

RISK POOLING AND COST ALLOCATION IN SPARE PARTS INVENTORIES

Mario Guajardo

Department of Finance and Management Science, NHH Norwegian School of Economics
Helleveien 30, N-5045 Bergen, Norway
mario.guajardo@nhh.no

Mikael Rönnqvist

Department of Finance and Management Science, NHH Norwegian School of Economics
Helleveien 30, N-5045 Bergen, Norway
mikael.ronnqvist@nhh.no

ABSTRACT

We present a problem on risk pooling and cost allocation in inventory of spare parts motivated in the oil and gas sector. Spare parts are held to facilitate operational conditions. Multiple inventory plants control their stock per separate. In practice, however, they supply each other when needed. Risk pooling is a source of savings in this situation. Under a centralized inventory solution, the plants should agree on how to share costs for taking advantage of the benefits of pooling. We report computations where the centralized inventory achieves significant savings over the decentralized solution, considering ordering and holding costs subject to a service level constraint. We then test five cost allocation methods and analyse their outcomes, using game theory principles. We show examples where a method results in stable allocations when the plants use the same target service levels, but not when at least one plant use a different target.

KEYWORDS. Inventory of spare parts. Risk pooling. Cost allocation.

P&G - PO na Área de Petróleo & Gás

L&T - Logística e Transportes

1. Introduction

Risk pooling has been traditionally used as a form of protecting against demand variability in inventory management. The principle is that by aggregating demand across different locations, high demand from one customer is likely to be offset by low demand from another. This reduction in variability allows a decrease in safety stock and average inventory. While risk pooling may result in lower total inventory costs comparing to the case when the locations act per separate, it is not straightforward to answer how the participants in the pool should share costs in order to achieve these savings. The difficulty arises, for example, because the locations can be exposed to different demand behaviours and, therefore, some of them may in practice perceive more benefits from the risk pooling situation than others.

The cost allocation problem in multi-location inventory systems has received attention during quite a few years (e.g. Gerchak and Gupta, 1991; Robinson, 1993; Hartman and Dror, 1996, 2000). However, in the specific context of spare parts inventories, the research in cost allocation is more sparse. The first related work was published by Wong et al. (2007), who use game theory principles to analyse cost allocation rules in spare parts inventories. Kilpi et al. (2009) discuss cooperative strategies for risk pooling of spare parts for aircrafts. More recently, Kartsen et al. (2009, 2011a, 2011b) and Karsten and Basten (2012) have studied a series of related problems.

Spare parts are relevant to facilitate operational conditions in contexts where equipments are subject to maintenance and failure. The consumption of spare parts is triggered by events subject to variability and, therefore, holding inventory to fulfil the demand requirements may be crucial. A low stock of spare parts could imply stockout and, in consequence, production downtime until the part is repaired or replenished. On the other hand, a high stock of spare parts could imply an important binding cost. There is a huge body of literature dealing with spare parts, as reviewed by Kennedy et al. (2002). An important part of this literature has been motivated in air force and aircraft contexts (e.g. Muckstadt, 2005; Sherbrooke, 2004). Applications in other contexts include the service support of IBM in the US (Cohen et al. 1990), a chemical plant in Belgium (Vereecke and Verstraeten, 1994), a white goods manufacturer in Italy (Kalchschmidt et al. 2003) and a distributor of castors and wheels in Greece (Nenes et al. 2010). Our research is motivated in the oil and gas sector, where we have found only one recent article on spare parts inventories, published by Porras and Dekker (2008). Characterized by high service levels (because of safety and production factors), customized equipment specifications with long lead times, and facility networks spread onshore and offshore, we believe the problem of inventory management of spare parts in this industry deserves special attention from the research community.

Together with the interest on inventories of spare parts, there is an extensive body of literature dealing with lateral transshipment and risk pooling, as reviewed by Paterson et al. (2011).

Although the so stated relevance of spare parts inventories and risk pooling, the literature lacks of evidence on how risk pooling in spare part inventories should be implemented and how the cost should be shared among the different players, in order to take advantage of the benefits of pooling. In this article, we describe a real-world situation where risk pooling may represent an important source of savings and we discuss the suitability of five cost allocation methods. We report numerical examples where these cost allocations methods can provide stable allocations in situations where the different players use the same target service levels. The same methods in the same numerical examples, however, may fail on that purpose when at least one of the players has a target service level different than the other players.

