supplier flexibility case than in the supplier redundancy case. For example, if the capacity is 25 and $(\alpha, \beta) = (0.001, 0.1)$, the second supplier in Study 6 provides a 97% savings but the second link in Study 7 provides only a 53% savings. This is because the single retailer gains all the benefit of the second supplier in Study 6, but multiple retailers share each supplier in Study 7, so the benefit is dampened.

We summarize our findings from this study in the following remarks.

Remark 11 *Under DU, supplier flexibility (in which multiple retailers share multiple suppliers) provides little or no benefit versus a system in which each retailer has a single dedicated supplier, assuming that the capacity of each supplier is sufficient to meet the demand most of the time. A second link (i.e., supplier) per retailer provides moderate*

savings if each supplier's capacity is tight (less than 2 SDs above the mean demand) and little benefit otherwise. A third or fourth link provides essentially no benefit, regardless of capacity.

Remark 12 *Under SU, supplier flexibility provides a substantial benefit since it provides both emergency capacity during and ramp-up capacity after disruptions. The marginal savings tends to increase as each supplier's capacity increases and to decrease as the suppliers' reliability decreases. A third link still provides some benefit, though less so than the second link, and a fourth link provides little benefit unless capacity is very tight.*

Remark 13 *Supplier flexibility provides less benefit than supplier redundancy (Study 6), since additional suppliers must be shared by multiple retailers.*

8. The Cost of Reliability

Most firms have designed their supply chains primarily with DU in mind, and many use strategies similar to the optimal strategies for DU discussed in the studies above. These firms may be reluctant to make large strategic changes to their supply chains to protect against SU—nor would such changes necessarily be advisable, since firms face a mix of SU and DU. However, we argue in this study that firms can greatly improve their reliability with respect to disruptions without large increases in cost.

Using the serial system from Study 4, we examine the tradeoff between the cost of a given solution under DU only and the service level of that solution under DU and SU. In other words, if a firm is accustomed to thinking about an objective function that considers DU only, how much of that objective does it need to sacrifice in order to achieve a given improvement in reliability when SU is considered? We find that the answer is usually “not much,” since the tradeoff curve between these two measures is generally “steep.”

8.1 Study 8: The Cost of Reliability

8.1.1 Base Model

Consider the serial system from Study 4 (Section 6.2). We have argued that it is preferable to hold inventory upstream under DU but downstream under SU. Even if

the firm is committed to holding inventory upstream, however, it may still benefit from holding a small amount of inventory downstream, since this inventory can be used to protect against both SU and DU.

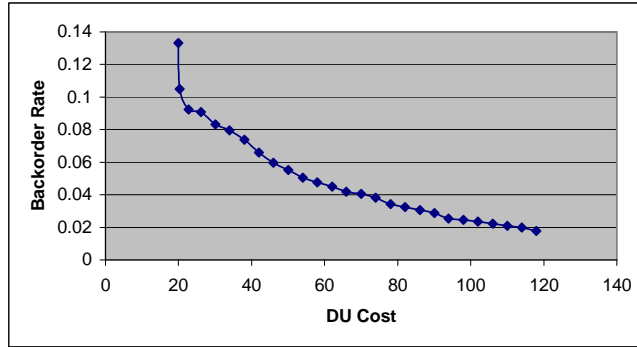
We simulated this system, keeping the base-stock level set to $S_3 = 20$ upstream and varying the base-stock level S_1 downstream, first considering DU only. In this case, the optimal base-stock level downstream (assuming $S_3 = 20$) is 8, with a mean cost of 19.9 per period if only DU is present. Now suppose that SU is present, as well, with $(\alpha, \beta) = (0.05, 0.5)$ at stage 2. Under the optimal DU solution ($S_1 = 8$), 13.3% of demands are backordered (i.e., the fill rate is 86.7%). By holding just a bit more inventory, however, we can achieve substantial reductions in the percentage of demands backordered. For example, if we increase S_1 from 8 to 10, the backorder rate decreases to 10.5% (a 21% reduction), while the cost increases only to 20.3 (a 2% increase in cost). If $S_1 = 12$, the backorder rate decreases 31% from its original level, with a cost increase of 14%.

