Online Demand Fulfillment Problem with Initial Inventory Placement: A Regret Analysis

Alessandro Arlotto, Irem Nur Keskin, Yehua Wei
The Fuqua School of Business, Duke University, Durham, North Carolina 27708
{aa249@duke.edu, iremnur.keskin@duke.edu, yehua.wei@duke.edu}

We investigate a joint inventory placement and online fulfillment problem. At the beginning, the inventory of a single item is distributed to different warehouses. Then, at each period, an order arrives from one of the demand regions, and the decision maker makes an irrevocable decision on whether to accept or reject the order. In our model, we propose the minimum-inventory regret, a notion that includes both the selection of initial inventories and the performance of the selected fulfillment policy. We consider two state-of-the-art fulfillment policies: probabilistic fulfillment and score-based fulfillment. We prove that probabilistic fulfillment has a minimum-inventory regret that scales with the square root of the time horizon. On the other hand, we show that the score-based fulfillment policy has a minimum-inventory regret bound that is independent of the time horizon and polynomial with respect to the number of warehouses and demand regions. Our results have the following implication: the score-based fulfillment policy, when paired with offline inventory placement, outperforms probabilistic fulfillment with *any* inventory placement, and the performance gap increases with the time horizon.

This version: April 24, 2024

1. Introduction

The e-commerce sector has expanded rapidly, with retail sales in 2021 accounting for approximately \$768 billion in the United States and \$4.9 trillion globally (Statista Research Department, 2021). However, this growth has led to skyrocketing fulfillment costs, exemplified by Amazon's \$76.7 billion in shipping costs in the same year (Statista Research Department, 2022). Consequently, e-retailers are actively seeking strategies to curtail these expenses. One strategy that has received wide interest is online demand fulfillment optimization (Acimovic and Farias, 2019). It is well-documented that an effective online demand fulfillment policy, which decides the warehouse from which to fulfill an order upon its arrival, can result in significant cost savings.

In general, determining the optimal fulfillment policy is intractable due to the curse of dimensionality. As a result, one stream of literature has focused on designing fulfillment policies with low regret. In this stream, there are two state-of-the-art low-regret policies, which we refer to as probabilistic fulfillment (e.g., from Balseiro et al., 2023) and score-based fulfillment (e.g., from Vera et al., 2021). The analysis for both policies was performed under the assumption that the initial inventory is fixed and exogenous. However, real-world e-commerce retailers face two significant challenges. Firstly, inventory placement decisions, which are made before the actual online fulfillment process, play a pivotal role in determining the effectiveness of fulfillment policies. As highlighted by Acimovic and Graves (2017), these placement decisions can substantially influence fulfillment costs. As a result, in some real-world situations, to fully comprehend the effectiveness of online demand fulfillment policies, it is essential to incorporate inventory placement decisions. Secondly, with the rise in e-commerce sales, there is a substantial increase in the number of demand locations and warehouses. Thus, the efficiency of an online fulfillment policy becomes increasingly tied to its performance with respect to the rising number of demand regions and warehouses.

With these two challenges in mind, we study a joint online demand fulfillment and inventory placement model with multiple warehouses and multiple demand regions. Specifically, we propose and study the minimum-inventory regret, a metric that incorporates the endogeneity of the inventory placement before the fulfillment process. This definition compares the expected cost of a fulfillment policy at the inventory placement that yields the lowest expected cost, with the optimal expected cost of the offline (full-information) fulfillment policy.

Our analysis addresses the question of which of the state-of-the-art fulfillment policies is preferable when one focuses on the minimum-inventory regret. We analyze how this regret metric interacts with three key parameters: the length of the time horizon T, the total number of warehouses, I, and the total number of demand regions, J. Our findings indicate that the probabilistic fulfillment policy fails to achieve a constant minimum-inventory regret, scaling instead at an order of $\sqrt{T}$. We also show that the score-based fulfillment policy attains a minimum-inventory regret that is independent of T and scales as $I^2J^2\log(J)$, with the number of warehouses I and the number of demand regions J. Our intuition for deriving a polynomial bound for score-based fulfillment stems from the online fulfillment settings where T tends to scale much faster than I and J. In these settings, a policy that has $\sqrt{T}$ regret can easily exceed regret of the policy which is constant with respect to T and polynomial in I and J. Therefore, our regret analysis suggests that, in the online fulfillment setting where T is much larger than I and J, the score-based fulfillment policy is preferable to the probabilistic fulfillment policy when the initial inventory in the warehouses is endogenized.

Our analysis of minimum-inventory regret also provides some insights into inventory placement decisions. The findings indicate that combining score-based fulfillment policy with the offline inventory placement, which is the placement decision that minimizes the expected cost of the offline policy, can achieve lower costs in the joint inventory placement and online fulfillment context. It is important to note that offline inventory placement can be efficiently approximated using sample average approximation. Hence, it is a practical inventory placement decision in real-world settings. Conversely, our lower-bound analysis for probabilistic fulfillment indicates that low regret is not always achievable with the probabilistic fulfillment policy, even when its optimal inventory placement is selected. We also conduct a set of numerical experiments to show that the score-based fulfillment policy outperforms the probabilistic fulfillment policy in non-asymptotic settings.

1.1. Literature Review

Several studies examine regret analysis in the context of general resource allocation problems which includes online fulfillment as a special case. The survey paper by Balseiro et al. (2023) studies a policy, which in our context is referred to as probabilistic fulfillment, that resolves the fluid approximation at each period and uses that solution to make probabilistic allocation decisions. When the fluid problem is finite-dimensional, Balseiro et al. (2023) provide a regret bound that remains constant with respect to T, whereas the regret in the infinite-dimensional setting grows logarithmically in T. This work generalizes the findings of Jasin and Kumar (2012), who showed that the probabilistic resolving policy achieves constant regret in a certain revenue management setting. Central to the results of Balseiro et al. (2023) and Jasin and Kumar (2012) is the assumption that the initial fluid approximation problem yields a nondegenerate solution.

Another constant regret policy for general allocation problems was suggested by Vera and Banerjee (2020); Vera et al. (2021, 2023) (see also Arlotto and Gurvich, 2019). The policy, which in our context is referred to as score-based fulfillment, uses fluid relaxations at each period, but the decisions are made based on the interpretation of the fluid solutions as scores. Their constant regret analysis holds for any finite-dimensional fulfillment problem, with or without the nondegeneracy assumption. It is worthwhile to note that Vera and Banerjee (2020); Vera et al. (2021) introduced a framework to achieve low regret which hinges on a novel "compensated coupling" technique that helps in quantifying the regret of online problems. Compensated coupling is also used under different problems. See e.g., Freund and Banerjee (2019) for its application on online bin packing, and Freund and Zhao (2021) for a network revenue management problem with no-shows. Interestingly, while the regret for the score-based policy is constant for finite-dimensional fulfillment problems, the recent work of Bray (2023) shows that for an infinite dimensional multi-secretary problem (which is a special case of our fulfillment problems), the regret scales at the rate of $\log(T)$ as $T \to \infty$.

Our work is also closely related to the growing literature on online fulfillment (see e.g., Acimovic and Farias, 2019). For instance, Acimovic and Graves (2015) delve into a problem that minimizes the total shipping cost of a single item and introduce a policy that uses the dual values of the fluid relaxation to predict future expected costs. Subsequently, Acimovic and Graves (2017) consider inventory allocation in online fulfillment and propose inventory replenishment policies combined with myopic fulfillment. Govindarajan et al. (2021) study the single-item joint inventory allocation and fulfillment problem in an omnichannel retail environment and also propose inventory replenishment methods. Chen and Graves (2021) formulate an integer program to choose the initial locations of the SKUs under deterministic demand.

The design of fulfillment networks has also received considerable attention. Asadpour et al. (2020) study the single-item fulfillment problem that aims to minimize the expected number of lost sales in the long-chain fulfillment network. One of their findings shows that a modified greedy policy achieves constant regret under the long chain. Xu et al. (2020) modify the policy of Asadpour et al. (2020) for more general fulfillment networks and incorporate initial inventory placement decisions. For more general fulfillment costs, DeValve et al. (2023) evaluate the performance of different fulfillment networks by simulating a number of fulfillment policies. They also propose a new class of spillover-limit fulfillment policies, which are shown to achieve a regret that scales as $\sqrt{T}$ as $T \to \infty$.

There are also studies on the worst-case performance guarantees of single-item online fulfillment. For an adversarial demand arrival model, Andrews et al. (2019) propose a primal-dual algorithm and derive an optimal competitive ratio for the single-item fulfillment model. In the joint online fulfillment and inventory placement setting, Chen et al. (2022) develop worst-case guarantees for a myopic policy coupled with an inventory placement method. Epstein and Ma (2024) also develop the worst-case guarantees considering myopic, offline and fluid inventory placements under two demand arrival models. In addition, recent work has quantified competitive ratios for multi-item fulfillment models. For a stochastic demand arrival model, Jasin and Sinha (2015) provide a policy for a multi-item fulfillment model with order size q and derive a finite asymptotic competitive ratio of $q/4$ as T goes to infinity. The asymptotic ratio is later improved by Ma (2023) to $1+\log(q)$. Xie et al. (2022) study an omnichannel fulfillment problem with in-store and online customers and study an adaptive booking limit algorithm with a tight competitive ratio. Zhao et al. (2022)

examine the multi-item fulfillment problem with a two-layer distribution network with one regional and one front distribution center and present tight competitive ratio bounds for a myopic policy.

In this paper, we analyze the regret of joint inventory placement and online demand fulfillment policies. To the best of our knowledge, Xu et al. (2020) is the only work that analyzes regret with inventory placement decisions. Their model aims to minimize the total number of lost sales while the fulfillment network satisfies a certain flexibility condition. In contrast, the model we study allows for general fulfillment costs and arbitrary fulfillment networks.

2. Joint Inventory Placement and Online Demand Fulfillment Model

In this section, we formally define the joint inventory placement and online demand fulfillment problem. We have a set $[I] := \{1, \cdots, I\}$ of warehouses and a set $[J] := \{1, \cdots, J\}$ of demand regions. Time is discrete with T periods and indexed by $t \in [T] := \{1, 2, \cdots, T\}$. Before the beginning of the first period, the decision maker places γ units of total inventory into the I warehouses. The initial inventory in warehouse i is denoted by κ_i, where $\gamma := \sum_{i\in[I]} \kappa_i$ and $\boldsymbol{\kappa} := [\kappa_i]_{i\in[I]}$.

After the initial inventory is placed, unitary demand arrives in each period $t \in [T]$. The probability that the arrival is from demand region j is λ_j and we have $\sum_{j=1}^{J} \lambda_j = 1$. We define the arrival probability vector as $\boldsymbol{\lambda} := [\lambda_j]_{j\in[J]}$. If the demand arrives from region j in period t, we set $D_j(t) = 1$ and $D_j(t) = 0$ otherwise. Hence, the demand vector $\boldsymbol{D}(t) := [D_j(t)]_{j\in[J]}$ has the generalized Bernoulli distribution with J categories and event probability vector $[\lambda_j]_{j\in[J]}$. We also use $\boldsymbol{D} := [\boldsymbol{D}(t)]_{t\in[T]}$ to denote the demand matrix over the entire time horizon with T periods.

When an order arrives from region j, a decision maker has to choose which warehouse to use to fulfill the order, or whether to lose the arriving order. The cost of using inventory from warehouse i to fulfill a region-j order is c_{ij}, and the cost of a lost sale from region j is $c_{\bullet j}$ with $\boldsymbol{c} := [c_{ij}]_{i\in[I]\cup\{\bullet\}, j\in[J]}$ as the cost matrix. The goal of the decision-maker is to minimize the total expected costs by making sequential fulfillment decisions. Throughout the paper, we let $\lambda_{\min} = \min_{j\in[J]} \lambda_j$ and $c_{\max} = \max_{i\in[I]\cup\{\bullet\}, j\in[J]} c_{ij}$.

A sequential policy π makes fulfillment decisions as soon as an order arrives and by using only the information available up to and including that time. If we receive an order from region j at time t, then we set $\ell_j^\pi(t) = 1$ if the order is rejected and $\ell_j^\pi(t) = 0$ otherwise. Upon accepting the incoming region-j order, we set $x_{ij}^\pi(t) = 1$ if the inventory from warehouse i is used to fulfill the incoming order, and $x_{ij}^\pi(t) = 0$ otherwise. Analogously to previous notations, we employ the vector $\boldsymbol{\ell}^\pi(t) := [\ell_j^\pi(t)]_{j\in[J]}$ and the matrix $\boldsymbol{x}^\pi(t) := [x_{ij}^\pi(t)]_{i\in[I], j\in[J]}$ to keep track of the period t decisions. An admissible policy π is then a sequence of random pairs $\xi^\pi(t) := (\boldsymbol{\ell}^\pi(t), \boldsymbol{x}^\pi(t))$ such that each $\xi^\pi(t)$ is measurable with respect to the σ-field generated by the history up to and including time t.

When the inventory placement vector $\boldsymbol{\kappa}$ is given and π is a sequential fulfillment policy, we let $c^\pi(\boldsymbol{\kappa}, \boldsymbol{D})$ be the corresponding (random) total cost. Let Π denote the set of all admissible sequential fulfillment policies and $U := \{\boldsymbol{\kappa} : \sum_{i\in[I]} \kappa_i \le \gamma\}$ denote the set of all feasible inventory placement vectors. Our joint placement-fulfillment dynamic optimization problem is thus given by

$$\min_{\boldsymbol{\kappa}\in U} \left\{ \min_{\pi\in\Pi} \mathbb{E}\left[c^\pi(\boldsymbol{\kappa}, \boldsymbol{D})\right] \right\}.$$

This dynamic optimization problem becomes intractable as the number of warehouses becomes large, even for a fixed inventory placement $\boldsymbol{\kappa}$. As a result, we propose heuristic policies and assess their performance in terms of their regret against an offline (full-information) solution to the problem. Our analysis begins with a discussion of such offline formulation.

2.1. Deterministic Demand Fulfillment Problem

In the deterministic version of this fulfillment problem, we are given an inventory placement vector $\boldsymbol{\kappa}$ and a demand vector $\boldsymbol{z} = [z_j]_{j\in[J]}$, and we consider the linear programming problem

$$\begin{aligned} F(\boldsymbol{\kappa}, \boldsymbol{z}) = \min_{\boldsymbol{x}, \boldsymbol{\ell}} & \sum_{i\in[I]} \sum_{j\in[J]} c_{ij} x_{ij} + \sum_{j\in[J]} c_{\bullet j} \ell_j \\ \text{s.t. } & \sum_{i\in[I]} x_{ij} + \ell_j \ge z_j && j \in [J] \\ & \sum_{j\in[J]} x_{ij} \le \kappa_i && i \in [I] \\ & x_{ij} \ge 0 && i \in [I],\ \ j \in [J] \\ & \ell_j \ge 0 && j \in [J]. \end{aligned} \tag{1}$$

In the linear programming formulation above, the decision variable x_{ij} corresponds to the amount of inventory in warehouse i that we use to fulfill the region-j demand; and the decision variable ℓ_j is the amount of lost sales from region j. The optimization problem $F(\cdot,\cdot)$ in (1) provides us with different insights that come from specifying different inventory and demand vectors. For instance, if $Z_j = \sum_{t=1}^T D_j(t)$ is the total demand that we receive from region j over T periods and $\boldsymbol{Z} = (Z_1, \cdots, Z_J)$, then the offline (full-information) minimal cost of our demand fulfillment problem with an inventory placement $\boldsymbol{\kappa}$ is given by $c^{\text{off}}(\boldsymbol{\kappa}, \boldsymbol{D}) := F(\boldsymbol{\kappa}, \boldsymbol{Z})$.

2.2. Minimum-Inventory Regret

The expected offline cost $\mathbb{E}[c^{\text{off}}(\boldsymbol{\kappa}, \boldsymbol{D})]$ represents our natural benchmark in quantifying the regret of different heuristic policies, but the initial inventory placement $\boldsymbol{\kappa}$ will play a crucial role both in defining and quantifying such regret. When a fulfillment policy π and the offline solution are both evaluated at the same inventory placement $\boldsymbol{\kappa}$, then we have the notion of fixed-inventory regret.

Definition 2.1. *The* ***fixed-inventory regret*** *of a fulfillment policy π given an inventory placement vector $\boldsymbol{\kappa}$ is*

$$\mathbb{E}\left[c^{\pi}(\boldsymbol{\kappa},\boldsymbol{D})\right]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa},\boldsymbol{D})].$$

We emphasize that for the fixed-inventory regret, both the online policy and the offline problem use an arbitrary inventory placement vector $\boldsymbol{\kappa}$. In practice, however, the initial inventory placement for the fulfillment problem is often endogenized to further reduce the expected costs (see e.g., Acimovic and Graves, 2017). This opens up the question of which are suitable inventory placement vectors to be used when quantifying the regret of a given fulfillment policy. We let $\boldsymbol{\kappa}^{\mathrm{off}}$ be an optimal inventory placement that minimizes the expected offline costs over the stochastic demand $\boldsymbol{D}$, i.e. $\boldsymbol{\kappa}^{\mathrm{off}} \in \underset{\boldsymbol{\kappa}\in U}{\arg\min}\ \mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa},\boldsymbol{D})]$. The optimality of the placement vector and the full-information property of the offline solution gives us, for any inventory vector $\boldsymbol{\kappa}$ and any admissible policy π, that

$$\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})]\leq\mathbb{E}\left[c^{\mathrm{off}}(\boldsymbol{\kappa},\boldsymbol{D})\right]\leq\mathbb{E}\left[c^{\pi}(\boldsymbol{\kappa},\boldsymbol{D})\right].$$

The expected offline cost with initial inventory placement $\boldsymbol{\kappa}^{\mathrm{off}}$ is then a lower bound for any admissible policy and any inventory placement. This motivates us to define the *minimum-inventory regret* of a given policy π as follows.

Definition 2.2. *Given a fulfillment policy π, let $\boldsymbol{\kappa}^{\pi}\in\arg\min_{\boldsymbol{\kappa}\in U}\mathbb{E}\left[c^{\pi}(\boldsymbol{\kappa},\boldsymbol{D})\right]$, then* ***minimum-inventory regret*** *of π is defined as*

$$\mathbb{E}\left[c^{\pi}(\boldsymbol{\kappa}^{\pi},\boldsymbol{D})\right]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})].$$

By optimality of $\boldsymbol{\kappa}^{\pi}$, we have that the minimum-inventory regret, which is computed as $\mathbb{E}\left[c^{\pi}(\boldsymbol{\kappa}^{\pi},\boldsymbol{D})\right]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})]$, is upper-bounded by $\mathbb{E}\left[c^{\pi}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})\right]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})]$. This represents the fixed-inventory regret of policy π with the inventory vector $\boldsymbol{\kappa}^{\mathrm{off}}$. Additionally, the online fulfillment problem can be seen as a specialized case of the broader online resource allocation problem studied in Balseiro et al. (2023) and Vera et al. (2021). Leveraging this connection, we can apply the results from the latter to identify fulfillment policies and calculate their corresponding fixed-inventory regret under inventory vector $\boldsymbol{\kappa}^{\mathrm{off}}$. This approach would immediately yield upper bounds for each of the policies. However, these upper bounds do not provide us with a sufficient understanding of the effectiveness of the two state-of-the-art fulfillment policies in the context of fulfillment with inventory placement.