2. Background

Our motivation comes from a real-world problem faced at an energy company producer of oil and gas. Headquartered in Norway and with presence in 42 countries, this company is one of the world's largest net sellers of crude oil and the second largest exporter of gas to Europe. It acts as operator of several offshore platforms with different ownership structure and it holds inventory of about 200,000 spare part items in a number of locations spread in the Scandinavian region. Some of these parts are highly critical to assure safety and production. At the same time, the store of spare parts means an important binding cost. The company has operational responsibility for seven warehouses located along the Norwegian coast, which serve offshore installations. Within this structure of warehouses, there are 24 inventory plants. The majority of these plants is owned by the company under study, but there is also a set of other companies owning a share of them. Each of the 24 plants holds its own inventory of spare parts and they are managed per separate. Each of the plants determines its own inventory control parameters separately, and is linked to one productive location (and vice-versa). When a spare part is required, it is provided from the corresponding plant. However, when a productive location requires certain spare part in such amount that it is not available at the corresponding plant, the requirement can be fulfilled (if agreed) from other plant that has availability on stock. This is especially realized when the required part is highly critical. This type of fulfilment occurs based on an informal agreement between inventory plants, but it is not considered when deciding the control parameters. It is, therefore, relevant to study the potential of savings that a risk pooling approach would generate if there would be a centralized decision on the inventory control parameters instead of the separate planning that the plants carry out today.

The importance of lateral resupply and pooling is illustrated in Muckstadt (2005), who gives examples of systems in which pooling exists that require roughly a third of the safety stock required when operating a completely decentralized system. A recent application by Kranenburg and van Houtum (2009) reports that under full pooling more than 50% of the no pooling cost can be saved in in the case of ASML, a Dutch equipment manufacturer. Before their implementation, ASML did not take lateral transshipment into account in the planning phase, i.e., the inventory in each local warehouse was planned separately. Nevertheless, in daily practice lateral transshipment was used, the same practice as realized at the company in our case. A similar situation is presented by Kukreja and Schmidt (2005), who deal with a large utility company having 29 generating plants. The plants do not consider the effect of being supplied from each other when deciding their own inventory control policies, but in practice they also collaborate. Experimentation by the authors including from two to five plants show that pooling lead to savings between 31% and 58.78%. Other works in lateral transshipments under emergency situations are reported in Archibald (2007), Wong et al. (2006) and Alfredsson and Verrijdt (1999).

A more detailed description about the context in our case can be found in Guajardo et al. (2012). Note that, while risk pooling can conduce to great savings, it is not clear how the different shareholders at the plants should share the costs. Exploring this problem is the main purpose of this paper. In the following section, we will define our problem settings and a procedure to find the optimal policy.

3. Problem setting and optimal base-stock policy

We consider a single-item problem under continuous review ($S-1, S$) or *base-stock policy*. In this policy, every demand event originates an order for an amount such that the base-stock level S is reached. For this reason, it is also called a *one-for-one replenishment policy*. The ($S-1, S$) policy is commonly assumed in inventory of spare part problems, where the items are usually expensive and slow-moving (Muckstadt, 2005; Sherbrooke, 2004). We

assume fixed lead time and demand per unit of time following a Poisson distribution. This distribution is also commonly assumed in spare parts inventory modelling.

There are traditionally three type of costs involved in inventory problems: order cost, holding cost and downtime cost. The order cost is a fixed cost per order, the holding cost is a variable cost of carrying inventory and the downtime cost is the cost incurred when a demand order is not satisfied because of a stockout. Finding the optimal inventory control parameter S could be pursued by minimizing the whole cost function considering these three cost elements. However, in practice, the downtime cost is usually hard to estimate. A popular strategy is to minimize carrying plus ordering costs, subject to a service level constraint. The service level criterion is generally easier to state and interpret by practitioners. A number of references point out this fact, such as Chen and Krass (2001), Bashyam and Fu (1998) and Cohen et al. (1989).

As service level measure, we utilize the *fill rate*, defined as the expected fraction of demand satisfied directly from stock on hand, which is usually utilized in spare parts inventory problems. Hence, we utilize a target service level β as lower bound that must be satisfied by the fill rate achieved for the control parameter S . We refer the reader to Silver et al. (1998) for more background on inventory costs and service levels, and for the following definitions and formulae.

Let A be the cost for an order and λ the average demand per unit of time. We compute the expected order cost as $C_{order} = A/\lambda$.

Let r be the carrying charge, v the unit value of the item and \bar{I} the average inventory on hand. We compute the expected holding cost as $C_{hold} = \bar{I}vr$. We calculate \bar{I} as $\kappa + 1/2$, where κ is the safety stock calculated as the reorder point $S-1$ minus the expected demand during the lead time λ_L . Let X_L be the demand during the lead time. Since the demand per unit of time follows a Poisson distribution with mean λ , then X_L follows a Poisson distribution with mean $\lambda_L = \lambda L$, where L is the lead time.

