MICHIGAN ROSS

Working Paper

Distribution-free Inventory Risk Pooling in a Multi-location Newsvendor

Aravind Govindarajan
Stephen M. Ross School of Business
University of Michigan

Amitabh Sinha
Amazon.com

Joline Uichanco
Stephen M. Ross School of Business
University of Michigan

Ross School of Business Working Paper
Working Paper No. 1389
January 2019

UNIVERSITY OF MICHIGAN

Distribution-free Inventory Risk Pooling in a Multi-location Newsvendor

Aravind Govindarajan, Joline Uichanco
Stephen M. Ross School of Business, University of Michigan, Ann Arbor, MI 48109,
arav, jolineu@umich.edu

Amitabh Sinha
Amazon.com, Seattle, WA 98109,
amitabsi@amazon.com

With rapidly increasing e-commerce sales, firms are leveraging the virtual pooling of online demands across customer locations in deciding the amount of inventory to be placed in each node in a fulfillment network. Such solutions require knowledge of the joint distribution of demands; however, in reality, only partial information about the joint distribution may be reliably estimated. We propose a distributionally robust multi-location newsvendor model for network inventory optimization where the worst-case expected cost is minimized over the set of demand distributions satisfying the known mean and covariance information. For the special case of two homogeneous customer locations with correlated demands, we show that a six-point distribution achieves the worst-case expected cost, and derive a closed-form expression for the optimal inventory decision. The general multi-location problem can be shown to be NP-hard. We develop a computationally tractable upper bound through the solution of a semidefinite program (SDP), which also yields heuristic inventory levels, for a special class of fulfillment cost structures, namely nested fulfillment structures. We also develop an algorithm to convert any general distance-based fulfillment cost structure into a nested fulfillment structure which tightly approximates the expected total fulfillment cost.

Key words: e-commerce, fulfillment, distribution-free optimization, inventory management

---

1. Introduction

As e-commerce continues to grow rapidly (Zaroban 2018), most retail firms are equipping themselves with the ability to fulfill online orders from multiple nodes (stores, fulfillment centers, etc.) in their network. In order to decide the amount of inventory to keep in each node to fulfill demands, traditional decentralized inventory planning solutions are no longer appropriate, as they do not take into account demand spillover from one node to the other in case of stockouts (Acimovic and Graves 2017). Network-based solutions can reduce the burden of carrying too much inventory, by taking into account this fulfillment flexibility which pools online demands across customer regions.

A fundamental assumption in inventory management is that the underlying distribution of the uncertainty, namely the demands, is known. In network inventory planning, it is imperative that the joint distribution of demands across customer regions is known reliably, as correlation of customer demands across locations affects the optimal inventory solutions due to movement of inventory across nodes in the network (Eppen 1979).

However, in the case of e-commerce demands, the exact distribution may be inaccessible due to high volatility in online customer behavior arising from factors such as competition, the use of dynamic price-matching strategies and flash promotions, recommendation engines that manipulate click-streams, etc. Additionally, only lower order moments may be available, as higher order moments require copious amounts of data for reliable estimation due to the scale of the problem.

A common workaround for this problem is to assume that the underlying demands follow a multivariate normal distribution. Such an assumption can help an inventory planner in two ways:

1. Describing a multivariate normal demand requires information about only the first two moments, namely the mean and covariance matrices, which can be reliably estimated.
2. The normal distribution lends itself to simple analytic solutions (e.g. Dong and Rudi 2004).

Despite its apparent advantages, the normal distribution assumption can lead to solutions that overestimate pooling benefits, if the true demand distribution is non-normal.[1] Eppen (1979) showed that for demand distributions that are of light-tailed nature (including the normal distribution), pooling can lead to savings in expected cost that scale with $\sqrt{n}$, where n is the number of random demands being pooled. However, Bimpikis and Markakis (2015) show that the pooling benefits can scale significantly lower than $\sqrt{n}$ when the demand distribution is heavy-tailed (such as lognormal, gamma, etc.). The predictability of the aggregate demand (which directly influences pooling benefits) depends on the distribution of the individual demands. As a result, incorrectly specifying a distribution can misestimate pooling benefits and lead to significant increases in cost. We illustrate through a simple example how assuming a normal distribution can lead to suboptimal outcomes.

Example 1. Consider a simple reactive-transshipment system of a firm serving customer demands from two locations through two warehouses. In region 1, warehouse W_1 serves its own demand D_1 with inventory y_1, and likewise for region 2. In a single period setting, in the event of a stockout at, say W_1, available inventory can be transshipped from W_2 to W_1 at a cost $s=1$. At the end of the period, the per-unit overage and underage costs are $h=1$ and $p=100$ respectively. Let the first and second moments of the joint demand distribution be $\mathbf{m}=\begin{pmatrix}10\\10\end{pmatrix}$, $\mathbf{\Sigma}=\begin{pmatrix}16 & 4\\4 & 16\end{pmatrix}$.

[1] There have been earlier studies that show evidence of real-life demands exhibiting non-normal distributions: Bimpikis and Markakis (2015) show empirical evidence of heavy-tailed demands for movies at Netflix and shoes at a major retailer, Agrawal and Smith (1996) show that the negative binomial distribution fits the sales data for men's slacks at a major retailer better than Poisson or Normal.

Support Points	Demand Realizations		Probability
	D_1	D_2	
1	9.35	9.35	0.9595
2	25.44	25.44	0.0171
3	9.35	41.37	0.0117
4	41.37	9.35	0.0117

Table 1 True joint distribution of demands

Assuming a multivariate normal distribution, the optimal inventory levels are given by $y^{*,N} = (17.4, 17.4)^\intercal$ (Dong and Rudi 2004, Equation 3), with an optimal expected cost of 15.7. Instead, if the true joint distribution was a discrete distribution as shown in Table 1, the expected cost while operating with inventory levels $y^{*,N}$ increases by 400% to 80.4. This is because the optimal inventory levels under the true distribution in Table 1 is $y^{*,D} = (25.4, 25.4)^\intercal$, which carries more than double the amount of safety stock as $y^{*,N}$, as the normal distribution deems extreme events as unlikely. Moreover, the reduction in expected cost by going from a system with no pooling to one with pooling is around 36% under the normal distribution, whereas it is only 20% under the discrete distribution, since the normal distribution overestimates the effect of pooling.

The distribution in Table 1 is actually one that leads to the worst-case expected cost for the inventory levels $y^{*,N}$ given the mean and covariance matrices (we shall show this in Section 3). Nevertheless, cases like this where events of imbalance (support points 3 and 4 in Table 1) have a significant impact on the expected cost motivate our search for a robust solution with a good worst-case performance.[2]

In this paper, we consider distributionally robust inventory optimization of a retail firm selling a single product through a fulfillment network consisting of multiple inventory nodes (such as warehouses, stores, etc.). We consider a single-period setting, where after demands are realized, network inventory can be re-routed through network fulfillment flows. We assume that the only information available are the first and second order moments of the vector of demands.

Main Results and Contributions

1. We extend the distributional robustness literature by analyzing the multi-location newsvendor with network flows, through which we characterize the impact of pooling on distributionally robust solutions. Previous studies like Scarf (1958), Gallego and Moon (1993) and Hanasusanto et al. (2015) do not consider any feature that pools inventories across locations or products.
2. We derive in closed-form the optimal inventory levels, and the probability distribution that achieves the worst-case expected cost given any inventory levels for the simple case of two

[2] For the purpose of comparison, the robust solution that we develop in this study has the solution $y^{*,R} = (25.8, 25.8)$, which has a worst-case expected cost of 63.4 over all possible distributions with mean $\mathbf{m}$ and covariance matrix $\mathbf{\Sigma}$.

homogeneous locations with correlated demands. We find that this distribution has six support points, with a set of probabilities over these support points achieving the same worst-case expected cost. We detail the effect of correlation and variance on the imbalance induced in the system and in turn, on the expected cost.

3. We describe a class of fulfillment cost structures, namely nested fulfillment structures, that reduces the recourse linear program which computes network flows into piecewise linear terms. We develop a simple algorithm based on hierarchical agglomerative clustering to approximate any general distance-based fulfillment cost structure as a nested fulfillment structure. We show empirically that this structure tightly approximates the expected total fulfillment cost under a variety of distributions. A nested fulfillment structure not only yields tractability for our robust problem, but also for stochastic systems.
4. For the NP-hard multi-location problem, we develop computationally tractable bounds and heuristics in the form of SDPs (semi-definite programs) under the nested fulfillment cost structure. The heuristic solutions provide inventory levels at each node in the network, and are computationally fast compared to the optimal solutions which require exponential number of SDP constraints. By means of numerical experiments, we find that including information about distribution asymmetry through partitioned statistics (studied by Natarajan et al. 2018) significantly reduces the worst-case expected cost when there is inventory pooling. We also show that the distributionally robust heuristic can lead to significant savings in expected cost as compared to stochastic solutions that assume an incorrect distribution.

1.1. Literature Review

The distributionally robust newsvendor problem has a long history, dating back to Scarf (1958) who considers the classic newsvendor problem with only mean, variance and support information available about the uncertain demand. He derives the optimal inventory levels, and a two-point worst-case distribution that minimizes expected profit given inventory levels. After more than 30 years, Gallego and Moon (1993) extended Scarf (1958) to multiple products with marginal moment information (mean and variance information). This setting was further extended to including cross-moment information (correlation among product demands) by Hanasusanto et al. (2015), who showed that the distributionally robust multi-item newsvendor problem given mean and covariance information is NP-hard.

Most of the above cited papers do not consider any recourse actions for the decision maker after observing the realization of demands. Gallego and Moon (1993) considered extension of the single product newsvendor to the case where a recourse action is available to place an additional order after demand realization. Mostard et al. (2005) considered resalable returns in a single period setting, where returned products can be resold if there is sufficient demand. In our study, recourse is available in the form of network flows which can be used to balance supply and demand. We note that such a recourse pools inventories across locations, a feature that is absent in previous studies on distributional robustness, to the best of our knowledge.

Our study can thus be related to the literature on inventory pooling. We only discuss papers that establish the importance of distributional properties of the uncertain demand in pooling. Eppen (1979) showed that if n normal and uncorrelated demands are pooled, the benefit from inventory pooling is $\sqrt{n}$, and the benefit decreases with increasing positive correlation among demands. Corbett and Rajaram (2006) extended Eppen's result to more general distributions. Yang and Schrage (2009) study various cases of 'inventory anomaly' (a situation where pooling leads to an increase in inventory as opposed to a reduction), one of which is for right-skewed demand distributions with product substitution. Berman et al. (2011) found through numerical simulations that the normal distribution misestimates the benefits of pooling stemming from a reduction in variance. An important result by Bimpikis and Markakis (2015) implies that the benefit from pooling under heavy-tailed demand distributions can be significantly lower than $\sqrt{n}$, which is predicted by normal demands. Specifically, they show that the benefit from pooling decreases as the tail of the demand becomes heavier. All these studies indicate that pooling benefits crucially depend on the distribution of the demands being pooled.

We study inventory pooling using network flows which balance demand and supply after realization of uncertainty. This setting is mathematically identical to reactive lateral transshipments, which have been discussed in great detail in the literature (for a review, refer to Paterson et al. 2011). Two features make the transshipment problem difficult to analyze. First, for more than two locations, analytically optimal solutions become elusive as a linear program recourse is needed to model the network flow problem among multiple locations (Robinson 1990). Second, in a multi-period setting with leadtimes, the optimal transshipment decisions are intractable due to complexity in the state space even for the two location problem, as it can be *ex-post* optimal to reserve inventory for future use rather than transshipping to another location (Tagaras and Cohen 1992). In this paper, we focus on the multi-location problem in a single period setting. Dong and Rudi (2004) consider a similar setting for the special case where the transshipment cost between any two locations is a constant. We introduce a more general case of fulfillment costs, namely nested fulfillment costs, which preserves tractability for both stochastic and distributionally robust settings.

Through the development of tractable heuristics to our problem, we relate to various studies in the distributional robustness literature. Burer (2009) showed that nonconvex mixed-integer quadratic programs can be expressed as completely positive programs. Natarajan et al. (2011) apply this to the case of robust mixed 0-1 linear programs with objective uncertainty. Using these results, Natarajan and Teo (2017a) develop a tractable heuristic in the form of a semi-definite program for the multi-item newsvendor problem with known mean and covariance matrices (without recourse actions) as considered by Hanasusanto et al. (2015). They achieve this by expressing piecewise linear terms through integer variables, and relaxing the equivalent completely positive program into a semi-definite program. They relax the integrality constraints through a boolean quadric polytope, previously studied by Padberg (1989). Natarajan et al. (2018) use similar techniques to derive tractable heuristics for the multi-item newsvendor with known mean, covariance and semivariance information, which additionally captures asymmetry in the distribution.

Finally, our study can also be related to the growing literature in e-commerce inventory and fulfillment optimization. Acimovic and Graves (2017) show that decentralized inventory solutions can lead to costly spillover effects, and perform poorly compared to network-based policies. Govindarajan et al. (2018) consider joint optimization of inventory and fulfillment decisions in an omnichannel setting where in-store demands cannot be flexibly fulfilled from other locations, whereas e-commerce demand can. Acimovic and Graves (2014) consider dynamic fulfillment decisions in an e-commerce setting, by minimizing immediate outbound shipping cost plus an estimate of future shipping costs. Given the volatility in online customer behavior, we contribute to this stream of literature by studying the distributionally robust inventory problem where only lower order moments of the random demands can be reliably estimated.

1.2. Preliminaries

For any integer n, we use notation $[n]$ to denote the set $\{1,2,\ldots,n\}$. We denote by $2^{[n]}$ the power set of $[n]$, defined as the set of all subsets of $[n]$ including the empty set. We denote by $\mathbf{e}_n$ the n-dimensional column vector of all ones, where we drop the subscript if the size is clear from the context. We denote by $\mathbf{I}_n$ the identity matrix of size $n \times n$, and $\mathbf{0}_{m,n}$ the zero matrix of size $m \times n$.

We denote by $\Re$ the set of real numbers, and by $\Re_{\geq 0}$ the set of nonnegative real numbers. We similarly denote by $\Re^n$ the set of n-dimensional vectors of real numbers, and $\Re^n_{\geq 0} := \{\mathbf{x} \in \Re^n | \mathbf{x} \geq 0\}$ its subset of nonnegative vectors. For a scalar variable $x \in \Re$, we define $x^+ := \max(0, x)$ as the positive part of x. For a column vector $\mathbf{x} = (x_i) \in \Re^n$, we define $\mathbf{x}^+ := (x_i^+)$ as the positive part of each element in $\mathbf{x}$. We write $\mathbf{A} \succeq 0$ if a square matrix $\mathbf{A}$ is symmetric positive semidefinite. We write $\mathbf{B} \geq 0$ if all entries of the matrix $\mathbf{B}$ are nonnegative.

2. The General Model

In this section, we describe the general model for distribution-free inventory risk pooling, which will serve as a foundation for the remainder of the paper. We begin by discussing the assumptions on the firm's cost parameters and the structure of the multi-location newsvendor problem with inventory risk pooling. Then, we discuss what type of information is available to the firm, and introduce the associated minmax robust formulation, which we refer to as the distributionally robust multi-location newsvendor with inventory risk pooling.

Consider a firm that has a network of inventory nodes (e.g., stores or warehouses) to support sales of a single product during a selling horizon. We assume that there is no inventory replenishment during the selling horizon (as is typically the case when the horizon is short compared to the procurement lead time), so the firm only needs to decide the initial inventory levels. We assume that there are n customer regions, and that $\tilde{d}_j$ denoting the demand in region $j \in [n]$ is random (a variable with a tilde placed on top refers to a random variable; the same variable without the tilde is a particular realization). The vector of stochastic demands in the n customer regions is $\tilde{\mathbf{D}} = (\tilde{d}_j)$ (unless otherwise stated, any vector is a column vector).

At the end of the horizon, unmet demand in any region incurs a per-unit penalty cost p, while unsold inventory in any node incurs a per-unit overage cost h. We also assume that the firm incurs a per-unit "fulfillment" cost (e.g. shipping cost and/or handling cost) for using an inventory node to fulfill demand in a customer region, and this cost depends on both the inventory node and the customer region. For simplicity of the model, we consider a network with n inventory nodes, where one node is located in each customer region. Our framework can easily be extended to general networks by modeling customer regions without a fulfillment node as zero-inventory nodes.

If a unit of demand in region $j \in [n]$ is met with the inventory from a node in the same location j, then the fulfillment cost is s_0; if met with inventory from a node $i \neq j$, then the fulfillment cost is s_{ij} $(= s_{ji})$. We assume that $s_{ij} > s_0$ for any $i \neq j$ (if $s_{ij} = s_0$ $\forall i, j$ then the problem is trivially equivalent to a classical newsvendor problem with inventory pooled in a single location with overage cost h and underage cost $p - s_0$). We also assume that $p + h > s_{ij}$ for all i, j; in other words, it is always beneficial to fulfill a unit of demand if there is available inventory in the network.[3] We denote the $n \times n$ matrix of per-unit fulfillment costs as $\mathbf{S} = (s_{ij})$, where $s_{ii} = s_0$ for all $i \in [n]$.

At the start of the selling horizon, the firm decides the vector of initial inventory levels $\mathbf{y} = (y_i)$ to fulfill demands that arrive throughout the selling season. We assume that fulfillment is done at the end of the selling horizon, so that the problem can be approximated as a single period problem.[4] The single-period approximation allows us to study pooling in a distribution-free context by side-stepping the complications that arise from dynamic fulfillment decision-making, and focusing on the inventory decisions. In Section 5, we show that the single-period assumption is a tight approximation for the dynamic setting under the common practice of myopic fulfillment.

[3] We can relax this assumption for nested fulfillment cost structures, as we discuss in Section 4.1.

[4] While taking into account dynamic fulfillment decisions is more realistic, the problem becomes complicated due to the well-known curse of dimensionality (see Tagaras and Cohen (1992) for a stochastic system, and Ben-Tal et al. (2004) for the robust system). Multi-location considerations often complicate the problem further by adding complexity to the action space, as a linear program recourse is needed to make fulfillment decisions.

In the single period setting, the vector of customer demands $\mathbf{D} = (d_j)$ is realized at the end of the period, and the firm fulfills the demand with the objective of minimizing the total newsvendor cost (i.e., penalty, overage, and fulfillment costs). Mathematically, if z_{ij} units of inventory from node i used to satisfy demand in region j, then the fulfillment quantities $\mathbf{Z} = (z_{ij})$ are determined by solving the following network flow problem:

$$\begin{aligned} C(\mathbf{y}, \mathbf{D}) := \underset{\mathbf{Z} \geq 0}{\text{minimize}} \quad & h \cdot \sum_{i \in [n]} \left(y_i - \sum_{j \in [n]} z_{ij} \right) + p \cdot \sum_{j \in [n]} \left(d_j - \sum_{i \in [n]} z_{ij} \right) + \sum_{i \in [n]} \sum_{j \in [n]} s_{ij} z_{ij} \\ \text{subject to} \quad & \sum_{j \in [n]} z_{ij} \leq y_i, \quad \forall i \in [n], \\ & \sum_{i \in [n]} z_{ij} \leq d_j, \quad \forall j \in [n], \end{aligned} \tag{1}$$

where the terms in the objective are the overage cost, the penalty cost, and the fulfillment cost, respectively. The first constraint specifies that the units used to fulfill demand from inventory node i should not exceed the initial stocking level y_i. The second constraint specifies that the total units used to fulfill demand in region j must not exceed d_j.

We can write the total cost (1) compactly in matrix notation as

$$\begin{aligned} C(\mathbf{y}, \mathbf{D}) = \underset{\mathbf{Z} \geq 0}{\text{minimize}} \quad & h \cdot \mathbf{e}^\top (\mathbf{y} - \mathbf{Z}\mathbf{e}) + p \cdot \mathbf{e}^\top \left(\mathbf{D} - \mathbf{Z}^\top \mathbf{e} \right) + \langle \mathbf{S}, \mathbf{Z} \rangle \\ \text{subject to} \quad & \mathbf{Z}\mathbf{e} \leq \mathbf{y}, \\ & \mathbf{Z}^\top \mathbf{e} \leq \mathbf{D}, \end{aligned} \tag{2}$$

where $\langle \cdot, \cdot \rangle$ is the Frobenius inner product and $\mathbf{e}$ is a vector of all ones with appropriate dimension. The compact notation will be useful for us in Section 4.

If the firm knew that the stochastic multi-location demand $\tilde{\mathbf{D}}$ has a joint distribution $f : \Re^n \mapsto \Re^+$, then the firm will choose the initial inventory vector $\mathbf{y}$ so as to minimize the expected total cost. Mathematically, the firm's problem is equivalent to solving the two-stage stochastic program:

$$\min_{\mathbf{y} \geq 0} \mathbb{E}_f \left[C(\mathbf{y}, \tilde{\mathbf{D}}) \right], \tag{3}$$

where $\mathbb{E}_f$ is the expectation operator under the joint probability distribution f. Note that the objective is to minimize the expected total cost, where $C(\mathbf{y}, \mathbf{D})$ is as described in (2). Since inventory in one location can be routed to meet demand in another location, we will refer to (3) as the *multi-location newsvendor problem with inventory risk pooling.* Note that the above problem can be numerically solved as a linear program either by sample average approximation using large enough number of samples, or by approximating the joint distribution by a discrete distribution.

In reality, however, firms do not have a complete description of the joint distribution of the multi-location demands. At best, the firm may only have partial information about the distribution. We will assume that the firm only has knowledge of the mean vector $\mathbf{m}$ and the covariance matrix $\mathbf{\Sigma}$. As previously illustrated in the Example 1, if the firm assumes a particular distribution for the demand vector, say a multivariate normal distribution with parameters $(\mathbf{m}, \mathbf{\Sigma})$, then the optimal decision resulting from the stochastic program (3) may be suboptimal with a high expected cost under the true (unknown) demand distribution. To protect against such cases, we adapt a minmax distributionally robust approach (Scarf 1958, Gallego and Moon 1993, Hanasusanto et al. 2015, Natarajan et al. 2018) that aims to choose inventory levels $\mathbf{y}$ to minimize the maximal expected cost over all demand distributions consistent with the information known to the firm.

To understand the minmax robust approach, assume that after the firm makes a decision on the inventory levels $\mathbf{y}$, an "adversary" is able to choose a joint distribution f that results in the highest expected cost $\mathbb{E}_f(\mathbf{y}, \tilde{\mathbf{D}})$ for the firm. However, the adversary cannot choose just any f; consistent with the known mean and covariance, it has to belong to the distribution set:

$$\mathcal{F}_{\geq 0} := \left\{ f : \Re^n_{\geq 0} \mapsto \Re_{\geq 0} \,|\, \mathbb{E}_f(1) = 1,\, \mathbb{E}_f\left(\tilde{\mathbf{D}}\right) = \mathbf{m},\, \mathbb{E}_f\left(\tilde{\mathbf{D}}\tilde{\mathbf{D}}^\top\right) = \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top \right\}, \tag{4}$$

which is the set of all joint probability distributions of the n-dimensional demand, whose support is nonnegative, where the sum of probabilities equal to 1, the expectation is $\mathbf{m}$, and the covariance is $\mathbf{\Sigma}$. The firm's best strategy against this adversary is to choose $\mathbf{y}$ that minimizes the "worst-case" expected cost (i.e., the maximum expected cost among distributions in $\mathcal{F}_{\geq 0}$). Mathematically, this is done by solving the following minmax robust problem

$$C^* := \min_{\mathbf{y} \geq 0} \sup_{f \in \mathcal{F}_{\geq 0}} \mathbb{E}_f\left[C(\mathbf{y}, \tilde{\mathbf{D}})\right]. \tag{5}$$

We denote the optimal value of (5) as C^* and its optimal solution as $\mathbf{y}^*$. If the firm chooses the initial inventory level $\mathbf{y}^*$, then it can be guaranteed that the expected cost is no larger than C^* under any joint demand distribution with mean-covariance $(\mathbf{m}, \mathbf{\Sigma})$. Since the inventory levels are chosen to be robust to any specific distribution, we also refer to (5) as the *distributionally robust* multi-location newsvendor problem with inventory risk pooling.

The remainder of the paper is devoted to the distributionally robust problem (5). For the special case of $s_{ij} > h + p$ for all $i \neq j$, it is never optimal to use cross-location fulfillment (i.e., $z_{ij} = 0$ for any $i \neq j$ and $z_{ii} = \min(d_i, y_i)$), so the cost reduces to n separable single-location newsvendor costs (particularly, $C(\mathbf{y}, \mathbf{D}) = \sum_{i \in [n]} [h(y_i - d_i)^+ + (p - s_0)(d_i - y_i)^+ + s_0 d_i]$). Note that in this special case, while the newsvendor cost is separable by location, the minmax robust problem (5) is not due to the joint constraints (4) on the joint probability distribution. Hanasusanto et al. (2015) proved that a minmax robust problem of n single-location newsvendor costs with joint mean and covariance information is NP-hard even in the absence of constraints on the distribution support.

In the general case where $s_{ij} \leq h + p$, the introduction of network flows only serve to complicate the problem further. Hence, the remainder of this paper is focused on developing computationally tractable heuristics for the minmax robust counterpart of the multi-location newsvendor problem.

The rest of the paper is organized as follows. In Section 3, we analyze the special case of two ($n = 2$) identical locations. Section 4 is devoted to developing computationally tractable heuristics for the multi-location case under the nested fulfillment cost structure. In Section 5, we analyze the multi-location heuristic solutions numerically to understand the effect of additional information and to test the performance of the heuristic solutions, as well as to illustrate the tightness of the nested fulfillment approximation. Extensions and future directions follow in Sections 6 and 7.