Specifically, Theorem 1 of Balseiro et al. (2023) implies that the upper bound on the regret of the probabilistic fulfillment policy is independent of T under the nondegeneracy assumption. However, such regret bound also scales with $1/\delta$, where δ is a measure of the degeneracy of $F(\boldsymbol{\kappa}^{\text{off}}, \boldsymbol{\lambda}T)$. Since $\boldsymbol{\kappa}^{\text{off}}$ is selected endogenously based on $\boldsymbol{D}$, $\mathbf{c}$, and T, there is no guarantee that $1/\delta$ does not grow with T. Similarly, Theorem 2 of Vera et al. (2021) suggests that the regret of our score-based fulfillment policy has an upper bound that is independent of T and $1/\delta$. This result, however, does not quantify the policy's dependency on parameters I and J.

While the above upper bounds on the fixed-inventory regret for probabilistic and scored-based fulfillment policies offer valuable perspectives, they do not provide a clear picture of which policy is more suitable for the online fulfillment problem with endogenized inventory placement. For the probabilistic fulfillment policy, the regret is not guaranteed to be good because δ, the measure of the degeneracy of $F(\boldsymbol{\kappa}^{\text{off}}, \boldsymbol{\lambda}T)$, can become close to zero as the inventory and T are scaled to infinity. In comparison, although the regret upper bound for the scored-based fulfillment policy does not depend on δ, it is not clear how it scales with parameters such as I and J. While it is commonly accepted that T is much larger than I and J in fulfillment settings, regrets that scale exponentially with I or J can easily dominate regrets that scale with $\sqrt{T}$.

The above discussion motivates us to perform a sharper analysis of the minimum-inventory regret of each of the aforementioned policies. Interestingly, as we will show in the subsequent sections, the minimum-inventory regret of the probabilistic fulfillment policy is $\Omega(\sqrt{T})$ [1], while the minimum-inventory regret of the score-based fulfillment policy exhibits an upper bound that is independent of T and polynomial in I and J.

3. Comparison between Probabilistic and Score-Based Policies

In this section, we formally introduce the probabilistic fulfillment and score-based fulfillment policies for our joint inventory placement and online demand fulfillment problem, and state our main results.

3.1. Probabilistic Fulfillment Policy (PF)

In each period t, the probabilistic fulfillment heuristic solves a linear programming problem that considers the remaining amount of inventory in the warehouses and the expected remaining demand (i.e., the expected value of the total demand that arrives between periods t and T). The (scaled) solution of that optimization problem is then interpreted as a probability vector. Heuristic actions are randomized on the basis of that probability vector and of the arrival type. More formally, in each period $t \in [T]$, the probabilistic fulfillment heuristic follows the steps listed below.

[1] Throughout the paper, we will study how regret scales with model parameters, including T, I, and J. We will often use big-O notations, where $O(\cdot)$ represents an upper bound up to a constant factor, $\Omega(\cdot)$ represents a lower bound up to a constant factor, $\tilde{O}(\cdot)$ represents an upper bound up to a logarithmic factor, and $\Theta(\cdot)$ represents both an upper and lower bound up to a constant factor.

1. Given the inventory vector $\boldsymbol{\kappa}(t)$ at the beginning of period t, solve the linear program $F(\boldsymbol{\kappa}(t), \boldsymbol{\lambda}(T-t+1))$ in (1) and let $(\widehat{\boldsymbol{\ell}}, \widehat{\boldsymbol{x}})$ be the solution;
2. If the demand arrives from region j, construct the probability vector

$$\Phi_j = [\phi_{1j}, \phi_{2j}, \cdots, \phi_{Ij}, \phi_{(I+1)j}] = \left[\frac{\widehat{x}_{1j}}{(T-t+1)\lambda_j}, \frac{\widehat{x}_{2j}}{(T-t+1)\lambda_j}, \cdots, \frac{\widehat{x}_{Ij}}{(T-t+1)\lambda_j}, \frac{\widehat{\ell}_j}{(T-t+1)\lambda_j}\right].$$

Each component $\frac{\widehat{x}_{ij}}{(T-t+1)\lambda_j}$ represents the probability the policy will choose warehouse i to fulfill the demand from region j; while $\frac{\widehat{\ell}_j}{(T-t+1)\lambda_j}$ is the probability that the policy will select no warehouse;

3. Draw a warehouse (or select no warehouse) according to the probability vector Φ_j. If no warehouse is selected, the demand is lost.

The fixed-inventory regret bound in Balseiro et al. (2023) for a general class of resource-constrained problems applies to the online demand fulfillment problem. From their work, it follows that under the so-called fluid scaling, where T is scaled to infinity and $\kappa_j = \bar{\kappa}_j T$ for some fixed $\bar{\boldsymbol{\kappa}}$, the PF policy admits fixed-inventory regret independent of T, provided that the fluid problem, $F(\bar{\boldsymbol{\kappa}}, \boldsymbol{\lambda})$, has a non-degenerate optimal solution. However, when the initial inventories are endogenized, i.e., $\boldsymbol{\kappa}^{\mathrm{PF}} \in \arg\min_{\boldsymbol{\kappa}\in U} \mathbb{E}\left[c^{\mathrm{PF}}(\boldsymbol{\kappa}, \boldsymbol{D})\right]$, then $\boldsymbol{\kappa}^{\mathrm{PF}}$ is no longer scaling with some fixed $\bar{\boldsymbol{\kappa}}$. More importantly, the minimum-inventory regret of probabilistic fulfillment policy, instead of being independent of T, scales as $\sqrt{T}$. This is formalized in the next theorem.

Theorem 3.1. *There exist problem instances such that the minimum-inventory regret under the probabilistic fulfillment is $\Omega(\sqrt{T})$. That is,*

$$\mathbb{E}\left[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})\right] - \mathbb{E}\left[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}}, \boldsymbol{D})\right] = \Omega(\sqrt{T}). \tag{2}$$

In general, it is difficult to characterize $\boldsymbol{\kappa}^{\mathrm{PF}}$ exactly, and as a result, while proving the lower bound in Theorem 3.1, we divide all possible values of $\boldsymbol{\kappa}^{\mathrm{PF}}$ into two regions and study the regret separately in each region via telescoping and compensated coupling. This essentially suggests that we can lower bound the performance of PF for all possible inventory placement vectors. Moreover, in Appendix A, we show that the minimum-inventory regret of PF is upper bounded by $O(I\sqrt{T})$ and thus, the lower-bound in Theorem 3.1 is tight with respect to T.

3.2. Score-based Fulfillment Policy (SF)

Similar to PF, the score-based fulfillment policy solves a linear optimization problem at the beginning of each period. The optimal solution to this linear program is then interpreted as a set of scores, and the heuristic decides on the arriving demand by choosing an action with the highest score. More precisely, the score-based fulfillment policy takes the following steps at each time period $t \in [T]$ when the inventory at the *beginning of period* t is $\boldsymbol{\kappa}(t)$:

1. Solve the linear optimization problem $F(\boldsymbol{\kappa}(t), \boldsymbol{\lambda}(T-t+1))$ defined in (1) and denote with $(\widehat{\boldsymbol{\ell}}, \widehat{\boldsymbol{x}})$ its optimal solution;
2. Decide on demand arriving from region j as follows:

—If $\max\{\widehat{\ell}_j, \widehat{x}_{ij}, \cdots, \widehat{x}_{Ij}\} = \widehat{\ell}_j$, then the arriving demand is rejected.

—If $\max\{\widehat{\ell}_j, \widehat{x}_{ij}, \cdots, \widehat{x}_{Ij}\} \neq \widehat{\ell}_j$, find $\widehat{i} \in \arg\max\{\widehat{x}_{ij}, \cdots, \widehat{x}_{Ij}\}$ and then fulfill the arriving demand with the inventory from warehouse $\widehat{i}$.

The next theorem shows that the minimum-inventory regret of the score-based fulfillment policy is independent of the length of the time horizon and polynomial with respect to the number of warehouses and demand regions.

Theorem 3.2. *The minimum-inventory regret of* SF *satisfies the bound*

$$\mathbb{E}\left[c^{\mathrm{SF}}(\boldsymbol{\kappa}^{\mathrm{SF}}, \boldsymbol{D})\right] - \mathbb{E}\left[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}}, \boldsymbol{D})\right] \leq c_{\max} \cdot \frac{J^2(I+1)^2}{\lambda_{\min}^2}(1 + 2\log(2J))$$

where $c_{\max} = \max_{ij} c_{ij}$ *and* $\lambda_{\min} = \min_j \lambda_j$.

Consistently with Vera et al. (2021), Theorem 3.2 shows the regret bound is independent of T. However, in contrast with previous work, our result also quantifies that the regret bound scales with I, the number of warehouses, and J, the number of demand regions, as $\tilde{O}(I^2J^2)$.

4. Regret Analysis

In this section, we formally prove the regret bounds of Section 3 for probabilistic and score-based fulfillment policies. In Section 4.1, we present an example showing that probabilistic fulfillment has a minimum-inventory regret of $\Omega(\sqrt{T})$. Then, in Section 4.2, we derive an upper bound for the regret of score-based fulfillment, which is independent of T and scales polynomially with I and J.

Central to our analysis is the compensated coupling idea introduced in Vera and Banerjee (2020); Vera et al. (2021), which divides the regret into per-period nonnegative quantities using a telescoping sum approach. Specifically, at time $1 \leq t \leq T$, let C_t^π be the cost incurred by policy π, $\boldsymbol{\kappa}(t)$ (which depends on π) be the inventory vector at the start of the time period, and $\boldsymbol{Z}(t)$ be the total arrivals from period t to T. We then have:

$$\begin{aligned}\mathbb{E}[c^\pi(\boldsymbol{\kappa}, \boldsymbol{D})] - \mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}, \boldsymbol{D})] &= \sum_{t=1}^T \mathbb{E}[C_t^\pi] - \sum_{t=1}^T \mathbb{E}[F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t)) - F(\boldsymbol{\kappa}(t+1), \boldsymbol{Z}(t+1))] \\ &= \sum_{t=1}^T \mathbb{E}[C_t^\pi - F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t)) + F(\boldsymbol{\kappa}(t+1), \boldsymbol{Z}(t+1))]. \end{aligned} \tag{3}$$

The first equality follows because we divide the total cost of the offline problem and the policy into per-period costs, and the second equality follows by the linearity of the expectation. Furthermore, we define:

$$\Delta_t^\pi := C_t^\pi - F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t)) + F(\boldsymbol{\kappa}(t+1), \boldsymbol{Z}(t+1)). \tag{4}$$

By definition, the overall regret is decomposed into a sum of Δ_t^π for t from 1 to T. Throughout the paper, Δ_t^π will be referred to as the t-th period regret (under π). One of the key features of compensated coupling, is that for any t, the t-th period regret is nonnegative almost surely, i.e.,

$$\mathbb{P}(\Delta_t^\pi \geq 0) = 1. \tag{5}$$

For a formal proof, see Lemma B.1 in the appendix.

4.1. Minimum-Inventory Regret Analysis for Probabilistic Fulfillment Policy

Theorem 3.1 indicates that the minimum-inventory regret of probabilistic fulfillment grows like the square root of T as $T \to \infty$. We will prove the lower bound by constructing a specific setting with $I = 2$ warehouses and $J = 2$ demand regions.

Example 4.1. *Consider a setting with two warehouses and two demand regions, and a cost vector such that $c_{22} < c_{11}$, $c_{\bullet 1} = c_{\bullet 2}$, $c_{11} < c_{21} < c_{\bullet 1}$, and $c_{22} < c_{\bullet 2} < c_{12}$. For any given horizon length T, we set the total inventory $\gamma = \theta T$, with $\lambda_2 < \theta < 1$.*

Under the conditions of Example 4.1, the offline policy, which knows all the arrivals in advance, fulfills the demand in the following order:

1. Fulfill the region-2 demand from warehouse 2. If there is not sufficient inventory in warehouse 2, the excess demand from region 2 is lost.

2. Fulfill the region-1 demand from warehouse 1. If there is still any unfulfilled demand from region 1, first check warehouse 2 and fulfill the demand as much as possible. If there is still unfulfilled demand, the remaining sales are lost.

In Lemma B.2, we formally present the solution to the offline problem at any time t, following the structure in Example 4.1.

We now describe a road map for our analysis that lower-bounds the minimum-inventory regret of probabilistic fulfillment. We will examine the difference $\mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] - \mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}}, \boldsymbol{D})]$ where $\boldsymbol{\kappa}^{\mathrm{PF}} \in \arg\min_{\boldsymbol{\kappa} \in U} \mathbb{E}\left[c^{\mathrm{PF}}(\boldsymbol{\kappa}, \boldsymbol{D})\right]$ and demonstrate that this difference is always bounded below by $C\sqrt{T}$ for some constant $C > 0$. Because the actual value of $\boldsymbol{\kappa}^{\mathrm{PF}}$ is difficult to characterize, we will separately study what happens when κ_2^{PF} and $\lambda_2 T$ are far apart or in close proximity. More specifically, we first study the case with κ_2^{PF} more than $2C^*\sqrt{T}$ away from $\lambda_2 T$ for some constant C^*, then the case with κ_2^{PF} at most $2C^*\sqrt{T}$ away from $\lambda_2 T$ [2].

To start our formal analysis, we first observe that, as Example 4.1 considers only $I = 2$ warehouses, the initial inventory γ is split between them: if $\kappa_2 \in [0, \gamma]$ is the amount of inventory placed

[2] We note that the fractional inventories are wasted at the end of the horizon.

in warehouse 2, then the initial inventory in warehouse 1 is $\kappa_1 = \gamma - \kappa_2$. Hence, the total expected cost $\mathbb{E}[c^{\text{off}}(\boldsymbol{\kappa}, \boldsymbol{D})] = \mathbb{E}[c^{\text{off}}((\gamma-\kappa_2, \kappa_2), \boldsymbol{D})]$ is a univariate function of κ_2, and we set

$$f(\kappa_2) := \mathbb{E}[c^{\text{off}}((\gamma-\kappa_2, \kappa_2), \boldsymbol{D})]. \tag{6}$$

Because $c^{\text{off}}(\boldsymbol{\kappa}, \boldsymbol{D})$ is the objective of a linear optimization problem, it is convex in $\boldsymbol{\kappa}$, implying that $c^{\text{off}}((\gamma-\kappa_2, \kappa_2), \boldsymbol{D})$ and hence $f(\kappa_2)$ is convex in κ_2.

In order to bound the difference between $\mathbb{E}\left[c^{\text{PF}}(\boldsymbol{\kappa}^{\text{PF}}, \boldsymbol{D})\right]$ and $\mathbb{E}\left[c^{\text{off}}(\boldsymbol{\kappa}^{\text{off}}, \boldsymbol{D})\right]$ in the case that κ_2^{PF} is "far" apart from $\lambda_2 T$, we focus on the offline difference component of the minimum-inventory regret, which is expressed as

$$\mathbb{E}\left[c^{\text{off}}(\boldsymbol{\kappa}^{\text{PF}}, \boldsymbol{D})\right] - \mathbb{E}\left[c^{\text{off}}(\boldsymbol{\kappa}^{\text{off}}, \boldsymbol{D})\right] = f(\kappa_2^{\text{PF}}) - f(\kappa_2^{\text{off}})$$

For this purpose, we next present a pair of bounds on the marginal changes in the expected offline costs when κ_2 differs from $\lambda_2 T$.

$$f(\kappa_2+1) - f(\kappa_2) \geq (c_{21}-c_{11}) - (c_{21}+c_{\bullet 1}) \exp\left\{-2\frac{(\kappa_2-\lambda_2 T)^2}{T}\right\} \qquad \text{for } \kappa_2 > \lambda_2 T; \tag{7}$$

and

$$f(\kappa_2+1) - f(\kappa_2) \leq -(c_{11}-c_{22}) + (c_{21}-c_{22}) \exp\left\{-2\frac{(\lambda_2 T-\kappa_2)^2}{T}\right\} \qquad \text{for } \kappa_2 < \lambda_2 T. \tag{8}$$

In the interest of space, we leave the proof for (7) and (8) as Lemma B.3 in the appendix. Next, we use this pair of inequalities to show an important characteristic of the offline inventory placement in Example 4.1: that κ_2^{off} and $\lambda_2 T$ does not differ much from each other.

Lemma 4.2. *Let C^* be a constant defined as*

$$C^* = \max\left\{\sqrt{\left|\log\left(\frac{c_{21}-c_{11}}{c_{21}+c_{\bullet 1}}\right)\right|}+1, \sqrt{\left|\log\left(\frac{c_{11}-c_{22}}{c_{21}-c_{22}}\right)\right|}\right\}. \tag{9}$$

Then, in the setting of Example 4.1, $\boldsymbol{\kappa}^{\text{off}}$ satisfies $|\kappa_2^{\text{off}} - \lambda_2 T| \leq C^*\sqrt{T}$ for all T where $\lambda_2 T \geq C^*\sqrt{T}$.

Proof. Recall from Example 4.1, that $0 < c_{22} < c_{11} < c_{21} < c_{\bullet 1}$. Let $C' = \sqrt{\left|\log\left(\frac{c_{21}-c_{11}}{c_{21}+c_{\bullet 1}}\right)\right|}+1$. If $\kappa_2 \geq \lambda_2 T + C'\sqrt{T}$, then we have $\kappa_2 - 1 \geq \lambda_2 T + (C'-1)\sqrt{T}$ as $\sqrt{T} \geq 1$. By Inequality (7), we have that whenever $\kappa_2 - 1 \geq \lambda_2 T + (C'-1)\sqrt{T}$,

$$f(\kappa_2) - f(\kappa_2-1) \geq (c_{21}-c_{11}) - (c_{21}+c_{\bullet 1}) \exp\left\{-2\frac{(\kappa_2-1-\lambda_2 T)^2}{T}\right\} > 0. \tag{10}$$

Similarly, let $C'' = \sqrt{\left|\log\left(\frac{c_{11}-c_{22}}{c_{21}-c_{22}}\right)\right|}$. By Inequality (8), if $\kappa_2 \leq \lambda_2 T - C''\sqrt{T}$, we have that

$$f(\kappa_2+1) - f(\kappa_2) \leq -(c_{11}-c_{22}) - (c_{21}+c_{22}) \exp\left\{-2\frac{(\kappa_2-\lambda_2 T)^2}{T}\right\} < 0. \tag{11}$$

Next, we will prove by contradiction that C^* is sufficient to show $|\kappa_2^{\text{off}} - \lambda_2 T| < C^*\sqrt{T}$. To see this, note that if $|\kappa_2^{\text{off}} - \lambda_2 T| > C^*\sqrt{T}$, then (10) and (11) imply that either

$$f(\kappa_2^{\text{off}}) - f(\kappa_2^{\text{off}} - 1) > 0 \text{ or } f(\kappa_2^{\text{off}} + 1) - f(\kappa_2^{\text{off}}) < 0,$$

contradicting the optimality of κ_2^{off}. Thus, we must have that $|\kappa_2^{\text{off}} - \lambda_2 T| \le C^*\sqrt{T}$. This completes the proof of the lemma. Q.E.D.

