European Journal of Operational Research 233 (2014) 94–104

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Production, Manufacturing and Logistics

Pooling of spare parts between multiple users: How to share the benefits?

Frank Karsten [a,*], Rob J.I. Basten [b]

[a] School of Industrial Engineering, Eindhoven University of Technology, P.O. Box 513, 5600 MB Eindhoven, The Netherlands
[b] Faculty of Engineering Technology, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands

ARTICLE INFO

Article history:
Received 4 September 2012
Accepted 19 August 2013
Available online 4 September 2013

Keywords:
Inventory
Cooperative game theory
Core
Inventory pooling
Spare parts

ABSTRACT

Companies that maintain capital goods (e.g., airplanes or power plants) often face high costs, both for holding spare parts and due to downtime of their technical systems. These costs can be reduced by pooling common spare parts between multiple companies in the same region, but managers may be unsure about how to share the resulting costs or benefits in a fair way that avoids free riders. To tackle this problem, we study several players, each facing a Poisson demand process for an expensive, low-usage item. They share a stock point that is controlled by a continuous-review base stock policy with full backordering under an optimal base stock level. Costs consist of penalty costs for backorders and holding costs for on-hand stock. We propose to allocate the total costs proportional to players' demand rates. Our key result is that this cost allocation rule satisfies many appealing properties: it makes all separate participants and subgroups of participants better off, it stimulates growth of the pool, it can be easily implemented in practice, and it induces players to reveal their private information truthfully. To obtain these game theoretical results, we exploit novel structural properties of the cost function in our $(S-1, S)$ inventory model.

1. Introduction

Capital goods such as trams, manufacturing systems, power plants, and airplanes form the backbone of much of our society. Users of such capital goods are often confronted with the difficult task of guaranteeing high availabilities of their expensive, technologically advanced systems. A commonly used strategy to prevent lengthy downtimes is to immediately replace any failed component with a functioning spare part. Obviously, this strategy functions only if a spare part is available when needed, but the required stocks tend to tie up a lot of capital. For instance, the commercial aviation industry has as much as $30 billion worth of spare engines on stock (Flint, 2006). More generally, the sale of spare parts and after-sales services has been pegged at $1 trillion every year in the United States alone, which represents 8% of its gross domestic product (Cohen, Agrawal, & Agrawal, 2006, pp. 129–130, & references therein). At the same time, being out of stock when a spare part is needed leads to downtime of capital goods, which is very expensive due to loss of operational continuity. For example, in the semiconductor industry, the opportunity costs for lost production are estimated to run into tens of thousands of euros per hour (Kranenburg, 2006, p. 17).

Because of the high costs involved, both spare parts holding costs and downtime costs, many companies in the capital goods industry are looking for ways to reduce these costs. Intuitively, it makes sense for companies in the same geographic area to pool common spare parts. Indeed, as stated by Cohen et al. (2006, p. 136): "The best way for companies to realize economies of scale is to pool spare parts". Tram operators in the Netherlands are a good example. In the Netherlands, the local public transport in the three largest cities (Amsterdam, Rotterdam, and The Hague, all of which are within an hour's driving distance of each other) is operated by a separate company per city. Although the operators use trams of different models, there is still a lot of commonality on the component level, enabling an excellent opportunity for inventory pooling. Another example setting is that of independently managed plants of a large energy company, as described in Guajardo, Ronnqvist, Halvorsenb, and Kallevik (2012): the plants currently hold their inventory separately, but annual savings of 44% may be obtainable if pooling is taken into account. While promising, this does raise the question of how the plants should share these benefits, which is mentioned by Guajardo et al. (2012) as an important research direction. Kukreja, Schmidt, and Miller (2001) describe a similar case of pooling possibilities between independently operating power-generating plants, for which substantial savings of 68% are achievable by pooling of common parts. Spare parts pools for multiple companies already exist in the

* Corresponding author. Tel.: +31 40 247 4432.
E-mail address: f.j.p.karsten@tue.nl (F. Karsten).

0377-2217/$ - see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.ejor.2013.08.029

airline industry (Flint, 2006) and in the military: a number of European air forces and navies are currently pooling their spare parts, and other countries have shown interest in joining the pool (Hale, 2011).

