Multi-Echelon Inventory Management for a Non-Stationary Capacitated Distribution Network

Alvaro Maggiar*
SCOT, Amazon.com
maggiara@amazon.com

Irene Song
SCOT, Amazon.com
songiren@amazon.com

Alp Muharremoglu
SCOT, Amazon.com
alpmuha@amazon.com

June 28, 2022

Abstract

We present an inventory management solution for a non-stationary capacitated multi-echelon distribution network involving thousands of products. Assuming backlogged sales, we revisit and leverage the seminal multi-echelon inventory management results in the literature to establish the structural properties of the problem, and derive an efficient and practical solution method. In particular, we describe how the additive separability nature of the objective function can be used in a dynamic programming framework to circumvent the usual curse of dimensionality inherent to high-dimensional inventory management problems. We incorporate features such as asynchronous replenishment intervals at the different echelons, accounting of perishability, capacity constraints (both volume and flow), stochastic lead times, reverse logistics, supply risk and risk-pooling. Illustrative examples demonstrate some of these features.

A version of this model was tested at a large grocery retailer in a two-layer network, resulting in a statistically significant increase of availability, sold units, and revenue, while lowering inventory volume at the hub.

1 Introduction

1.1 Setting

We consider a retailer operating a distribution network wherein products are purchased from vendors and manufacturers, and are sold to customers at spoke locations. The units may be consolidated and transferred at (potentially several layers of) hubs to eventually reach the spokes, or be purchased directly at the spoke locations. While inventory is periodically reviewed at each of the facilities, these review periods are dependent on the arc and the product under consideration, and may not be constant or synchronized. Each facility is subject to capacity constraints, which are both in terms of overall volume held at the facility at any point in time, or in terms of inbound quantities that can be processed in any given period. The asynchronicity of the decisions and timings associated with the different products further complicate the management of these constraints. We mostly assume deterministic lead times between the facilities and from the vendors to the facilities for simplicity of exposition, but also show how to accommodate stochastic lead times.

Figure 1 shows a partial and schematic view of the fullfillment network for a single product. For ease of notation, the facilities are called FCs (short for Fulfillment Center). A common route for a product is to originate at the vendor and be shipped to a hub that is then responsible for replenishing downstream stores where the units will be available for the customers to purchase. This is but one possible stream and we may have units being transported directly from the vendor to a spoke FC, or

*Corresponding author

facilities that serve both as a hub for some spoke FCs and as a customer-facing location. Additionally, while Figure 1 only presents at most a single hub for the FCs, our analysis does not require a restriction to such tree-like structure and a given FC could be fed by several upstream FC. It should also be noted that in some instances, such as that of a grocery system, several network graphs have to be drawn to account for different temperature zones, for products, although often flowing through the same FCs, are subject to separate storage means and thus specific and distinct constraints.

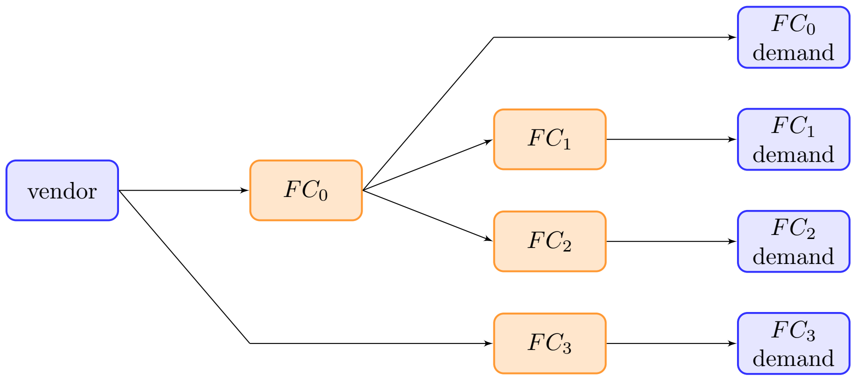

Figure 1: Partial and schematic view of a fullfillment network for a single product.

The purpose of our work is to derive and implement an inventory management solution for this problem so as to maximize the net present value of the business given operating constraints. A particular emphasis is placed on the practicality and scalability of the solution so as to allow for future extensions and modifications. Our approach is to decompose the problem into a collection of single serial lines for each product that are then aggregated to yield the product's network picture. These single-product results can then be themselves aggregated to evaluate network-level results across products in a manner similar to the decomposition-aggregation (DA) heuristic of Rong et al. [2017].

1.2 Model Overview: Partition and Aggregation

As mentioned above, our approach is to consider spoke FC/product combinations and the replenishment routes they follow. The network itself possesses an arborescent structure with vendors feeding FCs, which can themselves feed lower facilities and/or fulfill customer demand. A simple two-layer network example is depicted in Figure 1. A network is then decomposed into serial lines corresponding to the replenishment routes to each of the spoke FCs as illustrated in Figure 2, which shows the serial line corresponding to one of the spoke FCs in Figure 1. Similar serial lines exist for each spoke FC.

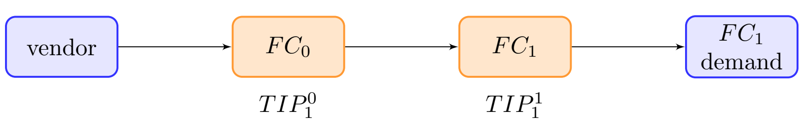

Figure 2: Serial Line corresponding to FC1.

We show that the optimal policies in our serial line formulation are given by echelon order-up-to policies, leading in each period to Target Inventory Positions (TIP) at each echelon. The serial lines can then be pieced back together, and their serial line echelon TIPs are aggregated to yield network echelon TIPs as described in Figure 3. The final FC_0 echelon TIP is equal to a function of the FC_0 echelon TIPs computed on each serial line, which in this case corresponds to the FC_0 TIP for its own demand TIP_0^0, the FC_0 TIP for demand at FC_1 TIP_1^0 and the FC_0 TIP for the demand at FC_2

TIP_2^0. An example of such function could simply be the sum of the echelon TIPs. The network TIP aggregation is covered in Section 4, with a particular interest in risk pooling.

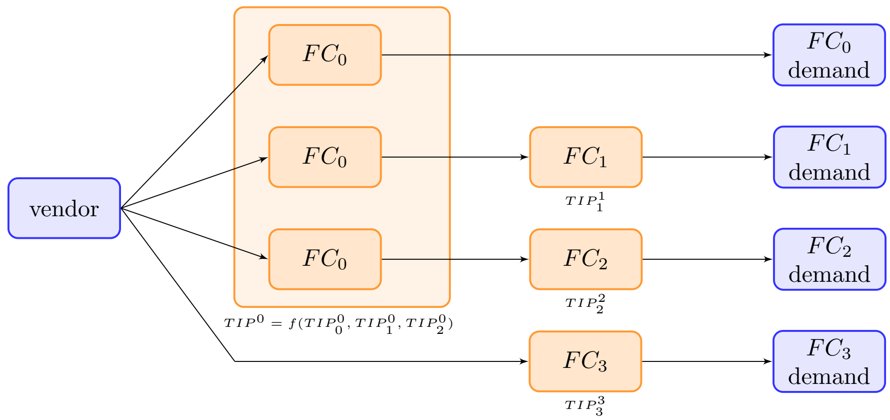

Figure 3: Aggregation of the serial lines.

In this framework, the network is hence composed of multiple serial lines that serve as building blocks for the overarching network to which they belong. This decomposition is appealing as it allows for a great level of parallelization and simplifications. The analysis of the problem is thus divided into two parts: we first consider the base serial line level model in Section 2, revisiting and extending existing results, for which we offer extensions in Section 3, and then consider the aggregation and network-level matters in Section 4, where we study in particular how to handle the capacity constraints that the inventory decisions must abide by.

1.3 Paper Structure

As mentioned above, the theoretical analysis of our model is developed in Sections 2, 3 and 4, at the serial line and network level, respectively. Section 5 details some of the practical considerations when implementing the model, which is then illustrated in Section 6 where we also present results from the implementation of a version of this model at a large grocery retailer.

2 Serial Line: Base Model

2.1 Notation

Serial Line Following the introduction of the problem, most of the work will be carried out at the serial line level. Serial lines are composed of nodes, some of them virtual, that represent the different installations along the serial line. These nodes and installations identify actual FCs, but also intermediate stages between them due to the lead times between FCs and from vendors to FCs, allowing us to keep track of the inventory in transit between physical locations. This, however, can lead to cumbersome notation and identification of the nodes. While it is more intuitive to number the physical (FC) installations only, we still need to number the intermediate virtual nodes to track in-transit inventory. Thus, letting n be the total number of installations, physical or virtual, we identify a node by an integer $i \in [1, n]$, 1 being the closest installation to customers (and necessarily an FC), and n being the most upstream installation.

There are also instances where it is desirable to identify a characteristic or parameter inherent to a physical FC, and we let N be the number of FCs in the serial line, and occasionally use $J \in [1, N]$ to label them. l'_J will denote FC J's installation lead time, in other words the lead time for sending units from FC $J + 1$ to FC J. (In the case of $J = N$, $N + 1$ will refer to the vendor). These concepts are illustrated in Figure 4 on an example with 3 FCs having installation lead times of $l_1^{'} = 4$, $l_2^{'} = 4$ and $l_3^{'} = 4$.

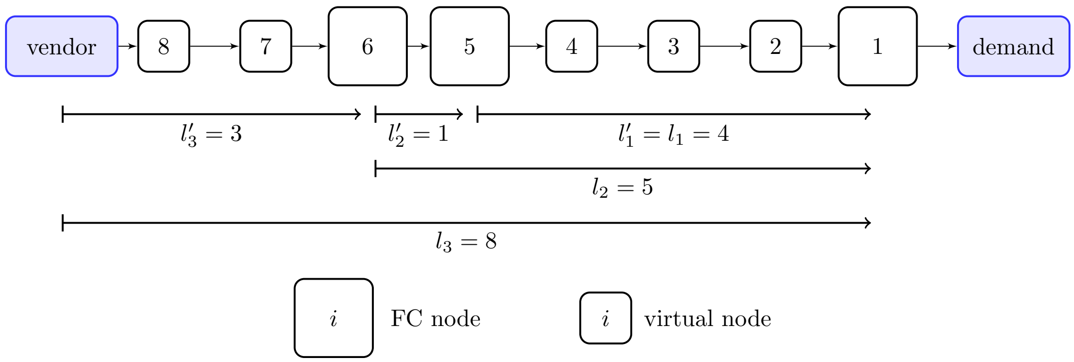

Figure 4: Illustration of the notations on an example with 3 FCs having installation lead times of $l_1^{'} = 4$, $l_2^{'} = 4$ and $l_3^{'} = 4$.

Echelons It is customary in the context of serial systems to consider the concept of *echelon* as opposed to installation. An echelon is a virtual aggregation of an installation and all the installations downstream from it, as illustrated in Figure 5, and many of the concepts or values we use admit both an installation and an echelon value. We will often use a prime symbol (') to refer to an installation value, as opposed to an echelon one. For example, we may define the echelon lead times l_J of FC J as $l_J := \sum_{k=1}^{J} l_k^{'}$.

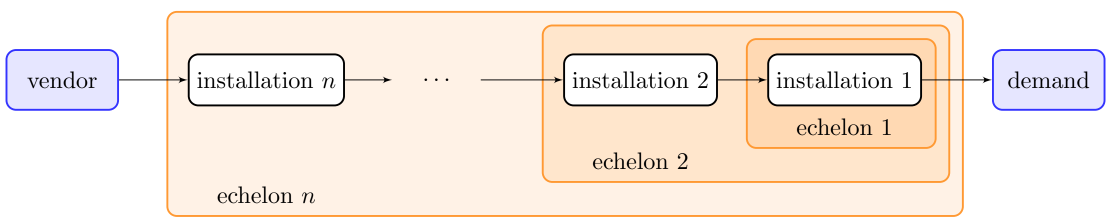

Figure 5: Illustration of the concept of echelon.

Inventory In their seminal paper Clark and Scarf [1960] demonstrated that the optimal policy in their serial system was given by an echelon base-stock policy. This result requires the consideration not of the installation inventory position x^i, but of the echelon inventory position $z^i = \sum_{j=1}^{i} x^j$ representing the sum of units held in all installations downstream from and including i. These two formulations are of course equivalent since we can recover x^i as $x^i = z^i - z^{i-1}$ with the convention that $z^0 = 0$, as illustrated in Figure 6. Working in terms of the echelon inventory position turns out to yield significant formulation benefits that are central to the derivation of the theoretical results as well as the numerical implementation of the model.

Costs The product is purchased from the vendor at a cost c and sold to the customers for a price p. When demand exceeds the inventory available at the customer-facing installation, i.e. installation 1, a backlogging cost b is incurred per unit of unmet demand in each period.

In addition to these product-level costs, there exist network-level costs that are incurred at the various installations. Units at installation i incur an installation holding cost h_i' per unit of volume v

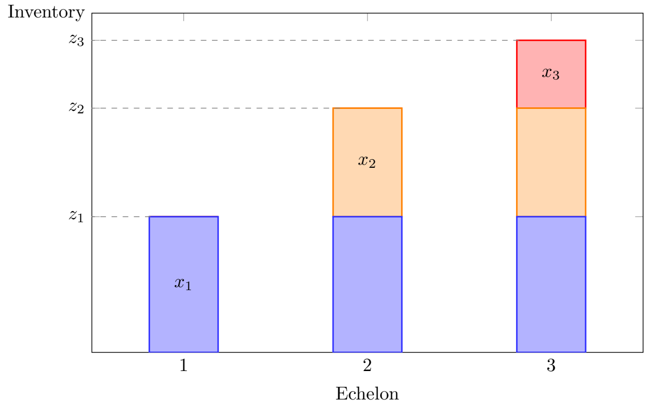

Figure 6: Illustration of the transformation from installation to echelon inventory position.

in each period. We may then define the echelon holding cost $h_i := h'_i - h'_{i+1}$ with the convention that $h_n = h'_n$. Similarly, we let w_i denote the cost of transferring one unit of volume from echelon $i + 1$ to echelon i. To unify the holding and backlogging costs across echelons, and defining D_t as the random demand in period t, we let:

$$H_t^i(z) := \begin{cases} h_i v z & \text{if } i > 1, \\ h_1 v z + (h'_1 v + b)\mathbb{E}\left[(D_t - z)^+\right] & \text{if } i = 1, \end{cases}$$

and then define

$$H_t(\mathbf{z}) := \sum_{i=1}^{n} H_t^i(z^i).$$

Remark 1

The time index on H is really only necessary for the first echelon that involves the backlogging cost and demand distribution at time t, but the notation hints at the fact that the model is readily extendable in the case of time-dependent parameters.

Given the multi-period nature of the problem, we must take into account the time value of money and let γ represent the per-period discounting factor.

Nodes The decomposition of the serial lines into nodes has already highlighted a difference between physical and virtual nodes, where physical nodes correspond to actual FCs and virtual ones to the intermediate steps in between them resulting from transfer lead times. We further characterize these nodes depending on whether a decision is made at their level. In our setting, a decision in a given period represents the setting of a target echelon inventory position $\hat{z}^i$ at node i from its current inventory position z^i, before demand is realized. Note that these decisions happen at nodes immediately following an FC, since once the units are in transit, no action can be taken until they reach their destination facility. When the lead time is greater than one, the decision is taken at a virtual node, while it is taken at an FC node if the lead time is exactly one. As a result, we end up with four types of nodes:

pipeline node: virtual node at which no decision is made,

pure FC node: physical node at which no decision is made, meaning that its lead time is greater than one,

pure decision node: virtual node at which a decision is made,

FC and decision node: physical node at which a decision is made, meaning that its lead time is exactly one.

These are illustrated in Figure 7.

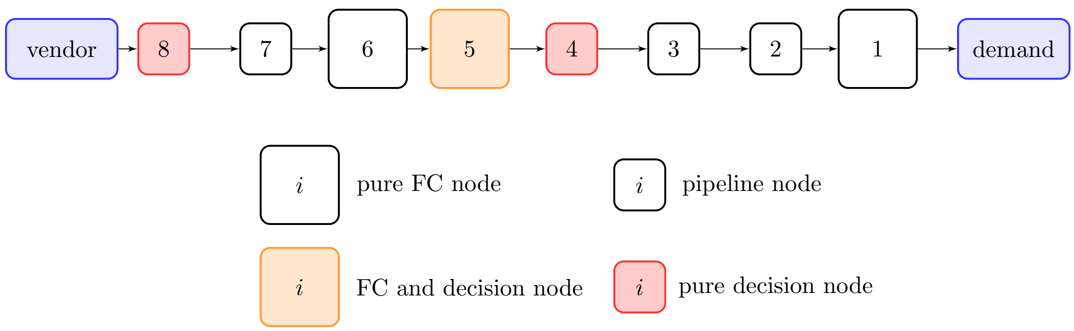

Figure 7: Illustration of the node types.

We define the following sets to identify which nodes fall in which category: $\mathcal{F}$ is the set of FC nodes (both pure and decision), $\mathcal{D}$ is the set of decision nodes (whether virtual or physical), and $\mathcal{P}$ is the set of pipeline nodes. This leads to $\mathcal{F} \setminus \mathcal{D}$ being the set of pure FC nodes, $\mathcal{F} \cap \mathcal{D}$ being the set of FC and decision nodes, and $\mathcal{D} \setminus \mathcal{F}$ being the set of pure decision nodes.

Part of our problem is handling the fact that replenishment orders/transfers cannot be requested in every period, and it is frequent that orders can only be placed to the vendors once a week, say on Sundays, while transfer runs only happen three times a week, on Tuesday, Thursday and Saturday for example. To capture these constraints, we let $\mathbb{T}_t \subset \mathcal{D}$ be the set of decision nodes that allow for a decision to be made at time t. Note that we do not make any periodicity or synchronicity assumption.

We finally let $\mathcal{N}_t := \{\mathcal{F}, \mathcal{D}, \mathcal{P}, \mathbb{T}_t\}$.

Policy An important type of policy in the context of serial systems is an *echelon base-stock policy* characterized by a vector of echelon inventory positions $\mathbf{s} = (s^1, \ldots, s^m)$, where $m = \text{card}(\mathbb{T})$. An echelon base-stock policy is such that each active decision node $i \in \mathbb{T}$ places orders as needed in order to bring its echelon inventory position to s^i.

Notation In most instances, hats will be used to qualify a variable as a decision variable in an optimization problem, such as $\hat{z}^i$ when optimizing for the optimal echelon inventory position. Boldface will refer to a vector, as in $\mathbf{z} = (z^1, \ldots, z^n)$, and $\mathbf{e}$ will represent a vector of ones of appropriate size.

$\mathbb{R}$ and $\mathbb{N}$ denote to the real and natural numbers, respectively. The notation $\mathbb{E}$ will be used to express the expectation of a random variable, e.g. $\mathbb{E}[D_t]$ as the mean of the demand in period t, while $\mathbb{P}$ represents the probability of an event, such as $\mathbb{P}[D_t \le z^1]$ to quantify the probability that the demand will be less than or equal to z^1 in period t.

We let the operators $(.)^+$ and $(.)^-$ refer to the non-negative and non-positive parts of a scalar, so that $x^+ = \max(x, 0)$ and $x^- = \max(-x, 0)$. We also define the operators $\wedge$ and $\vee$ as $x \wedge y = \min(x, y)$ and $x \vee y = \max(x, y)$.

The products in our model are subject to a random demand, and $D_{s,t}$ denotes the stochastic demand from period s to and excluding period t, with the convention that $D_{t,t} = 0$. In many cases we will use single period demands and we thus use the shorthand $D_t := D_{t,t+1}$ to refer to the demand in period t. We let $F_{s,t}$ be the cumulative density function of $D_{s,t}$, so that $F_{s,t}(z) = \mathbb{P}\left[D_{s,t} \le z\right]$.

Echelon inventory positions are by definition non-decreasing, leading the following set:

$$Z = \left\{ \mathbf{z} \in \mathbb{R}^n : \; z^i \ge z^{i-1}, i = 2, \ldots, n \right\}.$$

Related sets for the optimization problems are $\hat{Z}(\mathbf{z})$ and $\hat{Z}(\mathbf{z},\mathcal{N})$ for $z \in Z$ defined as:

$$\hat{Z}(\mathbf{z}) = \left\{\hat{z} \in Z : \; z^i \leq \hat{z}^i \leq z^{i+1}, \; i = 1, \ldots, n-1\right\},$$

$$\hat{Z}(\mathbf{z},\mathcal{N}) = \left\{ \hat{z} \in Z : \; \begin{cases} z^i \leq \hat{z}^i \leq z^{i+1}, & i \in \mathbb{T} \\ \hat{z}^i = z^{i+1}, & i \in \mathcal{P} \cup \mathcal{F} \setminus \mathcal{D} \\ \hat{z}^i = z^i, & i \in \mathcal{D} \setminus \mathbb{T} \end{cases} \right\},$$

with the convention that $z^{n+1} = +\infty$. Note that $\mathbb{T} \cup \mathcal{P} \cup \mathcal{F} \setminus \mathcal{D} \cup \mathcal{D} \setminus \mathbb{T} = \mathcal{P} \cup \mathcal{F} \cup \mathcal{D}$ and these sets are not overlapping, thus representing a partition of the node set.

2.2 Model

2.2.1 Formulation

We now formulate the problem at a serial line level, using the notation laid out in Section 2.1. As detailed above, units are purchased from a vendor and flow through intermediary FCs before reaching a customer facing facility that satisfies a random demand D_t for the product in period t. One major classical assumption that allows for the derivation of the structural results is that unsatisfied demand is backlogged.

> **Assumption 1**
> Unsatisfied demand at the customer-facing facility (node 1) is backlogged at a cost b per unit and per time period.