A simple consequence of Lemma 4.2 is that when the difference between κ_2^{PF} and $\lambda_2 T$ exceeds $2C^*\sqrt{T}$, then the difference between κ_2^{PF} and κ_2^{off} is at least $C^*\sqrt{T}$. Next, we use this fact, combined with the bounds on the marginal changes in the expected costs provided in (7) and (8) to show that when $|\kappa_2^{\text{PF}} - \lambda_2 T| \ge 2C^*\sqrt{T}$, the discrepancy between $\mathbb{E}[c^{\text{off}}(\boldsymbol{\kappa}^{\text{PF}}, \boldsymbol{D})]$ and $\mathbb{E}[c^{\text{off}}(\boldsymbol{\kappa}^{\text{off}}, \boldsymbol{D})]$ is at least on the order of $\sqrt{T}$.

Proposition 4.3. *Let C^* be the constant in Lemma 4.2. In the setting of Example 4.1, if $\boldsymbol{\kappa}^{\text{PF}}$ satisfies $|\kappa_2^{\text{PF}} - \lambda_2 T| \ge 2C^*\sqrt{T}$, then there exists a constant $C > 0$ such that*

$$f(\kappa_2^{\text{PF}}) - f(\kappa_2^{\text{off}}) = \mathbb{E}[c^{\text{off}}(\boldsymbol{\kappa}^{\text{PF}}, \boldsymbol{D})] - \mathbb{E}[c^{\text{off}}(\boldsymbol{\kappa}^{\text{off}}, \boldsymbol{D})] \ge C\sqrt{T}.$$

Consequently, we have

$$\mathbb{E}[c^{\text{PF}}(\boldsymbol{\kappa}^{\text{PF}}, \boldsymbol{D})] - \mathbb{E}[c^{\text{off}}(\boldsymbol{\kappa}^{\text{off}}, \boldsymbol{D})] \ge C\sqrt{T}.$$

Proof. Let's first consider the case where $\kappa_2^{\text{PF}} \ge \lambda_2 T + 2C^*\sqrt{T}$. By Lemma 4.2, this also implies $\kappa_2^{\text{PF}} \ge \kappa_2^{\text{off}} + C^*\sqrt{T}$. Recall that f is convex, which implies that it is monotonically increasing in κ_2 over the interval $[\kappa_2^{\text{off}}, \gamma]$ where γ is the total inventory to be distributed and κ_2^{off} is the inventory vector that minimizes $f(\kappa_2)$. Then, using $\kappa_2^{\text{PF}} \ge \lambda_2 T + 2C^*\sqrt{T}$ and the fact that $f(\kappa_2^{\text{off}}) \le f(\lambda_2 T + C^*\sqrt{T})$, we have

$$f(\kappa_2^{\text{PF}}) - f(\kappa_2^{\text{off}}) \ge f(\lambda_2 T + 2C^*\sqrt{T}) - f(\kappa_2^{\text{off}}) \ge f(\lambda_2 T + 2C^*\sqrt{T}) - f(\lambda_2 T + C^*\sqrt{T}). \tag{12}$$

We next divide the right-hand side of (12) using telescoping sum and then use (7) to bound the first differences:

$$\begin{aligned}
f(\lambda_2 T + 2C^*\sqrt{T}) - f(\lambda_2 T + C^*\sqrt{T}) &= \sum_{t=C^*\sqrt{T}+1}^{2C^*\sqrt{T}} \left(f(\lambda_2 T + t) - f(\lambda_2 T + t - 1)\right) \\
&\ge \sum_{t=1}^{C^*\sqrt{T}} (c_{21} - c_{11}) - (c_{21} + c_{\bullet 1}) \exp\left(-2(C^* + t/\sqrt{T})^2\right) \\
&\ge C^*\sqrt{T}\left((c_{21} - c_{11}) - (c_{21} + c_{\bullet 1}) \exp(-2C^{*2})\right).
\end{aligned}$$

By plugging in C^* defined in Lemma 4.2, we observe that $(c_{21}-c_{11})-(c_{21}+c_{\bullet 1})\exp(-2C^{*2})>0$. Hence, the difference $f(\kappa_2^{\mathrm{PF}})-f(\kappa_2^{\mathrm{off}})$ scales with $\sqrt{T}$.

The case where $\kappa_2^{\mathrm{PF}} \le \lambda_2 T - 2C^*\sqrt{T}$ builds on similar ideas. Note that, the function f is monotonically decreasing in κ_2 over the interval $[0,\kappa_2^{\mathrm{off}}]$ by its convexity. Then, by $\kappa_2^{\mathrm{PF}} \le \lambda_2 T - 2C^*\sqrt{T}$ and $f(\kappa_2^{\mathrm{off}}) \le f(\lambda_2 T - C^*\sqrt{T})$ we have that

$$f(\kappa_2^{\mathrm{PF}})-f(\kappa_2^{\mathrm{off}}) \ge f(\lambda_2 T - 2C^*\sqrt{T}) - f(\kappa_2^{\mathrm{off}}) \ge f(\lambda_2 T - 2C^*\sqrt{T}) - f(\lambda_2 T - C^*\sqrt{T}). \tag{13}$$

By rewriting the right-hand side of (13) with telescoping sum and using (8), we obtain

$$\begin{aligned} f(\lambda_2 T - 2C^*\sqrt{T}) - f(\lambda_2 T - C^*\sqrt{T}) &= \sum_{t=C^*\sqrt{T}+1}^{2C^*\sqrt{T}} \left(f(\lambda_2 T - t) - f(\lambda_2 T - t + 1)\right) \\ &\ge \sum_{t=1}^{C^*\sqrt{T}} -\left((c_{22}-c_{11}) + (c_{21}-c_{22})\exp\left(-2C^{*2}\right)\right) \\ &= -C^*\sqrt{T}\left((c_{22}-c_{11}) + (c_{21}-c_{22})\exp\left(-2C^{*2}\right)\right). \end{aligned}$$

Observe that the last term scales with $\sqrt{T}$ and is positive since $(c_{22}-c_{11})+(c_{21}-c_{22})\exp\left(-2C^{*2}\right)<0$ again by the definition of C^* in Lemma 4.2. By combining the two cases, and letting C to be the maximum of $C^*\left((c_{21}-c_{11})-(c_{21}+c_{\bullet 1})\exp(-2C^{*2})\right)$ and $-C^*\left((c_{22}-c_{11})+(c_{21}-c_{22})\exp\left(-2C^{*2}\right)\right)$, we obtain

$$f(\kappa_2^{\mathrm{PF}})-f(\kappa_2^{\mathrm{off}}) \ge C\sqrt{T}.$$

Finally, we have $\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{PF}},\boldsymbol{D})] \le \mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}},\boldsymbol{D})]$ because the expected offline cost is always a lower bound on the expected cost of PF at any inventory vector $\boldsymbol{\kappa}$. Therefore, we have

$$\mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}},\boldsymbol{D})]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})] \ge \mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{PF}},\boldsymbol{D})]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})] = f(\kappa_2^{\mathrm{PF}})-f(\kappa_2^{\mathrm{off}}) \ge C\sqrt{T}.$$

Q.E.D.

Having covered the case where $|\kappa_2^{\mathrm{PF}}-\lambda_2 T| \ge 2C^*\sqrt{T}$ in Proposition 4.3, we next analyze when $|\kappa_2^{\mathrm{PF}}-\lambda_2 T| \le 2C^*\sqrt{T}$. Before presenting the formal result and its analysis, we describe the intuition behind the key steps for our analysis.

The first step of our analysis is to carve out an event E with probability bounded away from zero, such that policy PF has $\Omega(\sqrt{T})$ regret conditioned on E. This event is somewhat parallel to the one illustrated by Bumpensanti and Wang (2020) in their network revenue management example, where the authors proved that an analogous policy to PF exhibits a regret of $\Omega(\sqrt{T})$.

However, our study diverges from the example in Bumpensanti and Wang (2020) in a critical aspect: the inventory level of warehouse-2 at time t under policy PF, $\kappa_2(t)$, does not conform to the typical martingale behavior conditioned on event E. This deviation leads us to forgo the conventional martingale concentration approach in the literature. Instead, we derive a novel analysis focusing on the decomposed per-period regret Δ_t^{PF} under compensated coupling. Although compensated coupling was initially introduced by Vera and Banerjee (2020) to establish an upper bound on regret, our application in determining a lower bound on regret for policy PF is novel and could be of significant independent interest for analyzing other policies.

Because of the non-standard behavior of $\kappa_2(t)$, we derive a lower bound for

$$\mathbb{E}\left[\Delta_t^{\mathrm{PF}} + \frac{\sum_{s=1}^{T}\Delta_s^{\mathrm{PF}}}{T'} \,|\, \kappa_2(t)\right],$$

for any value of $\kappa_2(t)$ where $T' = \Theta(T)$ and $T' \le T$. Our rationale is intuitive: since Δ_t^{PF} is non-negative for any t almost surely, we can establish a lower bound for $\mathbb{E}\left[\Delta_t^{\mathrm{PF}} \,|\, \kappa_2(t)\right]$ when $\kappa_2(t)$ closely aligns with the average offline trajectory, and a lower bound for $\mathbb{E}\left[\frac{\sum_{s=1}^{T}\Delta_s^{\mathrm{PF}}}{T'} \,|\, \kappa_2(t)\right]$ otherwise. Next, we present the formal statement of our result and its proof.

Proposition 4.4. *Let C^* be the constant in Lemma 4.2. If $\boldsymbol{\kappa}^{\mathrm{PF}}$ satisfies $|\kappa_2^{\mathrm{PF}} - \lambda_2 T| \le 2C^*\sqrt{T}$, then*

$$\mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] - \mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] = \Omega(\sqrt{T})$$

for the setting in Example 4.1.

Proof. Let $T' = \frac{(1-\theta)T}{2}$ and denote $B_2(t)$ to be the total number of region-2 arrivals from the beginning to period t, i.e., $B_2(t) = \sum_{s=1}^{t} D_2(s)$. Recall that in contrast, $Z_2(t)$ is defined as the number of region-2 arrivals from period t to T.

We divide the horizon into three distinct sub-intervals: $[0, T']$, $(T', 2T']$, and $(2T', T]$. We then define the events E_1, E_2, and E_3 to represent the arrivals from region 2 at each of these sub-intervals, respectively.

$$\begin{aligned}
E_1 &= \left\{B_2(T') - \lambda_2 T' \in [-7C^*\sqrt{T}, -6C^*\sqrt{T}]\right\},\\
E_2 &= \left\{B_2(t) - B_2(T') - \lambda_2(t - T') \in [-C^*\sqrt{T}, C^*\sqrt{T}], \forall t \in (T', 2T']\right\},\\
E_3 &= \left\{B_2(T) - B_2(2T') - \lambda_2(T - 2T') \in [10C^*\sqrt{T}, \infty)\right\}.
\end{aligned}$$

Let $E = E_1 \cap E_2 \cap E_3$. Under the event E, we have that the number of arrivals from region 2 is smaller than its expectation in the first interval (event E_1), approximately equal to its expectation *throughout* the second interval (event E_2), and larger than its expectation in the last interval (event

E_3). Moreover, for any arrival instance in E, we have that the number of total arrivals from region 2 exceeds κ_2^{PF}, as

$$Z_2 = B_2(T) \geq \lambda_2 T - 7C^*\sqrt{T} - C^*\sqrt{T} + 10C^*\sqrt{T} \geq \lambda_2 T + 2C^*\sqrt{T} \geq \kappa_2^{\mathrm{PF}}, \tag{14}$$

where the last inequality follows because $|\kappa_2^{\mathrm{PF}} - \lambda_2 T| \leq 2C^*\sqrt{T}$.

To analyze the fixed-inventory regret of PF, we use the compensated coupling equation described in (3), which decomposes the regret as:

$$\mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D}) - c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] = \sum_{t=1}^{T} \mathbb{E}[C_t^{\mathrm{PF}} - F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t)) + F(\boldsymbol{\kappa}(t+1), \boldsymbol{Z}(t+1))] = \sum_{t=1}^{T} \mathbb{E}[\Delta_t^{\mathrm{PF}}].$$

Recall that the t-th period regret, Δ_t^{PF}, is nonnegative almost surely. We also note that for T large enough, events E_1, E_2, and E_3 happen with a positive probability independent of T, as proved in Lemma B.4. Note that $\mathbb{P}(E) = \mathbb{P}(E_1) \cap \mathbb{P}(E_2) \cap \mathbb{P}(E_3) = \mathbb{P}(E_1)\mathbb{P}(E_2)\mathbb{P}(E_3)$, because E_1, E_2 and E_3 are independent. Hence, event E happens with positive probability as well. Thus, it is sufficient to show that the expected total regret conditioned on E, scales as $\sqrt{T}$. For the rest of the proof, we will only consider scenarios occurring in E, and whenever an expectation is taken, it would be the conditional expectation with respect to E.

To establish the desired lower-bound of $\sum_{t=1}^{T} \mathbb{E}[\Delta_t^{\mathrm{PF}}]$ (conditioned on E), we will show that for T large enough, any $T'+1 \leq t \leq 2T'$, and any $\kappa_2(t)$,

$$\mathbb{E}\left[\Delta_t^{\mathrm{PF}} + \frac{\sum_{s=1}^{T} \Delta_s^{\mathrm{PF}}}{T'} \mid \kappa_2(t)\right] \geq \frac{C}{\sqrt{T}}, \tag{15}$$

where $C > 0$ is a fixed constant independent of T.

We will establish (15) by considering three cases of $\kappa_2(t)$, for any fixed $t \in [T'+1, 2T']$. Before analyzing each case, we make the following observation under any fixed arrival sequence (conditioned on E). Recall that $x_{ij}^{\mathrm{PF}}(t)$ denotes the indicator for whether PF fulfills the region-j order from warehouse i at period t, and let $(\boldsymbol{\ell}^{\mathrm{PF}}, \boldsymbol{x}^{\mathrm{PF}})$ represent the aggregated fulfillment decisions over all periods for the probabilistic fulfillment policy. Then, there exists a cost parameter $c_l > 0$ such that

$$\Delta_t^{\mathrm{PF}} \geq c_l, \text{ if } x_{21}^{\mathrm{PF}}(t) = 1 \text{ and } \kappa_2(t) \leq Z_2(t) \text{ for some fixed } t > 0, \tag{16}$$

$$\sum_{t=1}^{T} \Delta_t^{\mathrm{PF}} \geq c_l \cdot (\kappa_2^{\mathrm{PF}} - x_{22}^{\mathrm{PF}}). \tag{17}$$

The intuition behind this observation is that when $\kappa_2(t) \leq Z_2(t)$, then the optimal offline policy will use inventory at warehouse 2 for region 2, and this incurs c_l regret at period t whenever $x_{21}^{\mathrm{PF}}(t) = 1$.

Also, since $Z_2 \geq \kappa_2$ under event E, the aggregate regret over time is at least c_l times the number of times PF failed to use inventory at warehouse 2 for region 2. For a formal proof of this observation, we refer readers to Lemma B.5.

We will prove (15) by dividing all possible values of $\kappa_2(t)$ into three cases. For the first case, we assume $\kappa_2(t) \in X_1 := [0, \kappa_2^{\mathrm{PF}} - \lambda_2(t-1) + 3C^*\sqrt{T}]$. Then, for any $t \in (T', 2T']$ and $\kappa_2(t) \in X_1$ we have

$$\begin{aligned}
\sum_{s=1}^{t-1} x_{21}^{\mathrm{PF}}(s) &= \kappa_2^{\mathrm{PF}} - \kappa_2(t) - \sum_{s=1}^{t-1} x_{22}^{\mathrm{PF}}(s) \\
&\geq \kappa_2^{\mathrm{PF}} - \kappa_2(t) - B_2(t-1) \\
&\geq -3C^*\sqrt{T} + 6C^*\sqrt{T} - C^*\sqrt{T} \\
&= 2C^*\sqrt{T},
\end{aligned}$$

where the first inequality follows because $\sum_{s=1}^{t-1} x_{22}^{\mathrm{PF}}(s) \leq B_2(t-1)$, and the second inequality follows from E_1 and E_2. This implies that the total fulfillment for region 1 from warehouse 2, x_{21}^{PF}, satisfies $x_{21}^{\mathrm{PF}} \geq 2C^*\sqrt{T}$, implying $x_{22}^{\mathrm{PF}} \leq \kappa_2^{\mathrm{PF}} - 2C^*\sqrt{T}$.

Therefore, by (17), we have $\sum_{s=1}^{T} \Delta_s^{\mathrm{PF}} \geq 2c_l C^*\sqrt{T}$, which implies that

$$\mathbb{E}\left[\Delta_t^{\mathrm{PF}} + \frac{\sum_{s=1}^{T} \Delta_s^{\mathrm{PF}}}{T'} \,|\, \kappa_2(t)\right] \geq \frac{C}{\sqrt{T}},$$

for some constant $C > 0$, for any $\kappa_2(t) \in X_1$. This completes our analysis for the first case.

Consider case 2 where $\kappa_2(t) \in X_2 := [\kappa_2^{\mathrm{PF}} - \lambda_2(t-1) + 3C^*\sqrt{T}, \kappa_2^{\mathrm{PF}} - \lambda_2(t-1) + 100C^*\sqrt{T}]$. Then, we have that

$$\frac{\kappa_2(t)}{T-t+1} > \frac{\kappa_2^{\mathrm{PF}} - \lambda_2(t-1) + 3C^*\sqrt{T}}{T-t+1} \geq \frac{\lambda_2(T-t+1) + C^*\sqrt{T}}{T-t+1} \geq \lambda_2 + \frac{C^*\sqrt{T}}{T-t+1} \geq \lambda_2 + \frac{C^*}{\sqrt{T}}, \quad (18)$$

where the second inequality follows from $|\kappa_2^{\mathrm{PF}} - \lambda_2 T| \leq 2C^*\sqrt{T}$. In addition, note that $\kappa_1(t) + \kappa_2(t) \leq \gamma = \theta T \leq T - t$ for $t \leq 2T'$, implying that the capacity constraint for warehouse 2 in $F(\boldsymbol{\kappa}(t), \boldsymbol{\lambda}(T-t+1))$ is binding. This implies that any optimal solution $(\hat{\boldsymbol{\ell}}, \hat{\boldsymbol{x}})$ for $F(\boldsymbol{\kappa}(t), \boldsymbol{\lambda}(T-t+1))$ must satisfy

$$\hat{x}_{22} + \hat{x}_{21} = \kappa_2(t), \, \hat{x}_{22} = \lambda_2(T-t+1),$$

which, combined with (18), imply that

$$\phi_{21} = \frac{\hat{x}_{21}}{T-t+1} \geq \frac{C^*}{\sqrt{T}}, \quad (19)$$

where ϕ_{21} is the probability PF chooses to fulfill from warehouse 2, given the order at period t is from region 1.

