Convex Optimization of Centralized Inventory Operations

Samuel Burer* Moshe Dror†

January 10, 2005

Abstract

Given a finite set of outlets with joint normally distributed demands and identical holding and penalty costs, inventory centralization induces a cooperative cost allocation game with nonempty core. It is well known that for this newsvendor inventory setting the expected cost of centralization can be expressed as a constant multiple of the standard deviation of the joint distribution. The lowering of the centralized cost without changing the mean and variance of demand at each outlet corresponds to a semidefinite optimization problem. This paper establishes a closed-form optimal solution of the semidefinite program and a core allocation of the cost at optimality. The issue of cost (and benefit) allocation separate from the optimization is also studied and it is shown that an exponential-size linear program can be approximated by a polynomial-size second-order program.

1 Introduction

A finite set of outlets with randomly fluctuating demand bands together to reduce costs by buying, storing and distributing their inventory jointly. This is termed *inventory centralization.* In order to keep the set of outlets in a grand coalition in the centralized inventory arrangement, the expected costs (and benefits) should be allocated among the outlets so that no subset has incentive to break off from the coalition – core allocation. The expected centralization cost can be lowered even further, without disrupting the demand behavior at individual outlets, by inducing the outlets to correlate their individual demands so that the random variations help to balance each other out.

There are many real-life examples where companies attempt to reduce inventory and manufacturing costs by manipulating correlated demand sources, or inducing demand patterns that balance each other out in an average sense. Hartman and Dror (2003) list a number of references on this topic and discuss strategies for adjusting correlations without changing individual means and variances. It is clear that companies understand well the

*Department of Management Sciences, University of Iowa, Iowa City, IA 52242-1000, USA. (Email: samuel-burer@uiowa.edu). This research was supported in part by NSF Grant CCR-0203426.

†OR Center, MIT & MIS Department, The University of Arizona.

benefits of manipulating correlations between demand points. The challenge is to provide computational strategies that determine an optimal covariance solution and allocate costs fairly.

In this paper, we restrict the analysis to outlets with single-period normally distributed, correlated individual demands and identical holding and penalty costs. We assume an infinite horizon setting with risk neutrality for costs and benefits. It is well known (Eppen, 1979; Hartman et al., 2000) that, in this inventory setting — known as the *newsvendor* setting — the expected cost of inventory centralization can be expressed as a constant multiple of the standard deviation. Furthermore, Hartman and Dror (2003) show that minimizing the centralization cost leads to a semidefinite optimization problem.

Our primary goal in this paper is to study the structure of the resulting semidefinite optimization problem. By extending some results from a branch of statistics called *minimum trace factor analysis*, we establish a closed-form optimal solution for the inventory centralization problem.

A secondary goal of this paper is to provide a fair allocation of costs to the outlets, where the precise meaning of "fair" is as described by Hartman and Dror (1996). We examine the case where the cost allocations are allowed to be negative (i.e., an outlet is given money to participate because its presence provides a significant benefit to the grand coalition), as well as the more typical case when all cost allocations are nonnegative. We show that, if an optimal inventory centralization has been achieved, then a fair cost allocation (either nonnegative or unrestricted) can be stated in a simple, closed form. For the general situation with outlet demands having arbitrary, non-optimized correlations, we discuss the theoretical difficulties of determining a fair allocation (or even whether one exists), the main challenge being the assessment of the feasibility of a linear program having an exponential number of constraints. As an alternative, we propose a heuristic approach that replaces this linear program with a polynomial-sized second-order program.

2 Elements of the Model

Let $U = \{1, \ldots, n\}$ be a set of outlets. Given a fixed duration time period, we assume that the demand at each outlet $i \in U$ is normally distributed with mean $\mu_i > 0$ and standard deviation $\sigma_i > 0$. We consider a static infinite horizon setting (i.e., an infinite number of time periods) where expected inventory cost is the metric. Inventory cost is composed of the holding cost h (per unit held) and the penalty cost p (per unit out of stock); we assume these numbers are constant for any assembly $S \subseteq U$ of outlets that bands together to centralize inventory. Even when the distributions for the individual outlets are correlated, the expected inventory cost for S is proportional to the standard deviation σ_S of their joint distribution (Eppen, 1979; Hartman et al., 2000), where the proportionality constant depends only on p, h, and the standard normal distribution and hence can be scaled to 1. Thus, the expected cost for any S can be represented by

$$c(S) = \sigma_S = \sqrt{\sum_{i \in S} \sum_{j \in S} \sigma_{ij}} = \sqrt{\sum_{i \in S} \sum_{j \in S} \sigma_i \sigma_j \rho_{ij}},$$

where σ_{ij} and ρ_{ij} represent the covariances and correlations of the joint distribution of the demand for S, which is normal since the individual outlet demands are normal.

Eppen (1979) and Hartman et al. (2000) showed that, in total, it is always cheaper for all outlets to participate in a single, centrally managed inventory system modelled as above, due to the inequality

$$c(U) = \sigma_U = \sqrt{\sum_{i=1}^{n}\sum_{j=1}^{n}\sigma_{ij}} = \sqrt{\sum_{i=1}^{n}\sigma_i^2 + 2\sum_{j>i}\sigma_i\sigma_j\rho_{ij}} \leq \sum_{i=1}^{n}\sigma_k = \sum_{i=1}^{n}c(\{i\}).$$

It is clear from the above inequality that low or negative correlations yield a lower cost to the coalition when the individual distributions of demand at the outlets are fixed. Said differently, a high degree of correlation between the individual demands forces the variance of the joint distribution higher. If it were possible to adjust correlations lower, while maintaining demands and variances at each outlet, then we could lower the coalition's overall expected cost.

Is it possible to adjust correlations in such a manner, however? Hartman and Dror (2003) discuss several real-life examples in which companies understand the benefit of — and implement strategies for — reducing these cross correlations. An assumption of this paper is that it is indeed possible to modify correlations in some interesting settings, and based upon this assumption, we present techniques for finding the correlations that minimize $c(U)$.

Another point of view is to examine the total benefits (or savings) received by a coalition S of outlets when it centrally manages its inventory. This benefit is measured as

$$v(S) = \sum_{i\in S} c(\{i\}) - c(S).$$

It is clear that minimizing $c(U)$ is equivalent to maximizing $v(U)$.

Once the correlations between the outlets have been determined, it is important to allocate the total cost $c(U)$ to all outlets in a "fair" manner. For example, one possible allocation would be to charge the entire cost $c(U)$ to outlet 1 and zero cost to outlets $2,\ldots,n$. While outlets $2,\ldots,n$ would see this allocation as beneficial to them, outlet 1 would undoubtedly judge this to be unfair. As a result, outlet 1 is likely to leave the coalition. In a similar fashion, it is important to distribute the benefits $v(U)$ fairly.

In order to define what a fair allocation is, we follow Hartman and Dror (2003) and interpret the centralization as a cooperative game, which is defined formally as follows: a *cooperative n-person game on U*, in characteristic function form, is an ordered pair $(U; f)$, where $f : 2^U \to \Re$ is a real-valued set-function on the collection 2^U of all subsets of U such that $f(\emptyset) = 0$. In our inventory model, we may take $f(\cdot) = c(\cdot)$ and $f(\cdot) = v(\cdot)$. The corresponding games are referred to as the *cost game* or *benefit game*, respectively.

Before defining our notion of a fair allocation, we first introduce a bit of notation. For any subset $A \subseteq U$ and vector $q \in \Re^n$, we use the notation $q(A)$ to define the sum of the components in q corresponding to the members in A, i.e., $q(A) = \sum_{i\in A} q_i$.