The total expected cost per item is $C = C_{order} + C_{hold} = A/\lambda + \bar{I}vr = A/\lambda + (S - \lambda_L - 1/2)vr$.

For a given base-stock level S , the fill rate β_S in our setting can be calculated as the probability of stockout $1 - P(X_L \leq S-1)$. Thus, $\beta_S = 1 - \sum_{k \leq S-1} \frac{(\lambda L)^k e^{-\lambda L}}{k!}$.

The optimization problem consists of finding the control parameter S that solves the problem below.

$$\begin{aligned} & \text{Min } C(S) \\ \text{s.t.} & \quad \beta_S \geq \beta \end{aligned}$$

Since the cost increases in S (in fact, note that C can be written as $C(S) = K + Svr$, where K does not depend on S), the problem reduces to find the minimum S such that the service level constraint $\beta_S \geq \beta$ is fulfilled. A simple and usual approach to find such optimal base-stock level S is to evaluate the fill rate β_S starting from $S = 1$ and consecutively increasing S by one, until the target β is fulfilled. Hence, if the service level achieved by $S = 1$ is less than β , then $S = 2$ is tried and so on, until the lowest S satisfying the constraint is found. For computational purposes, an upper bound on S could be predefined, so that the evaluation procedure is assured to finish. However, in practice the procedure will finish quickly, because in spare part problems usually a low value of S will be optimal.

4. Cost allocation

We consider a game consisting of a set N of all players (inventory plants), who decide whether to plan their inventory separately or to plan their inventory in a coalition with other players based on a risk pooling strategy. If a player stays alone, it will set its base-stock level S and run its inventory independent from the other players. If players act in

a coalition, they will set a base-stock level S under the same objective function and run a single inventory serving their whole demand. This last case corresponds to a pure cooperative strategy which in operational terms translates into a centralized inventory system. For a coalition $M \subseteq N$, we refer by C_M to the optimal expected cost if all players in coalition M would implement risk pooling. N is the *grand coalition* and C_N is its optimal expected cost, as if there would be a centralized inventory plant planning the inventory to satisfy demand from all the players. For a given cost allocation rule, we refer by u_j to the cost allocated to player j .

4.1 The core of the game

A cost allocation vector $u=(u_1, u_2, \dots, u_n)$ is said to be in the *core* of the game (Gillies, 1959) if it satisfies constraints (1)-(3) below.

$$u_j \leq C_{\{j\}} \quad \forall j \in N \quad (1)$$

$$\sum_{j \in M} u_j \leq C_M \quad \forall M \subset N \quad (2)$$

$$\sum_{j \in N} u_j = C_N \quad (3)$$

Constraint (1) corresponds to the *individually rational* condition, which states that the cost allocated to each player j must not be greater than its stand-alone cost. Constraint (2) corresponds to a *stability* condition, which states that there is no subset of players such that if they would form a coalition separate from the rest they would perceive less cost than the allocation u . Constraint (3) corresponds to the *efficiency* condition, which states that the sum of the costs allocated to all the players equals the optimal cost of the grand coalition, and thus it takes full advantage of risk pooling. The core of the game is the set of all vectors u satisfying constraints (1)-(3). In other words, a cost allocation vector in the core assures that the savings of risk pooling are achieved and makes all players to stay in the grand coalition, without incentives for a player to stay alone or within a smaller coalition.

4.2 Allocation methods

We will use the five cost allocation methods described below.

- **Cost allocation method 1: Egalitarian**

This method simply assigns equal cost shares to all the players, as follows:

$$u_j = \frac{C_N}{|N|} \quad \forall j \in N$$

- **Cost allocation method 2: Proportional to demand (or α Demand)**

For a player j , whose demand rate is λ_j , this method assigns a cost share proportional to the weight that this demand rate represents over the total demand rate of the pool.

$$u_j = \left(\frac{\lambda_j}{\sum_{i \in N} \lambda_i} \right) \cdot C_N \quad \forall j \in N$$

- **Cost allocation method 3: Altruistic**

For a player j , whose stand-alone cost is $C_{\{j\}}$, this method assigns a cost share proportional to the cost that $C_{\{j\}}$ represents over the total cost of the players if they all would act alone.

$$u_j = \left(\frac{C_{\{j\}}}{\sum_{i \in N} C_{\{i\}}} \right) \cdot C_N \quad \forall j \in N$$

- **Cost allocation method 4: Shapley value**

Shapley (1953) introduced an allocation based on the marginal costs in entering coalitions. The marginal cost of player j entering coalition M is $C_M - C_{M \setminus \{j\}}$. The Shapley cost allocation method assigns player j the average of the marginal costs he implies when entering all coalitions he can enter. The method considers all sizes of a coalition equally likely, thus the size $|M|$ is assigned a probability $1/|N|$. Consider player j and the $|N|-1$ remaining players. For a given size $|M|$, there are $\binom{|N|-1}{|M|-1}$ ways of choosing $|M|-1$ players in coalition M from the $|N|-1$ remaining players. Therefore, the probability of a particular coalition M is $\frac{1}{|N| \binom{|N|-1}{|M|-1}}$.

