Submitted to *Manufacturing & Service Operations Management*
manuscript MS-0000-0000.00

A General Model for Inventory Management with Dual Sources: Trading off Lead Time and Cost Differences

Awi Federgruen, Zhe Liu
Graduate School of Business, Columbia University, New York, NY 10027, af7@columbia.edu, zliu18@gsb.columbia.edu

Lijian Lu
China Innovation Fund, Beijing, China, ll2755@columbia.edu

Problem definition: We study a finite horizon, single product, periodic review inventory system with two supply sources and salvage options. A challenging trade-off exists between the two sources because the expedited supplier has a shorter lead time but charges a higher per-unit price. A further complication is the salvage option that allows for bilateral inventory adjustments. All inventory adjustments involve a fixed cost component in addition to variable costs or revenues and may be subject to capacity limits.

Academic/practical relevance: These strategic dilemmas arise whenever a company decides on off-shoring versus on-shoring, or a hybrid approach which utilizes both suppliers simultaneously.

Methodology: We analyze our model with dynamic programming and Markov Decision Process techniques.

Results: We show that an optimal policy first determines the size of an order with the expedited supplier, or the size of any salvage quantity, based, exclusively, on the regular full inventory position. Thereafter, any order with the regular supplier is determined as a function of the *adjusted* inventory position. Moreover, the dependence of the optimal order sizes and/or salvage quantity, on the period's starting inventory position follows a relatively simple structure. The above results apply to the special case where the lead times of the two suppliers differ by a single period. However, our structural results suggest effective heuristics for general lead time combinations.

Managerial implications: We evaluate our heuristics in a comprehensive numerical study. We focus on how various model primitives impact on the benefits of dual sourcing. For example, we show that the cost savings due to dual sourcing may be as large as 30%. For a given price differential, the benefits of dual sourcing grow as the lead time difference increases.

Key words: dual sourcing, capacity constraints, lead time, inventory management, convexity, dynamic programming

1. Introduction

Manufacturing companies and retail chains often have access to two alternative supply sources for component parts, product modules, finished goods or supply materials. One source is typically low cost but has long lead times, whereas the other provides quicker response but at a higher price. When designing its procurement process, the purchaser may select one of the two sources as its exclusive supplier. Alternatively, it may opt for a dual sourcing strategy which procures from both sources. In the latter case, the challenge is to determine how and when each of the sources is to be used, as a function of the dynamically evolving inventory information. The objective of this paper is to address this question within a general model that incorporates economies of scale with respect to the order costs, capacity limits for individual orders and opportunities to reduce inventory via salvage sales.

The above strategic dilemmas arise, first and foremost, when firms decide on offshoring vs onshoring options. In the past few decades, there has been a consistent trend to offshore. In 2010, PricewaterhouseCoopers (2010) reported that the share of products that are sourced from offshore locations ranges from 51% to 75%, in different European and American countries. At that point, the share was projected to grow more than 10%, annually, in the next five years. The survey was conducted based on interviews with 59 CEO in global retail and consumer companies in eight countries: Austria, Canada, China, France, Germany, India, the United Kingdom and the United States.

However, the trend has recently been reversed, as companies have come to realize that, along with other sourcing considerations, price savings associated with offshore options need to be traded off against increased inventory costs and stockout risks due to the larger lead times involved. Longer lead times translate into a need for larger safety stocks, under a given targeted service level, or inferior service levels under given inventory investments. In contrast, onshore procurement from a local or nearby market is fast but typically incurs a higher purchase price or manufacturing cost. In 2013, The Economist (2013) reported that for the time interval between 2012 and 2014, 19% of the surveyed companies declared an intention to reshore, almost the same percentage (23%) which intended to replace local suppliers by offshore alternatives. Another 33% intended to move between low- and high-cost countries, again revisiting their supplier base, see Figure 1 below.

Moreover, many companies have come to realize that a hybrid approach employing two or more suppliers, simultaneously, is, frequently, considerably more effective than one which relies on a *single* supplier (single sourcing). Scheller-Wolf et al. (2007), for example, compare the optimal procurement strategy under single sourcing with two heuristic policies that employ dual sourcing. The comparisons are conducted within a base model, i.e., a periodic-review, infinite horizon model with i.i.d. demands, in which uncapacitated orders can be placed with either supplier and order costs are linear, without any fixed costs. In their numerical study, the authors observed that the dual sourcing heuristics, on average, reduce costs by 9%, while, in some instances, the cost savings

Figure 1 The Economist's report on companies' outsourcing decisions

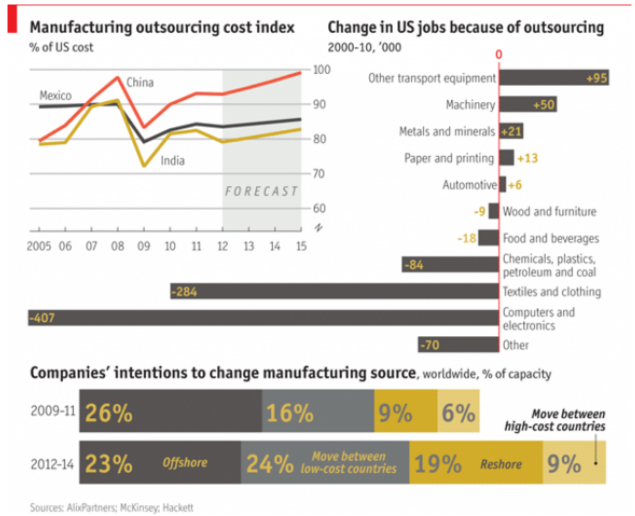

are as large as 25%! Moreover, the dual sourcing heuristics are guaranteed to outperform the optimal procurement strategy under single sourcing with either one of the two suppliers.

Within the operations management/research community, several authors have reported on specific company settings where dual sourcing is employed. These include Beyer and Ward (2002) reporting on Hewlett-Packard's strategy for manufacturing servers and Rao et al. (2000) on Caterpillar's for compact worktools sold in the North American market. Allon and Van Mieghem (2010a) report on a \$10 billion high-tech US manufacturer of wireless transmission components with two assembly plants, one in China and one in Mexico. The company implemented a dual sourcing strategy. The authors identify a (different) heuristic dual sourcing strategy and report that in their application the strategy saves up to 20% over the best single sourcing strategy. Based on this company setting, Allon and Van Mieghem (2010b) developed a teaching game, referred to as the "Mexico-China sourcing game", see also Van Mieghem and Allon (2015).

The study of periodic review, dual sourcing inventory models starts with four papers in the early sixties, shortly after the initiation of single source stochastic inventory models. These early papers addressed the above-mentioned periodic review model with independent demands, full backlogging of stockouts, and two suppliers with different lead times and different per-unit procurement prices. The "expedited" supplier charges the higher per-unit price, giving rise to a challenging trade-off problem. Other than the linear order costs, there are holding and backlogging costs assumed to be proportional with or convex in the end-of-the-period inventory and backlog levels, respectively.

The seminal paper by Fukuda (1964) showed that a so-called *single index dual base stock* policy is optimal in this model, as long as the lead time of the slower supplier is exactly *one* period longer than that of the expedited supplier. Under this policy, one starts by determining the size of the order to be placed with the *expedited* supplier, if any. This order is determined by a base-stock policy acting on the (full) inventory position = the inventory level plus all outstanding orders. In other words, an order is placed iff this inventory position is below a given base stock level; when applicable, it is sized to elevate the inventory position to this base stock level. After the order with the expedited supplier is added to this inventory position measure, a second base-stock policy is applied to determine the order size with the slower supplier (if any).

In general, lead times may, of course, differ by an *arbitrary* number of periods. For this general setting, it was shown, e.g., Whittemore and Saunders (1977), that no procurement strategy, acting on a single or even two inventory measures (so-called indices) needs to be optimal. As a consequence, many *heuristic* policies were proposed, mostly in the past decade, to handle the general lead time case. These are reviewed in Section 2.

However, to our knowledge, little progress was made to include important generalizations and complicating factors that arise in practice and have become part of standard single sourcing inventory models, by themselves or in various combinations; in particular:

(a) fixed order costs, and

(b) capacity limits associated with the orders with the two suppliers.

An additional generalization is

(c) the ability to *decrease* the inventory level by salvaging a given quantity through sales in a secondary channel (jobbers, discounters, outlet stores, etc).

As mentioned, optimal strategies that act on a single (or even a pair of) inventory indices, can only be expected when the lead times associated with the two suppliers differ by a single period. We, therefore, initially, confine ourselves to this case and show that a single index policy continues to be optimal for this very general model that incorporates the three complications (a)–(c), simultaneously. (The remaining model assumptions are standard and identical to those employed in the above-mentioned literature on dual sourcing models: full backlogging of stockouts, demands that are independent across time and holding and backlogging costs that are specified by convex functions of the end-of-the-period inventory levels.) In each period, one first determines the size of an order with the expedited supplier, if any, or the size of any salvage quantity, based, exclusively, on the regular full inventory position. Thereafter, the inventory position is adjusted upward (by the expedited supplier order) or downward (by the salvage quantity); any order with the regular supplier is then determined as a function of the *adjusted* inventory position.

Moreover, the dependence of the optimal order sizes and/or salvage quantity, on the period's starting inventory position follows a relatively simple structure. In the most general case, the optimal policy is characterized by fours critical threshold levels of the inventory position: $b^e \leq \bar{b}^e \leq \underline{s}^e \leq s^e$ partitioning the inventory position line into 5 consecutive regions. In the middle range $[\bar{b}^e, \underline{s}^e]$, it is optimal to forgo both an order with the expedited supplier, as well as any salvage activity. In the far-left (-right) region $(-\infty, b^e)$ $[(s^e, \infty)]$, it is optimal to place an order with the expedited supplier [to initiate a salvage sale], the size of which varies as a non-linear function of the inventory position. This leaves us with the remaining pair of intervals $[b^e, \bar{b}^e)$ and $(\underline{s}^e, s^e]$. In the former, i.e., when $b^e < \bar{b}^e$, it is possible that one alternates between sub-intervals with a positive order and those where it is optimal to stay put; however, salvaging is not to be considered in this interval $[b^e, \bar{b}^e)$. Similarly, when $\underline{s}^e < s^e$, one alternates on $(\underline{s}^e, s^e]$ between sub-intervals where it is optimal to stay put and those where it is optimal to initiate a salvage batch; however, ordering does not need to be considered.

As far as the second stage ordering decision with the *regular* supplier is concerned, the optimal policy is characterized by two threshold parameters, $b^r, \bar{b}^r$ partitioning the *adjusted inventory position* line in up to three regions. The structure of the optimal strategy within these regions is identical to those of the first stage decision in the three left most intervals mentioned there; see Figure 2.

Figure 2 **General optimal policy structure**

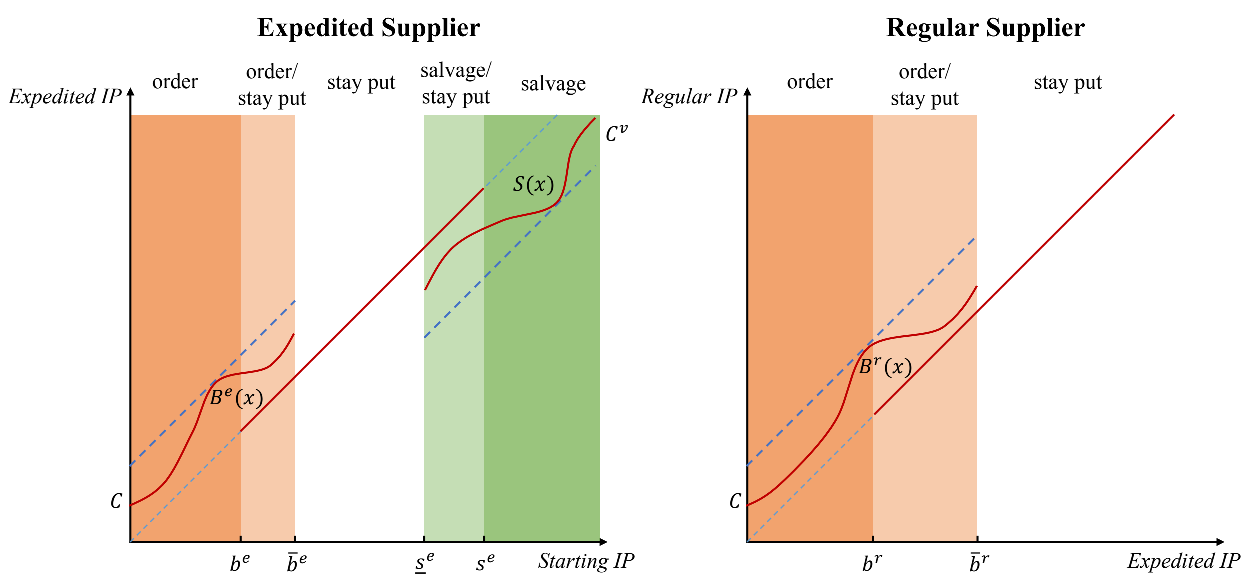

Even simpler and more pronounced structures arise in various special cases; where only a subset of the complications (a)–(c) prevails, or where some model parameters take on a specific value, for example, when either orders or salvage sales can be initiated without a fixed cost. These special cases are characterized in Section 5.

