Title
Inventory Models Motivated by Biopharmaceutical Manufacturing

Author
Wang, Yang

Publication Date
2017

Peer reviewed|Thesis/dissertation

Inventory Models Motivated by Biopharmaceutical Manufacturing

By

Yang Wang

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Philip Kaminsky, Chair
Professor Zuo-jun Shen
Professor Satish Rao

Fall 2017

Abstract

Inventory Models Motivated by Biopharmaceutical Manufacturing

by

Yang Wang

Doctor of Philosophy in Industrial Engineering and Operations Research

University of California, Berkeley

Professor Philip Kaminsky, Chair

The biopharmaceutical industry now is entering a more mature stage of its existence. In order to prepare for increasingly fierce competition, firms are beginning to shift their focus from developing new technology to improving manufacturing operations. However, the nature of biomanufacturing poses many unique challenges for the industry. In this thesis, we consider three closely related inventory models to address the challenges these firms are facing.

In the first model, we consider an inventory planning model with batch differentiation. Batch production is a natural constraint in biomanufacturing and the decision of how many batches to produce and how to differentiate the batches is crucial. In this work, we propose a series of heuristic algorithms using the idea of *certainty equivalent control* which converts a difficult stochastic integer programing to a deterministic problem. We show that our heuristic algorithms perform extremely well when the demand variation is small and we also explore how the trade-off between demand information and decision dynamics affects the performance.

In the second model we combine the inventory model with a queuing model to address the tactical level supply chain management and coordination issue between the manufacturer and a third party contractor. Small biopharma firms usually outsource the filling/labeling/packing operations to a third party contractor. As a consequence, the firm experiences large uncertainty in the time until the outsourced order is returned. We show that a *capacity reservation* contract that places an order at fixed intervals with a capacity constraint can outperform the traditional inventory management (r, q) policy by not only reducing the inventory cost at the firm, but also increasing the total profit of the entire system.

In the third model, we extend the traditional process flexibility literature by integrating inventory capabilities and give general guidelines on how to design an effective supply chain network. Given the expensive inventory storage capability in the biopharmaceutical industry, firms are often interested in incorporating *flexible* plants that can produce multiple types of product into an existing dedicated supply chain to better respond to demand uncertainty. We develop models that help the firm explore many important strategic issues, such as whether a new plant should be flexible or not, and whether and which plants should be modified to have inventory capabilities.

Contents

[1]This is a collaborated work with Professor Stefanus Jasin and Professor Hyun-Soo Ahn from Ross School of Businnes, University of Michigan. This work has been accepted for publication in *Operations Research*.

List of Figures

List of Tables

Acknowledgment

I would like to first acknowledge my advisor, Professor Philip Kaminsky, who has consistently supported me during my studies over the past five years. I would like to thank my co-authors, Professors Stefanus Jasin and Hyun-soo Ahn, for their contributions. I would also like to thank my committee members, Professor Zuo-jun Shen and Professor Satish Rao for their advice and suggestions.

I also want to express my deep appreciation to my wife, Mingge Wu, who encourages me when I am in a trough, and shares my happiness when I am on a peak. I also thank my best friend Xiao Liu, who always motivates and inspires me. Finally I thank my family, who made my study at Berkeley possible.

Chapter 1

Introduction

Biopharmaceuticals are medicines, including proteins, vaccines, nucleic acids, antibiotics, etc., that are manufactured using biological technology. In contrast to the direct chemical synthesis typically used in traditional pharmaceutical manufacturing, the majority of biopharmaceutical products are derived from live organisms. A typical biopharmaceutical manufacturing (abbreviated biomanufacturing) process consists of four main stages as illustrated in Figure 1.1.

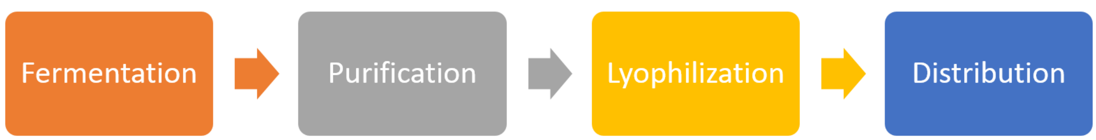

Figure 1.1: A typical biomanufacturing process

In the first fermentation stage, a solution of various materials is cultivated at a cell bank. An important feature of this stage is that it can not be started and stopped freely, which is distinct from the traditional manufacturing. After a warm-up period, the target product can be harvested and purified continuously over time. Then, the target product in intermediate form is lyophilized and stored in batches. These batches later are differentiated to meet final product specifications via filling, labeling and packing in the distribution stage.

The combination of fundamental characteristics in biomanufacturing make operations and supply chain management particularly challenging. We have summarized some of these challenges (see Kaminsky and Wang (2012)[38]) as follows:

- Biomanufacturing is highly variable in both output quality and quantity due to the nature of fermentation and purification processes.

- The biomanufacturing process is in campaign style: once a campaign is started, it will last for several weeks and cannot be stopped or altered and the output is naturally stored in batches. Thus, production planning is challenging and vulnerable to demand variation.

- The biomanufacturing process is highly regulated. Construction of new plants, expansion of existing plants, and improvement of production processes takes a significant amount of time.

- For small biopharmaceutical firms, some manufacturing steps are outsourced to a third party contractor, which makes the inventory management difficult and supply chain coordination important.

- Safe storage of biopharmaceuticals often requires sophisticated and expensive inventory storage facilities, and product shortage may significantly impact a firm's reputation, or even lead to fatalities. Thus, surgically accurate inventory management and demand prediction are desired.

- Research and development takes on average 10-15 years until approval from FDA and similar organizations in other countries, and commercialization of the final product, which leaves few years of effective patent protection. This makes the product selection crucial and long-run capacity planning challenging.

In Kaminsky and Wang (2012)[38] we categorizes these challenges into three levels: (1) strategic decisions, which are directly related to the firm's long-term objectives and performance. Such decisions are usually made by the highest level of management, and once executed, are extremely difficult and costly to change. Typical strategic decisions include new facility construction, existing facility modification, new product development and selection; (2) tactical decisions, which are relatively short term (on the order of months) decisions that support longer term strategic policy, and that define the approaches that firms use to achieve their objectives. Typical tactical decisions range from process and pipeline development to manufacturing planning, supply chain coordination, campaign scheduling, and capacity planning; (3) operational decisions, includes the detailed daily decision-making necessary to operate the biopharmaceutical supply chain. Typical operational decisions include operator assignment, detailed facility schedules, transportation plans, maintenance scheduling and inventory planning. In this thesis, we present three crucial and closely related inventory models to address the aforementioned challenges from each of these levels:

1. **Inventory planning with batch differentiation** (Chapter 2)

   In this chapter we focus on developing an inventory model to aid the firm in making daily operational decisions related to production planning and batch differentiation. The key idea here is to use *certainty equivalent control* which converts a difficult stochastic integer programing into an easy-to-solve deterministic problem. We show that such a scheme with various extensions performs extremely well when the demand variation is small and also explore how the trade-off between demand information and decision dynamics affects the performance.

2. **Outsourcing strategy for intermediate production steps** (Chapter 3)

   In this chapter we address the tactical-level supply chain management and coordination issue for a firm who outsources the last filling/labeling/packing operations to a third party contractor, and as a consequence must deal with uncertainty in the time until the outsourced order is returned. We demonstrate that a *capacity reservation* contract that places a capacity-constrained order at a fixed interval can outperform a traditional (r, q) policy, not only reducing inventory cost at the firm but also increasing the total profit of the entire system.

3. **Process flexibility with inventory** (Chapter 4)

   In this chapter, we concentrate on advising a firm on strategic decisions related to building new plants and modifying existing plants. Given the expensive nature of inventory storage capability in the industry, the firm is interested in incorporating *flexible* plants that can produce multiple types of product into the supply chain in order to better respond to demand uncertainty with limited inventory. We develop models that combine inventory along with flexibility and provide general guidelines on how to design an effective plant-product network in this setting.

Next, we present the details of each of these three topics, and in Chapter 5 we present our closing remarks.

Chapter 2

Inventory Planning with Batch Differentiation[1]

2.1 Introduction

Motivated by an inventory planning problem in a leading "orphan drug" biopharmaceutical firm with which we have worked, we consider the joint production and inventory allocation problem faced by a small biopharmaceutical firm that must differentiate batches of intermediate product into market-specific end products to meet demand over a finite horizon. This firm manufactures and distributes a biopharmaceutical treatment for a rare genetic disorder that leads to severe, potentially life threatening, symptoms in people who have the disorder: a so-called "orphan drug". There is a relatively small set of patients who can benefit from the drug, which results in a limited incentive for firms to invest in drug development. In many countries, the government gives certain tax and patent incentives to encourage firms to develop orphan drugs. However, due to the nature of biomanufacturing process the intermediate product is stored in huge discrete batches. The typical size of one batch is 16,000 vials, while the demand for different markets only ranges from several hundreds to thousands vials. In order to distribute the product to the market, the intermediate product must be packaged into different-sized vials, properly labeled, and packed into cartons for each country. Since the firm outsources its final filling/labeling/packing operations to a third party contractor (we'll address outsourcing strategy in Chapter 3), an entire batch must be filled and labeled at once. In addition, once a vial is labeled for a particular market, it is illegal to re-labeled it for another market. Therefore it is crucial for the firm to make the correct decision on (1) how many batches to manufacture, and (2) how to differential batches to satisfy uncertain demand in different markets. Failing to do so results in either large inventory holding cost or possibly in patient fatalities. The problem setting is illustrated in Figure 2.1.

[1]This is a collaborated work with Professor Stefanus Jasin and Professor Hyun-Soo Ahn from Ross School of Businnes, University of Michigan. This work has been accepted for publication in *Operations Research*.

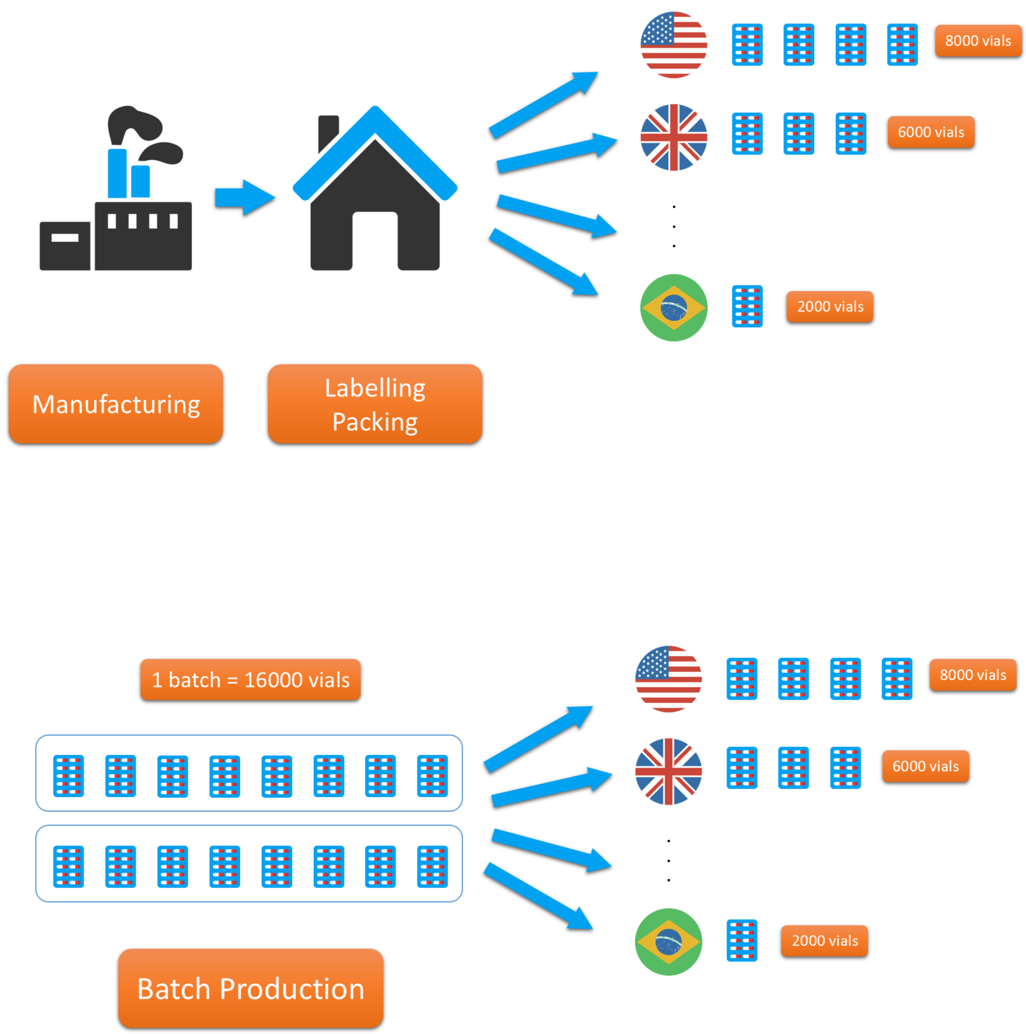

Figure 2.1: Batch production and differentiation

2.2 Literature Review

Our problem is essentially a centralized batch production and differentiation problem faced by a firm that must coordinate production of batches and allocations of batches to end products for different markets. Alternatively, it can be viewed as a multi-retailer inventory problem where, in each period, total orders from all retailers must be a multiple of an exogenously determined batch size. Various models related to batch production/ordering have been studied in the literature. Veinott (1965)[63] studies the problem of batch ordering for a single retailer. For the backorder case, he shows that a (R, nQ) policy, in which the inventory level is raised to at least R by ordering the smallest multiple of Q whenever it falls below R, is optimal for both the finite and infinite horizon problems. Axsäter (1993, 1995, 1998, 2000)[2, 3, 4, 5], Forsberg (1997)[24], and Cachon (2001)[11] study the batch ordering problem for two-echelon distribution systems with one warehouse and N retailers. The papers by Axsäter assume that retailers and warehouse use a continuous review (Q, R) policy. Axsäter (1993) assumes N identical retailers and proposes an exact and approximate methods to evaluate and optimize the performance of the system. Axsäter (1998) extends these results to 2 non-identical retailers, Forsberg (1997) extends these results to N non-identical retailers with Poisson demand, and Axsäter (1995, 2000) consider compound Poisson demand. In all of these continuous time models, there is no allocation issue. All demands are filled using a first-come-first serve approach. Cachon (2001) relaxes the Poisson demand assumption and models the periodic review version of the system, adopting a (R, nQ) policy with random allocation at the warehouse. He characterizes the optimal re-order point at each retailer given a re-order point at the warehouse, and then searches for the optimal warehouse re-order point. Chen (2000)[15] considers the batch production problem in a multi-echelon serial system (N stages) under periodic review and finds that a modified version of Veinott's (1965) (R, nQ) policy is optimal, and Chao (2009)[14] extends this model to allow fixed replenishment intervals (e.g., stage 1 can order new products every day, stage 2 can order new products every week, etc.). The authors show that the system achieves the minimum expected average cost when the ordering times for all of the stages are synchronized. In all of these work, however, the batch ordering restriction is imposed on each individual retailer or stage. In contrast, in our problem, the batch ordering restriction is placed on the total orders from all retailers. The manufacturer utilizes this information to decide how many batches to produce.

Our problem is also related to the discrete time multi-retailer inventory model under limited resources, in which the allocation issue is explicitly studied: Given a scarce resource such as production capacity shared by multiple retailers, each of whom sells unique products, the decision maker must decide how to allocate production capacity in each period. DeCroix and Arreola-Risa (1998)[20] characterize the optimal policy for homogeneous products and develop heuristic controls for the non-homogeneous case. Shaoxiang (2004)[53] extends these results to two non-homogeneous products, and Janakiraman et al. (2009)[33] further extend these results to more than two products and develop an asymptotically optimal control. In all these works, there is a single capacity constraint on the resource; in contrast, the constraint in our setting comes from the nature of batch production.

In contrast to most of the existing related literature, which primarily study the structure of the optimal ordering/allocation control, our primary objective in this chapter is to explore the performance of a simple yet commonly used heuristic control, and analyze different approaches for improving its performance. (For a special case, we also derive the structure of optimal allocation control.) First, we provide a theoretical performance bound for *Deterministic Control* (DC) relative to the optimal control in our setting. In the DC approach, all random variables (i.e., random demands) are replaced by their expected values, and the resulting deterministic optimization is solved to determine the operating control. In other words, the original stochastic dynamic problem is transformed into a deterministic optimization problem. Due to the challenges in estimating complete demand distribution in practice, as well as the challenges in solving the original stochastic dynamic problem, DC and its variants have became popular approaches for solving industrial scale inventory problem (Treharne and Sox (2002)[61], Calmon (2015)[12]). Indeed, based on our own limited interactions with industry practitioners, many inventory managers only use a combination of simple ordering policy (e.g., base-stock policy) and a form of spread-sheet optimization. The later is essentially a deterministic optimization and, more often than not, only utilizes a one-point approximation of demand.

Despite the prevalence of DC control, however, we are not aware of a rigorous analysis of the theoretical performance of DC in the inventory literature (see Sections 2.4 and 2.5 for more discussions). There is a deep literature focusing on deterministic inventory models, but this line of works tends to focus on solution approaches for these deterministic models and for the most part, the quality of a deterministic model as an approximation for the related stochastic model is not rigorously addressed. In this chapter, we show for our model that although DC can perform very poorly if the planning horizon is very long, under some technical conditions (see Theorems 2.5.2 and 2.5.3), periodic re-optimizations of DC (i.e., the so-called *Certainty Equivalent Control* (CEC)) improve the performance of DC by dampening the impact of the length of planning horizon on total costs. If, however, the size of demand variation is also relatively large, then re-optimizations only have a limited benefit and we need to apply more sophisticated heuristics control. These results shed light on the appropriateness of either DC or CEC approach in practice.

To address the limitation of CEC, we analyze two improvements of CEC: (1) *Open-Loop Feedback Control* (OLFC) and (2) *Multi-Point Approximation Control* (MPAC). Note that, to reduce the computational burden, in the CEC approach we deliberately de-emphasize two key elements of the original stochastic problem: (1) demand variation (because we ignore demand distribution information except for its expected value) and (2) decision dynamics (because we ignore the fact that future decisions will be contingent on current decisions, demand realizations, and system dynamics). Both OLFC and MPAC highlight the potential improvement of more explicitly considering these two elements. Unlike CEC, which only uses expected demand information, OLFC uses the complete demand distribution information for calculating batch production and product allocation. However, it only partially models decision dynamics in an ad-hoc manner via frequent re-optimizations. MPAC, on the other hand, explicitly models decision dynamics with a Dynamic Program (DP) but only uses partial demand distribution information via a multi-point approximation of the distribution. (See Table 2.1 for a comparison of all approaches analyzed in this chapter.)

We emphasize that both OLFC and MPAC do utilize more demand distribution information beyond the expected values used by DC and CEC. The crucial and practically relevant question that we attempt to answer is this: When designing a heuristic control for an inventory problem, is it more important to major in demand information (e.g., by exerting a lot of efforts in estimating the whole demand distribution curve) and minor in decision dynamics, or the other way around? That is, can we compensate the use of a less granular demand distribution approximation with a more sophisticated decision dynamic, and vice versa? We show that, for our model, although OLFC improves on the performance of CEC, the magnitude of this improvement can be limited, especially when the size of demand variation is large. Indeed, we show that a solution of OLFC can be identical to a solution of CEC in some cases. This suggests that, in our setting, the benefit of using a complete demand distribution information is already captured, at least partially, by periodic re-optimizations of a deterministic model—which is surprising, even if it only holds for some cases. In contrast to OLFC, with a carefully chosen demand approximation in each period, MPAC exhibits a stronger theoretical performance bound even with only a slightly more granular demand distribution approximation than expected value. This highlights the importance in our setting of effectively modeling decision dynamics–incorporating the complete demand distribution without a detailed enough model of decision dynamics is insufficient.

Since optimally solving a large-scale stochastic inventory problem is often intractable in practice, developing computationally efficient heuristic controls with analytic performance bounds becomes important. In addition to providing a guidance for constructing effective heuristic controls for the specific model that we are considering, the approach taken in this chapter, i.e., comparing the two heuristic controls— OLFC (capturing the fidelity of demand distribution information but simplifying the decision dynamics) and MPAC (capturing the decision dynamics but simplifying the demand distribution)— could potentially be applied to other inventory problems to lend insight into factors that drive heuristic performance, which can ultimately lead to more effective heuristic controls.

In Section 2.3, we formulate our model; in Sections 2.4 and 2.5, we introduce DC and CEC, and analyze their performances; in Section 2.6, we discuss OLFC and MPAC; in Section 2.7, we present results of our computational experiments; and finally, in Section 2.8, we conclude the chapter.

2.3 The Stochastic Planning Problem

We consider a discrete-time model where a firm (a centralized decision maker) must satisfy demands in multiple markets (which we call retailers) through joint production and allocation decisions. In each period, the sum of allocated units across all end-product markets must equal the number of units in the batches produced in that period. Specifically, we consider a model with T discrete periods and m retailers, where the time periods are indexed by $t \in \{1, \dots, T\}$ and the retailers are indexed by $i \in \{1, \dots, m\}$. Demands across different periods are assumed to be independent and stationary. We assume that lead time is zero and unsatisfied demands are backordered. In general, our basic solution approach is straightforward to extend to deterministic lead times.) The following notations are used throughout the chapter:

Notations:

$D_{t,i}$	Demand faced by retailer i in period t.
$F_i(\cdot)$	Cumulative demand distribution for retailer i.
μ_i	Expected demand in a period faced by retailer i.
$\Delta_{t,i}$	$D_{t,i} - \mu_i$
h_i	Per unit holding cost for retailer i.
p_i	Per unit penalty cost for retailer i.
c	Per unit production cost
B	Batch size (i.e., the number of units in a batch).
$I^{\pi}_{t,i}$	Starting inventory at retailer i at the beginning of period t under policy π.
$I_{1,i}$	Starting inventory at retailer i at the beginning of period 1.
$N^{\pi}_{t,i}$	Number of new units allocated to retailer i in period t under policy π.
Z^{π}_t	Number of new batches produced in period t under policy π.
C^{π}	Total costs under policy π.

Note that $\Delta_{t,i} = D_{t,i} - \mu_i$ is the difference between the actual and expected demand faced by retailer i in period t. Since total allocated units across all retailers must equal total units contained in the new batches, we must have: $\sum_{i=1}^{m} N^{\pi}_{t,i} = Z^{\pi}_t B$. For analytical tractability, although we require Z^{π}_t to be a non-negative integer for all t and i, we allow $N^{\pi}_{t,i}$ to be a non-negative real number. (In the context of biopharmaceutical problem that we consider, the magnitude of demand in each retailer (i.e., country) is typically on the order of hundreds to thousands. This justifies our relaxation of $N^{\pi}_{t,i}$ since a simple rounding error has a relatively negligible impact on the overall system performance.) Under the backorder assumption, the starting inventory level at retailer i at the beginning period $t+1$ under policy π is given by:

$$I^{\pi}_{t+1,i} = I^{\pi}_{t,i} + N^{\pi}_{t,i} - D_{t,i} = I_{1,i} + \sum_{s=1}^{t} N^{\pi}_{s,i} - \sum_{s=1}^{t} D_{s,i}. \tag{2.1}$$

At the beginning of period t, upon observing the available inventories at all retailers, the firm first decides how many new batches to produce. After producing the new batches, it must decide how many units of end product to allocate to each retailer (i.e., how many units to label and package for each market). Demands are then realized and inventories are consumed. Remaining units are held in inventory until the next period and unsatisfied demands are backordered. Let Π denote the set of non-anticipating controls, i.e., the set of controls that determine how many new batches to produce and how many units to allocate to each retailer in period t using only the accumulated information up to the beginning of period t. Let C^* denote the expected total costs under an optimal control $\pi^* \in \Pi$. We can write C^* as follows:

$$
\begin{aligned}
C^* &= \inf_{\pi\in\Pi} \sum_{t=1}^{T} \mathbf{E}\left[cZ_t^\pi B + \sum_{i=1}^{m} p_i(D_{t,i} - N_{t,i}^\pi - I_{t,i}^\pi)^+ + \sum_{i=1}^{m} h_i(N_{t,i}^\pi + I_{t,i}^\pi - D_{t,i})^+\right] && (2.2)\\
&\text{s.t.} \sum_{i=1}^{m} N_{t,i}^\pi = Z_t^\pi B,\ Z_t^\pi \in \mathbb{Z}^+,\ N_{t,i}^\pi \in \mathbb{R}^+ \qquad \forall t,\ i && (2.3)\\
&\quad I_{t+1,i}^\pi = I_{t,i}^\pi + N_{t,i}^\pi - D_{t,i} \qquad \forall t,\ i && (2.4)\\
&\quad I_{1,i}^\pi = I_{1,i} \qquad \forall i && (2.5)
\end{aligned}
$$

where all the constraints must be satisfied almost surely. Let $I_t = (I_{t,i})$, $N_t = (N_{t,i})$, and $D_t = (D_{t,i})$ denote the vector of starting inventory levels, allocated units, and realized demands in period t, respectively. We can write the optimal control problem (2.2) using Bellman's equation as follows:

$$
\begin{aligned}
&C_t(I_t) = \min_{Z_t\in\mathbb{Z}^+, N_t\in\Omega(Z_t)} \left\{cZ_tB + G(I_t+N_t) + \mathbf{E}\left[C_{t+1}(I_t+N_t-D_t)\right]\right\} \quad \text{for } t = 1,2,...,T\\
&C_{T+1}(I_{T+1}) = 0, && (2.6)
\end{aligned}
$$

where
$\Omega(Z) = \{N : N_i \in \mathbb{R}^+, \sum_{i=1}^m N_i = ZB\}$ and $G(y) = \mathbf{E}\left[\sum_i^m h_i(y-D)^+ + \sum_{i=1}^m p_i(D-y)^+\right]$.

In general, the joint production and allocation problem formulated in (2.6) is difficult to solve and its optimal control is challenging to characterize. As might be expected given the batch production requirements, the optimal expected cost $C_t(I_t)$ is not convex in starting inventory levels, so a simple base-stock style control is not likely to be optimal for this problem. It is also not difficult to find examples where the optimal decision as a function of inventory levels changes depending on the period for a given planning horizon. To calculate the optimal control, it is therefore necessary to explicitly solve the entire dynamic programming (2.6), which is generally intractable due to the problem size. As an illustration, if demand is discrete and integral, solving the full DP requires an exponential amount of space $O\left((I_{max} - I_{min})^m\right)$ to store the state information where I_{max} (I_{min}) is the maximum (minimum) possible inventory level. In addition, since the problem is not convex, to ensure global optimality, there is also an exponential number of decisions $O((Z_{max}B)^m)$ that need to be explored.

2.4 Deterministic Control

In this section, we analyze the performance of the simple non-adaptive heuristic control we introduced above—*Deterministic Control* (DC)—in which all random demand variables are replaced by their expected values and the resulting deterministic optimization problem is solved (Treharne and Sox (2002)[61]). Although not always known by that name (indeed, it is sometimes naively employed by managers without any name at all), DC is popular in practice because it addresses two complicating problems that arise when solving the original problem: (1) estimating the whole demand distribution curve is often challenging (for instance, the firm that motivates this project uses a one-point estimate of demand instead of the estimate of complete demand distribution— an approach which, in our experience, is common); (2) the optimal control problem, even if demand distribution can be fully estimated, is difficult to solve (for a typical industrial-scale problem, even solving a deterministic version of the problem is already quite challenging). Naturally, these concerns have often motivated practitioners to use a heuristic control that can be implemented with as little detailed demand information as possible.

Given the fact that a DC-like approach is widely used in practice, an interesting set of questions arises. In many settings, a deterministic model is a poor approximation of a stochastic system and its solution is typically suboptimal when applied to its stochastic counterpart. However, in a few canonical settings such as *Linear Quadratic Gaussian Control* problem (see Stengel (1994)[58]), it is known that DC performs quite well and can even be optimal. The important question for our multi-product inventory problem is: How much is actually lost if a heuristic control derived from a deterministic model such as DC is applied in a stochastic setting? Is there any setting in which a deterministic model is a good approximation of a stochastic system? If so, in what sense?

Our results in this section show that DC performs reasonably well when the magnitude of demand variation is relatively small relative to its mean (i.e., the coefficient of variation is small) and the length of planning horizon is relatively short. We characterize the performance of this deterministic heuristic and show that the performance of DC deteriorates at the rate of $T^{3/2}$ as the length of planning horizon increases.

To evaluate the performance of DC, we first define C^D as follows:

$$C^D = \min_{z,n} \sum_{t=1}^{T} \left[c\, z_t\, B + \sum_{i=1}^{m} p_i\, (\mu_i - x_{t,i} - n_{t,i})^+ + \sum_{i=1}^{m} h_i (x_{t,i} + n_{t,i} - \mu_i)^+ \right] \tag{2.7}$$

$$\text{s.t.} \sum_{i=1}^{m} n_{t,i} = z_t\, B,\ z_t \in \mathbb{Z}^+,\ n_{t,i} \in \mathbb{R}^+ \qquad \forall t,\ i \tag{2.8}$$

$$x_{t+1,i} = x_{t,i} + n_{t,i} - \mu_i \qquad \forall t,\ i \tag{2.9}$$

$$x_{1,i} = I_{1,i} \qquad \forall i \tag{2.10}$$

where the variables $\{x_{t,i}\}$ can be interpreted as the starting inventory levels for each retailer at the beginning of each period in a deterministic system with demands equal their expected values and under allocation $\{n_{t,i}\}$.

Note that (2.7) can be written as a mixed integer linear program (MILP). We first explore the relationship between C^D and C^*. In much of the DC literature (e.g., Jasin and Kumar

(2012)[35], Ciocan and Farias (2012)[19]), the optimal value of the deterministic problem often serves as either a lower or upper bound for the optimal value of the original stochastic control problem. This allows the optimal value of the deterministic problem to be used as a proxy for performance analysis of any feasible heuristic control. Unfortunately, this is not the case here due to the integrality of z_t, i.e., $C^D \not\leq C^*$. To see why, note that a standard argument for proving either $C^D \leq C^*$ or $C^D \geq C^*$, depending on the context, is to apply Jensen's inequality and replace all random variables with their expected values. Since we still require the number of batches to be integral, this strategy does not work. It is possible to further relax this assumption and allows z_t to be a real number, so we immediately get $C^D \leq C^*$. However, the resulting lower bound is too loose to be useful for performance benchmarking.

To describe DC, we let $z = z^D$ and $n = n^D$ denote an optimal solution to (2.7).

Deterministic Control - DC

---

1. At the beginning of period 1, solve C^D
2. For $t = 1, 2, ..., T$, do:
   (a) At the beginning of period t, produce exactly z_t^D new batches
   (b) After the new batches arrive, allocate exactly $n_{t,i}^D$ units to retailer i

---

As we define it, DC precludes any consideration of the starting inventory, or backorder, levels prior to making the production and allocation decisions in each period because both decisions are directly dictated by z^D and n^D, regardless of the actual demand realizations. We define $\mathbf{E}[C^{DC}]$ to be the expected total costs associated with implementing DC. The following result provides a bound for the regret introduced by implementing DC:

Theorem 2.4.1. *Let* $\sigma = \max_i \mathbf{E}[(D_{1,i} - \mu_i)^2]^{1/2}$. *Then,*

$$\mathbf{E}[C^{DC}] - C^* \;\leq\; 2\sigma(T+1)^{3/2} \left[\sum_{i=1}^{m}(p_i + h_i)\right].$$

The proof of Theorem 2.4.1 can be found in Appendix A.1.1. Two comments are in order. First, due to the non-differentiability of $(\cdot)^+ = \max(\cdot, 0)$ in (2.7), the optimal solution z^D and n^D may not be unique. However, the bound in Theorem 2.4.1 holds regardless of the choice of optimal solution. Second, the performance of DC is proportional to the size of demand variability as measured by σ. If $\sigma = 0$, then DC is optimal regardless of the length of planning horizon T. When $\sigma > 0$, however, the bound in Theorem 2.4.1 depends not only on σ but also on $T^{3/2}$. This scaling factor is not an artifact of the proof; our simulation results in Section 2.7 do show that the relative regret of DC quickly becomes worse as T

increases. Indeed, this is the reason why DC may perform very poorly for a multi-period inventory problem even when the magnitude of demand variation is relatively small compared to its mean. This is in contrast to the performance of DC-type heuristic controls in other application areas such as revenue management and dynamic pricing, where the regret scales with $\sqrt{T}$ instead of $T^{3/2}$ (Gallego and van Ryzin (1994)[28], Jasin (2014)[34], Jasin and Kumar (2013)[36]). Intuitively, in revenue management application, the penalty due to lost sales is only applied once to the cumulative sales at the end of period T; in contrast, in the multi-period inventory problem, the penalty due to under/over sales (i.e., holding/backorder cost) is applied in every period. This explains the difference in the performance of DC-type heuristic controls in the two application areas. The following result is a corollary of Theorem 2.4.1.

Corollary 2.4.2. *Suppose that demands are Poisson with $\mu_i = \mu^*$ for all i. Then, there exists a constant $M > 0$ independent of T and μ^* such that,*

$$\frac{\mathbf{E}[C^{DC}] - C^*}{C^*} \;\leq\; M\left(\frac{T}{\mu^*}\right)^{1/2}.$$

The proof of Corollary 2.4.2 can be found in Appendix A.1.2. The bound in Corollary 2.4.2 is proportional to $\sqrt{T/\mu^*}$. Thus, even for the case of Poisson demand, where the coefficient of variation goes to zero as the mean goes to infinity, T must be small relative to μ^* for DC to be reasonably effective. While this may not be an issue for instances with very large μ^*, this result shows that the applicability of DC is, at best, limited. To re-emphasize, in the context of our inventory problem, DC may perform poorly, even for instances with a small coefficient of variation, unless the planning horizon is also short. This paints a rather bleak picture on the usefulness of a deterministic approximation for multi-period stochastic inventory problems in general. The main culprit here is the manner in which randomness accumulates over time due to the per-period holding and penalty costs, which scale polynomially with T. This gives rise to an important question: Is it possible to construct an alternative heuristic control that retains the simplicity of DC and yet is more effective than DC, at least in the setting with a small coefficient of variation and typical industrial planning horizon (which is about 2 to 5 years, i.e., $T = 24$ to 60 if one period equals one month)? It turns out that simple re-optimization of DC at the beginning of every period significantly reduces the dependency of relative regret on T. This makes DC more practically appealing and amenable to problems with a longer planning horizon. We introduce and analyze this approach in the next section.

2.5 Improving Deterministic Control via Re-optimization

We next discuss several modifications to the basic DC that can improve its performance. For analytical tractability, we will focus our attention only on the class of controls that use z^D as the production control, but optimize the allocation control. (One might expect an additional improvement if the production control is also further optimized— this makes the analysis extremely challenging, but we computationally test this approach in Section 2.7.)

For a given production control, except for some special cases, the optimal allocation control is difficult to determine when we have multiple retailers. In the case of homogeneous retailers with identical cost structures and i.i.d demands, we are able to completely characterize an optimal allocation control. We then prove that, as long as the size of demand variation is not too large compared to its mean, a simple re-optimization of DC at the beginning of every period suffices to guarantee a significant improvement over DC. This result gives credence to the practice of re-optimization that is often employed in industry. Moreover, our analysis of re-optimization also suggests a natural inventory-balancing control that can be implemented in real-time. We discuss this at the end of this section.

Recall that z^D is computed at the beginning of the horizon by solving for C^D. Let $\tilde{\Pi}$ denote the set of non-anticipating controls that use z^D as the production control. Also, let J^π denote the total costs under control $\pi \in \tilde{\Pi}$ and J^* denote the total costs under the optimal allocation control for a given production control z^D. We can write:

$$J^* = \inf_{\pi \in \tilde{\Pi}} \sum_{t=1}^{T} \mathbf{E}\left[c z_t^D B + \sum_{i=1}^{m} p_i (D_{t,i} - N_{t,i}^\pi - I_{t,i}^\pi)^+ + \sum_{i=1}^{m} h_i (N_{t,i}^\pi + I_{t,i}^\pi - D_{t,i})^+ \right] \quad (2.11)$$

$$\text{s.t.} \sum_{i=1}^{m} N_{t,i}^\pi = z_t^D B, \; N_{t,i}^\pi \in \mathbb{R}^+ \qquad \forall t, i \quad (2.12)$$

$$I_{t+1,i}^\pi = I_{t,i}^\pi + N_{t,i}^\pi - D_{t,i} \qquad \forall t, i \quad (2.13)$$

$$I_{1,i}^\pi = I_{1,i} \qquad \forall i \quad (2.14)$$

Similar to C^* in Section 2.2, we can write J^* using Bellman's equation as follows:

$$J_t(I_t) = \min_{N_t \in \Omega(z_t^D)} \left\{ c z_t^D B + G(I_t + N_t) + \mathbf{E}\left[J_{t+1}(I_t + N_t - D_t) \right] \right\} \; t = 1, \ldots, T \quad (2.15)$$

where $J_{T+1}(I_{T+1}) = 0$ and $\Omega(\cdot)$ is as defined in Section 2.2.

2.5.1 Optimal Allocation Control

In general, the optimal allocation control that achieves J^* is challenging to compute; doing so requires solving a full DP using backward recursion, which is computationally intractable if T, m, or the support for demand distribution is large. However, if all retailers are homogeneous with identical cost structure (i.e., $h_i = h^*$ and $p_i = p^*$ for all i) and i.i.d. demands (i.e., $D_{t,i} \sim D$ for all t and i), a simple and easy-to-implement allocation control is optimal:

Theorem 2.5.1. *Suppose that $I_1 = 0$, $p_i = p^*$ and $h_i = h^*$ for all i, and demand at all retailers in all periods is i.i.d. Then, the optimal allocation control can be obtained as follows:*

(1) At the beginning of period t, sort all retailers from the smallest inventory level to the largest, $I_{t,1} \leq I_{t,2} \leq \ldots \leq I_{t,m}$. Let $\theta_t = \max\{k \in \mathbb{Z}^+ \mid \sum_{i=1}^{k} i(I_{t,i+1} - I_{t,i}) \leq z_t^D B\}$.

(2) Raise the inventory of retailers with $i \leq \theta_t$ to the same level, i.e.,

$$N_{t,i} + I_{t,i} = \frac{1}{\theta_t}\left(\sum_{i=1}^{\theta_t} I_{t,i} + z_t^D B\right) \qquad \forall i \leq \theta_t.$$

(3) Allocate nothing for all retailers with $i > \theta_t$ by setting $N_{t,i} = 0$, $\forall\, i > \theta_t$.

The proof of Theorem 2.5.1 can be found in Appendix A.1.3. In other words, when all retailers are homogeneous, the optimal allocation control is to balance the inventory levels in as many retailers as possible by allocating the new units starting with a retailer with the lowest inventory level. If the retailers are not homogeneous (either in demand distribution or cost parameters), this allocation control is no longer optimal. However, all is not lost. In Section 2.5.3, we will show that an inventory balancing control similar to the one described in Theorem 2.5.1 is near-optimal in the non-homogeneous setting.

2.5.2 Certainty Equivalent Control

Motivated by our discussion in Section 4.1, we now consider a simple heuristic control based on re-optimizing the deterministic counterpart of J^*, which we call *Certainty Equivalent Control* (CEC). Re-optimizations have been shown to significantly improve the performance of DC-type heuristic controls in many application areas (Jasin (2014)[34], Jasin and Kumar (2012)[35], Reiman and Wang (2008)[49]. and Ciocan and Farias (2012)[19]). In the context of assemble-to-order system, Plambeck and Ward (2006)[46] and Dogru et al. (2010)[22] propose adaptive controls that utilize some forms of re-optimization. However, their results do not carry over to our setting for at least two reasons: (1) In the assembly-to-order setting considered in these papers, the firm first observes demand before making a decision while, in our setting, the firm first makes a decision before observing demand; (2) in this assemble-to-order setting, the firm can make continuous adjustments while, in our setting, the firm is limited to making adjustments at the beginning of each period. In a standard inventory control setting, Secomandi (2008)[52] analyzes the impact of re-optimization on performance. He shows that re-optimization does not always improve the original solution and provides sufficient conditions for re-optimization to guarantee a better result. However, he does not provide a theoretical performance bound for his approach. The lack of existing results in the literature is quite surprising given the practicality and prevalence of re-optimization-based heuristic controls in industry. In fact, many companies with which have interacted employ a form of rolling horizon approach that periodically re-optimizes their planning models. Our results in this subsection contribute to the literature by characterizing the benefit of re-optimizations on model performance.

Define $J_t^D(I_t)$ as follows:

$$J_t^D(I_t) = \min_n \sum_{s=t}^{T} \left[cz_s^D B + \sum_{i=1}^{m} p_i(\mu_i - x_{s,i} - n_{s,i})^+ + \sum_{i=1}^{m} h_i(x_{s,i} + n_{s,i} - \mu_i)^+ \right] \quad (2.16)$$

$$\text{s.t.} \sum_{i=1}^{m} n_{s,i} = z_s^D B, \; n_{s,i} \in \mathbb{R}^+ \qquad \forall s, \, i \quad (2.17)$$

$$x_{s+1,i} = x_{s,i} + n_{s,i} - \mu_i \qquad \forall s, \, i \quad (2.18)$$

$$x_{t,i} = I_{t,i} \qquad \forall i \quad (2.19)$$

where, as before, the variables $\{x_{t,i}\}$ can be interpreted as the starting inventory levels for each retailer at the beginning of each period in a deterministic system where demands equal their expected values and under allocation $\{n_{t,i}\}$. Let $n^{*t} = (n_{s,i}^{*t})_{s \geq t, i \geq 1}$ denote an optimal solution of $J_t^D(I_t)$. (Note that n^{*t} is a function of I_t. We suppress its dependency on I_t for notational brevity.) The complete description of CEC is given below.

Certainty Equivalent Control - CEC

---

1. At the beginning of period 1, solve C^D

2. For $t = 1, 2, ..., T$, do:

   (a) At the beginning of period t, produce exactly z^D new batches

   (b) Solve $J_t^D(I_t)$ and allocate exactly $n_{t,i}^{*t}$ units to retailer i

   (c) Update $I_{t+1} = I_t + n_t^{*t} - D_t$

In contrast to DC, which is implemented independent of demand realizations, CEC incorporates realized demands and updated inventory/backorder level by re-optimizing the deterministic allocation problem at the beginning of every period. We now examine whether re-optimization is sufficient to significantly improve the performance of DC.

The impact of re-optimization for a problem with general holding and penalty costs is difficult to analyze, primarily due to the non-differentiability of the function $(\cdot)^+ = \max(\cdot, 0)$. Moreover, the optimal solution of $J_t^D(I_t)$ may not be unique. This makes the task of analyzing the evolution of the re-optimized solution analytically intractable (see also Remark 2.5.1 at the end of this subsection). However, we show that it is possible to theoretically characterize the benefit of re-optimization under a particular sequence of optimal solutions $n_1^{*1}, n_2^{*2}, ..., n_T^{*T}$ when either these solutions satisfy a certain condition (see Theorem 2.5.2) or all retailers are homogeneous with identical cost structure and i.i.d demands (see Theorem 2.5.3). In Section 4.3, we will argue that CEC can in fact be interpreted as a form of inventory-balancing control. This observation is useful and can be used to motivate the development of optimal inventory-balancing controls in other inventory control problems. Let $D_{1:T}$ denote the vector of all realized demands in T periods. Define the hindsight total costs $J^H(D_{1:T})$ as follows:

$$J^H = \min_n \sum_{t=1}^{T} \left[c z_t^D B + \sum_{i=1}^{m} p_i (D_{t,i} - x_{t,i} - n_{t,i})^+ + \sum_{i=1}^{m} h_i (x_{t,i} + n_{t,i} - D_{t,i})^+ \right] \quad (2.20)$$

$$\text{s.t.} \sum_{i=1}^{m} n_{t,i} = z_t^D B, \; n_{t,i} \in \mathbb{R}^+ \qquad \forall t, \; i \quad (2.21)$$

$$x_{t+1,i} = x_{t,i} + n_{t,i} - D_{t,i} \qquad \forall t, \; i \quad (2.22)$$

$$x_{1,i} = I_{1,i} \qquad \forall i \quad (2.23)$$

J^H is thus the optimal total cost incurred if the firm has perfect knowledge of all future demands. Since we obviously cannot do better than the perfect hindsight policy, we immediately have $\mathbf{E}[J^H] \leq J^*$. The result below gives us a sense of the level of improvement that may result from periodic re-optimizations.

Theorem 2.5.2. *Let $\sigma = \max_i \mathbf{E}[(D_{1,i} - \mu_i)^2]^{1/2}$. Suppose that $I_1 = 0$ and there exists an optimal solution n^{*1} and a constant $\varphi > 0$ such that $n_{t,i}^{*1} \geq \varphi$ and $\left|\sum_{s=1}^{t} n_{s,i}^{*1} - t\mu_i\right| \geq t\varphi$ for all t and i. Let J^{DC} and J^{CEC} be the expected total costs under DC and CEC, respectively, and define $\mathcal{A} = \{\sum_{i=1}^{m} |\Delta_{t,i}| < \varphi/2, \;\; \forall\, t\}$. Then,*

$$\mathbf{E}[(J^{DC} - J^H)\,\mathbf{1}\{\mathcal{A}\}] \;\leq\; 2\sigma(T+1)^{3/2} \left[\sum_{i=1}^{m} (p_i + h_i) \right].$$

Moreover, there exists a sequence of optimal solutions $n_2^{*2}, ..., n_T^{*T}$ such that

$$\mathbf{E}[(J^{CEC} - J^H)\,\mathbf{1}\{\mathcal{A}\}] \;\leq\; 2\sigma T \left[\sum_{i=1}^{m} (p_i + h_i) \right].$$

The proof of Theorem 2.5.2 can be found in Appendix A.1.4. Note that, under DC, we simply apply allocation control n_t^D during period t as in Theorem 1. Thus, the fact that the bound for DC is of order $T^{3/2}\sigma$ is not surprising. In contrast, the bound for CEC is only of order $T\sigma$, which means that re-optimizations improve the performance guarantee of DC by reducing the effect of the length of planning horizon on regret from $T^{3/2}$ to T, at least in the set $\mathcal{A}$ where the size of total demand variation during each period is relatively small compared to the number of allocated units. If σ is small compared to φ, then $\mathcal{A}$ happens with high probability. In this case, we can properly say that periodic re-optimization improves the performance of DC with a high probability. The conditions $n^{*1}_{t,i} \geq \varphi$ and $\left|\sum_{s=1}^{t} n^{*1}_{s,i} - t\mu_i\right| \geq t\varphi$ in Theorem 2.5.2 simply mean that in a deterministic world, we always allocate a positive number of products to each retailer at every period, and that the starting inventory level at each retailer at the beginning of period $t > 1$ (i.e., $\sum_{s=1}^{t-1} n^{*1}_{s,i} - (t-1)\mu_i$) is always either strictly positive or strictly negative. These conditions are not as strong as they appear; they can be easily satisfied, for example, in the case where there is a mismatch between the number of units per batch (or the number of units in an integer multiple of a batch) and the total expected demands at all retailers. (To illustrate, suppose that we have two identical retailers with $p_1 = p_2$, $h_1 = h_2$, $\mu_1 = \mu_2 = 980$, $B = 2000$ and $T = 10$. For simplicity, suppose that $c = 0$ (or any sufficiently small number). If the penalty cost for lost sales is sufficiently large, which is uncommon in the context of our biopharmaceutical problem, it is not difficult to show that it is optimal to produce exactly one batch in every period. Moreover, an optimal allocation is given by $n^{*1}_{t,1} = n^{*1}_{t,2} = 1000$ for all t, which yields the starting inventory levels $I_{t,1} = I_{t,2} = 20 \cdot (t-1)$ for all $t \leq 10$.) Finally, note that the probability of event $\mathcal{A}$ is a function of T. Without further assumptions on the cost structure and demand distribution, it is not immediately clear from Theorem 2.5.2 alone how long the planning horizon can be before the benefit of re-optimizations start to diminish due to the cumulative impact of demand randomness. Per our discussions in Section 2.4, T must be much smaller than μ^* for DC to perform sufficiently well. The next result shows that CEC clearly outperforms DC for a wide range of T values.

Theorem 2.5.3. *Suppose that the following conditions hold: $I_1 = 0$, $p_i = p^*$ and $h_i = h^*$ for all i, demands are i.i.d with mean μ^* and standard deviation σ^*, and $z^D_t > 0$ for all t. Define $\hat{\mathcal{A}} := \{\sum_{i=1}^{m} |\Delta_{t,i}| \leq B/(2m)\ \forall\, t\}$. There exists a sequence of optimal solutions $n^{*1}_1, n^{*2}_2, ..., n^{*T}_T$ such that*

$$\begin{aligned}\mathbf{E}[(J^{DC} - J^{H})\mathbf{1}\{\hat{\mathcal{A}}\}] &\leq 2m(p^* + h^*)(T+1)^{3/2}\sigma^* \quad \textit{and} \\ \mathbf{E}[(J^{CEC} - J^{H})\mathbf{1}\{\hat{\mathcal{A}}\}] &\leq 2m(p^* + h^*)T\sigma^*.\end{aligned}$$

Moreover, if demands are Poisson, $B \geq \sqrt{\mu^*}$, and $T = o(e^{\min\{\mu^*/5,\, B^2/(80m^4\mu^*)\}})$, there exists a constant $M > 0$ independent of T, B and μ^* such that, for all large B^2/μ^* and μ^*,

$$\frac{\mathbf{E}[J^{DC}] - J^*}{J^*} \leq \frac{M\sqrt{T}}{\sqrt{\mu^*} + \sqrt{T}} \quad \textit{and} \quad \frac{\mathbf{E}[J^{CEC}] - J^*}{J^*} \leq \frac{M}{\sqrt{\mu^*} + \sqrt{T}}.$$

The proof of Theorem 2.5.3 can be found in Appendix A.1.5. Some comments are in order. First, the setting in Theorem 2.5.3 is not a special case of the setting in Theorem 2.5.2. (In Theorem 2.5.3, we do not require $|\sum_{s=1}^{t} n^{*1}_{s,i} - t\mu_i| \geq t\varphi$ for some $\varphi > 0$ for all t and i; in fact, it is possible that $|\sum_{s=1}^{t} n^{*1}_{s,i} - t\mu_i| = 0$ for all t and i. Mathematically, we do not need this condition because we assume that all retailers are homogeneous and demands are i.i.d.) Thus, the result of Theorem 2.5.3 cannot be seen as a corollary of

Theorem 2.5.2. Second, the bound for DC in Theorem 2.5.3 is similar to the bound for DC in Corollary 2.4.2. Although the bound holds for $T = o(e^{\min\{\mu^*/5,\, B^2/(80m^4\mu^*)\}})$ (i.e., T can be very large, especially when B is proportional to μ^*, as in the case of our biopharmaceutical problem), T must be much smaller than μ^* to guarantee the effectiveness of DC. In contrast, the bound for CEC is almost independent of T— as long as $T = o(e^{\min\{\mu^*/5,\, B^2/(80m^4\mu^*)\}})$, which can be much larger than μ^*, the relative regret of CEC decreases to 0 at a rate that is (roughly speaking) proportional to $1/\sqrt{\mu^*}$ as $\mu^* \to \infty$. In contrast to DC, in the case of CEC, the additional $\sqrt{T}$ in the bound also helps speed up the convergence. However, since a typical planning horizon extends about 2 to 5 years, if one period is one month (i.e., $T = 24$ to 60) and μ^* is, at least, on the order of hundreds or thousands, the greatest impact on performance comes from $\sqrt{\mu^*}$ instead of $\sqrt{T}$. Practically, this means that re-optimization not only yields a stronger performance guarantee, but also allows for a much longer planning horizon. If T is larger than $o(e^{\min\{\mu^*/5,\, B^2/(80m^4\mu^*)\}})$, however, it becomes necessary to also re-optimize the production decision z_t in addition to the allocation decisions $\{n_{t,i}\}$ for good performance. This alternative re-optimization control, which essentially re-optimizes the whole integer program instead of a linear program, can also be used to address the case where demand variation is relatively large compared to its mean (in contrast to Poisson demand in Theorem 2.5.3). We computationally test the performance of this approach in Section 6.

Finally, the relative bounds in Theorem 2.5.3 hold for sufficiently large batch sizes. In particular, B must have at least the same order of magnitude as the standard deviation of Poisson demand for one retailer, $\sqrt{\mu^*}$. To see why this is necessary in general, note that if $B << \sqrt{\mu^*}$ and $z_t^D = 1$ for all t, then $\hat{\mathcal{A}}$ is unlikely to occur, and the benefit of re-optimizations are limited to a small set of demand realizations. In the context of our biopharmaceutical problem, the conditions of large B^2/μ^* and μ^* are easily satisfied since B is typically much larger than μ^* (i.e., the size of a batch is much larger than the expected demand at any one retailer) and μ^* ranges from about hundreds to thousands.

Remark 2.5.1. (On the Non-uniqueness of the Optimal Solution). Although we only prove the results in Theorems 2.5.2 and 2.5.3 for a particular choice of optimal solution, we conjecture that the non-uniqueness of the optimal solution is not detrimental to the performance of this approach (as in the context of revenue management; see Jasin and Kumar (2013)[36]). Indeed, if we use a differentiable convex cost functions for period holding and penalty costs instead of a linear holding and penalty cost function, the resulting deterministic problem is differentiable and its optimal solution is unique. In this setting, it can be shown that the bounds in Theorems 2.5.2 and 2.5.3 still hold. This suggests that the bounds in Theorems 2.5.2 and 2.5.3 are not simply an artifact of a particular choice of optimal solution.

Remark 2.5.2. (On the Use of Poisson Demand in Theorem 2.5.3). Instead of using Poisson demand in Theorem 2.5.3, it is also possible to use Normal demand to better highlight the impact of B, σ, and μ on performance. If demands are all Normal with mean μ^* and standard deviation σ^*. If $\sigma^* = o(B)$ and $T = o(e^{\min\{\Theta(\mu^*),\, \Theta((B/\sigma^*)^2)\}})$, it can be shown using arguments similar to those used in the proof of Theorem 2.5.3 that there exists a constant $M > 0$ independent of T, μ^*, σ^*, and B such that $\frac{\mathbf{E}[J^{CEC} - J^*]}{J^*} \leq \frac{M\sigma^*}{\mu^* + \sigma^*\sqrt{T}}$ for all large μ^* and B/σ^*.

2.5.3 Key Ideas in the Proof of Theorems 2.5.2 and 2.5.3

We now briefly discuss the key ideas behind the proofs of Theorems 2.5.2 and 2.5.3. We focus only on CEC, since the arguments for DC are similar. In addition, we will also draw a connection between the optimal allocation control derived in Theorem 2.5.1 and the proposed solution in Theorems 2.5.2 and 2.5.3.

The proof of Theorem 2.5.2 has roughly two major parts. In the first part, we show that the sequence $\{n^H_{t,i}\}$ defined as $n^H_{t,i} = n^{*1}_{t,i} + \Delta_{t,i} - \frac{1}{m}\sum_{j=1}^m \Delta_{t,j}$ for all t and i is optimal for the hindsight problem on $\mathcal{A}$. Let $I^H_{t,i}$ denote the starting inventory level for retailer i at the beginning of period t under hindsight policy. Simple algebra yields:

$$I^H_{t,i} = \sum_{s=1}^{t-1} n^{*1}_{s,i} - (t-1)\mu_i - \frac{1}{m}\sum_{s=1}^{t-1}\sum_{j=1}^{m}\Delta_{s,j} \quad \text{for all } t \text{ and } i.$$

In the second part of the proof, we study an optimal solution of $J^D_t(I_t)$ on $\mathcal{A}$. Define n^{*t} for $t > 1$ as follows:

$$n^{*t}_{t,i} = n^{*1}_{t,i} + \Delta_{t-1,i} - \frac{1}{m}\sum_{j=1}^{m}\Delta_{t-1,j} \quad \text{and} \quad n^{*t}_{s,i} = n^{*1}_{s,i} \quad \text{for } s > t.$$

We claim that if $n_s = n^{*s}_s$ for all $s \le t-1$, then:

(i) The starting inventory level for retailer i at the beginning of period $t > 1$ is given by

$$I^{CEC}_{t,i} = \sum_{s=1}^{t-1} n^{*1}_{s,i} - (t-1)\mu_i - \Delta_{t-1,i} - \frac{1}{m}\sum_{s=1}^{t-2}\sum_{j=1}^{m}\Delta_{s,j}$$

(ii) n^{*t} is an optimal solution for $J^D_t(I_t)$.

In the detailed proof in the Appendix, we prove this claim by induction. In the proof, we show that the proposed allocation is optimal by a duality argument. Specifically, we first re-write $J^D_t(I_t)$ as a linear program and then show that there exists a set of feasible dual and primal variables (that correspond to the proposed allocation) such that all the Karush-Kuhn-Tucker (KKT) conditions are satisfied. Since KKT conditions are both necessary and sufficient for optimality in linear program, we conclude that the proposed allocation is optimal.

Combining the results from both parts of the proof, we can write: $I^{CEC}_{t+1,i} = I^H_{t+1,i} - \Delta_{t,i} +$

$\frac{1}{m}\sum_{j=1}^{m}\Delta_{t,j}$ for all t and i. This allows us to bound $J^{CEC}-J^{H}$ on $\mathcal{A}$ as follows:

$$\begin{aligned} J^{CEC}-J^{H} &= \sum_{t=1}^{T}\sum_{i=1}^{m}\left[p_i(-I_{t+1,i}^{CEC})^{+}+h_i(I_{t+1,i}^{CEC})^{+}\right]-\sum_{t=1}^{T}\sum_{i=1}^{m}\left[p_i(-I_{t+1,i}^{H})^{+}+h_i(I_{t+1,i}^{H})^{+}\right] \\ &\le \sum_{t=1}^{T}\sum_{i=1}^{m}(p_i+h_i)\left|\Delta_{t,i}-\frac{1}{m}\sum_{j=1}^{m}\Delta_{t,j}\right|. \end{aligned}$$

The result of Theorem 2.5.2 follows because $\mathbf{E}[|\Delta_{t,i}-\frac{1}{m}\sum_{j=1}^{m}\Delta_{t,j}|\cdot\mathbf{1}\{\mathcal{A}\}]\le\mathbf{E}[|\Delta_{t,i}-\frac{1}{m}\sum_{j=1}^{m}\Delta_{t,j}|]\le\mathbf{E}[|\Delta_{t,i}|]+\frac{1}{m}\sum_{j=1}^{m}\mathbf{E}[|\Delta_{t,j}|]\le 2\sigma$.

The proof of Theorem 2.5.3 is more involved than the proof of Theorem 3 and can be broken down into four parts. In the first part, we again argue that the sequence $\{n_{t,i}^{H}\}$ defined as above is optimal for the hindsight problem on $\hat{\mathcal{A}}$. In the second part, we prove by induction that n^{*t} defined as above is also optimal for $J_t^D(I_t)$ (provided that we use $n_s=n_s^{*s}$ for all $s\le t-1$). However, there is one key difference. In the proof of Theorem 2.5.2, the induction uses a duality argument and constructs a set of feasible dual and primal variables that satisfy KKT conditions. In the proof of Theorem 2.5.3, this argument no longer works because we do not assume, as in Theorem 2.5.2, that $n_{t,i}^{*1}\ge\varphi$ and $\left|\sum_{s=1}^{t}n_{s,i}^{*1}-t\mu_i\right|\ge t\varphi$ for all t and i, for some $\varphi>0$; instead, we use a convexity argument (i.e., $\sum_{i=1}^{m}g(x_i)\ge\sum_{i=1}^{m}g(\frac{1}{m}\sum_{j=1}^{m}x_j)$ for any convex function $g(\cdot)$) to prove the optimality of n^{*t} for $J_t^D(I_t)$. This works because we assume that $p_i=p^*$ and $h_i=h^*$ for all i. For the first half of Theorem 2.5.3, the condition that demands are stationary over time and identically distributed among all retailers are not necessary. Since the same allocation is optimal in Theorem 2.5.3 as in Theorem 2.5.2, we still have:

$$J^{CEC}-J^{H} \le \sum_{t=1}^{T}\sum_{i=1}^{m}(p_i+h_i)\left|\Delta_{t,i}-\frac{1}{m}\sum_{j=1}^{m}\Delta_{t,j}\right|.$$

Thus, $\mathbf{E}[(J^{CEC}-J^{H})\cdot\mathbf{1}(\hat{\mathcal{A}})]\le 2mT(p^*+h^*)\sigma^*$. To complete the proof for the relative bounds with Poisson demands, we need two additional steps: we need to compute a bound for $P(\hat{\mathcal{A}})$ (this is the third part) and we need to compute a lower bound for J^* (this is the last part). We show in the electronic companion that $P(\hat{\mathcal{A}})\ge 1-2mTe^{-\min\{\mu^*,\,B^2/(16m^4\mu^*)\}}$. (This can be shown using a standard exponential tail argument via a moment generating function of a Poisson random variable.) Note that if both μ^* and B^2/μ^* are large, $\hat{\mathcal{A}}$ occurs with high probability, so $\mathbf{E}[J^{CEC}-J^{H}]$ approximately equals $\mathbf{E}[(J^{CEC}-J^{H})\cdot\mathbf{1}\{\hat{\mathcal{A}}\}]$ and we can (roughly) ignore $\mathbf{E}[(J^{CEC}-J^{H})\cdot\mathbf{1}\{\hat{\mathcal{A}}^c\}]$ when bounding $\mathbf{E}[J^{CEC}-J^{H}]$. As for a lower bound for J^*, we show in the electronic companion that there exists a constant $M'>0$ independent of $T>0$ such that, for all large μ^*, we have

$$J^* \ \ge\ cT\mu^*+M'\sqrt{\mu^*}\,T^{3/2}.$$

The proof of this bound utilizes a well-known Central Limit Theorem (CLT) result – the fact that $\mathbf{E}\left[|n\theta - \sum_{i=1}^{n} X_i|\right] = \Theta(\sigma\sqrt{n})$ for large n and i.i.d X_i's with mean θ and standard deviation σ. Combining these observations, we finally have:

$$\begin{aligned}\frac{\mathbf{E}[J^{CEC}] - J^*}{J^*} &\leq \frac{\mathbf{E}[J^{CEC} - J^H]}{J^*} \approx \frac{\mathbf{E}[(J^{CEC} - J^H)\cdot \mathbf{1}(\hat{\mathcal{A}})]}{J^*} \\ &\leq \frac{2mT(p^* + h^*)\sqrt{\mu^*}}{cT\mu^* + M'\sqrt{\mu^*}\, T^{3/2}} \leq \frac{M}{\sqrt{\mu^*} + \sqrt{T}}\end{aligned}$$

for some $M > 0$. This completes the proof.

There is a relationship between the optimal allocation control derived in Theorem 2.5.1 and the proposed solution used in Theorems 2.5.2 and 2.5.3. In both Theorems 2.5.2 and 2.5.3, the proposed optimal solution is of the form $n_{t,i}^{*t} = n_{t,i}^{*1} + \Delta_{t-1,i} - \frac{1}{m}\sum_{j=1}^{m}\Delta_{t-1,j}$ for all t and i. It is not difficult to check that this solution corresponds to a particular inventory-balancing policy. Let $I_{t,i}^D := \sum_{s=1}^{t-1} n_{s,i}^{*1} - (t-1)\mu$ denote the starting inventory level for retailer i at the beginning of period t under the deterministic system with $I_{1,i} = 0$. Also, let $\hat{I}_{t,i}^D := I_{t,i}^D + n_{t,i}^{*1}$ where $\hat{I}_{t,i}^D$ can be interpreted as the after-allocation target inventory level for retailer i in period t. Under CEC, we can write $I_{t,i}^{CEC} = \sum_{s=1}^{t-1} n_{s,i}^{*s} - \sum_{s=1}^{t-1} D_{s,i} = I_{t,i}^D - \frac{1}{m}\sum_{j=1}^{m}\sum_{s=1}^{t-2}\Delta_{s,j} - \Delta_{t-1,i}$. So, allocating $n_{t,i}^{*t}$ units to retailer i at period t immediately brings the inventory level to $I_{t,i}^{CEC} + n_{t,i}^{*t} = I_{t,i}^D + n_{t,i}^{*1} - \frac{1}{m}\sum_{j=1}^{m}\sum_{s=1}^{t-1}\Delta_{s,j} = \hat{I}_{t,i}^D - \frac{1}{m}\sum_{j=1}^{m}\sum_{s=1}^{t-1}\Delta_{s,j}$. This means that our proposed solution balances the inventory level at all retailers by the same offset relative to the deterministic target level. In the case of Theorem 2.5.1, since all retailers are homogeneous, there are uniform target levels; so, this is equivalent to bringing the inventory level at each retailer to the same value.

