MANAGEMENT SCIENCE
Vol. 52, No. 5, May 2006, pp. 639–657
ISSN 0025-1909 | EISSN 1526-5501 | 06 | 5205 | 0639

informs
DOI 10.1287/mnsc.1060.0515

On the Value of Mitigation and Contingency Strategies for Managing Supply Chain Disruption Risks

Brian Tomlin
Kenan-Flagler Business School, University of North Carolina at Chapel Hill,
Chapel Hill, North Carolina 27599-3490, brian.tomlin@unc.edu

We study a single-product setting in which a firm can source from two suppliers, one that is unreliable and another that is reliable but more expensive. Suppliers are capacity constrained, but the reliable supplier may possess volume flexibility. We prove that in the special case in which the reliable supplier has no flexibility and the unreliable supplier has infinite capacity, a risk-neutral firm will pursue a single disruption-management strategy: mitigation by carrying inventory, mitigation by single-sourcing from the reliable supplier, or passive acceptance. We find that a supplier's percentage uptime and the nature of the disruptions (frequent but short versus rare but long) are key determinants of the optimal strategy. For a given percentage uptime, sourcing mitigation is increasingly favored over inventory mitigation as disruptions become less frequent but longer. Further, we show that a mixed mitigation strategy (partial sourcing from the reliable supplier and carrying inventory) can be optimal if the unreliable supplier has finite capacity or if the firm is risk averse.

Contingent rerouting is a possible tactic if the reliable supplier can ramp up its processing capacity, that is, if it has volume flexibility. We find that contingent rerouting is often a component of the optimal disruption-management strategy, and that it can significantly reduce the firm's costs. For a given percentage uptime, mitigation rather than contingent rerouting tends to be optimal if disruptions are rare.

Key words: supply uncertainty; dual-sourcing; volume flexibility
History: Accepted by Candace A. Yano, operations and supply chain management; received August 1, 2003. This paper was with the author 1 year and 1 month for 4 revisions.

1. Introduction

In March 2000, lightning caused a fire that shut down the Philips Semiconductor plant in Albuquerque, New Mexico, for six weeks, leading to a shortage of components for both Ericsson and Nokia. According to *The Wall Street Journal*, "company officials say they [Ericsson] lost at least $400 million in potential revenue" and "when the company revealed the damage from the fire for the first time publicly last July, its shares tumbled 14% in just hours" (Latour 2001). In February 1997, a fire in a Toyota brake-supplier plant led directly to a two-week shut down of 18 Toyota plants in Japan, with a resulting cost of $195 million (Treece 1997). Fires, of course, are not the only cause of disruptions. Hurricane Mitch caused catastrophic damage to banana production in many parts of Central America in 1998. It took many growers over a year to recover, leading to a prolonged loss of supply for Dole and Chiquita (Griffy-Brown 2003). An earthquake in Taiwan severely disrupted supply of essential components to the personal-computer industry leading up to the 1999 holiday season (Burrows 1999).

It is informative to compare the supply chain strategies of companies and their resulting ability to cope with some of the above-mentioned disruptions. Nokia lost all of its supply from the Philips plant, but it was able to temporarily increase production at alternative suppliers during the disruption, and so suffered little financial impact. In contrast, Ericsson had been "weeding out backup suppliers for many parts" and, according to Jan Ahrenbring, Ericsson's marketing director for consumer goods, the company consequently "did not have a Plan B" (Latour 2001). Ericsson's single-source strategy caused it to lose over $400 million in potential revenue. A similar contrast appears in the Hurricane Mitch situation. Chiquita, although it lost significant supply, was able to temporarily increase production at some of its other (unaffected) suppliers in the region. Dole had no alternative suppliers in the region and lost 70% of its regional supply. Dole suffered a 4% decline in revenues (and lost over $100 million) while Chiquita increased revenues by 4% in the fourth quarter of 1998 (Griffy-Brown 2003). Firms do not need to rely exclusively on supply-side tactics during a disruption. Dell was able to leverage its demand-management capabilities to shift demand to alternative products that were less supply constrained during the 1999 Taiwanese earthquake disruption, whereas Apple, lacking the same demand-management capability, was much less able to cope with the disruption (Griffy-Brown 2003).

Table 1 Tactics for Managing Disruption Risks

Category	Tactic	Examples
Financial mitigation	Business interruption insurance	In the fourth quarter of 2003, Palm Inc. received a $6.4 million insurance settlement arising from an earlier fire at a supplier's factory.
Operational mitigation	Inventory	Playmates Toys mitigated the impact of the 2002 west-coast dock disruption by investing in inventory earlier in the year.
		The U.S. Strategic Petroleum Reserve protects the U.S. against interruptions in crude-oil supplies.
	Sourcing	Nokia's multiple-supplier sourcing strategy mitigated the impact of the Philips Semiconductor disruption in 2000.
		Chiquita's multiple-location sourcing strategy mitigated the impact of the 1998 Hurricane Mitch disruption.
Operational contingency	Rerouting	Nokia responded to the Philips Semiconductor disruption by temporarily increasing production at alternative suppliers.
		Chiquita responded to the Hurricane Mitch disruption by temporarily increasing production at alternative locations.
		New Balance responded to the west-coast dock disruption by rerouting ships to the east coast and by air freighting supplies.
		Chrysler responded to the air-traffic disruption in the immediate aftermath of September 11th by temporarily switching to ground transportation to move components from a U.S. supplier to the Dodge Ram assembly plant in Mexico.
	Demand management	Dell responded to the disruption in memory supply caused by the 1999 Taiwanese earthquake by shifting customer demand to lower-memory personal computers.

Firms can use a number of tactics to manage the risk of disruptions (see Table 1). Mitigation tactics are those in which the firm takes some action in advance of a disruption (and so incurs the cost of the action regardless of whether a disruption occurs). Contingency tactics are those in which a firm takes an action only in the event a disruption occurs. We note that a contingent-rerouting tactic is viable only if suppliers have volume flexibility, that is, the ability to temporarily increase their processing capacity. A firm is not limited to choosing a single tactic, and in many circumstances a combination of tactics might be the appropriate strategy for managing disruption risk. Mitigation and contingency actions are not free, and therefore passive acceptance of the disruption risk may be appropriate in certain circumstances. Passive acceptance is often the default strategy even when it is not appropriate. In a recent survey (Poirier and Quinn 2004), only 33% of firms responded that they paid "sufficient attention to supply chain vulnerability and risk mitigation actions." The focus of this paper is on operational tactics for managing disruption risk. In particular, we focus on the supply-side tactics of inventory, sourcing, and rerouting.

We study a single-product setting in which a firm can source from two suppliers, one that is unreliable and another that is reliable but more expensive. Suppliers are capacity constrained, but the reliable supplier may possess volume flexibility. In the special case where the reliable supplier has no flexibility, we prove that a risk-neutral firm will pursue a pure (i.e., not mixed) disruption-management strategy (mitigation by carrying inventory, mitigation by single-sourcing from the reliable supplier, or passive acceptance) if the unreliable supplier has infinite capacity. We find that supplier reliability and the nature of the disruptions (e.g., frequent but short versus rare but long) are key determinants of the optimal strategy. For a given supplier reliability, as measured by percentage uptime, sourcing mitigation tends to be favored over inventory mitigation as disruptions become less frequent but longer. We prove that a mixed mitigation strategy (partial sourcing from the reliable supplier and carrying inventory) can be optimal if the firm is risk averse or if the unreliable supplier has finite capacity. Finite capacity amplifies the effect of disruptions because the unreliable supplier cannot immediately recover from a disruption. We find that this issue of delayed recovery strongly influences the firm's optimal disruption-management strategy.

We also characterize the firm's optimal disruption-management strategies when the reliable supplier possesses volume flexibility. We introduce an operational definition of volume flexibility that is characterized by the amount of extra capacity that becomes available and the speed with which it becomes available. Volume flexibility allows contingent rerouting to be a tactic. We find that contingent rerouting is often a component of the optimal disruption-management strategy and that it can significantly reduce the firm's costs. One might expect that rare disruptions would favor a contingency tactic, as contingent costs are incurred only in the event of a disruption. Interestingly, for a given supplier reliability, we find that sourcing mitigation often becomes optimal as disruptions become less frequent.

The remainder of this paper is organized as follows. Section 2 surveys the existing literature. A general model is presented in §3. A restricted version of the model is analyzed in §4. Section 5 returns to the general model and conclusions are presented in §6. The proofs for this paper can be found in the online supplement (available on the *Management Science* website at http://mansci.pubs.informs.org/ecompanion.html) along with certain technical details that, while important building blocks, are of secondary importance to the main results. An unabridged version of this paper (available from the author upon request) contains an extended treatment of the risk-averse allocation decision considered in §5.1. We use the following mathematical notation in the paper: $[x]^+ = \max\{x, 0\}$; $[x]^- = \max\{-x, 0\}$; $\lfloor x \rfloor$ = largest integer less than or equal to x; $\lceil x \rceil$ = smallest integer greater than or equal to x.

2. Literature Review

While the mitigation and contingency framework seems like a natural one for supply-uncertainty problems, we are not aware of its usage in the existing supply-uncertainty literature. Nevertheless, we use this framework in positioning this paper in relation to the existing literature.

In supply-disruption models, a supplier (or resource) is either up or down. When the supplier is up it delivers an order in full and on time, but no order can be supplied when it is down. The supply process is characterized by the interfailure-time distribution and the repair-time distribution. The status of any supplier is always known to the firm in the disruption literature of which we are aware. The majority of supply-disruption papers focus on a single-supplier problem (Meyer et al. 1979, Bielecki and Kumar 1988, Parlar and Berkin 1991, Parlar and Perry 1995, Gupta 1996, Song and Zipkin 1996, Moinzadeh and Aggarwal 1997, Parlar 1997, Arreola-Risa and De Croix 1998). With no alternative source available, inventory mitigation is the only disruption-management strategy under consideration in these papers.

Parlar and Perry (1996) and Gürler and Parlar (1997) are, to the best of our knowledge, the only supply-disruption papers that consider more than one supplier. Both papers consider a firm that faces constant demand and sources from two identical-cost, infinite-capacity suppliers. The firm faces a fixed cost of ordering (although the fixed cost is only incurred once even if the order is split between suppliers). Interfailure and repair times are exponentially distributed for both suppliers in Parlar and Perry (1996). The authors propose a suboptimal ordering policy that is solved numerically. Gürler and Parlar (1997) extend the work of Parlar and Perry by considering the case of Erlang-k interfailure times and general repair times. They develop a cost expression that needs to be numerically evaluated and demonstrate how to numerically optimize the inventory for the case of Erlang-2 interfailure distributions and exponential repair times.

It is informative to consider the implications of the infinite-capacity and the identical-cost assumptions made by Parlar and Perry (1996) and Gürler and Parlar (1997) in relation to our mitigation-contingency framework. The identical-cost assumption removes any downside to sourcing mitigation: the firm is completely indifferent between suppliers if both suppliers are up when an order is placed. The combination of the identical-cost and infinite-capacity assumptions removes any downside to contingent rerouting: if one supplier is down, then the firm has no cost incentive to postpone ordering until that supplier is back up, nor is there any limitation on the order quantity it places with the other supplier. With no downside to sourcing mitigation or contingent rerouting, these two papers cannot (and do not purport to) offer any insights into the trade-offs between different disruption-management strategies; their contribution lies in the proposed solutions to the inventory-optimization problem. By considering finite-capacity suppliers that differ in cost, our paper goes beyond the existing literature by explicitly modeling the trade-offs and limitations inherent in mitigation and contingency strategies. This then enables us to provide insights into the structure of an optimal disruption-management strategy as well as the factors that make one strategy preferable over another. We note that our paper also allows for uncertain demand.