The successful collaborations in aerospace and defense are encouraging, and the potential for huge cost savings is attractive. Nevertheless, cooperative pooling of common spare parts between different companies, or between different business units, is not a common practice yet. One major obstacle appears to be the identification of a fair cost sharing mechanism. In our contacts with the capital goods industry, we find that practitioners are mainly hesitant to pool spare parts because they are not sure it will lead to a cost saving for themselves. Some of the commonly stated fears are that some group of firms may end up paying to subsidize the others, that the other companies may not disclose their private information truthfully, and that new members might take more benefit out of the pool than they bring in. In this paper, we will tackle these issues by applying concepts from game theory to determine an appealing cost sharing mechanism. This is practically relevant for decision makers that consider starting up a new spare parts pool and, additionally, it may aid participants in existing pools to decide on whether or not to adapt the cost allocation rules that are currently in use.

Before we can analyze cost sharing mechanisms, we first need a suitable model for a (shared) stock point operated by any number of players. The model should be realistic for parts for which pooling is especially interesting due to their large economic impact: expensive, low-demand spare parts with long lead times and no emergency supply flexibility. For such parts, whose demands typically occur in accordance with a Poisson process, a continuous-review base stock policy with one-for-one replenishments is appropriate. Therefore, we will analyze the resulting single-echelon $(S-1, S)$ inventory model, taking into account holding costs incurred while spare parts are in stock and penalty costs incurred while a capital good is down due to the unavailability of a spare part.

The cost and behavior of an inventory system greatly depend on what happens to demands when the system is out of stock. In practice, there are two common ways to deal with stock-outs: using backlogging or emergency procedures. An emergency procedure typically refers to the instantaneous delivery of a part from an alternative supplier. There are some real-life settings where such emergency deliveries are possible. In many other cases, backordering is the only option. This is often true for consumable parts that are produced by one supplier only and for repairable parts that are no longer in production (in which case one has to wait for a part to return from the repair shop). Backordering is in line with what happens in the real-life examples mentioned earlier; for instance, the stock planners of the Dutch tram operators plan for backorders. Backordering is also commonly assumed in the stream of literature on spare parts models, as reviewed in Section 2.

In the literature, the only previous analytical investigation of cost allocation mechanisms for spare parts pools with multiple players (Karsten, Slikker, & van Houtum, 2012) has focused on a model with emergency procedures. Yet, for the many real-life cases that lack an alternative supplier with a negligible lead time, the results of Karsten et al. (2012) do not apply. The present paper fills this gap by tackling the setting with backordering. Our analysis is drastically different from Karsten et al. (2012), as we discuss in Section 2. Besides that, we contribute to the literature by discussing implementation issues (in Sections 7.2 and 8) that were not considered by Karsten et al. (2012).

The $(S-1, S)$ inventory model with backordering that we consider is more generally applicable for expensive, low-demand items for which ordering costs are negligible compared with holding and shortage costs. Although this means that our analysis and results may also be relevant for other applications, such as inventory pooling of luxury cars, we use spare parts terminology in this paper to enable a concrete exposition and a concrete justification of assumptions. The general $(S-1, S)$ model has been studied extensively in the literature, due to its high practical relevance. As a result, the steady-state distributions of the number of items on hand and on backorder are well-known; the same holds for the average long-term costs and the behavior of these costs as a function of the base stock level. These results, however, do not directly help in identifying a suitable cost sharing mechanism for the problem at hand. For that, we need to understand how the average long-term costs behave when the demand rate varies (as a result of new players joining the pool). Therefore, we first derive new convexity and elasticity properties of the costs as a function of the demand rate in our $(S-1, S)$ inventory model, and we show that pooling the demand streams and inventory of a number of given stock points leads to reduced backorders, inventory, and costs.

After having formally shown that pooling is indeed cost effective from a system's point of view, we turn to our cost sharing problem. We focus on several players (e.g., companies, business units, or defense organizations) that are located geographically close together. They have identical cost structures and replenishment lead times. Players have the choice of either operating their own stock point (which behaves as an $(S-1, S)$ inventory system) or setting up a shared stock point from which the combined demand streams of the participants are fulfilled (also behaving as an $(S-1, S)$ inventory system, but with a higher demand rate and likely a higher optimal base stock level than for a single player).