For the case of general lead time combinations, we propose, as effective heuristics, single index policies of the same structure—but alternative policy parameters—as those shown to be optimal in the special case where the lead time difference is a single period. We illustrate this for several important cases, in particular, dual sourcing problems without salvage opportunities but with capacity limits or fixed costs for the orders.

Our structural results are obtained by showing that the value functions in a given dynamic programming formulation, satisfy the so-called (C_1K_1, C_2K_2)-convexity property, a generalization of the far better known convexity or K-convexity property, as well as many other generalized convexity properties employed in the *single sourcing* literature to handle specific subsets of the complications (a)–(c), above. We rely heavily on the analysis of the (C_1K_1, C_2K_2)-convexity property, conducted in Federgruen et al. (2017) to characterize the structure of an optimal strategy in the *single* sourcing model with the simultaneous presence of complications (a)–(c).

The remainder of this paper is organized as follows. In Section 2 we give a review of the literature on multi-sourcing procurement strategies. Section 3 introduces the model. In Section 4, we derive the structure of the optimal dual sourcing policy in the *single-period problem*, assuming the lead time difference equals one. Section 5 extends these results for a general finite (or infinite) planning horizon, both in the most general and in various special cases. Based on the structural results in Section 5 we develop, in Section 6, our proposed heuristics for the case of general lead time combinations, along with an illustrative numerical study. Section 7 contains concluding remarks.

2. Literature Review

In this section, we review the literature on dual sourcing stochastic inventory models. As mentioned in the Introduction, the 60-year-old literature has mostly focused on a base periodic review model with independent demands, full backlogging of stockouts, linear order, holding and backlogging costs, and *two* potential uncapacitated suppliers, differentiated by their per-unit procurement price and lead times. This workhorse model may be viewed as the direct extension of the *seminal* single source model by Arrow et al. (1951).

Three papers in the early sixties, i.e., Barankin (1961), Daniel (1963) and Neuts (1964) focused on the special case where the lead time of the expedited supplier is negligible (or zero), and that of the regular supplier exactly *one* period. Here, the optimal strategy employs *two* base-stock levels: first, the period's starting inventory level—including last period's regular order—is used by a simple base-stock policy to determine whether an order is to be placed with the expedited supplier. Thereafter, a second base-stock policy is used to determine the order with the regular supplier, if any, comparing the adjusted inventory level (inventory level plus the new order with the expedited supplier) with a second base-stock level. Fukuda (1964) extended this result to the case where the lead time of the regular supplier is arbitrary and that of the expedited supplier *one* period shorter. The same structural result continues to apply, except that the base-stock policies act on the period's starting *inventory position*, as opposed to the inventory level.

Whittemore and Saunders (1977) showed that no simple structure prevails for the general case with *non-consecutive* lead times, i.e., lead times that differ by more than a single period. In particular, the optimal policy needs to be based on more than one, or even any constant number of inventory measures, or indices.

With this negative insight, the development of dual sourcing models was interrupted for some 30 years. In the last decade, we have seen a plethora of papers suggesting and comparing *heuristic* policies for the infinite horizon, stationary base model with general lead times. This literature stream started with Veeraraghavan and Scheller-Wolf (2008) proposing the use of *dual index* base-stock polices. Here, the first stage order to the expedited supplier is determined by a base-stock policy acting on the so-called *expedited* inventory position, consisting of the inventory level plus all outstanding orders to arrive within the expedited supplier's lead time from the current period. The second stage order with the regular supplier is determined by a second base-stock policy acting on the *full inventory position*, the current period's order with the expedited supplier included. These authors compared the cost performance of the best policy within their proposed class, with the overall optimal policy. To this end, they conducted a numerical study assuming the demand distribution has support on a few values only., so that the set of possible inventory levels can be limited and the optimal policy can be found via dynamic programming, at least for small lead time values. The authors report that their proposed dual index policy, in the majority of cases, comes within 1% or 2% of optimality.

The authors also consider capacity limits for the two suppliers. In this case, base stock policies need to be replaced by "*modified base-stock policies*", where an order is placed to bring the relative inventory index as close as possible to the base-stock level. In the single sourcing literature, such policies were shown to be optimal, see e.g. Federgruen and Zipkin (1986a,b).

Scheller-Wolf et al. (2007) show that a *single-index* base-stock policy performs comparably or even better than the *dual index* base-stock policy in Veeraraghavan and Scheller-Wolf (2008). Under a single index base-stock policy, the order with the expedited supplier is determined by a base-stock policy acting on the full inventory position, which includes all outstanding orders, rather than the more limited "expedited" inventory position in Veeraraghavan and Scheller-Wolf (2008). In Scheller-Wolf et al. (2007)'s numerical study, the average difference between the cost values of the two heuristics is less than 0.5% and the maximum difference no more than 3%. While the cost performance is almost identical between the two heuristics, the authors point at several major advantages of the single index heuristic, including the fact that it is simpler and that the optimal base-stock levels can be computed *analytically*, at "25–60 times faster" computational times.

Sheopuri et al. (2010) develop six alternative heuristic polices, some of considerably more complex structure (for example, the vector-based base-stock policy and the "best weighted" index policy). Their numerical study shows that each of the heuristics has an average cost performance that comes within 1% of that of the dual index base-stock policy, and hence of the single index base-stock policy, as well. These authors also show that the single source *lost-sales* model can be viewed as a special case as the dual source model with backlogging. Several of the structural properties identified by Zipkin (2008) for the lost sales model can therefore be generalized in the study of dual sourcing problems. In the same spirit as Sheopuri et al. (2010), Hua et al. (2015) developed another class of heuristic policies, again with comparable cost performance, in their numerical study.

Allon and Van Mieghem (2010a) proposed a Tailored Based-Surge (TBS) heuristic in which a constant size order is placed with the regular supplier and a base-stock policy is used for the expedited supplier. Janakiraman et al. (2015) analyzed this class of policies, deriving optimality gap bounds when the demands consist of a regular base demand plus an infrequent surge demand. In a numerical study, similar to that in Veeraraghavan and Scheller-Wolf (2008), they observed that the optimality gap of the best TBS policy varies between 21% and 3.5%. Xin and Goldberg (2017) showed that a TBS policy is asymptotically optimal when the lead time difference goes to infinity. Xin et al. (2017) extend the results to a setting where the expedited supplier is unreliable, i.e., in any period the supplier is with a given probability unable to fill any order. The authors evaluate the performance of the TBS heuristic based on Walmart data. As mentioned in the Introduction, Allon and Van Mieghem (2010b) developed the Mexico-China teaching game, based on the industrial application which motivated their parallel paper. See also Chapter 7 in the prominent textbook on Operations Strategy by Van Mieghem and Allon (2015) for a treatment of the dual sourcing problem.

Sun and Van Mieghem (2017) developed a robust optimization approach for the base dual sourcing problem. Under this approach, the problem can be formulated as a mathematical program, avoiding the curse of dimensionality associated with standard dynamic programming formulations that aim at optimizing aggregate expected costs.

Very few papers address generalizations of the base model. A noted exception is Sethi et al. (2003) who incorporates fixed order costs into the model with the one period lead time difference. They show that, in each period, the order for the expedited supplier and that for the regular supplier are to be determined sequentially, both on the basis of an (s, S)-policy acting on the regular inventory position. We retrieve this result as one of the special cases in our general model. Fox et al. (2006) also allow for fixed order costs but assume zero lead time for the both suppliers. It is easily seen that this model reduces to a single sourcing problem with a piecewise linear concave order cost. Boute and Van Mieghem (2015) add capacity costs and order smoothing considerations into the base model.

A few papers have addressed dual sourcing problems in *continuous review* models. Moinzadeh and Schmidt (1991) and Song and Zipkin (2009) assumed demands are generated by a Poisson process and a dual-index base-stock policy is applied. The former paper assumes the lead times are deterministic while the latter allows for stochastic but exogenous lead times. Moinzadeh and Schmidt (1991) characterize the steady-state inventory level distribution and hence the long-run average cost for any (dual-index) base-stock policy. Song and Zipkin (2009) models the system as a queueing network with overflow by-passes and obtain the steady-state distribution of the network in product form. Very recently, Zhou and Yang (2016) extended their work by allowing for compound Poisson demands and fixed order costs. These authors propose, as heuristic, a *single index* pair of (R, nQ)-policies. At each demand epoch, one places an order with the expedited supplier iff the full inventory position is below a reorder level R^e; the order is sized as the smallest multiple of an order size Q^e which elevates the inventory position above R^e. With this order added to the inventory position, any order with the regular supplier is determined by a different (R^r, nQ^r)-policy.

All of the above continuous review papers make an upfront restriction to a specific class of heuristic policies. However, the recent paper by Song et al. (2017), characterizes the optimal policies in a setting with linear costs but *endogenously* determined stochastic lead times, under a certain condition, and propose near-optimal heuristic policies when the condition fails.

3. Model Formulation

We consider a single-item periodic review inventory system with two potential suppliers: a regular and a expedited supplier. The regular supplier produces the item at a lower (per-unit) cost, but a longer lead time is involved. We thus assume that the suppliers are differentiated in terms of their lead times and variable prices(, only). In particular, if orders are associated with fixed costs or capacity limits, we assume these to be the same for both suppliers. In each period, there may also be a salvage option to reduce inventory, generating a given per-unit revenue, possibly associated with a fixed cost and capacity limit. The lead time is l_e both for ordering from the expedited supplier and salvaging (when applicable), and $l_r > l_e$ for ordering from the regular supplier. In this section we focus on the case where $l_r = l_e + 1$. As explained in the Introduction, an optimal policy of a relatively simple structure can only be expected for this case. However, the structural results obtained for this special case suggest immediate extrapolations for the general model where $l_r \geq l_e + 2$, as heuristic policies. In Section 6, we design these heuristics and evaluate their performance.

We index the periods *backward* from 1 to N. The sequence of events in period n is as follows: At the beginning of period n, an order may arrive from either the expedited or the regular supplier, or both. Such deliveries are added to the inventory level. As in single sourcing models, we define the firm's *inventory position* as the aggregate of its inventory level along with all outstanding orders to be delivered in future periods, i.e., at the beginning of periods $n-1, n-2, \ldots, n-l_e$. Note that this covers *all* orders placed in the past that have not been delivered yet; even an order placed with the regular supplier in the prior period $n+1$, clearly the last to arrive, among all outstanding orders, will arrive at the beginning of period $n - l_e$ since $l_r = l_e + 1$. Based on its inventory position, the firm now decides on new order sizes to be placed with each of the two suppliers as well as any salvage quantity if it wants to *reduce* its inventory position. Stochastic demand is then realized and satisfied with on-hand inventory. At the end of the period, any unsatisfied demand is fully backlogged while leftover inventory is fully carried over to the next period. It is easily verified that it is never optimal to simultaneously place an order with the expedited supplier and to sell off a batch of inventory, since both inventory adjustments take effect in the very same period, i.e., l_e periods, hence.