The following is the order of events in period t:

1. The vector of echelon inventory positions $\mathbf{z}$ is observed,

2. A target echelon inventory position vector $\hat{\mathbf{z}} \in Z(\mathbf{z},\mathcal{N}_t)$ is set, triggering the transfer of $\hat{z}^i - z^i$ units from node $i+1$ to node i to be effective in the following period, on which we incur a transfer cost of γw_i per unit of volume, as well as a purchasing cost of γc per unit on the $\hat{z}^n - z^n$ units ordered from the vendor,

3. A realization of the random demand D_t is observed, and owing to the backlogging assumption, we sell D_t units at a price p,

4. We incur a holding cost of h'_i on each unit of volume left over at node $i \in [1, n]$, and possibly a backlogging penalty b on the $(D_t - z^1)^+$ backlogged units during the period,

5. The ordered and transferred units arrive and the vector of echelon inventory positions moves to $\hat{\mathbf{z}} - D_t\mathbf{e}$.

At the end of the horizon, the cost of acquiring, transferring and holding the remaining on-hand units is recouped or incurred depending on whether units are over or under stocked. This terminal value is defined as:

$$F_T(\mathbf{z}) := \gamma c z^n + \sum_{i=1}^{n} \gamma w_i v z^i - \sum_{i=1}^{n} h_i v z^i - (b + h'_1 v)(z^1)^-.$$

The objective is to minimize the expected discounted cost over a T period horizon (equivalently maximize profit), and the problem is at its core a control problem. Letting $C_t(\mathbf{z}_t)$ be the expected discounted cost from period t to T given an echelon inventory position vector $\mathbf{z}_t$, we aim at minimizing the cost $C_0(\mathbf{z}_0)$ given a starting position $\mathbf{z}_0 \in Z$ and the dynamics of the problem, which can be expressed as follows:

$$\begin{aligned} C_0(\mathbf{z}_0) = \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \mathbb{E}\Bigg[\sum_{t=0}^{T-1} \gamma^t \Big(\gamma c(\hat{z}_t^n - z_t^n) + b\left(z_t^1 - D_t\right)^- \\ + \sum_{i=2}^{n} h_i' v(z_t^i - z_t^{i-1}) + h_1' v(z_t^1 - D_t)^+ + \sum_{i=1}^{n} \gamma w_i v\left(\hat{z}_t^i - z_t^i\right)\Big) - \gamma^T F_T(\mathbf{z}_T)\Bigg] \\ \text{s.t. } \mathbf{z}_{t+1} = \hat{\mathbf{z}}_t - D_t \mathbf{e}. \end{aligned}$$

The first part of the objective function contains the purchasing and backlogging costs, while the second includes the holding and transfer costs. By following a traditional rearrangement of the terms (see e.g. Chao et al. [2015]), it is easily shown that the problem can be equivalently formulated as one without purchasing or transfer costs with redefined holding costs $h_i := h_i + (1-\gamma)w_i$ for $i < n$ and $h_n := h_n + (1-\gamma)w_n + (1-\gamma)\frac{c}{v}$. The transformation is detailed in Appendix A.1. As a result we will for the remainder of this paper consider the following formulation:

$$\begin{aligned} C_0(\mathbf{z}_0) = \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \mathbb{E}\left[\sum_{t=0}^{T-1} \gamma^t \left(\sum_{i=1}^{n} H_t^i(z_t^i)\right)\right] \\ \text{s.t. } \mathbf{z}_{t+1} = \hat{\mathbf{z}}_t - D_t \mathbf{e}, \end{aligned}$$

or equivalently:

$$\begin{aligned} C_0(\mathbf{z}_0) = \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \mathbb{E}\left[\sum_{t=0}^{T-1} \gamma^t H_t(\mathbf{z}_t)\right] \\ \text{s.t. } \mathbf{z}_{t+1} = \hat{\mathbf{z}}_t - D_t \mathbf{e}. \end{aligned} \tag{1}$$

Like most multi-period inventory management problem, this optimization problem lends itself quite naturally to a dynamic programming formulation, whose characterizing optimality equation is then given by Equation (2).

$$C_t(\mathbf{z}_t) = H_t(\mathbf{z}_t) + \gamma \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \mathbb{E}\left[C_{t+1}\left(\hat{\mathbf{z}}_t - D_t \mathbf{e}\right)\right]. \tag{2}$$

2.2.2 Structural Properties

We now jointly derive structural results of both the optimal policy and cost functions. The seminal work of Clark and Scarf [1960] established the optimality of echelon base-stock policies for serial systems, a result extended to an infinite time horizon by Federgruen and Zipkin [1984], and refined by Chen and Zheng [1994]. The result is established in the presence of replenishment intervals in Van Houtum et al. [2007], with the restrictive assumption that the replenishment interval at some echelon is a multiple of the replenishment interval at a lower echelon. We make no such restriction in our case and thus combine and generalize a number of the previous results.

The derivation of the optimal policy structure also reveals and hinges upon a characteristic feature of the cost functions that are shown to be *additive convex*, a property defined in Definition 2.1.

Definition 2.1
A function $C : V \subset \mathbb{R}^n \to \mathbb{R}$ is additive convex if:

$$C(\mathbf{z}) = \sum_{i=1}^{n} c^i(z^i), \quad \forall \mathbf{z} \in V$$

where the functions $c^i,\ i = 1, \ldots, n$ are univariate convex functions.

The additive convexity is not only central to the derivation of the results, but also allows for an efficient numerical implementation of the problem that circumvents the curse of dimensionality by growing linearly with the dimension of the problem, instead of exponentially. The proofs demonstrating additive convexity usually rely on a result by Karush [1959], stated in Lemma 2.1.

Lemma 2.1
Define the following problem:

$$f(a,b) = \min_{a \le z_1 \le \ldots \le z_n \le b} \sum_{i=1}^{n} f^n(z_n),$$

where the functions f^i are convex.
Then, there exist functions f^L and f^U, increasing and decreasing convex (respectively), such that

$$f(a,b) = f^L(a) + f^U(b).$$

Remark 2
Of particular interest is the case of a single dimensional problem $\min_{a \le z \le b} f(z)$. Letting s be the point at which f attains its minimum, we have that $f^L(a) = f(a \vee s)$ and $f^U(b) = f(b \wedge s) - f(s)$, as illustrated in Figure 8.

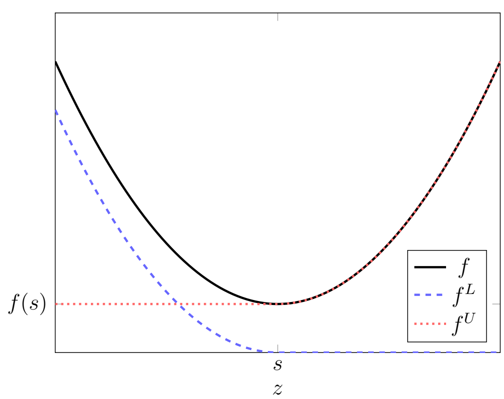

Figure 8: Illustration of Lemma 2.1 for the one-dimensional case.

We now state the main structural result.

Theorem 2.1
Consider the problem formulated in (1). Then:

1. The optimal policy is given by an echelon base-stock policy $\mathbf{s}_t$ for $t \in [1, T-1]$.
2. The cost functions are additive convex, i.e. $\forall t \in [1,T]$, there exist univariate convex functions $c_t^i,\ i = 1, \ldots, n$ such that:

$$C_t(\mathbf{z}) = \sum_{i=1}^{n} c_t^i(z^i).$$

3. For $i \in \mathbb{T}_t$, the components of the base stock policy s_t^i are given by:

$$s_t^i = \arg\min_z \mathbb{E}\left[c_{t+1}^i\left(z - D_t\right)\right].$$

Proof 1
The proof is by recursion.

The theorem is clearly valid at time T, where no decision takes place and we have:

$$C_T(\mathbf{z}) = 0,$$

leading to the following expressions for c_T^i:

$$c_T^i(z^i) = 0 \quad \forall i.$$

Suppose now that the statements are valid for $t+1$, $t \in [1, T-1]$, so that $C_{t+1}(\mathbf{z}) = \sum_{i=1}^n c_{t+1}^i(z^i)$ for some convex functions c_{t+1}^i. According to (2), we have:

$$C_t(\mathbf{z}) = \sum_{i=1}^n H_t^i(z^i) + \gamma \min_{\hat{\mathbf{z}} \in \hat{Z}(\mathbf{z}, \mathcal{N}_t)} \mathbb{E}\left[C_{t+1}\left(\hat{\mathbf{z}} - D_t \mathbf{e}\right)\right]. \tag{3}$$

Consider the optimization part of the expression, which can be expanded as:

$$\begin{aligned}
&\min_{\hat{\mathbf{z}} \in \hat{Z}(\mathbf{z}, \mathcal{N}_t)} \mathbb{E}\left[C_{t+1}\left(\hat{\mathbf{z}} - D_t \mathbf{e}\right)\right] && (4)\\
&= \min_{\hat{\mathbf{z}} \in \hat{Z}(\mathbf{z}, \mathcal{N}_t)} \sum_{i=1}^n \mathbb{E}\left[c_{t+1}^i\left(\hat{z}^i - D_t\right)\right]\\
&= \sum_{i \in \mathbb{T}_t} \min_{z^i \le \hat{z}^i \le z^{i+1}} \mathbb{E}\left[c_{t+1}^i\left(\hat{z}^i - D_t\right)\right] + \sum_{i \in \mathcal{P} \cup \mathcal{F} \setminus \mathcal{D}} \mathbb{E}\left[c_{t+1}^i\left(z^{i+1} - D_t\right)\right] + \sum_{i \in \mathcal{D} \setminus \mathbb{T}_t} \mathbb{E}\left[c_{t+1}^i\left(z^i - D_t\right)\right]\\
&= \sum_{i \in \mathbb{T}_t} \min_{z^i \le \hat{z}^i \le z^{i+1}} \tilde{c}_t^i\left(\hat{z}^i\right) + \sum_{i \in \mathcal{P} \cup \mathcal{F} \setminus \mathcal{D}} \tilde{c}_t^i\left(z^{i+1}\right) + \sum_{i \in \mathcal{D} \setminus \mathbb{T}_t} \tilde{c}_t^i\left(z^i\right), && (5)
\end{aligned}$$

where $\tilde{c}_t^i(z) := \mathbb{E}\left[c_{t+1}^i\left(z - D_t\right)\right]$, and we assume that $z^{n+1} = +\infty$.

We observe that the optimization problem has been decomposed into the sum of independent convex bounded univariate optimization problems that clearly lead to a echelon base-stock policy $\mathbf{s}_t = \left\{s_t^j : j \in \mathbb{T}_t\right\}$, where s_t^i is the minimizer of $\tilde{c}_t^i(z)$ for $i \in \mathbb{T}_t$. Using Remark 2.1, we can write:

$$\min_{z^i \le \hat{z}^i \le z^{i+1}} \tilde{c}_t^i\left(\hat{z}^i\right) = \tilde{c}_t^i(z^i \vee s_t^i) + \tilde{c}_t^i(z^{i+1} \wedge s_t^i) - \tilde{c}_t^i(s_t^i).$$

Plugging this equality into (5) and then back into (3), we find:

$$C_t(\mathbf{z}_t) = \sum_{i=1}^n H_t^i(z_i) + \gamma\left(\sum_{i \in \mathcal{P} \cup \mathcal{F} \setminus \mathcal{D}} \tilde{c}_t^i\left(z_t^{i+1}\right) + \sum_{i \in \mathcal{D} \setminus \mathbb{T}_t} \tilde{c}_t^i\left(z_t^i\right) + \sum_{i \in \mathbb{T}_t} \tilde{c}_t^i(z^i \vee s_t^i) + \tilde{c}_t^i(z^{i+1} \wedge s_t^i) - \tilde{c}_t^i(s_t^i)\right), \tag{6}$$

showing that $C_t(\mathbf{z})$ is additive convex as the sum of univariate convex functions. •

The additive convexity of the cost functions is a salient property of the model, and using the relations derived in the proof of Theorem 2.1, we can relate the univariate functions from one period to the next.

Theorem 2.2
Let $t \in [0, T-1]$, $\tilde{c}_t^i(z) := \mathbb{E}\left[c_{t+1}^i\left(z - D_t\right)\right]$, and $\mathbf{s}_t := \left\{s_t^i : i \in \mathbb{T}_t\right\}$ where:

$$s_t^i := \arg\min_z \tilde{c}_t^i(z).$$

Then:

$$c_t^i(z) = H_t^i(z) + \gamma \begin{cases} \tilde{c}_t^{i-1}(z \wedge s_t^{i-1}), & \text{if } (i-1) \in \mathbb{T}_t \\ \tilde{c}_t^{i-1}(z), & \text{else if } i \in \mathcal{P} \cup \mathcal{D} \end{cases} + \gamma \begin{cases} \tilde{c}_t^i(z \vee s_t^i) - \tilde{c}_t^i(s_t^i)\mathbf{1}_{i<n}, & i \in \mathbb{T}_t \\ \tilde{c}_t^i(z), & i \in \mathcal{D} \setminus \mathbb{T}_t \end{cases}. \tag{7}$$

Proof 2
Theorem 2.2 is obtained from (6) by observing that $(i-1) \in \mathcal{P} \cup \mathcal{F} \setminus \mathcal{D} \implies i \in \mathcal{P} \cup \mathcal{D}$. •

Proposition 2.1 – Solution Method
Theorem 2.2 provides with a method for solving the problem through a recursive approach:
For $t = T-1, \ldots, 0$:

1. For $j \in \mathbb{T}_t$, solve for $s_t^j = \arg\min_z \mathbb{E}\left[c_{t+1}^j(z - D_t)\right]$,
2. For $i = 1, \ldots, n$, evaluate c_t^i using the relations presented in Theorem 2.2:

$$c_t^i(z) = H_t^i(z) + \gamma \begin{cases} \tilde{c}_t^{i-1}(z \wedge s_t^{i-1}), & \text{if } (i-1) \in \mathbb{T}_t \\ \tilde{c}_t^{i-1}(z), & \text{else if } i \in \mathcal{P} \cup \mathcal{D} \end{cases} + \gamma \begin{cases} \tilde{c}_t^i(z \vee s_t^i) - \tilde{c}_t^i(s_t^i)\mathbf{1}_{i<n}, & i \in \mathbb{T}_t \\ \tilde{c}_t^i(z), & i \in \mathcal{D} \setminus \mathbb{T}_t \end{cases}.$$

Additionally, observe that the calculations in each step can be performed in parallel since they only require the knowledge of the functions at $t+1$ and do not depend on the results at other echelons at time t.

We expand on a dynamic programming implementation of the problem in Section 5.

2.3 Alternative Formulations and Solution Methods

2.3.1 Introduction

Proposition 2.1 offers a method to solve the problem. It requires the solving of (approximately) N optimization problems, where we recall that N is the number of FCs, and the evaluation of n functions. Using Theorem 2.2, we may derive alternative solution methods that generalize the formulation of Chen and Zheng [1994], also presented in Gallego and Zipkin [1999] and Shang and Song [2003], to a non-stationary and non-periodic setting. These alternative formulations still require solving N optimization problems, but only require the evaluation of as many functions (N instead of n) at the expense of more onerous evaluations. The formulations are also useful in deriving bounds (see Section 2.4) on the optimal policy. These formulations are equivalent, but differ in their accounting. While our initial dynamic programming formulation includes the costs incurred in the current period t, and which are actually unaffected by decisions made in t, we can alternatively contemplate an accounting scheme in which costs incurred because of a decision are affected to that decision, as originally introduced in Karlin and Scarf [1958] for the single echelon problem.

By rolling out the costs caused by an inventory decision and assigning them to the corresponding decision, we may reduce the components to keep track of only those at the decision nodes. Since there is exactly one decision node per FC (possibly confounded with the FC node if its lead time is 1), we use a labeling that refers to the FC that is under consideration. In other words, what we will refer to as node $J \in [1, N]$ corresponds to node l_J in the full serial line representation with virtual nodes. It will also be useful to let $D_t^J := D_{t,t+l'_J}$ be the demand during installation J's lead time, as well as to define the following constants:

$$\theta_J := \sum_{k=0}^{l'_J - 1} \gamma^k h_{l_J - k}, \qquad \delta_t^J := \sum_{k=0}^{l'_J - 1} \gamma^k h_{l_J - k} \mathbb{E}\left[D_{t,t+1+k}\right],$$

which we then use to define the expected time-discounted echelon cost $\Theta_t^J(z)$:

$$\begin{aligned} \Theta_t^J(z) &:= \mathbb{E}\left[\sum_{k=0}^{l'_J - 1} \gamma^k H_{t+k}^{l_J - k}(z - D_{t,t+1+k})\right] \\ &= \theta_J v z + \gamma^{l'_J - 1}(b + h'_1 v)\mathbb{E}\left[\left(z - D_t^1\right)^-\right] \mathbf{1}_{\{J=1\}} - \delta_t^J. \end{aligned}$$

This cost represents the discounted sum of all costs incurred at echelon J when a target position z is decided at time $t-1$. The resulting transferred units lead to a position $z - D_t$ at time t at node l_J, incurring a holding cost $H_t^{l_J}(z - D_t)$. These units then progress one (pipeline) node at a time until they reach FC J, incurring holding costs in each of those virtual nodes as well as the at the FC. In the absence of discounting, note that $\theta_J = h'_{l_{J-1}+1} - h'_{l_J+1}$, which is the difference in installation costs between consecutive FCs, and what is traditionally used as echelon cost in the literature.

We present two variations on this alternative formulation, which are obtained by rolling out the terms bracketed in Theorem 2.2. The first bracket in (7) contains costs incurred at a lower echelon and thus corresponds to a vertical unrolling of the costs (if we picture the serial line as a vertical), while the second bracket corresponds to costs incurred at the same echelon and thus leads to a horizontal unrolling of the costs. We first present the results obtained by unrolling only those costs in the first bracket, following which we unroll the costs from the second bracket. Figure 9 presents a schematic view of these formulations by highlighting the cost terms taken into account by each formulation.

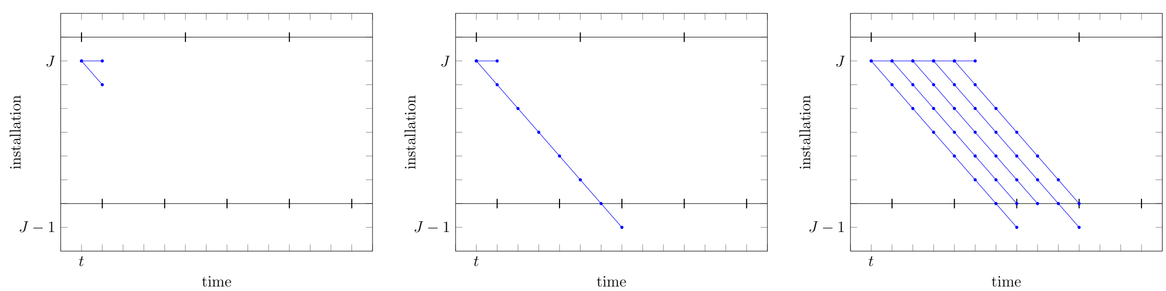

Figure 9: Schematic illustration of the costs taken into account by each formulation: (left) original formulation, (center) first alternative formulation, (right) second alternative formulation. The horizontal lines correspond to physical installations, and their ticks to periods in which a transfer can be requested from the lower installation.

2.3.2 First Alternative Formulation

Using the additional notation defined above, we may reach an alternative formulation of the problem, presented in Theorem 2.3.

Theorem 2.3
Let $t \in [0, T-1]$, $\tilde{c}_t^i(z) := \mathbb{E}\left[c_{t+1}^i\left(z - D_t\right)\right]$, and $\mathbf{s}_t := \left\{s_t^i : \ i \in \mathbb{T}_t\right\}$ where:

$$s_t^i := \arg\min_z \tilde{c}_t^i(z).$$

Then, for $J \in [1, N]$:

$$\tilde{c}_t^J(z) = \Theta_t^J(z) + \gamma^{l'_J} \begin{cases} \tilde{c}_{t+l'_J}^{J-1}\left((z - D_t^J) \wedge s_{t+l'_J}^{J-1}\right), & \text{if } J-1 \in \mathbb{T}_{t+l'_J} \\ 0, & \text{else} \end{cases}$$
$$+ \gamma \begin{cases} \tilde{c}_{t+1}^J(z - D_t \vee s_{t+1}^J) - \tilde{c}_t^J(s_{t+1}^J), & J \in \mathbb{T}_{t+1} \\ \tilde{c}_{t+1}^J(z - D_t), & J \in \mathcal{D} \setminus \mathbb{T}_{t+1} \end{cases}.$$

Proof 3
Observe first that when $l'_J = 1$, the formulation in Theorem 2.3 is equivalent to that from Theorem 2.2, since we are only concerned with decision nodes (and observe that Theorem 2.3 expresses $\tilde{c}$ directly, while Theorem 2.2 gives an expression for c).