If the players decide to operate a shared stock point, they should also decide on a rule to assign the resulting holding and backorder costs among the players, preferably in a way that is appealing from a practical perspective. Four relevant properties or requirements that an allocation rule might satisfy are that: (1) it gives a fair allocation of the total expected costs to the various players, (2) it stimulates growth by making it interesting for existing players to allow more players to join, (3) it is easy to understand and implement, and (4) it gives players an incentive to disclose all relevant information truthfully. One important notion of fairness from cooperative game theory – the core – requires that a cost allocation should not give any subgroup of players an incentive to split off and form a separate pool. We take this concept of the core as our guideline for the first requirement posited above, i.e., we aim to find an allocation under which each subgroup of players gets better off. This is not trivial: as is well-known in cooperative game theory, an overall lower cost is not necessarily a guarantee that a core allocation exists.

Taking this into consideration, identification of an allocation rule satisfying the first requirement, let alone all four requirements, may seem to be a complex problem. Nevertheless, we show that this problem does have a solution and a surprisingly simple one at that: the straightforward allocation of total costs *proportional* to player's demand rates satisfies all required properties! We see this as the main contribution of our paper. Interestingly, the expected cost allocations prescribed by this proportional rule coincide with the common practice of charging a fixed fee per flight hour for participation in an aircraft component pool (assuming that all costs are fully shared and that component failures rates per flight hour are the same across players). Thus, our results provide support for these flight-hour charges from a game theoretical perspective.

Implementation of this proportional cost allocation in practice is a next challenge, especially since *realized* costs in any period of time may differ greatly from expected costs. To deal with this, we propose a process to fairly allocate cost *realizations*, and discuss its implications for truthful information disclosure in the context of a non-cooperative game. These issues have been previously considered in the context of collaborating newsvendors with no inventory carryover: fair divisions of realized costs are studied by Dror, Guardiola, Meca, and Puerto (2008), Chen and Zhang (2009, Remark 3), and Kemahlıoğlu-Ziya and Bartholdi (2011, Section 6) while schemes that induce truthful revelation of private demand information are studied by Norde, Özen, and Slikker (2011). We, in contrast, tackle these practically relevant issues in an infinite-horizon continuous-review inventory model. Our approach differs from existing approaches in the newsvendor context. Specifically, we identify a process that fairly allocates costs as they materialize and, subsequently, we establish a link between this process and our proportional core allocation.

The remainder of this paper is organized as follows. We first discuss the related literature in Section 2 and subsequently describe our inventory model in Section 3. Next, we give some preliminaries on cooperative game theory in Section 4 and introduce our cooperative spare parts pooling games in Section 5. The next three sections form the main part of our work. We analyze the associated inventory model in Section 6. The results are used in Section 7 to show that the proportional rule always results in a core allocation and that adding extra players makes everyone better off. In Section 8, we show how this allocation rule can be implemented in practice and how it induces players to truthfully disclose their private information. We conclude in Section 9.

2. Related literature

There are two streams of related literature: the literature on spare parts inventory management and the literature on cooperative game theory applied to inventory or queueing systems.

Due to the high economic impact of spare parts all around the world, the amount of literature on spare parts inventory management is enormous. The first relevant paper is that of Feeney and Sherbrooke (1966), who derive the steady-state distribution of the number of items in resupply in the spare parts model that we consider. Several subsequent papers have studied pooling of spare parts between several locations, under the assumption that the system is owned by a single entity who decides whether or not to pool. Examples include Kukreja et al. (2001) and Wong, Cattrysse, and van Oudheusden (2005); they show that if there are multiple stock points at one echelon level, it is generally worthwhile to use lateral transshipments between these stock points in order to reduce costs or increase the service level. Our model is in line with the models used in these two examples, although we do not explicitly take transshipment costs into account in our present paper. Wong, van Houtum, Cattrysse, and van Oudheusden (2006) give an extensive overview of the literature on pooling in spare parts inventory models, and we refer to the books by Sherbrooke (2004) and Muckstadt (2005) for extensive overviews of the spare parts literature in general. Opposed to this literature, we will consider the setting with independent parties and address the issue of fair cost allocation.

