J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model

Joren Gijsbrechts
Universidade Católica Portuguesa, Católica Lisbon School of Business and Economics, Portugal

Robert N. Boute
Vlerick Business School, KU Leuven and Flanders Make, Belgium

Stephen M. Disney
University of Exeter Business School, United Kingdom

Jan A. Van Mieghem
Kellogg School of Management, Northwestern University, United States

June 10, 2021; revised December 9, 2021; revised January 28, 2022; accepted March 3, 2022.

We investigate how volume flexibility, defined by a sourcing cost premium beyond a base capacity, at a local responsive supplier impacts the decision to reshore supply. The buyer also has access to a remote supplier that is cheaper with no restrictions on volume flexibility. We show that with unit lead time difference between both suppliers, the optimal dual sourcing policy is a *modified dual base-stock* policy with three base-stock levels S_2^f, S_1^f, and S^s. The replenishment orders are generated by first placing a base order from the fast supplier of at most k units to raise the inventory position to S_1^f, if that is possible. After this base order, if the adjusted inventory position is still below S_2^f, additional units are ordered from the fast supplier at an overtime premium to reach S_2^f. Finally, if the adjusted inventory position is below S^s, an order from the slow supplier is placed to bring the final inventory position to S^s. Surprisingly, in contrast to single sourcing with limited volume flexibility, a more complex dual sourcing model often results in a "simpler" policy that replaces demand in each period. The latter allows analytical insights into the sourcing split between the responsive and the remote supplier. Our analysis shows how increased volume flexibility at the responsive supplier promotes the decision to reshore operations and effectively serves as a cost benefit. It also shows how investing in base capacity or additional volume flexibility act as strategic substitutes.

Key words: Dual Sourcing, Flexibility, Reshoring, Optimal Policy, Modified Dual Base-Stock

1. Introduction

Trade tensions—with the recent trade war between the United States and China as a prime example—and disruptions such as the Covid-19 pandemic, congestion in global shipping lanes, and truck driver shortages are forcing companies to rethink their offshore sourcing strategy in favor of more agility by sourcing locally. Sourcing simultaneously from two sources (or countries) is referred to as *dual sourcing*, a classic inventory problem that combines fast, responsive replenishment from a more expensive source with slow sourcing from a more economic one.

Most dual sourcing models assume linear sourcing costs for both suppliers. This implies both suppliers have perfect volume flexibility and can produce without any quantity limitations. These

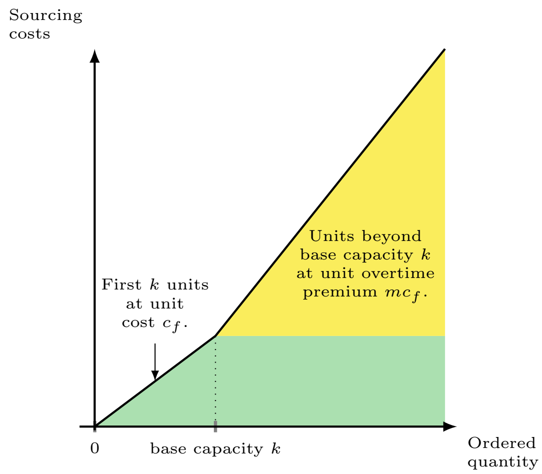

Figure (1) Piece-wise linear and convex sourcing costs. The first k units are charged a unit cost c_f; units exceeding k cost mc_f per unit (where $m \geq 1$ stands for the overtime multiplier).

assumptions may not hold in practice. Many companies have a workforce that can only be utilized within regular working hours, imposing a limit on the daily production quantity. Increasing the base workforce increases the supplier's capacity but also leads to higher unit labor cost if worker utilization is reduced. Alternatively, firms may produce beyond their daily limit by offering the existing workforce overtime or by hiring more expensive external temporary workers (agency staff). Both practices limit the volume flexibility to source any quantity at equal unit cost. Stronger legislation related to overtime awards labor working beyond regular working hours with an overtime premium; the additional flexibility of hiring temporary workers also comes at a cost premium.

Such *limited volume flexibility* at the local responsive supplier may impact the decision to reshore operations. We study the impact of the piece-wise linear sourcing cost at the responsive supplier, where sourcing beyond its base capacity incurs an overtime premium, as shown in Figure 1. The marginal cost of units sourced within the base capacity, k, is c_f. When the responsive supplier produces beyond its base capacity k, the marginal cost of units sourced increases to mc_f, with overtime multiplier, $m \geq 1$. The multiplier m acts as a flexibility *reduction* coefficient.

When $m = 1$, the responsive supplier's workforce has perfect flexibility as all units may be produced at a unit cost of c_f. Increasing m beyond unity limits volume flexibility. In the absence of a slow supplier (local single sourcing), it modifies the optimal replenishment policy of the buyer from a conventional *base-stock* policy, in which the inventory position before ordering is raised up to a base-stock level S [Karlin and Scarf 1958], to a *modified base-stock* policy [Porteus 1990, Martínez de Albéniz and Simchi-Levi 2005] (see top panels of Figure 2). The latter introduces an additional base-stock level: the base capacity is used to (try to) raise the inventory position up to the higher base-stock level S_1; if this raises the inventory position above the lower base-stock level

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

S_2, no overtime is used. Else, the base capacity and overtime are used to raise the inventory position up to S_2. This creates a region of *inaction* where demand is not fully replaced and only k units are ordered. In contrast to the conventional base-stock policy, where the order quantity always equals the demand of the past period when demand is stationary and iid, the single sourcing modified base-stock policy is not a demand replacement policy. As a result, the optimal policy parameters (and related costs) of the modified base-stock policy can only be obtained numerically.

When a buyer has access to two suppliers that both have linear sourcing costs (i.e, both have perfect volume flexibility) with the slow supplier charging a lower cost $c_s \leq c_f$, and the lead time difference between both suppliers equals one, a *dual base-stock* policy is optimal [Fukuda 1964]. Dual base-stock policies have two base-stock levels; one for the fast supplier, S^f, and one for the slow supplier, S^s. If the inventory position before ordering is below S^f, a fast order is placed to raise the inventory position up to S^f. After the fast supplier order is added to the inventory position, the order size with the slower source is determined in a similar way using S^s (see panel (c) of Figure 2).

We extend the above results by investigating a dual sourcing system where the responsive supplier has a base capacity with limited volume flexibility. We show that a *modified dual base-stock* policy with three base-stock levels (see panel (d) of Figure 2) minimizes a buyer's cost in such a dual sourcing system with unit lead time difference between both suppliers. Orders from the responsive supplier follow a modified base-stock policy, with base-stock levels S_1^f and S_2^f. If, after ordering from the fast supplier, the *adjusted* inventory position does not exceed the slow base-stock level S^s, the slow order is used to raise the inventory position up to the slow base-stock level. Interestingly, a more *complex* model often[1] results in a *simpler* optimal policy that *is* a demand replacement policy. This allows for analytical expressions of the optimal policy parameters and related costs. The analytic tractability reveals how limited volume flexibility delays—or postpones—replenishments from the responsive supplier to the slow supplier. In addition to quantifying the impact of volume flexibility on the optimal replenishment structure of the buyer, our analysis supports a deeper analysis on how additional base capacity or flexibility investments allows the responsive supplier to absorb more demand variability and reshore more supply. The magnitude of this increase grows with a higher demand uncertainty. Base capacity or volume flexibility at the responsive supplier are strategic substitutes: We show to what extent the base capacity should be increased for a given overtime premium (or vice versa) in order to reshore a target sourcing volume. We also demonstrate how more volume flexibility effectively functions as a (cost) benefit: Keeping the average sourcing cost per unit sourced identical, more will be reshored when the fast supplier has

[1] In §5 we derive the sufficient conditions.

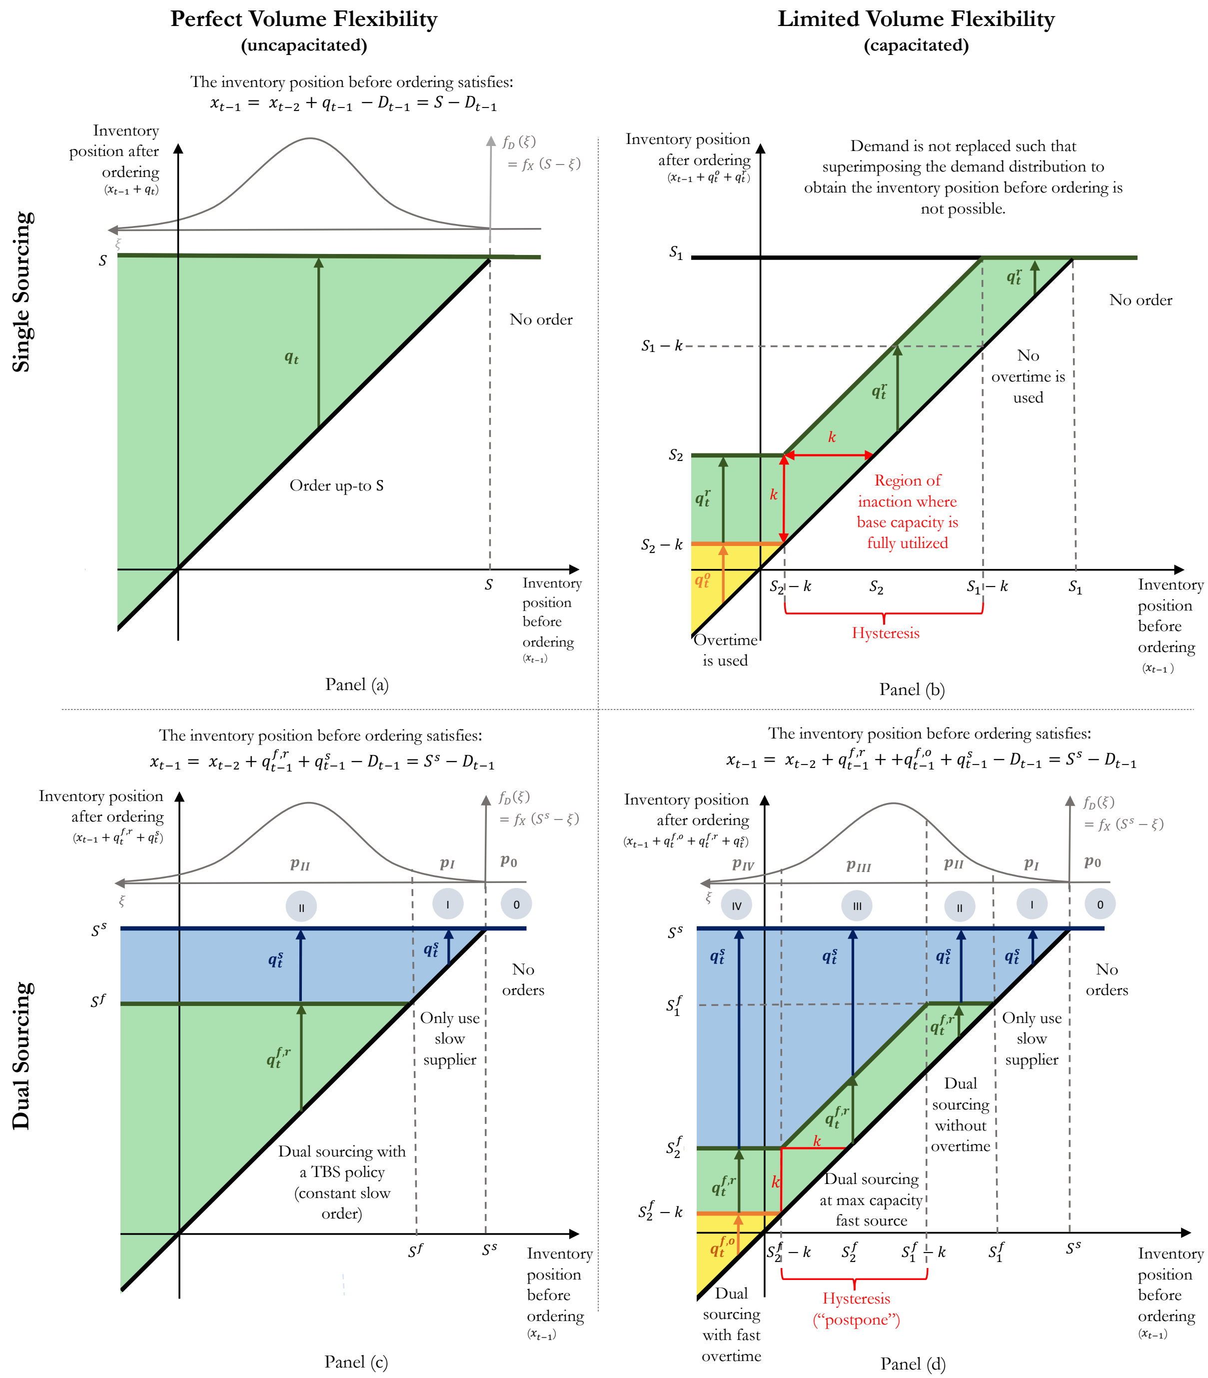

Figure (2) **The optimal single sourcing policy with linear costs (perfect volume flexibility) is a base-stock policy, panel (a). When the sourcing costs are piece-wise linear convex (limited volume flexibility), a modified base-stock policy is optimal, panel (b). The optimal dual sourcing policy with linear sourcing costs (perfect volume flexibility) at both suppliers and unit lead time difference between both suppliers is a dual base-stock policy, panel (c). When local sourcing costs are piece-wise linear and convex (limited volume flexibility), a modified dual base-stock policy is optimal, panel (d).**

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

linear sourcing costs compared to when its supply is characterized by a piece-wise linear sourcing cost. We acknowledge that reshoring decisions involve many issues, such as quality, lead time uncertainty due to port operations or import tariffs. However, the focal point of our paper is to study the impact of restricted volume flexibility that is often prevalent to local, reshored, supply.