We first introduce the notation for the model primitives, treating the per-unit revenue in any given period as a negative cost. For period $n = N, N-1, \ldots, 0$, denote:

$$\begin{aligned}
K_n, C_n &= \text{fixed cost and capacity limit for any order, respectively.} \\
c_n^e, c_n^r &= \text{unit price charged by the expedited and regular suppliers, respectively. } (c_n^e > c_n^r) \\
l_e, l_r &= \text{order lead time for the expedited and regular suppliers, respectively.} \\
K_n^v, c_n^v, C_n^v &= \text{fixed cost, unit revenue and capacity limit, respectively, when salvaging inventory.} \\
D_n &= \text{stochastic demand.} \\
\alpha &= \text{discount factor, } \alpha \in [0,1].
\end{aligned}$$

Assume $\{D_n, 1 \le n \le N\}$ are independent random variables with general distributions. Inventory and backlogging related costs will be introduced below. We impose the following restrictions on the cost parameters. These restrictions are innocuous and merely preclude arbitrage opportunities.

ASSUMPTION 1. *For any period n, the unit order costs, the unit salvage revenue and backlogging cost satisfy the restrictions: (i) $c_n^e \ge c_n^v$, and (ii) $c_n^r \ge \alpha c_{n-1}^v$.*

Assumption 1 (i) precludes the possibility of arbitrage opportunities where goods are procured from the expedited supplier, to be sold at a premium via the salvage channel. Similarly, Assumption 1 (ii) precludes the possibility of buying units form the regular supplier in period n, and initiating the sell-off of these units one period later, so that the purchased units can be sold off at a premium in the same period when they arrive, through the salvage channel.

The state and action variables are given by:

$x_n =$ the inventory position at the beginning of period n

$\quad =$ the inventory level at the beginning of period n, plus all outstanding orders from both suppliers, *before* the determination of the current orders.

$q_n^e =$ the size of the expedited-channel inventory adjustment in period n.

$q_n^r =$ the size of the order placed with the regular supplier in period n.

Note, $q_n^r \ge 0$ while $-\infty < q_n^e < \infty$. When $q_n^e \ge 0$, q_n^e represents the size of the order placed with the expedited supplier, while $q_n^e < 0$ indicates that in period n a batch of size $|q_n^e|$ is scheduled to be salvaged (l_e periods later).

Our ability to aggregate the order size placed with the expedited supplier, with the amount to be salvaged (a lead time later), follows, by Assumption 1, from the fact that it is never optimal to initiate both a negative and a positive expedited-channel inventory adjustment in the same period: by Assumption 1 (i), if the net inventory position adjustment is positive, one is better off reducing the order to the level of the *net* inventory adjustment and canceling the salvage batch; if the net inventory adjustment is negative, one is better off reducing the size of the salvage batch to this net adjustment level, while canceling the order with the expedited supplier.

The inventory position dynamics are given by

$$x_{n-1} = x_n + q_n^e + q_n^r - D_n.$$

Moreover, the *inventory level* I_{n-l_e} at the *end* of period $n - l_e$ is given by

$$I_{n-l_e} = x_n + q_n^e - D_{n,n-l_e}, \tag{1}$$

where $D_{n,m} = D_n + D_{n-1} + \cdots + D_m$ is the aggregate demand in the time interval $[n, m]$ with $m \le n$.

Assuming the inventory and backlogging related costs depend on the end-of-period inventory level sizes only—as in virtually all single-sourcing inventory models, the (single dimensional) inventory position x_n is a sufficient description of the state of the system. Instead of charging the actual inventory costs that arise at the end of period m ($m = 1, 2, \ldots$) to *that* period, one obtains an equivalent representation of the controllable part of the total expected discounted cost by charging its *expected value* to the *start* of period $m + l_e$, employing the stochastic identity (1). Thus for all $n = 1, 2, \ldots, N$, let

$$L_n(x_n + q_n^e) = \text{the expected value of all inventory and backlogging related costs}$$
$$\text{at the end of period } n - l_e \text{ discounted back to period } n$$

and impose a standard assumption regarding these functions, satisfied for most common cost structures.

Assumption 2. *The function $L_n(\cdot)$ is convex and $L_n(y) = O(|y|^p)$ for some $p \ge 1$, $n = 1, 2, \ldots, N$. Also, $\mathbb{E}(D_n^p) < \infty$ for $n = 1, 2, \ldots, N$.*

We are now ready to formulate the dynamic programming recursions for our model. Let

$$f_n(x) = \text{the optimal discounted expected total costs in the last } n \text{ periods of the planning horizon,}$$
$$\text{when period } n \text{ is started with an inventory position } x.$$

As mentioned, under Assumption 1, it is never optimal to place an order with the expedited supplier along with the initiation of a salvage order. Then, in any given period, the firm *either*

(i) places an order with the expedited supplier, possibly combined with an order for the regular supplier, or

(ii) initiates the sale of a salvage batch, possibly combined with an order for the regular supplier.

Therefore

$$f_n(x) = \min\{f_n^1(x), f_n^2(x)\}, \tag{2}$$

where

$$f_n^1(x) = \text{the minimum total expected cost in the last } n \text{ periods, when starting with an inventory position}$$
$$\text{of } x \text{ units, and assuming } \textit{no} \text{ inventory salvage is initiated in period } n,$$
$$f_n^2(x) = \text{the minimum total expected cost in the last } n \text{ periods, when starting with an inventory position}$$
$$\text{of } x \text{ units, and assuming a salvage sale } \textit{is} \text{ initiated in period } n.$$

These functions satisfy the recursions:

$$f_n^1(x) = \min_{q_n^e, q_n^r \in [0, C_n]} \{K_n \delta(q_n^e) + c_n^e q_n^e + K_n \delta(q_n^r) + c_n^r q_n^r + L_n(x + q_n^e) + \alpha \mathbb{E} f_{n-1}(x + q_n^e + q_n^r - D_n)\}, \tag{3}$$

$$f_n^2(x) = \min_{q_n^e \in [-C_n^v, 0], q_n^r \in [0, C_n]} \{K_n^v \delta(-q_n^e) + c_n^v q_n^e + K_n \delta(q_n^r) + c_n^r q_n^r + L_n(x + q_n^e) + \alpha \mathbb{E} f_{n-1}(x + q_n^e + q_n^r - D_n)\}, \tag{4}$$

where $\delta(u) = 1$ if $u > 0$ and $\delta(u) = 0$ otherwise.

In settings without salvage opportunities, $f_n(x) = f_n^1(x)$. However, to allow for a unified treatment, we model the "no-salvage" case as one in which $c_n^v = -M$, a large *negative* number such that salvaging is completely unattractive. Without loss of generality, we set $K_n^v = 0$ and $C_n^v = \infty$ in this case.

Since (3) and (4) share the last three terms, and substituting

$y_n = x + q_n^e =$ the beginning inventory position of period n, after inclusion of any expedited supplier order or salvage sale quantity,

$z_n = y_n + q_n^r = x_n + q_n^e + q_n^r =$ the beginning inventory position of period n, after inclusion of all inventory adjustments of period n,

we can rewrite (3) and (4) as

$$f_n^1(x) = \min_{y \in [x, x + C_n]} \{K_n \delta(y - x) + c_n^e (y - x) + g_n(y)\}, \tag{5}$$

$$f_n^2(x) = \min_{y \in [x - C_n^v, x]} \{K_n^v \delta(x - y) + c_n^v (y - x) + g_n(y)\}, \tag{6}$$

where

$$g_n(y) = L_n(y) + f_n^r(y), \tag{7}$$

$$f_n^r(y) = \min_{z \in [y, y + C_n]} \{K_n \delta(z - y) + c_n^r (z - y) + \alpha \mathbb{E} f_{n-1}(z - D_n)\}. \tag{8}$$

In other words, the order and salvage decisions in any given period may be thought of as occurring in two stages: first a new order with the expedited supplier/salvage quantity is determined based on x_n—the inventory position, followed by the choice of a supplementary order with the regular supplier based on the augmented inventory position y_n. Our second conclusion is that the *dual source* inventory planning problem is equivalent to a *single source* model with an "adjusted" future cost function $g_n(\cdot)$.

In the next section, we therefore summarize the structural properties of the optimal inventory policy in the *single sourcing* problem.

Beyond Assumption 1, we need a few restrictions on the fixed costs and capacity parameters:

Assumption 3.

$$K_n \geq \alpha K_{n-1}, \quad K_n^v \geq \alpha K_{n-1}^v,$$
$$C_n \leq C_{n-1}, \quad C_n^v \leq C_{n-1}^v.$$

In other words, capacity limits are assumed to weakly increase over time, an assumption satisfied by most practical applications. The inequalities $K_n \geq \alpha K_{n-1}$ and $K_n^v \geq \alpha K_{n-1}^v$ echo these in the basic single source inventory problem with fixed order costs, see Scarf (1960) or Zipkin (2000).

4. (C_1K_1, C_2K_2)-Convexity and the Single-Period Problem

Federgruen et al. (2017) introduced a key convexity concept, referred to as (C_1K_1, C_2K_2)-convexity, which is a generalization of all familiar convexity properties including K-convexity, see Scarf (1960), CK-convexity, see Gallego and Scheller-Wolf (2000), sym-K-convexity, see Chen and Simchi-Levi (2004), and weak (K_1, K_2)-convexity, see Semple (2007).

4.1. Definitions and Properties

Definition 1 ((C_1K_1, C_2K_2)-convexity). Given non-negative constants C_1, K_1 and C_2, K_2, a real-valued continuous function f is called *strongly* (C_1K_1, C_2K_2)*-convex* if for any $x \geq y, a \in [0, C_1]$ and $b \in (0, C_2]$,

$$f(x+a) + K_1 \geq f(x) + \frac{a}{b}\Big(f(y) - f(y-b) - K_2\Big). \tag{9}$$

Denote $SC_{C_1K_1,C_2K_2}$ as the set of all strongly (C_1K_1, C_2K_2)-convex functions. When (9) is required only for $x = y$, we refer to the property as *weak* (C_1K_1, C_2K_2)-convexity.

Figure 3 provides an intuitive way of understanding the strong (C_1K_1, C_2K_2)-convexity property. For any two points $y \leq x$, select any point $x + a$ with $a \in (0, C_1]$ and any point $y - b$ with $b \in (0, C_2]$. Raise the function value at point $x + a$ by K_1 and draw a ray from $(x, f(x))$ to $(x+a, f(x+a)+K_1)$. Similarly raise the function value at point $y - b$ by K_2 and draw a ray from $(y - b, f(y - b) + K_2)$ to $(y, f(y))$. Then f is strongly (C_1K_1, C_2K_2)-convex if the slope of the former ray is bigger than or equal to the slope of the latter ray.

Federgruen et al. (2017) show that (C_1K_1, C_2K_2)-convexity generalizes all familiar convexity structures as special cases:

- Simple convexity is $(\infty 0, \infty 0)$-convexity.
- K-convexity of Scarf (1960) is weak $(\infty K, \infty 0)$-convexity.
- Weak/strong CK-convexity of Gallego and Scheller-Wolf (2000) is weak/strong $(CK, \infty 0)$-convexity, respectively.
- Sym-K-convexity of Chen and Simchi-Levi (2004) is weak $(\infty K, \infty K)$-convexity.
- Weak (K_1, K_2)-convexity of Semple (2007) is weak $(\infty K_1, \infty K_2)$-convexity.