2.6 Beyond the Deterministic and Certainty Equivalent Controls

Erring on the side of simplicity, DC deliberately ignores two key elements of the original stochastic problem: (1) demand variation and (2) decision dynamics (adaptive decisions contingent on realized demands observations). The combined impact of these two elements shows up in the bound; the regret of DC scales linearly with σ and polynomially with T. In Section 2.5, we proved that simple periodic re-optimizations of DC (i.e., CEC) improve the performance of DC by reducing the dependency of its regret on T from polynomial to linear. This result highlights an important point that re-optimizations of a deterministic system are useful for addressing the lack of decision dynamics in DC, at least for our setting, which suggests several follow-up questions: Is it possible to further reduce the dependency of regret on either σ or T from linear to sublinear? Moving beyond the use of expected values, is it more useful to explicitly incorporate complete demand information and couple this with an ad-hoc modeling of decision dynamics via re-optimizations, or is it more useful to use a little bit more of demand distribution information beyond expected values and couple this with an explicit model of decision dynamics (via DP, for example)? That is, can we compensate for the lack of demand distribution information with a more sophisticated model of decision dynamics, and vice versa?

In this section, we analyze two approaches to improve the performance of CEC: (1) *Open-Loop Feedback Control* (OLFC) and (2) *Multi-Point Approximation Control* (MPAC). OLFC uses the complete demand distribution information but only partially incorporates decision dynamics in an ad-hoc manner via re-optimizations. In contrast, MPAC directly models the complete decision dynamics via a dynamic program, but approximates demand distribution using a multi-point approximation of demand. We show, at least in some cases, that an optimal solution under OLFC is also an optimal solution under CEC. (See also the discussions after the statement of Lemma 2.6.1 in Section 2.6.1.) Although OLFC improves the performance of CEC as shown by our numerical results in Section 6, in general its regret may still scale linearly with both σ and T. In contrast, with a proper choice of demand approximation in each period, the regret of MPAC scales sublinearly with T. This result has an important implication: To significantly reduce the dependency of regret on T (or potentially σ) from linear to sublinear in our setting, it appears to be necessary to include a more sophisticated model of decision dynamics.

2.6.1 Open-Loop Feedback Control

We proceed in two stages as follows: In stage 1, we solve C^D to calculate the number of new batches to produce at the beginning of each period; in stage 2, instead of re-optimizing $J_t^D(I_t)$ as in the case of CEC, we re-optimize $J_t^S(I_t)$ defined below:

$$
\begin{aligned}
J_t^S(I_t) \quad = \quad \min_n \sum_{\xi=t}^{T} \Bigg\{ c z_t^D B + \mathbf{E}\Bigg[\sum_{i=1}^{m} p_i \left(\sum_{s=t}^{\xi} D_{s,i} - \sum_{s=t}^{\xi} n_{s,i} - I_{t,i} \right)^+ \\
+ \sum_{i=1}^{m} h_i \left(I_{t,i} + \sum_{s=t}^{\xi} n_{s,i} - \sum_{s=t}^{\xi} D_{s,i} \right)^+ \Bigg] \Bigg\} \qquad (2.24)
\end{aligned}
$$

$$
\text{s.t.} \ \sum_{i=1}^{m} n_{s,i} = z_s^D B, \ n_{s,i} \in \mathbb{R}^+ \qquad \forall s, \ i \qquad (2.25)
$$

Observe that $J_t^S(I_t)$ is similar to a single period newsvendor problem. If demands are continuous and $F_i(\cdot)$ is differentiable and strictly positive on $(0, \infty]$ for all i, it is not difficult to show that the objective function in (2.24) is convex on $(0, \infty]^m$. Let $n^{St} = (n_{s,i}^{St})_{s \geq t, i \geq 1}$ denote the optimal solution of $J_t^S(I_t)$. (We suppress the notational dependency of n^{St} on I_t.) The complete description of OLFC is given below:

Open-Loop Feedback Control - OLFC

---

1. At the beginning of period 1, solve C^D
2. For $t = 1, 2, ..., T$, do:
   (a) At the beginning of period t, produce exactly z_t^D new batches
   (b) Solve $J_t^S(I_t)$ and allocate exactly $n_{t,i}^{St}$ units to retailer i
   (c) Update $I_{t+1} = I_t + n_t^{St} - D_t$

Although our numerical results in Section 2.7 show that OLFC consistently performs better than CEC, it is challenging to analytically characterize this improvement. Interestingly, it is possible to show that an optimal allocation under OLFC is sometimes also optimal under CEC (see Lemma 2.6.1 below). This suggests that the benefit of simply expoliting a complete demand distribution information without sufficient decision dynamics is already captured (at least, partially) by periodic re-optimizations of a deterministic system with a one-point demand approximation.

Lemma 2.6.1. *Suppose that $I_1 = 0$, $p_i = p^*$ and $h_i = h^*$ for all i, demands are i.i.d with mean μ^*, standard deviation σ^* and their common cdf is differentiable and strictly positive on $(0, \infty]$, and $z_t^D > 0$ for all t. Let $\hat{\mathcal{A}} := \{\sum_{i=1}^m |\Delta_{t,i}| \leq B/(2m)\ \forall t\}$. Define $n_{t,i}^{St} = \frac{z_t^D B}{m} + \Delta_{t-1,i} - \frac{1}{m}\sum_{j=1}^m \Delta_{t-1,j}$ for all t and i. Then, $\{n_{t,i}^{St}\}$ is an optimal allocation under OLFC on $\hat{\mathcal{A}}$.*

The proof of Lemma 2.6.1 can be found in Appendix A.1.6. Recall from Section 2.5.3 that $\frac{z_t^D B}{m} + \Delta_{t-1,i} - \frac{1}{m}\sum_{j=1}^m \Delta_{t-1,j}$ is the (constructed) optimal allocation under CEC for period t. Thus, despite the fact that OLFC uses complete demand distribution information and CEC only uses a one-point demand approximation, the regret associated with OLFC in general may still be $O(\sigma^* T)$ (because the bounds in Theorem 2.5.3 also hold for OLFC). Indeed, for the special case where $\frac{z_t^D B}{m} - \mu^* > \frac{B}{2m}$ for all t (which can happen if there is a mismatch between the number of units in a batch, or an integer multiple of a batch, and total expected demands in all retailers), it is not difficult to show that $J_t^S(I_t)$ is strictly convex on $\hat{\mathcal{A}}$ under the proposed solution. Since an interior optimal solution in a strongly convex problem is also the unique optimal solution to the problem, the regret of OLFC on the set $\hat{\mathcal{A}}$ in this special case is exactly the same as the regret of CEC on the same set under the proposed solution in Theorem 2.5.3.

2.6.2 Multi-Point Approximation Control

In the previous subsection, we saw that explicitly modeling the complete demand distribution and combining this with a simplified model decision dynamics is not sufficient to improve significantly on CEC. We now consider explicitly incorporating complete decision dynamics. The essence of MPAC is the use of a multi-point demand approximation to capture more demand distribution information beyond that captured by expected value, and using this multi-point approximation in a full DP to completely exploit decision dynamics. There are potentially many ways of doing this; here, we will only discuss one such approach. For each i, let S_i denote the support of $D_{t,i}$ and $\hat{D}_{t,i}$ denote the approximation of $D_{t,i}$. We consider an approximation of the following form: There exists a partition $\{\Omega^k_{t,i}\}$ (i.e., $\cup_k \Omega^k_{t,i} = S_i$) and a mapping $m_{t,i} : \{\Omega^k_{t,i}\} \to \{v^k_{t,i}\}$ such that $D_{t,i} \in \Omega^k_{t,i}$ is approximated (or represented) by $\hat{D}_{t,i} = v^k_{t,i}$. Note that, by construction, $\mathbf{P}(\hat{D}_{t,i} = v^k_{t,i}) = \mathbf{P}(D_{t,i} \in \Omega^k_{t,i})$. For example, if $v^k_{t,i} = \mu_i = \mathbf{E}[D_{t,i}]$ for all k, t, and i, then we have the one-point approximation used in DC and CEC.

Now, consider the following optimization problem:

$$\hat{C}^* = \inf_{\pi\in\Pi} \sum_{t=1}^{T} \mathbf{E}\left[cZ^\pi_t B + \sum_{i=1}^{m} p_i(\hat{D}_{t,i} - N^\pi_{t,i} - I^\pi_{t,i})^+ + \sum_{i=1}^{m} h_i(I^\pi_{t,i} + N^\pi_{t,i} - \hat{D}_{t,i})^+\right] \quad (2.26)$$

$$\text{s.t.} \sum_{i=1}^{m} N^\pi_{t,i} = Z^\pi_t B,\ Z^\pi_t \in \mathbb{Z}^+,\ N^\pi_{t,i} \in \mathbb{R}^+ \qquad \forall t,\ i \quad (2.27)$$

$$I^\pi_{t+1,i} = I^\pi_{t,i} + N^\pi_{t,i} - \hat{D}_{t,i} \qquad \forall t,\ i \quad (2.28)$$

$$I^\pi_{1,i} = I_{1,i} \qquad \forall i \quad (2.29)$$

where the expectation is taken with respect to the induced probability distribution for $\{\hat{D}_{t,i}\}$. Let $\hat{\pi}^* = (\hat{\pi}^*_1, \cdots, \hat{\pi}^*_T)$ denote the optimal control of $\hat{C}^*$. Since $\hat{\pi}^*$ is a control defined in a "virtual" world where demands are realized according to $\hat{D}$ instead of D, it is not immediately clear how to translate $\hat{\pi}^*$ into a control $\pi^R = (\pi^R_1, \cdots, \pi^R_T)$ to be implemented in a "real-world" setting where demands are realized according to D. (The superscript "R" stands for "real".) Here, we will focus on the following control translation scheme:

$$\pi^R_1 = \hat{\pi}^*_1 \quad \text{and} \quad \pi^R_t(D_1, ..., D_{t-1}) = \hat{\pi}^*_t(I^{\hat{\pi}^*}_t) \ \ \forall\, t > 1$$

where $I^{\hat{\pi}^*}_{t+1,i} = I^{\hat{\pi}^*}_{t,i} + N^{\hat{\pi}^*}_{t,i} - \hat{D}_{t,i}$ (i.e., $I^{\hat{\pi}^*}_t$ is the virtual starting inventory level at the beginning of period t under control $\hat{\pi}^*$ and demand realizations $\hat{D}_1, \ldots, \hat{D}_{t-1}$). Under control π, at the beginning of period t, we first calculate the virtual inventory level at each retailer; next, we produce exactly $Z^{\hat{\pi}^*}_t$ new batches and allocate exactly $N^{\hat{\pi}^*}_{t,i}$ units to retailer i. So, we respond as if demands are generated according to $\hat{D}$ instead of D. We state our result below.

Theorem 2.6.2. *Let $\theta_{t,i} := \mathbf{E}[(D_{t,i} - \hat{D}_{t,i})^2]^{1/2}$. Then,*

$$\mathbf{E}\left[C^{\pi^R}\right] - C^* \ \leq\ 2\sum_{i=1}^{m}(p_i + h_i)\left[\sum_{t=1}^{T}\left(\sum_{s=1}^{t}\theta^2_{s,i}\right)^{1/2}\right].$$

The proof of Theorem 2.6.2 can be found in Appendix A.1.7. Theorem 2.6.2 is the generalization of Theorem 2.4.1. (If $\hat{D}_{t,i} = \mu_i$, then $\theta_{t,i} \leq \sigma$, and we completely recover the bound in Theorem 2.4.1.) It highlights the value of information in a multi-period inventory control problem; in particular, it shows that it is most beneficial to use a more granular demand approximation during earlier instead of later periods. To illustrate this, suppose that $D_{t,i}$ is uniformly distributed on $[L_i, U_i]$. If we use a $(T-t+1)^{1/2+\alpha}-$demand approximation for $D_{t,i}$ for some $\alpha > 0$ (i.e., by using $(T-t+1)^{1/2+\alpha}$ points in $[L_i, U_i]$), then $\mathbf{E}[C^{\pi^R}] - C^* = O(\sigma T^{1-\alpha})$. Note that, as α becomes large, the regret decreases to 0.

Remark 2.6.1. (Computational Complexity of MPAC). Despite the promising bound in Theorem 2.6.2, MPAC solves a full DP. Thus, it is computationally much more intensive than either CEC or OLFC. One potential way to mitigate this computational burden is to use a form of rollout algorithm with limited lookahead (Bertsekas (2013)[8], Goodson et al. (2015)[30]). The analysis of a rollout algorithm for an undiscounted finite-horizon stochastic inventory problem is an open research problem. As it is possibly a very challenging task, we leave this for future research pursuit. Our purpose in this chapter is simply to highlight the potential benefit of including more decision dynamics in designing a heuristic control.

2.7 Computational Experiments

In this section, we test the performance of the heuristic controls discussed in previous sections (DC, CEC, OLFC, and MPAC) along with two others, CEC-IP and OLFC-IP, where CEC and OLFC respectively are modified so that both the production and allocation decisions are re-optimized at each iteration (see Table 2.1 for a summary). We run two sets of experiments. In the first set, we consider two non-homogeneous retailers and compare the performance all heuristic controls (as well as the optimal control from explicitly solving the DP). In the second set, we consider eight non-homogeneous retailers with industrial-scale demands. These problems are much larger, so we exclude MPAC and the optimal DP for computational reasons.

To assess the performance of our heuristic controls, we use the following performance benchmark (especially in the second set of experiments where we are not able to compute the optimal control):

$$\begin{aligned} C^H(D) \quad = \quad & \min_n \sum_{t=1}^{T} \left[cZ_tB + \sum_{i=1}^{m} p_i(D_{t,i} - x_{t,i} - n_{t,i})^+ + \sum_{i=1}^{m} h_i(x_{t,i} + n_{t,i} - D_{t,i})^+ \right] \\ & \text{s.t.} \sum_{i=1}^{m} n_{t,i} = Z_t\, B,\; Z_t^\pi \in \mathbb{Z}^+, n_{t,i} \in \mathbb{R}^+ \qquad \forall t,\, i \\ & \quad x_{t+1,i} = x_{t,i} + n_{t,i} - D_{t,i} \qquad \forall t,\, i \\ & \quad x_{1,i} = I_{1,i} \qquad \forall i \end{aligned}$$

As with J^H, C^H is the hindsight control for a given demand vector instance D, where the number of batch production in each period is also a decision variable. Clearly, $\mathbf{E}[C^H] \leq C^*$, so $\mathbf{E}[C^H]$ is a lower bound on the optimal expected total costs (as well as a lower bound on each of our heuristic controls). Let C^π denote the total costs under heuristic control π. We use three measures to evaluate the performance of each heuristic control:

1. Expected relative regret with respect to the hindsight lower bound

$$\mathbf{E}[R^\pi] = \frac{\mathbf{E}[C^\pi - C^H]}{\mathbf{E}[C^H]}$$

2. Coefficient of variation of the regret

$$\mathbf{cv}(R^\pi) = \frac{\sqrt{\mathbf{Var}(C^\pi - C^H)}}{\mathbf{E}[C^\pi] - \mathbf{E}[C^H]}$$

3. Running time to complete one instance

We estimate the first two measures using Monte-Carlo simulation; we generate a number of demand instances $D^s \in \mathcal{S}$ according to the relevant demand distribution F, and $\mathbf{E}[C^\pi]$ and $\mathbf{E}[C^H]$ are estimated using sample mean and sample variance respectively.

All experiments are run using MATLAB R2016a with Intel Core i7-5820K CPU and 16GB RAM. The Linear programs and mixed integer linear programs in DC, CEC, and CEC-IP are solved using CPLEX 12.6. The stochastic programs in OLFC and OLFC-IP are solved by interior-point method using the Optimization Toolbox in MATLAB 2016a.

Control	**Demand Information**	**Decision Dynamics**
DC	Use a single point estimate of demand	Solve production and allocation once
CEC	Use a single point estimate of demand	Solve production once, re-optimize allocation in every period
CEC-IP	Use a single point estimate of demand	Re-optimize production and allocation in every period
OLFC	Use complete demand distribution	Solve production once, re-optimize allocation in every period
OLFC-IP	Use complete demand distribution	Re-optimize production and allocation in every period
MPAC	Use multiple points estimate of demand	Solve an entire DP once

Table 2.1: Summary of all heuristic controls

2.7.1 Two Retailers

In this section we compare the performance of all heuristic controls (and also the optimal control, denoted as OPT) on a small example with two non-homogeneous retailers with mean demand $\mu = [10, 6]$, holding cost $h = [2, 3]$, penalty cost $p = [8, 5]$, and batch size $B = 10$. We run two different sets of experiments: In the first, demands are Binomial, modeling the case of "small" demand variation, and in the second, demands are Uniform, which represent the case of "large" demand variation. For each set of experiments, we generate 500 demand instances to estimate our three performance measures. The results are summarized in Tables 2.2-2.4.

Based on these experiments, the relative regret of the heuristic controls can be ordered from greatest to smallest as follows: DC > CEC > OLFC > CEC-IP > OLFC-IP > MPAC. As suggested by Theorems 2.4.1 and 2.5.3, despite the simplicity of DC, the regret associated with it gets worse rapidly as the length of planning horizon increases, e.g. the regret of DC over 10 periods is about 3 times worse than the regret associated with OPT. The performance of both CEC and OLFC is consistently better than DC as expected. Notice although OLFC consistently outperforms CEC by using complete demand information, it is only better by a small margin. This confirms the implication of Lemma 2.6.1 – OLFC leads to limited improvement over CEC. On the other hand, exploring more decision dynamics, such as re-optimizing the production decision in addition to the allocation decision (i.e., CEC-IP and OLFC-IP) significantly reduce the regret, especially as the length of planning horizon increases. In terms of relative regret, MPAC outperforms all other heuristic controls in both sets of experiments. Note that here, we only use a two-point approximation of the demand distribution. This suggests that, in many cases, we can achieve a small regret by slightly increasing the amount of demand information that we use, as long as we completely exploit the decision dynamics by solving a full DP. For practical implications, this means company should focus more on including decision dynamics in its production planning rather than estimating demand more precisely, given the estimation of the whole demand distribution is challenging in reality.

The coefficient of variation of the regret of the heuristic controls can be ranked as follows: OLFC > DC ≈ CEC > CEC-IP ≈ OLFC-IP ≈ MPAC. Note that re-optimizing both production and allocation decisions reduces the variability of regret. Interestingly, although OLFC consistently outperforms CEC on expected relative regret, it has the largest regret variability among all heuristic controls. In Table 2.4, we report the average running time on all instances with 10-period for each heuristic control and the optimal control for the case with Binomial demand – it can be ranked as follows: OPT > MPAC > OLFC-IP > CEC-IP > CEC > OLFC ≈ DC. Note that all heuristic controls run significantly faster than the the optimal control, but OLFC-IP and CEC-IP are particularly effective in terms of balancing the running time and the expected percentage regret.

Finally, note that in this case, since we are optimally solving the problem, we could have computed the relative regret with respect to the optimal solution. We chose not to do this to highlight the performance of our lower bound, since in the next section we are unable to solve the problem optimally. In particular, observe that in some cases, the relative regret of the heuristic control is largely due to the weakness of this bound. That said, our definition of regret with respect to hindsight control is still useful to highlight how different heuristic controls fare against each other.

	Percentage regret						
T	DC	CEC	CEC-IP	OLFC	OLFC-IP	MPAC	OPT
1	31.73%	31.73%	31.73%	25.77%	25.77%	25.77%	25.77%
2	44.01%	39.67%	32.78%	35.69%	28.51%	28.28%	27.61%
3	55.74%	47.53%	34.06%	43.92%	29.65%	29.63%	28.65%
4	69.63%	56.89%	34.82%	54.36%	30.65%	30.58%	29.43%
5	68.52%	52.09%	34.48%	48.76%	30.39%	30.45%	29.37%
6	76.93%	57.84%	34.59%	55.24%	30.72%	30.60%	29.13%
7	82.88%	60.50%	34.72%	58.02%	30.96%	30.96%	30.19%
8	91.08%	65.46%	34.71%	64.45%	31.03%	30.97%	30.02%
9	99.55%	70.56%	34.86%	70.58%	31.19%	31.11%	29.92%
10	101.79%	70.27%	34.67%	68.42%	31.00%	30.89%	29.69%
	c.v. of regret						
T	DC	CEC	CEC-IP	OLFC	OLFC-IP	MPAC	OPT
1	106.33%	106.33%	106.33%	124.96%	124.96%	124.96%	124.96%
2	81.04%	82.26%	77.60%	87.89%	80.93%	82.20%	81.92%
3	74.39%	70.96%	61.33%	79.62%	64.55%	65.72%	67.51%
4	74.25%	76.24%	55.74%	87.59%	59.55%	60.43%	59.55%
5	67.68%	67.67%	50.04%	78.83%	53.10%	54.09%	51.72%
6	62.69%	63.54%	44.62%	72.81%	46.63%	47.96%	46.73%
7	65.55%	68.19%	40.87%	81.15%	42.43%	43.41%	43.48%
8	63.85%	66.68%	38.39%	77.13%	40.61%	41.73%	40.77%
9	65.48%	69.96%	36.28%	81.06%	37.93%	39.76%	38.42%
10	63.01%	67.54%	34.62%	80.19%	36.13%	38.23%	35.56%

Table 2.2: Binomial demand

	Percentage regret						
T	DC	CEC	CEC-IP	OLFC	OLFC-IP	MPAC	OPT
1	107.31%	107.31%	107.31%	107.31%	107.31%	107.31%	107.31%
2	142.71%	133.54%	111.15%	130.94%	108.61%	107.43%	107.41%
3	174.81%	155.31%	113.46%	151.41%	110.72%	103.62%	101.44%
4	203.96%	174.72%	112.91%	170.57%	110.33%	104.64%	102.10%
5	212.79%	175.37%	113.32%	169.76%	110.45%	103.79%	100.95%
6	230.66%	186.68%	113.55%	180.76%	111.60%	102.84%	99.96%
7	249.01%	197.99%	113.86%	192.11%	111.45%	102.73%	99.26%
8	266.54%	209.60%	114.69%	203.62%	111.79%	102.46%	99.52%
9	285.86%	222.49%	115.81%	217.01%	113.04%	102.53%	100.17%
10	293.73%	226.22%	115.32%	219.35%	112.64%	102.60%	99.66%
	c.v. of regret						
T	DC	CEC	CEC-IP	OLFC	OLFC-IP	MPAC	OPT
1	61.02%	61.02%	61.02%	61.02%	61.02%	61.02%	61.02%
2	59.82%	55.71%	46.65%	46.76%	59.53%	45.46%	45.09%
3	59.92%	53.56%	39.18%	39.15%	60.00%	37.55%	37.86%
4	63.19%	58.70%	35.37%	33.86%	66.15%	31.91%	31.62%
5	59.89%	54.92%	31.83%	29.95%	61.32%	27.59%	28.97%
6	59.04%	54.07%	29.13%	27.64%	61.18%	24.90%	26.17%
7	60.43%	56.02%	26.65%	24.99%	63.03%	24.17%	24.38%
8	60.80%	55.41%	24.59%	23.14%	62.89%	22.06%	22.12%
9	61.35%	57.05%	24.27%	22.52%	64.57%	20.90%	20.79%
10	60.92%	56.62%	23.16%	21.48%	64.09%	19.58%	19.89%

Table 2.3: Uniform demand

	Binomial						
	DC	CEC	CEC-IP	OLFC	OLFC-IP	MPAC	OPT
Solving time (sec)	<0.01	0.0603	0.2240	<0.01	0.2417	24.4956	3425.3

Table 2.4: Average solution time for a 10-period instance

2.7.2 Eight Retailers

In this set of experiments, we compare DC, CEC, CEC-IP, OLFC and OLFC-IP on an industrial-size example to better highlight the potential performance of our heuristic controls. The test problem has eight non-homogeneous retailers with mean demand $\mu = [5000, 3000, 2000, 1000, 500, 300, 200, 100]$, holding cost $h = [1.8, 2.1, 2.3, 2.9, 3.1, 3.7, 4.3, 4.4]$, penalty cost $p = [8, 5.1, 9, 5.8, 9.6, 5.5, 11, 6.3]$, and batch size $B = 10,000$. We run four sets of experiments: In the first, demands are Poisson and, in the next three, demands are

Normal with standard deviations equal to 5%, 15%, and 25% of their mean, respectively. The first two sets of experiments represent the case of "small" demand variation and the last two experiments capture "large" demand variation. For each of the four experiments, we generate 50 demand instances to estimate our three performance measures. The results are summarized in Tables 2.5 to 2.9.

The expected relative regret is similar to that in the smaller experiments discussed above, with DC > CEC > OLFC > CEC-IP > OLFC-IP. However, as predicted in Corollary 2.4.2, when the size of demand variation is relatively small compared to its mean, the performance of the simple DC is much more satisfying, with only 13% to 25% regret compared to the hindsight lower bound over 20 periods. The performance of CEC appears to be quite stable for the case of Poisson and 5% Normal demands (i.e., the regret does not seem to be affected by the length of planning horizon). As the size of demand variation becomes large (e.g., 15% and 25% Normal), the regret of CEC slowly increases and the length of planning horizon starts to have a more noticeable impact on performance; this is also as expected. (Per Remark 2.6.1, as σ^* increases, $T = o(e^{B^2/(8m^4(\sigma^*)^2)})$ becomes smaller. So, it becomes necessary to also re-optimize the production decision.) Similar behavior is also observed with OLFC. Note that although OLFC consistently outperforms CEC, the regret reduction of OLFC over CEC appears marginal compared to the regret reduction of CEC over DC. This observation highlights our theoretical prediction in Section 5 that much of the benefit of using the complete demand distribution is already captured by periodic re-optimizations of a deterministic model with a one-point demand approximation. The same as in the smaller experiments in Section 6.1, incorporating more decision dynamics by re-optimizing the production decision, i.e. CEC-IP and OLFC-IP, reduces the regret of CEC and OLFC and also help stabilize the performance over longer planning horizon.

The coefficient of variation of regret also behaves similarly to the smaller experiments detailed above, with OLFC > DC ≈ CEC > OLFC-IP ≈ CEC-IP. We also report the average running time on a 20-period instance for each heuristic control in Table 2.9. Notice as the problem size increases, solving stochastic program in OLFC and OLFC-IP becomes increasingly more difficult than solving CEC and CEC-IP. Nevertheless the running time is by no means unreasonable for practical purpose. This highlights the efficiency of our heuristics and their potential for real-world application.

Managerial Insight: DC controls work well with small demand variation and a short planning horizon. In addition, incorporating more decision dynamics (e.g. in CEC-IP, MPAC) is more effective than using complete demand information (e.g. OLFC). Given that the estimation of full demand distribution is challenging and expensive in reality, this suggest that biopharmaceutical firms should focus on integrating more decision dynamics into their production planning.

	Percentage regret					c.v. of regret				
T	DC	CEC	CEC-IP	OLFC	OLFC-IP	DC	CEC	CEC-IP	OLFC	OLFC-IP
1	5.03%	5.03%	5.03%	3.85%	3.85%	48.32%	48.32%	48.32%	45.52%	45.52%
2	4.41%	3.64%	3.56%	2.41%	2.50%	37.96%	40.07%	37.80%	38.89%	42.80%
3	5.08%	3.67%	3.67%	2.13%	2.13%	41.47%	31.70%	31.70%	26.98%	26.98%
4	6.02%	3.90%	3.90%	2.18%	2.18%	44.43%	26.13%	26.13%	23.21%	23.21%
5	7.37%	4.28%	4.28%	2.57%	2.57%	34.36%	22.19%	22.19%	17.90%	17.90%
6	7.87%	4.21%	4.21%	2.61%	2.61%	30.55%	20.29%	20.29%	17.65%	17.65%
7	8.21%	4.49%	3.96%	2.90%	2.27%	31.21%	22.80%	18.43%	28.91%	19.59%
8	8.61%	4.45%	3.97%	2.80%	2.23%	32.69%	20.56%	16.87%	26.35%	15.96%
9	9.26%	4.62%	4.17%	2.77%	2.24%	34.16%	18.31%	15.86%	24.82%	15.11%
10	10.10%	4.71%	4.30%	2.90%	2.41%	29.44%	17.50%	15.58%	21.56%	13.61%
11	10.34%	4.58%	4.21%	2.89%	2.44%	27.54%	16.56%	13.54%	21.19%	12.73%
12	10.05%	4.41%	4.02%	2.63%	2.19%	31.75%	17.16%	15.76%	22.95%	14.48%
13	10.47%	4.42%	4.06%	2.60%	2.19%	32.82%	16.91%	15.40%	21.90%	13.77%
14	11.01%	4.55%	4.20%	2.63%	2.23%	32.91%	15.50%	13.47%	20.96%	12.66%
15	11.63%	4.56%	4.24%	2.70%	2.33%	29.51%	14.16%	12.22%	19.13%	12.09%
16	11.85%	4.51%	4.14%	2.74%	2.33%	28.37%	16.13%	11.41%	23.89%	12.01%
17	11.65%	4.39%	4.10%	2.51%	2.19%	31.15%	12.34%	10.54%	18.18%	11.94%
18	12.09%	4.46%	4.19%	2.49%	2.19%	31.95%	12.49%	11.12%	17.16%	11.39%
19	12.56%	4.54%	4.28%	2.56%	2.26%	31.41%	11.83%	10.58%	15.93%	10.68%
20	12.95%	4.52%	4.27%	2.60%	2.32%	29.25%	11.80%	10.32%	15.67%	10.64%

Table 2.5: Poisson demand

	Percentage regret					c.v. of regret				
T	DC	CEC	CEC-IP	OLFC	OLFC-IP	DC	CEC	CEC-IP	OLFC	OLFC-IP
1	8.87%	8.87%	8.87%	7.76%	7.76%	50.20%	50.20%	50.20%	42.94%	42.94%
2	8.53%	7.70%	7.14%	5.36%	5.21%	42.11%	51.20%	37.00%	57.54%	46.17%
3	9.04%	6.97%	6.97%	3.75%	3.75%	43.93%	32.89%	32.89%	34.59%	34.59%
4	10.96%	7.41%	7.41%	3.71%	3.71%	46.95%	30.60%	30.60%	30.65%	30.65%
5	13.99%	8.24%	8.24%	4.85%	4.85%	39.50%	29.36%	29.36%	28.86%	28.86%
6	15.35%	8.20%	8.16%	5.22%	5.15%	37.42%	27.14%	27.06%	24.57%	24.88%
7	16.62%	9.28%	7.68%	6.20%	4.56%	35.38%	28.70%	26.36%	40.41%	25.54%
8	17.14%	9.00%	7.68%	5.68%	4.21%	36.75%	27.26%	24.89%	41.20%	24.85%
9	18.40%	9.15%	7.93%	5.64%	4.27%	37.67%	26.35%	23.57%	41.60%	24.82%
10	19.60%	9.34%	8.21%	5.94%	4.69%	36.47%	23.68%	20.26%	36.61%	21.00%
11	20.22%	9.23%	7.96%	6.12%	4.71%	35.23%	25.34%	19.06%	39.56%	19.07%
12	19.88%	8.97%	7.68%	5.54%	4.34%	35.08%	20.74%	17.76%	31.40%	20.82%
13	20.53%	8.86%	7.71%	5.31%	4.18%	35.71%	19.16%	17.08%	29.97%	18.93%
14	21.55%	8.94%	7.84%	5.35%	4.27%	35.43%	18.99%	15.76%	29.75%	17.96%
15	22.64%	8.91%	7.88%	5.51%	4.50%	33.76%	18.60%	15.52%	27.52%	15.77%
16	24.44%	9.96%	7.78%	6.78%	4.46%	34.74%	33.59%	15.73%	50.06%	15.68%
17	23.15%	8.90%	7.60%	5.50%	4.20%	34.22%	22.46%	15.72%	34.32%	17.60%
18	23.63%	8.81%	7.60%	5.37%	4.15%	35.16%	21.15%	14.53%	33.67%	16.04%
19	24.47%	8.94%	7.78%	5.48%	4.31%	35.16%	20.45%	13.89%	32.82%	15.10%
20	25.70%	9.17%	7.81%	5.80%	4.44%	33.90%	23.91%	13.29%	37.75%	14.15%

Table 2.6: 5% Normal demand

	Percentage regret					**c.v. of regret**				
T	DC	CEC	CEC-IP	OLFC	OLFC-IP	DC	CEC	CEC-IP	OLFC	OLFC-IP
1	26.14%	26.14%	26.14%	21.18%	21.18%	47.06%	47.06%	47.06%	42.64%	42.64%
2	33.50%	29.37%	25.93%	21.11%	19.50%	42.97%	51.23%	43.97%	62.03%	50.88%
3	32.83%	25.10%	24.52%	14.75%	14.09%	38.37%	40.87%	35.08%	67.51%	57.79%
4	36.55%	24.57%	24.50%	14.44%	14.31%	37.92%	31.85%	31.31%	53.55%	51.56%
5	39.76%	24.81%	24.57%	15.85%	15.53%	35.29%	32.61%	30.48%	49.88%	45.97%
6	47.72%	28.65%	25.70%	20.96%	17.01%	39.46%	45.40%	30.16%	70.57%	40.64%
7	53.55%	32.24%	25.26%	23.68%	16.03%	39.25%	39.23%	27.39%	60.71%	42.24%
8	53.80%	29.93%	24.91%	21.04%	15.16%	39.25%	38.00%	25.39%	62.24%	39.18%
9	56.07%	29.48%	24.82%	20.77%	15.14%	39.29%	38.90%	21.83%	65.31%	33.49%
10	59.72%	30.32%	24.57%	22.28%	15.32%	41.18%	45.86%	19.95%	73.63%	27.60%
11	64.90%	33.34%	24.52%	26.20%	15.45%	43.40%	53.41%	18.05%	82.43%	25.73%
12	66.50%	32.37%	24.25%	23.95%	15.08%	40.14%	37.74%	16.96%	60.52%	26.36%
13	66.86%	30.93%	23.96%	22.48%	14.74%	39.79%	36.13%	14.85%	60.65%	23.80%
14	69.05%	31.01%	23.75%	22.79%	14.55%	40.70%	40.79%	14.41%	67.58%	21.26%
15	72.62%	32.44%	23.90%	24.73%	14.89%	41.78%	47.89%	13.84%	75.98%	21.98%
16	77.92%	35.57%	23.67%	28.87%	14.88%	42.20%	52.70%	13.25%	81.16%	19.89%
17	78.39%	34.34%	23.69%	26.37%	14.76%	39.09%	40.88%	12.64%	66.08%	20.00%
18	79.08%	33.72%	23.55%	25.88%	14.75%	39.63%	44.30%	11.92%	70.98%	20.37%
19	81.60%	34.41%	23.48%	26.79%	14.61%	40.85%	49.12%	11.15%	78.44%	18.61%
20	85.52%	36.23%	23.63%	29.35%	14.84%	41.99%	53.97%	10.40%	84.84%	17.53%

Table 2.7: 15% Normal demand

	Percentage regret					c.v. of regret				
T	DC	CEC	CEC-IP	OLFC	OLFC-IP	DC	CEC	CEC-IP	OLFC	OLFC-IP
1	44.68%	44.68%	44.68%	38.47%	38.47%	62.29%	62.29%	62.29%	77.57%	77.57%
2	63.36%	54.09%	48.91%	41.77%	39.58%	44.59%	45.70%	42.11%	56.86%	53.84%
3	58.77%	44.65%	42.78%	32.58%	31.98%	50.54%	48.59%	41.19%	63.90%	58.95%
4	63.60%	45.77%	44.12%	34.06%	33.45%	42.89%	41.53%	35.66%	53.78%	49.80%
5	69.90%	46.44%	43.50%	36.05%	32.37%	39.12%	39.76%	33.77%	52.60%	45.06%
6	83.18%	53.87%	43.73%	44.66%	31.45%	39.09%	44.16%	28.01%	62.70%	39.22%
7	95.51%	60.96%	43.68%	50.68%	30.70%	35.20%	34.98%	27.56%	48.78%	35.11%
8	96.48%	58.80%	43.21%	48.26%	30.57%	37.45%	36.89%	24.01%	52.61%	28.04%
9	99.71%	58.46%	43.06%	47.53%	29.77%	37.44%	40.26%	22.71%	57.79%	27.60%
10	105.98%	59.91%	42.05%	50.12%	29.11%	39.45%	45.18%	21.04%	64.68%	25.38%
11	115.99%	65.03%	42.42%	56.93%	29.25%	40.30%	48.07%	20.57%	68.93%	25.03%
12	116.85%	63.32%	42.98%	52.92%	30.17%	37.83%	37.86%	21.54%	53.59%	25.89%
13	120.85%	63.69%	43.21%	53.19%	30.01%	38.79%	39.71%	20.50%	56.42%	26.00%
14	125.84%	65.08%	43.40%	54.87%	30.24%	40.01%	42.26%	19.49%	61.33%	25.62%
15	133.13%	67.94%	43.31%	59.14%	29.89%	40.83%	46.31%	19.10%	66.58%	24.58%
16	142.33%	73.24%	43.18%	65.90%	30.03%	41.02%	49.37%	18.85%	69.49%	24.36%
17	142.75%	71.56%	43.24%	62.38%	29.84%	39.86%	44.62%	19.17%	64.00%	24.19%
18	147.14%	72.96%	43.59%	63.55%	29.99%	40.25%	46.93%	19.15%	67.30%	24.00%
19	151.73%	73.88%	43.31%	65.10%	29.97%	40.15%	49.25%	18.83%	70.80%	23.72%
20	157.22%	76.28%	43.29%	68.20%	29.70%	40.90%	52.71%	17.52%	75.06%	21.64%

Table 2.8: 25% Normal demand

	5% Normal				
	DC	CEC	CEC-IP	OLFC	OLFC-IP
Solving time (sec)	0.0146	0.3418	2.5976	41.9133	46.8846

Table 2.9: Average solution time for a 20-period instance

2.8 Conclusion

In practice, firms often solve planning problems by replacing random variables representing future demand with deterministic demand estimates and firms often use a rolling horizon approach to implement these solutions. In this chapter, we considered an inventory planning problem in a biopharmaceutical industry involving batch production and allocation, and analyzed a variety of heuristic controls that solve the deterministic version of this planning problem, both one time and in a rolling horizon setting. We characterized the performance of these heuristic controls, and found that the performance of this deterministic approximation deteriorates with coefficient of variation and horizon length, but that implementation of a rolling horizon re-optimization approach can significantly increase performance, at least in some cases (as described in the statement of Theorems 2.5.2 and 2.5.3).

We also explored the trade-off between (i) more demand information with less decision dynamics and (ii) more decision dynamics with less demand information. We found that using complete demand distribution information with limited decision dynamic (e.g. OLFC) has less value, while incorporating more decision dynamic via re-optimizing production decisions (e.g. CEC-IP) or using full decision dynamics with a slightly granular demand approximation beyond expected value (e.g. MPAC) potentially leads to a greater improvement. Given the estimation of the whole demand distribution is usually challenging in reality, this implies company should focus on including more decision dynamics into their production planning rather than estimating demand distribution more precisely. However, the benefit of using full decision dynamics (e.g. MPAC) comes at the expense of longer computational time.

Chapter 3

Outsourcing Strategy for Intermediate Production steps

3.1 Introduction

We consider an inventory control problem faced by a manufacturer that outsources the final step of its manufacturing process. This is motivated by our work with relatively small biopharmaceutical firms, which we observed often outsource their filling, labeling and packing of their final products to a third party contractor (which we refer to as the *outsourcer* in this chapter) to avoid the large capital investment associated with the equipment necessary for these operations. This requires the manufacturer to send intermediate products (which we refer to as an *order*) to the outsourcer for processing and wait for a certain amount of time until the outsourcer sends the order with finished product back. The time from the order placement until the items are delivered back to the manufacturer is called the *lead time* of an order. The manufacturer maintains inventory of finished product to meet demand from the end customer thus it is make-to-stock. On the other hand the outsourcer does not maintain any inventory and can only begin to process an order if one is initiated by the manufacturer, and therefore the outsourcer can be viewed as make-to-order.

In this chapter we compare the effectiveness of two inventory control policies at the manufacturer, both of which are widely adopted in industry. We explore the advantages of each policy, and consider coordination between the manufacturer and the outsourcer. The first policy (referred to as a *type (i) policy* in this chapter) is the classical (r, q) policy: the manufacturer continuously reviews the inventory of finished product, and when the inventory position (on-hand inventory - backorder + outstanding order) falls to a threshold r, an order of fixed size q is placed. Since order placement is driven by stochastic demand, the time between consecutive order placements is random. In addition, the outsourcer is in general shared by many other manufacturers, so orders placed via a type (i) policy need to join a queue at the outsourcer and wait for the outsourcer to finish processing orders for other manufacturers. This leads to intrinsically random order lead times, and biopharmaceutical manufacturers that adopt a type (i) policy indeed experience large uncertainty in lead times (Bu (2015)[10]).

To address this issue, some manufacturers turn to a capacity reservation (s, τ, u) policy (which we subsequently refer to as *type (ii) policy* in this chapter). In this policy, the manufacturer places an order every τ units of time that raises the inventory position as close to s as possible, given that the order size can not exceed u. Observe that order placement in a type (ii) policy is triggered by time, so even when the manufacturer is facing stochastic demand, the time between consecutive order placements is always deterministic at τ. In this setting, the manufacturer signs a contract with the outsourcer that reserves a capacity of u at the outsourcer every τ time units. Under a type (ii) policy, most of the uncertainty in the lead time is eliminated since the outsourcer is always immediately available when an order is placed. However this benefit comes with cost: (1) the manufacturer loses the flexibility to place an order at any time; (2) the order size is restricted by an upper limit; and (3) by the nature of capacity reservation contracts, the manufacturer needs to pay for the full capacity reserved, regardless of the actual order size.

To better analyze the trade-offs between type (i) and (ii) policies, we combine traditional inventory control models with a queuing approach to model the make-to-order outsourcer, and compare both policies both analytically and computationally. To better position our work in the existing literature, we briefly review the related topics in Section 3.2.

3.2 Literature Review

The type (i) policy has been extensively investigated in the literature. The use of an (r, q) policy in a continuous time model with Poisson demand dates back to the work of Galliher, Morse and Simond (1959)[29] and Hadley and Whitin (1963)[32], while the optimality of the (r, q) policy under a deterministic lead time is shown by Veinott (1965)[64], the convexity of the long-run average cost in both q and r is shown in Zipkin (1986)[66], and a distribution free upper bound on the long-run average cost is proposed by Gallego (1998)[27]. All of these models assume constant lead time. Zipkin (1986)[67] extends the continuous time model to allow random lead time. A key assumption in the model is that orders never cross in time, so the sequence of order placement is the same as the sequence of order delivery. Under this assumption, Zipkin shows that most results for the deterministic lead time case apply to the random lead time case with little modification. Later, Song (1994)[56] and Song et al. (2010)[57] study how randomness in the lead time affects long-run average cost and optimal r and q. They show that a stochastically larger lead time always results in higher long-run average cost.

In contrast, inventory control policies under which order placement is triggered by time received significantly less attention in the literature until early the 2000's. Rao (2003)[48] analyzes the (s, τ) policy which is a special case of the type (ii) policy when the capacity

u infinite. Under the (s, τ) policy, an order is placed every τ time to raise the inventory position back to s. Rao shows that the long-run average cost of the (s, τ) policy is always worse than that of the (r, q) policy, which demonstrates the benefit of (r, q) policy when the lead time is deterministic.

However, we are not aware of any work that addresses the core benefit of a type (ii) policy – the reduction in lead time variability, and the implication of this reduction for coordination of a make-to-order outsourcer and make-to-stock manufactures. In our model, we are able to show that when the lead time faced by a type (i) manufacturer is highly variable, the adoption of type (ii) policy can not only reduce the inventory cost of the manufacturer, but also increase the profit of the outsourcer.

3.3 Model and Assumptions

We consider a setting where a single make-to-order outsourcer processes orders from n make-to-stock manufacturers. Each manufacturer faces stochastic demand with arrival rate λ_i (see a discussion of the demand process in the next paragraph). The outsourcer manages a general queue to handle all orders placed by manufacturers. The aggregate arrival rate of orders from all manufacturers seen by the outsourcer is denoted by λ_0 and the outsourcer processes orders with service rate $\mu_0 > \lambda_0$. There are two types of manufacturers, distinguished by whether they use a type (i) or a type (ii) inventory control policy. An illustration of this structure is shown in Figure 3.1.

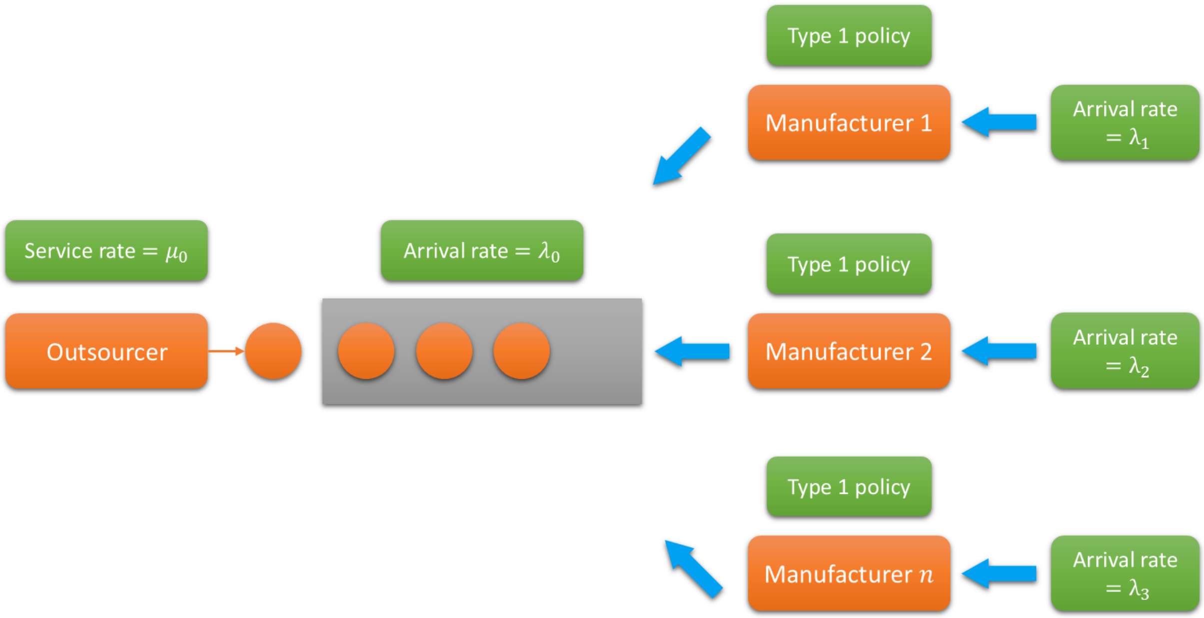

Figure 3.1: Structure of the outsourcer and manufacturers

We introduce the detailed model and assumptions below:

Demand process: By convention, we assume that demand processes faced by all manufacturers are independent stochastic processes and each process has independent and stationary increments. Let $D(t)$ be the cumulative demand occurring during a fixed period of time t, with p.d.f. $f_D(d,t)$ and c.d.f. $F_D(d,t)$. In this chapter, we restrict our consideration to two types of demand processes:

D1. $D(t)$ is a Poisson process with rate λt:

$$f_D(d,t) = \frac{(\lambda t)^d e^{-\lambda t}}{d!}$$

D2. $D(t)$ is a Brownian motion with drift with mean λt and variance $\sigma_D^2 t$:

$$f_D(d,t) = \phi\left(\frac{d - \lambda t}{\sigma_D \sqrt{t}}\right)$$

where $\phi(\cdot)$ is the p.d.f. of the standard normal distribution.

Though the demand processes D2 does not guarantee that $D(t)$ increases, when the time between consecutive order placements is sufficiently large, or demand rate λ is sufficiently larger than σ_D, D2 is non-decreasing with high probability (see Rao (2003)[48] for a discussion). In our analysis we primarily focus on demand process D1, but some of our results are applicable to D2 as well.

Manufacturer's policy: Under a continuous time demand process, type (i) manufacturers adopts the classical (r, q) policy: the manufacturer places an order of size q whenever the inventory position of the finished product at the manufacturer drops to r. Type (ii) manufacturers adopts the (s, τ, u) policy: the manufacturer places an order every τ units of time to raise the inventory position of finished product as close to s as possible, subject to a capacity constraint that the order size cannot exceed u. This type of policy is easy-to-implement and widely used in industry. Under both policies, we assume that unsatisfied demand is fully back-ordered.

Outsourcer's policy: We assume the outsourcer manages a queue under steady state and enforce a special property such that orders never cross in time. By this we mean that for a particular manufacturer the sequence of order placement is exactly the same as the sequence of order delivery. An order placed later in time never arrives before an earlier order. The majority of outsourcing activities in industry satisfy this property, which allows tractable analysis for type (i) policies (see Zipkin (1986)[67]).

Lead time process: The lead time of an order placed by a type (i) manufacturer is the order sojourn time (waiting time in the queue + processing time) at the outsourcer's facility. The lead time for type (i) manufacturers is intrinsically a random variable, denoted by L with mean μ_L, variance σ_L^2 and distribution $F_L(l)$. We assume the manufacturer can always estimate μ_L and σ_L^2 from historical data, but we in general assume the exact distribution of $F_L(l)$ is unknown. For type (ii) manufacturers, the lead time of an order is only the processing time at the outsourcer's facility since no order is ever waiting in the queue. In this chapter we assume the lead time faced by type (ii) manufacturers, denoted as l, is deterministic, which is a linear function of the reserved capacity u. It is equivalent to have an outsourcer with constant processing rate μ_0, so the lead time is $l = u/\mu_0$. Notice the lead time l must be less than or equal to the order cycle τ to have the type (ii) policy feasible.

Manufacturers' profit function: Both type (i) and (ii) manufacturers earn revenue from selling finished product at a fixed price γ per unit. Type (i) manufacturers cost has three components: (1) fixed cost k of placing an order regardless of the order size; (2) variable cost c per unit of an order; and (3) holding costs h per unit per unit time for on hand inventory and penalty cost p per unit per unit time for backorder. Type (ii) manufacturer's cost is similar, except the the variable cost c depends on the capacity reserved u, not the actual order size.

Outsourcer's profit function: The outsourcer's revenue comes from the k and c getting paid by the manufacturers, and the outsourcer is subject to three types of cost: (1) fixed cost k_0 of processing an order regardless of the order size; (2) variable cost c_0 per unit for an order; (3) waiting cost h_0 per unit per unit time for orders waiting to be processed. The waiting cost captures the cost of holding work-in-progress material at the outsourcer. By the nature of the type (ii) policy, the outsourcer does not incur any waiting cost.

Goal: Both the manufacturers and outsourcer focus on maximizing long-run average profit. We consider three scenarios in subsequent sections depending on the level of information sharing between the outsourcer and manufacturers:

- **The Decentralized System** (Sections 3.4 and 3.5)

  To reflect the situation faced by biopharmaceutical manufacturers, where other manufacturers that utilize the same outsourcer may come from completely different industries, in this setting we assume that the system is decentralized in a sense that the manufacturer manages its inventory without exact knowledge of the outsourcer and other manufacturers' control decisions. In addition, we assume the inventory control decisions made by a single manufacturer do not affect the steady state performance of the outsourcer's queue. Therefore, we focus solely on profit maximization at a single manufacturer. In Section 3.4, we assume there is no penalty cost associated with backorder, but instead a service level constraint is imposed. In Section 3.5, we remove the service level constraint and instead a penalty cost is incurred when backorder arises.

- **The Centralized System** (Section 3.6)

  In this setting, we assume that the outsourcer and all manufacturers belong to the same organization, so all information is shared. This scenario broadens the application of our model to a manufacturer that has a machine or special facility that serves as an "internal outsourcer", which is a shared resource for other manufacturing processes. The goal here is to maximize the long-run average profit of the entire system. Additional assumptions are discussed in detail in Section 3.6.

- **The Semi-centralized System** (Section 3.7)

  In this scenario, all manufacturers belong to the same organization, but no information is shared between this organization and the outsourcer. This setting models a manufacturer who has multiple plants that use the same outsourcer. In this setting, we assume that the outsourcer makes its control decisions first, and then each manufacturer follows and makes the corresponding optimal inventory control decisions. Additional assumptions are discussed in detail in Section 3.7.

3.4 The Decentralized System With Service Level Constraint

We first consider the decisions faced by the manufacturer in the decentralized setting. In this section, we assume no penalty cost is incurred for backorder, but instead an α-service level constraint is enforced, so that in the long-run only $1 - \alpha$ percent of the order cycles have backorder. To facilitate analysis, we also assume the inventory cost h is incurred even if the inventory level is negative (or equivalently, manufacturers earns revenue with stock-out).

This is a commonly used simplification that can greatly simplify the analysis and is known to be accurate if the service level α is close to one. In the next two subsections, we analyze the manufacturer's long-run average profit for each of the two policy types in this setting.

3.4.1 Type (i) Policy

The manufacturer implements a continuous time (r, q) policy inventory model for a setting with random lead time, were r is the re-order point and q is the order quantity. We adopt the following notation:

Notations:

γ	Selling price.
λ	Demand rate.
h	Holding cost.
α	Service level.
k	Fixed cost of an order.
c	Variable cost of an order.
L	Lead time.
	This is a random variable with c.d.ff F_L, mean μ_L and variance $\sigma_L^2 < \infty$.
$D(t)$	Demand during an interval of fixed length t.
	This is a random variable with c.d.f. $F_D(d, t)$.
$D(T)$	Demand during an interval of random length T.
	This is a random variable with c.d.f. $F_D(d, T)$.
$y^{(1)}$	Long-run average cost of the type (i) policy.
$z^{(1)}$	Long-run average profit of the type (i) policy.

Decision variables:

r	Reorder point.
q	Order quantity.

Given this notation, the long-run average profit of the service level model can be written as an optimization problem:

$$\max \quad z^{(1)}(r, q) = \lambda\gamma - \frac{k\lambda}{q} - \lambda c - h\left(r - \lambda\mu_L + \frac{q}{2}\right) \tag{3.1}$$

$$s.t. \quad F_D(r, L) \geq \alpha \tag{3.2}$$

where the first term $\lambda\gamma$ is the revenue from selling finished product at price γ. The second and third terms $\frac{k\lambda}{q} + \lambda c$ are the fixed and variable cost of an order, respectively. The fourth term $h(r - \lambda\mu_L + \frac{q}{2})$ is the average inventory holding cost. $F_D(r, L)$ is the distribution of lead time demand, the demand that occurs during a time of random length L. The constraint $F_D(r, L) \geq \alpha$ is the service-level constraint, which ensures that the probability of stock-out in an order cycle is less than $1 - \alpha$. It is easy to see that this profit maximization model is equivalent to a cost minimization model without the revenue $\lambda\gamma$ and variable cost λc terms, so that the service level model can be written as an minimization problem with long-run average cost $y^{(1)}(r,q)$

$$\begin{aligned} \min \quad & y^{(1)}(r,q) = h\left(r - \lambda\mu_L + \frac{q}{2}\right) + \frac{k\lambda}{q} && (3.3)\\ s.t. \quad & F_D(r,L) \geq \alpha && (3.4) \end{aligned}$$

Since the long-run average cost $y^{(1)}(r,q)$ is monotonically increasing in r, the optimal r is attained at $r_* = F_D^{-1}(\alpha)$. Rewriting $t = q/\lambda$, we have

$$\min \quad y^{(1)}(t) = h(r_* - \lambda\mu_L) + \frac{h\lambda}{2}t + \frac{k}{t}. \tag{3.5}$$

Since r_* is independent of t, we have

$$\begin{aligned} r_* &= F_D^{-1}(\alpha) && (3.6)\\ t_* &= \sqrt{\frac{2k}{h\lambda}} && (3.7)\\ y_*^{(1)} &= h(r_* - \lambda\mu_L) + \sqrt{2kh\lambda} && (3.8) \end{aligned}$$

To determine the optimal r_*, we need to analyze the lead time demand distribution $F_D(r,L)$. We next consider $F_D(r,L)$ in three cases.

Case 1: The outsourcer is modeled by a G/D/∞ queue
Though this case ignores many practical issues, the simplest model assumes the outsourcer side is G/D/∞ queue with service rate $\mu_0 = 1/\mu_L$. From the manufacturer's perspective, this equivalent to a deterministic lead time μ_L, so lead time demand is simply $D(\mu_L)$. The long-run average cost in this case, denoted by $\underline{y}^{(1)}(r,q)$, serves as a lower bound on $y^{(1)}(r,q)$ when the lead time is random.

If the demand process $D(t)$ is Poisson process D1, we can approximate the demand process with a normal distribution, then the optimal reorder point $\underline{r}_*$ and cost $\underline{y}_*^{(1)}$ is

$$\begin{aligned} \underline{r}_* &= \lambda\mu_L + \zeta_\alpha\sqrt{\lambda\mu_L} && (3.9)\\ \underline{y}_*^{(1)} &= \zeta_\alpha h\sqrt{\lambda\mu_L} + \sqrt{2kh\lambda}. && (3.10) \end{aligned}$$

If the demand process $D(t)$ is the drifted Brownian motion D2 with variance $\sigma_D^2 t$ then

$$\begin{aligned} \underline{r}_* &= \lambda\mu_L + \zeta_\alpha\sigma_D\sqrt{\mu_L} && (3.11)\\ \underline{y}_*^{(1)} &= \zeta_\alpha h\sigma_D\sqrt{\mu_L} + \sqrt{2kh\lambda} && (3.12) \end{aligned}$$

where $\zeta_\alpha = \Phi^{-1}(\alpha)$ and $\Phi(\cdot)$ is the c.d.f. of the standard Normal distribution.

Case 2: The outsourcer is modeled by an M/G/1 queue
In this special case, the exact distribution of the lead time demand can be obtained. To understand the relevance of this case, suppose that outsourcer signs contracts with n manufacturers. Each manufacturer faces Poisson demand, and manages inventory using (potentially different) (r, q) policies, so order arrivals from a single manufacturer form a renewal process with an inter-arrival time that follows the Erlang distribution. When n is large, the arrival process seen by the outsourcer is approximately a Poisson process (Palm–Khintchine theorem, see Khintchine (1969)[40]). If we assume the service time has a general distribution with mean $1/\mu_0$, then the outsourcer manages an M/G/1 queue and the lead time faced by a manufacturer is the sojourn time of the corresponding order in an M/G/1 queuing system.

In the special case of an M/M/1 queue, where the processing time at the outsourcer is exponentially distributed, the sojourn time is known to be exponentially distributed as well with rate $\beta = 1/\mu_L$. If the demand process faced by the manufacturer is the Poisson process D1, then the distribution of the lead time demand can be computed exactly as:

$$\mathbf{P}(D(L) = d) = \int \mathbf{P}(D(l) = d | L = l) f_L(l)\, dl \tag{3.13}$$

$$= \int_0^\infty \frac{(\lambda l)^d e^{-\lambda l}}{d!} \beta e^{-\beta l}\, dl \tag{3.14}$$

$$= \left(\frac{\lambda}{\lambda + \beta}\right)^d \left(\frac{\beta}{\lambda + \beta}\right) \tag{3.15}$$

which is the geometric distribution with parameter $\beta/(\lambda + \beta)$. Then

$$F_D(r, L) = \mathbf{P}(D(L) \leq r) = 1 - \left(\frac{\lambda}{\lambda + \beta}\right)^{r+1} = 1 - \left(\frac{\lambda \mu_L}{1 + \lambda \mu_L}\right)^{r+1} \tag{3.16}$$

Solving for r_* yields

$$r_* = \left\lfloor \frac{\log(1 - \alpha)}{\log(\frac{\lambda \mu_L}{1 + \lambda \mu_L})} \right\rfloor \tag{3.17}$$

In Section 3.6 and 3.7, we examine another special case, where the outsourcer is modeled by an M/D/1 queue, implying that each order has a deterministic processing time of $1/\mu_0$ at the outsourcer. The mean sojourn time of an M/D/1 queue is:

$$\mu_L = \frac{2\mu_0 - \lambda_0}{2\mu_0(\mu_0 - \lambda_0)}. \tag{3.18}$$

The distribution of the waiting time of an order in the M/D/1 queue, $\mathbf{P}(W < t)$, is given by Oliver (1968)[45] as

$$\mathbf{P}(W < t) = (1-\rho)\sum_{j=0}^{\lfloor t\mu_0 \rfloor} \frac{(\rho j - \lambda_0 t)^j}{j!} e^{-(\rho j - \lambda_0 t)} \quad t \geq 0 \tag{3.19}$$

$$\tag{3.20}$$

where $\rho = \lambda_0/\mu_0$. Then, the distribution of the sojourn time $\mathbf{P}(L < l)$ is

$$\mathbf{P}(L < l) \;=\; \begin{cases} (1-\rho)\sum_{j=0}^{\lfloor l\mu_0 - 1 \rfloor} \frac{(\rho j - \lambda_0 l + \rho)^j}{j!} e^{-(\rho j - \lambda_0 l + \rho)} & l \geq \frac{1}{\mu_0} \\ 0 & l < \frac{1}{\mu_0}. \end{cases} \tag{3.21}$$

The distribution of the lead time demand $\mathbf{P}(D(L) = d)$ can be computed numerically, then the optimal r_* can be found accordingly. However unlike the M/M/1 queue, the distribution of lead time demand explicitly requires the knowledge of μ_0 and λ_0, which is not available in our decentralized setting. We explore this M/D/1 queue more closely in conjunction with our centralized and semi-centralized settings in Sections 3.6 and 3.7.

Remark 3.4.1. (On the computation of the distribution of W) Computation of the distribution $\mathbf{P}(W < t)$ is not numerically stable due to possible overflow of $j!$ and $e^{\lambda_0 t}$ terms. Oliver (1968)[45] proposed a numerically stable algorithm, but there appear to be typographical errors in that presentation, so for completeness we present the correct algorithm here: we have

$$\mathbf{P}(W < t) = (1-\rho)\sum_{j=0}^{\lfloor t\mu_0 \rfloor} \frac{(\rho j - \lambda_0 t)^j}{j!} e^{-(\rho j - \lambda_0 t)} \quad t \geq 0. \tag{3.22}$$

Let $k = \lfloor t\mu_0 \rfloor$ and $\epsilon = t\mu_0 - \lfloor t\mu_0 \rfloor$. Then, $t = k/\mu_0 + \epsilon/\mu_0$ and if we let $W_k(\epsilon) = \mathbf{P}(W < t)$ we have

$$W_k(\epsilon) \;=\; (1-\rho)e^{\lambda_0 k/\mu_0 + \lambda_0 \epsilon/\mu_0} \sum_{j=0}^{k} \frac{(\rho j - \lambda_0 k/\mu_0 - \lambda_0 \epsilon/\mu_0)^j}{j!} e^{-\rho j} \tag{3.23}$$

$$\;=\; (1-\rho)e^{\rho k + \rho\epsilon} \sum_{j=0}^{k} \frac{(\rho j - \rho k - \rho\epsilon)^j}{j!} e^{-\rho j} \tag{3.24}$$

$$\;=\; (1-\rho)e^{\rho(k+\epsilon)} \sum_{j=0}^{k} \frac{(\rho e^{-\rho}(j - k - \epsilon))^j}{j!} \tag{3.25}$$

With a binomial expansion,we can write $e^{-\epsilon}W_k(\epsilon)$ as a polynomial in ϵ, so

$$W_k(\epsilon) = e^{\rho\epsilon}\sum_{j=0}^{k}\alpha_{k,j}(\rho\epsilon)^j \tag{3.26}$$

where $\alpha_{k,j}$ are the polynomial coefficients that can be computed iteratively as

$$\alpha_{0,0} = (1-\rho) \tag{3.27}$$

$$\alpha_{k+1,j} = -\frac{1}{j}\alpha_{k,j-1} \tag{3.28}$$

$$\alpha_{k+1,0} = e^{\rho}\sum_{j=0}^{k}\alpha_{k,j}\rho^j. \tag{3.29}$$

Case 3: The outsourcer is modeled by a general queuing system

If the queue managed by the outsourcer is far from an M/G/1 queue, the exact distribution of lead time F_L faced by the manufacturer typically depends on the queuing control policy at the outsourcer as well as the inventory control decisions of all of other manufacturers, which are impossible for an individual manufacturer to know under our assumption of a decentralized system. This models settings may happen if, for example, the total number of manufacturers n is small, or where the service time is not exponential or deterministic, if the outsourcer is served by a single queue with multi-class arrival and class dependent service time, or if the outsourcer manages the queue using other disciplines, such as a multi-queue polling system, or if the outsourcer has both type (i) and type (ii) manufacturers. If this is the case, our goal is to bound the long-run average cost $y^{(1)}(r,q)$ using only the first two moments of the lead time, which can be easily estimated using historical data.