3. The Two-Location Problem

We discuss in this section the special case of a two-location problem ($n = 2$), where the locations are identical: they have the same cost parameters (h, p, $s_{12} = s_{21} = s > s_0$) and the demands $\tilde{d}_1, \tilde{d}_2$, though correlated, have the same mean and standard deviation. For this special case, we are able to derive an analytic expression for the worst-case expected cost. The analytic expression allows us to not only gain insight into the demand scenarios that are most harmful to the retailer, but also understand how the worst-case distribution is affected by the cost and demand parameters.

Suppose that m and σ are the mean and standard deviation, respectively, of both demands. We denote the correlation coefficient by ρ. We assume that the joint distribution f belongs to the set:

$$\mathcal{F}^{m\sigma\rho} := \left\{ f : \Re^2 \mapsto \Re \;\middle|\; \begin{array}{l} \mathbb{E}_f(1) = 1, \\ \mathbb{E}_f(\tilde{d}_1) = \mathbb{E}_f(\tilde{d}_2) = m, \\ \mathbb{E}_f(\tilde{d}_1^2) = \mathbb{E}_f(\tilde{d}_2^2) = m^2 + \sigma^2, \\ \mathbb{E}_f(\tilde{d}_1 \tilde{d}_2) = m^2 + \rho\sigma^2, \end{array} \right\}.$$

The symmetry of the problem yields the following Lemma.

LEMMA 1. *For the identical two-location newsvendor problem with inventory pooling, there exists optimal inventory levels in each node that are also identical of the form* $\mathbf{y} = (y, y)^\intercal$ *for some* $y \geq 0$.

The proof is relegated to the Appendix, and relies on the convexity of the min-max robust problem in (5). This property aids us in simplifying the derivation, since we can restrict our analysis to only inventory decisions where the quantity is identical in both locations. Note that from the definition of $\mathcal{F}^{m\sigma\rho}$, we allow the demands to take all values in $\Re^2$, including negative values.[5] The later sections addressing the multi-locaton problem, however, assume a nonnegative support.

[5] If we assume a nonnegative support for the distribution f, the derivation of the analytic expression becomes more complicated; to simplify the derivation, we assume that the support is $\Re^2$ for the two-location problem.

Given our condition that $p+h>s$, it is optimal for the firm to use all its inventory towards meeting its demand to the maximum possible extent. Therefore, for any inventory level $\mathbf{y}=(y,y)$ and demand realization $\mathbf{D}=(d_1,d_2)$, we can write the newsvendor cost (2) as

$$C(\mathbf{y},\mathbf{D}) = p\cdot(d_1+d_2-2y)^+ + h\cdot(2y-d_1-d_2)^+ \\ + s_0\cdot\min(d_1,y) + s_0\cdot\min(d_2,y) + s\cdot\left[(d_1-y)^+ + (d_2-y)^+ - (d_1+d_2-2y)^+\right].$$

Note that the first term on the right-hand side is the penalty cost multiplied by the total unmet demand after using the $2y$ total inventory units. The second term is the overage cost multiplied by the total unsold inventory. The third and fourth terms are the in-location fulfillment costs, and the last term is the product of the cross-location fulfillment cost and the total units shipped to a cross-location. The number of cross-fulfilled units can be determined by the decrease in unmet demand by going from an unpooled system to a pooled system. Using the relationship $(a-b)^+ = a-\min(a,b)$, we can rewrite the newsvendor cost as:

$$C(\mathbf{y},\mathbf{D}) = p(d_1+d_2) + 2hy - (p+h-s)\min(d_1+d_2,2y) - \sum_{j=1,2}(s-s_0)\min(d_j,y). \tag{6}$$

Hence, when the cross-location fulfillment costs are identical between the two locations, it is possible to write the cost function as a piecewise linear convex function. This is critical in aiding our analysis of the worst-case expected cost, which allows us to prove the following theorem.

THEOREM 1. *For the identical two-location newsvendor problem with inventory pooling, if the cross-location fulfillment costs are constant and equal to s, and if $\sigma>0$ and $\rho\in(-1,1)$, then for any $\mathbf{y}=(y,y)$,*

$$\sup_{f\in\mathcal{F}^{m\sigma\rho}} \mathbb{E}_f[C(\mathbf{y},\tilde{\mathbf{D}})] \le \bar{C}(y) := 2s_0m - (p-h-s_0)(y-m) + (p+h-s_0)\sqrt{(y-m)^2+\gamma\sigma^2}, \tag{7}$$

where $\gamma := \frac{(p+h-s)(1+\rho)+s-s_0}{2(p+h)-s-s_0}$. Moreover, if $\gamma(\nu^2+1)\ge 2$, where $\nu := \frac{3(h+p-s_0)-2(s-s_0)}{h+p-s_0}$, then $\mathbb{E}_{f_y^*}[C(\mathbf{y},\tilde{\mathbf{D}})] = \bar{C}(y)$ for some probability distribution $f_y^*\in\mathcal{F}^{m\sigma\rho}$ with six support points.

An implication of Theorem 1 is that under some condition on the cost parameters, the worst-case expected cost of inventory level $\mathbf{y}=(y,y)$: $\max_{f\in\mathcal{F}^{m\sigma\rho}}\mathbb{E}_f[C(\mathbf{y},\tilde{\mathbf{D}})] = \bar{C}(y)$ is attained under a six-point distribution f_y^* (subscript y to emphasize that the worst-case distribution depends on the inventory level). We prove Theorem 1 in the appendix, however, we provide a discussion here.

Since the cost is equivalent to (6), under any distribution $f\in\mathcal{F}^{m\sigma\rho}$, the worst-case expected cost is bounded above by $2(pm+hy)-M(y)$, where $M(y)$ is the optimal value of the moment problem:

$$M(y) = \inf_{f\in\mathcal{F}^{m\sigma\rho}} \mathbb{E}_f\left[(p+h-s)\min\left(\tilde{d}_1+\tilde{d}_2,2y\right) + \sum_{j=1,2}(s-s_0)\min\left(\tilde{d}_j,y\right)\right] \tag{8}$$

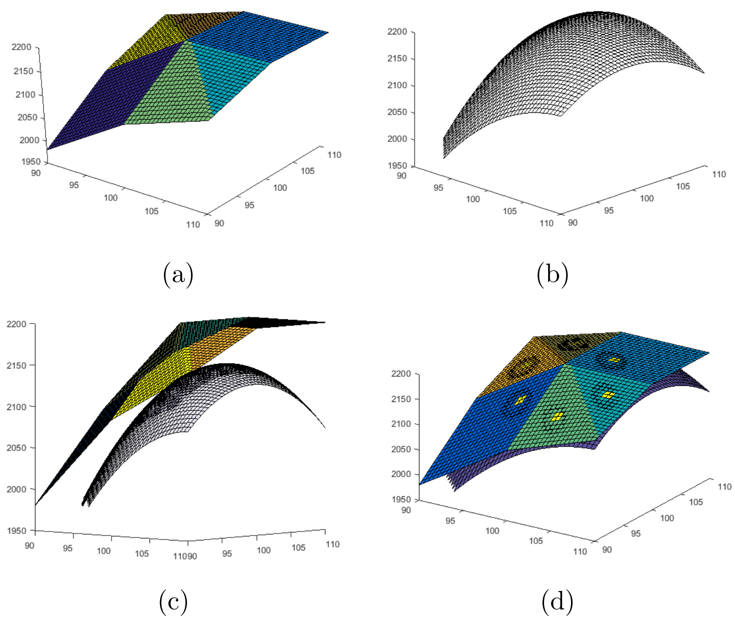

Figure 1 Illustration of the dual program. (a) The piecewise-planar function $q(d_1, d_2)$, (b) The quadratic function $g(d_1, d_2)$, (c) The functions corresponding to a dual feasible solution, (d) The functions corresponding to the dual optimal solution.

Then, to prove the first part of the theorem, we need to show that $M(y) \geq (h+p-s_0)(y+m-\sqrt{(y-m)^2+\gamma\sigma^2})$. We note that the right-hand side of (8) is a semi-infinite linear program since the distribution f corresponds to infinitely many variables. By weak duality, it is bounded below by its dual program, which takes the following simple form due to the moments of the marginal distributions being identical:

$$\begin{array}{lll} \max\limits_{t,u,r,v} & t + 2mr + 2(m^2+\sigma^2)u + (m^2+\rho\sigma^2)v & \\ \text{s.t.} & g(d_1,d_2;t,u,r,v) \leq q(d_1,d_2;t,u,r,v) \quad \forall (d_1,d_2) \in \Re^2 & \end{array} \tag{9}$$

where $g(d_1,d_2) := t + r(d_1+d_2) + u(d_1^2+d_2^2) + vd_1d_2$ and $q(d_1,d_2) := (h+p-s)\min(d_1+d_2, 2y) + (s-s_0)\min(d_1,y) + (s-s_0)\min(d_2,y)$.

Note that the dual program (9) has infinitely many constraints. The dual variables t,u,r,v are the parameters of a piecewise-planar function $q(d_1,d_2)$ (shown in Figure 1a) and a biquadratic function $g(d_1,d_2)$ (shown in Figure 1b). The dual variables are feasible if the biquadratic function is bounded above by the piecewise-planar function in all of $\Re^2$ (Figure 1c shows these functions corresponding to a dual feasible solution). The dual feasible solution corresponding to the two functions touching at exactly six points in $\Re^2$ (shown as the bright points in each face of the piecewise-planar function in Figure 1d) has a dual objective value of $(h+p-s_0)(y+m-\sqrt{(y-m)^2+\gamma\sigma^2})$, which is a lower bound on the dual optimal value, and by weak duality, is a lower bound on $M(y)$. If $\gamma(\nu^2+1) \geq 2$, then we can use the six intersection points to construct a feasible distribution of (8) whose primal

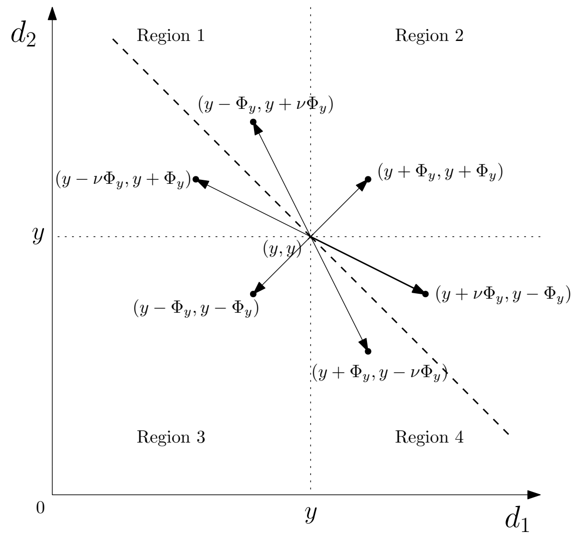

Figure 2 The six support points of the worst-case distribution that maximizes the expected cost of inventory levels $\mathbf{y} = (y, y)$. Each support point lies on one of three solid lines passing through (y, y), where the distance from this point increases in proportion to $\Phi_y = \sqrt{(y-m)^2 + \gamma\sigma^2}$. Demand realizations falling on the dashed line correspond to a perfect balance between demand and supply after cross-location fulfillment. Regions 1 and 4 (demarcated by dotted lines) induce imbalance in the system as, after in-location fulfillment, there is leftover inventory at one location, and unfulfilled demand at the other.

objective value is equal to $(h + p - s_0)(y + m - \sqrt{(y-m)^2 + \gamma\sigma^2})$. Hence, this distribution is the solution to the primal problem, and $M(y) = (h + p - s_0)(y + m - \sqrt{(y-m)^2 + \gamma\sigma^2})$.

Since we did not impose a restriction on the support of the distribution in $\mathcal{F}^{m\sigma\rho}$, then the dual feasible condition is that $g(d_1, d_2) \leq q(d_1, d_2)$ for all $(d_1, d_2) \in \Re^2$. If we only allowed demand to take nonnegative values, then this condition only needs to be true in $\Re^2_{\geq 0}$. While it is possible to use the same approach outlined above to analytically determine the support points under this restriction, it makes the derivation more cumbersome since it requires careful consideration of different cases of the dual variables. As demonstrated by the proof of Theorem 1 in the appendix, even for the simpler unrestricted case in the two-location problem, the proof is already cumbersome. This highlights the need for tractable approximations and heuristics for general cases beyond the identical two-location problem with unbounded support, which we develop in later sections of this paper.

3.1. The Effect of Cost and Demand Parameters on the Worst-case Distribution

We provide in the Appendix the exact expressions for the support points and the probabilities of the discrete distribution f^*_y which results in the worst-case expected cost when $\gamma(\nu^2 + 1) \geq 2$. Figure 2 shows the support points of f^*_y, given the order quantity (y, y), where $\Phi_y = \sqrt{(y-m)^2 + \gamma\sigma^2}$. From the figure, we observe that each support point lies on one of three solid lines that pass through point (y, y). When we either increase $|y - m|$ or increase σ or increase ρ, the distance of the support points to (y, y) increases proportionally in the direction indicated by the arrows in Figure 2.

Note that for any demand realization that falls on the single dashed line in the figure, there is no excess inventory nor unmet demand after pooling since $d_1 + d_2 = 2y$. Two of the solid lines converge to this dashed line as $\nu \to 1$, which occurs when $h + p$ decreases to the limit s. Instead, when the overage or the underage cost is high (or the fulfillment cost s is low), the solid lines pivot further away from the dashed line, and hence the support points of the worst-case distribution result in very large excess inventory or unmet demand after pooling.

We call a system imbalanced if, after in-location fulfillment, there is leftover inventory in one location and unfulfilled demand at the other location. In Figure 2, we divide the demand region into four quadrants (Regions 1 through 4 demarcated by the two dotted lines), and observe that when the demand realizations are in Regions 1 and 4, there is imbalance in the system. In these regions, we measure the magnitude of imbalance as the sum of the leftover inventory and unfulfilled demand. Mathematically, this is the L^1 distance between the support point and (y, y). We find that the magnitude of imbalance induced by the support points is increasing in $|y - m|$, σ and ρ.

The probability of imbalance (sum of probabilities of realizations in Regions 1 and 4), however, is no more than $\frac{\gamma\sigma^2}{\Phi_y^2}$. When the decision maker chooses inventory levels with high safety stock (as is typically the case in practice), the worst case distribution causes imbalance across locations with low probability, but of large magnitude. Thus retailers should strive to eliminate such low-probability extreme situations, which can be done by adopting various demand-shaping strategies for different customer locations, such as strategic product display, recommendations and flash promotions.

3.2. The Optimal Robust Inventory Solution

When $\gamma(\nu^2 + 1) \geq 2$, the worst-case expected cost of inventory levels $\mathbf{y} = (y, y)$ is given by $\bar{C}(y)$. Thus, a corollary of Theorem 1 is an analytic expression for the optimal inventory level under a minmax robust criterion.

PROPOSITION 1. *For the two-location newsvendor problem with inventory risk pooling, if cross-location fulfillment costs are constant and equal to s, if $\sigma > 0$ and $\rho \in (-1, 1)$, and if $\gamma(\nu^2 + 1) \geq 2$ where $\gamma := \frac{(p+h-s)(1+\rho)+s-s_0}{2(p+h)-s-s_0}$ and $\nu := \frac{3(h+p-s_0)-2(s-s_0)}{h+p-s_0}$, then the inventory levels that minimize the maximal expected cost over distributions in $\mathcal{F}^{m\sigma\rho}$ is $\mathbf{y}^* = (y^*, y^*)$, where*

$$y^* = m + \left(\frac{p - h - s_0}{2} \sqrt{\frac{\gamma}{h(p - s_0)}} \right) \sigma. \tag{10}$$

The minmax expected cost is $C^* = 2s_0 m + 2\sigma\sqrt{\gamma h(p - s_0)}$.

The proof is relegated to the appendix. The optimal worst-case cost $C^* = 2s_0 m + 2\sigma\sqrt{\gamma h(p - s_0)}$ consists of two terms – one corresponding to the case when there is no uncertainty, in which case the optimal cost would simply be the in-location fulfillment cost times the expected demand, and

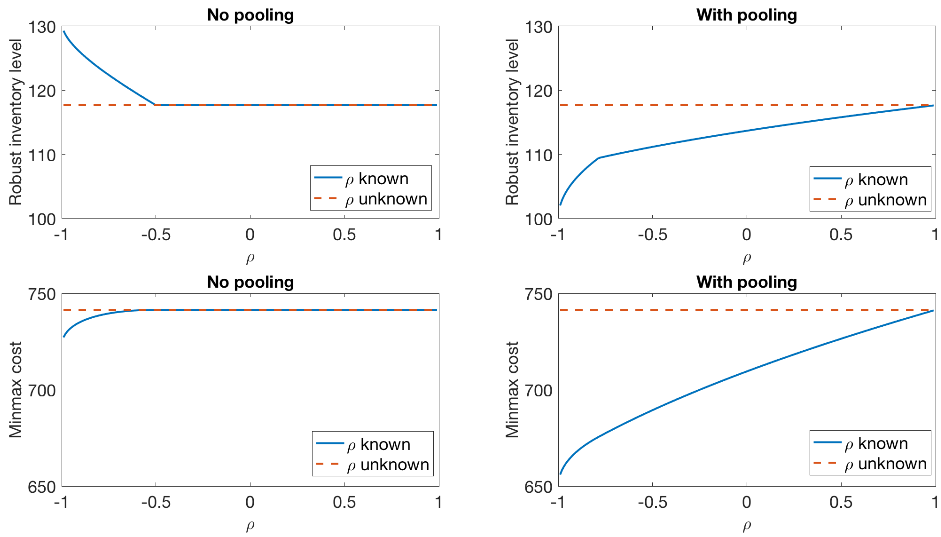

Figure 3 **The minmax cost (bottom row) and the robust inventory level (top row) for problems with and without additional knowledge of correlation ρ. The panels on the left are for the problem with no pooling, and the right panels are if cross-location fulfillment is allowed. Parameters are $m = 100$, $\sigma = 60$, $p = 5$, $h = 1$, $s_0 = 3$, $s = 4$.**

the other term arising as a result of uncertainty, which is independent of the average demand and depends only on the standard deviation.

Proposition 1 gives the optimal minmax quantity if the inventory risk of the two locations are pooled through cross-location fulfillment. On the other hand, if cross-location fulfillment is not allowed or if $s > p + h$, then inventory is decentralized with each location solving a classical single-location newsvendor problem. In the seminal work by Scarf (1958), an analytic solution to the classical minmax newsvendor problem with known mean and variance is found to be

$$y^*_{Sc} = m + \left(\frac{p - h - s_0}{2\sqrt{h(p - s_0)}} \right) \sigma. \tag{11}$$

Since $\gamma \leq 1$, it directly follows that $y^* \leq y^*_{Sc}$ whenever $p - s_0 \geq h$, and that $y^* \geq y^*_{Sc}$ whenever $p - s_0 \leq h$. Since $p - s_0$ is the underage cost without pooling, then when $p - s_0 \geq h$, the decentralized solution y^*_{Sc} is large due to the high underage cost. On the other hand, in a centralized system, unmet demand can be fulfilled by inventory from any location, so the solution y^* is lower when $p - s_0 \geq h$. For a similar reason, when $p - s_0 \leq h$, the decentralized solution is low due to the high overage cost. In a centralized system, excess inventory in one location can be used elsewhere, so the solution y^* is higher when $p - s_0 \leq h$. Hence, the fact that $|y^* - m| \leq |y^*_{Sc} - m|$ is because of inventory risk sharing resulting from pooling, mirroring similar results from stochastic systems. Indeed, as $\rho \to 1$, we have $\gamma \to 1$, and as a result, y^* converges to the decentralized solution y^*_{Sc}.

3.3. The Role of Correlation

We next discuss the role of correlation information in the minmax robust inventory problem with multiple locations. Consider first the inventory problem when cross-location fulfillment is not allowed (i.e., no pooling). The left panels of Figure 3 shows the minmax cost and the robust inventory levels under two situations: if ρ is known, and if ρ is unknown. The horizontal axis corresponds to the actual value of ρ. For each value of ρ, we solve for the minmax cost and the robust inventory level exactly, which can be done through solving a semidefinite program. Note that when ρ is unknown, then the robust inventory level is equal to y^*_{Sc}, the Scarf (1958) inventory level. We observe from the figure that when ρ is greater than -0.55, then there is no difference (in minmax cost or robust inventory levels) if the firm learned the true value of the correlation. Since inventory risk is not shared between locations, additional correlation information does not reduce the worst-case expected cost, which from the figure, we surmise to occur for a set of demand distributions whose correlation coefficient is greater than -0.55. This is consistent with the observation from Figure 9 of Natarajan et al. (2018) that having additional covariance information does not significantly change the robust inventory levels since there is no pooling.

In the multi-item newsvendor without pooling (as considered by Natarajan et al. 2018), information on the joint distribution is only used to narrow down the distribution set, which would result in less conservative solutions. So, as we see from the figure, the effect of joint information on the minmax cost is nominal. However, when cross-location fulfillment is allowed (right panels of Figure 3), correlation information can significantly reduce the minmax cost, as cross-fulfillment critically depends on correlation between demands, and in turn affects the inventory solutions. We observe that with pooling, if the correlation coefficient is unknown, the minmax cost and robust inventory level match the case with no pooling. This is because the adversary is going to choose a distribution where inventory pooling would have the least benefit for the firm, i.e., when demands in the two locations are perfectly correlated. With pooling, correlation information can reduce the minmax cost, since inventory risk pooling has the most benefit when locations are negatively correlated and there is high likelihood of inventory imbalance that can be corrected by cross-location fulfillment. This explains why the robust inventory levels decrease as correlation decreases.

This discussion serves to motivate the importance of knowing the covariance of demands between locations when inventory risk is pooled. Covariance information not only reduces the minmax cost, but also allows the firm to remain distributionally robust but still utilize the covariance information to benefit from pooling by reducing the safety stock. While we demonstrate this in the two-location case, these effects are amplified if there are more locations.

4. The Multi-Location Problem

In the previous section, if there are only two locations that are identical, then an analytic expression for the worst-case expected cost can be derived. However, in the general case with n locations, even if pooling is not allowed, the minmax robust problem with mean and covariance is known to be NP-hard (Hanasusanto et al. 2015). In this section, we will develop a computationally tractable upper bound to the worst-case expected cost in a multi-location newsvendor problem with pooling. Unlike the previous section, we consider a non-negative support, and we do not make any assumption on whether the locations have identical marginal demand distributions. We demonstrate the upper bound to be empirically tight, and it also yields a heuristic for the robust inventory levels.

In developing the upper bound, we utilize an observation from the two-location case. That is, if there is a special structure to the cost parameters, then the cost $C(\mathbf{y},\mathbf{D})$, which is the optimal value of a network flow problem, can be written analytically as a piecewise-linear convex function. In this section, we explore this further for a more a general class of fulfillment cost structures, namely nested fulfillment structures. We demonstrate that the worst-case expected cost can be bounded by the optimal value of a semidefinite program (SDP). We also develop an algorithm in Section 4.3 that approximates a general distance-based shipping cost structure to the nested fulfillment cost structure, which we empirically show to approximate the expected total shipping cost well under a variety of distributions in Section 5.

4.1. The Nested Fulfillment Cost Structure

In general, the minmax robust problem under a non-identical fulfillment cost structure is intractable. However, for special cases which we refer to as *nested fulfillment structures*, we can derive a computationally tractable and empirically tight upper bound on the worst-case expected cost. Specifically, for this cost structure, we will show that we can write $C(\mathbf{y},\mathbf{D})$ as the sum of only linear terms and functions of the form $(d_i - y_i)^+$ or $\left(\mathbf{e}^\top\mathbf{D} - \mathbf{e}^\top\mathbf{y}\right)^+$.

To introduce the nested fulfillment structure, consider the following sequential stages of fulfilling a demand realization $\mathbf{D} = (d_i)$.

1. In-location fulfillment (Level 0): Each inventory node $i \in [n]$ first fulfills demand within the same location i, incurring a per-unit fulfillment cost s_0.
2. Zone fulfillment (Level 1): We assume that the n customer locations can be partitioned into n_1 zones (i.e., each location belongs to exactly one zone). Let $\mathcal{I}_k \subseteq [n]$ be the set of customer locations mapped to zone $k \in [n_1]$, where $\cup_{k\in[n_1]}\mathcal{I}_k = [n]$. The second fulfillment stage is for remaining inventory in a zone to fulfill unmet demand at customer locations in the same zone, incurring a per-unit fulfillment cost s_1.
3. Network fulfillment (Level 2): The final stage is for any unfulfilled demand after the second stage to be fulfilled from any remaining inventory from any node, with a per-unit cost s_2.