A *cost allocation* is a vector $a = (a_1,\ldots,a_n)$ that assigns to each outlet its portion of the cost $c(U)$. Note that a may have negative entries. Likewise, a *benefit allocation* is a vector $x = (x_1, \ldots, x_n)$ whose coordinates are the benefits received by the members of U. We consider cost and benefit allocations a and x to be fair if they satisfy the following conditions:

1. *Efficiency* (all of the costs and benefits of the inventory system are distributed out):

$$a(U) = c(U) \qquad \text{and} \qquad x(U) = v(U);$$

2. *Core*, or *stability* (no subset of outlets has an incentive to leave the coalition, either because it feels it is subsidizing the remaining outlets by paying too much of the cost or because it could benefit more by setting up its own inventory management system):

$$a(S) \leq c(S) \qquad \text{and} \qquad x(S) \geq v(S), \qquad \forall \;\; S \subseteq U;$$

3. *Justifiability* (an individual outlet's allocated cost plus allocated benefit should equal the cost incurred if the outlet were managing its own inventory):

$$a_i + x_i = c(\{i\}), \qquad \forall \;\; i \in U.$$

A few remarks regarding these fairness conditions are in order. First, allocations satisfying the core condition have been proven to exist for newsvendor centralization problems like ours, under the assumption that the demand at all outlets is of the same distribution type. This was first conjectured by Hartman et al. (2000) and proven by Muller et al. (2002). Our additional assumption of normal distributions at each outlet further enables the expected cost to be represented analytically, which is not true for most other distributions. Second, if one also imposes nonnegativity on a and x, then a fair (core) allocation may not exist. Third, because $c(\{i\}) = \sigma_i$ in our model, the justifiability condition can be expressed compactly as $a + x = \sigma$. Moreover, because the core condition for singleton subsets implies $a \leq \sigma$, it is not difficult to see that an efficient, core allocation a gives rise to an efficient, core allocation x by the formula $x = \sigma - a$, and as such, a and x are justifiable as well.

There may be multiple allocations that satisfy the above fairness conditions. Relative to the efficiency and core conditions and specific to a single game (either cost or benefit), Schmeidler (1969) described how to compute a specific allocation (either nonnegative or unrestricted) that "makes the least-well-off coalition as well-off as possible" (Young, 1985); this allocation is called the *nucleolus*. Thus, from now on we focus on the nucleolus as the candidate cost allocation scheme even though many other equally viable allocation schemes in the core may exist. The nucleolus can be calculated by solving a sequence of n successively refined linear programs (LPs), each having roughly 2^n constraints. (In Section 4, we provide a full description of this procedure.) In the case of our cost game, the initial LP that one must solve is as follows:

$$\begin{aligned} \max \quad & \varepsilon \\ \text{s.t.} \quad & c(S) - a(S) \geq \varepsilon \qquad \forall \;\; \emptyset \neq S \subsetneq U \\ & a(U) = c(U) \\ & [\, a \geq 0 \,] \end{aligned} \tag{1}$$

(The brackets around the constraint $a \geq 0$ indicate that it can be enforced if desired.) Note that any feasible solution (a, ε) to this LP automatically satisfies the efficiency condition. Moreover, the core condition is met if and only if $\varepsilon \geq 0$. In this sense, solving (1) can be interpreted as finding an allocation, which is as fair as possible, i.e., an allocation for which ε is as large as possible. Moreover, the comments of the previous paragraph imply that, if nonnegativity is enforced, then there may be no feasible solution with $\varepsilon \geq 0$; if not, then $\varepsilon \geq 0$ is always attainable. Finally, one can also define a similar initial LP for calculating the nucleolus with respect to the benefit game.

In the remaining sections of this paper, we consider three issues: (i) adjusting correlations so as to minimize $c(U)$; (ii) calculating a fair allocation a after $c(U)$ has been minimized; and (iii) approximating a fair allocation a even when $c(U)$ has not been optimized.

3 Adjusting The Correlations – The Semidefinite Optimization

Since we do not wish to change the distributions at the individual outlets, we assume that μ_i and σ_i are fixed for all $i \in U$. Letting R denote the correlation matrix consisting of correlations ρ_{ij}, we know from Section 2 that our optimization problem is to find R that minimizes $c(U)$, or equivalently minimizes

$$c(U)^2 = \sum_{i=1}^{n} \sigma_i^2 + 2 \sum_{j>i} \sigma_i \sigma_j \rho_{ij} = \sigma^T R \sigma.$$

Since the set of correlation matrices can be described as all symmetric, positive semidefinite matrices with ones on the diagonal, we can express our optimization as

$$\begin{aligned} \min \quad & \sigma^T R \sigma \\ \text{s.t.} \quad & \operatorname{diag}(R) = e \\ & R \succeq 0, \end{aligned} \tag{2}$$

where e is the all-ones vector and the constraint $R \succeq 0$ indicates that R is symmetric and positive semidefinte. As the objective and diagonal constraints are linear in R, (2) is a (linear) semidefinite program (SDP), which is a type of convex programming problem that can be solved up to any desired accuracy in polynomial-time using interior-point methods (Wolkowicz et al., 2000).

In the context of inventory centralization, the SDP (2) was introduced by Hartman and Dror (2003), but SDPs having the same constraint structure have been studied for some time. In particular, the set of correlation matrices serves as the basis for relaxations of the maximum-cut problem (Goemans and Williamson, 1995), which refers to the decomposition of a graph by deletion of edges of maximum weight. In addition, optimization over correlations matrices has been investigated in the statistics literature on minimum trace factor analysis. In fact, based on the specific structure of the objective function of (2), Shapiro

(1982) has provided a partial classification of the optimal solutions of (2). To state the theorem, we make the following assumption without loss of generality:

Assumption. The components of σ are sorted in nonincreasing order.

Then the theorem of Shapiro (1982) is as follows:

Theorem 3.1 *Let $n \geq 2$, and let $v \in \Re^n$ be the vector having all -1's except for a 1 in the first position. It holds that:*

(a) if $v^T\sigma > 0$, then $R = vv^T$ is the unique optimal solution of (2) with optimal value $(v^T\sigma)^2$;

(b) if $v^T\sigma = 0$, then $R = vv^T$ is an optimal solution of (2) with optimal value 0; and

(c) if $v^T\sigma < 0$, then the optimal value of (2) is 0, so that all optimal solutions satisfy the equation $R\sigma = 0$.

Theorem 3.1 provides valuable information about (2) — and does so practically, since one can easily check the sign of $v^T\sigma$. It is important to point out, however, that Theorem 3.1 does not specifically provide an optimal solution when $v^T\sigma < 0$, i.e., we still must calculate an optimal solution on our own. In the following subsection, we examine simple ways to calculate such a solution.

3.1 Calculating an optimal solution when $v^T\sigma < 0$

In this subsection, we assume $v^T\sigma < 0$, where v is as in Theorem 3.1; note that $v^T\sigma < 0$ implies $n \geq 3$. By Theorem 3.1, the optimal value of the inventory centralization is 0. We wish to calculate an optimal correlation matrix R satisfying $\sigma^T R\sigma = 0$ or, equivalently, $R\sigma = 0$.

Our analysis will be based on the fundamental fact that any symmetric positive semidefinite R can be factored as $R = VV^T$ for some $V \in \Re^{n\times n}$. (Note that the factorization is not unique in general.) Thus, our search for a correlation matrix R can be cast as a search for V with unit-length rows such that $\sigma^T VV^T\sigma = \|\sigma^T V\|^2 = 0$, or equivalently, $\sigma^T V = 0$.