W.l.o.g suppose that the lead time L has finite support on $1,\ldots,m$ with $\mathbf{P}(L=l)=x_l$ for $l=1,\ldots,m$ such that

$$\mathbf{E}[L] = \sum_{l=1}^{m} l x_l = \mu_L \tag{3.30}$$

$$\mathbf{Var}(L) = \sum_{l=1}^{m} l^2 x_l - (\sum_{l=1}^{m} l x_l)^2 = \sigma_L^2 < \infty \tag{3.31}$$

The c.d.f. of the lead time demand conditioning on $L=l$ is

$$\mathbf{P}(D(l)\le d\,|\,L=l) \sim F_D(d,l) \tag{3.32}$$

and then the unconditional probability

$$\mathbf{P}(D(L)\le d) = \sum_{l=1}^{m} F_D(d,l)x_l. \tag{3.33}$$

Given a $\delta \geq 0$, we try to minimize $\mathbf{P}(D(L) \leq \lambda\mu_L + \delta)$ by solving the following parametric (in δ) linear program:

$$\begin{aligned} \min \quad & \sum_{l=1}^{m} x_l F_D(\lambda\mu_L + \delta, l) && (3.34) \\ s.t. \quad & \sum_{l=0}^{m} l x_l = \mu_L && (3.35) \\ & \sum_{l=0}^{m} l^2 x_l = \sigma_L^2 + \mu_L^2 && (3.36) \\ & \sum_{l=0}^{m} x_l = 1 && (3.37) \\ & x_l \geq 0 \quad \forall l = 0, 1, \ldots, m && (3.38) \end{aligned}$$

The optimal x_l represents the worst lead time distribution that minimizes the service level. Next, we want to find the smallest δ_* such that $\mathbf{P}(D(L) \leq \lambda\mu_L + \delta^*) = \alpha$. Since $F_D(\lambda\mu_L + \delta, l)$ is monotonic in δ, the optimal δ_* that achieves the α-service level can be found very quickly via bisection search. Then $\bar{r} = \lambda\mu_L + \delta_*$ is the upper bound of the re-order point and the upper bound on the long-run average cost $\bar{y}^{(1)}$ is $\bar{y}^{(1)} = h(\bar{r} - \lambda\mu_L) + \sqrt{2kh\lambda}$.

Remark 3.4.2. (On the tightness of $\bar{y}^{(1)}$) Several remarks on the upper bound $\bar{y}^{(1)}$: (1) The bound is tight. When $\sigma_L = 0$, $\bar{y}^{(1)}$ reduces to the long-run average cost $\underline{y}^{(1)}(r, q)$ in Case 1 where the lead time is deterministic. When $\sigma_L > 0$, there always exists a lead time distribution such that the bound is achieved; (2) The bound does not require the demand process to be a Poisson process. Since only the conditional distribution $F_D(d, l)$ is required in the parametric linear program, this bound also applies to the demand process D2; (3) We can obtain another upper bound by using Chebyshev's inequality. To do so, observe that we can compute the mean and variance of the lead time demand $D(L)$ as

$$\begin{aligned} \mathbf{E}[D(L)] &= \mathbf{E}[\mathbf{E}[D|L]] = \lambda\mu_L && (3.39) \\ \mathbf{Var}(D(L)) &= \mathbf{E}[\mathbf{Var}[D|L]] + \mathbf{Var}[\mathbf{E}[D|L]] = \mathbf{E}[\mathbf{Var}[D|L]] + \lambda^2\sigma_L^2 && (3.40) \end{aligned}$$

and the long-run average cost can be roughly bounded using the one-sided Chebyshev's inequality

$$\mathbf{P}(D(L) \geq r) = \mathbf{P}(D(L) \geq \lambda\mu_L + \delta) \leq \frac{\mathbf{Var}(D(L))}{\mathbf{Var}(D(L)) + \delta^2} \leq 1 - \alpha \quad (3.41)$$

which implies

$$r = \lambda\mu_L + \delta = \lambda\mu_L + \sqrt{\mathbf{Var}(D(L))}\sqrt{\frac{\alpha}{1-\alpha}}. \quad (3.42)$$

For a Poisson demand process $\mathbf{Var}(D(L)) = \lambda\mu_L + \lambda^2\sigma_L^2$, then plug it in we have

$$\bar{r} = \lambda\mu_L + \sqrt{\lambda\mu_L + \lambda^2\sigma_L^2}\sqrt{\frac{\alpha}{1-\alpha}} \tag{3.43}$$

$$\bar{y}^{(1)} = h\sqrt{\lambda\mu_L + \lambda^2\sigma_L^2}\sqrt{\frac{\alpha}{1-\alpha}} + \sqrt{2kh\lambda}. \tag{3.44}$$

This bound is in general very loose, but it has a closed form for both $\bar{r}$ and $\bar{y}^{(1)}$.

Example 3.4.1. Assume that the demand process is a Poisson process and let $\lambda = 10, \mu_L = 5$, and $\sigma_L = 5$. With $\alpha = 0.98$, the bound obtained via the parametric linear program (3.34) is $\delta_* = 307.2$ while the Chebyshev bound is

$$\delta = \sqrt{50 + 100 \cdot 25}\sqrt{0.98/0.02} = 353.48.$$

Our proposed bound is tighter.

□

3.4.2 Type (ii) Policy

In this section we consider the type (ii) policy capacity reservation policy. As with the type (i) policy, holding and penalty cost is incurred continuously in time.We adopt the following notation:

Notations:

γ	Selling revenue.
λ	Demand rate.
h	Holding cost.
α:	Service level.
k	Fixed cost.
c	Variable cost.
l	Lead time.
$D(t)$	Demand during an interval of fixed length t. This is a random variable.
$y^{(2)}$	Long-run average cost.
$z^{(2)}$	Long-run average profit.

Decision variables:

s	Order-up-to level.
τ	Cycle time.
u	Reserved capacity.

Recall that the type (ii) manufacturer manages inventory using a (s, τ, u) policy: place an order every τ time units that raises the inventory position as close to s as possible, subject to a constraint u on maximum order size. Mean demand during a cycle of length τ is $\lambda\tau$. An illustration of type (ii) policy with a sample path of inventory position and inventory level are shown in Figure 3.2

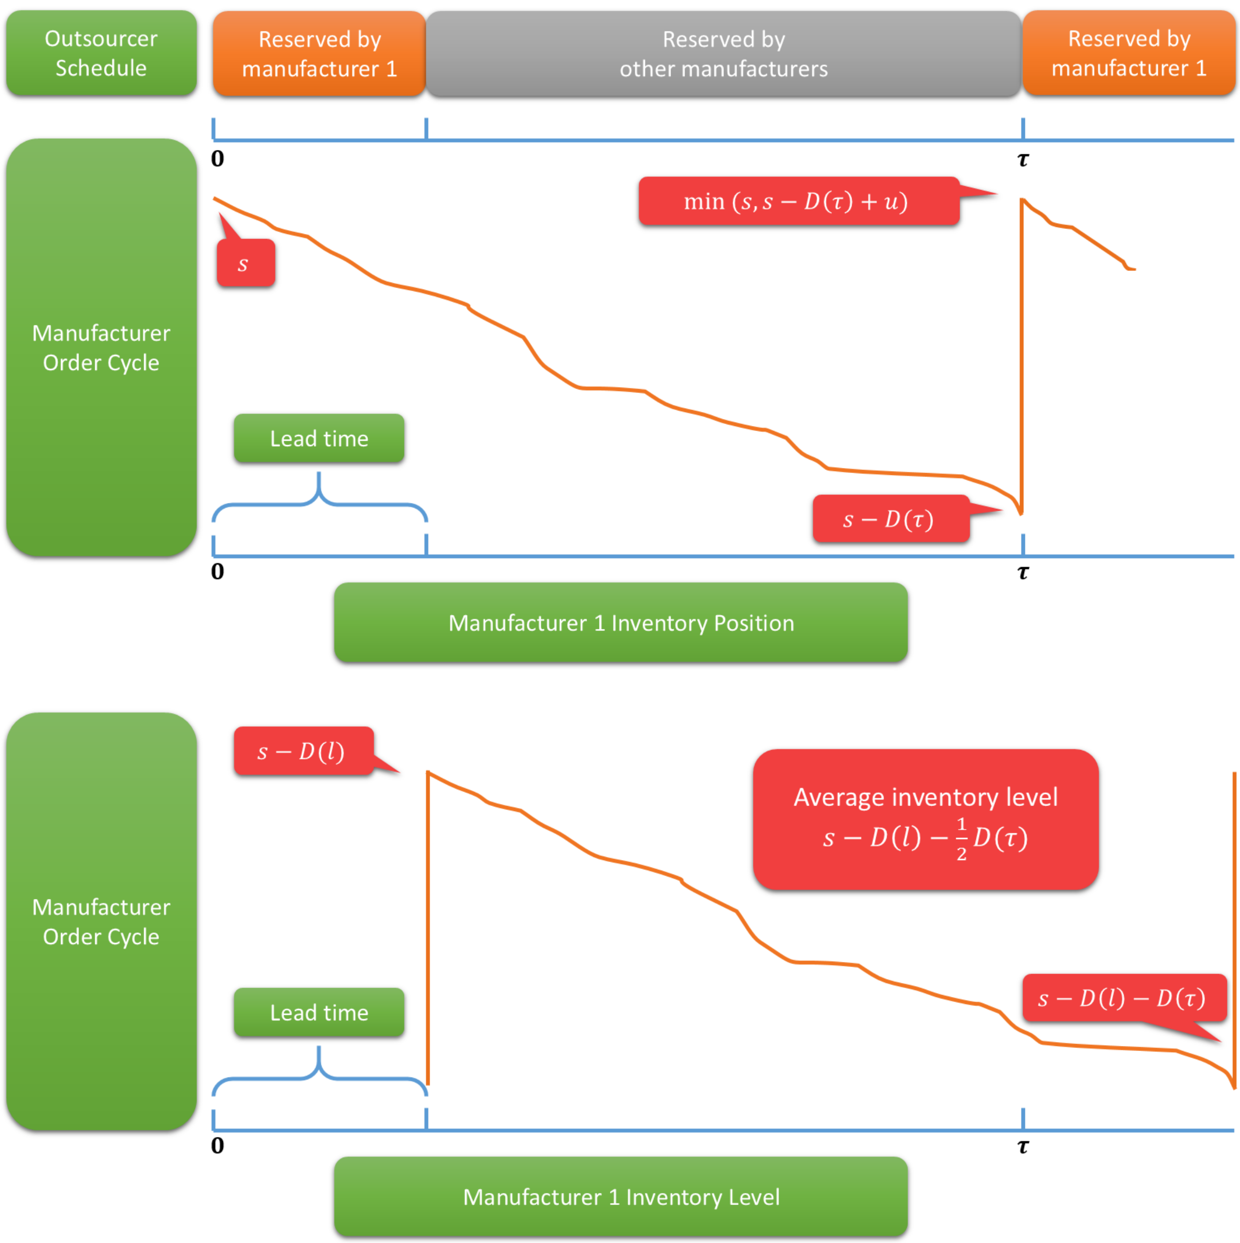

Figure 3.2: A sample path of type (ii) policy

In Figure 3.2 we assume the inventory position at the beginning of the cycle is s, and the inventory position at the end of the cycle just before placing an order is $s - D(\tau)$. Then the inventory position at the beginning of the next cycle is $\min\{s, s - D(\tau) + u\}$ since the order size can not exceed the capacity u. Recall that we assume no lead time for the order to begin processing (since capacity is reserved), and a fixed, deterministic lead time l after the order is placed until it arrives, independent of order size. On the other hand, inventory level at the beginning of the cycle is $s - D(l)$, and the inventory level at the end of the cycle just before the next order arrives is $s - D(l) - D(\tau)$. Then the average inventory level within the cycle is $s - D(l) - \frac{1}{2}D(\tau)$ so the inventory holding cost can be computed. Next we consider two cases depending on the variable cost c.

Case 1: $c = 0$
First consider a special case where the variable cost c is 0, so given the specifications of our model, the manufacturer will reserve $u = \infty$ capacity under a (s, τ, u) policy. As a consequence the inventory position at the beginning of each ordering cycle will always be s. The long-run average profit can be expressed as the following constrained optimization problem:

$$\max \quad z^{(2)}(s, \tau, \infty) = \lambda\gamma - \frac{k}{\tau} - h\left(s - \lambda l - \frac{1}{2}\lambda\tau\right) \tag{3.45}$$

$$s.t. \quad \mathbf{P}(D(l+\tau) \le s) > \alpha \tag{3.46}$$

Here the first term $\lambda\gamma$ is the revenue, the second term $\frac{k}{\tau}$ is the fixed cost, the third term $h(s - \lambda l - \frac{1}{2}\lambda\tau)$ is the average holding cost and the constraint $\mathbf{P}(D(l+\tau) \le s) > \alpha$ is the service level constraint. This profit maximization model is equivalent to a cost minimization model without the revenue term, which can be written as follows:

$$\min \quad y^{(2)}(s, \tau, \infty) = \frac{k}{\tau} + h\left(s - \lambda l - \frac{1}{2}\lambda\tau\right) \tag{3.47}$$

$$s.t. \quad \mathbf{P}(D(l+\tau) \le s) > \alpha \tag{3.48}$$

The optimal order-up-to level s is determined uniquely by $D(l+\tau)$. Approximating the Poisson demand using Normal distribution, we get $s_* = \lambda(l+\tau) + \zeta_\alpha\sqrt{\lambda(l+\tau)}$. Substituting back into the objective function, we get

$$y^{(2)}(\tau) = \frac{1}{2}h\lambda\tau + h\zeta_\alpha\sqrt{\lambda}\sqrt{l+\tau} + \frac{k}{\tau} \tag{3.49}$$

so then

$$\tau_* = \arg\min y^{(2)}(\tau) \tag{3.50}$$

$$s_* = \lambda(l+\tau_*) + \zeta_\alpha\sqrt{\lambda}\sqrt{l+\tau_*} \tag{3.51}$$

$$y_*^{(2)} = \frac{1}{2}h\lambda\tau_* + h\zeta_\alpha\sqrt{\lambda}\sqrt{l+\tau_*} + \frac{k}{\tau_*} \tag{3.52}$$

Unfortunately, the closed form expression for τ_* is complex. A simple lower bound on $y^{(2)}(\tau)$ can be obtained by setting the lead time $l = 0$, so $\underline{y}^{(2)}(\tau) = \frac{1}{2}h\lambda\tau_* + h\zeta_\alpha\sqrt{\lambda}\sqrt{\tau_*} + \frac{k}{\tau_*}$.

Case 2: $c > 0$
When the variable cost c is positive, reserved capacity u will be finite. Let S_t be the starting inventory position at the beginning of the t^{th} order cycle. The sequence of $\{S_t\}$ forms a discrete time Markov chain on discrete state $s, s-1, \ldots$ with state transition $S_{t+1} = \min\{s, S_t - D_t(\tau) + u\}$, as the demand during an order cycle is $D(\tau)$ so inventory position in the $(t+1)^{th}$ order cycle is the minimum between s and $S_t - D_t(\tau) + u$ given the order size cannot exceed u. The transition probability matrix is

$$\begin{bmatrix} \mathbf{P}(D(\tau) \leq u) & \mathbf{P}(D(\tau) = u+1) & \mathbf{P}(D(\tau) = u+2) & \mathbf{P}(D(\tau) = u+3) & \cdots \\ \mathbf{P}(D(\tau) \leq u-1) & \mathbf{P}(D(\tau) = u) & \mathbf{P}(D(\tau) = u+1) & \mathbf{P}(D(\tau) = u+2) & \cdots \\ \mathbf{P}(D(\tau) \leq u-2) & \mathbf{P}(D(\tau) = u-1) & \mathbf{P}(D(\tau) = u) & \mathbf{P}(D(\tau) = u+1) & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix} \quad (3.53)$$

If $u > \lambda\tau$, the Markov chain is recurrent so the stationary distribution of S_t exists. Let S_∞ be a random variable denoting the stationary inventory position at the beginning of an order cycle. Observe that the distribution of S_∞ is a function of s, τ and u. To ensure the α-service level, we need

$$\mathbf{P}\left(D(l+\tau) \leq S_\infty\right) > \alpha. \quad (3.54)$$

Given a particular cycle τ and capacity u, the minimum order-up-to level s_* which ensures that the corresponding S_∞ satisfies this service level constraint can be obtained. Let $\psi(s_*, \tau, u)$ be the expected value of S_∞, and then the long-run average profit $z^{(2)}$ is

$$z^{(2)}(s_*, \tau, u) = \lambda\gamma - \frac{k}{\tau} - \frac{cu}{\tau} - h\left(\psi(s_*, \tau, u) - \lambda l - \frac{1}{2}\lambda\tau\right) \quad (3.55)$$

which is equivalent to minimizing

$$y^{(2)}(s_*, \tau, u) = \frac{k}{\tau} + \frac{cu}{\tau} + h\left(\psi(s_*, \tau, u) - \lambda l - \frac{1}{2}\lambda\tau\right). \quad (3.56)$$

Clearly $\underline{y}^{(2)}(s, \tau, \infty)$ from Case 1 above when $u = \infty$ serves as a lower bound on $y^{(2)}(s_*, \tau, u < \infty)$.

3.4.3 Comparison Between Two Policies

In this section we compare the performance of type (i) and (ii) policies computationally by comparing the expected long-run cost y and profit z. Our goal is to highlight the performance of type (ii) policy when there is a large lead time uncertainty under type (i) policy. To do so, we consider a test problem with parameters $k = 500, \lambda = 10, \alpha = 0.98, h = 1, \gamma = 25, c = 5$, as we found in preliminary computation that other choices of these parameters do not materially impact our conclusions. Then, we fix the expected lead time of the type (i) policy at $\mu_L = 5$ (we will study the effect of changing in expected lead time in the centralized system in Section 3.6) and change the lead time variability σ_L, and compare the performance to the type (ii) policy with various lead time l and decision controls τ and u. The results are shown in Table 3.1.

Policy	Cases	Lead time	Parameters			Cost	Profit
			Changing σ_L				
(i)	Exact	$\sigma_L = 0$	$r_* = 64.52$	$t_* = 10$		114.52	85.48
(i)	Upper bound	$\sigma_L = 0.25$	$r_* = 65.72$	$t_* = 10$		115.72	84.28
(i)	Upper bound	$\sigma_L = 0.5$	$r_* = 71.08$	$t_* = 10$		121.08	78.92
(i)	Upper bound	$\sigma_L = 0.75$	$r_* = 84.72$	$t_* = 10$		134.72	65.28
(i)	Upper bound	$\sigma_L = 1$	$r_* = 99.58$	$t_* = 10$		149.58	50.42
(i)	Upper bound	$\sigma_L = 2$	$r_* = 162.02$	$t_* = 10$		212.02	-12.02
(i)	Upper bound	$\sigma_L = 3$	$r_* = 226.28$	$t_* = 10$		276.28	-76.28
(i)	Upper bound	$\sigma_L = 4$	$r_* = 291.50$	$t_* = 10$		341.50	-141.50
(i)	M/D/1 $\lambda_0 = 0.5$	$\sigma_L = 4.8$	$r_* = 166$	$t_* = 10$		216.00	-16.00
(i)	M/M/1	$\sigma_L = 5$	$r_* = 196.55$	$t_* = 10$		246.55	-46.55
(i)	Upper bound	$\sigma_L = 5$	$r_* = 357.20$	$t_* = 10$		407.20	-207.20
(i)	M/D/1 $\lambda_0 = 1$	$\sigma_L = 5.5$	$r_* = 180$	$t_* = 10$		230.00	-30.00
			Changing l				
(ii)	Exact	$l = 0$	$s_* = 110.26$	$\tau_* = 9.07$	$u = \infty$	120.04	-
(ii)	Exact	$l = 1$	$s_* = 121.78$	$\tau_* = 9.11$	$u = \infty$	121.08	-
(ii)	Exact	$l = 2$	$s_* = 133.18$	$\tau_* = 9.15$	$u = \infty$	122.08	-
(ii)	Exact	$l = 3$	$s_* = 144.49$	$\tau_* = 9.18$	$u = \infty$	123.03	-
(ii)	Exact	$l = 4$	$s_* = 155.71$	$\tau_* = 9.21$	$u = \infty$	123.94	-
(ii)	Exact	$l = 5$	$s_* = 166.87$	$\tau_* = 9.24$	$u = \infty$	124.82	-
			Changing τ				
(ii)	Exact	$l = 5$	$s_* = 175.15$	$\tau = 10$	$u = \infty$	125.15	-
(ii)	Exact	$l = 5$	$s_* = 229.04$	$\tau = 15$	$u = \infty$	137.37	-
(ii)	Exact	$l = 5$	$s_* = 282.47$	$\tau = 20$	$u = \infty$	157.47	-
(ii)	Exact	$l = 5$	$s_* = 335.57$	$\tau = 25$	$u = \infty$	180.57	-
(ii)	Exact	$l = 5$	$s_* = 388.42$	$\tau = 30$	$u = \infty$	205.09	-
			Changing u				
(ii)	Exact	$l = 5$	$s_* = 176$	$\tau = 10$	$u = 115$	125.61	66.89
(ii)	Exact	$l = 5$	$s_* = 179$	$\tau = 10$	$u = 110$	127.66	67.34
(ii)	Exact	$l = 5$	$s_* = 192$	$\tau = 10$	$u = 105$	136.57	60.93
(ii)	Exact	$l = 5$	$s_* = 246$	$\tau = 10$	$u = 102$	176.20	22.80

Table 3.1: Comparison of type (i) and (ii) for service level model policy

In the first section of the table, which explores the performance of the type (i) policy, when $\sigma_L = 0$, i.e. the lead time is deterministic, the optimal controls r_* and t_* can be computed exactly as in Case 1 in Section 3.4.1 when the outsourcer is modeled as a G/D/∞ queue and the corresponding cost/profit serves as a lower/upper bound. When $\sigma_L > 0$, when the exact distribution of the lead time can be obtained (M/M/1 and M/D/1 with known λ_0), the optimal controls r_* and t_* as well as the cost/profit are computed as in Case 2 in Section 3.4.1 where the outsourcer is modeled as an M/G/1 queue. In other instances where the exact distribution of the lead time is impossible to obtain, as in Case 3 in Section 3.4.1, the outsourcer is modeled as a general queue and the upper/lower bound of the cost/profit are computed. In the second section of the table, which explores the performance of the type (ii) policy, we study the performance by varying the type (ii) policy lead time l, the order cycle τ and the reserved capacity u. When $u = \infty$, the cost are computed according to Case 1 in Section 3.4.2. When $u < \infty$, the cost/profit are computed according to Case 2 in Section 3.4.2.

From Table 3.1, we can see that when the lead time is deterministic, type (i) policy has the lowest cost of 114.52 which outperforms the cost of type (ii) policy, even for the case of the lower bound with zero lead time $l = 0$ and infinite reserved capacity $u = \infty$. This phenomenon holds true as long as the optimal order cycle of the type (ii) policy τ_* is larger than the expected lead time of the type (i) policy μ_L, as shown in Theorem 3.4.1:

Theorem 3.4.1. *If $\tau_* \geq \mu_L$, the lower bound of type (ii) policy long-run average cost $\underline{y}^{(2)}$ is strictly larger than the type (i) policy long-run average cost with deterministic lead time $\underline{y}^{(1)}$.*

The proof of Theorem 3.4.1 can be found in Appendix A.2.1. Theorem 3.4.1 shows the superiority of the type (i) policy when the lead time is deterministic. However, for more realistic cases described in Case 2 and Case 3 in Section 3.4.1 when the lead time is random, type (ii) (s, τ, u) policy can dominate. As σ_L increases, the long-run average cost of the type (i) policy increases rapidly. For example in the instance which the outsourcer is modeled as an M/M/1 queue, the optimal long-run average cost of type (i) policy is 247.55, which is significantly larger than the type (ii) policy, even with suboptimal $\tau = 30 \approx 3\tau_*$ which leads to a long-run average cost of 205.09, or extremely constrained capacity $u = 102 = \lambda\tau + 2$, which leads to a long-run average cost of only 176.20. The intuition is that for type (i) manufacture, the reorder point r_* that ensures the α service level is mainly determined by the unlucky long lead time when the lead time variability is high. Higher lead time variability leads to higher r_*, which increases the long-run average cost. On the other hand, under type (ii) policy the lead time variability is eliminated so the cost is much stabler: the cost does not change much with different lead time from $l = 0$ to $l = 5$. In addition the long-run average cost when $u < \infty$ is close to the cost with infinite capacity as long as the capacity u is slightly larger than the mean demand in a cycle. This is because the probability of starting inventory position smaller than the order-up-to level s, $\mathbf{P}(S_\infty < s)$ decreases exponentially fast as u increases, which is shown in Theorem 3.4.2.

Theorem 3.4.2. *The probability of the inventory position at the beginning of an order cycle less than the order-up-to level, $\mathbf{P}(S_\infty < s)$ decreases exponentially as capacity u increases.*

The proof of Theorem 3.4.2 can be found in Appendix A.2.2. Theorem 3.4.2 suggests that the capacity constraint in the type (ii) policy is not very restrictive. If a manufacturer reserves capacity u that is slightly larger than the mean demand during a cycle $\lambda\tau$, the cost is close the case of infinite capacity.

Managerial Insight: The comparison between the type (i) and (ii) policies crucially depends on the lead time process at the outsourcer. When the lead time experienced by a type (i) manufacturer is deterministic or has very small variance, which is common if the outsourcer primarily works for the manufacturer (or even is dedicated to the manufacturer) so that no other manufacturers compete for the same outsourcer, then the type (i) policy yields lower cost. When the lead time experienced by a type (i) manufacturer is highly variable, which is common if the manufacturer is only a small customer of the outsourcer (e.g. the biopharmaceutical industry) so that most of the lead time is due to the waiting for the availability of the outsourcer, then it is better to switch to a type (ii) policy that reserves the outsourcer capacity in advance, which eliminates the variability in the lead time and therefore results in lower cost. In addition, under type (ii) policy, the capacity constraint is not restrictive as long as the reserved capacity u is slightly larger than the mean demand during a cycle $\lambda\tau$.

3.5 Decentralized System With Penalty Cost

In this section, we briefly analyze the same setting as in the previous section, except with a penalty cost rather than a service level constraint. Specifically, a penalty cost p per unit of time is incurred for each unit of backorder. We analyze both type (i) and type (ii) policies, and show that the same conclusion can be drawn as in the service level model.

3.5.1 Type (i) Policy

We adopt the following notation:

Notations:

γ	Selling price.
λ	Demand rate.
h	Holding cost.
p	Penalty cost.
k	Fixed cost of an order.
c	Variable cost of an order.
q	Order size.
L	Lead time.
	This is a random variable with c.d.f. F_L, mean μ_L and variance $\sigma_L^2 < \infty$.
$D(t)$	Demand during an interval of constant length t.
	This is a random variable with c.d.f. $F_D(d,t)$.
$D(T)$	demand during an interval of random length T. Random variable with c.d.f. $F_D(d,T)$.
$z^{(1)}(r,q)$	Long-run average profit.
$g(y)$	$= h\mathbf{E}[(y - D(L))^+] + p\mathbf{E}[(D(L) - y)^+]$
	$= (h+p)\sum_{j=0}^{y-1} F_D(j,L) + p\lambda\mu_L - py$

Decision variables:

r	Re-order level.
q	Order quantity.

If $F_D(d,L)$ and μ_L are known, the long-run average profit of the manufacturer is

$$z^{(1)}(r,q) = \lambda\gamma - \frac{k\lambda}{q} - \lambda c - \frac{1}{q}\sum_{y=r+1}^{r+q} g(y) \quad (3.57)$$

under the Poisson demand process D1, and minimizing this is equivalent to minimizing cost:

$$y^{(1)}(r,q) = \frac{k\lambda}{q} + \frac{1}{q}\sum_{y=r+1}^{r+q} g(y). \quad (3.58)$$

To compute $y^{(1)}(r,q)$ we need the exact distribution of lead time demand F_D, which in general is difficult to obtain. As in Section 3.4, we consider three cases:

Case 1: The outsourcer is modeled by a G/D/∞ queue
In this case the lead time is deterministic, so this serves as a general lower bound on the long-run average cost $y^{(1)}$. The optimal r and q can be computed efficiently but in general there is no closed form solution for the optimal r_*, q_* and $y^{(1)}(r_*, q_*)$. (see Federgruen and Zheng (1992)[23]).

Case 2: The outsourcer is modeled by an M/M/1 queue
As we have discussed in Section 3.4.1, when the outsourcer manages an M/M/1 queue, the lead time demand faced by the manufacturers $F_D(d,L)$ follows a geometric distribution with

parameter $\lambda/(\lambda+\beta)$:

$$F_D(j,L) \quad = \quad \mathbf{P}(D \leq j) = 1 - \left(\frac{\lambda}{\lambda+\beta}\right)^{j+1} \tag{3.59}$$

where $\beta = 1/\mu_L$, so

$$g(y) \quad = \quad (h+p)\sum_{j=0}^{y-1}\left(1-\left(\frac{\lambda}{\lambda+\beta}\right)^{j+1}\right) + \frac{p\lambda}{\beta} - py \tag{3.60}$$

$$= \quad hy - (h+p)\sum_{j=1}^{y}\left(\frac{\lambda}{\lambda+\beta}\right)^{j} + \frac{p\lambda}{\beta} \tag{3.61}$$

$$= \quad hy - (h+p)\frac{\lambda}{\beta}\left(1-\left(\frac{\lambda}{\lambda+\beta}\right)^{y}\right) + \frac{p\lambda}{\beta} \tag{3.62}$$

$$= \quad hy + (h+p)\frac{\lambda}{\beta}\left(\frac{\lambda}{\lambda+\beta}\right)^{y} - \frac{h\lambda}{\beta} \tag{3.63}$$

Then,

$$\sum_{y=r+1}^{y=r+q} g(y) \quad = \quad \frac{hq(2r+q+1)}{2} + (h+p)\left(\frac{\lambda}{\beta}\right)^{2}\left(\frac{\lambda}{\lambda+\beta}\right)^{r}\left(1-\left(\frac{\lambda}{\lambda+\beta}\right)^{q}\right) - \frac{\lambda hq}{\beta}$$

so for a given r and q, we have the explicit formula for $y^{(1)}(r,q)$:

$$y^{(1)}(r,q) \quad = \quad \frac{k\lambda}{q} + \frac{h(2r+q+1)}{2} + \frac{h+p}{q}\left(\frac{\lambda}{\beta}\right)^{2}\left(\frac{\lambda}{\lambda+\beta}\right)^{r}\left(1-\left(\frac{\lambda}{\lambda+\beta}\right)^{q}\right) - \frac{\lambda h}{\beta}.$$

Case 3: The outsourcer side is modeled by a general queuing system

As in Section 3.4.1, an upper bound on the long-run average cost $y^{(1)}(r,q)$ can be found using only the first two moments of the lead time μ_L, σ_L^2. We follow the approach in Gallego (1998)[27]. We first bound $\sum_{y=r+1}^{r+q} g(y)$ as

$$\sum_{y=r+1}^{r+q} g(y) \quad \leq \quad \int_r^{r+q} g(y)dy \tag{3.64}$$

$$= \quad \int_r^{r+q}\int_0^{\infty}\left[h(y-x)^+ + p(x-y)^+\right] dF_D(x,L)\, dy \tag{3.65}$$

$$= \quad \int_0^{\infty}\int_r^{r+q}\left[h(y-x)^+ + p(x-y)^+\right] dy\, dF_D(x,L) \tag{3.66}$$

$$\tag{3.67}$$

Let $s(x;q,r) = \int_r^{r+q}\left[h(y-x)^+ + p(x-y)^+\right] dy$; then to compute $s(x;q,r)$ we have three cases:

- if $x \leq r$

$$\int_r^{r+q} \left[h(y-x)^+ + p(x-y)^+\right] dy = h\int_r^{r+q}(y-x)dy \tag{3.68}$$
$$= -hq(x-r-\frac{1}{2}q) \tag{3.69}$$

- if $x \geq r+q$

$$\int_r^{r+q} \left[h(y-x)^+ + p(x-y)^+\right] dy = p\int_r^{r+q}(x-y)dy \tag{3.70}$$
$$= pq(x-r-\frac{1}{2}q) \tag{3.71}$$

- if $r+1 \leq x \leq r+q-1$

$$\int_r^{r+q} \left[h(y-x)^+ + p(x-y)^+\right] dy = p\int_r^{x}(x-y)dy + h\int_x^{r+q}(x-y)dy \tag{3.72}$$
$$= \frac{1}{2}p(x-r)^2 + \frac{1}{2}h(x-r-q)^2 \tag{3.73}$$

Combining these three cases, we see that

$$s(x;q,r) \leq \frac{1}{2}h(x-r-q)^2 + \frac{1}{2}p(x-r)^2 \tag{3.74}$$
$$\mathbf{E}[s(x;q,r)] \leq \frac{1}{2}(h+p)\sigma_D^2 + \frac{1}{2}h(\mu_D-r-q)^2 + \frac{1}{2}p(\mu_D-r)^2 \tag{3.75}$$

where

$$\mu_D = \lambda\mu_L \tag{3.76}$$
$$\sigma_D^2 = \lambda\mu_L + \lambda^2\sigma_L^2 \tag{3.77}$$

are the first two moments of the lead time demand as derived in the remark in Section 3.4. If we let $s(q,r) = \mathbf{E}[s(x;q,r)]$, we have

$$s(q,r) = \frac{1}{2}(h+p)(\lambda\mu_L + \lambda^2\sigma_L^2) + \frac{1}{2}h(\lambda\mu_L - r - q)^2 + \frac{1}{2}p(\lambda\mu_L - r)^2 \tag{3.78}$$
$$\bar{y}^{(1)}(q,r) = \frac{1}{q}\lambda k + \frac{1}{q}s(q,r). \tag{3.79}$$

3.5.2 Type (ii) Policy

We adopt the same notation as in Section 3.4.2 except that a penalty cost p is incurred for each backorder per unit of time.

Notations:

γ	Selling price.
l	Lead time.
λ	Demand rate.
h	Holding cost.
p	Penalty cost.
k	Fixed cost of an order.
c	Variable cost of an order.
$D(t)$	Demand during an interval of fixed length t. This is a random variable with c.d.f. $F_D(d,t)$.
$z^{(2)}(s,\tau,u)$	Long-run average profit.
$g(s,t)$	$= h\mathbf{E}[(s-D(t))^+] + p\mathbf{E}[(D(t)-s)^+]$ $(h+p)\sum_{j=0}^{s} F_D(j,t) + p\lambda t - ps$

Decision variables:

s	order-up-to level.
τ	cycle time.
u	reserved capacity.

Case 1: $c = 0$
Under the assumption of zero variable cost, or equivalently infinite capacity, the model reduces to the (s,t) policy as described in Rao (2003)[48]. The long-run average profit can be written as

$$z^{(2)}(s,\tau,\infty) = \lambda\gamma - \frac{k}{\tau} - \frac{1}{\tau}\int_{l}^{l+\tau} g(s,t)dt \tag{3.80}$$

or in cost minimization form

$$y^{(2)}(s,\tau,\infty) \quad = \quad \frac{k}{\tau} + \frac{1}{\tau}\int_{l}^{l+\tau} g(s,t)dt \tag{3.81}$$

Rao has shown that $y^{(2)}(s,\tau)$ is jointly convex in s and τ.

Case 2: $u < \infty$
As with the service level model, consider S_t to be the starting inventory position of the t^{th} cycle. Given S_t, the cycle average profit is

$$z^{(2)}(s,\tau,u|S_t) \quad = \quad \lambda\gamma - \frac{k}{\tau} - \frac{cu}{\tau} - \frac{1}{\tau}\int_{l}^{l+\tau} g(S_t,t)dt \tag{3.82}$$

then

$$z^{(2)}(s,\tau,u) = \lambda\gamma - \frac{k}{\tau} - \frac{cu}{\tau} - \frac{1}{\tau}\int_{l}^{l+\tau} \mathbf{E}[g(S_\infty,t)]dt \tag{3.83}$$

where S_∞ is the stationary inventory position and the equality holds since $g(S_\infty,t) > 0$ thus Tonelli's theorem applies. For both cases, the closed form solutions for s, τ, u and long-run average profit $z^{(2)}$ generally do not exist so numerical analysis is necessary.

3.5.3 Comparison Between Two Policies

We computationally compare the performance of the two policies using the same test problem as in Section 3.4 with the service level model, except we replace the service level constraint $\alpha = 0.98$, with a penalty cost $p = 9$. The results are shown in Table 3.2:

Policy	Cases	Lead time	Parameters			Cost	Profit
			Changing σ_L				
(i)	Exact	$\sigma_L = 0$	$r_* = 39$	$q_* = 108$		97.1634	102.8366
(i)	Upper bound	$\sigma_L = 0.25$	$r_* = 39$	$q_* = 109$		97.5023	102.4977
(i)	Upper bound	$\sigma_L = 0.5$	$r_* = 39$	$q_* = 109$		98.3624	101.6376
(i)	Upper bound	$\sigma_L = 0.75$	$r_* = 39$	$q_* = 111$		99.7815	100.2185
(i)	Upper bound	$\sigma_L = 1$	$r_* = 39$	$q_* = 113$		101.7389	98.2611
(i)	Upper bound	$\sigma_L = 2$	$r_* = 37$	$q_* = 127$		114.2402	85.7598
(i)	Upper bound	$\sigma_L = 3$	$r_* = 35$	$q_* = 146$		132.4863	67.5137
(i)	Upper bound	$\sigma_L = 4$	$r_* = 33$	$q_* = 171$		154.4357	45.5643
(i)	M/M/1 Exact	$\sigma_L = 5$	$r_* = 57$	$q_* = 151$		166.9695	33.0305
(i)	M/M/1 Exact	$\sigma_L = 5$	$r = 30$	$q = 199$		173.1334	26.8666
(i)	Upper bound	$\sigma_L = 5$	$r_* = 30$	$q_* = 199$		178.7462	21.2538
			Changing l				
(ii)	Exact	$l = 0$	$s_* = 95$	$\tau_* = 10.5$	$u = \infty$	108.0585	-
(ii)	Exact	$l = 1$	$s_* = 105$	$\tau_* = 10.5$	$u = \infty$	108.4389	-
(ii)	Exact	$l = 2$	$s_* = 116$	$\tau_* = 10.6$	$u = \infty$	108.8200	-
(ii)	Exact	$l = 3$	$s_* = 126$	$\tau_* = 10.6$	$u = \infty$	109.1871	-
(ii)	Exact	$l = 4$	$s_* = 137$	$\tau_* = 10.7$	$u = \infty$	109.5574	-
(ii)	Exact	$l = 5$	$s_* = 146$	$\tau_* = 10.7$	$u = \infty$	109.9128	-
			Changing τ				
(ii)	Exact	$l = 5$	$s_* = 149$	$\tau = 11$	$u = \infty$	109.9536	-
(ii)	Exact	$l = 5$	$s_* = 185$	$\tau = 15$	$u = \infty$	115.5574	-
(ii)	Exact	$l = 5$	$s_* = 229$	$\tau = 20$	$u = \infty$	129.5029	-
(ii)	Exact	$l = 5$	$s_* = 274$	$\tau = 25$	$u = \infty$	146.8500	-
(ii)	Exact	$l = 5$	$s_* = 319$	$\tau = 30$	$u = \infty$	165.9054	-
			Changing u				
(ii)	Exact	$l = 5$	$s = 149$	$\tau = 11$	$u = 125$	110.2885	82.8933
(ii)	Exact	$l = 5$	$s = 149$	$\tau = 11$	$u = 120$	110.9347	84.5198
(ii)	Exact	$l = 5$	$s = 149$	$\tau = 11$	$u = 115$	116.1427	81.5846
(ii)	Exact	$l = 5$	$s = 149$	$\tau = 11$	$u = 112$	162.1651	36.9258

Table 3.2: Comparison of type (i) and (ii) for penalty cost model policy

Table 3.2 leads to conclusions similar to those drawn in Section 3.4. When the lead time variability is small, the long-run average cost of the type (i) policy outperforms the type (ii) policy. As lead time variability increases, the performance of type (i) policy deteriorates rapidly. The performance of the type (ii) policy is much more stable. Even under suboptimal $\tau = 30 \approx 3\tau_*$ or extremely constrained capacity $u = 112 = \lambda\tau + 2$, the long-run average cost of the type (ii) policy is still better than that of the type (i) policy.

Managerial Insight: Our conclusion in Section 3.4 is not an artifact of the choice of the service level constraint. The penalty cost model leads to the same conclusion as the service level model: the type (ii) policy outperforms when the lead time variability is high.

3.6 Centralized System

We now turn to the analysis of a centralized system in which both the manufacturers and the outsourcer belong to the same organization and the goal is to maximize the total profit of the entire system. This setting may apply to, for example, a firm who manufactures several similar products that share a single machine. Such a machine may serve as an "outsourcer" as in the decentralized system. If the system is operated under type (i) policy, the demand for a product is accumulated until a pre-determined batch size is reached (which is analogous to an order with size q in (r, q) policy), and then the entire batch is sent to the machine; the batch joins a queue and waits for processing. If the system is operated under type (ii) policy, then the schedule of the machine is divided into several time slots, one slot dedicated to each type of product (which is analogous to the reserved capacity u in a type (ii) policy). Given these sets, we explore which policy can lead to higher profit for the firm. Moreover, this centralized system also serves as a baseline for the decentralized system, which allows us to compare the performance of type (i) and (ii) policies in a fair setting that avoids explicit assumptions about the lead time distribution. To be consistent through out the chapter, we still refer to the machine as the "outsourcer" even though the outsourcer and manufacturers are part of the same organization. To study how type (i) and (ii) policies affect the profit of the entire system, we consider two cases in the next two subsections: (1) all manufacturers are type (i); and (2) all manufacturers are type (ii).

3.6.1 Case 1: All Manufacturers Are Type (i)

We first consider the setting where all manufacturers are type (i), who manage their inventory using an (r, q) policy. We adopt the following notation in this section:

Notations:

Manufacturer

n	Total number of manufacturers.
γ	Selling revenue.
λ	Arrival rate.
h	Holding cost.
k	Fixed cost of placing an order.
c	Variable cost of an order.
α	Service level.
L	Stationary lead time of an order.
(r, q)	Inventory control decisions.
$y^{(1)}(q, n)$	Stationary cost rate for type (i) manufacturer.
$z^{(1)}(q, n)$	Stationary profit rate for type (i) manufacturer.

Outsourcer

μ	Service rate (in demand units).
λ_0	Arrival rate.
μ_0	Service rate (in order units), $\mu_0 = \mu/q$.
k_0	Fixed cost of processing an order.
k	Fixed revenue of processing an order.
c_0	Variable cost of processing an unit.
c	Variable revenue of processing an unit.
h_0	Waiting cost per unit per unit of time.
W	Stationary waiting time of an order.
$z_0^{(1)}(k, c, q, n)$	Stationary profit rate when all manufacturers are type (i).

System:

$y_c^{(1)}(q, n)$	Stationary cost rate for the type (i) centralized system.
$z_c^{(1)}(q, n)$	Stationary profit rate for the type (i) centralized system.

We model the system using the queuing setting shown in Figure 3.3: all incoming demand at manufacturers is Poisson with rate λ. All manufacturers manage inventory using an (r, q) policy and all orders joins a single queue at the outsourcer with FIFO service discipline. The outsourcer has a deterministic processing time $1/\mu_0$ and the arrival rate of incoming orders seen by the outsourcer is λ_0. Next we discuss the underlying assumptions of our model:

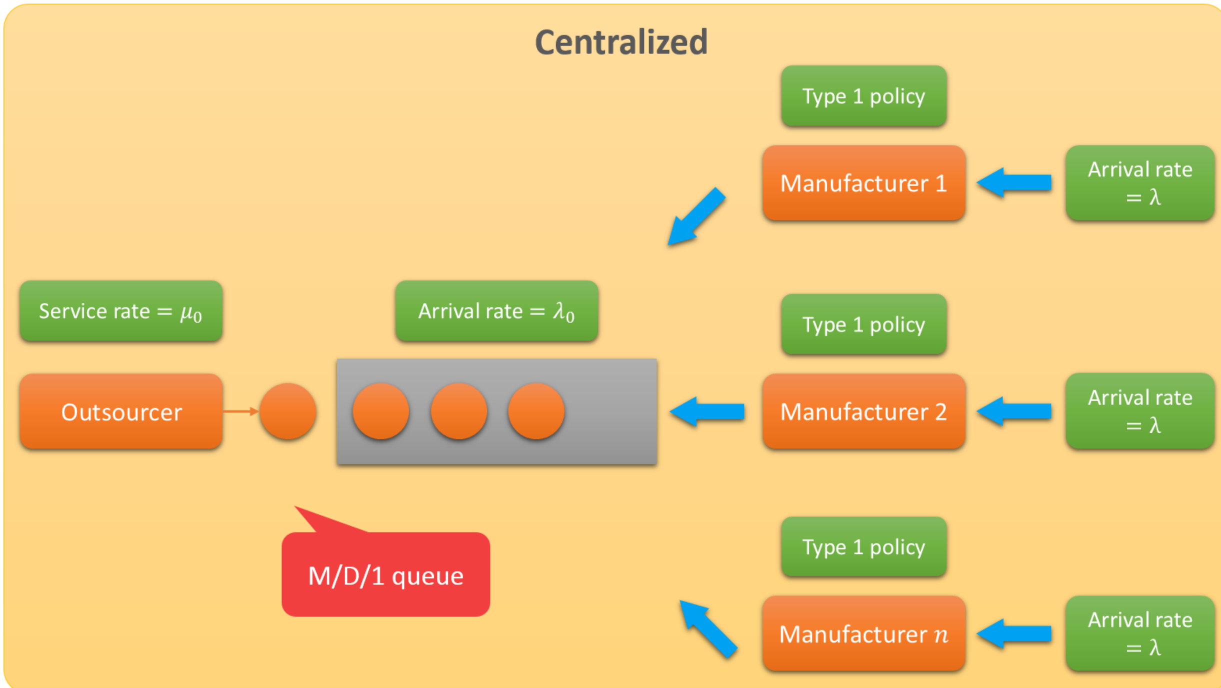

Figure 3.3: A centralized system with all type (i) manufacturers

Assumptions:

A1. The outsourcer manages a single queue with a first-in-first-out (FIFO) service discipline

A1 may not accurately capture real world settings; for example, the outsourcer may use other disciplines such as a polling model to manage all incoming orders (see Takagi (1986)[59]). Since our goal is to compare the effectiveness of two types of outsourcing polices, rather than the impact of advanced queue scheduling, we adopt this concise model to simplify analysis. Relaxing A1 would mostly like require simulation-based analysis.

A2. The service time is deterministic $1/\mu_0$ where $1/\mu_0$ is a linear function of order size q such that $1/\mu_0 = q/\mu$

A2 enables us to fairly compare type (i) and type (ii) policies, since for the latter we assume that processing time during a reservation is deterministic.

A3. $\mu_0 > \lambda_0$

A3 ensures that all manufacturers and the outsourcer have finite long-run average profit.

A4. The total number of manufacturers n is large

As we discussed in Section 3.4, since an order from a manufacturer arrives at the outsourcer according to a renewal process with rate λ/q, A4 allows us to aggregate all of the renewal processes approximately into a Poisson process with rate $n\lambda/q$. Together with the A2, this means that the outsourcer manages an M/D/1 queue with arrival rate $\lambda_0 = n\lambda/q$ and service rate $\mu_0 = \mu/q$.

A5. All manufacturers are homogeneous, i.e. with identical

- Poisson demand process with rate λ
- inventory control policy (r, q)
- service level α
- holding cost h, fixed cost k and variable cost c

A5 may fail to capture certain real-world settings, but otherwise the model is intractable for theoretical analysis. In addition, we assume the demand process for each manufacturer is Poisson process D1. This assumption can be relaxed to allow for other demand processes.

A6. The manufacturer faces a service level constraint as described in Section 3.4

We make this assumption, but our analysis applies equally to the penalty cost model described in Section 3.5.

A7. The system is incurred a waiting cost h_0 per unit per unit time for each order waiting in the queue.

In A7 we assume a linear waiting cost, which captures the notion that the system must pay the inventory holding cost for work-in-progress inventory. Our analysis generalizes to non-linear waiting cost. Just for convenience, we include this cost in the profit calculation of the outsourcer $z_0^{(1)}$.

The total profit of the entire system, denoted by $z_c^{(1)}$ (c stands for centralized), consists of the profit of the outsourcer $z_0^{(1)}$ and the profit of the manufactures $z^{(1)}$. The long-run average profit rate of the outsourcer is

$$z_0^{(1)}(k, c, q, n) = \lambda_0 (k + cq) - \lambda_0 (k_0 + c_0 q) - h_0 \lambda_0 q \mathbf{E}[W] \tag{3.84}$$

$$= n\lambda \left(\frac{k}{q} + c \right) - n\lambda \left(\frac{k_0}{q} + c_0 \right) - h_0 n\lambda \mathbf{E}[W] \tag{3.85}$$

where the first term is the revenue from processing orders, the second term is the cost of processing orders, and the third term is the waiting cost. For an M/D/1 queue, the expected waiting time of an order is

$$\mathbf{E}[W] = \frac{1}{2\mu_0}\frac{\lambda_0}{\mu_0-\lambda_0} \tag{3.86}$$

$$= \frac{n\lambda}{2\mu}\frac{q}{\mu-n\lambda} \tag{3.87}$$

so $z_0(k,c,q,n)$ can be written as

$$z_0(k,c,q,n) = n\lambda\left(\frac{k}{q}+c\right) - n\lambda\left(\frac{k_0}{q}+c_0\right) - \frac{h_0 n^2\lambda^2}{2\mu(\mu-n\lambda)}q. \tag{3.88}$$

The long-run average profit rate of a single manufacturer can be written as

$$z^{(1)}(q,r) = \lambda\gamma - \lambda\left(\frac{k}{q}+c\right) - h(r-\lambda\mathbf{E}[L]+\frac{q}{2}) \tag{3.89}$$

where the first term is the revenue from selling finished product. The second term is the cost of placing orders. The third term is the inventory holding cost. For an M/D/1 queue the expected sojourn time of an order is

$$\mathbf{E}[L] = \mathbf{E}[W] + \frac{1}{\mu_0} = \frac{2\mu-n\lambda}{2\mu(\mu-n\lambda)}q. \tag{3.90}$$

Notice the expected lead time is proportional to the order size q. The re-order point r equals $F_L^{-1}(\alpha)$, and so it is determined by the lead time distribution F_L, which is a function of order quantity q and the total number of manufacturers n. We let $r(q,n)$ denote the smallest re-order point that satisfies the service-level constraints given q and n, so $z(q,r)$ can be written as

$$z^{(1)}(q,n) = \lambda\gamma - \lambda\left(\frac{k}{q}+c\right) - h\left(r(q,n) - \frac{\lambda(2\mu-n\lambda)}{2\mu(\mu-n\lambda)}q + \frac{q}{2}\right). \tag{3.91}$$

Then the total profit of the system is

$$z_c^{(1)}(k,c,q,n) \tag{3.92}$$

$$= z_0^{(1)}(k,c,q,n) + nz^{(1)}(q,n) \tag{3.93}$$

$$= n\lambda\gamma - hn\left(r(q,n) - \frac{\lambda(2\mu-n\lambda)}{2\mu(\mu-n\lambda)}q + \frac{q}{2}\right) - \frac{k_0 n\lambda}{q} - n\lambda c_0 - h_0\frac{n^2\lambda^2}{2\mu(\mu-n\lambda)}q. \tag{3.94}$$

Observe that the prices offered by the outsourcer k and c do not affect the total profit in a centralized system, so we can omit k and c in $z_c^{(1)}(k,c,q,n)$ and so we have $z_c^{(1)}(q,n)$. Maximizing the total profit $z_c^{(1)}(q,n)$ is equivalent to minimizing the cost:

$$y_c^{(1)}(q,n) = \frac{k_0\lambda}{q} + h\left(r(q,n) - \frac{\lambda(2\mu-n\lambda)}{2\mu(\mu-n\lambda)}q + \frac{q}{2}\right) + \frac{n\lambda}{2\mu}\frac{h_0\lambda}{\mu-n\lambda}q \tag{3.95}$$

Let q_c be the global minimizer of $y_c^{(1)}(q,n)$, then the total profit of the centralized system with all type (i) manufacturers is $z_c^{(1)}(q_c,n)$.

The difficulty of optimizing $y_c^{(1)}(q,n)$ is mainly due to the implicit function $r(q,n)$. A change in the order quantity q or total number of manufacturers n will change the distribution of the lead time, which will further affect the re-order point for all manufacturers. Since we assume the demand process at each manufacturer is a Poisson process, $r(q,n)$ is naturally discrete on integral values, so the function $y_c^{(1)}(q,n)$ is not a continuous function. However even with a continuous demand process such as a Brownian motion with drift, since the lead time distribution F_L must be discretized computationally, the resulting $y_c^{(1)}(q,n)$ is still discontinuous. Next we show an example of computing the profit of a centralized system with all type (i) manufactures.

Example 3.6.1. Consider an example with $n = 95$ manufacturers with the following parameters:

Manufacturer:

$$\begin{aligned} \lambda &= 10 \\ \gamma &= 25 \\ h &= 1 \\ \alpha &= 0.98 \end{aligned}$$

Outsourcer:

$$\begin{aligned} \mu &= 1000 \\ k_0 &= 800 \\ c_0 &= 5 \\ h_0 &= 1 \end{aligned}$$

The traffic density at the outsourcer is $\rho = n\lambda/\mu = 95\%$. The optimal centralized order quantity q_c can be found by minimizing the centralized cost function $y_c^{(1)}(q,n)$. A plot of the function versus order quantity q is shown in Figure 3.4. The horizontal axis is the order quantity q from 75 to 130 in 0.1 increment and the vertical axis is the cost $y^{(1)}(q,n)$. The cost function is discontinuous as expected.

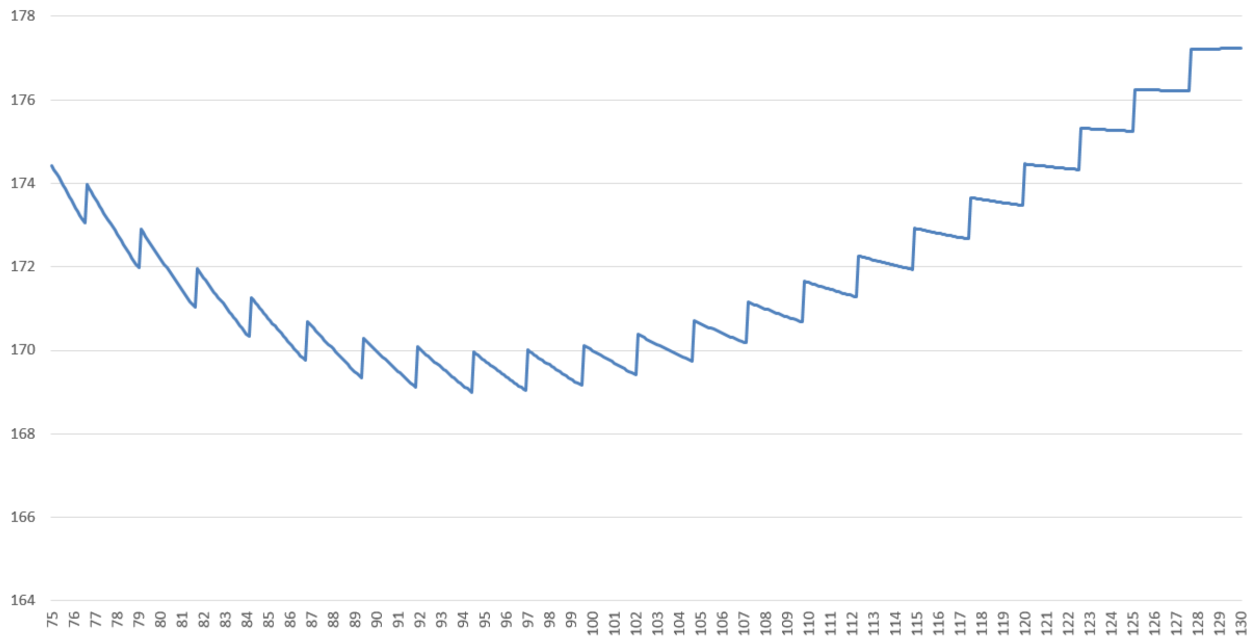

Figure 3.4: Cost function $y^{(1)}(q, n)$ versus q

In this example, the optimal order quantity is $q_c = 94.4$ and the corresponding re-order point $r(q_c, n)$ is 38.

Under this order quantity, the arrival and service rates seen by the outsourcer are

$$\lambda_0 = \frac{n\lambda}{q} = 10.0636 \tag{3.96}$$

$$\mu_0 = \frac{\mu}{q} = 10.5932 \tag{3.97}$$

and thus the expected waiting time and the expected lead time of an order in the corresponding M/D/1 queue are

$$\mathbf{E}[W] = \frac{1}{2\mu_0} \frac{\lambda_0}{\mu_0 - \lambda_0} = 0.8968 \tag{3.98}$$

$$\mathbf{E}[L] = \frac{1}{\mu_0} + \mathbf{E}[W] = 0.9912. \tag{3.99}$$

Assuming the outsourcer offers price $k = 500, c = 10$, the profit of each manufacturer is

$$z^{(1)}(q_c, n) = \lambda\gamma - \lambda \left(\frac{k}{q_c} + c \right) - h \left(r(q_c, n) - \lambda \mathbf{E}[L] + \frac{q_c}{2} \right) = 21.7459 \tag{3.100}$$

and the profit of the outsourcer is

$$z_0^{(1)}(k, c, q_c, n) = \lambda_0 (k + cq_c) - \lambda_0 (k_0 + c_0 q_c) - h_0 \lambda_0 q_c \mathbf{E}[W] = 878.9722. \tag{3.101}$$

Thus, the profit of the entire system is

$$z_c^{(1)} = z_0(k, c, q_c, n) + nz(q_c, n) = 2944.80. \tag{3.102}$$

□

3.6.2 Case 2: All Manufacturers Are Type (ii)

In this subsection we analyze the case where all manufacturers are type (ii). We adopt the following notation in this subsection:

Notations:

Manufacturer:

n	Total number of manufacturers.
γ	Selling revenue.
λ	Arrival rate.
h	Holding cost.
k	Fixed cost of placing an order.
c	Variable cost of an order.
α	Service level.
(s, τ, u)	Inventory control parameters.
$z^{(2)}(s, \tau, u)$	Stationary profit rate for type (ii) manufacturer.

Outsourcer:

μ	Service rate (in the unit of a demand).
k_0	Fixed cost of processing an order.
k	Fixed revenue of processing an order.
c_0	Variable cost of processing an unit.
c	Variable revenue of processing an unit.
$z_0^{(2)}(k, c, u)$	Stationary profit rate.

System:

$y_c^{(2)}(s, \tau, u)$	Stationary cost rate for the all type (ii) centralized system.
$z_c^{(2)}(s, \tau, u)$	Stationary profit rate for the all type (ii) centralized system.

We make the same set of assumptions A1 - A7, except now all manufacturers are type (ii). In this setting, even with $\mu > n\lambda$, the system may not be stable. The stability condition depends on the reserved capacity u for each manufacturer. For a manufacturer to have finite long-run average revenue we need $u > \lambda\tau$. For the outsourcer to be able to accommodate all capacity reservations from all n manufacturers, $\mu\tau > nu$ is required. This suggests that the outsourcer may not have the capacity for all manufacturers if the reserved capacity u is larger than a threshold. To simplify the analysis we make the following additional assumption:

Assumptions:

A8. The outsourcer allocates an equal amount of capacity to each manufacturer, with $u = \frac{\mu\tau}{n}$.

Under this assumption, a simple non-idle cyclic schedule can be adopted at the outsourcer: capacity is reserved by each manufacturer cyclically in sequence $1, 2, ..., n, 1, ...$ where each reservation lasts for τ/n long. Such a schedule and a resulting sample path of inventory position of one manufacturer is shown in Figure 3.5. In addition, with this assumption, the capacity reserved u is no longer a decision variable for the manufacturer.

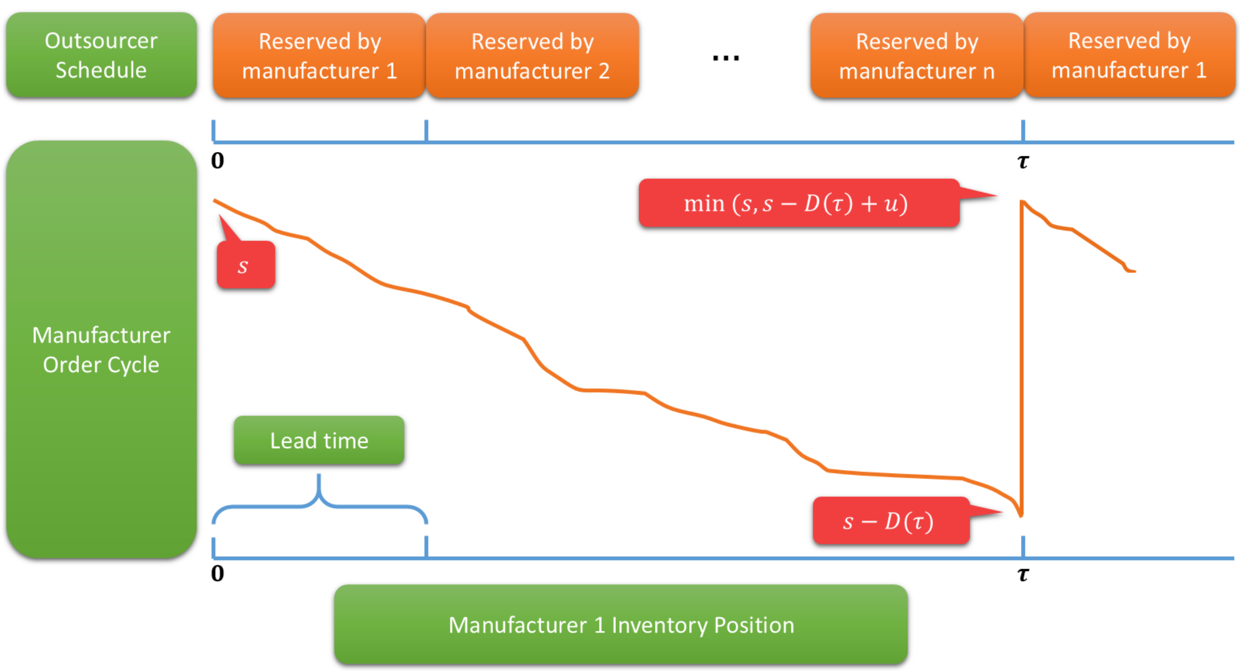

Figure 3.5: Outsourcer adopts a non-idle cyclic schedule with all type (ii) manufacturers

With the assumption A8 the profit of the outsourcer is:

$$z_0^{(2)}(k, c, n) = \frac{n}{\tau}(k + cu) - \frac{n}{\tau}(k_0 + c_0 u). \tag{3.103}$$

Since no manufacturer is ever waiting for an order to be processed, there is no waiting cost incurred for the system. Given a fixed service level α and capacity u, the manufacturer finds the smallest order-up-to level s_* that satisfies the service level constraint and the profit is

$$z^{(2)}(s_*, \tau, n) = \lambda\gamma - y^{(2)}(s_*, \tau, u) \tag{3.104}$$
$$= \lambda\gamma - h\left(\psi(s_*, \tau, u) - \lambda l - \frac{1}{2}\lambda\tau\right) - \frac{k}{\tau} - \frac{cu}{\tau}. \tag{3.105}$$

The total profit of the entire system is

$$z_c^{(2)}(k, c, s_*, \tau, n) = z_0^{(2)}(k, c, u) + nz^{(2)}(s_*, \tau, u) \tag{3.106}$$
$$= n\lambda\gamma - \frac{n}{\tau}(k_0 + c_0 u) - nh\left(\psi(s_*, \tau, u) - \lambda l - \frac{1}{2}\lambda\tau\right) \tag{3.107}$$

which is equivalent to minimizing the cost $y_c^{(2)}$

$$y_c^{(2)}(\tau, s_*, n) = \frac{n}{\tau}(k_0 + c_0 u) + nh\left(\psi(s_*, \tau, u) - \lambda l - \frac{1}{2}\lambda\tau\right). \tag{3.108}$$

Notice with the assumption A8 the system ends up paying variable cost for all of its processing capacity $c_0\frac{nu}{\tau} = c_0\mu$, instead of the variable cost $c_0 n\lambda$ as in the type (i) policy. If $\mu \gg n\lambda$, then the choice of u under the assumption A8 may be suboptimal. This means that our computation of the system profit is conservative compared to the optimal type (ii) system with capacity u as a decision variable. As with the cost function $y_c^{(1)}$ in Case 1, once again the prices k and c do not affect the total profit of the system. To minimize $y_c^{(2)}$, it is crucial to understand the implicit function $\psi(s_*, \tau, u)$, which is the expected inventory position at the beginning of an order cycle given τ, s_* and u, where s_* is the minimum order-up-to level at which the service level constraint is satisfied. Unfortunately, function $\psi(s_*, \tau, u)$ has to be evaluated numerically and is naturally discontinuous due to demand process D1. We will let τ_c denote the optimal τ that minimize $y_c^{(2)}$.