Yield-uncertainty models differ from supply-disruption models in that there is uncertainty at the time of order placement as to the fraction of the order that will be delivered. Much of the literature on yield uncertainty is focused on single-supplier models. Our attention in this survey is restricted to multiple-supplier models. We refer the interested reader to Yano and Lee (1995) for a comprehensive review of the yield-uncertainty literature. There is a limited literature on multiple-sourcing in the context of yield uncertainty. Gerchak and Parlar (1990), Agrawal and Nahmias (1997), Gurnani et al. (2000), Dada et al. (2003), Tomlin and Wang (2005), Anupindi and Akella (1993), Parlar and Wang (1993), and Swaminathan and Shanthikumar (1999) all investigate supplier diversification in the presence of yield uncertainty, with the latter three being the only ones to consider multiperiod problems. The focus of all these papers is on inventory and sourcing mitigation. In the single-period problems, contingent rerouting does not arise because no sourcing actions are allowed after uncertainty has been resolved; in the multiperiod problems, the state of a supplier is the same at the start of each period and so there is no reason to reroute supply.

The literature on random capacity has a random upper bound on production in each period and has typically focused on single-supplier (or machine) problems, e.g., Ciarallo et al. (1994), Khang and Fujiwara (2000), and Bollapragada et al. (2004a, b). One exception is the paper of Kouvelis and Milner (2002), which allows for multiple suppliers in the context of outsourcing.

All of the multiple-sourcing papers cited above assume identical lead-time suppliers. Even in the absence of supply uncertainty, the use of multiple suppliers can be beneficial if the suppliers differ in lead times and demand is uncertain (see Fukuda 1964, Moinzadeh and Nahmias 1988, Scheller-Wolf and Tayur 1998, Sethi et al. 2003, and Feng et al. 2004).

We now turn our attention to the literature on flexibility. We refer the reader to Sethi and Sethi (1990), Gerwin (1993), and Suarez et al. (1995) for reviews of flexibility. Mix flexibility, whereby a resource can produce multiple products, has been widely studied in the literature, including in Fine and Freund (1990), Jordan and Graves (1995), Van Mieghem (1998), Kouvelis and Vairaktarakis (1998), Graves and Tomlin (2003), and Tomlin and Wang (2005). Volume flexibility, whereby a resource can temporarily alter its capacity, has received much less attention.

So-called quantity-flexible or options contracts in the newsvendor contracting literature (e.g., Barnes-Schuster et al. 2002) are somewhat related to volume flexibility. A firm commits to the purchase of a certain number of units and sets an option reservation quantity. After some demand uncertainty is resolved, the firm has the flexibility of exercising some fraction of the reserved options. The firm pays a higher marginal price in total for reserved units than it does for committed units. One can think of committed units as being mitigation inventory and the exercising of reserved options as being a contingent action. A key aspect of these models is that the firm itself chooses the option reservation quantity, and so the limitation placed on its contingent action is a direct result of its own earlier decision. Note that there is no supply uncertainty in these models. The definition of volume flexibility as the ability to temporarily adjust capacity fits more naturally in a multiperiod setting. Tsay and Lovejoy (1999) is the only multiperiod, quantity-flexible contract paper of which we are aware. Again, there is no supply uncertainty in their model. Tsay and Lovejoy consider a model in which a firm can revise previous orders placed with a supplier. Order revisions (upwards or downwards) are bounded, and looser bounds are associated with higher supplier flexibility. Presumably a supplier with more volume flexibility can offer a higher degree of revision flexibility in the contract.

Our paper makes a key contribution to the literature on flexibility by introducing an operational definition of volume flexibility that directly captures a supplier's ability to temporarily adjust capacity. Our model captures two critical dimensions of volume flexibility—the magnitude of the capacity increase, and the time it takes for the extra capacity to become available.

Increased attention is being paid to the consideration of risk in operational decisions, especially in the context of a single-product newsvendor, e.g., Eeckhoudt et al. (1995), Agrawal and Seshadri (2000), Schweitzer and Cachon (2000), and Caldentey and Haugh (2004). Sourcing strategies are likely to be strongly influenced by the firm's attitudes toward risk, and so we consider both risk-neutral and risk-averse decision making in this paper. To the best of our knowledge, the only other supply-uncertainty paper that considers risk aversion is Tomlin and Wang (2005). However, that paper investigates a fundamentally different setting: a single-period, yield-uncertainty problem in which the firm faces trade-offs between mix flexibility and dual-sourcing. Contingency actions are not considered in that paper.

3. The Model

In our model, the firm operates an infinite-horizon, periodic-review inventory system with complete backlogging of unmet demand. Demand in period t, D_t, is drawn from a stationary distribution with strictly positive support. On-hand inventory at the end of a period costs h per unit, and back orders at the end of a period cost p per unit.

The firm has two suppliers (U and R) available to it. Supplier U is unreliable in the sense that it is either up or down in a period. We assume that supplier U's failure and repair transitions are such that supplier U can be modeled as a discrete-time Markov process. Further assumptions and definitions regarding the Markov process are provided at the end of this section. Supplier U is assumed to have a constant capacity of ν_u per period. Production is instantaneous, but there is a transit lead time of $L \geq 0$ periods before production arrives at the firm. Supplier R is completely reliable and also has instantaneous production and a transit lead time of L periods. The firm chooses an allocation $0 \leq w \leq 1$, such that it orders wD_t in every period from supplier R. For its given allocation w, supplier R provides sufficient capacity to produce wD_t each period. It cannot, however, instantaneously ramp up capacity during a disruption to supplier U. As with Nokia and Chiquita's suppliers, we assume that supplier R can potentially provide volume flexibility during a disruption. The function $\delta(\tau)$ denotes the volume flexibility profile; if given τ periods' notice, supplier R can increase its capacity by $\delta(\tau) \geq 0$. We assume that $\delta(\tau)$ is nondecreasing in τ, and we assume that volume flexibility is only made available during a disruption. In particular we will assume the following structure for the volume flexibility function: $\delta(\tau) = 0$ for $\tau < \theta_r$ and $\delta(\tau) = \delta > 0$ for $\tau \geq \theta_r$. We refer to δ as the flexibility magnitude and $\theta_r \geq 0$ as the flexibility response time. Our flexibility profile is then parameterized by (θ_r, δ). Magnitude and response time are two key parameters of volume flexibility: "To replace more than two million power amplifiers, they [Nokia] asked one Japanese and one U.S. supplier of the same chip to make millions more each. Largely because Nokia is such an important customer, both took the additional order with only five days of lead time" (Latour 2001).

Units ordered from supplier U cost c_u per unit. Supplier R charges c_r per unit for its normal allocation and charges $c_f \geq c_r$ per unit provided from its volume-flexible capacity. We assume that $c_f \geq c_r$ to reflect that volume flexibility may be associated with a higher marginal cost. We also assume that $c_r \geq c_u$ so as to ignore the trivial case where single-sourcing from R is clearly optimal. While the motivating examples in the introduction discussed disruptions at suppliers external to the firm, our model makes no assumption as to the ownership of either supplier—either or both suppliers could be internal or external to the firm.

All events during period t occur in the following sequence:

- The state of supplier U is observed at the beginning of the period.
- Demand is observed.
- Ordering decisions are made.
- Units produced by a supplier in period $t - L$ arrive.
- Demand is filled (if possible) and holding/back-order costs are incurred.
- Supplier U's state transition occurs at the end of the period.

A number of papers (e.g., Kalymon 1971, Song and Zipkin 1993, Parlar et al. 1995, and Song and Zipkin 1996) study Markovian inventory systems in the context of single-sourcing. The paper by Song and Zipkin (1996) is particularly relevant, and we leverage a number of their results in certain proofs. The sequence of events above mirrors that proposed by Song and Zipkin, with the exception that we use the convention (as in Graves 1988, 1999) that demand is observed before an order is placed, and so in effect our lead time is shorter by one time unit.

We define the following variables:

q_u^t: the order placed on supplier U in period t. We note that $q_u^t = 0$ if supplier U is down, as there is no point in placing an order.

q_r^t: the order placed on supplier R's regular capacity in period t. By assumption, this is $q_r^t = wD_t$ where $0 \leq w \leq 1$ is the allocation chosen by the firm.

q_f^t: the order placed on supplier R's flexible capacity in period t. We note that $q_f^t = 0$ if supplier U is up, because flexible capacity is made available only during a disruption.

x_t: on-hand inventory level (positive or negative) at the end of period t.

z_t: inventory position (on hand, on order, and in transit) at the end of period t.

The ordering and inventory/back-order costs in period t are then given by

$$c(q_u^t, q_r^t, q_f^t) = c_u q_u^t + c_r q_r^t + c_f q_f^t \quad \text{and} \qquad \hat{C}(x_t) = p[-x_t]^+ + h[x_t]^+, \tag{1}$$

respectively. We note that an extension to the case in which the firm pays inventory-holding costs for units in transit is easily accommodated. We assume that there is no cost associated with a volume-flexibility request, and so the firm will request such a capacity increase during a disruption. As evidenced by the Ericsson case, however, a firm may not respond immediately to a disruption at a supplier. We therefore assume that the firm does not respond with a capacity-increase request until the start of period θ_f of a disruption, where $\theta_f \geq 1$. For example, if $\theta_f = 1$, then the firm responds instantaneously to a disruption and places a capacity-increase request at the start of the first period of a disruption, whereas if $\theta_f = 2$, then the firm doesn't place a capacity-increase request until the start of the second period. We refer to θ_f as the firm's response time. Given our above assumption about the supplier's volume-flexibility profile, the effective flexible capacity available to the firm in period i of a disruption is then 0 if $i < \theta_f + \theta_r$ and δ if $i \geq \theta_f + \theta_r$. We define the supply chain response time as $\theta_{SC} = \theta_f + \theta_r$.

While ordering decisions are tactical in nature, the allocation decision is strategic. An allocation of $w = 0$ or $w = 1$ indicates the firm single-sources from supplier U or supplier R, respectively. An allocation of $0 < w < 1$ indicates the firm dual-sources. We note that even if $w = 0$, the firm may avail itself of supplier R's volume flexibility if it so chooses. An extension to the case in which supplier R only makes flexibility available if $w \geq w_{\min} > 0$ is easily handled. Ordering decisions and allocation decisions are likely made at different levels of the organization, with the allocation decisions made at a higher level. For a given allocation w, we assume that ordering decisions are made to minimize the long-run average cost. The firm, however, may exhibit risk aversion in making the allocation decision. We consider both a mean-variance and a conditional value-at-risk (CVaR) approach in considering risk. Further details regarding these approaches are given in §5.1.

We now specify our assumptions about supplier U's discrete-time Markov process. Let $i = 0, 1, \ldots, \infty$ denote the number of periods (including the current period) for which supplier U has been down. In other words, $i = 0$ if supplier U is up and $i = 1, \ldots, \infty$ if supplier U is down and was down in each of the previous $i - 1$ periods. The probability of a failure $1 - \lambda(0)$ is assumed to be independent of the number of periods for which U has been up. The probability $\lambda(i)$ of a disruption ending after i periods for which U has been down is assumed to depend only on i. With these assumptions, supplier U's failure and repair process can be modeled as a Markov chain with state space $i = 0, 1, \ldots, \infty$. Let i_+ denote the state after a transition from state $i \geq 0$; then either $i_+ = 0$, with probability $\lambda(i)$, or $i_+ = i + 1$, with probability $1 - \lambda(i)$. We note that for $i = 1, \ldots, \infty$, $\lambda(i)$ is the probability that a disruption ends after i periods conditional on it having lasted $i - 1$ periods, i.e., $\lambda(i)$ is the hazard rate for the repair process. We assume that the $\lambda(i)$ are nondecreasing in i for $i > 0$, i.e., the longer a disruption has lasted the more likely it is to end in the current period, a reasonable model for many disruptions. The residual life $r(i)$ at the start of the ith period ($i > 0$) of a disruption is the remaining number of periods until the disruption ends. The mean residual life $\bar{r}(i)$ is the expected remaining number of periods until the disruption ends. We note that $\bar{r}(i) \leq 1/\lambda(i)$ and that the increasing hazard rate assumption implies that $\bar{r}(i)$ is nonincreasing in $i = 1, \ldots, \infty$. Steady-state probabilities for the Markov chain are denoted by $\pi(i)$. The cumulative distribution function for the steady-state probabilities is given by $F[i] = \sum_{\tau=0}^{i} \pi(\tau)$.