Regrouping terms and considering all possible subsets of N , this method allocates the costs as follows:

$$u_j = \sum_{M \subseteq N: j \in M} \left(\frac{(|N| - |M|)! (|M| - 1)!}{|N|!} \right) \cdot (C_M - C_{M \setminus \{j\}}) \quad \forall j \in N$$

- **Cost allocation method 5: Equal Profit Method (EPM)**

EPM looks for a stable allocation such that the relative savings are as similar as possible for all participants. Such allocation u is solution of the following linear programming model:

$$\begin{aligned} & \min f \\ \text{s.t.} & \\ & f \geq \frac{u_i}{C_{\{i\}}} - \frac{u_j}{C_{\{j\}}} \quad \forall i, j \in N \\ & \sum_{j \in M} u_j \leq C_M \quad \forall M \subset N \\ & \sum_{j \in N} u_j = C_N \\ & u_j \geq 0 \quad \forall j \in N; f \in \mathbb{R}. \end{aligned}$$

The egalitarian method is referred by Tijs and Driessen (1986) as the simplest cost allocation method. The α Demand and the Shapley methods correspond to what Wong et al. (2007) refer to as *cost allocation policies 3* and *4*, respectively. The altruistic method has been used by Audy et al. (2012) in a forestry problem. The EPM, recently proposed by Frisk et al. (2010), has been used in collaborative forest transportation and always results in a stable allocation, if the core is not empty. In addition to these methods, of course, a number of other methods have been proposed in the literature (see, for instance, Tijs and Driessen (1986) who present a survey of cost allocation methods and game theory, and propose their own method too). Wong et al. (2007) used four methods in computational experiments, but in their problem they consider a downtime cost, while we explicitly use a service level constraint instead. Other difference in our setting is that we do not consider lateral transshipment costs that they do. In our oil and gas context, all the plants we are dealing with lie relatively close to each other. Moreover, some of the inventory plants are physically located within the same warehouse. Also, the most extreme emergency cases are those when a main part in a platform off-shore need to be replaced and its corresponding on-shore plant does not have the part on stock. If the part is available from other plant, the part does not need to be sent from one to another on-shore plant before being sent off-shore, but it is just sent straight from the on-shore plant supplying it to the off-shore location.

5. Numerical examples

Next, we illustrate the allocations generated by the five methods using real-world data of three items in three plants. The data has been taken from the company in the oil and gas industry that we introduced in the background section. We use the historical demand data to estimate the demand distribution parameters in each plant. When computing the parameters for a coalition of plants, we aggregate the historical demand data of the corresponding plants.

5.1. Equal target service levels

We first compute results for the case when, for a same item, the three plants have the same target service levels. The targets are $\beta_{item1}=0.97$, $\beta_{item2}=0.95$, $\beta_{item3}=0.99$. Other parameter values are shown in Table 1. For confidentiality reasons, in the following analysis we do not reveal details on the cost data and when showed in the results, currency units and conversion factors have been omitted.

Table 1: Data on lead times and average demand for each item at each plant.

Plant	Item 1	Item 2	Item 3
Lead time	6.67	1.17	0.80
λ Plant P1	0.0448	0.0448	0.0597
λ Plant P2	0.0149	0.0448	0.0149
λ Plant P3	0.0299	0.0149	0.0597

For all the instances solved in this article, the optimal base-stock levels were quickly computed, in less than a second. The results observed for S are usually in the range from 1 to 3. A summary of results on safety stocks, average inventory on hand and total expected costs in the case of equal target service levels for all the plants is shown in Table 2. It can be observed that risk pooling implies decreases of more than 50% in safety stocks and average inventory on hand for all items, when comparing to the sum of the results when the plants control their inventory per separate. The cost reduction is 51%, 38% and 45% for items 1, 2 and 3, respectively.

Table 2: Safety stock (κ), average inventory on hand (\bar{I}) and total cost (C) results for three items in three equal service level plants (the first three rows show stand-alone costs, the fourth row their sum and the last row the cost when pooled in the grand coalition N).