Example 3.6.2. Consider the the example shown in Case 1 in Section 3.6.1 with the same parameters, except that all $n = 95$ manufacturers are now type (ii) manufactures. Assume the order cycle for all manufacturers is $\tau = 12$, then the allocated capacity for each manufacturer u is

$$u = \frac{\mu\tau}{n} = 126.32. \tag{3.109}$$

which is larger than the expected cycle demand $\lambda\tau = 120$. The lead time of each order is

$$l = \frac{u}{\mu} = 0.1263. \tag{3.110}$$

As we can see, the lead time is deterministic and is much shorter than the expected lead time in Case 1 with $\mathbf{E}[L] = 0.9912$.

Using the same $k = 500, c = 10$ as in Case 1, with given τ and u, the smallest order-up-to level s_* that ensures the α-service level can be found numerically as described in Section 3.4 which gives $s_* = 161$, and then the profit of each manufacturer is

$$z^{(2)}(s_*, \tau, u) = \lambda\gamma - y^{(2)}(s_*, \tau, u) - \frac{cu}{\tau} = 8.43 \tag{3.111}$$

and the profit of the outsourcer is

$$z_0^{(2)}(k, c, u) = \frac{n}{\tau}(k + cu) - \frac{n}{\tau}(k_0 + c_0 u) = 2625.00. \tag{3.112}$$

The total profit of the entire system is

$$z_c^{(2)} = z_0^{(2)}(k, c, u) + nz^{(2)}(s_*, \tau, u) = 3425.40. \tag{3.113}$$

Thus, the total profit of Case 2 with all type (ii) manufacturers is about 16.32% higher than that of Case 1 with all type (i) manufacturers.

□

3.6.3 Computational Results

In this section, we extend the example shown in Case 1 in Section 3.6.1 and Case 2 in Section 3.6.2 to explore how changes in parameter values affects total profit of the centralized system. We adopt the following parameters:

Manufacturer:

$$\begin{aligned} \lambda &= 10 \\ \gamma &= 25 \\ h &= 1 \\ \alpha &= 0.98 \end{aligned}$$

Outsourcer:

$$\begin{aligned} \mu &= 1000 \\ k_0 &= 800 \\ c_0 &= 5 \\ h_0 &= 1 \end{aligned}$$

First we show how fixed cost k and variable cost c affect the profit of each party and the system. If all manufacturers are type (i), results are summarized in Table 3.3.

Changing k and c							
k	c	q_c	r	$\mathbf{E}[L]$	$z^{(1)}$	$z_0^{(1)}$	$z_c^{(1)}$
600	5	94.4	38	0.9912	61.15	-2864.70	2944.80
600	10	94.4	38	0.9912	11.15	1885.30	2944.80
600	15	94.4	38	0.9912	-38.85	6635.30	2944.80
800	5	94.4	38	0.9912	39.97	-851.96	2944.80
800	10	94.4	38	0.9912	-10.03	3898.00	2944.80
800	15	94.4	38	0.9912	-60.03	8648.00	2944.80
1000	5	94.4	38	0.9912	18.78	1160.80	2944.80
1000	10	94.4	38	0.9912	-31.22	5910.80	2944.80
1000	15	94.4	38	0.9912	-81.22	10661.00	2944.80

Table 3.3: Centralized with all type (i) manufacturers

The first two columns are the fixed cost k and variable cost c. The third and fourth columns are the optimal inventory control parameter q_c and r. The fifth column is the expected lead time of an order. The last three columns are the profit of manufacturer $z^{(1)}$, profit of the outsourcer $z_0^{(1)}$ and the total profit of the system $z_c^{(1)}$. As we can see from the table, k and c only affect the profit of each party, but not the profit of the system which is always 2944.80. Similar conclusions can be drawn for the Case 2 with all type (ii) manufacturers.

Managerial Insight: In a centralized system, the fixed cost k and variable cost c control the allocation of profit between the outsourcer and the manufacturers, but not the total system profit.

Next we focus on the type (ii) policy and study impact of the order cycle τ on the total profit. We fix $k = 500$ and $c = 10$, and the results are shown in Table 3.4.

Changing τ							
τ	u	s_*	l	$y_c^{(2)}$	$z^{(2)}$	$z_0^{(2)}$	$z_c^{(2)}$
8	84.21	120	0.0842	135.83	8.91	1437.50	2284.10
9	94.74	134	0.0947	136.84	7.90	1833.30	2584.00
10	105.26	140	0.1053	133.51	11.22	2150.00	3216.20
11	115.79	154	0.1158	137.14	7.60	2409.10	3130.80
12	126.32	161	0.1263	136.31	8.43	2625.00	3425.40
13	136.84	174	0.1368	140.42	4.32	2807.70	3217.70
14	147.37	182	0.1474	141.45	3.29	2964.30	3276.60
15	157.89	194	0.1579	145.48	-0.75	3100.00	3029.10
16	168.42	203	0.1684	148.04	-3.30	3218.80	2904.80

Table 3.4: Centralized with all type (ii) manufacturers

In Table 3.4 the first three columns are inventory control decisions τ, u and s_*. The capacity u is first computed with a given τ then the smallest order-up-to level s_* is computed for a given τ and u. The fourth column is the lead time of an order. Keep in mind that under an type (ii) policy, this lead time is deterministic. The fifth column is the optimal cost function $y_c^{(2)}$ and the last three columns are the profit of the manufacturer $z^{(2)}$, profit of the outsourcer $z_0^{(2)}$ and profit of the system $z_c^{(2)}$. As we can see, the optimal order cycle is around $\tau = 12$ with profit 3425.40, which is 16.32% better than the profit of Case 1 with all type (i) manufacturers. This shows that type the (ii) policy can outperforms the type (i) policy under some choices of the order cycles τ.

Next we study how the traffic density at the outsourcer affects the performance of two policies. To do so, we vary the number of manufactures from $n = 98$ to $n = 80$. When n is large, the outsourcer has a higher traffic density thus is busier and vice versa. The results are shown in Table 3.5.

	Changing n								
	Type (i) policy				Type (ii) policy				
n	q_c	r	$\mathbf{E}[L]$	$z_c^{(1)}$	τ	u	s_*	l	$z_c^{(2)}$
98	73.4	77	1.8717	-2151.90	12	122.45	235	0.1224	-1687.30
97	82.7	55	1.4197	750.96	12	123.71	197	0.1237	1017.70
96	89.8	45	1.1674	2103.50	12	125.00	168	0.1250	3013.50
95	94.4	38	0.9912	2944.80	12	126.32	161	0.1263	3425.40
92	106.0	27	0.7155	4194.10	12	130.43	150	0.1304	3874.30
89	112.0	21	0.5651	4689.50	12	134.83	146	0.1348	3847.60
86	114.0	17	0.4641	4899.00	12	139.53	145	0.1395	3596.90
83	112.3	14	0.3864	4958.00	12	144.58	145	0.1446	3287.60
80	111.7	12	0.3351	4931.70	12	150.00	145	0.1500	2987.90

Table 3.5: Comparison between type (i) and (ii) policies with varying n

From Table 3.5 we can see as the number of manufacturers n increases, the expected lead time of type (i) policy increases, and the reserved capacity u of type (ii) policy decreases. As a results, type (i) policy requires a higher re-order point r and type (ii) policy requires a higher order-up-to level s_* to guarantee the service level constraint. Observe that as n becomes larger and larger, the performance of type (ii) policy begins to outperform the type (i) policy. This is because when n is large, type (i) manufacture suffers from the long and highly variable lead time so large waiting cost is incurred along with the expensive inventory cost. On the other hand when n is small, the lead time faced by the type (i) manufactures is short and less variable, therefore type (i) policy generate a higher total profit for the system.

Managerial Insight: The results for the centralized system confirm the conclusion drawn from the computational results in Section 3.4 and 3.5: when the lead time is highly variable, the type (ii) policy can generate higher profit than the type (i) policy for the entire system.

3.7 Semi-centralized System

In this section, we briefly analyze the setting where all manufacturers belong to a single organization, but the outsourcer does not. All the inventory control decisions are transparent across all manufacturers. The organization tries to maximize its total profit instead of the maximizing the total profit of the entire system. We call such a system as a "semi-centralized" system. This setting may apply to, for example, a firm with multiple plants which share the same outside outsourcer. Under type (i) policy, each plant is operated independently and places order that joins a queue at the outsourcer. Under type (ii) policy, the firm proposes a fixed capacity reservation scheme that all plants follow such that no order is ever waiting at the outsourcer. We again consider two case: (1) all manufacturers are type (i); (2) all manufacturers are type (ii) and we show that the semi-centralized system can be coordinated such that the total profit of the system equals to the centralized system.

3.7.1 Case 1: All Manufacturers Are Type (i)

We continue making the same set of assumptions A1-A8 as discussed in the centralized system in Section 3.6, except that in the semi-centralized system we assume the outsourcer has larger power, and so first selects the fixed charge k and variable charge c, and then each manufacturer picks the optimal (r, q) to maximize profit:

$$z^{(1)}(q,n) \;=\; \lambda\gamma - \lambda\left(\frac{k}{q} + c\right) - h\left(r(q,n) - \frac{\lambda(2\mu - n\lambda)}{2\mu(\mu - n\lambda)}q + \frac{q}{2}\right) \tag{3.114}$$

which is equivalent to minimizing

$$y_s^{(1)}(q,n) \;=\; \frac{\lambda k}{q} + h\left(r(q,n) - \frac{\lambda(2\mu - n\lambda)}{2\mu(\mu - n\lambda)}q + \frac{1}{2}q\right). \tag{3.115}$$

Let $q_s(k)$ be the minimizer of $y^{(1)}(q,n)$ for a given k where subscript s stands for "semi-centralized". Since by the assumption A5 all manufacturers are homogeneous, $nz^{(1)}(q_s(k),n)$ is the maximum profit of all manufacturers. Then, the total profit of the system is:

$$z_s^{(1)}(k,c,n) \;=\; z_0^{(1)}(k,c,q_s(k),n) + nz_s^{(1)}(q_s(k),n). \tag{3.116}$$

If there exists a price k_s such that the optimal order quantity of each manufacture in a semi-centralized system $q_s(k_s)$ equals the optimal order quantity in a centralized system q_c (defined in Section 3.6.1), then the system is coordinated. In other words, if the outsourcer charges a fixed cost k_s, the manufacturer will order $q_s = q_c$, and then the total profit of the system is the same as the centralized system. Notice that the cost function of the semi-centralized system $y_s^{(1)}$ has the same term $r(q,n)$ as in the cost function of the centralized system $y_c^{(1)}$ (defined in Section 3.6.1), and thus $y_s^{(1)}$ has the same discontinuous property as well. The existence of such k_s is difficult to show due to this discontinuity. Instead, we'll show the existence of k_s under demand process D2: Brownian motion with drift, in which both $y_s^{(1)}$ and $y_c^{(1)}$ are continuous:

Theorem 3.7.1. *If the demand process faced by each manufacture is the D2 Brownian motion with drift, then the semi-centralized system with all type (i) manufacturers can be coordinated, i.e. there exists a fixed cost k_s such that $q_s(k_s) = q_c$ with total profit $z_s^{(1)} = z_c^{(1)}$, and there exists a variable cost c_s such that for any $k \neq k_s$ and c,*

$$z_0^{(1)}(k_s, c_s) \geq z_0^{(1)}(k, c) \quad \textit{and} \quad z^{(1)}(k_s, c_s) \geq z^{(1)}(k, c).$$

The proof of Theorem 3.7.1 can be found in Appendix A.2.3. Theorem 3.7.1 states that there is a special fixed cost k_s and variable cost c_s, such that the total profit of the semi-centralized system equals to the centralized system and the profit for each party is at least as good as any other choices of $k \neq k_s$ and c. This gives incentive to the outsourcer to choose such a k_s. The variable cost c serves as a lever that controls the proportion of the profit that each party earns. There exists a $c = \bar{c}$ such that the outsourcer earns all the profit, and there exists a $c = \underline{c}$ such that the manufacturer earns all the profit.

Example 3.7.1. Consider an example with the same parameter as in the Example 3.6.1 in Case 1 Section 3.6.1. Though the existence of k_s is not shown under a Poisson demand process, we can still numerically find one of the possible $k_s = 700$, and then the optimal order quantity of each manufacturer $q_s = 94.4 = q_c$ and the total profit $z_s^{(1)} = 2944.8 = z_c^{(1)}$.

Notice that k_s is smaller than k_0, so to earn a positive profit the outsourcer must charge $c > c_0$. For this example, the smallest $\underline{c}$ such that the outsourcer earns a positive profit is $\underline{c} = 6.96$. The largest $\bar{c}$ such that the manufacturer to earn a positive profit is $\bar{c} = 10.05$. In other words, if the outsourcer charges fixed cost $k = k_s = 700$ and variable cost $c \in (\underline{c}, \bar{c}) = (6.96, 10.05)$ then both parties earn positive profit with the total profit equals to the centralized system profit so the system is coordinated.

$\square$

3.7.2 Case 2: All Manufacturers Are Type (ii)

With a fixed charge k and variable charge c offered by the outsourcer, each manufacturer chooses the optimal order cycle τ_s (s stands for semi-centralized) that minimizes the cost function

$$y_s^{(2)}(s_*, \tau, n) \quad = \quad \frac{1}{\tau}(k + cu) + h\left(\psi(s_*, \tau, u) - \lambda l - \frac{1}{2}\lambda\tau\right). \tag{3.117}$$

Similar to the case where all manufacturers are type (i), the system with all type (ii) manufacturers can also be coordinated if there exists a special fixed cost k_s such that each manufacturer selects the optimal order cycle $\tau_s(k_s)$ that equals the optimal order cycle in a centralized system τ_c (defined in Section 3.6.2). As a consequence, the total profit of the semi-centralized system equals to the profit of the centralized system. Due to the same discontinuous property of $\psi(s_*, \tau, u)$ in the cost function $y_s^{(2)}$ and $y_c^{(2)}$, we show such k_s exists under demand process D2: Brownian motion with drift.

Theorem 3.7.2. *If the demand process faced by each manufacture is the D2 Brownian motion with drift, then the semi-centralized system with all type (ii) manufacturers can be coordinated, i.e. there exists a fixed cost k_s such that $\tau_s(k_s) = \tau_c$ with total profit $z_s^{(2)} = z_c^{(2)}$, and there exists a variable cost c_s such that for any $k \neq k_s$ and c,*

$$z_0^{(2)}(k_s, c_s) \geq z_0^{(2)}(k, c) \quad \textit{and} \quad z^{(2)}(k_s, c_s) \geq z^{(2)}(k, c).$$

The proof is similar to the proof of Theorem 3.7.1 and can be found in Appendix A.2.3.

3.7.3 Computational Results

We again use the same example from the centralized system in Section 3.6.3 to explore how changes in k impact the profit of the semi-centralized system as a whole, as well as the profit of each party in the chain. For brevity, we only show the results for the case with all type (i) manufacturers in Section 3.7.1. The results are summarized in Table 3.6.

Changing k									
k	c	$q_s(k)$	r	$\mathbf{E}[L]$	$y_s^{(1)}$	$z_s^{(1)}$	$z_0^{(1)}$	$z_s^{(1)}$	% change
500	10	79.0	32	0.8295	126.4961	23.5039	429.43	2662.30	-9.59%
550	10	84.1	34	0.8830	132.6178	17.3822	1167.00	2818.30	-4.30%
600	10	86.7	35	0.9103	138.4507	11.5493	1776.10	2873.30	-2.43%
650	10	89.3	36	0.9376	144.0619	5.9381	2348.30	2912.40	-1.10%
700	10	94.4	38	0.9912	149.4405	0.5595	2891.70	2944.80	0.00%
750	10	96.9	39	1.0175	154.6749	-4.6749	3385.30	2941.20	-0.12%
800	10	99.5	40	1.0447	159.7045	-9.7045	3852.00	2930.10	-0.50%
850	10	104.6	42	1.0983	164.5790	-14.5790	4260.10	2875.10	-2.37%
900	10	107.1	43	1.1245	169.3381	-19.3381	4670.40	2833.30	-3.79%
700	7.5	94.4	38	0.9912	149.4405	25.5595	516.68	2944.80	0.00%

Table 3.6: Semi-centralized with all type (i) manufacturers

In Table 3.6, in the last column we record the percentage gap compare to the total profit of the centralized system $z_c^{(1)}$. Under the semi-centralized system, the total profit $z_s^{(1)}$ varies with different fixed costs k. When $k = 700$ the optimal order quantity $q_s(k)$ equals the centralized order quantity $q_c = 94.4$ with the total profit $z_s^{(1)} = z_c^{(1)}$ so $k_s = 700$. To show how the variable cost c affects the profit earned by each party, in the last row of Table 3.6, we propose a special $c_s = 7.5$ and compare the profit of each party to the first row where $k = 500$ and $c = 10$. As we can see, with $k = k_s = 700$ and $c_s = 7.5$, the profit of the manufacturer is $25.5595 > 23.5039$ in first row and profit of the outsourcer is $516.68 > 429.43$ in the first row so both parties earn larger profit with k_s and c_s. Thus, the outsourcer has incentive to choose $k_s = 700$ and the system is coordinated.

Managerial Insight: The semi-centralized system can be coordinated so that the total profit equals that of the centralized system.

3.8 Conclusion

Small biopharmaceutical manufacturers often suffer from long and highly random lead times when outsourcing some of their production steps. In this chapter, we analyzed a common capacity reservation inventory control policy in industry that addresses this issue, and compare its performance to that of a classical inventory control (r, q) policy. We show that when the lead time in the (r, q) policy setting is long and variable, the capacity reservation policy not only results in a higher profit for the manufacturer, but also increases the profit of the entire system.

Chapter 4

Process Flexibility with Inventory

4.1 Introduction

Given a fixed total amount of capacity, there are two levers that a manufacturer with multiple plants and multiple products can manipulate to minimize the cost of supply/demand mismatch over time: inventory and process flexibility. Inventory is the traditional way for a manufacturing system to deal with demand fluctuation. When demand for a product exceeds capacity, the system can still satisfy demand using inventory accrued in earlier periods. Process flexibility, or the "the ability to produce different products within the same plant" (Jordan and Graves (1995)[37]), provides another lever. When demand for one product is high while demand for another product is low, a flexible plant can adjust the allocation of its production capacity so that more demand can be satisfied.

Although both inventory and process flexibility can play a role in matching supply with demand, they take effect in different dimensions and require different resources. Inventory allows the utilization of excess capacity over time, and requires an investment in storage, handling, and other inventory-related costs. Process flexibility, on the other hand, coordinates unused capacity throughout the network of plants to quickly respond to demand spikes, and requires upfront capital investment in flexible equipment. In this chapter, we explore the trade-off between inventory and flexibility.

This work was motivated by our experience modeling the supply chain of a major biopharmaceutical firm. In this industry, both the flexibility and inventory capabilities of a manufacturing facility require significant capital investment and long lead time, whether in manufacturing equipment or storage capacity. Thus, the decision of whether a new plant should be flexible, or have significant inventory capability, or both, is a strategic decision made during the drug approval process. Once the decision is made, in general it is difficult to alter in the future, so the correct decision is crucial, and the trade-off between flexibility and inventory is in particular interesting (Kaminsky and Wang (2015) [38]).

At the time we worked with them, this firm's supply chain consisted primarily of dedicated (i.e. non-flexible) facilities, and as such it relied on inventory to cope with demand variation. The firm was exploring the possibility of incorporating flexible manufacturing facilities into its supply chain. By doing so, the firm could accrue a variety of benefits; not only increased ability to meet demand, but increased flexibility to adapt the supply chain to make products that share similar intermediates, or to new products with relatively small demand that makes it difficult to justify new capacity. The management of this firm was hoping to better understand the interplay between inventory and flexibility in their supply chain and to better understand whether investments in process flexibility would be worthwhile.

These questions motivated us to explore a variety of related managerial questions:

- **What are guidelines that might be useful in deciding whether a plant should have flexibility or inventory capability, or both?**

  This broad question leads to a variety of more detailed questions relevant to the biopharmaceutical firm mentioned above:

  - If the firm wants to build a new plant while keeping the current production dedicated as before, should the new plant be built with flexible manufacturing capacity, or inventory capability, or both?
  - If the firm wants to add flexible manufacturing capacity to several existing plants in their current dedicated network, are there general guidelines that will give insight into how many and which plants it would be most beneficial to add this capacity to?
  - Similarly, if a firm wants to add inventory capability to several existing plants in a flexible network without inventory capability, are there general guidelines that will give insight into how many and which plants it would be most beneficial to add this inventory capability to?

- **Are inventory capability and flexibility substitutes or complements?**

  When both inventory capability and flexibility capacity are present, do they complement each other, or does each one decrease the marginal benefit of the other? Is mixing flexibility and inventory in general a good idea?

- **Is it possible to use flexibility to eliminate inventory?**

  In settings such as in the biopharmaceutical industry where inventory storage can be very costly, how does a supply chain that features flexible capacity rather than inventory compare to a dedicated network?

To answer these questions, we extend the class of models most common in the process flexibility literature (see Chou (2008)[18] for a survey of these models) to incorporate both inventory and flexibility. We were motivated by a setting where both the capacity to hold inventory and manufacturing technology enabling flexibility are expensive. Clearly, the relative cost of these two options plays a role in determining which of these or what combination of these two is most appropriate. As is typical in this literature, however, we do not explicitly aim for cost minimization in this chapter. Instead, we focus on maximizing the demand that can be satisfied, as we explain in detail in later section.

To better position our research in the context of existing research, we postpone the discussion of relevant literature (in Section 4.3) until after we introduce notation and the topology of our extended network, which we do in the next section (Section 4.2). In Section 4.4 we introduce our model with the underlying assumptions. In Section 4.5 we analyze a special type of network which a plant can either hold inventory or be flexible but not both. In Section 4.6 we relax this constraint to allow a plant to be both flexible and can hold inventory. In Section 4.7 we summarize our findings.

4.2 Network Topology and Notation

We begin by introducing a graph-based representation of the inventory and flexibility capabilities of a supply chain over time, extending traditional graph-based representations of manufacturing network flexibility. We model a firm designing a set of (capacity-constrained) plants that need to produce a set of products to meet demand over time. Plants can either be dedicated to producing a single product, or can have some degree of flexibility, so that they produce two or more products. In addition, plants may be designed to hold inventory, so that production in earlier periods can be used meet demand in later periods.

Specifically, consider a multi-period, discrete time model of such a manufacturing network with N plants (indexed by i), N products (indexed by j) and T periods (indexed by t). As is typical for this literature, we assume a *balanced network* so that the number of plants equals the number of products, and we assume that each plant has a primary product, so that if plant i is dedicated then it only makes its primary product, so $j = i$. The manufacturing capabilities of this network – which products can be made at which plants, which periods they can be made, and what inventory can be carried at any time – can be represented by a graph with NT plants nodes and NT product nodes. Let v_{it} denote the node representing plant i in period t, and w_{js} denote the node representing (demand for) product j in period s. Although we postpone introducing assumptions and model details until Section 4.4, we introduce the following notation now so we can present our literature in the context of our notation: let c_{it} be the total production capacity of plant i in period t, and let D_{js} be a random variable representing the random demand for product j in period s. Also, let d_{js} be a particular realization of D_{js} and $d = \{d_{js}, \forall j, s\}$ be an instance of a realized demand for all products in all periods. Finally, let $\mathcal{D}$ denote the set of all possible demand realizations.

In our graph-based representation, an arc between a plant node and a product node (v_{it}, w_{js}) means plant i is capable of producing product j at time t, and this production can be used to satisfy the demand for product j at time s. By definition, if $t < s$, so that product produced at time t is used to satisfy the demand in period s, then plant i can hold inventory of that product. We always assume that there is an arc between (v_{it}, w_{it}) – a plant i has the capability of manufacturing its primary product $j = i$ in all periods. If there is an arc (v_{it}, w_{jt}) where $j \neq i$, then plant i is flexible. Let $\mathcal{A}$ denote the set of all arcs between plants and products, and thus $\mathcal{A}$ captures the inventory and flexibility configuration of the manufacturing system.

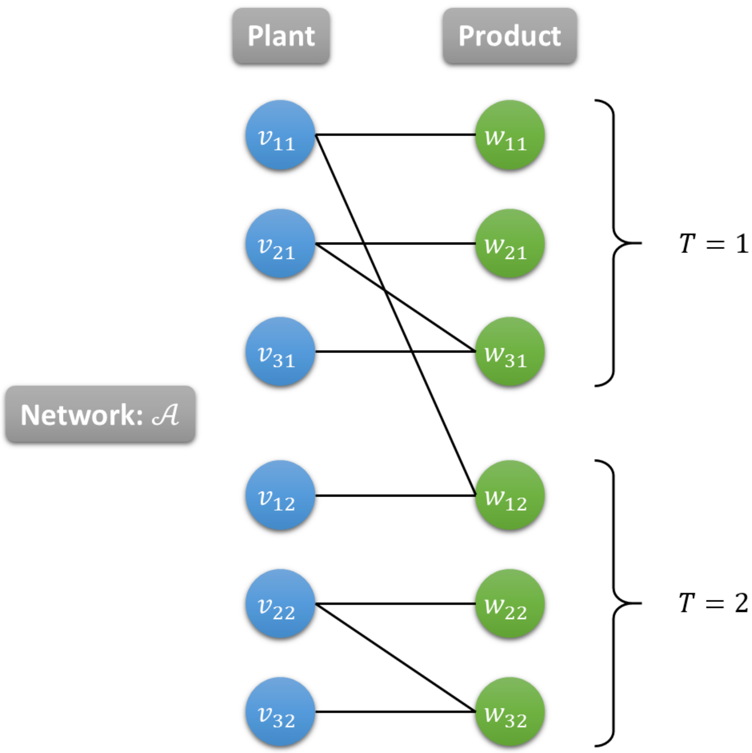

Figure 4.1: An example with 3 plants, 3 products and 2 periods

An example with 3 plants, 3 products and 2 periods is shown in Figure 4.1. Plant 1 can produce product 1 and can hold inventory. Plant 2 is flexible, able to produce both product 2 and 3, but plant 2 can not hold inventory. Plant 3 is neither flexible nor can hold inventory. Observe that the graph represents the capability of the network in terms of production and inventory – in later sections we specify models using the notation and topology described in this section, and then characterize the performance of the network given these model specifications.

Researchers have identified a variety of standard flexibility configurations (see Section 4.3 for a literature review), and we introduce additional inventory configurations. Note that in all of these, we assume that the inventory and flexibility configurations are static over time.

In an *open chain* configuration, $\mathcal{L}_{1\to k}$, plant 1 can produce product 1 and 2, plant 2 can produce product 2 and 3 etc., up to plant $k-1$, and plants k through n are not flexible.

Using the notation defined above,

$$\mathcal{L}_{1\to k} = \{(v_{it}, w_{it})|i = 1, \ldots, N\ \forall t\} \cup \{(v_{it}, w_{i+1,t})|i = 1, \ldots, k-1, \forall t\}.$$

A *closed chain* configuration $\mathcal{C}$ is an open chain $\mathcal{L}_{1\to N}$ with one addition arc – plant N can produce product N and 1. In our notation, a closed chain

$$\mathcal{C} = \mathcal{L}_{1\to N} \cup \{(v_{Nt}, w_{1t})|\ \forall t\}.$$

In the *full flexibility* configuration $\mathcal{F}$, each plant can produce all products, so

$$\mathcal{F} = \{(v_{it}, w_{jt})|\forall i, \forall j, \forall t\}.$$

For any given set $S \in \{1, 2, ..., N\}$, *inventory configuration* $\mathcal{I}_S$ indicates that plant $i \in S$ can hold inventory of the product i and plants not in S hold no inventory, so

$$\mathcal{I}_S = \{(v_{it}, w_{is})|i \in S\ , \forall s > t\}.$$

In a slight abuse of notation, we will use $S = l \to k$ to indicate plants $l, l+1, \ldots, k$.

We define a tuple (a, b, N, T) to describe a general flexibility and inventory configuration $\mathcal{A}$ where a and b represent the flexibility structure and inventory configuration respectively. If $a = \emptyset$, then no plant is flexible (called the dedicated configuration). If $b = \emptyset$, then no plant can hold inventory. N is the number of plant-product pair and T is the number of time periods and when the context is clear we may omit N and T. Examples of flexibility and inventory configurations are shown in Figure 4.2: (i) $(\mathcal{F}, \emptyset, 3, 2)$ (ii) $(\mathcal{L}_{1\to 3}, \mathcal{I}_3, 3, 2)$ (iii) $(\mathcal{L}_{1\to 2}, \mathcal{I}_{2\to 3}, 3, 2)$ (iv) $(\mathcal{C}, \emptyset, 3, 2)$.

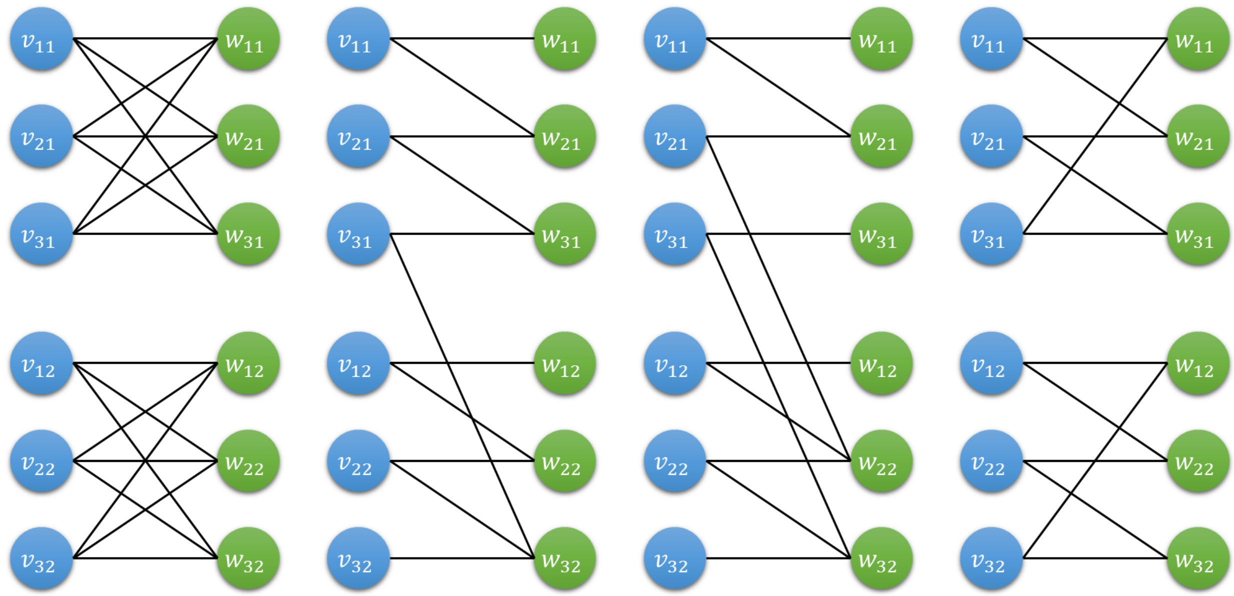

Figure 4.2: Example of flexibility and inventory configurations

Indices:	
$i \in \{1, 2, \ldots, N\}$	The index of plants
$j \in \{1, 2, \ldots, N\}$	The index of products
$t \in \{1, 2, \ldots, T\}$	The index of time periods
Notations:	
$\{v_{it}\}$	The set of plant nodes
$\{w_{jt}\}$	The set of product nodes
c	Capacity of plant
D p_k	Discrete random variable for demand Probability of demand equals to k, i.e. $\mathbf{P}(D = k)$
d_{jt} d $\mathcal{D}$:	Demand for product j in period t in the demand instance d An instance of demand realization for all products in all periods. $d = \{d_{jt}\}$ The set of all possible demand instances
Classes of flexibility and inventory configurations:	
$\mathcal{C}$	Closed chain (also called closed chain)
$\mathcal{L}_k$	Open chain from plant 1 to k
$\mathcal{F}$	Full flexibility
$\mathcal{I}_S$	Inventory can be hold at plant $i \in S$.
$\mathcal{A}$	A general flexibility and inventory configuration

Table 4.1: Summary of notation

4.3 Literature Review

We first review the process flexibility literature that focuses on a single period, and for brevity, notation related to time t and inventory $\mathcal{I}$ is dropped. When no inventory can be held, full flexibility $\mathcal{F}$ clearly maximizes the amount of demand that can be met. However, full flexibility is typically expensive and impractical to implement, so we are interested in exploring so-called sparse designs. A design is sparse if each plant can produce just a few different types of product, typically two or three, so that in sparse design $\mathcal{A}$, the number of arcs is significantly less than that of a fully flexible design, while ideally most of the benefits of full flexibility are maintained. Among the many possible sparse flexibility designs, Jordan and Graves (1995)[37] first noticed the power of the closed chain $\mathcal{C}$. Specifically, they assumed that the decision maker observes demand before making production decisions. Let $\Phi(\mathcal{A}, N; d)$ be the maximum demand that can be satisfied using network $\mathcal{A}$ given observed demand $d = \{d_j, \forall j\}$. $\Phi(\mathcal{A}, N; d)$ can be modeled easily as a maximum flow problem:

$$\Phi(\mathcal{A}, N; d) = \max \sum_{i,j|(i,j)\in\mathcal{A}} X_{ij} \tag{4.1}$$

$$\text{s.t.} \sum_{i|(i,j)\in\mathcal{A}} X_{ij} \leq d_j \quad \forall j \tag{4.2}$$

$$\sum_{j|(i,j)\in\mathcal{A}} X_{ij} \leq c_i \quad \forall i \tag{4.3}$$

$$X_{ij} \geq 0 \quad \forall i, j \tag{4.4}$$

where X_{ij} is the flow on arc (i, j). To measure the effectiveness of a network $\mathcal{A}$, they use the *expected maximum flow* $\mathbf{E}[\Phi(\mathcal{A}, N)]$ where the expectation is taken over the set of all possible demand realizations $\mathcal{D}$. For a balanced (N plants N products) and symmetric (identical capacity c and *i.i.d.* demand) manufacturing system, Jordan and Graves find via simulation that the closed chain $\mathcal{C}$ can achieve almost the same benefit as the full flexibility design, i.e. $\mathbf{E}[\Phi(\mathcal{C}, N)]$ is close to $\mathbf{E}[\Phi(\mathcal{F}, N)]$. Although they don't not give an explicit closed-form expression for the performance of a closed chain, they demonstrate that given a realized demand, the probability that the performance of closed chain is worse than that of full flexibility is small.

In the decades after the innovative work of Jordan and Graves, researchers have attempted to fully understand the power of the closed chain. Chou et al. (2010)[16] demonstrated the superiority of a closed chain by comparing the expected performance of a closed chain and a fully flexible chain as system size goes to infinity:

$$\lim_{N\to\infty} \frac{\mathbf{E}[\Phi(\mathcal{C}, N)]}{\mathbf{E}[\Phi(\mathcal{F}, N)]}. \tag{4.5}$$

The closed chain performs worse as N increases, but in the worst case a closed chain can still achieve 89.6% of the performance of full flexibility if the demand is a two point distribution, taking the values 0 and $2c$ with equal probability. On the other hand, a closed chain performs better if the demand variance is small. To obtain the asymptotic average performance of a closed chain, Chou et al. (2010)[16] first observed that a closed chain is equivalent to an open chain if $N = \infty$:

$$\lim_{N\to\infty} \mathbf{E}[\Phi(\mathcal{C}, N)] = \lim_{N\to\infty} \mathbf{E}[\Phi(\mathcal{L}_{1\to N}, N)] \tag{4.6}$$

The performance of an open chain can be easily calculated using a greedy algorithm. The authors modeled the flow on flexible arcs as a generalized random walk with boundaries 0 and c, and this makes it relatively straightforward to calculate $\lim_{N\to\infty} \mathbf{E}[\Phi(\mathcal{C}, N)]$ for any discrete distribution that is symmetric around the mean.

Later, Chou et al. (2011)[17] studied the worse-case performance of a closed chain. For the balanced and symmetric case, they employed the concept of graph expander to show that there always exists a sparse design with only $O(N)$ arcs that is within the ϵ-optimality of full flexibility for all demand scenarios if demand is bounded around its mean. Specifically,

$$\Phi(\mathcal{C}, N; d) \geq (1-\epsilon)\Phi(\mathcal{F}, N; d) \; \forall d \in \mathcal{D} \tag{4.7}$$

if $D \leq \lambda \mathbf{E}[D]$ almost surely. Surprisingly, for the case in which each plant can produce more than two products, the best sparse design is no longer a chain structure. Chou also extended the graph expander concept to a more general unbalanced and non-symmetric case in which the number of plants and products as well as the capacity and demand distribution may differ. They provided a heuristic and guidelines to find good sparse design for non-symmetric cases.

Simchi-Levi and Wei (2012)[55] were able to characterize the performance of closed chain of any size N. They utilized the concept of supermodularity to show that the flexible arcs in closed chain are complements. The existence of a flexible arc increases the benefit of other flexible arcs, and thus each flexible arc added in the closed chain has increasing marginal benefit:

$$\mathbf{E}[\Phi(\mathcal{L}_{1\to 2}, N)] - \mathbf{E}[\Phi(\emptyset, N)] \leq \ldots \leq \mathbf{E}[\Phi(\mathcal{C}, N)] - \mathbf{E}[\Phi(\mathcal{L}_{1\to N}, N)]. \tag{4.8}$$

To compute the average performance of a closed chain, Simchi-Levi and Wei first showed that $\mathbf{E}[\Phi(\mathcal{C}, N)]$ can be decomposed into parts:

$$\mathbf{E}[\Phi(\mathcal{C}, N)] = N(\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, N)] - \mathbf{E}[\Phi(\mathcal{L}_{1\to N-1}, N-1)]). \tag{4.9}$$

Based on this, they developed a matrix multiplication algorithm that runs in $O(nN^2)$ to calculate $\mathbf{E}[\Phi(\mathcal{C}, N)]$. They also revealed several insights into closed chains. They showed that the per node performance of a closed chain is non-decreasing as system size increases:

$$\frac{\mathbf{E}[\Phi(\mathcal{C}, N)]}{N} \leq \frac{\mathbf{E}[\Phi(\mathcal{C}, N+1)]}{N+1}. \tag{4.10}$$

However, the benefit decreases exponentially fast:

$$\frac{\mathbf{E}[\Phi(\mathcal{C}, N+1)]}{N+1} - \frac{\mathbf{E}[\Phi(\mathcal{C}, N)]}{N} \leq Ke^{-cN} \tag{4.11}$$

which implies that several smaller closed chains performs as well a single closed chain if N is large. In addition, they showed the difference between the fill rate of full flexibility and that of closed chain is non-decreasing:

$$\frac{\mathbf{E}[\Phi(\mathcal{F}, N)]}{N} - \frac{\mathbf{E}[\Phi(\mathcal{C}, N)]}{N} \leq \frac{\mathbf{E}[\Phi(\mathcal{F}, N+1)]}{N+1} - \frac{\mathbf{E}[\Phi(\mathcal{C}, N+1)]}{N+1} \tag{4.12}$$

which means as the system size increases, a closed chain is increasingly less effective than a fully flexible system.

Mak and Shen (2009)[43] and Deng and Shen (2012)[21] relax the assumption of a balanced and symmetric system. Mak and Shen formulated the problem as a two-stage stochastic program where (1)the first stage selects a flexibility configuration $\mathcal{A}$ and (2)the second stage observes the demand and then assigns capacity to maximize the profit under the flexibility configuration chosen in the first stage. Their formulation allows distinct demand distributions, and non-identical selling prices, and flexibility investment cost and production cost are incorporated. They solved the problem heuristically using Lagrangian relaxation with state-independent multipliers. They find that if the demand variation is very low or large, their stochastic programming approach can find a better flexibility structure than the closed chain. Deng and Shen focused on the analysis of the symmetric but unbalanced system (N plants M products, $M > N$) with $2M$ edges. They converted the classical bipartite representation to a circular representation and proposed several new flexibility design guidelines for unbalanced system.

All of papers mentioned above contain fundamental models and results for studying process flexibility, but are restricted to single period setting. Graves and Tomlin (2003)[31] extended the notion of process flexibility to a multi-stage setting. In their model, each product must be processed through all stages before satisfying the demand. They showed that even reasonable flexibility designs can result in inefficiencies such that the multistage supply chain performs worse than the bottleneck stage. Francas et at. (2009)[25] studied multi-period process flexibility under non-stationary demand. The problem is formulated as two-stage stochastic program similar to that of Mak and Shen (2009)[43]. The first stage selects the flexibility configuration $\mathcal{A}$ and the second stage determines operational decisions given the realization of demand in all periods. They found that chaining principles still remain superior to other flexible configurations. In these two papers, flexibility in different stages or time periods is "independent" in the sense that neither inventory nor backorder is incorporated into the model.

As far as we are aware, only one recent paper, by Cong, Wei and Zhong (2015) [54] considers a multi-period process flexibility model while different periods are related via backorder. In their stochastic make-to-order model, in each period the demand for that period is observed first, and then production decisions are made based on the process flexibility structure. Unsatisfied demand is backlogged into the next period but no inventory is allowed in the model. Their work tries to answer two questions: First, should a sparse structure be designed (such as the closed chain in single-period model) in order to satisfy most of the demand? Second, how should the production policy be designed so the actual production plan can be implemented? To answer the first question, they come up with a Generalized Chaining Condition(GCC) structure. Any structure that is GCC will perform close to full flexibility. For the second question, they propose a maximum weight policy, which solves a weighted maximum flow problem in each period that prioritizes the products with higher backorders. However in their model, all products can be back-ordered by default. This means backorder is not part of the network structure, while in our problem, inventory capability affects the network structure, i.e. whether a plant can hold inventory or not is part of the "strategic decision".

To our best knowledge, this is the first attempt to combine process flexibility modeling with traditional inventory based supply chain modeling to develop insight into the relationship between flexibility and inventory. In the next section we will officially introduce our model with the underlying assumptions.

4.4 The Model and Assumptions

In this section, we build on the network representation in Section 4.2 to develop a model that extends models of single period flexibility in the literature to incorporate inventory. We first review the implicit and explicit assumptions built into our model, and then present the performance measure we focus on for the remainder of this chapter.

The network structure introduced in Section 4.2 can capture many possible inventory, flexibility, capacity, and demand scenarios. In that section, we introduced several restrictions on network structure. In particular, we have already restricted the model to reflect the following restrictions on network topology:

A1. There are an equal number of plants and products.

A2. Plant i is always capable of producing product $j = i$.

A3. If plant i is capable of producing product j, then plant i can produce j in any period t.

A4. If plant i can carry the inventory of product j, then plant i can carry the inventory of product j to any period t after the current period.

These assumptions imply that the manufacturing system is balanced, symmetric, and time invariant. Recall from Section 4.3 that starting with Jordan and Graves (1995)[37], the performance of a flexible network has been measured by determining expected maximum flow on the appropriate graph. In our setting, if we further assume that:

A5. Capacity is invariant in time and between plants, i.e. $c_{it} = c$.

A6. Demand for all products in all periods D_{js} is observed before making production.

then to extend the formulation (4.2) defined by Jordan and Graves (1995)[37], given any set of demand realizations d across products and the time horizon for network configuration $(\mathcal{A}, N, T)$, we can define $\Phi(\mathcal{A}, N, T; d)$ to be the maximum demand that can be satisfied. $\Phi(\mathcal{A}, N, T; d)$ can be formulated as a maximum flow problem as shown in Figure 4.3.

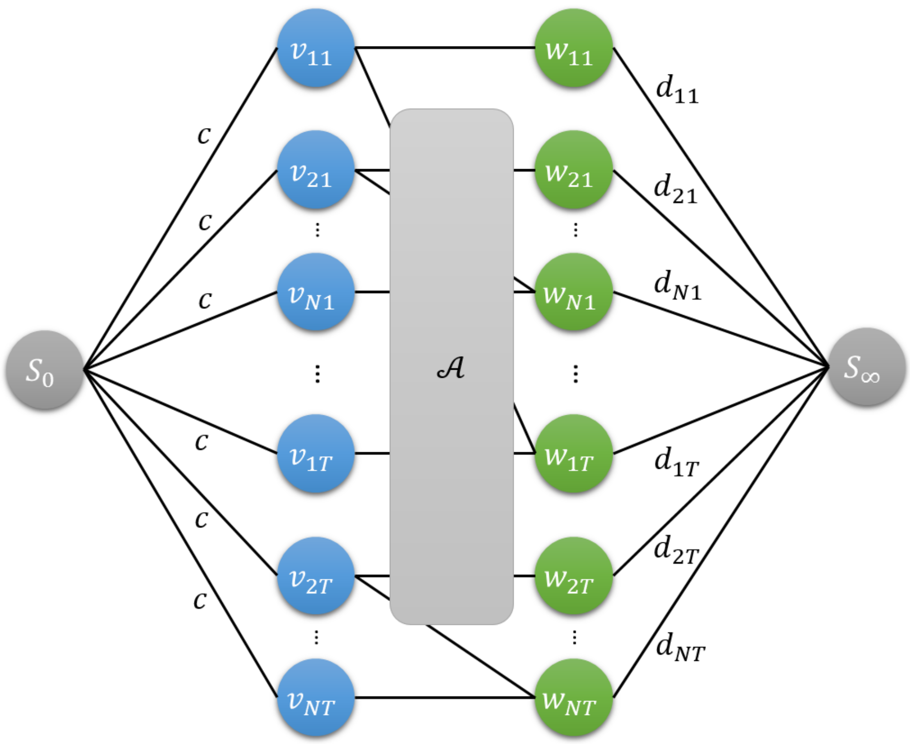

Figure 4.3: Maximum flow formulation of $\Phi(\mathcal{A}, N, T; d)$

Arcs (S_0, v_{it}) have capacity c, arcs (w_{jt}, S_∞) have capacity d_{jt}, and all other arcs have infinite capacity. There is a $\Phi(\mathcal{A}, N, T; d)$ associated with each demand instance $d \in \mathcal{D}$. Therefore $\Phi(\mathcal{A}, N, T)$ is a random variable on a natural probability space. We define the *performance* of a inventory and flexibility configuration $\mathcal{A}$ as the expected maximum demand that can be satisfied over the set of all demand scenarios $\mathbf{E}[\Phi(\mathcal{A}, N, T)]$.

For the remainder of the chapter, we assume:

A7. Demand D is an *i.i.d.* discrete random variable on support $\{0, 1, 2, \dots, 2c\}$ with probability $p_0, p_1, \dots, p_{2c}$ and $\mathbf{E}[D] = c$.

Given this characterization of demand, the performance of a particular inventory and flexibility configuration $\mathbf{E}[\Phi(\mathcal{A}, N, T)]$ can naively be evaluated by averaging the optimal objective function value of $(2c+1)^{NT}$ deterministic linear programs $\Phi(\mathcal{A}, N, T; d)$. Despite the fact that there are an exponential number of demand scenarios, later, in Section 4.5 and 4.6, we show that $\mathbf{E}[\Phi(\mathcal{A}, N, T)]$ can be calculated in polynomial time for certain types of configurations. Ultimately, our goal is to analyze and compare $\mathbf{E}[\Phi(\mathcal{A}, N, T)]$ for different flexibility and inventory configurations $\mathcal{A}$.

We note that this approach implicitly incorporates the assumptions that:

A8. There is no obligation to satisfy all demand. Unsatisfied demand is lost.

A9. The effectiveness of the network can be characterized by the maximum amount of demand that can be met. No capital, holding, penalty or production costs are explicitly considered.

Assumptions A1–A5 describe a manufacturing system that is balanced and symmetric. These are standard assumption in much of the flexibility literature, facilitating tractable analysis. Assumption A6 is also a standard assumption in the literature, but for our multi-period problem, A6 means that the decision maker has a perfect demand forecast over T periods. In many cases (at least in the settings that motivate us), this assumption is not as restrictive as it initially appears. First, in the biopharmaceutical industry, demand over the relatively short term can be estimated based on a deep understanding of the patient population. In addition, demand is often met by regional distributors, who commit to orders several periods ahead of time. Assumptions A8 and A9 are also standard in this literature – in the conclusion to this chapter we discuss the potential cost implications of our analysis.

In order to develop guidelines for designing flexibility and inventory configurations $\mathcal{A}$, we develop algorithms for computing and analyzing $\mathbf{E}\left[\Phi(\mathcal{A}, N, T)\right]$ for a variety of special cases and settings. In Section 4.5 we study the trade-off between inventory and flexibility by comparing the performance of a series of networks $\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2 \ldots$, with an increasing number of inventory arcs but a decreasing number of flexibility arcs, under the constraint that each plant has either flexibility or inventory capability, but not both. In Section 4.6 we relax this constraint and allow a plant to be both flexible and carry inventory.

4.5 Plant with Either Flexibility or Inventory

In this section, we explore the performance of configurations when both flexibility and inventory are present, but with a special constraint: a plant is either capable of holding inventory, or has the capability to be flexible, but not both.

Consider a sequence of inventory and flexibility configurations $\mathbb{A} = \{\mathcal{A}_0, \mathcal{A}_1, \mathcal{A}_2, \ldots \mathcal{A}_N\}$ constructed as follows; each has N plants, N products and T periods. $\mathbb{A}$ starts with a closed chain without inventory $\mathcal{A}_0 = (\mathcal{C}, \emptyset, N, T)$. Then, in each subsequent configuration, $\mathcal{A}_k$, the flexibility arc at the last flexible plant $i = N - k + 1$ is removed and an inventory arc at that plant is added. After N configurations are constructed, a dedicated configuration with inventory $\mathcal{A}_N = (\emptyset, \mathcal{I}_{1\to N}, N, T)$ is obtained. For a case of $N = 3, T = 2$, the set of configurations $\mathbb{A}$ is shown in Figure 4.4:

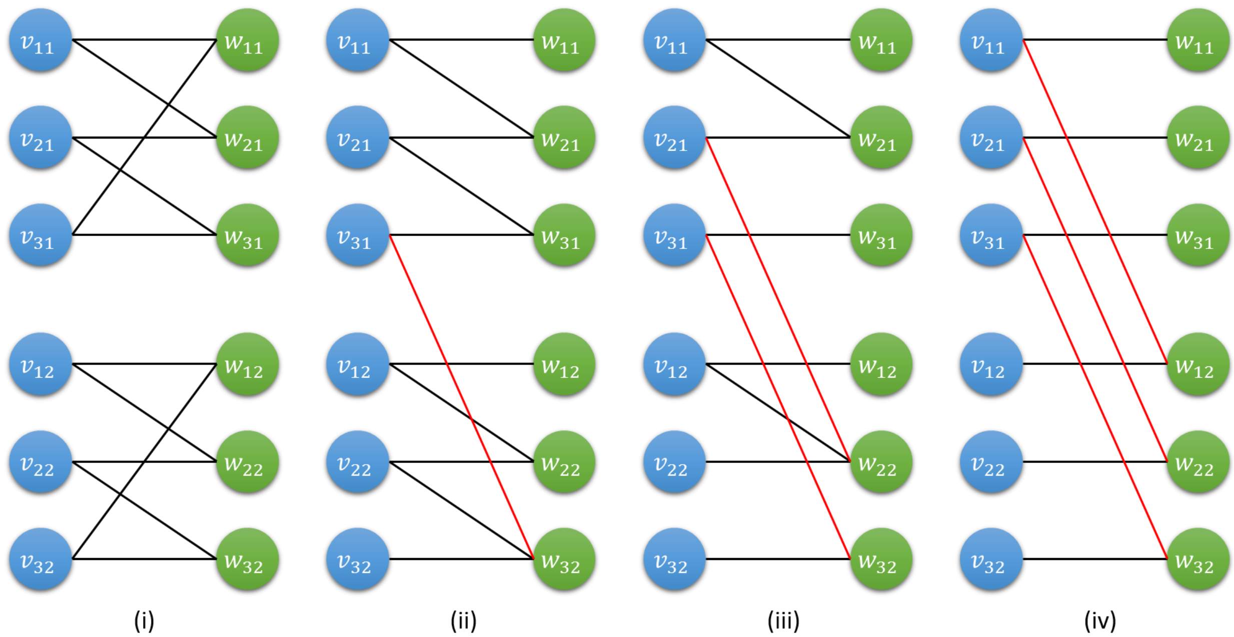

Figure 4.4: The sequence of configurations in $\mathbb{A}$

We are interested in comparing the expected maximum flow in each of the configuration in $\mathbb{A}$, to develop insight into how substituting inventory for flexibility impacts the performance of the system. Specifically, we compare the following:

$$\mathbf{E}[\Phi(\mathcal{C}, \emptyset, N, T)], \mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)], \mathbf{E}[\Phi(\mathcal{L}_{1\to N-1}, \mathcal{I}_{N-1\to N}, N, T)], ..., \mathbf{E}[\Phi(\emptyset, \mathcal{I}_{1\to N}, N, T)]$$

Notice that in $\mathbb{A}$, whenever an inventory arc is added, a flexibility arc is removed. Thus, our goal is to study how changing the *quantity* of inventory and flexibility arcs affects the performance of a network. In a setting where a firm has a limited budget and is forced to choose between flexibility and inventory at a particular plant, it is useful to understand the impact of this choice on the performance of the network.

Performance of all configuration in $\mathbb{A}$ can be analyzed using a relatively straightforward extension to the methodology introduced by Simchi-levi and Wei [55], but the analysis requires several building blocks. In Section 4.5.1 we analyze the performance of the dedicated configuration $\mathbf{E}[\Phi(\emptyset, \mathcal{I}_{1\to N}, N, T)$. In Section 4.5.2 we analyze the performance of a modified dedicated configuration $\mathbf{E}[\tilde{\Phi}(\emptyset, \mathcal{I}_{1\to N}, N, T)$ which serves as an important subproblem in the next Section 4.5.3, which we analyze the performance of an open chain with only the last plants can hold inventory $\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)]$. Finally in Section 4.5.4 we show how all of these building blocks can be used to analyze the performance of any configuration in $\mathbb{A}$. The results of a computational analysis are presented in Section 4.5.5 and other related configurations that can be analyzed in the same way are shown in Section 4.5.6.

4.5.1 A Dedicated Configuration with Inventory

This section focuses on evaluating the performance of a dedicated configuration with inventory $\mathbf{E}[\Phi(\emptyset, \mathcal{I}_{1\to N}, N, T)]$. With the assumption of *i.i.d.* demand, we have

$$\mathbf{E}[\Phi(\emptyset, \mathcal{I}_{1\to N}, N, T)] = N\,\mathbf{E}[\Phi(\emptyset, \mathcal{I}_1, 1, T)] \quad (4.13)$$

From now on we will focus on a single plant-product pair with inventory. For brevity, we use $\mathbf{E}[\Phi(\mathcal{I}, 1, T)]$ to denote $\mathbf{E}[\Phi(\emptyset, \mathcal{I}_1, 1, T)]$. Let $\Phi(\mathcal{I}, 1, T; d)$ denote the performance given a particular demand realization $d \in \mathcal{D}$. Finding $\Phi(\mathcal{I}, 1, T; d)$ is equivalent to solving a maximum flow problem on a simpler graph as shown in Figure 4.5: one node per period from 1 to T plus a source s_0 and a sink s_∞. In the figure, the number next to an arc is the arc capacity.

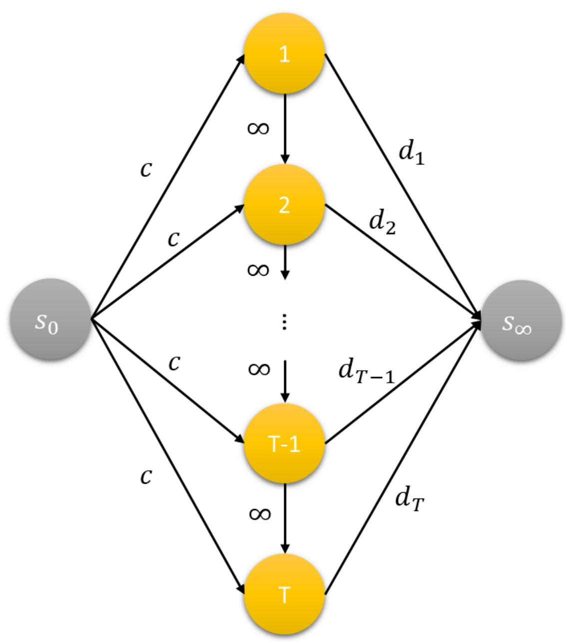

Figure 4.5: Maximum flow formulation of $\Phi(\mathcal{I}, 1, T)$

For the network shown in Figure 4.5, let X_t represent flow on arc (s_0, t), Y_t represent the flow on arc (t, s_∞), and I_t represent the flow on arc $(t, t+1)$. Then $\Phi(\mathcal{I}, 1, T; d)$ can be written as a linear program:

$$\Phi(\mathcal{I},1,T;d) = \max \sum_{t=1}^{T} Y_t \tag{4.14}$$

$$s.t. \quad X_t \leq c \quad \forall t \tag{4.15}$$

$$Y_t \leq d_t \quad \forall t \tag{4.16}$$

$$I_{t-1} + X_t = Y_t + I_t \quad \forall t \tag{4.17}$$

$$X_t, Y_t, I_t \geq 0 \quad \forall t \tag{4.18}$$

Since $\Phi(\mathcal{I},1,T;d)$ has a simple structure, the optimal solution can be obtained by a greedy algorithm in $O(T)$ described in the following lemma.

Lemma 4.5.1.
The maximum flow $\Phi(\mathcal{I},1,T;d)$ can be obtained through the following procedure:
For $t = 1,\ldots,T$, assign the production capacity at time t to satisfy the first unmet demand from period t to T.

Proof: If we apply the procedure given in Lemma 4.5.1 to find a feasible flow from $t = 1$ to T, we reach one of three cases: (1) capacity in all periods is fully utilized, or (2) demand in all periods are satisfied, or (3) there is unused capacity in some periods and unmet demand in some periods, but the first unused capacity occurs later than the last unmet demand.

If (1) and/or (2) apply, clearly the feasible flow is the maximum flow. If (3) applies, W.l.o.g. suppose we have a unused capacity c' in period t and an unmet demand d' in period s where $s < t$. If we are forced to satisfy the demand in period s, then we must borrow the capacity in period k such that $k < s$ because there is no unused capacity before period s, which creates an unmet demand in period k. This can be repeated until the unmet demand happens in period $k = 1$, which cannot be met by any feasible flow. Therefore we cannot increase the flow, so the feasible flow found by applying Lemma 4.5.1 is the maximum flow.

□

The computation of $\mathbf{E}[\Phi(\mathcal{I},1,T)]$ can be done recursively based on Lemma 4.5.1. Let Q^t be the total unmet demand for $\Phi(\mathcal{I},1,t)$. Clearly Q^t is a random variable on support $(0,1,2,\ldots,tc-1,tc)$. Because the demand is *i.i.d.*, we can add one more time period at the very beginning, which leads to two cases: (1) if the demand in the first period $D > c$, then by Lemma 4.5.1 $X_1 = c, Y_1 = c$, and everything from period 2 to $t+1$ is isomorphic to $\Phi(\mathcal{I},1,t)$, or (2) if the demand in the first period $D \leq c$, then there is unused capacity $c - D$ in the first period. By Lemma 4.5.1, $X_1 = \min\{c, Q^t + D\}$ and Q^{t+1} can be updated as $Q^{t+1} = (Q^t + D - c)^+$. Therefore $\mathbf{E}[\Phi(\mathcal{I},1,T+1)]$ can be calculated recursively as

$$\mathbf{E}[\Phi(\mathcal{I},1,t+1)] = \mathbf{E}[\Phi(\mathcal{I},1,t)] + \mathbf{E}[(\min\{c, Q^t + D\})] \tag{4.19}$$

A sketch of the algorithm follows:

Algorithm 4.5.1: Computation $\mathbf{E}[\Phi(\mathcal{I},1,T)]$

Step 1: Initialization

$$\mathbf{q}_1 = \left[\sum_{i=0}^{c} p_i \; p_{c+1} \; p_{c+2} \cdots \; p_{2c} \right]^\top \tag{4.20}$$

$$\mathbf{P}_1 = \begin{bmatrix} \sum_{i=0}^{c} p_i & \sum_{i=0}^{c-1} p_i & \sum_{i=0}^{c-2} p_i & \cdots & p_0 \\ p_{c+1} & p_c & p_{c-1} & \cdots & p_1 \\ p_{c+2} & p_{c+1} & p_c & \cdots & p_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ p_{2c} & p_{2c-1} & p_{2c-2} & \cdots & p_c \\ 0 & p_{2c} & p_{2c-1} & \cdots & p_{c+1} \\ 0 & 0 & p_{2c} & \cdots & p_{c+2} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & p_{2c} \end{bmatrix} \tag{4.21}$$

Step 2: Recursive Expression

$$\mathbf{E}[\Phi(\mathcal{I},1,t+1)] = \mathbf{E}[\Phi(\mathcal{I},1,t)] + c\,\mathbf{e}^\top \overline{\mathbf{p}} + \underline{\mathbf{p}}^\top \mathbf{C}\tilde{\mathbf{q}}_\mathbf{t} \tag{4.22}$$
$$\mathbf{q_{t+1}} = \mathbf{P_t q_t} \tag{4.23}$$
$$\tag{4.24}$$

where

Q^t: Total unmet demand of $\Phi(\mathcal{I},1,T)$. $Q^{t+1} = (Q^t + D - c)^+$

q_k^t: Probability of $Q^t = k$

$\mathbf{q_t}$: $(tc+1) \times 1$ vector represents the p.m.f of Q^t

$= \left[q_0^t \; q_1^t \; q_2^t \; \ldots \; q_{tc-1}^t \; q_{tc}^t \right]^\top$

$\tilde{\mathbf{q}}_\mathbf{t} = \left[q_0^t \; q_1^t \; \ldots \; q_{c-1}^t \; \sum_{i=c}^{tc} q_i^t \right]^\top$

$\mathbf{P_t}$: $((t+1)c+1) \times (tc+1)$ constant matrix

$\underline{\mathbf{p}}$: $(c+1) \times 1$ constant vector

$\overline{\mathbf{p}}$: $c \times 1$ constant vector

$\mathbf{e}$: $c \times 1$ constant vector

$\mathbf{C}$: $(c+1) \times (c+1)$ constant matrix

Complexity: Complexity of computing $\mathbf{E}[\Phi(\mathcal{I},1,T)]$ is $O(T^2c^2)$.

The details of Algorithm 4.5.1 and the proof of correctness can be found in Appendix A.3.1.