We position our work in §2, formulate our model in §3, characterize and prove the optimal policy and parameters in §4 and §5 respectively, and provide insights on how volume flexibility at the responsive supplier impacts the sourcing split and costs in §6. Throughout this article we adopt the following notation. Arguments of functions are given in square brackets $[\cdot]$. The cumulative distribution function (CDF) of a random variable X is given by $F_X[x] = \mathbb{P}[X \leq x]$; its density by $f_X[x]$. The tail distribution is given by $\bar{F}_X[x] = 1 - F_X[x]$ while the inverse CDF, $F_X^{-1}[p]$, corresponds to the value x for which $F_X[x] = p$ holds. To avoid notational clutter we drop the subscript when it is unambiguous which random variable is used. The minimum and maximum operators are given by $a \wedge b = \min\left\{a, b\right\}$ and $a \vee b = \max\left\{a, b\right\}$, respectively. The positive part operator $\left[a\right]^+ = \max\left\{a, 0\right\}$ and by definition, $\sum_{i=1}^{0} = 0$.

2. Related literature

The seminal work of Karlin and Scarf [1958] considers a single sourcing setting with sourcing costs that are linear in the ordered volume, inventory mismatch costs that are convex in the net stock levels, and unmet demand that can be backlogged. They show a base-stock policy minimizes expected costs. When sourcing costs are convex, Karlin [1958] shows the optimal base-stock level becomes state-dependent and is non-decreasing with respect to the inventory levels before order placement. Intuitively, when inventory is low it is better to postpone (a part of) the order, hereby incurring an additional backlog cost instead of fully replacing the demand at the more expensive unit cost. Porteus [1990, p662] refers to this as a *generalized base-stock* policy and shows there are finitely many base-stock levels if the ordering cost is piece-wise linear and convex. We adopt the more recent moniker of a *modified base-stock policy* (MBS) [Martínez de Albéniz and Simchi-Levi 2005].

Whereas the optimal single sourcing policy structure is known for models with convex sourcing costs, the optimal base-stock levels cannot be captured in closed form, and have to be obtained through numerical analysis. Lu and Song [2014] use dynamic programming to obtain the optimal base-stock levels. Martínez de Albéniz and Simchi-Levi [2005] demonstrate how to reduce the computational effort by limiting the search space. Henig et al. [1997] investigate a single sourcing model where a buyer commits to base volume (capacity) in advance and the first k units can be purchased at no cost. The optimal inventory policy of their model satisfies a MBS policy. As they employ a single sourcing system, they have to numerically obtain the base-stock levels and expected cost to assess whether the investment in capacity offsets the benefits of having contracted volume at no additional cost (even under a zero lead assumption).

The optimal policy structure of the dual sourcing problem is only known in limited settings. During the early 1960s, several works characterize and prove the optimality of *Dual Base-Stock* policies when the lead times of the fast and slow supplier equal zero and one, respectively [Barankin 1961, Neuts 1964, Bulinskaya 1964a,b]. They adopt conventional assumptions; the sourcing costs of both suppliers are linear in the sourcing quantities and the inventory mismatch function is convex. Fukuda [1964] extends this result towards general consecutive lead times $l_s = l_f + 1$, with l_s and l_f being the lead time of the slow and fast supplier, respectively.[2]

Whittemore and Saunders [1977] show the optimal policy structure is no longer a (dual) base-stock type policy for larger lead time gaps. To circumvent the complexity of tracking the entire pipeline vector, the single index policy of Scheller-Wolf et al. [2003] simply tracks total inventory, which is the sum of the on-hand inventory and units in-transit. Dual index policies include two inventory positions, the sum of the inventory on hand and outstanding orders within the lead time from the fast and slow source respectively [Veeraraghavan and Scheller-Wolf 2008]. Sun and Van Mieghem [2019] show that the *Capped Dual Index* policy is robustly optimal for general lead times, i.e., it minimizes the worst case performance across a range of deterministic demand scenarios, and show numerically this policy also performs well in a stochastic setting. Their policy is a dual base-stock dual index policy that places a cap on slow orders that naturally smooths the variability of the slow orders. Xin and Goldberg [2017] show that when the lead time gap grows to infinity, the *Tailored Base-Surge* (TBS) policy of Allon and Van Mieghem [2010] is asymptotically optimal. The TBS policy places a constant order from the slow source while a base-stock policy is used to control the fast source. TBS is a capped dual index policy.

Studies that extend the linear sourcing cost assumption in dual sourcing systems are scarce. Federgruen et al. [2021] find optimal dual sourcing policies for capacitated dual sourcing systems using (C_1, C_2, K_1, K_2)-convexity. They show when the fast source is capacitated, a capped base-stock policy is optimal for the fast supplier: order up to the base-stock level if you can, else order the cap. We extend this case by showing the value of adding volume flexibility to the fast supplier. Tomlin [2006] includes volume flexibility in a dual sourcing system but assumes equal lead times with one unreliable supplier and one reliable (but more expensive) supplier with volume flexibility. Our model focuses on the impact of volume flexibility at responsive suppliers in reshoring decisions, while other aspects are left out of scope to retain focus. Chen and Hu [2017], for instance, analyses the impact of offshore dependence of the onshore supplier.

[2] We will denote the local responsive supplier the *fast* supply with subscript (or superscript) f and the remote supplier the *slow* supply with subscript (or superscript) s.

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

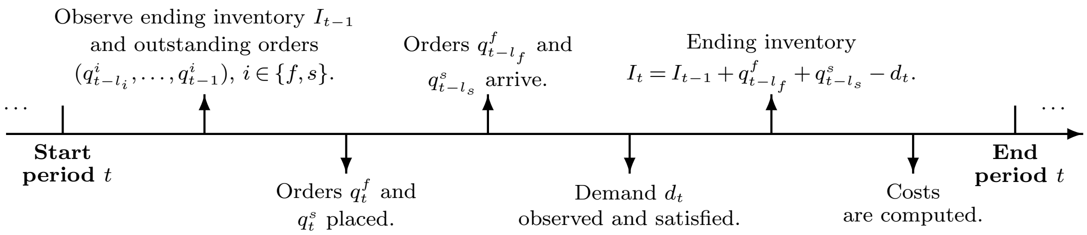

Figure (3) Sequence of events in dual-source supply chains with lead times l_f/l_s for fast/slow supply.

Boute and Van Mieghem [2014] use the MBS policy in a single sourcing setting to motivate a linear smoothing heuristic in a dual sourcing model where the fast source has an installed capacity (i.e., a capacity for which one always has to pay regardless of the ordered quantity). In the same setting, Boute et al. [2021] adopt a constant order policy for the slow supplier, combined with a proportional base-stock policy for the fast supply. The proportional base-stock policy enables one to smooth the orders to avoid expensive overtime production beyond the installed capacity. Although their heuristic is not optimal, it is analytically tractable, even for correlated and non-stationary demand. In contrast, we identify the optimal policy (for consecutive lead times). Whereas Boute and Van Mieghem [2014] and Boute et al. [2021] impose a sunk cost kc_f and set the marginal cost of units produced within the base capacity equal to zero, we assume local units sourced below the base capacity k scale linearly in the ordered volume, with $c_f \geq c_s$.

3. Model Formulation

Let I_{t-1} denote the buyer's net stock at the end of period $t-1$, which is the inventory on-hand minus the backorders. The sequence of events in period t is visualized in Figure 3. First, based on the previous period's ending inventory, I_{t-1}, and outstanding orders $(q^i_{t-l_i}, \ldots, q^i_{t-1})$, $i \in \{f, s\}$, we decide how much to order from the fast supplier, q^f_t, and from the slow supplier, q^s_t. Second, orders placed l_f periods ago from the fast supplier $q^f_{t-l_f}$, and l_s periods ago from the slow supplier $q^s_{t-l_s}$, arrive, and are added to the inventory. Third, demand ξ_t is observed and satisfied, resulting in the following inventory balance equation:

$$I_t = I_{t-1} + q^f_{t-l_f} + q^s_{t-l_s} - \xi_t. \tag{1}$$

Finally, we tally the costs of period t. These consist of the sourcing costs paid to both suppliers and the inventory mismatch costs. Although we assume the sourcing cost of the fast supply to be piece-wise linear (as in Figure 1), sourcing costs for the slow supplier are linear[3]:

$$C_t[I_t, q^s_t, q^f_t] \triangleq q^s_t c_s + c_f\big(q^f_t \wedge k\big) + mc_f[q^f_t - k]^+ + h[I_t]^+ + b[-I_t]^+. \tag{2}$$

[3] The latter is motivated by the fact that one can source *any* quantity as long as one is willing to wait. We assume the slow supplier has, effectively, an infinite capacity, perhaps facilitated by multiple outsourcing opportunities.

Here, $q_t^s c_s$ is the cost of purchasing q_t^s units from the slow supplier at marginal cost c_s, $c_f\left(q_t^f \wedge k\right)$ is the cost of producing units from the fast supplier's base capacity at marginal cost c_f, and $mc_f[q_t^f - k]^+$ is the cost of producing units beyond the fast supplier's base capacity at marginal cost mc_f (with overtime multiplier, $m \geq 1$). As conventionally assumed, the inventory mismatch costs are linear in both positive and negative inventory, $h[I_t]^+ + b[-I_t]^+$, where h and b denote the unit inventory holding and backlogging cost respectively. We also define the expected inventory mismatch in period t as

$$L_t[y_t] \triangleq \begin{cases} \int_0^{y_t} h(y_t - \xi_t) f[\xi_t] d\xi_t + \int_{y_t}^{\infty} b(\xi_t - y_t) f[\xi_t] d\xi_t & \text{if } \ y_t > 0 \\ \int_0^{\infty} b(\xi_t - y_t) f[\xi_t] d\xi_t & \text{if } \ y_t \leq 0 \end{cases}, \tag{3}$$

where $y_t = I_{t-1} + q_{t-l_f}^f + q_{t-l_s}^s$ represents the net inventory after order arrival but before the demand is observed, and the demand in period t, ξ_t, is a random variable with density $f[\xi_t]$.

The objective of our dual sourcing problem is to find the purchasing policy that minimizes the long-run average cost. Let Π denote the set of all feasible ordering policies. A feasible policy, π, consists of a sequence of mappings $f_t^\pi : \mathbb{R}^{l_f + l_s + 1} \mapsto \mathbb{R}^2, t \geq 1$. That is, based on the net ending inventory and all outstanding orders, the quantity sourced from each supplier, is determined via[4]

$$(q_t^f, q_t^s) = f_t^\pi[q_{t-l_s}^s, \ldots, q_{t-1}^s, q_{t-l_f}^f, \ldots, q_{t-1}^f, I_{t-1}]. \tag{4}$$

Let C_t^π denote the cost of policy π in period t. The long-run average cost of policy π is given by:

$$C[\pi] \triangleq \lim_{T \to \infty} \frac{1}{T} \sum_{t=1}^{T} \mathbb{E}\left[C_t^\pi\right]. \tag{5}$$

The objective of our dual sourcing problem is to find the policy $\pi \in \Pi$ that minimizes the long-run average cost:

$$C^{\text{OPT}} \triangleq \inf_{\pi \in \Pi} C[\pi]. \tag{6}$$

We shall first derive the policy that minimizes the finite-horizon discounted cost

$$C[\pi, T, \alpha, c_{T+1}] \triangleq \sum_{t=1}^{T} \alpha^t \mathbb{E}\left[C_t^\pi\right], \tag{7}$$

with discount factor $0 < \alpha < 1$ and salvage value for ending inventory of c_{T+1}, and then invoke standard results to show that the policy also minimizes the long-run average cost.

Whittemore and Saunders [1977] show that (4) is a complex function when the lead times of both sources are non-consecutive, i.e., $l_s - l_f > 1$. When lead times are consecutive, however, the problem is greatly simplified. If $l_s - l_f = 1$, the state space of the corresponding dynamic program becomes one-dimensional, representing the dimension of the inventory position (i.e., the sum of inventory on hand and all outstanding orders at the end of period $t-1$).

[4] The optimal orders are also dependent on the distributions of the demands from period t until the end of the horizon T, $(D_t, \ldots, D_T)$, but we omit them here as most of our analysis later on assumes continuous and iid demand.

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

4. Optimal Policy

When the lead time difference between the fast supplier and the slow supplier is one, we will show that the policy that minimizes (7) is a modified dual base-stock policy with a sequence of time-dependent order-up-to levels $\{(S^f_{2,t}, S^f_{1,t}, S^s_t) : t \in \mathbb{N}\}$. The modified dual base-stock policy functions as follows. First, place a base order from the fast supplier of, at most, k units to raise the inventory position to $S^f_{1,t}$, if that is possible. If, after this base order, the adjusted inventory position is still below $S^f_{2,t}$, then order additional units from the fast supplier at an overtime premium to reach $S^f_{2,t}$. Finally, if the adjusted inventory position does not exceed S^s_t, place an order from the slow supplier to bring the final inventory position to S^s_t.