Assume then that $l'_J > 1$, meaning that $(l'_J - 1) \notin \mathcal{D}$. Theorem 2.2 gives:

$$c_t^{l_J}(z) = H_t^{l_J}(z) + \gamma \begin{cases} \tilde{c}_t^{l_J}(z \vee s_t^{l_J}) - \tilde{c}_t^i(s_t^{l_J}), & l_J \in \mathbb{T}_t \\ \tilde{c}_t^{l_J}(z), & l_J \in \mathcal{D} \setminus \mathbb{T}_t \end{cases} + \gamma \tilde{c}_t^{l_J - 1}(z).$$

Since $\tilde{c}_t^{l_J-1}(z) = \mathbb{E}\left[c_{t+1}^{l_J-1}(z - D_t)\right]$. Using Theorem 2.2, we have:

$$\tilde{c}_t^{l_J-1}(z) = \mathbb{E}[H_{t+1}^{l_J-1}(z - D_t)] + \gamma \begin{cases} \tilde{c}_{t+1}^{l_J-2}(z - D_t \wedge s_{t+1}^{l_J-2}), & \text{if } (l_J - 2) \in \mathbb{T}_{t+1} \\ \tilde{c}_{t+1}^{l_J-2}(z - D_t), & \text{else if } l_J - 1 \in \mathcal{P} \cup \mathcal{D} \end{cases},$$

which can then be substituted back in the above equality to yield:

$$\begin{aligned} c_t^{l_J}(z) = H_t^{l_J}(z) + \gamma \mathbb{E}[H_{t+1}^{l_J-1}(z - D_t)] + \gamma^2 &\begin{cases} \tilde{c}_{t+1}^{l_J-2}(z - D_t \wedge s_{t+1}^{l_J-2}), & \text{if } (l_J - 2) \in \mathbb{T}_{t+1} \\ \tilde{c}_{t+1}^{l_J-2}(z - D_t), & \text{else if } l_J - 1 \in \mathcal{P} \cup \mathcal{D} \end{cases} \\ + \gamma &\begin{cases} \tilde{c}_t^{l_J}(z \vee s_t^{l_J}) - \tilde{c}_t^i(s_t^{l_J}), & l_J \in \mathbb{T}_t \\ \tilde{c}_t^{l_J}(z), & l_J \in \mathcal{D} \setminus \mathbb{T}_t \end{cases}. \end{aligned}$$

If $l'_J = 2$ we obtain the desired formulation (since in that case $(l_J - 1) \in \mathcal{F}$ and not $\mathcal{P} \cup \mathcal{D}$), else the final term in the summation is $\tilde{c}_{t+1}^{l_J-2}(z - D_t)$, which we can similarly expand, and iteratively repeat the process l'_J times, yielding the required expression. •

Proposition 2.2 – First Alternative Solution Method

Theorem 2.3 provides with an alternative method for solving the problem through a recursive approach:
For $t = T - 1, \dots, 0$, and for $J \in [1, N]$:

1. Evaluate $\tilde{c}_t^J$ using the relations presented in Theorem 2.3:

$$\begin{aligned} \tilde{c}_t^J(z) = \Theta_t^J(z) + \gamma^{l'_J} &\begin{cases} \tilde{c}_{t+l'_J}^{J-1}\left((z - D_t^J) \wedge s_{t+l'_J}^{J-1}\right), & \text{if } J - 1 \in \mathbb{T}_{t+l'_J} \\ 0, & \text{else} \end{cases} \\ + \gamma &\begin{cases} \tilde{c}_{t+1}^J(z - D_t \vee s_{t+1}^J) - \tilde{c}_t^J(s_{t+1}^J), & J \in \mathbb{T}_{t+1} \\ \tilde{c}_{t+1}^J(z - D_t), & J \in \mathcal{D} \setminus \mathbb{T}_{t+1} \end{cases}. \end{aligned}$$

2. if $J \in \mathbb{T}_t$, solve for $s_t^j = \arg\min_z \tilde{c}_t^J(z)$,

2.3.3 Second Alternative Formulation

We now reach a second alternative formulation by rolling out the second bracket in (7), or equivalently the second bracket in Theorem 2.3. To that end, it will be convenient to define r_t^J at time t and for echelon J as the time until the next decision at that echelon. In other words, $r_t^J := \min\{t' : t' > t, J \in \mathbb{T}_{t'}\}$.

Theorem 2.4

Let $t \in [0, T-1]$, $\tilde{c}_t^i(z) := \mathbb{E}\left[c_{t+1}^i(z - D_t)\right]$, and $\mathbf{s}_t := \left\{s_t^i : \ i \in \mathbb{T}_t\right\}$ where:

$$s_t^i := \arg\min_z \tilde{c}_t^i(z).$$

Then, for $J \in [1, N]$:

$$\begin{aligned} \tilde{c}_t^J(z) = \sum_{k=0}^{r_t^J - 1} \gamma^k &\left(\Theta_{t+k}^J(z - D_{t,t+k}) + \gamma^{l'_J} \begin{cases} \tilde{c}_{t+l'_J+k}^{J-1}\left((z - D_{t+k}^J) \wedge s_{t+l'_J+k}^{J-1}\right), & \text{if } J - 1 \in \mathbb{T}_{t+l'_J+k} \\ 0, & \text{else} \end{cases}\right) \\ + \gamma^{r_t^J} &\left(\tilde{c}_{t+r_t^J}^J\left(z - D_{t+r_t^J} \vee s_{t+r_t^J}^J\right) - \tilde{c}_{t+r_t^J}^J\left(s_{t+r_t^J}^J\right)\right). \end{aligned}$$

Proposition 2.3 – Second Alternative Solution Method
Theorem 2.4 provides with an alternative method for solving the problem through a recursive approach:
For $t = T-1, \dots, 0$, and for $J \in \mathbb{T}_t$:

1. Evaluate $\tilde{c}_t^J$ using the relations presented in Theorem 2.4:

$$\tilde{c}_t^J(z) = \sum_{k=0}^{r_t^J-1} \gamma^k \left(\Theta_{t+k}^J(z - D_{t,t+k}) + \gamma^{l'_J} \begin{cases} \tilde{c}_{t+l'_J+k}^{J-1}\left((z - D_{t+k}^J) \wedge s_{t+l'_J+k}^{J-1}\right), & \text{if } J-1 \in \mathbb{T}_{t+l'_J+k} \\ 0, & \text{else} \end{cases} \right)$$
$$+ \gamma^{r_t^J} \left(\tilde{c}_{t+r_t^J}^J\left(z - D_{t+r_t^J} \vee s_{t+r_t^J}^J\right) - \tilde{c}_{t+r_t^J}^J\left(s_{t+r_t^J}^J\right)\right).$$

2. solve for $s_t^j = \arg\min_z \tilde{c}_t^J(z)$.

2.3.4 Comparison of the Formulations

Propositions 2.1, 2.2 and 2.3 are extremely similar in nature and offer different ways of solving the problem. In Proposition 2.1, we make use of the entire discretization of the serial line, evaluating both physical and virtual nodes; whereas in Propositions 2.2 and 2.3, the same number of optimization sub-problems are solved, but only one evaluation is required per FC. The trade-off is in the cost associated with the evaluations carried out in these methods. The evaluations in Proposition 2.2 and 2.3 require that expectations over several periods be evaluated. These can be numerically expensive, and forecast distributions for spans greater than one may not be readily available. On the other hand, Proposition 2.1 only requires single-period forecasts, but that many more functions be evaluated and stored. It is likely that Proposition 2.1 will be favored in cases of problems with shorter lead times, while Propositions 2.2 and 2.3 might be more beneficial in the presence of longer lead times.

Similarly, comparing Proposition 2.3 and 2.2, we observe that the evaluations required in Proposition 2.3 are more expensive than the ones in Proposition 2.2. However, not only does Proposition 2.3 only require that only one function be evaluated by echelon, it also only requires that it only be evaluated at times when a decision is to be made at the echelon. Thus, if decision schedules are such that decisions are far apart, Proposition 2.3 can become more advantageous than Proposition 2.2.

Nonetheless, the benefits of the full discretization of the serial line extend beyond the above considerations. In particular, they allow us to keep track of the inventory levels at a more granular level, which is especially salient when it comes to evaluating on-hand levels and inbound flows at the different installations when capacity considerations come into play (see Section 4).

2.3.5 Stationary Case

The formulation we have presented are very general and do not require that demand be stationary, or that ordering schedules abide by any periodicity or synchronicity, as is usually customary in the literature. We briefly show here how these particular cases can be recovered from our general formulations. We thus assume in this section that demand is stationary and then first consider the case where orders can be placed in every period for all echelons, and then the case where the review period at any echelon is a multiple of the review period at a lower echelon. Because of the stationarity assumption, we temporarily drop the time indexing of variables.

Single Period Review When inventory decisions are possible in every period and all echelons, the function evaluation in Proposition 2.2 reduces to:

$$\tilde{c}^J(z) = \Theta^J(z) + \gamma^{l'_J} \tilde{c}^{J-1}\left((z - D^J) \wedge s^{J-1}\right) + \gamma\left(\tilde{c}^J(z - D \vee s^J) - \tilde{c}^J(s^J)\right).$$

The order-up-to level s^J minimizes this function. It is easily seen that in this stationary setting, the last term can be done away with for it does not affect the procedure, yielding the following simplified expression:

$$\tilde{c}^J(z) = \Theta^J(z) + \gamma^{l'_J} \tilde{c}^{J-1}\left((z - D^J) \wedge s^{J-1}\right).$$

When plugged back into Proposition 2.2, the solution method corresponds exactly to the classical one initially presented in Chen and Zheng [1994] and detailed in [Shang and Song, 2003, Section 2].

Multiple Period Review Similarly to single period review scenario, the stationarity of the problem in a multiple period review case (where review periods at an upper echelon is a multiple of the review period at its lower echelon) allows us to simplify Proposition 2.3 as follows:

$$\tilde{c}^J(z) = \sum_{k=0}^{r^J-1} \gamma^i \left(\Theta^J(z - D_{0,k}) + \gamma^{l'_J} \begin{cases} \tilde{c}^{J-1}\left((z - D^J) \wedge s^{J-1}\right), & \text{if } J-1 \in \mathbb{T}_{l'_J+k} \\ 0, & \text{else} \end{cases}\right).$$

2.4 Bounds

We derive in this section bounds on the base-stock policies $\mathbf{s}_t$, as well as on the gradient of the cost functions $\tilde{c}_t^J$, where we use the same shorthand notation as in Section 2.3, namely $J := l_J$. These bounds can be used to verify numerical implementations and construct heuristics to approximate the optimal policies. Such bounds have been derived for serial systems in different contexts. For example, in the case of a stationary system in which decisions can be made in each period, Shang and Song [2003] derive newsvendor-type bounds, and Gallego and Özer [2005] propose a similar newsvendor heuristic. In the case of a stationary system in which decisions at each echelon are made periodically with the constraint that any echelon's review period be a multiple of its immediately preceding echelon review period, Van Houtum et al. [2007] present closed form solutions for the base-stock policies. Our setting is more general, requiring neither stationarity, nor periodicity/synchronicity of the order schedule. A consequence of these relaxed conditions is that the bounds we derive tend to be weaker than in other cases.

2.4.1 First Echelon Decision

We start by considering the first echelon decision, which happens at node $i = l_1$.

Theorem 2.5
Let:

$$\underline{\tilde{c}}_t^{l_1}(z) := \sum_{k=0}^{r_t^1-1} \gamma^k \Theta_{t+k}^1(z - D_{t,t+k}),$$

and for t such that $l_1 \in \mathbb{T}_t$, let $\bar{s}_t^{l_1} := \arg\min_z \underline{\tilde{c}}_t^{l_1}(z)$, which satisfies

$$\left(\sum_{k=0}^{r_t^1-1} \gamma^k\right) \theta_1 + (b + h'_1) \sum_{s=l_1}^{l_1+r_t^1-1} \gamma^{s-1} \left(F_{t,t+s}(\bar{s}_t^{l_1}) - 1\right) = 0.$$

Then:

1. The gradient of $\underline{\tilde{c}}_t^{l_1}$ is a lower bound on the gradient of $\tilde{c}_t^{l_1}(z)$, i.e. $\frac{\partial \underline{\tilde{c}}_t^{l_1}}{\partial z}(z) \leq \frac{\partial \tilde{c}_t^{l_1}}{\partial z}(z)$ for all z,
2. $\bar{s}_t^{l_1}$ is an upper bound on $s_t^{l_1}$.

Proof 4
By Theorem 2.4, we have:

$$\tilde{c}_t^{l_1}(z) := \sum_{k=0}^{r_t^1-1} \gamma^k \Theta^1_{t+k}(z - D_{t,t+k}) + \gamma^{r_t^1} \left(\tilde{c}^{l_1}_{t+r_t^1} \left(z - D_{t+r_t^1} \vee s^{l_1}_{t+r_t^1} \right) - \tilde{c}^{l_1}_{t+r_t^1} \left(s^{l_1}_{t+r_t^1} \right) \right).$$

The first term in this equality corresponds to $\underline{\tilde{c}}_t^{l_1}(z)$, while the second is a non-decreasing function of z so that its derivative is non-negative. Consequently, we have:

$$\frac{\partial \underline{\tilde{c}}_t^{l_1}}{\partial z}(z) \leq \frac{\partial \tilde{c}_t^{l_1}}{\partial z}(z).$$

Given the convexity of the function and the unequality on its derivative derived given by the above equation, it follows immediately that the minimizer $\bar{s}_t^{l_1}$ of $\underline{\tilde{c}}_t^{l_1}$ is a upper bound of $s_t^{l_1}$. •

Remark 3
Two remarks can be made in the stationary case:

1. $\bar{s}^{l_1}$ is not only an upper bound, but is actually equal to the optimal order-up-to level s^{l_1} since the second term in the expression of $\tilde{c}^{l_1}$ in the proof can actually be ignored as in Section 2.3.5,
2. Assuming a discount factor of 1, the optimality condition on s^{l_1} reads:

$$\frac{1}{r^1} \sum_{s=l_1}^{l_1+r^1-1} F_{t,t+s}\left(\bar{s}_t^{l_1}\right) = \frac{b + h'_{l_1+1}}{b + h'_1},$$

whence we recover the formula for s^{l_1} presented in [Van Houtum et al., 2007, Corollary 6].

2.4.2 Upper Echelon Decisions

Using the same logic as for the first echelon, we can derive the following bounds at upper echelons.

Theorem 2.6
For $J \in [2, N]$, let:

$$\underline{\tilde{c}}_t^J(z) = \sum_{k=0}^{r_t^J-1} \gamma^k \left(\Theta^J_{t+k}(z - D_{t,t+k}) + \gamma^{l'_J} \begin{cases} \underline{\tilde{c}}^{J-1}_{t+l'_J+k}\left((z - D^J_{t+k}) \wedge \bar{s}^{J-1}_{t+l'_J+k} \right), & \text{if } J-1 \in \mathbb{T}_{t+l'_J+k} \\ 0, & \text{else} \end{cases} \right),$$

and for t such that $J \in \mathbb{T}_t$, let $\bar{s}_t^J := \arg\min_z \underline{\tilde{c}}_t^J(z)$, which satisfies

$$\left(\sum_{k=0}^{r_t^J-1} \gamma^k \right) \theta_J + \sum_{k=0}^{r_t^J-1} \gamma^{l'_J} \begin{cases} \frac{\partial}{\partial z} \underline{\tilde{c}}^{J-1}_{t+l'_J+k}\left(\left(\bar{s}_t^J - D^J_{t+k} \right) \wedge \bar{s}^{J-1}_{t+l'_J+k} \right), & \text{if } J-1 \in \mathbb{T}_{t+l'_J+k} \\ 0, & \text{else} \end{cases} = 0.$$

Then:

1. The gradient of $\underline{\tilde{c}}_t^J$ is a lower bound on the gradient of $\tilde{c}_t^J(z)$, i.e. $\frac{\partial \underline{\tilde{c}}_t^J}{\partial z}(z) \leq \frac{\partial \tilde{c}_t^J}{\partial z}(z)$ for all z,
2. $\bar{s}_t^J$ is an upper bound on s_t^J.

Remark 4
Assuming stationarity and that the review periods at upper echelons are multiples of the ones at lower echelons, we find that the upper bounds in Theorem 2.6 are actually tight, and we can recover the optimality conditions presented in [Van Houtum et al., 2007, Corollary 6].

3 Serial Line: Extensions

We present in this section some examples to the base model developed in Section 2. These cover: stochastic lead times in Section 3.1, perishability in Section 3.2, reverse logistics in Section 3.3 and random supply in Section 3.4.

3.1 Stochastic Lead Times

While we consider stochastic demands, lead times have so far been considered deterministic. In practice, shipment times are subject to variability and their arrivals cannot be claimed to be known exactly. A more realistic approach should thus allow for stochasticity in the lead times as well. Stochastic lead times introduce an additional degree of complexity that is usually handled following the early work of Kaplan [1970] where the first results incorporating lead times are presented. Central to the derivation of these results are the assumptions that orders do not cross and that their arrival probabilities are independent of the number and size of outstanding orders. Nahmias [1979] uses the same assumptions to analyze a lost-sales model and provides an alternative interpretation of the lead times as resulting from a sequence of auxiliary random variables that describe the age of the youngest order arriving at an installation in a given period. This interpretation is then used by Ehrhardt [1984] to revisit the results of Kaplan [1970]. In the multi-echelon setting, the single-unit decomposition approach of Muharremoglu and Tsitsiklis [2008] allows for a lead time model similar to Kaplan [1970].

To allow for stochastic lead times, we consider a lead time model similar to Kaplan's. In other words, we assume that orders and transfers do not cross and that lead times are independent of the number and size of outstanding orders. Our model however allows for dependencies of the lead times between different echelons in the same period, but we will assume independence for ease of exposition. Lead times are bounded above and we let L'_J be the stochastic lead time into installation J and l'_J be here an upper bound on L'_J. Kaplan's model, as interpreted by Nahmias, arises from a sequence of random variables $\{A_t^J\}$ assuming values in $\{1, \dots, l'_J\}$ with $p_k^J = \mathbb{P}\left[A_t^J = k\right]$. A_t^J characterizes the age of the oldest order arriving at the end of period t. As a result, the event $\{A_t^J = 1\}$ corresponds to a realization of a lead time of 1, and more generally, the event $\{A_t^J > 0, A_{t+1}^J > 1, \dots, A_{t+k-1}^J > k-1, A_{t+k}^J \leq k\}$ corresponds to the event $\{L'_J = k\}$.

The stochastic lead times modify the state transition. While in the deterministic lead time case we had $\mathbf{z}_{t+1} = \hat{\mathbf{z}}_t - D_t\mathbf{e}$, the transition is now stochastic and depends on the realization of the random variables $\{A_t^J\}$. Let $k \in [0, l'_J - 1]$ and $J \in [1, N]$, and consider node $l_J - k$, which stands between the J-th and $J+1$-th physical FCs. Letting:

$$\breve{z}_t^{l_J - k} := \begin{cases} \hat{z}_t^{l_J}, & A_t^J = 0 \\ \hat{z}_t^{l_J - 1}, & A_t^J = 1 \\ \vdots & \\ \hat{z}_t^{l_J - k + 1}, & A_t^J = k-1 \\ \hat{z}_t^{l_J - k}, & A_t^J \geq k \end{cases}, \tag{8}$$

we now have:

$$\mathbf{z}_{t+1} = \breve{\mathbf{z}}_t - D_t\mathbf{e}.$$

Figure 10 illustrates the state transition on an example where the maximum lead time is 5 and the realization of the random variable A_t is 3.

The optimization problem now takes a form similar to Equation (1), except with the transition function replaced by the one just derived, and the expectation extending to the random variables $\{A_t^J\}$,

$$C_0(\mathbf{z}_0) = \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t H_t(\mathbf{z}_t)\right]$$
$$\text{s.t. } \mathbf{z}_{t+1} = \breve{\mathbf{z}}_t - D_t\mathbf{e}.$$

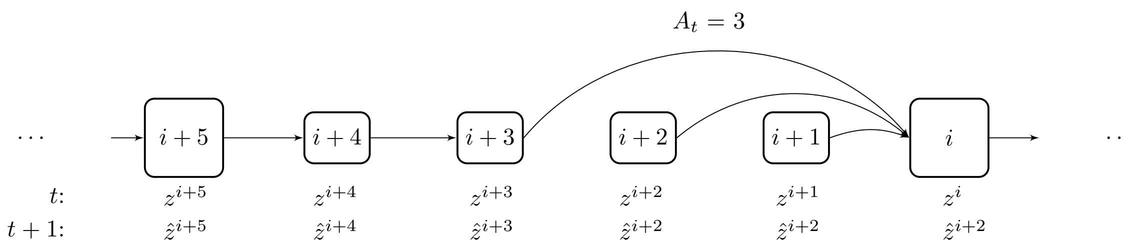

Figure 10: Illustration of the stochastic lead time transition on an example with maximum lead time of 5 and a realization of A_t equal to 3.

The dynamic programming formulation of the problem similarly reads:

$$C_t(\mathbf{z}_t) = H_t(\mathbf{z}_t) + \gamma \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \mathbb{E}\left[C_{t+1}\left(\check{\mathbf{z}}_t - D_t\mathbf{e}\right)\right]. \tag{9}$$

It is not hard to see that all the structural properties derived in Section 2.2.2 hold in the stochastic lead time case. We present the equivalent of Theorem 2.2, which in turn allows us to construct a solution method for the problem. For notational purposes, we let $q^J_k := \mathbb{P}[A^J \geq k]$.