The factor V can in fact be taken to have rank(R) columns. However, we have no prior knowledge of this rank except for the following: rank(R) $\leq n - 1$ because σ is a nonzero vector in the null space of R. So without loss of generality, we can restrict our search to $V \in \Re^{n\times(n-1)}$. (In the results below, we will in fact demonstrate an optimal R such that rank(R) ≤ 2.)

3.1.1 The case for $n = 3$

We consider first the case when $n = 3$. Given $\sigma = (\sigma_1, \sigma_2, \sigma_3)^T$, define

$$R^* = \begin{pmatrix} 1 & \rho_{12} & \rho_{13} \\ \rho_{12} & 1 & \rho_{23} \\ \rho_{13} & \rho_{23} & 1 \end{pmatrix} \quad \text{and} \quad V^* = \begin{pmatrix} 1 & 0 \\ \rho_{12} & \sqrt{1-\rho_{12}^2} \\ \rho_{13} & -\sqrt{1-\rho_{13}^2} \end{pmatrix}, \tag{3}$$

where

$$\rho_{12} = \left(\sigma_3^2 - \sigma_1^2 - \sigma_2^2\right) / \left(2\,\sigma_1\sigma_2\right)$$
$$\rho_{13} = \left(\sigma_2^2 - \sigma_1^2 - \sigma_3^2\right) / \left(2\,\sigma_1\sigma_3\right)$$
$$\rho_{23} = \left(\sigma_1^2 - \sigma_2^2 - \sigma_3^2\right) / \left(2\,\sigma_2\sigma_3\right).$$

We have the following proposition:

Proposition 3.1 *Suppose $\sigma \in \Re^3$ satisfies $v^T\sigma < 0$, and define R^* and V^* by (3). Then $R^* = V^*(V^*)^T$, and R^* is the unique optimal solution of (2). Specifically, $\sigma^T R^* \sigma = 0$ and $\sigma^T V^* = 0$.*

Proof. We remark that the assumption $v^T\sigma < 0$ implies any optimal solution to (2) must have rank equal to 2, which is consistent with the proposed structure for V^* (i.e., 2 columns).

We first show that V^* is well defined, i.e., that $|\rho_{12}| \le 1$ and $|\rho_{13}| \le 1$. In particular, we show $|\rho_{12}| < 1$ and $|\rho_{13}| < 1$. Since ρ_{12} and ρ_{13} are defined similarly, we give the argument for ρ_{12} only. If the numerator of ρ_{12} is nonnegative, then it suffices to show $\rho_{12} < 1$, which is equivalent to

$$\sigma_3^2 - \sigma_1^2 - \sigma_2^2 < 2\,\sigma_1\sigma_2 \quad\Longleftrightarrow\quad \sigma_3^2 < (\sigma_1 + \sigma_2)^2 \quad\Longleftrightarrow\quad \sigma_3 < \sigma_1 + \sigma_2,$$

which is true because σ is sorted in descending order. On the other hand, if the numerator is negative, then we require $\rho_{12} > -1$, which is equivalent to

$$\sigma_3^2 - \sigma_1^2 - \sigma_2^2 > -2\,\sigma_1\sigma_2 \quad\Longleftrightarrow\quad \sigma_3^2 > (\sigma_1 - \sigma_2)^2 \quad\Longleftrightarrow\quad \sigma_3 > \max\{\sigma_1 - \sigma_2, \sigma_2 - \sigma_1\},$$

which is also true because $v^T\sigma < 0$ and because σ is sorted. It follows that $|\rho_{12}| < 1$.

We next argue $R^* = V^*(V^*)^T$. It is clear that $\operatorname{diag}\left(V^*(V^*)^T\right) = e$, $(V^*(V^*)^T)_{12} = \rho_{12}$, and $(V^*(V^*)^T)_{13} = \rho_{13}$, and so it remains to show

$$\rho_{12}\rho_{13} - \sqrt{1-\rho_{12}^2}\sqrt{1-\rho_{13}^2} = \rho_{23}. \tag{4}$$

We first remark that one can readily show $\sigma_2^2\,(1-\rho_{12}^2) = \sigma_3^2\,(1-\rho_{13}^2)$, which implies the left-hand side of (4) equals $\rho_{12}\rho_{13} - (\sigma_2/\sigma_3)\,(1-\rho_{12}^2)$. This expression in turn simplifies to ρ_{23}.

To complete the proof of the proposition, we demonstrate that any $V \in \Re^{3\times 2}$ having unit-length rows and satisfying $\sigma^T V = 0$ must yield $R^* = VV^T$. We note that, for all orthogonal $Q \in \Re^{2\times 2}$: (i) $VV^T = (VQ)(VQ)^T$; and (ii) the equation $\sigma^T V = 0$ holds if and only if $\sigma^T VQ = 0$. As a result, we may assume without loss of generality that a rotation has been applied to the rows of V so that $V_{1\cdot} = (1, 0)$. We write

$$V = \begin{pmatrix} 1 & 0 \\ a & b \\ c & d \end{pmatrix}$$

and require

$$a^2 + b^2 = 1 \qquad c^2 + d^2 = 1 \qquad \sigma_1 + \sigma_2 a + \sigma_3 c = 0 \qquad \sigma_2 b + \sigma_3 d = 0.$$

These equations imply

$$a^2 + (\sigma_3/\sigma_2)^2 d^2 = 1 \qquad (1/\sigma_3)^2 (\sigma_1 + \sigma_2 a)^2 + d^2 = 1,$$

which in turn imply $a = \rho_{12}$ after substituting for d^2. A similar argument yields $c = \rho_{13}$. Since the equation $\sigma_2 b + \sigma_3 d = 0$ implies that b and d have opposite signs, we have

$$(b, d) = \pm \left(\sqrt{1 - \rho_{12}^2}, -\sqrt{1 - \rho_{13}^2} \right).$$

In either case, the product VV^T equals R^*. ■

Proposition 3.1 provides the optimal solution for the case when $n = 3$, and although this result may seem specialized, we show in the next subsection that it actually allows us to solve (2) easily for arbitrary n.

3.1.2 The case for arbitrary $n \geq 3$

Let $\mathcal{S}$ be a bi- or tri-partition of $\{1, \ldots, n\}$; we write $\mathcal{S} = \{S_1, S_2\}$ or $\mathcal{S} = \{S_1, S_2, S_3\}$ as appropriate. Define

$$\bar{\sigma} \in \Re^{|\mathcal{S}|} \qquad \bar{\sigma}_j = \sigma(S_j) \qquad 1 \leq j \leq |\mathcal{S}|. \tag{5}$$

We say that $\mathcal{S}$ is *balanced* if

$$\bar{\sigma}_1 = \bar{\sigma}_2 \qquad \text{when } |\mathcal{S}| = 2$$

or if

$$\bar{\sigma}_1 < \bar{\sigma}_2 + \bar{\sigma}_3 \qquad \bar{\sigma}_2 < \bar{\sigma}_1 + \bar{\sigma}_3 \qquad \bar{\sigma}_3 < \bar{\sigma}_1 + \bar{\sigma}_2 \qquad \text{when } |\mathcal{S}| = 3. \tag{6}$$

Intuitively, $\mathcal{S}$ divides the outlets into groups, where the total amount of standard deviation of demand in each group is equal — or not excessively high. The following result, the proof of which is constructive, establishes that σ admits a balanced partition.