The formal proofs are relegated to the Appendix. We highlight the most important results below, starting with establishing finite-horizon discounted cost optimality:

Theorem 1 *In a dual sourcing setting with consecutive lead times where the fast supplier has a piece-wise linear sourcing cost and the slow supplier's sourcing costs are linear, the policy that minimizes the finite-horizon discounted cost $C[\pi, T, \alpha, c_{T+1}]$ is a modified dual base-stock policy with a sequence of time-dependent order-up-to levels $\{(S^f_{2,t}, S^f_{1,t}, S^s_t) : t = 1, \cdots, T\}$ that define the optimal order quantities at $t = 1, \cdots, T$. The fast order follows a modified base-stock policy:*

$$q^{f,\star}_t = \begin{cases} 0 & \text{if } S^f_{1,t} \le x_{t-1}, \\ S^f_{1,t} - x_{t-1} & \text{if } S^f_{1,t} - k \le x_{t-1} < S^f_{1,t}, \\ k & \text{if } S^f_{2,t} - k \le x_{t-1} < S^f_{1,t} - k, \\ S^f_{2,t} - x_{t-1} & \text{if } x_{t-1} < S^f_{2,t} - k, \end{cases} \tag{8}$$

after which the slow order raises the updated inventory position $(x_{t-1} + q^{f,\star}_t)$ up to S^s_t:

$$q^{s,\star}_t = \left[S^s_t - x_{t-1} - q^{f,\star}_t\right]^+. \tag{9}$$

Theorem 1 is proven by backward induction for a finite horizon of length T (extending Xin and Van Mieghem [2021]). In particular, we obtain the policy with optimal discounted cost-to-go $v_t[x_{t-1}]$ from t until T starting with inventory level x_{t-1} that satisfies the well-known Bellman equation for every period $t \in \{1, \cdots, T\}$:

$$v_t[x_{t-1}] = \min_{q^f_t, q^s_t \ge 0} \left\{ m c_f [q^f_t - k]^+ + c_f (q^f_t \wedge k) + c_s q^s_t + L_t[x_{t-1} + q^f_t] + \alpha \mathbb{E}[v_{t+1}[x_{t-1} + q^f_t + q^s_t - D_t]] \right\}, \tag{10}$$

and terminal value function $v_{T+1}[x_T] = -c_{T+1} x_T$, where $c_{T+1} \ge 0$.

We first show that the value function in period T is jointly convex in the inventory before ordering x_{T-1}, and the order-up-to-level, y^f_T:

Lemma 1 *The value function in period T, $v_T[\cdot]$, is jointly convex in x_{T-1} and y^f_T.*

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

The novel part of our proof is to demonstrate a modified dual base-stock policy is optimal in the penultimate period $T-1$, leveraging convexity of $v_T[\cdot]$. We identify four cases dependent on whether the fast and slow order are positive or zero: the *interior case* where both optimal quantities are positive: $q^{f,\star}_{T-1} > 0$ and $q^{s,\star}_{T-1} > 0$ (Case A). Then we consider the three boundary cases: $q^{\star,f}_{T-1} = 0$ while $q^{s,\star}_{T-1} > 0$ (Case B); $q^{f,\star}_{T-1} > 0$ while $q^{s,\star}_{T-1} = 0$ (Case C); and finally $q^{f,\star}_{T-1} = q^{s,\star}_{T-1} = 0$ (Case D). Determining when the boundary cases arise will have important implications on the possibility to determine the policy parameters explicitly in §5. Combining the four cases reveals that a modified dual base-stock policy is optimal in period $T-1$.

In all four cases, the value function preserves convexity:

Lemma 2 *If the value function in period T is convex and a modified dual base-stock policy is optimal in period $T-1$, then the value function in period $T-1$ remains convex.*

We can follow the same line of argument for period $T-2$ and repeating until period t reveals a modified dual base-stock policy is optimal in every period.

In general, the order quantities are time-dependent. Later we demonstrate that when demand is iid, the horizon is infinite (or a specific salvage cost is chosen for the finite horizon problem), and when the the modified dual base-stock policy is demand replacement, the base-stock levels are stationary and can be obtained easily. While the first two pre-requisites are common assumptions to express base-stock levels in inventory management, the demand replacement property is typically fulfilled implicitly. We will show that for our model, a sufficient condition for demand replacement is that $S^f_{1,t} < S^s_t$, which holds for most problem settings (we provide more details in §5 and Appendix C).

Appendix 2 demonstrates that the necessary conditions listed by Huh et al. [2011] hold for our model such that our results hold when we minimize the infinite-horizon average cost:

Theorem 2 *The modified dual base-stock policy is optimal in the infinite horizon average-cost setting.*

In the remainder of this paper we will characterize the optimal policy for the case with consecutive lead times for both sources. For notational simplicity, we choose a system where $l_f = 0$ and $l_s = 1$. Appendix B details how to retrieve the results for general consecutive lead times.

5. Policy parameters

In this section, we characterize the time-dependent order-up-to levels $\{S^f_{2,t}, S^f_{1,t}, S^s_t\}$, for each period $t \in \{1, \cdots, T\}$. Henceforth, we assume demand is continuous and iid, but we note the same results hold when the demand distribution of each period weakly dominates the distribution of the previous one, i.e., when for each period t: $\forall \xi \in (-\infty, \infty)$ $F_{D_t}[\xi] \leq F_{D_{t+1}}[\xi]$. A sufficient condition to obtain expressions of the base-stock levels is that $S^f_{2,t} < S^f_{1,t} < S^s_t$ (see also Appendix C for a more elaborate discussion). In this case we always order up to the slowest base-stock level so the actual demand of the previous period ξ_{t-1} is always replaced, i.e., $q^f_t + q^s_t = \xi_{t-1}$. The fast base-stock levels then satisfy fractile solutions, according to part A of the proof of Theorem 1, whereas the slow base-stock level then is obtainable by taking first order conditions on the Bellman equation.

Formally, we define two myopic order-up-to levels $S^f_{2,\mathcal{M}}$ and $S^f_{1,\mathcal{M}}$:

$$S^f_{2,\mathcal{M}} \triangleq F_D^{-1}\left[\frac{b-(mc_f-c_s)}{b+h}\right] \leq S^f_{1,\mathcal{M}} \triangleq F_D^{-1}\left[\frac{b-(c_f-c_s)}{b+h}\right], \tag{11}$$

and consider the function $g(x)$ for $x \geq S^f_{1,\mathcal{M}}$:

$$g[x] = (c_s - \alpha c_s) + (c_s - mc_f)\bar{F}_D[x - S^f_{2,\mathcal{M}} + k] + \int_{x-S^f_{1,\mathcal{M}}+k}^{x-S^f_{2,\mathcal{M}}+k} L'[x-\xi+k]f[\xi]d\xi +$$
$$(c_s - c_f)(\bar{F}_D[x - S^f_{1,\mathcal{M}}] - \bar{F}_D[x - S^f_{1,\mathcal{M}} + k]) + \int_0^{x-S^f_{1,\mathcal{M}}} L'[S^s - \xi]f[\xi]d\xi. \tag{12}$$

The function g is the derivative of the optimal value function and, given the value function is convex, g is increasing. Therefore, if a solution exists where $g[x] = 0$, it is unique and quickly found numerically (e.g., using the bisection, Regula-Falsi, or Newton-Rhapson method). If there exists an $x_0 \geq S^f_{1,\mathcal{M}}$ for which $g[x_0] = 0$ holds, then we define the slow order-up-to level $S^s = x_0$. (Otherwise, part C in the proof applies: there may be periods where only the fast supplier is used. Then the optimal policy may not be demand-replacing in every period such that computing the probability of the ending inventory in period t would become dependent on the ending inventory in period $t-1$, complicating analytical tractability.)

We will show, when the terminal salvage cost, $c_{T+1} = c_s/\alpha$ and $S^f_{2,\mathcal{M}} \leq S^f_{1,\mathcal{M}} < S^s$, then $S^f_{2,\mathcal{M}}$, $S^f_{1,\mathcal{M}}$ and S^s are the optimal base-stock levels in every period. We express this result in Corollary 1:

Corollary 1 *If demand is continuous and iid, and $S^f_{2,\mathcal{M}} \leq S^f_{1,\mathcal{M}} < S^s$, then $S^f_{2,\mathcal{M}}$, $S^f_{1,\mathcal{M}}$ and S^s are the optimal base stock levels in each period when the terminal salvage cost is given by $c_{T+1} = c_s/\alpha$.*

Corollary 1 extends to infinite horizon and average cost. Informally, given the policy parameters are independent of the horizon T, they hold for any T and α. They will thus also hold in the limit $T \to \infty$ and $\alpha \to 1$. As with any limiting arguments, a rigorous proof would be much more involved and is beyond the scope of this paper.

A special case of Corollary 1, when $\bar{F}[S^s - S^f_{1,\mathcal{M}}] = 0$, exists when we never source fast. Using (12), determining the slow order then reduces to a newsvendor critical fractile over two periods of demand such that for any period t when $\alpha = 1$: $S_t^s = F_{D_{t,t+1}}^{-1}[b/(h+b)]$, with $D_{t,t+1}$ representing the random variable of two periods of demand.

In what follows we shortly discuss the impact when $S_{2,t}^f < S_{1,t}^f < S_t^s$ no longer holds. First, when $S_t^s < S_{2,t}^f$ the slow supplier will never be used and our model becomes a single sourcing model with piece-wise linear sourcing costs. For this setting we know that no simple expression exists to obtain the base-stock levels, even when demand is iid (see e.g., Porteus [1990], Martínez de Albéniz and Simchi-Levi [2005]). Alternatively, when the slow base-stock level would fall in between the fast base-stock levels ($S_{2,t}^f < S_t^s < S_{1,t}^f$) it is also difficult to obtain an expression for the policy parameters in an expression as the fast orders also impacts the costs in the next period. Yet, the latter scenario is a rather pathological case that rarely occurs. We show this numerically in Appendix C.

To summarise, Theorem 1 may result in four potential ordering strategies: (1) single sourcing from the slow supplier; (2) dual sourcing with demand replacement; (3) dual sourcing without demand replacement; and (4) single sourcing from the fast supplier. Corollary 1 shows that the first two strategies allow for a characterization of the policy parameters under specific conditions. Strategy 3 and 4 require numerical analysis as analytical solutions are not available.

6. The impact of volume flexibility at the responsive supplier

The characterization of the optimal policy (parameters) facilitates a deeper analysis of the impact of the responsive supply's volume flexibility on a buyer's decision to reshore production. In what follows we focus our analysis on the regular dual sourcing settings where a modified dual base-stock policy with demand replacement (and stationary base-stock levels $S_2^f \leq S_1^f < S^s$) is optimal.

We will analyze the sourcing split between the responsive and offshore supply, as well as their impact on the order variability at both suppliers. We also show how the base capacity k and the volume flexibility, measured through the overtime premium m, act as strategic substitutes. Finally, we provide a sensitivity analysis by relaxing one of the main assumptions of our model by numerically investigating the performance of the modified dual base-stock policy when the lead time difference between both sources exceeds one period.

We denote the expected order quantities sourced from the slow supply at cost c_s, the expected order quantities sourced from the fast supply at regular cost c_f, and the expected order quantities sourced from the fast supply at overtime premium mc_f, respectively as $\mu^s \triangleq \mathbb{E}[q^s]$, $\mu_r^f \triangleq \mathbb{E}[q_r^f]$ and $\mu_o^f \triangleq \mathbb{E}[q_o^f]$ (such that $\mu^s + \mu_r^f + \mu_o^f = \mu$ and $\mu_r^f + \mu_o^f = \mu^f$), and the standard deviation of the order quantities placed at the slow and fast supply by σ_{q^s} and σ_{q^f}, respectively. The coefficient of variation of orders placed to both suppliers is given by: $CoV_{q^s} \triangleq \sigma_{q^s}/\mu_s$ and $CoV_{q^f} \triangleq \sigma_{q^f}/\mu_f$. The *sourcing split* is defined as the long-run or average fraction sourced from each supply option:

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

(1) the fraction sourced from the slow supply, $\rho^s \triangleq \mu^s/\mu$, (2) the fraction sourced from the fast supply at regular cost, $\rho^f_r \triangleq \mu^f_r/\mu$, and (3) the fraction sourced from the fast supply at overtime cost $\rho^f_o \triangleq \mu^f_o/\mu$. See also Appendix B for the explicit expressions of the sourcing splits.

Throughout the remainder of our analysis we will use p_i to refer to the probability that the inventory position before ordering $(S^s - D)$ falls in region i, with $i \in \{0, \text{I}, \text{II}, \text{III}, \text{IV}\}$. Because we investigate the modified dual base-stock policy with demand replacement, superimposing the demand distribution over the regions (see panel (d) of Figure 2) results in the expected order quantities and related costs. In particular, when region i has lower and upper bound $S^s - a_i$ and $S^s - b_i$, respectively, then $p_i = F[a_i] - F[b_i]$. The base scenario for our numerical experiments, which visually supports our analysis, is always: $h = 1$, $b = 9$, $c_f = 4$, $c_s = 3.8$, $m = 1.1$, and $k = 1$. Demand is normally distributed with mean $\mu = 10$ and standard deviation $\sigma = 2.5$ (similar to Boute et al. [2021]).