Theorem 3.1

Let $t \in [0, T-1]$, $\tilde{c}^i_t(z) := \mathbb{E}\left[c^i_{t+1}\left(z - D_t\right)\right]$, and $\mathbf{s}_t := \left\{s^i_t : \ i \in \mathbb{T}_t\right\}$ where:

$$s^i_t := \arg\min_z \tilde{c}^i_t(z).$$

Then:

$$\begin{aligned} c^i_t(z) = H^i_t(z) &+ \gamma \begin{cases} \tilde{c}^{i-1}_t(z \wedge s^{i-1}_t), & \text{if } (i-1) \in \mathbb{T}_t \\ q^J_{l_J - i+1}\tilde{c}^{i-1}_t(z), & \text{else if } i \in \mathcal{P} \cup \mathcal{D} \end{cases} + \gamma \begin{cases} \tilde{c}^i_t(z \vee s^i_t) - \tilde{c}^i_t(s^i_t), & i \in \mathbb{T}_t \\ \tilde{c}^i_t(z), & i \in \mathcal{D} \setminus \mathbb{T}_t \end{cases} \\ &+ \gamma \begin{cases} p^{J-1}_0 \sum_{j=l_{J-2}+1}^{i-2} \tilde{c}^j_t(z \wedge s^{i-1}_t), & \text{if } (i-1) \in \mathbb{T}_t \\ p^J_{l_J-i+1} \sum_{j=l_{J-1}+1}^{i-2} \tilde{c}^j_t(z), & \text{else if } i \in \mathcal{P} \cup \mathcal{D} \end{cases} \\ &+ \gamma \begin{cases} p^J_0 \left(\sum_{j=l_{J-1}+1}^{l_J-1} \tilde{c}^j_t(z \vee s^i_t) - \tilde{c}^j_t(s^i_t)\right), & i \in \mathbb{T}_t \\ p^J_0 \sum_{j=l_{J-1}+1}^{l_J-1} \tilde{c}^i_t(z), & i \in \mathcal{D} \setminus \mathbb{T}_t \end{cases}. \end{aligned}$$

Proof 5

The proof of Theorem 3.1 is similar to the one of Theorem 2.2. Assume C_{t+1} is additively convex, so that $C_{t+1}(\mathbf{z}) = \sum_{i=1}^n c^i_t(z^i)$. According to Equation (9), we have:

$$\begin{aligned} C_t(\mathbf{z}_t) &= H_t(\mathbf{z}_t) + \gamma \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \mathbb{E}\left[C_{t+1}\left(\check{\mathbf{z}}_t - D_t\mathbf{e}\right)\right] \\ &= \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \sum_{J=1}^{N} \sum_{i=l_{J-1}+1}^{l_J} H^i(z^i) + \gamma \tilde{c}^i_t(\check{z}^i) \\ &= \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \sum_{J=1}^{N} \sum_{i=l_{J-1}+1}^{l_J} H^i(z^i) + \gamma \left[\left(\sum_{k=l_J-i+1}^{l_J-1} p_k\right) \tilde{c}^i_t(\hat{z}^i) + \sum_{k=0}^{l_J-i} p_k \tilde{c}^i_t(\hat{z}^{l_J-k})\right], \end{aligned}$$

where the last equality uses the state transition (8). Further observe that $\sum_{k=l_J-i+1}^{l_J-1} p_k = q_{l_J-i+1}$, yielding:

$$C_t(\mathbf{z}_t) = \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \sum_{J=1}^{N} \sum_{i=l_{J-1}+1}^{l_J} H^i_t(z^i) + \gamma \left[q_{l_J-i+1}\tilde{c}^i_t(\hat{z}^i) + \sum_{k=0}^{l_J-i} p_k \tilde{c}^i_t(\hat{z}^{l_J-k})\right]. \tag{10}$$

Consider now the latter part of Equation (10), which we transform as follows:

$$\begin{aligned}\sum_{i=l_{J-1}+1}^{l_J}\sum_{k=0}^{l_J-i} p_k \tilde{c}_t^i(\hat{z}^{l_J-k}) &= \sum_{i=l_{J-1}+1}^{l_J}\sum_{j=i}^{l_J} p_{l_J-j}\tilde{c}_t^i(\hat{z}^j) \\ &= \sum_{j=l_{J-1}+1}^{l_J}\sum_{i=l_{J-1}+1}^{j} p_{l_J-j}\tilde{c}_t^i(\hat{z}^j) \\ &= \sum_{i=l_{J-1}+1}^{l_J} p_{l_J-i}\left(\sum_{j=l_{J-1}+1}^{i}\tilde{c}_t^j(\hat{z}^i)\right),\end{aligned}$$

where we only switched indices in the last equality. Substituting back into (10), we obtain:

$$C_t(\mathbf{z}_t) = \min_{\hat{\mathbf{z}}_t\in\hat{Z}(\mathbf{z}_t,\mathcal{N}_t)} \sum_{J=1}^{N}\sum_{i=l_{J-1}+1}^{l_J} H_t^i(z^i) + \gamma\left[q_{l_J-i+1}\tilde{c}_t^i(\hat{z}^i) + p_{l_J-i}\sum_{j=l_{J-1}+1}^{i}\tilde{c}_t^j(\hat{z}^i)\right].$$

Using the fact that for $i \neq l_J$ we have $\hat{z}^i = z^{i+1}$, we find:

$$\begin{aligned}C_t(\mathbf{z}_t) &= \min_{\hat{\mathbf{z}}_t\in\hat{Z}(\mathbf{z}_t,\mathcal{N}_t)} \sum_{J=1}^{N}\Bigg\{\sum_{i=l_{J-1}+1}^{l_J} H_t^i(z^i) + \sum_{i=l_{J-1}+1}^{l_J-1}\gamma\left[q_{l_J-i+1}\tilde{c}_t^i(z^{i+1}) + p_{l_J-i}\sum_{j=l_{J-1}+1}^{i}\tilde{c}_t^j(z^{i+1})\right] \\ &\qquad + \gamma\left[q_1\tilde{c}_t^i(\hat{z}^{l_J}) + p_0\sum_{j=l_{J-1}+1}^{l_J}\tilde{c}_t^j(\hat{z}^{l_J})\right]\Bigg\} \\ &= \min_{\hat{\mathbf{z}}_t\in\hat{Z}(\mathbf{z}_t,\mathcal{N}_t)} \sum_{J=1}^{N}\Bigg\{\sum_{i=l_{J-1}+2}^{l_J} H_t^i(z^i) + \gamma q_{l_J-i+1}\tilde{c}_t^{i-1}(z^i) + p_{l_J-i+1}\sum_{j=l_{J-1}+1}^{i-2}\tilde{c}_t^j(z^i) \\ &\qquad + H_t^{l_{J-1}+1}(z^{l_{J-1}+1}) + \gamma\left[\tilde{c}_t^i(\hat{z}^{l_J}) + p_0\sum_{j=l_{J-1}+1}^{l_J-1}\tilde{c}_t^j(\hat{z}^{l_J})\right]\Bigg\}.\end{aligned}$$

Finally, applying Remark 2 yields the result. •

Theorem 3.1 is the stochastic lead time equivalent of Theorem 2.2. We can similarly derive alternative formulations akin to Theorems 2.3 and 2.4, although their expressions are much less appealing and we thus leave them out.

3.2 Perishability

3.2.1 Introduction

In a large number of instances, products are subject to a finite shelf-life and expire and/or can no longer be sold after some given period of time. This is especially true in industries such as the grocery or pharmaceutical industries, where spoilage can represent a significant cost and jeopardize the profitability of the business. The perishability of a product translates into a heightened overage cost since higher inventory levels translate into a higher risk of shrinkage. In the absence of perishability, there usually is an asymmetry in the underage/overage trade-off, with overage costs typically much lower than underage costs, resulting in high service levels and very flat objective functions around the optimal inventory decisions. In the case of a perishable product, the trade-off becomes much more salient and it is paramount to effectively tackle this characteristic feature, lest we incur high losses of inventory.

Even in the case of a single echelon with no lead time, an exact study of perishability inventory management would require that the state of the problem keep track of how many units on-hand have a given shelf-life left. Letting m be the shelf life of the product, the dimension of the problem would be $m-1$, leading to the usual "curse of dimensionality". As a result, most of the literature on perishable products focuses on two areas: structural properties and results, and heuristics and approximations. The former has theoretical value but is often impractical, requiring that efficient but approximate implementations be devised.

An early myopic approximation policy can be found in Nahmias [1976], further reviewed and expanded on in Nahmias [1982] and Nandakumar and Morton [1993]. More recently, Chen et al. [2014] have tackled the structural properties of the joint inventory and pricing problem for perishable products, and Li et al. [2016] considered perishability in the context of clearance sales and segregation. An approximation algorithm with worst-case performance guarantee was also presented in Chao et al. [2015].

We here develop a heuristic for our multi-echelon use-case, while aiming at maintaining the important structural properties of the formulation, and in particular the additive convexity of the cost functions that yields the desirable order-up-to policies encountered so far.

3.2.2 Heuristic

A common way of incorporating the perishability risk into the formulation is to assign to any decision the expected spoilage cost it will cause, which is a combination of the loss of the product cost c and an additional removal cost θ.

To motivate our heuristic, consider the simple case of a single-echelon problem with a null lead time. At the beginning of period t, the inventory position is given by z^0, and we raise it to $\hat{z}^0$, thus bringing in $\hat{z}^0 - z^0$ units. Letting m be the shelf life of the product, the expected spoilage cost they will incur m periods hence is given by $(c+\theta)O_t^m(\hat{z}^0, z^0)$, where $O_t^m(\hat{z}^0, z^0)$ is the expected number of the $\hat{z}^0 - z^0$ units that will perish and on which we will incur a cost. Letting $G_t^m(z) := \mathbb{E}\left[(z - D_{t,t+m})^+\right]$, it was shown in Nahmias [1976] that $O_t^m(\hat{z}, z)$ can be bounded as follows:

$$G_t^m(\hat{z} - z) \leq O_t^m(\hat{z}, z) \leq G_t^m(\hat{z}) - G_t^m(z).$$

The upper bound is usually argued to be a tighter bound and used as an approximation of O_t^m. Consequently, the cost of spoilage $S_t^0(\hat{z}, z)$ associated with the decision $\hat{z}^0$ at time t is approximated by:

$$S_t^0(\hat{z}, z) \approx \gamma^m (c+\theta)\left(G_t^m(\hat{z}) - G_t^m(z)\right). \tag{11}$$

We now seek to extend this approximation to a multi-echelon setting in which decisions are made at each installation, whether they be virtual or physical. Consider for a moment node i at time t, and its echelon position z_t^i and inventory decision $\hat{z}_t^i$, bringing in $\hat{z}_t^i - z_t^i$ at the end of the period, or equivalently at the beginning of the following period. Suppose we know what the remaining life m_{t+1}^i of these units is when they arrive at the beginning of period $t+1$ (implicitly also assuming that all arriving units have the same remaining shelf-life); the expected cost of spoilage $S_t^i(\hat{z}_t^i, z_t^i)$ is thus given by:

$$S_t^i(\hat{z}_t^i, z_t^i) = \gamma^{m_{t+1}^i + 1}(c+\theta)\left(G_{t+1}^{m_{t+1}^i}(\hat{z}_t^i - D_t) - G_{t+1}^{m_{t+1}^i}(z_t^i - D_t)\right).$$

Note that this expression differs slightly from the earlier one (11) due to the lead time of 1 period between decision and arrival of the units. Observing that $\hat{z}_t^i - D_t = z_{t+1}^i$ and that $G_{t+1}^{m_{t+1}^i}(z_t^i - D_t) = G_t^{m_{t+1}^i + 1}(z_t^i)$, we can overload the above expression as:

$$S_t^i(z_{t+1}^i, z_t^i) = \gamma^{m_{t+1}^i + 1}(c+\theta)\left(G_{t+1}^{m_{t+1}^i}(z_{t+1}^i) - G_t^{m_{t+1}^i + 1}(z_t^i)\right).$$

The next step is to evaluate how these additional shrinkage costs affect the formulation of the problem. The following expression corresponds to the additional terms to be added to the objective function of the multi-period problem presented in Equation (1):

$$\sum_{t=0}^{T-1}\sum_{i=1}^{n}\gamma^t S_t^i(z_{t+1}^i, z_t^i) = (c+\theta)\sum_{t=0}^{T-1}\sum_{i=1}^{n}\gamma^{t+m_{t+1}^i+1}\left(G_{t+1}^{m_{t+1}^i}(z_{t+1}^i) - G_t^{m_{t+1}^i+1}(z_t^i)\right)$$
$$= (c+\theta)\sum_{t=0}^{T-1}\sum_{i=1}^{n}\gamma^{t+m_t^i}\left(G_t^{m_t^i}(z_t^i) - \gamma^{m_{t+1}^i-m_t^i+1}G_t^{m_{t+1}^i+1}(z_t^i)\right)$$
$$+\sum_{i=1}^{n}\gamma^{T+m_T^i}G_T^{m_T^i}(z_T^i) - \gamma^{m_1^i+1}G_0^{m_1^i+1}(z_0^i).$$

By subtracting $\sum_{i=1}^n \gamma^{T+m_T^i}G_T^{m_T^i}(z_T^i)$ from the terminal conditions, and observing that the term $\gamma^{m_1^i+1}G_0^{m_1^i+1}(z_0^i)$ is independent of the decision variables and can thus be discarded, we can reach a formulation of the problem characterized by the following optimality equation.

$$C_t(\mathbf{z}_t) = H_t(\mathbf{z}_t) + \tilde{P}_t(\mathbf{z}_t) + \gamma \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \mathbb{E}\left[C_{t+1}\left(\hat{\mathbf{z}}_t - D_t\mathbf{e}\right)\right], \tag{12}$$

where:

$$\tilde{P}_t(\mathbf{z}_t) = \sum_{i=1}^{n}\tilde{P}_t^i(z_t^i), \qquad \tilde{P}_t^i(z) = \gamma^{m_t^i}(c+\theta)\left(G_t^{m_t^i}(z) - \gamma^{m_{t+1}^i-m_t^i+1}G_t^{m_{t+1}^i+1}(z)\right).$$

This formulation is however not practical, for $\tilde{P}_t^i$ need not be convex. It has been shown that the term is quasi-convex in some instances (see Nahmias [1976]), but while Karush's lemma in Remark 2 carries over to quasi-convex functions, the fact that the sum of quasi-convex functions is not necessarily quasi-convex prevents our results from holding. We must consequently address two points: find convex approximations of the $\tilde{P}_t^i$, and expressions for the remaining shelf lives m_t^i.

Shelf Lives: In deriving the formulation above, we assumed the remaining shelf life of units arriving at any node i in period t to be known. In a single echelon setting, this is quite straightforward since units move from an upper (virtual) installation to the next in every period, so that we immediately have $m_t^i = m - (n-i+1)$, where $n-i+1$ is the time it takes units to reach installation i from the vendor. In a multi-echelon setting, units may sit at a physical installation for some time before being transferred to a lower installation. This is especially true in the presence of transfer schedules that do not allow for units to be transferred in every single period.

We define l_t^i as the shortest lead time units may have taken to reach installation i in period t, illustrated in Figure 11. Note that some paths are "infeasible" in the sense that there exists no schedule that would allow for units to arrive at installation i in period t. In that case we define l_t^i to be equal to the most recent such feasible lead time. We may then define m_t^i as:

$$m_t^i = m - l_t^i.$$

When the review period is of one period at every installation, this leads to m_t^i being independent of t and equal to $m_t^i = m^i = m - (n-i+1)$. Additionally, it is easy to see from our definition that we must have the following inequality:

$$m_t^i \leq m_{t+1}^i. \tag{13}$$

Convex Approximation of $\tilde{P}_t^i$: The optimality equation (12) contains terms that approximate the spoilage costs in the form of $\tilde{P}_t$. While $\tilde{P}_t$ is additively separable, its univariate components $\tilde{P}_t^i$ need not be convex; a deficiency that must be resolved in order to preserve our structural results. Recall that we have:

$$\tilde{P}_t^i(z) = \gamma^{m_t^i}(c+\theta)\left(G_t^{m_t^i}(z) - \gamma^{m_{t+1}^i-m_t^i+1}G_t^{m_{t+1}^i+1}(z)\right).$$

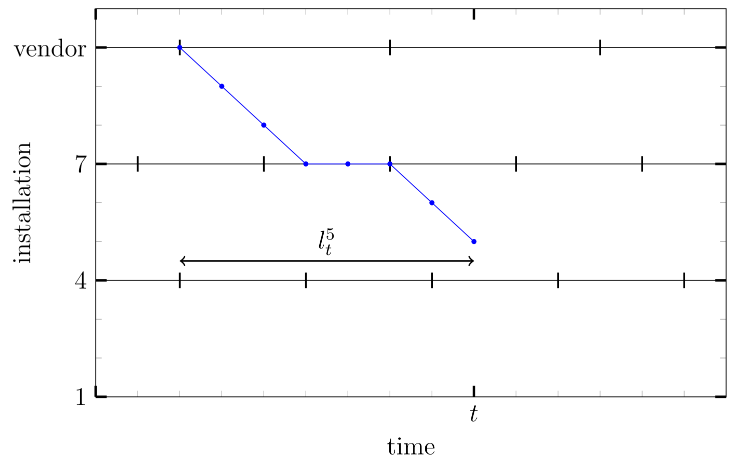

Figure 11: Illustration of the shortest path taken by units to reach installation 5 in a 3-echelon serial network with lead times of 3 between all physical installations. Ticks at physical locations denote times when transfers can be triggered.

Note that because of (13), we also have $m_{t+1}^i + 1 > m_t^i$. As a result, the following bounds on $G_t^{m_{t+1}^i+1}(z)$ hold and are illustrated in Figure 12:

$$0 \leq G_t^{m_{t+1}^i+1}(z) \leq G_t^{m_t^i}(z).$$

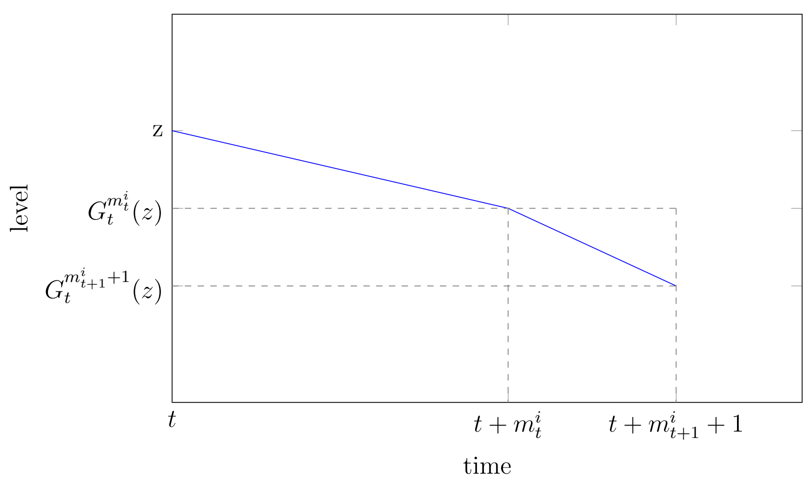

Figure 12: Illustration of the bounds on $G_t^{m_{t+1}^i+1}$.

Using these bounds, we can approximate $G_t^{m_{t+1}^i+1}$ as a convex combination of 0 and $G_t^{m_t^i}$, i.e. as $G_t^{m_{t+1}^i+1} = \alpha_t^i G_t^{m_t^i}$ with $\alpha_t^i \in [0, 1]$. We suggest $\alpha_t^i = \frac{m_t^i}{m_{t+1}^i+1}$ as a reasonable value. Using this approximation, we can use P_t^i in lieu of $\tilde{P}_t^i$, where P_t^i is defined as:

$$P_t^i(z) = \gamma^{m_t^i}(c + \theta)\left(1 - \alpha_t^i \gamma^{m_{t+1}^i - m_t^i + 1}\right) G_t^{m_t^i}(z). \tag{14}$$

We may now express the final formulation of the problem in the presence of perishability.

Formulation of the problem with perishability

$$C_t(\mathbf{z}_t) = H_t(\mathbf{z}_t) + P_t(\mathbf{z}_t) + \gamma \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \mathbb{E}\left[C_{t+1}\left(\hat{\mathbf{z}}_t - D_t \mathbf{e}\right)\right], \tag{15}$$

where:

$$P_t(\mathbf{z}_t) = \sum_{i=1}^{n} P_t^i(z_t^i), \qquad P_t^i(z) = \gamma^{m_t^i}(c+\theta)\left(1 - \alpha_t^i \gamma^{m_{t+1}^i - m_t^i + 1}\right) G_t^{m_t^i}(z).$$

In this formulation, P_t is additively convex, allowing for all our structural results to continue to hold.

3.3 Reverse Logistics

Inventory management problems are most often concerned with optimizing replenishment decisions. Additional possible revenue streams such as returns (to the vendor) and liquidations, or demand shaping options such as pricing allow for a richer framework. They allow for a finer control of revenue, but more importantly of capacity. In the single echelon case, Maggiar and Sadighian [2017] study a joint pricing and removal problem. In the multi-echelon framework, Angelus [2011] considers secondary market sales and Angelus and Özer [2018] incorporate reverse logistics in their analysis.

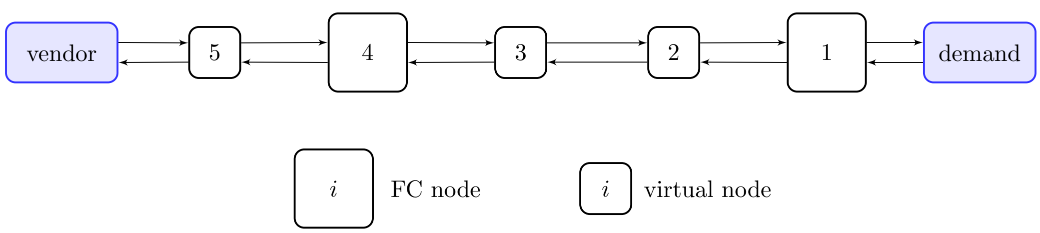

Figure 13: Illustration of the reverse logistics setting.

Reverse logistics refer to the bi-directional flow of products in the network. Whereas we have so far been concerned with the forward flow of units, from upstream nodes to more downstream ones, reverse logistics reflect the ability to send units back up in the network (see Figure 13). This is especially interesting in a capacity management setting where seasonal products might no longer be needed at spoke locations and can be sent back up to a less constrained facility, or even returned to the vendor or liquidated if those options are available.

The incorporation of these additional decisions require that we somewhat modify our assumptions to accommodate them. We assume in this section the following.