Figure 3 **Geometric illustration of strongly** (C_1K_1, C_2K_2)**-convex functions.**

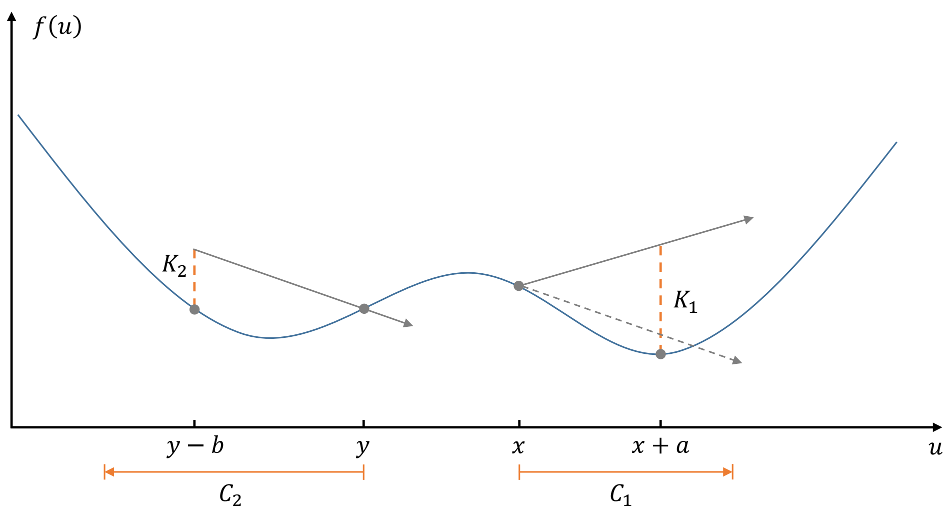

4.2. Single-Period, Bilateral-Adjustment Inventory Problem

In Section 3, we showed that the *dual sourcing* inventory problem may be viewed as equivalent to a single sourcing problem with an adjusted salvage value function. To derive the structure of an optimal strategy we therefore review from Federgruen et al. (2017) the structure of the optimal policy in the following *single sourcing* problem. Assume the firm starts the period with an inventory level (or position) of x units, and has the ability to elevate the level, instantaneously, via an order of up to C_1 units, or to decrease the level, by selling up to C_2 units. An order (salvage sale) involves a fixed cost K_1 (K_2) and a variable per-unit cost (revenue) β_1 (β_2). The expected cost during the period is given by a function $g(y)$ of the inventory level after any adjustment (order or salvage quantity). This gives rise to the optimization problem:

$$g_1(x) = \min_{y\in[x,x+C_1]}\{K_1\delta(y-x)+\beta_1(y-x)+g(y)\}, \tag{10}$$

$$g_2(x) = \min_{y\in[x-C_2,x]}\{K_2\delta(x-y)+\beta_2(y-x)+g(y)\}, \tag{11}$$

$$g_0(x) = \min\{g_1(x), g_2(x)\}, \tag{12}$$

Note that, in each period n, the value functions $f_n^1(\cdot), f_n^2(\cdot)$ satisfy a single stage problem of the structure, given by (10) and (11), with $g_1(\cdot) = f_n^1(\cdot), g_2(\cdot) = f_n^2(\cdot)$ and $g(\cdot) = g_n(\cdot)$. Federgruen et al. (2017) show that the optimal inventory adjustment strategy has a relatively simple structure when the function $g(\cdot)$ is strongly (C_1K_1, C_2K_2)-convex. In the next section, we will inductively show that in each period n, the function $g_n(\cdot)$ has the strong (C_1K_1, C_2K_2)-convexity property for the proper parameters C_1, K_1, C_2 and K_2.

Define auxiliary functions

$$\widetilde{g}_1(x) = K_1 + \min_{y\in[x,x+C_1]}\{\beta_1(y-x)+g(y)\}, \tag{13}$$

$$\widetilde{g}_2(x) = K_2 + \min_{y\in[x-C_2,x]}\{\beta_2(y-x)+g(y)\}, \tag{14}$$

as counterparts of $g_1(x)$ and $g_2(x)$, under *definitive* inventory adjustment, i.e., *definitively* incurring fixed costs for ordering or salvaging, respectively, and let $A_i(x) = \widetilde{g}_i(x) - g(x)$ be the increase in minimal cost if forced to order (for $i=1$) or salvage (for $i=2$).

To characterize the structure of an optimal policy we need to define some critical points as follows, with the convention that the infimum (supremum) of an empty set equals $+\infty$ ($-\infty$).

Definition 2. (Critical Points) For a continuous function $g(\cdot)\in SC_{C_1K_1,C_2K_2}$ and any β_1,β_2, define

$$B = \inf\Big\{\arg\min_y\{\beta_1 y + g(y)\}\Big\}, \qquad b=\inf\{x:\ A_1(x)\ge 0\}, \qquad \bar{b}=\sup\{x:\ A_1(x)<0\}, \tag{15}$$

$$S = \sup\Big\{\arg\min_y\{\beta_2 y + g(y)\}\Big\}, \qquad s=\sup\{x:\ A_2(x)\ge 0\}, \qquad \underline{s}=\inf\{x:\ A_2(x)<0\}. \tag{16}$$

These critical points play important roles in the structure of the optimal strategy. By its definition, B is the (smallest) global minimizer of $\widetilde{g}_1(x)$ if $C_1=\infty$, i.e., the smallest order-up-to level for sufficiently small x if ordering is better than staying put. Similarly, S is the (largest) global minimizer of $\widetilde{g}_2(x)$ if $C_2=\infty$, .i.e., the biggest salvage-down-to level for sufficiently large x if salvaging is better than staying put; b is the smallest among all inventory levels where ordering is not better than staying put; $\bar{b}$ is the largest among all inventory levels where ordering is better than staying put; s is the largest among all inventory levels where salvaging is not better than staying put; $\underline{s}$ is the smallest among all inventory levels where salvaging is better than staying put.

Note that by definition we have

$$g_1(x) = \min\{g(x),\ \widetilde{g}_1(x)\}, \qquad A_1(x)<0 \quad \forall x<b, \qquad A_1(x)\ge 0 \quad \forall x>\bar{b}, \tag{17}$$

$$g_2(x) = \min\{g(x),\ \widetilde{g}_2(x)\}, \qquad A_2(x)<0 \quad \forall x>s, \qquad A_2(x)\ge 0 \quad \forall x<\underline{s}. \tag{18}$$

The lemma below characterizes the ranking of the critical points, which is important when developing the optimal policy structure.

Lemma 1 **(Critical Points)**. *Assume $\beta_1\ge\beta_2$ and $g(\cdot)\in SC_{C_1K_1,C_2K_2}$, then*

(i) $-\infty\le b\le\bar{b}\le\underline{s}\le s\le\infty$*;*

(ii) $-\infty\le b\le B\le S\le s\le\infty$*;*

(iii) *If $C_2=\infty$ and $K_1\ge K_2$, then $\bar{b}\le B$; if $C_1=\infty$ and $K_1\le K_2$, then $S\le\underline{s}$;*

(iv) *If $C_1 = \infty$ and $K_2 = 0$, then $b = \bar{b}$; if $C_2 = \infty$ and $K_1 = 0$, then $\underline{s} = s$. If $C_1 = C_2 = \infty$ and $K_1 = K_2 = 0$, then $b = \bar{b} = B, S = \underline{s} = s$.*

Proof. See Proof of Proposition 2 in Federgruen et al. (2017).

In this lemma, (i) ranks four critical points. (ii) ranks and locates the global minimizers B and S between b and s. (iii) and (iv) lead to simple policy structures, in certain special cases, which will be discussed later.

With the definitions and lemmas developed so far, we now proceed to the optimal single-period policy structure, in the following proposition.

PROPOSITION 1 **(Optimal Policy Structure)**. *Assume $\beta_1 \geq \beta_2$ and $g(\cdot) \in SC_{C_1K_1,C_2K_2}$, then $g_0(x)$ and the corresponding minimizer $y^*(x)$ are characterized by Table 1, in which $\widetilde{g}_1(\cdot)$ and $\widetilde{g}_2(\cdot)$ are defined by (13) and (14), respectively. If $y^*(x)$ is specified as a two-element set $\{\cdot,\cdot\}$, either one of the two elements may apply. Let*

$$B(x) = \inf\{\underset{x \leq y \leq x+C_1}{\arg\min}\{\beta_1 y + g(y)\}\}, \quad S(x) = \sup\{\underset{x-C_2 \leq y \leq x}{\arg\min}\{\beta_2 y + g(y)\}\}$$

denote minimizers of $\widetilde{g}_1(x)$ and $\widetilde{g}_2(x)$, respectively.

Table 1 Optimal policy structure

x	$(-\infty, b)$	$[b, \bar{b})$	$[\bar{b}, \underline{s}]$	$(\underline{s}, s]$	(s, ∞)
$g_0(x)$	$\widetilde{g}_1(x)$	$\min\{\widetilde{g}_1(x), g(x)\}$	$g(x)$	$\min\{\widetilde{g}_2(x), g(x)\}$	$\widetilde{g}_2(x)$
$y^*(x)$	$B(x)$	$\{B(x), x\}$	x	$\{S(x), x\}$	$S(x)$

Proof. See Proof of Theorem 2 in Federgruen et al. (2017).

In other words, four critical points partition the inventory position line into 5 regions. In the two extreme regions, $(-\infty, b)$ and (s, ∞), a positive order or salvage transaction needs to be initiated, respectively; in the middle region, $[\bar{b}, \underline{s}]$, it is optimal to stay put; in the second region, $[b, \bar{b})$, it is optimal to either order or to stay put, and in the fourth region, $(\underline{s}, s]$, it is optimal to either initiate a salvage transaction or to stay put. Within the latter two regions, it is possible that the optimal policy alternates several times between ordering or salvaging versus staying put, a phenomenon already discovered in simpler models without salvage opportunities, see e.g., Chen and Lambrecht (1996) and Chen (2004).

4.3. Preservation of (C_1K_1, C_2K_2)-Convexity under Value Iteration

The (C_1K_1, C_2K_2)-convexity is preserved under the minimization operations specified by (10)–(12). This enables us to extend the structural results, above, to general multi-period planning horizons.

PROPOSITION 2 **(Preservation)**. *Assuming $\beta_1 \geq \beta_2$, if $g(\cdot)$ is strongly (C_1K_1, C_2K_2)-convex, then*

$$
\begin{aligned}
g_1(x) &= \min_{y\in[x,x+C_1']}\{K_1\delta(y-x)+\beta_1(y-x)+g(y)\},\\
g_2(x) &= \min_{y\in[x-C_2',x]}\{K_2\delta(x-y)+\beta_2(y-x)+g(y)\},\\
g_0(x) &= \min\{g_1(x), g_2(x)\},
\end{aligned}
$$

are also strongly (C_1K_1, C_2K_2)-convex for any $C_1' \geq C_1, C_2' \geq C_2$.

Proof. See Proof of Proposition 3 in Federgruen et al. (2017).

5. The Full Planning Horizon: Optimal Policy Structures

Given the results of Section 4, the structure of the optimal policy in each period $n = 1, \ldots, N$ follows easily by showing that the value functions $\{f_n(x)\}$ satisfy the required strong (C_1K_1, C_2K_2)-convexity property.

THEOREM 1. *Assume $f_0(x) \in SC_{C_0K_0, C_0^vK_0^v}$ and $f_0(x) = O(|x|^p)$ for some integer $p \geq 1$. Then $f_n(x) \in SC_{C_nK_n, C_n^vK_n^v}$ and $f_n(x) = O(|x|^p)$ for $n = N, N-1, \ldots, 0$.*

Proof. We prove this theorem by induction. By our assumption, the theorem holds for $n = 0$. Suppose the result holds for period $n-1$, i.e., $f_{n-1}(\cdot) \in SC_{C_{n-1}K_{n-1}, C_{n-1}^vK_{n-1}^v}$ and $f_{n-1}(x) = O(|x|^p)$. We first prove that $f_n(x) = O(|x|^p)$. Since $f_{n-1}(x) = O(|x|^p)$, there exists a constant $A > 0$ such that $|f_{n-1}(x)| \leq A|x|^p$; so that $|\mathbb{E}f_{n-1}(z-D_n)| \leq A\mathbb{E}|z-D_n|^p \leq A\mathbb{E}(|z|+D_n)^p = A\sum_{l=0}^p \binom{p}{l}\mathbb{E}D_n^{p-l}|z|^l \leq B\max\{|z|^p, 1\}$ for some constant $B > 0$. Let $z^*(y)$ achieve the minimum in (8), then $|f_n^r(y)| \leq K_n + c_n^r|z^*| + \alpha B|z^*|^p \leq K_n + c_n^r(|y|+C_n) + \alpha B\max\{1, (|y|+C_n)^p\} = O(|y|^p)$. By the same argument, and since $L_n(y) = O(|y|^p), g_n(y) = O(|y|^p)$ and $f_n^1(x), f_n^2(x)$ and $f_n(x)$ are all $O(|x|^p)$.