6.1. Sensitivity with respect to base capacity k

The local supplier's volume flexibility is characterized by its base capacity k, defining the volume that can be sourced at regular cost, as well as the overtime multiplier m, defining the cost premium for units sourced beyond base capacity. We first investigate the impact of increasing the base capacity, for instance through a larger local supplier's workforce. When $k = 0$, sourcing costs of both supplies are linear with all units sourced locally at cost mc_f and the optimal policy is a classic dual base-stock policy (cf. panel (c) of Figure 2). An increase in base capacity to $k > 0$ shifts the optimal policy to a modified dual base-stock policy (panel (d) of Figure 2). Our analysis captures the impact of increasing k on the sourcing split by computing the gradient of the shares ρ^s, ρ^f_r, and ρ^f_o with respect to k. Using Leibniz's rule for differentiation under the integral we find that increasing k reshores volumes sourced locally at regular cost, ρ^f_r. The substitution arises from a reduction of both slow supply, ρ^s and fast supply at overtime cost, ρ^f_o. The substitution effects are as follows. The fraction sourced from the fast supply at regular cost increases according to

$$\frac{d\rho^f_r}{dk} = \frac{1}{\mu}\left(p_{\text{IV}} + p_{\text{III}} - p_{\text{II}}\frac{dS^s}{dk}\right) \geq 0, \tag{13}$$

by reducing both the volumes sourced from slow supply and from the fast supply at overtime cost:

$$\frac{d\rho^s}{dk} = \frac{1}{\mu}\left((p_{\text{IV}} + p_{\text{II}})\frac{dS^s}{dk} - p_{\text{III}}\right) \leq 0 \qquad \text{and} \qquad \frac{d\rho^f_o}{dk} = \frac{1}{\mu}\left(-p_{\text{IV}}\frac{dS^s}{dk} - p_{\text{IV}}\right) \leq 0. \tag{14}$$

Here $dS^s/dk \in [-1, 0]$ represents how much the slow base-stock level reduces when we increase k, see (42) in the Appendix D for a more explicit expression. The fraction reshored is largest when $-dS^s/dk \to 1$. The value of dS^s/dk is dependent on the underlying demand distribution, as demonstrated in Appendix D. Therefore, without knowledge of the demand distribution we

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

cannot show it is concave (or convex) in k. More so, for many distributions (such as the normal distribution), it will have inflection points. However, in the limit, when $k \to \infty$, the marginal return of a capacity addition equals zero. In this case, the modified base-stock policy converges again to a classic dual base-stock policy with all units sourced locally occurring a cost of c_f without overtime production because dS^s/dk, p_{IV}, and p_{III} all reduce to zero for any demand distribution with finite variance as

$$\lim_{k\to\infty} \frac{dS^s}{dk} = \lim_{k\to\infty} p_{\mathrm{IV}} = \lim_{k\to\infty} p_{\mathrm{III}} = 0.$$

While the fraction sourced from the fast supply ρ^f is non-decreasing wrt k, the marginal increase vanishes towards zero when k increases to infinity, $\lim_{k\to\infty} d\rho^f/dk = 0$. Proposition 1 captures these results:

Proposition 1 *Increasing the local base capacity k reshores sourcing volumes from remote to local suppliers and reallocates volume sourced locally from overtime to regular production. The incremental volumes reshored decrease with a higher k and vanish to zero in the limit as $k \to \infty$.*

Panel (a) of Figure 4 shows the impact of these substitution effects on the sourcing splits. The hatched area quantifies how much is reshored from the slow to the fast supply due to a larger base capacity k. It also shows the diminishing marginal returns of additional capacity investments.

In addition to reshoring orders, increasing the base capacity k also increases the standard deviation of the orders placed to the fast supplier, while it reduces the standard deviation of the orders placed at the slow supplier, see also panel (d) of Figure 4. More volume flexibility allows the buyer to leverage one of the key strengths of the responsive supplier: to absorb more of the demand variability. While the slow supplier is used to leverage the sourcing cost advantage, the fast supplier's responsiveness is used to offset the inventory increase due to the slow supplier's longer lead time. A higher base capacity reduces the coefficient of variation of both suppliers' orders, as visualized in panel (g) of Figure 4. The reduction is largest for the fast supplier and converges for higher values of k.

6.2. Sensitivity with respect to the volume flexibility reduction coefficient m

The overtime premium m captures the level of volume flexibility of the responsive supply. It functions as a flexibility reduction coefficient: A higher overtime premium reduces the volume flexibility, as volume sourced beyond the base capacity k (at overtime cost mc_f) becomes more expensive. This reduces the volumes reshored to the local supply, both at regular and overtime cost:

$$\frac{d\rho_r^f}{dm} = \frac{1}{\mu}\left(-p_{\mathrm{II}}\frac{dS^s}{dm}\right) \leq 0, \qquad \text{and} \qquad \frac{d\rho_o^f}{dm} = \frac{1}{\mu}\left(-p_{\mathrm{IV}}\left(\frac{dS^s}{dm} - \frac{dS_2^f}{dm}\right)\right) \leq 0, \tag{15}$$

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

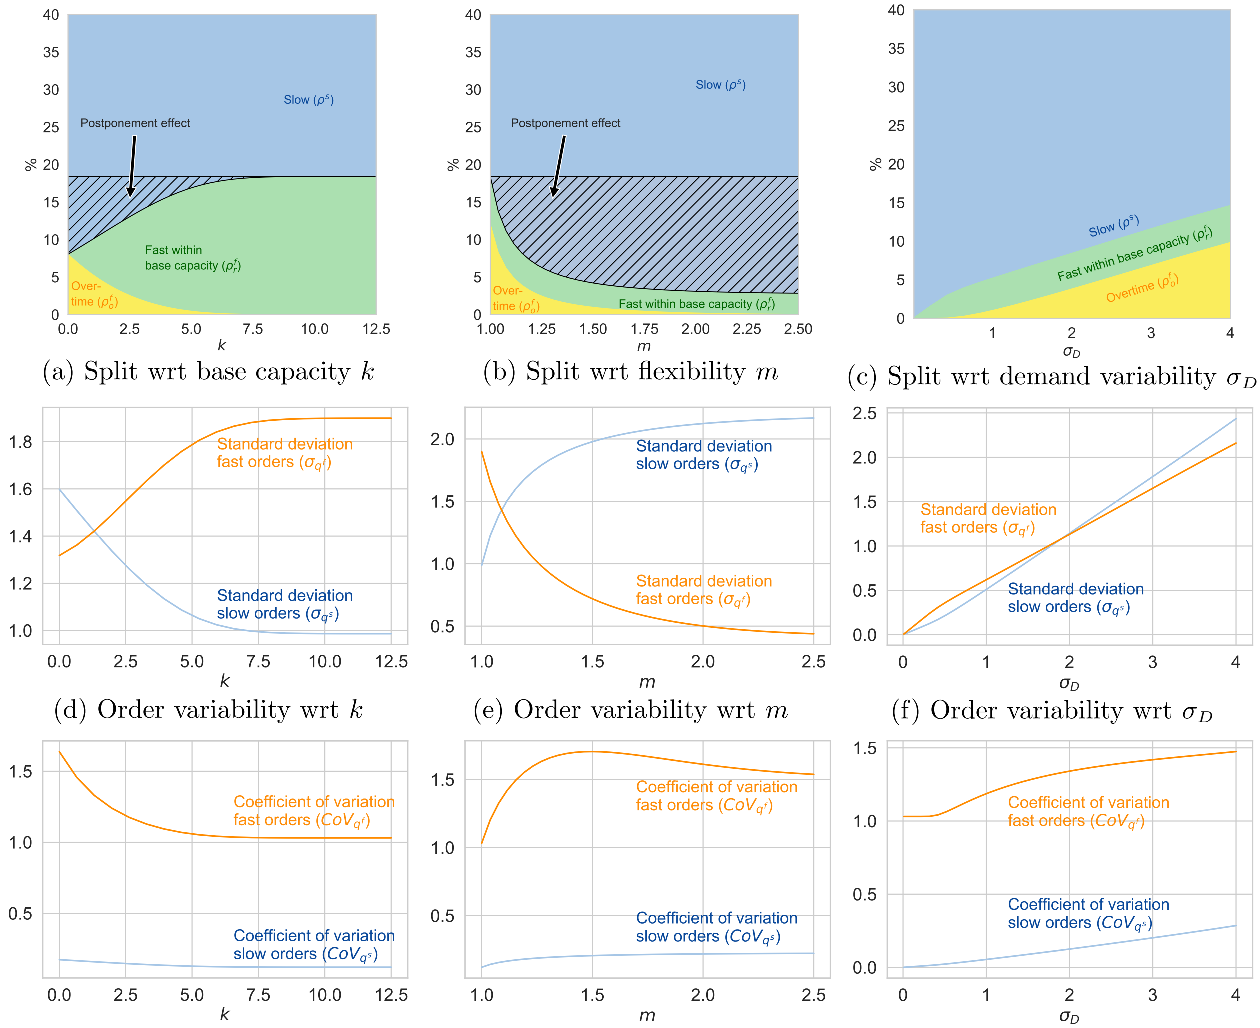

(a) Split wrt base capacity k (b) Split wrt flexibility m (c) Split wrt demand variability σ_D

(d) Order variability wrt k (e) Order variability wrt m (f) Order variability wrt σ_D

(g) CoV orders wrt k (h) CoV orders wrt m (i) CoV orders wrt σ_D

Figure (4) Impact of base capacity, overtime premium and demand uncertainty on the sourcing split (top row), the order variance (middle row) and the coefficient of variation of orders (bottom row). Base parameters: $h=1, b=9, c_f=4, c_s=3.8, k=1, m=1.1,$ $d_t \overset{d}{\sim} \mathcal{N}(\mu=10, \sigma=2.5)$

in favor of a higher fraction sourced from the slow supply:

$$\frac{d\rho^s}{dm} = \frac{1}{\mu}\left(p_{\text{II}}\frac{dS^s}{dm} + p_{\text{IV}}\left(\frac{dS^s}{dm} - \frac{dS_2^f}{dm}\right)\right) \geq 0, \tag{16}$$

with dS^s/dm and dS_2^f/dm representing the sensitivity of the S^s and S_2^f with respect to m. The expressions for $(dS^s/dm) \geq 0$ and $dS_2^f/dm = -c_f f[S^s - S_2^f]/(b+h) \leq 0$, are provided in Appendix D.

The reduction in reshoring increases with higher values of dS^s/dm and $-dS_2^f/dm$. Their values depend again on the underlying demand distribution, as shown in Appendix D. The proportion reshored to the local supply is non-increasing in m. In the limit dS^s/dm and dS_2^f/dm reduce to zero for any demand distribution with finite variance, so that nothing is produced in overtime when $m \to \infty$:

$$\lim_{m\to\infty} \frac{dS^s}{dm} = \lim_{m\to\infty} \frac{dS_2^f}{dm} = 0. \tag{17}$$

We capture these results in Proposition 2:

Proposition 2 *A higher overtime premium, m, reduces the reshored volumes at regular and overtime cost in favor of more offshored supply ρ^s. The share reshored to the fast supply is non-increasing but the gradient vanishes towards zero for large m.*

Panel (b) of Figure 4 visualizes how a higher overtime premium m reduces the fraction reshored to local supply. Compared to perfect volume flexibility with $m = 1$, orders are postponed when $m > 1$ by shifting from the fast to the slow supply. Extending the delivery by one additional period avoids the expensive overtime premium. When $m \to +\infty$, our policy reduces to a dual base-stock policy with one base-stock level S_1^f for the fast supply (and $S_2^f = -\infty$) *and* a cap on the fast order at k, similar to the findings of Federgruen et al. [2021].

A higher overtime premium m also reduces the variance of the reshored orders, while it increases the variance of the orders placed to the slow supply, see panel (e) of Figure 4. The increased premium reduces the attractiveness of the fast supply such that more of the variance is taken by the slow supplier. This effect is similar to an increase in base capacity. In §6.4 we show how k and m act as strategic substitutes for the fast supplier when targeting a specific sourcing split. Whereas the coefficients of variation of the order quantities of the slow source increase with an increasing overtime premium m, the impact on the coefficient of variation of the fast supplier's orders is not monotone.

6.3. Impact of demand variability

We also investigate how demand variability impacts the orders placed under limited volume flexibility of the responsive supply. Although it cannot be analytically proven, we conjecture (and confirm through numerical results, see panel (c) of Figure 4) that more is reshored to the responsive supply when demand is more uncertain: in these cases the responsive source covers the more variable part of the demand. We find that a higher demand variability, σ, results in a shift between fast (overtime) supply:

$$\frac{d\rho_r^f}{d\sigma} = \frac{1}{\mu}\left(-p_{\mathrm{II}}\left(\frac{dS^s}{d\sigma} - z_1\right)\right), \qquad \text{and} \qquad \frac{d\rho_o^f}{d\sigma} = \frac{1}{\mu}\left(-p_{\mathrm{IV}}\left(\frac{dS^s}{d\sigma} - z_2\right)\right), \tag{18}$$

and slow supply:

$$\frac{d\rho^s}{d\sigma} = \frac{1}{\mu}\left(p_{\mathrm{II}}\left(\frac{dS^s}{d\sigma} - z_1\right) + p_{\mathrm{IV}}\left(\frac{dS^s}{d\sigma} - z_2\right)\right). \tag{19}$$

Unfortunately it cannot be shown analytically whether the shift is from fast to slow or vice versa, which depends on the signs of $(dS^s/d\sigma - z_1)$ and $(dS^s/d\sigma - z_2)$. However, all our numerical experiments confirm that volume is shifted from slow supply towards fast (overtime) supply. We find this effect is largest for perfect volume flexibility, which follows from our results visualized in panels (a) and (b) of Figure 4: more volume flexibility reshores a larger fraction of demand.