Assumption 2

1. All nodes are decision nodes,
2. Nodes that are not directly under or over an FC can make forward and backward decisions in every single period,
3. Nodes directly under an FC, which were until now the decision nodes, are still subject to schedule constraints in terms of forward decisions,
4. Nodes directly above an FC are subject to schedule constraints in terms of backwrad decisions,

5. We may return units to the vendor, or liquidate them. However, this is only feasible at the most upstream node.

The first and second assumptions do not mean that we assume each node to be a physical facility. They instead mean that we may effectively affect the lead times between facilities and may freely move units up and down or have them sit still between physical locations without reaching them. This assumption is actually somewhat realistic: when FCs are full, trucks are indeed made to wait outside of facilities. This is also a motivation between the explicit use of virtual nodes and the accounting of their holding and transfer costs. The third and fourth assumptions simply state that the procurement and removal of inventory from FCs might be subject to schedule constraints, and just like we may not be able to purchase or transfer units in every single period, we may not be able to remove units from an FC and ship them upstream in every single period.

To formalize the problem with the inclusion of reverse logistics, we define some following additional notation. Recall that $\mathbb{T}_t$ refers to the set of nodes at which a procurement decision can effectively be made in period t. In the reverse logistics setting, this will include all nodes not directly under an FC at all times, and the latter in the periods in which transfers and purchases can occur. We similarly defined $\check{\mathbb{T}}_t$ to be the set of nodes at which a backward decision can be made in period t. It will here always include all nodes not directly above an FC, and the latter in periods in which units may be removed from an FC. We then extend the set $\mathcal{N}_t$ to include this new decision set: $\mathcal{N}_t = \{\mathcal{F}, \mathcal{D}, \mathcal{P}, \mathbb{T}_t, \check{\mathbb{T}}_t\}$, where we could also dispense with the pipeline node set since it is now empty. Note that because of the assumption that all nodes are decision nodes, the constraint set $\hat{Z}(z, \mathcal{N})$ simplifies to:

$$\hat{Z}(\mathbf{z}, \mathcal{N}) = \left\{ \hat{z} \in Z : \begin{cases} z^i \le \hat{z}^i \le z^{i+1}, & i \in \mathbb{T} \\ \hat{z}^i = z^i, & i \notin \mathbb{T} \end{cases} \right\}.$$

We similarly define a constraint set for the backward decisions:

$$\check{Z}(\mathbf{z}, \mathcal{N}) := \left\{ \check{z} \in Z : \begin{cases} z^{i-1} \le \check{z}^i \le z^i, & i \in \mathbb{T} \\ \check{z}^i = z^i, & i \notin \mathbb{T} \end{cases} \right\}.$$

The addition of backward and return/liquidation decisions requires that we specify the new sequence of events, which is similar to the one presented in Section 2.2.1 with the introduction of removal decisions:

1. The vector of echelon inventory positions $\mathbf{z}$ is observed,

2. Removal decisions are made, leading to a new echelon inventory position $\check{\mathbf{z}} \in \check{Z}(\mathbf{z}, \mathcal{N}_t)$, and initiating the transfer of $z^i - \check{z}^i$ units from node i to node $i+1$ on which we incur a transfer cost $\check{w}_i$ per unit of volume, as well as a return/liquidation cost of s such that $s \le \gamma c$ per unit on the $z^n - \check{z}^n$ units return to the vendor, or liquidated,

3. A target echelon inventory position vector $\hat{\mathbf{z}} \in Z(\check{\mathbf{z}}, \mathcal{N}_t)$ is set, triggering the transfer of $\hat{z}^i - \check{z}^i$ units from node $i+1$ to node i to be effective in the following period, on which we incur a transfer cost of γw_i per unit of volume, as well as a purchasing cost of γc per unit on the $\hat{z}^n - \check{z}^n$ units ordered from the vendor,

4. A realization of the random demand D_t is observed, and owing to the backlogging assumption, we sell D_t units at a price p,

5. We incur a holding cost of h'_i on each unit of volume at left over at node $i \in [1, n]$, and possibly a backlogging penalty b on the $(D_t - z^1)^+$ backlogged units during the period,

6. The ordered and transferred units arrive and the vector of echelon inventory positions moves to $\hat{\mathbf{z}} - D_t\mathbf{e}$.

The formulation of the problem then becomes the following:

$$\begin{aligned} C_0(\mathbf{z}_0) = \quad & \min_{\substack{\check{\mathbf{z}}_t \in \check{Z}(\mathbf{z}_t, \mathcal{N}_t) \\ \hat{\mathbf{z}}_t \in \hat{Z}(\check{\mathbf{z}}_t, \mathcal{N}_t)}} \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t \left(\gamma c(\hat{z}_t^n - \check{z}_t^n) - s(z_t^n - \check{z}_t^n) \right. \right. \\ & \left. \left. + H_t(\mathbf{z}_t) + \sum_{i=1}^{n} \gamma w_i v \left(\hat{z}_t^i - \check{z}_t^i \right) + \sum_{i=1}^{n} \gamma \breve{w}_i v \left(z_t^i - \check{z}_t^i \right) \right) \right] \\ & \text{s.t. } \mathbf{z}_{t+1} = \hat{\mathbf{z}}_t - D_t \mathbf{e}. \end{aligned}$$

The first line of the objective function contains the purchasing and return/liquidation costs, while the second line contains the holding, backlogging and transfer costs.

Operating transformations almost identical to the ones performed in Section 2.2.1, we may express the problem through it Dynamic Programming formulation as:

$$\begin{aligned} C_t(\mathbf{z}_t) = & H_t(\mathbf{z}_t) + (c - s) z_t^n + \sum_{i=1}^{n} (\breve{w}_i + w_i) v z_t^i \\ & + \min_{\check{\mathbf{z}}_t \in \check{Z}(\mathbf{z}_t, \mathcal{N}_t)} \left\{ -(\gamma c - s) \check{z}_t^n - \sum_{i=1}^{n} (\gamma w_i + \breve{w}_i) v \check{z}_t^i + \gamma \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\check{\mathbf{z}}_t, \mathcal{N}_t)} \mathbb{E}\left[C_{t+1}(\hat{\mathbf{z}}_t - D_t \mathbf{e}) \right] \right\}. \end{aligned} \tag{16}$$

Using this formulation we may prove that the cost function is additively convex, and that the optimal policy follows an *interval-stock* policy defined below and illustrated in Figure 14.

Definition 3.1
An interval-stock policy is a policy characterized by two levels, an order-up-to (base-stock) level $\underline{s}$ and a remove-down-to level $\bar{s}$ such that $\underline{s} \leq \bar{s}$, and such that the inventory position z is raised to $\underline{s}$ if $z \leq \underline{s}$, it is lowered to $\bar{s}$ if $z \geq \bar{s}$, and it is left unmodified if $\underline{s} \leq z \leq \bar{s}$. In other words, the updated inventory position z_+ is given by:

$$z_+(z) = \max\left(\underline{s}, \min(\bar{s}, z)\right).$$

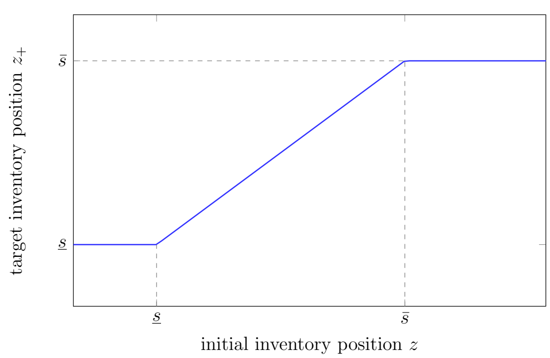

Figure 14: Illustration of the interval-stock policy.

Theorem 3.2
Consider the problem defined in (16). Then:

1. The optimal policy is given by an interval-stock policy at all echelons and time periods,
2. The cost functions are additive convex.

Proof 6
The proof is by induction. At time T, the cost C_T is given by $C_T(\mathbf{z}) = \sum_{i=1}^n H^i(z^i)$, which is clearly additive convex.
Assume then that C_{t+1} is additive convex. According to (16), we have:

$$C_t(\mathbf{z}) = H_t(\mathbf{z}) + (c-s)z^n + \sum_{i=1}^n (\check{w}_i + w_i)vz^i + \min_{\check{\mathbf{z}} \in \check{Z}(\mathbf{z},\mathcal{N}_t)} \left\{ -(\gamma c - s)\check{z}^n - \sum_{i=1}^n (\gamma w_i + \check{w}_i)v\check{z}^i + \gamma \min_{\hat{\mathbf{z}} \in \hat{Z}(\check{\mathbf{z}},\mathcal{N}_t)} \mathbb{E}\left[C_{t+1}(\hat{\mathbf{z}} - D_t\mathbf{e})\right] \right\}.$$

The problem consists of two nested optimization subproblems, where the first decision is on the number of units to ship backward and remove, followed by a decision on the number of units to transfer and procure. Using the assumption that C_{t+1} is additive convex, letting as before $\tilde{c}_t^i(z) := \mathbb{E}\left[c_{t+1}^i(z - D_t)\right]$ and $\underline{s}_t^i := \arg\min_z \tilde{c}_t^i(z)$, and using Remark 2, we may rewrite this problem as:

$$C_t(\mathbf{z}) = H_t(\mathbf{z}) + (1-\gamma)cz^n + \sum_{i=1}^n (1-\gamma)w_i v z^i + \min_{\check{\mathbf{z}} \in \check{Z}(\mathbf{z},\mathcal{N}_t)} \left\{ (\gamma c - s)(z^n - \check{z}^n) + \sum_{i=1}^n (\gamma w_i + \check{w}_i)v(z^i - \check{z}^i) + \gamma \sum_{i=1}^n \tilde{c}_t^i(\check{z}^i \vee \underline{s}_t^i) + \tilde{c}_t^i(\check{z}^{i+1} \wedge \underline{s}_t^i) - \tilde{c}_t^i(\underline{s}_t^i) \right\},$$

where we abusively let $\underline{s}_t^i = -\infty$ for $i \notin \mathbb{T}_t$.
The problem can then be broken up into univariate optimization problems, which for $i \in \check{\mathbb{T}}_t$ and $i < n$ are of the form:

$$\min_{z^{i-1} \leq \check{z}^i \leq z^i} \check{c}_t^i(\check{z}^i),$$

where we define $\check{c}_t^i$ as:

$$\check{c}_t^i(\check{z}^i) := (\gamma w_i + \check{w}_i)v(z^i - \check{z}^i) + \gamma \left(\tilde{c}_t^i(\check{z}^i \vee \underline{s}_t^i) + \tilde{c}_t^{i-1}(\check{z}^i \wedge \underline{s}_t^{i-1}) - \tilde{c}_t^i(\underline{s}_t^i) \right). \tag{17}$$

Similar functions can be defined for other values of i.
The objective function (17) being optimized is clearly additive convex, and applying Remark 2 once again yields that the resulting function is additive convex as a function of z^i and z^{i-1}, which in turn shows that $C_t(\mathbf{z})$ is additive convex.
Let then $\bar{s}_t^i := \arg\min_z \check{c}_t^i(z)$. It remains to show that $\bar{s}_t^i \geq \underline{s}_t^i$. Suppose that is not the case, i.e. $\bar{s}_t^i < \underline{s}_t^i$, and let y such that $\bar{s}_t^i < y < \underline{s}_t^i$. We then have:

$$\begin{aligned}\check{c}_t^i(y) = (\gamma w_i + \check{w}_i)v(z^i - y) + \gamma\tilde{c}_t^{i-1}(y \wedge \underline{s}_t^{i-1}) &< (\gamma w_i + \check{w}_i)v(z^i - \bar{s}_t^i) + \gamma\tilde{c}_t^{i-1}(y \wedge \underline{s}_t^{i-1}) \\ &\leq (\gamma w_i + \check{w}_i)v(z^i - \bar{s}_t^i) + \gamma\tilde{c}_t^{i-1}(\bar{s}_t^i \wedge \underline{s}_t^{i-1}) \\ &= \check{c}_t^i(\bar{s}_t^i).\end{aligned}$$

The first inequality follows from the fact that $(\gamma w_i + \check{w}_i) \geq 0$. (A similar argument would apply in the case when $i = n$ since $\gamma c > s$). The second inequality is due to the fact that $\tilde{c}_t^{i-1}$ is non-increasing. This result contradicts the optimality of $\bar{s}_t^i$, whence $\bar{s}_t^i \geq \underline{s}_t^i$. •

3.4 Random Vendor Supply

Incorporating an awareness of the procurement risk, and potentially its fluctuations, is also an important factor in hedging against this risk and anticipating especially heightened periods of supply risk. This risk can be modeled in a number of ways and a rich body of literature is concerned with it, often modeling it either through a random vendor capacity, or a stochastic yield. We opt for the latter option in this work, as it lends itself more naturally to our formulation. Such a model can be found for example in Henig and Gerchak [1990] or Bollapragada and Morton [1999]. More specifically, we consider a supply risk modeled by a Bernoulli distribution that characterizes whether an order to a vendor is fully fulfilled, or not at all. Such a model might appear overly simplistic, but it empirically quite realistic, and from a modeling perspective, it often leads to the preservation of structural results, namely the order-up-to policy, as has been remarked in Özekici and Parlar [1999].

Let then A_t be a Bernoulli distributed random variable in period t with mean α_t. In period t, the order placed to the vendor is given by $\hat{z}_t^n - z_t^n$, and the yield on that order is thus given by $A_t(\hat{z}_t^n - z_t^n)$. As a result, the transition function for the uppermost node, i.e. the vendor facing node n, is modified to read $z_{t+1}^n = z_t^n + A_t(\hat{z}_t^n - z_t^n) - D_t$. Additionally, all the expectations previously taken with respect to only the demands D_t in the base model, must now also be taken with respect to the yield rates A_t, which we assume to be independent from the demand. Nonetheless, the Bernoulli structure of this random variable allows us, to obtain a formulation of the problem equivalent to the one in Equation (1):

$$C_0(\mathbf{z}_0) = \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \mathbb{E}_{D,A}\left[\sum_{t=0}^{T-1} \gamma^t H_t(\mathbf{z}_t)\right]$$
$$\text{s.t. } \mathbf{z}_{t+1} = \tilde{\mathbf{z}}_t - D_t \mathbf{e},$$

where $\tilde{\mathbf{z}}_t$ is defined as:

$$\tilde{\mathbf{z}}_t = \begin{cases} \hat{z}_t^i, & i < n \\ A_t \hat{z}_t^n + (1 - A_t) z_t^n \end{cases}.$$

By replacing $\hat{\mathbf{z}}$ with $\tilde{\mathbf{z}}$ in the transformation in Appendix A.1, we can see that the exact same reformulation as Equation (1) holds, with the only difference that expectations are taken over both the demand and the random yield. As a result, the optimality equation is slightly amended to read:

$$C_t(\mathbf{z}_t) = H_t(\mathbf{z}_t) + \gamma \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \mathbb{E}_{D_t, A_t}\left[C_{t+1}\left(\tilde{\mathbf{z}}_t - D_t \mathbf{e}\right)\right]. \tag{18}$$

The following Proposition shows that the results obtained for the base models carry over to the model with random supply.

Proposition 3.1
Let the vendor supply be random in period t and let it be characterized by a Bernoulli random variable A_t with mean α_t. Then the results in Theorem 2.1 hold in the random vendor supply scenario:

1. The optimal policy is given by an echelon base-stock policy $\mathbf{s}_t$ for $t \in [1, T-1]$.
2. The cost functions are additive convex, i.e. $\forall t \in [1, T]$, there exist univariate convex functions $c_t^i, \ i = 1, \ldots, n$ such that:

$$C_t(\mathbf{z}) = \sum_{i=1}^{n} c_t^i(z^i).$$

3. For $i \in \mathbb{T}_t$, the components of the base stock policy s_t^i are given by:

$$s_t^i = \arg\min_z \mathbb{E}\left[c_{t+1}^i\left(z - D_t\right)\right].$$

Proof 7
The proof is by recursion.

The theorem is clearly valid at time T, where no decision takes place and we have:

$$C_T(\mathbf{z}) = 0,$$

leading to the following expressions for c_T^i:

$$c_T^i(z^i) = 0 \quad \forall i.$$

Suppose now that the statements are valid for $t+1$, $t \in [1, T-1]$, so that $C_{t+1}(\mathbf{z}) = \sum_{i=1}^{n} c_{t+1}^i(z^i)$ for some convex functions c_{t+1}^i. The optimality equation (18) in the presence of random supply reads:

$$C_t(\mathbf{z}) = \sum_{i=1}^{n} H_t^i(z^i) + \gamma \min_{\hat{\mathbf{z}} \in \hat{Z}(\mathbf{z}, \mathcal{N}_t)} \mathbb{E}_{D_t, A_t}\left[\sum_{i=1}^{n-1} c_{t+1}^i\left(\hat{z}^i - D_t\right) + c_t^n\left(z^n + A_t(\hat{z}^n - z^n) - D_t\right)\right]. \tag{19}$$

Consider the optimization part of the expression, which can be expressed as:

$$\begin{aligned}
&\min_{\hat{\mathbf{z}}\in\hat{Z}(\mathbf{z},\mathcal{N}_t)} \mathbb{E}_{D_t,A_t}\left[\sum_{i=1}^{n-1} c_{t+1}^i\left(\hat{z}^i - D_t\right) + c_{t+1}^n\left(z^n + A_t(\hat{z}^n - z^n) - D_t\right)\right] \\
&= \min_{\hat{\mathbf{z}}\in\hat{Z}(\mathbf{z},\mathcal{N}_t)} \mathbb{E}_{D_t}\left[\sum_{i=1}^{n-1} c_{t+1}^i\left(\hat{z}^i - D_t\right) + \alpha_t c_{t+1}^n(\hat{z}^n - D_t) + (1-\alpha_t)c_{t+1}^n(z^n - D_t)\right] \\
&= \sum_{i\in\mathbb{T}_t, i\neq n} \min_{z^i\leq\hat{z}^i\leq z^{i+1}} \mathbb{E}_{D_t}\left[c_{t+1}^i\left(\hat{z}^i - D_t\right)\right] + \sum_{i\in\mathcal{P}\cup\mathcal{F}\setminus\mathcal{D}} \mathbb{E}\left[c_{t+1}^i\left(z^{i+1} - D_t\right)\right] + \sum_{i\in\mathcal{D}\setminus\mathbb{T}_t} \mathbb{E}\left[c_{t+1}^i\left(z^i - D_t\right)\right] \\
&\quad + \mathbf{1}_{\{n\in\mathbb{T}_t\}}\left[\alpha_t \min_{z^{n-1}\leq\hat{z}^n} c_{t+1}^n(\hat{z}^n - D_t) + (1-\alpha_t)c_{t+1}^n(z^n - D_t)\right] \\
&= \sum_{i\in\mathbb{T}_t, i\neq n} \min_{z^i\leq\hat{z}^i\leq z^{i+1}} \tilde{c}_t^i\left(\hat{z}^i\right) + \mathbf{1}_{\{n\in\mathbb{T}_t\}}\left(\alpha_t \min_{z^{n-1}\leq\hat{z}^n} \tilde{c}_t^n(\hat{z}^n) + (1-\alpha_t)\tilde{c}_t^n(z^n)\right) \\
&\quad + \sum_{i\in\mathcal{P}\cup\mathcal{F}\setminus\mathcal{D}} \tilde{c}_t^i\left(z^{i+1}\right) + \sum_{i\in\mathcal{D}\setminus\mathbb{T}_t} \tilde{c}_t^i\left(z^i\right), \qquad (20)
\end{aligned}$$

where $\tilde{c}_t^i(z) := \mathbb{E}\left[c_{t+1}^i\left(z - D_t\right)\right]$.

We observe that the optimization problem has been decomposed into the sum of independent convex bounded univariate optimization problems that clearly lead to a echelon base-stock policy $\mathbf{s}_t = \left\{s_t^j : j\in\mathbb{T}_t\right\}$, where s_t^i is the minimizer of $\tilde{c}_t^i(z)$ for $i\in\mathbb{T}_t$. Using Remark 2.1, we can write:

$$\min_{z^i\leq\hat{z}^i\leq z^{i+1}} \tilde{c}_t^i\left(\hat{z}^i\right) = \tilde{c}_t^i(z^i\vee s_t^i) + \tilde{c}_t^i(z^{i+1}\wedge s_t^i) - \tilde{c}_t^i(s_t^i).$$

Plugging this equality into (20) and then back into (19), we find that $C_t(\mathbf{z})$ is additive convex as the sum of univariate convex functions. •

4 Network

4.1 Aggregation and Risk Pooling

Section 2 covered the serial line model we use as a building block for the network model. Once the network has been broken into serial lines for all products and spoke locations, and the serial line TIPs have been computed, we must aggregate the results to yield the network-level echelon and installation target positions. This task must also take into account one of the major benefits of distribution networks, risk pooling. Figure 15 illustrates the state prior to the aggregation step. The TIPs at the spoke locations remain unchanged, but the serial line TIPs at FC_0, which is here a demand-facing hub, need to be somehow pieced together to yield the network value.

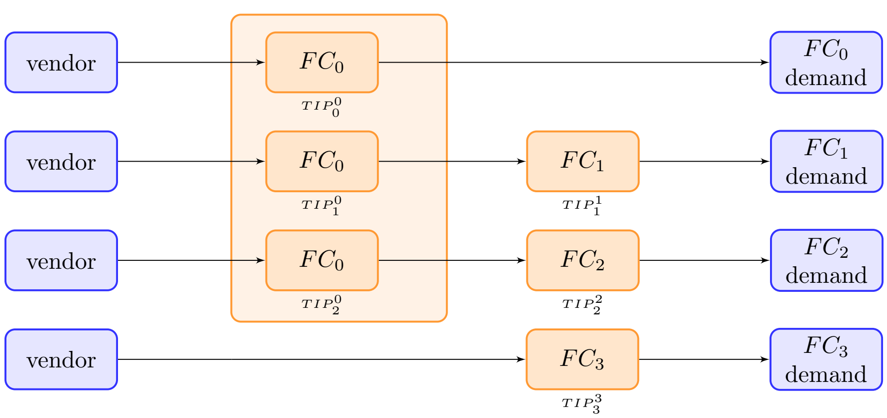

Figure 15: Illustration of the aggregation step.