The complete tradeoff curve is pictured in Figure 22. The horizontal axis displays the mean cost per period assuming DU only, while the vertical axis displays the percent of demands backordered assuming SU is present as well. Therefore, the graph displays the required change in the cost the firm is used to optimizing (the DU-only cost) in order to achieve a given change in the backorder rate. For example, the left-most point represents the optimal DU solution: $S_1 = 8$. This solution has a DU cost of 19.9 per period and a backorder rate (under DU and SU) of 13.3%. The next point represents the solution $S_1 = 10$, which has a DU cost of 20.3 and a backorder rate of 10.5%. In general, the steepness of the left-most part of the curve suggests that, if we start at the optimal DU solution, large improvements in reliability (measured by backorder rate) are possible with only small increases in DU cost. We demonstrate that this result tends to hold in general in the next section.

8.1.2 Extended Study

We generated tradeoff curves for a number of additional test instances, created from the base case in Section 8.1.1 by varying h_1 , p_1 , σ , and (α_2, β_2) . Note that, unlike in previous tests, (α, β) were varied independently of the other parameters, rather than testing five values of (α, β) for each set of parameters. This resulted in a total of 17 instances.

Figure 22: Tradeoff curve: backorder rate vs. DU cost (Study 8).



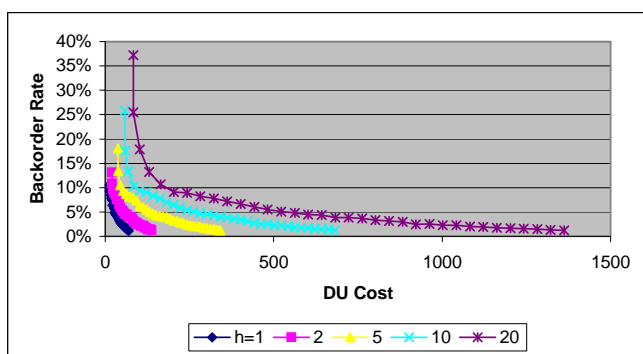
For each instance, we first simulated the system under DU only ($\alpha = \beta = 0$) with $S_3 = 20$ and $S_1 = 0, 2, 4, \dots, 70$. We then selected the value S_3^* that minimized the mean cost per period (typically, $0 < S_3^* < 10$) and simulated the system under both DU and SU for $S_3 = S_3^*, S_3^* + 2, \dots, 70$ in order to evaluate the change to the DU cost and the DU–SU backorder rate as S_3 increases from its optimal value under DU.

The results are displayed in Figure 23. Each graph contains several tradeoff curves, one for each value of the changing parameter. Figure 23(a) contains tradeoff curves for varying values of h_1 , Figure 23(b) contains curves for p_1 , and so on.

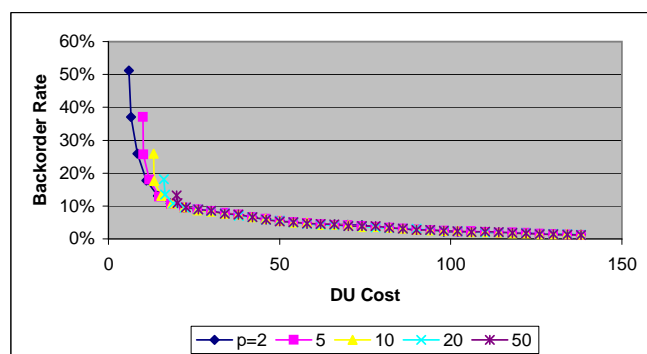
In general, the tradeoff curves are steep, indicating that large improvements in backorder rate are possible with only small increases in cost. On average, the second solution represents a 21.1% decrease in backorder rate versus the left-most solution, with only a 4.4% increase in expected DU cost per period. (These averages are taken across all tradeoff curves in all figures.) The third solution represents an average 33.3% decrease in backorder rate with a 20.4% increase in cost, while the fourth represents an average 40.9% decrease in backorder rate with a 42.3% increase in cost. This suggests that the benefit from increasing inventory at stage 1 is large at first but decreases quickly.