We also observe the order variability of both suppliers, measured by both standard deviation and coefficient of variation, increases with higher demand uncertainty (panel (f) of Figure 4). The magnitude of the results obtained in §6.1 and §6.2 are thus larger when demand is more variable.

6.4. Base capacity and flexibility as strategic substitutes

The above analyses reveal how more volume is reshored by increasing the base capacity or by lowering the overtime premium at the local supply. The local supplier thus has two levers to influence the fraction reshored. These two levers interact with each other. We quantify to what extent: a) the base capacity should be increased for a given overtime premium to retain target sourcing volumes, and b) the overtime premium should be reduced to induce a reduction in base capacity while still reshoring a target fraction of demand. Figure 5 illustrates, for the same parameters as in Figure 4, how the base capacity and overtime premium interact to retain a given reshored fraction.[5]

We observe that when base capacity and overtime premium are high (top left corner), a significant reduction in m is needed to compensate for a reduction in k. Likewise, a significant increase of the base capacity, k is needed to compensate for an increase in m when the base capacity and overtime premium are low (bottom right).

6.5. Comparison between perfect vs. limited volume flexibility

A higher base capacity k and lower overtime premium m promote the decision to reshore volumes to the responsive supply. Figure 6, panel (a) visualizes the reshored volumes for various values of k and m. In both cases, the volumes sourced locally at the regular cost c_f and at overtime premium mc_f are impacted. This in turn will affect the resulting *average* sourcing cost per unit sourced locally, which we define by $\bar{c}_f \triangleq (c_f \mu_r^f + mc_f \mu_o^f)/(\mu_r^f + \mu_o^f)$, such that $c_f \leq \bar{c}_f \leq mc_f$.

Panel (b) of Figure 6 illustrates how the resulting $\bar{c}_f$ changes with k and m for the same parameter values as panel (a). It shows how the average unit sourcing cost $\bar{c}_f$ decreases with higher values of k (as more units can be sourced at the lower cost c_f); and for lower values of m (as sourcing overtime becomes less expensive). To compare perfect volume flexibility versus limited volume flexibility of the responsive supply *under the same average unit sourcing cost* $\bar{c}_f$, panel (c) shows how much volume is reshored when sourcing costs are linear with local sourcing cost $\bar{c}_f$ for each unit sourced.

[5] Results are plotted using a logarithmic scale with base 1.25.

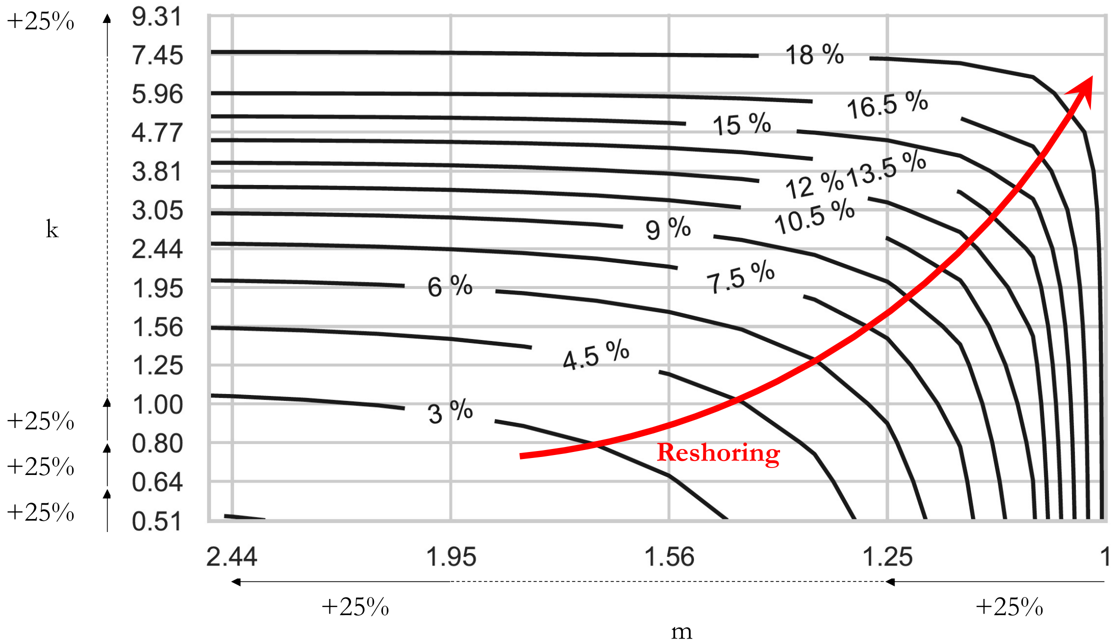

Figure (5) **Isograph of the fixed fast sourcing fraction (or reshored fraction) in log-log scale. These indifference curves quantify the substitution effect between base capacity and overtime premium in percentage terms. The contours reveal how an increase in k can reduce the overtime premium (or vice versa) to maintain the same fraction reshored. It also shows how increasing volume flexibility (through adding base capacity or reducing the overtime premium) is largest for inflexible systems with low values of m and vanishes to zero as the fast supply is more flexible for large m.**

Panel (d) plots the difference between (a) and (c): it shows, for the same average unit cost differential, how much less is reshored when the responsive supply has limited flexibility for different values of k and m. More volume flexibility at the local responsive source encourages more reshoring and serves as a (cost) benefit compared to the remote source.

6.6. Performance under non-consecutive lead times

A key condition to prove the optimality of the modified dual base-stock policy is that both sources have consecutive lead times. We numerically investigate the impact of relaxing this assumption. Under non-consecutive lead times, the optimal order quantity is no longer of a base-stock type, but becomes a complicated, unknown function of the net inventory and all outstanding orders, i.e.,

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

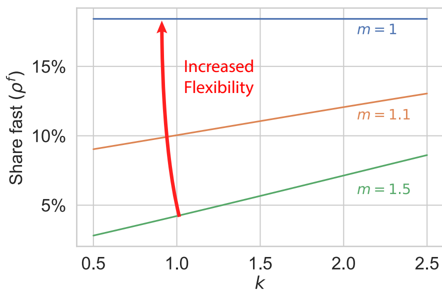

(a) Re-shoring (Piece-wise linear)

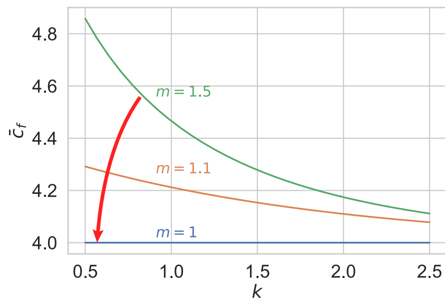

(b) Average unit sourcing cost $\bar{c}_f$

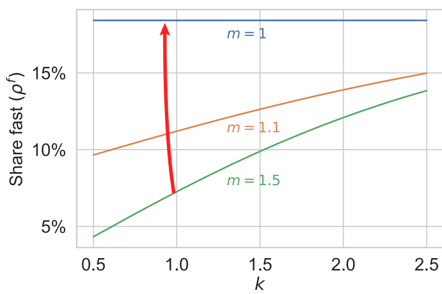

(c) Re-shoring (Linear at $\bar{c}_f$)

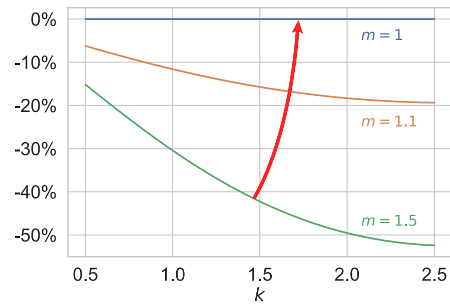

(d) Reduction due to limited flexibility

Figure (6) **The top left panel plots the average share reshored for different values of k and m. We use these results to compute the average fast sourcing cost per unit, $\bar{c}_f$ (top right panel). We then show the share reshored assuming linear sourcing cost $\bar{c}_f$ for the fast supply (bottom left panel) and demonstrate (bottom right panel) that a piece-wise linear sourcing cost functions works as a (cost) disadvantage: less volume is reshored when costs are piece-wise linear (for an equal average cost per unit). Base parameters:** $h=1, b=9, c_f=4, c_s=3.8, k=1, m=1.1$, $d_t \stackrel{d}{\sim} \mathcal{N}(\mu=10, \sigma=2.5)$

the pipeline inventory vector.[6] To operate our modified dual base-stock policy we will use both a single index policy that keeps track of one inventory position that aggregates the net inventory and all outstanding orders such as in Scheller-Wolf et al. [2003], and a dual index policy that keeps track of two inventory positions that aggregate the net inventory and all outstanding orders to arrive within the lead time of the fast and slow source respectively, such as in Veeraraghavan and Scheller-Wolf [2008].

We adopt 36 of the numerical parameter settings of Scheller-Wolf et al. [2003] with a small uniform demand support for numerical tractability. Their model does not include a base capacity or overtime premium. We set the base capacity to either $k=1$, such that only limited units are available within the base capacity, or to $k=+\infty$ indicating there is ample capacity and our model is equivalent to conventional linear dual sourcing. The overtime premium is set to either $m=1.1$

[6] Whittemore and Saunders [1977] have shown this result under linear sourcing costs; our numerical results indicate that in our setting when the sourcing costs of the fast supply are piece-wise linear and convex, the optimal policy is also no longer a base-stock policy.

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

in the base case, to $m = +\infty$ to reflect the case with no volume flexibility, or to $m = 1$ to reflect perfect volume flexibility (which is equivalent to $k = +\infty$).

Table 1 reports the optimality gap of the single and dual index modified base-stock policies. The cost of the optimal policies is obtained using a linear programming formulation of the underlying dynamic program, and the cost of the modified dual base-stock policies by obtaining the steady-state cost of the underlying Markov Chain with the best policy parameters found through an exhaustive search.

We find that the optimality gaps under limited volume flexibility ($k = 1$ and $m = 1.1$) are of the same order of magnitude as the scenario with no volume flexibility ($k = 1$ and $m = +\infty$), as well as the scenario with linear sourcing costs ($k = \infty$) with the dual index policies outperforming the single index policies (similar to the results of Scheller-Wolf et al. [2003] under linear sourcing costs). The performance of both policies decreases as the lead time difference between both sources gets longer, where the performance of the single index tends to decrease faster compared to the dual index. As also outlined by Scheller-Wolf et al. [2003], this may be attributed to the fact that, particularly when the demand support is small, a single index policy has no fine control. Finally, our numerical results confirm Proposition 1 that limited volume flexibility reduces the fraction re-shored and thus effectively works as a cost disadvantage. In other words, our key findings continue to hold under non-consecutive lead times.

7. Conclusion

We investigate dual-source supply chains in which fast supply is not only more expensive, it also is less flexible. Volumes sourced locally beyond a base capacity are charged an overtime premium. We proved a modified dual base-stock policy is optimal when both supply options have consecutive lead times. When sourcing costs are piece-wise linear, the parameters of the optimal *single* sourcing policy cannot be captured analytically as the optimal policy is not a demand replacement policy. When combined with a remote supplier that has linear sourcing costs, the resulting optimal *dual* sourcing policy, while being more *complex*, is often a demand replacement policy. This *simplifies* the characterization of the optimal policy parameters. Our analysis shows how a local supplier becomes more competitive and captures more volume by increasing its base capacity or lowering its overtime premium. Both levers act as strategic substitutes, of which we show their marginal substitution rates. In practice, both actions require investments and, depending on the marginal investment costs, one will be preferred over the other. Our key recommendation to local suppliers is to consider investments that improve volume flexibility to compete with cheaper offshore supply.