In their presentation of the Decomposition-Aggregation method, Rong et al. [2017] suggest a backorder matching procedure to compute the network-level installation TIPs, which provides with a risk pooling benefit compared to a straightforward addition of the serial line installation target inventory positions. Their procedure defines installation base-stock levels as the difference between echelon base-stock levels, which does not quite follow the optimal policy, which is set at the echelon level and must account for the inventory in the serial line. Setting a base-stock level at the installation level independently of the inventory picture is suboptimal. Consider for example a stationary two-echelon network with optimal echelon levels s^1 and s^2. If the actual inventory level at the first echelon is z^1, the optimal policy is to bring in $(s^2 - z^1)^+$ units, which can be greater or smaller than the suggested installation quantity $s^2 - s^1$. We also need to account for the possibly different schedules at the different serial lines.

Consider a physical facility f at which different serial lines corresponding to the same product pass, such as FC_0 in Figure 15. We index by $d \in D$ the different serial lines, corresponding to different streams of demand at their respective spoke locations, and let c_d be the next downstream physical installation after f in the serial line corresponding to demand stream d. More precisely, we let f_d be the decision node corresponding to facility f in the serial line indexed by d, and c_d be the decision node corresponding to the next downstream installation in that serial line. For each serial line d, the quantity $u_t^{f_d}$ to be targeted at f in serial line d and time period t is given by:

$$u_t^{f_d} := \left(s_t^{f_d} - \hat{z}_t^{c_d}\right)^+.$$

Note that because of the ordering schedules, $s_t^{f_d}$ might not be properly defined if $f_d \notin \mathbb{T}_t^d$. To cover these cases, we let $p(t) := \max\left\{t' \leq t : \; f_d \in \mathbb{T}_{t'}^d\right\}$ be the most recent ordering period before (and including) t. We then define $s_t^{f_d}$ for all t as follows:

$$s_t^{f_d} := s_{p(t)}^{f_d} - D_{p(t),t}^d.$$

A first additive aggregation logic that does not take risk-pooling into account consists in summing up the various serial line level quantities at f.

Additive Aggregation:

$$u_t^f = \sum_{d \in D} u_t^{f_d}.$$

A backorder matching procedure similar to the one presented in Rong et al. [2017] allows us to incorporate a notion of risk pooling. Taken independently, each serial line (demand stream) generates an expected backlog quantity at installation f during its lead time $L^{d'}$, given by:

$$Q_t^d(u_t^{f_d}) := \mathbb{E}\left[\left(D_{t,t+L^{d'}}^d - (\hat{z}_t^{c_d} + u_t^{f_d} - s_t^{c_d})\right)^+\right].$$

In the event that the updated echelon position at the lower echelon $\hat{z}_t^{c_d}$ is lower than the target echelon position at the current echelon $s_t^{f_d}$, we find $\hat{z}_t^{c_d} + u_t^{f_d} - s_t^{c_d} = s_t^{f_d} - s_t^{c_d}$, which would lead to the same expressions as in Rong et al. [2017]. However, if lower echelons are overstocked to the point that $\hat{z}_t^{c_d}$ is greater than $s_t^{f_d}$, the expected backlog is going to be lower, as reflected by the new installation position $\hat{z}_t^{c_d} + u_t^{f_d} - s_t^{c_d} = \hat{z}_t^{c_d} - s_t^{c_d}$. The idea is then to pull these backlogs together and find the installation quantity to be targeted to yield the same expected total backlog. Let Q_t be defined as:

$$Q_t(u) := \mathbb{E}\left[\left(\sum_d \left(D_{t,t+L^{d'}}^d - (\hat{z}_t^{c_d} - s_t^{c_d})\right) - u\right)^+\right].$$

Backorder Matching Aggregation:

$$u_t^f = Q_t^{-1}\left(\sum_d Q_t^d(u_t^{f_d})\right),$$

where $Q_t^{-1}(y) = \min\{u : Q_t(u) \leq y\}$.

This formulation differs slightly from the one in Rong et al. [2017] by taking into account possible overstock at lower installations.

4.2 Capacity Constraints

One of the motivations behind multi-echelon systems is the ability to handle increasingly constrained facilities as we move down a serial line. The network allows for a reduction of lead times feeding into any given facility, thus reducing the required amount to be held to satisfy a given service level, a crucial feature considering that most spoke FCs will suffer from capacity limitations. Additionally, they allow for the pooling of risk across FCs, a benefit we explore in Section 4.1.

We focus in this section on capacity constraints, which can be of two types:

Volumetric: This refers to the amount of volume an FC can physically hold,

Flow: This refers to the quantity (in units or volume) that an FC can process in any given period, which can be either in terms of inbound or outbound flow (or both).

4.2.1 Volumetric Constraints

FCs can only store a finite amount of volume. Given these constraints, it is necessary to adjust our buying policies to ensure that we do not violate capacities at the various installations, and prevent our decisions from leading to too large inventory volumes. At a single product level, such guarantees are virtually impossible given the stochasticity of demand, and could only reasonably be expressed in terms of chance constraints. On the other hand, because we consider the distribution system of a large retailer and the planning of thousands of products, we can leverage some form of the law of large numbers to consider the constraints in expectation instead.

Suppose then that we are managing inventory for a set $\mathcal{A}$ of n_a products. For simplicity of notation, assume that all products flow through the same serial line, although the results port immediately to the more general distribution network. Using a to index products in $\mathcal{A}$, the objective is to minimize the total expected discounted cost aggregated across products, subject to a vector of capacity constraints $\mathbf{K}$ such that the expected inventory volume in any given period t and any node i must be less than or equal to K_i. While these capacity constraints technically only apply at the physical node levels, i.e. for $i \in \mathcal{F}$, we let them be valid at any node for the sake of a more uniform notation, resulting in infinite capacities at the virtual nodes.

Let $C_0^a(\mathbf{z}_0^a)$ be the expected discounted cost at time 0 of product $a \in \mathcal{A}$ given its initial inventory positions $\mathbf{z}_0^a$, as expressed by (1), and $C_0(\mathbf{z}_0^{a_1}, \ldots, \mathbf{z}_0^{a_{n_a}})$ be the total expected discounted cost across ASINs. The expected inventory volume at node i and period t is given by

$$\sum_{a\in\mathcal{A}} \mathbb{E}\left[\left(z_t^{a,i} - z_t^{a,i-1}\right)^+ v^a\right].$$

Our goal is thus to solve the following problem:

$$\begin{aligned} C_0(\mathbf{z}_0^{a_1}, \ldots, \mathbf{z}_0^{a_{n_a}}) = \min \sum_{a\in\mathcal{A}} \sum_{t=0}^{T} \gamma^t H_t(\mathbf{z}_t^a) \\ \text{s.t. } \mathbf{z}_{t+1}^a = \hat{\mathbf{z}}_t^a - D_t^a \mathbf{e} \\ \gamma^t \sum_{a\in\mathcal{A}} \mathbb{E}\left[\left(z_t^{a,i} - z_t^{a,i-1}\right)^+ v^a\right] \leq \gamma^t K^i,\ i \in [1, n], t \in [0, T], \end{aligned}$$

where the inclusion of the discount factor in the constraint is for notational convenience.

The capacity constraints link the products together, preventing us from solving the product-level problems independently. By dualizing these constraints using dual values $\lambda_t^{i'}$, we may recover separability across products at the expense of having to solve an outer optimization problem to find the optimal penalties $\boldsymbol{\lambda}'$. These are installation level penalties, which can equivalently be expressed at the echelon level by letting $\boldsymbol{\lambda}$ be such that $\lambda_t^i := \lambda_t^{i'} - \lambda_t^{i+1'}$. We then obtain the following formulation:

$$C_0(\mathbf{z}_0^{a_1}, \dots, \mathbf{z}_0^{a_{n_a}}) = \max_{\boldsymbol{\lambda}' \geq 0} \left\{ -\sum_{t=0}^{T} \gamma^t \mathbf{K}^T \boldsymbol{\lambda}'_t + \sum_{a \in \mathcal{A}} C_0^a(\mathbf{z}_0^a; \boldsymbol{\lambda}) \right\},$$

where $C_0(\mathbf{z}_0; \boldsymbol{\lambda})$ is a penalized form of $C_0(\mathbf{z}_0)$, whose objective function is now given by:

$$\mathbb{E}\left[\sum_{t=0}^{T} \gamma^t H_t(\mathbf{z}_t) + \sum_{t=0}^{T} \gamma^t \lambda_t^{i'} \left(z_t^{a,i} - z_t^{a,i-1}\right)^+ v^a\right].$$

Rearranging the terms in a manner identical to the one that was performed on the initial unconstrained problem, we find:

$$\mathbb{E}\left[\sum_{t=0}^{T} \gamma^t H_t(\mathbf{z}_t) + \sum_{t=0}^{T} \gamma^t \lambda_t^{i'} \left(z_t^{a,i} - z_t^{a,i-1}\right)^+ v^a\right] = \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t H(\mathbf{z}_t; \boldsymbol{\lambda}_t)\right],$$

where $H_t(\mathbf{z}; \boldsymbol{\lambda})$ are updated holding and backlogging costs that account for the capacity penalties:

$$H_t(\mathbf{z}; \boldsymbol{\lambda}) := \sum_{i=1}^{n} H_t^i(z^i; \lambda^i), \qquad H_t^i(z^i; \lambda^i) := (h_i + \lambda^i) v z^i + (b + (h'_1 + \lambda^{1'}) v) \mathbb{E}\left[(z - D_t)^-\right] \mathbf{1}_{\{i=1\}}.$$

These definitions make apparent the fact that the dual variables associated with the capacity constraints can be viewed as supplemental holding costs that are added to the regular holding costs h_i. The capacity constrained network problem can thus be tackled by solving an outer optimization problem that optimizes a function made up of the sum of independent product-level inner problems.

Outer Problem:

$$C_0(\mathbf{z}_0^{a_1}, \dots, \mathbf{z}_0^{a_{n_a}}) = \max_{\boldsymbol{\lambda}' \geq 0} \left\{ -\sum_{t=0}^{T} \gamma^t \mathbf{K}^T \boldsymbol{\lambda}'_t + \sum_{a \in \mathcal{A}} C_0^a(\mathbf{z}_0^a; \boldsymbol{\lambda}) \right\}. \tag{21}$$

Inner Problems:

$$\begin{aligned} C_0(\mathbf{z}_0; \boldsymbol{\lambda}) = \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t, \mathcal{N}_t)} \ & \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t H_t(\mathbf{z}_t; \boldsymbol{\lambda}_t)\right] \\ \text{s.t. } & \mathbf{z}_{t+1} = \hat{\mathbf{z}}_t - D_t \mathbf{e}. \end{aligned}$$

4.2.2 Flow Constraints

Flow constraints refer to the limited capacity of a facility in processing inbound or outbound units in any given period due to labor and operating constraints, which often necessitate that replenishment orders be spread out over time. They are related to production capacities in the sense that production capacities place limit on the flow at the decision nodes, while we are more interested in the flow at the FC nodes. Nonetheless, because we model the problem in a general way, both types of constraints are equally handled. This applies in particular to capacity constrained vendors. Parker and Kapuscinski [2004] study a two-echelon capacitated serial line where there is a smaller capacity at the downstream installation, and show that a Modified Echelon Base Stock (MEBS) policy is optimal. The objective function is however shown to lose its additive separability property, which would severely hinder our ability to efficiently implement a solution. We are nonetheless able to overcome this issue by formulating the problem in expectation across all products, in a manner similar to the way volumetric constraints were handled in Section 4.2.1.

We here again consider the problem at the serial line level for ease of notation, and let $\mathbf{F}$ be a vector of constraints on the number of units (equivalently volume) that can be inbounded in any period t at node i. The expected number of units to arrive in period $t > 0$ at node i is given by:

$$\mathbb{E}\left[\sum_{a\in\mathcal{A}}\left(\hat{z}^{a,i}_{t-1} - z^{a,i}_{t-1}\right)\right].$$

As a result, the flow constrained formulation of the problem reads:

$$\begin{aligned} C_0(\mathbf{z}_0^{a_1},\dots,\mathbf{z}_0^{a_{n_a}}) = \min & \sum_{a\in\mathcal{A}}\sum_{t=0}^{T}\gamma^t H_t(\mathbf{z}_t^a) \\ \text{s.t. } & \mathbf{z}^a_{t+1} = \hat{\mathbf{z}}^a_t - D^a_t\mathbf{e} \\ & \gamma^{t+1}\sum_{a\in\mathcal{A}}\mathbb{E}\left[\left(\hat{z}^{a,i}_{t-1} - z^{a,i}_{t-1}\right)\right] \le \gamma^{t+1}F^i,\ i\in[1,n], t\in[1,T]. \end{aligned}$$

Many of the variables we have used in the multi-echelon formulation possess an installation and an echelon interpretation, such as the holding costs or capacity penalties, which highlight the *vertical* trade-off between installation costs in any given period t. We for example defined the echelon capacity penalty λ^i_t as the difference between consecutive installation penalties $\lambda^i_t = \lambda^{i'} - \lambda^{i+1'}$. The dualization of the flow constraints with dual variables $\boldsymbol{\beta}'$ introduces a similar concept at a *horizontal* level, where we compare costs not across installations in a given period, but at a given installation across periods. This leads to the definition of a time-shifted installation cost $\beta^i_t := {\beta^i_t}' - \gamma {\beta^i_{t+1}}'$ that compares (discounted) costs at a given installation i in consecutive periods. The problem then reads:

$$C_0(\mathbf{z}_0^{a_1},\dots,\mathbf{z}_0^{a_{n_a}}) = \max_{\boldsymbol{\beta}'\ge 0}\left\{\sum_{t=0}^{T}\gamma^{t+1}\left(\mathbb{E}[D_t]\mathbf{e} - \mathbf{F}\right)^T\boldsymbol{\beta}'_t + \sum_{a\in\mathcal{A}}\tilde{C}^a_0(\mathbf{z}^a_0;\boldsymbol{\beta})\right\},$$

where $\tilde{C}_0(\mathbf{z}_0;\boldsymbol{\beta})$ is a penalized form of $C_0(\mathbf{z}_0)$, whose objective function is now given by:

$$\mathbb{E}\left[\sum_{t=0}^{T}\gamma^t\tilde{H}_t(\mathbf{z}_t;\boldsymbol{\beta}_t)\right],$$

where $\tilde{H}_t(\mathbf{z}_t;\boldsymbol{\beta}_t)$ are updated holding and backlogging costs that account for the flow penalties:

$$\tilde{H}_t(\mathbf{z};\boldsymbol{\beta}) := \sum_{i=1}^{n}\tilde{H}^i_t(z^i;\beta^i), \qquad \tilde{H}^i_t(z^i;\beta^i) := (h_i v + \beta^i)z^i + (b + h'_1 v)\mathbb{E}\left[(z - D_t)^-\right]\mathbf{1}_{\{i=1\}},$$

and where we recall that:

$$\beta^i_t = {\beta^i_t}' - \gamma{\beta^i_{t+1}}'.$$

The interpretation of the flow penalties resembles that of the capacity (volumetric) penalties. While the latter were akin to heightened holding costs, the flow penalties have the interpretation of supplemental transfer costs to be added to the original transfer costs w_i.

Similarly to the volume constraints, we can tackle the flow constrained problem by solving an outer optimization problem that optimizes a function made up of the sum of independent product-level inner problems.

Outer Problem:

$$C_0(\mathbf{z}_0^{a_1},\dots,\mathbf{z}_0^{a_{n_a}}) = \max_{\boldsymbol{\beta}'\ge 0}\left\{\sum_{t=0}^{T}\gamma^{t+1}\left(\mathbb{E}[D_t]\mathbf{e} - \mathbf{F}\right)^T\boldsymbol{\beta}'_t + \sum_{a\in\mathcal{A}}\tilde{C}^a_0(\mathbf{z}^a_0;\boldsymbol{\beta})\right\}. \tag{22}$$

Inner Problems:

$$C_0(\mathbf{z}_0;\boldsymbol{\beta}) = \min_{\hat{\mathbf{z}}_t \in \hat{Z}(\mathbf{z}_t,\mathcal{N}_t)} \mathbb{E}\left[\sum_{t=0}^{T} \gamma^t \tilde{H}_t(\mathbf{z}_t;\boldsymbol{\beta}_t)\right]$$
$$\text{s.t. } \mathbf{z}_{t+1} = \hat{\mathbf{z}}_t - D_t\mathbf{e}.$$

Remark 5

Given the finite-horizon nature of our formulation, we have taken little care in ensuring that the problem is realistically feasible, letting the terminal conditions wipe out any accrued backlog. A major risk in the presence of backlogs is that too strict a constraint can lead to an unstable behavior where the inventory position keeps drifting downwards because of the impossibility to bring in sufficient inventory to meet demand and backlogs. The stability of a capacitated multi-echelon network was studied in Glasserman and Tayur [1994] and Huh et al. [2010]. To render the problem stable, we would need a condition of the form $\sum_{a\in\mathcal{A}} \mathbb{E}[D_t^a] \leq \min_i F^i,\ \forall t$ for example.

5 Implementation

5.1 Overview

We present in this section some implementation details of the solution method for the capacitated distribution network described in Sections 2 and 4. The implementation issues to be considered concern the two levels of optimization that we are to solve: (i) at the serial line level, and (ii) at the network level. These two classes of problems correspond to the inner and outer loops in the dual formulation of the problem in the presence of capacity constraints, as detailed in Section 4.2.

At the serial line level, and as is frequent in inventory management given the often Markovian nature of the formulations, the problem lends itself to a Dynamic Programming (DP) implementation. This corresponds to the recursive solution of the optimality equation (2). The state dimension of the serial lines can be relatively large as it is roughly equal to the length of the lead time from the vendor to the customer facing facility, which in general would render any practical implementation computationally prohibitively expensive due to the "curse of dimensionality". This is were the structural properties of the cost functions derived in Section 2, and in particular their additive convexity, are especially important as they curb the growth of the problem's complexity making it linear instead of exponential, allowing for a practical implementation.

At the network level, we must solve for the optimal dual variables associated with the capacity (volume and flow) constraints. Solving at the network level requires that the serial lines that it comprises be iteratively solved for different values of these penalties as they are optimized over. This process can become somewhat computationally costly and effective distributed methods to solve them must be employed.

At either level, we must also consider the impact of some of the assumptions, which while necessary for the derivation of the strong theoretical results we rely on, may lack some realism and be improved upon by careful modifications to the optimization procedures. The following sections address these main points.

5.2 Primal Problem: Serial Line

5.2.1 Solving the Primal Problem

The primal problem corresponds to solving the serial line formulations laid out in Section 2, which also represent the inner problems in the dual representation of the constrained network formulation of Section 4. These are expressed by the optimality equation (2) and are to be solved using the methodology laid out in Proposition 2.1, which suggests a Dynamic Programming approach to the problem. This approach relies on a discretization of the random variables and the state space, as well as on the additive convexity property of the cost functions that makes a practical implementation of the DP feasible.

5.2.2 Discretization

Demand Solving the serial lines requires the optimization, and thus evaluation, of functions expressed as expectations over random variables, most notably the random demands. In order to evaluate these expectations, and in the absence of closed form expressions thereof, we must often resort to a discretization of a distribution that might initially have been expressed as continuous. A common approach in practice is the representation of the distribution through a set of K equally probably quantiles (see for example Raz and Porteus [2006]).

States The state space used to represent the serial line consists of the echelon inventory positions at the various nodes, whether physical or virtual, and thus corresponds to $\mathbb{R}^L$, where L is the lead time from the vendor to the customer facing FC. For implementation purposes, this space needs to be bounded and discretized, and we must choose a minimum position m, a maximum position M and a discretization grain $N+1$ corresponding to the number of points $\{y_i\}_i$ chosen in that interval. Letting the points be equidistant, we have $y_i = m + i\frac{M-m}{N}$, and the resulting discretized space given by $\{y_0, \dots, y_N\}^L$. Given such a grid discretization of the state space, the DP can be implemented by evaluating and storing the value of the cost functions at points $\mathbf{z} \in \{y_0, \dots, y_N\}^L$, from which approximations of the values at non-grid points $\mathbf{z}$ are obtained by interpolation/extrapolation.

5.2.3 Additive Convexity

Curse of Dimensionality Following the state space discretization, as described in the preceding section, the recursive solving of the optimality equation (2) necessitates that we store values $C_t(\mathbf{z})$ at each discretized point $\mathbf{z} \in \{y_0, \dots, y_N\}^L$. In general, this would require the optimization and storage of $O\left((N+1)^L\right)$ values in each period, which is prohibitively expensive even for small values of L. As an illustration, for a discretization of $N+1 = 100$ and $L = 6$, storing those values as floats would require 4TB of data in each time period, let alone the actual computational burden itself.

The additive convexity property of the cost functions (Definition 2.1) allows us to condense the necessary amount of information to keep track of. Instead of having to keep track of the values of $C_t(\mathbf{z})$ at each discretized point $\mathbf{z} \in \{y_0, \dots, y_N\}^L$, we may only keep track of the values of L univariate functions $\{c_t^i\}_{i=1,\dots,L}$ at the $N+1$ points $\{y_0, \dots, y_N\}$. Any value $C_t(\mathbf{z})$ for $\mathbf{z} \in \{y_0, \dots, y_N\}^L$ can then be recreated as $C_t(\mathbf{z}) = \sum_{i=1}^{L} c_t^i(z^i)$. The benefits of this separability is illustrated in Figure 16 on a two dimensional example. Going back to the example above with $N+1 = 100$ and $L = 6$, we would only require the solving and storing of $(N+1)L$ problems and points, which in this case would amount to 4MB.