Proposition 3.2 *Let $n \geq 3$, and suppose $v^T\sigma < 0$. Then there exists a balanced bi-partition $\mathcal{S} = \{S_1, S_2\}$ or a balanced tri-partition $\mathcal{S} = \{S_1, S_2, S_3\}$ of $\{1, \ldots, n\}$ with respect to σ.*

Proof. Because σ is sorted in nonincreasing order, it is trivial to determine whether there exists k such that $\sigma_1 + \cdots + \sigma_k = \sigma_{k+1} + \cdots + \sigma_n$. If so, then define $\mathcal{S} = \{S_1, S_2\}$, where $S_1 = \{1, \ldots, k\}$ and $S_2 = \{k + 1, \ldots, n\}$; $\mathcal{S}$ is the required bi-partition.

Otherwise, let k be the index such that

$$\begin{aligned} \sigma_1 + \cdots + \sigma_{k-1} &< \sigma_k + \sigma_{k+1} + \cdots + \sigma_n \\ \sigma_1 + \cdots + \sigma_{k-1} + \sigma_k &> \sigma_{k+1} + \cdots + \sigma_n \end{aligned}$$

We certainly have $k \geq 2$ because $v^T\sigma < 0$; we also have $k \leq n - 1$ because the components of σ are sorted and because $n \geq 3$. Now define $\mathcal{S}$ by

$$S_1 = \{1, \ldots, k - 1\} \qquad S_2 = \{k\} \qquad S_3 = \{k + 1, \ldots, n\},$$

and also define $\bar{\sigma}$ according to (5). We claim that $\mathcal{S}$ is balanced. The first inequality above shows $\bar{\sigma}_1 < \bar{\sigma}_2 + \bar{\sigma}_3$, while the second inequality shows $\bar{\sigma}_3 < \bar{\sigma}_1 + \bar{\sigma}_2$. Finally, we have

$$\bar{\sigma}_2 = \sigma_k \leq \sigma_1 < (\sigma_1 + \cdots + \sigma_{k-1}) + (\sigma_{k+1} + \cdots + \sigma_n) = \bar{\sigma}_1 + \bar{\sigma}_3,$$

which follows because σ is sorted and because $n \geq 3$. So $\mathcal{S}$ is the desired tri-partition. ∎

We remark that, in the case of a balanced tri-partition, we may assume without loss of generality that $\mathcal{S}$ has been defined so that the components of $\bar{\sigma}$ are sorted in nonincreasing order. Moreover, if we let $\bar{v} = (1, -1, -1)^T$ be as in Theorem 3.1 relative to $\bar{\sigma}$, then the three conditions (6) are equivalent to the single inequality $\bar{v}^T \bar{\sigma} < 0$.

Using Propositions 3.1 and 3.2, the following theorem shows how to construct an optimal solution of (2) easily.

Theorem 3.2 *Let $n \geq 3$, and suppose that $\sigma \in \Re^n$ satisfies $v^T\sigma < 0$. Let $\mathcal{S}$ be any balanced bi- or tri-partition relative to σ, and define $\bar{\sigma}$ by (5). Define $\bar{V}^* \in \Re^{|\mathcal{S}|\times(|\mathcal{S}|-1)}$ as follows:*

(i) if $|\mathcal{S}| = 2$, then $\bar{V}^* = (1, -1)^T$;

(ii) if $|\mathcal{S}| = 3$, then $\bar{V}^*$ is as defined by (3) with respect to $\bar{\sigma}$.

In addition, define $V^* \in \Re^{n\times(|\mathcal{S}|-1)}$ row-by-row as

$$V^*_{i\cdot} = \bar{V}^*_{j\cdot} \quad \text{if } \ i \in S_j \qquad (1 \leq i \leq n, \ 1 \leq j \leq |\mathcal{S}|). \tag{7}$$

and $R^* = V^*(V^*)^T$. Then R^* is an optimal solution of (2). In particular, $\sigma^T R^* \sigma = 0$ and $\sigma^T V^* = 0$.

Proof. By construction, each row of V^* has unit norm so that R^* is feasible for (2). It remains to show that R^* is optimal, which is implied by the following equation:

$$\sigma^T V^* = \sum_{i=1}^{n} \sigma_i V^*_{i\cdot} = \sum_{j=1}^{|\mathcal{S}|}\sum_{i\in S_j} \sigma_i V^*_{i\cdot} = \sum_{j=1}^{|\mathcal{S}|}\sum_{i\in S_j} \sigma_i \bar{V}^*_{j\cdot} = \sum_{j=1}^{|\mathcal{S}|} \sigma(S_j)\bar{V}^*_{j\cdot} = \sum_{j=1}^{|\mathcal{S}|} \bar{\sigma}_j \bar{V}^*_{j\cdot} = \bar{\sigma}^T \bar{V}^* = 0.$$

∎

We remark that Theorem 3.2 provides only a single optimal solution R^* of (2) and that R^* is dependent on the partition $\mathcal{S}$. In general, (2) may have multiple optimal solutions. However, R^* is "minimal" in the sense that its rank, which equals 1 or 2, is the smallest of all optimal solutions.

4 Computing Fair Allocations For Optimized Correlations

In this section, we assume that we have optimized the total cost $c(U)$ (or maximized the total benefit $v(U)$) by adjusting correlations between the outlets, without affecting the demand distribution at each outlet (see Section 3). We now turn our attention to allocating the costs and benefits in a fair fashion (see Section 2).

By Theorem 3.1, we know that the optimal cost is positive if $v^T\sigma > 0$ and 0 otherwise. In the second case, $a = 0$ is an obvious fair cost allocation. When only nonnegative allocations are considered, $a = 0$ is in fact the only feasible allocation and so is also the cost nucleolus.

Defining $x = \sigma - a = \sigma$, we get that (a, x) is a fair pair of allocations, i.e., satisfying efficiency, core, and justifiability.

For the case when $c(U) = v^T\sigma > 0$, a fair allocation — specifically, the cost nucleolus — can be computed in a closed form, as shown in Theorem 4.1 below. We remark that the theorem considers both nonnegative and unrestricted allocations simultaneously and provides a nonnegative allocation that is the nucleolus in both situations. For the statement of the theorem, note that e_1 is the first coordinate vector.

Theorem 4.1 *If, according to Theorem 3.1, $c(U) = v^T\sigma > 0$, then the nucleolus of the cost game is $(v^T\sigma)e_1$.*

An interpretation of the result is as follows. Because outlet 1 has a high standard deviation of demand compared to all other outlets (i.e., $v^T\sigma > 0$), the overall cost $c(U)$ is positive; if σ_1 were smaller (specifically if $v^T\sigma$ were nonpositive), then the overall cost would be zero. Hence, the positive cost can be attributed to outlet 1, and so in fairness, outlet 1 should absorb all costs. It is important to keep in mind that, although outlet 1 absorbs all the cost, it still benefits from the coalition in the amount $\sigma_2 + \cdots + \sigma_n$.

We prove Theorem 4.1 in several steps. First, we discuss the procedure of Schmeidler (1969) for determining the nucleolus. Recall that the nucleolus is calculated by solving a sequence of n linear programs (LPs), starting with (1). Generally, once the k-th LP has been solved, the $(k+1)$-st LP is constructed as follows. Let ε_k be the optimal value of the k-th LP, and let $\mathcal{P}^k$ denote the collection of proper subsets $S \notin \mathcal{P}^1 \cup \cdots \cup \mathcal{P}^{k-1}$ of $\{1, \ldots, n\}$ for which the inequality

$$c(S) - a(S) \geq \varepsilon$$

is active in all optimal solutions of the k-th LP. Then the $(k+1)$-st LP is

$$\begin{aligned} \max \quad & \varepsilon \\ \text{s.t.} \quad & c(S) - a(S) \geq \varepsilon \qquad \forall\ S \notin \mathcal{P}^1 \cup \cdots \cup \mathcal{P}^k \quad \emptyset \neq S \subsetneq U \\ & c(S) - a(S) = \varepsilon_\ell \qquad \forall\ S \in \mathcal{P}^\ell \quad \ell = 1, \ldots, k \\ & a(U) = c(U) \\ & [\, a \geq 0 \,]. \end{aligned} \tag{8}$$

It is proven by Schmeidler (1969) that the n-th LP is guaranteed to have a unique optimal solution, which is the nucleolus.