The aim of this research is to provide insights into the factors that influence a firm's optimal disruption-management strategy. As such, the model should not be viewed as being designed for decision support. We have made a number of simplifying assumptions that merit discussion. We ignore any fixed cost of ordering. In essence, we assume that either fixed costs are negligible or the underlying base time unit is large enough that it is sensible to place an order in every period. This latter interpretation is reasonable only so long as the time scale for disruptions is larger than the time scale for orders. Our model is appropriate for firms that order on a daily or weekly basis and are concerned about disruptions (such as discussed in the introduction) that may last on the order of weeks or even months. If fixed costs are such that ordering once a month is optimal but disruptions are on the scale of days, then our model is inappropriate. The model is appropriate for monthly ordering as long as the firm is primarily faced with the risk of catastrophic disruptions (e.g., Hurricane Mitch), which can cause losses of supply that may last for months.

We assume equal lead times for both suppliers. In certain circumstances this is a reasonable approximation of reality. For example, one of Nokia's alternative suppliers to the U.S.-based Philips plant was also located in the United States. In other circumstances, lead times may differ significantly. As mentioned earlier, it has been established in the literature that multiple-sourcing is beneficial if suppliers differ in their lead times. We note that many of these papers either make a restrictive assumption that lead times differ by one unit (e.g., Fukuda 1964) or else resort to heuristics (e.g., Moinzadeh and Nahmias 1988) or simulation-based optimization (Scheller-Wolf and Tayur 1998) to solve models in which supplier lead times are not consecutive. Feng et al. (2005) show that the optimal policy structure is sensitive to the consecutive lead-time assumption and that a base-stock policy is no longer optimal when the assumption is relaxed. A goal of our research is to provide insight into disruption-management strategies, and we assume equal lead times so as to remove the nonidentical lead-time motive for dual-sourcing. All of the dual-sourcing, supply-uncertainty papers cited earlier make a similar assumption. A decision-support model might need to incorporate fixed costs and general lead-time suppliers. Solving such a model would likely require heuristics or simulation-based optimization.

In closing this section, we introduce a key lemma used in a number of later proofs. Consider the following inventory system. Demand in each period is stochastic but stationary. The demand random variable, denoted by D, has a strictly positive support. Supply is completely reliable, with a guaranteed lead time of $L \geq 0$. Ordering costs are linear but state dependent, with $c(0) = c_u$ and $c(i) = c_f$, $i > 0$. The state space and state transitions are identical to that described above for supplier U. Using a long-run average cost criterion, the following results hold for this inventory system.

Lemma 1. *A state-dependent, base-stock policy is optimal. The optimal base-stock levels $y^*(i)$ are such that $y^*(i) \leq y^M$ for $i \geq 1$, where y^M minimizes $E[\hat{C}(y - D^{(L)})]$, and $D^{(L)}$ is the L-fold convolution of D. $y^*(i)$ is nonincreasing in i. If $D_t = d$ with probability 1 (i.e., deterministic demand), then $y^*(0) \geq Ld$, $y^*(i) = Ld$ for $0 < i \leq i_{\mathrm{crit}}$, and $y^*(i) = -\infty$ for $i > i_{\mathrm{crit}}$, where i_{crit} is the maximum value of i such that $\bar{r}(i) > (c_f - c_u)/p$ if $\bar{r}(1) > (c_f - c_u)/p$ and $i_{\mathrm{crit}} = 0$ otherwise. If demand is $D_t' = kD_t$ where $k \geq 0$, then the optimal base-stock levels and the optimal cost are $ky^*(i)$ and kV^* respectively, where V^* is the optimal cost when $k = 1$.*

4. A Restricted Model

In this section, we focus on a restricted version of the model in which we assume that (i) the firm is risk neutral, (ii) demand is deterministic (equal to 1 without loss of generality), and (iii) supplier U has infinite capacity. Each of these assumptions will be relaxed in §5.

4.1. The Optimal Ordering Policy

On a tactical level, the firm needs to determine the optimal ordering policy. For a given supplier R allocation w, the firm must decide the quantity $q_u(0)$ to order from supplier U when it is up, and the quantity $q_f(i)$ to order from supplier R's volume flexibility when supplier U is down, i.e., the timing and quantity of contingent-rerouting orders.

We first consider two extreme cases of volume flexibility: (1) supplier R offers no flexibility, that is, $\delta(\tau)=0$ for $\tau=0,\ldots,\infty$; and (2) supplier R offers instantaneous and infinite volume flexibility, that is, $\delta(\tau)=\infty$ for $\tau=0,\ldots,\infty$ (in this case we also assume that the firm responds to a disruption immediately, i.e., $\theta_f=1$). We will use the abbreviation *II-flexibility* to refer to the instantaneous and infinite-flexibility case.

THEOREM 1. *In the zero-flexibility case, a base-stock policy is optimal for orders placed with supplier U when supplier U is up. Furthermore, $y_Z^*(0,w)\geq L$, where $y_Z^*(0,w)$ is the optimal base-stock level. In the II-flexibility case, a base-stock policy is optimal for orders placed with supplier U when supplier U is up, and a state-dependent base-stock policy is optimal for the contingent-rerouting orders when supplier U is down. Furthermore, $y_{II}^*(0,w)\geq L$, $y_{II}^*(i,w)=L$ for $0<i\leq i_{\text{crit}}$, and $y_{II}^*(i,w)=-\infty$ for $i>i_{\text{crit}}$, where $y_{II}^*(i,w)$ is the optimal base-stock level in state i.*

Demand in each period is 1 and the firm always orders w from supplier R's regular capacity. Therefore, in the II-flexibility case, $q_f(i)=[y_{II}^*(i,w)-(z(i_-)-(1-w))]^+$ is the rerouting quantity required in state i to bring the ending inventory position to at least $y_{II}^*(i,w)$, where i_- denotes the state immediately prior to state i and $z(i_-)$ is the ending inventory position in the prior period. We note that i_{crit} captures the trade-off between rerouting and back orders and that i_{crit} increases as the rerouting cost decreases relative to the back-order cost. The contingent-rerouting policy has the following properties.

- There exists a threshold value i_{crit} such that if the disruption has lasted more than i_{crit} periods, then it is optimal to wait until the disruption is over rather than to reroute production, that is, $q_f(i)=0$ for all inventory positions if $i>i_{\text{crit}}$.
- If the disruption has lasted less than or equal to i_{crit} periods and the inventory position is high enough to prevent a back order L periods from now, then the firm should not reroute any production.
- If the disruption has lasted less than or equal to i_{crit} periods but the inventory position is not high enough to prevent a back order L periods from now, then the firm should reroute a sufficient quantity to prevent any back orders but not so large a quantity as to build inventory, i.e., it should reroute just enough to bring its inventory position to L.

We now consider the case of partial flexibility or a delayed firm response. In this case the firm faces constraints on the quantity it can reroute. In particular, $\delta(i-\theta_{SC})=0$ for $i<\theta_{SC}$ and $\delta(i-\theta_{SC})=\delta$ for $i\geq\theta_{SC}$, that is, it cannot reroute any quantity if $i<\theta_{SC}$ and can reroute at most δ per period if $i\geq\theta_{SC}$. Recall that $\theta_{SC}=\theta_f+\theta_r$.

THEOREM 2. *In the partial-flexibility (or delayed-firm-response) case, a base-stock policy is optimal for orders placed with supplier U when supplier U is up, and a modified state-dependent base-stock policy is optimal for the contingent-rerouting orders when supplier U is down. If the firm's inventory position is below the state-dependent base-stock level $y_{PF}^*(i,w)$, then it should reroute*

$$q_f(i)=\min\{[y_{PF}^*(i,w)-(z(i-1)-(1-w))]^+,\delta(i-\theta_f)\},$$

where $y_{PF}^*(i,w)=L+n(i)[1-w-\delta]^+$ for $0<i\leq i_{\text{crit}}$; $y_{PF}^*(i,w)=-\infty$ for $i>i_{\text{crit}}$; $n(i)$ is the maximum integer $n\geq 0$ such that $c_f-c_u<M(i,n)$ and $n(i)$ is defined only for $0<i\leq i_{\text{crit}}$; $M(i,n)=(p+h)P[r(i)\geq n]\bar{r}(i+n)-h\bar{r}(i)$; and $P[r(i)\geq n]$ is the probability that the residual life is at least n periods.

Note that $n(i)\geq 0$ as $i\leq i_{\text{crit}}$ and that the $n(i)$ are nonincreasing in i. This theorem states that during a disruption, the firm attempts to reach a state-dependent base-stock level, but it may not be able to do so because of the volume-flexibility capacity constraint. We see again that it is not optimal to reroute if $i>i_{\text{crit}}$. In the II-flexibility case, the optimal base-stock level is L for $0<i\leq i_{\text{crit}}$, the logic being that contingent rerouting is perfectly reliable, demand is constant at 1, and the lead time is L. With partial flexibility, this result still holds if $\delta\geq 1-w$, the intuition being that if the inventory position is above L and the flexible capacity has come online, the firm should not yet reroute as it has sufficient capacity to always keep the inventory position at L if it so chooses. The rerouting policy differs, however, if $\delta<1-w$. In this case, the extra capacity δ is not sufficient to compensate for the lost supply $1-w$, and so the firm may reroute lost production before its inventory position falls below L so as to prevent future back orders that result from the rerouting-capacity constraint. Rerouting before the inventory position falls below L means that the firm will incur additional inventory costs until the inventory position falls below L but will reduce the back orders incurred after that point. This trade-off is reflected in the $M(i,n)$ function.

If $i_{\text{crit}} = 0$, then the firm never reroutes lost production in any state, i.e., the contingency strategy is never used to manage disruption risk. However, $i_{\text{crit}} > 0$ does not imply that a contingency strategy is necessarily used; the optimal base-stock level in state 0 might be sufficient to last beyond i_{crit}, in which case rerouting will not occur.

4.2. The Optimal Base-Stock Level When Supplier U Is Up

Having characterized the optimal ordering policy, we now proceed to develop the long-run average cost expression as a function of $y(0, w)$ and w. Recall that $y(0, w)$ is the base-stock level when supplier U is up, and w is supplier R's allocation. Letting $y(0, w) = I_0 + L$, we will develop the long-run average cost expression as a function of I_0 and w. We will then optimize for the base-stock level and supplier R's allocation. In what follows, we will at times use the term "average" in place of "long-run average." The average cost can be written as

$$C_{LRA} = c_u q^u_{LRA}(I_0, w) + c_r q^r_{LRA}(I_0, w) + c_f q^f_{LRA}(I_0, w) + hI^+_{LRA}(I_0, w) + pI^-_{LRA}(I_0, w), \qquad (2)$$

where for a given base stock $y(0, w) = I_0 + L$ and supplier R allocation w,

- $q^u_{LRA}(I_0, w)$ is the average quantity per period sourced from supplier U,
- $q^r_{LRA}(I_0, w)$ is the average quantity per period sourced from supplier R's regular capacity,
- $q^f_{LRA}(I_0, w)$ is the average quantity per period sourced from supplier R's flexible capacity,
- $I^+_{LRA}(I_0, w)$ is the average (nonnegative) on-hand inventory level, and
- $I^-_{LRA}(I_0, w)$ is the average back order level.