In order to apply (16) to derive a lower bound for Δ_t^{PF}, we also need the condition $\kappa_2(t) \leq Z_2(t)$ in addition to (19). For this purpose, define event $E_{3,A}$ as

$$E_{3,A} = \left\{ B_2(T) - B_2(2T') \geq \lambda_2(T - 2T' + 1) + 104C^*\sqrt{T} \right\}.$$

Note that $E_{3,A} \subset E_3$, and by Lemma B.4, $\mathbb{P}(E_{3,A})$ is bounded away from 0 for T large enough, and conditioned on $E_{3,A}$, we have that $\kappa_2(t) \leq Z_2(t)$, as

$$\begin{aligned}
Z_2(t+1) &= B_2(T) - B_2(2T') + B_2(2T') - B_2(t) \\
&\geq \lambda_2(T - 2T' + 1) + 104C^*\sqrt{T} + (B_2(2T') - \lambda_2 T') - (B_2(t) - \lambda_2(t - T')) - \lambda_2(t - 2T') \\
&= \lambda_2(T - t + 1) + 104C^*\sqrt{T} + (B_2(2T') - \lambda_2 T') - (B_2(t) - \lambda_2(t - T')) \\
&\geq \lambda_2 T - \lambda_2(t-1) + 102C^*\sqrt{T} \\
&\geq \kappa_2^{\mathrm{PF}} - \lambda_2(t-1) + 100C^*\sqrt{T} \\
&\geq \kappa_2(t)
\end{aligned}$$

where the second inequality holds because both $B_2(2T') - \lambda_2 T'$ and $B_2(t-1) - \lambda_2(t-1-T')$ lie within the interval $[B_2(T') - C^*\sqrt{T}, B_2(T') + C^*\sqrt{T}]$, and the third inequality holds as $|\kappa_2^{\mathrm{PF}} - \lambda_2 T| \leq 2C^*\sqrt{T}$.

As $\kappa_2(t) \leq Z_2(t)$ under $E_{3,A}$, by (16), we have that conditioned on event $E_{3,A}$,

$$\mathbb{E}[\Delta_t^{\mathrm{PF}} \mid \kappa_2(t)] \geq \lambda_2 \phi_{21} \cdot \mathbb{P}(E_{3,A}) \cdot c_l,$$

whenever $\kappa_2(t) \in X_2$. By our lower bound for ϕ_{21} in (19), we obtain

$$\mathbb{E}[\Delta_t^{\mathrm{PF}} \mid \kappa_2(t)] \geq \lambda_2 \mathbb{P}(E_{3,A}) c_l \cdot \frac{C^*}{\sqrt{T}} > \frac{C}{\sqrt{T}},$$

for some constant $C > 0$. Thus, we have that

$$\mathbb{E}\left[\Delta_t^{\mathrm{PF}} + \frac{\sum_{s=1}^{T} \Delta_s^{\mathrm{PF}}}{T'} \mid \kappa_2(t)\right] \geq \frac{C}{\sqrt{T}} \text{ for any } \kappa_2(t) \in X_2.$$

In the last case, we assume $\kappa_2(t) \in X_3 := [\kappa_2^{\mathrm{PF}} - \lambda_2(t-1) + 100C^*\sqrt{T}, \infty)$. Then, it follows that

$$\sum_{s=1}^{t-1} x_{22}^{\mathrm{PF}}(s) \leq \kappa_2^{\mathrm{PF}} - \kappa_2(t) \leq \lambda_2(t-1) - 100C^*\sqrt{T} \leq B_2(t-1) - 92C^*\sqrt{T}, \tag{20}$$

where the last inequality holds due to E_1 and E_2.

Next, consider another event $E_{3,B}$, defined as

$$E_{3,B} = \left\{ \lambda_2(T-2T') + 10C^*\sqrt{T} \le B_2(T) - B_2(2T') \le \lambda_2(T-2T') + 11C^*\sqrt{T} \right\}.$$

Note that $E_{3,B} \subset E_3$, and by Lemma B.4, $\mathbb{P}(E_{3,B})$ is bounded away from zero for any T large enough. Conditioned on $E_{3,B}$, we obtain

$$\sum_{s=2T'+1}^{T} x_{22}^{\mathrm{PF}}(s) \le B_2(T) - B_2(2T') \le \lambda_2(T-2T') + 11C^*\sqrt{T}. \tag{21}$$

Finally, due to E_2, we have that

$$\sum_{s=t}^{2T'} x_{22}^{\mathrm{PF}}(s) \le B_2(2T') - B_2(t-1) \le \lambda_2(2T'-t+1) + 2C^*\sqrt{T}. \tag{22}$$

Combining (20), (21) and (22) suggests that

$$x_{22}^{\mathrm{PF}} = \sum_{s=1}^{T} x_{22}^{\mathrm{PF}}(s) \le B_2(t-1) - 79C^*\sqrt{T}.$$

Therefore, by (17), we have $\sum_{s=1}^{T} \Delta_s^{\mathrm{PF}} \ge 79c_l C^*\sqrt{T}$, which implies that

$$\mathbb{E}\left[\Delta_t^{\mathrm{PF}} + \frac{\sum_{s=1}^{T}\Delta_s^{\mathrm{PF}}}{T'} \mid \kappa_2(t)\right] \ge \mathbb{P}(E_{3,B}) \cdot 79c_l C^*\sqrt{T} \ge \frac{C}{\sqrt{T}},$$

for $\kappa_2(t) \in X_3$ and some constant $C > 0$. This completes our analysis for the last case.

Now that we established (15) for any given $\kappa_2(t)$, by the law of total expectation, we have

$$2\sum_{t=1}^{T} \mathbb{E}[\Delta_t^{\mathrm{PF}}] \ge \sum_{t=T'+1}^{2T'} \mathbb{E}\left[\Delta_t^{\mathrm{PF}} + \frac{\sum_{s=1}^{T}\Delta_s^{\mathrm{PF}}}{T'}\right] \ge T'\frac{C}{\sqrt{T}} = \frac{C(1-\theta)}{2}\cdot\sqrt{T},$$

for any κ_2^{PF} satisfying $|\kappa_2^{\mathrm{PF}} - \lambda_2 T| \le 2C^*\sqrt{T}$ and T large enough.

Q.E.D.

We next complete the proof for Theorem 3.1, which immediately follows from Propositions 4.3 and 4.4.

Proof of Theorem 3.1. In Example 4.1, by Proposition 4.3 and Proposition 4.4, there exists $C > 0$ and $C^* > 0$ such that

$$\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] - \mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}}, \boldsymbol{D})] \ge C\sqrt{T}, \text{ if } |\kappa_2^{\mathrm{PF}} - \lambda_2 T| \ge 2C^*\sqrt{T},$$
$$\mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] - \mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] \ge C\sqrt{T}, \text{ if } |\kappa_2^{\mathrm{PF}} - \lambda_2 T| \le 2C^*\sqrt{T}.$$

Because $\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] - \mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}}, \boldsymbol{D})] \ge 0$, $\mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] - \mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] \ge 0$. Therefore, we have the minimum-inventory regret, defined as $\mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] - \mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}}, \boldsymbol{D})]$, is at least $C\sqrt{T}$. Thus, Example 4.1 is a set of problem instances such that the minimum-inventory regret under PF is $\Omega(\sqrt{T})$. This completes the proof of Theorem 3.1. Q.E.D.

4.2. Minimum-Inventory Regret Analysis of Score-based Fulfillment Policy

To show Theorem 3.2, we will derive an upper bound on the fixed inventory regret of the score-based policy for *any* starting inventory vector $\boldsymbol{\kappa}$. Note that such a bound immediately implies an upper bound on the minimum-inventory regret, denoted $\min_{\boldsymbol{\kappa}\in U}\mathbb{E}[c^{\mathrm{SF}}(\boldsymbol{\kappa},\boldsymbol{D})]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})]$, because

$$\min_{\boldsymbol{\kappa}\in U}\mathbb{E}[c^{\mathrm{SF}}(\boldsymbol{\kappa},\boldsymbol{D})]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})]\le \mathbb{E}[c^{\mathrm{SF}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})].$$

Our analysis applies the compensated coupling technique introduced in (Vera and Banerjee, 2020; Vera et al., 2021). Let $c_{\max}:=\max_{ij} c_{ij}$. Recall that the fixed inventory regret given $\boldsymbol{\kappa}$ can be decomposed as

$$\mathbb{E}[c^{\mathrm{SF}}(\boldsymbol{\kappa},\boldsymbol{D})]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa},\boldsymbol{D})]=\sum_{t=1}^{T}\mathbb{E}[\Delta_t^{\mathrm{SF}}]\le c_{\max}\sum_{t=1}^{T}\mathbb{P}(\Delta_t^{\mathrm{SF}}>0). \tag{23}$$

In the analysis of Vera et al. (2021), it was demonstrated that the probability $\mathbb{P}(\Delta_t^{\mathrm{SF}}>0)$ is bounded by $C(T-t)^{-2}$, with C being a constant that does not depend on $T-t$ in a general resource allocation setting. However, the authors did not specify how C can be bounded in terms of other model parameters. Recently, it has been observed that C can scale exponentially with the dimension of the constraint matrix in a network revenue management problem (Jiang et al., 2022). In contrast, we derive an upper bound on $\mathbb{P}(\Delta_t^{\mathrm{SF}}>0)$ that decreases exponentially with respect to $T-t$, and explicit polynomial dependencies with I, J, and $\lambda_{\min}$. The upper bound is subsequently used to show that the regret of the score-based policy is not only independent with respect to T but also upper-bounded by $\tilde{O}(I^2J^2)$. Our analysis relies crucially on the network flow structure inherent in $F(\boldsymbol{\kappa},\boldsymbol{z})$.

To simplify our analysis, we start by rewriting the linear program in (1) by letting x_{0j} replace ℓ_j, c_{0j} replace $c_{\bullet j}$ for all $j\in[J]$:

$$\begin{aligned}
F(\boldsymbol{\kappa},\boldsymbol{z})=\min_{\boldsymbol{x}}\ & \sum_{i\in[I]\cup\{0\}}\sum_{j\in[J]}c_{ij}x_{ij} & \\
\text{s.t. } & \sum_{i\in[I]\cup\{0\}}x_{ij}\ge z_j & j\in[J]\\
& \sum_{j\in[J]}x_{ij}\le \kappa_i & i\in[I]\\
& x_{ij}\ge 0 & i\in[I]\cup\{0\},\ \ j\in[J].
\end{aligned} \tag{24}$$

A critical component of our analysis is the distance between the optimal solutions of (24) when $\boldsymbol{z}$ changes. Specifically, let $\bar{\boldsymbol{x}}$ be an optimal solution of (24). For any $\boldsymbol{\delta}\in\mathbb{R}^J$, we claim that there exists an optimal solution $\boldsymbol{x}^{\boldsymbol{\delta}}$ of $F(\boldsymbol{\kappa},\boldsymbol{z}+\boldsymbol{\delta})$ such that

$$|x_{ij}^{\boldsymbol{\delta}}-\bar{x}_{ij}|\le\sum_{j=1}^{J}|\boldsymbol{\delta}_j|=\|\boldsymbol{\delta}\|_1,\ \forall i\in[I]\cup\{0\},j\in[J]. \tag{25}$$

More general bounds on the sensitivity of the optimal solutions with respect to right-hand-side have been well established in the literature (see, e.g., Hoffman, 2003; Mangasarian and Shiau, 1987). The inequality we derived in (25) is obtained by applying the specific network flow structure of (24), and is thus tighter. In the interest of space, we leave the derivation of (25) as proof of Lemma B.7 in the appendix.

Next, we provide a bound on the probability that the per-period regret incurred by SF is greater than 0 for each $1 \le t \le T$.

Proposition 4.5. *For any $1 \le t \le T$, we have*

$$\mathbb{P}(\Delta_t^{\mathrm{SF}} > 0) \le 2J \exp\left(-2\left(\frac{\lambda_{\min}(T-t+1)}{J(I+1)} - \frac{1}{J}\right)^2 (T-t+1)^{-1}\right)$$

where $\lambda_{\min} = \min_{j \in [J]} \lambda_j$.

Proof. First, we observe that whenever $x_{ij}^{\mathrm{SF}}(t) = 1$ for some i and j, the regret Δ_t^{SF} is guaranteed to be zero as long as an optimal solution to $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t))$, denoted as $\bar{\boldsymbol{x}}$, satisfies $\bar{x}_{ij} \ge 1$ (see Lemma B.8 in appendix for a formal analysis). As a result, we have that

$$\begin{aligned}\mathbb{P}(\Delta_t^{\mathrm{SF}} > 0) &\le \max_{\boldsymbol{\kappa}, j} \mathbb{P}(\Delta_t^{\mathrm{SF}} > 0 \,|\, \boldsymbol{\kappa}(t) = \boldsymbol{\kappa}, D_j(t) = 1) \\ &\le \max_{\boldsymbol{\kappa}, j} \mathbb{P}(x_{\hat{i}j}^{\mathrm{SF}}(t) = 1, \bar{x}_{\hat{i}j} < 1 \,|\, \boldsymbol{\kappa}(t) = \boldsymbol{\kappa}, D_j(t) = 1) \end{aligned} \tag{26}$$

where the maximization is taken over all of the plausible $\boldsymbol{\kappa}$, $j \in [J]$, and $\hat{i}$ is the warehouse selected given that $\boldsymbol{\kappa}(t) = \boldsymbol{\kappa}$ and the demand at period t is from region j.

Now, fix any $\boldsymbol{\kappa}$, $j \in [J]$ and let E be the event such that $\boldsymbol{\kappa}(t) = \boldsymbol{\kappa}$ and $D_j(t) = 1$. We will now bound $\mathbb{P}(x_{\hat{i}j}^{\mathrm{SF}}(t) = 1, \bar{x}_{\hat{i}j} < 1 \,|\, E)$. To this end, let $\boldsymbol{\delta}(t) = \boldsymbol{Z}(t) - \boldsymbol{\lambda}(T-t+1)$. By (25), we have $|\hat{x}_{\hat{i}j} - \bar{x}_{\hat{i}j}| \le \|\boldsymbol{\delta}(t)\|_1$, where $\bar{\boldsymbol{x}}$ is an optimal solution to $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t))$. This consequently implies that $\bar{x}_{\hat{i}j} \ge 1$ if $\hat{x}_{\hat{i}j} \ge \|\boldsymbol{\delta}(t)\|_1 + 1$. Hence, we can further bound the probability that $x_{\hat{i}j}^{\mathrm{SF}}(t) = 1$ and $\bar{x}_{\hat{i}j} < 1$ given event E as follows:

$$\mathbb{P}\left(x_{\hat{i}j}^{\mathrm{SF}}(t) = 1, \bar{x}_{\hat{i}j} < 1 \,|\, E\right) \le \mathbb{P}\left(\hat{x}_{\hat{i}j} < \|\boldsymbol{\delta}(t)\|_1 + 1 \,|\, E\right). \tag{27}$$

Now, recall that the score-based fulfillment selects $\hat{x}_{\hat{i}j} = \max_{n \in I \cup \{0\}} \hat{x}_{nj}$, which implies that

$$\hat{x}_{\hat{i}j} \ge \sum_{n \in I \cup \{0\}} \hat{x}_{nj}/(I+1) = \lambda_j(T-t+1)/(I+1) \text{ almost surely.}$$

Therefore, the right-hand-side of (27) can be bounded as

$$\begin{aligned}\mathbb{P}\left(\hat{x}_{\hat{i}j} < \|\boldsymbol{\delta}(t)\|_1 + 1 \,|\, E\right) &\le \mathbb{P}\left(\frac{\lambda_j(T-t+1)}{I+1} < \|\boldsymbol{\delta}\|_1 + 1 \,|\, E\right) &(28) \\ &\le \mathbb{P}\left(\frac{\lambda_{\min}(T-t+1)}{I+1} < \|\boldsymbol{\delta}\|_1 + 1\right). &(29)\end{aligned}$$

Finally, by union bound and Hoeffding's inequality, we can bound (27), as

$$\begin{aligned}
\mathbb{P}\left(\frac{\lambda_{\min}(T-t+1)}{I+1} < ||\boldsymbol{\delta}(t)||_1 + 1\right) &\leq \mathbb{P}\left(\delta_j(t) > \frac{\lambda_{\min}(T-t+1)}{J(I+1)} - \frac{1}{J}, \forall j \in J\right) \\
&\leq \sum_{j=1}^{J} \mathbb{P}\left(\delta_j(t) > \frac{\lambda_{\min}(T-t+1)}{J(I+1)} - \frac{1}{J}\right) \\
&\leq 2J \exp\left(-2\left(\frac{\lambda_{\min}(T-t+1)}{J(I+1)} - \frac{1}{J}\right)^2 (T-t+1)^{-1}\right).
\end{aligned}$$

Q.E.D.

Next, we complete the proof of Theorem 3.2.