Plant	Item 1			Item 2			Item 3		
	κ	\bar{I}	C	K	\bar{I}	C	K	\bar{I}	C
P1	1.70	2.20	1,113	0.95	1.45	562	0.95	1.45	4,824
P2	0.90	1.40	646	0.95	1.45	562	0.99	1.49	4,619
P3	0.80	1.30	673	-0.02	0.48	187	0.95	1.45	4,824
P1+P2+P3	3.40	4.90	2,432	1.88	3.38	1,310	1.94	2.94	9,443
N	1.40	1.90	1,196	0.88	1.38	810	0.89	1.39	5,167

Table 3 presents the cost allocation results obtained by the five different methods. For items 1 and 3, the five methods conducted to cost allocations in the core. This is easily verifiable when using the cost allocations values from Table 3 in constraints (1) - (3). For item 2, all the methods conducted to cost allocations in the core, except for the egalitarian one. The cost allocated to plant 3 by this method is 270 (highlighted cell in Table 3), which is greater than the cost 187 of this plant for item 2 when it was planned per separate, thus constraint (1) is violated and the resulting allocation does not belong to the core.

Table 3: Cost allocations by the five methods for three items in three plants, same target service levels.

Plant	Item 1					Item 2					Item 3				
	Egalit.	αD	Altruistic	Shapley	EPM	Egalit.	αD	Altruistic	Shapley	EPM	Egalit.	αD	Altruistic	SAM	EPM
P1	399	598	547	564	547	270	347	347	353	347	1,722	2,296	1,747	1,791	1,747
P2	399	199	317	302	317	270	347	347	353	347	1,722	574	1,673	1,585	1,673
P3	399	399	331	330	331	270	116	116	104	116	1,722	2,296	1,747	1,791	1,747
Total	1,196	1,196	1,196	1,196	1,196	810	810	810	810	810	5,167	5,167	5,167	5,167	5,167

Note that the upper bounds for constraints (1) correspond to the C_T values in the first three rows of Table 2, while the right-hand side of equation (3) corresponds to the optimal cost of the grand coalition given in the last row of Table 2. The upper bounds for constraints (2), namely the optimal cost for two-players coalitions are given in Table 4.

Table 4: Optimal costs for coalitions of two players, same target service levels.

Coalition	Item 1	Item 2	Item 3
{P1, P2}	1,141	748	4,893
{P1, P3}	1,168	624	5,098
{P2, P3}	1,113	624	4,893

5.2. Different target service levels

Now we consider the case when the three plants have different target service levels for a same item. This may happen, for instance, because the same items are utilized in different contexts in the different plants, or because the service levels at different plants have been set regarding different evaluation basis. Moreover, Porras and Dekker (2008) and Guajardo et al. (2012) have reported oil related problems where the same item can be utilized with different service levels even at a single inventory location.

We define plant P1 as a *high safe* plant, in the sense that it requires relatively high target service levels, while plant P2 is *medium safe* and plant P3 *low safe*.

For plant P1, the targets are $\beta^{P1}_{item1}=0.97$, $\beta^{P1}_{item2}=0.95$, $\beta^{P1}_{item3}=0.99$.

For plant P2, the targets are $\beta^{P2}_{item1}=0.87$, $\beta^{P2}_{item2}=0.85$, $\beta^{P2}_{item3}=0.89$.

For plant P3, the targets are $\beta^{P3}_{item1}=0.77$, $\beta^{P3}_{item2}=0.75$, $\beta^{P3}_{item3}=0.79$.

When plants act in coalition, the target service level that we consider is the maximum of the target values of the plants involved in the coalition. Rounding-up target service levels in this way is a common practice (although not necessarily optimal, see e.g. threshold rationing policies by Dekker et al. 1998 and Deshpande et al. 2003; and a computational study by Guajardo and Rönnqvist 2012). In our context, it is arguable if a player would be willing to join a coalition under a target service level lower than its own target; on the other hand, a player with a lower target would still satisfy its own target when joining a coalition with a higher target value.

Table 5: Safety stock (κ), average inventory on hand (\bar{I}) and total cost (C) results for three items in three different service level plants (the first three rows show stand-alone costs, the fourth row their sum and the last row the cost when pooled in the grand coalition N).