We note that $q^u_{LRA}(I_0, w) + q^r_{LRA}(I_0, w) + q^f_{LRA}(I_0, w) = 1$ and that $q^r_{LRA}(I_0, w) = w$ by definition. Therefore, Equation (2) can be rewritten as

$$C_{LRA} = c_u + (c_r - c_u)w + (c_f - c_u)q^f_{LRA}(I_0, w) + hI^+_{LRA}(I_0, w) + pI^-_{LRA}(I_0, w). \qquad (3)$$

As detailed in Appendix A2 (in the online supplement), we can use the renewal reward theorem to derive the following expressions for the average quantities:

$$q^f_{LRA}(I_0, w) = \sum_{i=1}^{\infty} q_f(i)\pi(i), \qquad (4)$$

$$I^+_{LRA}(I_0, w) = I_0\pi(0) + \sum_{i=1}^{\infty}\left[I_0 - i(1-w) + \sum_{k=1}^{i} q_f(k)\right]^+ \pi(i), \qquad (5)$$

$$I^-_{LRA}(I_0, w) = \sum_{i=1}^{\infty}\left[I_0 - i(1-w) + \sum_{k=1}^{i} q_f(k)\right]^- \pi(i). \qquad (6)$$

These long-run average expressions are in fact the steady-state expected values, and so the long-run average cost is equal to the steady-state expected cost. The quantity rerouted in state i, $q_f(i)$, depends on the volume-flexibility profile. Expressions for $q^f_{LRA}(I_0, w)$, $I^+_{LRA}(I_0, w)$, and $I^-_{LRA}(I_0, w)$ tailored to the volume-flexibility profile assumed in this paper can be found in Appendix A2 (in the online supplement).

We now define two key variables that influence the optimal base-stock level:

$$i_1^* = F^{-1}\left[\frac{p}{p+h}\right],$$

$$i_S = \begin{cases} 0, & \text{if } (c_f - c_u)\pi(1) < hF[0] \\ \text{maximum } i \text{ such that } (c_f - c_u)\pi(i) \geq hF[i-1], & \text{otherwise.} \end{cases}$$

Both i_1^* and i_S refer to numbers of periods and capture different trade-offs facing the firm. The trade-off between inventory and back orders is captured by i_1^*. We note that i_1^* increases as inventory becomes cheaper relative to back orders. One can think of i_1^* as a newsvendor-type variable. In fact, as shown below, i_1^* is the number of periods' coverage provided by the optimal base stock in the zero-flexibility case. The trade-off between inventory and rerouting is captured by i_S. It is cheaper for the firm to carry an additional unit of inventory than to reroute in period i of a disruption if $i \leq i_S$. Note that i_S increases as inventory becomes cheaper relative to rerouting. As noted earlier, the trade-off between back orders and rerouting is captured by i_{crit}. For a given allocation w, the optimal base-stock level $y^*(0, w) = I_0^*(w) + L$ is specified by the following theorem for all values of θ_f, θ_r, and δ. Recall that $\theta_{SC} = \theta_f + \theta_r$.

Theorem 3. *If $w = 1$, then $I_0^*(w) = 0$. If $w < 1$, then*

$$i_{\text{crit}} < \max\{\theta_{SC}, i_1^*\} \Rightarrow I_0^*(w) = i_1^*(1-w),$$

$$i_{\text{crit}} \geq \max\{\theta_{SC}, i_1^*\},\ \delta \leq 1-w$$

$$\Rightarrow I_0^*(w) = \begin{cases} i_1^*(1-w), & \text{if } i_1^* < \theta_{SC}, \\ i_R(1-w) + n(i_R)(1-w-\delta) - \delta, & \text{otherwise,} \end{cases}$$

$$i_{\text{crit}} \geq \max\{\theta_{SC}, i_1^*\},\ \delta > 1-w$$

$$\Rightarrow I_0^*(w) = \begin{cases} i_S(1-w), & \text{if } i_S \geq \theta_{SC} - 1, \\ i_1^*(1-w), & \text{if } i_S < \theta_{SC} - 1 \text{ and } i_1^* \leq \bar{i}(w), \\ \bar{I}_0(w), & \text{otherwise,} \end{cases}$$

where $i_R = \max\{\theta_{SC}, i_S + 1\}$,

$$\bar{I}_0(w) = \begin{cases} \left[\begin{array}{l}(\theta_{SC}-1)(1-w) \\ \quad -(\delta-(1-w))(i_{\text{crit}}-\theta_{SC})\end{array}\right]^+, \\ \quad \text{if } \left(\begin{array}{l}(h+p)F[\bar{i}(w)] - pF[i_{\text{crit}}-1] \\ \quad -(c_f-c_u)\pi(i_{\text{crit}})\end{array}\right) \geq 0, \\ (\theta_{SC}-1)(1-w), \\ \quad \text{if } \left(\begin{array}{l}(h+p)F[\theta_{SC}-1] - pF[\theta_{SC}] \\ \quad -(c_f-c_u)\pi(\theta_{SC}+1)\end{array}\right) \leq 0, \\ \text{maximum } I_0 \text{ such that} \\ \quad \left(\begin{array}{l}(h+p)F[T(I_0,w)] - pF[N(I_0,w)-1] \\ \quad -(c_f-c_u)\pi(N(I_0,w))\end{array}\right) \geq 0, \\ \quad \text{otherwise,} \end{cases}$$

and

$$T(I_0, w) = \left\lfloor \frac{I_0}{1-w} \right\rfloor,$$

$$N(I_0, w) = \begin{cases} \theta_{SC} - 1 + \left\lceil \dfrac{(\theta_{SC}-1)(1-w) - I_0}{\delta-(1-w)} \right\rceil, & T(I_0,w) < \theta_{SC}-1 \\ T(I_0,w)+2, & T(I_0,w) \geq \theta_{SC}-1, \end{cases}$$

$$\bar{i}(w) = \left\lfloor \frac{[(\theta_{SC}-1)(1-w) - (\delta-(1-w))(i_{\text{crit}}-\theta_{SC})]^+}{(1-w)} \right\rfloor.$$

4.3. The Optimal Sourcing Strategy

We now proceed to determine the optimal sourcing strategy. Recall that $w^* = 1$ implies that the firm single-sources from supplier R, $0 < w^* < 1$ implies the firm dual-sources, and $w^* = 0$ implies the firm single-sources from supplier U.

We first consider the two extreme cases of volume flexibility: II-flexibility and zero flexibility.

Theorem 4. *Single-sourcing is optimal in both cases, that is,* $w^*_{II} \in \{0, 1\}$ *and* $w^*_Z \in \{0, 1\}$. *If single-sourcing from supplier U is optimal for the zero-flexibility case, then it is also optimal for the II-flexibility case, i.e.,* $w^*_Z = 0 \Rightarrow w^*_{II} = 0$. *If single-sourcing from supplier R is optimal for the II-flexibility case, then it is also optimal for the zero-flexibility case, i.e.,* $w^*_{II} = 1 \Rightarrow w^*_Z = 1$.

This theorem tells us that an extreme sourcing strategy is optimal at both ends of the flexibility profile. Does this extreme result hold for intermediate flexibility regimes? To answer this question, we now consider the case in which supplier R offers only partial flexibility and/or the firm does not respond instantaneously to a disruption. Define $E_1(i) = \sum_{\tau=0}^{i} \tau\pi(\tau)$, $E_2(i) = \sum_{\tau=i+1}^{\infty} \tau\pi(\tau)$, $K_1(i) = E_1(i) - iF[i]$, and $K_2(i) = E_2(i) - i(1 - F[i])$. The optimal sourcing strategy for the case of partial flexibility (or delayed firm response) is given in the following theorem.

Theorem 5. *If* $i_1^* > i_{\text{crit}}$ *(that is, rerouting is too expensive relative to inventory), then volume flexibility is never used and single-sourcing is optimal, i.e.,* $w^* \in \{0, 1\}$, *with*

$$w^* = 0 \iff c_r \geq c_u - hK_1(i_1^*) + pK_2(i_1^*). \tag{7}$$

If $i_1^* \leq i_{\text{crit}}$, *but* $\theta_{SC} > i_{\text{crit}}$ *(that is, the supply chain response time is too slow and/or rerouting is too expensive relative to back orders), then volume flexibility is never used and single-sourcing is optimal, i.e.,* $w^* \in \{0, 1\}$, *with Equation (7) again determining the optimal supplier choice. If* $i_1^* \leq i_{\text{crit}}$, *and* $\theta_{SC} \leq i_{\text{crit}}$ *(that is, rerouting is a viable option both in terms of cost and supply chain response time), then the optimal sourcing strategy depends on both the supply chain response time* θ_{SC} *and the flexibility magnitude* δ *in the manner specified by Table 2.*

If dual-sourcing is optimal, then $w^* \geq [1-\delta]^+$, that is, the firm chooses an allocation such that the magnitude of flexibility is sufficient to prevent (if it chooses to) any further increase in back orders once flexibility becomes available. We segmented the sourcing strategy by the flexibility parameters, as this gave the cleanest segmentation. Of course, other parameters influence the optimal allocation, either indirectly through their influence on i_1^*, i_S, and i_{crit}, or directly through their influence on the purchasing costs. In numerical tests, the behavior of w^* with respect to model parameters was as one would expect, e.g., supplier R's allocation was increasing in the back-order cost p and inventory-holding cost h and decreasing in its relative cost c_r/c_u and supplier U's reliability [as measured by $\pi(0)$, the steady-state probability that U is up]. Supplier R's allocation was found to be particularly sensitive to relative cost and reliability. We also found that supplier R's allocation decreased in the level of flexibility it offered, as flexibility allows the firm to engage in contingent rerouting rather than mitigation sourcing (i.e., routine sourcing from supplier R). Suppliers are therefore advised to be cautious about offering flexibility.

4.4. The Optimal Disruption-Management Strategy

Having characterized the optimal rerouting, inventory, and sourcing decisions for the firm, we now proceed to use these results to characterize the set

Table 2 Optimal Sourcing Strategy When $i_1^* \leq i_{\text{crit}}$ and $\theta_{SC} \leq i_{\text{crit}}$

	Flexibility magnitude	
Response time	Low (i.e., $\delta < 1$)	High (i.e., $\delta \geq 1$)
Intermediate response (i.e., $\theta_{SC} > i_S + 1$)	Single- or dual-source $w^* = 0$ or $1 - \delta \leq w^* \leq 1$	Single- or dual-source $0 \leq w^* \leq 1$
Fast response (i.e., $\theta_{SC} \leq i_S + 1$)	Single- or dual-source $w^* \in \{0, 1-\delta, 1\}$	Single-source $w^* \in \{0, 1\}$

of possible optimal disruption-management strategies for the three flexibility cases: zero flexibility, II-flexibility, and partial flexibility (see Table 3). Note that the contingency strategy (rerouting) is not available to the firm in the zero-flexibility case, and so mitigation (either through inventory or supplier R sourcing) is the only option for actively managing the disruption risk in that situation. We proved above that single-sourcing is optimal in the zero-flexibility and II-flexibility cases, and therefore a mixed mitigation strategy of partially sourcing from supplier R and carrying inventory cannot be optimal in those cases.

The impact of the various costs on the attractiveness of a given strategy are as one would expect. Mitigation (through either sourcing from supplier R or carrying inventory) becomes less attractive as the relative cost of supplier R (c_r/c_u) increases or as the cost of inventory h increases. Contingent rerouting becomes less attractive as the premium c_f/c_r paid for volume-flexible units increases. Of more interest is the impact that supplier reliability and expected disruption length have on the optimal disruption-management strategy. We measure supplier U reliability by the percentage of time that supplier U is up, i.e., $\pi(0)$. A given percentage uptime can result from frequent but short disruptions or from rare but long disruptions. We therefore classify disruption processes by their expected lengths. We assume the following structure for the disruption process in the numeric examples that follow: a disruption lasts for a minimum of M periods, after which there is a constant probability λ_{du} of the disruption ending in each period. In other words, a disruption length is the sum of a constant and a geometric random variable. Table 4 describes the labeling scheme used in Figures 1 to 3.