Proof of Theorem 3.2. By Proposition 4.5, we have

$$\mathbb{P}(\Delta_t^{\mathrm{SF}} > 0) \leq 2J \exp\left(-2\left(\frac{\lambda_{\min}(T-t+1)}{J(I+1)} - \frac{1}{J}\right)^2 (T-t+1)^{-1}\right).$$

For any $t \leq T - \frac{2(I+1)}{\lambda_{\min}} + 1$, we have $\frac{\lambda_{\min}(T-t+1)}{2J(I+1)} \geq \frac{1}{J}$. Thus,

$$\mathbb{P}(\Delta_t^{\mathrm{SF}} > 0) \leq 2J \exp\left(-\frac{\lambda_{\min}^2(T-t+1)}{J^2(I+1)^2}\right), \text{ for } t \leq T - \frac{2(I+1)}{\lambda_{\min}} + 1. \tag{30}$$

Furthermore, as $2\log(2) \geq 1$ and $\lambda_{\min} < 1$, if $t \leq T - \frac{2J^2(I+1)^2\log(2J)}{\lambda_{\min}^2} + 1$, then (30) implies that $\mathbb{P}(\Delta_t^{\mathrm{SF}} > 0) \leq 1$. Next, define

$$a := \frac{2J^2(I+1)^2\log(2J)}{\lambda_{\min}^2}, \text{ and } b := \frac{\lambda_{\min}^2}{J^2(I+1)^2}.$$

Considering the summation of $\mathbb{P}(\Delta_t^{\mathrm{SF}} > 0)$ from 1 to $T-a+1$, we have

$$\sum_{t=1}^{T-a+1} \mathbb{P}(\Delta_t^{\mathrm{SF}} > 0) \leq \sum_{s=a}^{T} 2J \exp\left(-\frac{\lambda_{\min}^2 s}{J^2(I+1)^2}\right) = \sum_{s=a}^{T} 2J\exp(-bs).$$

Note that $ab = 2\log(2J)$. Moreover, since $\lambda_{\min} \leq 1$, $I \geq 1$ and $J \geq 1$ it follows that $b \leq \frac{1}{4}$, which implies $e^b \leq 2$. Then, we have that

$$\sum_{t=1}^{T-a+1} \mathbb{P}(\Delta_t^{\mathrm{SF}} > 0) \leq \sum_{s=a}^{T} 2J\exp(-bs) = 2J\frac{e^{b-ab}}{e^b-1} = \frac{1}{2J}\frac{e^b}{e^b-1} \leq \frac{1}{J}\frac{1}{e^b-1} \leq \frac{b^{-1}}{J} \leq b^{-1} \tag{31}$$

where the second inequality is established by $e^b \leq 2$, the third inequality follows from $\exp(b) \geq b+1$ for any $b \neq 0$, and the last inequality holds because $J \geq 1$.

Recall that, by (23) the regret is bounded by $\sum_{t=1}^{T} c_{\max}\mathbb{P}(\Delta_t^{\mathrm{SF}} > 0)$. We finish the proof by dividing the sum for $t \leq T-a+1$ and $t > T-a+1$. From equation (31) and given that $\mathbb{P}(\Delta_t^{\mathrm{SF}} > 0) \leq 1$ for all t, the regret bound can be expressed as:

$$\sum_{t=1}^{T-a+1} c_{\max}\mathbb{P}(\Delta_t^{\mathrm{SF}} > 0) + \sum_{t=T-a+2}^{T} c_{\max}\mathbb{P}(\Delta_t^{\mathrm{SF}} > 0) \leq c_{\max}(b^{-1} + a) = c_{\max} \cdot \frac{J^2(I+1)^2}{\lambda_{\min}^2}(1 + 2\log(2J)).$$

This concludes the proof of Theorem 3.2. Q.E.D.

5. Numerical Examples

The analysis in Section 4 suggests that SF has smaller regret than PF when T grows to infinity. In this section, we test the performance of SF and PF over two sets of examples with finite values of T. For each example, we vary $T \in \{100, 200, 300, 400, 500\}$ while scaling the initial inventory placement as $\gamma = \theta T = 0.8T$, and for each value of T, we run 1000 simulations. Throughout this section, instead of calculating $\boldsymbol{\kappa}^{\mathrm{SF}}$, we simulate the performance of SF at the inventory placement $\boldsymbol{\kappa}^{\mathrm{off}}$ which already provides an upper-bound on the minimum-inventory regret of SF. We start by introducing our first example, which satisfies the conditions of Example 4.1.

Example 5.1. *Let the number of warehouses and the number of demand regions be $I = 2$, and $J = 2$, respectively. Let $c_{11} = 2$, $c_{21} = 3$, $c_{12} = 5$, $c_{22} = 1$, $c_{\bullet 1} = c_{\bullet 2} = 4$, and let the arrival probabilities satisfy $\lambda_1 = \lambda_2 = \frac{1}{2}$.*

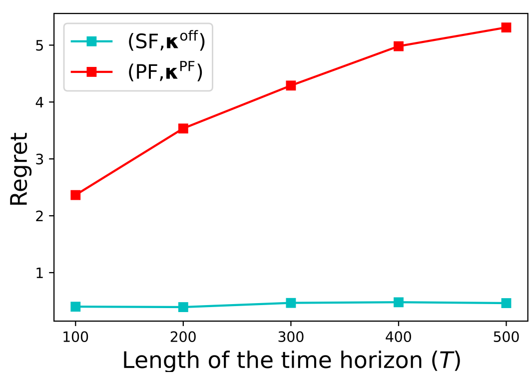

(a) Regret of $(\mathrm{SF}, \boldsymbol{\kappa}^{\mathrm{off}})$ and $(\mathrm{PF}, \boldsymbol{\kappa}^{\mathrm{PF}})$ with respect to different values of T for Example 5.1

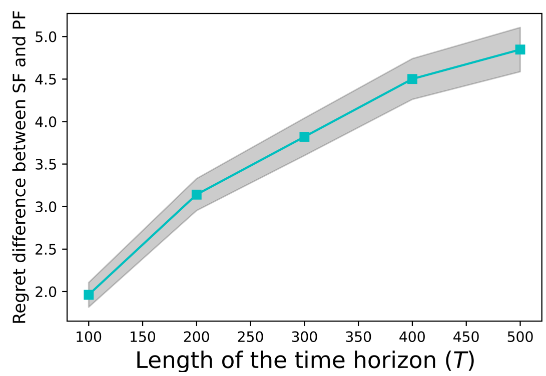

(b) The difference $\mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D}) - c^{\mathrm{SF}}(\boldsymbol{\kappa}^{\mathrm{off}}, \boldsymbol{D})]$ for varying values of T for Example 5.1. The gray shaded region specifies the 95% confidence band given the 1000 simulations.

Figure 1 **Regret of $(\mathrm{SF}, \boldsymbol{\kappa}^{\mathrm{off}})$ and $(\mathrm{PF}, \boldsymbol{\kappa}^{\mathrm{PF}})$ (left) and the regret difference between the policies SF and PF (right) for Example 5.1**

In our simulation, we first identify the value for $\boldsymbol{\kappa}^{\mathrm{PF}}$ that minimizes the expected cost for policy PF, by computing the expected cost for all feasible integral inventory placement vectors. Subsequently, we compute the difference in the expected costs of PF at $\boldsymbol{\kappa}^{\mathrm{PF}}$ and the expected cost of the offline policy at $\boldsymbol{\kappa}^{\mathrm{off}}$, termed as the minimum-inventory regret of PF. Similarly, we find the difference between the expected cost of SF under placement vector $\boldsymbol{\kappa}^{\mathrm{off}}$ and the expected cost of the offline policy at $\boldsymbol{\kappa}^{\mathrm{off}}$ which upper bounds the minimum-inventory regret of SF. To approximate $\boldsymbol{\kappa}^{\mathrm{off}}$, we utilize the sample average approximation method and solve $\min_{\boldsymbol{\kappa} \in U} \mathbb{E}[F(\boldsymbol{\kappa}, \boldsymbol{z})]$ by generating 1000 demand scenarios.

The graphical representation of the regrets for both PF and SF policies is illustrated in Figure 1a. This figure shows that, unlike the regret of the score-based fulfillment policy (SF), which remains stable as T increases, the regret for the probabilistic fulfillment policy (PF) grows with T as a concave function. Additionally, Figure 1b presents the difference in the regrets between PF and SF, including a 95% pointwise confidence band, illustrating that the difference between PF and SF is statistically significant for different values of T. We note that for this particular numerical example, $\boldsymbol{\kappa}^{\mathrm{PF}} = \boldsymbol{\kappa}^{\mathrm{off}}$ for all selected T.

For the next example, we generate locations for the demand regions and the warehouses at random within the unit square (See Figure 2). The cost structure is simply determined based on the distances between the demand locations and warehouses.

Example 5.2. *Let the number of warehouses and demand regions be $I = 3$, and $J = 5$, respectively. We let the lost sales cost be $c_{\bullet j} = 2$ for all $j \in [J]$. Set also arrival probabilities to be the same for any demand region, i.e., $\lambda_j = \frac{1}{5}$, for all $j \in [J]$.*

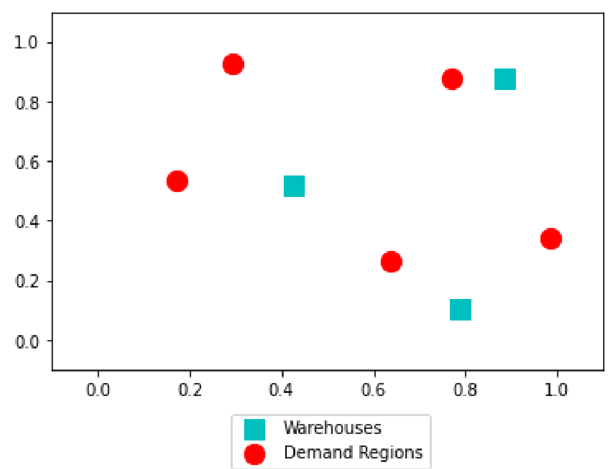

Figure 2 **Warehouse and Demand region locations in Example 5.2.**

For $I \geq 3$, precisely determining $\boldsymbol{\kappa}^{\mathrm{PF}}$ becomes computationally challenging due to the significantly larger search space. Therefore, in our numerical example, we resort to optimizing the inventory placement for PF over a smaller search space with a set of randomly generated inventory placement vectors. More specifically, we select 100 inventory placement vectors randomly over the region $[\boldsymbol{\kappa}_{\mathrm{off}} - \mathbf{1}\sqrt{T}, \boldsymbol{\kappa}_{\mathrm{off}} + \mathbf{1}\sqrt{T}]$ and also include $\boldsymbol{\kappa}^{\mathrm{off}}$ in our randomly generated set. By evaluating for T values in the set $\{100, 200, 300, 400, 500\}$, we obtain an approximate measure of the minimum-inventory regret for PF.

Similar to the first numerical example, Figure 3a suggests that the regret of SF remains stable with increasing T, whereas the regret of PF increases with T as a concave function. Moreover,

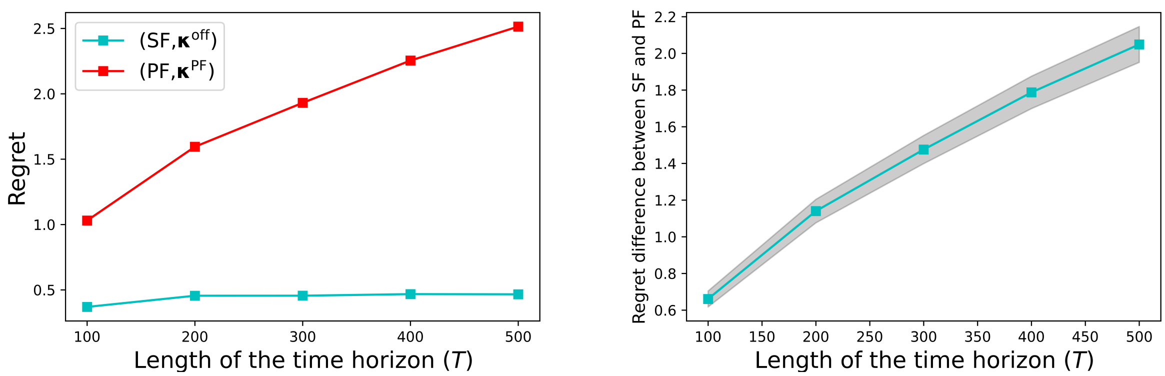

(a) "Approximate" minimum-inventory regret of SF and PF for Example 5.2

(b) The difference $\mathbb{E}[c^{\text{PF}}(\boldsymbol{\kappa}^{\text{PF}}, \boldsymbol{D}) - c^{\text{SF}}(\boldsymbol{\kappa}^{\text{off}}, \boldsymbol{D})]$ for varying values of T for Example 5.2. The gray shaded region specifies the 95% pointwise confidence band given the 1000 simulations.

Figure 3 Regret of $(\text{SF}, \boldsymbol{\kappa}^{\text{off}})$ and $(\text{PF}, \boldsymbol{\kappa}^{\text{PF}})$ (left) and the regret difference between the policies SF and PF (right) for Example 5.2

Figure 3b displays the difference in regrets between PF and SF, complete with a 95% pointwise confidence band, underscoring the statistically significant difference across all evaluated T values.

The aggregate numerical results from both examples align with our theoretical findings from the asymptotic analysis, highlighting that the SF fulfillment policy generally surpasses PF when considering endogenized inventory placements.

6. Conclusion

We study the performance of two state-of-the-art policies, probabilistic fulfillment and score-based fulfillment, on joint inventory placement and online demand fulfillment problem. We introduce a new regret metric called minimum-inventory regret, which factors in the inventory placement vector while evaluating the performance of the policies.

Our analysis shows that the probabilistic fulfillment policy has a minimum-inventory regret that scales with $\sqrt{T}$ as $T \to \infty$. Meanwhile, the minimum-inventory regret of score-based fulfillment policy is independent of T. Moreover, by utilizing the network flow structure of our problem, we show that the regret of score-based fulfillment is also polynomial with respect to the number of warehouses and demand regions. Our analysis suggests that utilizing the offline inventory placement with score-based fulfillment policy can outperform probabilistic fulfillment policy at its optimal inventory placement. Therefore, the score-based fulfillment policy is preferable in our joint inventory placement and online demand fulfillment problem setting when T is much larger than I and J. In addition to theoretical analysis, we also confirm our findings in several numerical simulations.

References

Jason Acimovic and Vivek F. Farias. The fulfillment-optimization problem. *Operations Research*, pages 218–237, 2019.

Jason Acimovic and Stephen C Graves. Making better fulfillment decisions on the fly in an online retail environment. *Manufacturing & Service Operations Management*, 17(1):34–51, 2015.

Jason Acimovic and Stephen C Graves. Mitigating spillover in online retailing via replenishment. *Manufacturing & Service Operations Management*, 19(3):419–436, 2017.

Ravindra K Ahuja, Thomas L Magnanti, and James B Orlin. *Network flows*. Cambridge, Mass.: Alfred P. Sloan School of Management, Massachusetts ..., 1988.

John M Andrews, Vivek F Farias, Aryan I Khojandi, and Chad M Yan. Primal–dual algorithms for order fulfillment at urban outfitters, inc. *INFORMS Journal on Applied Analytics*, 49(5): 355–370, 2019.

Alessandro Arlotto and Itai Gurvich. Uniformly bounded regret in the multisecretary problem. *Stochastic Systems*, 9(3):231–260, 2019.

Arash Asadpour, Xuan Wang, and Jiawei Zhang. Online resource allocation with limited flexibility. *Management Science*, 66(2):642–666, 2020.

Santiago R Balseiro, Omar Besbes, and Dana Pizarro. Survey of dynamic resource-constrained reward collection problems: Unified model and analysis. *Operations Research*, 2023.

Patrick Billingsley. *Probability and measure*. John Wiley & Sons, 2017.

Stéphane Boucheron, Gábor Lugosi, and Olivier Bousquet. Concentration inequalities. In *Summer school on machine learning*, pages 208–240. Springer, 2003.

Robert L Bray. Logarithmic regret in multisecretary and online linear programs with continuous valuations. 2023.

Pornpawee Bumpensanti and He Wang. A re-solving heuristic with uniformly bounded loss for network revenue management. *Management Science*, 66(7):2993–3009, 2020.

Annie I Chen and Stephen C Graves. Item aggregation and column generation for online-retail inventory placement. *Manufacturing & Service Operations Management*, 23(5):1062–1076, 2021.

Xingxing Chen, Jacob Feldman, Seung Hwan Jung, and Panos Kouvelis. Approximation schemes for the joint inventory selection and online resource allocation problem. *Production and Operations Management*, 31(8):3143–3159, 2022.

Levi DeValve, Yehua Wei, Di Wu, and Rong Yuan. Understanding the value of fulfillment flexibility in an online retailing environment. *Manufacturing & service operations management*, 25(2): 391–408, 2023.

Boris Epstein and Will Ma. Optimizing inventory placement for a downstream online matching problem. *arXiv preprint arXiv:2403.04598*, 2024.

Daniel Freund and Siddhartha Banerjee. Good prophets know when the end is near. *Available at SSRN 3479189*, 2019.

Daniel Freund and Jiayu Zhao. Overbooking with bounded loss. In *Proceedings of the 22nd ACM Conference on Economics and Computation*, pages 477–478, 2021.

Aravind Govindarajan, Amitabh Sinha, and Joline Uichanco. Joint inventory and fulfillment decisions for omnichannel retail networks. *Naval Research Logistics (NRL)*, 68(6):779–794, 2021.

Alan J Hoffman. On approximate solutions of systems of linear inequalities. In *Selected Papers Of Alan J Hoffman: With Commentary*, pages 174–176. World Scientific, 2003.

Stefanus Jasin and Sunil 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.

Stefanus Jasin and Amitabh Sinha. An LP-based correlated rounding scheme for multi-item ecommerce order fulfillment. *Operations Research*, 63(6):1336–1351, 2015.

Jiashuo Jiang, Will Ma, and Jiawei Zhang. Degeneracy is ok: Logarithmic regret for network revenue management with indiscrete distributions. *arXiv preprint arXiv:2210.07996*, 2022.

Will Ma. Order-optimal correlated rounding for fulfilling multi-item e-commerce orders. *Manufacturing & Service Operations Management*, 2023.

Olvi L Mangasarian and T-H Shiau. Lipschitz continuity of solutions of linear inequalities, programs and complementarity problems. *SIAM Journal on Control and Optimization*, 25(3):583–595, 1987.

Irina Shevtsova. On the absolute constants in the berry-esseen type inequalities for identically distributed summands. *arXiv preprint arXiv:1111.6554*, 2011.

Statista Research Department. Global retail e-commerce sales 2014-2025, 2021. https://www.statista.com/statistics/379046/worldwide-retail-e-commerce-sales/.

Statista Research Department. United states: retail e-commerce revenue 2017-2025, 2022. https://www.statista.com/chart/17207/amazon-shipping-and-fulfillment-costs/.

Alberto Vera and Siddhartha Banerjee. The bayesian prophet: A low-regret framework for online decision making. *Management Science*, 2020.

Alberto Vera, Siddhartha Banerjee, and Itai Gurvich. Online allocation and pricing: Constant regret via bellman inequalities. *Operations Research*, 69(3):821–840, 2021.

Alberto Vera, Alessandro Arlotto, Itai Gurvich, and Eli Levin. Dynamic resource allocation: the geometry and robustness of constant regret. *Mathematics of Operations Research*, forthcoming, 2023.

Yaqi Xie, Will Ma, and Linwei Xin. The benefits of delay to online decision-making. *Available at SSRN*, 2022.

Zhen Xu, Hailun Zhang, Jiheng Zhang, and Rachel Q Zhang. Online demand fulfillment under limited flexibility. *Management Science*, 66(10):4667–4685, 2020.

Yanyang Zhao, Xinshang Wang, and Linwei Xin. Multi-item online order fulfillment in a two-layer network. *Chicago Booth Research Paper*, (20-41), 2022.