From Figure 23(a), it is evident that the first few points in the tradeoff curves represent larger improvements in backorder rate when h is large. This is because when h is large, the optimal DU solution results in a large backorder rate, so small increases in inventory result in large service level improvements. On the other hand, these changes also result in larger cost increases for larger values of h . The right-hand portions of the curves remain quite steep for smaller values of h and flatten out more quickly for

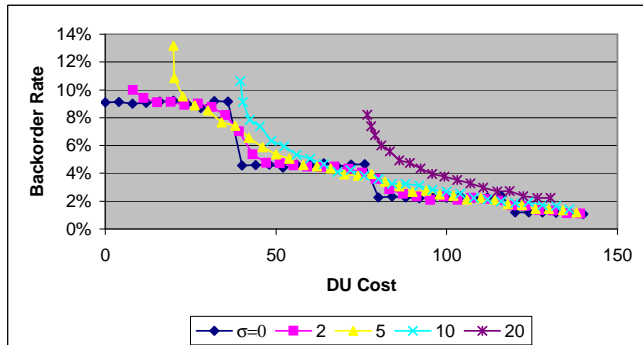
Figure 23: Tradeoff curve (backorder rate vs. DU cost) for varying values of (a) h_1 , (b) p_1 , (c) σ , and (d) (α_2, β_2) .



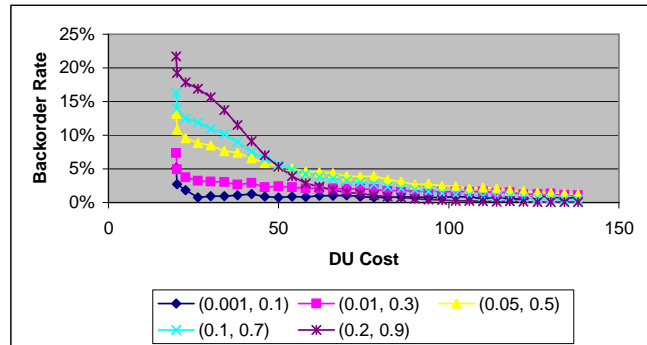
(a)



(b)



(c)



(d)

larger values.

Similarly, from part (b), the first few solutions represent larger improvements in backorder rate when p is small, because these solutions, too, have poor service levels in the optimal DU solution. On the other hand, the percent cost increase is larger when p is smaller, because when p is large, inventory increases provide additional protection against costly stockouts even under DU. As the inventory level increases, the curves for differing values of p coincide, because these solutions all generate very few stockouts, so p becomes irrelevant. The same effect does not happen as h increases: these solutions still involve holding some inventory, because $h < p$ in all instances.

The tradeoff curves in Figure 23(c) are quite flat for small values of σ . When $\sigma = 0$, extra inventory provides no benefit, only cost, under DU, so the DU cost increases are sharper and the SU-DU backorder improvements are more modest. Sharp improvements in backorder rate occur when S_1 is a multiple of $\mu = 20$: If $S_1 < 20$, then any disruption incurs some backorders, while if $S_1 \geq 20$, only disruptions lasting 2 periods or more cause backorders. Similar effects occur at larger multiples of 20. This effect is still visible, though less pronounced, when $\sigma = 2$. For $\sigma > 2$, the DU smoothes out the SU effects; increases in inventory have a smaller impact on cost and a large impact on service level.

As α increases in Figure 23(d), so does the backorder rate. As we stated earlier, for our values of (α, β) , the DOWN probability increases as α increases. Therefore, the curves for larger α values result in more disrupted periods and more stockouts. These curves remain steeper longer, suggesting that inventory is a more valuable tool for protecting against frequent, short stockouts than against rare, long ones. This finding echoes those from Study 4 (Section 6.2.2) suggesting that the optimal base-stock levels increase for large values of α , whereas for smaller values, the optimal strategy is simply to accept the disruption risk without increasing inventory.

We summarize our findings from this study in the following remark:

Remark 14 *In a serial system, a small increase in inventory levels from the optimal DU solution can provide substantial improvements in the service level when SU is considered with only small increases in DU cost. This effect is more pronounced when h is small (because inventory increases are less costly), when p is small (because the optimal DU solution generated a low service level), when σ is large (because extra inventory*

Table 5: Summary of results.