We call the above structure as the *3-level nested fulfillment cost structure.* Naturally, we need $s_0 < s_1 < s_2$, so that a higher level fulfillment is used only if fulfillment in a lower level is not possible due to lack of inventory. Moreover, we assume that $s_2 < p + h$ so that it is cheaper to use a unit of inventory towards fulfilling an unmet demand in any location than to keep in an inventory

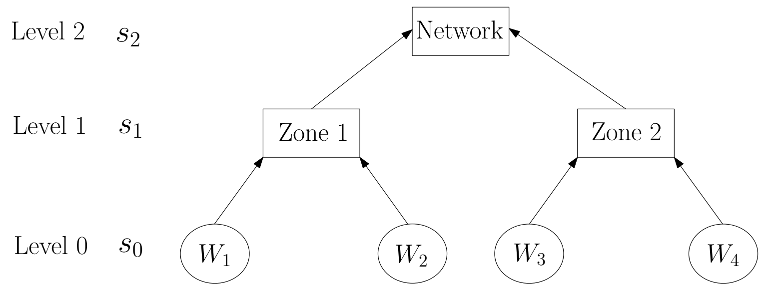

Figure 4 Example of a 3-level nested fulfillment cost structure with four warehouses $W_1, .., W_4$**.**

node. In general, the nested structure of the fulfillment network aids in modeling the cost function without the need for a linear program. We demonstrate this with a simple example.

EXAMPLE 2. Consider the 3-level nested fulfillment structure shown in Figure 4, where W_1 and W_2 are warehouses in Zone 1, and W_3 and W_4 are warehouses in Zone 2. Hence, $n = 4$ and $n_1 = 2$. Suppose the starting inventory in each warehouse is 10 units ($y_i = 10$ for all i), and the demand realization in the four locations are $d_1 = 15$, $d_2 = 8$, $d_3 = 3$, $d_4 = 14$. The sequence implied by the fulfillment cost structure is as follows. First, each warehouse tries to fulfill its own demand with shipping cost s_0. Second is zone-based fulfillment. There are 5 units of unmet demand at location 1, of which 2 units are fulfilled from the remaining inventory at W_2, at a per-unit cost of s_1. In Zone 2, there are 4 unmet demand units (from location 4), which are all fulfilled by the excess 7 inventory units at W_3. Third is the network fulfillment. The remaining 3 units of unfulfilled demand in Zone 1 are routed to the network, and can be fulfilled from Zone 2 which has excess inventory of 3 units, at a per-unit cost of s_2.

The total overage cost and the total penalty cost are straightforward, as they are equal to $h \cdot (\mathbf{e}^\top \mathbf{y} - \mathbf{e}^\top \mathbf{D})^+$ and $p \cdot (\mathbf{e}^\top \mathbf{D} - \mathbf{e}^\top \mathbf{y})^+$, respectively. Using the identity $(a-b)^+ = a - \min(a,b)$, the fulfillment costs can be computed as:

- In-location fulfillment: $s_0 \sum_{i\in[n]} \min(d_i, y_i) = s_0 \cdot \left[\mathbf{e}^\top \mathbf{D} - \sum_{i\in[n]} (d_i - y_i)^+\right]$,
- Zone 1 fulfillment: $s_1 \cdot [(d_1 - y_1)^+ + (d_2 - y_1)^+ - (d_1 + d_2 - y_1 - y_2)^+]$,
- Zone 2 fulfillment: $s_1 \cdot [(d_3 - y_3)^+ + (d_4 - y_4)^+ - (d_3 + d_4 - y_3 - y_4)^+]$,
- Network fulfillment $s_2 \cdot [(d_1 + d_2 - y_1 - y_2)^+ + (d_3 + d_4 - y_3 - y_4)^+ - (\mathbf{e}^\top \mathbf{D} - \mathbf{e}^\top \mathbf{y})^+]$.

Note that the units shipped in each zone fulfillment is the reduction in unmet demand going from pooling locations in the zone to no pooling. The units shipped in the network fulfillment is the difference between unmet demand after zone fulfillment to unmet demand after pooling all locations. Thus, the total cost has only linear terms or functions of the form $g(x) = x^+$.

We can generalize the intuition from the example to nested fulfillment structures with more than three levels, and with non-identical fulfillment costs within each level. To do so, we introduce some notation. Consider an L-level nested fulfillment cost structure. In each level ℓ, where $\ell = 0, 1, \ldots, L-1$, assume that the locations in $[n]$ are partitioned into n_ℓ sets, where the level ℓ partition is $\{\mathcal{I}_1^{(\ell)}, \mathcal{I}_2^{(\ell)}, \ldots, \mathcal{I}_{n_\ell}^{(\ell)}\}$. That is, $\mathcal{I}_k^{(\ell)} \cap \mathcal{I}_m^{(\ell)} = \emptyset$ if $k \neq m$, and $\cup_{k\in[n_\ell]} \mathcal{I}_k^{(\ell)} = [n]$. If two locations i_1 and i_2 are in $\mathcal{I}_k^{(\ell)}$ for some $k \in [n_\ell]$, then fulfillment between the two locations incurs a per-unit cost $s_{\ell,k}$. We also assume that each set $\mathcal{I}_k^{(\ell)}$ is a union of sets in level $\ell-1$. That is, $\mathcal{I}_k^{(\ell)} = \cup_{m\in\mathcal{K}_k^{(\ell)}} \mathcal{I}_m^{(\ell-1)}$ where $\mathcal{K}_k^{(\ell)}$ is the level $\ell-1$ components of set k. Note that $n_0 > n_1 > \cdots > n_{L-1}$, since there are fewer sets in higher order levels. We assume that for level 0, $n_0 = n$ and $\mathcal{I}_i^{(0)} = \{i\}$ for all $i \in [n]$. For the final level $L-1$, $n_{L-1} = 1$ and $\mathcal{I}_1^{(L-1)} = [n]$. Any L-level nested fulfillment structure can be represented as a L-level tree similar to Figure 4.

We further assume that it is less costly to fulfill demand from location $i \in [n]$ using lower level fulfillment. Mathematically, if $k^{(\ell)}(i)$ is the level ℓ set index of location i, then we assume that $s_{0,k^{(0)}(i)} \leq s_{1,k^{(1)}(i)} \leq \cdots \leq s_{L-1,k^{(L-1)}(i)}$, with at least one of the inequalities holding strictly such that $s_{L-1,k^{(L-1)}(i)} > s_{0,k^{(0)}(i)}$ (as otherwise, the problem reduces to a single location problem). Note that all the inequalities need not be strict, as some components may not be pooled in subsequent levels.[6] Finally, we also have $s_{L-1,k^{(L-1)}(i)} < p + h$.[7]

We define the level ℓ assignment matrix $\mathbf{E}_\ell$, where $\ell = 0, 2, \ldots, L-1$, which is a binary matrix of size $n_\ell \times n$ where the (k,i) entry is equal to 1 if and only if $i \in \mathcal{I}_k^{(\ell)}$. Note that $\mathbf{E}_0$ is the $n \times n$ identity matrix, and that $\mathbf{E}_{L-1}$ is the row vector of all ones. The usefulness of the nested fulfillment structure is demonstrated by the simplicity in modeling the cost function through piecewise linear terms, which leads to the following Lemma.

LEMMA 2. *Given an L-level nested fulfillment cost structure, the distributionally robust multi-location newsvendor problem with inventory risk pooling (5) is equivalent to*

$$C^* = \min_{\mathbf{y}} \left(h \cdot \mathbf{e}^\top (\mathbf{y} - \mathbf{m}) + \mathbf{s}_0^\top \mathbf{m} + M_L(\mathbf{y}) \right), \tag{12}$$

where

$$M(\mathbf{y}) = \max_{f\in\mathcal{F}_{\geq 0}} \quad \mathbb{E}_f \left[\sum_{\ell=0}^{L-1} \eta_\ell^\top \left(\mathbf{E}_\ell \tilde{\mathbf{D}} - \mathbf{E}_\ell \mathbf{y} \right)^+ \right]. \tag{13}$$

where $\eta_\ell = (\eta_{\ell,k})_{k\in[n_\ell]}$, and $\eta_{\ell,k} = s_{\ell+1,m^{(\ell+1)}(k)} - s_{\ell,k}$ for $\ell = 0, \ldots, L-2$, with $\eta_{L-1,1} = p + h - s_{L-1,1}$, and $\mathbf{E}_\ell$ is the level ℓ assignment matrix.

[6] If for some ℓ and k we have $\mathcal{K}_k^{(\ell)} = 1$, then $\mathcal{I}_k^{(\ell)} = \mathcal{I}_{m^{(\ell+1)}(k)}^{(\ell+1)}$, where $m^{(\ell+1)}(k) \in [n_{\ell+1}]$ is the level $\ell+1$ parent of set $k \in [n_\ell]$. Then, for all $i \in \mathcal{I}_k^{(\ell)}$, we will have $s_{\ell,k^{(\ell)}(i)} = s_{\ell+1,k^{(\ell+1)}(i)}$. In essence, the structure allows for dummy nodes, which is illustrated in Figure 6b.

[7] We note that this assumption can be relaxed for nested fulfillment structures. If for any i, ℓ we have $s_{\ell,k^{(\ell)}(i)} > h + p$, then all children nodes of $k^{(\ell)}(i)$ can be removed to form separate networks as it is never optimal to fulfill demands at level ℓ for these nodes. The total cost is simply the sum of piecewise linear costs of all the networks.

Note that $\eta_{\ell,k} \geq 0$, and can be interpreted as the marginal benefit of fulfillment in level ℓ instead of level $\ell+1$ of any demand occuring in $\mathcal{I}_k^{(\ell)}$. Similarly, $\eta_{L-1,1}$ is the marginal benefit using a unit of inventory for fulfillment with the highest cost instead of holding onto the unit. We provide a sketch of the proof of the Lemma below.

To obtain the fulfillment cost of a demand realization $\mathbf{D}$ in an L-level nested fulfillment cost structure, we sum the fulfillment costs in each level. The total fulfillment cost in level 0 is $\sum_{i\in[n]} s_{0,i} \cdot \min(d_i, y_i) = \mathbf{s}_0^\top \mathbf{D} - \sum_{i\in[n]} s_{0,i} \cdot (d_i - y_i)^+$, where $\mathbf{s}_0 = (s_{0,i})_{i\in[n]}$. The total number of units fulfilled in level ℓ, where $\ell = 1, 2, \ldots, L-1$, after pooling inventory in set $\mathcal{I}_k^{(\ell)}$, where $k \in [n_\ell]$, is

$$\underbrace{\sum_{m\in\mathcal{K}_k^{(\ell)}} \left(\sum_{i\in\mathcal{I}_m^{(\ell-1)}} d_i - \sum_{i\in\mathcal{I}_m^{(\ell-1)}} y_i \right)^+}_{\text{unmet demand in } \mathcal{I}_k^{\ell} \text{ after level } \ell-1} - \underbrace{\left(\sum_{i\in\mathcal{I}_k^{(\ell)}} d_i - \sum_{i\in\mathcal{I}_k^{(\ell)}} y_i \right)^+}_{\text{unmet demand in } \mathcal{I}_k^{(\ell)} \text{ after level } \ell}, \tag{14}$$

for a per-unit fulfillment cost $s_{\ell,k}$. Since $p+h$ is strictly greater than all fulfillment costs, then the penalty cost is $p \cdot (\mathbf{e}^\top \mathbf{D} - \mathbf{e}^\top \mathbf{y})^+$, and the overage cost is $h \cdot (\mathbf{e}^\top \mathbf{y} - \mathbf{e}^\top \mathbf{D})^+$. Therefore, the total cost (overage, penalty and fulfillment) is equal to

$$\begin{aligned} C(\mathbf{y}, \mathbf{D}) = h \cdot \mathbf{e}^\top (\mathbf{y} - \mathbf{D}) + \mathbf{s}_0^\top \mathbf{D} + (p + h - s_{L-1,1}) \cdot \left(\mathbf{e}^\top \mathbf{D} - \mathbf{e}^\top \mathbf{y}\right)^+ \\ + \sum_{\ell=0}^{L-2} \sum_{k\in[n_\ell]} (s_{\ell+1, m^{(\ell+1)}(k)} - s_{\ell,k}) \cdot \left(\sum_{i\in\mathcal{I}_k^{(\ell)}} d_i - \sum_{i\in\mathcal{I}_k^{(\ell)}} y_i \right)^+, \end{aligned} \tag{15}$$

where $m^{(\ell+1)}(k) \in [n_{\ell+1}]$ is the level $\ell+1$ parent of set $k \in [n_\ell]$. Substituting (15) in (5), and using compact notation with the parameters η_ℓ and assignment matrices $\mathbf{E}_\ell$, we obtain the Lemma.

4.2. Tractable Upper Bound and Heuristic Solution

We now focus on deriving a tractable upper bound and heuristic solution for the NP-hard distributionally robust multi-location newsvendor problem with inventory risk pooling under a nested fulfillment cost structure, as described in Lemma 2. First, while allowing for negative demand outcomes is often unrealistic, it is worth discussing because it allows for an exact solution to the minmax robust problem in Lemma 2 through a semidefinite program (SDP) with 2^{N+1} semidefinite constraints, where $N := \sum_{\ell=0}^{L-1} n_\ell$ is the total number of nodes in the tree representation of the nested fulfillment structure (similar to Figure 4). We show this in the following lemma.

Lemma 3. *For the n-location newsvendor problem under inventory risk pooling with a nested L-level fulfillment cost structure, if $\mathcal{F}$ is the set of all joint probability distributions of the n-location demand $\tilde{\mathbf{D}}$ with mean vector $\mathbf{m}$ and covariance matrix $\mathbf{\Sigma} \succ 0$, then for any $\mathbf{y} \in \Re^n$,*

$$\begin{aligned} \sup_{f\in\mathcal{F}} \mathbb{E}_f[C(\mathbf{y}, \tilde{\mathbf{D}})] = \underset{t,\mathbf{r},\mathbf{Y},\mathbf{y}}{\text{minimize}} \quad & h \cdot \mathbf{e}^\top (\mathbf{y} - \mathbf{m}) + s_0 \cdot \mathbf{e}^\top \mathbf{m} + t + \mathbf{r}^\top \mathbf{m} + \langle \mathbf{Y}, \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top \rangle \\ \text{subject to} \quad & \begin{pmatrix} \mathbf{Y} & \frac{1}{2}(\mathbf{r} - \mathbf{a}) \\ \frac{1}{2}(\mathbf{r} - \mathbf{a})^\top & t + \mathbf{a}^\top \mathbf{y} \end{pmatrix} \succeq 0, \quad \forall \mathbf{a} \in \mathcal{L}, \end{aligned} \tag{16}$$

where $\mathcal{L} := \left\{\mathbf{a} \,|\, \mathbf{a} = \sum_{\ell=0}^{L-1} \mathbf{E}_\ell^\top \left(\eta_\ell \odot \mathbf{e}_{A_\ell}\right) \text{ for some } (A_0, A_1, \cdots, A_{L-1}) \in 2^{[n_0]} \times 2^{[n_1]} \times \cdots \times 2^{[n_{L-1}]}\right\}$, $\odot$ *is the element-wise product operator,* $\mathbf{E}_\ell$ *is the* $n_\ell \times n$ *binary assignment matrix in level* ℓ*, and* $\mathbf{e}_{A_\ell}$ *(where* $A_\ell \in 2^{[n_\ell]}$*) is an* n_ℓ*-dimensional binary vector where* $e_{A_\ell,k} = 1$ *if and only if* $k \in A_\ell$*.*

The proof is relegated to the Appendix. Similar to the proof of the closed-form expression for the two-location case, the proof of Lemma 3 relies on taking the dual of the moment problem (13). Since $\mathbf{\Sigma} \succ 0$, then the condition for strong duality of moment problems is met (Smith 1995). The dual program has a linear objective, and the constraints are that an n-dimensional quadratic function, $g(\mathbf{D})$, should be bounded below by a concave piecewise-linear function for all $\mathbf{D} \in \Re^n$. The piecewise-linear function is in fact the maximum of 2^N linear functions, $\max_{k\in[2^N]}\{q_k(\mathbf{D})\}$. Therefore, this is true if $g(\mathbf{D}) - q_k(\mathbf{D}) \geq 0$, for all $\mathbf{D} \in \Re^n$ and $k \in [2^N]$. Note that $g(\mathbf{D}) - q_k(\mathbf{D})$ is a quadratic function, which is nonnegative for all $\mathbf{D} \in \Re^n$ if and only if a matrix constructed from its coefficients is positive semidefinite. Thus, by strong duality, the moment problem is equivalent to a semidefinite program with 2^N SDP constraints. In the worst case, the number of nodes could be $N = \frac{n(n+1)}{2}$ if the number of levels is $L = n$ and $n_\ell = n - \ell$. While we cannot construct the worst-case joint probability distribution, it is known from Smith (1995) that it has support in at most $\frac{n(n+3)}{2} + 1$ discrete points. For the two location case ($n = 2$), there are at most six discrete support points, in agreement with our results from the previous section.

The implication of Lemma 3 is that the distributionally robust inventory level $\mathbf{y}^*$ under knowledge of mean vector $\mathbf{m}$ and covariance matrix $\mathbf{\Sigma}$, and without restriction on the distribution support, can be found by simply setting $\mathbf{y}$ as a decision variable in the semidefinite program in (16). Semi-definite programs, much like linear programs, can be solved through interior point methods which have polynomial time worst-case complexity (Vandenberghe and Boyd 1996).

While Lemma 3 provides an exact method to solve for the optimal distributionally robust inventory levels, the program is not computationally tractable beyond small values for n, since it involves 2^N constraints of the type that an $(n+1) \times (n+1)$ matrix is positive semidefinite. This motivates the need for tractable approximations to the distributionally robust problem for larger values of n that also work when the support is restricted to nonnegative values.

In the following proposition, we show that the worst-case expected cost under distribution set $\mathcal{F}_{\geq 0}$ is bounded above by the optimal value $\bar{C}(\mathbf{y})$ of a semidefinite program with a *single* semidefinite constraint of size $(N + n + 1) \times (N + n + 1)$. Hence, a heuristic for the minmax newsvendor problem with inventory pooling is to find the inventory levels which minimize the upper bound $\bar{C}(\mathbf{y})$.

Proposition 2. *For the n-location newsvendor problem under inventory risk pooling with an L-level nested fulfillment structure, we have* $\sup_{f\in\mathcal{F}_{\geq 0}} \mathbb{E}_f[C(\mathbf{y},\tilde{\mathbf{D}})] \leq \bar{C}(\mathbf{y})$ *for any* $\mathbf{y}\in\Re^n$*, where*

$$\begin{aligned}
\bar{C}(\mathbf{y}) := \min_{\substack{t_0,\mathbf{t},\mathbf{Y},\mathbf{u},\\ \mathbf{B},\mathbf{W},\mathbf{U},\mathbf{V}}} \quad & h\cdot\mathbf{e}^\top(\mathbf{y}-\mathbf{m}) + \mathbf{s}_0^\top\mathbf{m} + t_0 + \mathbf{t}^\top\mathbf{m} + \langle \mathbf{Y}, \mathbf{\Sigma}+\mathbf{m}\mathbf{m}^\top\rangle + \mathbf{e}^\top\mathbf{B}\mathbf{e} \\
\text{s.t.} \quad & \begin{pmatrix} t_0 & \frac{1}{2}\mathbf{t}^\top & \frac{1}{2}\mathbf{u}^\top \\ \frac{1}{2}\mathbf{t} & \mathbf{Y} & -\frac{1}{2}\mathbf{V}^\top \\ \frac{1}{2}\mathbf{u} & -\frac{1}{2}\mathbf{V} & \mathbf{U} \end{pmatrix} \succeq 0, \\
& \mathbf{u} = -\mathbf{W}\mathbf{e} + \left(\mathbf{B}+\mathbf{B}^\top\right)\mathbf{e} + \mathbf{P}\mathbf{y}, \\
& \mathbf{V} \geq \mathbf{P}, \\
& \mathbf{U} \leq \mathbf{W} - \mathbf{B}, \\
& \mathbf{W}, \mathbf{B} \geq 0, \\
& t_0\in\Re,\ \mathbf{t}\in\Re^n,\ \mathbf{u}\in\Re^N,\ \mathbf{Y}\in\Re^{n\times n},\ \mathbf{B},\mathbf{W},\mathbf{U}\in\Re^{N\times N},\ \mathbf{V}\in\Re^{N\times n},
\end{aligned} \tag{17}$$

with $\mathbf{P} := \left(\mathbf{E}_{L-1}^\top \mathrm{diag}(\eta_{L-1})\ \mathbf{E}_{L-2}^\top \mathrm{diag}(\eta_{L-2})\ \cdots\ \mathbf{E}_0^\top \mathrm{diag}(\eta_0)\right)^\top \in \Re^{N\times n}$.

In order to prove the proposition, we reformulate the moment problem in (13) as follows:

$$M(\mathbf{y}) = \max_{f\in\mathcal{F}_{\geq 0}} \ \mathbb{E}_f\left[\max_{\substack{\mathbf{x}^{(0)}\in\{0,1\}^{n_0},\\ \mathbf{x}^{(1)}\in\{0,1\}^{n_1},\cdots,\\ \mathbf{x}^{(L-1)}\in\{0,1\}^{n_{L-1}}}} \sum_{\ell=0}^{L-1} \left(\mathbf{x}^{(\ell)}\odot\eta_\ell\right)^\top \left(\mathbf{E}_\ell\tilde{\mathbf{D}} - \mathbf{E}_\ell\mathbf{y}\right)\right], \tag{18}$$

where $\mathbf{x}^{(\ell)}$ is an n_ℓ-dimensional binary vector, for $\ell = 0,1,\ldots,L-1$, and $\odot$ is the element-wise product operator. Given a demand realization $\mathbf{D}$, let $\{\mathbf{x}^{(\ell)}(\mathbf{D}),\ \ell=0,1,\ldots,L-1\}$ be the maximizer of the innermost function in (18). Since the value of the maximizer depends on the specific realization of a stochastic demand, the maximizers are random variables, which we will denote as $\{\tilde{\mathbf{x}}^{(\ell)},\ \ell=0,1,\ldots,L-1\}$. Note that for any ℓ and $k\in n_\ell$, since $\eta_{\ell,k}\geq 0$, we have $\tilde{x}_k^{(\ell)} = 1$ if and only if $(\mathbf{E}_\ell\tilde{\mathbf{D}} - \mathbf{E}_\ell\mathbf{y})_k > 0$. Hence, we have

$$M(\mathbf{y}) = \sup_{f\in\mathcal{F}_{\geq 0}} \mathbb{E}_f\left(\sum_{\ell=0}^{L-1}\left(\tilde{\mathbf{x}}^{(\ell)}(\tilde{\mathbf{D}})\odot\eta_\ell\right)^\top\left(\mathbf{E}_\ell\tilde{\mathbf{D}} - \mathbf{E}_\ell\mathbf{y}\right)\right), \tag{19}$$

where we use tilde on the binary variables to emphasize that they are stochastic variables.

Note that reformulation (19) has cross products of random variables, where the underlying uncertainty has a joint distribution f which is unknown but belongs to $\mathcal{F}_{\geq 0}$. A method resulting in a linear relaxation of a bilinear program is to introduce new variables, which *lifts* the problem to a higher dimensional space (see for instance Sherali and Alameddine 1992). This method was used in developing heuristics for the distributionally robust multi-item newsvendor problem without inventory risk pooling in Natarajan and Teo (2017a) and in Natarajan et al. (2018). When inventory is pooled across multiple locations, the cost function is more complicated since it is the optimal value of a network flow problem. However, due to the nested fulfillment cost structure, the cost function takes on a simple piecewise-linear form, allowing us to derive a relaxation via lifting, which we show to be tight in empirical experiments.

Consider the N-dimensional random vector $\tilde{\mathbf{x}} := \left(\tilde{\mathbf{x}}^{(L-1)\top}\ \tilde{\mathbf{x}}^{(L-2)\top}\ \cdots\ \tilde{\mathbf{x}}^{(0)\top}\right)^\top$. As in the previous section, we linearize the objective by lifting the problem to one that is higher dimensional by the introduction of the following new variables:

$$\mathbf{x} := \mathbb{E}_f\left(\tilde{\mathbf{x}}\right) \in \Re^N, \tag{20}$$

$$\mathbf{Q} := \mathbb{E}_f\left(\tilde{\mathbf{x}}\tilde{\mathbf{D}}^\top\right) \in \Re^{N\times n}, \tag{21}$$

$$\mathbf{R} := \mathbb{E}_f\left(\tilde{\mathbf{x}}\tilde{\mathbf{x}}^\top\right) \in \Re^{N\times N}. \tag{22}$$

Defining $N_{L-1} := 0$ and $N_\ell := \sum_{m=0}^{L-\ell-2} n_{L-1-m}$ for $\ell = 0, 1, \ldots, L-2$, we have that the objective of (18) is equivalent to

$$\sum_{\ell=0}^{L-1} \sum_{k\in[n_\ell]} \eta_{\ell,k} \cdot \left(\sum_{i\in\mathcal{I}_k^{(\ell)}} Q_{N_\ell+k,i} - \sum_{i\in\mathcal{I}_k^{(\ell)}} x_{N_\ell+k} \cdot y_i \right). \tag{23}$$

The relaxation is found by developing several necessary conditions met by any $f \in \mathcal{F}_{\geq 0}$, resulting in an semidefinite program. First, since the mean and covariance under f are $\mathbf{m}$ and $\boldsymbol{\Sigma}$, respectively, then it also follows that

$$\mathbb{E}_f\left(\begin{pmatrix} 1 \\ \tilde{\mathbf{D}} \\ \tilde{\mathbf{x}} \end{pmatrix} \begin{pmatrix} 1 \\ \tilde{\mathbf{D}} \\ \tilde{\mathbf{x}} \end{pmatrix}^\top \right) = \begin{pmatrix} 1 & \mathbf{m}^\top & \mathbf{x}^\top \\ \mathbf{m} & \boldsymbol{\Sigma} + \mathbf{m}\mathbf{m}^\top & \mathbf{Q}^\top \\ \mathbf{x} & \mathbf{Q} & \mathbf{R} \end{pmatrix} \succeq 0.$$

Second, since the support of f is nonnegative, then it must follow that $\mathbf{Q} \geq \mathbf{0}$. A third necessary condition follows from the fact that $\tilde{\mathbf{x}} \in \{0,1\}^N$, implying that

$$\begin{pmatrix} 1 & \mathbf{x}^\top \\ \mathbf{x} & \mathbf{R} \end{pmatrix} := \mathbb{E}_f\left(\begin{pmatrix} 1 \\ \tilde{\mathbf{x}} \end{pmatrix} \begin{pmatrix} 1 \\ \tilde{\mathbf{x}} \end{pmatrix}^\top \right) \in \operatorname{conv}\left\{ \begin{pmatrix} 1 \\ \mathbf{w} \end{pmatrix} \begin{pmatrix} 1 \\ \mathbf{w} \end{pmatrix}^\top : \mathbf{w} \in \{0,1\}^N \right\}. \tag{24}$$

That is, the left-hand side matrix is a convex combination of boolean matrices where each entry is a product of boolean variables. The convex hull of such matrices is often refered to as the Boolean quadric polytope (Padberg 1989). Computing the convex hull is a difficult problem since unconstrained binary quadratic programming is NP-hard in general. We instead use a simple linear relaxation of this polytope. Hence, a relaxation of $M(\mathbf{y})$ is a semidefinite program (see proof in appendix for the SDP), the dual of which yields the final form of the SDP in the proposition.