Appendix A: Minimum-inventory Regret Upper Bound for PF

In this section, we show that the minimum-inventory regret of PF is $O(I\sqrt{T})$ in our online fulfillment problem which follows closely the arguments in Bumpensanti and Wang (2020) within the network revenue management setting, where they provide a fixed-inventory regret bound scaling with $\sqrt{T}$ for a policy analogous to PF.

Proposition A.1. *The minimum-inventory regret of* PF *satisfies the bound* $O(I\sqrt{T})$*. Specifically,*

$$\mathbb{E}\left[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}},\boldsymbol{D})\right]-\mathbb{E}\left[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})\right]=O(I\sqrt{T}). \tag{32}$$

Proof. Let $c^{\mathrm{DLP}}(\boldsymbol{\kappa},\boldsymbol{D})$ be the cost of the problem $T\cdot F(\boldsymbol{\kappa}/T,\boldsymbol{\lambda})$, which is formulated as

$$\begin{aligned}
\min_{\boldsymbol{x},\boldsymbol{\ell}}\ & T\left(\sum_{i\in[I]}\sum_{j\in[J]}c_{ij}x_{ij}+\sum_{j\in[J]}c_{\bullet j}\ell_j\right) & \\
\text{s.t. } & \sum_{i\in[I]}x_{ij}+\ell_j\geq\lambda_j & j\in[J]\\
& \sum_{j\in[J]}x_{ij}\leq\kappa_i/T & i\in[I]\\
& x_{ij}\geq 0 & i\in[I],\ \ j\in[J]\\
& \ell_j\geq 0 & j\in[J].
\end{aligned} \tag{33}$$

Note that $c^{\mathrm{DLP}}(\boldsymbol{\kappa},\boldsymbol{D})$ is an upper bound on the expected offline cost. Specifically,

$$\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})]=\mathbb{E}[F(\boldsymbol{\kappa},\boldsymbol{z})]\leq\mathbb{E}[F(\boldsymbol{\kappa},\boldsymbol{\lambda}T)]=TF(\boldsymbol{\kappa}/T,\boldsymbol{\lambda})=c^{\mathrm{DLP}}(\boldsymbol{\kappa},\boldsymbol{D}).$$

Since we can upper-bound the minimum-inventory regret of PF by the fixed inventory regret at $\boldsymbol{\kappa}^{\mathrm{off}}$, it follows that

$$\mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}},\boldsymbol{D})]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})]\leq\mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})]-\mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})]\leq\mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D})]-c^{\mathrm{DLP}}(\boldsymbol{\kappa}^{\mathrm{off}},\boldsymbol{D}). \tag{34}$$

Now, recall that, at each time t, PF solves $F(\boldsymbol{\kappa}(t),\boldsymbol{\lambda}(T-t+1))$. We can equivalently solve the problem $F(\boldsymbol{\kappa}(t)/(T-t+1),\boldsymbol{\lambda})$ which finds the expected per-period fulfillment cost between periods t and T instead. Let the optimal solution to this problem be $(\hat{x}(t),\hat{\ell}(t))$. Furthermore, let $(x^{\mathrm{DLP}},\ell^{\mathrm{DLP}})$ denote the optimal solution of the problem $T\cdot F(\boldsymbol{\kappa}/T,\lambda)$ given in (33). By Lemma A.2, the cost difference between $F(\boldsymbol{\kappa}/T,\lambda)$ and $F(\boldsymbol{\kappa}(t)/(T-t+1),\boldsymbol{\lambda})$ can be bounded as

$$\sum_{i\in[I]}\sum_{j\in[J]}c_{ij}\hat{x}_{ij}(t)+\sum_{j\in[J]}c_{\bullet j}\hat{\ell}_j(t)-\sum_{i\in[I]}\sum_{j\in[J]}c_{ij}x_{ij}^{\mathrm{DLP}}(t)-\sum_{j\in[J]}c_{\bullet j}\ell_j^{\mathrm{DLP}}(t)\leq\sum_{i\in[I]}c_{\max}\left|\frac{\kappa_i}{T}-\frac{\kappa_i(t)}{T-t+1}\right|. \tag{35}$$

We can also bound the expectation of $\left|\frac{\kappa_i}{T} - \frac{\kappa_i(t)}{T-t+1}\right|$ as follows:

$$\mathbb{E}\left[\left|\frac{\kappa_i}{T} - \frac{\kappa_i(t)}{(T-t+1)}\right|\right] \le \sqrt{\sum_{s=1}^{t-1} \frac{1}{(T-s)^2}}, \tag{36}$$

See Lemma A.3 for the proof of Inequality (36). Because the inequalities (35) and (36) hold for all $\boldsymbol{\kappa} \in U$, by combining (34), (35), and (36), we can bound the minimum-inventory regret of PF as follows:

$$\begin{aligned}
&E[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{PF}}, \boldsymbol{D})] - \mathbb{E}[c^{\mathrm{off}}(\boldsymbol{\kappa}^{\mathrm{off}}, \boldsymbol{D})] \le \mathbb{E}[c^{\mathrm{PF}}(\boldsymbol{\kappa}^{\mathrm{off}}, \boldsymbol{D})] - c^{\mathrm{DLP}}(\boldsymbol{\kappa}^{\mathrm{off}}, \boldsymbol{D}) \\
&= \mathbb{E}\left[\sum_{t\in[T]}\left(\sum_{i\in[I]}\sum_{j\in[J]} c_{ij}\hat{x}_{ij}(t) + \sum_{j\in[J]} c_{\bullet j}\hat{\ell}_j(t)\right) - \sum_{t\in[T]}\left(\sum_{i\in[I]}\sum_{j\in[J]} c_{ij}x_{ij}^{\mathrm{DLP}}(t) + \sum_{j\in[J]} c_{\bullet j}\ell_j^{\mathrm{DLP}}(t)\right)\right] \\
&\le \mathbb{E}\left[\sum_{t\in[T]}\sum_{i\in[I]} c_{\max}\left|\frac{\kappa_i}{T} - \frac{\kappa_i(t)}{T-t+1}\right|\right] \\
&\le \sum_{t\in[T]}\sum_{i\in[I]} c_{\max}\sqrt{\sum_{s=1}^{t-1}\frac{1}{(T-s)^2}} \\
&\le \sum_{t\in[T]}\sum_{i\in[I]} c_{\max}\sqrt{\int_1^{t-1}\frac{1}{(T-s)^2}ds} \\
&\le \sum_{i\in[I]} c_{\max}\sum_{t\in[T]}\sqrt{\frac{1}{T-t+1}} \\
&\le \sum_{i\in[I]} c_{\max}\int_1^T\sqrt{\frac{1}{T-t+1}}ds \\
&\le I c_{\max}(2\sqrt{T}-2).
\end{aligned}$$

Q.E.D.

Lemma A.2. *Let $\boldsymbol{\delta} \in \mathbb{R}^I$ such that $\boldsymbol{\kappa} + \boldsymbol{\delta} \ge 0$. Then,*

$$F(\boldsymbol{\kappa}, \boldsymbol{z}) - F(\boldsymbol{\kappa} + \boldsymbol{\delta}, \boldsymbol{z}) \le c_{\max}\sum_{i\in[I]}|\delta_i|.$$

Proof. First, we write the dual formulations for the problems $F(\boldsymbol{\kappa}, \boldsymbol{z})$ and $F(\boldsymbol{\kappa}+\boldsymbol{\delta}, \boldsymbol{z})$. Let $(\boldsymbol{\mu}, \boldsymbol{\nu})$, where $\boldsymbol{\mu} = [\mu]_{j\in[J]}$ and $\boldsymbol{\nu} = [\nu]_{i\in[I]}$ be the dual variables corresponding to the demand, and inventory constraints, respectively. The dual problems are then defined as follows:

$$D(\boldsymbol{\kappa}, \boldsymbol{z}) = \left\{\max\left(\sum_{j\in[J]} z_j\mu_j - \sum_{i\in[I]}\kappa_i\nu_i\right) \text{ s.t. } \mu_j - \nu_i \le c_{ij},\quad \mu_j \le c_{\bullet j},\quad \mu_j \ge 0,\quad \nu_i \ge 0\right\}$$

and

$$D(\boldsymbol{\kappa}+\boldsymbol{\delta}, \boldsymbol{z}) = \left\{\max\left(\sum_{j\in[J]} z_j\mu_j - \sum_{i\in[I]}(\kappa_i+\delta_i)\nu_i\right) \text{ s.t. } \mu_j - \nu_i \le c_{ij},\quad \mu_j \le c_{\bullet j},\quad \mu_j \ge 0,\quad \nu_i \ge 0\right\}.$$

Let $(\boldsymbol{\mu}^*, \boldsymbol{\nu}^*)$ be an optimal solution of the problem $D(\boldsymbol{\kappa}, \boldsymbol{z})$. Since the feasible regions of $D(\boldsymbol{\kappa}, \boldsymbol{z})$ and $D(\boldsymbol{\kappa}+\boldsymbol{\delta}, \boldsymbol{z})$ are the same, $(\boldsymbol{\mu}^*, \boldsymbol{\nu}^*)$ is feasible for $D(\boldsymbol{\kappa}+\boldsymbol{\delta}, \boldsymbol{z})$. Then,

$$F(\boldsymbol{\kappa}+\boldsymbol{\delta}, \boldsymbol{z}) \ge \sum_{j\in[J]} z_j\mu_j^* - \sum_{i\in[I]}(\kappa_i+\delta_i)\nu_i^* = F(\boldsymbol{\kappa}, \boldsymbol{z}) - \sum_{i\in[I]}\delta_i\nu_i^*,$$

where the inequality follows from the weak duality and the equality follows from the strong duality. Then, we have

$$F(\boldsymbol{\kappa},\boldsymbol{z}) - F(\boldsymbol{\kappa}+\boldsymbol{\delta},\boldsymbol{z}) \leq \sum_{i\in[I]} \delta_i \nu_i^* \leq \max_{i\in[I]} \nu_i^* \sum_{i\in I} |\delta_i|. \tag{37}$$

Next, we will bound ν_i^*, for all $i \in [I]$. Since $(\boldsymbol{\mu}^*, \boldsymbol{\nu}^*)$ is an optimal solution to $D(\boldsymbol{\kappa},\boldsymbol{z})$, we can plug in $\boldsymbol{\mu}^*$ to $D(\boldsymbol{\kappa},\boldsymbol{z})$ and solve the following optimization problem to find $\boldsymbol{\nu}^*$:

$$\left\{ \max \left(\sum_{j\in[J]} z_j \mu_j^* - \sum_{i\in[I]} \kappa_i \nu_i \right) \text{ s.t. } \mu_j^* - \nu_i \leq c_{ij}, \quad \nu_i \geq 0 \right\}.$$

Since $\kappa_i \geq 0$ and $\nu_i \geq 0$, an optimal solution to this problem is $\nu_i^* = \max_{j\in[J]} (\mu_j^* - c_{ij})^+$ for all $i \in [I]$. By the feasibility of μ_j^*, we have $0 \leq \mu_j^* \leq c_{\bullet j}$ which implies

$$\nu_i^* = \max_{j\in[J]} (\mu_j^* - c_{ij})^+ \leq \max_{j\in[J]} |c_{\bullet j} - c_{ij}| \leq c_{\max},$$

for all $i \in [I]$. Combining this with Inequality (37), we obtain

$$F(\boldsymbol{\kappa},\boldsymbol{z}) - F(\boldsymbol{\kappa}+\boldsymbol{\delta},\boldsymbol{z}) \leq c_{\max} \sum_{i\in[I]} |\delta_i|.$$

Q.E.D.

Lemma A.3. *Let $\kappa_i(t)$ be the remaining inventory in warehouse i in the beginning of time t, under* PF *policy. Then,*

$$\mathbb{E}\left[\left| \frac{\kappa_i}{T} - \frac{\kappa_i(t)}{(T-t+1)} \right| \right] \leq \sqrt{\sum_{s=1}^{t} \frac{1}{(T-s)^2}}.$$

Proof. By telescoping sum, we have that

$$\left| \frac{\kappa_i}{T} - \frac{\kappa_i(t)}{T-t+1} \right| \leq \left| \sum_{s=1}^{t-1} \left(\frac{\kappa_i(s)}{T-s+1} - \frac{\kappa_i(s+1)}{T-s} \right) \right|. \tag{38}$$

Note that $\frac{1}{T-s+1} = \frac{1}{T-s} - \frac{1}{(T-s+1)(T-s)}$. Hence, we can write (38) as

$$\left| \frac{\kappa_i}{T} - \frac{\kappa_i(t)}{T-t+1} \right| = \left| \sum_{s=1}^{t-1} \left(\frac{\kappa_i(s) - \kappa_i(s+1)}{T-s} - \frac{\kappa_i(s)}{(T-s)(T-s+1)} \right) \right|.$$

We have $\sum_{j\in[J]} x_{ij}^{\mathrm{PF}}(t) = \kappa_i(s) - \kappa_i(s+1)$, and the optimal solution to $F(\boldsymbol{\kappa}(t)/(T-t+1), \boldsymbol{\lambda})$ satisfies $\sum_{j\in[J]} \hat{x}_{ij}(t) \leq \kappa_i(t)/(T-t+1)$. Then,

$$\left| \frac{\kappa_i}{T} - \frac{\kappa_i(t)}{T-t+1} \right| \leq \left| \sum_{s=1}^{t-1} \left(\frac{\sum_{j\in[J]} x_{ij}^{\mathrm{PF}}(s)}{T-s} - \frac{\sum_{j\in[J]} \hat{x}_{ij}(s)}{T-s} \right) \right|.$$

By taking the expectation on both sides and applying Cauchy-Shwarz inequality, we obtain

$$\mathbb{E}\left[\left| \frac{\kappa_i}{T} - \frac{\kappa_i(t)}{T-t+1} \right| \right] \leq \sqrt{\mathbb{E}\left[\left(\sum_{s=1}^{t-1} \frac{\sum_{j\in[J]} x_{ij}^{\mathrm{PF}}(s) - \sum_{j\in[J]} \hat{x}_{ij}(s)}{T-s} \right)^2 \right]}. \tag{39}$$

Let us define $\frac{\sum_{j\in[J]} x_{ij}^{\mathrm{PF}}(s) - \sum_{j\in[J]} \hat{x}_{ij}(s)}{T-s} := a_s$. Furthermore, let $(\mathcal{F}_s)_{s\in[T]}$ denote the filtration generated by the history up to and *not* including time s. Then, we can rewrite (39) as

$$\begin{aligned}
\mathbb{E}\left[\left| \frac{\kappa_i}{T} - \frac{\kappa_i(t)}{T-t+1} \right| \right] \leq \sqrt{\mathbb{E}\left[\left(\sum_{s=1}^{t-1} a_s \right)^2 \right]} &= \sqrt{\sum_{s=1}^{t-1} \mathbb{E}\left[a_s^2\right] + 2 \sum_{s=1}^{t-1} \sum_{l=s+1}^{t-1} \mathbb{E}\left[a_s a_l\right]} \\
&= \sqrt{\sum_{s=1}^{t-1} \mathbb{E}\left[\mathbb{E}\left[a_s^2 | \mathcal{F}_s\right]\right] + 2 \sum_{s=1}^{t-1} \sum_{l=s+1}^{t-1} \mathbb{E}\left[\mathbb{E}\left[a_s a_l | \mathcal{F}_l\right]\right]},
\end{aligned}$$

where the last equality follows by the law of total expectation. Note that, for any $s \in [T]$, $\mathbb{E}[\sum_{j\in[J]} x_{ij}^{\mathrm{PF}}(s)|\mathcal{F}_s] = \sum_{j\in[J]} \hat{x}_{ij}(s)$ which follows by the definition of PF policy. This suggests that $\mathbb{E}[a_s|\mathcal{F}_s] = 0$ for all $s \in [T]$. Consequently, it follows that $\mathbb{E}[a_s a_l|\mathcal{F}_l] = a_s \mathbb{E}[a_l|\mathcal{F}_l] = 0$. Then,

$$
\begin{aligned}
\mathbb{E}\left[\left|\frac{\kappa_i}{T} - \frac{\kappa_i(t)}{T-t+1}\right|\right] &\leq \sqrt{\sum_{s=1}^{t-1} \mathbb{E}\left[\mathbb{E}\left[a_s^2|\mathcal{F}_s\right]\right]} \\
&= \sqrt{\mathbb{E}\left[\sum_{s=1}^{t-1} \mathbb{E}\left[\left(\frac{\sum_{j\in[J]} x_{ij}^{\mathrm{PF}}(s) - \sum_{j\in[J]} \hat{x}_{ij}(s)}{T-s}\right)^2 \middle| \mathcal{F}_s\right]\right]} \\
&= \sqrt{\mathbb{E}\left[\sum_{s=1}^{t-1} \frac{1}{(T-s)^2} \mathbb{E}\left[\left(\sum_{j\in[J]} x_{ij}^{\mathrm{PF}}(s) - \sum_{j\in[J]} \hat{x}_{ij}(s)\right)^2 \middle| \mathcal{F}_s\right]\right]} \\
&= \sqrt{\sum_{s=1}^{t-1} \frac{1}{(T-s)^2} \mathbb{E}\left[\mathrm{Var}\left[\left(\sum_{j\in[J]} x_{ij}^{\mathrm{PF}}(s) - \sum_{j\in[J]} \hat{x}_{ij}(s)\right) \middle| \mathcal{F}_s\right]\right]} \\
&\leq \sqrt{\sum_{s=1}^{t-1} \frac{1}{(T-s)^2}},
\end{aligned}
$$

where the third equality follows by the linearity of the expectation, the fourth equality follows since $\mathbb{E}[\sum_{j\in[J]} x_{ij}^{\mathrm{PF}}(s) - \sum_{j\in[J]} \hat{x}_{ij}(s)|\mathcal{F}_s] = 0$ and the last inequality follows because $\mathrm{Var}(\sum_{j\in[J]} x_{ij}^{\mathrm{PF}}(s) - \sum_{j\in[J]} \hat{x}_{ij}(s)|\mathcal{F}_s) \leq 1$.

Q.E.D.

Appendix B: Additional Lemmas for Section 4

Lemma B.1. $\mathbb{P}(\Delta_t^\pi \geq 0) = 1$.