Plant	Item 1			Item 2			Item 3		
	κ	\bar{I}	C	K	\bar{I}	C	K	\bar{I}	C
P1	1.70	2.20	1,113	0.95	1.45	562	0.95	1.45	4,824
P2	-0.10	0.40	234	-0.05	0.45	311	-0.05	0.45	1,791
P3	-0.20	0.30	261	-0.02	0.48	187	-0.01	0.49	1,585
P1+P2+P3	1.40	2.90	1,608	0.88	2.38	1,060	0.89	2.39	8,200
N	1.40	1.90	1,196	0.88	1.38	810	0.89	1.39	5,167

As it was in the case of equal service level plants, Table 5 reveals that risk pooling implies important savings in total costs. Although in this case safety stocks are the same comparing to when the plants act per separate, the reduction of average inventory on hand translates into cost reductions of 26%, 24% and 37% for items 1, 2 and 3, respectively.

Table 6 presents the cost allocation results obtained by the five different methods. In contrast to the case when all plants used the same service level, now for all items most of the cost allocation derived from the methods do not belong to the core. We have highlighted all the cases where constraint (1) is violated.

Table 6: Cost allocations by the five methods for three items in three plants, different target service levels.

Plant	Item 1					Item 2					Item 3				
	Egalit.	αD	Altruistic	Shapley	EPM	Egalit.	αD	Altruistic	Shapley	EPM	Egalit.	αD	Altruistic	SAM	EPM
P1	399	598	828	838	828	270	347	429	478	436	1,722	2,296	3,039	3,813	3,307
P2	399	199	174	165	174	270	347	238	228	233	1,722	574	999	574	986
P3	399	399	194	193	194	270	116	143	104	140	1,722	2,296	1,128	780	873
Total	1,196	1,196	1,196	1,196	1,196	810	810	810	810	810	5,167	5,167	5,167	5,167	5,167

The only methods that produce allocations in the core for all the items are the Shapley value method and the EPM.

The egalitarian method for all items leaves either plants P2 or P3 or both of them with a more expensive allocation than if each of them would act alone without any pooling strategy and, therefore, the resulting allocation violates the individually rational condition (1) thus it does not belong to the core.

For the same reason, the cost allocation produced by the α Demand method does not belong to the core.

The altruistic method conduces to a cost allocation that satisfies constraints (1) for all items. However, only for item 1 constraints (2) and (3) also hold and, therefore, for this item the resulting allocation belongs to the core. For items 2 and 3, the stability constraint (2) for coalition {P2, P3} is violated. We present the upper bounds for constraints (2) in Table 7. It can be verified that for item 2 the altruistic method produces $u_2+u_3 = 381$, which is greater than the cost 373 of coalition {P2, P3} for this item. It means that under this cost allocation, plants P2 and P3 would have incentives to apply risk pooling between them, excluding plant P1 and, therefore, the cost allocation u is not in the core. The same situation is observed for item 3, where the altruistic method produces $u_2+u_3 = 2,127$, which is greater than the cost 1,859 of coalition {P2, P3} for this item.

Table 7: Optimal costs for coalitions of two players, different target service levels.

Coalition	Item 1	Item 2	Item 3
{P1, P2}	1,141	748	4,893
{P1, P3}	1,168	624	5,098
{P2, P3}	701	373	1,859

In addition, we have built other instances with different target service levels such that the Shapley method, as well as the proportional methods, produce non-stable allocations. Table 8 shows the results of the cost allocations derived from the Shapley value method for an instance with target service levels set as follows:

For plant P1, the targets are $\beta^{P1}_{item1} = 0.99$, $\beta^{P1}_{item2} = 0.9999$, $\beta^{P1}_{item3} = 0.9999$.

For plant P2, the targets are $\beta^{P2}_{item1} = 0.89$, $\beta^{P2}_{item2} = 0.90$, $\beta^{P2}_{item3} = 0.59$.

For plant P3, the targets are $\beta^{P3}_{item1} = 0.79$, $\beta^{P3}_{item2} = 0.90$, $\beta^{P3}_{item3} = 0.59$.

Note that this service level setting considers *very high* targets for plant P1 and lower targets for plants P2 and P3.

Table 8: Verification of constraints (1)-(3) for cost allocations generated by Shapley value method in additional instances with different target service levels.

	Item 1		Item 2		Item 3	
u_1	1,044	≤ 1,113	854	> 812	8,363	> 7,857
u_2	165	≤ 234	353	> 311	574	≤ 1,585
u_3	399	> 261	104	≤ 187	2,296	> 1,791
u_1+u_2	1,209	> 1,141	1,207	≤ 1,248	8,937	> 7,926
u_1+u_3	1,443	≤ 1,580	957	> 874	10,659	≤ 11,165
u_2+u_3	564	≤ 701	457	> 373	2,870	> 1,859
$u_1+u_2+u_3$	1,608	= 1,608	1,310	= 1,310	11,233	= 11,233

In Table 8, we have highlighted all constrains of type (1) and (2) violated by the Shapley value allocation. Note that for items 2 and 3, we set equal target service levels for plants P2 and P3. Still in these cases, where only plant P1 used different targets, the resulting allocations are not stable. Moreover, by running the linear programming associated to the EPM we obtain infeasibility, thus we are able to verify that the core in these three instances is empty.