Study	Conclusion
1	The cost of not protecting against uncertainty is greater under SU than DU.
2	Batch ordering is preferred under SU, even when batch and one-for-one ordering are equivalent under DU.
3	Diversification of inventory sites is preferred under SU, while consolidation is preferred under DU.
4	It is preferable to hold inventory downstream under SU but upstream under SU.
5	Point-to-point networks are preferred under SU, while hub-and-spoke networks are preferred under DU.
6	Backup suppliers are significantly more valuable under SU than under DU.
7	Supplier flexibility is significantly more valuable under SU than under DU.
8	Supply chain resilience can be improved significantly without large increases in cost.

helps alleviate DU, too) and when α is large (since inventory is a more valuable tool for mitigating SU when disruptions are more frequent).

9. Conclusions

In this paper we investigate the differences between DU and SU using simulation. We consider SU in the form of supply disruptions, during which a portion of the supply chain is completely inoperative. To the best of our knowledge, this is the first paper to systematically examine the difference between these two types of uncertainty, and to highlight the different strategies that are appropriate. In addition, ours is one of the first papers to consider disruptions in multi-echelon models, which we believe is important because disruptions are not local, but rather tend to propagate downstream through a supply chain.

Our main findings are summarized in Table 5; please refer to the individual studies above for details of the assumptions that led to these conclusions and a more thorough explanation of why they hold.

Our simulation results show that the two types of uncertainty have different optimal strategies in terms of ordering frequency, inventory placement, and supply chain structure. In fact, the optimal strategy for dealing with supply uncertainty is, in many cases, the exact opposite from that for demand uncertainty. However, we are not arguing that firms should radically change their systems and their strategies to deal with these uncertainties. Rather, we are trying to point out these two competing tendencies.

In practice, both DU and SU exist, and the optimal strategy should be determined by considering the mix of DU and SU.

We demonstrate that the cost of unreliability is greater under SU than under DU, in the sense that, if the firm fails to plan for uncertainty, the cost under a given level of SU is greater than that under the equivalent level of DU. This suggests that planning for SU is crucial. The strategies discussed in this paper can reduce the cost under SU, but this cost may still be substantial. Further research should explore strategies for mitigating SU in multi-echelon supply chains.

We also analyze the trade-off between cost and reliability, and show that it is important for a company to find the right trade-off between cost and reliability. The cost difference between the reliability-minimization solution and the reliability-maximization solution can be quite large. We also show that significant improvements in reliability can often be achieved at relatively low cost.

We believe it is an interesting research problem to consider both SU and DU in the same model and determine the optimal strategy to deal with overall uncertainty. We also plan to derive theoretical results confirming some claims in this paper.

References

- Arreola-Risa, Antonio, Gregory A. DeCroix. 1998. Inventory management under random supply disruptions and partial backorders. *Naval Research Logistics* **45** 687–703.
- Axsäter, Sven. 2000. *Inventory Control*. Kluwer, Boston.
- Berk, Emre, Antonio Arreola-Risa. 1994. Note on “Future supply uncertainty in EOQ models”. *Naval Research Logistics* **41** 129–132.
- Berman, Oded, Dmitry Krass, Mozart B. C. Menezes. 2004. Reliability issues, strategic co-location and centralization in m-median problems. *Operations Research*, forthcoming .
- Berman, Oded, Dmitry Krass, Mozart B. C. Menezes. 2005. MiniSum with imperfect information: Trading off quantity for reliability of locations. Working paper, Rotman School of Management, University of Toronto.
- Çakanyildirim, Metin, James H. Bookbinder. 1999. Random lead times and expedited orders in (Q, r) inventory systems. *European Journal of Operational Research* **115** 300–313.
- Chopra, Sunil, Giles Reinhardt, Usha Mohan. 2005. The importance of decoupling recurrent and disruption risks in a supply chain. Working paper, Northwestern University.
- Church, Richard L., Maria P. Scaparra. 2005. Protecting critical assets: The r -