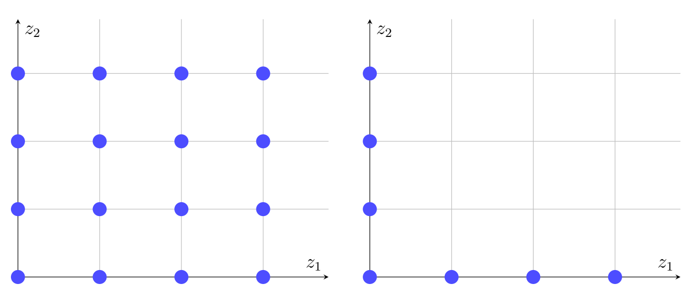

Figure 16: States that require solving and storing in the cases where the objective function is: (left) nonseparable, (right) separable.

Function Representation More concretely, the functions $\{c_t^i\}$ are approximated by a piecewise linear representation that is obtained by interpolating between the stored values $c_t^i(y_j)$ for $j = 0, \dots, N+1$. (If values outside of the discretized range $[m, M]$ are needed, they are obtained by extrapolating outside of the range using the slope of the approximating function at the edge). This representation can be interpreted as storing T matrices whose rows correspond to univariate functions c_t^i, and the columns to the point y_j, as illustrated in Figure 17.

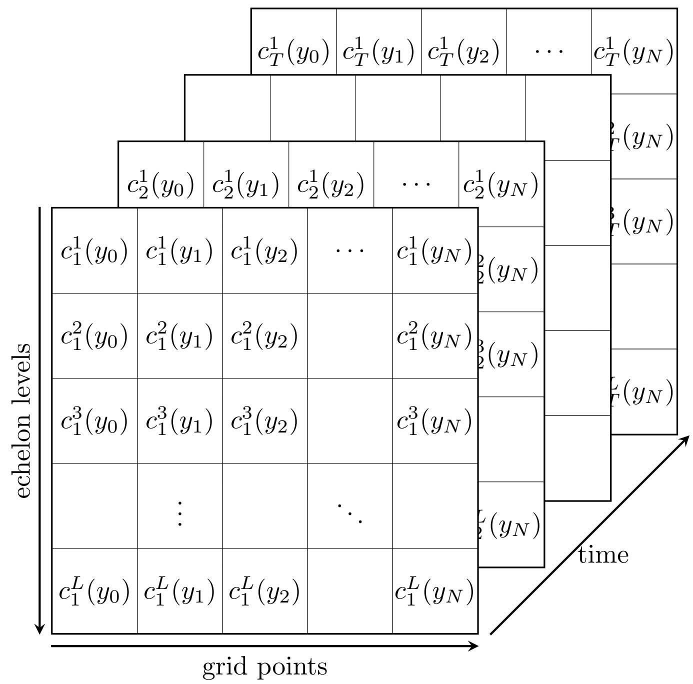

Figure 17: Illustration of the matrix representation of the cost functions.

5.3 Dual Problem: Network

5.3.1 Solving the Dual Problem

At the network level, in the presence of capacity or flow constraints (or both), we are to solve the dual problem represented by Equations (21) or (22) (or a combination of both). While solving the serial line through dynamic programming only requires the optimization of univariate convex functions (see Proposition 2.1), which can be efficiently solved using simple algorithms such as the Golden Section search (see e.g. [Luenberger and Ye, 2016, Ch.8]), the dual problem is a multivariate problem that involves the dual variables associated with the constraints at several installations and time periods. The evaluation of the serial line for fixed values of these dual variables itself can result onerous given since we must solve a serial line problem for each SKU-Customer facing installation pair. Given that these combinations can number in the millions, it is of paramount importance that we are able to solve the dual problem as efficiently as possible, leading us to consider distributed derivative based methods. The latter qualifier (derivative-based) is necessary for fast convergence in terms of the number of iterations, while the former (distributed) is crucial in order to speed up the computational time of a single iteration.

5.3.2 Estimating the Gradients

In order to implement a gradient-based method, we must first be able to evaluate the gradient of the function we are to optimize. Considering only the volume constraints for ease of exposition, the dual problem to be solved is given by Equation (21):

$$\max_{\boldsymbol{\lambda}' \geq 0} \phi(\boldsymbol{\lambda}'),$$

where

$$\phi(\boldsymbol{\lambda}') := \sum_{a \in \mathcal{A}} C_0^a(\mathbf{z}_0^a; \boldsymbol{\lambda}) - \sum_{t=0}^{T} \gamma^t \mathbf{K}^T \boldsymbol{\lambda}'_t.$$

The gradient $\nabla_{\boldsymbol{\lambda}'} \phi$ is however readily obtained since it is directly given by the violation of the constraints [Luenberger and Ye, 2016, Ch.14]. In other words, for each dual variable $\lambda_t^{i\prime}$, the partial derivative $\frac{\partial \phi}{\partial \lambda_t^{i\prime}}$ is given by:

$$\frac{\partial \phi}{\partial \lambda_t^{i\prime}} = \gamma^t \left(\sum_{a \in \mathcal{A}} \mathbb{E}\left[\left(z_t^{a,i} - z_t^{a,i-1} \right)^+ v^a \right] - K^i \right). \tag{23}$$

The calculation of the gradient of ϕ is hence no more expensive than the evaluation of the network for fixed dual values.

In order to estimate (23), we must however derive the values of the $z_t^{a,i}$, the echelon inventory position for each serial line, in each time period, and at each echelon. It is important to recall at this point that these sample path estimations carry little value at an individual serial line level, other than being the "expected" inventory positions, but that we only consider them insofar as their aggregate values are more meaningful and, provided we deal with sufficient serial lines, good estimates of the installation level resource utilization.

Given the initial inventory positions $\mathbf{z}_0^a$, we iteratively apply the deterministic approximation of the transition function to estimate $\mathbf{z}_t^a$ for $t > 0$:

$$z_{t+1}^i = \min\left(z_t^{i+1}, \hat{z}_t^i\right) - \mathbb{E}[D_t], \tag{24}$$

and use these values to estimate the gradients (23).

Remark 6
The "sample path estimation" above abides by the backordering assumption we have made throughout. While necessary to derive all the theoretical results, there are practical issues with this approach. For example, backordered units will not carry any holding cost penalty at the customer facing facility, while even if the assumption were true, units would still need to transit through it and be, even if only momentarily, stored. We discuss in Section 5.4 a variation on this estimation to take into account the more realistic lost sales scenario and to adjust the gradient computation and better manage resources.

5.3.3 Optimization

The dual network problem is separable into serial line level subproblems. A common strategy is thus to solve these subproblems in parallel for fixed values of the dual variables, gather the results, compute the gradient of the dual function, and update the dual variables using some first order method update. Such a method works relatively well in many instances but can result in somewhat slow convergence, and more importantly may suffer from instability in the absence of strict convexity, a phenomenon which we observe. To remedy these issues, the Alternative Direction Method of Multipliers (ADMM) suggests considering the augmented Lagrangian [Boyd et al., 2011]. To preserve the separability of the network problem into serial line subproblems, two approaches may be considered:

- The traditional Gauss-Seidel approach of ADMM where two groups of variables are sequentially updated,
- A Jacobi approach where each block (serial line in our case) is solved in parallel.

The Gauss-Seidel approach requires the introduction of auxiliary variables, leading to a Variable Splitting ADMM [Wang et al., 2013]. The Jacobi approach is a more direct one that avoids the addition of new variables and solves the subproblems in parallel by fixing all other primal variables in each subproblem. The convergence of this approach is not guaranteed unless an additional proximal term is added [Deng et al., 2017]. We have found good results with the latter approach.

5.4 Backlogging vs Lost-Sales

There often exists a disconnect between the assumptions made for theoretical or practical purposes, and those that are empirically realistic. Our work does not escape this pitfall as we relied, as is often necessary in inventory management, on the assumption that missed sales are backlogged, as opposed to lost. This assumption was paramount in the derivation of our structural properties, notable among them the additive convexity of the cost functions, which in turn proved central to an efficient solution implementation. It was required in order for the transition function to not depend on the inventory level at the customer facing facility. Although by setting appropriate penalties for backlogged units it is possible to approximate the lost sales scenario quite well, especially in the higher service levels regimes, such an accounting can prove deceitful in the presence of capacity or flow constraints. Backlogged units will often not be applied the capacity penalties, which is not realistic as even if that assumption were true, units would still have to be temporarily stored. Similarly, such a behavior could shift units from downstream to upstream facilities and unrealistically inflate storage at those less constrained nodes, when in reality such units would not be necessary since their corresponding demand has already been lost.

As a result, some mitigating measures can be taken to address the issue. In particular, the transition function (24) in the evaluation of the expected inventory positions can be modified to yield more realistic results in practice. Specifically, in each period, we can compute the "expected drain" as $\mu_t = \mathbb{E}[(D_t - z_t^0)^+]$ and replace $\mathbb{E}[D_t]$ with μ_t in (24). This is possible in the forward pass that evaluates the expected inventory positions, and eliminates some spurious behaviors associated with the backlogging assumption.

6 Examples and Results

6.1 Illustrative Examples

We illustrate in this section the results of the model described in Sections 2, 4 and whose implementation was discussed in Section 5. We consider simple scenarios to demonstrate some of the main features of our model, such as the asynchronicity of the review periods, or the handling of capacity constraints.

6.1.1 Simple Serial Line

We first consider a 4 echelon serial line for a single product with the characteristics presented in Table 1. The demand we consider for this product is almost stationary, with a Gamma distribution with mean

Table 1: Parameters of the product, serial line, and DP in the simple serial line example.

Product	
c	6
b	5
v	1

	FC1	FC2	FC3	FC4
Lead Time	1	2	3	4
h	2	1	0.5	0.3

DP	
γ	0.99
quantiles	99

50 in every period except for a "peak" around period 23, with mean 75 in periods 22 and 24 and 150 in period 23. The coefficient of variation of the demand is 0.5 in every period.

Single period review period We first consider the case where the review periods are the same for all echelons and equal to one period. The resulting Target Inventory Positions in each period, along with the mean demand and shaded area between the 25th and 75th quantiles of the demand, are plotted in Figure 18. We observe the impact of the lead times and how the peak demand affects upper echelons earlier than lower ones as units need to flow down the serial line in order to be at the spoke node in time to meet demand.

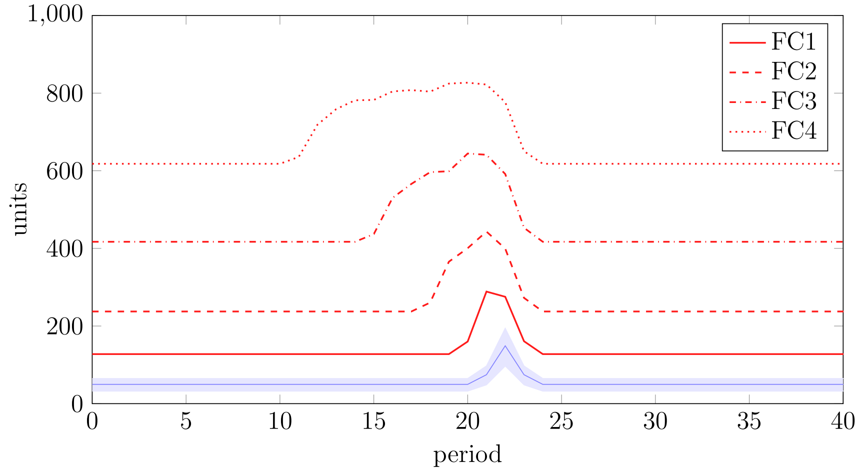

Figure 18: Target Inventory Position for the simple serial line example in the case of a unit review period at all echelons.

Asynchronous review periods We now consider the case of asynchronous review periods where the review periods are neither the same, nor necessarily a multiple of the lower echelon, as shown in Table 2.

Table 2: Review periods used in the asynchronous scenario.

	FC1	**FC2**	**FC3**	**FC4**
Review Period	2	3	4	5

Similarly to the case of the single, the results are presented in Figure 19. The interpretation of the results requires a bit more thought than in the previous scenario. The interplay between the demand shape, lead times and asynchronous review periods produces results that are more arduous to analyze. Still, we observe that in order to find a path that allows for units to reach the spoke node in time to satisfy demand, the Target Inventory Positions need to be increased earlier for most echelons. These results nonetheless illustrate the flexibility of our model, which does not require any synchronicity assumption in terms of review periods at different echelons.

6.1.2 Volume Constrained FC

We consider in this example a simple two-echelon network where the demand facing FC is volume constrained. We will let the volume constraint vary and observe how the TIPs evolve for 4 economically identical products who only differ in their volume. The parameters of the problem are shown in Table 3,

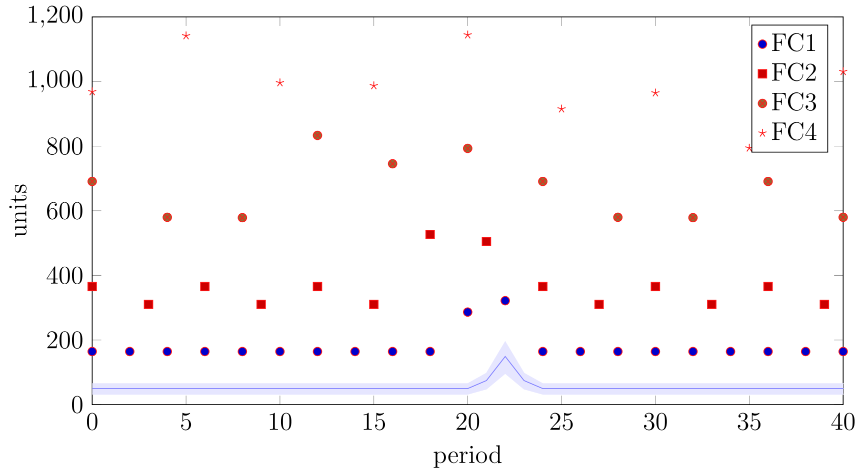

Figure 19: Target Inventory Position for the simple serial line example in the case of a non-homogeneous review periods at all echelons.

and we further assume stationary and identically Gamma distributed demands for all 4 products, with means 50 and coefficient of variation of 1.

Table 3: Parameters of the product, serial line, and DP in the simple serial line example.

Product i	
c	6
b	5
v	i

	FC1	**FC2**
h	0	0
Lead Time	1	2
Review Period	1	1

DP	
γ	0.99
quantiles	99

The example is purely illustrative as we recall that the assumption that motivated our formulation of the problem is that we dealt with a large enough number of products that the law of large numbers was approximately valid. Because we are in a simple stationary setting, we may consider the service levels of the products in any given period. We plot in Figure 20 the service levels of all 4 products as a function of the capacity at the demand facing FC. When the capacity constraint is almost non-existent, i.e. for high values of the FC capacity, the service levels for all 4 products are essentially identical and converge to 1. However, as the capacity constraint becomes tighter, we start seeing a differentiation in the service levels of the products where less voluminous products are favored at the expense of more voluminous ones, as we would expect. Such mechanism is paramount in properly handling capacity constraints and in trading off the cost of capacity opportunity against the profitability of the products.

6.1.3 Inbound Constrained FC

We here consider the other type of constraint explored in our modeled, namely inbound constraints (Section 4.2.2. It is common for many businesses to have "peak period" such as around Thanksgiving/Black Friday or the end of the year holidays. The inventory for these peak periods usually needs to be gradually taken in and their arrivals to be spread over a period of time so as to not overflow the processing capacity of the facilities, especially when it might already be split between heavy inbound and outbound operations.

We thus consider a simple and slightly contrived example of a product whose demand is mostly stationary, with a mean demand of 50 units except for a peak period with a mean demand of 1,000 units shouldered by two periods of mean demand of 100 units. We assume Gamma distributions with

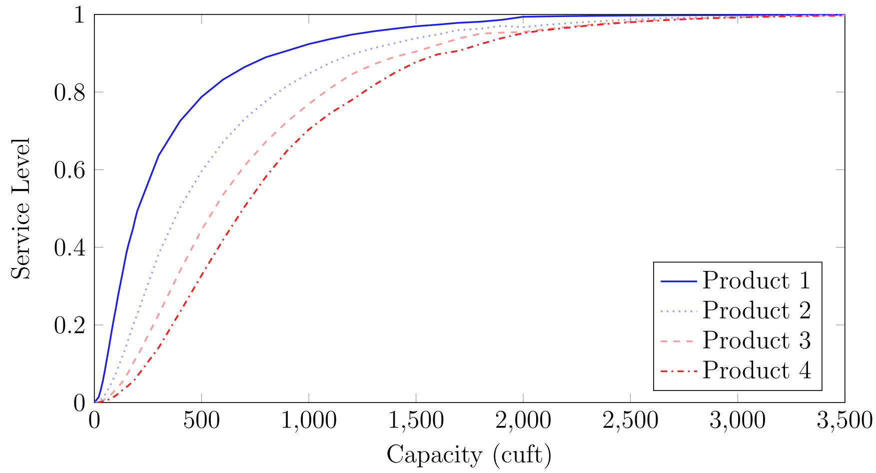

Figure 20: Service levels as a function of demand facing FC capacity for the four products considered in the volume constrained example.

coefficient of variations of 0.5. The parameters of the product, serial line and DP are shown in Table 4 and are very similar to the ones in the volume constrained example above.

Table 4: Parameters of the product, serial line, and DP in the simple serial line example.

Product	
c	6
b	5
v	1

	FC1	FC2
h	0.5	0.2
Lead Time	1	2
Review Period	1	1

DP	
γ	0.99
quantiles	99

We plot in Figure 21 the Target Inventory Positions at the demand facing FC in two cases: one in which there is no constraint on the inbound quantity, and another one where the inbound flow is restricted to at most 150 units per period. As a result, we observe that in the constrained case, the Target Inventory Positions increase earlier in order to ramp up inventory so as to hold the required units in time for the peak period, as would be expected. The inbound flow penalties (dual variables) are thus seen to perform their role in propagating the constraint information through time and triggering the right anticipatory buying behavior.

6.1.4 Time-Dependent Costs

It is common for vendors to periodically offer discounts to retailers on some of their products. This can happen for a range of reasons such as nearing the end season of a seasonal product, or trying to secure a large buying commitment. The presence of these periodic discounts will impact our decisions, especially in those periods where the discounts apply. It is then especially pertinent to leverage a multi-period model that is able to trade-off the savings generated by purchasing units at a discount and the additional holding costs they will generate, by taking into account non-stationary demand forecasts, as well as knowledge, or anticipation, of the periodicity at which these discounts occur. Conversely, vendors might request that the retailer commit to purchasing a given number of units, in which case the retailer will need to figure out an appropriate discount for such quantity.

As previously mentioned, time-dependent costs are easily incorporated in our model. In the present scenario, the purchasing cost c_t of the product in period t, heretofore constant, is time dependent. Such dependence on time only requires a minor adjustment to our model. Recall from Section 2.2.1 that following the the reformulation of the problem (as detailed in Appendix A.1), we defined the echelon

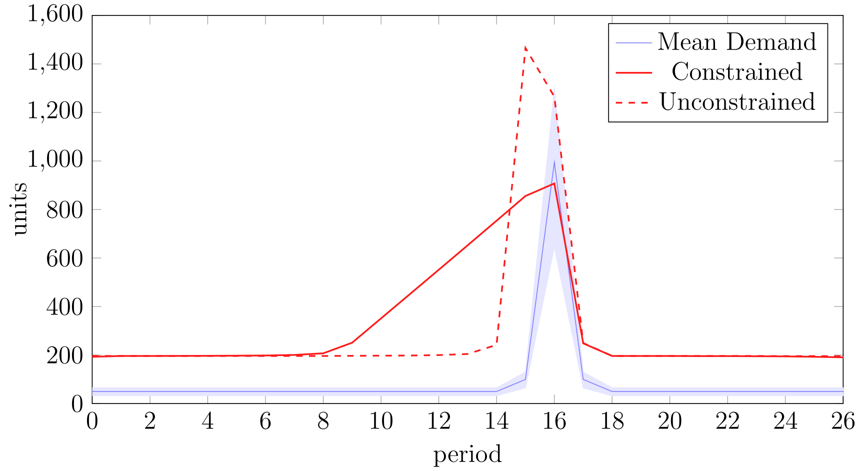

Figure 21: Target Inventory Positions at the demand facing FC in the inbound constrained example.

holding cost at the uppermost echelon as $h_n := h_n + (1-\gamma)\frac{c}{v}$ (discarding the transfer cost to avoid cumbersome expressions). In the presence of a time-dependent cost, the exact same transformation carries by letting the new echelon cost be $h_n^t := h_n + (c_{t-1} - \gamma c_t)\frac{1}{v}$ in period t.

To illustrate the impact of these periodic discounts on the optimal TIPs in the discounted periods, we look at the impact of two factors: 1) the discount periodicity, and 2) the discount amount. To simplify the presentation of the results, we consider a single echelon problem with the parameters presented in Table 5, and a stationary Gamma distributed demand with mean demand 50 and coefficient of variation equal to 0.5.

Table 5: Parameters of the product, serial line, and DP in the simple serial line example.

Product			**FC1**	**DP**	
c	6	h	0.05	γ	0.99
b	5	Lead Time	3	quantiles	99
v	1	Review Period	1		

We first investigate the impact of the discount periodicity on the TIPs in periods when the discount applies. Intuitively, we should expect the TIPs to increase with longer times in between possible discounted purchases: if we are presented with a discount opportunity every other period, we should at most be setting TIPs to cover two periods, but the wider the spacing, the larger the opportunity for savings. Eventually, we will reach a point where the marginal expected holding cost will overtake the savings from the discounted cost. Figure 22 shows the evolution of the TIP in discounted periods with the periodicity of the discount, for a discount of 20%. As expected, the TIP increases with the period before reaching a plateau. The results also underline the importance of taking into account the possibility of a discount in the future, as opposed to assuming that opportunity is a one-off event.