We then prove that $f_n(x) \in SC_{C_nK_n, C_n^vK_n^v}$. Since $f_{n-1}(\cdot) \in SC_{C_{n-1}K_{n-1}, C_{n-1}^vK_{n-1}^v}$, by Lemma 1 (iii) and (iv) in Federgruen et al. (2017), and Assumption 3,

$$
\alpha\mathbb{E}f_{n-1}(z-D_n) \in SC_{C_{n-1}(\alpha K_{n-1}), C_{n-1}^v(\alpha K_{n-1}^v)} \subset SC_{C_nK_n, C_n^vK_n^v}. \tag{19}
$$

It then follows from Proposition 2 and Assumption 3 that $f_n^r(y) \in SC_{C_nK_n, C_n^vK_n^v}$. Since $L_n(\cdot)$ is convex, we have

$$
g_n(y) = L_n(y) + f_n^r(y) \in SC_{C_nK_n, C_n^vK_n^v} \tag{20}
$$

by Lemma 1 (iii) in Federgruen et al. (2017). Finally by Proposition 2 again, $f_n^1(x), f_n^2(x), f_n(x) \in SC_{C_nK_n, C_n^vK_n^v}$. $\square$

Theorem 1 guarantees strong $(C_nK_n, C_n^vK_n^v)$-convexity of the value function $f_n(x)$ in any period, which allows us to obtain the optimal policy structure using the results in Section 4.2. Define auxiliary functions

$$\begin{aligned}
\widetilde{g}_n^1(x) &= K_n + \min_{y\in[x,x+C_n]}\{c_n^e(y-x)+g_n(y)\}, & A_n^1(x) &= \widetilde{g}_n^1(x) - g_n(x),\\
\widetilde{g}_n^2(x) &= K_n^v + \min_{y\in[x-C_n^v,x]}\{c_n^v(y-x)+g_n(y)\}, & A_n^2(x) &= \widetilde{g}_n^2(x) - g_n(x),\\
\widetilde{g}_n^r(y) &= K_n + \min_{z\in[y,y+C_n]}\{c_n^r(z-y)+\alpha\mathbb{E}f_{n-1}(z-D_n)\}, & A_n^r(y) &= \widetilde{g}_n^r(y) - \alpha\mathbb{E}f_{n-1}(y-D_n),
\end{aligned}$$

where $\widetilde{g}_n^1(x)$ represents the optimal expected cost from period n on assuming one is committed to place an order with the expedited supplier in this period (possibly combined with an order with the regular supplier). $A_n^1(x) = \widetilde{g}_n^1(x) - g_n(x)$ denotes the cost differential with the optimal expected cost from period n on, assuming one is committed to *forgo* any order with the expedited supplier or a salvage batch in this period. Similarly, $\widetilde{g}_n^2(x)$ represents the optimal expected cost from period n on assuming one is committed to initiate a salvage batch in this period (possibly combined with an order with the regular supplier). $A_n^2(x) = \widetilde{g}_n^2(x) - g_n(x)$ denotes the cost differential with the same benchmark $g_n(\cdot)$ used in the definition of $A_n^1(\cdot)$. Finally, given an inventory position y in period n, $\widetilde{g}_n^r(y)$ denotes the expected optimal cost until the end of the planning horizon, assuming one is committed to place an order with the regular supplier, while $A_n^r(y)$ denotes the cost differential vis-à-vis the optimal expected cost when forgoing an order with the regular supplier in this period n.

Define the following critical points for period n with implications as mentioned after Definition 2:

$$B_n^e = \inf\{\arg\min_y\{c_n^e y + g_n(y)\}\}, \qquad b_n^e = \inf\{x : A_n^1(x) \geq 0\}, \qquad \bar{b}_n^e = \sup\{x : A_n^1(x) < 0\}, \tag{21}$$

$$S_n = \sup\{\arg\min_y\{c_n^v y + g_n(y)\}\}, \qquad s_n = \sup\{x : A_n^2(x) \geq 0\}, \qquad \underline{s}_n = \inf\{x : A_n^2(x) < 0\}, \tag{22}$$

$$B_n^r = \inf\{\arg\min_z\{c_n^r z + \alpha\mathbb{E}f_{n-1}(z-D_n)\}\}, \; b_n^r = \inf\{x : A_n^r(x) \geq 0\}, \; \bar{b}_n^s = \sup\{x : A_n^r(x) < 0\}. \tag{23}$$

Using these critical points, the following theorem follows immediately from Proposition 1.

THEOREM 2. *The optimal solution to $f_n(x_n)$ is characterized by Table 2, in which certain $y_n^*(x_n)$ and $z_n^*(y_n^*)$ are either uniquely determined or take a value in a bi-valued set $\{\cdot,\cdot\}$, and*

$$\begin{aligned}
B_n^e(x) &= \inf\{\arg\min_{x\leq y\leq x+C_n}\{c_n^e y + g_n(y)\}\},\\
S_n(x) &= \sup\{\arg\min_{x-C_n^v\leq y\leq x}\{c_n^v y + g_n(y)\}\},\\
B_n^r(y) &= \inf\{\arg\min_{y\leq z\leq y+C_n}\{c_n^r z + \alpha\mathbb{E}f_{n-1}(z-D_n)\}\}
\end{aligned}$$

are minimizers of $\widetilde{g}_n^1(x), \widetilde{g}_n^2(x)$ and $\widetilde{g}_n^r(y)$, respectively.

Table 2 Optimal policy structure for systems with fixed costs and capacity limits

(a) Expedited supplier and salvage

x_n	$(-\infty, b_n^e)$	$[b_n^e, \bar{b}_n^e)$	$[\bar{b}_n^e, \underline{s}_n^v]$	$(\underline{s}_n^v, s_n]$	(s_n, ∞)
$f_n(x_n)$	$\widetilde{g}_n^1(x_n)$	$\min\{\widetilde{g}_n^1(x_n), g(x_n)\}$	$g(x_n)$	$\min\{\widetilde{g}_n^2(x_n), g(x_n)\}$	$\widetilde{g}_n^2(x_n)$
$y_n^*(x_n)$	$B_n^e(x_n)$	$\{B_n^e(x_n), x_n\}$	x_n	$\{S_n(x_n), x_n\}$	$S_n(x_n)$

(b) Regular supplier

y_n^*	$(-\infty, b_n^r)$	$[b_n^r, \bar{b}_n^r)$	$[\bar{b}_n^r, \infty)$
$f_n^r(y_n^*)$	$\widetilde{g}_n^r(y_n^*)$	$\min\{\widetilde{g}_n^r(y_n^*), \alpha\mathbb{E}f_{n-1}(y_n^* - D_n)\}$	$\alpha\mathbb{E}f_{n-1}(y_n^* - D_n)\}$
$z_n^*(y_n^*)$	$B_n^r(y_n^*)$	$\{B_n^r(y_n^*), y_n^*\}$	y_n^*

Proof. By Assumption 1, $c_n^e \geq c_n^v$. By (20) in the proof of Theorem 1, $g_n(y) \in SC_{C_nK_n, C_n^vK_n^v}$. Applying Proposition 1 with properly defined critical points, we immediately obtain the optimal policy structure for ordering with the expedited supplier and salvaging, as given by Table 2 (a). For the regular supplier, since $\alpha\mathbb{E}f_{n-1}(z - D_n) \in SC_{C_nK_n, C_n^vK_n^v}$ by (19) in the proof of Theorem 1, we can apply Proposition 1 again and use Corollary 2 in Federgruen et al. (2017) to obtain the optimal policy structure given by Table 2 (b). $\square$

Based on Theorem 2, the optimal order and salvage decisions, in period n, can be determined in two steps. In the first step, determine the optimal adjustment target $y_n^*(x_n)$ for the expedited channel (i.e., the optimal order amount from the expedited supplier or salvage amount) using Table 2 (a) after computing all necessary auxiliary functions and critical points; this also involves the calculation of the value function $f_n^r(y)$. In the second step, with y_n^* obtained from the first step, determine the optimal adjustment target $z_n^*(y_n^*)$ for the regular channel (which implies the optimal order amount from the regular supplier) using Table 2 (b).

5.1. Special Settings

In this subsection, we consider several special settings where the optimal policy structure takes on simpler forms.

5.1.1. No Fixed Costs, No Capacity Limits, Salvage Opportunities

In this setting $K_n = K_n^v = 0$ and $C_n = C_n^v = \infty$. Based on Table 3 (c) in Corollary 3 in Federgruen et al. (2017), the optimal policy structure for systems without fixed costs and capacity limits is summarized by Table 3. The optimal inventory decision for the expedited channel (ordering from the expedited supplier or salvaging) follows a double "base stock"–type policy. The inventory position line is partitioned into three consecutive regions. In the left and right most regions, it is optimal to order up to a base-stock level B_n^e from the expedited supplier or to salvage down to a level S_n, respectively; in the middle region it is optimal to stay put. The optimal inventory decision for the regular supplier is a simple base stock policy with an order-up-to level B_n^r. Note that if salvaging is not allowed, the first stage policy is a simple base stock policy, i.e., $S_n = \infty$, eliminating the last column in subtable (a).

Table 3 Optimal policy structure for systems without fixed costs and capacity limits

(a) Expedited supplier and salvage

x_n	$(-\infty, B_n^e)$	$[B_n^e, S_n]$	(S_n, ∞)
$y_n^*(x_n)$	B_n^e	x_n	S_n

(b) Regular supplier

y_n^*	$(-\infty, B_n^r)$	$[B_n^r, \infty)$
$z_n^*(y_n^*)$	B_n^r	y_n^*

5.1.2. Fixed Costs, No Capacity Limits, No Salvage. In this setting $K_n > 0, K_n^v = 0$, $C_n = C_n^v = \infty$. Based on Table 3 (a) from Corollary 3 in Federgruen et al. (2017), the optimal policy structure is summarized by Table 4. Both channels take on the classical "(s, S)"–type policy. More specifically, the expedited channel adopts a (b_n^e, B_n^e) policy and the regular channel follows a (b_n^r, B_n^r) policy.

Table 4 Optimal policy structure for systems with fixed costs but without capacity limits

(a) Expedited supplier

x_n	$(-\infty, b_n^e)$	$[b_n^e, \infty)$
$y_n^*(x_n)$	B_n^e	x_n

(b) Regular supplier

y_n^*	$(-\infty, b_n^r)$	$[b_n^r, \infty)$
$z_n^*(y_n^*)$	B_n^r	y_n^*

5.1.3. No Fixed Costs, Capacity Limits, No Salvage. In this setting, $K_n = K_n^v = 0$ and $C_n < \infty, C_n^v = \infty$. Based on Table 4 (c) from Corollary 4 in Federgruen et al. (2017), under the optimal policy, both channels adopt a modified base-stock policy as in Table 5.

Table 5 Optimal policy structure for systems without salvage option

(a) Expedited supplier

x_n	$(-\infty, \bar{b}_n^e - C_n)$	$[\bar{b}_n^e - C_n, \bar{b}_n^e)$	$[\bar{b}_n^e, \infty)$
$y_n^*(x_n)$	$x_n + C_n$	$\bar{b}_n^e$	x_n

(b) Regular supplier

y_n^*	$(-\infty, \bar{b}_n^r - C_n)$	$[\bar{b}_n^r - C_n, \bar{b}_n^r)$
$z_n^*(y_n^*)$	$y_n^* + C_n$	$\bar{b}_n^r$

5.1.4. Fixed Costs and Capacity Limits, No Salvage. In this setting $K_n > 0, K_n^v = 0$ and $C_n < \infty, C_n^v = \infty$. Based on Table 4 (a) from Corollary 4 in Federgruen et al. (2017), the optimal policy structure for systems with fixed costs and capacity limits but without salvage option is summarized by Table 6, where

$$\mathbf{1}^+_{b_n^e} = \mathbf{1}(b_n^e > \bar{b}_n^e - C_n), \quad \mathbf{1}^-_{b_n^e} = \mathbf{1}(b_n^e < \bar{b}_n^e - C_n); \qquad \mathbf{1}^+_{b_n^r} = \mathbf{1}(b_n^r > \bar{b}_n^r - C_n), \quad \mathbf{1}^-_{b_n^r} = \mathbf{1}(b_n^r < \bar{b}_n^r - C_n).$$

The optimal policy structures for the expedited and regular suppliers are similar in form. There are four regions. In the far-left region, the optimal policy is to order as much as possible—up to the capacity limit, while in the far-right region it is optimal to stay put. In the remaining two intermediate regions, the optimal policy is less categorical: in the third region, it is optimal to either stay put or to place an order; in the second region, one of two cases prevails: in one case it is optimal to either order as much as possible or stay put and in the other it is optimal to order a quantity specified by the aforementioned functions, $\bar{B}_n^e(\cdot)$ or $\bar{B}_n^r(\cdot)$.