Note that while $\bar{C}(\mathbf{y})$ is an upper bound to the worst-case expected cost, through computational experiments, we find that this bound is empirically tight in the neighborhood of the minmax robust

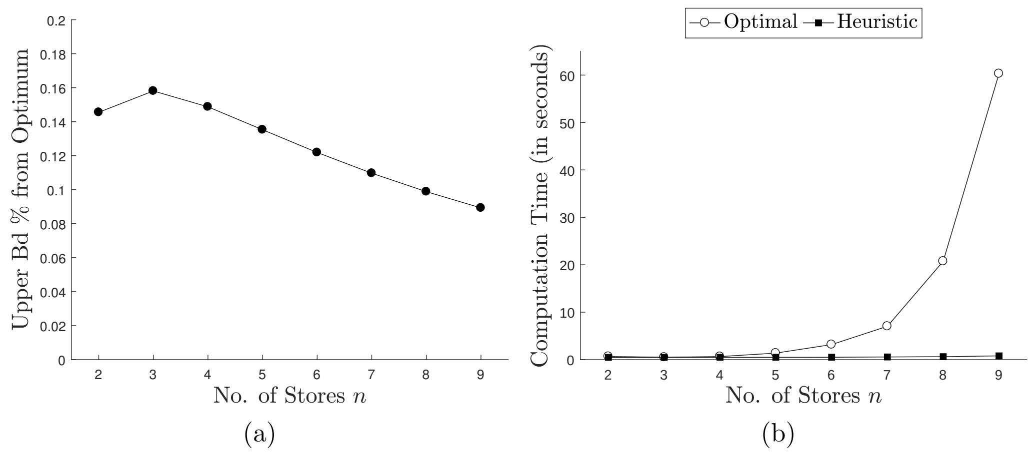

Figure 5 Experiments with the optimal minmax cost and the optimized upper bound cost. (a) The gap between the two optimized costs is small, revealing that optimizing the upper bound is a good heuristic for the minmax robust problem. (b) The computational time of the heuristic is significantly smaller than the optimal SDP.

inventory levels. Since we can solve for the exact minmax cost if there is no restriction on the support of the distribution (see Lemma 3), in our experiments, we compare the minmax cost to the upper bound when demand is allowed to be negative. Note that this means the sign of $\mathbf{q}$ and $\mathbf{Q}$ are now unrestricted, so the only change in (17) is that the third constraint is an equality constraint.

For this purpose, we consider a simple 2-level structure: in level 0, there is in-location fulfillment at a per-unit cost s_0, and in level 1, inventory from any location in the network can be used to meet unfulfilled demand in level 0 at a per-unit cost s. Figure 5 shows the results of the experiment as the number of locations is varied, with cost parameters $h = 1$, $p = 100$, $s = 1$, $s_0 = 0$, with identical marginal distributions (mean $m = 100$ and standard deviation $\sigma = 50$), and with each pair of locations having a correlation coefficient of $\rho = 0.25$. Figure 5a shows the gap between the minmax cost, which is the optimal value of (16) after setting $\mathbf{y}$ as a decision variable, and the similarly optimized upper bound (17). The plot shows that the optimized upper bound is close to the minmax cost (within 0.2%) and, in general, this gap decreases with the number of locations. Therefore, the upper bound is empirically tight in the neighborhood of the robust inventory levels. This reveals that a good heuristic for approximating the robust inventory levels is to solve for the inventory levels that optimize the SDP upper bound (17).

Figure 5b shows the computational tractability of this heuristic compared to solving for the robust inventory levels through (16). Note that, as opposed to the exponential number of SDP constraints in (16), there is only a single SDP constraint with an order of n^2 variables in (17). Hence, the computational time of the heuristic is significantly smaller compared to the optimal SDP, as demonstrated in the figure.[8]

[8] Even for $n = 9$, the optimal robust solution required around 20,000 unique SDP variables involved in the SDP constraints, whereas the heuristic required only 20 unique SDP variables, and the disparity is clearly seen in the

We thus have a tractable upper bound with a single SDP constraint of size $(N+n+1)\times(N+n+1)$, which approximates an optimal semidefinite program with 2^N SDP constraints of size $(n+1)\times(n+1)$. As a result, the above heuristic solves a general class of problems with nested fulfillment cost structures, with significantly less computational burden than the optimal SDP.

We note that the nested fulfillment structure also provides tractability for stochastic systems. Due the piecewise-linear representation of the cost function, it is straightforward to derive first order conditions for the single-period problem if the underlying distribution is known, thus preventing the need for solutions based on sample average approximations which can be difficult to solve for larger problems. In the following section, we discuss how to approximate a general distance-based fulfillment cost structure by a nested fulfillment structure.

4.3. Approximation of the Fulfillment Cost Structure

Let $\mathcal{R}=(r_{ij})_{ij}$ be the distance matrix. We make the assumption that the fulfillment cost is a function of the distance between locations ($s_{ij}=f(r_{ij})$, $\forall i,j$), and let $\mathbf{S}=(s_{ij})_{ij}$ be the corresponding fulfillment cost matrix. A straightforward decomposition of a general cost structure with n locations into an n-level nested fulfillment structure is done by hierarchical agglomerative clustering, which has been extensively studied in literature, dating back to Johnson (1967). We outline the procedure in Algorithm 1.

Algorithm 1 Hierarchical Agglomerative Clustering Algorithm

1: Let $\mathcal{S}=\{1,2,\dots,n\}$. Set $\bar{\mathcal{R}}=(\bar{r}_{ij})_{i,j\in\mathcal{S}}=\mathcal{R}$.
2: **while** $|\mathcal{S}|>1$ **do**
3: Choose the two closest nodes $i^*,j^*=\text{argmin}_{i,j\in\mathcal{S}}\,\bar{r}_{ij}$.
4: Cluster i^*,j^* into a single node: $\mathcal{S}\leftarrow\mathcal{S}+\{i^*,j^*\}-\{i^*\}-\{j^*\}$
5: Recalculate distance matrix $\bar{\mathcal{R}}=(\hat{R}_{ij})_{i,j\in\mathcal{S}}$

The algorithm proceeds by progressively clustering two closest nodes into a single node starting from the n leaf nodes corresponding to locations, until there remains only one cluster node which encompasses all the locations. In each step of the algorithm, the number of nodes is reduced by 1. Note that in order to choose the two closest nodes in each step, we need a notion of distance between clusters of nodes. Traditionally, a variety of measures have been considered to define the distance between two clusters, namely the minimum or maximum or average distance between the nodes in the two clusters, Ward's method, distance between the center of masses of the two

computational times. The heuristic could solve up to $n=100$ in under an hour, whereas the optimal solution could not be evaluated even for $n=10$ due to memory constraints.

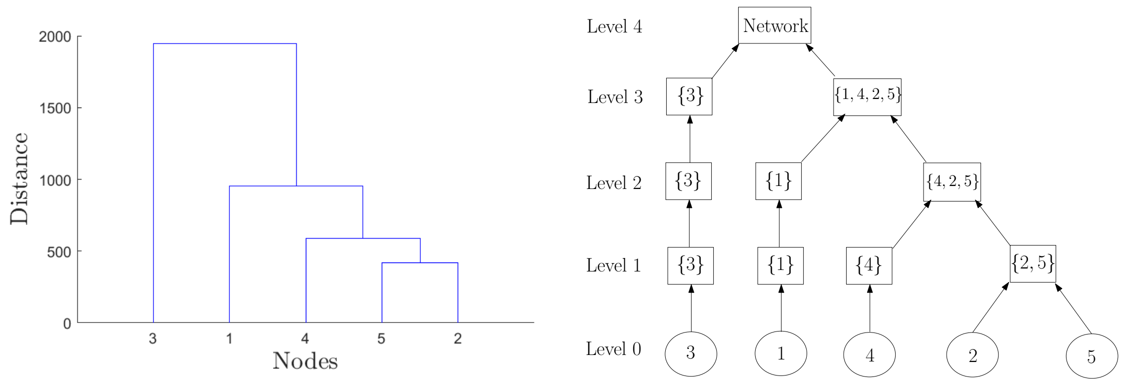

(a) (b)

Figure 6 Figure showing the approximation of a general distance-based fulfillment cost structure by a nested fulfillment structure. The figure on the left shows the dendrogram obtained by hierarchical clustering based on the distance matrix $\mathcal{R}$. The figure on the right shows the corresponding 5-level nested fulfillment structure.

clusters, etc. For the purpose of this study we consider the average distance measure, namely the Unweighted Pair Group Method with Arithmetic Mean (UPGMA), because it is simple to understand and implement. The UPGMA distance is defined as follows: if $i \in \mathcal{I}_1$ and $j \in \mathcal{I}_2$ are the indices that denote the set of nodes in each cluster, the distance between the two clusters is $\hat{r}_{\mathcal{I}_1,\mathcal{I}_2} = \frac{1}{|\mathcal{I}_1|} \cdot \frac{1}{|\mathcal{I}_2|} \cdot \sum_{i\in\mathcal{I}_1, j\in\mathcal{I}_2} r_{ij}$.

Consider the following distance matrix among 5 nodes, and the corresponding fulfillment cost matrix constructed by the equation $s_{ij} = 10 + 0.005 \cdot r_{ij}$:

$$\mathcal{R} = \begin{bmatrix} 0 & 1,220 & 1,411 & 770 & 872 \\ 1,220 & 0 & 2,404 & 624 & 420 \\ 1,411 & 2,404 & 0 & 1,785 & 2,187 \\ 770 & 624 & 1,785 & 0 & 557 \\ 872 & 420 & 2,187 & 557 & 0 \end{bmatrix}, \mathbf{S} = \begin{bmatrix} 10.0 & 16.1 & 17.1 & 13.8 & 14.4 \\ 16.1 & 10.0 & 22.0 & 13.1 & 12.1 \\ 17.1 & 22.0 & 10.0 & 18.9 & 20.9 \\ 13.8 & 13.1 & 18.9 & 10.0 & 12.8 \\ 14.4 & 12.1 & 20.9 & 12.8 & 10.0 \end{bmatrix}$$

Applying Algorithm 1, and using UPGMA as the distance metric, we obtain the dendrogram shown in Figure 6a. A dendrogram depicts the clustering in each step of the algorithm, and the distance between the entities being clustered: for instance, in the second step of the algorithm, the nodes $\{4\}$ and $\{\{5\},\{2\}\}$ are clustered at the UPGMA distance $\hat{r}_{4,\{5,2\}} = 590.6$. This means that any fulfillment between nodes 4 and 5, or between nodes 4 and 2 takes place at a cost $\hat{s}_{4,\{5,2\}} = 10 + 0.005 \cdot \hat{r}_{4,\{5,2\}}$. Thus we have an approximation of the distance and fulfillment cost matrices:

$$\hat{\mathcal{R}} = \begin{bmatrix} 0 & 954 & 1,947 & 954 & 954 \\ 954 & 0 & 1,947 & 591 & 420 \\ 1,947 & 1,947 & 0 & 1,947 & 1,947 \\ 954 & 591 & 1,947 & 0 & 591 \\ 954 & 420 & 1,947 & 591 & 0 \end{bmatrix}, \hat{\mathbf{S}} = \begin{bmatrix} 10.0 & 14.8 & 19.7 & 14.8 & 14.8 \\ 14.8 & 10.0 & 19.7 & 13.0 & 12.1 \\ 19.7 & 19.7 & 10.0 & 19.7 & 19.7 \\ 14.8 & 13.0 & 19.7 & 10.0 & 13.0 \\ 14.8 & 12.1 & 19.7 & 13.0 & 10.0 \end{bmatrix}$$

Notice that we now have an $L = n$ level nested fulfillment structure (see Figure 6b), where $l = 0$ corresponds to the actual nodes $\{3,1,4,5,2\}$, $l = 1$ corresponds to the nodes: $\{3,1,4,\{5,2\}\}$ and so on, with $l = n-1$ corresponding to the cluster of all nodes.

For the numerical analyses, we will take $L = n$, as this gives the best approximation. However, we can generate the nested fulfillment structure for any general L. We do this by cutting the dendrogram at $L-2$ places across the y-axis: the l^{th} line from the bottom gives rise to a partition of the set of locations by cutting through links that cluster these partitioned sets further. This partition gives us the clusters at level l. We provide a detailed example in Appendix C.

The nested fulfillment structure is a good approximation whenever the geographical region inherently contains this hierarchical cluster structure, where inter-cluster distances are much higher than intra-cluster distances. As we move to higher levels, we ascribe a single fulfillment cost to cross-location fulfillment between a large number of nodes, which can be a potential source of error. However the number of units being shipped in these higher levels are very small, as more demands are being pooled within each component, and as a result the error in approximation is small.[9]

We also note that the nested fulfillment cost structure resembles a tree metric, which is an approximation for a general metric on n nodes derived from an edge-weighted rooted tree with these n nodes as leaves. The distance between two nodes in the tree is defined as the sum of edge weights on the unique path from one node to another. A metric tree is commonly used to model evolutionary processes, as the rooted tree models evolution from a common ancestor. Approximation of a general metric by a tree metric has been studied in the Computer Science literature (see Bartal 1998, Fakcharoenphol et al. 2004), where probabilistic approximations of $\mathcal{O}(\log n)$ are available. In our case, the nested fulfillment cost structure has a rooted tree structure (see Figure 4), where roots of sub-trees correspond to cluster of nodes. In Section 5, we empirically show that the n-level nested fulfillment structure developed above tightly approximates the actual expected total shipping cost under a variety of distributions.

5. Numerical Analysis

We conduct multiple experiments on the proposed heuristic solutions. First, we compare the heuristic solution to stochastic optimal solutions in the 2-level structure (constant cross-location fulfillment costs) for various distributions to understand the expected value of additional information (EVAI) in a pooling context, which is defined as the loss incurred due to incomplete information about the distribution, as the heuristic solution only uses mean and covariance information. We then conduct experiments on simulated data to illustrate the superiority of the robust solution compared to stochastic solutions, as well as to study the nested fulfillment structure approximation.

[9] Govindarajan et al. (2018, Proposition 5) showed that the error from assuming a constant fulfillment cost diminishes to zero in the asymptotic case where there are infinite number of locations while holding positive safety stock. This is because, as the number of inventory nodes in a given area increases, the chance that a unit of unfulfilled demand from one location is fulfilled from a close-by location is high, and hence the fulfillment cost for every unit of demand is arbitrarily close to the in-location fulfillment cost. In our case, a similar intuition applies: the probability that fulfillment happens in higher levels is low, as there is enough supply to fulfill the pooled demand in lower levels. As a result, the error contribution to the expected total shipping cost from higher levels of fulfillment is low.

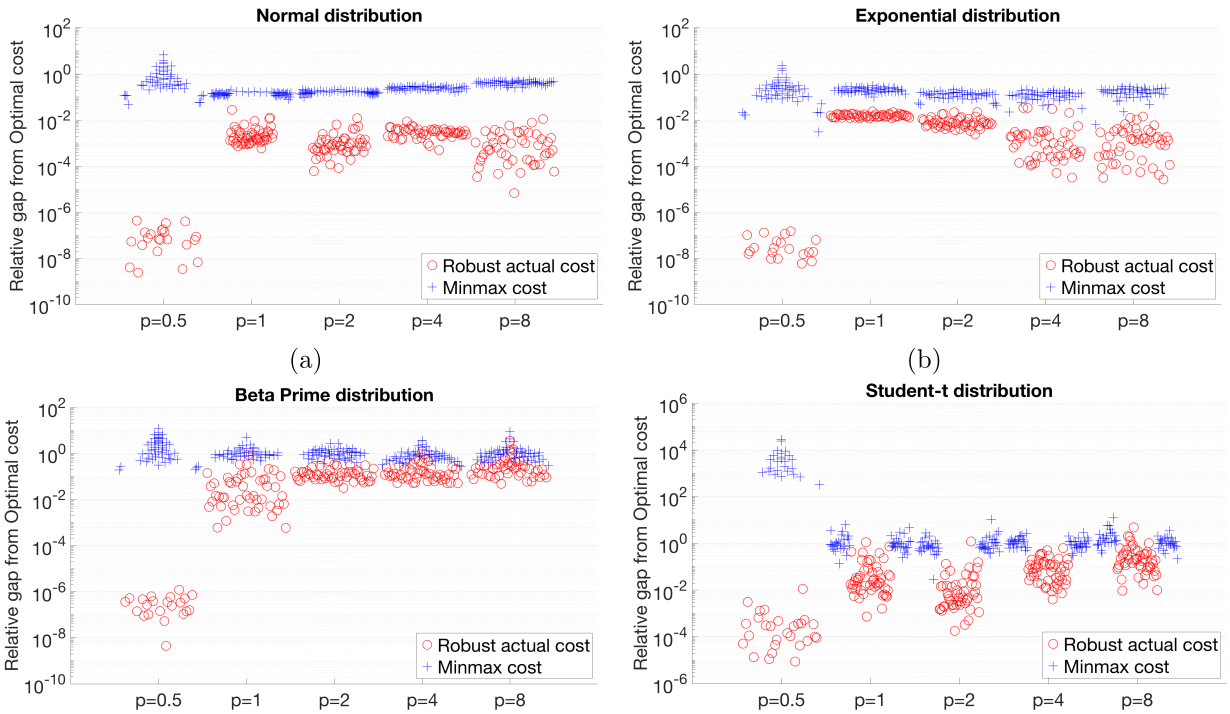

(a) (b) (c) (d)

Figure 7 **Relative gap $(C_f^{\mathrm{H}} - C_f^*)/C_f^*$ between the expected cost (under a specific distribution) of the robust inventory levels and the optimal expected cost (circle markers). The plots also show the relative gap $(\bar{C}^* - C^*)/C^*$ between the minmax cost and the optimal expected cost under the distribution (plus markers). The x-axis spread of the data around each value of p is solely for visual clarity.**

5.1. Experiments on the 2-Level Heuristic

We study the performance of the distributionally robust inventory levels under various demand distributions. The 2-level structure is useful because we can isolate the effect of pooling from the network structure, as the cross-location fulfillment costs are constant ($s_{ij} = s > s_0$ for all $i \neq j$), which simplifies the network flow problem.

We randomly generate distribution parameters for the following parametric families: Normal, exponential, beta prime, and Student-t (the details for the parameter generation are given in Appendix D). Given a specific joint demand distribution f, we estimate the optimal expected cost $C_f^* := \min_{\mathbf{y}\geq 0} \mathbb{E}_f \left[C(\mathbf{y}, \tilde{\mathbf{D}})\right]$ using sample average approximation with 10^4 samples of the demand vector. Given the mean and the covariance of the random demand, we use our heuristic to approximate the robust inventory levels with $\mathbf{y}^{\mathrm{H}}$ which is the minimizer of the upper bound to the worst-case expected cost, $\bar{C}(\mathbf{y})$, defined in (17) for $L = 2$. We denote the minimal value of this bound by $\bar{C}^*$, which is an upper bound to the minmax cost $C^* := \min_{\mathbf{y}\geq 0} \sup_{f\in\mathcal{F}} \mathbb{E}_f \left[C(\mathbf{y}, \tilde{\mathbf{D}})\right]$. We then compute the expected cost of the heuristic solution under the known true distribution f (C_f^{H}).

We first test how the heuristic performs for different levels of the underage penalty cost p. We choose $n = 5$, and set parameters $h = 1$, $s = 1$, and $s_0 = 0.5$. We conduct the experiment on

$p \in \{0.5, 1, 2, 4, 8\}$. Figure 7 presents the results. Each circle marker represents the gap between C_f^{H} and C_f^* for a specific distribution f, while the plus marker is the gap between $\bar{C}^*$ and C_f^*. Clearly, the plus markers are above the circle markers since $C_f^{\mathrm{H}} \le \sup_{f \in \mathcal{F}} \mathbb{E}_f \left[C(\mathbf{y}^{\mathrm{H}}, \tilde{\mathbf{D}}) \right] \le \bar{C}^*$. We observe that the performance of the heuristic (illustrated by the circle markers) depends on p, seen from the small optimality gap for small values of p. If the distribution is either Normal or exponential, the heuristic has an actual expected cost that is close to optimal even for high values of p, with relative gaps in the order of .1% or 1%.

For the remaining parametric distribution families, we observe that the relative gap can be as high as the order of 10%. Since for these distribution families, the circle markers are close to the plus markers, we can conclude that the expected cost under these distributions in the neighborhood of $\mathbf{y}^{\mathrm{H}}$ is close to the worst-case expected cost. We next discuss how the performance of a distributionally robust heuristic can be improved for these cases. Note that since there are multiple joint ditributions in $\mathcal{F}$, then the range of values in $\left\{ \mathbb{E}_f \left[C(\mathbf{y}, \tilde{\mathbf{D}}) \right], \ f \in \mathcal{F} \right\}$ for a given $\mathbf{y}$ could potentially be wide. This ambiguity may result in the true optimal solution to be different from the robust solution under some distributions (e.g. under beta prime or Student-t).

A way to reduce ambiguity is by further restricting the distribution set, which can be accomplished by adding more information to $\mathcal{F}$. In fact, we do this with *partitioned statistics* information, specifically, the mean and covariance of random vector $(\tilde{\mathbf{D}}^+, \tilde{\mathbf{D}}^-)$ whose i^{th} elements are $(\tilde{d}_i - m_i)^+$ and $(m_i - \tilde{d}_i)^+$, respectively. Partitioned statistics measures asymmetry of the distribution that is not represented by covariance alone (Natarajan et al. 2018). Moreover, we can utilize the techniques from this section, hence adapt Proposition 2, for a distributionally robust heuristic under this additional information (see Appendix E for the complete formulation). Note that while Natarajan et al. (2018) similarly uses partitioned statistics, they only do so for the simpler newsvendor problem without inventory pooling. Figure 8 demonstrates that additional partitioned statistics information could significantly reduce the expected cost of the distributionally robust inventory levels. It is no surprise that asymmetry information is important to estimate the impact of pooling, as this is in line with results from the pooling literature, specifically Yang and Schrage (2009) who show that right-skewed demand distributions can cause inventory levels to increase rather than decrease under pooling.

5.2. Experiments on Simulated Data

In this section, we use a fictitious network of warehouses located in mainland US to study the performance of our heuristic solutions. In particular, we illustrate the cost of misspecifying the demand distribution, and the empirical tightness of the nested fulfillment structure in approximating the expected total fulfillment cost.

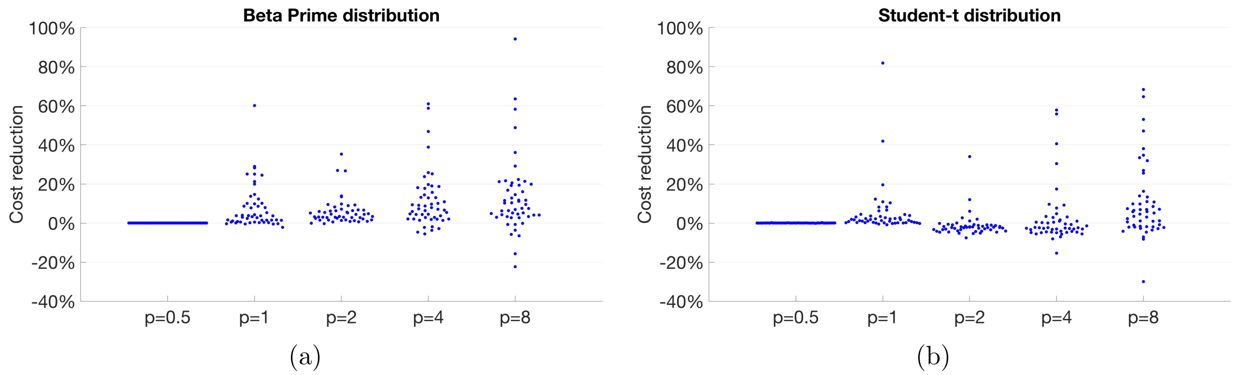

Figure 8 **Reduction in expected cost (under the true distribution) of the robust inventory levels with partitioned statistics information.**

5.2.1. Network Setup. We build a fictitious network of fulfillment centers based on publicly available data from Chen (2017), who use unofficial data of a US-based online retailer's fulfillment center network. The dataset contains information about locations of the fulfillment centers, population of the US by zipcode, and estimated shipping costs based on UPS Ground and UPS Next Day Air from the fulfillment centers to customer locations.