4.5.2 Residual Problem

Before analyzing more complex configurations, we consider an extension of the case discussed in the previous section: a dedicated configuration with inventory where the capacity is no longer a constant c, but is drawn randomly from an *i.i.d.* distribution correlated with demand. This configuration, which we call the *residual problem* $\tilde{\Phi}$ (for reasons that will become clear in the next section), turns out to be an important subproblem for our later analysis of the open chain configuration with inventory. Let $\tilde{C}$ be the random capacity (residual capacity) and $\tilde{D}$ be the random demand (residual demand), and given a realization of capacity and demand $\tilde{c}_t$ and $\tilde{d}_t$, let $\tilde{\Phi}(\emptyset,\mathcal{I}_1,1,T;\tilde{c},\tilde{d})$ be

$$\begin{aligned}
\tilde{\Phi}(\emptyset,\mathcal{I}_1,1,T;\tilde{c},\tilde{d}) = \max \quad & \sum_{t=1}^{T} Y_t && (4.25)\\
\text{s.t.} \quad & X_t \le \tilde{c}_t \quad \forall t && (4.26)\\
& Y_t \le \tilde{d}_t \quad \forall t && (4.27)\\
& I_{t-1} + X_t = Y_t + I_t \quad \forall t && (4.28)\\
& X_t, Y_t, I_t \ge 0 \quad \forall t && (4.29)
\end{aligned}$$

The joint probability mass function of capacity and demand is

$$\begin{aligned}
\mathbf{P}(\tilde{C}=c,\tilde{D}=0) &= \gamma_c\\
\mathbf{P}(\tilde{C}=c-1,\tilde{D}=0) &= \gamma_{c-1}\\
&\vdots\\
\mathbf{P}(\tilde{C}=1,\tilde{D}=0) &= \gamma_1\\
\mathbf{P}(\tilde{C}=0,\tilde{D}=0) &= \gamma_0+\delta_0-1=\zeta_0\\
\mathbf{P}(\tilde{C}=0,\tilde{D}=1) &= \delta_1\\
&\vdots\\
\mathbf{P}(\tilde{C}=0,\tilde{D}=c) &= \delta_c
\end{aligned}$$

with

$$\boldsymbol{e}^\top\boldsymbol{\gamma} = 1 \qquad (4.30)$$

$$\boldsymbol{e}^\top\boldsymbol{\delta} = 1 \qquad (4.31)$$

where $\boldsymbol{\gamma} = [\ \gamma_0\ \gamma_1\ \ldots\ \gamma_c\ ]^\top, \boldsymbol{\delta} = [\ \delta_0\ \delta_1\ \ldots\ \delta_c\ ]^\top$. Notice that $\tilde{C}$ and $\tilde{D}$ can't be strictly positive at the same time. As in the previous section, $\mathbf{E}[\tilde{\Phi}(\mathcal{I}, 1, t)]$ can be computed recursively as

$$\mathbf{E}[\tilde{\Phi}(\mathcal{I}, 1, t+1)] \quad = \quad \mathbf{E}[\tilde{\Phi}(\mathcal{I}, 1, t)] + \mathbf{E}[(\min\{\tilde{C}, Q^t + \tilde{D}\})] \tag{4.32}$$

A sketch of the algorithm follows:

Algorithm 4.5.2: Computation $\mathbf{E}[\tilde{\Phi}(\mathcal{I}, 1, T)]$

Step 1: Determine $\boldsymbol{\gamma}$ and $\boldsymbol{\delta}$ (See Algorithm 4.5.3)

Step 2: Initialization

$$\mathbf{q_1} \quad = \quad [\ \delta_0\ \delta_1\ \delta_2 \ldots\ \delta_c\ ]^\top \tag{4.33}$$

$$\tilde{\mathbf{P}}_1 \quad = \quad \begin{bmatrix} \delta_0 & \sum_{i=1}^{c} \gamma_i & \sum_{i=2}^{c} \gamma_i & \cdots & \gamma_c \\ \delta_1 & \zeta_0 & \gamma_1 & \cdots & \gamma_{c-1} \\ \delta_2 & \delta_1 & \zeta_0 & \cdots & \gamma_{c-2} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \delta_c & \delta_{c-1} & \delta_{c-2} & \cdots & \zeta_0 \\ 0 & \delta_c & \delta_{c-1} & \cdots & \delta_1 \\ 0 & 0 & \delta_c & \cdots & \delta_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & \delta_c \end{bmatrix} \tag{4.34}$$

Step 3: Recursive expression

$$\mathbf{E}[\tilde{\Phi}(\mathcal{I}, 1, t+1)] \quad = \quad \mathbf{E}[\tilde{\Phi}(\mathcal{I}, 1, t)] + \boldsymbol{\gamma}^\top \tilde{\mathbf{C}} \tilde{\mathbf{q}}_\mathbf{t} \tag{4.35}$$

$$\mathbf{q_{t+1}} \quad = \quad \tilde{\mathbf{P}}_\mathbf{t}^\top \mathbf{q_t} \tag{4.36}$$

$$\tag{4.37}$$

where

Q^t: Total unmet demand of $\tilde{\Phi}(\mathcal{I}, 1, T)$. $Q^{t+1} = (Q^t + \tilde{D} - \tilde{C})^+$

q_k^t: Probability of $Q^t = k$

$\mathbf{q_t}$: $(tc \times 1)$ vector represents the p.m.f. of Q^t

$\quad = \quad [\ q_0^t\ q_1^t\ q_2^t\ \ldots\ q_{tc-1}^t\ q_{tc}^t\ ]^\top$

$$\tilde{\mathbf{q}}_\mathbf{t} \quad = \quad \left[q_0^t\ q_1^t\ \ldots\ q_{c-1}^t\ \sum_{i=c}^{tc} q_i^t \right]^\top$$

$\boldsymbol{\gamma} \quad = \quad [\ \gamma_0\ \gamma_1\ \ldots\ \gamma_c\ ]^\top$

$\boldsymbol{\delta} \quad = \quad [\ \delta_0\ \delta_1\ \ldots\ \delta_c\ ]^\top$

$\tilde{\mathbf{P}}_\mathbf{t}$: $((t+1)c+1)\times(tc+1)$ constant matrix
$\tilde{\mathbf{C}}$: $(c+1)\times(c+1)$ constant matrix

Complexity: The complexity of computing $\mathbf{E}[\tilde{\Phi}(\mathcal{I},1,T)]$ is $O(T^2c^2)$.

The details of Algorithm 4.5.2 and the proof of correctness can be found in Appendix A.3.2.

4.5.3 Open Chain with Inventory

In this section we analyze the performance of the open chain configuration in which only the last non-flexible plant can hold inventory: $\mathbf{E}[\Phi(\mathcal{L}_{1\to N},\mathcal{I}_N,N,T)]$. The idea is to decompose the computation of $\Phi(\mathcal{L}_{1\to N},\mathcal{I}_N,N,T)$ in two steps: step 1 solves a problem with only open chain flexibility which is an extension of the approach developed in Simchi-Levi and Wei (2012)[55], and step 2 solves a residual problem which only involves inventory as in Section 4.5.2.

To be more specific, consider a single period open chain without inventory $\Phi(\mathcal{L}_{1\to N},\emptyset,N,1)$. Let $\tilde{C}_N$ be the unused capacity at plant N. $\tilde{C}_N$ is a random variable on support $[\,0\ 1\ 2\ \dots\ c]^\top$ and let $\boldsymbol{\gamma}_N=[\,\gamma_{N_0}\ \gamma_{N_1}\ \dots\ \gamma_{N_c}\,]^\top$ be the vector representing its probability mass function. Similarly, let $\tilde{D}_N$ be the unmet demand for product N. Since demand D is at most $2c$, $\tilde{D}_N$ also lies on support $[\,0\ 1\ 2\ \dots\ c]^\top$. Let $\boldsymbol{\delta}_N=[\,\delta_{N_0}\ \delta_{N_1}\ \dots\ \delta_{N_c}\,]^\top$ be the vector representing its probability mass function. Now consider the open chain configuration with inventory $\Phi(\mathcal{L}_{1\to N},\mathcal{I}_N,N,T)$, which can be viewed as T copies of single period open chain without inventory $\Phi(\mathcal{L}_{1\to N},\emptyset,N,1)$ plus the inventory capability at plant N. Then, we immediately know

$$\mathbf{E}[\Phi(\mathcal{L}_{1\to N},\mathcal{I}_N,N,T)]>\mathbf{E}[\Phi(\mathcal{L}_{1\to N},\emptyset,N,T)] \tag{4.38}$$

because whenever there is unused capacity $\tilde{C}_N>0$ in some period t and unmet demand $\tilde{D}_N$ in some other periods s for $s>t$, the inventory capacity in $\Phi(\mathcal{L}_{1\to N},\mathcal{I}_N,N,T)$ can satisfy more demand by utilizing the unused capacity in period s and unmet demand in period t. Define *residual performance* of a configuration $(\mathcal{L}_{1\to N},\mathcal{I}_N,N,T)$ as the difference

$$\mathbf{E}[\tilde{\Phi}(\mathcal{L}_{1\to N},\mathcal{I}_N,N,T)]=\mathbf{E}[\Phi(\mathcal{L}_{1\to N},\mathcal{I}_N,N,T)]-\mathbf{E}[\Phi(\mathcal{L}_{1\to N},\emptyset,N,T)] \tag{4.39}$$

so residual performance is just the marginal benefit of adding the extra inventory arc. If residual performance can be analyzed, then immediately we know $\mathbf{E}[\Phi(\mathcal{L}_{1\to N},\mathcal{I}_N,N,T)]$. Later, in the proof of Algorithm 4.5.3, we show that the residual performance is exactly equal to $\tilde{\Phi}$ discussed in Section 4.5.2 with $\boldsymbol{\gamma}=\boldsymbol{\gamma}_N,\boldsymbol{\delta}=\boldsymbol{\delta}_N$ (see Figure 4.6 for an illustration).

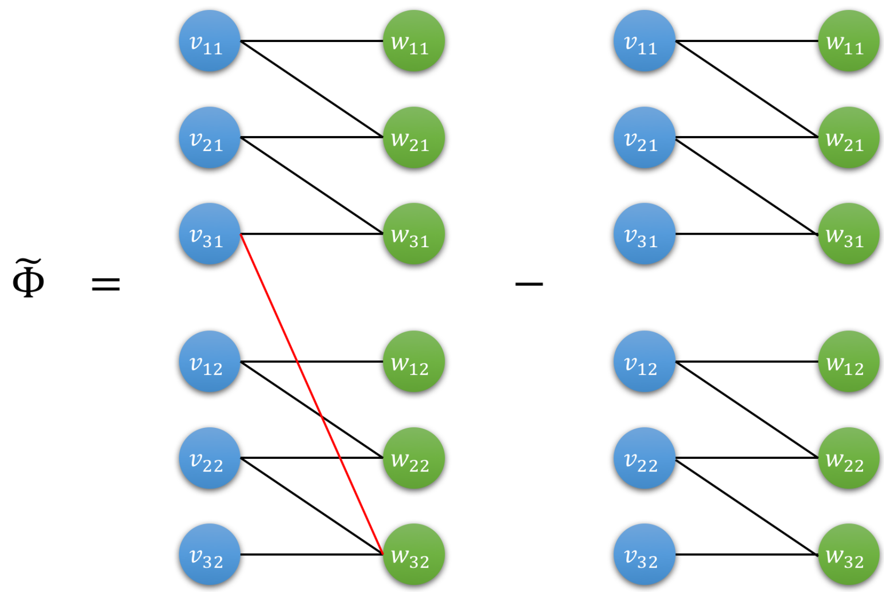

Figure 4.6: Definition of the residual performance $\tilde{\Phi}(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)$

Sketch of the algorithm for computing $\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)]$ is shown below. Since Simchi-Levi and Wei (2012) did not explicitly explain how to compute $\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \emptyset, N, 1)]$, we show the computation of this quantity as well.

Algorithm 4.5.3: Computation of $\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)]$

Initialization:

$$\gamma_0 = [\, 1\, 0\, 0 \ldots\, 0\,]^\top \tag{4.40}$$

Step 1: Compute the probability of unused capacity at plant N and unmet demand of product N: $\boldsymbol{\gamma}_N, \boldsymbol{\delta}_N$ recursively as

$$\boldsymbol{\gamma}_N = \mathbf{A}\boldsymbol{\gamma}_{N-1} \tag{4.41}$$
$$\boldsymbol{\delta}_N = \mathbf{B}\boldsymbol{\gamma}_{N-1} \tag{4.42}$$

where

A: $(c+1)\times(c+1)$ constant matrix
B: $(c+1)\times(c+1)$ constant matrix

Step 2: Compute the performance of open chain without inventory $\mathbf{E}[\Phi(\mathcal{L}_N, \emptyset, N, T)]$ recursively as

$$\mathbf{E}[\Phi(\mathcal{L}_N, \emptyset, N, 1)] = \mathbf{E}[\Phi(\mathcal{L}_{N-1}, \emptyset, N-1, 1)] + \mathbf{E}[\min\{c + \tilde{C}_{N-1}, D\}] \tag{4.43}$$
$$= \mathbf{E}[\Phi(\mathcal{L}_{N-1}, \emptyset, N-1, 1)] + \boldsymbol{\gamma}_{N-1}^\top \mathbf{K}\mathbf{p} \tag{4.44}$$

where

$\mathbf{K}$: $(c+1) \times (2c+1)$ constant matrix
$\mathbf{p} = [\, p_0\, p_1\, \dots\, p_{2c}\,]^\top$

Step 3: Compute residual performance $\mathbf{E}[\tilde{\Phi}(\mathcal{L}_N, \mathcal{I}_N, N, T)]$ with $\boldsymbol{\gamma} = \boldsymbol{\gamma}_N, \boldsymbol{\delta} = \boldsymbol{\delta}_N$ using Algorithm 4.5.2.

Step 4: $\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)] = \mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \emptyset, N, T)] + \mathbf{E}[\tilde{\Phi}(\mathcal{L}_{1\to N}, \mathcal{I}_N, 1, T)]$

Complexity: The complexity of computing $\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)]$ is $O(Nc^2 + T^2c^2)$

Details of Algorithm 4.5.3 and a proof of correctness can be found in Appendix A.3.3.

4.5.4 Performance of Configurations in $\mathcal{P}$

Recall that the goal of this section is to compare the performance of the sequence of configurations $\mathbb{A} = \{\mathcal{A}_0, \dots, \mathcal{A}_N\}$ presented at the beginning of Section 4.5:

$$\mathbf{E}[\Phi(\mathcal{C}, \emptyset, N, T)], \mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)], \mathbf{E}[\Phi(\mathcal{L}_{1\to N-1}, \mathcal{I}_{N-1\to N}, N, T)], \dots, \mathbf{E}[\Phi(\emptyset, \mathcal{I}_{1\to N}, N, T)]$$

Using the building blocks from Sections 4.5.1 to 4.5.3, we complete the comparison between flexibility and inventory in this section. Under the assumption of *i.i.d.* demand, observe that:

$$\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_{k\to N}, N, T)] = \mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_k, k, T)] + (N-k)\mathbf{E}[\Phi(\emptyset, \mathcal{I}_1, 1, T)] \tag{4.45}$$

which means $\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_{k\to N}, N, T)]$ can be decomposed into two parts: (1) $(N-k)$ copies of a dedicated configuration with inventory as in Section 4.5.1, and (2) the open chain configuration with inventory at the last non-flexible plant as in Section 4.5.3. An example for $N = 3, T = 2, k = 2$ is shown in Figure 4.7:

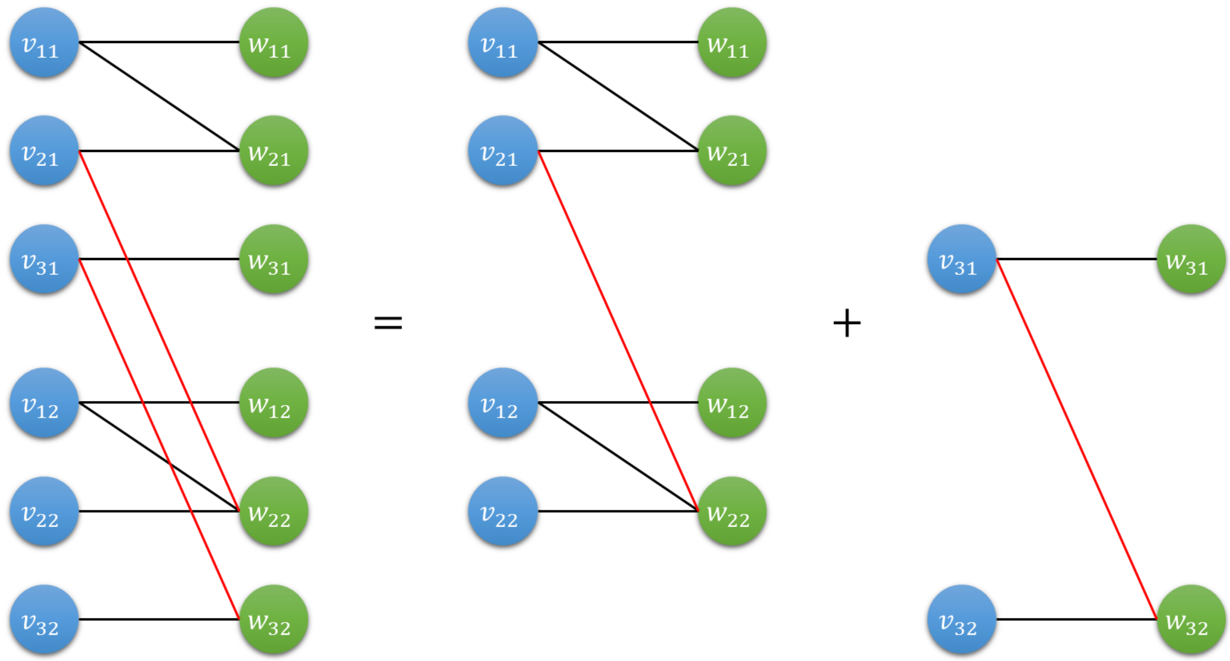

Figure 4.7: The decomposition of $\mathbf{E}[\Phi(\mathcal{L}_{1\to 2}, \mathcal{I}_{2\to 3}, 3, 2)]$

The algorithm of computing $\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_{k\to N}, N, T)]$ is summarized as following:

Algorithm 4.5.4: Computation of $\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_{k\to N}, N, T)]$

Step 1: Compute $\mathbf{E}[\Phi(\emptyset, \mathcal{I}_1, 1, T)]$ using Algorithm 4.5.1.

Step 2: Compute $\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_k, k, T)]$ using Algorithm 4.5.3.

Step 3: Compute $\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_{k\to N}, N, T)]$ as

$$\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_{k\to N}, N, T)] = \mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_k, k, T)] + (N-k)\mathbf{E}[\Phi(\emptyset, \mathcal{I}_1, 1, T)] \tag{4.46}$$

Complexity: The complexity of computing $\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_{k\to N}, N, T)]$ is $O(Nc^2 + T^2c^2)$

4.5.5 Computational Analysis

In this section we numerically test our algorithm developed in previous sections to give insight into one of the questions posed in the introduction to this chapter: *If a firm wants to build a new plant while keeping the current production network dedicated as before, should the new plant be flexible?*

Specifically, we consider a setting with $N = 4$ plants and products with plant capacity $c = 5$ and product demand $\mathbf{E}[D] = c = 5$. We vary the planning horizon from $T = 2$ to $T = 12$. We set parameters to capture the current situation of the firm with which we worked: few products and a 12 month forecast horizon. We consider three types of demand: Bernoulli, Uniform, Binomial which has decreasing variability. For brevity, with fixed N and T we replace the notation $\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_{k\to N}, N, T)]$ with $\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_{k\to N})]$. The detailed results are shown in Tables 4.2 to 4.4.

	Bernoulli				
T	$\mathbf{E}[\Phi(\mathcal{C}, \emptyset)]$	$\mathbf{E}[\Phi(\mathcal{L}_{1\to 4}, \mathcal{I}_4)]$	$\mathbf{E}[\Phi(\mathcal{L}_{1\to 3}, \mathcal{I}_{3\to 4})]$	$\mathbf{E}[\Phi(\mathcal{L}_{1\to 2}, \mathcal{I}_{2\to 4})]$	$\mathbf{E}[\Phi(\emptyset, \mathcal{I}_{1\to 4})]$
2	**30.00**	28.13	26.88	25.63	25.00
3	**45.00**	42.66	41.41	40.16	40.00
4	**60.00**	57.30	56.37	55.43	56.25
5	**75.00**	72.02	71.40	70.77	72.50
6	**90.00**	86.79	86.63	86.48	89.38
7	105.00	101.59	101.91	102.22	**106.25**
8	120.00	116.43	117.31	118.19	**123.52**
9	135.00	131.29	132.73	134.18	**140.78**
10	150.00	146.16	148.24	150.32	**158.32**
11	165.00	161.05	163.77	166.48	**175.86**
12	180.00	175.96	179.36	182.76	**193.60**

Table 4.2: $\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_{k\to N}, N, T)]$ with Bernoulli demand

	Uniform				
T	$\mathbf{E}[\Phi(\mathcal{C}, \emptyset)]$	$\mathbf{E}[\Phi(\mathcal{L}_{1\to 4}, \mathcal{I}_4)]$	$\mathbf{E}[\Phi(\mathcal{L}_{1\to 3}, \mathcal{I}_{3\to 4})]$	$\mathbf{E}[\Phi(\mathcal{L}_{1\to 2}, \mathcal{I}_{2\to 4})]$	$\mathbf{E}[\Phi(\emptyset, \mathcal{I}_{1\to 4})]$
2	**34.69**	33.01	32.09	31.33	30.91
3	**52.04**	49.82	48.86	48.16	47.95
4	**69.39**	66.70	65.80	65.27	65.40
5	**86.74**	83.60	82.85	82.57	83.12
6	**104.08**	100.54	99.98	100.02	101.05
7	**121.43**	117.48	117.16	117.57	119.13
8	**138.78**	134.44	134.39	135.20	137.33
9	**156.12**	151.40	151.66	152.90	155.64
10	173.47	168.37	168.96	170.66	**174.04**
11	190.82	185.34	186.28	188.46	**192.51**
12	208.17	202.32	203.62	206.31	**211.05**

Table 4.3: $\mathbf{E}[\Phi(\mathcal{L}_{1\to k}, \mathcal{I}_{k\to N}, N, T)]$ with Uniform demand

	Binomial				
T	$\mathbf{E}[\Phi(\mathcal{C},\emptyset)]$	$\mathbf{E}[\Phi(\mathcal{L}_{1\to 4},\mathcal{I}_4)]$	$\mathbf{E}[\Phi(\mathcal{L}_{1\to 3},\mathcal{I}_{3\to 4})]$	$\mathbf{E}[\Phi(\mathcal{L}_{1\to 2},\mathcal{I}_{2\to 4})]$	$\mathbf{E}[\Phi(\emptyset,\mathcal{I}_{1\to 4})]$
2	**37.49**	36.71	36.29	35.96	35.78
3	**56.24**	55.22	54.76	54.45	54.33
4	**74.98**	73.75	73.32	73.07	73.08
5	**93.73**	92.31	91.93	91.79	91.96
6	**112.48**	110.87	110.58	110.58	110.93
7	**131.22**	129.44	129.26	129.42	129.98
8	**149.97**	148.01	147.96	148.30	149.09
9	**168.72**	166.59	166.68	167.22	168.25
10	**187.46**	185.17	185.41	186.16	187.46
11	206.21	203.75	204.16	205.13	**206.70**
12	224.95	222.33	222.91	224.12	**225.97**

Table 4.4: $\mathbf{E}[\Phi(\mathcal{L}_{1\to k},\mathcal{I}_{k\to N},N,T)]$ with Binomial demand

Tables 4.2 to 4.4 answer the question raised by the firm: if the firm want to build a new plant for a new drug while keep the current plants as dedicated before, should the new plant has flexibility? This question is depicted in Figure 4.8 for a case of $N=3,T=2$ and plant product 1 is treated as the new plant product pair (represented by the outlined node with dashed arcs).

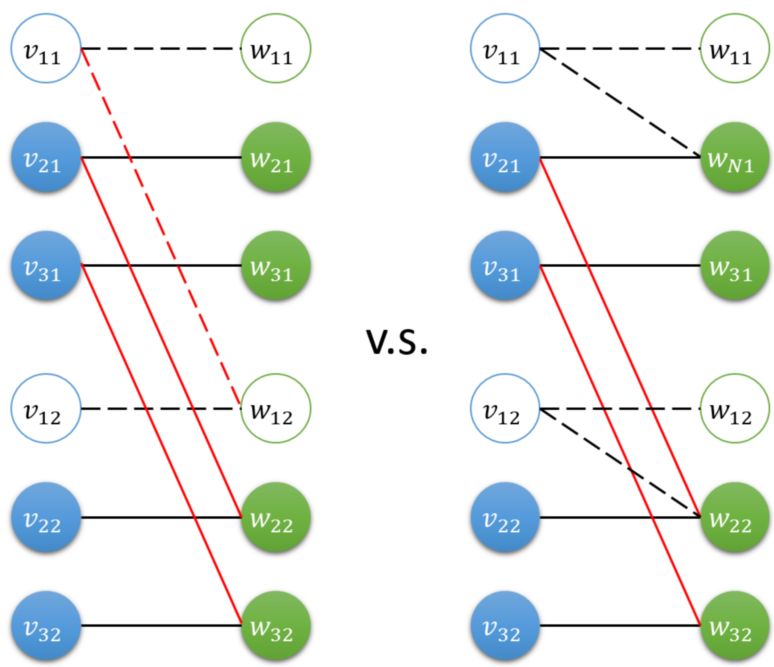

Figure 4.8: Should the new plant has inventory or flexibility?

Then the question above translates to comparing $\mathbf{E}[\Phi(\emptyset,\mathcal{I}_{1\to 4})]$ and $\mathbf{E}[\Phi(\mathcal{L}_{1\to 2},\mathcal{I}_{2\to 4})]$, and as shown in the Table 4.2-4.4 we see when T is small the new plant should be flexible, while when T is large the new plant should feature inventory capability. This conclusion generalizes to other configurations as well: in general, when T is small, the best configuration features flexibility without inventory $(\mathcal{C}, \emptyset)$, while when T is large, the dedicated structure with inventory $(\emptyset, \mathcal{I}_{1\to N})$ outperforms. The critical switching T increases as demand variation decreases. Intermediate configurations featuring some plants with flexibility and the other plants with inventory is never optimal.

Managerial Insight: These computational tests suggest that if the firm has a reliable multi-period demand forecast (i.e. T is large), it should stick to a dedicated design employing inventory rather than sacrificing inventory for flexibility. In other words, if each plant is limited to flexibility or inventory, but not both, flexibility and inventory acts as substitutes. The best configuration is either flexibility without inventory or inventory without flexibility for all plants.

4.5.6 Other Related Configurations

Consider a case where the firm's the current configuration is the closed chain without inventory $(\mathcal{C}, \emptyset)$ and has reliable demand forecast. From our conclusion in previous section, under this situation a dedicated network with inventory may have better performance so the firm want to switch. However assume the firm only has limited budget so only a subset of plants $k < N$ is chosen, in which each plant adds inventory capability while removes its flexibility, then which k plants should the firm choose? To answer this question, we consider another sequence of configurations $\mathbb{A}' = \{\mathcal{A}_1, \ldots, \mathcal{A}_{N/2}\}$ starting with $\mathcal{A}_1 = (\mathcal{L}_{1\to N-1}, \mathcal{I}_{N-1\to N}, N, T)$, then gradually shifting the inventory arc from plant $N-1$ to plant $N-2, N-3, \ldots$ until it reaches plant $N/2$. An example of $\mathbb{A}'$ with $N = 6$ is shown in Figure 4.9.

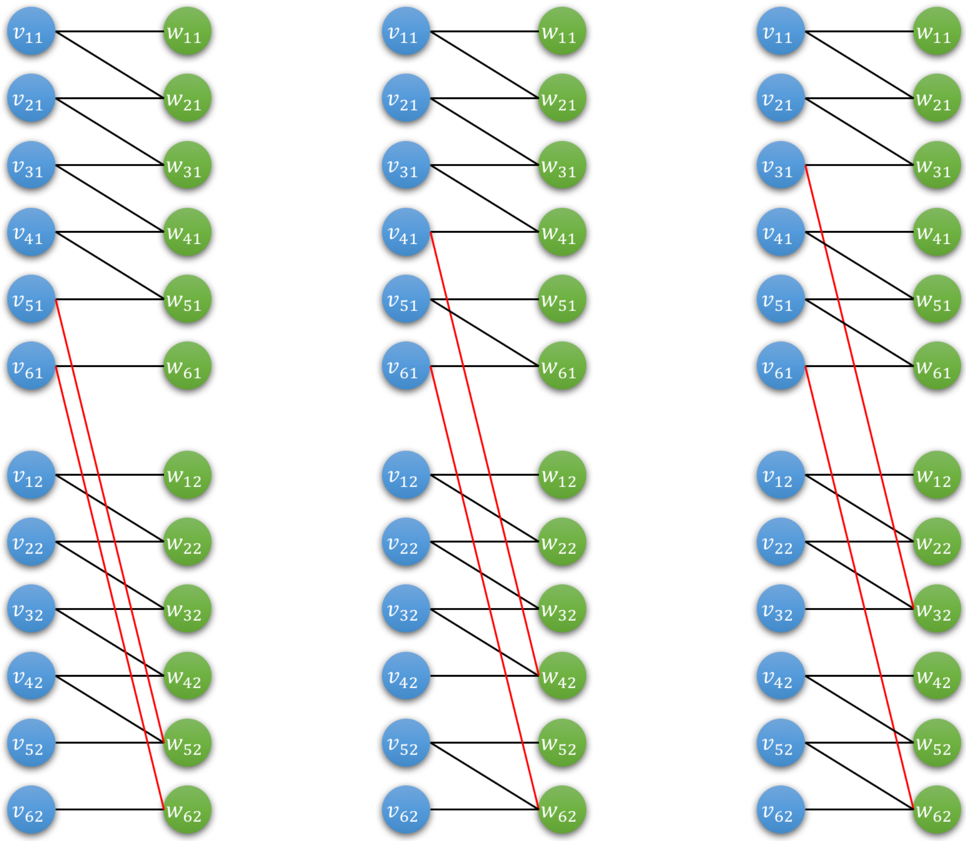

Figure 4.9: The sequence of configurations in $\mathbb{A}'$

In contrast to the sequence $\mathbb{A}$ in the previous section which we study how changing the *quantity* of inventory and flexibility arcs affects the performance, in this subsection we use the sequence $\mathbb{A}'$ to study how changing the *position* of inventory and flexibility arcs affects performance. It is easy to see that each configuration in $\mathbb{A}'$ can be decomposed into two parts, both of which have been studied in previous sections. Specifically, we want to compare the performance of

$$
\begin{aligned}
\mathbf{E}[\Phi(\mathcal{A}_1, N, T)] &= \mathbf{E}[\Phi(\mathcal{L}_{1\to N-1}, \mathcal{I}_{N-1}, N-1, T)] + \mathbf{E}[\Phi(\emptyset, \mathcal{I}_1, 1, T)] \\
\mathbf{E}[\Phi(\mathcal{A}_2, N, T)] &= \mathbf{E}[\Phi(\mathcal{L}_{1\to N-2}, \mathcal{I}_{N-2}, N-2, T)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 2}, \mathcal{I}_2, 2, T)] \\
\mathbf{E}[\Phi(\mathcal{A}_3, N, T)] &= \mathbf{E}[\Phi(\mathcal{L}_{1\to N-3}, \mathcal{I}_{N-3}, N-3, T)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 3}, \mathcal{I}_3, 3, T)] \\
& \dots \\
\mathbf{E}[\Phi(\mathcal{A}_{N/2}, N, T)] &= \mathbf{E}[\Phi(\mathcal{L}_{1\to N/2}, \mathcal{I}_{N/2}, N/2, T)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to N/2}, \mathcal{I}_{N/2}, N/2, T)]
\end{aligned}
$$

We test the performance of the configurations in $\mathbb{A}'$ using an instance with $N = 12$, $T = 4$ or $T = 12$. The results are summarized in Tables 4.5 and 4.6 and we see $\mathbf{E}[\Phi(\mathcal{A}_1, N, T)]$ always dominates.

$T = 4$	Bernoulli	Uniform	Binomial
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 11}, \mathcal{I}_{11}, 11, 4)] + \mathbf{E}[\Phi(\emptyset, \mathcal{I}_1, 1, 4)$	176.37	205.96	225.34
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 10}, \mathcal{I}_{10}, 10, 4)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 2}, \mathcal{I}_2, 2, 4)]$	174.61	204.58	224.42
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 9}, \mathcal{I}_9, 9, 4)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 3}, \mathcal{I}_3, 3, 4)]$	174.61	203.86	223.78
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 8}, \mathcal{I}_8, 8, 4)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 4}, \mathcal{I}_4, 4, 4)]$	174.61	203.51	223.37
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 7}, \mathcal{I}_7, 7, 4)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 5}, \mathcal{I}_5, 5, 4)]$	174.61	203.36	223.13
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 6}, \mathcal{I}_6, 6, 4)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 6}, \mathcal{I}_6, 6, 4)]$	174.61	203.32	223.05

Table 4.5: Comparison of configurations in $\mathbb{A}'$ when $T = 4$

$T = 12$	Bernoulli	Uniform	Binomial
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 11}, \mathcal{I}_{11}, 11, 12)] + \mathbf{E}[\Phi(\emptyset, \mathcal{I}_1, 1, 12)$	539.36	623.60	678.29
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 10}, \mathcal{I}_{10}, 10, 12)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 2}, \mathcal{I}_2, 2, 12)]$	531.92	618.82	675.42
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 9}, \mathcal{I}_9, 9, 12)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 3}, \mathcal{I}_3, 3, 12)]$	531.92	616.09	673.26
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 8}, \mathcal{I}_8, 8, 12)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 4}, \mathcal{I}_4, 4, 12)]$	531.92	614.77	671.86
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 7}, \mathcal{I}_7, 7, 12)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 5}, \mathcal{I}_5, 5, 12)]$	531.92	614.20	671.06
$\mathbf{E}[\Phi(\mathcal{L}_{1\to 6}, \mathcal{I}_6, 6, 12)] + \mathbf{E}[\Phi(\mathcal{L}_{1\to 6}, \mathcal{I}_6, 6, 12)]$	531.92	614.04	670.80

Table 4.6: Comparison of configurations in $\mathbb{A}'$ when $T = 12$

Managerial Insight: This implies if there is a budget constraint so that the firm can only choose a subset of plants to replace flexibility with inventory capability, these plants should be adjacent to each other so that the network can maintain a longer chain for the flexibility part.

4.6 Plant with both Flexibility and Inventory

In previous sections, we consider a setting where each plant can have either inventory capability or flexible capacity, but not the both. Here, we study the performance of a more general configuration: a closed chain with inventory when a plant can be both flexible and hold inventory, i.e. $\mathbf{E}[\Phi(\mathcal{C}, \mathcal{I}_S, N, T)]$. For brevity, we assume that flexible plant i can only hold inventory for its primary product i, and we call this dedicated inventory capability. Our analysis only requires minor modification to capture the more general flexible inventory case.

To begin with, we consider a sequence of configurations $\mathbb{B} = \{\mathcal{B}_0, \mathcal{B}_1, \ldots, \mathcal{B}_N\}$, with the initial configuration $\mathcal{B}_0$ features a closed chain configuration without inventory capability $(\mathcal{C}, \emptyset, N, T)$. Subsequent configurations the inventory to plant 1, then plant 2, then plant 3 and so on until plant N are added. For the case where $N = 3, T = 2$, the set of configurations in $\mathbb{B}$ is shown in Figure 4.10. Our goal is compare the performance of

$$\mathbf{E}[\Phi(\mathcal{C}, \emptyset, N, T)], \mathbf{E}[\Phi(\mathcal{C}, I_1, N, T)], \mathbf{E}[\Phi(\mathcal{C}, I_{1\to 2}, N, T)], \ldots, \mathbf{E}[\Phi(\mathcal{C}, I_{1\to N}, N, T)]$$

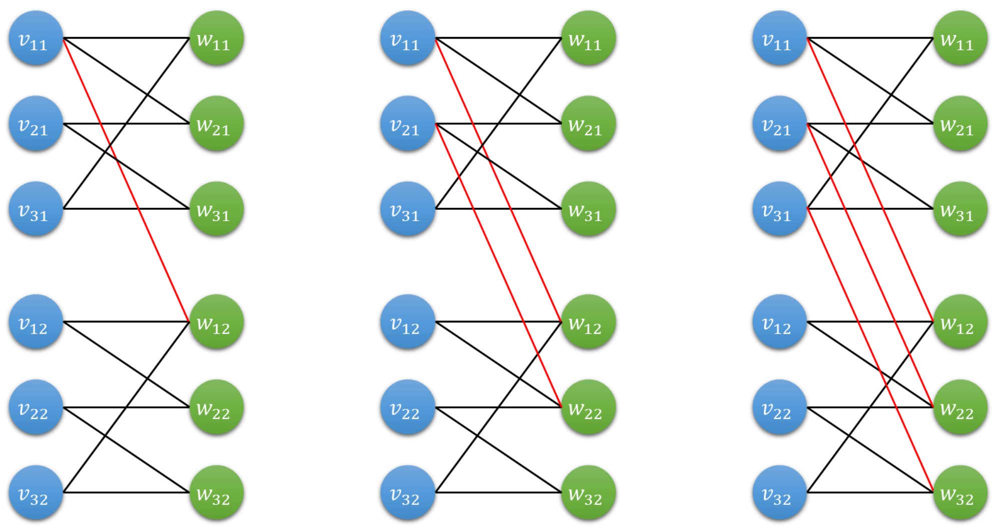

Figure 4.10: The sequence of configurations in $\mathbb{B}$

In contrast to the sequence of configurations in Section 4.5, in this section whenever an inventory arc is added, the corresponding flexibility arc is not removed. This scenario is in particular useful for a firm that currently has a flexibility dense network and want to alter existing plants by adding inventory capabilities. Analyzing configurations in $\mathbb{B}$ enables us to study the marginal benefit of adding inventory arcs. Unfortunately when a plant can have both inventory and flexibility, the analysis seems to be significantly more challenging. In Section 4.6.1 we focus on a special case when demand is *i.i.d.* Bernoulli distributed with planning horizon $T = 2$, and we derive close form expression for computing $\mathbf{E}[\Phi(\mathcal{A}_k, N, T = 2)]$. In Section 4.6.2 we show indirect evidence that there may not have a polynomial time algorithm for computing $\mathbf{E}[\Phi(\mathcal{A}_k, N, T)]$ for general demand or $T > 2$. In Section 4.6.3 we turn to simulation based analysis for the general case.

4.6.1 A Special Case

In this section, we focus on the analysis of configurations in $\mathbb{B}$ under a special case: demand for all products are *i.i.d.* Bernoulli between 0 and $2c$ and planning horizon $T = 2$. W.l.o.g. assume $c = 1$. As in Section 4.5, we first define the residual performance of a configuration as

$$\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)] = \mathbf{E}[\Phi(\mathcal{C}, I_{1\to k}, N, 2)] - \mathbf{E}[\Phi(\mathcal{C}, \emptyset, N, 2)]$$

which is the marginal benefit of allowing plant 1 through k to hold inventory compared to a closed chain without inventory. Next we show there is closed form expression for

$\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$ under the special case:

Closed form expression for $\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$

Assume demand for all products are *i.i.d.* Bernoulli (0 *w.p.* $\frac{1}{2}$, 2 *w.p.* $\frac{1}{2}$). Assume $N \geq 3$.
For $k = 1, 2, \ldots, N-2$,

$$\mathbf{E}[\tilde{\Phi}(\mathcal{C}, \mathcal{I}_{1\to k}, N, 2)] = -\frac{1}{72}\left(\frac{1}{4}\right)^k + \frac{1}{12}k + \frac{1}{72} \tag{4.47}$$

For $k = N-1$,

$$\mathbf{E}[\tilde{\Phi}(\mathcal{C}, \mathcal{I}_{1\to N-1}, N, 2)] = \mathbf{E}[\tilde{\Phi}(\mathcal{C}, \mathcal{I}_{1\to N-1}, N+1, 2)] - \left(\frac{1}{2}\right)^{2N+1} \tag{4.48}$$

$$= -\frac{5}{9}\left(\frac{1}{4}\right)^N + \frac{1}{12}N - \frac{5}{72} \tag{4.49}$$

For $k = N$,

$$\mathbf{E}[\tilde{\Phi}(\mathcal{C}, \mathcal{I}_{1\to N}, N, 2)] = -\frac{N}{3}\left(\frac{1}{4}\right)^N + \frac{1}{12}N \tag{4.50}$$

The detailed derivation of this result can be found in the Appendix A.3.4. Here, we sketch the derivation. Notice the residual performance $\mathbf{E}[\tilde{\Phi}(\mathcal{A}_k, N, 2)]$ is closely related to the residual capacity in the first period and residual demand in the second period. To analyze the residual capacity we first introduce some simplified notations. Let $\tilde{C}_i$ denote the residual capacity of plant i in the first period when there is no inventory arcs, and let

$$\mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix} \tag{4.51}$$

denote the joint probability of $\tilde{C}_i = b_i$ for all i from 1 to k given there are N plants in total. Since demand is Bernoulli, clearly b_i can only take on value 0 or 1 (b stands for binary). The Bernoulli demand ensures that given an instance of demand $(d_{11}, \ldots, d_{N1})$ in the first period, there is a *unique* residual capacity vector $(b_1, \ldots, b_k)$. This means the probability measure $\mathbf{P}^N$ is well defined: $\mathbf{P}^N$ is just the sum of the probabilities of all demand scenarios that yield the residual capacity $(b_1, \ldots, b_k)$. Here $\mathbf{P}^N$ says nothing about the inventory capability of plant 1 to k. Plant i may or may not be able to hold inventory. An example illustrating the concept of residual capacity given an instance of demand in the first period is shown in Figure 4.11.

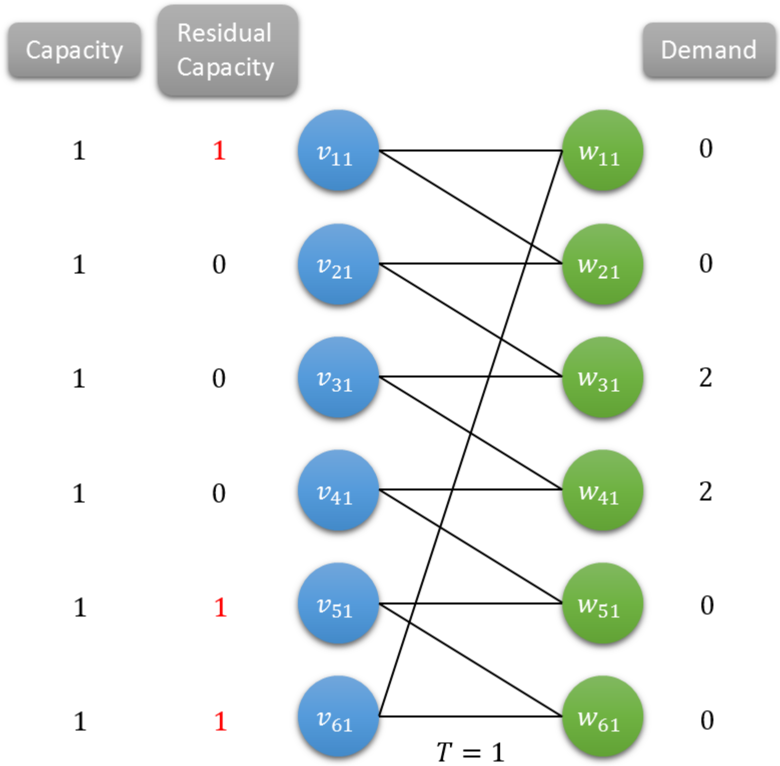

Figure 4.11: Residual capacity given an instance of demand in the first period

In comparison the residual demand given an instance of demand in the second period is not unique. For example, suppose the demand in the second period is $(d_{12} = 0, d_{22} = 2, d_{32} = 2)$, then the residual demand can be either $(0, 1, 0)$ or $(0, 0, 1)$, depending on the residual capacity in the first period and the inventory capability of each plant. An illustration of this non-uniqueness is shown in the Figure 4.12. Thus, in contrast to the case discussed in Section 4.5, analyzing the distribution of residual demand in the second period independently and pairing it with the residual capacity in the first period does not help.

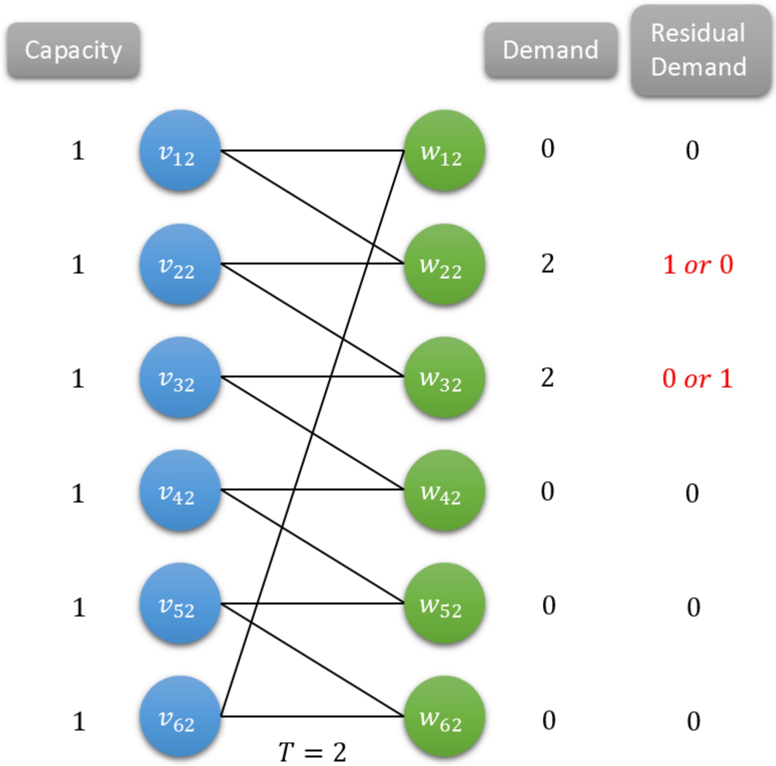

Figure 4.12: Residual demand given an instance of demand in the second period

To address this issue, we define *conditional residual performance* π

$$\pi^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix} \tag{4.52}$$

as the residual performance on $(\mathcal{C}, I_{1\to k}, N, 2)$ given that the residual capacity in the first period is $(b_1, \ldots, b_k)$, i.e. $\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2) \,|\, (\tilde{C}_1, \ldots, \tilde{C}_k) = (b_1, \ldots, b_k)]$. This conditional residual performance can be treated as the difference between two closed chain networks $\mathcal{C}'$ and $\mathcal{C}$:

$$\pi^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix} = \mathbf{E}[\Phi(\mathcal{C}', \emptyset, N, 2)] - \mathbf{E}[\Phi(\mathcal{C}, \emptyset, N, 2)]$$

where $\mathcal{C}$ is the regular single-period closed chain of size N, and $\mathcal{C}'$ is a single-period modified closed chain where each plant has capacity c_i defined as

$$c_i \quad = \quad \begin{cases} 1 + b_i & i = 1, \ldots, k \\ 1 & i = k+1, \ldots, N \end{cases} \tag{4.53}$$

In other words, the capacity of plants in the modified closed chain $\mathcal{C}'$ is

- $1 + b_i$ for plants 1 to k which have inventory capability
- 1 as before for plants $k+1$ to N which do not have inventory capability

An illustration of $\mathcal{C}'$ with $N = 6, k = 3$ is shown in Figure 4.13.

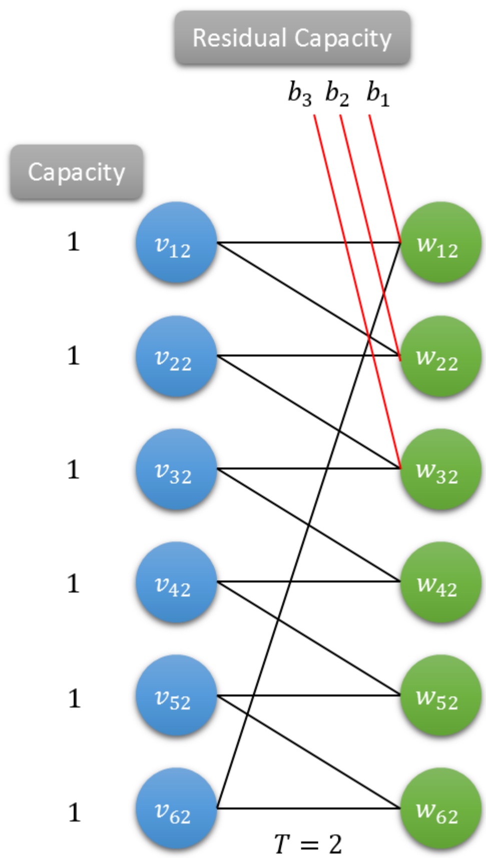

Figure 4.13: Illustration of $\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to 3}, 6, 2) \,|\, (\tilde{C}_1, \ldots, \tilde{C}_3) = (b_1, \ldots, b_3)]$

The intuition behind the definition of $\mathcal{C}'$ is simple. The benefit of inventory shows up only in the second period and it is equivalent to having additional capacity b_i at plants with inventory capability. If we could analyze the probability distribution b_i and the performance of the closed chain given additional capacity we could analyze the residual performance we originally interested in $\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$.

Given $\mathbf{P}$ and π, the unconditional residual performance of a network $\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$ can be expressed as the sum of 2^k terms:

$$\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)] \quad = \quad \sum_{b_1} \cdots \sum_{b_k} \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix} \pi^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix}$$

Despite the fact that there are an exponential number of terms in the expression, it turns out that $\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$ follows a recurrence relation, and so only a polynomial number of $\mathbf{P}$ and π need to be evaluated and a closed form expression can be obtained.

Now goes back to the closed form expression for the residual performance 4.6.1. Define *average residual performance* as the residual performance per inventory arc $\frac{1}{k}\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$. Given a fixed N, we can see the average residual performance decreases as the number of inventory arcs k increase, especially for the last two inventory arcs $k = N - 1$ and $k = N$. The intuition is that the benefit of adding an inventory arc for plant i comes from occasionally high demand for product i and $i + 1$ in the second period, so the benefit of two adjacent inventory arcs overlaps with each other. This phenomenon holds not only for this special case, but also can be generalized to other cases where $T > 2$ and demand is not Bernoulli, as shown in the following lemma:

Lemma 4.6.1. *Adding inventory arcs in closed chain with inventory* $(\mathcal{C}, \mathcal{I}_{1\to k}, N, T)$ *has diminishing marginal benefit, i.e.*

$$\mathbf{E}[(\mathcal{C}, \mathcal{I}_{1\to k+1}, N, T)] - \mathbf{E}[(\mathcal{C}, \mathcal{I}_{1\to k}, N, T)] \leq \mathbf{E}[(\mathcal{C}, \mathcal{I}_{1\to k}, N, T)] - \mathbf{E}[(\mathcal{C}, \mathcal{I}_{1\to k-1}, N, T)]$$

for $k = 2, \ldots, N - 1$

Proof: Using similar argument as in Simchi-Levi and Wei (2012)[55], it can be easily checked that any pair of inventory arcs are in parallel as defined in Gale and Politof (1981)[26], therefore two inventory arcs are substitutes thus have diminishing benefit.

□

Managerial Insight: This suggests that it is unnecessary to incorporate inventory capability at every plant if the current network is already a closed chain. When only a few inventory arcs are added, the guideline is discussed in Section 4.6.4.

On the other hand, given fixed number of inventory arcs k, the average residual performance is exactly the same for all N if $N \geq k + 2$. In other words, the benefit of adding inventory arcs is independent of the network size. However this fact is only true for Bernoulli demand.

4.6.2 The General Case

The approach we used to derive $\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$ for the special case relies on two major assumptions: *i.i.d.* Bernoulli demand and planning horizon $T = 2$. When the planning horizon $T > 2$ or demand is more general than Bernoulli, the expected maximum flow is much more difficult to calculate. The following two examples that illustrate the difficulty:

Example 1: $T = 2$ but demand is general
When demand is general, the probability that the residual capacity equals $\mathbf{b}$ defined in the previous subsection $\mathbf{P}(\tilde{\mathbf{C}} = \mathbf{b})$ also depends on the demand in the second period. To see this, assume $c = 1, N = 3$ and demand is discrete uniformly distributed between 0,1 and 2. Let demand in the first period be $(d_{11} = 0, d_{21} = 1, d_{31} = 0)$, then the residual capacity vector $\mathbf{b}$ can be either $(1, 0, 1)$ or $(0, 1, 1)$. Which of these is the actual residual capacity in the optimal solution depends on the demand realization in the second period. This precludes the possibility of extending the methodology in the previous section to this case.

Example 2: Demand is Bernoulli but $T > 2$
When planning horizon $T > 2$, consider an example with $T = 3$ shown in Figure 4.14. Demand is shown adjacent to product nodes w_{ij}. An optimal flow is shown with solid arcs (with flow 1) and dashed arcs (with flow 0). Observe that without inventory arcs, the flow on the closed chain in period 2 is not optimal for the closed chain. It is easy to verify that this remains true for other optimal flows, which means analyzing the residual capacity and demand without inventory arcs may provide little insight.

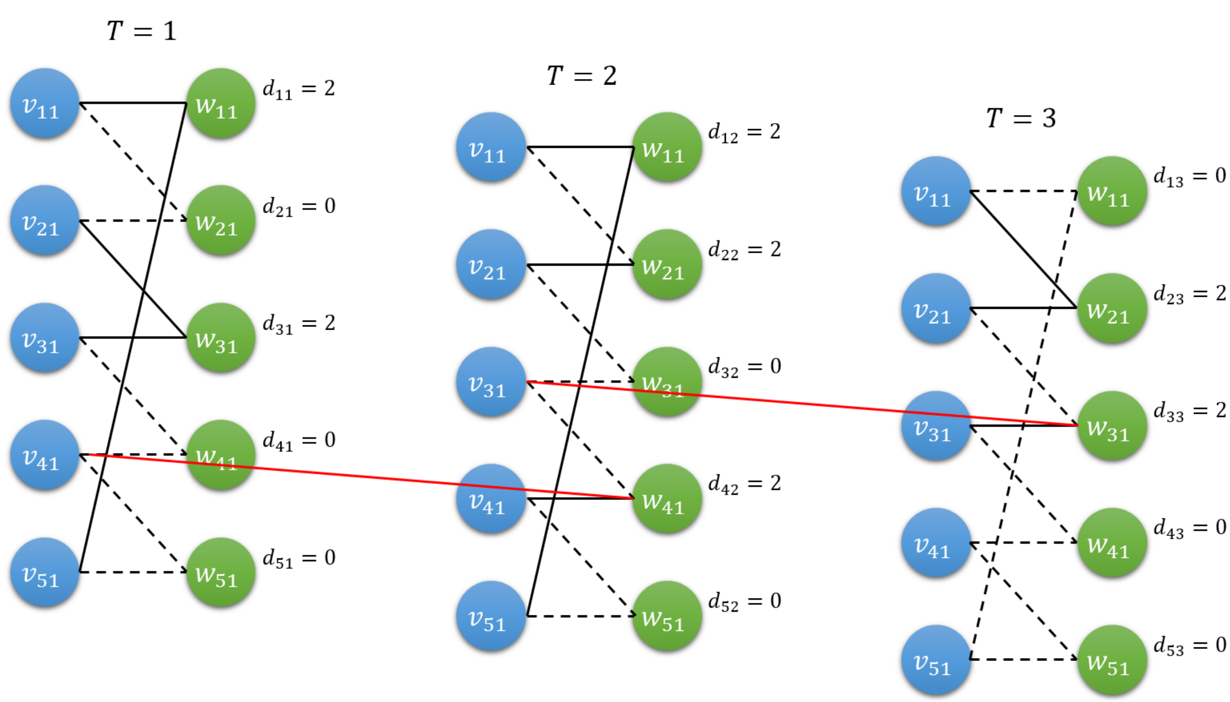

Figure 4.14: An example of $T > 2$

Our problem is closely related to the network reliability problem in the literature given a capacitated network subjected to random arc failure (equivalent to Bernoulli demand in our problem). There are two possible measures of the reliability of a network from source s to sink t: #1 expected maximum flow from s to t (the same as our performance measure) , or #2 the probability that s and t are connected. It is easy to see that measure #2 is a special case of the measure #1. Ball (1986)[6] showed that #2 on a general graph is NP-hard by showing the relationship between #2 and counting the number of minimum cardinality cutset, which is a known NP-hard problem (see Valiant (1979)[62]). Carey and Hendrickson (1984)[13] provide lower and upper bound on the measure #1 and Nagamochi and Ibaraki (1991)[44] shows the necessary and sufficient condition on when the lower bound is tight. Roughly speaking, the lower bound of #1 is the *maximum expected flow* which means no re-routing the flow is allowed and flows on path with failed arcs are lost. The upper bound of #1 is *maximum flow on expected network* which the capacity of all arcs are replaced with its mean. It can be easily checked both bounds are not tight in our problem (even in simpler open chain or closed chain without inventory). There are efficient algorithms to compute #1 and #2 on graphs with special structure such as bipartite or series-parallel graph (see Carey and Hendrickson (1984)[13], Satyanarayana and Wood (1985)[51]) but our problem lies in neither of the categories. In addition, measure #2 remains NP-hard on planar graph and directed acyclic graph (Provan and Ball (1983)[47]).

We conjecture that the computation of expected maximum flow $\mathbf{E}[\Phi]$ when either $T > 2$ or demand in general is NP-hard. When the computation of $\mathbf{E}[\Phi]$ can be done in polynomial time, the topology of the graph, capacity and demand process must preserve strong symmetry. To summarize, $\mathbf{E}[\Phi]$ can be computed in polynomial time when

- the network topology is a closed chain without inventory $(\mathcal{C}, \emptyset, N, T)$, the polynomial time computation of $\mathbf{E}[\Phi]$ requires (1) equal number of plants and products; (2) Identical capacity; (3) Exchangeable demand. (see Simchi-Levi and Wei (2012)[55]).

- the network topology is an open chain with inventory $(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)$, the polynomial time computation of $\mathbf{E}[\Phi]$ requires (1) equal number of plants and products; (2) Identical capacity; (3) *i.i.d.* demand. (see Section 4.5).

- the network topology is closed chain with inventory $(\mathcal{C}, \mathcal{I}_{1\to k}, N, T)$, the polynomial time computation of $\mathbf{E}[\Phi]$ requires (1) equal number of plants and products; (2) Identical capacity; (3) *i.i.d.* Bernoulli demand; (4) $T = 2$ (see Section 4.6).

Observe that assumptions become more restrictive as the symmetry of the network decreases. Whether or not $\mathbf{E}[\Phi]$ can be computed in polynomial time when any of these three conditions are violated remains an open question. However, since $\Phi(\cdot; d)$ is the maximum flow of a graph given arc capacity vector d, which has the Lipschitz property in term of d, i.e. a small change in d results in a small change in the maximum flow. This suggests that Monte Carlo simulation may perform well to estimate $\mathbf{E}[\Phi]$ because Φ is concentrated around its mean (by McDiarmid's inequality). Therefore, we turn to simulation-based analysis for the general case. See Alexopoulos and Fishman (1991)[1] for a discussion of using Monte Carlo simulation to compute network reliability.

4.6.3 Computational Analysis

In this section, we numerically analyze the performance of a closed chain with inventory configuration. For the special case mentioned in Section 4.6.1 we'll use the derived closed form solution. For general cases we estimate the performance via Monte Carlo simulation. To do this, we generate large number of instances which in each instance the demand d_{jt} is sampled from a demand distribution. Next we solve each instance as a maximum flow problem. Then, the performance of a network $\mathbf{E}[\Phi(\mathcal{A})]$ can be estimated using the average optimal objective value of all instances.

For the first set of tests, we fix the planning horizon $T = 2$ and study the effect of changing the number of inventory arcs for various N with three demand distributions: Bernoulli, Discrete Uniform and Binomial. The average residual performance $\frac{1}{k}\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$ is shown in the Table 4.7.

	Bernoulli				
k	$N=4$	$N=5$	$N=6$	$N=7$	$N=8$
1	0.4688	0.4688	0.4688	0.4688	0.4688
2	0.4492	0.4492	0.4492	0.4492	0.4492
3	0.4362	0.4395	0.4395	0.4395	0.4395
4	0.4102	0.4333	0.4340	0.4340	0.4340
5		0.4150	0.4304	0.4305	0.4305
6			0.4163	0.4282	0.4282
7				0.4166	0.4266
8					0.4166
	Uniform				
k	$N=4$	$N=5$	$N=6$	$N=7$	$N=8$
1	0.3515	0.6528	0.1433	0.2705	0.5155
2	0.3289	0.4381	0.4494	0.4137	0.2889
3	0.2574	0.2798	0.3681	0.4251	0.2802
4	0.2280	0.2636	0.2532	0.2834	0.2655
5		0.1859	0.2242	0.2845	0.2309
6			0.2827	0.2537	0.2697
7				0.1423	0.2321
8					0.2096
	Binomial				
k	$N=4$	$N=5$	$N=6$	$N=7$	$N=8$
1	0.5015	0.2777	0.4403	0.3536	0.4585
2	0.1841	0.1730	0.1528	0.2407	0.2414
3	0.0824	0.1226	0.1778	0.1461	0.1582
4	0.1066	0.1319	0.1444	0.1010	0.1331
5		0.1299	0.0786	0.0856	0.1055
6			0.0532	0.0795	0.0958
7				0.0742	0.0724
8					0.0562

Table 4.7: Average residual performance $\frac{1}{k}\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$

As we see in Table 4.7, given a fixed N the benefit of adding inventory capability decreases as the number of inventory arcs k increases, which confirms Lemma 4.6.1. In addition, the benefit further decreases as the demand variation decreases.

Managerial Insight: This suggests that if demand variation is small, a closed chain without inventory itself can already satisfy most of the demand, so unless the capital investment of inventory capability is inexpensive there is no need to incorporate inventory.

For the second set of tests we fix the number of plants product pairs $N = 4$ and vary the planning horizon from $T = 2$ to $T = 6$. The average residual performance $\frac{1}{k(T-1)}\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$ is shown in Table 4.8

	Bernoulli				
k	$T=2$	$T=3$	$T=4$	$T=5$	$T=6$
1	0.4688	0.6372	0.6753	0.5208	0.8076
2	0.4492	0.6146	0.6289	0.6622	0.7203
3	0.4362	0.4740	0.6539	0.6668	0.7495
4	0.4102	0.5586	0.5790	0.5979	0.6560
	Uniform				
k	$T=2$	$T=3$	$T=4$	$T=5$	$T=6$
1	0.3515	0.6532	0.7527	0.7197	0.7704
2	0.3289	0.4118	0.4576	0.4810	0.4917
3	0.2574	0.2876	0.3816	0.3938	0.4212
4	0.2280	0.2651	0.2632	0.2991	0.3233
	Binomial				
k	$T=2$	$T=3$	$T=4$	$T=5$	$T=6$
1	0.5015	0.4485	0.5053	0.4878	0.5834
2	0.1841	0.2223	0.2513	0.2779	0.2885
3	0.0824	0.1694	0.1698	0.1862	0.2102
4	0.1066	0.1126	0.1257	0.1381	0.1475

Table 4.8: Average residual performance $\frac{1}{k(T-1)}\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$

As in Section 4.5, from Table 4.8 the average residual performance on inventory arcs increases as the planning horizon T increases.

Managerial Insight: Inventory capability is most useful when there is a reliable, long horizon forecast.

4.6.4 Other Related Configurations

In previous subsections, for configurations in $\mathbb{B}$ we gradually add inventory arcs to a closed chain to analyze the marginal benefit of inventory capability. In other words, we study how changing the *quantity* of inventory arcs affects the performance. Similar to Section 4.5.6, we can also fix the number inventory arcs and study how changing the *position* of inventory arcs influences the performance. Consider another sequence of configurations $\mathbb{B}' = \{\mathcal{B}_1, \mathcal{B}_2, \ldots, \mathcal{B}_{N/2}\}$, each configuration has 2 inventory arcs. In the initial configuration $\mathcal{A}_1 = (\mathcal{C}, \mathcal{I}_{1,2})$, which plant 1 and 2 are able to hold inventory, and in subsequent configurations, the inventory arc at the second plant is moved to the next plant in each step, so $\mathcal{B}_k = (\mathcal{C}, \mathcal{I}_{1,k+1})$ until $k+1 \geq N/2$. An example of $\mathbb{B}'$ with $N=6$ is shown in Figure 4.15.