For the zero-flexibility case, we present the three optimal disruption-management strategies as a function of supplier U's percentage uptime and the expected disruption length for the following set of parameters: $c_u = 1$, $c_r = 1.05$, $p = 0.15$, $h = 0.0015$, $\lambda_{du} = 0.1$, $L = 0$ (Figure 1). For a given percentage uptime, the expected disruption length was varied by changing M. An increase in the expected disruption length reduces the probability (or frequency) of a disruption. Disruptions are frequent but short at the bottom left of the figure, and very rare but very long at the top right. While the relative sizes of the various regions changed as we varied the model parameters, the relative positioning of the regions did not. We see in Figure 1 that acceptance is the optimal strategy if supplier U has a very high percentage uptime. It can be shown in general that acceptance is optimal if and only if $\pi(0) \geq p/(p+h)$ and $c_r \geq c_u + p\sum_{\tau=1}^{\infty} \tau\pi(\tau)$; that is, the percentage uptime is at least as large as

Table 3 Possible Optimal Disruption-Management Strategies

Disruption-management strategy	Flexibility profile: Zero	II	Partial
Acceptance	Yes	Yes	Yes
Mitigation only			
Inventory	Yes	Yes	Yes
Sourcing exclusively from R	Yes	Yes	Yes
Contingency only			
Rerouting	No	Yes	Yes
Mitigation and contingency			
Inventory and rerouting	No	Yes	Yes
Inventory and partially sourcing from R	No	No	Yes
Rerouting and partially sourcing from R	No	No	Yes
Inventory, rerouting, and partially sourcing from R	No	No	Yes

Table 4 Labeling Scheme for Disruption-Management Strategies

Label	Strategy	Description
A	Acceptance	The firm passively accepts the disruption risk. It sources exclusively from the unreliable supplier and carries no inventory.
IM	Inventory mitigation	The firm sources exclusively from the unreliable supplier but carries some inventory to mitigate disruptions.
SM	Sourcing mitigation	The firm sources exclusively from the reliable supplier.
CR	Contingent rerouting	The firm sources exclusively from the unreliable supplier when that supplier is up. The firm carries no inventory, but it reroutes to the reliable supplier during a disruption.
IMCR	Inventory mitigation and contingent rerouting	The firm sources exclusively from the unreliable supplier when that supplier is up. The firm carries some inventory to mitigate disruptions, but during a disruption it may also reroute production to the reliable supplier.

Figure 1 Optimal Disruption-Management Strategies for Zero-Flexibility Case

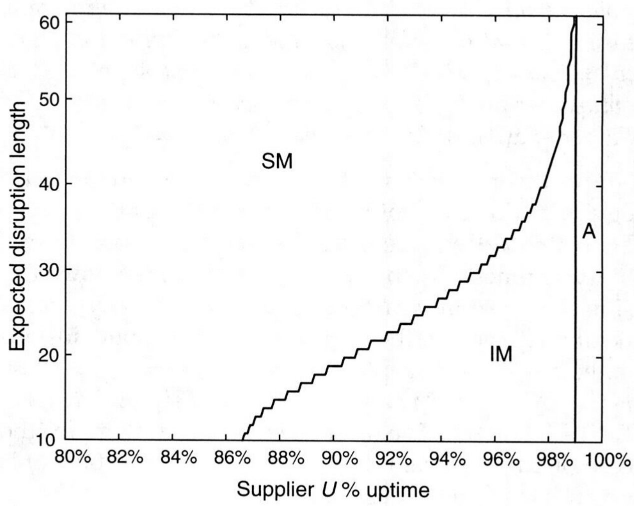

Figure 2 **Optimal Disruption-Management Strategies for II-Flexibility Case** ($c_f = 2.5c_r$)

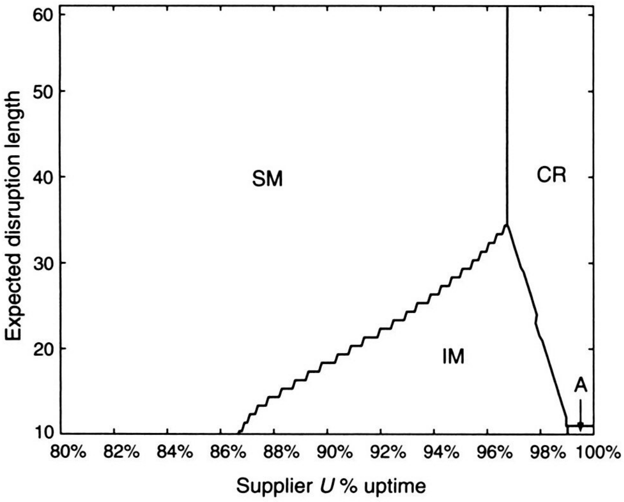

a newsvendor fractile and the reliable-supplier cost is at least as large as the unreliable-supplier cost plus expected back-order costs. Acceptance is therefore favored in environments of high supplier reliability, with the definition of "high" being dependent on the costs of mitigation and the cost of back orders. For a given percentage uptime, inventory mitigation is favored if disruptions tend to be short (and hence more frequent), but sourcing mitigation is favored if disruptions tend to be long (and hence less frequent). The reason that sourcing mitigation is favored over inventory mitigation as disruptions become longer and less frequent is that more inventory is required in such environments, and so inventory mitigation becomes less attractive. Carrying high amounts of inventory for rare events is not an attractive strategy if another mitigation option (in our case, sourcing from supplier R) is available.

Figure 3 **Optimal Disruption-Management Strategies for II-Flexibility Case** ($c_f = 1.25c_r$)

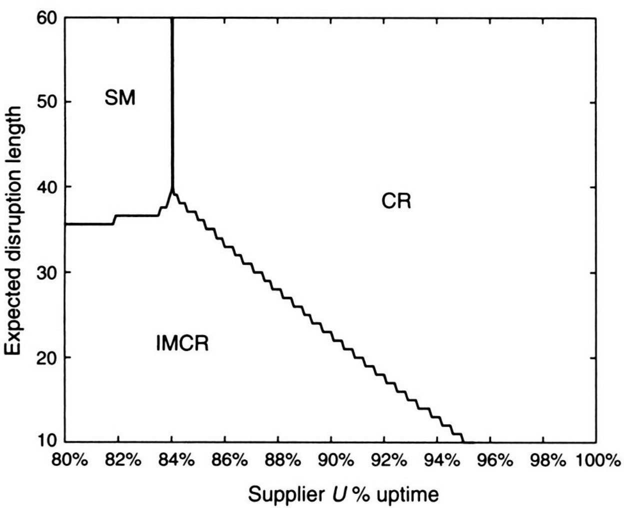

Mitigation is the only available strategy in the zero-flexibility case, but contingent rerouting is available when supplier R provides volume flexibility. While contingent-rerouting costs more than regular production, i.e., $c_f \geq c_r > c_u$, the benefit of the contingency strategy is that the firm only incurs the higher contingent-rerouting cost during a disruption. With mitigation strategies, the firm incurs a cost (either inventory or the higher supplier R sourcing cost) even when supplier U is up. One might therefore expect that a contingency strategy would be preferred if disruptions are rare events, i.e., if the probability of supplier failure is very low. The story, however, is somewhat more nuanced than this. Using the same parameters as above, we present the optimal disruption-management strategy for the II-flexibility case when the relative rerouting cost is very high, i.e., $c_f = 2.5c_r$ (see Figure 2) and when it is lower, i.e., $c_f = 1.25c_r$ (see Figure 3). We see from Figure 2 that a contingency strategy can be optimal even for a very high rerouting cost. As the rerouting cost decreases, contingent rerouting (either in isolation or combined with inventory mitigation) is preferred over a larger region. Again, while the relative sizes of the various regions vary as parameters change, the relative positioning of the regions does not. Contingent rerouting becomes less attractive as supplier U's percentage uptime decreases (for a given expected disruption length), the reason being that the benefit of the contingent-rerouting strategy decreases as the higher contingency-rerouting cost is incurred a greater percentage of the time. Sourcing mitigation is never optimal in the II-flexibility case if $c_f = c_r$, as contingent rerouting is no more expensive than mitigation sourcing but is incurred less frequently. Contingent rerouting also becomes less attractive as the expected disruption length increases (for a given supplier U percentage uptime). Eventually mitigation, either through inventory or routine sourcing from supplier R, becomes optimal as the disruption length increases. This indicates that a mitigation strategy is preferred to the contingency strategy when the frequency of disruption is very low, a result that might seem counterintuitive. The reason for this again lies with the fraction of time the contingent cost is incurred. Whereas the percentage downtime is constant for a given percentage uptime, the percentage of time spent rerouting is not constant; it depends on the expected disruption length. Recall from Theorems 1 and 2 that the use of contingent rerouting depends on the residual life of a disruption. As the expected disruption length increases, the firm spends more time in states in which rerouting is optimal, and hence the relative benefit of the contingency strategy decreases until a point is reached at which mitigation becomes optimal.

We see from Figures 1, 2, and 3 that the optimal disruption-management strategy depends on whether volume flexibility is available. Volume flexibility enables a contingency strategy, and a contingency strategy can be preferable to a mitigation (or acceptance) strategy. Does volume flexibility significantly reduce a firm's long-run average costs? In Figures 4 and 5 we illustrate the relative cost reduction achieved by II-flexibility (over zero flexibility) as a function of supplier U's percentage uptime and the expected disruption length using the same parameters as in Figures 2 and 3. The relative cost reduction is zero in regions where contingent rerouting is not an optimal strategy, and so the three-dimensional surface graphs reflect the regions illustrated in Figures 2 and 3. In the regions where contingent rerouting is optimal, we see that II-flexibility can result in a cost reduction of 3%–4%. This is a reduction in the overall costs (including procurement) and not simply a reduction in inventory and back-order costs. As such, this is a significant reduction. In fact, if one could make supplier U perfectly reliable in the zero-flexibility case, then the maximum savings from doing this would be 4.76% (as the reliable-supplier cost of 1.05 places an upper bound on the possible long-run average cost when supplier U is unreliable). We therefore see that the availability of II-flexibility is almost as beneficial as perfect supplier reliability even when the cost of exercising that flexibility is expensive. Note that as the rerouting cost decreases to $c_f = c_r$, the relative cost reduction increases and the region in which those reductions occur also grows.