Proof. Let $j^* \in [J]$ represent the region where the demand arrives at time t, and let $i^* \in I \cup \{\bullet\}$ denote the fulfillment decision made by policy π for that time period (where $\{\bullet\}$ refers to the lost sales). This implies that the cost incurred by policy π at time t is $C_t^\pi = c_{i^*j^*}$. Moreover, let $(\bar{x}(t+1), \bar{\ell}(t+1))$ be the optimal solution to $F(\boldsymbol{\kappa}(t+1), \boldsymbol{Z}(t+1))$. Let us define a solution $\tilde{x}$ such that $\tilde{x}_{i^*j^*} = \bar{x}_{i^*j^*}(t+1) + 1$ and $\tilde{x}_{ij} = \bar{x}_{ij}(t+1)$ otherwise. Then, $\tilde{x}$ is feasible for $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t))$ since $\kappa_{i^*}(t) = \kappa_{i^*}(t+1) + 1$ and $Z_{j^*}(t) = Z_{j^*}(t+1) + 1$. Then, the t-th period regret can be bounded as:

$$
\begin{aligned}
\Delta_t^\pi &= C_t^\pi - F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t)) + F(\boldsymbol{\kappa}(t+1), \boldsymbol{Z}(t+1)) \\
&\geq C_t^\pi - \left(\sum_{i\in[I]}\sum_{j\in[J]} c_{ij}\tilde{x}_{ij}(t+1) + \sum_{j\in[J]} c_{\bullet j}\tilde{\ell}_j(t+1)\right) + \sum_{i\in[I]}\sum_{j\in[J]} c_{ij}\bar{x}_{ij}(t+1) + \sum_{j\in[J]} c_{\bullet j}\bar{\ell}_j(t+1) \\
&= C_t^\pi - \left(\sum_{i\in[I]}\sum_{j\in[J]} c_{ij}\bar{x}_{ij}(t+1) + \sum_{j\in[J]} c_{\bullet j}\bar{\ell}_j(t+1) + c_{i^*j^*}\right) + \sum_{i\in[I]}\sum_{j\in[J]} c_{ij}\bar{x}_{ij}(t+1) + \sum_{j\in[J]} c_{\bullet j}\bar{\ell}_j(t+1) \\
&= C_t^\pi - c_{i^*j^*} \\
&= 0.
\end{aligned}
$$

Q.E.D.

B.1. Lemmas for Section 4.1

Lemma B.2. *Under Example 4.1, the optimal solution to $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t))$ is unique and satisfies the following:*

1. *If $Z_1(t) > \kappa_1(t)$, and $Z_2(t) > \kappa_2(t)$:*

$$\bar{x}_{11} = \kappa_1(t),\ \bar{x}_{12} = 0,\ \bar{x}_{21} = 0,\ \bar{x}_{22} = \kappa_2(t),\ \bar{x}_{\bullet 1} = Z_1(t) - \kappa_1(t),\ \bar{x}_{\bullet 2} = Z_2(t) - \kappa_2(t). \tag{40}$$

2. *If $Z_1(t) \le \kappa_1(t)$, and $Z_2(t) > \kappa_2(t)$:*

$$\bar{x}_{11} = Z_1(t),\ \bar{x}_{12} = 0,\ \bar{x}_{21} = 0,\ \bar{x}_{22} = \kappa_2(t),\ \bar{x}_{\bullet 1} = 0,\ \bar{x}_{\bullet 2} = Z_2(t) - \kappa_2(t). \tag{41}$$

3. *If $Z_1(t) > \kappa_1(t)$, and $Z_2(t) \le \kappa_2(t)$:*
   - *3a. When $\kappa_2(t) - Z_2(t) > Z_1(t) - \kappa_1(t)$:*

$$\bar{x}_{11} = \kappa_1(t),\ \bar{x}_{12} = 0,\ \bar{x}_{21} = Z_1(t) - \kappa_1(t),\ \bar{x}_{22} = Z_2(t),\ \bar{x}_{\bullet 1} = 0,\ \bar{x}_{\bullet 2} = 0. \tag{42}$$

   - *3b. When $\kappa_2(t) - Z_2(t) \le Z_1(t) - \kappa_1(t)$:*

$$\bar{x}_{11} = \kappa_1(t),\ \bar{x}_{12} = 0,\ \bar{x}_{21} = \kappa_2(t) - Z_2(t),\ \bar{x}_{22} = Z_2(t),\ \bar{x}_{\bullet 1} = Z_1(t) + Z_2(t) - \kappa_1(t) - \kappa_2(t),\ \bar{x}_{\bullet 2} = 0. \tag{43}$$

4. *If $Z_1(t) \le \kappa_1(t)$, and $Z_2(t) \le \kappa_2(t)$:*

$$\bar{x}_{11} = Z_1(t),\ \bar{x}_{12} = 0,\ \bar{x}_{21} = 0,\ \bar{x}_{22} = Z_2(t),\ \bar{x}_{\bullet 1} = 0,\ \bar{x}_{\bullet 2} = 0. \tag{44}$$

Proof. For each of the cases, it is easy to check that the solutions provided in the lemma are feasible. Next, we will show that the solution is optimal and unique using the cost structure in Example 4.1 via complementary slackness. To show this, we will start by writing the dual of Example 4.1:

$$\begin{aligned}
\max\ \ & y_1 Z_1(t) + y_2 Z_2(t) - y_3 \kappa_1(t) - y_4 \kappa_2(t) \\
\text{s.t. } & y_1 - y_3 \le c_{11} \\
& y_1 - y_4 \le c_{21} \\
& y_2 - y_3 \le c_{12} \\
& y_2 - y_4 \le c_{22} \\
& y_1 \le c_{\bullet 1} \\
& y_2 \le c_{\bullet 2} \\
& y_1, y_2, y_3, y_4 \ge 0.
\end{aligned}$$

First, consider Case 1. One can check that the following is a dual feasible solution:

$$\bar{y}_1 = c_{\bullet 1},\ \bar{y}_2 = c_{\bullet 2}, \bar{y}_3 = c_{\bullet 1} - c_{11},\ \bar{y}_4 = c_{\bullet 2} - c_{22}. \tag{45}$$

By complementary slackness, we have the following:

$$\bar{x}_{11}(\bar{y}_1 - \bar{y}_3 - c_{11}) = 0, \qquad \bar{x}_{21}(\bar{y}_1 - \bar{y}_4 - c_{21}) = 0, \qquad \bar{x}_{12}(\bar{y}_2 - \bar{y}_3 - c_{12}) = 0, \tag{46}$$

$$\bar{x}_{22}(\bar{y}_2 - \bar{y}_4 - c_{22}) = 0, \qquad \bar{x}_{\bullet 1}(\bar{y}_1 - c_{\bullet 1}) = 0, \qquad \bar{x}_{\bullet 2}(\bar{y}_2 - c_{\bullet 2}) = 0, \tag{47}$$

and

$$\bar{y}_1(Z_1(t) - \bar{x}_{11} - \bar{x}_{21} - \bar{x}_{\bullet 1}) = 0, \qquad \bar{y}_2(Z_2(t) - \bar{x}_{22} - \bar{x}_{12} - \bar{x}_{\bullet 2}) = 0, \tag{48}$$

$$\bar{y}_3(\kappa_1(t) - \bar{x}_{11} - \bar{x}_{12}) = 0, \qquad \bar{y}_4(\kappa_2(t) - \bar{x}_{22} - \bar{x}_{21}) = 0. \tag{49}$$

Note that $\bar{x}$ provided in (40) and $\bar{y}$ provided in (45) satisfy the complementary slackness conditions provided above. Hence, $\bar{x}$ and $\bar{y}$ are optimal for primal and dual problems, respectively. To show the uniqueness of $\bar{x}$, we will show that $\bar{y}$ is a nondegenerate optimal dual solution. Observe that the second and third constraints of the dual problem are not binding since $\bar{y}_1 - \bar{y}_4 = c_{\bullet 1} - c_{\bullet 2} + c_{22} = c_{22} < c_{21}$ and $y_2 - y_3 = c_{\bullet 2} - c_{\bullet 1} + c_{11} = c_{11} < c_{12}$. This implies that only four constraints of the dual problem are binding which suggests that $\bar{y}$ is a nondegenerate dual solution. This shows the uniqueness of $\bar{x}$. The uniqueness and optimality of $\bar{x}$ in Case 2 can be shown by considering the following dual solution:

$$\bar{y}_1 = c_{11},\ \bar{y}_2 = c_{\bullet 2},\ \bar{y}_3 = 0,\ \bar{y}_4 = c_{\bullet 2} - c_{22}. \tag{50}$$

While the demonstration of optimality and uniqueness for Case 3 and Case 4 have been omitted, their proofs can be established using analogous methods.

Q.E.D.

Lemma B.3. *Under the settings of Example 4.1, the cost function f satisfies the following first-difference properties:*

$$f(\kappa_2+1) - f(\kappa_2) \geq (c_{21} - c_{11}) - (c_{21} + c_{\bullet 1}) \exp\left\{-2\frac{(\kappa_2 - \lambda_2 T)^2}{T}\right\} \qquad \text{for } \kappa_2 > \lambda_2 T; \tag{51}$$

and

$$f(\kappa_2+1) - f(\kappa_2) \leq -(c_{11} - c_{22}) + (c_{21} - c_{22}) \exp\left\{-2\frac{(\lambda_2 T - \kappa_2)^2}{T}\right\} \qquad \text{for } \kappa_2 < \lambda_2 T. \tag{52}$$

Proof. We first focus on proving that (51) holds. We begin by observing that the cost function f defined in (6) can be written as

$$\begin{aligned} f(\kappa_2) =& c_{22}\mathbb{E}[\min\{Z_2, \kappa_2\}] + c_{11}\mathbb{E}[\min\{Z_1, \gamma - \kappa_2\}] + c_{21}\mathbb{E}[\min\{(Z_1 - \gamma + \kappa_2)^+, (\kappa_2 - Z_2)^+\}] \\ &+ c_{\bullet 1}\mathbb{E}[\max\{Z_1 - \gamma + \kappa_2 - (\kappa_2 - Z_2)^+), 0\}] + c_{\bullet 2}\mathbb{E}[(Z_2 - \kappa_2)^+]. \end{aligned} \tag{53}$$

In the above expression, the first and second terms respectively correspond to the cost of using warehouse-2 inventory to fulfill region-2 demand and warehouse-1 inventory to fulfill region-1 demand. The third term accounts for the scenario in which region-1 demand is fulfilled with the remaining inventory of warehouse 2. The fourth and fifth terms represent the cost of expected lost sales from region 1 and region 2. Notably, there is no term involving c_{12}; Example 4.1 assumes that $c_{\bullet 2} \leq c_{12}$, so it is less costly to lose the demand from region 2 rather than fulfilling it from warehouse 1.

Because $\gamma = \theta T \leq T$ and $Z_1 + Z_2 = T$, we have that $\kappa_2 - Z_2 = \kappa_2 - T + Z_1 \leq \kappa_2 - \gamma + Z_1$, so we see that the third term in (53) is equal to its right minimand: $\mathbb{E}[\min\{(Z_1 - \gamma + \kappa_2)^+, (\kappa_2 - Z_2)^+\}] = \mathbb{E}[(\kappa_2 - Z_2)^+]$.

Moreover, as $c_{\bullet 1} = c_{\bullet 2}$, we can rewrite the sum of the last two terms in (53) as

$$\begin{aligned} &c_{\bullet 1}\mathbb{E}[\max\{Z_1 - \gamma + \kappa_2 - (\kappa_2 - Z_2)^+, 0\}] + c_{\bullet 2}\mathbb{E}[(Z_2 - \kappa_2)^+] \\ &\qquad = c_{\bullet 1}\mathbb{E}[\max\{Z_1 - \gamma + \kappa_2 - (\kappa_2 - Z_2)^+, 0\} + (Z_2 - \kappa_2)^+] \\ &\qquad = c_{\bullet 1}\mathbb{E}[\max\{Z_1 - \gamma + \kappa_2 - (\kappa_2 - Z_2), (Z_2 - \kappa_2)^+\}] \\ &\qquad = c_{\bullet 1}\left(T - \gamma + \mathbb{E}[(Z_2 - \kappa_2 - T + \gamma)^+)]\right) \\ &\qquad = c_{\bullet 1}\left(T - \gamma + \mathbb{E}[(\gamma - \kappa_2 - Z_1)^+)]\right). \end{aligned}$$

Hence, after rearranging, the cost function $f(\kappa_2)$ can be written as

$$\begin{aligned} f(\kappa_2) =& c_{22}\mathbb{E}[\min\{Z_2, \kappa_2\}] + c_{11}\mathbb{E}[\min\{Z_1, \gamma - \kappa_2\}] \\ &+ c_{21}\mathbb{E}[(\kappa_2 - Z_2)^+] + c_{\bullet 1}\left(T - \gamma + \mathbb{E}[(\gamma - \kappa_2 - Z_1)^+)]\right). \end{aligned} \tag{54}$$

The representation (54) allows us to study the difference $f(\kappa_2+1)-f(\kappa_2)$ by examining how much each term on the right-hand side changes when κ_2 increases by one. For the first term of (54), we have

$$\begin{aligned}
c_{22}\mathbb{E}[\min\{Z_2,\kappa_2+1\}]-c_{22}\mathbb{E}[\min\{Z_2,\kappa_2\}] &= \sum_{z=0}^{\infty}(\min\{z,\kappa_2+1\}-\min\{z,\kappa_2\})\,\mathbb{P}(Z_2=z)\\
&= \sum_{z=\kappa_2+1}^{\infty}\mathbb{P}(Z_2=z)\\
&= c_{22}\mathbb{P}(Z_2>\kappa_2).
\end{aligned}$$

Similarly, for the remaining terms, we also get

$$\begin{aligned}
c_{11}\mathbb{E}[\min\{Z_1,\gamma-(\kappa_2+1)\}]-c_{11}\mathbb{E}[\min\{Z_1,\gamma-\kappa_2\}] &= -c_{11}\mathbb{P}(Z_1\geq\gamma-\kappa_2);\\
c_{21}\mathbb{E}[(\kappa_2+1-Z_2)^+]-c_{21}\mathbb{E}[(\kappa_2-Z_2)^+] &= c_{21}\mathbb{P}(Z_2\leq\kappa_2);\ \text{and}\\
c_{\bullet 1}\left(T-\gamma+\mathbb{E}[(\gamma-(\kappa_2+1)-Z_1)^+)]\right)-c_{\bullet 1}\left(T-\gamma+\mathbb{E}[(\gamma-\kappa_2-Z_1)^+)]\right) &= -c_{\bullet 1}\mathbb{P}(Z_1<\gamma-\kappa_2).
\end{aligned}$$

By recalling (54) and adding up the four terms above, we obtain

$$f(\kappa_2+1)-f(\kappa_2)=c_{22}\mathbb{P}(Z_2>\kappa_2)-c_{11}\mathbb{P}(Z_1\geq\gamma-\kappa_2)+c_{21}\mathbb{P}(Z_2\leq\kappa_2)-c_{\bullet 1}\mathbb{P}(Z_1<\gamma-\kappa_2). \tag{55}$$

Next, we obtain a sequence of lower bounds for $f(\kappa_2+1)-f(\kappa_2)$ that will eventually lead to (51). Dropping the first addend on the right-hand side of (55), and replacing $\mathbb{P}(Z_1\geq\gamma-\kappa_2)$ with 1, we get

$$f(\kappa_2+1)-f(\kappa_2)\geq -c_{11}+c_{21}\mathbb{P}(Z_2\leq\kappa_2)-c_{\bullet 1}\mathbb{P}(Z_1<\gamma-\kappa_2).$$

Continuing from this inequality, recall that $Z_1=T-Z_2$ and $\gamma\leq T$, we have $\mathbb{P}(Z_1<\gamma-\kappa_2)=\mathbb{P}(T-Z_2<\gamma-\kappa_2)=\mathbb{P}(Z_2>\kappa_2+T-\gamma)\leq\mathbb{P}(Z_2>\kappa_2)$, which implies

$$\begin{aligned}
f(\kappa_2+1)-f(\kappa_2) &\geq -c_{11}+c_{21}\{1-\mathbb{P}(Z_2>\kappa_2)\}-c_{\bullet 1}\mathbb{P}(Z_2>\kappa_2).\\
&= (c_{21}-c_{11})-c_{21}\mathbb{P}(Z_2>\kappa_2)-c_{\bullet 1}\mathbb{P}(Z_2>\kappa_2).
\end{aligned}$$

Because $\kappa_2-\lambda_2T>0$, Hoeffding's inequality (see, e.g., Boucheron et al., 2003) gives us that

$$\begin{aligned}
f(\kappa_2+1)-f(\kappa_2) &\geq (c_{21}-c_{11})-c_{21}\mathbb{P}(Z_2-\lambda_2T>\kappa_2-\lambda_2T)-c_{\bullet 1}\mathbb{P}(Z_2-\lambda_2T>\kappa_2-\lambda_2T)\\
&\geq (c_{21}-c_{11})-(c_{21}+c_{\bullet 1})\exp\left\{-2\frac{(\kappa_2-\lambda_2T)^2}{T}\right\},
\end{aligned} \tag{56}$$

completing the proof of (51).

For the proof of (52), we begin by recalling that $c_{11}<c_{\bullet 1}$, so we also have that $c_{11}\leq c_{11}\mathbb{P}(Z_1\geq\gamma-\kappa_2)+c_{\bullet 1}\mathbb{P}(Z_1<\gamma-\kappa_2)$. If we plug this last lower bound in the decomposition (55) and rearrange, we obtain the upper bound

$$f(\kappa_2+1)-f(\kappa_2)\leq c_{22}\mathbb{P}(Z_2>\kappa_2)+c_{21}\mathbb{P}(Z_2\leq\kappa_2)-c_{11}=-(c_{11}-c_{22})+(c_{21}-c_{22})\mathbb{P}(Z_2\leq\kappa_2).$$

Applying the Hoeffding inequality together with $\lambda_2T-\kappa_2>0$ give us the further upper bound

$$\begin{aligned}
f(\kappa_2+1)-f(\kappa_2) &\leq -(c_{11}-c_{22})+(c_{21}-c_{22})\mathbb{P}(Z_2-\lambda_2T\leq-(\lambda_2T-\kappa_2))\\
&\leq -(c_{11}-c_{22})+(c_{21}-c_{22})\exp\left\{-2\frac{(\lambda_2T-\kappa_2)^2}{T}\right\},
\end{aligned}$$

completing the proof of (52). Q.E.D.

Lemma B.4. *Consider independent random variables X_1, X_2, $\cdots$, X_T that are Bernoulli distributed with parameter $\lambda \in (0,1)$. For $-\infty \le C_1 < C_2 \le \infty$ and $C_3 \ge 1$, let E and E' be the events respectively defined as*

$$E = \left\{ \sum_{s=1}^{T} (X_s - \lambda) \in [C_1\sqrt{T}, C_2\sqrt{T}] \right\},$$

and

$$E' = \left\{ \sum_{s=1}^{t} (X_s - \lambda) \in [-C_3\sqrt{T}, C_3\sqrt{T}], \forall t \in (0,T] \right\}.$$

Then, for T large enough, the probabilities $\mathbb{P}(E)$ and $\mathbb{P}(E')$ are bounded away from 0.