We now turn to the relation between TIP and discount percentage by fixing the periodicity of the discounts to 10 periods. Figure 23 shows the evolution of the TIP as a function of the discount amount for different values of the holding cost. As anticipated, the TIP increases rapidly with the discount amount before stagnating once the marginal savings from the discount approach the marginal holding cost. Similarly, higher holding costs require deeper discounts to reach the same TIP levels.

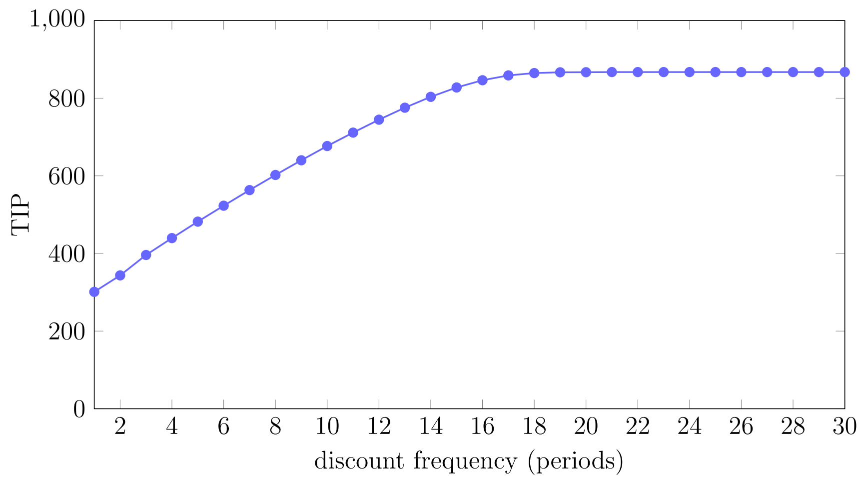

Figure 22: Evolution of the TIP in discounted periods with the periodicity of a 20% discount.

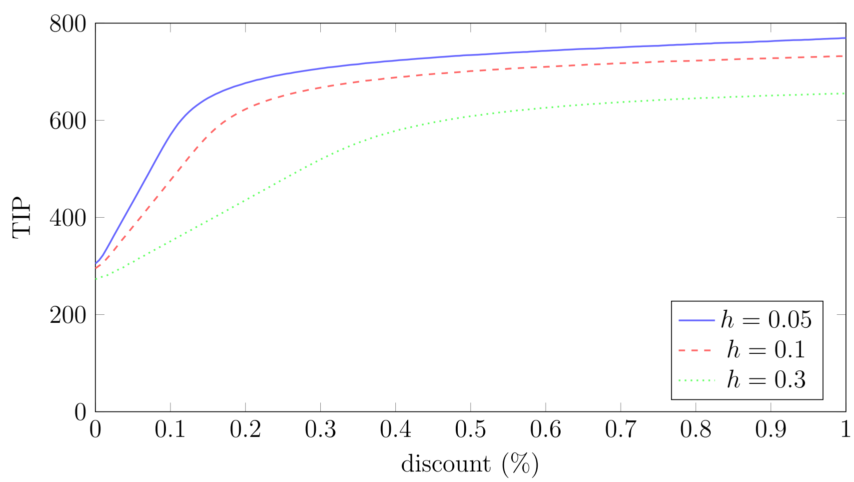

Figure 23: Evolution of the TIP in discounted periods with the discount percentage for a discount period of 10.

6.1.5 Random Vendor Supply

We illustrate in this section the impact of random vendor supply on the order-up-to levels, following the discussion in Section 3.4. We consider a single echelon problem whose parameters are described in Table 6, and a stationary Gamma distributed demand with mean 50 and coefficient of variation of 0.5.

Table 6: Parameters of the product, serial line, and DP in the random vendor supply example.

Product	
c	6
b	5
v	1

	FC1
h	0.5
Lead Time	4
Review Period	1

DP	
γ	0.99
quantiles	99

We then let the expected vendor yield vary and observe the resulting order-up-to policy, recalling that the vendor yield is modeled through a Bernoulli random variable. Figure 24 shows the evolution of the order-up-to level as a function of the expected vendor yield. We observe, as expected, that the more uncertain the yield, the higher the TIP since orders need to be inflated so that larger orders make up for the ones that are not confirmed.

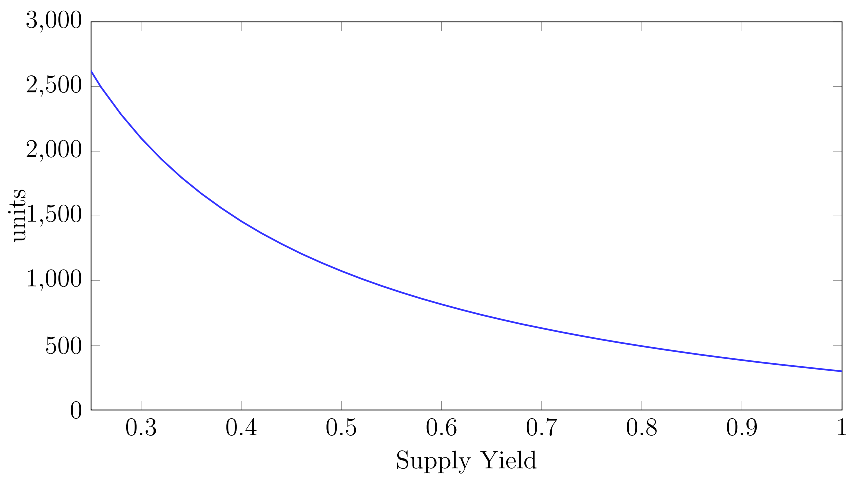

Figure 24: Evolution of the TIP as a function of the vendor yield.

6.2 Results

A version of this model was implemented at a large grocery retailer, and A/B tested in production in 2019. The test was performed in a two-layer network with a hub feeding four downstream and demand facing facilities, with products either flowing through the hub, or also potentially being delivered directly at the FCs. The products were further differentiated by their storage temperature requirements and divided into 3 categories: ambient, chilled and frozen.

The existing solution against which our model was tested set its target inventory positions on a model based on periods of cover. While yielding high levels of demand-weighted in-stockness, the previous model caused volume capacity issues at the hub and lacked capacity management capabilities, especially in non-ambient temperature zones. Although capacity constraints are expected to be reached at retail facilities that are usually of limited size due to their proximity to customers, such should not in general be the case at the hub level.

The A/B test ran over 3 months and yielded the statistically significant results shown in Table 7 (at a 0.05 p-value level). Our model was able to:

- Increase the already high service levels of the supply chain,
- Increase sales and revenue,
- Reduce hub inventory volume.

In addition, while not significant at the 0.05 level, the hub inventory volume for ambient products was also directionally lower.

Table 7: Statistically significant results of the A/B test.

Metric	Impact
Ordered Product Sales ($)	+4.37%
Sold Units	+4.63%
Demand-Weighted Availability	+1.75%
Non-Ambient Hub Inventory Volume (cuft)	-10.77%

7 Conclusion and Future Research

We presented in this paper a general inventory management solution for a non-stationary capacitated multi-echelon distribution network. Our approach leveraged and extended existing results in the literature to contribute both to further the theoretical understanding of multi-echelon networks, as well as to developing practical solutions. The practical benefits of this work were demonstrated at a large grocery retailer where its implementation in an A/B test resulted in statistically significant increases in sales, revenue, and availability, while reducing the volume of inventory.

The framework developed in this paper readily lends itself to a number of extensions, such as stochastic yields from the vendors to incorporate supply risk, time-dependent costs, or additional cross-product constraints for example. One avenue of future research we are to explore is to extend the setting of the problem to allow for demand to not be restricted to particular FCs. The current setting can somewhat be associated with physical stores where demand manifests at the stores and can only be fulfilled from that particular location. While such setting is particularly well suited in the context of a grocery retailer, other businesses, especially online or multi-channel retailers, do not necessarily have a mapping of demand to a particular FC, but can instead fulfill demand from several locations, albeit at possibly different costs depending on the proximity of these FCs to the demand node. In future work, we will extend our work to incorporate this type of "demand spillover" effect.

Acknowledgments

We would like to thank Anna Wang, Kevin Hajissa and Yiyan Zhu, whose contribution made the implementation of the work in production possible. We would also like to acknowledge Ali Sadighian and Nimesh Panchal for their help and valuable feedback.

References

Alexandar Angelus. A Multiechelon Inventory Problem with Secondary Market Sales. *Management Science*, 57(12):2145–2162, 2011. 23

Alexandar Angelus and Özalp Özer. When Variability Trumps Volatility: Optimal Control and Value of Reverse Logistics in Supply Chains with Multiple Flows of Product. *Mays Business School Research Paper*, (3071398), 2018. 23

Srinivas Bollapragada and Thomas E Morton. Myopic Heuristics for the Random Yield Problem. *Operations Research*, 47(5):713–722, 1999. 26

Stephen Boyd, Neal Parikh, and Eric Chu. *Distributed Optimization and Statistical Learning via the Alternating Direction Method of Multipliers*. Now Publishers Inc, 2011. 36

Xiuli Chao, Xiting Gong, Cong Shi, and Huanan Zhang. Approximation Algorithms for Perishable Inventory Systems. *Operations Research*, 63(3):585–601, 2015. 8, 20

Fangruo Chen and Yu-Sheng Zheng. Lower Bounds for Multi-Echelon Stochastic Inventory Systems. *Management Science*, 40(11):1426–1443, 1994. 8, 11, 15

Xin Chen, Zhan Pang, and Limeng Pan. Coordinating Inventory Control and Pricing Strategies for Perishable Products. *Operations Research*, 62(2):284–300, 2014. 20

Andrew J Clark and Herbert Scarf. Optimal Policies for a Multi-Echelon Inventory Problem. *Management Science*, 6(4):475–490, 1960. 4, 8

Wei Deng, Ming-Jun Lai, Zhimin Peng, and Wotao Yin. Parallel Multi-Block ADMM with O(1/k) Convergence. *Journal of Scientific Computing*, 71(2):712–736, 2017. 37

Richard Ehrhardt. (s, S) Policies for a Dynamic Inventory Model with Stochastic Lead Times. *Operations Research*, 32(1):121–132, 1984. 17

Awi Federgruen and Paul Zipkin. Computational issues in an infinite-horizon, multiechelon inventory model. *Operations Research*, 32(4):818–836, 1984. 8

Guillermo Gallego and Özalp Özer. A New Algorithm and a New Heuristic for Serial Supply Systems. *Operations Research Letters*, 33(4):349–362, 2005. 15

Guillermo Gallego and Paul Zipkin. Stock Positioning and Performance Estimation in Serial Production-Transportation Systems. *Manufacturing & Service Operations Management*, 1(1):77–88, 1999. 11

Paul Glasserman and Sridhar Tayur. The Stability of a Capacitated, Multi-Echelon Production-Inventory System under a Base-Stock Policy. *Operations Research*, 42(5):913–925, 1994. 33

Mordechai Henig and Yigal Gerchak. The Structure of Periodic Review Policies in the Presence of Random Yield. *Operations Research*, 38(4):634–643, 1990. 26

Woonghee Tim Huh, Ganesh Janakiraman, and Mahesh Nagarajan. Capacitated Serial Inventory Systems: Sample Path and Stability Properties under Base-Stock Policies. *Operations Research*, 58 (4-part-1):1017–1022, 2010. 33

Robert S Kaplan. A Dynamic Inventory Model with Stochastic Lead Times. *Management Science*, 16 (7):491–507, 1970. 17

Samuel Karlin and Herbert Scarf. Inventory Models of the Arrow-Harris-Marschak Type with Time Lag. In Kenneth J. Arrow, Samuel Karlin, and Herbert Scarf, editors, *Studies in the Mathematical Theory of Inventory and Production*, chapter 10, pages 155–178. Stanford University Press, Stanford, CA, 1958. 11

William Karush. A Theorem in Convex Programming. *Naval Research Logistics (NRL)*, 6(3):245–260, 1959. 9

Qing Li, Peiwen Yu, and Xiaoli Wu. Managing Perishable Inventories in Retailing: Replenishment, Clearance Sales, and Segregation. *Operations Research*, 64(6):1270–1284, 2016. 20

David G. Luenberger and Yinyu Ye. *Linear and Nonlinear Programming*. Springer, 4th edition, 2016. ISBN 3319188410. 35, 36

Alvaro Maggiar and Ali Sadighian. Joint Inventory and Revenue Management with Removal Decisions, 2017. URL https://ssrn.com/abstract=3018984. 23

Alp Muharremoglu and John N Tsitsiklis. A Single-Unit Decomposition Approach to Multiechelon Inventory Systems. *Operations Research*, 56(5):1089–1103, 2008. 17

Steven Nahmias. Myopic Approximations for the Perishable Inventory Problem. *Management Science*, 22(9):1002–1008, 1976. 20, 21

Steven Nahmias. Simple Approximations for a Variety of Dynamic Leadtime Lost-Sales Inventory Models. *Operations Research*, 27(5):904–924, 1979. 17

Steven Nahmias. Perishable Inventory Theory: A Review. *Operations research*, 30(4):680–708, 1982. 20

Purushottaman Nandakumar and Thomas E Morton. Near Myopic Heuristics for the Fixed-Life Perishability Problem. *Management Science*, 39(12):1490–1498, 1993. 20

Süleyman Özekici and Mahmut Parlar. Inventory Models with Unreliable Suppliersin a Random Environment. *Annals of Operations Research*, 91:123–136, 1999. 26

Rodney P Parker and Roman Kapuscinski. Optimal Policies for a Capacitated Two-Echelon Inventory System. *Operations Research*, 52(5):739–755, 2004. 31

Gal Raz and Evan L Porteus. A Fractiles Perspective to the Joint Price/Quantity Newsvendor Model. *Management science*, 52(11):1764–1777, 2006. 34

Ying Rong, Zümbül Atan, and Lawrence V Snyder. Heuristics for Base-Stock Levels in Multi-Echelon Distribution Networks. *Production and Operations Management*, 26(9):1760–1777, 2017. 2, 29, 30

Kevin H Shang and Jing-Sheng Song. Newsvendor Bounds and Heuristic for Optimal Policies in Serial Supply Chains. *Management Science*, 49(5):618–638, 2003. 11, 15

Geert-Jan Van Houtum, Alan Scheller-Wolf, and Jinxin Yi. Optimal Control of Serial Inventory Systems with Fixed Replenishment Intervals. *Operations Research*, 55(4):674–687, 2007. 8, 15, 16

Xiangfeng Wang, Mingyi Hong, Shiqian Ma, and Zhi-Quan Luo. Solving Multiple-Block Separable Convex Minimization Problems Using Two-Block Alternating Direction Method of Multipliers, 2013. URL https://arxiv.org/abs/1308.5294. 37

A Reformulations

A.1 Reformulation of the Optimization Problem

Recall that the objective function of the original optimization problem reads:

$$\mathbb{E}\left[\sum_{t=0}^{T-1}\gamma^t\left(\gamma c(\hat{z}_t^n - z_t^n) + b\left(z_t^1 - D_t\right)^- + \sum_{i=2}^{n} h_i' v(z_t^i - z_t^{i-1}) + h_1' v(z_t^1 - D_t)^+ \right.\right.$$
$$\left.\left. + \sum_{i=1}^{n}\gamma w_i v\left(\hat{z}_t^i - z_t^i\right)\right) - \gamma^T F_T(\mathbf{z}_T)\right]. \tag{25}$$

- Consider first the terms corresponding to the holding and backlogging costs. Using the fact that $h_i = h_i' - h_{i+1}'$ (with the convention that $h_n = h_n'$) and that $z = (z)^+ - (z)^-$, we can rearrange the terms to obtain:

$$b\left(z_t^1 - D_t\right)^- + \sum_{i=2}^{n} h_i' v(z_t^i - z_t^{i-1}) + h_1' v(z_t^1 - D_t)^+ = \sum_{i=1}^{n} h_i v z_t^i + (b + h_1' v)\left(z_t^1 - D_t\right)^- - h_1' v D_t. \tag{26}$$

- Consider now the purchasing and transfer costs, and recall that $\hat{z}_t^i = z_{t+1}^i + D_t$. We thus have:

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=0}^{T-1}\gamma^t\left(\gamma c(\hat{z}_t^n - z_t^n) + \sum_{i=1}^{n}\gamma w_i v\left(\hat{z}_t^i - z_t^i\right)\right)\right] \\
&= \mathbb{E}\left[\sum_{t=0}^{T-1}\gamma^t\left(\gamma c(z_{t+1}^n + D_t) + \sum_{i=1}^{n}\gamma w_i v\left(z_{t+1}^i + D_t\right)\right) - \sum_{t=0}^{T-1}\gamma^t\left(\gamma c z_t^n + \sum_{i=1}^{n}\gamma w_i v z_t^i\right)\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T}\gamma^{t-1}\left(\gamma c z_t^n + \sum_{i=1}^{n}\gamma w_i v z_t^i\right) - \sum_{t=0}^{T-1}\gamma^t\left(\gamma c z_t^n + \sum_{i=1}^{n}\gamma w_i v z_t^i\right) + \sum_{t=0}^{T-1}\gamma^{t+1}(c + w_i v) D_t\right] \\
&= \mathbb{E}\left[\sum_{t=1}^{T-1}\gamma^t(1-\gamma)\left(c z_t^n + \sum_{i=1}^{n} w_i v z_t^i\right) + \gamma^T\left(c z_T^n + \sum_{i=1}^{n} w_i v z_T^i\right) - \gamma\left(c z_0^n + \sum_{i=1}^{n} w_i v z_0^i\right)\right] \qquad (27) \\
&\quad + \sum_{t=0}^{T-1}\gamma^{t+1}(c + w_i v)\mathbb{E}\left[D_t\right].
\end{aligned}$$

We then add and subtract terms to yield:

$$\gamma^T\left(c z_T^n + \sum_{i=1}^{n} w_i v z_T^i\right) = \gamma^T(1-\gamma)\left(c z_T^n + \sum_{i=1}^{n} w_i v z_T^i\right) + \gamma^T\left(\gamma c z_T^n + \sum_{i=1}^{n}\gamma w_i v z_T^i\right),$$

$$\text{and} \quad -\gamma\left(c z_0^n + \sum_{i=1}^{n} w_i v z_0^i\right) = (1-\gamma)\left(c z_0^n + \sum_{i=1}^{n} w_i v z_0^i\right) - \left(c z_0^n + \sum_{i=1}^{n} w_i v z_0^i\right).$$

Plugging these two equalities back into Equation (27) yields:

$$\begin{aligned}
&\mathbb{E}\left[\sum_{t=0}^{T-1}\gamma^t\left(\gamma c(\hat{z}_t^n - z_t^n) + \sum_{i=1}^{n}\gamma w_i v\left(\hat{z}_t^i - z_t^i\right)\right)\right] \\
&= \mathbb{E}\left[\sum_{t=0}^{T-1}\gamma^t(1-\gamma)\left(c z_t^n + \sum_{i=1}^{n} w_i v z_t^i\right) + \gamma^T\left(\gamma c z_T^n + \sum_{i=1}^{n}\gamma w_i v z_T^i\right)\right] \\
&\quad - \left(c z_0^n + \sum_{i=1}^{n} w_i v z_0^i\right) + \sum_{t=0}^{T-1}\gamma^{t+1}(c + w_i v)\mathbb{E}\left[D_t\right].
\end{aligned} \tag{28}$$

We may now substitute Equations (27) and (28) into Equation (25):

$$\mathbb{E}\left[\sum_{t=0}^{T-1}\gamma^t\left(\gamma c(\hat{z}_t^n - z_t^n) + b\left(z_t^1 - D_t\right)^- + \sum_{i=2}^{n} h_i' v(z_t^i - z_t^{i-1}) + h_1' v(z_t^1 - D_t)^+ \right.\right.$$
$$\left.\left. + \sum_{i=1}^{n}\gamma w_i v\left(\hat{z}_t^i - z_t^i\right)\right) - \gamma^T F_T(\mathbf{z}_T)\right]$$
$$= \mathbb{E}\left[\sum_{t=0}^{T-1}\gamma^t\left(\sum_{i=1}^{n-1}(h_i + (1-\gamma)w_i)v z_t^i + \left(h_n + (1-\gamma)\left(w_n + \frac{c}{v}\right)\right)v z_t^n + (b + h_1' v)(z_t^1)^-\right)\right]$$
$$- \left(c z_0^n + \sum_{i=1}^{n} w_i v z_0^i\right) + \sum_{t=0}^{T-1}\gamma^t\left(\gamma\left(c + \sum_{i=1}^{n} w_i\right)v - h_1' v\right)\mathbb{E}[D_t].$$

Redefining $h_i := h_i + (1-\gamma)w_i$ for $i < n$ and $h_n := h_n + (1-\gamma)\left(w_n + \frac{c}{v}\right)$, we may rewrite this last equality as:

$$\mathbb{E}\left[\sum_{t=0}^{T}\gamma^t\left(\gamma c(\hat{z}_t^n - z_t^n) + b\left(z_t^1\right)^- + \sum_{i=2}^{n} h_i' v(z_t^i - z_t^{i-1}) + h_1' v(z_t^1)^+ + \sum_{i=1}^{n}\gamma w_i v\left(\hat{z}_t^i - z_t^i\right)\right) - \gamma^T F_T(\mathbf{z}_T)\right]$$
$$= \mathbb{E}\left[\sum_{t=0}^{T}\gamma^t H_t(\mathbf{z}_t)\right] - \left(c z_0^n + \sum_{i=1}^{n} w_i v z_0^i\right) + \sum_{t=0}^{T-1}\gamma^t\left(\gamma\left(c + \sum_{i=1}^{n} w_i\right)v - h_1' v\right)\mathbb{E}[D_t].$$

Finally, the last two terms are constants that can be pulled out from the minimization problem, yielding the simplified optimization problem described in Equation (1).