II-flexibility is unlikely to exist for most firms, either because suppliers do not offer it (e.g., Nokia's suppliers took five days to increase production) or because the firm does not respond instantaneously to a supplier disruption (e.g., Ericsson did not respond

Figure 4 **Cost Reduction Arising from II-Flexibility** ($c_f = 2.5c_r$)

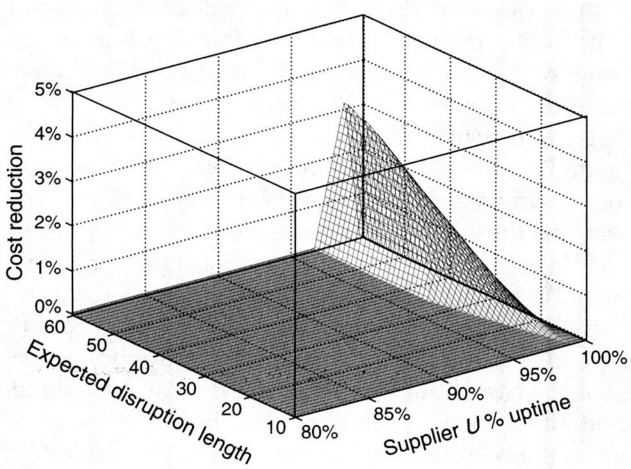

Figure 5 **Cost Reduction Arising from II-Flexibility** ($c_f = 1.25c_r$)

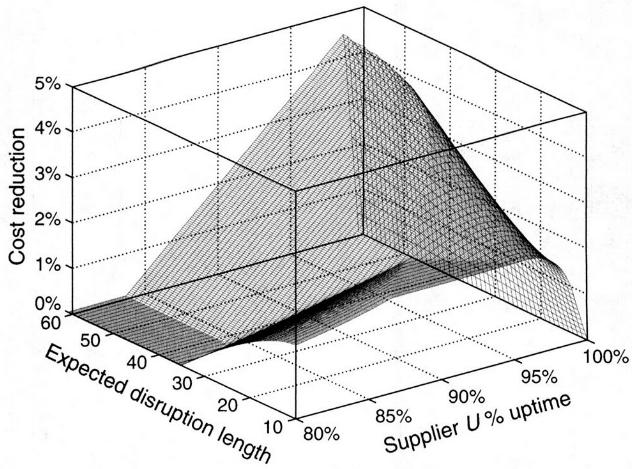

with contingent-rerouting requests until several of weeks after the disruption). Partial flexibility is a more reasonable model of the reality most firms face. It is of interest to understand how quickly the benefits of II-flexibility decrease as flexibility decreases. Recall that we parameterize supplier flexibility by the magnitude δ and the response time θ_r, and that flexibility decreases as the magnitude decreases or as the response time increases. In Figure 6 we present exchange curves (or contours) for magnitude and response time for three different expected disruption lengths (expressed in numbers of periods). The contours specify the percentage (from 10% to 90%) of the II-flexibility benefit that is delivered by partial flexibility. An uptime percentage of 97% was chosen and the same cost parameters as in Figure 3 were used.

Two observations are noteworthy, as they were observed in many other numeric examples. First, for very high magnitudes, the performance of partial flexibility is more sensitive to response-time degradation when the expected disruption length is low. This reflects the fact that the extra capacity is of little value if a disruption is nearly over by the time capacity becomes available. Second, the tradeoff between magnitude and response is more extreme when disruptions are short (and hence more frequent) in nature: for high response times, a large increase in magnitude is needed to offset a small degradation in response time (again this reflects the fact that magnitude is of little value if it becomes available when a disruption is nearly over), but for low response times, a very small increase in magnitude offsets a large degradation in response time (this reflects the fact that inventory can compensate for an increase in response time, and as discussed above, carrying inventory is relatively less expensive when disruptions are short and frequent rather than long and rare). Measuring response time relative to the expected disruption length, we found that for very high magnitudes

Figure 6 Partial Flexibility Exchange Curves for a 97% Uptime Percentage as Expected Disruption Length (EDL) Increases

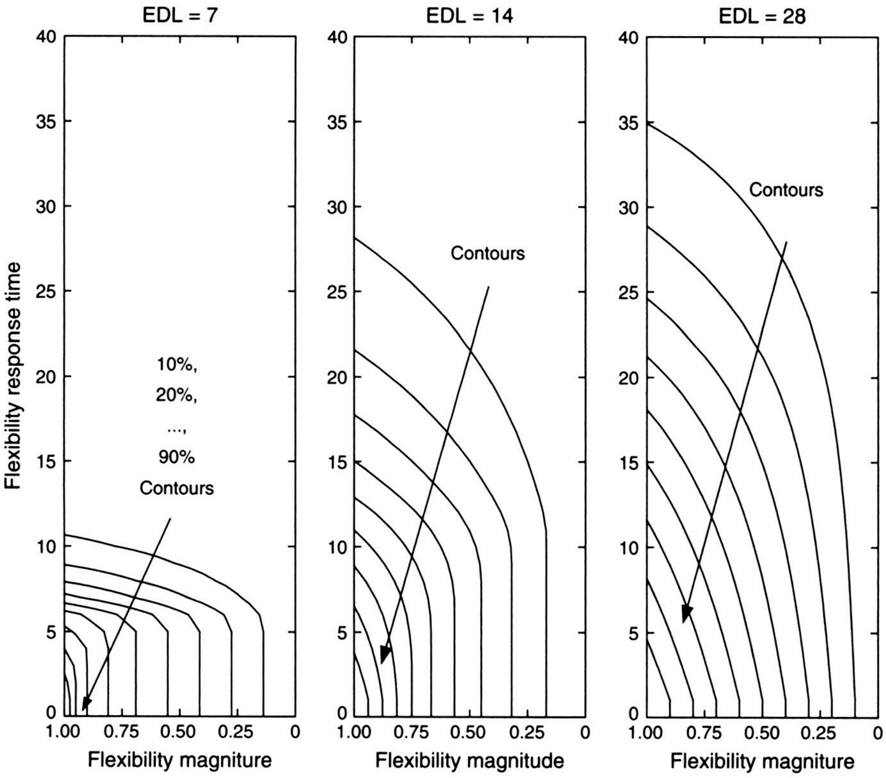

the performance of partial flexibility is more sensitive to the relative response time when the expected disruption length is greater. This observation again reflects the fact that inventory can compensate for an increase in response time. As inventory is unattractive when disruptions are long and rare, inventory is less effective at compensating for an increase in relative response time.

In closing we mention that all the above figures assume instantaneous firm response time, i.e., $\theta_f = 0$. For a noninstantaneous firm response time, the flexible capacity becomes available after the sum of the firm's response time and the supplier's response time, i.e., $\theta_f + \theta_r$, and so one can still use the above figures by interpreting the response time to be the combined response time.

4.5. The Impact of Misestimating Supplier Reliability

Our model assumes that the firm can accurately characterize the reliability of a supplier. In practice, a firm may have to estimate a supplier's reliability when choosing a disruption-management strategy. We conducted a numeric study to investigate the impact of parameter misestimation on the long-run average cost. We assumed the same disruption structure as in the earlier numeric examples, namely a disruption length that is the sum of a constant plus a geometric random variable. Recall that the model assumes a constant probability of failure when supplier U is up. In our study, we focused on the misestimation of the failure probability, and assumed that the firm could accurately estimate repair probabilities. There is a one-to-one relationship between the failure probability and supplier U's uptime percentage (for a given disruption process), and we therefore use the uptime percentage as the parameter that the firm estimates. For each problem instance we assumed that the true uptime percentage was 97% but that the firm's estimate of the uptime percentage could deviate from the true value by as much as ±3%. We solved for the firm's optimal strategy given its estimate of the uptime percentage and then calculated the actual long-run average cost of the resulting strategy given the true uptime percentage. The impact of uptime misestimation was then measured by the percentage increase in long-run average cost resulting from the incorrect estimate. We fixed the holding cost h at 0.0015, the unreliable-supplier cost c_u at 1, the flexibility premium c_f/c_r at 1.05, the firm response time θ_f at 1, and the lead time L at 1. We conducted a full-factorial study for the parameters listed in Table 5.

In Table 6 we present the percentage cost increase (over all the relevant instances) resulting from supplier uptime misestimations for zero flexibility and II-flexibility ($\delta = 1$, $\theta_r = 0$). While misestimation can significantly increase costs (with the cost increase

Table 5 Parameters Used in Misestimation Study

Parameter	Values used in study				
Flexibility magnitude, δ	0	0.25	0.50	0.75	1.00
Flexibility response time, θ_r	0	3	6	9	12
Reliable-supplier cost, c_r	1.02	1.04	1.06		
Back-order cost, p	0.05	0.10	0.15		
Expected disruption length	11	16	31		

being as high as 7.36% in one case), II-flexibility significantly reduces the negative effect of misestimation. In addition to reducing a firm's long-run average cost, volume flexibility results in a disruption-management strategy that tends to be more robust to parameter misestimation. We observed the following results in the numeric study. The cost increase due to overestimation was typically higher than the cost increase due to underestimation. There were, however, instances for which the reverse was true—underestimation tended to result in a larger cost increase than overestimation when the back-order cost was low. The cost increase arising from misestimation increased as the back-order cost increased and as the expected disruption length increased. The cost increase typically decreased as the flexibility magnitude increased and as the flexibility response time decreased, but there were some instances in which flexibility amplified the misestimation penalty.

5. Extensions to the Restricted Model

We now relax three key assumptions of the restricted model: risk neutrality, deterministic demand, and infinite supplier U capacity. Unless otherwise stated, we assume instantaneous lead times ($L = 0$) in all that follows.

5.1. Risk Aversion in the Allocation Decision

Up to this point we have assumed that the allocation decision w is made with an objective of minimizing the long-run average cost. We now relax this assumption by allowing for risk aversion in the allocation decision.

Table 6 Average and Maximum Percentage Cost Increases Due to Uptime Misestimation

	Percentage uptime misestimate						
	−3.0%	−2.0%	−1.0%	0%	1.0%	2.0%	3.0%
Zero flexibility							
Average	0.25	0.18	0.05	0.00	0.13	1.07	1.74
Maximum	0.81	0.81	0.17	0.00	1.03	5.00	7.36
II-flexibility							
Average	0.02	0.01	0.01	0.00	0.01	0.40	0.53
Maximum	0.12	0.07	0.07	0.00	0.07	0.84	1.14

Using a mean-variance approach, we consider a firm that is concerned with both the expectation $C(w)$ and variance $V(w)$ of the steady-state distribution of costs for a given allocation w. Recall that the long-run average cost was shown to equal the steady-state expected cost per period. A common mean-variance objective is to minimize $C(w) + \eta V(w)$ where $\eta \geq 0$ is a variance penalty. Multiplying this objective by $1/(1+\eta)$, we obtain an equivalent objective function

$$\Pi(w, \beta) = (1-\beta)C(w) + \beta V(w), \quad 0 \leq \beta \leq 1, \qquad (8)$$

where $\beta = \eta/(1+\eta)$ is the relative variance penalty. At $\beta = 0$ the firm minimizes the expected cost per period, and at $\beta = 1$ the firm minimizes the variance of costs. Let $w^*(\beta)$ minimize $\Pi(w, \beta)$.

THEOREM 6. *For any flexibility profile (θ_r, δ), $w^*(\beta)$ is nondecreasing in β, that is, a higher allocation is sourced from supplier R as the relative variance penalty increases.*

We proved above that single-sourcing is optimal in the zero-flexibility and II-flexibility cases for a risk-neutral firm. This result is no longer true under a mean-variance objective. We observed in numerical tests that the optimal allocation, and hence the optimal disruption-management strategy, was sensitive to the variance penalty. In Figure 7 we present the optimal disruption-management strategies for both the risk-neutral and mean-variance objectives as a function of supplier U's percentage uptime and the disruption length for the zero-flexibility case using the following set of parameters: $c_u = 1$, $c_r = 1.05$, $p = 0.15$, $h = 0.0015$, $\lambda_{du} = 1$, and $\beta = 0.05$. We use the same labeling scheme as described earlier in Table 4, but there is now an additional strategy, labeled MPSI, which is mitigation through partial sourcing from supplier R and carrying inventory. In this strategy, the firm sources from both suppliers if there is no disruption and also carries inventory to mitigate disruptions. There is a large region in which inventory mitigation (IM) is the optimal strategy in the risk-neutral setting. In contrast, MPSI is the dominant strategy for the mean-variance objective. We note that the fraction sourced from supplier R increases in the disruption length and decreases in supplier U's uptime. As with the risk-neutral case, single-sourcing from supplier R becomes optimal as disruptions become longer (and hence more rare) for a given supplier U percentage uptime.

Arguably, a mean-variance approach is inappropriate for firms facing situations in which there are relatively small probabilities of severe events, i.e., very rare but very long disruptions. In such cases, the firm might wish to minimize expected cost subject to some constraint on downside risk. A CVaR approach can be used to handle such situations. We refer the reader to the unabridged version of this paper for a treatment of the CVaR approach.