Our next step is to specialize the first LP (1) to the case of Theorem 4.1, i.e., when vv^T is the unique optimal solution of (2) and the optimal cost $c(U)$ equals $v^T\sigma > 0$. For this, the first important observation is that

$$c(S) = \left| \sum_{i \in S} \sigma_i v_i \right| = \begin{cases} \sigma_1 - \sigma(S \setminus 1) & \text{if } 1 \in S \\ \sigma(S) & \text{if } 1 \notin S. \end{cases}$$

Defining $\mathcal{S}_i$ to be the collection of all proper subsets S of $\{1, \ldots, n\}$ having $i \in S$ and $\mathcal{S}_i^c$ to be the collection of S such that $i \notin S$, (1) becomes

$$\begin{aligned} \max \;\; & \varepsilon \\ \text{s.t.} \;\; & \sigma_1 - \sigma(S \setminus 1) - a(S) \geq \varepsilon \quad && \forall \;\; S \in \mathcal{S}_1 \\ & \sigma(S) - a(S) \geq \varepsilon && \forall \;\; S \in \mathcal{S}_1^c \\ & a(U) = v^T \sigma \\ & [\, a \geq 0 \,]. \end{aligned} \tag{9}$$

Lemma 4.1 *The vector* $(a, \varepsilon) = \big((v^T\sigma)e_1, \sigma_n\big)$ *is an optimal solution of (9).*

Proof. We first show that the proposed vector (a, ε) is feasible. Clearly $a(U) = c(U)$ and $a \geq 0$. For convenience, let $[n]$ denote the set $\{1, \ldots, n\}$. Then, for $S \in \mathcal{S}_1$, we have

$$\begin{aligned} \sigma_1 - \sigma(S \setminus 1) - a(S) &= \sigma_1 - \sigma(S \setminus 1) - v^T\sigma \\ &= \sigma_1 - \sigma(S \setminus 1) - (\sigma_1 - \sigma([n] \setminus 1)) \\ &= \sigma([n] \setminus 1) - \sigma(S \setminus 1) \\ &= \sigma([n] \setminus S) \\ &\geq \sigma_n, \end{aligned}$$

where the last inequality follows because $[n]\setminus S$ is nonempty and σ_n is the smallest component in σ. Next, for $S \in \mathcal{S}_1^c$, we have

$$\sigma(S) - a(S) = \sigma(S) \geq \sigma_n,$$

which follows because S is nonempty and σ_n is smallest. Overall, we see that the proposed (a, ε) is feasible.

We can also show that any feasible (a, ε) has $\varepsilon \leq \sigma_n$, which will prove the result. Taking $S = \{n\}$ and considering the constraint $\sigma(S) - a(S) \geq \varepsilon$, we see $\sigma_n - a_n \geq \varepsilon$. In addition, taking $S = \{1, \ldots, n-1\}$ and considering the constraint $\sigma_1 - \sigma(S \setminus 1) - a(S) \geq \varepsilon$, we see that

$$\sigma_1 - (\sigma_2 + \cdots + \sigma_{n-1}) - (a_1 + \cdots + a_{n-1}) \geq \varepsilon.$$

Adding these two inequalities and using the fact that $a(U) = v^T\sigma = \sigma_1 - (\sigma_2 + \cdots + \sigma_{n-1})$, we conclude $2\sigma_n \geq 2\varepsilon$, as desired. ∎

In accordance with the procedure for calculating the nucleolus, our next step is to determine which inequalities are active in every optimal solution of (9). We do so using a duality argument. Introducing a dual variable $y_S \geq 0$ for each inequality constraint of (9) and a free dual variable λ for the equality constraint, the dual of (9) is

$$\begin{aligned} \min \quad & \sum_{S \in \mathcal{S}_1} [\sigma_1 - \sigma(S \setminus 1)]\, y_S + \sum_{S \in \mathcal{S}_1^c} \sigma(S)\, y_S - (v^T\sigma)\lambda \\ \text{s.t.} \quad & \lambda = [\, \leq \,] \sum_{S \in \mathcal{S}_j} y_S \qquad \forall \;\; j = 1, \ldots, n \\ & \sum_{\emptyset \neq S \subsetneq U} y_S = 1 \qquad\qquad y \geq 0. \end{aligned} \tag{10}$$

To facilitate our discussion, we also introduce the following notation: $N = \{i : \sigma_i = \sigma_n\}$.

Lemma 4.2 *Consider the optimal solution* $(a, \varepsilon) = ((v^T\sigma)e_1, \sigma_n)$ *of (9). An inequality corresponding to* S *is active if and only if* S *is of the form* $\{i\}$ *or* $[n] \setminus i$ *for some* $i \in N$*. As a result, in every optimal dual solution* (y, λ) *of (10), it holds that* $y_S = 0$ *for all other sets* S*.*

Proof. The proof of Lemma 4.1 shows that, with the optimal solution (a, ε), the inequalities of (9) reduce to

$$\begin{aligned} \sigma([n] \setminus S) &\geq \sigma_n \qquad \forall \ S \in \mathcal{S}_1 \\ \sigma(S) &\geq \sigma_n \qquad \forall \ S \in \mathcal{S}_1^c. \end{aligned}$$

From this, it is not difficult to see that equality is attained — and can only be attained — for sets S of the form $\{i\}$ or $[n] \setminus i$ for some $i \in N$. Complementary slackness proves the second part of the lemma. ∎

Proposition 4.1 *Regarding (9), the inequalities which are active in every optimal solution correspond precisely to those* S *of the form* $\{i\}$ *or* $[n] \setminus i$ *for some* $i \in N$*.*

Proof. We know from Lemma 4.2 that those inequalities corresponding to S, which are not of the specified form, are inactive in the specific optimal solution provided by Lemma 4.1. It remains to show that those inequalities, where S *is* of the specified form, are active in all optimal solutions of (9).

Consider the dual (10). Lemma 4.2 implies that, at optimality, the dual simplifies to

$$\begin{aligned} \min \quad & \sum_{i \in N} [\sigma_1 - \sigma([n] \setminus \{1, i\})] \, y_{[n]\setminus i} + \sum_{i \in N} \sigma_i \, y_i - (v^T\sigma)\lambda \\ \text{s.t.} \quad & \lambda = [\leq] \sum_{i \in N} y_{[n]\setminus i} \\ & \lambda = [\leq] \, y_j + \sum_{i \in N, i \neq j} y_{[n]\setminus i} \qquad \forall \ j \in N \\ & \sum_{i \in N} \left(y_{[n]\setminus i} + y_i \right) = 1 \qquad\qquad y \geq 0. \end{aligned}$$

Note that the objective function simplifies to

$$\sum_{i \in N} [\sigma_1 - (\sigma_2 + \cdots + \sigma_{n-1})] \, y_{[n]\setminus i} + \sum_{i \in N} \sigma_n \, y_i - (v^T\sigma)\lambda.$$