We consider networks of size $n = 10$, by choosing n random fulfillment centers from the 87 fulfillment centers available in the data. Since the results of our experiments depend on network characteristics, we take a sample of 10^2 networks and conduct our experiments for each sample, reporting the mean values over all the networks for the metrics considered.

For each network, the mean demands at customer locations (approximated by zipcodes) are the population in millions, and each customer location demand is assigned to the nearest fulfillment center. That is, the fulfillment centers can fulfill demands from their assigned customer locations at the in-location fulfillment cost. The coefficient of variation is taken to be equal to 1. We generate a random correlation matrix based on Numpacharoen and Atsawarungruangkit (2012), such that the maximum correlation coefficient has an absolute value less than 0.4. We take 10^3 samples of the demand vector for sample average approximations.

Similar to Jasin and Sinha (2015) and Lei et al. (2018), we take the fulfillment costs to be linear functions of the distance. Specifically we have $s_{ij} = s_0 + \lambda \cdot r_{ij}$, where r_{ij} is the distance in miles, $\lambda = 0.005$ is the distance sensitivity factor, with in-location fulfillment done at cost $s_0 = \$10$. This gives fulfillment costs in the range of $[\$10, \$23.6]$ for the entire network. The overage and underage cost parameters are taken to be: $h = \$10$, $p = \$50$. We use Algorithm 1 to generate the L-level nested fulfillment structure with $L = n$ as the base case.

5.2.2. Misspecifying the Distribution. We first study the effect of misspecifying demand distributions that was illustrated in Example 1. In particular, we compare the expected costs under the nested fulfillment cost structure achieved by the following two inventory solutions:

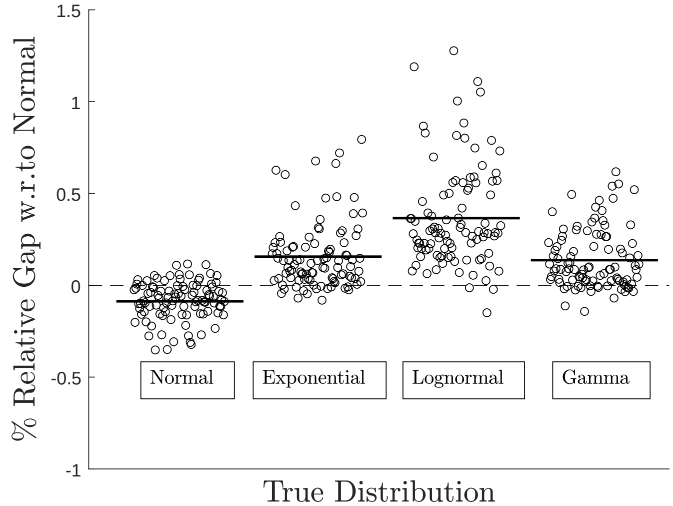

Figure 9 **Figure showing the relative gap in expected costs under the true distribution, of the robust solution $y^{*,R}$ with respect to the normal solution $y^{*,N}$. The x-axis spread of the data is solely for visual clarity.**

1. The robust inventory solution derived from the SDP: $y^{*,R} = \text{argmin}_{\mathbf{y}\geq 0} \sup_{f\in\mathcal{F}} \mathbb{E}_f \left[C(\mathbf{y}, \tilde{\mathbf{D}})\right]$,
2. The stochastic inventory solution derived by sample average approximation assuming a normal distribution $y^{*,N} = \text{argmin}_{\mathbf{y}\geq 0}\mathbb{E}_N \left[C(\mathbf{y}, \tilde{\mathbf{D}})\right]$

The expected costs $\mathbb{E}_f \left[C(\mathbf{y}, \tilde{\mathbf{D}})\right]$ are calculated under the true distributions, $f \in$ {normal, exponential, lognormal, gamma} (details are provided in Appendix D), and the results are shown in Figure 9. Each circle corresponds to a randomly chosen network of size $n = 10$.

Indeed, if the true demand distribution were normal, then $y^{*,N}$ will be the true optimal solution, in which case the relative gap in expected cost achieved by the robust solution is negative. However, this is not usually the case in reality, as the real joint distribution of demands can seldom be accurately predicted. We see that for certain networks, when the true distribution is non-normal, significant savings in expected costs can be realized by using the robust solution $y^{*,R}$ instead of the normal solution $y^{*,N}$. The savings are likely to be higher for larger networks as the normal distribution perceives higher pooling benefits which may not be the case under the true distribution.

5.2.3. Fulfillment Structure Approximation. In Figure 10, we study the approximation of the expected total fulfillment cost by a nested fulfillment structure. In particular, we calculate the relative gap in expected total fulfillment costs assuming the nested structure with respect to the actual structure. Given inventory levels, if the relative gap is found to be small, we can use the nested structure to calculate the optimal inventory levels as it provides tractability.

Indeed, we see that the nested cost structure tightly approximates the expected total fulfillment cost for a variety of distributions such as normal, exponential, lognormal and gamma (Figure

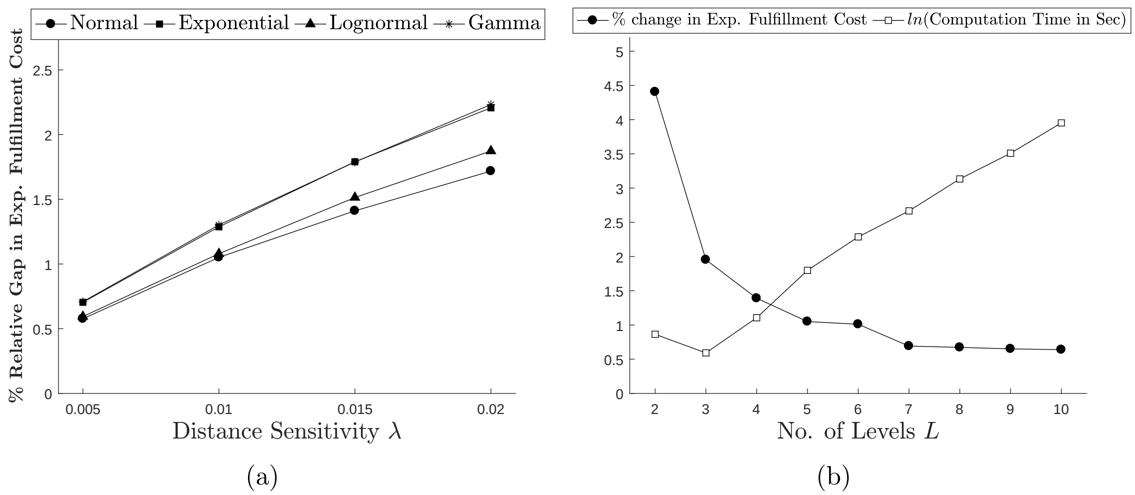

Figure 10 Figure showing the approximation of expected total fulfillment cost by the nested fulfillment structure. The figure on the left shows that the relative gap in expected total fulfillment cost under a variety of distributions as fulfillment costs become more sensitive to distance. The figure on the right shows the quality of approximation and computational time as the number of levels in the nested structure is varied.

10a). The approximation is detriorating as fulfillment costs become more sensitive to distance, however, even for high values of distance sensitivity factor λ, the relative gap in expected total fulfillment costs is less than 3%.[10] As a result, the nested fulfillment structure can serve as a good approximation for most distance-based shipping alternatives seen in practice.

So far, we considered $L = n = 10$, which gives us the best approximation for the general fulfillment costs. However, the computational requirements are high, as the size of the SDP constraint is $(N+n+1)\times(N+n+1)$ where $N = \frac{n\cdot(n+1)}{2}$. In this experiment, we test the performance of lower values of L with respect to quality of approximation of the expected total fulfillment cost (under the exponential distribution) and computational time, and the results are shown in 10b. As expected, increasing L improves approximation quality, albeit marginally for $L > 7$, however, the computational time increases exponentially. Hence it is better to choose nominal values of L, as good approximations can be achieved in relatively shorter time.

5.2.4. Dynamic Myopic Fulfillment. We study the quality of the single-period assumption made in our study by considering a dynamic setting where demands arrive at random. We model random arrivals in the following fashion: we generate the single period demand vector, and randomize the sequence in which each unit of demand arrives. The decision on which location fulfills an incoming unit of demand is taken myopically – the nearest location with available inventory is chosen to fulfill the demand, which is the fulfillment norm in practice.

[10] We note that the distance sensitivity factors for the UPS Ground shipping is 0.0005 as estimated by Jasin and Sinha (2015), which is ten times lower than the lowest value of distance sensitivity considered in this experiment.

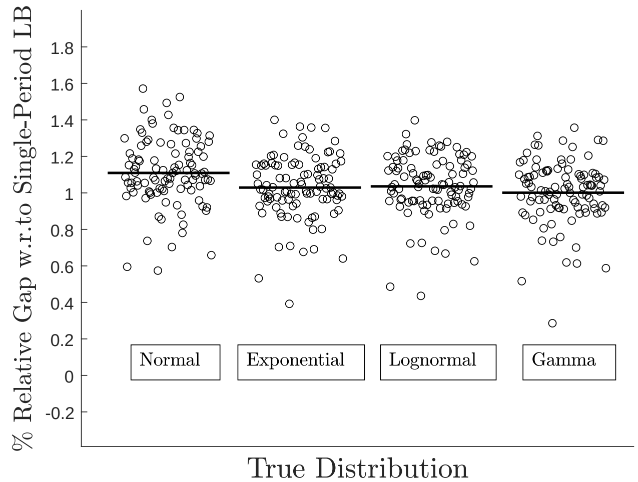

Figure 11 **Figure showing the relative gap in expected costs between the dynamic setting under myopic fulfillment and the single-period lower bound for various distributions.**

The starting inventory levels are set by the robust heuristic. Note that given any inventory levels, the single period cost is a *hindsight optimal* lower bound for the cost under the dynamic setting. We see in Figure 11 that the relative gap in expected costs between the single-period and dynamic settings is less than 2%, and hence the single-period expected cost can serve as a good approximation for the expected cost under a dynamic setting. Note that the myopic strategy need not be the optimal fulfillment strategy in a dynamic setting, and hence the actual relative gap in expected costs will be less than 2%. Similar results were observed when the nested fulfillment structure was used in place of the actual fulfillment costs.

6. Extensions

6.1. Location-specific Demand Classes

In the previous sections, we made the assumption that all demands can be fulfilled by inventory in any node, regardless of the demand location or the inventory node location. However, in some settings, there may be classes of demand that cannot be fulfilled by inventory nodes in a different location. One example is an omni-channel store network; in each location, there are two types of customers: those purchasing from the local brick-and-mortar store, and those placing an order through the online store. Demand from store customers can only be met with inventory that is located in the local store. On the other hand, demand from an online customer can be fulfilled from any store location, through what is known in the retail industry as *ship-from-store* fulfillment.

Note that there are two different types of inventory risk pooling involved here. First, within a location, store demand and local online demand are pooled since they deplete from the same store inventory. Second, online demand across locations are pooled since they are fulfilled from inventory in the store network. While ignoring the first type (for instance, by keeping a separate inventory for store customers) simplifies the problem to one explored in the previous section, it results in suboptimal inventory levels since it is likely that local demands are highly correlated.

Detailed analyses of this setting can be found in Appendix F. We show that this problem can also be analyzed in a similar fashion to Section 4, except that there are nested piecewise linear terms of the form $(x-(y-z)^+)^+$ in the objective. We can still linearize these expressions using integer variables similar to what was done in (23), however we encounter products of integer variables. We deal with this complication by introducing new integer variables to replace these product terms, and we obtain the heuristic in the form of an SDP of increased size.

6.2. Uncertainty in Moment Information

As e-commerce demands are highly volatile, there may be uncertainty in the moment information estimated from the data. Such uncertainty may be in the form of confidence intervals constructed around the moment information through empirical estimation from the data, or in the form of more complicated uncertainty sets. These can be incorporated easily into our models by simply including the constraint $(\mu,\Sigma)\in\mathcal{U}$, where $\mathcal{U}$ is a non-empty, closed and convex uncertainty set for the estimated mean and covariance matrices, and allowing μ and Σ to be variables that are constrained in the above fashion, rather than parameters (Natarajan et al. 2011).[11]

Natarajan et al. (2011) provide two examples of uncertainty set representations:

1. Linear: $\mathcal{U}=\{(\mu,\Sigma):\mu_L\le\mu\le\mu_U,\Sigma_L\le\Sigma\le\Sigma_U\}$. This can simply be incorporated as linear constraints, for which the dual can be taken easily.
2. Ellipsoidal (Delage and Ye 2010): $(\mu-\mu_0)^\intercal\Sigma_0^{-1}(\mu-\mu_0)\le\gamma_1,\quad \Sigma-2\mu\mu_0^\intercal+\mu_0\mu_0^\intercal\le\gamma_2\Sigma_0$, where μ_0,Σ_0 are the estimated mean and covariance matrices, and γ_1,γ_2 are parameters. Notice that the first constraint is non-linear, but can be expressed as the following semi-definite constraint:

$$\begin{pmatrix}\gamma_1 & (\mu-\mu_0)^\intercal\\ (\mu-\mu_0) & \Sigma_0\end{pmatrix}\succeq 0$$

We note that any uncertainty set that can be characterized by linear or semi-definite constraints can be included, as they easily yield themselves to dualizing.

7. Conclusion

Robust strategies are gaining importance in retail due to the increase in complexity arising from innovations. Particularly for e-commerce demands, incorrect forecasting may lead to disastrous results, as inventory planning is done at the network level. We provide a framework to analyze the distributionally robust newsvendor network problem where there are network flows after realization

[11] Note that this modification is to be made before taking the dual SDP of the inner robust problem. For example, in the case of identical cross-location fulfillment costs, the constraint $(\mu,\Sigma)\in\mathcal{U}$, where $\mathcal{U}$ is included in the SDP relaxation of the moment problem in maximization form in (68) before taking the dual.

of uncertainty. We solve the two-location setting in closed form, and derive insights on the role of imbalance and correlation on the expected worst-case cost.

For the multi-location case, we provide a heuristic approximation and upper bound for the case where the fulfillment costs exhibit a nested fulfillment structure, where the cost function can be written as the sum of piecewise linear terms. We show how any general fulfillment cost structure can be approximated by this nested fulfillment structure through simple agglomerative clustering algorithms, and that the approximation of the expected total fulfillment cost is empirically tight for commonly seen distance-based shipping cost structures under various distributions.

Following Natarajan et al. (2018), we show that the value of asymmetry information is significant for a system with pooling, which also echoes results from pooling literature which state that the shape of the distributions have a significant effect on pooling benefits. We also demonstrate that a distributionally robust solution can significantly outperform stochastic inventory solutions that assumes a particular demand distribution.

Multiple directions for future work exist. A multi-period formulation can be considered, where actions in the current period affect the future state. While tractable formulations can elude us, we can approximate future stages through an affine approximation, where the future actions are restricted to be affine functions of the corresponding data (Ben-Tal et al. 2004). Under such settings, robust fulfillment decisions can be analyzed which can yield helpful tools for practitioners to fulfill online demands. Our heuristic also yields the probability of stockout at the end of the period for each node in the nested fulfillment structure, which can also be used to guide dynamic fulfillment. Another natural extension is to consider how the network should look like in the first place: the solution from the inventory optimization can inform network design decisions, which is an important unexplored area in e-commerce.

Acknowledgments

This work was started while the author Dr. Amitabh Sinha was at the Ross School of Business at University of Michigan, and was completed independently of his employment at Amazon. This material is based upon work supported by the National Science Foundation under Grant Number 1561791. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

Acimovic J, Graves SC (2014) Making better fulfillment decisions on the fly in an online retail environment. *Manufacturing & Service Oper. Management* 17(1):34–51.

Acimovic J, Graves SC (2017) Mitigating spillover in online retailing via replenishment. *Manufacturing & Service Oper. Management* 19(3):419–436.

Agrawal N, Smith SA (1996) Estimating negative binomial demand for retail inventory management with unobservable lost sales. *Naval Res. Logist.* 43(6):839–861.

Bartal Y (1998) On approximating arbitrary metrices by tree metrics. *Sympos. Theory Comput. (STOC)*, 161–168 (ACM, New York).

Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A (2004) Adjustable robust solutions of uncertain linear programs. *Math. Programming* 99(2):351–376.

Berman O, Krass D, Mahdi Tajbakhsh M (2011) On the benefits of risk pooling in inventory management. *Production Oper. Management* 20(1):57–71.

Bimpikis K, Markakis MG (2015) Inventory pooling under heavy-tailed demand. *Management Sci.* 62(6):1800–1813.

Burer S (2009) On the copositive representation of binary and continuous nonconvex quadratic programs. *Math. Programming* 120(2):479–495.

Chen AIA (2017) *Large-scale optimization in online-retail inventory management.* Ph.D. thesis, Massachusetts Institute of Technology.

Corbett CJ, Rajaram K (2006) A generalization of the inventory pooling effect to nonnormal dependent demand. *Manufacturing & Service Oper. Management* 8(4):351–358.

Delage E, Ye Y (2010) Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.* 58(3):595–612.

Dong L, Rudi N (2004) Who benefits from transshipment? exogenous vs. endogenous wholesale prices. *Management Sci.* 50(5):645–657.

Eppen GD (1979) Noteeffects of centralization on expected costs in a multi-location newsboy problem. *Management Sci.* 25(5):498–501.

Fakcharoenphol J, Rao S, Talwar K (2004) A tight bound on approximating arbitrary metrics by tree metrics. *J. Comput. Sys. Sci.* 69(3):485–497.

Gallego G, Moon I (1993) The distribution free newsboy problem: review and extensions. *J. Oper. Res. Soc.* 44(8):825–834.

Govindarajan A, Sinha A, Uichanco J (2018) Joint inventory and fulfillment decisions for omnichannel retail networks, Working Paper, University of Michigan.

Hanasusanto GA, Kuhn D, Wallace SW, Zymler S (2015) Distributionally robust multi-item newsvendor problems with multimodal demand distributions. *Math. Programming* 152(1-2):1–32.

Isii K (1964) Inequalities of the types of Chebyshev and Cramér-Rao and mathematical programming. *Annals of the Institute of Statistical Math.* 16(1):277–293.

Jasin S, Sinha A (2015) An LP-based correlated rounding scheme for multi-item e-commerce order fulfillment. *Oper. Res.* 63(6):1336–1351.

Johnson SC (1967) Hierarchical clustering schemes. *Psychometrika* 32(3):241–254.

Lei Y, Jasin S, Sinha A (2018) Joint dynamic pricing and order fulfillment for e-commerce retailers. *Manufacturing & Service Oper. Management* 20(2):269–284.

Mostard J, De Koster R, Teunter R (2005) The distribution-free newsboy problem with resalable returns. *Internat. J. Production Econom.* 97(3):329–342.

Natarajan K, Sim M, Uichanco J (2018) Asymmetry and ambiguity in newsvendor models. *Management Sci.* 64(7):3146–3167.

Natarajan K, Song M, Teo CP (2009) Persistency model and its applications in choice modeling. *Management Sci.* 55(3):453–469.

Natarajan K, Teo CP (2017a) On reduced semidefinite programs for second order moment bounds with applications. *Math. Programming* 161(1-2):487–518.

Natarajan K, Teo CP, Zheng Z (2011) Mixed 0-1 linear programs under objective uncertainty: A completely positive representation. *Oper. Res.* 59(3):713–728.

Numpacharoen K, Atsawarungruangkit A (2012) Generating correlation matrices based on the boundaries of their coefficients. *PLoS One* 7(11):e48902.

Padberg M (1989) The boolean quadric polytope: some characteristics, facets and relatives. *Math. programming* 45(1-3):139–172.

Paterson C, Kiesmüller G, Teunter R, Glazebrook K (2011) Inventory models with lateral transshipments: A review. *Eur. J. Oper. Res.* 210(2):125–136.

Robinson LW (1990) Optimal and approximate policies in multiperiod, multilocation inventory models with transshipments. *Oper. Res.* 38(2):278–295.

Scarf H (1958) A min max solution of an inventory problem. *Studies in the Mathematical Theory of Inventory and Production* 201–209.

Sherali HD, Alameddine A (1992) A new reformulation-linearization technique for bilinear programming problems. *J. Global Optimization* 2(4):379–410.

Smith JE (1995) Generalized chebychev inequalities: theory and applications in decision analysis. *Oper. Res.* 43(5):807–825.

Tagaras G, Cohen MA (1992) Pooling in two-location inventory systems with non-negligible replenishment lead times. *Management Sci.* 38(8):1067–1083.

Vandenberghe L, Boyd S (1996) Semidefinite programming. *SIAM Rev.* 38(1):49–95.

Yang H, Schrage L (2009) Conditions that cause risk pooling to increase inventory. *Eur. J. Oper. Res.* 192(3):837–851.

Zaroban S (2018) U.S. e-commerce sales grow 16.0% in 2017. *InternetRetailer.* Retrieved from: https://www.digitalcommerce360.com/article/us-ecommerce-sales/.

Appendices

Appendix A: Proofs

A.1. Proof of Lemma 1

The identical two location problem is described as:

$$\begin{aligned} C^* &= \min_{y_1, y_2 \geq 0} \sup_{f \in \mathcal{F}^{m\sigma\rho}} \mathbb{E}_f \left[C(y_1, y_2, \tilde{\mathbf{D}}) \right] \\ &:= \min_{y_1, y_2 \geq 0} G(y_1, y_2) \end{aligned}$$

It is easy to see that $\mathbb{E}_f \left[C(y_1, y_2, \tilde{\mathbf{D}}) \right]$ is jointly convex in y_1, y_2, as C can be expressed as a linear program, and expectation preserves convexity. Note that C is also symmetric with respect to y_1 and y_2 when the locations are identical. Thus, we have $G(y_1, y_2) = G(y_2, y_1)$. If we show that G is jointly convex in y_1, y_2, we are done, because if (y_1^*, y_2^*) is an optimal solution to C^*, then so is (y_2^*, y_1^*), and so is $\left(\frac{y_1^*+y_2^*}{2}, \frac{y_1^*+y_2^*}{2}\right)$.

To show joint convexity of G, consider two points: $(\hat{y}_1, \hat{y}_2)$ and $(\bar{y}_1, \bar{y}_2)$. Let $\lambda \in [0,1]$. We have:

$$\begin{aligned} &G\left(\lambda\hat{y}_1 + (1-\lambda)\bar{y}_1, \lambda\hat{y}_2 + (1-\lambda)\bar{y}_2\right) \\ &= \sup_{f \in \mathcal{F}^{m\sigma\rho}} \mathbb{E}_f \left[C\left(\lambda\hat{y}_1 + (1-\lambda)\bar{y}_1, \lambda\hat{y}_2 + (1-\lambda)\bar{y}_2, \tilde{\mathbf{D}}\right) \right] \\ &\leq \sup_{f \in \mathcal{F}^{m\sigma\rho}} \left(\lambda \cdot \mathbb{E}_f \left[C(\hat{y}_1, \hat{y}_2, \tilde{\mathbf{D}}) \right] + (1-\lambda) \cdot \mathbb{E}_f \left[C(\bar{y}_1, \bar{y}_2, \tilde{\mathbf{D}}) \right] \right) \\ &= \lambda \cdot \mathbb{E}_{f^*} \left[C(\hat{y}_1, \hat{y}_2, \tilde{\mathbf{D}}) \right] + (1-\lambda) \cdot \mathbb{E}_{f^*} \left[C(\bar{y}_1, \bar{y}_2, \tilde{\mathbf{D}}) \right] \\ &\leq \lambda \cdot \sup_{f \in \mathcal{F}^{m\sigma\rho}} \left(\mathbb{E}_f \left[C(\hat{y}_1, \hat{y}_2, \tilde{\mathbf{D}}) \right] \right) + (1-\lambda) \cdot \sup_{f \in \mathcal{F}^{m\sigma\rho}} \left(\mathbb{E}_f \left[C(\bar{y}_1, \bar{y}_2, \tilde{\mathbf{D}}) \right] \right) \\ &= \lambda G(\hat{y}_1, \hat{y}_2) + (1-\lambda) G(\bar{y}_1, \bar{y}_2) \end{aligned}$$

where f^* is the solution to $\sup_{f \in \mathcal{F}^{m\sigma\rho}} \left(\lambda \cdot \mathbb{E}_f \left[C(\hat{y}_1, \hat{y}_2, \tilde{\mathbf{D}}) \right] + (1-\lambda) \cdot \mathbb{E}_f \left[C(\bar{y}_1, \bar{y}_2, \tilde{\mathbf{D}}) \right] \right)$. The first inequality follows from joint convexity of $\mathbb{E}_f \left[C(y_1, y_2, \tilde{\mathbf{D}}) \right]$.

Q.E.D.