Figure 7 Optimal Disruption Strategies for Risk-Neutral and Mean-Variance Objectives

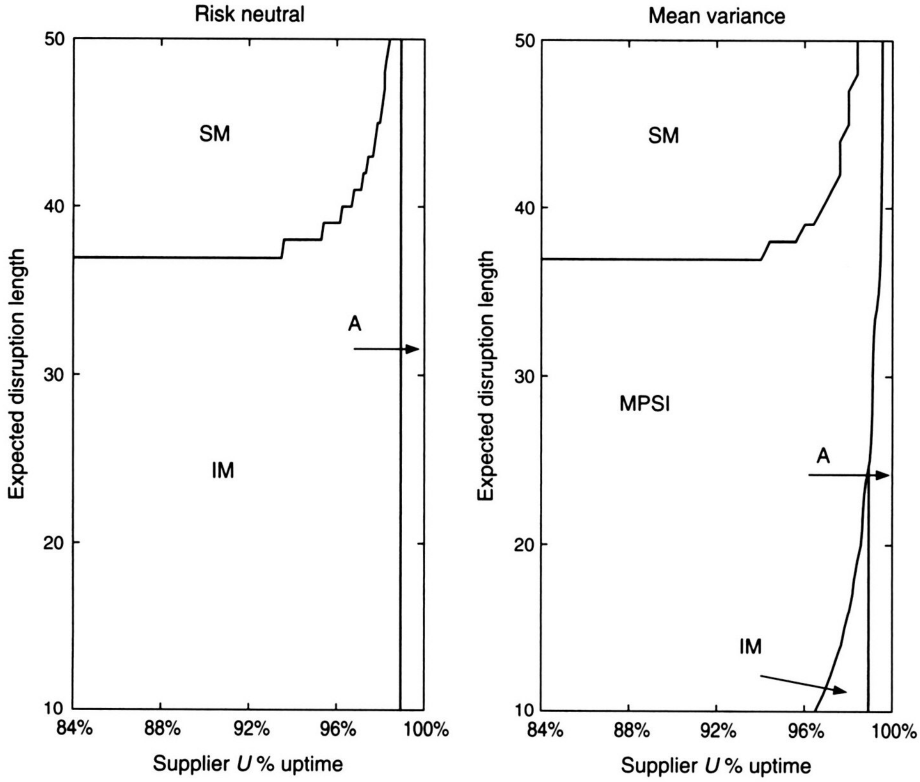

5.2. Stochastic Demand

Because single-sourcing is optimal in the zero-flexibility and II-flexibility cases for the restricted model, a mixed mitigation strategy of partial sourcing from supplier R and carrying inventory cannot be optimal for these cases. Demand is deterministic in the restricted model, and this raises the possibility that a mixed mitigation strategy might be optimal if demand is stochastic. Because we already know that partial sourcing can be optimal for a risk-averse firm, we restrict attention to a risk-neutral firm to focus attention on the impact of demand uncertainty on the optimal sourcing (and hence disruption-management) strategy.

THEOREM 7 (II-FLEXIBILITY). (a) *For a given supplier R allocation w, a state-dependent base-stock policy is optimal. The optimal base-stock levels are* $y^*_{II}(i, w) = (1 - w)y^*(i)$, *where* $y^*(i)$ *is the same as in Lemma* 1. *The* $y^*_{II}(i, w)$ *are nonincreasing in i.* (b) *Single-sourcing is optimal, that is,* $w^* \in \{0, 1\}$.

A zero-flexibility system is equivalent to an II-flexibility system with $c_f = \infty$. This observations leads directly to the following theorem.

THEOREM 8 (ZERO FLEXIBILITY). (a) *For a given supplier R allocation w, a base-stock policy is optimal for supplier U orders when it is up. The optimal base-stock level is* $y^*_Z(0, w) = (1 - w)y^*_\infty(0)$, *where* $y^*_\infty(0)$ *is the base-stock level in Lemma* 1 *if* $c_f = \infty$. (b) *Single-sourcing is optimal, that is,* $w^* \in \{0, 1\}$.

The single-sourcing result for the zero-flexibility and II-flexibility cases therefore still holds even if demand is stochastic. The fact that the single-sourcing result is preserved under stochastic demand means that a mixed disruption-management strategy (partial sourcing from supplier R and carrying inventory) cannot be optimal for these two extreme flexibility cases in the stochastic-demand case. Note that if $L > 0$, then the firm might hold inventory even if it single-sources from supplier R. Such inventory would, however, be held for the sole purpose of protecting against demand uncertainty as there is no supply uncertainty if the firm single-sources from supplier R.

In closing we note that the optimal supplier choice is not necessarily the same for the zero-flexibility and II-flexibility cases. Define w^*_{II} and w^*_Z as the optimal supplier R allocation for the II-flexibility and zero-flexibility cases respectively.

COROLLARY 1. $w^*_Z = 0 \Rightarrow w^*_{II} = 0$. $w^*_{II} = 1 \Rightarrow w^*_Z = 1$.

So, if single-sourcing from supplier U is optimal for the zero-flexibility case, then it is optimal for the II-flexibility case. Likewise, if single-sourcing from supplier R is optimal for the II-flexibility case, then it is optimal for the zero-flexibility case.

5.3. Finite Capacity at Supplier U

Supplier U capacity is assumed to be infinite in the restricted model, with the consequence that supplier U can immediately recover any lost production once a disruption ends. This in turn means that the firm is able to fully replenish its mitigation inventory before any new disruption can occur. The results of the restricted model all go through if supplier U has sufficient capacity to immediately recover any lost production. Immediate recovery is, however, a strong assumption, and in this section we investigate the impact of supplier U's capacity ν_u on the firm's optimal disruption-management strategy. We focus attention on the zero-flexibility case, for which single-sourcing was proven to be optimal when supplier U had infinite capacity. We make the following assumptions:

- Risk neutrality in the allocation decision, as we know that risk aversion can make dual-sourcing optimal.
- Zero fixed cost of ordering, and so the base time unit is infinitesimal (continuous-time model).
- Constant demand rate d ($d = 1$ without loss of generality), as we know that stochastic demand does not alter the single-sourcing result for the zero-flexibility case.
- Exponentially distributed uptimes, with rate λ_{ud} (up to down).
- Exponentially distributed downtimes, with rate λ_{du} (down to up).

The allocation w is the order rate placed with supplier R. Supplier R provides a capacity ν_r so that it can produce at a rate of w and no higher (because there is zero flexibility). As supplier R is perfectly reliable, we are then left with a single unreliable-supplier system with demand rate of $1-w$. The long-run average cost is then given by

$$C_{LRA}(w) = c_r w + c_u(1-w) + J^*_{\nu_u}(w), \tag{9}$$

where $J^*_{\nu_u}(w)$ is the optimal long-run average inventory/back-order cost for a single-supplier system with demand $1-w$ and capacity ν_u.

Theorem 9. *The firm must source a minimum of*

$$w \geq \left[1 - \nu_u\left(\frac{\lambda_{ud}}{\lambda_{du}+\lambda_{ud}}\right)\right]^+$$

from supplier R for the system to be stable. The optimal ordering policy is a modified base-stock policy. If supplier U is up, then the optimal order rate for supplier U is

$$q_t^u = \begin{cases} 1-w, & \text{if } x_t = y^*(0, w) \\ \nu_u, & \text{if } x_t < y^*(0, w), \end{cases}$$

where x_t is the inventory at time t. The optimal base-stock level $y^*(0, w)$ is given by

$$y^*(0,w) = \begin{cases} 0, & \text{if } \dfrac{\lambda_{ud}}{\lambda_{du}+\lambda_{ud}}\left(1+\dfrac{p}{h}\right)\dfrac{\nu_u}{\nu_u-(1-w)} \leq 1 \\[2ex] \dfrac{\log\left(\dfrac{\lambda_{ud}}{\lambda_{du}+\lambda_{ud}}\left(1+\dfrac{p}{h}\right)\dfrac{\nu_u}{\nu_u-(1-w)}\right)}{\dfrac{\lambda_{du}}{1-w}-\dfrac{\lambda_{ud}}{\nu_u-(1-w)}}, & \text{otherwise.} \end{cases}$$

The optimal inventory/back-order cost is given by

$$J^*_{\nu_u}(w) = \begin{cases} \dfrac{\lambda_{ud}}{\lambda_{du}+\lambda_{ud}}\dfrac{pv(1-w)}{\lambda_{du}\nu_u-(\lambda_{du}+\lambda_{ud})(1-w)}, & \text{if } \dfrac{\lambda_{ud}}{\lambda_{du}+\lambda_{ud}}\left(1+\dfrac{p}{h}\right)\dfrac{\nu_u}{\nu_u-(1-w)} \leq 1 \\[2ex] \dfrac{h(1-w)}{\lambda_{du}+\lambda_{ud}} + \dfrac{h\log\left(\dfrac{\lambda_{ud}}{\lambda_{du}+\lambda_{ud}}\left(1+\dfrac{p}{h}\right)\dfrac{\nu_u}{\nu_u-(1-w)}\right)}{\dfrac{\lambda_{du}}{1-w}-\dfrac{\lambda_{ud}}{\nu_u-(1-w)}}, & \text{otherwise.} \end{cases}$$

The optimal allocation w^* is nonincreasing in ν_u.

We proved above that single-sourcing is optimal in the zero-flexibility case when supplier U has infinite capacity. This result is no longer true when supplier U has limited capacity. We observed in numerical tests that the optimal allocation, and hence the optimal disruption-management strategy, was sensitive to the supplier U's capacity. In Figure 8 we present the optimal disruption-management strategies for both the infinite-capacity and finite-capacity cases as a function of percentage uptime and expected disruption length for the following set of parameters: $c_u = 1$, $c_r = 1.1$, $p = 0.15$, $h = 0.0015$, $\lambda_{du} = 0.1$, and $\nu_u = 1.5$. We use the same labeling scheme as before, but there is now an additional strategy, labeled MPS, which is mitigation through partial sourcing from supplier R. In this strategy, the firm sources from both suppliers if there is no disruption but does not carry inventory to mitigate disruptions. There is a large region in which IM is the optimal strategy for the infinite-capacity case. In contrast, mitigation through partial sourcing from supplier R and carrying inventory (MPSI) is the dominant

Figure 8 Optimal Disruption Strategies for Infinite-Capacity and Finite-Capacity Cases

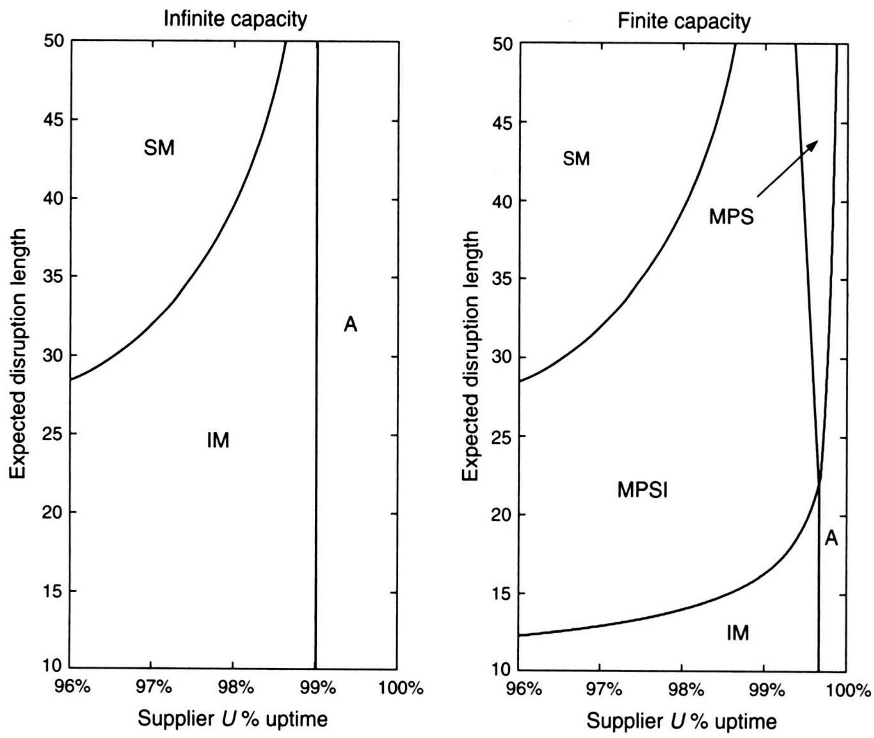

strategy in the finite-capacity case. Capacity influences a supplier's ability to recover from a disruption, and a supplier's ability to recover has a large impact on the firm's optimal disruption-management strategy. We observed that the fraction sourced from supplier *R* increases in the expected disruption length for a fixed supplier *U's* percentage uptime. This reflects the fact that recovery takes longer when disruptions are longer. We also observed that the fraction sourced from supplier *R* decreases in supplier *U* percentage uptime for a fixed expected disruption length. This reflects the fact that supplier *U* spends less time in recovery mode as its percentage uptime increases.