Table 6 Optimal policy structure for systems without salvage option

(a) Expedited supplier

x_n	$(-\infty, \min\{\bar{b}_n^e - C_n, b_n^e\})$	$[\min\{\bar{b}_n^e - C_n, b_n^e\}, \max\{\bar{b}_n^e - C_n, b_n^e\})$	$[\max\{\bar{b}_n^e - C_n, b_n^e\}, \bar{b}_n^e)$	$[\bar{b}_n^e, \infty)$
$y_n^*(x_n)$	$x_n + C_n$	$\{x_n + C_n, x_n\}\mathbf{1}^-_{b_n^e} + \bar{B}_n^e(x_n)\mathbf{1}^+_{b_n^e}$	$\{\bar{B}_n^e(x_n), x_n\}$	x_n

(b) Regular supplier

y_n^*	$(-\infty, \min\{\bar{b}_n^r - C_n, b_n^r\})$	$[\min\{\bar{b}_n^r - C_n, b_n^r\}, \max\{\bar{b}_n^r - C_n, b_n^r\})$	$[\max\{\bar{b}_n^r - C_n, b_n^r\}, \bar{b}_n^r)$	$[\bar{b}_n^r, \infty)$
$z_n^*(y_n^*)$	$y_n^* + C_n$	$\{y_n^* + C_n, y_n^*\}\mathbf{1}^-_{b_n^r} + \bar{B}_n^r(y_n^*)\mathbf{1}^+_{b_n^r}$	$\{\bar{B}_n^r(y_n^*), y_n^*\}$	y_n^*

6. General Lead Times

In this Section, we address the general case, with arbitrary lead time combinations $l_r > l_e$, rather than the *special* setting where $l_r - l_e = 1$. As shown in Whittemore and Saunders (1977) and, more recently, in Sheopuri et al. (2010), when the difference between the lead times $\Delta l \equiv l_r - l_e \geq 2$, it is no longer possible to collapse the state space into a one-dimensional space. Instead, Sheopuri et al. (2010) show that a *minimal* state description is of dimension Δl, even in the simplest of models, i.e., in the absence of any fixed order costs, capacity limits or salvage opportunities. The following is one such "minimal" state descriptions:

$$\mathcal{I}_t \equiv (I_t^{l_e}, I_t^{l_e+1}, \ldots, I_t^{l_r-1}),$$

where $I_t^l =$ the net inventory level at the beginning of period t plus all outstanding orders that will arrive by the beginning of period $t + l$. While of dimension Δl, this state description is already, a major simplification, beyond the straightforward state description, which includes the inventory level and each of l_e outstanding orders with the expedited supplier and the l_r orders with the regular supplier.

Note that $x_t = I_t^{l_r-1}$, represents the regular full inventory position at the beginning of period t, including *all* outstanding orders with the expedited and regular suppliers.

The exact optimal procurement strategy is a function of the full Δl-dimensional state vector $\mathcal{I}_t$. For example, Whittemore and Saunders (1977) already showed that, even in the absence of fixed costs, capacity limits or salvage opportunities, the optimal order policy for the regular supplier fails to be a base-stock policy based on the full inventory position, see also Theorem 3.2 in Sheopuri et al. (2010). This fully optimal strategy fails to have an elegant structure, as in the case $\Delta l = 1$ developed in Section 5. It is prohibitively difficult to compute, in particular when $\Delta l \geq 3$, say, and even if computable, it would be prohibitively difficult to implement.

For $\Delta l \geq 2$, we therefore need to resort to heuristics, and propose the following procedure. For the sake of notational simplicity, we confine ourselves to infinite horizon models with stationary inputs (parameters and demand distributions).

Two-Step Procedure:

Step 1: Set $\hat{l}_r = l_e + 1$. Identify the structure of the optimal policy from the results in Section 5, and calculate the minimal optimal cost value $\underline{C}$, a lower bound for the true optimal cost C^*. Similarly, an upper bound $\bar{C} \geq C^*$ is obtained when increasing l_e to $l_r - 1$, leaving l_r unchanged.

Step 2: Within the identified class of policies for expedited and regular orders, search for the optimal combination of policy parameters using $x_t = I_t^{l_r-1}$, the full inventory position measure, as the inventory index.

The procedure is most easily explained in the base model without fixed costs, capacity limits or salvage opportunities: when $l_e = l_r - 1$, it is known since Fukuda (1964)—and we showed it as a simple corollary of our general framework, see Table 3 and the adjacent discussion—that it is optimal to use a single index pair of base-stock policies. In Step 2, we follow the same order rules, using the full inventory position as the single inventory index, but search for the pair of base-stock levels $B^r > B^e$ under which the long-run average cost value is minimal along all such combined base-stock policies. Let

$$IP_1 = x_t = I_t^{l_r-1} = \text{full inventory position at the beginning of period } t,$$

$$IP_2 = IP_1 + q_t^e,$$

$$IP_3 = IP_1 + q_t^e + q_t^r = IP_2 + q_t^r.$$

As mentioned in Section 2, Scheller-Wolf et al. (2007) already proposed this class of single index policies in the base model. They show how the optimal base-stock level pair can be found *analytically*, with a special algorithm. In contrast, Veeraraghavan and Scheller-Wolf (2008) proposed a class of combined base-stock policies where the order with the expedited supplier is based on $I_t^{l_e}$ rather than $I_t^{l_r-1}$, while the order for the regular supplier continues to be based on $IP_2 = IP_1 + q_t^e$. The authors refer to such policies as "*dual index policies*".

Clearly, alternative dual index policies can be entertained by basing the expedited supplier order on *any one* of the components of the full state vector $\mathcal{I}_t \equiv (I_t^{l_e}, I_t^{l_e+1}, \ldots, I_t^{l_r-1})$—not just $I_t^{l_r-1} = x_t$ or $I_t^{l_e}$ as in Veeraraghavan and Scheller-Wolf (2008)—or a weighted average of those components, as in Sheopuri et al. (2010). As discussed in Section 2, in the base model, all of these various choices exhibit very comparable cost performance. This motivates our restriction to *single index* policies for the general model, with fixed costs, capacity limits and/or salvage opportunities.

6.1. Numerical Study

We have conducted a numerical study to explain how and to what extent dual sourcing can result in system-wide cost savings, compared with an optimal *single* sourcing policy (with either the slow or the fast supplier). We also investigate how the optimal policy parameters and the optimal cost value depend on various model primitives. As mentioned, in analyzing the dual sourcing setting, we follow the Two-Step Procedure, described above.

We investigate two sets of problem instances, with 25 and 30 problem instances, respectively. The first set relates to the base model, without fixed order costs, capacity limits or salvage opportunities. Here, the Two-Step Procedure suggests the use of a single index pair of base-stock policies, and we compute the best such pair with the algorithm in Scheller-Wolf et al. (2007). In the second set of instances, we add a fixed order cost. As a consequence, the Two-Step Procedure suggests the use of a single index pair of (s, S)-policies, both acting on the full inventory position. Such a pair of policies is characterized by 4 parameters (s_e, S_e, s_r, S_r); we determine the optimal quadruple of policy parameters via simulation and an exhaustive search over an integer grid (To reduce the computational times, we employed the parameter restriction $S_r - s_r = S_e - s_e$, i.e., we assumed an identical gap between the reorder level and order-up-to level for both suppliers.) Every problem instance was evaluated for every policy parameter condition, with a simulation over a horizon of 100,000 periods. All problem instances assume the demand distribution is Normal.

Table 7 exhibits the results for the first set of instances. The first 9 columns specify the input parameters, in particular: (i) the expedited lead time l_e; (ii) the regular lead time l_r; (iii) the *differential* between the per unit cost rates charged by the expedited and regular suppliers, c;

Table 7 Numerical results for settings without fixed cost

	Parameters								Dual Sourcing						Single - Expedited		Single - Regular	
Varying parameters	l_e	l_r	c	h	p	SL	μ	σ	Cost*	% Savings	S_e	S_r	% Expedited	Expedite Freq	Cost*	S_e	Cost*	S_r
Lead time difference l_r-l_e	0	1	10	5	495	99%	5	1	17.20	8.9%	6.9	12.7	2.2%	19.8%			18.89	13.3
	0	2	10	5	495	99%	5	1	19.28	16.9%	12.1	17.5	4.6%	34.6%			23.21	19.1
	0	3	10	5	495	99%	5	1	20.77	21.8%	16.9	22.1	6.8%	45.0%	63.49	7.4	26.54	24.7
	0	4	10	5	495	99%	5	1	21.79	27.2%	21.4	26.4	8.7%	52.6%			29.93	30.2
	0	5	10	5	495	99%	5	1	22.91	30.2%	25.8	30.6	10.2%	58.1%			32.85	35.7
Lead times l_r, l_e (fixing l_r-l_e)	0	2	10	5	495	99%	5	1	19.28	16.9%	12.1	17.5	4.6%	34.6%	63.49	7.4	23.21	19.1
	1	3	10	5	495	99%	5	1	24.14	9.0%	18.0	23.5	3.7%	29.5%	68.89	13.3	26.54	24.7
	2	4	10	5	495	99%	5	1	28.00	6.5%	23.7	29.3	3.4%	27.6%	73.07	19.0	29.93	30.2
	3	5	10	5	495	99%	5	1	31.33	4.6%	29.3	34.9	3.0%	25.1%	76.70	24.7	32.85	35.7
	4	6	10	5	495	99%	5	1	34.19	3.2%	34.7	40.6	2.2%	20.0%	79.93	30.2	35.32	41.1
Unit cost c	2	4	1	5	495	99%	5	1	24.55	12.5%	23.0	27.1	20.1%	81.8%	28.07	19.0		
	2	4	2.5	5	495	99%	5	1	25.71	14.1%	23.4	28.1	11.8%	63.2%	35.57	19.0		
	2	4	5	5	495	99%	5	1	26.79	10.5%	23.6	28.7	6.8%	45.2%	48.07	19.0	29.93	30.2
	2	4	10	5	495	99%	5	1	28.00	6.5%	23.7	29.3	3.4%	27.6%	73.07	19.0		
	2	4	20	5	495	99%	5	1	29.00	3.1%	23.6	29.8	1.2%	11.8%	123.07	19.0		
Service level $SL = p/(h+p)$	2	4	10	5	12	70%	5	1	12.95	0.4%	19.5	26.1	0.4%	5.0%	60.06	15.9	13.00	26.2
	2	4	10	5	20	80%	5	1	15.54	1.1%	20.5	26.7	1.0%	10.5%	62.17	16.5	15.72	26.9
	2	4	10	5	45	90%	5	1	19.21	2.6%	21.6	27.5	1.9%	17.4%	65.25	17.2	19.72	27.9
	2	4	10	5	95	95%	5	1	22.25	4.0%	22.3	28.2	2.3%	20.7%	67.93	17.9	23.18	28.7
	2	4	10	5	495	99%	5	1	28.00	6.5%	23.7	29.3	3.4%	27.6%	73.07	19.0	29.93	30.2
Demand volatility σ	2	4	10	5	495	99%	5	0.1	2.80	6.5%	20.4	25.4	0.3%	27.7%	52.31	15.4	2.99	25.5
	2	4	10	5	495	99%	5	0.5	14.00	6.5%	21.9	27.2	1.7%	27.6%	61.54	17.0	14.97	27.6
	2	4	10	5	495	99%	5	1.0	28.00	6.5%	23.7	29.3	3.4%	27.6%	73.07	19.0	29.93	30.2
	2	4	10	5	495	99%	5	1.5	42.00	6.5%	25.6	31.5	5.1%	27.6%	84.61	21.1	44.90	32.8
	2	4	10	5	495	99%	5	2.0	56.00	6.5%	27.4	33.6	6.8%	27.6%	96.20	23.1	59.87	35.4

Note. Setting: Normal demands for 100,000 periods; full backlogging.