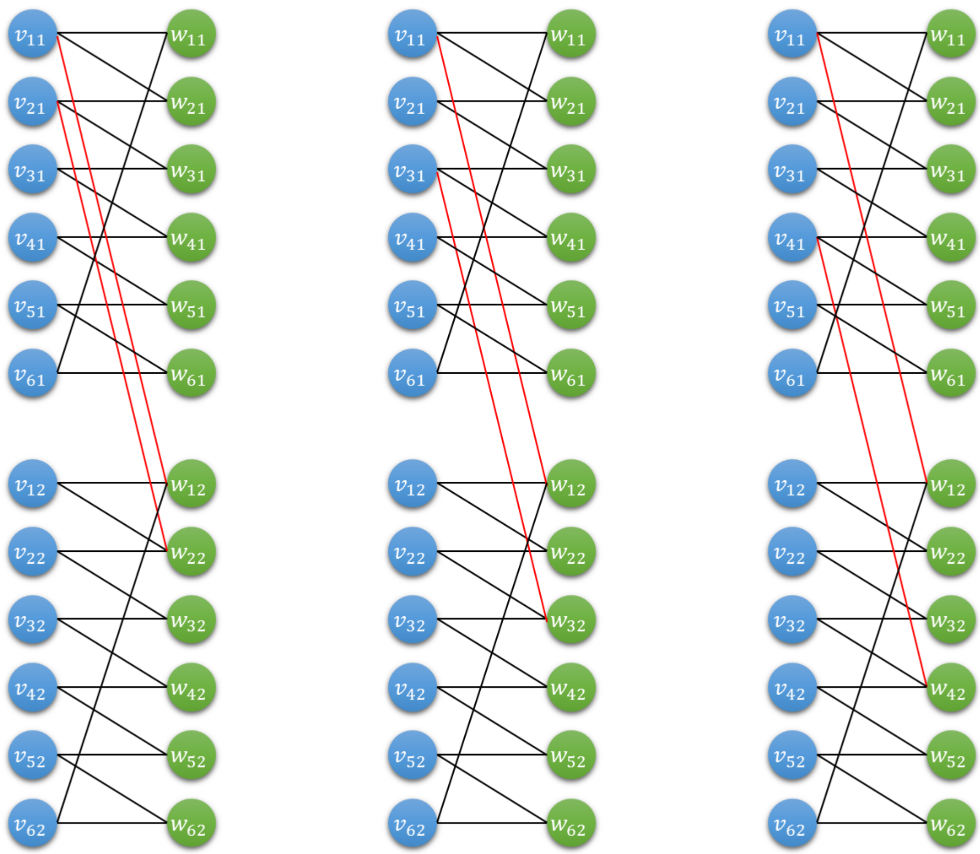

Figure 4.15: The sequence of configurations in $\mathbb{B}'$

Our approach that derives the closed form solution for the special case in Section 4.6.1 still applies to $\mathbb{B}'$ if demand is Bernoulli and planning horizon $T=2$. For general cases we again use Monte Carlo simulation to estimate the performance. It turns out the residual performance follows the relation

$$\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{\{1,2\}}, N, 2)] < \mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{\{1,3\}}, N, 2)] \leq \cdots \leq \mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{\{1,N/2\}}, N, 2)] \tag{4.54}$$

where the equality holds for the special case. In other words, the marginal benefit of inventory arcs increases if we spread them further apart. This is intuitive as we discussed at the end of Section 4.6.1: the benefit of an inventory arc may overlap with another inventory arc so by spreading out inventory arcs the overlap is minimized. This again confirms the conclusion in Lemma 4.6.1 that inventory capability on a closed chain network are substitutes.

Managerial Insight: If a firm currently has a flexible network without inventory and wants to add inventory capability for some plants while keeping flexibility, then these inventory-capable plants should be spread as far as possible to maximize the benefit.

4.7 Conclusion

Motivated by a biopharmaceutical firm whose managers wanted to incorporate process flexibility into their dedicated supply chain network, we have extended the flexibility literature to model inventory capability at individual plants in the network. Following the standard approach in the flexibility literature, we defined the performance of a network to be the expected maximum demand that can be satisfied and compare this performance among various types of networks.

We have developed the following managerial guidelines to help with effective network design in this setting:

- Inventory capability is most useful when there is a reliable demand forecast. In other situations, flexibility is more valuable.
- If a firm currently has a dedicated network with inventory, the length of its reliable forecast horizon dictates whether or not is worth sacrificing some inventory capability for flexibility in this network.
- If a firm currently has a flexible network with a chain configuration but without inventory, and wants to replace flexibility with inventory capability for some plants, then these converted plants should be adjacent to each other in the chain configuration so that the network can maintain a longer chain.
- If a firm currently has a flexible network with a closed chain configuration but without inventory, and wants to add inventory capability to some plants while maintaining the flexibility configuration, then these plants should be separated as much as possible in the chain.
- For a firm that has a flexible network, adding inventory capability to a few plants can capture most of the benefit of supply chain inventory, because adding inventory capability to additional plants in a closed chain configuration has diminishing marginal benefit.

Although we followed the standard approach in the flexibility literature and thus did not incorporate cost directly into our model, we believe the conclusions and insights in this work can help guide firms as the make strategic decisions related to flexibility and inventory capability, as our performance measure can be easily translated (at least approximately) to revenue gain.

Chapter 5

Closing Remarks

Motivated by the many supply chain management challenges faced by biopharmaceutical firms with which we worked, in this thesis we presented three inventory models that can help these firms make various crucial decisions.

In the first model, we studied inventory planning with batch differentiation. We characterized the performance of deterministic control and its extensions, approaches that are widely applied in industry already for managing batch production and differentiation. We showed that the regret of such control can be reduced to $O(T)$ via re-optimization in each period. We also explored the trade-off between (1) more accurate demand information and (2) more decision dynamics. We found that decision dynamics plays a more important role in reducing the regret. However the benefit of using full decision dynamics comes at the expense of larger computational burden. A natural way to mitigate this is to implement a roll-out algorithm with limited look ahead (see Bertesakes (2013)[8], Goodson et al. (2015)[30]). For example, we can solve a single deterministic control from period 2 to T to quickly estimate the cost of meeting future demand and then solve a single period dynamic program for the current period to capture the decision dynamics. We are not aware of previous results in the literature that focus on the theoretical performance of this type of algorithm, and this is one direction for future research.

In the second model, we addressed the issue of highly variable lead time faced firms that outsource their final production step to a third party contractor. We showed that a capacity reservation contract can not only reduce the inventory cost of the firm, but also increase the total profit of the entire system. We envision three possible tracks for future research on this topic if some assumptions in the analysis are relaxed: (1) Extend the study to a discrete time model. Under a discrete-time model, we need to compare the capacity reservation policy to an (s, S) policy with random lead time (see Kapalan (1970)[39], Sahin (1983)[50], Bashyam and Fu (1998)[7] for a discussion of the (s, S) policy with random lead time). Our preliminary simulation analysis leads to similar observations to those we have made for the continuous time model; (2) Allow the outsourcer to serve a mixture of type (i) and (ii) manufacturers. When both types of manufacturer are present, the outsourcer may manage a priority queue where the type (i) manufacturers are of low priority and join the queue based on a FIFO service discipline, while type (ii) manufacturers who have reserved capacity are of high priority and preempt any ongoing processing of type (i) manufacturer orders. The theoretical analysis of such a priority queue seems to be difficult. Our preliminary simulation analysis reveals that the mixture of type (i) and type (ii) may be worse than all type (i) manufacturers or all type (ii) manufacturers, depending on the traffic density of the system; (3) Allow the outsourcer to use other more sophisticated control schemes, such as a polling model (see Takagi (1986, 2000)[59, 60], Levy and Sidi (1990)[41], Boon et al. (2011)[9]) or a multi-class priority queue (see Maglaras (2006)[42]). The analysis of both type (i) and (ii) policies at the manufacturer is likely to be challenging in such settings.

In the third model, we explored a firm's strategic decisions around building new plants and adding process flexibility to existing plants. We defined the performance of a plant-product network as the expected maximum demand that can be satisfied, and proposed algorithms to evaluate the performance for some specific network structures. There are two ways that this model can be further extended to reveal more insight: (1) Incorporate cost into the model. This will change the underlying maximum flow problem to minimum cost network flow problem. The performance on a general network is expected to be difficult to obtain, but results for some particular networks may be a plausible direction. (2) Restrict the firm to make production plans before observing the demand. This will completely change the nature of the problem to a stochastic dynamic setting as in the batch differentiation problem in Chapter 2. The production policy is likely to be difficult to characterize (see Shaoxiang (2004)[53], Janakiraman et al. (2009)[33] for similar settings), and the nature of effective supply chain network design guidelines is an interesting open problem.

This thesis addresses several key challenges faced by biopharmaceutical firms, but it by no means comprehensively covers all the challenges in the industry. For example Yuen (2012)[65] focuses on the capacity planning problem, and Bu (2015)[10] studies production inventory control under uncertain yield – both of which are problems we outlined in Chapter 1. There are many remaining challenges in the industry that could benefit from the attention of academia.

Chapter 6

Bibliography

[1] C. Alexopoulos and G. S. Fishman. Characterizing stochastic flow networks using the monte carlo method. *Networks*, 21(7):775–798, 1991.

[2] S. Axsäter. Exact and approximate evaluation of batch-ordering policies for two-level inventory systems. *Operations Research*, 41(4):777–785, 1993.

[3] S. Axsäter. Approximate evaluation of batch-ordering policies for a one-warehouse, n non-identical retailer system under compound poisson demand. *Naval Research Logistics*, 42(5):807–819, 1995.

[4] S. Axsäter. Evaluation of installation stock based (r, q)-policies for two-level inventory systems with poisson demand. *Operations Research*, 46(3-supplement-3):S135–S145, 1998.

[5] S. Axsäter. Exact analysis of continuous review (r, q) policies in two-echelon inventory systems with compound poisson demand. *Operations Research*, 48(5):686–696, 2000.

[6] M. O. Ball. Computational complexity of network reliability analysis: An overview. *IEEE Transactions on Reliability*, 35(3):230–239, 1986.

[7] S. Bashyam and M. C. Fu. Optimization of (s, s) inventory systems with random lead times and a service level constraint. *Management Science*, 44(12-part-2):S243–S256, 1998.

[8] D. P. Bertsekas. Rollout algorithms for discrete optimization: A survey. In *Handbook of combinatorial optimization*, pages 2989–3013. Springer, 2013.

[9] M. A. Boon, R. Van der Mei, and E. M. Winands. Applications of polling systems. *Surveys in Operations Research and Management Science*, 16(2):67–82, 2011.

[10] D. Bu. Production planning and inventory control in pharmaceutical manufacturing process. *PhD Thesis, University of California Berkeley*, 2015.

[11] G. P. Cachon. Exact evaluation of batch-ordering inventory policies in two-echelon supply chains with periodic review. *Operations research*, 49(1):79–98, 2001.

[12] A. d. P. Calmon. Reverse logistics for consumer electronics: forecasting failures, managing inventory, and matching warranties. *PhD Thesis, Massachusetts Institute of Technology*, 2015.

[13] M. Carey and C. Hendrickson. Bounds on expected performance of networks with links subject to failure. *Networks*, 14(3):439–456, 1984.

[14] X. Chao and S. X. Zhou. Optimal policy for a multiechelon inventory system with batch ordering and fixed replenishment intervals. *Operations Research*, 57(2):377–390, 2009.

[15] F. Chen. Optimal policies for multi-echelon inventory problems with batch ordering. *Operations research*, 48(3):376–389, 2000.

[16] M. C. Chou, G. A. Chua, C.-P. Teo, and H. Zheng. Design for process flexibility: Efficiency of the long chain and sparse structure. *Operations research*, 58(1):43–58, 2010.

[17] M. C. Chou, G. A. Chua, C.-P. Teo, and H. Zheng. Process flexibility revisited: The graph expander and its applications. *Operations Research*, 59(5):1090–1105, 2011.

[18] M. C. Chou, C.-P. Teo, and H. Zheng. Process flexibility: design, evaluation, and applications. *Flexible Services and Manufacturing Journal*, 20(1-2):59–94, 2008.

[19] D. F. Ciocan and V. Farias. Model predictive control for dynamic resource allocation. *Mathematics of Operations Research*, 37(3):501–525, 2012.

[20] G. A. DeCroix and A. Arreola-Risa. Optimal production and inventory policy for multiple products under resource constraints. *Management Science*, 44(7):950–961, 1998.

[21] T. Deng and Z.-J. M. Shen. Process flexibility design in unbalanced networks. *Manufacturing & Service Operations Management*, 15(1):24–32, 2013.

[22] M. K. Doğru, M. I. Reiman, and Q. Wang. A stochastic programming based inventory policy for assemble-to-order systems with application to the w model. *Operations research*, 58(4-part-1):849–864, 2010.

[23] A. Federgruen and Y.-S. Zheng. An efficient algorithm for computing an optimal (r, q) policy in continuous review stochastic inventory systems. *Operations research*, 40(4):808–813, 1992.

[24] R. Forsberg. Exact evaluation of (r, q)-policies for two-level inventory systems with

poisson demand. *European Journal of Operational Research*, 96(1):130 – 138, 1997.

[25] D. Francas, M. Kremer, S. Minner, and M. Friese. Strategic process flexibility under lifecycle demand. *International Journal of Production Economics*, 121(2):427–440, 2009.

[26] D. Gale and T. Politof. Substitutes and complements in network flow problems. *Discrete Applied Mathematics*, 3(3):175–186, 1981.

[27] G. Gallego. New bounds and heuristics for (q, r) policies. *Management Science*, 44(2):219–233, 1998.

[28] G. Gallego and G. Van Ryzin. A multiproduct dynamic pricing problem and its applications to network yield management. *Operations research*, 45(1):24–41, 1997.

[29] H. Galliher, P. M. Morse, and M. Simond. Dynamics of two classes of continuous-review inventory systems. *Operations Research*, 7(3):362–384, 1959.

[30] J. C. Goodson, B. W. Thomas, and J. W. Ohlmann. A rollout algorithm framework for heuristic solutions to finite-horizon stochastic dynamic programs. *European Journal of Operational Research*, 258(1):216–229, 2017.

[31] S. C. Graves and B. T. Tomlin. Process flexibility in supply chains. *Management Science*, 49(7):907–919, 2003.

[32] G. Hadley and T. M. Whitin. Analysis of inventory systems. *Englewood Cliffs, NJ Prentice-Hall*, 1963.

[33] G. Janakiraman, M. Nagarajan, and S. Veeraraghavan. Simple policies for capacitated multi-product inventory systems. *submitted to Management Science*, 2008.

[34] S. Jasin. Reoptimization and self-adjusting price control for network revenue management. *Operations Research*, 62(5):1168–1178, 2014.

[35] S. Jasin and S. Kumar. A re-solving heuristic with bounded revenue loss for network revenue management with customer choice. *Mathematics of Operations Research*, 37(2):313–345, 2012.

[36] S. Jasin and S. Kumar. Analysis of deterministic lp-based booking limit and bid price controls for revenue management. *Operations Research*, 61(6):1312–1320, 2013.

[37] W. C. Jordan and S. C. Graves. Principles on the benefits of manufacturing process flexibility. *Management Science*, 41(4):577–594, 1995.

[38] P. Kaminsky and Y. Wang. Analytical models for biopharmaceutical operations and supply chain management: a survey of research literature. *Pharmaceutical Bioprocessing*, 3(1):61–73, 2015.

[39] R. S. Kaplan. A dynamic inventory model with stochastic lead times. *Management Science*, 16(7):491–507, 1970.

[40] A. Khintchine. Mathematical methods in the theory of queueing (translation dm andrews and mh quenouille, with additional notes by e. wolman). *Griffin, London*, 1969.

[41] H. Levy and M. Sidi. Polling systems: applications, modeling, and optimization. *IEEE Transactions on communications*, 38(10):1750–1760, 1990.

[42] C. Maglaras. Revenue management for a multiclass single-server queue via a fluid model analysis. *Operations Research*, 54(5):914–932, 2006.

[43] H.-Y. Mak and Z.-J. M. Shen. Stochastic programming approach to process flexibility design. *Flexible services and manufacturing journal*, 21(3-4):75–91, 2009.

[44] H. Nagamochi and T. Ibaraki. Maximum flows in probabilistic networks. *Networks*, 21(6):645–666, 1991.

[45] R. M. Oliver. Table of the waiting time distribution for the constant service queue (m/d/1). *International Journal of Computer Mathematics*, 2(1-4):35–56, 1968.

[46] E. L. Plambeck and A. R. Ward. Optimal control of a high-volume assemble-to-order system. *Mathematics of Operations Research*, 31(3):453–477, 2006.

[47] J. S. Provan and M. O. Ball. The complexity of counting cuts and of computing the probability that a graph is connected. *SIAM Journal on Computing*, 12(4):777–788, 1983.

[48] U. S. Rao. Properties of the periodic review (r, t) inventory control policy for stationary, stochastic demand. *Manufacturing & Service Operations Management*, 5(1):37–53, 2003.

[49] M. I. Reiman and Q. Wang. An asymptotically optimal policy for a quantity-based network revenue management problem. *Mathematics of Operations Research*, 33(2):257–282, 2008.

[50] I. Sahin. On the continuous-review (s, s) inventory model under compound renewal demand and random lead times. *Journal of Applied Probability*, 20(1):213–219, 1983.

[51] A. Satyanarayana and R. K. Wood. A linear-time algorithm for computing k-terminal reliability in series-parallel networks. *SIAM Journal on Computing*, 14(4):818–832, 1985.

[52] N. Secomandi. An analysis of the control-algorithm re-solving issue in inventory and revenue management. *Manufacturing & Service Operations Management*, 10(3):468–483, 2008.

[53] C. Shaoxiang. The optimality of hedging point policies for stochastic two-product flexible manufacturing systems. *Operations Research*, 52(2):312–322, 2004.

[54] C. Shi, Y. Wei, and Y. Zhong. Process flexibility for multi-period production systems. *Under review*, 2015.

[55] D. Simchi-Levi and Y. Wei. Understanding the performance of the long chain and sparse designs in process flexibility. *Operations research*, 60(5):1125–1141, 2012.

[56] J.-S. Song. The effect of leadtime uncertainty in a simple stochastic inventory model. *Management Science*, 40(5):603–613, 1994.

[57] J.-S. Song, H. Zhang, Y. Hou, and M. Wang. The effect of lead time and demand uncertainties in (r, q) inventory systems. *Operations Research*, 58(1):68–80, 2010.

[58] R. F. Stengel. *Optimal control and estimation*. Dover Publications; Reissue edition, 1994.

[59] H. Takagi. Analysis of polling systems. *MIT press*, 1986.

[60] H. Takagi. Analysis and application of polling models. *Performance Evaluation: Origins and Directions*, pages 423–442, 2000.

[61] J. T. Treharne and C. R. Sox. Adaptive inventory control for nonstationary demand and partial information. *Management Science*, 48(5):607–624, 2002.

[62] L. G. Valiant. The complexity of enumeration and reliability problems. *SIAM Journal on Computing*, 8(3):410–421, 1979.

[63] A. F. Veinott Jr. The optimal inventory policy for batch ordering. *Operations Research*, 13(3):424–432, 1965.

[64] A. F. Veinott Jr. Optimal policy for a multi-product, dynamic, nonstationary inventory problem. *Management Science*, 12(3):206–222, 1965.

[65] M. K. Yuen. Initial production capacity investments for commercializing pharmaceutical products. *PhD Thesis, University of California Berkeley*, 2012.

[66] P. Zipkin. Inventory service-level measures: convexity and approximation. *Management Science*, 32(8):975–981, 1986.

[67] P. Zipkin. Stochastic leadtimes in continuous-time inventory models. *Naval Research Logistics (NRL)*, 33(4):763–774, 1986.

Appendix A

Proofs and Analysis

A.1 Proofs of Results Appearing in Chapter 2

A.1.1 Proof of Theorem 2.4.1

For any feasible policy $\pi \in \Pi$, define a sequence $\{W^{\pi}_{t,i}\}$ where $W^{\pi}_{t+1,i} = W^{\pi}_{t,i} + N^{\pi}_{t,i} - \mu_i$ for $t \geq 1$ and $W^{\pi}_1 = I^{\pi}_1 = I_1$. Observe that $W^{\pi}_{t,i} = I_{1,i} + \sum_{s=1}^{t-1} N^{\pi}_{s,i} - \sum_{s=1}^{t-1} \mu_i$ for $t \geq 2$. Since $I^{\pi}_{t,i} = I_{1,i} + \sum_{s=1}^{t-1} N^{\pi}_{s,i} - \sum_{s=1}^{t-1} D_{s,i}$, we can write: $I^{\pi}_{t,i} = W^{\pi}_{t,i} - \sum_{s=1}^{t-1} \Delta_{s,i}$, where $\Delta_{s,i} = D_{s,i} - \mu_i$. We now proceed to prove Theorem 1 in three steps: in Step 1, we compute an upper bound for $C^D - C^*$; in Step 2, we compute an upper bound for $\mathbf{E}[C^{DC}] - C^D$; and, in Step 3, we put together the bounds from Steps 1 and 2 to get a bound for $\mathbf{E}[C^{DC}] - C^*$ (note that $\mathbf{E}[C^{DC}] - C^* = \mathbf{E}[C^{DC}] - C^D + C^D - C^*$).

Step 1

We first compute an upper bound for $C^D - C^*$. We claim that:

$$C^* \geq C^D - \mathbf{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{m}(p_i + h_i)\left(\sum_{s=1}^{t}\Delta_{s,i}\right)^+\right] \tag{A.1}$$

This is not difficult to show. For any policy $\pi \in \Pi$, we can bound:

$$
\begin{aligned}
&\sum_{t=1}^{T}\left[c\,Z_t^{\pi}\,B+\sum_{i=1}^{m}p_i(D_{t,i}-N_{t,i}^{\pi}-I_{t,i}^{\pi})^{+}+\sum_{i=1}^{m}h_i(N_{t,i}^{\pi}+I_{t,i}^{\pi}-D_{t,i})^{+}\right]\\
&\quad=\sum_{t=1}^{T}\left[c\,Z_t^{\pi}\,B+\sum_{i=1}^{m}(p_i+h_i)(I_{t,i}^{\pi}+N_{t,i}^{\pi}-D_{t,i})^{+}-\sum_{i=1}^{m}p_i(I_{t,i}^{\pi}+N_{t,i}^{\pi}-D_{t,i})\right]\\
&\quad=\sum_{t=1}^{T}\left[c\,Z_t^{\pi}\,B+\sum_{i=1}^{m}(p_i+h_i)(I_{t+1,i}^{\pi})^{+}-\sum_{i=1}^{m}p_iI_{t+1,i}^{\pi}\right]\\
&\quad\geq\sum_{t=1}^{T}\left[c\,Z_t^{\pi}\,B+\sum_{i=1}^{m}(p_i+h_i)(W_{t+1,i}^{\pi})^{+}-\sum_{i=1}^{m}p_iW_{t+1,i}^{\pi}\right]\\
&\qquad\quad-\sum_{t=1}^{T}\sum_{i=1}^{m}(p_i+h_i)\left(\sum_{s=1}^{t}\Delta_{s,i}\right)^{+}+\sum_{t=1}^{T}\sum_{i=1}^{m}\sum_{s=1}^{t}p_i\Delta_{s,i}\\
&\quad\geq C^{D}-\sum_{t=1}^{T}\sum_{i=1}^{m}(p_i+h_i)\left(\sum_{s=1}^{t}\Delta_{s,i}\right)^{+}+\sum_{t=1}^{T}\sum_{i=1}^{m}\sum_{s=1}^{t}p_i\Delta_{s,i}
\end{aligned}
\tag{A.2}
$$

where the first inequality holds because the identity $I_{t,i}^{\pi}=W_{t,i}^{\pi}-\sum_{s=1}^{t-1}\Delta_{s,i}$ implies $(I_{t,i}^{\pi})^{+}\geq(W_{t,i}^{\pi})^{+}-(\sum_{s=1}^{t-1}\Delta_{s,i})^{+}$ and the second inequality follows by the definition of C^{D}. Taking expectation on both sides of (A.2) and minimizing the sum in the left side of the inequality over $\pi\in\Pi$ yields (A.1).

Step 2

We now compute an upper bound for $\mathbf{E}[C^{DC}]-C^{D}$. We claim that:

$$
\mathbf{E}[C^{DC}]-C^{D}\leq\mathbf{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{m}(p_i+h_i)\left(-\sum_{s=1}^{t}\Delta_{s,i}\right)^{+}\right]
\tag{A.3}
$$

This can be shown using similar arguments as in Step 1. Define two sequences $\{I_{t,i}\}$ and $\{x_{t,i}\}$ as follows: $I_{t+1,i}=I_{t,i}+n_{t,i}^{D}-D_{t,i}$ and $x_{t+1,i}=x_{t,i}+n_{t,i}^{D}-\mu_i$ for $t\geq1$, where $x_1=I_1$. (Note that $I_{t,i}$ is simply the starting inventory for retailer i at the beginning of period t under DC control.)

Since $I_{t,i} = x_{t,i} - \sum_{s=1}^{t-1} \Delta_{s,i}$ for $t \geq 2$, we can bound:

$$
\begin{aligned}
\mathbf{E}[C^{DC}] &= \sum_{t=1}^{T} \mathbf{E}\left[c\, z_t^D\, B + \sum_{i=1}^{m} h_i (I_{t+1,i})^+ + \sum_{i=1}^{m} p_i (-I_{t+1,i})^+ \right] \\
&= \sum_{t=1}^{T} \mathbf{E}\left[c\, z_t^D\, B + \sum_{i=1}^{m} (p_i + h_i)(I_{t+1,i})^+ - \sum_{i=1}^{m} p_i I_{t+1,i} \right] \\
&\leq \sum_{t=1}^{T} \mathbf{E}\left[c\, z_t^D\, B + \sum_{i=1}^{m} (p_i + h_i)(x_{t+1,i})^+ - \sum_{i=1}^{m} p_i x_{t+1,i} \right] \\
&\qquad + \mathbf{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i) \left(-\sum_{s=1}^{t} \Delta_{s,i} \right)^+ \right] + \mathbf{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{m} \sum_{s=1}^{t} p_i \Delta_{s,i} \right] \\
&= C^D + \mathbf{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i) \left(-\sum_{s=1}^{t} \Delta_{s,i} \right)^+ \right]
\end{aligned}
$$

The inequality follows because $I_{t,i} = x_{t,i} - \sum_{s=1}^{t-1} \Delta_{s,i}$ implies $(I_{t,i})^+ \leq (x_{t,i})^+ + (-\sum_{s=1}^{t-1} \Delta_{s,i})^+$.

Step 3

Putting the bounds from Steps 1 and 2 together, we conclude that

$$
\begin{aligned}
\mathbf{E}[C^{DC}] - C^* &= \mathbf{E}[C^{DC}] - C^D + C^D - C^* \\
&\leq \mathbf{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i) \left(-\sum_{s=1}^{t} \Delta_{s,i} \right)^+ \right] + \mathbf{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i) \left(\sum_{s=1}^{t} \Delta_{s,i} \right)^+ \right] \\
&\leq 2 \sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i)\, \mathbf{E}\left[\left(\sum_{s=1}^{t} \Delta_{s,i} \right)^2 \right]^{1/2} \\
&= 2 \sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i)\, \sigma_i \sqrt{t} \\
&\leq 2\,(T+1)^{3/2} \left[\sum_{i=1}^{m} (p_i + h_i)\, \sigma_i \right]
\end{aligned}
$$

where the second equality follows from $\mathbf{E}[\Delta_{t,i}^2] = \sigma_i^2$ for all $t \geq 1$ (since demands on different periods are independent, $\mathbf{E}[\Delta_{t,i}\Delta_{s,i}] = 0$ for all $s \neq t$ and i) and the last inequality follows from integral comparison $\sum_{t=1}^{T} \sqrt{t} \leq \int_1^{T+1} \sqrt{x} \cdot dx \leq (T+1)^{3/2}$.

□

A.1.2 Proof of Corollary 2.4.2

The key is to note that each fulfilled demand incurs at least an ordering cost and each unfulfilled demand incurs at least a penalty cost; so, we can roughly bound C^* from below as follows: $C^* \geq \mathbf{E}[\sum_{t=1}^{T}\sum_{i=1}^{m}\min\{c,p_i\}\ D_{t,i}] = T\sum_{i=1}^{m}\ \min\{c,p_i\}\mu_i$. Putting this together with the bound in Theorem 1 and the fact that $\sigma_i = \sqrt{\mu^*}$ for all i (because demand is Poisson), we get:

$$\frac{\mathbf{E}[C^{DC}] - C^*}{C^*} \;\leq\; \frac{2(T+1)^{3/2}\left[\sum_{i=1}^{m}(p_i+h_i)\right]\sqrt{\mu^*}}{T\left[\sum_{i=1}^{m}\min\{c,p_i\}\right]\mu^*} \;=\; O\left(\left(\frac{T}{\mu^*}\right)^{1/2}\right)$$

□

A.1.3 Proof of Theorem 2.5.1

The proof proceeds in two steps. The first step shows the general structure of the allocation policy: raise the inventory for a subset of retailers to the same level and allocate nothing to the rest of them. The second step shows the retailers we raise to the same level are the retailers with lowest inventory.

Step 1

Since the ordering cost $cz_t^D B$ is just a constant, we'll ignore the ordering cost in $J_t(I_t)$. Let $g(y) = \mathbf{E}_D\left[(h^*(y-D)^+ + p^*(y-D)^-)\right]$ be the single retailer inventory cost function so $G(I_t+N_t) = \sum_{i=1}^{m} g(I_{t,i}+N_{t,i})$. Notice the objective function in Bellman's equation (15) can be separated by retailers as $G(I_t+N_t)+\mathbf{E}[J_{t+1}(I_t+N_t-D_t)] = \sum_{i=1}^{m}\gamma_t(I_{t,i}+N_{t,i})$. This can be seen using simple induction since $J_{T+1}=0$ and $J_T = \sum_{i=1}^{m} g(I_{T,i}+N_{T,i})$ so $\gamma_T(I_{T,i}+N_{T,i}) = g(I_{T,i}+N_{T,i})$. Since

$$\begin{aligned} J_t(I_t) &= \min_{N_{t,i}\in\Omega_t}\sum_{i=1}^{m} g(I_{t,i}+N_{t,i}) + \mathbf{E}_D\left[\min_{N_{t+1,i}\in\Omega_{t+1}}\sum_{i=1}^{m}\gamma_{t+1}(I_{t,i}+N_{t,i}-D_{t,i}+N_{t+1,i})\right] \\ &= \min_{N_{t,i}\in\Omega_t}\sum_{i=1}^{m} g(I_{t,i}+N_{t,i}) + \sum_{i=1}^{m}\mathbf{E}_D\left[\gamma_{t+1}(I_{t,i}+N_{t,i}-D_{t,i}+N_{t+1,i}^{\pi*})\right] \end{aligned}$$

where $N_{t+1,i}^{\pi*}$ is the optimal allocation in period $t+1$ which is a function of D_t, then $\gamma_t(I_{t,i}+N_{t,i}) = g(I_{t,i}+N_{t,i}) + \mathbf{E}_D\left[\gamma_{t+1}(I_{t,i}+N_{t,i}-D_{t,i}+N_{t+1,i}^{\pi*})\right]$. Notice γ_t is independent of index i since all retailers are homogeneous and demand is *i.i.d.* by assumption. In a word, $\sum_{i=1}^{m}\gamma_t(I_{t,i}+N_{t,i})$ is the cost if the allocation decision in period t is $\{N_{t,i}\}$ and the allocation decisions in period $t+1$ to T are the optimal allocation given $\{N_{t,i}\}$. Let $Y_{t,i} = I_{t,i}+N_{t,i}$, given I_t and z_t^D, to find the optimal allocation in period t we need to solve

$$\begin{aligned}
\min_{Y_{t,i}} \quad & \sum_{i=1}^{m} \gamma_t(Y_{t,i}) \\
s.t. \quad & \sum_{i=1}^{m} (Y_{t,i} - I_{t,i}) = z_t^D B \\
& Y_{t,i} - I_{t,i} \geq 0 \quad \forall i
\end{aligned}$$

Since $J_{T+1} = 0$ and γ_t is just sum of strictly convex functions g and the expectation preserves convexity, γ_t is strictly convex in $Y_{t,i}$, and in Ω the equality constraint is affine and inequality constraints are convex as well, therefore the first-order conditions are sufficient for optimality: there exists λ, μ_i satisfies

$$\gamma_t'(Y_{t,i}) + \lambda - \mu_i = 0 \ \forall i \tag{A.4}$$
$$(Y_{t,i} - I_{t,i})\mu_i = 0 \ \forall i \tag{A.5}$$
$$\mu_i \geq 0 \ \forall i \tag{A.6}$$
$$\sum_{i=1}^{m} (Y_{t,i} - I_{t,i}) = z_t^D B \tag{A.7}$$
$$Y_{t,i} - I_{t,i} \geq 0 \ \forall i \tag{A.8}$$

To characterize the structure of the optimal solution, we divide the retailers be into two subsets: those with $Y_{t,i} > I_{t,i}$ and those with $Y_{t,i} = I_{t,i}$. Let $\mathcal{A} = \{i \,|\, Y_{t,i} > I_{t,i}\}$ and $k = |\mathcal{A}|$.

Case 1: For retailers $i \in \mathcal{A}^c$ with $Y_{t,i} = I_{t,i}$
Allocate nothing for those retailers.

Case 2: For retailers $i \in \mathcal{A}$ with $Y_{t,i} > I_{t,i}$
By (A.5) immediately we have $\mu_i = 0$, then (A.4) - (A.8) for $i \in \mathcal{A}$ reduce to:

$$\gamma_t'(Y_{t,i}) + \lambda = 0 \quad \forall i \in \mathcal{A} \tag{A.9}$$
$$\sum_{i \in \mathcal{A}} (Y_{t,i} - I_{t,i}) = z_t^D B \tag{A.10}$$
$$Y_{t,i} - I_{t,i} \geq 0 \quad \forall i \in \mathcal{A} \tag{A.11}$$

By (A.9) we have $\gamma_t'(Y_{t,i}) = \gamma_t'(Y_{t,j})$ for any $i, j \in \mathcal{A}$. Since γ_t is strictly convex, γ_t' is strictly increasing, which implies $Y_{t,i} = Y_{t,j}$, i.e. raising inventory of all retailers in $\mathcal{A}$ to the same level.

Step 2

Next, we'll show that the set of retailers in $\mathcal{A}$ is indeed the θ retailers with lowest starting inventory level, i.e. $\mathcal{A} = \{i \,|\, i \leq \theta\}$ and $k = \theta$. From now on, we will drop the subscript t if it is not ambiguous. Assume retailers have pre-allocation inventory level $\{I_1, I_2, \ldots, I_m\}$ and remember these levels are sorted from smallest to largest. Let Y^* denote the order-up-to level in Theorem 2: $Y^* = \frac{1}{\theta}\left(\sum_{i=1}^{\theta} I_i + zB\right) \ \ \forall i \leq \theta$ and let the cost of the policy that raise $\{I_1, I_2, \ldots, I_\theta\}$ to the same level Y^* to be P^*. Now consider two cases:

Case 1: Raise $\theta' \neq \theta$ retailers with the lowest starting inventory to the same level. If $\theta' > \theta$ then (A.7) is infeasible. Now, without loss of generality, assume $\theta - 1$ retailers with inventory $\{I_1, I_2, \ldots, I_{\theta-1}\}$ are raised to the same level $\bar{Y}$. Let the cost of this policy to be P'.

$$\begin{aligned} P' &= (\theta-1)\gamma(\bar{Y}) + \gamma(I_\theta) + \sum_{i>\theta} \gamma(I_i) \\ &\geq \theta\,\gamma(Y^*) + \sum_{i>\theta} \gamma(I_i) \\ &= P^* \end{aligned}$$

where the inequality holds because of the convexity of $\gamma(\cdot)$. So P^* dominates P'.

Case 2: Raise inventory for θ arbitrarily selected retailers to the same level, Without loss of generality, suppose we skip retailer j and raise the retailers with inventory $\{I_1, I_2, \ldots, I_{j-1}, I_{j+1}, \ldots, I_\theta, I_{\theta+1}\}$ to the same level $\bar{Y}$. It is easy to check we have $\bar{Y} \geq I_{\theta+1} \geq Y^* \geq I_j$. Let the cost of this policy to be P'. By the convexity $\gamma(\cdot)$, we have

$$\begin{aligned} P' &= \theta\gamma(\bar{Y}) + \gamma(I_j) + \sum_{i>\theta+1} \gamma(I_i) \\ &\geq \theta\gamma(Y^*) + \gamma(I_{\theta+1}) + \sum_{i>\theta+1} \gamma(I_i) \\ &= P^* \end{aligned}$$

So, P^* dominates P'. Thus, raising θ retailers with the lowest inventory level to the same level dominates any other possible set of retailers, so we have $\mathcal{A} = \{i \,|\, i \leq \theta\}$. By (A.8), this level is

$$N_{t,i} + I_{t,i} = \frac{1}{k}\left(\sum_{i\in\mathcal{A}} I_{t,i} + z_t^D B\right) \quad \forall i \in \mathcal{A}$$

Since there is no other allocation outperforms the one in Theorem 2, the expected total costs under the allocation policy in Theorem 2 must be J^*.

□

A.1.4 Proof of Theorem 2.5.2

We proceed in several steps. In Step 1, we characterize an optimal solution of $J^H(D_{1:T})$; in Step 2, we use the result in Step 1 to prove $\mathbf{E}[(J^{DC} - J^H)\,\mathbf{1}\{\mathcal{A}\}] \leq 2\sigma(T+1)^{3/2}\,[\sum_{i=1}^m (p_i + h_i)]$; in Step 3, we characterize an optimal solution of $J_t^D(I_t)$; finally, in Step 4, we use the results from previous Steps to prove $\mathbf{E}[(J^{CEC} - J^H)\,\mathbf{1}\{\mathcal{A}\}] \leq 2\sigma T\,[\sum_{i=1}^m (p_i + h_i)]$.

Step 1

Define a sequence $\{n_{t,i}^H\}$ as follows: $n_{t,i}^H = n_{t,i}^{*1} + \Delta_{t,i} - \frac{1}{m}\sum_{j=1}^m \Delta_{t,j}$. We claim that if $D_{1:T} \in \mathcal{A}$, then n^H is an optimal allocation for the hindsight problem. First note that, since $x_{t+1,i} = I_{1,i} + \sum_{s=1}^{t} n_{s,i} - t\mu_i$ (because $x_{s+1,i} = x_{s,i} + n_{s,i} - \mu_i$), optimization $J_1^D(I_1)$ can be written as follows:

$$
\begin{aligned}
J_1^D(I_1) \quad &= \quad \min_{n,y,\theta} \sum_{t=1}^{T} \left[c\, z_t^D\, B + \sum_{i=1}^{m} p_i\, y_{t,i} + \sum_{i=1}^{m} h_i \theta_{t,i} \right] && \text{(A.12)} \\
&\text{s.t.} \quad y_{t,i} \geq t\mu_i - \sum_{s=1}^{t} n_{s,i} - I_{1,i} \qquad \forall t,\, i && \text{(A.13)} \\
&\qquad y_{t,i} \geq 0 \qquad \forall t,\, i && \text{(A.14)} \\
&\qquad \theta_{t,i} \geq I_{1,i} + \sum_{s=1}^{t} n_{s,i} - t\mu_i \qquad \forall t,\, i && \text{(A.15)} \\
&\qquad \theta_{t,i} \geq 0 \qquad \forall t,\, i && \text{(A.16)} \\
&\qquad \sum_{i=1}^{m} n_{t,i} = z_t^D B \qquad \forall t && \text{(A.17)} \\
&\qquad n_{t,i} \geq 0 \qquad \forall t,\, i && \text{(A.18)}
\end{aligned}
$$

By Karush-Kuhn-Tucker (KKT) conditions, there exist dual variables $\lambda_{t,i}^1$, $\lambda_{t,i}^2$, $\lambda_{t,i}^3$, $\lambda_{t,i}^4$, ξ_t, and $\Omega_{t,i}$ corresponding to constraints (A.13)-(A.18) respectively such that the following hold:

$$
\begin{aligned}
p_i &= \lambda_{t,i}^1 + \lambda_{t,i}^2 \qquad \forall t,\, i && \text{(A.19)} \\
h_i &= \lambda_{t,i}^3 + \lambda_{t,i}^4 \qquad \forall t,\, i && \text{(A.20)} \\
0 &= -\sum_{s=t}^{T} \lambda_{t,i}^1 + \sum_{s=t}^{T} \lambda_{t,i}^3 + \xi_t - \Omega_{t,i} \qquad \forall t,\, i && \text{(A.21)} \\
0 &= \lambda_{t,i}^1 \left[y_{t,i} - t\mu_i + \sum_{s=1}^{t} n_{s,i} + I_{1,i} \right] \qquad \forall t,\, i && \text{(A.22)} \\
0 &= \lambda_{t,i}^2\, y_{t,i} \qquad \forall t,\, i && \text{(A.23)} \\
0 &= \lambda_{t,i}^3 \left[\theta_{t,i} - I_{1,i} - \sum_{s=1}^{t} n_{s,i} + t\mu_i \right] \qquad \forall t,\, i && \text{(A.24)} \\
0 &= \lambda_{t,i}^4\, \theta_{t,i} \qquad \forall t,\, i && \text{(A.25)} \\
0 &= \Omega_{t,i}\, n_{t,i} \qquad \forall t,\, i && \text{(A.26)} \\
\lambda_{t,i}^1 &\geq 0,\ \ \lambda_{t,i}^2 \geq 0,\ \ \lambda_{t,i}^3 \geq 0,\ \ \lambda_{t,i}^4 \geq 0 \qquad \forall t,\, i. && \text{(A.27)}
\end{aligned}
$$

The last five equalities represent complementary slackness conditions. (Since we assume that $n_{t,i}^{*1} > \varphi > 0$, by $\Omega_{t,i}\, n_{t,i} = 0$, we immediately have $\Omega_{t,i} = 0$ for all t and i.) To show that n^H is optimal for the hindsight problem on $\mathcal{A}$, it is sufficient that we show: (C1) $n_{t,i}^H \geq 0$ for all t and i, (C2) $\sum_{i=1}^{m} n_{t,i}^H = z_t^D B$ for all t, and (C3) $I_{1,i} + \sum_{s=1}^{t} n_{s,i}^H - \sum_{s=1}^{t} D_{s,i}$ has the same sign (i.e., either strictly positive or strictly negative) as $I_{1,i} + \sum_{s=1}^{t} n_{s,i}^{*1} - t\mu_i$ for all t and i. If these conditions are satisfied, then we can use the dual variables for $J_1^D(I_1)$ as dual variables for the hindsight problem. Indeed, it is not difficult to check that the combination of n^H and these dual variables satisfy the KKT conditions for the hindsight problem. (The KKT conditions for hindsight problem are exactly the same as the KKT conditions above, with two exceptions: (A.22) becomes $0 = \lambda^1_{t,i}\,[y_{t,i} - \sum_{s=1}^t D_{s,i} + \sum_{s=1}^t n_{s,i} + I_{1,i}]$ and (A.24) becomes $0 = \lambda^3_{t,i}\,[\theta_{t,i} - I_{1,i} - \sum_{s=1}^t n_{s,i} + \sum_{s=1}^t D_{s,i}]$. If condition (C3) is satisfied, then we can use the same $\lambda^1_{t,i}$, $\lambda^3_{t,i}$ and $\theta_{t,i}$ as in above, and both (A.21) and (A.24) are satisfied for the hindsight problem at $n_{t,i} = n^H_{t,i}$.) Since KKT conditions are both necessary and sufficient for optimality in linear program, we can then conclude that n^H is optimal for $J^H(D_{1:T})$ on $\mathcal{A}$. But, conditions (C1)-(C3) immediately follow from the definition of n^H and $\mathcal{A}$, together with the assumptions that $I_{1,i} = 0$ and $|\sum_{s=1}^t n^{*1}_{s,i} - t\mu_i| \geq t\varphi$ for all $t \geq 1$ and i. This completes the proof.

Step 2

We now prove that $\mathbf{E}[(J^{DC} - J^H)\,\mathbf{1}\{\mathcal{A}\}] \leq 2\sigma(T+1)^{3/2}\,[\sum_{i=1}^m (p_i + h_i)]$. This is not difficult to show. Using n^H (defined in Step 1) as our optimal solution for the hindsight problem, the starting inventory level for retailer i at the beginning of period t is given by $I^H_{t+1,i} = \sum_{s=1}^t n^H_{t,i} - \sum_{s=1}^t D_{s,i} = \sum_{s=1}^t n^{*1}_{s,i} + \sum_{s=1}^t \Delta_{s,i} - \frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m \Delta_{s,j} - \sum_{s=1}^t D_{s,i}$. Moreover, the starting inventory level under DC is $I^{DC}_{t+1,i} = \sum_{s=1}^t n^{*1}_{s,i} - \sum_{s=1}^t D_{s,i}$. So, on $\mathcal{A}$, we can bound $J^{DC} - J^H$ as follows:

$$
\begin{aligned}
& J^{DC} - J^H \\
= \;& \sum_{t=1}^T \sum_{i=1}^m \left[p_i(-I^{DC}_{t+1,i})^+ + h_i(I^{DC}_{t+1,i})^+\right] \;-\; \sum_{t=1}^T \sum_{i=1}^m \left[p_i(-I^{H}_{t+1,i})^+ + h_i(I^{H}_{t+1,i})^+\right] \\
= \;& \sum_{t=1}^T \sum_{i=1}^m p_i \left[(-I^{DC}_{t+1,i})^+ - (-I^{H}_{t+1,i})^+\right] \;+\; \sum_{t=1}^T \sum_{i=1}^m h_i \left[(I^{DC}_{t+1,i})^+ - (I^{H}_{t+1,i})^+\right] \\
= \;& \sum_{t=1}^T \sum_{i=1}^m p_i \left[\left(\sum_{s=1}^t D_{s,i} - \sum_{s=1}^t n^{*1}_{s,i}\right)^+ - \left(\sum_{s=1}^t D_{s,i} - \sum_{s=1}^t n^{*1}_{s,i} - \sum_{s=1}^t \Delta_{s,i} + \frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m \Delta_{s,j}\right)^+\right] \\
& + \sum_{t=1}^T \sum_{i=1}^m h_i \left[\left(\sum_{s=1}^t n^{*1}_{s,i} - \sum_{s=1}^t D_{s,i}\right)^+ - \left(\sum_{s=1}^t n^{*1}_{s,i} + \sum_{s=1}^t \Delta_{s,i} - \frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m \Delta_{s,j} - \sum_{s=1}^t D_{s,i}\right)^+\right] \\
\leq \;& \sum_{t=1}^T \sum_{i=1}^m (p_i + h_i) \left|\sum_{s=1}^t \Delta_{s,i} - \frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m \Delta_{s,j}\right|
\end{aligned}
$$

where the inequality holds since $a^+ - (a-b)^+ \leq |b|$ for all a and b. The result immediately follows because $\mathbf{E}[|\Delta_{t,i}|\,\mathbf{1}\{\mathcal{A}\}] \leq \mathbf{E}[|\Delta_{t,i}|] \leq \sigma_i \leq \sigma$ for all t and i, and

$$
\mathbf{E}\left[\left|\sum_{s=1}^t \Delta_{s,i} - \frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m \Delta_{s,j}\right|\right] \;\leq\; \mathbf{E}\left[\left|\sum_{s=1}^t \Delta_{s,i}\right|\right] + \mathbf{E}\left[\left|\frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m \Delta_{s,j}\right|\right] \;\leq\; 2\sigma\sqrt{t}
$$

Step 3

We now characterize an optimal solution of $J^D_t(I_t)$. Since it may not be unique, we will only focus on a particular sequence of optimal solution $\{n^{*t}_t\}$. Define n^{*t} for $t > 1$ as follows: $n^{*t}_{t,i} = n^{*1}_{t,i} + \Delta_{t-1,i} - \frac{1}{m}\sum_{j=1}^m \Delta_{t-1,j}$ and $n^{*t}_{s,i} = n^{*1}_{s,i}$ for $s > t$. Suppose that $D_{1:T} \in \mathcal{A}$. In this step, we will prove the following statement: If we use $n_s = n_s^{*s}$ for all $s \leq t-1$, then

(i) The starting inventory level for retailer i at the beginning of period $t > 1$ is given by

$$I_{t,i} = \sum_{s=1}^{t-1} n_{s,i}^{*1} - (t-1)\mu_i - \Delta_{t-1,i} - \frac{1}{m}\sum_{s=1}^{t-2}\sum_{j=1}^{m}\Delta_{s,j}$$

(ii) n^{*t} is an optimal solution for $J_t^D(I_t)$.

Both (i) and (ii) can be proved by induction. We start with $t = 2$ as our base case. (The case $t = 1$ is trivially true.) The starting inventory level for retailer i at the beginning of period 2 is $I_{2,i} = I_{1,i} + n_{1,i}^{*1} - D_{1,i} = n_{1,i}^{*1} - \mu_i - \Delta_{1,i}$. So, (i) holds. We now check (ii). At the beginning of period 2, we have to solve the following linear program:

$$J_2^D(I_2) = \min_{n,y,\theta} \sum_{t=2}^{T}\left[c\, z_t^D\, B + \sum_{i=1}^{m} p_i\, y_{t,i} + \sum_{i=1}^{m} h_i \theta_{t,i}\right] \tag{A.28}$$

$$\text{s.t.} \quad y_{t,i} \geq (t-1)\mu_i - \sum_{s=2}^{t} n_{s,i} - I_{2,i} \qquad \forall t \geq 2,\ i \tag{A.29}$$

$$y_{t,i} \geq 0 \qquad \forall t \geq 2,\ i \tag{A.30}$$

$$\theta_{t,i} \geq I_{2,i} + \sum_{s=2}^{t} n_{s,i} - (t-1)\mu_i \qquad \forall t \geq 2,\ i \tag{A.31}$$

$$\theta_{t,i} \geq 0 \qquad \forall t \geq 2,\ i \tag{A.32}$$

$$\sum_{i=1}^{m} n_{t,i} = z_t^D B \qquad \forall t \geq 2 \tag{A.33}$$

$$n_{t,i} \geq 0 \qquad \forall t \geq 2,\ i \tag{A.34}$$

By Karush-Kuhn-Tucker (KKT) conditions, there exist dual variables $\hat{\lambda}_{t,i}^1$, $\hat{\lambda}_{t,i}^2$, $\hat{\lambda}_{t,i}^3$, $\hat{\lambda}_{t,i}^4$, $\hat{\xi}_t$, and

$\hat{\Omega}_{t,i}$ corresponding to constraints (A.29)-(A.34) respectively such that the following hold:

$$p_i = \hat{\lambda}^1_{t,i} + \hat{\lambda}^2_{t,i} \qquad \forall t \geq 2,\, i \tag{A.35}$$

$$h_i = \hat{\lambda}^3_{t,i} + \hat{\lambda}^4_{t,i} \qquad \forall t \geq 2,\, i \tag{A.36}$$

$$0 = -\sum_{s=t}^{T} \hat{\lambda}^1_{t,i} + \sum_{s=t}^{T} \hat{\lambda}^3_{t,i} + \hat{\xi}_t - \hat{\Omega}_{t,i} \qquad \forall t \geq 2,\, i \tag{A.37}$$

$$0 = \hat{\lambda}^1_{t,i} \left[y_{t,i} - t\mu_i + \sum_{s=1}^{t} n_{s,i} + I_{1,i} \right] \qquad \forall t \geq 2,\, i \tag{A.38}$$

$$0 = \hat{\lambda}^2_{t,i}\, y_{t,i} \qquad \forall t \geq 2,\, i \tag{A.39}$$

$$0 = \hat{\lambda}^3_{t,i} \left[\hat{\theta}_{t,i} - I_{1,i} - \sum_{s=1}^{t} n_{s,i} + t\mu_i \right] \qquad \forall t \geq 2,\, i \tag{A.40}$$

$$0 = \hat{\lambda}^4_{t,i}\, \hat{\theta}_{t,i} \qquad \forall t \geq 2,\, i \tag{A.41}$$

$$0 = \hat{\Omega}_{t,i}\, n_{t,i} \qquad \forall t \geq 2,\, i \tag{A.42}$$

$$\hat{\lambda}^1_{t,i} \geq 0,\ \ \hat{\lambda}^2_{t,i} \geq 0,\ \ \hat{\lambda}^3_{t,i} \geq 0,\ \ \hat{\lambda}^4_{t,i} \geq 0 \qquad \forall t \geq 2,\, i. \tag{A.43}$$

The last five equalities represent complementary slackness conditions. To show that n^{*2} is an optimal solution for $J^D_2(I_2)$, we use similar arguments as in Step 1. We claim that, on $\mathcal{A}$, all the three conditions listed in the last paragraph in Step 1 still hold: (C1) $n^{*2}_{t,i} \geq 0$ for all $t \geq 2$ and i, (C2) $\sum_{i=1}^{m} n^{*2}_{t,i} = z^D_t B$ for all $t \geq 2$, and (C3) $\sum_{s=2}^{t} n^{*2}_{s,i} - (t-1)\mu_i$ has the same sign (i.e., either strictly positive or strictly negative) as $\sum_{s=2}^{t} n^{*1}_{s,i} - (t-1)\mu_i$ for all $t \geq 2$ and i. The first two conditions are straight-forward to check; the last condition holds because $|\sum_{s=2}^{t} n^{*2}_{s,i} - \sum_{s=2}^{t} n^{*1}_{s,i}| = |n^{*2}_{2,i} - n^{*1}_{2,i}| = |\Delta_{1,i} - \frac{1}{m}\sum_{j=1}^{m} \Delta_{1,j}| \leq 2\sum_{j=1}^{m} |\Delta_{1,i}| < \varphi$ and $|\sum_{s=2}^{t} n^{*1}_{s,i} - (t-1)\mu_i| \geq (t-1)\varphi$ on $\mathcal{A}$. As in Step 1, these three conditions allow us to use the same dual variables for the corresponding constraints in both $J^D_2(I_1 + n^{*1}_1 - \mu)$ and $J^D_2(I_2)$. To be precise, first, note that the constraints indexed with $t \geq 2$ in $J^D_1(I_1)$ are exactly identical with the constraints indexed with $t \geq 2$ in $J^D_2(I_1 + n^{*1}_1 - \mu)$. Since $J^D_2(I_1 + n^{*1}_1 - \mu)$ is a sub-problem of $J^D_1(I_1)$ and n^{*1}_1 is optimal for $J^D_1(I_1)$, the optimal dual variables for the constraints indexed with $t \geq 2$ in $J^D_1(I_1)$ are also optimal dual variables for the corresponding constraints in $J^D_2(I_1 + n^{*1}_1 - \mu)$. Since conditions (C1)-(C3) hold on $\mathcal{A}$, we can use $\hat{\lambda}^1_{t,i} = \lambda^1_{t,i}$, $\hat{\lambda}^2_{t,i} = \lambda^2_{t,i}$, $\hat{\lambda}^3_{t,i} = \lambda^3_{t,i}$, $\hat{\lambda}^4_{t,i} = \lambda^4_{t,i}$, $\hat{\xi}_t = \xi_t$, and $\hat{\Omega}_{t,i} = \Omega_{t,i}$ for $t \geq 2$ (note: $\lambda^1_{t,i}$, $\lambda^2_{t,i}$, $\lambda^3_{t,i}$, $\lambda^4_{t,i}$, ξ_t, and $\Omega_{t,i}$ are the dual variables for $J^D_1(I_1)$ in Step 1) and the KKT conditions (A.35)-(A.43) are all satisfied at $n_{t,i} = n^{*2}_{t,i}$. Thus, by the sufficiency of KKT conditions for optimality in linear program, we conclude that n^{*2} is optimal for $J^D_2(I_2)$. This is our base case.

Now, suppose that both (i) and (ii) hold for all $s \leq t-1$. We want to show that they still hold for $s = t$. By induction hypothesis, we can write: $I_{t,i} = I_{t-1,i} + n^{*\,t-1}_{t-1,i} - D_{t-1,i} = (\sum_{s=1}^{t-2} n^{*1}_{s,i} - (t-2)\mu_i - \Delta_{t-2,i} - \frac{1}{m}\sum_{s=1}^{t-3}\sum_{j=1}^{m} \Delta_{s,j}) + (n^{*1}_{t-1,i} + \Delta_{t-2,i} - \frac{1}{m}\sum_{j=1}^{m} \Delta_{t-2,j}) - D_{t-1,i} = \sum_{s=1}^{t-1} n^{*1}_{s,i} - (t-1)\mu_i - \Delta_{t-1,i} - \frac{1}{m}\sum_{s=1}^{t-2}\sum_{j=1}^{m} \Delta_{s,j}$. So, (i) holds. We now check (ii). At the beginning of period t, we have to solve the following linear program:

$$\begin{aligned}
J_t^D(I_t) &= \min_{n,y,\theta} \sum_{s=t}^{T} \left[c\, z_s^D\, B + \sum_{i=1}^{m} p_i\, y_{s,i} + \sum_{i=1}^{m} h_i \theta_{s,i} \right] \\
& \text{s.t.} \quad y_{t',i} \geq (t'-t+1)\mu_i - \sum_{s=t}^{t'} n_{s,i} - I_{t,i} \qquad \forall t' \geq t,\, i \\
& \quad\quad y_{t',i} \geq 0 \qquad \forall t' \geq t,\, i \\
& \quad\quad \theta_{t',i} \geq I_{t,i} + \sum_{s=t}^{t'} n_{s,i} - (t'-t+1)\mu_i \qquad \forall t' \geq t,\, i \\
& \quad\quad \theta_{t',i} \geq 0 \qquad \forall t' \geq t,\, i \\
& \quad\quad \sum_{i=1}^{m} n_{s,i} = z_s^D B \qquad \forall s \geq t \\
& \quad\quad n_{s,i} \geq 0 \qquad \forall s \geq t,\, i
\end{aligned}$$

By similar arguments as before, it is not difficult to check that, on $\mathcal{A}$, we have: (C1) $n_{s,i}^{*t} \geq 0$ for all $s \geq t$ and i, (C2) $\sum_{i=1}^{m} n_{s,i}^{*t} = z_s^D B$ for all $s \geq t$, and (C3) $\sum_{s=t}^{t'} n_{s,i}^{*t} - (t'-t+1)\mu_i$ has the same sign (i.e., either strictly positive or strictly negative) as $\sum_{s=t}^{t'} n_{s,i}^{*1} - (t'-t+1)\mu_i$ for all $t' \geq t$ and i. (The last condition holds because $|\sum_{s=t}^{t'} n_{s,i}^{*t} - \sum_{s=t}^{t'} n_{s,i}^{*1}| = |n_{t,i}^{*t} - n_{t,i}^{*1}| = |\Delta_{t-1,i} - \frac{1}{m}\sum_{j=1}^{m} \Delta_{t-1,j}| \leq 2\sum_{j=1}^{m} |\Delta_{t-1,i}| < \varphi$ and $|\sum_{s=t}^{t'} n_{s,i}^{*1} - (t'-t+1)\mu_i| \geq \varphi$ on $\mathcal{A}$.) These allow us to use the same dual variables for the corresponding constraints in both $J_t^D(I_1 + \sum_{s=1}^{t-1} n_s^{*1} - (t-1)\mu)$ for $J_t^D(I_t)$; hence, by the sufficiency of KKT conditions, we conclude that n^{*t} is optimal for $J_t^D(I_t)$. This completes the induction.

Step 4

We now make two important observations: under CEC (i.e., using $n_s = n_s^{*s}$ for $s \leq t$), the starting inventory level for retailer i at the beginning of period t is given by $I_{t+1,i}^{CEC} = \sum_{s=1}^{t} n_{s,i}^{*1} - t\mu_i - \Delta_{t,i} - \frac{1}{m}\sum_{s=1}^{t-1}\sum_{j=1}^{m} \Delta_{s,j}$. In contrast, under the perfect hindsight policy discussed in Step 1, it is not difficult to check that the starting inventory level for retailer i at the beginning of period t is given by $I_{t+1,i}^{H} = \sum_{s=1}^{t} n_{s,i}^{*1} - t\mu_i - \frac{1}{m}\sum_{s=1}^{t}\sum_{j=1}^{m} \Delta_{s,j}$ (because $n_{t,i}^{H} = n_{t,i}^{*1} + \Delta_{t,i} - \frac{1}{m}\sum_{j=1}^{m} \Delta_{t,j}$ and $I_{t+1,i}^{H} = \sum_{s=1}^{t} n_{s,i}^{H} - \sum_{s=1}^{t} D_{s,i}$). So, $I_{t+1,i}^{CEC} = I_{t+1,i}^{H} - \Delta_{t,i} + \frac{1}{m}\sum_{j=1}^{m} \Delta_{t,j}$, which implies:

$$J^{CEC} - J^{H} \leq \sum_{t=1}^{T}\sum_{i=1}^{m} (p_i + h_i) \left| \Delta_{t,i} - \frac{1}{m}\sum_{j=1}^{m} \Delta_{t,j} \right|$$

The result follows because $\mathbf{E}[|\Delta_{t,i}|\, \mathbf{1}\{\mathcal{A}\}] \leq \mathbf{E}[|\Delta_{t,i}|] \leq \sigma_i \leq \sigma$ for all t and i.

□

A.1.5 Proof of Theorem 2.5.3

We proceed in several steps. In Step 1, we prove $\mathbf{E}[(J^{DC} - J^H)\,\mathbf{1}\{\hat{\mathcal{A}}\}] \leq 2m(p^* + h^*)(T+1)^{3/2}\sigma^*$; in Step 2, we prove $\mathbf{E}[(J^{CEC} - J^H)\,\mathbf{1}\{\hat{\mathcal{A}}\}] \leq 2m(p^* + h^*)T\sigma^*$. The arguments in Steps 1 and 2 are similar to those in the proof of Theorem 3; however, instead of using duality arguments, we will use convexity arguments. In Step 3, we compute a bound for $P(\hat{\mathcal{A}})$; in Step 4, we put all the results from Steps 1-3 together to get a bound for each $\frac{\mathbf{E}[J^{DC}]-J^*}{J^*}$ and $\frac{\mathbf{E}[J^{CEC}]-J^*}{J^*}$.

Step 1

Define a sequence $\{n^H_{t,i}\}$ as follows: $n^H_{t,i} = \frac{z^D_t B}{m} + \Delta_{t,i} - \frac{1}{m}\sum_{j=1}^m \Delta_{t,j}$. We claim that if $D_{1:T} \in \hat{\mathcal{A}}$, then n^H is an optimal allocation for the hindsight problem. To see this, simply note that, under any feasible allocation $\{n_{t,i}\}$, we must have:

$$
\begin{aligned}
&\sum_{t=1}^T \left[c\, z^D_t\, B + \sum_{i=1}^m p^*\,(D_{t,i} - x_{t,i} - n_{t,i})^+ + \sum_{i=1}^m h^*(x_{t,i} + n_{t,i} - D_{t,i})^+ \right] \\
= \; & \sum_{t=1}^T \left[c\, z^D_t\, B + \sum_{i=1}^m p^* \left(\sum_{s=1}^t D_{s,i} - \sum_{s=1}^t n_{s,i} \right)^+ + \sum_{i=1}^m h^* \left(\sum_{s=1}^t n_{s,i} - \sum_{s=1}^t D_{s,i} \right)^+ \right] \\
\geq \; & \sum_{t=1}^T \left[c\, z^D_t\, B + \sum_{i=1}^m p^* \left(\frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m D_{s,j} - \frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m n_{s,j} \right)^+ \right. \\
& \left. + \sum_{i=1}^m h^* \left(\frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m n_{s,j} - \frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m D_{s,j} \right)^+ \right] \\
= \; & \sum_{t=1}^T \left[c\, z^D_t\, B + \sum_{i=1}^m p^* \left(\frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m D_{s,j} - \frac{1}{m}\sum_{s=1}^t z^D_s B \right)^+ \right. \\
& \left. + \sum_{i=1}^m h^* \left(\frac{1}{m}\sum_{s=1}^t z^D_s B - \frac{1}{m}\sum_{s=1}^t \sum_{j=1}^m D_{s,j} \right)^+ \right]
\end{aligned}
$$

where the first equality holds because $x_{t+1} = x_t + n_t - D_t$ for all $t \geq 1$, the first inequality follows by the convexity of $(\cdot)^+$, and the last inequality follows from $\sum_{i=1}^m n_{t,i} = z^D_t B$. Since the initial inventory level for retailer i at the beginning of period $t+1$ is given by $I_{t+1,i} = \sum_{s=1}^t n_{s,i} - \sum_{s=1}^t D_{s,i}$, it is not difficult to check that the above lower bound (the mathematical expression after the last equality) is achieved by setting $n_t = n^H_t$. Moreover, on $\hat{\mathcal{A}}$, we have $n^H_t > 0$ for all t. So, n^H is an optimal feasible solution for $J^H(D_{1:T})$.