Next, setting $y_i = y_{[n]\setminus i} = 1/2|N|$ and $\lambda = 1/2$, it is straightforward to check that we have a feasible solution with objective value

$$\frac{1}{2} [\sigma_1 - (\sigma_2 + \cdots + \sigma_{n-1})] + \frac{1}{2} \sigma_n - \frac{1}{2} v^T\sigma = \sigma_n.$$

Hence, the constructed solution is optimal for the dual, and since $y_i = y_{[n]\setminus i} > 0$, complementary slackness implies that corresponding inequalities in (9) are active in all optimal solutions of (9). This proves the desired result. ■

Proposition 4.1 allows us to construct the second nucleolus LP:

$$
\begin{aligned}
\max\ & \varepsilon && && (11)\\
\text{s. t.}\ & \sigma_1 - \sigma(S \setminus 1) - a(S) \geq \varepsilon && \forall\ \ S \in \mathcal{S}_1 \setminus \{[n] \setminus i : i \in N\}\\
& \sigma(S) - a(S) \geq \varepsilon && \forall\ \ S \in \mathcal{S}_1^c \setminus \{\{i\} : i \in N\}\\
& \sigma_1 - \sigma(S \setminus 1) - a(S) = \sigma_n && \forall\ \ S \in \{[n] \setminus i : i \in N\}\\
& \sigma(S) - a(S) = \sigma_n && \forall\ \ S \in \{\{i\} : i \in N\}\\
& a(U) = v^T \sigma\\
& [\, a \geq 0\,].
\end{aligned}
$$

Note that the third and fourth constraints simplify to

$$
\begin{aligned}
&\sigma_1 - (\sigma_2 + \cdots + \sigma_{n-1}) - (a_1 + \cdots + a_{i-1} + a_{i+1} + \cdots a_n) = \sigma_n\\
&\sigma_n - a_i = \sigma_n.
\end{aligned}
$$

Here, the second equality implies $a_i = 0$ so that the first is implied by $a(U) = v^T\sigma$. As a result, (11) simplifies to

$$
\begin{aligned}
\max\ & \varepsilon && && (12)\\
\text{s. t.}\ & \sigma_1 - \sigma(S \setminus 1) - a(S) \geq \varepsilon && \forall\ \ S \in \mathcal{S}_1 \setminus \{[n] \setminus i : i \in N\}\\
& \sigma(S) - a(S) \geq \varepsilon && \forall\ \ S \in \mathcal{S}_1^c \setminus \{\{i\} : i \in N\}\\
& a(U) = v^T \sigma\\
& a_i = 0 \quad \forall\ \ i \in N\\
& [\, a \geq 0\,].
\end{aligned}
$$

It is possible to simplify (12) further. Let $i \in N$, and consider $S \in \mathcal{S}_1 \setminus \{[n] \setminus i : i \in N\}$ such that $i \in S$. Using that $a_i = 0$, the corresponding inequality constraint of (12) can be expressed as

$$\sigma_1 - \sigma(S \setminus \{1, i\}) - a(S \setminus i) \geq \varepsilon + \sigma_n,$$

which implies the constraint corresponding to the set $S \setminus i$. So, without loss of generality, we may eliminate all constraints corresponding to $S \in \mathcal{S}_1 \setminus \{[n] \setminus i : i \in N\}$ for which $i \notin S$. In a similar fashion, the constraint corresponding to any $S \in \mathcal{S}_1^c \setminus \{\{i\} : i \in N\}$, for which $i \in S$, simplifies to

$$\sigma(S \setminus i) - a(S \setminus i) \geq \varepsilon - \sigma_n,$$

which is implied by the constraint corresponding to the set $S \setminus i$. We thus may eliminate all constraints corresponding to $S \in \mathcal{S}_1^c \setminus \{\{i\} : i \in N\}$ for which $i \in S$.

Putting the two above observations together and defining

$$
\begin{aligned}
\mathcal{S}_{1N} &= \mathcal{S}_1 \cap \bigcap_{i \in N} \mathcal{S}_i\\
\mathcal{S}_{1N}^c &= \mathcal{S}_1^c \cap \bigcap_{i \in N} \mathcal{S}_i^c
\end{aligned}
$$

we see that (12) is equivalent to

$$\begin{aligned}
\max\ & \varepsilon \\
\text{s.t.}\ & \sigma_1 - \sigma(S \setminus 1) - a(S) \geq \varepsilon \qquad \forall\ S \in \mathcal{S}_{1N} \\
& \sigma(S) - a(S) \geq \varepsilon \qquad \forall\ S \in \mathcal{S}_{1N}^c \\
& a(U) = v^T \sigma \\
& a_i = 0 \quad \forall\ i \in N \\
& [\, a \geq 0\,].
\end{aligned} \tag{13}$$

It is now not difficult to see that (13) is equivalent to the first nucleolus LP (9) for the problem given by

$$\bar{\sigma} = \left(\sigma_1 - \sum_{i \in N} \sigma_i, \sigma_2, \ldots, \sigma_{n-|N|} \right)^T \in \Re^{n-|N|}.$$

A simple inductive argument can thus be applied to show that the nucleolus of the cost game with respect to σ has $a_2 = \cdots = a_n = 0$, which immediately implies Theorem 4.1.

5 Computing Fair Allocations In General

In the previous section, we exploited specific properties of the optimal correlation matrix — as well as characteristics of the cost function $c(S)$ for our inventory centralization — to establish a fair allocation of costs and benefits once an optimal centralization had been achieved. In this section, we present an idea for computing a fair allocation without assuming an optimal centralization. Our motivation is that, in real-life situations, it may be challenging to implement optimal correlations, or the coalition may only be able to achieve near-optimal correlations, and yet a fair allocation of costs is still necessary.

In the discussion that follows, we attempt to handle both unrestricted and nonnegative allocations simultaneously. Specific properties of the two cases will be highlighted as necessary.

Given a correlation matrix R relating outlets demands, whether a nonnegative fair allocation exists can be determined by solving (1); existence is ensured if and only if the optimal value is nonnegative. As before, note that dropping the nonnegativity requirements for allocation guarantees the existence of a fair allocation (Muller et al. (2002)). In either case, one may continue with the procedure outlined in Section 4 for calculating the nucleolus if a fair allocation exists — or for calculating a "pseudo-nucleolus," if a fair allocation does not exist. Such a pseudo-nucleolus may still be interpreted as an allocation that is as fair as possible in that it "makes the least-well-off coalition as well-off as possible."

However, optimizing (1) and its associated sequence of LPs is theoretically difficult since the number of constraints in each LP is roughly 2^n, i.e., the LPs have exponential size. Practically, the optimization may be reasonable for small values of n, but the solution time will grow quickly as n becomes larger. In addition, the structures of the LPs do not seem to be amenable to reformulation (or other manipulation) that could reduce their size.

Thus, it is of theoretical and practical interest to develop a technique for approximating (1), which will also provide an approximately fair allocation. (Here, we use the notion of "approximation" loosely. The ideas that we present are more similar to heuristic techniques rather than approximation algorithms, which have guarantees of solution quality.) In what follows, we exploit the characteristics of the cost function $c(S)$ to calculate approximately a fair allocation via a polynomial-sized second-order program.