The literature on cooperative games in inventory systems has recently been reviewed by Fiestras-Janeiro, García-Jurado, Meca, and Mosquera (2011) and Dror and Hartman (2011). Four lines of research can be distinguished. First, the literature on games in which players face deterministic demand, use economic order quantity policies, and cooperate by using joint replenishments; see, e.g., Anily and Haviv (2007). In contrast to this line of research, we consider a model with *stochastic* demand. Second, the vast literature on single-period newsvendor games, in which players face stochastic demand and may cooperate by coordinating orders and pooling inventory; see, e.g., Hartman, Dror, and Shaked (2000), Slikker, Fransoo, and Wouters (2001), Müller, Scarsini, and Shaked (2002), Dror et al. (2008), Özen, Fransoo, Norde, and Slikker (2008), Chen and Zhang (2009), and Kemahlıoğlu-Ziya and Bartholdi (2011). In this paper, we study an *infinite*-horizon model. Surprisingly, the resulting class of spare parts pooling games turns out to coincide with the class of (single-period) newsvendor games where all players have Poisson distributed demand. Yet, as we discuss in Section 7.3, our focus on this specific subclass of newsvendor games – a subclass that has never been explicitly studied before – enables us to find novel results that do extend to newsvendor games in general. The third line of research is on inventory centralization games in a continuous-review setting with stochastic demand and penalty costs per backorder occurrence independent of duration. Hartman and Dror (1996) and references therein study such games via approximate evaluation. We, in contrast, perform exact evaluation of a setting in which backorder costs are paid for each unit of time a part is lacking.

The fourth, relatively scarce line of research on cooperative games in inventory systems is motivated by spare parts applications. Wong, van Oudheusden, and Cattrysse (2007) are the first to study a multi-location, continuous-review, infinite-horizon setting with several players who cooperate by pooling their parts. They propose various cost allocation policies and numerically illustrate them, but their work lacks structural results. Karsten et al. (2012) derive structural results for cooperative games in which resources, such as spare parts, can be pooled. As mentioned in the introduction (Section 1), their model differs from ours in one key aspect: we assume full backordering if a demand cannot be fulfilled immediately, whereas they assume that there is an emergency option, which results in lost sales for the inventory system under study. It is well-known in the inventory literature (see, e.g. Feeney & Sherbrooke, 1966) that results obtained for a model with lost sales need not carry over to a model with backordering, or vice versa, and that the two models require different analysis. Indeed, to show the existence of a core allocation for their games, Karsten et al. (2012) use properties of (new extensions of) the Erlang loss formula (i.e., the blocking probability of an $M/G/s/s$ queueing system), which has no direct relation with the model considered in the present paper.

Finally, we briefly mention the literature that applies cooperative game theory to analyze resource pooling in queueing facilities. A recent overview is provided by Karsten, Slikker, and van Houtum (2011b), who themselves study a model in which several $M/M/s$ queues join forces. Özen, Reiman, and Wang (2011) study the core of similar queueing games. The stream of literature on queueing games is relevant because there is a correspondence between the pipeline stock in our spare parts inventory model and the number of busy servers in an $M/G/\infty$ queueing model. Nevertheless, we have an inventory buffer confounding our analysis and the $M/G/\infty$ queue behaves fundamentally different from the $M/M/1$ and $M/M/s$ queues that have been considered in existing queueing games.

3. The $(S-1, S)$ inventory model with backlogging

We consider a single location that stocks one item. Initially, there are $S \in \mathbb{N}_0 (\mathbb{N}_0 = \mathbb{N} \cup \{0\})$ parts on stock. The demand process is a Poisson process with stationary rate $\lambda > 0$. A demand is immediately fulfilled from stock if a part is available. Otherwise, it is backordered and fulfilled first come first serve. In either case, an order for a new part is instigated immediately. This means that the stock point operates under a continuous-review base stock policy with base stock level S and one-for-one replenishments.