6. Concluding remarks

Motivated from a case in the oil and gas sector, we have presented a problem where risk pooling has potential of significant savings. Our work highlights the importance of finding cost allocations which would allow the implementation of risk pooling in inventories of spare parts. Although risk pooling is a traditional approach in inventory management, in the specific context of spare parts the problem about how to allocate costs between the different players in the pool has received little attention from the literature.

We have performed a series of computations, using five cost allocations methods. We showed examples where a same cost allocation method could lead to solutions in the core when the players have the same target service levels, while under a round-up policy and the same setting but at least one player with different target service level the same method leads to an allocation which does not belong to the core or, moreover, the core becomes empty. Wong et al. (2007) also presented illustrative examples of cost allocations in spare part inventories and analysed them using game theory principles, but a main difference of our setting is that we consider a target service level explicitly when finding the optimal base-stock levels instead of a downtime cost. As well as in their work though, it remains as pending work finding structural results in order to characterize the allocations and the core of the game. This is not a straightforward task, however, because of the untractability of the expressions involved in the computation. Specifically when trying to find the optimal base-stock level S subject to a service level constraint, an iterative procedure as the one used in our article is the common approach, which does not facilitate closed forms to characterize S in terms of the parameters of the problem. Structural properties for cost allocation in spare parts problems have been found by Karsten et al. (2009, 2011a, 2011b) and Karsten and Basten (2012), using penalty costs for backorders instead of service level constraints. In practice, a service level constraints instead of penalty costs is generally easier to state and interpret by practitioners (Chen and Krass, 2001; Bashyam and Fu, 1998; Cohen et al. 1989).

As an extension to our article, we are designing a new allocation method which deals with diferent target service levels. It is based on computing a referential cost allocation for each player derived from the scenario where its target is the maximum allowed for all players. Then, in the actual scenario, where the targets of all players are allowed, we use a linear programming model which minimizes the maximum difference between the actual allocation and the referential cost. We are also studying the cost

allocation problem in threshold rationing policies (Dekker et al. 1998, Deshpande et al. 2003) as an alternative to the round-up policy.

In future research, it would be interesting to test more cost allocation methods and to approach the problem with other demand distributions or other control policies, such as the (s, S) policy. Finally, incorporating other practical issues, such as where to actually locate the parts when implementing a centralized solution for players originally decentralized or the implications of costs in a network with lateral transshipments instead of a centralized solution complement the agenda for future work. At the time of writing, we are collaborating with the company presented in this article in order to extend our analysis in these and other issues by incorporating a large set of plants and items.

References

- Alfredsson, P. and J. Verrijdt** (1999). Modeling emergency supply flexibility in a two-echelon inventory system. *Management Science* 45, 1416–1431.
- Archibald, T. W.** (2007). Modelling replenishment and transshipment decisions in periodic review multilocation inventory systems. *Journal of the Operational Research Society* 58, 948–956.
- Audy, J., S. D'Amours and M. Rönnqvist** (2012). An empirical study on coalition formation and cost/savings allocation. *International Journal of Production Economics* 136, 13–27.
- Bashyam, S. and M. C. Fu** (1998). Optimization of (s, S) inventory systems with random lead times and a service level constraint. *Management Science* 44, S243–S256.
- Chen, F. Y. and D. Krass** (2001). Inventory models with minimal service level constraints. *European Journal of Operational Research* 134, 120–140.
- Cohen, M. A., P. R. Kleindorfer and H. L. Lee** (1989). Near-optimal service constrained stocking policies for spare parts. *Operations Research* 37, 104–117.
- Cohen, M., P. V. Kamesam, P. Kleindorfer and H. Lee** (1990). Optimizer: Ibm's multi-echelon inventory system for managing service logistics. *Interfaces* 20, 65–82.
- Dekker, R., M. J. Kleijn and P. J. de Rooij** (1998). A spare parts stocking policy based on equipment criticality. *International Journal of Production Economics* 56-57, 69–77.
- Deshpande, V., M. A. Cohen and K. Donohue** (2003). A threshold inventory rationing policy for service differentiated demand classes. *Management Science* 49, 683–703.
- Frisk, M., M. Göthe-Lundgren, K. Jörnsten and M. Rönnqvist** (2010). Cost allocation in collaborative forest transportation. *European Journal of Operational Research* 205, 448–458.
- Gerchak, Y. and D. Gupta** (1991). On apportioning costs to customers in centralized continuous review inventory systems. *Journal of Operations Management* 10, 546–551.
- Gillies, D. B.** (1959). Solutions to general non-zero-sum games. In A. W. Tucker and R. D. Luce (Eds.) *Contributions to the Theory of Games IV (Annals of Mathematics Studies 40)*, Princeton University Press, 47–85.
- Guajardo, M., M. Rönnqvist, A. M. Halvorsen and S. I. Kallevik** (2012). Inventory management of spare parts in an energy company. *Working paper NHH Norwegian School of Economics*.
- Guajardo, M. and M. Rönnqvist** (2012). Comparing reorder points in $(S-1, S)$ inventory systems for items with multiple criticality indexes. *Working paper NHH Norwegian School of Economics*.
- Hartman, B. C. and M. Dror** (1996). Cost allocation in continuous-review inventory models. *Naval Research Logistics* 43, 549–561.
- Hartman, B. C. and M. Dror** (2000). Cores of inventory centralization games. *Games and Economic Behavior* 31, 26–49.