- interdiction median problem with fortification. forthcoming in *Geographical Analysis* .
- Clark, Andrew J., Herbert Scarf. 1960. Optimal policies for a multi-echelon inventory problem. *Management Science* **6**(4) 475–490.
- Dada, Maqbool, Nicholas C. Petruzzi, Leroy B. Schwarz. 2003. A newsvendor model with unreliable suppliers. Working paper, Krannert Graduate School of Management, Purdue University.
- Eppen, Gary D. 1979. Effects of centralization on expected costs in a multi-location newsboy problem. *Management Science* **25**(5) 498–501.
- Fishman, George S. 2001. *Discrete-Event Simulation: Modeling, Programming, and Analysis*. Springer-Verlag, New York.
- Gupta, Diwakar. 1996. The (Q, r) inventory system with an unreliable supplier. *INFOR* **34**(2) 59–76.
- Gürler, Ülkü, Mahmut Parlar. 1997. An inventory problem with two randomly available suppliers. *Operations Research* **45**(6) 904–918.
- Hariharan, Rema, Paul Zipkin. 1995. Customer-order information, leadtimes, and inventories. *Management Science* **41**(10) 1599–1607.
- Hendricks, Kevin B., Vinod R. Singhal. 2003. The effect of supply chain glitches on shareholder wealth. *Journal of Operations Management* **21**(5) 501–522.
- Hendricks, Kevin B., Vinod R. Singhal. 2005a. Association between supply chain glitches and operating performance. *Management Science* **51**(5) 695–711.
- Hendricks, Kevin B., Vinod R. Singhal. 2005b. An empirical analysis of the effect of supply chain disruptions on long-run stock price performance and equity risk of the firm. *Production and Operations Management* **14**(1) 35–52.
- Hillman, Mark, Fenella Sirkisoon. 2006. Pandemic readiness study. Report, AMR Research, www.amrresearch.com/research/reports/images/2006/0605AMR-A-19413.pdf.
- Hopp, Wallace J., Zigeng Yin. 2006. Protecting supply chain networks against catastrophic failures. working paper, Northwestern University.
- Jordan, William C., Stephen C. Graves. 1995. Principles on the benefits of manufacturing process flexibility. *Management Science* **41**(4) 577–594.
- Kim, Hyoungtae, Jye-Chyi Lu, Paul H. Kvam. 2005. Ordering quantity decisions considering uncertainty in supply-chain logistics operations. Working paper, Georgia Institute of Technology, Atlanta, GA.
- Lee, Hau L. 2004. Supply chain security: Are you ready? White paper, Stanford Global Supply Chain Management Forum.
- Lewis, Brian M., Alan L. Erera, Chelsea C. White. 2005. An inventory control model with possible border disruptions. Working paper, Georgia Institute of Technology, Atlanta, GA.
- Li, Zhaolin, Susan H. Xu, Jack Hayya. 2004. A periodic-review inventory system with supply interruptions. *Probability in the Engineering and Informational Sciences* **18** 33–53.
- Lynn, Barry C. 2005. *End of the Line: The Rise and Coming Fall of the Global Corporation*. Doubleday.