A.2. Proof of Theorem 1

Note that for any inventory level $\mathbf{y} = (y, y)$, $\sup_{f \in \mathcal{F}^{m\sigma\rho}} \mathbb{E}_f[C(\mathbf{y}, \tilde{\mathbf{D}})] = 2(pm + hy) - M(y)$, where $M(y)$ is the optimal value of the moment problem

$$\begin{aligned} M(y) := \quad & \inf_f \mathbb{E}_f \left[(p+h-s) \min\left(\tilde{d}_1 + \tilde{d}_2, 2y\right) + \sum_{j=1,2} (s - s_0) \min\left(\tilde{d}_j, y\right) \right] \\ & \text{s.t. } \mathbb{E}_f(1) = 1, \\ & \quad \mathbb{E}_f(\tilde{d}_j) = m, \qquad j = 1, 2, \\ & \quad \mathbb{E}_f(\tilde{d}_j^2) = m^2 + \sigma^2, \qquad j = 1, 2, \\ & \quad \mathbb{E}_f(\tilde{d}_1 \tilde{d}_2) = m^2 + \rho\sigma^2, \\ & \quad f(\mathbf{D}) \geq 0, \qquad \forall \mathbf{D} \in \Re^2. \end{aligned} \tag{25}$$

The dual of the semi-infinite linear program (25) is as follows:

$$\begin{aligned} \sup_{t, u_1, u_2, r_1, r_2, v} \quad & t + m(r_1 + r_2) + (m^2 + \sigma^2)(u_1 + u_2) + (m^2 + \rho\sigma^2)v \\ \text{s.t.} \quad & t + r_1 d_1 + r_2 d_2 + u_1 d_1^2 + u_2 d_2^2 + v d_1 d_2 \\ & \leq (h + p - s) \min(d_1 + d_2, 2y) + (s - s_0) \min(d_1, y) + (s - s_0) \min(d_2, y), \qquad \forall (d_1, d_2) \in \Re^2. \end{aligned} \tag{26}$$

Facet i	$\nabla g(d_1^*, d_2^*) = \nabla f_i(d_1^*, d_2^*)$	$g(d_1^*, d_2^*) = f_i(d_1^*, d_2^*)$
1	$(d_1^*, d_2^*) = \left(\frac{h+p-s_0-r}{2u+v}, \frac{h+p-s_0-r}{2u+v}\right)$	$t = \frac{(h+p-s_0-r)^2}{2u+v}$
2	$(d_1^*, d_2^*) = \left(\frac{h+p-s_0-r}{2u+v} + \frac{(s-s_0)v}{4u^2-v^2}, \frac{h+p-s_0-r}{2u+v} - \frac{2(s-s_0)u}{4u^2-v^2}\right)$	$y(4u^2-v^2) + (s-s_0)u - (h+p-s_0-r)(2u-v) = 0$
3	$(d_1^*, d_2^*) = \left(\frac{h+p-s_0-r}{2u+v} - \frac{2(s-s_0)u}{4u^2-v^2}, \frac{h+p-s_0-r}{2u+v} + \frac{(s-s_0)v}{4u^2-v^2}\right)$	$y(4u^2-v^2) + (s-s_0)u - (h+p-s_0-r)(2u-v) = 0$
4	$(d_1^*, d_2^*) = \left(\frac{-r}{2u+v}, \frac{-r}{2u+v}\right)$	$t = \frac{r^2}{2u+v} + 2(h+p-s_0)y$
5	$(d_1^*, d_2^*) = \left(\frac{-r}{2u+v} + \frac{2(s-s_0)u}{4u^2-v^2}, \frac{-r}{2u+v} - \frac{(s-s_0)v}{4u^2-v^2}\right)$	$t = \frac{r(r-s+s_0)(2u-v)+(s-s_0)^2u}{4u^2-v^2} + (h+p-s)(2y) + (s-s_0)y$
6	$(d_1^*, d_2^*) = \left(\frac{-r}{2u+v} - \frac{(s-s_0)v}{4u^2-v^2}, \frac{-r}{2u+v} + \frac{2(s-s_0)u}{4u^2-v^2}\right)$	$t = \frac{r(r-s+s_0)(2u-v)+(s-s_0)^2u}{4u^2-v^2} + (h+p-s)(2y) + (s-s_0)y$

Table 2 Points of contact of biquadratic with each facet, and conditions on (t,u,r,v) for biquadratic and facet to touch at exactly one point.

A result by Smith (1995) is that strong duality holds for moment problems if the moment vector is an interior point of the set of feasible moments. For $\mathcal{F}^{m\sigma\rho}$, this is true for $\sigma > 0$ and $\rho \in (-1,1)$.

Note that because $\tilde{d}_1$ and $\tilde{d}_2$ are interchangeable in the primal, r_1 and r_2 must be interchangeable in the dual. The same argument applies for u_1 and u_2 as well. This implies, $r_1 = r_2 = r$, and $u_1 = u_2 = u$. Thus, we have the following dual formulation:

$$\begin{aligned}
\sup_{t,u,r,v} \quad & t + 2mr + 2(m^2+\sigma^2)u + (m^2+\rho\sigma^2)v \\
\text{s.t.} \quad & t + r(d_1+d_2) + u(d_1^2+d_2^2) + vd_1d_2 \\
& \le (h+p-s)\min(d_1+d_2, 2y) + (s-s_0)\min(d_1, y) + (s-s_0)\min(d_2,y), \quad \forall (d_1,d_2) \in \Re^2.
\end{aligned} \tag{27}$$

The right hand side of the constraint is a piecewise linear function in $\Re^2$. For notational brevity, define the quadratic function $g(d_1,d_2;t,u,r,v) = t + r(d_1+d_2) + u(d_1^2+d_2^2) + vd_1d_2$. Hence, the dual formulation can be equivalently reformulated as

$$\begin{aligned}
\sup_{t,u,r,v} \quad & t + 2mr + 2(m^2+\sigma^2)u + (m^2+\rho\sigma^2)v \\
\text{s.t.} \quad & g(d_1,d_2;t,u,r,v) \le (h+p-s_0)(d_1+d_2), && \forall d_1 \le y,\ d_2 \le y \\
& g(d_1,d_2;t,u,r,v) \le (h+p-s)(d_1+d_2) + (s-s_0)(d_1+y), && \forall d_1 \le y \le d_2,\ d_1+d_2 \le 2y \\
& g(d_1,d_2;t,u,r,v) \le (h+p-s)(d_1+d_2) + (s-s_0)(y+d_2), && \forall d_2 \le y \le d_1,\ d_1+d_2 \le 2y \\
& g(d_1,d_2;t,u,r,v) \le (h+p-s_0)(2y), && \forall d_1 \ge y,\ d_2 \ge y \\
& g(d_1,d_2;t,u,r,v) \le (h+p-s)(2y) + (s-s_0)(d_1+y), && \forall d_1 \le y \le d_2,\ d_1+d_2 \ge 2y \\
& g(d_1,d_2;t,u,r,v) \le (h+p-s)(2y) + (s-s_0)(y+d_2), && \forall d_2 \le y \le d_1,\ d_1+d_2 \ge 2y.
\end{aligned} \tag{28}$$

Note that the dual feasible set is the set of all bi-quadratic functions $g(x_1,x_2)$ that are bounded above by a piecewise linear function with six facets (one for each constraint). Let $q_i(x_1,x_2)$ denote the linear function for facet i, i.e., the right hand side of the constraint i in model (28).

Let us consider the case where $g(d_1,d_2)$ touches the piecewise linear function at exactly 6 points, one on each facet. We will later show that this case corresponds to the dual optimal solution. To find these points, for each i, we equate $\nabla g(d_1,d_2) = \nabla q_i(d_1,d_2)$ and solve for (d_1^*,d_2^*) as a function of the dual variables t,u,r,v. Then, setting $g(d_1^*,d_2^*) = q_i(d_1^*,d_2^*)$ gives us a condition on the dual variables for which the two functions touch at exactly one point. We for now ignore the ranges of d_1, d_2 in which each constraint is valid (we will later use these ranges to establish constraints on the dual variables). Table 2 gives, for each facet, the points of contact and the condition on dual variables t, u, r, v. Note that we have the following four equations that need to be satisfied for $g(d_1, d_2)$ to touch all six facets of the piecewise linear function:

$$t = \frac{(h+p-s_0-r)^2}{2u+v}, \tag{29}$$

$$y(4u^2 - v^2) + (s-s_0)u - (h+p-s_0-r)(2u-v) = 0, \tag{30}$$

$$t = \frac{r^2}{2u+v} + 2(h+p-s_0)y, \tag{31}$$

$$t = \frac{r(r-s+s_0)(2u-v) + (s-s_0)^2 u}{4u^2 - v^2} + (h+p-s)(2y) + (s-s_0)y. \tag{32}$$

We use the following transformation of variables:

$$\theta = 2u - v \tag{33}$$

$$\phi = 2u + v \tag{34}$$

We convert all the dual variables into functions of θ and ϕ. It directly follows that: $u = \frac{1}{4}(\phi+\theta)$ and $v = \frac{1}{2}(\phi - \theta)$.

From (29) and (31), we have:

$$r = \frac{h+p-s_0}{2} - y\phi. \tag{35}$$

Using (35) in (29), we have:

$$t = \frac{\left(y\phi + \frac{h+p-s_0}{2}\right)^2}{\phi}. \tag{36}$$

Using (35) and (36) in (32), we have the following:

$$\phi = \theta\left(\frac{2}{s-s_0}(h+p-s_0) - 1\right). \tag{37}$$

Note that we have not used (30) yet, but substituting (35)–(37) into (30), we find that (30) is satisfied already. That is, of the four equations (29)–(32), one of them is linearly dependent on other three.

We can now write all the dual variables t, u, v, r as a function of θ, summarized as follows:

$$r = \frac{h+p-s_0}{2} - y\theta\left(\frac{2}{s-s_0}(h+p-s_0) - 1\right), \tag{38}$$

$$t = \frac{\left(y\theta\left(\frac{2}{s-s_0}(h+p-s_0)-1\right) + \frac{h+p-s_0}{2}\right)^2}{\theta\left(\frac{2}{s-s_0}(h+p-s_0)-1\right)}, \tag{39}$$

$$u = \frac{\theta(h+p-s_0)}{2(s-s_0)}, \tag{40}$$

$$v = \frac{\theta(h+p-s)}{s-s_0}. \tag{41}$$

Thus, we know that the dual variables need to be of this form so that the biquadratic touches all six facets. We still need to check whether the points at which the biquadratic touches each facet satisfies the corresponding ranges of d_1, d_2 in (28). Substituting the values (38)–(41) of the dual variables into the touching points in

Facet i	**Contact points**	**Condition on θ**
1	$(d_1^*, d_2^*) = \left(y + \frac{\frac{h+p-s_0}{2}}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)}, y + \frac{\frac{h+p-s_0}{2}}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)} \right)$	$d_1^* \le y,\ d_2^* \le y \Leftrightarrow \theta < 0$
2	$(d_1^*, d_2^*) = \left(y + \frac{\frac{3}{2}(h+p-s_0)-s+s_0}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)}, y - \frac{\frac{h+p-s_0}{2}}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)} \right)$	$d_1^* \le y \le d_2^*,\ d_1^* + d_2^* \le 2y \Leftrightarrow \theta < 0$
3	$(d_1^*, d_2^*) = \left(y - \frac{\frac{h+p-s_0}{2}}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)}, y + \frac{\frac{3}{2}(h+p-s_0)-s+s_0}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)} \right)$	$d_2^* \le y \le d_1^*,\ d_1^* + d_2^* \le 2y \Leftrightarrow \theta < 0$
4	$(d_1^*, d_2^*) = \left(y - \frac{\frac{h+p-s_0}{2}}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)}, y - \frac{\frac{h+p-s_0}{2}}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)} \right)$	$d_1^* \ge y,\ d_2^* \ge y \Leftrightarrow \theta < 0$
5	$(d_1^*, d_2^*) = \left(y + \frac{\frac{h+p-s_0}{2}}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)}, y - \frac{\frac{3}{2}(h+p-s_0)-s+s_0}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)} \right)$	$d_1^* \le y \le d_2^*,\ d_1^* + d_2^* \ge 2y \Leftrightarrow \theta < 0$
6	$(d_1^*, d_2^*) = \left(y - \frac{\frac{3}{2}(h+p-s_0)-s+s_0}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)}, y + \frac{\frac{h+p-s_0}{2}}{\theta\left(\frac{2(h+p)-s-s_0}{s-s_0}\right)} \right)$	$d_2^* \le y \le d_1^*,\ d_1^* + d_2^* \ge 2y \Leftrightarrow \theta < 0$

Table 3 **Condition on θ so that the points of contact of biquadratic with each facet occurs in the required range.**

Table 2, and observing that $h+p > s > s_0$, we find that the dual variables are feasible (i.e., the touching points are in the required range) for any $\theta < 0$ (see Table 3).

Thus, we consider the following optimization program:

$$\sup_{\theta<0} \frac{1}{4(s-s_0)(2(h+p)-s-s_0)} \left[a + b\theta + \frac{c}{\theta}\right] \tag{42}$$

where,

$$a = 4(y+m)(h+p-s_0)(s-s_0)\left(2(h+p)-s-s_0\right), \tag{43}$$

$$b = 4\left(2(h+p)-s-s_0\right)^2 \left[(y-m)^2 + \sigma^2 \left(\frac{(p+h-s)(1+\rho)+s-s_0}{2(h+p)-s-s_0}\right)\right], \tag{44}$$

$$c = (h+p-s_0)^2(s-s_0)^2, \tag{45}$$

Note that the objective function is the objective of a dual feasible solution (38)–(41) parameterized by θ. The supremum is achieved at $\theta^* = -\sqrt{\frac{c}{b}}$, where $b > 0$ since we have that $h+p > s > s_0$, $\sigma > 0$, and $\rho \in (-1,1)$. Let $\gamma := \frac{(p+h-s)(1+\rho)+s-s_0}{2(h+p)-s-s_0} \in (0,1]$. The optimal θ^* is given by:

$$\theta^* = -\frac{(h+p-s_0)(s-s_0)}{2\left(2(h+p)-s-s_0\right)\sqrt{(y-m)^2+\gamma\sigma^2}}. \tag{46}$$

The corresponding dual value is $\frac{1}{4(s-s_0)(2(h+p)-s-s_0)}(a-2\sqrt{bc})$, where a,b,c are according to the equations (43)–(45), which simplifies to:

$$(h+p-s_0)\left(y+m-\sqrt{(y-m)^2+\gamma\sigma^2}\right) \le M(y), \tag{47}$$

where the inequality follows from weak duality. Hence, an upper bound for the expected cost of any distribution $f \in \mathcal{F}^{m\sigma\rho}$ is $2(pm+hy) - M(y)$, which is bounded above by $\bar{C}(y) := 2s_0 m - (p-h-s_0)(y-m) + (p+h-s_0)\sqrt{(y-m)^2+\gamma\sigma^2}$.

All that is left to prove the theorem is to show that, if $\gamma(1+\nu^2)\geq 2$ where $\nu := \frac{3(h+p-s_0)-2(s-s_0)}{h+p-s_0}\in(1,3)$, then there exists a six-point distribution $f_y^*\in\mathcal{F}^{m\sigma\rho}$ for which the expected cost is equal to $\bar{C}(y)$. This is equivalent to the task of finding a $f_y^*\in\mathcal{F}^{m\sigma\rho}$ such that

$$\mathbb{E}_{f_y^*}\left[(p+h-s)\min(\tilde{d}_1+\tilde{d}_2,2y)+\sum_{j=1,2}(s-s_0)\min(\tilde{d}_j,y)\right]=(h+p-s_0)\left(y+m-\sqrt{(y-m)^2+\gamma\sigma^2}\right). \tag{48}$$

To construct the distribution, we use the contact points of the biquadratic to each facet as the support points. Define $z_y := (y-m)/\sigma$ and $\Phi(z_y) := \sqrt{z_y^2+\gamma}$, where we note that $\Phi(z_y) > z_y$. If we use the optimal θ^*, defined in (46), to find the associated contact points in Table 3, where we use the fact that $\theta^*\left(\frac{2(h+p)-s-s_0}{s-s_0}\right) = -\frac{h+p-s_0}{2\Phi(z_y)\sigma}$, we get the following six support points of f_y^*:

$$\begin{aligned}\mathbf{D}^{(1)} &= \begin{bmatrix} m+(z_y-\Phi(z_y))\sigma \\ m+(z_y-\Phi(z_y))\sigma\end{bmatrix}, \mathbf{D}^{(2)} = \begin{bmatrix} m+(z_y-\nu\Phi(z_y))\sigma \\ m+(z_y+\Phi(z_y))\sigma\end{bmatrix}, \mathbf{D}^{(3)} = \begin{bmatrix} m+(z_y+\Phi(z_y))\sigma \\ m+(z_y-\nu\Phi(z_y))\sigma\end{bmatrix} \\ \mathbf{D}^{(4)} &= \begin{bmatrix} m+(z_y+\Phi(z_y))\sigma \\ m+(z_y+\Phi(z_y))\sigma\end{bmatrix}, \mathbf{D}^{(5)} = \begin{bmatrix} m+(z_y-\Phi(z_y))\sigma \\ m+(z_y+\nu\Phi(z_y)\end{bmatrix}, \mathbf{D}^{(6)} = \begin{bmatrix} m+(z_y+\nu\Phi(z_y)\sigma \\ m+(z_y-\Phi(z_y))\sigma\end{bmatrix}\end{aligned} \tag{49}$$

We next construct probabilities for the distribution f_y^* to ensure that it is a feasible distribution in $\mathcal{F}^{m\sigma\rho}$. In particular, we find the probabilities $\pi_1,\pi_2,\ldots,\pi_6$ such that the following relationships are true:

$$\sum_{i=1}^{6}\pi_i = 1 \tag{50}$$

$$\sum_{i=1}^{6}\pi_i\mathbf{D}^{(i)} = \begin{bmatrix} m \\ m\end{bmatrix} \tag{51}$$

$$\sum_{i=1}^{6}\pi_i\mathbf{D}^{(i)}\odot\mathbf{D}^{(i)} = \begin{bmatrix} m^2+\sigma^2 \\ m^2+\sigma^2\end{bmatrix} \tag{52}$$

$$\sum_{i=1}^{6}\pi_i d_1^{(i)}d_2^{(i)} = m^2+\rho\sigma^2, \tag{53}$$

where $\mathbf{a}\odot\mathbf{b}=(a_ib_i)$ denotes element-wise multiplication of vectors $\mathbf{a},\mathbf{b}$.

From the equalities (50)–(53), we have the following system of linear equations: (where for notational brevity, we drop the subscript on z_y and drop the dependence of Φ on z_y)

$$\begin{aligned}
&\pi_1 + \pi_2 + \pi_3 + \pi_4 + \pi_5 + \pi_6 = 1\\
&(z-\Phi)\pi_1 + (z-\nu\Phi)\pi_2 + (z+\Phi)\pi_3 + (z+\Phi)\pi_4 \ (z-\Phi)\pi_5 + (z+\nu\Phi)\pi_6 = 0\\
&(z-\Phi)\pi_1 + (z+\Phi)\pi_2 + (z-\nu\Phi)\pi_3 + (z+\Phi)\pi_4 + (z+\nu\Phi)\pi_5 + (z-\Phi)\pi_6 = 0\\
&(z-\Phi)^2\pi_1 + (z-\nu\Phi)^2\pi_2 + (z+\Phi)^2\pi_3 + (z+\Phi)^2\pi_4 \ (z-\Phi)^2\pi_5 + (z+\nu\Phi)^2\pi_6 = 1\\
&(z-\Phi)^2\pi_1 + (z+\Phi)^2\pi_2 + (z-\nu\Phi)^2\pi_3 + (z+\Phi)^2\pi_4 + (z+\nu\Phi)^2\pi_5 + (z-\Phi)^2\pi_6 = 1\\
&(z-\Phi)^2\pi_1 + (z-\nu\Phi)(z+\Phi)\pi_2 + (z-\nu\Phi)(z+\Phi)\pi_3 + (z+\Phi)^2\pi_4 + (z+\nu\Phi)(z-\Phi)\pi_5 + (z+\nu\Phi)(z-\Phi)\pi_6 = \rho
\end{aligned}$$

By simple row operations, we can show that the last equation is linearly dependent on the others. Additionally, it is easy to see that if we interchange π_2 and π_3 as well as π_5 and π_6, the equations remain unaltered, thus $\pi_2=\pi_3$ and $\pi_5=\pi_6$. Thus, the new system of equations are:

$$\begin{aligned}
&\pi_1 + 2\pi_2 + \pi_4 + 2\pi_5 = 1\\
&(z-\Phi)\pi_1 + (2z+(1-\nu)\Phi)\pi_2 + (z+\Phi)\pi_4 + (2z-(1-\nu)\Phi)\pi_5 = 0\\
&(z-\Phi)^2\pi_1 + ((z-\nu\Phi)^2+(z+\Phi)^2)\pi_2 + (z+\Phi)^2\pi_4 \ ((z-\Phi)^2+(z+\nu\Phi)^2)\pi_5 = 1
\end{aligned}$$

Facet i	**Support Point**	**Probability**
1	$(d_1^*, d_2^*) = (m + (z_y - \Phi(z_y))\sigma,\ m + (z_y - \Phi(z_y))\sigma)$	$\frac{1}{2} + \frac{2(h+p-s)\pi}{h+p-s_0} + \frac{z_y}{2\sqrt{z_y^2+\gamma}} - \frac{(1-\gamma)}{2(\nu-1)(z_y^2+\gamma)}$
2	$(d_1^*, d_2^*) = (m + (z_y - \nu\Phi(z_y))\sigma,\ m + (z_y + \Phi(z_y))\sigma)$	$-\pi + \frac{(1-\gamma)}{(\nu^2-1)(z_y^2+\gamma)}$
3	$(d_1^*, d_2^*) = (m + (z_y + \Phi(z_y))\sigma,\ m + (z_y - \nu\Phi(z_y))\sigma)$	$-\pi + \frac{(1-\gamma)}{(\nu^2-1)(z_y^2+\gamma)}$
4	$(d_1^*, d_2^*) = (m + (z_y + \Phi(z_y))\sigma,\ m + (z_y + \Phi(z_y))\sigma)$	$\frac{1}{2} - \frac{2(h+p-s)\pi}{h+p-s_0} - \frac{z_y}{2\sqrt{z_y^2+\gamma}} - \frac{(1-\gamma)(3-\nu)}{2(\nu^2-1)(z_y^2+\gamma)}$
5	$(d_1^*, d_2^*) = (m + (z_y - \Phi(z_y))\sigma,\ m + (z_y + \nu\Phi(z_y))$	π
6	$(d_1^*, d_2^*) = (m + (z_y + \nu\Phi(z_y))\sigma,\ m + (z_y - \Phi(z_y))\sigma)$	π

Table 4 **The support points and the corresponding probabilities in a worst-case probability distribution $f^*_{y,\pi}$, where $\max(0,\alpha_1) \le \pi \le \min(\beta_1, \beta_2)$.**

Since we have three equations and four unknowns, we use parameter $\pi_5 = \pi$. Then, the solution to the set of equations are as follows:

$$\begin{aligned}
\pi_1 &= \frac{1}{2} + \frac{2(h+p-s)\pi}{h+p-s_0} + \frac{z}{2\sqrt{z^2+\gamma}} - \frac{(1-\gamma)}{2(\nu-1)(z^2+\gamma)} \\
\pi_2 = \pi_3 &= -\pi + \frac{(1-\gamma)}{(\nu^2-1)(z^2+\gamma)} \\
\pi_4 &= \frac{1}{2} - \frac{2(h+p-s)\pi}{h+p-s_0} - \frac{z}{2\sqrt{z^2+\gamma}} - \frac{(1-\gamma)(3-\nu)}{2(\nu^2-1)(z^2+\gamma)} \\
\pi_5 = \pi_6 &= \pi
\end{aligned} \tag{54}$$

We need to ensure that the probabilities lie in $[0,1]$ (they already sum up to one because of (50)), which can be accomplished by putting restrictions on the value of π. Defining:

$$\alpha_1(z) := \left(\frac{h+p-s_0}{2(h+p-s)}\right)\left(-\frac{1}{2} - \frac{z}{2\sqrt{z^2+\gamma}} + \frac{(1-\gamma)(\nu+1)}{2(\nu^2-1)(z^2+\gamma)}\right), \tag{55}$$

$$\beta_1(z) := \frac{1-\gamma}{(\nu^2-1)(z^2+\gamma)}, \tag{56}$$

$$\beta_2(z) := \left(\frac{h+p-s_0}{2(h+p-s)}\right)\left(\frac{1}{2} - \frac{z}{2\sqrt{z^2+\gamma}} + \frac{(1-\gamma)(\nu-3)}{2(\nu^2-1)(z^2+\gamma)}\right), \tag{57}$$

we have that the probabilities are nonnegative for $\max(0,\alpha_1) \le \pi \le \min(\beta_1,\beta_2)$. Note that $\beta_1 \ge 0$. If $\beta_2 < 0$, $\beta_1 < \alpha_1$, or $\beta_2 < \alpha_2$, then the set of feasible values for π is empty. However, according to the following lemma, if $\gamma(\nu^2+1) \ge 2$ then this set is non-empty for all values of z.