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

Parameters					**Optimality Gap**		**Fraction reshored**		
c_f	l_s	b	k	m	Single Index	Dual Index	Single Index	Dual Index	Optimal
1020	2	95	1	1.1	4,8%	1,6%	0,0%	1,7%	5,0%
				$+\infty$	4,8%	1,6%	0,0%	1,7%	5,0%
			$+\infty$	1	5,2%	0,6%	0,0%	6,1%	7,6%
		495	1	1.1	3,0%	1,0%	10,0%	5,0%	2,5%
				$+\infty$	3,0%	1,0%	10,0%	5,0%	2,5%
			$+\infty$	1	12,7%	1,2%	10,0%	13,9%	9,3%
	3	95	1	1.1	10,6%	2,7%	10,0%	6,0%	7,4%
				$+\infty$	10,6%	2,7%	10,0%	6,0%	7,4%
			$+\infty$	1	14,3%	2,9%	10,0%	14,3%	15,1%
		495	1	1.1	12,4%	3,4%	10,0%	3,2%	4,5%
				$+\infty$	12,4%	4,4%	10,0%	3,1%	4,5%
			$+\infty$	1	27,4%.	2,8%	30,0%	14,3%	15,9%
1050	2	95	1	1.1	0,0%	0,0%	0,0%	0,0%	0,0%
				$+\infty$	0,0%	0,0%	0,0%	0,0%	0,0%
			$+\infty$	1	0,0%	0,0%	0,0%	0,0%	0,0%
		495	1	1.1	8,4%	0,9%	0,0%	1,7%	2,5%
				$+\infty$	8,4%	0,9%	0,0%	1,7%	2,5%
			$+\infty$	1	8,4%	0,9%	0,0%	1,7%	2,5%
	3	95	1	1.1	2,7%	0,4%	0,0%	3,1%	2,3%
				$+\infty$	2,7%	0,4%	0,0%	3,1%	2,3%
			$+\infty$	1	3,1%	0,8%	0,0%	1,5%	2,9%
		495	1	1.1	12,1%	1,0%	0,0%	3,2%	4,1%
				$+\infty$	12,1%	1,7%	0,0%	3,1%	4,1%
			$+\infty$	1	17,7%	1,8%	0,0%	3,9%	4,9%
1100	2	95	1	1.1	0,0%	0,0%	0,0%	0,0%	0,0%
				$+\infty$	0,0%	0,0%	0,0%	0,0%	0,0%
			$+\infty$	1	0,0%	0,0%	0,0%	0,0%	0,0%
		495	1	1.1	1,2%	0,0%	0,0%	1,7%	1,7%
				$+\infty$	1,2%	0,0%	0,0%	1,7%	1,7%
			$+\infty$	1	1,2%	0,0%	0,0%	1,7%	1,7%
	3	95	1	1.1	0,0%	0,0%	0,0%	0,0%	0,0%
				$+\infty$	0,0%	0,0%	0,0%	0,0%	0,0%
			$+\infty$	1	0,0%	0,0%	0,0%	0,0%	0,0%
		495	1	1.1	5,0%	0,7%	0,0%	1,3%	1,7%
				$+\infty$	5,0%	0,7%	0,0%	1,3%	1,7%
			$+\infty$	1	5,5%	1,1%	0,0%	1,5%	1,9%

Table (1) **Performance of the modified dual base-stock policy for a larger numerical experiment. The fixed numerical parameters are $l_f = 0$, $h = 5$, and $c_s = 1000$. Demand is discrete and uniform: $D \sim U[0, 4]$.**

References

Allon G, Van Mieghem JA (2010) Global Dual Sourcing: Tailored Base-Surge Allocation to Near- and Offshore Production. *Mgt Sci* 56(1):110–124.

Barankin EW (1961) A delivery-lag inventory model with an emergency provision (the single-period case). *Naval Research Logistics Quarterly* 8(3):285–311.

Boute R, Disney S, Gijsbrechts J, Van Mieghem J (2021) Dual Sourcing and Smoothing Under Non-Stationary Demand Time Series: Re-Shoring with Speedfactories. Forthcoming in Management Science URL `http://dx.doi.org/10.1287/mnsc.2020.3951`.

Boute R, Van Mieghem JA (2014) Global Dual Sourcing and Order Smoothing: The Impact of Capacity and Leadtimes. *Mgt Sci* 9(61):2080–2099.

Bulinskaya EV (1964a) Some results concerning optimum inventory policies. *Theory Probab. Appl.* 9(3):389–403.

Bulinskaya EV (1964b) Steady-State Solutions in Problems of Optimum Inventory Control. *Theory Probab. Appl.* 9(3):502–507.

Chen L, Hu B (2017) Is Reshoring Better Than Offshoring? The Effect of Offshore Supply Dependence. *Manufacturing Service Oper. Management* 19(2):166–184.

Federgruen A, Liu Z, Lu L (2021) Dual sourcing: Creating and utilizing flexible capacities with a second supply source. Working Paper, Columbia University.

Fukuda Y (1964) Optimal Policies for the Inventory Problem with Negotiable Leadtime. *Mgt Sci* 10(4):690–708.

Henig M, Gerchak Y, Ernst R, Pyke DF (1997) An Inventory Model Embedded in Designing a Supply Contract. *Mgt Sci* 43(2):389–403.

Huh W, Janakiraman G, Nagarajan M (2011) Average Cost Single-Stage Inventory Models: An Analysis Using a Vanishing Discount Approach. *Ops Res* 59(1):143–155.

Karlin S (1958) *Studies in the Mathematical Theory of Inventory and Production* (Stanford University Press).

Karlin S, Scarf H (1958) Chapter 10: Inventory models of the arrow-harris-marschak type with time lag. *Studies in the Mathematical Theory of Inventory and Production* (Stanford University Press).

Lu Y, Song M (2014) Inventory Control with a Fixed Cost and a Piecewise Linear Convex Cost. *Production Oper. Management* 23(11):1966–1984.

Martínez de Albéniz V, Simchi-Levi D (2005) A portfolio approach to procurement contracts. *Production Oper. Management* 14(1):90–114.

Neuts MF (1964) An inventory model with an optional time lag. *Journal of the Society for Industrial and Applied Mathematics* 12(1):179–185.

Porteus E (1990) Stochastic inventory theory. *Handbooks in Operations Research and Management Science* 2:605–652.

Porteus E (2002) *Foundations of Stochastic Inventory Theory* (Stanford University Press).

Scheller-Wolf A, Veeraraghavan S, van Houtum GJ (2003) Effective dual sourcing with a single index policy. Working Paper, Wharton at the University of Pennsylvania.

Sheopuri A, Janakiraman G, Seshadri S (2010) New Policies for the Stochastic Inventory Control Problem with Two Supply Sources. *Ops Res* 58(3):734–745.

Sun J, Van Mieghem JA (2019) Robust Dual Sourcing Inventory Management: Optimality of Capped Dual Index Policies. *Manufacturing Service Oper. Management* 21(4):713–948.

Tomlin B (2006) On the Value of Mitigation and Contingency Strategies for Managing Supply Chain Disruption Risks. *Mgt Sci* 52(5):639–657.

Veeraraghavan S, Scheller-Wolf A (2008) Now or Later: A Simple Policy for Effective Dual Sourcing in Capacitated Systems. *Ops Res* 56(4):850–864.

Whittemore AS, Saunders SC (1977) Optimal Inventory Under Stochastic Demand With Two Supply Options. *J on Appl Math* 32(2):293–305.

Xin L, Goldberg DA (2017) Asymptotic Optimality of Tailored Base-Surge Policies in Dual-Sourcing Inventory Systems. *Mgt Sci* 64(1):437–452.

Xin L, Van Mieghem JA (2021) Dual Sourcing, Dual Mode Dynamic Inventory Models: A Review. Working Paper, UC Chicago.

Appendix A: Proofs

Proof of Theorem 1 We prove discounted-cost optimality of the modified dual base-stock policy by backward induction for a finite horizon of length T (extending Xin and Van Mieghem [2021]). In particular, we aim to find the policy with optimal discounted cost-to-go $v_t[x_{t-1}]$ from t until T starting with inventory level x_{t-1} that satisfies the well-known Bellman equation for every period $t \in \{1, \cdots, T\}$:

$$v_t[x_{t-1}] = \min_{q_t^f, q_t^s \geq 0} \left\{ mc_f[q_t^f - k]^+ + c_f(q_t^f \wedge k) + c_s q_t^s + L_t[x_{t-1} + q_t^f] + \alpha \mathbb{E}[v_{t+1}[x_{t-1} + q_t^f + q_t^s - D_t]] \right\}. \quad (20)$$

and terminal value function $v_{T+1}[x_T] = -c_{T+1} x_T$, where $c_{T+1} \geq 0$.

The cost in the last period T consists of the ordering cost of the fast source (both at regular and overtime cost), the expected inventory mismatch, and the units left over at the end of the period may be returned for a per-unit of revenue c_{T+1} (likewise, backlogged demand is fulfilled at per-unit cost c_{T+1}). Later, we will determine the specific salvage cost that ensures the optimal policy is stationary under certain conditions. No order from the slow supplier will be placed in the last period, $q_T^{s,\star} = 0$, as the ordered units will not arrive in time to influence the inventory mismatch of period T. It only remains to determine q_T^f based on the ending inventory of the previous period I_{T-1} and the slow order q_{T-1}^s that arrives in period T. Without loss of generalization, we only need to consider their sum, the inventory position x_{T-1}.

The value function in the last period T thus satisfies:

$$v_T[x_{T-1}] = \min_{q_T^f \geq 0} \left\{ mc_f[q_T^f - k]^+ + c_f(q_T^f \wedge k) + L_T[x_{T-1} + q_T^f] - \alpha \mathbb{E}[c_{T+1}(x_{T-1} + q_T^f - D_T)] \right\}. \tag{21}$$

By a change of variables, $y_T^f \triangleq x_{T-1} + q_T^f$, the minimization becomes

$$v_T[x_{T-1}] = \min_{y_T^f \geq x_{T-1}} \left\{ mc_f[y_T^f - x_{T-1} - k]^+ + c_f([y_T^f - x_{T-1}]^+ \wedge k) + L_T[y_T^f] - \alpha \mathbb{E}[c_{T+1}(y_T^f - D_T)] \right\}, \tag{22}$$

which is jointly convex in x_{T-1} and y_T^f as per Lemma 1.

Let $y_T^{f,\star}[x_{T-1}] \triangleq S_T^f[x_{T-1}]$ be a minimizer. Including it in (22) yields:

$$\begin{aligned} v_T[x_{T-1}] =& mc_f[S_T^f[x_{T-1}] - x_{T-1} - k]^+ + c_f([S_T^f[x_{T-1}] - x_{T-1}]^+ \wedge k) + \\ & L_T[S_T^f[x_{T-1}]] - \alpha \mathbb{E}[c_{T+1}(S_T^f[x_{T-1}] - D_T)]. \end{aligned} \tag{23}$$

By convexity, established in Lemma 1, it is then optimal to order fast up to $S_T^f[x_{T-1}]$, with:

$$S_T^f[x_{T-1}] = \begin{cases} x_{T-1} & \text{if } S_{1,T}^f \leq x_{T-1}, \\ S_{1,T}^f & \text{if } S_{1,T}^f - k \leq x_{T-1} < S_{1,T}^f, \\ x_{T-1} + k & \text{if } S_{2,T}^f - k \leq x_{T-1} < S_{1,T}^f - k, \\ S_{2,T}^f & \text{if } x_{T-1} < S_{2,T}^f - k, \end{cases} \tag{24}$$

where the two fast order-up-to levels in period T are:

$$S_{1,T}^f = F_{D_T}^{-1}\left[\frac{b - (c_f - \alpha c_{T+1})}{b+h}\right] \geq S_{2,T}^f = F_{D_T}^{-1}\left[\frac{b - (mc_f - \alpha c_{T+1})}{b+h}\right], \tag{25}$$

which satisfy the sufficient first-order conditions on (22). As $q_T^{f,\star} = [S_T^f[x_{T-1}] - x_{T-1}]^+$, using (24), the fast order becomes

$$q_T^{f,\star}[x_{T-1}] = \begin{cases} 0 & \text{if } S_{1,T}^f \leq x_{T-1}, \\ S_{1,T}^f - x_{T-1} & \text{if } S_{1,T}^f - k \leq x_{T-1} < S_{1,T}^f, \\ k & \text{if } S_{2,T}^f - k \leq x_{T-1} < S_{1,T}^f - k, \\ S_{2,T}^f - x_{T-1} & \text{if } x_{T-1} < S_{2,T}^f - k. \end{cases} \tag{26}$$

Proceeding backwards to period $T-1$:

$$\begin{aligned} v_{T-1}[x_{T-2}] = \min_{q_{T-1}^f, q_{T-1}^s \geq 0} & \left\{ mc_f[q_{T-1}^f - k]^+ + c_f(q_{T-1}^f \wedge k) + c_s q_{T-1}^s + \right. \\ & \left. L_{T-1}[x_{T-2} + q_{T-1}^f] + \alpha \mathbb{E}[v_T[x_{T-2} + q_{T-1}^f + q_{T-1}^s - D_{T-1}]] \right\}. \end{aligned} \tag{27}$$

We solve optimization (27) first for the *interior case* where both optimal quantities are positive: $q_{T-1}^{f,\star} > 0$ and $q_{T-1}^{s,\star} > 0$ (Case A). Then we consider the three boundary cases: $q_{T-1}^{\star,f} = 0$ while $q_{T-1}^{s,\star} > 0$ (Case B); $q_{T-1}^{f,\star} > 0$ while $q_{T-1}^{s,\star} = 0$ (Case C); and finally $q_{T-1}^{f,\star} = q_{T-1}^{s,\star} = 0$ (Case D). Determining when the boundary cases arise will have important implications on the possibility to determine the policy parameters explicitly in §5.