(iv) K, the fixed cost incurred for any order; (v) h, the holding cost rate for any unit carried in inventory, in any given period; (vi) p, the backlogging cost rate incurred for any unit backlogged, per period; (vii) $SL = p/(h+p)$, the service level in a single source setting when governed by a base-stock level; (viii) $\mu =$ the mean of the Normally distributed one-period demand; (ix) $\sigma =$ the standard deviation of the Normally distributed one-period demand.

The remaining *ten* columns in the table display various output measures as follows:

(a) Cost* = the optimal long-run average cost value within our considered class of dual sourcing policies;

(b) % Savings = the percentage savings realized by the above dual sourcing policy, compared with the optimal single sourcing policy;

(c) S_e and S_r = the optimal base-stock level for the fast and slow supplier, respectively;

(d) % Expedited = the *percentage* of sales that is procured from the expedited supplier;

(e) *Expedite Freq*: the percentage of periods in which an order is placed with the fast supplier (by itself or in combination with a *regular* order).

(f) *Single–Expedited*: the optimal long-run average cost (Cost*) and base-stock level(S_e), assuming the *fast* supplier is used as the *single* supplier;

(g) *Single–Regular*: the optimal long-run average cost (Cost*) and base-stock level(S_r), assuming the *regular* supplier is used as the single supplier.

The first panel in the table investigates the impact of a widening gap between the two lead times, assuming the expedited orders are delivered instantaneously. With a one-period difference in the lead times, only 2.2% of demands are covered by procurements from the fast supplier. Nevertheless, the availability of this fast supply source allows for a cost saving of up to 8.9%. The availability of the expedited supply source, allows one to reduce the base-stock level for the regular supplier from 13.3 to 12.7. Moreover, it allows one to avoid many stockouts under unusually large demands. Orders with the fast supplier are placed in about 1 out of 5 periods.

As the lead time difference increases from one to five periods, the cost savings grows to 30.2%. This extensive cost saving is accomplished, even though, even in this case, little more than 10% of sales are sourced from the fast supplier (and order is placed with the fast supplier in less than 60% of the periods). Across the board, it is significantly cheaper to use the regular supplier as an exclusive source, as opposed to the fast supplier, as long as the difference in the per-unit prices amounts to $10.

In the second panel, we fix the lead time difference to be $\Delta l = l_r - l_e = 2$, but increase l_e from 0 to 4, in increments of 1. The cost saving achieved by dual sourcing is largest in the first instance (16.9%) and monotonically decreases to 3.2%. This pattern is intuitive: in the single source setting, the long-run average inventory-related cost is proportional with the square root of the lead time; therefore, the *differential* in the inventory related costs under exclusive usage of the regular versus the expedited supplier is proportional with $\sqrt{l_e + \Delta l} - \sqrt{l_e} = \Delta l / (\sqrt{l_e + \Delta l} + \sqrt{l_e})$, which decreases monotonically with l_e for a fixed value of Δl. Again, remarkably, the 16.9% cost saving is achieved even while assigning less than 5% of the business to the fast supplier.

In the third panel, we investigate the impact of the cost differential c. When $c = 1$, it is cheaper to use the fast, as compared to the slow, supplier as an exclusive source. However, the opposite is true when $c \geq 2.5$. Either way, significant savings—up to 14.1% in this set of instances—can be achieved with a dual sourcing approach. Note that when $c = 1$, the optimal dual sourcing policy uses the fast supplier for only 20% of the total procurement value and only 81.8% of the time, even though she is the preferred choice in a single sourcing strategy.

In the fourth panel, we vary the value of the backlogging cost rate p, so that the corresponding service level index $SL = p/(h + p)$ varies between 70% and 99%. The cost savings due to dual sourcing increases monotonically from 0.4% to 6.5% as a higher service level is targeted; this is consistent with our explanation above, i.e., that the primary benefit of an expedited supply source

Table 8 Numerical results for settings with fixed cost

	Parameters									Dual Sourcing								Single - Expedited			Single - Regular		
Varying parameters	l_e	l_r	c	K	h	p	SL	μ	σ	Cost*	% Savings	s_e	S_e	s_r	S_r	% Expedited	Expedite Freq	Cost*	s_e	S_e	Cost*	s_r	S_r
Lead time difference l_r-l_e	**0**	**1**	10	10	5	495	99%	5	1	28.13	4.5%	6	7	12	13	0.6%	2.3%				29.46	12	13
	0	**2**	10	10	5	495	99%	5	1	31.89	3.5%	12	12	18	18	1.7%	15.9%				33.05	18	19
	0	**3**	10	10	5	495	99%	5	1	33.95	7.8%	17	18	22	23	4.9%	15.9%	74.42	7	7	36.83	23	25
	0	**4**	10	10	5	495	99%	5	1	35.65	10.1%	22	23	27	28	4.9%	15.9%				39.66	29	30
	0	**5**	10	10	5	495	99%	5	1	37.37	12.2%	27	28	32	33	4.9%	15.9%				42.55	35	36
Lead times l_r, l_e (fixing l_r-l_e)	**0**	**2**	10	10	5	495	99%	5	1	31.89	3.5%	12	12	18	18	1.7%	15.9%	74.42	7	7	33.05	18	19
	1	**3**	10	10	5	495	99%	5	1	35.97	2.4%	17	18	23	24	0.6%	2.3%	79.51	12	13	36.83	23	25
	2	**4**	10	10	5	495	99%	5	1	38.83	2.1%	23	24	29	30	0.6%	2.3%	83.10	18	19	39.66	29	30
	3	**5**	10	10	5	495	99%	5	1	42.49	0.2%	28	29	35	36	0.0%	0.1%	86.88	23	25	42.55	35	36
	4	**6**	10	10	5	495	99%	5	1	44.61	0.8%	34	35	40	41	0.6%	2.3%	89.71	29	30	44.95	40	41
Unit cost c	2	4	**1**	10	5	495	99%	5	1	38.05	0.0%	18	19	18	19	100.0%	100.0%	38.05	18	19			
	2	4	**2.5**	10	5	495	99%	5	1	38.47	3.0%	23	24	28	29	4.9%	15.9%	45.56	18	19			
	2	4	**5**	10	5	495	99%	5	1	38.67	2.5%	23	24	29	30	0.6%	2.3%	58.07	18	19	39.66	29	30
	2	4	**10**	10	5	495	99%	5	1	38.83	2.1%	23	24	29	30	0.6%	2.3%	83.10	18	19			
	2	4	**20**	10	5	495	99%	5	1	39.16	1.3%	23	24	29	30	0.6%	2.3%	133.14	18	19			
Fixed cost K	2	4	10	**0**	5	495	99%	5	1	28.41	4.2%	24	24	30	30	1.7%	15.9%	73.10	18	19	29.66	29	30
	2	4	10	**10**	5	495	99%	5	1	38.83	2.1%	23	24	29	30	0.6%	0.02	83.10	18	19	39.66	29	30
	2	4	10	**20**	5	495	99%	5	1	49.07	1.2%	23	24	29	30	0.6%	0.02	93.09	17	19	49.66	29	30
	2	4	10	**30**	5	495	99%	5	1	56.62	0.0%	20	27	28	35	0.0%	0.00	101.08	17	24	56.63	28	35
	2	4	10	**40**	5	495	99%	5	1	61.63	0.0%	20	27	28	35	0.0%	0.00	106.09	17	24	61.65	28	35
Service level $SL = p/(h+p)$	2	4	10	10	5	**7.5**	**60%**	5	1	20.86	0.0%	14	17	23	26	0.0%	0.0%	68.70	13	16	20.86	23	26
	2	4	10	10	5	**20**	**80%**	5	1	25.65	0.0%	19	21	25	27	0.1%	0.1%	72.64	15	16	25.65	25	27
	2	4	10	10	5	**45**	**90%**	5	1	29.59	0.1%	20	22	26	28	0.1%	0.1%	75.37	16	17	29.62	26	28
	2	4	10	10	5	**95**	**95%**	5	1	33.14	0.1%	21	23	27	29	0.1%	0.1%	77.94	16	18	33.18	27	29
	2	4	10	10	5	**495**	**99%**	5	1	38.83	2.1%	23	24	29	30	0.6%	2.3%	83.10	18	19	39.66	29	30
Demand volatility σ	2	4	10	10	5	495	99%	5	**0.1**	14.99	0.0%	20	20	26	26	0.0%	0.0%	65.00	12	16	14.99	22	26
	2	4	10	10	5	495	99%	5	**0.5**	25.52	0.0%	21	22	27	28	0.0%	0.0%	71.55	15	17	25.53	26	28
	2	4	10	10	5	495	99%	5	**1.0**	38.83	2.1%	23	24	29	30	0.6%	2.3%	83.10	18	19	39.66	29	30
	2	4	10	10	5	495	99%	5	**1.5**	53.09	2.2%	25	26	31	32	3.1%	9.2%	94.61	20	21	54.29	31	33
	2	4	10	10	5	495	99%	5	**2.0**	66.97	2.9%	26	28	33	35	1.5%	2.8%	105.98	21	23	68.99	33	36

Note. Setting: Normal demands for 100,000 periods; full backlogging.

is to avoid stockouts without acquiring excessive regular inventory investments. Similarly, both the optimal base-stock levels S_e and S_r increase monotonically with SL. Comparing the two extreme instances in this panel, when an SL level of 70% is targeted, an order is placed with the fast supplier in only 5% of the periods as opposed to 27.6% of the periods when $SL = 99\%$.

Dependence on the demand volatility is assessed by varying the coefficient of variation of the demand distribution from 2% to 40%, with the base value of 20% as its midpoint. In this case, the relative cost saving remains almost constant, at the approximate level of 6.5%. Similarly, the frequency with which orders are placed with the fast supplier is approximately 27.6% across all instances, but the average order *quantity* with the fast supplier increases (somewhat sublinearly) with the value of σ.

We now turn to Table 8 where a fixed order cost K is added in each of the evaluated scenarios. The format of the table is identical to the previous one, except that we now report the optimal values of the reorder and order-up-to level (s_e, S_e, s_r, S_r) are now reported. We choose $K = 10$ as the base value for the fixed order cost.

In general, we note that the average cost values increase significantly when comparing corresponding instances in the two tables. As a consequence, the *relative* cost savings due to dual sourcing are invariably lower when the fixed order costs are added.

In the first panel, we again vary l_r, the lead time for the regular supplier from $l_r = 1$ to $l_r = 5$. The long-run average cost increases monotonically from 28.13 to 37.37, as does the cost savings due to dual sourcing which increases from 4.5% to 12.2%. (In the absence of fixed order costs, these cost savings vary from 8.9% to 30.2%, in the corresponding first panel of Table 7.) In the presence of fixed order cost, there is a reduced incentive to place parallel orders with both suppliers. For example, when $l_r = 1$, the fast supplier is retained in only 2.3% of the periods, as opposed to 19.3% of the time when no fixed costs are involved. When $l_r = 5$, the corresponding frequencies are 15.9% versus 58.1%.

In the second panel, we maintain a lead time difference $\Delta l = 2$, but increase l_e (and hence l_r) by one period at a time. Quite intuitively, the long-run average cost value as well as all 4 policy parameters (s_e, S_e, s_r, S_r) increase monotonically with the lead times. The relative cost savings due to dual sourcing now decrease from 3.5% to 0.8%, as l_e and $l_r = l_e + 2$ are increased. These cost saving *percentages* are, again, significantly lower than for the corresponding instances in the second panel of Table 7 (16.9% and 3.2%, respectively).