5.1 Background on second-order cone programming

Second-order cone programming (SOCP) is an extension of linear programming that has received considerable attention in the past few years, where research has focused on both applications and algorithms. In its simplest form, the standard form SOCP is

$$\min \left\{ c^T x : Ax = b,\ x \in K_d \right\}, \tag{14}$$

where A, b, c are the problem data, x is the vector variable, and

$$K_d = \left\{ v \in \Re^d : v_1 \geq \|(v_2, \ldots, v_d)\| \right\}$$

is the second-order cone. It is not difficult to show that K_d is a convex set, so that (14) is a convex optimization problem. The dual of (14) is

$$\max \left\{ b^T y : c - A^T y \in K_d \right\}. \tag{15}$$

When presented with an SOCP, which is not in standard form, one can apply standard techniques for constructing the dual. It is well-known that strong duality holds between (14) and (15) if both problems have feasible interior points. The interior of K_d is defined to be those $v \in K_d$ such that $v_1 > \|(v_2, \ldots, v_d)\|$.

The constraint $x \in K_d$ differentiates (14) from standard form linear programming. In fact, one can think of $x \in K_d$ as a new notion of nonnegativity, which differs from the standard notion $x \geq 0$. Nevertheless, many of the algorithms for LP (specifically those based on interior-point methods) can be extended to solve (14) and (15) to any desired accuracy in polynomial time. Practically speaking, the algorithms are quite effective, and several high-quality software packages exist for solving second-order programs.

We refer the reader to the survey paper by Lobo et al. (1998) for further background on second-order cone programming.

5.2 The nucleolus LP using second-order cones

Let R be the correlation matrix relating the demands of the n outlets, and let $V \in \Re^{n \times r}$ satisfy $R = VV^T$. Here, r is the rank of R.

For our presentation, we will find it convenient to encode each proper subset S of $\{1, \ldots, n\}$ as a vector $z \in \{0,1\}^n$ such that $z_i = 1$ if and only if $i \in S$. Then the entire collection of proper subsets can be expressed as

$$D = \{z \in \{0,1\}^n : 1 \leq e^T z \leq n-1\},$$

where e is the all-ones vector. The centralization cost incurred by a specific S can be expressed as

$$c(S) = \left\| \sum_{i \in S} \sigma_i V_{i\cdot} \right\| = \left\| z^T \mathrm{Diag}\,(\sigma) V \right\| = \left\| V^T \mathrm{Diag}\,(\sigma) z \right\|.$$

As a result, the first nucleolus LP (1) can be written as

$$\begin{aligned} \max \quad & \min\left\{\|V^T \mathrm{Diag}\,(\sigma) z\| - z^T a : z \in D\right\} \\ \text{s. t.} \quad & e^T a = \|\sigma^T V\| \\ & [\, a \geq 0\,], \end{aligned} \tag{16}$$

for which the max-min structure is clearly the distinguishing feature. A trivial rewriting of the "min" portion is

$$\min\left\{\zeta - z^T a : z \in D,\ \zeta \geq \|V^T \mathrm{Diag}\,(\sigma) z\|\right\}, \tag{17}$$

where z and ζ are the variables of the optimization and a is considered fixed. Since ζ and z are only linked by the constraint $\zeta \geq \|V^T \mathrm{Diag}\,(\sigma) z\|$ and since we are minimizing, it is clear that this constraint will be active in any optimal solution. Moreover, the connection with second-order cones is clear since

$$\zeta \geq \|V^T \mathrm{Diag}\,(\sigma) z\| \qquad \Longleftrightarrow \qquad \big(\zeta, V^T \mathrm{Diag}\,(\sigma) z\big) \in K_{1+r}.$$

Overall, we have

$$\begin{aligned} \max \quad & \min\left\{\zeta - z^T a : z \in D,\ \big(\zeta, V^T \mathrm{Diag}\,(\sigma) z\big) \in K_{1+r}\right\} \\ \text{s. t.} \quad & e^T a = \|\sigma^T V\| \\ & [\, a \geq 0\,]. \end{aligned} \tag{18}$$

5.3 The nucleolus LP tightened as a second-order cone program

Even though (18) is an exact rewriting of (1) using second-order cones, it is still difficult to optimize, in part because D is a discrete set of exponential size. We suggest replacing D in (18) with its convex hull

$$\mathrm{conv}(D) = \{0 \leq z \leq e : 1 \leq e^T z \leq n-1\}.$$

The effect is a tightening, in the sense that the optimal value of the new problem is a lower bound on the optimal value of (18). Using duality, we now argue that this new problem can be solved explicitly as the second-order cone program (21) below.

To see this, note that after replacing D with its convex hull, the new "min" portion of the max-min structure is

$$\begin{aligned} \min \quad & \zeta - z^T a \\ \text{s. t.} \quad & 0 \leq z \leq e \quad 1 \leq e^T z \leq n-1 \\ & \big(\zeta, V^T \mathrm{Diag}\,(\sigma) z\big) \in K_{1+r}, \end{aligned} \tag{19}$$

which is a polynomial-sized second-order cone program. It is not difficult to verify that the dual of this SOCP is

$$\begin{aligned} \max \quad & \alpha - (n-1)\beta - e^T q \\ \text{s. t.} \quad & \mathrm{Diag}\,(\sigma) V p + q + (\beta - \alpha) e \geq a \\ & \alpha, \beta \geq 0 \quad q \geq 0 \quad (1, p) \in K_{1+r}. \end{aligned} \tag{20}$$

Moreover, both primal and dual have interior points so that strong duality holds. Hence, replacing (19) with (20), we have a max-max structure that can be written explicitly as

$$
\begin{aligned}
\max \quad & \alpha - (n-1)\beta - e^T q \\
\text{s.t.} \quad & \operatorname{Diag}(\sigma) V p + q + (\beta - \alpha) e \geq a \\
& \alpha, \beta \geq 0 \quad q \geq 0 \quad (1, p) \in K_{1+r} \\
& e^T a = \|\sigma^T V\| \\
& [\, a \geq 0 \,].
\end{aligned} \tag{21}
$$

The second-order cone program (21) specifies a cost allocation that heuristically approximates the solution of (1) and its associated LPs, and although it seems difficult to establish a theoretical guarantee of the quality of the approximation, we have found on randomly generated instances that the approximation is quite reasonable. Of course, (21) has the advantage of polynomial size and so can be solved easily for large coalition sizes n.

Table 1 gives the optimal values and timings for calculating the nucleolus (or pseudo-nucleolus, as the case may be) and the SOCP (21) on six randomly generated instances of the inventory centralization problem for various values of n. The instances where generated as follows: a random matrix $\hat{V} \in \Re^{n \times n}$ with entries uniform in $[-1, 1]$ was generated; then the rows of $\hat{V}$ were scaled to norm 1, yielding $V \in \Re^{n \times n}$; and finally the test correlation matrix R was formed as VV^T. The first six results correspond to unrestricted allocations, while the second six correspond to nonnegative allocations. For the LPs, the optimal value of (1) is reported since it is directly comparable with the optimal value of (21). All experiments were done on a Pentium 4 2.4 GHz with CPLEX 9.0 (ILOG, Inc., 2003) as the LP solver and SeDuMi 1.05 (Sturm, 1999) as the SOCP solver.

As predicted, in each case the optimal value of (1) is larger, and the optimization of (21) takes considerably less time, especially for larger values of n. Moreover, it is interesting to note that, for the instances with unrestricted allocations, the SOCP delivered an optimal value of 0, which is a certificate that the allocation a provide by (21) is a fair allocation (though not the nucleolus).