Case A: $q_{T-1}^{f,\star} > 0$ **and** $q_{T-1}^{s,\star} > 0$. As both orders are strictly positive, optimization (27) can be rearranged such that the minimization of the slow order becomes an inner minimization:

$$\begin{aligned} v_{T-1}[x_{T-2}] = \min_{q_{T-1}^f > 0} & \left\{ mc_f[q_{T-1}^f - k]^+ + c_f(q_{T-1}^f \wedge k) + L_{T-1}[x_{T-2} + q_{T-1}^f] + \right. \\ & \left. \min_{q_{T-1}^s > 0} \left\{ c_s q_{T-1}^s + \alpha \mathbb{E}[v_T[x_{T-2} + q_{T-1}^f + q_{T-1}^s - D_{T-1}]] \right\} \right\}, \end{aligned} \tag{28}$$

and by a change of variables: $y^s_{T-1} \triangleq x_{T-2} + q^f_{T-1} + q^s_{T-1}$ becomes[7]:

$$v_{T-1}[x_{T-2}] = \min_{q^f_{T-1}>0} \Big\{ (mc_f - c_s)[q^f_{T-1} - k]^+ + (c_f - c_s)(q^f_{T-1} \wedge k) + L_{T-1}[x_{T-2} + q^f_{T-1}] + \min_{y^s_{T-1} > x_{T-2}+q^f_{T-1}} \big\{ c_s(y^s_{T-1} - x_{T-2}) + \alpha \mathbb{E}[v_T[(y^s_{T-1} - D_{T-1})]] \big\} \Big\}. \quad (29)$$

Let $y^{s,*}_{T-1} \triangleq S^s_{T-1}$ be a minimizer of the inner minimization. Then:

$$v_{T-1}[x_{T-2}] = \min_{q^f_{T-1}>0} \big\{ (mc_f - c_s)[q^f_{T-1} - k]^+ + (c_f - c_s)(q^f_{T-1} \wedge k) + L_{T-1}[x_{T-2} + q^f_{T-1}] + c_s(S^s_{T-1} - x_{T-2}) + \alpha \mathbb{E}[v_T[S^s_{T-1} - D_{T-1}]] \big\}, \quad (30)$$

and it is optimal to raise the inventory position to S^s_{T-1} using both orders: $q^{s,\star}_{T-1} = S^s_{T-1} - x_{T-2} - q^{f,\star}_{T-1}$, yet it remains to determine $q^{f,\star}_{T-1}$. Given that $S^s_{T-1} > x_{T-2} + q^{f,\star}_{T-1}$ in Case A, S^s_{T-1} is independent of the fast order, $q^{f,\star}_{T-1}$. The latter then minimizes the sum of the first three terms in (30) which are myopic, i.e., independent of the future cost $v_T[\cdot]$. In that case, after a change of variables, $y^f_{T-1} \triangleq x_{T-2} + q^f_{T-1}$, we denote a minimizer of this myopic cost by:

$$U_{\mathcal{M}}[x_{T-2}] = \underset{y^f_{T-1} \geq x_{T-2}}{\arg\min} \big\{ (mc_f - c_s)[y^f_{T-1} - x_{T-2} - k]^+ + (c_f - c_s)([y^f_{T-1} - x_{T-2}]^+ \wedge k) + L_{T-1}[y^f_{T-1}] \big\}. \quad (31)$$

The latter objective function is convex so that the sufficient first-order conditions reveal it is optimal to order up to $U_{\mathcal{M}}[x_{T-2}]$, with

$$U_{\mathcal{M}}[x_{T-2}] = \begin{cases} S^f_{1,T-1} & \text{if } S^f_{1,T-1} - k \leq x_{T-2} < S^f_{1,T-1}, \\ x_{T-2} + k & \text{if } S^f_{2,T-1} - k \leq x_{T-2} < S^f_{1,T-1} - k, \\ S^f_{2,T-1} & \text{if } x_{T-2} < S^f_{2,T-1} - k, \end{cases} \quad (32)$$

where

$$S^f_{1,T-1} = F^{-1}_{D_{T-1}} \left[\frac{b - (c_f - c_s)}{b+h} \right] \geq S^f_{2,T-1} = F^{-1}_{D_{T-1}} \left[\frac{b - (mc_f - c_s)}{b+h} \right]. \quad (33)$$

Hence: $q^{f,\star}_T = U_{\mathcal{M}}[x_{T-2}] - x_{T-2}$ where:

$$q^{f,\star}_{T-1}[x_{T-2}] = \begin{cases} S^f_{1,T-1} - x_{T-2} & \text{if } S^f_{1,T-1} - k \leq x_{T-2} < S^f_{1,T-1}, \\ k & \text{if } S^f_{2,T-1} - k \leq x_{T-2} < S^f_{1,T-1} - k, \\ S^f_{2,T-1} - x_{T-2} & \text{if } x_{T-2} < S^f_{2,T-1} - k. \end{cases} \quad (34)$$

We conclude by noting that, in support of our analysis in §5, Case A only holds when both orders are are strictly positive. The inventory position thus needs to be below both base-stock levels: $x_{T-2} < S^s_{T-1}$ and $x_{T-2} < U_{\mathcal{M}}[x_{T-2}]$.

Case B: $q^{f,\star}_{T-1} = 0$ **while** $q^{s,\star}_{T-1} > 0$. Optimization (27) simplifies to

$$v_{T-1}[x_{T-2}] = \min_{q^s_{T-1}>0} \big\{ c_s q^s_{T-1} + L_{T-1}[x_{T-2}] + \alpha \mathbb{E}[v_T[x_{T-2} + q^s_{T-1} - D_{T-1}]] \big\}. \quad (35)$$

Similar to Case A, let $y^{s,*}_{T-1} \triangleq S^s_{T-1}$ be a minimizer of (35). The optimal slow order quantity is then to order up to S^s_{T-1}, hence: $q^{s,\star}_{T-1} = S^s_{T-1} - x_{T-2}$. As no fast order is placed in Case B ($q^{f,\star}_{T-1} = 0$) the optimal fast order-up-to level simply equals x_{T-2}. Note, Case B applies when the inventory position is at or above the highest fast base-stock level of Case A (else ordering fast would reduce costs), and below the slow base-stock

[7] Note that $\forall k$, $q^f_{T-1} = [q^f_{T-1} - k]^+ + (q^f_{T-1} \wedge k)$.

level (else not ordering slow would reduce costs), i.e., $S^f_{1,T-1} \leq x_{T-2} < S^s_{T-1}$ where $S^f_{1,T-1}$ is defined as in (33).

Case C: $q^{f,\star}_{T-1} > 0$ **while** $q^{s,\star}_{T-1} = 0$. Optimization (27) can then be expressed as:

$$v_{T-1}[x_{T-2}] = \min_{q^f_{T-1} \geq 0} \big\{ mc_f[q^f_{T-1} - k]^+ + c_f(q^f_{T-1} \wedge k) + L_{T-1}[x_{T-2} + q^f_{T-1}] + \alpha \mathbb{E}[v_T[x_{T-2} + q^f_{T-1} - D_{T-1}]] \big\}. \tag{36}$$

In this case, the fast order also impacts future periods. After a change of variables $y^f_{T-1} \triangleq x_{T-2} + q^f_{T-1}$, we denote a (non-myopic) minimizer in period $T-1$ by:

$$U_{\overline{\mathcal{M}}}[x_{T-2}] = \operatorname*{arg\,min}_{y^f_{T-1} \geq x} \big\{ mc_f[y^f_{T-1} - x_{T-2} - k]^+ + c_f([y^f_{T-1} - x_{T-2}]^+ \wedge k) + L_{T-1}[y^f_{T-1}] + \alpha \mathbb{E}[v_T[y^f_{T-1} - D_{T-1}]] \big\}. \tag{37}$$

Again, the objective function is convex so that the sufficient first-order conditions yield that it is optimal to order up to $U_{\overline{\mathcal{M}}}[x_{T-2}]$, with

$$U_{\overline{\mathcal{M}}}[x_{T-2}] = \begin{cases} S^f_{1,T-1} & \text{if } S^f_{1,T-1} - k \leq x_{T-2} < S^f_{1,T-1}, \\ x_{T-2} + k & \text{if } S^f_{2,T-1} - k \leq x_{T-2} < S^f_{1,T-1} - k, \\ S^f_{2,T-1} & \text{if } x_{T-2} < S^f_{2,T-1} - k, \end{cases} \tag{38}$$

yet, in contrast to case A, no simple fractile expression on the demand in period $T-1$ can be determined for the fast base-stock levels, $S^f_{1,T-1}$ and $S^f_{2,T-1}$, as the fast order also impacts the cost in period T. The optimal fast order satisfies $q^{f,\star}_T = U_{\overline{\mathcal{M}}}[x_{T-2}] - x_{T-2}$ where:

$$q^{f,\star}_{T-1}[x_{T-2}] = \begin{cases} S^f_{1,T-1} - x_{T-2} & \text{if } S^f_{1,T-1} - k \leq x_{T-2} < S^f_{1,T-1}, \\ k & \text{if } S^f_{2,T-1} - k \leq x_{T-2} < S^f_{1,T-1} - k, \\ S^f_{2,T-1} - x_{T-2} & \text{if } x_{T-2} < S^f_{2,T-1} - k. \end{cases} \tag{39}$$

and, as per the boundary condition of Case C: $q^{s,\star}_{T-1} = 0$. Note, a necessary condition for case C to occur is that the inventory position must exceed the slow base-stock level (as established in cases A and B) and must not exceed the fast base-stock level: $S^s_{T-1} \leq x_{T-2} < U_{\overline{\mathcal{M}}}[x_{T-2}]$. In settings with stationary demand, the latter condition will mostly hold when only the fast supplier is used, as we will show in §5 and Appendix C.

Case D: $q^{f,\star}_{T-1} = q^{s,\star}_{T-1} = 0$. Optimization (27) simplifies to

$$v_{T-1}[x_{T-2}] = L_{T-1}[x_{T-2}] + \alpha \mathbb{E}[v_T[x_{T-2} - D_{T-1}]], \tag{40}$$

which is again convex using similar arguments as above. The optimal order quantities are $q^{s,\star}_{T-1} = q^{f,\star}_{T-1} = 0$ as per the boundary condition such that $S^s_{T-1} \leq x_{T-2}$. If $S^s_{T-1} < U_{\overline{\mathcal{M}}}[x_{T-2}]$ such that there exist inventory positions for which case C may occur, then the following constraint on x_{T-2} must hold to enforce no fast orders are made: $U_{\overline{\mathcal{M}}}[x_{T-2}] < x_{T-2}$. If $S^s > U_{\overline{\mathcal{M}}}[x_{T-2}]$ such that case C never appears, then $U_{\mathcal{M}}[x_{T-2}] < x_{T-2}$ must hold to enforce no fast orders. The latter conditions show that Case D only occurs whenever the inventory position exceeds the highest base-stock levels of all possible Cases A, B and C.

Combining the four cases reveals that a modified dual base-stock policy is optimal in period $T-1$ such that (8) and (9) apply when $t = T-1$. We can follow the same line of argument for period $T-2$ (i.e., repeating the same arguments to show that $v_{T-2}[\cdot]$ is convex as $v_{T-1}[\cdot]$ is convex and so are all costs incurred in period $T-2$, and showing that a modified dual base-stock policy is again optimal in period $T-2$). (Note that only convexity of $v_{T-1}[\cdot]$ is required to ensure a modified dual base-stock policy is optimal; to express the policy parameters, which we do in the next section, we also use that a modified dual base-stock policy is optimal in period $T-1$.) Repeating until period t reveals a modified dual base-stock policy is optimal in every period. ■

Proof of Theorem 2 We rely on Huh et al. [2011] who derived sufficient conditions. In particular, we demonstrate how all the conditions to invoke Theorem 3.1 in Huh et al. [2011] are satisfied: They first state two assumptions which are clearly satisfied as our problem is very close to the problems they study: (1) our way of modeling the inventory problem and demand uncertainty does not deviate from their mild assumptions on the underlying MDP (see their Assumption 1); (2) our piece-wise linear ordering cost and convex inventory mismatch function satisfies the desired structure (see their Assumption 2). In addition to these assumptions, they state a sufficient condition which implies that states with inventory positions outside of a range $[\underline{M}, \overline{M}]$ are dominated by states that fall within the range. Our state space is one-dimensional (consisting of the inventory position). A natural candidate for $\underline{M}$ is zero: all states with lower inventory positions are dominated by $\underline{M} = 0$ as long as infinite backlogging is sub-optimal. For the upper bound $\overline{M}$ we can set its value to the first level at which the expected marginal increase in holding cost exceeds the cost of ordering using the fast source (see also Sheopuri et al. [2010] for a more detailed dual sourcing example). This concludes that the modified dual base-stock policy is infinite-horizon average-cost optimal. ■

Proof of Corollary 1 From Theorem 1 we know that a modified dual base-stock policy is optimal in every period. We focus on period $T-1$ and will show that if $S^f_{2,\mathcal{M}} \leq S^f_{1,\mathcal{M}} < S^s$, which we henceforth assume, then these are the optimal base-stock levels in period $T-1$: $S^f_{2,T-1} = S^f_{2,\mathcal{M}} \leq S^f_{1,T-1} = S^f_{1,\mathcal{M}} < S^s_{T-1} = S^s$. As long as the slow base-stock level is strictly larger than the fast base-stock levels, boundary case C of our proof in which only the fast supplier is used ($q^{s,\star}_{T-1} = 0$ while $q^{f,\star}_{T-1} > 0$) will never occur. As a consequence, the fast base-stock levels are always myopic. From our proof, it is clear that in case A, the fast base-stock levels are given by (33) such that $S^f_{2,T} = S^f_{2,T-1} = S^f_{2,\mathcal{M}} \leq S^f_{1,T} = S^f_{1,T-1} = S^f_{1,\mathcal{M}}$ directly follows from assuming that $c_{T+1} = c_s/\alpha$. In cases B and D, no fast order is placed as these cases correspond to inventory positions above the highest fast base-stock level $S^f_{1,T-1} = S^f_{1,\mathcal{M}} \leq x_{T-2}$.