In the third panel, we gauge the impact of the unit price differential. As in the case where $K = 0$, see Table 7, when $c = 1$, the fast supplier is the better single source. Moreover, with a fixed order cost $K = 10$, no cost savings can be achieved with a dual sourcing policy, while a 12.5% cost improvement is possible when $K = 0$. When $c \geq 2.5$, the "regular" supplier is cheaper than the fast supplier in a single source setting, but in these cases additional cost savings of up to 3% are possible with a dual sourcing strategy.

In the fourth panel, we vary the value of K, the fixed order cost, from $K = 0$ to $K = 40$, in increments of \$10, all while maintaining $l_e = 2$ and $l_r = 4$. When $K = 0$, the cost savings due to dual sourcing amount to 4.2%. With $K = 10$, the cost savings reduce to 2.1%. When $K \geq 30$, the fixed order costs make it unattractive to split orders, *any time*. In this case, the optimal dual sourcing policy relies exclusively on the regular supplier and achieves no cost savings, whatsoever. Clearly, as K increases, so does the long-run average cost value; the same monotonicity applies to $\kappa = S_r - s_r = S_e - s_e$.

The impact of the backlogging cost rate p is gauged in the fifth panel, where $SL = p/(p+h)$ is varied between 60% and 99%. When a level of $SL = 60\%$ is targeted, there is no use for the expedited supplier whose main function is to protect against surges in the demand stream. Usage of the fast supplier increases monotonically with SL (and hence with p). In the base case where $SL = 99\%$, this implies a 2.1% cost saving due to dual sourcing.

Finally, the sixth and last panel assess the impact of the demand volatility, as measured by the coefficient of variation of the demand distribution. Clearly, the optimal cost value increases monotonically with σ, as does the cost saving percentage due to dual sourcing. With a low demand volatility, i.e., coefficient of variation ≤ 0.1, there is no use for the expedited supplier. In other words, the risk of large demand surges is not large enough to justify placing orders, ever, with the fast supplier. Recall that, in this panel, the cost savings due to dual sourcing are 6.5% when $K = 0$; they are less than 2.9% when $K = 10$.

7. Concluding Remarks

This paper addresses a general periodic review model, to identify an optimal procurement policy in the presence of *two* suppliers, differentiated by their lead time and per-unit cost price. The model allows for salvage opportunities, capacity limits and fixed costs associated with orders and salvage transactions. We have provided a full characterization of the optimal procurement strategy in the general model as well as various special cases that arise when only part of the above complications prevail. Our exact results are confined to the case where the lead times of the two suppliers differ by a single period only (Even in the base model without fixed cost, capacity limits or salvage opportunities, it is well known that only a one-period lead time difference allows for an optimal policy that acts on a single inventory measure.) However, our structural results for this special case suggest effective heuristics for general lead time combinations, as demonstrated in Section 6.

Indeed, significant cost savings can be achieved with dual sourcing. Remarkably, such savings can be obtained even when the expedited supplier is only used for a small part, say 5% of all procurement. The availability of the fast(er) supplier allows one to forgo major inventory investments to achieve a given service level or to prevent costly stockouts under a given inventory investment.

In practice, a second expedited supplier may only be available if a minimum sales volume can be guaranteed, or a minimum frequency with which orders are placed. Future work should address such “participation” constraints.

It is also important to extend our results to settings with more than two suppliers. Generalizations of our structural results are possible when the lead time of any given supplier differs from that of the next fast supplier by a single period. This extension was carried out by Feng et al. (2005) in the base model. Moreover, we believe that these results continue to suggest effective heuristic strategies under general lead time differences.

References

Allon G, Van Mieghem JA (2010a) Global dual sourcing: Tailored base-surge allocation to near- and offshore production. *Management Science* 56(1):110–124, URL http://dx.doi.org/10.1287/mnsc.1090.1099.

Allon G, Van Mieghem JA (2010b) The mexico-china sourcing game: Teaching global dual sourcing. *INFORMS Transactions on Education* 10(3):105–112, URL http://dx.doi.org/10.1287/ited.1100.0045.

Arrow KJ, Harris T, Marschak J (1951) Optimal inventory policy. *Econometrica* 19(3):250–272, ISSN 00129682, 14680262, URL http://www.jstor.org/stable/1906813.

Barankin EW (1961) A delivery-lag inventory model with an emergency provision (the single-period case). *Naval Research Logistics Quarterly* 8(3):285–311, ISSN 1931-9193, URL http://dx.doi.org/10.1002/nav.3800080310.

Beyer D, Ward J (2002) *Network Server Supply Chain at HP: A Case Study*, chapter 8, 257–282 (Boston, MA: Springer US), ISBN 978-1-4757-6635-6, URL http://dx.doi.org/10.1007/978-1-4757-6635-6_8.

Boute RN, Van Mieghem JA (2015) Global dual sourcing and order smoothing: The impact of capacity and lead times. *Management Science* 61(9):2080–2099, URL http://dx.doi.org/10.1287/mnsc.2014.1992.

Chen S (2004) The infinite horizon periodic review problem with setup costs and capacity constraints: A partial characterization of the optimal policy. *Operations Research* 52(3):409–421, URL http://dx.doi.org/10.1287/opre.1030.0104.

Chen S, Lambrecht M (1996) X-y band and modified (s, s) policy. *Operations Research* 44(6):1013–1019, URL http://dx.doi.org/10.1287/opre.44.6.1013.

Chen X, Simchi-Levi D (2004) Coordinating inventory control and pricing strategies with random demand and fixed ordering cost: The finite horizon case. *Operations Research* 52(6):887–896, URL http://dx.doi.org/10.1287/opre.1040.0127.

Daniel KH (1963) A delivery-lag inventory model with emergency. *Multistage Inventory Models and Techniques* .

Federgruen A, Liu Z, Lu L (2017) Synthesis and generalization of structural results in inventory management: Generalized convexity properties, URL https://ssrn.com/abstract=3096735, working paper.

Federgruen A, Zipkin P (1986a) An inventory model with limited production capacity and uncertain demands i. the average-cost criterion. *Mathematics of Operations Research* 11(2):193–207, URL http://dx.doi.org/10.1287/moor.11.2.193.

Federgruen A, Zipkin P (1986b) An inventory model with limited production capacity and uncertain demands ii. the discounted-cost criterion. *Mathematics of Operations Research* 11(2):208–215, URL http://dx.doi.org/10.1287/moor.11.2.208.

Feng Q, Gallego G, Sethi SP, Yan H, Zhang H (2005) Periodic-review inventory model with three consecutive delivery modes and forecast updates. *Journal of Optimization Theory and Applications* 124(1):137–155, ISSN 1573-2878, URL http://dx.doi.org/10.1007/s10957-004-6469-6.

Fox EJ, Metters R, Semple J (2006) Optimal inventory policy with two suppliers. *Operations Research* 54(2):389–393, URL http://dx.doi.org/10.1287/opre.1050.0229.

Fukuda Y (1964) Optimal policies for the inventory problem with negotiable leadtime. *Management Science* 10(4):690–708, URL http://dx.doi.org/10.1287/mnsc.10.4.690.

Gallego G, Scheller-Wolf A (2000) Capacitated inventory problems with fixed order costs: Some optimal policy structure. *European Journal of Operations Research* 126:603–613.

Hua Z, Yu Y, Zhang W, Xu X (2015) Structural properties of the optimal policy for dual-sourcing systems with general lead times. *IIE Transactions* 47(8):841–850, URL http://dx.doi.org/10.1080/0740817X.2014.982839.

Janakiraman G, Seshadri S, Sheopuri A (2015) Analysis of tailored base-surge policies in dual sourcing inventory systems. *Management Science* 61(7):1547–1561, URL http://dx.doi.org/10.1287/mnsc.2014.1971.

Moinzadeh K, Schmidt CP (1991) An (s 1, s) inventory system with emergency orders. *Operations Research* 39(2):308–321, URL http://dx.doi.org/10.1287/opre.39.2.308.

Neuts MF (1964) An inventory model with an optional time lag. *Journal of the Society for Industrial and Applied Mathematics* 12(1):179–185, URL http://dx.doi.org/10.1137/0112016.

PricewaterhouseCoopers (2010) Global sourcing: Shifting strategies. URL https://www.pwc.at/en/publikationen/branchen-und-wirtschaftsstudien/global-sourcing-shifting-strategies.pdf.

Rao U, Scheller-Wolf A, Tayur S (2000) Development of a rapid-response supply chain at caterpillar. *Operations Research* 48(2):189–204, URL http://dx.doi.org/10.1287/opre.48.2.189.12380.

Scarf H (1960) The optimality of (s, S) policies in the dynamic inventory problem. Arrow K, Karlin S, Suppes P, eds., *Mathematical Methods in the Social Sciences*, 196–202 (Stanford University Press, Stanford, CA).

Scheller-Wolf A, Veeraraghavan S, van Houtum GJ (2007) Effective dual sourcing with a single index policy. *Working Paper, Tepper School of Business, Carnegie Mellon University, Pittsburgh.*

Semple J (2007) Note: Generalized notions of concavity with an application to capacity management. *Operations Research* 55(2):284–291, URL http://dx.doi.org/10.1287/opre.1060.0357.

Sethi SP, Yan H, Zhang H (2003) Inventory models with fixed costs, forecast updates, and two delivery modes. *Operations Research* 51(2):321–328, URL http://dx.doi.org/10.1287/opre.51.2.321.12777.

Sheopuri A, Janakiraman G, Seshadri S (2010) New policies for the stochastic inventory control problem with two supply sources. *Operations Research* 58(3):734–745, URL http://dx.doi.org/10.1287/opre.1090.0799.

Song JS, Xiao L, Zhang H, Zipkin P (2017) Optimal policies for a dual-sourcing inventory problem with endogenous stochastic lead times. *Operations Research* 65(2):379–395, URL http://dx.doi.org/10.1287/opre.2016.1557.

Song JS, Zipkin P (2009) Inventories with multiple supply sources and networks of queues with overflow bypasses. *Management Science* 55(3):362–372, URL http://dx.doi.org/10.1287/mnsc.1080.0941.

Sun J, Van Mieghem JA (2017) Robust dual sourcing inventory management: Optimality of capped dual index policies, URL https://ssrn.com/abstract=2991250, available at SSRN: https://ssrn.com/abstract=2991250.

The Economist (2013) Here, there and everywhere. *The Economist*, volume Jan 19, 2013, URL https://www.economist.com/news/special-report/21569572-after-decades-sending-work-across-world-companies-are-rethinking-their-offshoring.

Van Mieghem JA, Allon G (2015) *Operations Strategy: Principles and Practice* (Dynamic Ideas LLC), 2nd edition, ISBN 978-0-9899108-6-6.

Veeraraghavan S, Scheller-Wolf A (2008) Now or later: A simple policy for effective dual sourcing in capacitated systems. *Operations Research* 56(4):850–864, URL http://dx.doi.org/10.1287/opre.1080.0552.

Whittemore AS, Saunders SC (1977) Optimal inventory under stochastic demand with two supply options. *SIAM Journal on Applied Mathematics* 32(2):293–305, URL http://dx.doi.org/10.1137/0132023.

Xin L, Goldberg DA (2017) Asymptotic optimality of tailored base-surge policies in dual-sourcing inventory systems. *Management Science* URL http://dx.doi.org/10.1287/mnsc.2016.2607.

Xin L, He L, Bewli J, Bowman J, Feng H, Qin Z (2017) On the performance of tailored base-surge policies: theory and application at walmart.com, URL https://ssrn.com/abstract=3090177, working paper.

Zhou SX, Yang C (2016) Continuous-review (r, nq) policies for inventory systems with dual delivery modes. *Operations Research* 64(6):1302–1319, URL http://dx.doi.org/10.1287/opre.2016.1538.

Zipkin P (2000) *Foundations of Inventory Management* (McGraw-Hill).

Zipkin P (2008) On the structure of lost-sales inventory models. *Operations Research* 56(4):937–944, URL http://dx.doi.org/10.1287/opre.1070.0482.