We will now prove that $\mathbf{E}[(J^{DC} - J^H)\,\mathbf{1}\{\hat{\mathcal{A}}\}] \leq 2m(p^* + h^*)(T+1)^{3/2}\sigma^*$. This is not difficult to show. Note that, by similar convexity arguments as above, it can be shown that $n^D_{t,i} = z^D_t B/m$ for all i is an optimal solution of $J^D_1(I_1)$. The initial inventory level for retailer i at the beginning of period $t+1$ under DC is thus given by $I^{DC}_{t+1} = \sum_{s=1}^{t} n^D_s - \sum_{s=1}^{t} D_s$, which implies:

$$J^{DC} = \sum_{t=1}^{T}\left[c\, z^D_t\, B + \sum_{i=1}^{m} p^*\left(\sum_{s=1}^{t} D_{s,i} - \frac{1}{m}\sum_{s=1}^{t} z^D_s B\right)^+ + \sum_{i=1}^{m} h^*\left(\frac{1}{m}\sum_{s=1}^{t} z^D_s B - \sum_{s=1}^{t} D_{s,i}\right)^+\right]$$

So, on $\hat{\mathcal{A}}$, we can bound $J^{DC} - J^H$ as follows:

$$\begin{aligned} J^{DC} - J^H &= \sum_{t=1}^{T}\sum_{i=1}^{m} p^*\left[\left(\sum_{s=1}^{t} D_{s,i} - \frac{1}{m}\sum_{s=1}^{t} z^D_s B\right)^+ - \left(\frac{1}{m}\sum_{s=1}^{t}\sum_{j=1}^{m} D_{s,j} - \frac{1}{m}\sum_{s=1}^{t} z^D_s B\right)^+\right] \\ &\quad + \sum_{t=1}^{T}\sum_{i=1}^{m} h^*\left[\left(\frac{1}{m}\sum_{s=1}^{t} z^D_s B - \sum_{s=1}^{t} D_{s,i}\right)^+ - \left(\frac{1}{m}\sum_{s=1}^{t} z^D_s B - \frac{1}{m}\sum_{s=1}^{t}\sum_{j=1}^{m} D_{s,j}\right)^+\right] \\ &\leq \sum_{t=1}^{T}\sum_{i=1}^{m}(p^* + h^*)\left[\left|\sum_{s=1}^{t}\Delta_{s,i}\right| + \left|\frac{1}{m}\sum_{s=1}^{t}\sum_{j=1}^{m}\Delta_{s,j}\right|\right] \end{aligned}$$

where the inequality follows because $\mu_i = \mu^*$ for all i and $(a-c)^+ - (b-c)^+ \leq |a| + |b|$ for all a, b, and c (note that $\sum_{s=1}^{t} D_{s,i} = t\mu^* + \sum_{s=1}^{t}\Delta_{s,i}$ and $\frac{1}{m}\sum_{s=1}^{t}\sum_{j=1}^{m} D_{s,j} = t\mu^* + \frac{1}{m}\sum_{s=1}^{t}\sum_{j=1}^{m}\Delta_{s,j}$). The bound $\mathbf{E}[(J^{DC} - J^H)\,\mathbf{1}\{\hat{\mathcal{A}}\}] \leq 2m(p^* + h^*)(T+1)^{3/2}\sigma^*$ follows because $\mathbf{E}[|\Delta_{t,i}|\,\mathbf{1}\{\hat{\mathcal{A}}\}] \leq \mathbf{E}[|\Delta_{t,i}|] \leq \sigma^*$ for all t and i.

Step 2

We now argue that $\mathbf{E}[(J^{CEC} - J^H)\,\mathbf{1}\{\hat{\mathcal{A}}\}] \leq 2m(p^* + h^*)T\sigma^*$. Define n^{*t} as follows: $n^{*1}_{t,i} = \frac{z^D_t B}{m}$ for all t, $n^{*t}_{t,i} = \frac{z^D_t B}{m} + \Delta_{t-1,i} - \frac{1}{m}\sum_{j=1}^{m}\Delta_{t-1,j}$, and $n^{*t}_{s,i} = \frac{z^D_s B}{m}$ for all $s > t$. Suppose that $D_{1:T} \in \hat{\mathcal{A}}$. We will prove the following statement: If we use $n_s = n^{*s}_s$ or all $s \leq t-1$, then

(i) The starting inventory level for retailer i at the beginning of period t is given by

$$I_{t,i} = \frac{1}{m}\sum_{s=1}^{t-1} z^D_s B - (t-1)\mu^* - \Delta_{t-1,i} - \frac{1}{m}\sum_{s=1}^{t-2}\sum_{j=1}^{m}\Delta_{s,j}$$

(ii) n^{*t} is an optimal solution of $J^D_t(I_t)$.

As in Step 3 in the proof of Theorem 3, both (i) and (ii) can be proved by induction. We start with $t=2$. (The case $t=1$ is trivial.) At the beginning of period 2, we have $I_{2,i} = I_{1,i} + n^{*1}_{1,i} - D_{1,i} = \frac{z^D_1 B}{m} - \mu^* - \Delta_{1,i}$. So, (i) holds. To show that n^{*2} is optimal for $J^D_2(I_2)$, by convexity of $(\cdot)^+$, for any feasible allocation $\{n_{t,i}\}$, we can bound:

$$\begin{aligned}
& \sum_{t=2}^{T}\left[c\, z_t^D\, B + \sum_{i=1}^{m} p^*\,(\mu^* - I_{t,i} - n_{t,i})^+ + \sum_{i=1}^{m} h^*(I_{t,i} + n_{t,i} - \mu^*)^+\right] \\
= & \sum_{t=2}^{T}\left[c\, z_t^D\, B + \sum_{i=1}^{m} p^* \left((t-1)\mu^* - \sum_{s=2}^{t} n_{s,i} - I_{2,i}\right)^+ + \sum_{i=1}^{m} h^* \left(I_{2,i} + \sum_{s=2}^{t} n_{s,i} - (t-1)\mu^*\right)^+\right] \\
\geq & \sum_{t=2}^{T}\left[c\, z_t^D\, B + \sum_{i=1}^{m} p^* \left((t-1)\mu^* - \frac{1}{m}\sum_{s=2}^{t}\sum_{j=1}^{m} n_{s,j} - \frac{1}{m}\sum_{j=1}^{m} I_{2,j}\right)^+ \right. \\
& \left. + \sum_{i=1}^{m} h^* \left(\frac{1}{m}\sum_{j=1}^{m} I_{2,j} + \frac{1}{m}\sum_{s=2}^{t}\sum_{j=1}^{m} n_{s,j} - (t-1)\mu^*\right)^+\right] \\
= & \sum_{t=2}^{T}\left[c\, z_t^D\, B + \sum_{i=1}^{m} p^* \left((t-1)\mu^* - \frac{1}{m}\sum_{s=2}^{t} z_s^D B - \frac{1}{m}\sum_{j=1}^{m} I_{2,j}\right)^+ \right. \\
& \left. + \sum_{i=1}^{m} h^* \left(\frac{1}{m}\sum_{j=1}^{m} I_{2,j} + \frac{1}{m}\sum_{s=2}^{t} z_s^D B - (t-1)\mu^*\right)^+\right]
\end{aligned}$$

where the inequality holds because $\frac{1}{m}\sum_{i=1}^{m} x_i^+ \geq (\frac{1}{m}\sum_{i=1}^{m} x_i)^+$ for all x_i's, which implies $\sum_{i=1}^{m} x_i^+ \geq \sum_{i=1}^{m}(\frac{1}{m}\sum_{j=1}^{m} x_i)^+$. Recursively solving $(t-1)\mu^* - \sum_{s=2}^{t} n_{s,i} - I_{2,i} = (t-1)\mu^* - \frac{1}{m}\sum_{s=2}^{t} z_s^D B - \frac{1}{m}\sum_{j=1}^{m} I_{2,j}$ yields $n_{s,i} = n_{s,i}^{*2}$ for all $s \geq 2$ and i. Since n^{*2} exactly achieves the lower bound, it must be optimal. Since $z_t^D \geq 1$ for all t, we have $n^{*2} > 0$ on $\hat{\mathcal{A}}$, so it is a feasible optimal solution. This is our base case.

Now, suppose that both (i) and (ii) hold for all $s \leq t-1$. We want to show that they also hold for $s = t$. By induction hypothesis, we have:

$$\begin{aligned}
I_{t,i} &= I_{t-1,i} + n_{t-1,i}^{*\,t-1} - D_{t-1,i} \\
&= \left(\frac{1}{m}\sum_{s=1}^{t-2} z_s^D B - (t-2)\mu^* - \Delta_{t-2,i} - \frac{1}{m}\sum_{s=1}^{t-3}\sum_{j=1}^{m}\Delta_{s,j}\right) \\
&\quad + \left(\frac{z_{t-1}^D B}{m} + \Delta_{t-2,i} - \frac{1}{m}\sum_{j=1}^{m}\Delta_{t-2,j}\right) - D_{t-1,i} \\
&= \frac{1}{m}\sum_{s=1}^{t-1} z_s^D B - (t-1)\mu^* - \Delta_{t-1,i} - \frac{1}{m}\sum_{s=1}^{t-2}\sum_{j=1}^{m}\Delta_{s,j}
\end{aligned}$$

so, (i) holds. As for (ii), using similar convexity arguments as above, it is not difficult to check that n^{*t} is optimal for $J_t^D(I_t)$ on $\hat{\mathcal{A}}$. This completes the induction.

Putting our results together, the starting inventory level for retailer i at the beginning of period t under CEC is given by $I_{t+1,i}^{CEC} = \frac{1}{m}\sum_{s=1}^{t} z_s^D B - t\mu^* - \Delta_{t,i} - \frac{1}{m}\sum_{s=1}^{t-1}\sum_{j=1}^{m}\Delta_{s,j}$ and the starting inventory level under hindsight policy (from Step 1) is given by $I_{t+1,i}^{H} = \frac{1}{m}\sum_{s=1}^{t} z_s^D B - t\mu^* -$

$\frac{1}{m}\sum_{s=1}^{t}\sum_{j=1}^{m}\Delta_{s,j}$. So, $I^{CEC}_{t+1,i} = I^{H}_{t+1,i} - \Delta_{t,i} + \frac{1}{m}\sum_{j=1}^{m}\Delta_{t,j}$. This implies:

$$J^{CEC} - J^{H} \leq \sum_{t=1}^{T}\sum_{i=1}^{m}(p^* + h^*)\left|\Delta_{t,i} - \frac{1}{m}\sum_{j=1}^{m}\Delta_{t,j}\right|$$

Taking expectation yields $\mathbf{E}[(J^{CEC} - J^{H})\,\mathbf{1}\{\hat{\mathcal{A}}\}] \leq 2m(p^* + h^*)T\sigma^*$.

Step 3

We claim that $P(\hat{\mathcal{A}}) \geq 1 - 2mTe^{-\min\{\mu^*,\, B^2/(16m^4\mu^*)\}}$ for all sufficiently large B^2/μ^* and μ^*. To see this, first, note that we can bound $P(\hat{\mathcal{A}})$ using sub-additive property of probability as follows:

$$P(\hat{\mathcal{A}}) \;=\; \prod_{t=1}^{T} P\left(2m\sum_{i=1}^{m}|\Delta_{t,i}| < B\right) \geq \prod_{t=1}^{T}\left[1 - \sum_{i=1}^{m} P\left(|\Delta_{t,i}| \geq \frac{B}{2m^2}\right)\right]$$

By an exponential tail bound for Poisson random variable, it can be shown that $P(|\Delta_{t,i}| \geq B/2m^2) \leq 2e^{-\min\{\mu^*,\, B^2/(16m^4\mu^*)\}}$ (see below for proof). Applying Bernoulli's inequality, we get $P(\hat{\mathcal{A}}) \geq (1 - 2m\; e^{-\min\{\mu^*,\, B^2/(16m^4\mu^*)\}})^T \geq 1 - 2mTe^{-\min\{\mu^*,\, B^2/(16m^4\mu^*)\}}$ for all large B^2/μ^* and μ^*. This proves our claim.

Proof of an exponential tail bound for Poisson random variable. By Markov's inequality, $P(|\Delta_{t,i}| > B/2m^2) \leq \mathbf{E}[e^{r|\Delta_{t,i}|}]/e^{rB/2m^2}$ for $r > 0$. By moment generating function of Poisson distribution, as long as $r \leq 1$, $\mathbf{E}[e^{r|\Delta_{t,i}|}] \leq \mathbf{E}[e^{r\Delta_{t,i}}] + \mathbf{E}[e^{-r\Delta_{t,i}}] = e^{\mu^*(e^r - 1 - r)} + e^{\mu^*(e^{-r} - 1 + r)} \leq 2e^{\mu^* r^2}$. (The last inequality holds because $e^r - 1 - r \leq r^2$ for all $0 \leq r \leq 1$.) This implies $P(|\Delta_{t,i}| > B/2m^2) \leq 2e^{\mu^*(r^2 - (B/2\mu^* m^2)r)}$ for all $0 \leq r \leq 1$. Minimizing the bound over $r > 0$, yields $r = B/(4\mu^* m^2)$. If $B/(4\mu^* m^2) \leq 1$, we can use the above bound and get $P(|\Delta_{t,i}| > B/2m^2) \leq 2e^{-B^2/(16m^4\mu^*)}$. If, on the other hand, $B/(4\mu^* m^2) > 1$, we can simply bound: $P(|\Delta_{t,i}| > B/2m^2) \leq \mathbf{E}[e^{r|\Delta_{t,i}|}]/e^{rB/2m^2} \leq (\mathbf{E}[e^{r\Delta_{t,i}}] + \mathbf{E}[e^{-r\Delta_{t,i}}])/e^{rB/2m^2} = (e^{\mu^*(e^r - 1 - r)} + e^{\mu^*(e^{-r} - 1 + r)})/e^{rB/2m^2}$ and use $r = 1$, which yields $P(|\Delta_{t,i}| > B/2m^2) \leq 2e^{\mu^* - (B/2m^2)} \leq 2e^{-\mu^*}$. Putting the bounds for the case $B/(4\mu^* m^2) \leq 1$ and $B/(4\mu^* m^2) > 1$ together, we conclude that $P(|\Delta_{t,i}| > B/2m^2) \leq 2e^{-\min\{\mu^*,\, B^2/(16m^4\mu^*)\}}$.

Step 4

We now put the results in Steps 1-3 together. From Step 1, we have:

$$\begin{aligned} J^* \;&\geq\; \mathbf{E}[J^H] \\ &\geq\; \sum_{t=1}^{T}\mathbf{E}\left[cz_t^D B + \sum_{i=1}^{m} p^* \left(\frac{1}{m}\sum_{s=1}^{t}\sum_{j=1}^{m} D_{s,j} - \frac{1}{m}\sum_{s=1}^{t} z_s^D B\right)^+ \right. \\ &\qquad \left. + \sum_{i=1}^{m} h^* \left(\frac{1}{m}\sum_{s=1}^{t} z_s^D B - \frac{1}{m}\sum_{s=1}^{t}\sum_{j=1}^{m} D_{s,j}\right)^+ \right] \end{aligned}$$

Note that $\sum_{t=1}^{T} z_t^D B \geq TB \geq T\mu^*$ (because we assume $z_t^D > 0$). For each t, if $\sum_{s=1}^{t} z_s^D B \geq tm\mu^*$, we can bound:

$$\left(\frac{1}{m} \sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j} - \frac{1}{m} \sum_{s=1}^{t} z_s^D B \right)^+ \geq 0 \quad \text{and}$$
$$\left(\frac{1}{m} \sum_{s=1}^{t} z_s^D B - \frac{1}{m} \sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j} \right)^+ \geq \left(t\mu^* - \frac{1}{m} \sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j} \right)^+$$

If, on the other hand, $\sum_{s=1}^{t} z_s^D B < tm\mu^*$, we can bound:

$$\left(\frac{1}{m} \sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j} - \frac{1}{m} \sum_{s=1}^{t} z_s^D B \right)^+ \geq \left(\frac{1}{m} \sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j} - t\mu^* \right)^+ \quad \text{and}$$
$$\left(\frac{1}{m} \sum_{s=1}^{t} z_s^D B - \frac{1}{m} \sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j} \right)^+ \geq 0$$

In either case, we can further bound J^* as follows:

$$J^* \geq cT\mu^* + h^* \sum_{t=1}^{T} \min \left\{ \mathbf{E}\left[\left(\sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j} - tm\mu^* \right)^+ \right], \mathbf{E}\left[\left(tm\mu^* - \sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j} \right)^+ \right] \right\}$$

Since $\mathbf{E}[\sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j}] = tm\mu^*$, by Central Limit Theorem (CLT), the expectations inside the $\min\{\cdot\}$ operator are both of order $\sqrt{tm\mu^*}$ (for large μ^*), i.e.,

$$\mathbf{E}\left[\left(\sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j} - tm\mu^* \right)^+ \right] = \Theta(\sqrt{tm\mu^*}) \quad \text{and}$$
$$\mathbf{E}\left[\left(tm\mu^* - \sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j} \right)^+ \right] = \Theta(\sqrt{tm\mu^*})$$

where $\Theta(\cdot)$ is the usual big-Θ notation. (Intuitively, the order magnitude of these two expectations are of the same size as the order magnitude of the standard deviation of $\sum_{s=1}^{t} \sum_{j=1}^{m} D_{s,j}$.) This means that there exists a constant $M' > 0$ independent of $T > 0$ such that, for all large μ^*, we have

$$J^* \geq cT\mu^* + M' \sqrt{\mu^*}\, T^{3/2}$$

where the inequality follows because $\sum_{t=1}^{T} \sqrt{t}$ is of order $T^{3/2}$.

Now, for any feasible policy π, on $\hat{\mathcal{A}}^c$, we can loosely bound J^π as follows:

$$J^\pi \;\leq\; \sum_{t=1}^{T} c z_t^D B + \sum_{t=1}^{T}\sum_{i=1}^{m} T(p^* + h^*) D_{t,i}$$

In the above, we simply assume that each unit of demand incurs both the holding and penalty cost T times. This is the maximum possible total costs that can be incurred per unit inventory. Since $\mathbf{E}[D_{t,i}\mathbf{1}\{\hat{\mathcal{A}}^c\}] \leq \mathbf{E}[D_{t,i}^2]^{1/2}\, \mathbf{E}[\mathbf{1}\{\hat{\mathcal{A}}^c\}^2]^{1/2} \leq 2\mu^* P(\hat{\mathcal{A}}^c)^{1/2}$ (because $\mathbf{E}[D_{t,i}^2] = Var(D_{t,i}) + \mathbf{E}[D_{t,i}]^2 = \mu^* + (\mu^*)^2 \leq 4(\mu^*)^2$ for $\mu^* \geq 1$), using the exponential tail bound proved in Step 3, we can bound:

$$\begin{aligned}
\mathbf{E}[(J^\pi - J^*)\,\mathbf{1}\{\hat{\mathcal{A}}^c\}] &\leq \mathbf{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{m} T(p^* + h^*) D_{t,i}\mathbf{1}\{\hat{\mathcal{A}}^c\}\right] \\
&\leq 2\mu^* m\, T^2\, (p^* + h^*) P(\hat{\mathcal{A}}^c)^{1/2} \\
&\leq 2\sqrt{2}\mu^* m^{3/2}\, T^{5/2}\, (p^* + h^*) e^{-\min\{\mu^*/2,\, B^2/(32 m^4 \mu^*)\}} \\
&\leq M''
\end{aligned}$$

for some $M'' > 0$ independent of T, B, and μ^*, for all sufficiently large B^2/μ^* and μ^*, and $T = o(e^{\min\{\mu^*/5,\, B^2/(80 m^4 \mu^*)\}})$. (The first inequality follows because, under both π and the optimal policy, we always incur the purchasing costs $\sum_{t=1}^{T} c z_t^D B$ (so, this term cancels out), and we simply ignore the total holding and penalty costs under the optimal policy.) Putting this together with the bounds in Steps 1 and 2, we conclude that, for all large B^2/μ^* and μ^*, and $T = o(e^{\min\{\mu^*/5,\, B^2/(80 m^4 \mu^*)\}})$,

$$\frac{\mathbf{E}[J^{DC}] - J^*}{J^*} \;\leq\; \frac{2m(p^* + h^*)(T+1)^{3/2}\sqrt{\mu^*} + M''}{cT\mu^* + M'\sqrt{\mu^*}\; T^{3/2}} \;\leq\; \frac{M\sqrt{T}}{\sqrt{\mu^*} + \sqrt{T}} \quad \text{and}$$

$$\frac{\mathbf{E}[J^{CEC}] - J^*}{J^*} \;\leq\; \frac{2m(p^* + h^*)T\sqrt{\mu^*} + M''}{cT\mu^* + M'\sqrt{\mu^*}\; T^{3/2}} \;\leq\; \frac{M}{\sqrt{\mu^*} + \sqrt{T}}$$

for some $M > 0$ independent of T, B, and μ^*.

□

A.1.6 Proof of Lemma 2.6.1

We prove Lemma 2.6.1 by induction. As in the proof of Theorem 2.5.3, the idea is again to characterize a candidate optimal solution using KKT conditions. Let $F_{t,t'}(\cdot)$ denote the c.d.f. of $\sum_{\xi=t}^{t'} D_{\xi,i}$, defined on an extended space $(-\infty, \infty)$ as follows:

$$F_{t,t'}(x) \;=\; \begin{cases} P\left(\sum_{\xi=t}^{t'} D_{\xi,i} \leq x\right) & \text{for } x > 0 \\ 0 & \text{for } x \leq 0 \end{cases}$$

Also, define:

$$G_t(n_t; I_t) \;:=\; \sum_{\xi=t}^{T}\sum_{i=1}^{m}\mathbf{E}\left[p^*\left(\sum_{s=t}^{\xi} D_{s,i} - \sum_{s=t}^{\xi} n_{s,i} - I_{t,i}\right)^+ + \; h^*\left(I_{t,i} + \sum_{s=t}^{\xi} n_{s,i} - \sum_{s=t}^{\xi} D_{s,i}\right)^+\right]$$

The following expression of partial derivative of G_t is useful for the proof:

$$\frac{\partial G_t}{\partial n_{s,i}} \;=\; \sum_{k=s}^{T}\left[(h^*+p^*)F_{t,k}\left(\sum_{\xi=t}^{k} n_{\xi,i} + I_{t,i}\right) - p^*\right] \quad \text{for all } s \geq t. \tag{A.44}$$

We now proceed in two steps. In Step 1, we show that the lemma is true for $t = 1$ (this is our induction base case); in Step 2, we show that the lemma is true for $t > 1$ by induction hypothesis.

<u>Step 1</u>

Consider $J_1^S(0)$. By (A.44) and KKT conditions, there exist dual variables $v_{s,i}^1 \geq 0$ and w_s^1 corresponding to constraints $n_{s,i} \geq 0$ and $\sum_{i=1}^{m} n_{n_{s,i}} = z_s^D B$, respectively, such that:

$$\frac{\partial G_1}{\partial n_{T,i}} \;=\; (h^*+p^*)F_{1,T}\left(\sum_{\xi=1}^{T} n_{\xi,i}\right) - p^* \;=\; v_{T,i}^1 + w_T^1 \tag{A.45}$$

$$\frac{\partial G_1}{\partial n_{T-1,i}} \;=\; \sum_{k=T-1}^{T}\left[(h^*+p^*)F_{1,k}\left(\sum_{\xi=1}^{k} n_{\xi,i}\right) - p^*\right] \;=\; v_{T-1,i}^1 + w_{T-1}^1 \tag{A.46}$$

$$\vdots$$

$$\frac{\partial G_1}{\partial n_{1,i}} \;=\; \sum_{k=1}^{T}\left[(h^*+p^*)F_{1,k}\left(\sum_{\xi=1}^{k} n_{\xi,i}\right) - p^*\right] \;=\; v_{1,i}^1 + w_1^1 \tag{A.47}$$

$$v_{s,i}^1 \cdot n_{s,i} \;=\; 0 \quad \forall\, s \geq 1 \tag{A.48}$$

Equivalently, (A.45)-(A.47) can be written as:

$$(h^*+p^*)F_{1,T}\left(\sum_{\xi=1}^{T} n_{\xi,i}\right) - p^* \;=\; v_{T,i}^1 + w_T^1 \tag{A.49}$$

$$(h^*+p^*)F_{1,T-1}\left(\sum_{\xi=1}^{T-1} n_{\xi,i}\right) - p^* \;=\; v_{T-1,i}^1 + w_{T-1}^1 - (v_{T,i}^1 + w_T^1) \tag{A.50}$$

$$\vdots$$

$$(h^*+p^*)F_{1,1}(n_{1,i}) - p^* \;=\; v_{1,i}^1 + w_1^1 - (v_{2,i}^1 + w_2^1) \tag{A.51}$$

$$v_{s,i}^1 \cdot n_{s,i} \;=\; 0 \quad \forall\, s \geq 1 \tag{A.52}$$

Let $\theta_{T,i}^1 = v_{T,i}^1 + w_T^1$ and $\theta_{s,i}^1 := v_{s,i}^1 + w_s^1 - (v_{s+1,i}^1 + w_{s+1}^1)$ for $s \leq T-1$. We claim that

$n_{s,i} = z_s^D B/m$ is the unique optimal solution of $J_1^S(0)$. To prove this, note that, if we set $v_{s,i}^1 = 0$ for all s and i, the variables $\{\theta_{s,i}^1\}$ are independent of i. By abuse of notation, let $\theta_{s,i}^1 = \theta_s^1$ for all s and i. By (A.49), $\sum_{\xi=1}^T n_{\xi,i} = F_{1,T}^{-1}\left(\frac{p^*+\theta_T^1}{p^*+h^*}\right)$. Taking the sum over all i's, and using $\sum_{i=1}^m n_{\xi,i} = z_\xi^D B$ for all $\xi \geq 1$, gives $\sum_{\xi=1}^T z_\xi^D B = m\, F_{1,T}^{-1}\left(\frac{p^*+\theta_T^1}{p^*+h^*}\right)$, which implies:

$$\sum_{\xi=1}^{T} n_{\xi,i} \quad = \quad \sum_{\xi=1}^{T} \frac{z_\xi^D B}{m}$$

Similarly, by (A.50) and (A.51), for $s \leq T-1$, we have $\sum_{\xi=1}^s n_{\xi,i} = F_{1,s}^{-1}\left(\frac{p^*+\theta_s^1}{p^*+h^*}\right)$. Taking the sum over all i's gives $\sum_{\xi=1}^s z_\xi^D B = m\, F_{1,s}^{-1}\left(\frac{p^*+\theta_s^1}{p^*+h^*}\right)$, which implies:

$$\sum_{\xi=1}^{s} z_\xi^D B = m\, F_{1,s}^{-1}\left(\frac{p^*+\theta_s^1}{p^*+h^*}\right) \quad \text{and} \quad \sum_{\xi=1}^{s} n_{\xi,i} = \sum_{\xi=1}^{s} \frac{z_\xi^D B}{m}.$$

We conclude that:

$$n_{s,i} = \frac{z_s^D B}{m} \quad \text{and} \quad \theta_s^1 = (p^*+h^*) F_{1,s}\left(\sum_{\xi=1}^{s} \frac{z_\xi^D B}{m}\right) - p^* \quad \text{for all } s$$

from which the constants $w_1^1, w_2^1, ..., w_T^1$ can be calculated properly. Note that, since $p^*+\theta_s^1 > 0$ and $p^*+\theta_s^1 < p^*+h^*$, the term $F_{1,s}^{-1}\left(\frac{p^*+\theta_s^1}{p^*+h^*}\right)$ is well-defined. Thus, we have just shown that there exist dual variables $v_{s,i}^1 \geq 0$ and w_s^1 that not only satisfy KKT conditions but also yield $n_{s,i} = z_s^D B/m$ for all s and i. Since the objective function in $J_1^S(0)$ is convex, our result for $t=1$ follows by the sufficiency of KKT conditions for optimality in convex optimization.

Step 2

Now, suppose that the the lemma holds for all $t \leq t'$. We want to show that it also holds for $t = t'+1$. Since the formula for n^{St} is exactly the same as the constructed optimal solution in Theorem 4, by the same arguments as in the proof of Theorem 4, we have:

$$I_{t'+1,i} = \frac{1}{m}\sum_{s=1}^{t'} z_s^D B - t'\mu^* - \Delta_{t',i} - \frac{1}{m}\sum_{s=1}^{t'-1}\sum_{j=1}^{m} \Delta_{s,j} \tag{A.53}$$

Consider $J_{t'+1}^S(I_{t'+1})$. By KKT conditions, there exist dual variables $v_{s,i}^{t'+1} \geq 0$ and $w_s^{t'+1}$ corresponding to constraints $n_{s,i} \geq 0$ and $\sum_{i=1}^m n_{n_{s,i}} = z_s^D B$, respectively, such that:

$$(h^* + p^*)F_{t'+1,T}\left(I_{t'+1,i} + \sum_{\xi=t'+1}^{T} n_{\xi,i}\right) - p^* \;=\; v_{T,i}^{t'+1} + w_T^{t'+1} \tag{A.54}$$

$$(h^* + p^*)F_{t'+1,T-1}\left(I_{t'+1,i} + \sum_{\xi=t'+1}^{T-1} n_{\xi,i}\right) - p^* \;=\; v_{T-1,i}^{t'+1} + w_{T-1}^{t'+1} - (v_{T,i}^{t'+1} + w_T^{t'+1}) \tag{A.55}$$

$$\vdots$$

$$(h^* + p^*)F_{t'+1,t'+1}\left(I_{t'+1,i} + n_{t'+1,i}\right) - p^* \;=\; v_{t'+1,i}^{t'+1} + w_{t'+1}^{t'+1} - (v_{t'+2,i}^{t'+1} + w_{t'+2}^{t'+1}) \tag{A.56}$$

$$v_{s,i}^{t'+1} \cdot n_{s,i} \;=\; 0 \;\; \forall\, s \tag{A.57}$$

Let $\theta_{T,i}^{t'+1} = v_{T,i}^{t'+1} + w_T^{t'+1}$ and $\theta_{s,i}^{t'+1} := v_{s,i}^{t'+1} + w_s^{t'+1} - (v_{s+1,i}^{t'+1} + w_{s+1}^{t'+1})$ for $s \leq T-1$. Arguing as in Step 1, set $v_{s,i}^{t'+1} = 0$ for all $s \geq t'+1$ and i. So, $\{\theta_{s,i}^{t'+1}\}$ become independent of i. By abuse of notation, we simply write: $\theta_{s,i}^{t'+1} = \theta_s^{t'+1}$. Now, by (A.54), $I_{t'+1,i} + \sum_{\xi=t'+1}^{T} n_{\xi,i} = F_{t'+1,T}^{-1}\left(\frac{p^*+\theta_T^{t'+1}}{p^*+h^*}\right)$. Taking the sum over all i's gives $\sum_{i=1}^{m} I_{t'+1,i} + \sum_{\xi=t'+1}^{T} z_\xi^D B = mF_{t'+1,T}^{-1}\left(\frac{p^*+\theta_T^{t'+1}}{p^*+h^*}\right)$, which implies:

$$I_{t'+1,i} + \sum_{\xi=t'+1}^{T} n_{\xi,i} \;=\; \sum_{j=1}^{m} \frac{I_{t'+1,j}}{m} + \sum_{\xi=t'+1}^{T} \frac{z_\xi^D B}{m}$$

Applying similar arguments to (A.55)-(A.56) yields:

$$I_{t'+1,i} + \sum_{\xi=t'+1}^{s} n_{\xi,i} \;=\; \sum_{j=1}^{m} \frac{I_{t'+1,j}}{m} + \sum_{\xi=t'+1}^{s} \frac{z_\xi^D B}{m} \quad \text{for all } s \geq t'+1 \text{ and } i$$

Solving for $\{n_{s,i}\}$, a combination of simple algebra and (A.53) give:

$$n_{t'+1,i} \;=\; \frac{z_{t'+1}^D B}{m} - I_{t'+1,i} + \sum_{j=1}^{m} \frac{I_{t'+1,j}}{m} \;=\; \frac{z_{t'+1}^D B}{m} + \Delta_{t',i} - \frac{1}{m}\sum_{j=1}^{m} \Delta_{t',j} \quad \text{and}$$

$$n_{s,i} \;=\; \frac{z_s^D B}{m} \quad \text{for all } s > t'+1$$

Note that $n_{s,i} > 0$ for all $s \geq t'+1$ on $\hat{\mathcal{A}}$; so, $\{n_{s,i}\}$ is feasible. Moreover, we also have:

$$\theta_s^{t'+1} \;=\; (p^* + h^*)F_{t'+1,s}\left(\sum_{j=1}^{m} \frac{I_{t'+1,j}}{m} + \sum_{\xi=t'+1}^{s} \frac{z_\xi^D B}{m}\right) - p^* \quad \text{for } s \geq t'+1$$

from which $w_{t'+1}^{t'+1}, \ldots, w_T^{t'+1}$ can be calculated properly. Since $0 < p^* + \theta_s^{t'+1} < p^* + h^*$, the term $F_{t'+1,T}^{-1}\left(\frac{p^*+\theta_s^{t'+1}}{p^*+h^*}\right)$ is well-defined on $\hat{\mathcal{A}}$. Thus, we have just shown that there exist dual variables and primal solution $\{n_{s,i}\}$ that satisfy KKT conditions.

□

A.1.7 Proof of Theorem 2.6.2

The proof of Theorem 2.6.2 is similar to the proof of Theorem 2.4.1. For any policy $\pi \in \Pi$, define a sequence $\{W^\pi_{t,i}\}$ where $W^\pi_{t+1,i} = W^\pi_{t,i} + N^\pi_{t,i} - \hat{D}_{t,i}$ for $t \geq 1$ and $W^\pi_1 = I^\pi_1 = I_1$. Observe that we can write: $W^\pi_{t,i} = I_{1,i} + \sum_{s=1}^{t-1} N^\pi_{s,i} - \sum_{s=1}^{t-1} \hat{D}_{s,i}$ and $I^\pi_{t,i} = I_{1,i} + \sum_{s=1}^{t-1} N^\pi_{s,i} - \sum_{s=1}^{t-1} D_{s,i}$. So, $I^\pi_{t,i} = W^\pi_{t,i} - \sum_{s=1}^{t-1} \tilde{\Delta}_{s,i}$, where $\tilde{\Delta}_{s,i} = D_{s,i} - \hat{D}_{s,i}$. We now proceed in three steps. In Step 1, we compute an upper bound for $\hat{C}^* - C^*$; in Step 2, we compute an upper bound for $\mathbf{E}[C^{\pi^R}] - \hat{C}^*$; and, in Step 3, we put together the results from Steps 1 and 2 to get a bound for $\mathbf{E}[C^{\pi^R}] - C^*$.

Step 1

We first compute an upper bound for $\hat{C}^* - C^*$. We claim that:

$$C^* \quad \geq \quad \hat{C}^* - \mathbf{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{m}(p_i + h_i)\left(\sum_{s=1}^{t}\tilde{\Delta}_{s,i}\right)^+\right]$$

For any policy $\pi \in \Pi$, we can bound:

$$\begin{aligned}
&\sum_{t=1}^{T}\left[c\, Z^\pi_t\, B + \sum_{i=1}^{m} p_i(D_{t,i} - N^\pi_{t,i} - I^\pi_{t,i})^+ + \sum_{i=1}^{m} h_i(N^\pi_{t,i} + I^\pi_{t,i} - D_{t,i})^+\right] \\
&\quad = \sum_{t=1}^{T}\left[c\, Z^\pi_t\, B + \sum_{i=1}^{m}(p_i + h_i)(I^\pi_{t,i} + N^\pi_{t,i} - D_{t,i})^+ - \sum_{i} p_i(I^\pi_{t,i} + N^\pi_{t,i} - D_{t,i})\right] \\
&\quad = \sum_{t=1}^{T}\left[c\, Z^\pi_t\, B + \sum_{i=1}^{m}(p_i + h_i)(I^\pi_{t+1,i})^+ - \sum_{i=1}^{m} p_i I^\pi_{t+1,i}\right] \\
&\quad \geq \sum_{t=1}^{T}\left[c\, Z^\pi_t\, B + \sum_{i=1}^{m}(p_i + h_i)(W^\pi_{t+1,i})^+ - \sum_{i=1}^{m} p_i W^\pi_{t+1,i}\right] \\
&\qquad - \sum_{t=1}^{T}\sum_{i=1}^{m}(p_i + h_i)\left(\sum_{s=1}^{t}\tilde{\Delta}_{s,i}\right)^+ + \sum_{t=1}^{T}\sum_{i=1}^{m}\sum_{s=1}^{t} p_i \tilde{\Delta}_{s,i}
\end{aligned}$$

where the first inequality holds because the identity $I^\pi_{t,i} = W^\pi_{t,i} - \sum_{s=1}^{t-1} \tilde{\Delta}_{s,i}$ implies $(I^\pi_{t,i})^+ \geq$

$(W^{\pi}_{t,i})^+ - (\sum_{s=1}^{t-1} \tilde{\Delta}_{s,i})^+$. Taking expectation on both sides of the inequality, we have:

$$\begin{aligned}
&\sum_{t=1}^{T} \mathbf{E}\left[c\, Z^{\pi}_t\, B + \sum_{i=1}^{m} p_i (D_{t,i} - N^{\pi}_{t,i} - I^{\pi}_{t,i})^+ + \sum_{i=1}^{m} h_i (N^{\pi}_{t,i} + I^{\pi}_{t,i} - D_{t,i})^+\right] \\
&\geq \sum_{t=1}^{T} \mathbf{E}\left[c\, Z^{\pi}_t\, B + \sum_{i=1}^{m} (p_i + h_i)(W^{\pi}_{t+1,i})^+ - \sum_{i=1}^{m} p_i W^{\pi}_{t+1,i}\right] - \mathbf{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{m}(p_i + h_i)\left(\sum_{s=1}^{t} \tilde{\Delta}_{s,i}\right)^+\right] \\
&\geq \hat{C}^* - \mathbf{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{m}(p_i + h_i)\left(\sum_{s=1}^{t} \tilde{\Delta}_{s,i}\right)^+\right]
\end{aligned}$$

The first inequality follows because, by definition, $\mathbf{E}[\tilde{\Delta}_{s,i}] = 0$ for all s and i, and the second inequality follows from the definition of $\hat{C}^*$ (by construction, $W^{\pi}_{t,i}$ is the starting inventory level for retailer i at the beginning of period t under π in the virtual world, and the optimal policy is $\hat{\pi}^*$, which gives expected total costs $\hat{C}^*$). Since the lower bound in the above inequality holds for all $\pi \in \Pi$, minimizing the expectation in the left side of the inequality over $\pi \in \Pi$ immediately yields $C^* \geq \hat{C}^* - \mathbf{E}[\sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i)(\sum_{s=1}^{t} \tilde{\Delta}_{s,i})^+]$ (by construction, $I^{\pi}_{t,i}$ is the starting inventory level for retailer i at the beginning of period t under π in the real world, and the corresponding optimal policy and expected total costs are π^* and C^*).

Step 2

We now compute an upper bound for $\mathbf{E}[C^{\pi^R}] - \hat{C}^*$. We claim that:

$$\mathbf{E}[C^{\pi^R}] - \hat{C}^* \;\leq\; \mathbf{E}\left[\sum_{t=1}^{T}\sum_{i=1}^{m}(p_i + h_i)\left(-\sum_{s=1}^{t} \tilde{\Delta}_{s,i}\right)^+\right]$$

This can be shown using similar arguments as in Step 1. Define a sequence $\{x_{t,i}\}$ where $x_{t+1,i} = x_{t,i} + N^{\pi^R}_{t,i} - \hat{D}_{t,i}$ for $t \geq 1$ and $x_1 = I_1$. Since $I^{\pi^R}_{t+1,i} = I^{\pi^R}_{t,i} + N^{\pi^R}_{t,i} - D_{t,i}$ and $I^{\pi^R}_{1,i} = I_{1,i}$, we can write $I^{\pi^R}_{t,i} = x_{t,i} - \sum_{s=1}^{t-1} \tilde{\Delta}_{s,i}$. We can then bound:

$$
\begin{aligned}
\mathbf{E}[C^{\pi^R}] &= \sum_{t=1}^{T} \mathbf{E}\left[c\, z_t^{\pi^R} B + \sum_{i=1}^{m} h_i (I^{\pi^R}_{t+1,i})^+ + \sum_{i}^{m} p_i (-I^{\pi^R}_{t+1,i})^+ \right] \\
&= \sum_{t=1}^{T} \mathbf{E}\left[c\, z_t^{\pi^R} B + \sum_{i=1}^{m} (p_i + h_i) (I^{\pi^R}_{t+1,i})^+ - \sum_{i}^{m} p_i I^{\pi^R}_{t+1,i} \right] \\
&\leq \sum_{t=1}^{T} \mathbf{E}\left[c\, z_t^{\pi^R} B + \sum_{i=1}^{m} (p_i + h_i) (x_{t+1,i})^+ - \sum_{i}^{m} p_i x_{t+1,i} \right] \\
&\qquad + \mathbf{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i) \left(-\sum_{s=1}^{t} \tilde{\Delta}_{s,i} \right)^+ \right] + \mathbf{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{m} \sum_{s=1}^{t} p_i \tilde{\Delta}_{s,i} \right] \\
&= \hat{C}^* + \mathbf{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i) \left(-\sum_{s=1}^{t} \tilde{\Delta}_{s,i} \right)^+ \right]
\end{aligned}
$$

The inequality follows since $I^{\pi^R}_{t,i} = x_{t,i} - \sum_{s=1}^{t-1} \tilde{\Delta}_{s,i}$ implies $(I^{\pi^R}_{t,i})^+ \leq (x_{t,i})^+ + (-\sum_{s=1}^{t-1} \tilde{\Delta}_{s,i})^+$ and the last equality follows from the definition of $\hat{C}^*$ and the fact that $z_t^{\pi^R} = z_t^{\hat{\pi}^*}$.

Step 3

Putting together the bounds from Steps 1 and 2, we conclude that

$$
\begin{aligned}
\mathbf{E}[C^{\pi^R}] - C^* &= \mathbf{E}[C^{\pi^R}] - \hat{C}^* + \hat{C}^* - C^* \\
&\leq \mathbf{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i) \left(-\sum_{s=1}^{t} \tilde{\Delta}_{s,i} \right)^+ \right] + \mathbf{E}\left[\sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i) \left(\sum_{s=1}^{t} \tilde{\Delta}_{s,i} \right)^+ \right] \\
&\leq 2 \sum_{t=1}^{T} \sum_{i=1}^{m} (p_i + h_i) \mathbf{E}\left[\left(\sum_{s=1}^{t} \tilde{\Delta}_{s,i} \right)^2 \right]^{1/2} \\
&\leq 2 \sum_{i=1}^{m} (p_i + h_i) \left[\sum_{t=1}^{T} \left(\sum_{s=1}^{t} \theta_{s,i}^2 \right)^{1/2} \right]
\end{aligned}
$$

□

A.2 Proofs of Results Appearing in Chapter 3

A.2.1 Proof of Theorem 3.4.1

Consider a general form of the cost function for the type (ii) policy

$$y^{(2)}(\tau) = a\tau + b\sqrt{\tau} + \frac{c}{\tau}, \tau > 0 \tag{A.58}$$

where a, b, and c are constants. Taking the derivative, we have

$$\frac{d}{d\tau}\, y^{(2)}(\tau) \quad = \quad a + \frac{1}{2} b\tau^{-\frac{1}{2}} - \frac{c}{\tau^2}. \tag{A.59}$$

Solving for $\frac{d}{d\tau}\, y^{(2)}(\tau) = 0$ yields

$$a\tau + \frac{1}{2} b\tau^{\frac{1}{2}} - \frac{c}{\tau} \quad = \quad 0. \tag{A.60}$$

The left hand side is strictly increasing, when $\tau \to 0$ it goes to $-\infty$, and when $x \to +\infty$ it goes to $+\infty$, so by Intermediate Value Theorem there is an unique τ_* such that $y'(\tau_*) = 0$. Since $y(0) = y(+\infty) = +\infty$, τ_* must be the unique minimizer of $y^{(2)}(\tau)$.

Similarly, the lower bound of the cost function of type (i) policy can be written as

$$y^{(1)}(t) = at + b\sqrt{\mu_L} + \frac{c}{t}, t > 0. \tag{A.61}$$

Let t_* be the minimizer of $y^{(1)}(t)$ which minimizes $a\tau + c/\tau$ in $y^{(2)}(\tau)$, then $\frac{d}{d\tau}\, y^{(2)}(t_*) = \frac{1}{2} b t_*^{-1/2} > 0$, therefore $t_* > \tau_*$. Then, by direct comparison of $y^{(1)}$ and $y^{(2)}$, we can see that if $\tau_* > \mu_L$, $y^{(2)} > y^{(1)}$.

□

A.2.2 Proof of Theorem 3.4.2

Consider a Markov chain on two states: $S = s$ and $S < s$ where S is the starting inventory position of an order cycle. The transition probability of this Markov chain is

$$P \quad = \quad \begin{bmatrix} \mathbf{P}(D(\tau) \leq u) & \mathbf{P}(D(\tau) > u) \\ p(u) & 1 - p(u) \end{bmatrix} \tag{A.62}$$

where $p(u)$ is a function of u representing the probability of transiting from state $S < s$ to $S = s$. Conditioning on $S = s - j$, we have

$$p(u) \quad = \quad \sum_{j=1}^{\infty} \mathbf{P}(S = s - j) \mathbf{P}(D(\tau) \leq u - j). \tag{A.63}$$

Clearly, $p(u)$ is an increasing function of u.

Let π_1, π_2 denote the stationary probability of the Markov chain. After some straightforward algebra, we can show

$$\pi_2 = \frac{\mathbf{P}(D(\tau) \geq u)}{p(u) + \mathbf{P}(D(\tau) \geq u)} \tag{A.64}$$

$$\leq \frac{\mathbf{P}(D(\tau) \geq u)}{p(\lambda\tau) + \mathbf{P}(D(\tau) \geq \lambda\tau)} \tag{A.65}$$

where the inequality holds since $p(u)$ is increasing in u. Under demand process D1, a Poisson process, or D2, a Brownian motion with drift, the numerator $\mathbf{P}(D(\tau) \geq u)$ has an exponential tail, and the denominator $p(\lambda\tau) + \mathbf{P}(D(\tau) \geq \lambda\tau)$ is a constant. Therefore π_2 decreases exponentially in u, so the probability of the starting inventory position S_∞ less than s decreases exponentially in u.

□

A.2.3 Proof of Theorem 3.7.1 and 3.7.2

W.l.o.g. we only prove the proposition for a fixed n so for brevity all function arguments with respect to n are dropped. The optimal centralized order quantity q_c is the minimizer of

$$y_c(q) = \frac{k_0\lambda}{q} + h\left(r(q) - \frac{\lambda(2\mu - n\lambda)}{2\mu(\mu - n\lambda)}q + \frac{q}{2}\right) + \frac{n\lambda}{2\mu}\frac{h_0\lambda}{\mu - n\lambda}q \tag{A.66}$$

$$= \frac{\lambda k_0}{q} + hr(q) + c_1 q \tag{A.67}$$

where c_1 is a constant.

Similarly, the optimal semi-centralized order quantity $q_s(k)$ is the minimizer of

$$y_s(q) = \frac{\lambda k}{q} + h\left(r(q) - \frac{\lambda(2\mu - n\lambda)}{2\mu(\mu - n\lambda)}q + \frac{q}{2}\right) \tag{A.68}$$

$$= \frac{\lambda k}{q} + hr(q) + c_2 q \tag{A.69}$$

for a given k and c_2 is a constant with $c_2 < c_1$.

Our goal is to show there exists a special $k = k_s$ such that $q_s(k_s) = q_c$. To proceed we'll first show several useful lemmas.

Lemma A.2.1. *$r(q)$ is continuous and strictly increasing in q under demand process D2 Brownian motion with drift.*

Proof:

From the Remark 3.4.1, the waiting time W has the following distribution

$$F_W(t) = (1-\rho)\sum_{j=0}^{\lfloor t\mu/q\rfloor} \frac{(\rho j - n\lambda t/q)^j}{j!} e^{-(\rho j - n\lambda t/q)} \quad t \geq 0 \tag{A.70}$$

For a fixed $\lfloor t\mu/q\rfloor$ the distribution function $F_W(t)$ is strictly decreasing in q. In other word, larger q results in stochastically larger waiting time.

From another form of $F_W(t)$ we have

$$W_k(\epsilon) \quad = \quad (1-\rho)e^{\rho(k+\epsilon)}\sum_{j=0}^{k} \frac{(\rho e^{-\rho}(j-k-\epsilon))^j}{j!} \tag{A.71}$$

where $k = \lfloor t\mu_0 \rfloor$ and $\epsilon = t\mu_0 - \lfloor t\mu_0 \rfloor = t\mu/q - \lfloor t\mu/q \rfloor$. Since $W_{k-1}(1) = W_k(0)$, $W_k(\epsilon)$ is continuous ϵ which implies the distribution of the waiting time $F_W(t)$ is continuous in q.

The lead time demand $D(L)$ can be decomposed as the demand during waiting time $D(W)$ plus the demand during deterministic process time $D(q/\mu)$. $D(W)$ has the distribution

$$F_D(d, W) = \int_0^\infty F_D(d, l)\; dF_W(l)$$

and $D(q/\mu)$ has the distribution $F_D(d, q/\mu)$, which both are strictly decreasing in q. In other words, larger q results in stochastically larger lead time demand. Under demand process D2, demand during fixed interval $D(l)$ is normally distributed thus $F_D(d, l)$ is continuous. Therefore $F_D(d, L)$ is continuous and strictly decreasing in q. Then inverse function of the lead time demand distribution $r(q) = F_D^{-1}(\alpha)$ is a continuous and strictly increasing in q.

□

Lemma A.2.2. $q_s(0) = 0$ *and* $q_s(k_0) > q_c$

Proof: The derivative of the centralized cost function is

$$y_c'(q) \quad = \quad -\frac{\lambda k_0}{q^2} + hr'(q) + c_1 \tag{A.72}$$

The derivative of the semi-centralized cost function is

$$y_s'(q) \quad = \quad -\frac{\lambda k_0}{q^2} + hr'(q) + c_2 \tag{A.73}$$

If $k = 0$ then the optimal solution is $q_s(0) = 0$. If $k = k_0$ then by the fact that $c_2 < c_1$ and the monotonicity of $r(q)$ from the lemma A.2.1, we have $q_s(k_0) > q_c$.

□

The cost function $y_s(q)$ can be viewed as a parametric optimization problem with parameter k. By the lemma A.2.1 the parametric optimization problem satisfies the Berge Maximum Theorem and with the lemma A.2.2 the optimal semi-centralized order quantity $q_s(k)$ is a continuous function in k on the interval $[0, q_s(k_0)]$. Then by the Intermediate Value Theorem there must exists a $k = k_s$ such that $q_s(k_s) = q_c$.

The proof of Theorem 3.7.2 is similar. The optimal centralized order cycle τ_c is the minimizer of

$$y_c^{(2)}(\tau, s_*, n) \quad = \quad \frac{1}{\tau}(k_0 + c_0 u) + nh\left(\psi(s_*, \tau, u) - \lambda l - \frac{1}{2}\lambda\tau\right) \tag{A.74}$$

and the optimal semi-centralized order cycle τ_s is the minimizer of

$$y_s^{(2)}(\tau, s_*, u) \quad = \quad \frac{1}{\tau}(k + cu) + h\left(\psi(s_*, \tau, u) - \lambda l - \frac{1}{2}\lambda\tau\right) \tag{A.75}$$

The key is to show the function $\psi(s_*, \tau, u)$ is continuous under demand process D2 Brownian motion with drift. With demand process D2, the starting inventory position at the beginning of an order cycle S_t forms a continuous state space Markov chain on $(-\infty, s]$, with state transition the same as $S_{t+1} = \min\{s, S_t - D(\tau) + u\}$. Under demand process D2, $D(\tau)$ is a normal random variable with mean $\lambda\tau$ and variance $\sigma_D^2\tau$. It can be verified that if the reserved capacity $u > \lambda\tau$, the Markov chain is recurrent and the stationary distribution S_∞ can be obtained, which is a continuous random variable on $(-\infty, s]$. In consequence, $\psi(s_*, \tau, u)$, the expectation of S_∞, is a continuous function in τ.

□

A.3 Proofs of Results Appearing in Chapter 4

A.3.1 Proof of Algorithm 4.5.1

Algorithm 4.5.1: Algorithm to compute E$[\Phi(\mathcal{I}, 1, T)]$

Step 1: Initialization

$$\mathbf{q}_1 = \left[\sum_{i=0}^{c} p_i \; p_{c+1} \; p_{c+2} \ldots \; p_{2c} \right]^\top \tag{A.76}$$

$$\mathbf{P}_1 = \begin{bmatrix} \sum_{i=0}^{c} p_i & \sum_{i=0}^{c-1} p_i & \sum_{i=0}^{c-2} p_i & \cdots & p_0 \\ p_{c+1} & p_c & p_{c-1} & \cdots & p_1 \\ p_{c+2} & p_{c+1} & p_c & \cdots & p_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ p_{2c} & p_{2c-1} & p_{2c-2} & \cdots & p_c \\ 0 & p_{2c} & p_{2c-1} & \cdots & p_{c+1} \\ 0 & 0 & p_{2c} & \cdots & p_{c+2} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & p_{2c} \end{bmatrix} \tag{A.77}$$

Step 2: Recursive

$$\mathbf{E}[\Phi(\mathcal{I}, 1, t+1)] = \mathbf{E}[\Phi(\mathcal{I}, 1, t)] + c\,\mathbf{e}^\top \overline{\mathbf{p}} + \underline{\mathbf{p}}^\top \mathbf{C}\tilde{\mathbf{q}}_\mathbf{t} \tag{A.78}$$

$$\mathbf{q}_{\mathbf{t+1}} = \mathbf{P}_\mathbf{t}\mathbf{q}_\mathbf{t} \tag{A.79}$$

$$\tag{A.80}$$

where

Q^t: Total unmet demand of $\Phi(\mathcal{I}, 1, T)$. $Q^{t+1} = (Q^t + D - c)^+$

q_k^t: Probability of $Q^t = k$

$\mathbf{q}_\mathbf{t}$: $(tc+1) \times 1$ vector represents the p.m.f of Q^t

$= \left[\, q_0^t \; q_1^t \; q_2^t \; \ldots \; q_{tc-1}^t \; q_{tc}^t \, \right]^\top$

$\mathbf{P_t}$ $\quad ((t+1)c+1) \times (tc+1)$ matrix

$$
= \left[\begin{array}{ccccccccccccccc}
\sum_{i=0}^{c} p_i & \sum_{i=0}^{c-1} p_i & \sum_{i=0}^{c-2} p_i & \cdots & p_0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\
p_{c+1} & p_c & p_{c-1} & \cdots & p_1 & p_0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\
p_{c+2} & p_{c+1} & p_c & \cdots & p_2 & p_1 & p_0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
p_{2c} & p_{2c-1} & p_{2c-2} & \cdots & p_c & p_{c-1} & p_{c-1} & \cdots & p_0 & 0 & 0 & \cdots & 0 & 0 & 0 \\
0 & p_{2c} & p_{2c-1} & \cdots & p_{c+1} & p_c & p_{c-1} & \cdots & p_1 & p_0 & 0 & \cdots & 0 & 0 & 0 \\
0 & 0 & p_{2c} & \cdots & p_{c+2} & p_{c+1} & p_c & \cdots & p_2 & p_1 & p_0 & \cdots & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & p_2 & p_1 & p_0 \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & p_3 & p_2 & p_1 \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & p_4 & p_3 & p_2 \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & p_{c+2} & p_{c+1} & p_c \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & p_{c+3} & p_{c+2} & p_{c+1} \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & p_{c+4} & p_{c+3} & p_{c+2} \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & p_{2c} & p_{2c-1} & p_{2c-2} \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & p_{2c} & p_{2c-1} \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & p_{2c}
\end{array}\right]
$$

$$\underline{\mathbf{p}} = [\, p_0 \; p_1 \; \ldots \; p_c \,]^\top$$

$$\overline{\mathbf{p}} = [\, p_{c+1} \; p_{c+2} \; \ldots \; p_{2c} \,]^\top$$

$$\mathbf{e} = [\, 1 \; 1 \; \ldots \; 1 \,]^\top$$

$$\tilde{\mathbf{q}} = \left[q_0^t \; q_1^t \; \ldots \; q_{c-1}^t \; \sum_{i=c}^{tc} q_i^t \right]^\top$$

$$
\mathbf{C} = \left[\begin{array}{cccccc}
0 & 1 & \cdots & c-2 & c-1 & c \\
1 & 2 & \cdots & c-1 & c & c \\
2 & 3 & \cdots & c & c & c \\
\vdots & \vdots & \vdots & \vdots & \vdots & \\
c-1 & c & \cdots & c & c & c \\
c & c & \cdots & c & c & c
\end{array}\right]
$$

Complexity: Complexity of computing $\mathbf{E}[\Phi(\mathcal{I}, 1, T)]$ is $O(T^2c^2)$.

Consider $t = 1$, by the definition of Q^t, clearly $\mathbf{P}(Q^1 = 0) = \mathbf{P}(D^1 \leq c) = \sum_{i=0}^{c} p_i$ and $\mathbf{P}(Q^1 = k) = \mathbf{P}(D^1 = c + k) = p_{c+k}$ for $k = 1, \ldots, c$, then we have $\mathbf{q_1}$ as defined in the algorithm.

Now consider the p.m.f. of Q^{t+1}. By definition $Q^{t+1} = (Q^t + D^t - c)^+$. It can be easily verified that $\mathbf{q_{t+1}} = \mathbf{P_t q_t}$ with $\mathbf{P_t}$ defined in the algorithm.

Given $\mathbf{q_t}$, to compute $\mathbf{E}[\min\{c, Q^t + D\}]$, if $D > c$ we have $\min\{c, Q^t + D\} = c$ therefore

$$\mathbf{E}[\min\{c, Q^t + D\}|D > c]\,\mathbf{P}(D > c) = c\mathbf{e}^\top \overline{\mathbf{p}}$$

If $D \leq c$, it is easy to check that

$$\mathbf{E}[\min\{c, Q^t + D\}|D \leq c]\,\mathbf{P}(D \leq c) = \underline{\mathbf{p}}^\top \mathbf{C}\tilde{\mathbf{q}}_\mathbf{t}$$

The complexity of computing $\mathbf{q_t}$ is $O(Tc^2)$ because $\mathbf{P_t}$ is sparse. The complexity of computing $\underline{\mathbf{p}}^\top \mathbf{C}\tilde{\mathbf{q}}_\mathbf{t}$ is $O(c^2)$. Overall, recurring T times the complexity is $O(T^2c^2)$.

□

A.3.2 Proof of Algorithm 4.5.2

Algorithm 4.5.2: Algorithm to compute $\mathbf{E}[\tilde{\Phi}(\mathcal{I}, 1, T)]$

Step 1: Determine $\boldsymbol{\gamma}$ and $\boldsymbol{\delta}$ (See Algorithm 4.5.3)

Step 2: Initialization

$$\mathbf{q_1} = [\ \delta_0\ \delta_1\ \delta_2 \ldots\ \delta_c\ ]^\top \tag{A.81}$$

$$\tilde{\mathbf{P}}_1 = \begin{bmatrix} \delta_0 & \sum_{i=1}^{c}\gamma_i & \sum_{i=2}^{c}\gamma_i & \cdots & \gamma_c \\ \delta_1 & \zeta_0 & \gamma_1 & \cdots & \gamma_{c-1} \\ \delta_2 & \delta_1 & \zeta_0 & \cdots & \gamma_{c-2} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \delta_c & \delta_{c-1} & \delta_{c-2} & \cdots & \zeta_0 \\ 0 & \delta_c & \delta_{c-1} & \cdots & \delta_1 \\ 0 & 0 & \delta_c & \cdots & \delta_2 \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \cdots & \delta_c \end{bmatrix} \tag{A.82}$$

Step 3: Recursive

$$\mathbf{E}[\tilde{\Phi}(\mathcal{I}, 1, t+1)] = \mathbf{E}[\tilde{\Phi}(\mathcal{I}, 1, t)] + \boldsymbol{\gamma}^\top \tilde{\mathbf{C}}\tilde{\mathbf{q}}_\mathbf{t} \tag{A.83}$$

$$\mathbf{q_{t+1}} = \tilde{\mathbf{P}}_\mathbf{t}^\top \mathbf{q_t} \tag{A.84}$$

$$\tag{A.85}$$

where

Q^t: Total unmet demand of $\tilde{\Phi}(\mathcal{I}, 1, T)$. $Q^{t+1} = (Q^t + \tilde{D} - \tilde{C})^+$

q_k^t: Probability of $Q^t = k$

$\mathbf{q_t}$: $(tc \times 1)$ vector represents the p.m.f. of Q^t

$$= \left[\, q_0^t \; q_1^t \; q_2^t \; \ldots \; q_{tc-1}^t \; q_{tc}^t \,\right]^\top$$

$$\boldsymbol{\gamma} = \left[\, \gamma_0 \; \gamma_1 \; \ldots \; \gamma_c \,\right]^\top$$

$$\boldsymbol{\delta} = \left[\, \delta_0 \; \delta_1 \; \ldots \; \delta_c \,\right]^\top$$

$\tilde{\mathbf{P}}_\mathbf{t}$: $((t+1)c+1) \times (tc+1)$ matrix

$$= \begin{bmatrix}
\delta_0 & \sum_{i=1}^{c} \gamma_i & \sum_{i=2}^{c} \gamma_i & \cdots & \gamma_c & 0 & 0 & \cdots & 0 & 0 & 0 \\
\delta_1 & \zeta_0 & \gamma_1 & \cdots & \gamma_{c-1} & \gamma_c & 0 & \cdots & 0 & 0 & 0 \\
\delta_2 & \delta_1 & \zeta_0 & \cdots & \gamma_{c-2} & \gamma_{c-1} & \gamma_c & \cdots & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \\
\delta_c & \delta_{c-1} & \delta_{c-2} & \cdots & \zeta_0 & \gamma_1 & \gamma_2 & \cdots & 0 & 0 & 0 \\
0 & \delta_c & \delta_{c-1} & \cdots & \delta_1 & \zeta_0 & \gamma_1 & \cdots & 0 & 0 & 0 \\
0 & 0 & \delta_c & \cdots & \delta_2 & \delta_1 & \zeta_0 & \cdots & 0 & 0 & 0 \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
0 & 0 & 0 & \cdots & \delta_c & \delta_{c-1} & \delta_{c-2} & \cdots & \delta_2 & \delta_1 & \zeta_0 \\
0 & 0 & 0 & \cdots & 0 & \delta_c & \delta_{c-1} & \cdots & \delta_3 & \delta_2 & \delta_1 \\
0 & 0 & 0 & \cdots & 0 & 0 & \delta_c & \cdots & \delta_4 & \delta_3 & \delta_2 \\
\vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & \delta_c & \delta_{c-1} & \delta_{c-2} \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & \delta_c & \delta_{c-1} \\
0 & 0 & 0 & \cdots & 0 & 0 & 0 & \cdots & 0 & 0 & \delta_c
\end{bmatrix}$$

$$\tilde{\mathbf{C}} = \begin{bmatrix}
0 & 0 & 0 & \cdots & 0 & 0 \\
0 & 1 & 1 & \cdots & 1 & 1 \\
0 & 1 & 2 & \cdots & 2 & 2 \\
\vdots & \vdots & \vdots & \vdots & \vdots & \\
0 & 1 & 2 & \cdots & c-1 & c-1 \\
0 & 1 & 2 & \cdots & c-1 & c
\end{bmatrix}$$

Complexity: The complexity of computing $\mathbf{E}[\tilde{\Phi}(\mathcal{I}, 1, T)]$ is $O(T^2 c^2)$.

The proof is very similar to the proof of Algorithm 4.5.1, with constant capacity c replaced by $\tilde{C}$ and original random demand D replaced by $\tilde{D}$.