It is also interesting to examine the allocations produced by the two algorithms. For $n = 8$, the unrestricted allocations are

$$
a_{\text{LP}} = \begin{pmatrix} 0.0363 \\ -0.0600 \\ 0.2800 \\ 0.0036 \\ 0.2570 \\ -0.0388 \\ 0.6532 \\ 0.0322 \end{pmatrix} \qquad a_{\text{SOCP}} = \begin{pmatrix} 0.0975 \\ -0.0580 \\ 0.2555 \\ -0.0718 \\ 0.2466 \\ -0.0125 \\ 0.6944 \\ 0.0117 \end{pmatrix},
$$

Table 1: Optimal values and timings (in seconds) for calculating the nucleolus or pseudo-nucleolus (via the sequence of LPs starting from (1)) and the SOCP (21) on six randomly generated instances of the inventory centralization problem. The first six results correspond to unrestricted allocations, while the second six results correspond to nonnegative allocations.

n	Opt Val		Time (s)	
	LP	SOCP	LP	SOCP
2	0.0297	0.0000	0.80	0.86
4	0.0029	0.0000	3.65	0.85
6	0.0005	0.0000	9.05	0.90
8	0.0263	0.0000	31.29	0.87
10	0.0023	0.0000	138.55	0.92
12	0.0000	0.0000	722.18	0.92
2	0.0297	0.0000	0.82	1.00
4	-0.0075	-0.0436	3.72	0.94
6	0.0005	-0.0156	9.13	0.99
8	0.0263	-0.0848	31.42	0.98
10	-0.2665	-0.4099	141.83	0.97
12	-0.3073	-0.3824	720.06	0.96

and the nonnegative allocations are

$$a_{\text{LP}} = \begin{pmatrix} \\ 0.2174 \\ 0.0505 \\ 0.2570 \\ \\ 0.6064 \\ 0.0322 \end{pmatrix} \qquad a_{\text{SOCP}} = \begin{pmatrix} \\ 0.2828 \\ \\ 0.2197 \\ \\ 0.6610 \\ \\ \end{pmatrix}.$$

For $n = 12$, we have

$$a_{\text{LP}} = \begin{pmatrix} 0.0066 \\ 0.4995 \\ 0.0517 \\ 0.1705 \\ 0.3370 \\ -0.0628 \\ 0.1450 \\ 0.1360 \\ 0.3850 \\ 0.6560 \\ -0.4648 \\ 0.0908 \end{pmatrix} \qquad a_{\text{SOCP}} = \begin{pmatrix} 0.0066 \\ 0.4951 \\ 0.0621 \\ 0.1790 \\ 0.3403 \\ -0.0662 \\ 0.1442 \\ 0.1328 \\ 0.3849 \\ 0.6432 \\ -0.4646 \\ 0.0933 \end{pmatrix}$$

and

$$a_{\text{LP}} = \begin{pmatrix} 0.0079 \\ 0.3805 \\ 0.0391 \\ 0.1842 \\ 0.3356 \\ \\ 0.1055 \\ \\ 0.3282 \\ 0.5696 \\ \\ \\ \end{pmatrix} \qquad a_{\text{SOCP}} = \begin{pmatrix} 0.0075 \\ 0.4100 \\ \\ 0.1223 \\ 0.2945 \\ \\ 0.1221 \\ \\ 0.3685 \\ 0.6017 \\ \\ 0.0240 \end{pmatrix}.$$

(Here, an entry has been left blank if it is equal to 0 in the first four decimal places.) Although we have not been able to derive a theoretical relationship between a_{LP} and a_{SOCP}, we believe that the approximation of a_{LP} by a_{SOCP} is surprisingly close. Certainly the order of magnitude of the entries matches nicely.

Overall, we propose the SOCP (21) as a viable alternative to calculating the nucleolus or pseudo-nucleolus directly. While the size of the associated LPs makes the calculation slow for moderate values of n, the solution of (21) scales much better. For example, a randomly generated instance having $n = 150$ required approximately 3 seconds for the solution of the SOCP. Moreover, for small cases that we can verify directly, the allocation produced by the SOCP approximates the nucleolus or pseudo-nucleolus quite well.

6 Summary

The main purpose of this paper is to examine the cost optimization problem implied by newsvendor centralization arrangement when individual outlets experience normally distributed demand and their inventory cost parameters are identical. In addition, it is paramount that the cost of the centralized arrangement be fairly assessed to the individual outlets. The cost optimization problem has the form of a semidefinite program of Goemans and Williamson (1995) – optimization over correlation matrices. Tracing back to statistical literature (Shapiro, 1982) allows us to obtain a partial classification of the optimal solutions to our SDP (Theorem 3.1). We complete the classification of the optimal solutions by constructing the optimal correlation matrix for any balanced partition relative to σ (Theorem 3.2).

This leads to the problem of calculating a fair cost allocation, expressed here as the calculation of the nucleolus for the corresponding cooperative game. In principle, it is a tedious process of solving n linear programs with an exponential number of constraints. However, for the optimal solution of the corresponding SDP we prove that this can be done easily — computed in a closed form — in our case (Theorem 4.1). Unfortunately, this does not imply easy computation of a fair cost allocation given a general correlation relation between the demands of the outlets. In practice one cannot assure that the manipulation of outlets' demands can be always lined up with that of the optimal correlation solution. Thus, the issue of calculating a fair allocation in the more general newsvendor centralization case is addressed by approximating the nucleolus computation with a second-order cone program. It results in a heuristic solution obtained in a fraction of time with encouraging empirical results “close” to the nucleolus solution. This fairs well for larger problems.

From the current study, there are some interesting open questions and possible extensions. For example, newsvendor centralization problems with identical cost parameters are known to have a nonempty core, while the correlation optimization problem examined in this paper is restricted to normally distributed demands. For what other interesting demand distributions can analysis like that of the paper be extended? In addition, this paper has provided the solution of the “unconstrained” correlation optimization, i.e., where we are allowed to shift correlations arbitrarily so as to minimize cost. In reality, the ability to alter correlations could be limited in some manner, which would constrain the SDP (2). It would be interesting to investigate such situations, for which a general purpose SDP solver is likely to be required.

Acknowledgement

The authors are in debt to Alexander Shapiro for pointing out Theorem 3.1.

References

G.D. Eppen. Effects of centralization on expected costs in a multi-location newsboy problem. *Management Science*, 25(5):298–501, 1979.

M.X. Goemans and D.P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of ACM*, 42: 1115–1145, 1995.

B.C. Hartman and M. Dror. Cost allocation in continuous review inventory models. *Naval Research Logistics*, 43:549–561, 1996.

B.C. Hartman and M. Dror. Optimizing centralized inventory operations in a cooperative game theory setting. *IIE Transactions on Operations Engineering*, 35:243–257, 2003.

B.C. Hartman, M. Dror, and M. Shaked. Cores of inventory centralization games. *Games Econom. Behav.*, 31(1):26–49, 2000.

ILOG, Inc. *ILOG CPLEX 9.0, User Manual*, 2003.

M. Lobo, L. Vandenberghe, S. Boyd, and H. Lebret. Applications of second-order cone programming. *Linear Algebra and its Applications*, 284:193–228, 1998.

A. Muller, M. Scarsini, and M. Shaked. The newsvendor game has a nonempty core. *Games Econom. Behav.*, 38:118–126, 2002.

D. Schmeidler. The nucleolus of a characteristic function game. *SIAM J. Appl. Math.*, 17: 1163–1170, 1969.

A. Shapiro. Weighted minimum trace factor analysis. *Psychometrika*, 47:243–264, 1982.

Jos F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optim. Methods Softw.*, 11/12(1-4):625–653, 1999. Interior point methods.

H. Wolkowicz, R. Saigal, and L. Vandenberghe, editors. *Handbook of Semidefinite Programming*. Kluwer Academic Publishers, 2000.

H. Peyton Young. *Cost Allocation: Methods, Principles, Applications*, chapter 1, pages 3–29. Elsevier Science Publishers B.V., North-Holland, 1985.