Lemma 4. *If $\gamma(\nu^2+1) \ge 2$, then $\beta_2(z) \ge 0$, $\alpha_1(z) \le \beta_1(z)$, and $\alpha_1(z) \le \beta_2(z)$ for all $z \in \Re$.*

Since $\nu - 1 = \frac{2(h+p-s)}{h+p-s_0}$, we can rewrite the following:

$$\alpha_1 := \left(\frac{1}{2(\nu-1)}\right)\left(-1 - \frac{z}{\Phi(z)} + \frac{(1-\gamma)(\nu+1)}{(\nu^2-1)\Phi^2(z)}\right), \tag{58}$$

$$\beta_1 := \frac{1-\gamma}{(\nu^2-1)\Phi^2(z)}, \tag{59}$$

$$\beta_2 := \left(\frac{1}{2(\nu-1)}\right)\left(1 - \frac{z}{\Phi(z)} + \frac{(1-\gamma)(\nu-3)}{(\nu^2-1)\Phi^2(z)}\right), \tag{60}$$

Note that $\beta_2(z) \ge 0$ if and only if: $\Phi(z)(\Phi(z)-z) \ge \frac{(1-\gamma)(3-\nu)}{(\nu^2-1)}$. Let $w(z) = \Phi(z)(1-z)$. Then $w'(z) = 2z - \frac{z^2}{\sqrt{z^2+\gamma}} - \sqrt{z^2+\gamma}$, and $w''(z) = 2 - \frac{3z}{\sqrt{z^2+\gamma}} + \frac{z^3}{(z^2+\gamma)^{\frac{3}{2}}}$. Note that $w''(z)$ can be shown to be non-negative (we can prove: $-2 \le -\frac{3z}{\sqrt{z^2+\gamma}} + \frac{z^3}{(z^2+\gamma)^{\frac{3}{2}}} \le 2$), implying that $w(z)$ is a convex function minimized at $z = 0$

(from equating $w'(z)=0$). Thus, whenever $\gamma \geq \frac{(1-\gamma)(3-\nu)}{(\nu^2-1)}$, we have $\beta_2(z)\geq 0$ for all z. The sufficient condition translates to: $\gamma \geq \frac{3-\nu}{\nu^2-\nu+2}$.

$\beta_2(z)\geq \alpha_1(z)$ if and only if: $\frac{2(1-\gamma)}{(\nu^2-1)\Phi^2(z)} \leq 1$, which simplifies to: $z^2 \geq \frac{2-\gamma(\nu^2+1)}{\nu^2-1}$. Thus, a sufficient condition is given by: $0 \geq \frac{2-\gamma(\nu^2+1)}{\nu^2-1}$, which translates to: $\gamma(\nu^2+1)\geq 2$. Note that the condition $\gamma \geq \frac{(1-\gamma)(3-\nu)}{(\nu^2-1)}$ is implied by $\gamma(\nu^2+1)\geq 2$.

$\beta_1(z)\geq \alpha_1(z)$ if and only if: $\Phi(z)(\Phi(z)+z)\geq \frac{(1-\gamma)(3-\nu)}{(\nu^2-1)}$. The left hand side can be shown to be a convex function minimized at $z=0$ from the same argument in the case $\beta_2(z)\geq 0$. Thus, the sufficient condition is the same as the case $\beta_2(z)\geq 0$. ■

Let us define $f^*_{y,\pi}$ as the six-point distribution that is summarized in Table 4 for some valid π. Note that the probabilities of $f^*_{y,\pi}$ only sure that the distribution has the appropriate moments to belong in $\mathcal{F}^{m\sigma\rho}$. We also need to ensure that the strong duality condition (48) is true. The left-hand side of (48) evaluates to

$$\begin{aligned}
P(y) = &\left(\frac{1}{2}+\frac{2(h+p-s)\pi}{h+p-s_0}+\frac{z}{2\sqrt{z^2+\gamma}}-\frac{(1-\gamma)}{2(\nu-1)(z^2+\gamma)}\right)\Big(2\left(h+p-s_0\right)\left(m+\left(z-\Phi(z)\right)\sigma\right)\Big)\\
&+2\left(-\pi+\frac{(1-\gamma)}{(\nu^2-1)(z^2+\gamma)}\right)\Big((h+p-s_0)(2m+\sigma(2z+(1-\nu)\Phi(z)))-(s-s_0)\sigma\Phi(z)\Big)\\
&+\left(\frac{1}{2}-\frac{2(h+p-s)\pi}{h+p-s_0}-\frac{z}{2\sqrt{z^2+\gamma}}+\frac{(1-\gamma)(\nu-3)}{2(\nu^2-1)(z^2+\gamma)}\right)\Big(2(h+p-s_0)y\Big)\\
&+2\left(\pi\right)\Big((h+p-s_0)(2y)-(s-s_0)\sigma\Phi(z)\Big)
\end{aligned} \tag{61}$$

The coefficient of π in (61) is given by:

$$\begin{aligned}
&=\frac{2(h+p-s)}{h+p-s_0}\left(2(h+p-s_0)\left[m+\left(z-\Phi(z)\right)\sigma-y\right]\right)+2\left((h+p-s_0)\left[2y-2m-\sigma(2z+(1-\nu)\Phi(z))\right]\right)\\
&=4(h+p-s)(-\sigma\phi)+2(h+p-s_0)(-\sigma(1-\nu)\Phi(z)) && (y=m+z\sigma)\\
&=4(h+p-s)(-\sigma\phi)+2(h+p-s_0)(-\sigma(1-\nu)\Phi(z)) && (1-\nu=-2(h+p-s)/(h+p))\\
&=0
\end{aligned}$$

Hence, for any π, the left-hand side of (48) with $f^*_y=f^*_{y,\pi}$ simplifies to

$$\begin{aligned}
P(y) &= (h+p-s_0)(2y-(z+\Phi(z))\sigma)=(h+p-s_0)(y+m-\Phi(z)\sigma)\\
&=(h+p-s_0)(y+m-\sqrt{(y-m)^2+\gamma\sigma^2})
\end{aligned}$$

which is equal to the right-hand side of (48). Therefore, the expected cost is equal to the upper bound $\bar{C}(y)$ for any of the distributions in distribution set $\{f^*_{y,\pi} \mid \max(0,\alpha_1)\leq \pi \leq \min(\beta_1,\beta_2)\}$. ■

A.3. Proof of Proposition 1

Let $\mathbf{y}^*=(y_1^*,y_2^*)$ be the optimal solution of the distributionally robust problem (5). Since the locations are identical, we have that $y_1^*=y_2^*=y^*$ for some y^*. Hence, we need only consider the subset of inventory levels $\mathbf{y}=(y,y)$, for which we derive an analytic expression of the worst-case cost as $\bar{C}(y)$ in (7). Thus, the distributionally robust problem (5) is equivalent to $\min_y \bar{C}(y)$. The first two derivatives of $\bar{C}(y)$ are

$$\begin{aligned}
\bar{C}'(y) &= -(p-h-s_0)+\frac{(p+h-s_0)(y-m)}{\sqrt{(y-m)^2+\gamma\sigma^2}}\\
\bar{C}''(y) &= \frac{(p+h-s_0)\gamma\sigma^2}{((y-m)^2+\gamma\sigma^2)^{\frac{3}{2}}}
\end{aligned} \tag{62}$$

Since $\gamma>0$, $\bar{C}(y)$ is convex in y, and the optimal solution is given by the first-order condition $\bar{C}'(y^*)=0$, which gives y^* as the right-hand side of (10). ■

A.4. Proof of Lemma 3

Note that (13) is equivalent to the following semi infinite linear program

$$\begin{aligned} M(\mathbf{y}) = \max_{f} \quad & \mathbb{E}_f\left[\sum_{\ell=0}^{L-1} \eta_\ell^\top (\mathbf{E}_\ell \mathbf{D} - \mathbf{E}_\ell \mathbf{y})^+\right] \\ \text{s.t.} \quad & \mathbb{E}_f(1) = 1, \\ & \mathbb{E}_f\left(\tilde{\mathbf{D}}\right) = \mathbf{m}, \\ & \mathbb{E}_f\left(\tilde{\mathbf{D}}\tilde{\mathbf{D}}^\top\right) = \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top, \\ & f(\mathbf{D}) \geq 0, \quad \forall \mathbf{D} \in \Re^n. \end{aligned} \tag{63}$$

Since $\Sigma \succ 0$, then the moments $(\mathbf{m}, \mathbf{\Sigma})$ are strictly in the interior of the feasible moment cone. Hence, strong duality of moment problems holds (Smith 1995). The dual of the moment problem is

$$\begin{aligned} M(\mathbf{y}) = \min_{t,\mathbf{r},\mathbf{Y}} \quad & t + \mathbf{r}^\top \mathbf{m} + \langle \mathbf{Y}, \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top \rangle \\ \text{s.t.} \quad & t + \mathbf{r}^\top \mathbf{x} + \mathbf{x}^\top \mathbf{Y} \mathbf{x} \geq \sum_{\ell=0}^{L-1} \eta_\ell^\top (\mathbf{E}_\ell \mathbf{x} - \mathbf{E}_\ell \mathbf{y})^+, \quad \forall \mathbf{x} \in \Re^n \end{aligned} \tag{64}$$

We can reformulate the dual as the following semi infinite linear program:

$$\begin{aligned} M(\mathbf{y}) = \min_{t,\mathbf{r},\mathbf{Y}} \quad & t + \mathbf{r}^\top \mathbf{m} + \langle \mathbf{Y}, \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top \rangle \\ \text{s.t.} \quad & t + \mathbf{r}^\top \mathbf{x} + \mathbf{x}^\top \mathbf{Y} \mathbf{x} \geq \sum_{\ell=0}^{L-1} (\eta_\ell \odot \mathbf{e}_{A_\ell})^\top (\mathbf{E}_\ell \mathbf{x} - \mathbf{E}_\ell \mathbf{y}), \quad \forall \mathbf{x} \in \Re^n, \\ & \qquad \forall (A_0, A_1, \cdots, A_{L-1}) \in 2^{[n_0]} \times 2^{[n_1]} \times 2^{[n_{L-1}]} \end{aligned} \tag{65}$$

where $\odot$ is the element-wise product operator, and $\mathbf{e}_{A_\ell}$ is an n_ℓ-dimensional binary vector whose k^{th} element is 1 if and only if $k \in A_\ell$. For simplicity, we can write the right-hand side as $\mathbf{a}_k^\top \mathbf{x} + \mathbf{b}_k^\top \mathbf{y}$ for $k \in [2^N]$, where $N = \sum_{\ell=0}^{L-1} n_\ell$. The constraint now becomes: $\mathbf{x}^\top \mathbf{Y} \mathbf{x} + (\mathbf{r} - \mathbf{a_k})^\top \mathbf{x} + t - \mathbf{b}_k^\top \mathbf{y} \geq 0, \forall \mathbf{x}$. This is true if and only if

$$\begin{bmatrix} \mathbf{Y} & \frac{1}{2}(\mathbf{r} - \mathbf{a_k}) \\ \frac{1}{2}(\mathbf{r} - \mathbf{a_k})^\top & t - \mathbf{b}_k^\top \mathbf{y} \end{bmatrix} \succeq 0, \quad \forall k. \quad \blacksquare$$

A.5. Proof of Proposition 2

With the introduction of the new variables (20)–(22), the objective of (18) is equivalent to

$$\sum_{\ell=0}^{L-1} \sum_{k \in [n_\ell]} \eta_{\ell,k} \cdot \left(\sum_{i \in \mathcal{I}_k^{(\ell)}} Q_{N_\ell + k, i} - \sum_{i \in \mathcal{I}_k^{(\ell)}} x_{N_\ell + k} \cdot y_i \right). \tag{66}$$

To develop the relaxation of (18), we discuss several necessary conditions for any $f \in \mathcal{F}_{\geq 0}$. First, we have that

$$\mathbb{E}_f\left(\begin{pmatrix} 1 \\ \tilde{\mathbf{D}} \\ \tilde{\mathbf{x}} \end{pmatrix} \begin{pmatrix} 1 \\ \tilde{\mathbf{D}} \\ \tilde{\mathbf{x}} \end{pmatrix}^\top \right) = \begin{pmatrix} 1 & \mathbf{m}^\top & \mathbf{x}^\top \\ \mathbf{m} & \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top & \mathbf{Q}^\top \\ \mathbf{x} & \mathbf{Q} & \mathbf{R} \end{pmatrix} \succeq 0.$$

The non-negative support of f also implies that $\mathbf{Q} \geq \mathbf{0}$. Since $\tilde{\mathbf{x}} \in \{0,1\}^N$, we also have that:

$$\begin{pmatrix} 1 & \mathbf{x}^\top \\ \mathbf{x} & \mathbf{R} \end{pmatrix} := \mathbb{E}_f\left(\begin{pmatrix} 1 \\ \tilde{\mathbf{x}} \end{pmatrix} \begin{pmatrix} 1 \\ \tilde{\mathbf{x}} \end{pmatrix}^\top \right) \in \operatorname{conv}\left\{ \begin{pmatrix} 1 \\ \mathbf{w} \end{pmatrix} \begin{pmatrix} 1 \\ \mathbf{w} \end{pmatrix}^\top : \mathbf{w} \in \{0,1\}^N \right\}. \tag{67}$$

We further relax the last condition. Note that if the matrix on the left-hand side belongs to the Boolean quadric polytope, then it must follow that $R_{ii} = x_i$ for all $i \in [N]$, that $R_{ij} \le x_i$ for all $i, j \in [N]$, that $R_{ij} \ge x_i + x_j - 1$ for all $i, j \in [N]$, and that $\mathbf{R} \ge \mathbf{0}$.

Therefore, we have that $M(\mathbf{y})$ is bounded above by the following SDP relaxation

$$\begin{aligned}
\bar{M}(\mathbf{y}) := \max_{z,\mathbf{x},\mathbf{q},\mathbf{Q},\mathbf{r},\mathbf{R}} \quad & \sum_{\ell=0}^{L-1} \sum_{k \in [n_\ell]} \eta_{\ell,k} \cdot \left(\sum_{i \in \mathcal{I}_k^{(\ell)}} Q_{N_\ell + k, i} - \sum_{i \in \mathcal{I}_k^{(\ell)}} x_{N_\ell + k} \cdot y_i \right) \\
\text{s.t.} \quad & \begin{pmatrix} 1 & \mathbf{m}^\top & \mathbf{x}^\top \\ \mathbf{m} & \mathbf{\Sigma} + \mathbf{m}\mathbf{m}^\top & \mathbf{Q}^\top \\ \mathbf{x} & \mathbf{Q} & \mathbf{R} \end{pmatrix} \succeq 0, \\
& R_{ii} = x_i \qquad\qquad i \in [N], \\
& R_{ij} \le x_i \qquad\qquad i \in [N], j \in [N], \\
& R_{ij} \ge x_i + x_j - 1 \qquad i \in [N], j \in [N], \\
& \mathbf{Q}, \mathbf{R} \ge 0
\end{aligned} \tag{68}$$

Note that the dual of the SDP gives the final form in the proposition. The dual variables $t_0, \mathbf{t}, \mathbf{u}, \mathbf{Y}, \mathbf{V}, \mathbf{U}$ are introduced for the semidefinite constraint. The dual variables $\mathbf{W}$ are introduced for the constraints $R_{ij} \le x_i$. The dual variables $\mathbf{B}$ are introduced for the constraints $R_{ij} \ge x_i + x_j - 1$. ∎

Appendix B: Optimal Inventory Solutions for Two-Locations Systems

We have four cases for which the distributionally robust solution needs to be calculated: pooling/no pooling (P/NP), and known/unknown correlation (C/NC). Note that we restrict the search to identical solutions of the form (y, y).

1. **No pooling, ρ unknown:** This is the same setting as Scarf (1958), and the optimal inventory and worst-case cost are given by:

$$y^{NP,NC} = m + \frac{p - h - s_0}{2\sqrt{h(p - s_0)}} \cdot \sigma \tag{69}$$

$$C^{NP,NC} = 2m(s_0 - h) + 2hy + (p + h - s_0)(m - y + \sqrt{\sigma^2 + (m - y)^2}) \tag{70}$$

2. **No pooling, ρ known:** This is the same setting as Natarajan and Teo (2017a), and the solutions are given through an SDP.

$$\begin{aligned}
C^{NP,C} := \min_{t_0, r, u, v, y} \quad & 2(s_0 - h)m + 2hy + (p + h - s_0)(t_0 + 2rm + 2u(m^2 + \sigma^2) + v(m^2 + \rho\sigma^2)) \\
\text{s.t.} \quad & \begin{pmatrix} t_0 + 2y & \frac{1}{2}(r-1) & \frac{1}{2}(r-1) \\ \frac{1}{2}(r-1) & u & \frac{1}{2}v \\ \frac{1}{2}(r-1) & \frac{1}{2}v & u \end{pmatrix} \succeq 0, \\
& \begin{pmatrix} t_0 + y & \frac{1}{2}(r-1) & \frac{1}{2}r \\ \frac{1}{2}(r-1) & u & \frac{1}{2}v \\ \frac{1}{2}r & \frac{1}{2}v & u \end{pmatrix} \succeq 0, \\
& \begin{pmatrix} t_0 + y & \frac{1}{2}r & \frac{1}{2}r \\ \frac{1}{2}r & u & \frac{1}{2}v \\ \frac{1}{2}r & \frac{1}{2}v & u \end{pmatrix} \succeq 0, \\
& \begin{pmatrix} t_0 & \frac{1}{2}r & \frac{1}{2}r \\ \frac{1}{2}r & u & \frac{1}{2}v \\ \frac{1}{2}r & \frac{1}{2}v & u \end{pmatrix} \succeq 0, \\
& y \ge 0
\end{aligned} \tag{71}$$

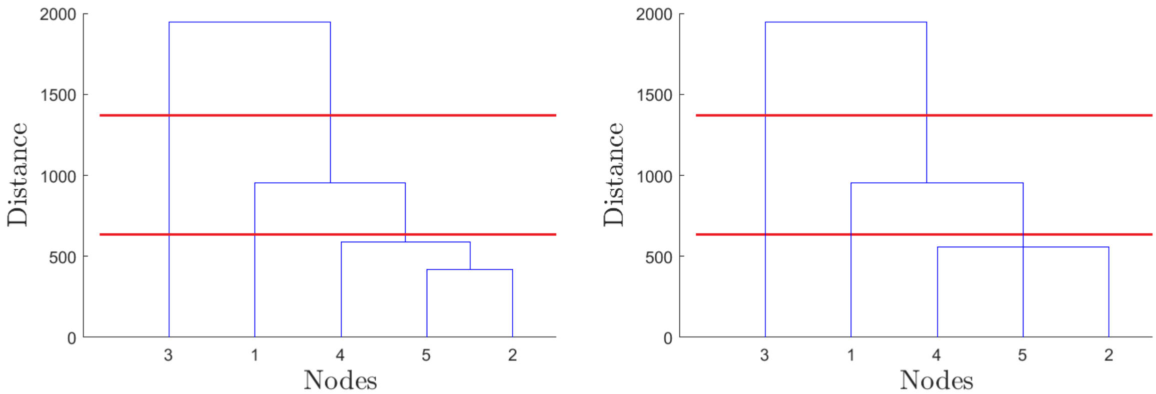

(a) (b)

Figure 12 Creating a nested fulfillment structure with $L=4$ from a dendrogram.

3. **With pooling, ρ unknown:** This is simply an extension of our setting where only marginal information (m,σ) is known, and cross-moment information (ρ) is unknown. The solutions are given through an SDP.

$$
\begin{aligned}
C^{P,NC} := \min_{t_0,r,u,y} \quad & h(2y-2m)+2s_0m+t_0+2rm+2u(m^2+\sigma^2) \\
\text{s.t.} \quad & \\
& \begin{pmatrix} t_0 & \frac{1}{2}r & \frac{1}{2}r \\ \frac{1}{2}r & u & \frac{1}{2}v \\ \frac{1}{2}r & \frac{1}{2}v & u \end{pmatrix} \succeq 0, \\
& \begin{pmatrix} t_0+(s-s_0)y & \frac{1}{2}r & \frac{1}{2}(r-s+s_0) \\ \frac{1}{2}r & u & 0 \\ \frac{1}{2}(r-s+s_0) & 0 & u \end{pmatrix} \succeq 0, \\
& \begin{pmatrix} t_0+(s-s_0)y & \frac{1}{2}(r-s+s_0) & \frac{1}{2}r \\ \frac{1}{2}(r-s+s_0) & u & 0 \\ \frac{1}{2}r & 0 & u \end{pmatrix} \succeq 0, \\
& \begin{pmatrix} t_0+2(p+h-s_0)y & \frac{1}{2}(r-p-h+s_0) & \frac{1}{2}(r-p-h+s_0) \\ \frac{1}{2}(r-p-h+s_0) & u & 0 \\ \frac{1}{2}(r-p-h+s_0) & 0 & u \end{pmatrix} \succeq 0, \\
& \begin{pmatrix} t_0+y(s-s_0)+2y(p+h-s) & \frac{1}{2}(r-p-h+s) & \frac{1}{2}(r-p-h+s-s+s_0) \\ \frac{1}{2}(r-p-h+s) & u & 0 \\ \frac{1}{2}(r-p-h+s-s+s_0) & 0 & u \end{pmatrix} \succeq 0, \\
& \begin{pmatrix} t_0+y(s-s_0)+2y(p+h-s) & \frac{1}{2}(r-p-h+s-s+s_0) & \frac{1}{2}(r-p-h+s) \\ \frac{1}{2}(r-p-h+s-s+s_0) & u & 0 \\ \frac{1}{2}(r-p-h+s) & 0 & u \end{pmatrix} \succeq 0, \\
& y \geq 0
\end{aligned}
\tag{72}
$$

4. **With pooling, ρ known:** This is the setting considered by our paper, and the solutions $y^{P,C}, C^{P,C}$ are given in closed-form in Proposition 1.

Appendix C: Example for Generating Nested Fulfillment Structure with $L<n$

Example 3. Consider Figure 12, where a nested fulfillment structure with $L=4$ is created from the dendrogram in Figure 6a. Here, the range of distances are partitioned into three quantiles by the two lines drawn on the dendrogram. In Figure 12a, the lower line gives rise to three connected components:

$\{\{3\},\{1\},\{4,5,2\}\}$. The nodes in each connected component are considered to be a single cluster in level $l=1$, and the UPGMA distances are recalculated for the new clusters. The upper line gives rise to two connected components: $\{\{3\},\{1,4,5,2\}\}$, which form the two components at level $l=2$, resulting in Figure 12b.

Appendix D: Details for Numerical Experiments

D.1. Constant Fulfillment Heuristic

For $n=5$, the marginal distribution parameters for the four distributions (Normal, Exponential, BetaPrime and Student-t) in the following way:

1. Normal: the means are identical with $m=300$, and the standard deviation is chosen at random from $[100,800]$.
2. Exponential: the mean of the exponential distribution is chosen at random from [100,500]. The standard deviation is equal to the mean.
3. BetaPrime: the mean is fixed at $m=2$. The parameters α and $beta$ are chosen as follows. β is chosen at random from $[2,3]$, and $\alpha=m\cdot(\beta-1)$.
4. Student-t: the parameter ν is chosen at random from $[2,3]$.

We generate 50 such instances of marginal distribution parameters. We generate a random correlation matrix based on Numpacharoen and Atsawarungruangkit (2012). Then, using the method of Gaussian copula, we generate 5000 correlated random demand samples for each distribution, and report the sample average approximations.

D.2. Nested Fulfillment Heuristic

The mean and covariance matrices are calculated based on the populations for each fulfillment center. Given mean m and variance v for a demand distribution, we calculate the marginal distribution parameters for four distributions (Normal, Exponential, BetaPrime and Pareto) as follows:

1. Normal: Mean $\mu=m$, Variance $\sigma^2=v$
2. Exponential: Mean $\frac{1}{\lambda}=m=\sqrt{v}$
3. BetaPrime: $\beta=2+\frac{m\cdot(m+1)}{v}$, $\alpha=m\cdot(\beta-1)$
4. Generalized Pareto: $k=-\frac{v}{m^2}+\sqrt{\frac{v^2}{m^4}+\frac{v}{m^2}}$, $\sigma=m\cdot k\cdot(1-k)$, $\theta=\frac{\sigma}{k}$.

We generate a random correlation matrix based on Numpacharoen and Atsawarungruangkit (2012) such that the correlation coefficients do not exceed .4 in magnitude. We then use the Gaussian copula to generate 10^3 training samples of correlated random vectors. The stochastic solutions are calculated based on a sample average approximation linear program using these training samples, and the robust solution is calculated based on the partitioned statistics estimated from the training samples. The inventory solutions are then evaluated through simulations based on 10^3 test samples generated in a similar fashion to the training samples.

Appendix E: Asymmetry Information

Based on Natarajan et al. (2018), we incorporate into our robust models the partitioned statistics information. Specifically, the mean and covariance of random vector $(\tilde{\mathbf{D}}^+, \tilde{\mathbf{D}}^-)$ whose i^{th} elements are $(\tilde{d}_i - m_i)^+$ and $(m_i - \tilde{d}_i)^+$, respectively, are defined to be:

$$\mathbb{E}\left[\begin{pmatrix}\tilde{\mathbf{D}}^+\\ \tilde{\mathbf{D}}^-\end{pmatrix}\right] =: \bar{\mathbf{m}} \qquad \mathbb{E}\left[\begin{pmatrix}\tilde{\mathbf{D}}^+\\ \tilde{\mathbf{D}}^-\end{pmatrix}\begin{pmatrix}\tilde{\mathbf{D}}^+\\ \tilde{\mathbf{D}}^-\end{pmatrix}^{\intercal}\right] =: \bar{\mathbf{Q}} \tag{73}$$

The set of distributions that the random demand can take is defined as $\bar{\mathcal{F}}_{\geq 0}$, which specifies that the random demand has non-negative support, with mean $\mathbf{m}$, and with mean and covariance of the partitioned statistics given in (73). We follow the same approach as in Theorem 4.3 in Natarajan et al. (2018) to adapt Proposition 2 to derive the following upper bound including the partitioned statistics information. We omit the proof to avoid repetition.