□

A.3.3 Proof of Algorithm 4.5.3

Algorithm 4.5.3: Algorithm to compute $\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)]$

Initialization:

$$\boldsymbol{\gamma}_0 = [\,1\,0\,0 \ldots\ 0\,]^\top \tag{A.86}$$

Step 1: Compute the probability of unused capacity at plant N and unmet demand of product N: $\boldsymbol{\gamma}_N, \boldsymbol{\delta}_N$ recursively as

$$\boldsymbol{\gamma}_N = \mathbf{A}\boldsymbol{\gamma}_{N-1} \tag{A.87}$$
$$\boldsymbol{\delta}_N = \mathbf{B}\boldsymbol{\gamma}_{N-1} \tag{A.88}$$

where

$$\mathbf{A} = \begin{bmatrix} \sum_{i=c}^{2c} p_i & \sum_{i=c+1}^{2c} p_i & \sum_{i=c+2}^{2c} p_i & \cdots & p_{2c} \\ p_{c-1} & p_c & p_{c+1} & \cdots & p_{2c-1} \\ p_{c-2} & p_{c-1} & p_c & \cdots & p_{2c-2} \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ p_0 & \sum_{i=0}^{1} p_i & \sum_{i=0}^{2} p_i & \cdots & \sum_{i=0}^{c} p_i \end{bmatrix}$$

$$\mathbf{B} = \begin{bmatrix} \sum_{i=0}^{c} p_i & \sum_{i=0}^{c+1} p_i & \cdots & \sum_{i=0}^{2c-2} p_i & \sum_{i=0}^{2c-1} p_i & 1 \\ p_{c+1} & p_{c+2} & \cdots & p_{2c-1} & p_{2c} & 0 \\ p_{c+2} & p_{c+3} & \cdots & p_{2c} & 0 & 0 \\ p_{c+3} & p_{c+4} & \cdots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ p_{2c} & 0 & \cdots & 0 & 0 & 0 \end{bmatrix}$$

Step 2: Compute the performance of open chain without inventory $\mathbf{E}[\Phi(\mathcal{L}_N, \emptyset, N, T)]$ recursively as

$$\mathbf{E}[\Phi(\mathcal{L}_N, \emptyset, N, 1)] = \mathbf{E}[\Phi(\mathcal{L}_{N-1}, \emptyset, N-1, 1)] + \mathbf{E}[\min\{c + \tilde{C}_{N-1}, D\}] \tag{A.89}$$
$$= \mathbf{E}[\Phi(\mathcal{L}_{N-1}, \emptyset, N-1, 1)] + \boldsymbol{\gamma}_{N-1}^\top \mathbf{K}\mathbf{p} \tag{A.90}$$

where

$$\mathbf{K} = \begin{bmatrix} 0 & 1 & \cdots & c & c & c & \cdots & c \\ 0 & 1 & \cdots & c & c+1 & c+1 & \cdots & c+1 \\ 0 & 1 & \cdots & c & c+1 & c+2 & \cdots & c+2 \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\ 0 & 1 & \cdots & c & c+1 & c+2 & \cdots & 2c \end{bmatrix}$$

$$\mathbf{p} \;=\; [\, p_0\, p_1\; \ldots\; p_{2c}\,]^\top$$

Step 3: Compute residual performance $\mathbf{E}[\tilde{\Phi}(\mathcal{L}_N, \mathcal{I}_N, N, T)]$ with $\boldsymbol{\gamma} = \boldsymbol{\gamma}_N, \boldsymbol{\delta} = \boldsymbol{\delta}_N$ using Algorithm 4.5.2.

Step 4: $\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)] = \mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \emptyset, N, T)] + \mathbf{E}[\tilde{\Phi}(\mathcal{L}_{1\to N}, \mathcal{I}_N, 1, T)]$

Complexity: The complexity of computing $\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)]$ is $O(Nc^2 + T^2c^2)$

We first show the residual performance $\tilde{\Phi}$ is exactly the one defined in Section 4.5.2.

$\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T; d)$ can be formulated as a maximum flow problem on graph $\mathcal{G}$ as shown in Figure A.1. Arcs (s_0, v_{it}) have capacity c. Arcs (v_{it}, w_{js}) have capacity ∞. Arcs (w_{js}, s_∞) have capacity d_{js}. Given a realization of demand d_{js}, $\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T; d)$ is the maximum flow from s_0 to s_∞.

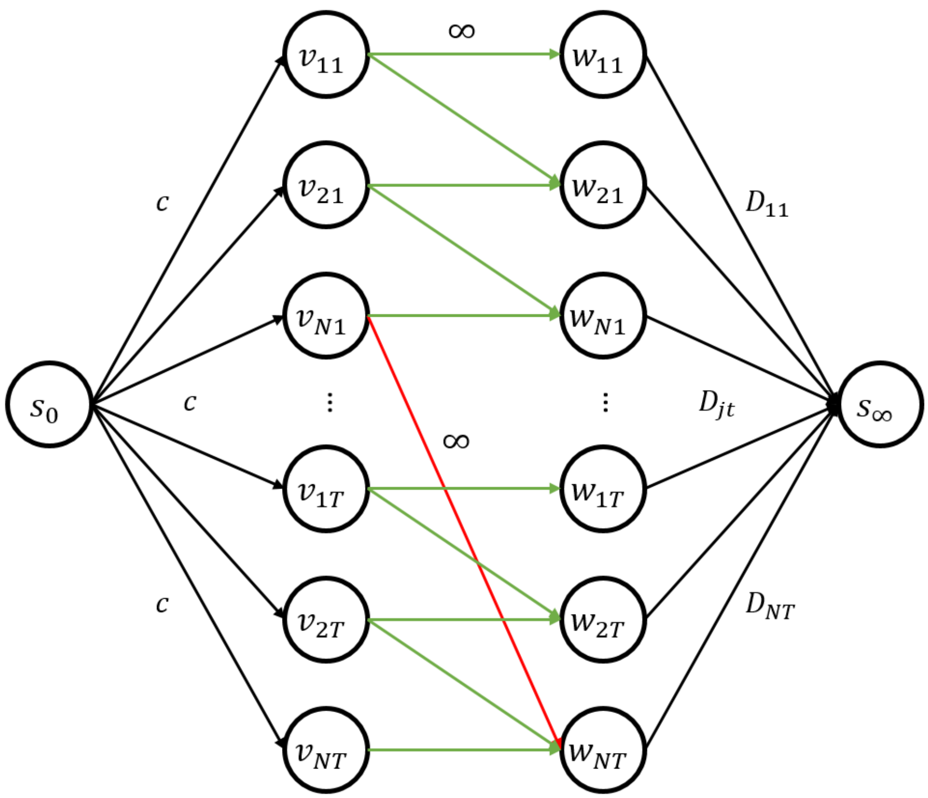

Figure A.1: Maximum flow formulation of $\Phi(\mathcal{L}_N, \mathcal{I}_1, N, T)$

First find a maximum flow from s_0 to s_∞ without the red arcs. The solution must be optimal to $\Phi(\mathcal{L}_{1\to N}, \emptyset, N, T; d)$, which can be viewed as T subproblems of $\Phi(\mathcal{L}_{1\to N}, \emptyset, N, 1; d_t)$ where $d_t = \{d_{jt}, \forall j\}$. Let the optimal solution of each subproblem denote as $z(\mathcal{L}_{1\to N}, \emptyset, N, 1; d_t)$, or in short z^t. Let $\tilde{c}_{Nt}$ be the unused capacity of plant N in subproblem t and let $\tilde{d}_{Nt}$ be the unmet demand of product N in subproblem t under demand scenario d. If there are multiple optimal solutions z^t for one subproblem, we focus on the solutions that maximize the unused capacity $\tilde{c}_{Nt}$ or maximize the unmet demand $\tilde{d}_{Nt}$. These two sub-objectives are not contradict against each other because in any optimal solution z^t, $\tilde{c}_{Nt}$ and $\tilde{d}_{Nt}$ can't be both strictly positive.

Next, consider adding red arcs, which are the inventory arcs between plant N and product N. Clearly the set of solutions $z = \{I_t, \forall t\}$ is feasible on $\mathcal{G}$. Now find augmenting path on the residual graph $\mathcal{G}_z$. Notice there is no augmenting path using arcs (w_{jt}, s_∞) for $j \neq N$, otherwise z^t cannot be optimal for subproblem $\Phi(\mathcal{L}_{1\to N}, \emptyset, N, 1; d_t)$. Also, there is no augmenting path using arcs (v_{Nt}, w_{Nt}) or (w_{Nt}, v_{Nt}) because the particular z^t we choose (i.e., maximize $\tilde{c}_{Nt}$ or $\tilde{d}_{Nt}$). Since v_{Nt} is only connected with w_{Ns} for $s \geq t$, augmenting paths can only use arcs $(s_0, v_{Nt}), (w_{Nt}, s_\infty)$ and (v_{Nt_1}, w_{Nt_2}) for $t_1 < t_2$, which is equivalent of solving

$$\max \quad \sum_{t_1<t_2} f_{v_{Nt_1}, w_{Nt_2}} \tag{A.91}$$

$$\text{s.t.} \quad \sum_{t_2>t} f_{v_{Nt}, w_{Nt_2}} \leq \tilde{c}_{Nt} \quad \forall t \tag{A.92}$$

$$\sum_{t_1<t} f_{v_{Nt_1}, w_{Nt}} \leq \tilde{d}_{Nt} \quad \forall t \tag{A.93}$$

$$f_{v_{Nt_1}, w_{Nt_2}} \geq 0 \tag{A.94}$$

By easy transformation, the above problem is equivalent to the residual problem $\tilde{\Phi}(\emptyset, \mathcal{L}_1, 1, T; \tilde{c}, \tilde{d})$ defined in Section 4.5.2 (See Figure A.2).

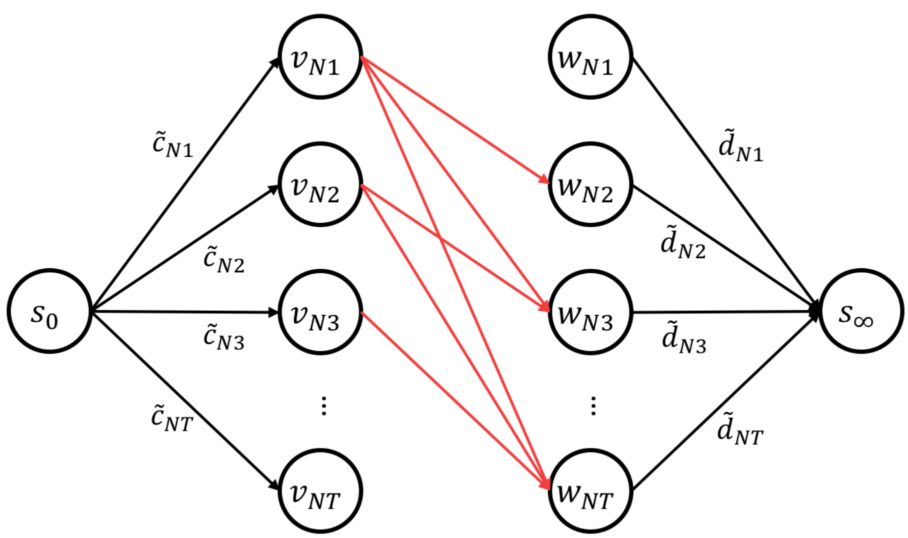

Figure A.2: Finding augmenting paths is equivalent to $\tilde{\Phi}(\emptyset, \mathcal{I}_1, 1, T; \tilde{c}, \tilde{d})$

Therefore we have

$$\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T; d) = \Phi(\mathcal{L}_{1\to N}, \emptyset, N, T; d) + \tilde{\Phi}(\emptyset, \mathcal{I}_1, 1, T; \tilde{c}, \tilde{d}) \tag{A.95}$$

Since $\tilde{c}$ and $\tilde{d}$ is unique for a given d and A.95 holds for all $d \in \mathcal{D}$, taking expectation on both sides, we have

$$\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)] = \mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \emptyset, N, T)] + \mathbf{E}[\tilde{\Phi}(\emptyset, \mathcal{I}_1, 1, T)] \tag{A.96}$$

Next we'll show how to actually compute $\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)]$. First we need to show $\boldsymbol{\gamma} = \boldsymbol{\gamma}_N, \boldsymbol{\delta} = \boldsymbol{\delta}_N$ as in Step 1. Consider $N = 1$, then clearly $\boldsymbol{\gamma}_1 = [\sum_{i=c}^{2c} p_i, p_{c-1}, p_{c-2}, \ldots, p_0]$ and $\boldsymbol{\delta}_1 = [\sum_{i=0}^{c} p_i, p_{c+1}, p_{c+2}, \ldots, p_{2c}]$, then immediately we have $\boldsymbol{\gamma}_1 = \mathbf{A}\boldsymbol{\gamma}_0$ and $\boldsymbol{\delta}_1 = \mathbf{B}\boldsymbol{\gamma}_0$ as in the Step 1. To compute $\boldsymbol{\gamma}_N$ and $\boldsymbol{\delta}_N$ given $\boldsymbol{\gamma}_{N-1}$, by Simchi-Levi and Wei (2012) we have $\tilde{C}_N = (\tilde{C}_{N-1} + c - D)^+$ and $\tilde{D}_N = (D - \tilde{C}_{N-1} - c)^+$, it is easy to verify that $\boldsymbol{\gamma}_N = \mathbf{A}\boldsymbol{\gamma}_{N-1}$, $\boldsymbol{\delta}_N = \mathbf{B}\boldsymbol{\gamma}_{N-1}$ and $\mathbf{E}[\min\{\tilde{C}_{N-1} + c, D\}] = \boldsymbol{\gamma}_{N-1}^\top \mathbf{K}\mathbf{p}$ with $\mathbf{A}$, $\mathbf{B}$ and $\mathbf{K}$ given as in the Step 1 and 2.

The complexity of computing $\boldsymbol{\gamma}_N$ and $\boldsymbol{\delta}_N$ is $O(Nc^2)$. The complexity of computing $\mathbf{E}[\Phi(\mathcal{L}_N, \emptyset, N, 1)]$ is $O(c^2)$. Overall the complexity of computing $\mathbf{E}[\Phi(\mathcal{L}_{1\to N}, \mathcal{I}_N, N, T)]$ is $O(Nc^2 + T^2c^2)$.

□

A.3.4 Proof of Closed Form 4.5.4

Case 1: $1 \leq k \leq N-2$

Notice the residual performance $\tilde{\Phi}$ is closely related to the residual capacity in the first period and residual demand in the second period. To begin with, we introduce some simplified notations and several important lemmas. First ignore inventory arcs and let $\tilde{C}_i$ denote the residual capacity of plant i in the first period, and let

$$\mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix} \tag{A.97}$$

denote the joint probability of $\tilde{C}_i = b_i$ for all i from 1 to k given there are N plants in total. Since demand is Bernoulli, clearly b_i can only take on value 0 or 1 (b stands for binary). The Bernoulli demand gives us a very useful lemma

Lemma A.3.1. *Given an instance of demand $(d_{11}, \ldots, d_{N1})$ in the first period, there is a unique residual capacity vector $(b_1, \ldots, b_k)$*

Proof: It can be easily checked that for a closed chain, b_i can be determined in an unique way:
If $d_i = 2$ then $b_i = 0$
If $d_i = 0$ and $d_{i+1} = 0$, then $b_i = 1$
If $d_i = 0$ and $d_{i+1} = 2$, then $b_i = 0$

□

Lemma A.3.1 ensures the probability measure $\mathbf{P}^N$ is well defined: $\mathbf{P}^N$ is just the sum of the probabilities of all demand scenarios that yield the residual capacity $(b_1, \ldots, b_k)$. Here $\mathbf{P}^N$ says nothing about the inventory capability of plant 1 to k. Plant i may or may not be able to hold inventory.

In comparison the residual demand in the second period given an instance of demand is not unique. For example, suppose the demand in the second period is $(d_{12} = 0, d_{22} = 2, d_{32} = 2)$, then the residual demand can be either $(0, 1, 0)$ or $(0, 0, 1)$. Therefore analyzing the residual demand in the second period independent of the first period does not help. To circumvent this issue, we define a *conditional residual performance* π

$$\pi^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix} \tag{A.98}$$

as the residual performance on $(\mathcal{C}, I_{1\to k}, N, 2)$ given the residual capacity in the first period is $(b_1, \ldots, b_k)$, i.e. $\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2) \,|\, (\tilde{C}_1, \ldots, \tilde{C}_k) = (b_1, \ldots, b_k)]$. By the definition of $\tilde{\Phi}$, π can be treated as the difference between two parts:

$$\pi^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix} \quad = \quad \mathbf{E}[\mathcal{C}_N^{'}] - \mathbf{E}[\mathcal{C}_N]$$

where $\mathbf{E}[\mathcal{C}_N]$ is the performance of a single-period closed chain of size N, and $\mathbf{E}[\mathcal{C}_N^{'}]$ is the performance of a single-period modified closed chain with

- Plant i from 1 to k have capacity $1 + b_i$, and b_i can only be used to satisfied the demand of product i.
- Plants i from $k + 1$ to N have capacity 1.

Let $\tilde{\mathbf{C}}$ be the vector of $(\tilde{C}_1, \ldots, \tilde{C}_k)$ and $\mathbf{b}$ be the vector of $(b_1, \ldots, b_k)$. For the case of $k \leq N-2$, we have the following lemma

Lemma A.3.2. *For* $1 \leq k \leq N - 1$,

$$\mathbf{P}^N(\tilde{\mathbf{C}} = \mathbf{b}) = \mathbf{P}^{N+l}(\tilde{\mathbf{C}} = \mathbf{b})$$

for $l = 0, 1, 2, \ldots$

For $1 \leq k \leq N - 2$,

$$\pi^N(\tilde{\mathbf{C}} = \mathbf{b}) = \pi^{N+l}(\tilde{\mathbf{C}} = \mathbf{b})$$

for $l = 0, 1, 2, \ldots$

Proof: The first part is true by Lemma A.3.1 because given a demand instance in the second period $(d_{12}, \ldots, d_{N2})$, $\mathbf{P}^N$ only depends on the demand from d_{12} up to $d_{k+1,2}$.

In the second part, it can be easily verified that given a demand instance in the second period $(d_{12}, \ldots, d_{N2})$, whether the residual capacity b_i is used only depends on the demand $(d_{i-1,2}, d_{i,2}, d_{i+1,2})$. Then $\pi^N(b_1, \ldots, b_k)$ can be purely determined by $k+2$ demand: d_1 to d_{k+1} and d_N.

□

Lemma A.3.2 shows another merit of Bernoulli demand, which says the benefit of an inventory arc at plant i only affects the product $i-1, i$, and $i+1$. Here the benefit means the residual capacity b_i is used to satisfy the demand which can not be met without b_i. In other words, as long as $N - k \geq 2$, which it is the case in this subsection, $\mathbf{P}^N$ and π^N is independent of N, so we can drop the superscript N.

In the next lemma, we will show several useful properties of $\mathbf{P}$ and π. In the lemma, if there is no further clarification, we always require the subscript of $\tilde{C}_i$ is between 1 and $N-2$.

Lemma A.3.3. *Some useful properties of $\mathbf{P}$ and π:*

(1) No single zero between two ones:

$$\mathbf{P}\begin{pmatrix} \tilde{C}_{i-1} & = & 1 \\ \tilde{C}_i & = & 0 \\ \tilde{C}_{i+1} & = & 1 \end{pmatrix} = 0$$

(2) Starting point invariant

$$\mathbf{P}\begin{pmatrix} \tilde{C}_i & = & b_i \\ \tilde{C}_{i+1} & = & b_{i+1} \\ & \vdots & \\ \tilde{C}_{i+k} & = & b_{i+k} \end{pmatrix} = \mathbf{P}\begin{pmatrix} \tilde{C}_j & = & b_i \\ \tilde{C}_{j+1} & = & b_{i+1} \\ & \vdots & \\ \tilde{C}_{j+k} & = & b_{i+k} \end{pmatrix}$$

(3) Add a row

$$\mathbf{P}\begin{pmatrix} \tilde{C}_i & = & b_i \\ \tilde{C}_{i+1} & = & b_{i+1} \\ & \vdots & \\ \tilde{C}_{i+k} & = & b_{i+k} \end{pmatrix} = \mathbf{P}\begin{pmatrix} \tilde{C}_i & = & b_i \\ \tilde{C}_{i+1} & = & b_{i+1} \\ & \vdots & \\ \tilde{C}_{i+k} & = & b_{i+k} \\ \tilde{C}_{i+k+1} & = & 0 \end{pmatrix} + \mathbf{P}\begin{pmatrix} \tilde{C}_i & = & b_i \\ \tilde{C}_{i+1} & = & b_{i+1} \\ & \vdots & \\ \tilde{C}_{i+k} & = & b_{i+k} \\ \tilde{C}_{i+k+1} & = & 1 \end{pmatrix}$$

$$= \mathbf{P}\begin{pmatrix} \tilde{C}_{i-1} & = & 0 \\ \tilde{C}_i & = & b_i \\ \tilde{C}_{i+1} & = & b_{i+k+1} \\ & \vdots & \\ \tilde{C}_{i+k} & = & b_k \end{pmatrix} + \mathbf{P}\begin{pmatrix} \tilde{C}_{i-1} & = & 1 \\ \tilde{C}_i & = & b_i \\ \tilde{C}_{i+1} & = & b_{i+k+1} \\ & \vdots & \\ \tilde{C}_{i+k} & = & b_k \end{pmatrix}$$

(4) Redundant zeros at two ends

$$\pi\begin{pmatrix} \tilde{C}_i & = & b_i \\ \tilde{C}_{i+1} & = & b_{i+1} \\ & \vdots & \\ \tilde{C}_{i+k-1} & = & b_{i+k-1} \\ \tilde{C}_{i+k} & = & 0 \end{pmatrix} = \pi\begin{pmatrix} \tilde{C}_i & = & b_i \\ \tilde{C}_{i+1} & = & b_{i+1} \\ & \vdots & \\ \tilde{C}_{i+k-1} & = & b_{i+k-1} \end{pmatrix}$$

(5) Split

$$\pi\begin{pmatrix} \tilde{C}_i & = & b_i \\ \tilde{C}_{i+1} & = & b_{i+1} \\ & \vdots & \\ \tilde{C}_{i+j} & = & 0 \\ \tilde{C}_{i+j+1} & = & 0 \\ & \vdots & \\ \tilde{C}_{i+k} & = & b_{i+k} \end{pmatrix} = \pi\begin{pmatrix} \tilde{C}_i & = & b_i \\ \tilde{C}_{i+1} & = & b_{i+1} \\ & \vdots & \\ \tilde{C}_{i+j-1} & = & b_{i+j-1} \end{pmatrix} + \pi\begin{pmatrix} \tilde{C}_{i+j+2} & = & b_{i+j+2} \\ & \vdots & \\ \tilde{C}_{i+k} & = & b_k \end{pmatrix}$$

Proof:

(1) $b_{i-1} = 1$ requires $d_{i-1} = d_i = 0$. $b_{i+1} = 1$ requires $d_{i+1} = d_{i+2} = 0$. Then since $d_i = d_{i+1} = 0$, b_i must be 1 as well.

(2) This is true because the closed chain is rotational symmetric.

(3) This is true because of the definition of the marginal probability distribution

(4) Residual capacity of plant $i+k$ is 0 is equivalent to delete the corresponding inventory arc at plant $i+k$.

(5) This is true by Lemma A.3.2. Since $\tilde{C}_{i+j}$ and $\tilde{C}_{i+j+1}$ are both 0, we can treat there is not inventory arc at plant $i+j$ and $i+j+1$. The benefit of inventory arcs at plant i to $i+j-1$ only affect products $i-1$ to $i+j$. The benefit of inventory arcs at plant $i+j+2$ to $i+k$ only affects products $i+j+1$ to $i+k+1$, which is independent of the products $i-1$ to $i+j$.

□

To simplify the notation of $\mathbf{E}[\tilde{\Phi}(\mathcal{C}, I_{1\to k}, N, 2)]$, define $f(k)$ as the unconditional residual performance of a network with k inventory arcs, then

$$f(k) \quad = \quad \sum_{b_1} \cdots \sum_{b_k} \mathbf{P}\begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix} \pi \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix}$$

Now we show $f(k)$ follows a linear recurrence relation.

Lemma A.3.4. *Linear recurrence relation between $f(k+1), f(k)$ and $f(k-1)$*

$$f(k+1) - 2f(k) + f(k-1)$$

$$= \mathbf{P}\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{k+1} & = & 1 \end{pmatrix} \left[\pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{k+1} & = & 1 \end{pmatrix} - 2\,\pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix} + \pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{k-1} & = & 1 \end{pmatrix} \right]$$

Proof: We first manipulate $f(k)$

$$
\begin{aligned}
f(k) &= \sum_{b_1}\cdots\sum_{b_k}\mathbf{P}\begin{pmatrix}\tilde{C}_1 &=& b_1\\ \tilde{C}_2 &=& b_2\\ &\vdots& \\ \tilde{C}_k &=& b_k\end{pmatrix}\pi\begin{pmatrix}\tilde{C}_1 &=& b_1\\ \tilde{C}_2 &=& b_2\\ &\vdots& \\ \tilde{C}_k &=& b_k\end{pmatrix}\\
&= \sum_{b_1}\cdots\sum_{b_k}\left[\mathbf{P}\begin{pmatrix}\tilde{C}_1 &=& 0\\ \tilde{C}_2 &=& b_1\\ &\vdots& \\ \tilde{C}_{k+1} &=& b_k\end{pmatrix}+\mathbf{P}\begin{pmatrix}\tilde{C}_1 &=& 1\\ \tilde{C}_2 &=& b_1\\ &\vdots& \\ \tilde{C}_{k+1} &=& b_k\end{pmatrix}\right]\pi\begin{pmatrix}\tilde{C}_1 &=& b_1\\ \tilde{C}_2 &=& b_2\\ &\vdots& \\ \tilde{C}_k &=& b_k\end{pmatrix}\\
&= \sum_{b_2}\cdots\sum_{b_{k+1}}\mathbf{P}\begin{pmatrix}\tilde{C}_1 &=& 0\\ \tilde{C}_2 &=& b_2\\ &\vdots& \\ \tilde{C}_{k+1} &=& b_{k+1}\end{pmatrix}\pi\begin{pmatrix}\tilde{C}_1 &=& b_2\\ \tilde{C}_2 &=& b_3\\ &\vdots& \\ \tilde{C}_k &=& b_{k+1}\end{pmatrix}+\\
&\quad \sum_{b_2}\cdots\sum_{b_{k+1}}\mathbf{P}\begin{pmatrix}\tilde{C}_1 &=& 1\\ \tilde{C}_2 &=& b_2\\ &\vdots& \\ \tilde{C}_{k+1} &=& b_{k+1}\end{pmatrix}\pi\begin{pmatrix}\tilde{C}_1 &=& b_2\\ \tilde{C}_2 &=& b_3\\ &\vdots& \\ \tilde{C}_k &=& b_{k+1}\end{pmatrix}
\end{aligned}
$$

where the first equality holds because of the Properties (3), the second equality holds because of the Properties (4) and in the third equality we just change the name of $b_1, \ldots, b_k$ to $b_2, \ldots, b_{k+1}$.

For $f(k+1)$, expanding along b_1 we have

$$
\begin{aligned}
f(k+1) &= \sum_{b_2}\cdots\sum_{b_{k+1}} \mathbf{P}\begin{pmatrix} \tilde{C}_1 &=& 0 \\ \tilde{C}_2 &=& b_2 \\ & \vdots & \\ \tilde{C}_{k+1} &=& b_{k+1} \end{pmatrix} \pi \begin{pmatrix} \tilde{C}_1 &=& 0 \\ \tilde{C}_2 &=& b_2 \\ & \vdots & \\ \tilde{C}_{k+1} &=& b_{k+1} \end{pmatrix} + \\
& \sum_{b_2}\cdots\sum_{b_{k+1}} \mathbf{P}\begin{pmatrix} \tilde{C}_1 &=& 1 \\ \tilde{C}_2 &=& b_2 \\ & \vdots & \\ \tilde{C}_{k+1} &=& b_{k+1} \end{pmatrix} \pi \begin{pmatrix} \tilde{C}_1 &=& 1 \\ \tilde{C}_2 &=& b_2 \\ & \vdots & \\ \tilde{C}_{k+1} &=& b_{k+1} \end{pmatrix} \\
&= \sum_{b_2}\cdots\sum_{b_{k+1}} \mathbf{P}\begin{pmatrix} \tilde{C}_1 &=& 0 \\ \tilde{C}_2 &=& b_2 \\ & \vdots & \\ \tilde{C}_{k+1} &=& b_{k+1} \end{pmatrix} \pi \begin{pmatrix} \tilde{C}_2 &=& b_2 \\ \tilde{C}_3 &=& b_3 \\ & \vdots & \\ \tilde{C}_{k+1} &=& b_{k+1} \end{pmatrix} + \\
& \sum_{b_2}\cdots\sum_{b_{k+1}} \mathbf{P}\begin{pmatrix} \tilde{C}_1 &=& 1 \\ \tilde{C}_2 &=& b_2 \\ & \vdots & \\ \tilde{C}_{k+1} &=& b_{k+1} \end{pmatrix} \pi \begin{pmatrix} \tilde{C}_1 &=& 1 \\ \tilde{C}_2 &=& b_2 \\ & \vdots & \\ \tilde{C}_{k+1} &=& b_{k+1} \end{pmatrix}
\end{aligned}
$$

where Property (4) is used in the equality.

By Property (2) we can see the first term in $f(k)$ and $f(k+1)$ can be canceled out, so we have

$$
\begin{aligned}
& f(k+1) - f(k) \\
= & \sum_{b_2}\cdots\sum_{b_{k+1}} \mathbf{P}\begin{pmatrix} \tilde{C}_1 &=& 1 \\ \tilde{C}_2 &=& b_2 \\ & \vdots & \\ \tilde{C}_{k+1} &=& b_{k+1} \end{pmatrix} \left[\pi \begin{pmatrix} \tilde{C}_1 &=& 1 \\ \tilde{C}_2 &=& b_2 \\ & \vdots & \\ \tilde{C}_{k+1} &=& b_{k+1} \end{pmatrix} - \pi \begin{pmatrix} \tilde{C}_1 &=& b_2 \\ \tilde{C}_2 &=& b_3 \\ & \vdots & \\ \tilde{C}_k &=& b_{k+1} \end{pmatrix} \right]
\end{aligned}
$$

Similarly we have

$$
\begin{aligned}
& f(k) - f(k-1) \\
= & \sum_{b_2}\cdots\sum_{b_k} \mathbf{P}\begin{pmatrix} \tilde{C}_1 &=& 1 \\ \tilde{C}_2 &=& b_2 \\ & \vdots & \\ \tilde{C}_k &=& b_k \end{pmatrix} \left[\pi \begin{pmatrix} \tilde{C}_1 &=& 1 \\ \tilde{C}_2 &=& b_2 \\ & \vdots & \\ \tilde{C}_k &=& b_k \end{pmatrix} - \pi \begin{pmatrix} \tilde{C}_1 &=& b_2 \\ \tilde{C}_2 &=& b_3 \\ & \vdots & \\ \tilde{C}_{k-1} &=& b_k \end{pmatrix} \right]
\end{aligned}
$$

Now we try to compute $f(k+1) - 2f(k) + f(k-1)$. We first manipulate $f(k+1) - f(k)$ by expanding along b_{k+1} such that

$$f(k+1) - f(k) = g_1 + g_2 + g_3$$

where g_1 is defined as

$$\begin{aligned}
g_1 &= \sum_{b_2} \cdots \sum_{b_k} \mathbf{P} \begin{pmatrix} \tilde{C}_1 &= 1 \\ \tilde{C}_2 &= b_2 \\ & \vdots \\ \tilde{C}_{k+1} &= 0 \end{pmatrix} \left[\pi \begin{pmatrix} \tilde{C}_1 &= 1 \\ \tilde{C}_2 &= b_2 \\ & \vdots \\ \tilde{C}_{k+1} &= 0 \end{pmatrix} - \pi \begin{pmatrix} \tilde{C}_1 &= b_2 \\ \tilde{C}_2 &= b_3 \\ & \vdots \\ \tilde{C}_k &= 0 \end{pmatrix} \right] \\
&= \sum_{b_2} \cdots \sum_{b_k} \mathbf{P} \begin{pmatrix} \tilde{C}_1 &= 1 \\ \tilde{C}_2 &= b_2 \\ & \vdots \\ \tilde{C}_{k+1} &= 0 \end{pmatrix} \left[\pi \begin{pmatrix} \tilde{C}_1 &= 1 \\ \tilde{C}_2 &= b_2 \\ & \vdots \\ \tilde{C}_k &= b_k \end{pmatrix} - \pi \begin{pmatrix} \tilde{C}_1 &= b_2 \\ \tilde{C}_2 &= b_3 \\ & \vdots \\ \tilde{C}_{k-1} &= b_{k-1} \end{pmatrix} \right]
\end{aligned}$$

where the equality holds because the Property (4). So g_1 represents the part that $b_{k+1} = 0$.

Then define g_2 as

$$g_2 = \sum_{b_2} \cdots \sum_{b_k} \mathbf{P} \begin{pmatrix} \tilde{C}_1 &= 1 \\ \tilde{C}_2 &= b_2 \\ & \vdots \\ \tilde{C}_j &= 0 \\ \tilde{C}_{j+1} &= 0 \\ & \vdots \\ \tilde{C}_{k+1} &= 1 \end{pmatrix} \left[\pi \begin{pmatrix} \tilde{C}_1 &= 1 \\ \tilde{C}_2 &= b_2 \\ & \vdots \\ \tilde{C}_j &= 0 \\ \tilde{C}_{j+1} &= 0 \\ & \vdots \\ \tilde{C}_{k+1} &= 1 \end{pmatrix} - \pi \begin{pmatrix} \tilde{C}_1 &= b_2 \\ \tilde{C}_2 &= b_3 \\ & \vdots \\ \tilde{C}_{j-1} &= 0 \\ \tilde{C}_j &= 0 \\ & \vdots \\ \tilde{C}_k &= 1 \end{pmatrix} \right]$$

g_2 represents the part that $b_{k+1} = 1$ and there are some $b_j = 0$ for $j = 2, \ldots, k$. By Property (1), we can assume W.L.O.G that $b_j = b_{j+1} = 0$ for some $j = 2, \ldots, k-1$. By the Property (5) the two π terms in g_2 can be split as

$$\pi \begin{pmatrix} \tilde{C}_1 &= 1 \\ \tilde{C}_2 &= b_2 \\ & \vdots \\ \tilde{C}_j &= 0 \\ \tilde{C}_{j+1} &= 0 \\ & \vdots \\ \tilde{C}_{k+1} &= 1 \end{pmatrix} = \pi \begin{pmatrix} \tilde{C}_1 &= 1 \\ \tilde{C}_2 &= b_2 \\ & \vdots \\ \tilde{C}_{j-1} &= b_{j-1} \end{pmatrix} + \pi \begin{pmatrix} \tilde{C}_{j+2} &= b_{j+2} \\ \tilde{C}_{j+3} &= b_{j+3} \\ & \vdots \\ \tilde{C}_{k+1} &= 1 \end{pmatrix}$$

and

$$\pi\begin{pmatrix} \tilde{C}_1 & = & b_2 \\ \tilde{C}_2 & = & b_3 \\ & \vdots & \\ \tilde{C}_{j-1} & = & 0 \\ \tilde{C}_j & = & 0 \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix} = \pi\begin{pmatrix} \tilde{C}_1 & = & b_2 \\ \tilde{C}_2 & = & b_3 \\ & \vdots & \\ \tilde{C}_{j-2} & = & b_{j-1} \end{pmatrix} + \pi\begin{pmatrix} \tilde{C}_{j+1} & = & b_{j+2} \\ \tilde{C}_{j+2} & = & b_{j+3} \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix}$$

$$= \pi\begin{pmatrix} \tilde{C}_1 & = & b_2 \\ \tilde{C}_2 & = & b_3 \\ & \vdots & \\ \tilde{C}_{j-2} & = & b_{j-1} \end{pmatrix} + \pi\begin{pmatrix} \tilde{C}_{j+2} & = & b_{j+2} \\ \tilde{C}_{j+3} & = & b_{j+3} \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix}$$

where the Property (2) is applied. By canceling out the second term we have

$$g_2 = \sum_{b_2}\cdots\sum_{b_k}\mathbf{P}\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_j & = & 0 \\ \tilde{C}_{j+1} & = & 0 \\ & \vdots & \\ \tilde{C}_{k+1} & = & 1 \end{pmatrix} \left[\pi\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_{j-1} & = & b_{j-1} \end{pmatrix} - \pi\begin{pmatrix} \tilde{C}_1 & = & b_2 \\ \tilde{C}_2 & = & b_3 \\ & \vdots & \\ \tilde{C}_{j-2} & = & b_{j-1} \end{pmatrix}\right]$$

Finally define g_3 as

$$g_3 = \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{k+1} & = & 1 \end{pmatrix} \left[\pi\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{k+1} & = & 1 \end{pmatrix} - \pi\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix}\right]$$

g_3 represents the part that all $b_i = 1$ for $i = 1, \ldots, k+1$.

Similary, manipulate $f(k) - f(k-1)$ by applying the Property (3) such that

$$f(k) - f(k-1) = h_1 + h_2 + h_3$$

Define h_1 as

$$h_1 = \sum_{b_2}\cdots\sum_{b_k}\mathbf{P}\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \\ \tilde{C}_{k+1} & = & 0 \end{pmatrix} \left[\pi\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix} - \pi\begin{pmatrix} \tilde{C}_1 & = & b_2 \\ \tilde{C}_2 & = & b_3 \\ & \vdots & \\ \tilde{C}_{k-1} & = & b_k \end{pmatrix}\right]$$

so h_1 represents the part that $b_{k+1} = 0$. Immediately we found $g_1 = h_1$.

Define h_2 as

$$h_2 = \sum_{b_2} \cdots \sum_{b_k} \mathbf{P} \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_j & = & 0 \\ \tilde{C}_{j+1} & = & 0 \\ & \vdots & \\ \tilde{C}_k & = & b_k \\ \tilde{C}_{k+1} & = & 1 \end{pmatrix} \left[\pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_j & = & 0 \\ \tilde{C}_{j+1} & = & 0 \\ & \vdots & \\ \tilde{C}_k & = & b_k \end{pmatrix} - \pi \begin{pmatrix} \tilde{C}_1 & = & b_2 \\ \tilde{C}_2 & = & b_3 \\ & \vdots & \\ \tilde{C}_{j-1} & = & 0 \\ \tilde{C}_j & = & 0 \\ & \vdots & \\ \tilde{C}_{k-1} & = & b_{k-1} \end{pmatrix} \right]$$

h_2 represents the part that $b_{k+1} = 1$ with $b_j = b_{j+1} = 0$. Similar to g_2, we can split h_2 and since the tail with b_{j+2} to b_k involved can be canceled out, we have

$$h_2 = \sum_{b_2} \cdots \sum_{b_k} \mathbf{P} \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_j & = & 0 \\ \tilde{C}_{j+1} & = & 0 \\ & \vdots & \\ \tilde{C}_k & = & b_k \\ \tilde{C}_{k+1} & = & 1 \end{pmatrix} \left[\pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_{j-1} & = & b_{j-1} \end{pmatrix} - \pi \begin{pmatrix} \tilde{C}_1 & = & b_2 \\ \tilde{C}_2 & = & b_3 \\ & \vdots & \\ \tilde{C}_{j-2} & = & b_{j-1} \end{pmatrix} \right]$$

and again we have $g_2 = h_2$.

Finally define h_3 as

$$h_3 = \mathbf{P} \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_k & = & 1 \\ \tilde{C}_{k+1} & = & 1 \end{pmatrix} \left[\pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix} - \pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{k-1} & = & 1 \end{pmatrix} \right]$$

h_3 represents the part that all $b_i = 1$ for $i = 1, \ldots, k$. Since h_3 and g_3 are the only parts that cannot be canceled out, we have

$$\begin{aligned} & f(k+1) - 2f(k) + f(k-1) \\ = \; & g_3 - h_3 \\ = \; & \mathbf{P} \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{k+1} & = & 1 \end{pmatrix} \left[\pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{k+1} & = & 1 \end{pmatrix} - 2\,\pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix} + \pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{k-1} & = & 1 \end{pmatrix} \right] \end{aligned}$$

□

From Lemma 5, notice the computation of $f(k)$ only depends on a single type of $\mathbf{P}$ and π: when all b_i are 1. In the next Lemma we'll show the explicit formula for $\mathbf{P}$ and π with all $b_i = 1$.

Lemma A.3.5.

$$\mathbf{P}\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix} = \left(\frac{1}{2}\right)^{k+1}$$

$$\pi\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{k+1} & = & 1 \end{pmatrix} - 2\,\pi\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix} + \pi\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{k-1} & = & 1 \end{pmatrix} = -\left(\frac{1}{2}\right)^{k+1}$$

Proof:

$$\mathbf{P}\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix} = \left(\frac{1}{2}\right)^{k+1}$$

is obvious.

To compute π, now consider a modified closed chain $\mathcal{C}_N^{'}$ and open chain $\mathcal{L}_N^{'}$ with the following properties:

- Plant 1 and N have capacity 1.
- Plants $i = 2, \ldots, N-1$ has capacity 2. One of the capacity of plant i is flexible, i.e. can be used to satisfy the demand of either product i or $i+1$. The other capacity of plant i is dedicated, i.e. can only be used to satisfy the demand of product i.

Let $\mathbf{E}[\mathcal{C}_N^{'}]$ and $\mathbf{E}[\mathcal{L}_N^{'}]$ denote the expected performance of this modified closed and open chain respectively, then by the definition of π we have

$$\pi\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix} \quad = \quad \mathbf{E}[\mathcal{C}_{k+2}^{'}] - \mathbf{E}[\mathcal{C}_{k+2}]$$

By Algorithm 4.5.2 (extension to Simchi-Levi and Wei (2012)) $\mathbf{E}[\mathcal{C}_k^{'}]$ can be computed as

$$\begin{aligned} \mathbf{E}[\mathcal{C}_k^{'}] &= \mathbf{E}[\mathcal{L}_k^{'}] + \mathbf{P}\{\tilde{D}_1 = 1\}\mathbf{P}\{\tilde{C}_k = 1\} && \text{(A.99)} \\ &= \mathbf{E}[\mathcal{L}_k^{'}] + \frac{1}{4} && \text{(A.100)} \end{aligned}$$

The performance of the modified open chain $\mathbf{E}[\mathcal{L}_k^{'}]$ can be computed recursively. Let $\tilde{D}_k$ denote the residual demand of product k, then

$$\begin{aligned} \mathbf{E}[\mathcal{L}_{k+1}^{'}] &= (\mathbf{E}[\mathcal{L}_k^{'}] + 1)\mathbf{P}\{\tilde{D}_k = 0\} + (\mathbf{E}[\mathcal{L}_k^{'}] + \frac{3}{2})\mathbf{P}\{\tilde{D}_k = 1\} \\ \mathbf{P}\{\tilde{D}_{k+1} = 1\} &= \frac{1}{2}\mathbf{P}\{\tilde{D}_k = 1\} \\ \mathbf{E}[\mathcal{L}_3^{'}] &= \frac{19}{8} \\ \mathbf{P}\{\tilde{D}_3 = 1\} &= \frac{1}{8} \end{aligned}$$

Arrange the above terms we have

$$\mathbf{E}[\mathcal{L}_{k+1}^{'}] - 2\mathbf{E}[\mathcal{L}_k^{'}] + \mathbf{E}[\mathcal{L}_{k-1}^{'}] = -\left(\frac{1}{2}\right)^{k+1}$$

We have shown the performance of a regular open chain $\mathbf{E}[\mathcal{L}_k]$ with Bernoulli demand is

$$\mathbf{E}[\mathcal{L}_k] = \frac{3k-1}{4}$$

and immediately we have

$$\mathbf{E}[\mathcal{L}_{k+1}] - 2\mathbf{E}[\mathcal{L}_k] + \mathbf{E}[\mathcal{L}_{k-1}] = 0$$

Combine above results we have the desired results in Lemma A.3.5.

□

Now the correctness of the closed form expression 4.6.1 can be proved using Lemma A.3.4 and A.3.5:

Proof of the correctness of Algorithm 4.6.1

With Lemma A.3.4 and A.3.5 we have

$$\frac{f(k+2) - 2f(k+1) + f(k)}{f(k+1) - 2f(k) + f(k-1)} = \frac{1}{4}$$

Given $f(1) = \frac{3}{32}, f(2) = \frac{23}{128}, f(3) = \frac{135}{512}$ we can obtain the general formula of $f(k)$:

$$\begin{aligned} f(k) &= \frac{1}{72} - \frac{1}{72}\left(\frac{1}{4}\right)^k + \frac{1}{12}k \\ &= \mathbf{E}[\tilde{\Phi}(\mathcal{C}, \mathcal{I}_{1\to k}, N, 2)] \end{aligned}$$

□

Case 2: $k = N - 1$

Since the underlying flexibility structure is a closed chain, when $k \geq N - 1$ the benefit of add the $(N-1)^{th}$ and N^{th} inventory arc will overlap with the benefit of the first inventory arc, the computation of $\mathbf{E}[\tilde{\Phi}(\mathcal{C}, \mathcal{I}_{1\to k}, N, 2)]$ needs special treatment. For $k = N - 1$, by Lemma 3 we still have $\mathbf{P}$ holds, but π now depends on N which is different from Case 1. As defined in Case 1, let

$$\pi^N \begin{pmatrix} \tilde{C}_1 = b_1 \\ \tilde{C}_2 = b_2 \\ \vdots \\ \tilde{C}_{N-1} = b_{N-1} \end{pmatrix} \tag{A.101}$$

denote expected residual demand can be satisfied given $(b_1, \ldots, b_{N-1})$. The superscript N reminds that there are N plants and products and π^N is different from π. Then similarly we define $f^N(N-1)$ as the unconditional expected residual demand satisfied with dedicated inventory from plant 1 to $N - 1$, then we have

$$f^N(N-1) = \sum_{b_1} \cdots \sum_{b_{N-1}} \mathbf{P} \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_{N-1} & = & b_{N-1} \end{pmatrix} \pi^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_{N-1} & = & b_{N-1} \end{pmatrix}$$

Then the correctness of closed form expression 4.6.1 can be proved by examining $f(N-1) - f^N(N-1)$

$$\begin{aligned} f(N-1) - f^N(N-1) = & \sum_{b_1} \cdots \sum_{b_{N-1}} \mathbf{P} \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_{N-1} & = & b_{N-1} \end{pmatrix} \pi \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_{N-1} & = & b_{N-1} \end{pmatrix} - \\ & \sum_{b_1} \cdots \sum_{b_{N-1}} \mathbf{P} \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_{N-1} & = & b_{N-1} \end{pmatrix} \pi^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_{N-1} & = & b_{N-1} \end{pmatrix} \end{aligned}$$

If there exists some i such that $b_i = b_{i+1} = 0$, it is equivalent to think as if there is no inventory arc for plant i and $i+1$, then by definition π^N is reduced to π as in the Case 1:

$$\pi^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_{N-1} & = & b_{N-1} \end{pmatrix} = \pi \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_{N-1} & = & b_{N-1} \end{pmatrix}$$

Therefore

$$\begin{aligned} f(N-1) - f^N(N-1) &= \mathbf{P}\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-1} & = & 1 \end{pmatrix} \pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-1} & = & 1 \end{pmatrix} - \\ & \mathbf{P}\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-1} & = & 1 \end{pmatrix} \pi^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-1} & = & 1 \end{pmatrix} \end{aligned}$$

By Lemma A.3.4 we have

$$\mathbf{E}[\mathcal{L}_k^{'}] = -\left(\frac{1}{2}\right)^k + k - \frac{1}{2}$$

therefore

$$\begin{aligned} \pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix} &= \mathbf{E}[\mathcal{L}_{k+2}^{'}] - \mathbf{E}[\mathcal{L}_{k+2}] \\ &= -\left(\frac{1}{2}\right)^{k+2} + \frac{1}{4}k + \frac{1}{4} \end{aligned}$$

To compute π^N, similarly to $\mathcal{L}_k^{'}$, we define a modified open chain $\mathcal{L}_k^{''}$ and closed chain $\mathcal{C}_k^{''}$ such that plant k has capacity 1 and plant 1 to $k-1$ has capacity 2. Then by the definition of π^N we have

$$\pi^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-1} & = & 1 \end{pmatrix} = \mathbf{E}[\mathcal{C}_N^{''}] - \mathbf{E}[\mathcal{C}_N]$$

$\mathbf{E}[\mathcal{C}_N^{''}]$ can be computed using $\mathbf{E}[\mathcal{L}_N^{''}]$ since we have $\mathbf{E}[\mathcal{C}_N^{''}] = \mathbf{E}[\mathcal{L}_N^{''}]$

$$
\begin{aligned}
\mathbf{E}[\mathcal{L}''_{N+1}] &= (\mathbf{E}[\mathcal{L}''_N]+1)\mathbf{P}\{\tilde{D}_N = 0\} + (\mathbf{E}[\mathcal{L}''_N]+\frac{3}{2})\mathbf{P}\{\tilde{D}_N = 1\} \\
\mathbf{P}\{\tilde{D}_{N+1} = 1\} &= \frac{1}{2}\mathbf{P}\{\tilde{D}_N = 1\} \\
\mathbf{E}[\mathcal{L}''_2] &= \frac{23}{8} \\
\mathbf{P}\{\tilde{D}_2 = 1\} &= \frac{1}{8}
\end{aligned}
$$

therefore

$$
\mathbf{E}[\mathcal{C}''_N] = -\left(\frac{1}{2}\right)^N + N
$$

Then

$$
\pi^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-1} & = & 1 \end{pmatrix} = -\left(\frac{1}{2}\right)^N + \frac{1}{4}N
$$

Then

$$
\pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_k & = & 1 \end{pmatrix} - \pi^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-1} & = & 1 \end{pmatrix} = \left(\frac{1}{2}\right)^{N+1}
$$

since

$$
\mathbf{P} \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-1} & = & 1 \end{pmatrix} = \left(\frac{1}{2}\right)^N
$$

therefore

$$
f^N(N-1) = f(N-1) - \left(\frac{1}{2}\right)^{2N+1}
$$

□

Case 3: $k = N$

When $k = N$, now $\mathbf{P}$ also depends on N. Define

$$\mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \end{pmatrix} \tag{A.102}$$

as the joint probability of $\tilde{C}_i = b_i$ for all i from 1 to N, with N plants and products. Then $f^N(N)$ can be expressed using $\mathbf{P}^N$ and π^N:

$$f^N(N) \quad = \quad \sum_{b_1} \cdots \sum_{b_N} \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \end{pmatrix} \pi^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \end{pmatrix}$$

By definition $f^N(N)$ is the residual performance on all N inventory arcs. Since our problem is homogeneous, the residual performance on each inventory arc must be the same. We can define the residual performance on a single inventory arc as $\bar{f}^N(N)$. Then immediately we have $f^N(N) = N\,\bar{f}^N(N)$. W.l.o.g, we can assume $\bar{f}^N(N)$ is the residual performance on the first inventory arc, then $\bar{f}^N(N)$ can be expressed as

$$\bar{f}^N(N) \quad = \quad \sum_{b_2} \cdots \sum_{b_N} \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \end{pmatrix} \bar{\pi}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \end{pmatrix}$$

where $\bar{\pi}^N$ is the residual performance on the first inventory arc given the residual capacity of plant 1 to N are $(1, b_2, \ldots, b_N)$. Here are some properties of $\mathbf{P}^N$ and $\bar{\pi}^N$

Lemma A.3.6. *Some useful properties of* $\mathbf{P}^N$ *and* $\bar{\pi}^N$*:*

(1) No single zero between two ones

(2) Starting point invariant

(3) Add a row DOES NOT HOLD:

$$\mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \end{pmatrix} \quad \neq \quad \mathbf{P}^{N+1} \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \\ \tilde{C}_{N+1} & = & 0 \end{pmatrix} + \mathbf{P}^{N+1} \begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \\ \tilde{C}_{N+1} & = & 1 \end{pmatrix}$$

(4) Relationship between $\mathbf{P}^N$ *and* $\mathbf{P}$*:*

$$\mathbf{P}\begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \end{pmatrix} = \mathbf{P}^{N+1}\begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \\ \tilde{C}_{N+1} & = & 0 \end{pmatrix} + \mathbf{P}^{N+1}\begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \\ \tilde{C}_{N+1} & = & 1 \end{pmatrix}$$

(5) Rotational symmetry:

$$\mathbf{P}^N\begin{pmatrix} \tilde{C}_1 & = & b_1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_{N-1} & = & b_{N-1} \\ \tilde{C}_N & = & b_N \end{pmatrix} = \mathbf{P}^N\begin{pmatrix} \tilde{C}_1 & = & b_2 \\ \tilde{C}_2 & = & b_3 \\ & \vdots & \\ \tilde{C}_{N-1} & = & b_N \\ \tilde{C}_N & = & b_1 \end{pmatrix}$$

(6) Condition of $\mathbf{P}^N = \mathbf{P}$*:*

$$\mathbf{P}^N\begin{pmatrix} \tilde{C}_1 & = & b_1 \\ & \vdots & \\ \tilde{C}_{j-1} & = & b_{j-1} \\ \tilde{C}_j & = & 0 \\ \tilde{C}_{j+1} & = & 0 \\ \tilde{C}_{j+2} & = & 1 \\ & \vdots & \\ \tilde{C}_{N-2} & = & 1 \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix} = \mathbf{P}\begin{pmatrix} \tilde{C}_1 & = & b_1 \\ & \vdots & \\ \tilde{C}_{j-1} & = & b_{j-1} \\ \tilde{C}_j & = & 0 \\ \tilde{C}_{j+1} & = & 0 \\ \tilde{C}_{j+2} & = & 1 \\ & \vdots & \\ \tilde{C}_{N-2} & = & 1 \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix}$$

(7) Relationship between $\bar{\pi}^N$*,* π^N *and* π*:*

If there are some i, j *such that* $\tilde{C}_{j+1} = \tilde{C}_{j+2} = \tilde{C}_{N-i} = \tilde{C}_{N-i-1} = 0$ *and* $N - i - 1 \geq j + 1$*, then*

$$\bar{\pi}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ & \vdots & \\ \tilde{C}_j & = & 1 \\ \tilde{C}_{j+1} & = & 0 \\ \tilde{C}_{j+2} & = & 0 \\ \tilde{C}_{j+3} & = & b_{j+3} \\ & \vdots & \\ \tilde{C}_{N-i-2} & = & b_{i-3} \\ \tilde{C}_{N-i-1} & = & 0 \\ \tilde{C}_{N-i} & = & 0 \\ \tilde{C}_{N-i+1} & = & 1 \\ & \vdots & \\ \tilde{C}_N & = & 1 \end{pmatrix} = \frac{1}{i+j}\,\pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{i+j} & = & 1 \end{pmatrix}$$

If $\tilde{C}_i = 1$ *for all* $i = 1, \ldots, N$*, then*

$$\bar{\pi}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_N & = & 1 \end{pmatrix} = \frac{1}{N}\,\pi^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_N & = & 1 \end{pmatrix}$$

Now we can prove the correctness of closed form expression 4.6.1 by examining $\bar{f}^N(N) - \bar{f}^{N-1}(N-1)$. We have

$$\bar{f}^N(N) = \sum_{b_2} \cdots \sum_{b_N} \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \end{pmatrix} \bar{\pi}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & b_2 \\ & \vdots & \\ \tilde{C}_N & = & b_N \end{pmatrix}$$

Now we classify all 2^{N-1} terms in $\bar{f}^N(N)$ into N groups: each group $q = 1, \ldots, N-1$ contains all terms with $\bar{\pi}^N$ that can be reduced using Property (7) to

$$\frac{1}{q}\,\pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_q & = & 1 \end{pmatrix}$$

and group N contains the term with $\bar{\pi}^N$ that can be reduced using Property (7) to

$$\frac{1}{N}\,\pi^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_N & = & 1 \end{pmatrix}$$

Define $\bar{f}_q^N(N)$ as the sum of all terms in group q, for example, sum of all terms in group 1 is

$$\begin{aligned} \bar{f}_1^N(N) &= \sum_{b_4}\cdots\sum_{b_{N-2}} \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 0 \\ \tilde{C}_3 & = & 0 \\ \tilde{C}_4 & = & b_4 \\ & \vdots & \\ \tilde{C}_{N-2} & = & b_{N-2} \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix} \bar{\pi}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 0 \\ \tilde{C}_3 & = & 0 \\ \tilde{C}_4 & = & b_4 \\ & \vdots & \\ \tilde{C}_{N-2} & = & b_{N-2} \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix} \\ &= \sum_{b_4}\cdots\sum_{b_{N-2}} \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 0 \\ \tilde{C}_3 & = & 0 \\ \tilde{C}_4 & = & b_4 \\ & \vdots & \\ \tilde{C}_{N-2} & = & b_{N-2} \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix} \pi\begin{pmatrix} \tilde{C}_1 & = & 1 \end{pmatrix} \end{aligned}$$

From Property (7) we can see

$$\bar{\pi}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 0 \\ \tilde{C}_3 & = & 0 \\ \tilde{C}_4 & = & b_4 \\ & \vdots & \\ \tilde{C}_{N-2} & = & b_{N-2} \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix} = \pi\begin{pmatrix} \tilde{C}_1 & = & 1 \end{pmatrix}$$

therefore all terms above are indeed in group 1.

Now we manipulate the terms in group 1. First we expand over b_4 and use the Property (5)

$$\bar{f}_1^N(N) = \left(\sum_{b_5} \cdots \sum_{b_{N-2}} \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & b_5 \\ & \vdots & \\ \tilde{C}_{N-6} & = & b_{N-2} \\ \tilde{C}_{N-5} & = & 0 \\ \tilde{C}_{N-4} & = & 0 \\ \tilde{C}_{N-3} & = & 1 \\ \tilde{C}_{N-2} & = & 0 \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix} + \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & b_5 \\ & \vdots & \\ \tilde{C}_{N-6} & = & b_{N-2} \\ \tilde{C}_{N-5} & = & 0 \\ \tilde{C}_{N-4} & = & 0 \\ \tilde{C}_{N-3} & = & 1 \\ \tilde{C}_{N-2} & = & 0 \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 1 \end{pmatrix} \right) \pi \left(\tilde{C}_1 = 1 \right)$$

Using Property (4) we have

$$\bar{f}_1^N(N) = \sum_{b_5} \cdots \sum_{b_{N-2}} \mathbf{P} \begin{pmatrix} \tilde{C}_1 & = & b_5 \\ & \vdots & \\ \tilde{C}_{N-6} & = & b_{N-2} \\ \tilde{C}_{N-5} & = & 0 \\ \tilde{C}_{N-4} & = & 0 \\ \tilde{C}_{N-3} & = & 1 \\ \tilde{C}_{N-2} & = & 0 \\ \tilde{C}_{N-1} & = & 0 \end{pmatrix} \pi \left(\tilde{C}_1 = 1 \right)$$

Using Property (5) again and rename b_5 to b_{N-2} we can see

$$\begin{aligned} \bar{f}_1^N(N) &= \sum_{b_4} \cdots \sum_{b_{N-3}} \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 0 \\ \tilde{C}_3 & = & 0 \\ \tilde{C}_4 & = & b_4 \\ & \vdots & \\ \tilde{C}_{N-3} & = & b_{N-3} \\ \tilde{C}_{N-2} & = & 0 \\ \tilde{C}_{N-1} & = & 0 \end{pmatrix} \pi \left(\tilde{C}_1 = 1 \right) \\ &= \bar{f}_1^{N-1}(N-1) \end{aligned}$$

which is exactly the group 1 of $\bar{f}^{N-1}(N-1)$.

Now for group q from 2 to $N-4$, each group will have q terms. For example in group 2 we have

2 terms

$$\bar{f}_2^N(N) = \frac{1}{2}\sum_{b_5}\cdots\sum_{b_{N-2}} \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ \tilde{C}_3 & = & 0 \\ \tilde{C}_4 & = & 0 \\ \tilde{C}_5 & = & b_5 \\ & \vdots & \\ \tilde{C}_{N-2} & = & b_{N-2} \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix} \pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \end{pmatrix} +$$

$$\frac{1}{2}\sum_{b_4}\cdots\sum_{b_{N-3}} \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 0 \\ \tilde{C}_3 & = & 0 \\ \tilde{C}_4 & = & b_4 \\ & \vdots & \\ \tilde{C}_{N-3} & = & b_{N-3} \\ \tilde{C}_{N-2} & = & 0 \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 1 \end{pmatrix} \pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \end{pmatrix}$$

By Property (5) these 2 terms are exactly the same. Similar argument can be made for all groups from 3 to $N-4$, therefore w.l.o.g. we can write

$$\bar{f}_q^N(N) = \sum_{b_{q+3}}\cdots\sum_{b_{N-2}} \mathbf{P}^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ & \vdots & \\ \tilde{C}_q & = & 1 \\ \tilde{C}_{q+1} & = & 0 \\ \tilde{C}_{q+2} & = & 0 \\ \tilde{C}_{q+3} & = & b_{q+3} \\ & \vdots & \\ \tilde{C}_{N-2} & = & b_{N-2} \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix} \pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_q & = & 1 \end{pmatrix}$$

For q from 2 to $N-4$, similar manipulation can be made as we manipulate group 1, i.e. we have

$$\bar{f}_q^N(N) = \bar{f}_q^{N-1}(N-1)$$

for $q = 1, \ldots, N-4$.

Now we will look at group $N-3$ to N one by one. For group N, we have

$$\bar{f}_N^N(N) \;=\; \frac{1}{N}\mathbf{P}^N\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_N & = & 1 \end{pmatrix}\pi^N\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_N & = & 1 \end{pmatrix}$$

For group $N-1$, using Property (5) we can easily see

$$\bar{f}_{N-1}^N(N) \;=\; \mathbf{P}^N\begin{pmatrix} \tilde{C}_1 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-1} & = & 1 \\ \tilde{C}_N & = & 0 \end{pmatrix}\pi^N\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-1} & = & 1 \end{pmatrix}$$

By Property (1) we have

$$\bar{f}_{N-1}^N(N) \;=\; 0$$

For group $N-2$, using Property (5) we have

$$\bar{f}_{N-2}^N(N) \;=\; \mathbf{P}^N\begin{pmatrix} \tilde{C}_1 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-2} & = & 1 \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix}\pi\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-2} & = & 1 \end{pmatrix}$$

For group $N-3$, using Property (5) we have

$$\bar{f}_{N-3}^N(N) \;=\; \mathbf{P}^N\begin{pmatrix} \tilde{C}_1 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-3} & = & 1 \\ \tilde{C}_{N-2} & = & 0 \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix}\pi\begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-3} & = & 1 \end{pmatrix}$$

It can be easily checked that

$$\mathbf{P}^N\begin{pmatrix} \tilde{C}_1 & = & 1 \\ & \vdots & \\ \tilde{C}_N & = & 1 \end{pmatrix} = \mathbf{P}^N\begin{pmatrix} \tilde{C}_1 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-2} & = & 1 \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix} = \mathbf{P}^N\begin{pmatrix} \tilde{C}_1 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-3} & = & 1 \\ \tilde{C}_{N-2} & = & 0 \\ \tilde{C}_{N-1} & = & 0 \\ \tilde{C}_N & = & 0 \end{pmatrix} = \left(\frac{1}{2}\right)^N$$

To compute π^N, similar to the previous case, we define a modified closed chain $\mathcal{C}_k'''$ such that all plants from 1 to N has capacity 2. Then by the definition of π^N we have

$$\pi^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_N & = & 1 \end{pmatrix} = \mathbf{E}[\mathcal{C}_N'''] - \mathbf{E}[\mathcal{C}_N]$$

Since all demand can be satisfied under $\mathcal{C}_k'''$, immediately we have $\mathbf{E}[\mathcal{C}_k'''] = N$, therefore

$$\pi^N \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_N & = & 1 \end{pmatrix} = \frac{1}{4}N$$

Now we can compute the difference

$$\begin{aligned} & \bar{f}^N(N) - \bar{f}^{N-1}(N-1) \\ = & \sum_{q=1}^{N} \bar{f}_q^N(N) - \sum_{q=1}^{N-1} \bar{f}_q^{N-1}(N-1) \\ = & \sum_{q=N-3}^{N} \bar{f}_q^N(N) - \sum_{q=N-3}^{N-1} \bar{f}_q^{N-1}(N-1) \\ = & -\left(\frac{1}{2}\right)^N \pi \begin{pmatrix} \tilde{C}_1 & = & 1 \\ \tilde{C}_2 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-3} & = & 1 \end{pmatrix} + \left(\frac{1}{2}\right)^N \pi \begin{pmatrix} \tilde{C}_2 & = & 1 \\ \tilde{C}_1 & = & 1 \\ & \vdots & \\ \tilde{C}_{N-2} & = & 1 \end{pmatrix} - \frac{1}{4}\left(\frac{1}{2}\right)^N \end{aligned}$$

then combine the results and simplify the above expression we have

$$\bar{f}^N(N) - \bar{f}^{N-1}(N-1) = \left(\frac{1}{4}\right)^N$$

Given the smallest instance $\bar{f}^3(3) = \frac{5}{64}$, we have the general form of $\bar{f}^N(N)$ as

$$\bar{f}^N(N) = -\frac{1}{3}\left(\frac{1}{4}\right)^N + \frac{1}{12}$$

then using the fact that $f^N(N) = N\,\bar{f}^N(N)$ we have

$$f^N(N) = -\frac{N}{3}\left(\frac{1}{4}\right)^N + \frac{1}{12}N$$

□