We now focus on determining S^s_{T-1} which, in contrast to the fast order, does influence the costs in period T. Because $S^f_{1,T-1} = S^f_{1,\mathcal{M}} < S^s_{T-1}$ as per our assumption, the inventory balance equation in period $T-1$: $x_{T-1} = x_{T-2} + q^f_{T-1} + q^s_{T-1} - D_{T-1}$ is equivalent to $x_{T-1} = S^s_{T-1} - D_{T-1}$. Formulating the gradient of the value function in period $T-1$ wrt S^s_{T-1} yields:

$$\frac{dv_{T-1}[x_{T-2}]}{dS^s_{T-1}} = \begin{cases} c_s + \alpha \frac{dv_T[S^s_{T-1} - D_{T-1}]}{dS^s_{T-1}} & \text{if } \; x_{T-2} \leq S^s_{T-1}, \text{[Case A and B of Theorem 1's proof.]} \\ \alpha \frac{dv_T[x_{T-2} - D_{T-1}]}{dS^s_{T-1}} & \text{if } \; S^s_{T-1} \leq x_{T-2}, \text{[Case D of Theorem 1's proof.]} \end{cases} \tag{41}$$

with

$$\begin{aligned}\frac{dv_T[S^s_{T-1} - D_{T-1}]}{dS^s_{T-1}} =& -\alpha c_s + (c_s - mc_f)\bar{F}_{D_{T-1}}[S^s_{T-1} - S^f_{T,2} + k] + \\ & \int_{S^s_{T-1} - S^f_{T,1} + k}^{S^s_{T-1} - S^f_{T,2} + k} L'_T[S^s_{T-1} - \xi_{T-1} + k] f[\xi_{T-1}] d\xi_{T-1} + \end{aligned}$$

$$(c_s - c_f)(\bar{F}_{D_{T-1}}[S^s_{T-1} - S^f_{T,1}] - \bar{F}_{D_{T-1}}[S^s_{T-1} - S^f_{T,1} + k]) +$$
$$\int_0^{S^s_{T-1} - S^f_{T,1}} L'_T[S^s_{T-1} - \xi_{T-1}] f[\xi_{T-1}] d\xi_{T-1} = 0,$$

with ξ_{T-1} representing the random demand in period $T-1$ with density $f[\xi_{T-1}]$ and the gradient of the inventory mismatch, $L'_T[\cdot]$, equal to $hF_{D_T}[\cdot] - b\bar{F}_{D_T}[\cdot]$. The latter function is similar to results obtained in dual sourcing systems with linear sourcing costs (see e.g., Bulinskaya [1964a,b], Whittemore and Saunders [1977]) but includes two new regions, i.e., regions III and IV in panel (d) of Figure 2, that emerge due to the piece-wise linear cost function we adopt. As we assume demand is continuous and iid, the case where the starting inventory position exceeds the slow base-stock level, $S^s_{T-1} < x_{T-2}$ is transient (Case D of Theorem 1's proof). Once the inventory position falls below the slow base-stock level, and assuming demand is non-negative, the inventory position will not exceed the slow base-stock levels in future periods.

Following the same arguments for the preceding periods shows that both the fast base-stock levels and the slow base-stock level remain stationary in all periods $t \in \{1, \cdots, T-2\}$, given by the values specified in (11) and (12), in which we drop the time indices as the base-stock levels are stationary (with the exception of the final period T where no slow order is placed). ■

Proof of Lemma 1 The ordering cost and inventory loss function are convex, and so is their sum (positive superposition) while subtracting the linear convex terminal value function preserves convexity. The objective function is thus convex in y for each x_{T-1}. Lastly, using Theorem A.4 of Porteus [2002] (minimization preserves convexity under specific conditions which a base-stock policy satisfies) we can conclude that also $v_T[\cdot]$ is convex. ■

Proof of Lemma 2 The ordering cost and inventory loss function are convex, and so is their sum (positive superposition) whereas the value function in period T is convex as per Lemma 1. The proof of Theorem 1 identifies four cases. In Case A, the inner minimization of the fast order is convex as minimization preserves convexity (Theorem A.4 of Porteus [2002]). Also for the slow order minimization, Theorem A.4 of Porteus [2002] can be invoked such that $v_{T-1}[\cdot]$ is convex in Case A. The latter argument (minimization preserves convexity) directly holds for cases B, C and D, concluding that convexity of the value function is preserved. ■

Appendix B: Applying our model for general consecutive lead times

To apply our results in a system with general consecutive lead times where $l_f = l$ and $l_s = l+1$, we use the same approach as Fukuda [1964] (see their Eqs 42-46). By defining $L_{t+i}[x] \triangleq \alpha^i \int_0^\infty L_{t+i-1}[x - \xi_{t+i}] f[\xi_{t+i}] d\xi_{t+i}$ for $i \geq 1$, equations containing the inventory mismatch of period t, $L_t[\cdot]$, may be rewritten using the inventory mismatch in period $t + l_f$, $L_{t+l_f}[\cdot]$. All inventory mismatch costs prior to period $t - 1 + l_f$ are 'sunk' as they cannot be influenced by the current order quantities. Nonetheless, the analysis to obtain the optimal policy and parameter remains the same, except that all results are obtained by using convolutions of the demand.

Appendix C: Comment on when the policy parameters can be expressed explicitly

In §5 we demonstrated how the base-stock levels can be obtained whenever the slow base-stock level strictly exceeds the highest fast myopic base-stock level: $S^f_{1,\mathcal{M}} < S^s$. The latter will hold when there exists a solution for (12) where $g[x] = 0$. The function $g[x]$ is a non-decreasing function (as it is the gradient on the convex value function). It will therefore have a solution where $g[x]=0$ when $g[S^f_{1,\mathcal{M}}]<0$. The latter is easily satisfied for a vast amount of numerical dual sourcing settings (we only focus on settings where both sources are used). The only possible case where this condition would not hold is when the policy parameters are ordered as follows: $S^f_2 < S^s < S^f_1$. We tried to numerically construct an example of such a scenario which appeared particularly difficult. A feasible, yet very extreme, scenario would be when demand is normally distributed with mean $\mu = 10$ and standard deviation $\sigma = 2.5$, as in the base scenario of the previous section, while the unit holding cost is also kept fixed at $h = 1$. We then adapt the parameters as follows to create an extreme case where $S^f_2 < S^s < S^f_1$. We set the backlog cost to a small number: $b = 0.01$ such that the service level is smaller than 1% (as in this case sourcing slow is discouraged to avoid holding costs). Subsequently we set the cost gap between both sources equal to zero: $c_s = c_f = 3.8$ (again to encourage sourcing fast). The overtime premium is set infinitely high, $m = +\infty$, to create a large region of inaction in which only k units are ordered. It remains to determine k: For values up to $k = 30$, the slow base-stock level exceeds the fast myopic base-stock level. For values above $k = 30$, we do not find a slow base-stock level that satisfies Corollary 1. Yet, the case is so extreme that it is highly unlikely that a slow order will ever be placed, as the demand needs to exceed the mean demand by 8 standard deviations.

Appendix D: Base-stock sensitivities

Figure 7 visualizes the impact of changing the base capacity k (left panel) and the overtime premium m (middle panel) on the three base-stock levels. The sensitivity of the slow base-stock level S^s wrt the base capacity, k, is obtained by using implicit differentiation on (12) wrt k of both sides while solving for dS^s/dk. Using Leibniz's rule we obtain:

$$\begin{aligned}\frac{dS^s}{dk} &= -\frac{\int_{S^s-S^f_1+k}^{S^s-S^f_2+k} L''_{t+1}[S^s-\xi+k]f[\xi]d\xi}{\int_{S^s-S^f_1+k}^{S^s-S^f_2+k} L''_{t+1}[S^s-\xi+k]f[\xi]d\xi + \int_0^{S^s-S^f_1} L''_{t+1}[S^s-\xi]f[\xi]d\xi},\\ &= -\frac{\int_{S^s-S^f_1+k}^{S^s-S^f_2+k} f[S^s-\xi+k]f[\xi]d\xi}{\int_{S^s-S^f_1+k}^{S^s-S^f_2+k} f[S^s-\xi+k]f[\xi]d\xi + \int_0^{S^s-S^f_1} f[S^s-\xi]f[\xi]d\xi} \leq 0.\end{aligned} \tag{42}$$

We observe partial substitution: $-1 \leq (dS^s/dk) \leq 0$. Similarly, we can show the slow base-stock level S^s increases by taking the derivative with respect to m of both sides while solving for dS^s/dm, which will result in:

$$\frac{dS^s}{dm} = \frac{c_f(1-F[S^s-S^f_2+k])}{\int_{S^s-S^f_1+k}^{S^s-S^f_2+k} L''_{t+1}[S^s-\xi+k]f[\xi]d\xi} = \frac{c_f(1-F[S^s-S^f_2+k])}{(h+b)\int_{S^s-S^f_1+k}^{S^s-S^f_2+k} f[S^s-\xi+k]f[\xi]d\xi} \geq 0. \tag{43}$$

The impact of m on S^s is strictly non-negative but reduces to zero when m increases to $+\infty$. When demand is normally distributed we find the slow base-stock level S^s is non-decreasing in σ:

$$\begin{aligned}\frac{dS^s}{d\sigma} &= \frac{1}{\int_0^{S^s-S^f_1} L''_{t+1}[S^s-\xi]f[\xi]d\xi + \int_{S^s-S^f_1+k}^{S^s-S^f_2+k} L''_{t+1}[S^s-\xi+k]f[\xi]d\xi}\\ &= \frac{1}{(h+b)(\int_0^{S^s-S^f_1} f[S^s-\xi]f[\xi]d\xi + \int_{S^s-S^f_1+k}^{S^s-S^f_2+k} f[S^s-\xi+k]f[\xi]d\xi)} \geq 0.\end{aligned}$$

J. Gijsbrechts, R.N. Boute, S.M. Disney, and J.A. Van Mieghem, (2022),
Volume Flexibility at Responsive Suppliers in Reshoring Decisions: Analysis of a Dual Sourcing Inventory Model,
Accepted for publicaiton in Production and Operations Management, 3/3/2022.

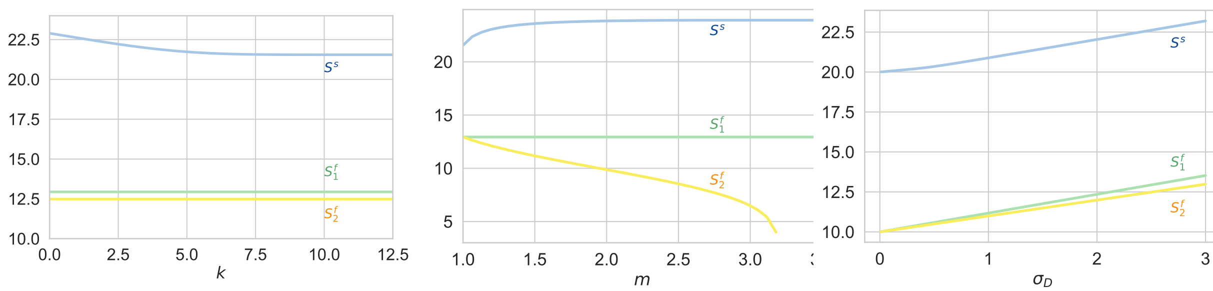

(a) Base capacity, k (b) Overtime premium, m (c) Standard deviation, σ

Figure (7) Impact of base capacity, overtime premium and variability on the base-stock levels. Parameters: $h=1, b=9, c_f=4, c_s=3.8, k=1, m=1.1$, $d_t \stackrel{d}{\sim} \mathcal{N}(\mu=10, \sigma=2.5)$

We also show the impact of the standard deviation when demand is normally distributed (right panel). We see that S_1^f is independent of k and m, which follows directly from (11). The lower fast base-stock level, S_2^f, is independent of the base capacity but does depend on m. The gradient is: $dS_2^f/dm = -c_f f[S^s - S_2^f]/(b+h) \leq 0$, which drops to $-\infty$ when $m \to \infty$ as visually supported in panel (b) of Figure 7. This eliminates the need for overtime; all orders are capped at the base capacity k. To capture the impact of the standard deviation we assume demand is normally distributed with mean μ and standard deviation σ; its CDF and pdf are denoted by Φ and ϕ, respectively. Denote

$$z_1 = \Phi^{-1}\left[\frac{b-(c_f-c_s)}{b+h}\right], \text{ and } z_2 = \Phi^{-1}\left[\frac{b-(mc_f-c_s)}{b+h}\right], \tag{44}$$

then the fast base-stock levels satisfy $S_1^f = \mu + z_1\sigma$ and $S_2^f = \mu + z_2\sigma$. We observe that the fast base-stock levels scale linearly wrt the standard deviation: $dS_1^f/d\sigma = z_1$ and $dS_2^f/d\sigma = z_2$ (see panel c of Figure 7).

Appendix E: Expressions for the sourcing splits

Here we provide explicit expressions to compute the sourcing shares of the fast supply at regular cost:

$$\rho_r^f = \frac{\int_{S^s-S_1^f+k}^{\infty} kf[\xi]d\xi + \int_{S^s-S_1^f}^{S^s-S_1^f+k} (\xi-(S^s-S_1^f))f[\xi]d\xi}{\mu}; \tag{45}$$

for the share of fast supply at overtime cost:

$$\rho_r^o = \frac{\int_{S^s-S_2^f+k}^{\infty} \left(\xi-(S^s-S_2^f+k)\right) f[\xi]d\xi}{\mu}; \tag{46}$$

and for the share of slow supply:

$$\rho^s = \frac{\int_{S^s-S_2^f+k}^{\infty} \left(S^s-S_2^f\right) f[\xi]d\xi + \int_{S^s-S_1^f+k}^{S^s-S_2^f+k} (\xi-k) f[\xi]d\xi + \int_{S^s-S_1^f}^{S^s-S_1^f+k} \left(S^s-S_1^f\right) f[\xi]d\xi + \int_0^{S^s-S_1^f} \xi f[\xi]d\xi}{\mu}. \tag{47}$$