PROPOSITION 3. *For the n-location newsvendor problem under inventory risk pooling with a L-level nested fulfillment cost structure, we have $\sup_{f\in\bar{\mathcal{F}}_{\geq 0}} \mathbb{E}_f[C(\mathbf{y},\tilde{\mathbf{D}})] \leq \bar{\bar{C}}_L(\mathbf{y})$ for any $\mathbf{y}\in\Re^n$, where*

$$\begin{aligned}
\bar{\bar{C}}_L(\mathbf{y}) := \min_{\substack{t_0,\mathbf{t},\mathbf{Y},\mathbf{u},\\ \mathbf{B},\mathbf{W},\mathbf{U},\mathbf{V}}} \quad & h\cdot\mathbf{e}^\top(\mathbf{y}-\mathbf{m}) + \mathbf{s}_0^\top\mathbf{m} + t_0 + \mathbf{t}^\top\bar{\mathbf{m}} + \langle \mathbf{Y},\bar{\mathbf{Q}}\rangle + \mathbf{e}^\top\mathbf{B}\mathbf{e}\\
\text{s.t.} \quad & \begin{pmatrix} t_0 & \frac{1}{2}\mathbf{t}^\top & \frac{1}{2}\mathbf{u}^\top\\ \frac{1}{2}\mathbf{t} & \mathbf{Y} & -\frac{1}{2}\mathbf{V}^\top\\ \frac{1}{2}\mathbf{u} & -\frac{1}{2}\mathbf{V} & \mathbf{U}\end{pmatrix} \succeq 0,\\
& \mathbf{u} = -\mathbf{W}\mathbf{e} + \left(\mathbf{B}+\mathbf{B}^\top\right)\mathbf{e} + \mathbf{P}(\mathbf{y}-\mathbf{m}),\\
& \mathbf{V} \geq \bar{\mathbf{P}},\\
& \mathbf{U} \leq \mathbf{W} - \mathbf{B},\\
& \mathbf{W},\mathbf{B} \geq 0,\\
& t_0\in\Re,\ \mathbf{t}\in\Re^{2n},\ \mathbf{u}\in\Re^N,\ \mathbf{Y}\in\Re^{2n\times 2n},\ \mathbf{B},\mathbf{W},\mathbf{U}\in\Re^{N\times N},\ \mathbf{V}\in\Re^{N\times 2n},
\end{aligned} \tag{74}$$

with $\mathbf{P} := \left(\mathbf{E}_{L-1}^\top \text{diag}(\eta_{L-1})\ \mathbf{E}_{L-2}^\top\text{diag}(\eta_{L-2})\ \cdots\ \mathbf{E}_0^\top\text{diag}(\eta_0)\right)^\top \in \Re^{N\times n}$, and $\bar{\mathbf{P}} = \begin{bmatrix}\mathbf{P} & -\mathbf{P}\end{bmatrix}\in\Re^{N\times 2n}$.

The heuristic solution can be similarly obtained by setting $\mathbf{y}$ as a decision variable, constrained by $\mathbf{y}\geq 0$.

Appendix F: Multiple Demand Channels

To simplify our discussion, we consider a two-level nested fulfillment cost structure for the online demand (i.e., where cross-location fulfillment cost is constant), though the technique can be generalized to an L-level structure. Let p_b and p_o be the penalty cost of unmet brick-and-mortar store demand and online demand, respectively. The per-unit overage cost is h. We normalize the cost for meeting store demand to zero. As before, the cost of in-location fulfillment of online demand is s_0, and the cost of cross-location fulfillment is s, where $s > s_0$. For a customer region $j\in[n]$, let $\tilde{d}_j^o$ and $\tilde{d}_j^b$ be the stochastic online demand and the stochastic store demand, respectively. We denote the vector of online demands as $\tilde{\mathbf{D}}^o = (\tilde{d}_j^o)$ and the vector of store demands as $\tilde{\mathbf{D}}^b = (\tilde{d}_j^b)$. We let $\tilde{\mathbf{D}} = (\tilde{\mathbf{D}}^b, \tilde{\mathbf{D}}^o)$ as the vector of all demands with a mean vector $\mathbf{m} = (\mathbf{m}^b,\ \mathbf{m}^o)$ and covariance matrix $\mathbf{\Sigma}$.

Store demand can only be met with inventory from the same location. However, online demand can be fulfilled from inventory from any location. We assume that $p_o + h > s$, that $p_b + h > s$, and that $p_b + s_0 > p_o$. Given our assumptions on the cost parameters, it is optimal for each local store to first meet the store demands to the maximum extent possible, then for excess inventory to be used to fulfill in-location online demand, before cross-location fulfillment is used. To see why, note that since $p_b + s_0 > p_o$, then it is cheaper to use an inventory unit to meet store demand than to fulfill a local online demand. Moreover, the assumptions imply that $p_b + h + s_0 > s$, so it is cheaper to use cross-location fulfillment on an online demand than to use in-location fulfillment and not meet a store demand. Therefore, we can write the cost as

$$\begin{aligned} C(\mathbf{y},\mathbf{D}) = h \cdot \left(\sum_{j\in[n]} \left(y_j - d_j^b\right)^+ - \sum_{j\in[n]} d_j^o \right)^+ + p_o \cdot \left(\sum_{j\in[n]} d_j^o - \sum_{j\in[n]} \left(y_j - d_j^b\right)^+ \right)^+ \\ + p_b \cdot \sum_{j\in[n]} \left(d_j^b - y_j\right)^+ + s_0 \cdot \left(\sum_{j\in[n]} d_j^o - \sum_{j\in[n]} \left(d_j^o - \left(y_j - d_j^b\right)^+\right)^+ \right) \\ + s \cdot \left(\sum_{j\in[n]} \left(d_j^o - \left(y_j - d_j^b\right)^+\right)^+ - \left(\sum_{j\in[n]} d_j^o - \sum_{j\in[n]} \left(y_j - d_j^b\right)^+ \right)^+ \right) \end{aligned}$$

We observe that, due to the presence of store demand which is prioritized due to its lower cost of fulfillment, the cost structure is more complicated than before. In particular, the last term in the cost function has a composition of a function $f(x) = (a-x)^+$ and $g(\mathbf{x}) = \sum_j x_j^+$. This requires a careful treatment in developing the tractable SDP heuristic. We first simplify the cost function by reducing the number of such terms using the relationship that if $a \geq 0$, then $(a - (b-c)^+) = (a+c-b)^+ - (c-b)^+$. Also using the fact that $(c-b)^+ = b - c + (c-b)^+$, we can simplify the cost function to

$$\begin{aligned} C(\mathbf{y},\mathbf{D}) = h \cdot \mathbf{e}^\top \left(\mathbf{y} - \mathbf{D}^o - \mathbf{D}^b\right) + s_0 \cdot \mathbf{e}^\top \mathbf{D}^o + (h + p_b + s_o - s) \cdot \sum_{j\in[n]} \left(d_j^b - y_j\right)^+ \\ + (s - s_0) \cdot \sum_{j\in[n]} \left(d_j^o + d_j^b - y_j\right)^+ + (h + p_o - s) \cdot \left(\sum_{j\in[n]} \left(d_j^o + d_j^b - y_j\right) - \sum_{j\in[n]} \left(d_j^b - y_j\right)^+ \right)^+ \end{aligned}$$

We define the constants $\gamma := h + p_b + s_o - s$, $\eta_0 := s - s_0$, and $\eta_1 := h + p_o - s$. Hence, the minmax expected cost under the omni-channel demand is equivalent to

$$C_o^* := \min_{\mathbf{y}} \left((s_0 - h) \cdot \mathbf{e}^\top \mathbf{m}_o + h \cdot \mathbf{e}^\top (\mathbf{y} - \mathbf{m}_s) + M_o(\mathbf{y}) \right) \tag{75}$$

where $M_o(\mathbf{y})$ is the optimal value of the moment problem

$$\begin{aligned} M_o(\mathbf{y}) := \max_{f\in\mathcal{F}} \quad \mathbb{E}_f \Bigg[\gamma \cdot \sum_{j\in[n]} \left(\tilde{d}_j^b - y_j\right)^+ + \eta_0 \cdot \sum_{j\in[n]} \left(\tilde{d}_j^o + \tilde{d}_j^b - y_j\right)^+ \\ + \eta_1 \cdot \left(\sum_{j\in[n]} \left(\tilde{d}_j^o + \tilde{d}_j^b - y_j\right) - \sum_{j\in[n]} \left(\tilde{d}_j^b - y_j\right)^+ \right)^+ \Bigg]. \end{aligned} \tag{76}$$

We can write the moment problem as

$$\begin{aligned} M_o(\mathbf{y}) = \max_{f\in\mathcal{F}_{\geq 0}} \quad \mathbb{E}_f \Bigg[\max_{\mathbf{x}^{(0)}\in\{0,1\}^n, x^{(1)}\in\{0,1\}, \mathbf{z}\in\{0,1\}^n} \gamma \cdot \mathbf{z}^\top \left(\tilde{\mathbf{D}}^b - \mathbf{y}\right) + \eta_0 \cdot \mathbf{x}^{(0)\top} \left(\tilde{\mathbf{D}}^o + \tilde{\mathbf{D}}^b - \mathbf{y}\right), \\ + \ \eta_1 \cdot x^{(1)} \cdot \left(\mathbf{e}^\top \left(\tilde{\mathbf{D}}^o + \tilde{\mathbf{D}}^b - \mathbf{y}\right) - \mathbf{z}^\top \left(\tilde{\mathbf{D}}^b - \mathbf{y}\right) \right) \Bigg] \end{aligned} \tag{77}$$

To see why, note that the coefficient of z_j is equal to $\left(\gamma - \eta_1 x^{(1)}\right) \cdot \left(\tilde{d}_j^b - y_j\right)$. Based on our assumptions on the cost parameters, we have that $\gamma = h + p_b + s_o - s > 0$, and $\gamma - \eta_1 = p_b - p_o + s_0 > 0$. Therefore, z_j is equal to 1 if and only if $d_j^b - y_j \geq 0$. Note that unlike in the previous section where the newly introduced variables only interact with other constants or the random demand, we have cross interactions between the new variables from the term $x^{(1)} \cdot \mathbf{z}$. Hence, we introduce a new n-dimensional vector $\mathbf{w} = x^{(1)} \cdot \mathbf{z}$.

Consider the $(3n+1)$-dimensional random vector $\tilde{\mathbf{x}} := \left(\tilde{x}^{(1)\top}\ \tilde{\mathbf{x}}^{(0)\top}\ \tilde{\mathbf{z}}^\top\ \tilde{\mathbf{w}}\right)^\top$, which collects all the new binary variables into a single vector. We again have the following transformation

$$\mathbf{x} := \mathbb{E}_f\left(\tilde{\mathbf{x}}\right) \in \Re^{3n+1}, \tag{78}$$

$$\mathbf{Q} := \mathbb{E}_f\left(\tilde{\mathbf{x}}\tilde{\mathbf{D}}^\top\right) \in \Re^{(3n+1)\times(2n)}, \tag{79}$$

$$\mathbf{R} := \mathbb{E}_f\left(\tilde{\mathbf{x}}\tilde{\mathbf{x}}^\top\right) \in \Re^{(3n+1)\times(3n+1)}. \tag{80}$$

Therefore, we have linearized the objective to

$$\gamma \cdot \sum_{j\in[n]} \left(Q_{1+n+j,j} - x_{1+n+j} \cdot y_j\right) + \eta_0 \cdot \sum_{j\in[n]} \left(Q_{1+j,j} + Q_{1+j,n+j} - x_{1+j} \cdot y_j\right) \tag{81}$$

$$+ \eta_1 \cdot \sum_{j\in[n]} \left(Q_{1,j} + Q_{1,n+j} - x_1 \cdot y_j - Q_{2n+1+j,j} + x_{2n+1+j} \cdot y_j\right) \tag{82}$$

The constraints are the same as before, but with the addition of a few other constraints that follow from the fact that $\tilde{w}_j = \tilde{x}^{(1)} \cdot \tilde{z}_j$ for all $j \in [n]$. In particular, note that

$$R_{1,n+1+j} = x_{1+2n+j}, \qquad \forall j \in [n], \tag{83}$$

$$R_{1+n+i,1+2n+j} = R_{1+2n+1,1+2n+j}, \quad \forall i \in [n], j \in [n]. \tag{84}$$

The first constraint follows since the left-hand side is by definition equal to $\mathbb{E}_f\left(\tilde{x}^{(1)} \cdot \tilde{z}_j\right)$, and the right-hand side is $\mathbb{E}_f\left(\tilde{w}_j\right)$. In the second constraint, the left-hand side is equal to $\mathbb{E}_f\left(\tilde{z}_i \cdot \tilde{w}_j\right) = \mathbb{E}_f\left(\tilde{x}^{(1)}\tilde{z}_i\tilde{z}_j\right)$. The right-hand side is equal to $\mathbb{E}_f\left(\tilde{w}_i \cdot \tilde{w}_j\right) = \mathbb{E}_f\left(\tilde{x}^{(1)}\tilde{x}_i\tilde{x}_j\right)$ since $(\tilde{x}^{(1)})^2 = \tilde{x}^{(1)}$. Due to the nonnegativity of demand, aside from the constraint that $\mathbf{Q} \geq 0$, we also have that $\mathbf{z} \leq \mathbf{x}^{(0)}$. This is because $d_j^b - y_j \geq 0$ impies that $d_j^b + d_j^o - y_j \geq 0$, which is equivalent to the condition that $\tilde{z}_j \leq \tilde{x}_j^{(0)}$.

Proposition 4. *For the n-location newsvendor problem under inventory risk pooling with online and store demand in each location, if the cross-location fulfillment costs of online demand are all equal to s, then*

$\sup_{f\in\mathcal{F}_{\geq 0}} \mathbb{E}_f[C(\mathbf{y},\tilde{\mathbf{D}})] \leq \bar{C}_o(\mathbf{y})$ *for any* $\mathbf{y}\in\Re^n$*, where*

$$\begin{aligned}
\bar{C}_{\mathrm{o}}(\mathbf{y}) := \min_{\substack{t_0,\mathbf{t},\mathbf{Y},\mathbf{u},\\ \mathbf{B},\mathbf{W},\mathbf{U},\mathbf{V},\\ \mathbf{g},\mathbf{h},\mathbf{H}}} \quad & h\cdot\mathbf{e}^\top\left(\mathbf{y}-\mathbf{m}^o-\mathbf{m}^b\right)+s_0\cdot\mathbf{e}^\top\mathbf{m}^o+t_0+\mathbf{t}^\top\mathbf{m}+\langle\mathbf{Y},\boldsymbol{\Sigma}+\mathbf{m}\mathbf{m}^\top\rangle+\mathbf{e}^\top\mathbf{B}\mathbf{e}\\
\text{s.t.} \quad & \begin{pmatrix} t_0 & \frac{1}{2}\mathbf{t}^\top & \frac{1}{2}\mathbf{u}^\top\\ \frac{1}{2}\mathbf{t} & \mathbf{Y} & -\frac{1}{2}\mathbf{V}^\top\\ \frac{1}{2}\mathbf{u} & -\frac{1}{2}\mathbf{V} & \mathbf{U}\end{pmatrix}\succeq 0,\\
& \mathbf{u}=-\mathbf{W}\mathbf{e}+\left(\mathbf{B}+\mathbf{B}^\top\right)\mathbf{e}+\begin{pmatrix}\mathbf{0}_{1,n}\\ -\mathbf{I}_n\\ \mathbf{I}_n\\ \mathbf{0}_{n,n}\end{pmatrix}\mathbf{g}+\begin{pmatrix}\mathbf{0}_{1,n}\\ \mathbf{0}_{n,n}\\ \mathbf{0}_{n,n}\\ -\mathbf{I}_n\end{pmatrix}\mathbf{h}+\begin{pmatrix}\eta_1\cdot\mathbf{e}_n^\top\\ \eta_0\cdot\mathbf{I}_n\\ \gamma\cdot\mathbf{I}_n\\ -\eta_0\cdot\mathbf{I}_n\end{pmatrix}\mathbf{y},\\
& \mathbf{V}\geq\begin{pmatrix}\eta_1\cdot\mathbf{e}_n^\top & \eta_1\cdot\mathbf{e}_n^\top\\ \eta_0\cdot\mathbf{I}_n & \eta_0\cdot\mathbf{I}_n\\ \gamma\cdot\mathbf{I}_n & \mathbf{0}_{n,n}\\ -\eta_0\cdot\mathbf{I}_n & \mathbf{0}_{n,n}\end{pmatrix},\\
& \mathbf{U}\leq\mathbf{W}-\mathbf{B}+\begin{pmatrix}\mathbf{0}_{1,n+1} & \mathbf{h}^\top & \mathbf{0}_n^\top\\ \mathbf{0}_{n,n+1} & \mathbf{0}_{n,n} & \mathbf{0}_{n,n}\\ \mathbf{0}_{n,n+1} & \mathbf{0}_{n,n} & \mathbf{H}\\ \mathbf{0}_{n,n+1} & \mathbf{0}_{n,n} & -\mathbf{H}\end{pmatrix},\\
& \mathbf{g},\mathbf{W},\mathbf{B}\geq 0,\\
& t_0\in\Re,\ \mathbf{g},\mathbf{h}\in\Re^n,\ \mathbf{t}\in\Re^{2n},\ \mathbf{u}\in\Re^{3n+1},\ \mathbf{H}\in\Re^{n\times n},\\
& \mathbf{Y}\in\Re^{2n\times 2n},\ \mathbf{B},\mathbf{W},\mathbf{U}\in\Re^{(3n+1)\times(3n+1)},\ \mathbf{V}\in\Re^{(3n+1)\times 2n}.
\end{aligned} \tag{85}$$

Proof: Suppose that $(z(\mathbf{D}),\mathbf{x}(\mathbf{D}),\mathbf{w}(\mathbf{D}))$ are the optimal recourse variables for demand realization $\mathbf{D}$. Let us define the following variables

$$\begin{pmatrix} z\\ \mathbf{x}\\ \mathbf{w}\end{pmatrix}=\mathbb{E}_f\left(\begin{pmatrix} z(\tilde{\mathbf{D}})\\ \mathbf{x}(\tilde{\mathbf{D}})\\ \mathbf{w}(\tilde{\mathbf{D}})\end{pmatrix}\right) \tag{86}$$

$$\begin{pmatrix}\mathbf{y}_{zs}^\top & \mathbf{y}_{zo}^\top\\ \mathbf{Y}_{xs} & \mathbf{Y}_{xo}\\ \mathbf{Y}_{ws} & \mathbf{Y}_{wo}\end{pmatrix}=\mathbb{E}_f\left(\begin{pmatrix} z(\tilde{\mathbf{D}})\\ \mathbf{x}(\tilde{\mathbf{D}})\\ \mathbf{w}(\tilde{\mathbf{D}})\end{pmatrix}\begin{pmatrix}\tilde{\mathbf{D}}_s\\ \tilde{\mathbf{D}}_o\end{pmatrix}^\top\right) \tag{87}$$

$$\bar{\mathbf{X}}=\mathbb{E}_f\left(\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top\right) \tag{88}$$

$$\hat{\mathbf{X}}=\mathbb{E}_f\left(z(\mathbf{D})\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top\right). \tag{89}$$

Also define the constants

$$\boldsymbol{\Sigma}+\mathbf{m}\mathbf{m}^\top=\begin{pmatrix}\mathbf{Q}_{ss} & \mathbf{Q}_{so}^\top\\ \mathbf{Q}_{so} & \mathbf{Q}_{oo}\end{pmatrix}. \tag{90}$$

Note that

$$\begin{pmatrix}1\\ \mathbf{D}_s\\ \mathbf{D}_o\\ z(\mathbf{D})\\ \mathbf{x}(\mathbf{D})\\ \mathbf{w}(\mathbf{D})\end{pmatrix}\begin{pmatrix}1\\ \mathbf{D}_s\\ \mathbf{D}_o\\ z(\mathbf{D})\\ \mathbf{x}(\mathbf{D})\\ \mathbf{w}(\mathbf{D})\end{pmatrix}^\top=\begin{pmatrix}
1 & \mathbf{D}_s^\top & \mathbf{D}_o^\top & z(\mathbf{D}) & \mathbf{x}(\mathbf{D})^\top & \mathbf{w}(\mathbf{D})^\top\\
\mathbf{D}_s & \mathbf{D}_s\mathbf{D}_s^\top & \mathbf{D}_s\mathbf{D}_o^\top & \mathbf{D}_s z(\mathbf{D}) & \mathbf{D}_s\mathbf{x}(\mathbf{D})^\top & \mathbf{D}_s\mathbf{w}(\mathbf{D})^\top\\
\mathbf{D}_o & \mathbf{D}_o\mathbf{D}_s^\top & \mathbf{D}_o\mathbf{D}_o^\top & \mathbf{D}_o z(\mathbf{D}) & \mathbf{D}_o\mathbf{x}(\mathbf{D})^\top & \mathbf{D}_o\mathbf{w}(\mathbf{D})^\top\\
z(\mathbf{D}) & z(\mathbf{D})\mathbf{D}_s^\top & z(\mathbf{D})\mathbf{D}_o^\top & z(\mathbf{D}) & \mathbf{w}(\mathbf{D})^\top & \mathbf{w}(\mathbf{D})^\top\\
\mathbf{x}(\mathbf{D}) & \mathbf{x}(\mathbf{D})\mathbf{D}_s^\top & \mathbf{x}(\mathbf{D})\mathbf{D}_o^\top & \mathbf{w}(\mathbf{D}) & \mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top & z(\mathbf{D})\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top\\
\mathbf{w}(\mathbf{D}) & \mathbf{w}(\mathbf{D})\mathbf{D}_s^\top & \mathbf{w}(\mathbf{D})\mathbf{D}_o^\top & \mathbf{w}(\mathbf{D}) & z(\mathbf{D})\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top & z(\mathbf{D})\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top
\end{pmatrix} \tag{91}$$

where we use the fact that $z(\mathbf{D})^2 = z(\mathbf{D})$, $z(\mathbf{D})\mathbf{x}(\mathbf{D}) = \mathbf{w}(\mathbf{D})$, $\mathbf{z}(\mathbf{D})\mathbf{w}(\mathbf{D}) = z(\mathbf{D})^2\mathbf{x}(\mathbf{D}) = z(\mathbf{D})\mathbf{x}(\mathbf{D}) = \mathbf{w}(\mathbf{D})$, $\mathbf{x}(\mathbf{D})\mathbf{w}(\mathbf{D})^\top = z(\mathbf{D})\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top$, $\mathbf{w}(\mathbf{D})\mathbf{w}(\mathbf{D})^\top = z(\mathbf{D})^2\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top = z(\mathbf{D})\mathbf{x}(\mathbf{D})\mathbf{x}(\mathbf{D})^\top$. Taking the expectation on both sides, we have that

$$\begin{pmatrix} 1 & \mathbf{m}_s^\top & \mathbf{m}_o^\top & z & \mathbf{x}^\top & \mathbf{w}^\top \\ \mathbf{m}_s & \mathbf{Q}_{ss} & \mathbf{Q}_{so}^\top & \mathbf{y}_{zs} & \mathbf{Y}_{xs}^\top & \mathbf{Y}_{ws}^\top \\ \mathbf{m}_o & \mathbf{Q}_{so} & \mathbf{Q}_{oo} & \mathbf{y}_{zo} & \mathbf{Y}_{xo}^\top & \mathbf{Y}_{wo}^\top \\ z & \mathbf{y}_{zs}^\top & \mathbf{y}_{zo}^\top & z & \mathbf{w}^\top & \mathbf{w}^\top \\ \mathbf{x} & \mathbf{Y}_{xs} & \mathbf{Y}_{xo} & \mathbf{w} & \bar{\mathbf{X}} & \hat{\mathbf{X}} \\ \mathbf{w} & \mathbf{Y}_{ws} & \mathbf{Y}_{wo} & \mathbf{w} & \hat{\mathbf{X}} & \hat{\mathbf{X}} \end{pmatrix} \succeq 0, \tag{92}$$

and that

$$\begin{pmatrix} 1 & z & \mathbf{x}^\top & \mathbf{w}^\top \\ z & z & \mathbf{w}^\top & \mathbf{w}^\top \\ \mathbf{x} & \mathbf{w} & \bar{\mathbf{X}} & \hat{\mathbf{X}} \\ \mathbf{w} & \mathbf{w} & \hat{\mathbf{X}} & \hat{\mathbf{X}} \end{pmatrix} \in \text{BQP}. \tag{93}$$

Note that the a linear relaxation of the BQP constraints is the following:

$$\mathbf{w} \le \mathbf{x}, \tag{94}$$

$$\mathbf{w} \le z \cdot \mathbf{e}, \tag{95}$$

$$-\mathbf{w} + \mathbf{x} + z \cdot \mathbf{e} \le 1, \tag{96}$$

$$\bar{X}_{ii} = x_i, \tag{97}$$

$$\bar{X}_{ij} \le x_i, \tag{98}$$

$$-\bar{X}_{ij} + x_i + x_j \le 1, \tag{99}$$

$$\hat{X}_{ii} = w_i, \tag{100}$$

$$\hat{X}_{ij} \le w_i, \tag{101}$$

$$-\hat{X}_{ij} + w_i + w_j \le 1, \tag{102}$$

$$\hat{X}_{ij} \le x_i, \tag{103}$$

$$-\hat{X}_{ij} + x_i + w_j \le 1. \tag{104}$$

Removing redundant constraints and taking the dual of this SDP gives the Lemma. Q.E.D.