Proof. Observe that $\sum_{s=1}^t X_s$ is binomial with parameters (t,λ). Let F_t be the cumulative distribution function of $\frac{\sum_{s=1}^t (X_s-\lambda)}{\sqrt{\lambda(1-\lambda)t}}$ for all $t \in [T]$. Then,

$$\begin{aligned}
\mathbb{P}(E) &= \mathbb{P}\left(\sum_{s=1}^{T} (X_s - \lambda) \in \left[C_1\sqrt{T}, C_2\sqrt{T}\right] \right) \\
&= \mathbb{P}\left(\frac{C_1}{\sqrt{\lambda(1-\lambda)}} \le \frac{\sum_{s=1}^T (X_s-\lambda)}{\sqrt{\lambda(1-\lambda)T}} \le \frac{C_2}{\sqrt{\lambda(1-\lambda)}} \right) \\
&= \mathbb{P}\left(\frac{\sum_{s=1}^T (X_s-\lambda)}{\sqrt{\lambda(1-\lambda)T}} \le \frac{C_2}{\sqrt{\lambda(1-\lambda)}} \right) - \mathbb{P}\left(\frac{\sum_{s=1}^T (X_s-\lambda)}{\sqrt{\lambda(1-\lambda)T}} \le \frac{C_1}{\sqrt{\lambda(1-\lambda)}} \right) \\
&= \Phi\left(\frac{C_2}{\sqrt{\lambda(1-\lambda)}}\right) - \Phi\left(\frac{C_1}{\sqrt{\lambda(1-\lambda)}}\right) + F_T\left(\frac{C_2}{\sqrt{\lambda(1-\lambda)}}\right) - \Phi\left(\frac{C_2}{\sqrt{\lambda(1-\lambda)}}\right) \\
&\quad - F_T\left(\frac{C_1}{\sqrt{\lambda(1-\lambda)}}\right) + \Phi\left(\frac{C_1}{\sqrt{\lambda(1-\lambda)}}\right) \\
&\ge \Phi\left(\frac{C_2}{\sqrt{\lambda(1-\lambda)}}\right) - \Phi\left(\frac{C_1}{\sqrt{\lambda(1-\lambda)}}\right) - \left| F_T\left(\frac{C_2}{\sqrt{\lambda(1-\lambda)}}\right) - \Phi\left(\frac{C_2}{\sqrt{\lambda(1-\lambda)}}\right) \right| \\
&\quad - \left| F_T\left(\frac{C_1}{\sqrt{\lambda(1-\lambda)}}\right) - \Phi\left(\frac{C_1}{\sqrt{\lambda(1-\lambda)}}\right) \right|.
\end{aligned}$$

By the Berry-Esseen theorem (see Shevtsova, 2011), we have $|F_t(x) - \Phi(x)| \le \frac{0.4847\lambda}{(\lambda(1-\lambda))^{3/2}\sqrt{t}}$ for all x and t. Hence, for large enough T, $\mathbb{P}(E) > 0$.

It is easy to check that $\sum_{s=1}^t (X_s - \lambda)$ is a martingale with respect to X_i, $i = 1, \cdots, t$. Now, we will bound $\mathbb{P}(E')$ using Kolmorogov's inequality (see, e.g Billingsley, 2017). For any $T > 0$, we have

$$\begin{aligned}
&\mathbb{P}\left(\sum_{s=1}^{t} (X_s - \lambda) \in [-C_3\sqrt{T}, C_3\sqrt{T}], \forall t \in (0,T] \right) \\
&= \mathbb{P}\left(\max_{t\in(0,T]} \left| \sum_{s=1}^{t} (X_s - \lambda) \right| \le C_3\sqrt{T} \right) \\
&= 1 - \mathbb{P}\left(\max_{t\in(0,T]} \left| \sum_{s=1}^{t} (X_s - \lambda) \right| \ge C_3\sqrt{T} \right) \\
&\ge 1 - \frac{\mathrm{Var}(\sum_{s=1}^t X_s)}{C_3^2 T} \\
&= 1 - \frac{\lambda(1-\lambda)t}{C_3^2 T} \\
&\ge 1 - \frac{\lambda(1-\lambda)}{C_3^2} \\
&> 0,
\end{aligned}$$

where the first inequality follows Kolmorogov's inequality and the last inequality follows since $C_3 \geq 1$ and $\lambda < 1$. Q.E.D.

Lemma B.5. *Consider the setting in Example 4.1, under policy* PF*. Suppose that $Z_2(t) \geq \kappa_2(t)$, then there exists a constant $c_l > 0$ which is determined by $\{c_{ij}\}$ and $\{c_{\bullet j}\}$ such that*

$$\Delta_t^{\mathrm{PF}} \geq c_l, \text{ if } x_{21}^{\mathrm{PF}}(t) = 1,$$
$$\sum_{s=t}^{T} \Delta_t^{\mathrm{PF}} \geq c_l \left(\kappa_2(t) - \sum_{s=t}^{T} x_{22}^{\mathrm{PF}}(s)\right).$$

In particular, when $t = 1$ and $Z_2 = Z_2(1) \geq \kappa_2 = \kappa_2(1)$, we have

$$\sum_{t=1}^{T} \Delta_t^{\mathrm{PF}} \geq c_l \left(\kappa_2 - \sum_{t=1}^{T} x_{22}^{\mathrm{PF}}(t)\right).$$

Proof. Let $(\bar{\boldsymbol{\ell}}, \bar{\boldsymbol{x}})$ be a solution for $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t))$. By Lemma B.2, the optimal solution for $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t))$ is unique and $\bar{x}_{22} = \kappa_2(t)$. Because $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t))$ can be formulated as a min-cost network flow problem, we have that the cost of any augmenting path or cycle for $(\bar{\boldsymbol{\ell}}, \bar{\boldsymbol{x}})$ is positive. Let c_l be the minimum cost of all the possible augmenting paths or cycles. Then, $c_l > 0$ is determined by $\{c_{ij}\}$ and $\{c_{\bullet j}\}$ and any feasible solution $(\boldsymbol{\ell}', \boldsymbol{x}')$ with $x'_{22} < \bar{x}_{22} = \kappa_2(t)$ must have an objective that is at least $c_l(\kappa_2(t) - x'_{22})$.

Suppose $x_{21}^{\mathrm{PF}}(t) = 1$. Let $(\boldsymbol{\ell}^{t+1}, \boldsymbol{x}^{t+1})$ be the optimal offline solution for $F(\boldsymbol{\kappa}(t+1), \boldsymbol{Z}(t+1))$. Note that $(\boldsymbol{\ell}^{t+1}, \boldsymbol{x}^{t+1} + \boldsymbol{x}^{\mathrm{PF}}(t))$ is a feasible solution for $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t))$, and

$$x_{22}^{t+1} + x_{22}^{\mathrm{PF}}(t) = x_{22}^{t+1} \leq \kappa_2(t+1) = \kappa_2(t) - 1. \tag{57}$$

Thus,

$$\begin{aligned}
\Delta_t^{\mathrm{PF}} &= \sum_{i,j} c_{ij} \left(x_{ij}^{\mathrm{PF}}(t) + x_{ij}^{t+1}\right) + \sum_j c_{\bullet j} l_j^{t+1} - \left(\sum_{i,j} c_{ij} \bar{x}_{ij} + \sum_j c_{\bullet j} \bar{l}_j\right) \\
&\geq c_l(\kappa_2(t) - x_{22}^{\mathrm{PF}}(t) - x_{22}^{t+1}) \\
&= c_l.
\end{aligned}$$

where the first inequality follows by the definition of c_l and uniqueness of $\bar{x}$, the second equality follows because $x_{22}^{\mathrm{PF}}(t) = 0$ and $\bar{x}_{22} = \kappa_2(t)$, and the second inequality follows from (57).

Let $x_{ij}^{\mathrm{PF}} = \sum_{s=t}^{T} x_{ij}^{\mathrm{PF}}(s)$ and $l_j^{\mathrm{PF}} = \sum_{s=t}^{T} l_j^{\mathrm{PF}}(s)$, and observe that $(\boldsymbol{\ell}^{\mathrm{PF}}, \boldsymbol{x}^{\mathrm{PF}})$ is a feasible solution for $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t))$. Thus, by definition of Δ_t^{PF}, we have

$$\begin{aligned}
\sum_{s=t}^{T} \Delta_t^{\mathrm{PF}} &= \sum_{i,j} c_{ij} \sum_{s=t}^{T} x_{ij}^{\mathrm{PF}}(s) + \sum_j c_{\bullet j} \sum_{s=t}^{T} l_j^{\mathrm{PF}}(s) - \left(\sum_{i,j} c_{ij} \bar{x}_{ij} + \sum_j c_{\bullet j} \bar{l}_j\right) \\
&= \sum_{i,j} c_{ij} x_{ij}^{\mathrm{PF}} + \sum_j c_{\bullet j} l_j^{\mathrm{PF}} - \left(\sum_{i,j} c_{ij} \bar{x}_{ij} + \sum_j c_{\bullet j} \bar{l}_j\right) \\
&\geq c_l(\bar{x}_{22} - x_{22}^{\mathrm{PF}}) \\
&= c_l(\kappa_2(t) - x_{22}^{\mathrm{PF}}).
\end{aligned}$$

where the inequality holds by uniqueness of $\bar{x}$, and definition of c_l. Q.E.D.

B.2. Lemmas for Section 4.2

Lemma B.6. *Let $\bar{\boldsymbol{x}}^1$ be an optimal solution of $F(\boldsymbol{\kappa}, \boldsymbol{z}^1)$. Consider a vector $\boldsymbol{z}^2$ is such that $z_n^2 \neq z_n^1$ for some $n \in [J]$, and $z_j^2 = z_j^1$ for $j \in [J] \setminus \{n\}$. Then, there exists $\bar{\boldsymbol{x}}^2$ such that*

$$\bar{\boldsymbol{x}}^2 \text{ is an optimal solution of } F(\boldsymbol{\kappa}, \boldsymbol{z}^2), \tag{58}$$

$$|\bar{x}_{ij}^2 - \bar{x}_{ij}^1| \le |z_n^1 - z_n^2| \text{ for all } i \in [I] \cup \{0\}, j \in [J]. \tag{59}$$

Proof. Observe that $F(\boldsymbol{\kappa}, \boldsymbol{z})$ can be converted to a minimum-cost flow problem by including additional source and sink nodes. Specifically, define dummy nodes s (source node) and t (sink node). Let $N = \{s\} \cup [I] \cup \{0\} \cup [J] \cup \{t\}$ be the set of nodes, and A be the set of arcs. Then, $F(\boldsymbol{\kappa}, \boldsymbol{z})$ can be written in the following form.

$$\begin{aligned}
\min \quad & \sum_{i \in [I] \cup \{0\}, j \in [J]} c_{ij} f_{ij} & \\
\text{s.t.} \quad & \sum_{i \in [I] \cup \{0\}} f_{ij} = f_{jt} & j \in [J] \\
& \sum_{j \in [J]} f_{ij} = f_{si} & i \in [I] \cup \{0\} \\
& \sum_{j \in [J]} f_{jt} = \sum_{i \in [I] \cup \{0\}} f_{si} & \\
& 0 \le f_{si} \le \kappa_i & i \in I \cup \{0\} \\
& z_j \le f_{jt} & j \in [J] \\
& 0 \le f_{ij} & i \in [I] \cup \{0\}, j \in [J].
\end{aligned}$$

It is easy to check that any optimal solution $\boldsymbol{x}^*$ of $F(\boldsymbol{\kappa}, \boldsymbol{z})$ in the original formulation (24) has a corresponding optimal network flow $\mathbf{f}^*$ in the new formulation, where $x_{ij}^* = f_{ij}^*$ for any $i \in [I] \cup \{0\}$, and $j \in [J]$.

Recall that $\bar{\boldsymbol{x}}^1$ is the optimal solution of $F(\boldsymbol{\kappa}, \boldsymbol{z}^1)$, and let $\bar{\mathbf{f}}^1$ be the corresponding optimal network flow. Let G be the residual graph corresponding to flow vector $\bar{\mathbf{f}}^1$. By the optimality of $\bar{\mathbf{f}}^1$, it must satisfy the reduced cost optimality condition, that is, there exists a $|N|$ dimensional potential vector π^1 such that

$$c_{ij} - \pi_i + \pi_j \ge 0, \text{ for every arc } (i, j) \text{ in } G.$$

If $z_n^2 - z_n^1 = 1$, let p be the shortest path between s and n in the residual graph G. By Lemma 9.12 of Ahuja et al. (1988) on the property of the successive shortest path algorithm, we can send a unit flow along path p to obtain a feasible flow $\bar{\mathbf{f}}^2$ that satisfies the reduced cost optimality condition with some potential vector π^2. Let $\bar{\boldsymbol{x}}^2$ be the corresponding optimal solution for $F(\boldsymbol{\kappa}, \boldsymbol{z}^2)$, we then have that

$$|\bar{f}_{ij}^2 - \bar{f}_{ij}^1| = |\bar{x}_{ij}^2 - \bar{x}_{ij}^1| \le 1 = |z_n^1 - z_n^2| \text{ for all } i \in [I] \cup \{0\}, j \in [J].$$

If $z_n^2 - z_n^1 = -1$, let p be the shortest path between n and s in the residual graph G. Again by Lemma 9.12 of Ahuja et al. (1988) we can send a unit flow along path p to obtain a feasible flow $\bar{\mathbf{f}}^2$ that satisfies the reduced cost optimality condition with some potential vector π^2. Let $\bar{\boldsymbol{x}}^2$ be the corresponding optimal solution for $F(\boldsymbol{\kappa}, \boldsymbol{z}^2)$. We obtain that

$$|\bar{x}_{ij}^2 - \bar{x}_{ij}^1| \le 1 = |z_n^1 - z_n^2| \text{ for all } i \in [I] \cup \{0\}, j \in [J].$$

Finally, for the case where $|z_n^2 - z_n^1| > 1$, we can simply apply the successive shortest path algorithm for multiple iterations until we reach an optimal solution $\bar{\mathbf{f}}^2$. By induction, the optimal solution $\bar{\boldsymbol{x}}^2$ corresponding to $\bar{\mathbf{f}}^2$ satisfies that

$$|\bar{x}_{ij}^2 - \bar{x}_{ij}^1| \le |z_n^1 - z_n^2| \text{ for all } i \in [I] \cup \{0\}, j \in [J].$$

Q.E.D.

Lemma B.7. *Let $\hat{\boldsymbol{x}}$ be an optimal solution of (24). For any $\boldsymbol{\delta} \in \mathbb{R}^J$, there exists $\boldsymbol{x}^{\boldsymbol{\delta}}$ such that*

$$\boldsymbol{x}^{\boldsymbol{\delta}} \text{ is an optimal solution of } F(\boldsymbol{\kappa}, \boldsymbol{z}+\boldsymbol{\delta}),$$
$$|x^{\boldsymbol{\delta}}_{ij} - \hat{x}_{ij}| \le \sum_{j=1}^{J} |\boldsymbol{\delta}_j| = \|\boldsymbol{\delta}\|_1, \forall i \in [I] \cup \{0\}, j \in [J].$$

Proof. The proof follows from Lemma B.6 where we show that for any $\boldsymbol{z}'$ that only differs from $\boldsymbol{z}$ in the n^{th} component, then $F(\boldsymbol{\kappa}, \boldsymbol{z}')$ has an optimal solution $\hat{x}'$ such that $|\hat{x}'_{ij} - \hat{x}_{ij}| \le |z'_n - z_n|$.

First, define $\boldsymbol{\delta}(n) = [\delta_1, \delta_2, \cdots, \delta_n, 0, \cdots, 0]$. In other words, $\boldsymbol{\delta}(n)$ is the truncation of $\boldsymbol{\delta}$ with just the first n terms. Let $\boldsymbol{x}^n$ be the optimal solution to $F(\boldsymbol{\kappa}, \boldsymbol{z}+\boldsymbol{\delta}(n))$. By Lemma B.6, there exists some $\boldsymbol{x}^{n+1}$ which is an optimal solution to $F(\boldsymbol{\kappa}, \boldsymbol{z}^1 + \boldsymbol{\delta}(n+1))$ such that $|x^n_{ij} - x^{n+1}_{ij}| \le |\delta_{n+1}|$ because $\boldsymbol{\delta}(n)$ and $\boldsymbol{\delta}(n+1)$ only differ by their $(n+1)^{\text{th}}$ component. Hence, by triangle inequality,

$$|\hat{x}_{ij} - x^J_{ij}| \le |\hat{x}_{ij} - x^1_{ij}| + \cdots + |x^{J-1}_{ij} - x^J_{ij}| \le \sum_{j=1}^{J} |\delta_j|.$$

Denote $\boldsymbol{x}^{\boldsymbol{\delta}} = \boldsymbol{x}^J$. Then, $\boldsymbol{x}^{\boldsymbol{\delta}}$ is the optimal solution to $F(\boldsymbol{\kappa}, \boldsymbol{z}+\boldsymbol{\delta})$ which satisfies $|x^{\boldsymbol{\delta}}_{ij} - \hat{x}_{ij}| \le \|\boldsymbol{\delta}\|_1$.
Q.E.D.

Lemma B.8. *Suppose that $x^{\text{SF}}_{mn}(t) = 1$ for any $t > 0$, $m \in [I]$ and $n \in [J]$. If $\bar{\boldsymbol{x}}$ is the optimal solution of $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t))$, and $\bar{x}_{mn} \ge 1$, then Δ^{SF}_t, defined as $C^{\text{SF}}_t - F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t)) + F(\boldsymbol{\kappa}(t+1), \boldsymbol{Z}(t+1))$, is equal to zero.*

Proof. Since the optimal solution of $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t))$ satisfies $\bar{x}_{mn} \ge 1$, we have that

$$\begin{aligned}
F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t)) = \min\ & c_{mn} + \sum_{i \in [I] \cup \{0\}} \sum_{j \in [J]} c_{ij} x_{ij} \\
\text{s.t.}\ & \sum_{i \in [I] \cup \{0\}} x_{ij} \ge Z_j(t) - I(j=n) && j \in [J] \\
& \sum_{j \in [J]} x_{ij} \le \kappa_i(t) - I(i=m) && i \in [I] \\
& x_{ij} \ge 0 && i \in [I] \cup \{0\},\quad j \in [J],
\end{aligned}$$

where I is the indicator function. Then, as $Z_j(t+1) = Z_j(t) - I(j=n)$ for $j \in [J]$ and $\kappa_i(t+1) = \kappa_i(t) - I(i=m)$ for $i \in [I]$, we have that $F(\boldsymbol{\kappa}(t), \boldsymbol{Z}(t)) = c_{mn} + F(\boldsymbol{\kappa}(t+1), \boldsymbol{Z}(t+1)) = C^{\text{SF}}_t + F(\boldsymbol{\kappa}(t+1), \boldsymbol{Z}(t+1))$, which implies $\Delta^{\text{SF}}_t = 0$. Q.E.D.