- Moinzadeh, Kamran, Prabhu Aggarwal. 1997. Analysis of a production/inventory system subject to random disruptions. *Management Science* **43**(11) 1577–1588.
- Nahmias, Stephen. 2005. *Production and Operations Analysis*. 5th ed. McGraw-Hill/Irwin.
- Nelson, Barry L. 1995. *Stochastic Modeling: Analysis & Simulation*. McGraw-Hill, New York.
- Parlar, M., D. Berkin. 1991. Future supply uncertainty in EOQ models. *Naval Research Logistics* **38** 107–121.
- Parlar, Mahmut. 1997. Continuous-review inventory problem with random supply interruptions. *European Journal of Operational Research* **99** 366–385.
- Parlar, Mahmut, David Perry. 1995. Analysis of a (Q, r, T) inventory policy with deterministic and random yields when future supply is uncertain. *European Journal of Operational Research* **84** 431–443.
- Parlar, Mahmut, David Perry. 1996. Inventory models of future supply uncertainty with single and multiple suppliers. *Naval Research Logistics* **43** 191–210.
- Poirier, Charles C., Francis J. Quinn. 2006. Survey says: Solid gains. *Supply Chain Management Review* **10**(1) 34–41.
- Porteus, Evan L. 2002. *Foundations of Stochastic Inventory Theory*. Stanford University Press, Stanford, CA.
- Qi, Lian, Zuo-Jun Max Shen. 2005. A supply chain design model with unreliable supply. Working paper, University of Florida.
- Qi, Lian, Zuo-Jun Max Shen, Lawrence V. Snyder. 2006. The effect of supply uncertainty on supply chain design decisions. Working paper, University of California, Berkeley, CA.
- Rice, James B., Federico Caniato, Jonathan Fleck, Deena Disraelly, Don Lowtan, Reshma Lensing, Chris Pickett. 2003. Supply chain response to terrorism project: Creating resilient and secure supply chains. Interim report, MIT Center for Transportation and Logistics.
- Scaparra, Maria P., Richard L. Church. 2005. An optimal approach for the interdiction median problem with fortification. Working paper 78, Kent Business School, Canterbury, England.
- Schmitt, Amanda, Lawrence V. Snyder. 2006. Infinite-horizon models for inventory control under yield uncertainty and disruptions. Working paper, Lehigh University, Bethlehem, PA.
- Sheffi, Yossi. 2001. Supply chain management under the threat of international terrorism. *International Journal of Logistics Management* **12**(2) 1–11.
- Sheffi, Yossi. 2005. *The Resilient Enterprise: Overcoming Vulnerability for Competitive Advantage*. MIT Press, Cambridge, MA.
- Simchi-Levi, David, Xin Chen, Julien Bramel. 2005. *The Logic of Logistics*. 2nd ed. Springer.
- Simchi-Levi, David, Lawrence V. Snyder, Mike Watson. 2002. Strategies for uncertain times. *Supply Chain Management Review* **6**(1) 11–12.
- Snyder, Lawrence V. 2006a. BaseStockSim software v2.4. www.lehigh.edu/lvs2/software.html.

- Snyder, Lawrence V. 2006b. A tight approximation for a continuous-review inventory model with supplier disruptions. Working paper, Lehigh University.
- Snyder, Lawrence V., Mark S. Daskin. 2005. Reliability models for facility location: The expected failure cost case. *Transportation Science* **39**(3) 400–416.
- Snyder, Lawrence V., M. Paola Scaparra, Mark L. Daskin, Richard C. Church. 2006. Planning for disruptions in supply chain networks. H. J. Greenberg, ed., *TutORials in Operations Research*. INFORMS.
- Song, Jing-Sheng, Paul H. Zipkin. 1996. Inventory control with information about supply conditions. *Management Science* **42**(10) 1409–1419.
- Stauffer, David. 2003. Risk: The weak link in your supply chain. *Harvard Management Update* **8**(3) 3–5.
- Tomlin, Brian, Yimin Wang. 2004. On the value of mix flexibility and dual sourcing in unreliable newsvendor networks. Working paper, Kenan-Flagler Business School, UNC-Chapel Hill.
- Tomlin, Brian T. 2004. Supplier diversification: The impact of supplier-reliability forecasting on supplier selection and sourcing quantities. Working paper, Kenan-Flagler Business School, UNC-Chapel Hill.
- Tomlin, Brian T. 2006. On the value of mitigation and contingency strategies for managing supply chain disruption risks. *Management Science* **52**(5) 639–657.
- Tomlin, Brian T., Lawrence V. Snyder. 2006. On the value of a threat advisory system for managing supply chain disruptions. Working paper, Lehigh University.
- Wein, Lawrence M., Yifan Liu. 2005. Analyzing a bioterror attack on the food supply: The case of botulinum toxin in milk. *Proceedings of the National Academy of Sciences* **102**(28) 9984–9989.
- Yano, Candace A., Hau A. Lee. 1995. Lot sizing with random yield: A review. *Operations Research* **43**(2) 311–334.
- Zipkin, Paul H. 2000. *Foundations of Inventory Management*. Irwin/McGraw-Hill.