- Kalchschmidt, M., G. Zotteri and R. Verganti** (2003). Inventory management in a multi-echelon spare parts supply chain. *International Journal of Production Economics* 81-82, 397–413.
- Karsten, F., M. Slikker and G. V. Houtum** (2009). Spare parts inventory pooling games. *BETA Working Paper Eindhoven University of Technology* 300.
- Karsten, F., M. Slikker and G. V. Houtum** (2011a). Analysis of resource pooling games via a new extension of the Erlang loss function. *BETA Working Paper Eindhoven University of Technology* 344.
- Karsten, F., M. Slikker and G. V. Houtum** (2011b). Resource pooling and cost allocation among independent service providers. *BETA Working Paper Eindhoven University of Technology* 352.
- Karsten, F. and R. Basten** (2012). Spare parts inventory pooling: how to share the benefits? *BETA Working Paper Eindhoven University of Technology* 372.
- Kennedy, W. J., J. W. Patterson and L. D. Fredendall** (2002). An overview of recent literature on spare parts inventories. *International Journal of Production Economics* 76, 201–215.
- Kranenburg, A. A. and G. J. van Houtum** (2009). A new partial pooling structure for spare parts networks. *European Journal of Operational Research* 199, 908–921.
- Muckstadt, J. A.** (2005). *Analysis and Algorithms for Service Parts Supply Chain* (First ed.). New York, USA: Springer Series in Operations Research and Financial Engineering.
- Nenes, G., S. Panagiotidou and G. Tagaras** (2010). Inventory management of multiple items with irregular demand: A case study. *European Journal of Operational Research* 205, 313–324.
- Paterson, C., G. Kiesmüller, R. Teunter and K. Glazebrook** (2011). Inventory models with lateral transshipments: A review. *European Journal of Operational Research* 210, 125–136.
- Porras, E. and R. Dekker** (2008). An inventory control system for spare parts at a refinery: An empirical comparison of different re-order point methods. *European Journal of Operational Research* 184, 101–132.
- Robinson, L. W.** (1993). A comment on Gerchak and Gupta's "On apportioning costs to customers in centralized continuous review inventory systems". *Journal of Operations Management* 11, 99–102.
- Shapley, L. S.** (1953). A Value for n-person games. *Annals of Mathematical Studies* 28, 307–317.
- Sherbrooke, C. C.** (2004). *Optimal Inventory Modeling of Systems Multi-Echelon Techniques* (Second ed.). Boston, USA: Kluwer Academic Publishers.
- Silver, E. A., D. F. Pyke and R. Peterson** (1998). *Inventory Management and Production Planning and Scheduling* (Third ed.). United States: John Wiley & Sons.
- Tijs, S. H. and T. S. H. Driessen** (1986). Game theory and cost allocation problems. *Management Science* 32, 1015–1028.
- Vereecke, A. and P. Verstraeten** (1994). An inventory management model for an inventory consisting of lumpy items, slow movers and fast movers. *International Journal of Production Economics* 35, 379–389.
- Wong, H., D. V. Oudheusden and D. Cattrysse** (2007). Cost allocation in spare parts inventory pooling. *Transportation Research Part E* 43, 370–386.
- Wong, H., G. J. van Houtum, D. Cattrysse and D. van Oudheusden** (2006). Multi-item spare parts systems with lateral transshipment and waiting time constraints. *European Journal of Operational Research* 171, 1071–1093.