6. Conclusions

An effective disruption-management strategy is a necessary component of a firm's overall supply chain strategy. Firms that passively accept the risk of disruptions leave themselves open to the danger of severe financial and market-share loss, as evidenced by the Philips Semiconductor and Hurricane Mitch disruptions discussed in the introduction. Active disruption-management strategies rely on mitigation and/or contingency actions. In this paper, we have focused on the supply-side tactics available to a firm: sourcing mitigation, inventory mitigation, and contingent rerouting. We established that, along with cost, supplier characteristics such as percentage uptime, disruption length, capacity, and flexibility, and firm characteristics such as risk tolerance, play a large role in determining the firm's optimal disruption-management strategy. Percentage uptime and disruption length influence the optimal disruption-management strategy through their impact on the frequency and level at which mitigation and contingency costs are incurred. Capacity plays an important role through its effect on a supplier's recovery time in the aftermath of a disruption. Volume flexibility can substantially benefit the firm, as it enables contingent rerouting to be an element of the firm's strategy, and this can significantly reduce the firm's costs. We showed that inventory mitigation was not an attractive strategy in an environment of rare but long disruptions, as significant quantities of inventory need to be carried for extended periods without a disruption. This result is at least partly driven by the assumption of a constant probability of failure. In certain circumstances, for example labor disputes, a firm may have advance warning that a disruption is more likely. Such advance information then allows the firm to build mitigation inventory in advance of a potential disruption rather than carrying mitigation inventory continuously. The role of advance information in disruption management is the focus of ongoing research.

In closing, we note that the operations literature has devoted significantly more attention to mix flexibility than to volume flexibility. Interestingly, a recent survey revealed that "more than 50% [of respondents] identified volume flexibility within supply chain management and operations as the key area for improvement in 2002/2003" (Sheppard and Kent 2002, p. 40).

Volume flexibility provides an alternative to inventory in managing temporary imbalances in supply and demand, which can arise because of supply-side disruptions or temporary shifts in demand. We believe that our model of volume flexibility, parameterized by magnitude and response time, provides a foundation for future research into the benefits of volume flexibility in contexts other than disruption management.

An online supplement to this paper is available on the *Management Science* website (http://mansci.pubs.informs.org/ecompanion.html).

Acknowledgments

This research was supported in part by an award from the Frank H. Kenan Institute of Private Enterprise. The author is grateful for the invaluable comments and suggestions provided by the referees, the associate editor, and the department editor on earlier versions of this paper.

References

Agrawal, N., S. Nahmias. 1997. Rationalization of the supplier base in the presence of yield uncertainty. *Production Oper. Management* **6** 291–308.

Agrawal, V., S. Seshadri. 2000. Impact of uncertainty and risk aversion on price and order quantity in the newsvendor problem. *Manufacturing Service Oper. Management* **2** 410–423.

Anupindi, R., R. Akella. 1993. Diversification under supply uncertainty. *Management Sci.* **39** 944–963.

Arreola-Risa, A., G. A. De Croix. 1998. Inventory management under random disruptions and partial back-orders. *Naval Res. Logist.* **45** 687–703.

Barnes-Schuster, D., Y. Bassok, R. Anupindi. 2002. Coordination and flexibility in supply contracts with options. *Manufacturing Service Oper. Management* **4** 171–207.

Bielecki, T., P. R. Kumar. 1988. Optimality of zero-inventory policies for unreliable manufacturing systems. *Oper. Res.* **36** 532–541.

Bollapragada, R., U. S. Rao, J. Zhang. 2004a. Managing two-stage serial inventory systems under demand and supply uncertainty and customer service level requirements. *IIE Trans.* **36** 73–85.

Bollapragada, R., U. S. Rao, J. Zhang. 2004b. Managing inventory and supply performance in assembly systems with random capacity and demand. *Management Sci.* **50** 1729–1743.

Burrows, P. 1999. The vice trapping PC makers. *Business Week Online* (October 25).

Caldentey, R., M. Haugh. 2004. Optimal control and hedging of operations in the presence of financial markets. *Math. Oper. Res.* Forthcoming.

Ciarallo, F. W., R. Akella, T. E. Morton. 1994. A periodic review, production planning model with uncertain capacity and uncertain demand—optimality of extended myopic policies. *Management Sci.* **40** 320–332.

Dada, M., N. C. Petruzzi, L. B. Schwarz. 2003. A newsvendor model with unreliable suppliers. Working paper, Krannert Graduate School of Management, Purdue University, West Lafayette, IN.

Eeckhoudt, L., C. Gollier, H. Schlesinger. 1995. The risk-averse (and prudent) newsboy. *Management Sci.* **41** 786–794.

Feng, Q., S. P. Sethi, H. Yan, H. Zhang. 2005. Are base-stock policies optimal in inventory problems with multiple delivery modes? *Oper. Res.* Forthcoming.

Fine, C. H., R. M. Freund. 1990. Optimal investment in product-flexible manufacturing capacity. *Management Sci.* **36** 449–466.

Fukuda, Y. 1964. Optimal policies for the inventory problem with negotiable lead time. *Management Sci.* **10** 690–708.

Gerwin, D. 1993. Manufacturing flexibility: A strategic perspective. *Management Sci.* **39** 395–410.

Gerchak, Y., M. Parlar. 1990. Yield randomness, cost trade-offs and diversification in the EOQ model. *Naval Res. Logist.* **37** 341–354.

Graves, S. C. 1988. Safety stocks in manufacturing systems. *J. Manufacturing Oper. Management* **1** 67–101.

Graves, S. C. 1999. A single-item inventory model for a nonstationary demand process. *Manufacturing Service Oper. Management* **1** 50–61.

Graves, S. C., B. T. Tomlin. 2003. Process flexibility in supply chains. *Management Sci.* **49** 907–919.

Griffy-Brown, C. 2003. Just-in-time to just-in-case. *Graziadio Bus. Rep.* **6**.

Gupta, D. 1996. The (Q, r) inventory system with an unreliable supplier. *INFOR* **34** 59–76.

Gürler, Ü., M. Parlar. 1997. An inventory problem with two randomly available suppliers. *Oper. Res.* **45** 904–918.

Gurnani, H., R. Akella, J. Lehoczky. 2000. Supply management in assembly systems with random yield and random demand. *IIE Trans.* **32** 701–714.

Jordan, W. C., S. C. Graves. 1995. Principles on the benefits of manufacturing flexibility. *Management Sci.* **41** 577–594.

Kalymon, B. 1971. Stochastic prices in a single-item inventory purchasing model. *Oper. Res.* **19** 1434–1458.

Khang, D. B., O. Fujiwara. 2000. Optimality of myopic ordering policies for inventory model with stochastic supply. *Oper. Res.* **48** 181–184.

Kouvelis, P., J. M. Milner. 2002. Supply chain capacity and outsourcing decisions: The dynamic interplay of demand and supply uncertainty. *IIE Trans.* **34** 717–728.

Kouvelis, P., G. Vairaktarakis. 1998. Flowshops with processing flexibility across production stages. *IIE Trans.* **30** 735–746.

Latour, A. 2001. Trial by fire: A blaze in Albuquerque sets off major crisis for cell-phone giants. *Wall Street Journal* (January 29).

Meyer, R. R., M. H. Rothkopf, S. A. Smith. 1979. Reliability and inventory in a production-storage system. *Management Sci.* **25** 799–807.

Moinzadeh, K., P. Aggarwal. 1997. Analysis of a production/inventory system subject to random disruptions. *Management Sci.* **43** 1577–1588.

Moinzadeh, K., S. Nahmias. 1988. A continuous review model for an inventory system with two supply modes. *Management Sci.* **34** 761–773.

Parlar, M. 1997. Continuous review inventory problem with random supply interruptions. *Eur. J. Oper. Res.* **99** 366–385.

Parlar, M., D. Berkin. 1991. Future supply uncertainty in EOQ models. *Naval Res. Logist.* **38** 107–121.

Parlar, M., D. Perry. 1995. Analysis of a (Q, r, T) inventory policy with deterministic and random yields when future demand is uncertain. *Eur. J. Oper. Res.* **84** 431–443.

Parlar, M., D. Perry. 1996. Inventory models of future supply uncertainty with single and multiple suppliers. *Naval Res. Logist.* **43** 191–210.

Parlar, M., D. Wang. 1993. Diversification under yield randomness in inventory models. *Eur. J. Oper. Res.* **66** 52–64.

Parlar, M., Y. Wang, Y. Gerchak. 1995. A periodic review inventory model with Markovian supply availability. *Internat. J. Production Econom.* **42** 131–136.

Poirier, C. C., F. J. Quinn. 2004. How are we doing? A survey of supply chain progress. *Supply Chain Management Rev.* **8** 24–31.

Scheller-Wolf, A., S. Tayur. 1998. A Markovian dual-source production-inventory model with order bands. Working paper #1998-E200, GSIA, Carnegie Mellon University, Pittsburgh, PA.

Schweitzer, M. E., G. P. Cachon. 2000. Decision bias in the newsvendor problem with a known demand distribution: Experimental evidence. *Management Sci.* **46** 404–420.

Sethi, A. K., S. P. Sethi. 1990. Flexibility in manufacturing systems: A survey. *Internat. J. Flexible Manufacturing Systems* **2** 289–328.

Sethi, S., H. Yan, H. Zhang. 2003. Inventory models with fixed costs, multiple delivery modes and forecast updates. *Oper. Res.* **51** 321–328.

Sheppard, I., D. Kent. 2002. Executive interviews reveal weak spots in supply chain management. *Semiconductor Fabtech* **16** 37–40.

Song, J.-S., P. H. Zipkin. 1993. Inventory control in a fluctuating demand environment. *Oper. Res.* **41** 351–370.

Song, J.-S., P. H. Zipkin. 1996. Inventory control with information about supply conditions. *Management Sci.* **42** 1411–1419.

Suarez, F. F., M. A. Cusumano, C. H. Fine. 1995. An empirical study of flexibility in manufacturing. *Sloan Management Rev.* **37** 25–32.

Swaminathan, J. M., J. G. Shanthikumar. 1999. Supplier diversification: Effect of discrete demand. *Oper. Res. Lett.* **24** 213–221.

Tomlin, B., Y. Wang. 2005. On the value of mix flexibility and dual sourcing in unreliable newsvendor networks. *Manufacturing Service Oper. Management* **7** 37–57.

Treece, J. B. 1997. Just-too-much single sourcing spurs Toyota purchasing review. *Automotive News* (March 3).

Tsay, A. A., W. S. Lovejoy. 1999. Quantity flexible contracts and supply chain performance. *Manufacturing Service Oper. Management* **1**(2) 89–111.

Van Mieghem, J. A. 1998. Investment strategies for flexible resources. *Management Sci.* **44** 1071–1078.

Yano, C. A., H. L. Lee. 1995. Lot sizing with random yield: A review. *Oper. Res.* **43** 311–334.