The stock point orders parts at an external, uncapacitated supplier. The time that elapses between demand occurrence and receival of the new part is called the lead time. Lead times are assumed to be independent and identically distributed (i.i.d.)

according to some general distribution function, and we assume without loss of generality (by rescaling time) that its mean is 1 time period.

In the remainder, we will analyze the resulting inventory system in steady-state. First, we consider the number of parts on order, the so-called pipeline stock, denoted by $X(\lambda)$. By Palm (1938), $X(\lambda)$ is Poisson distributed with mean λ, i.e., for all $x \in \mathbb{N}_0$ it holds that

$$\mathbb{P}[X(\lambda) = x] = \frac{\lambda^x}{x!} e^{-\lambda}. \tag{1}$$

We are mainly interested in the number of backorders, $B(S, \lambda)$, and the stock on hand, $I(S, \lambda)$, both as functions of the base stock level S and the demand rate λ. Backorders exist if the pipeline stock is larger than the base stock level, so $B(S, \lambda) = \max\{X(\lambda) - S, 0\}$. Similarly, $I(S, \lambda) = \max\{S - X(\lambda), 0\}$. Thus, using Eq. (1), we can obtain the distributions and expectations of the number of backorders and the total stock on hand. For instance, the distribution of the number of backorders, $B(S, \lambda)$, is given by

$$\mathbb{P}[B(S,\lambda) = x] = \begin{cases} \sum_{y=0}^{S} \mathbb{P}[X(\lambda) = y] & \text{if } x = 0; \\ \mathbb{P}[X(\lambda) = x + S] & \text{if } x \in \mathbb{N}. \end{cases} \tag{2}$$

Accordingly, the expected number of backorders is

$$\begin{aligned} \mathbb{E}B(S,\lambda) &= \sum_{x=S+1}^{\infty} (x - S)\mathbb{P}[X(\lambda) = x] \\ &= \sum_{x=0}^{\infty} x \cdot \mathbb{P}[X(\lambda) = x] - \sum_{x=0}^{\infty} S \cdot \mathbb{P}[X(\lambda) = x] \\ &\quad + \sum_{x=0}^{S} (S - x)\mathbb{P}[X(\lambda) = x] = \lambda - S \\ &\quad + \sum_{x=0}^{S} (S - x)\mathbb{P}[X(\lambda) = x]. \end{aligned} \tag{3}$$

Similarly, the expected on-hand stock is

$$\mathbb{E}I(S,\lambda) = \sum_{x=0}^{S} (S - x)\mathbb{P}[X(\lambda) = x] = \mathbb{E}B(S,\lambda) - \lambda + S. \tag{4}$$

We consider holding costs $h > 0$ per unit time per spare part in the on-hand stock. These costs encompass warehousing, insurance, and interest costs on the capital tied up by the inventory. Furthermore, we consider penalty costs $b > 0$ per unit time per backordered demand. (We disregard the part procurement price or holding costs for pipeline stock because these cost factors would represent constant terms, unaffected by decisions on base stock level or collaboration.) The long-term average costs per unit time are given by

$$K(S,\lambda) = h \cdot \mathbb{E}I(S,\lambda) + b \cdot \mathbb{E}B(S,\lambda). \tag{5}$$

Although the above-described model is general and might also apply for, e.g., inventories of luxury cars, its formulation and underlying assumptions are driven by inventory systems of expensive, low-demand spare parts meant for technologically advanced capital goods. Indeed, the critical assumptions (the stationary Poisson demand process, the continuous-review one-for-one replenishment strategy, and the independence of successive lead times) are justifiable for this spare parts setting (see, e.g., Wong et al., 2006, 2012).

The assumption regarding the demand process in particular is standard in the spare parts literature. It is justified because a stock point typically serves multiple high-tech machines, whose merged stream of component failure processes corresponds very well to a Poisson process. Although no failures occur in a machine that is down due to a backordered part, the work for a broken machine is often largely taken over by the functional machines (albeit at penalty costs b per unit time) and, moreover, the expected number of broken machines is typically small relative to the total number of machines. Together, this implies that the total component failure rate remains close to constant.

In a typical spare parts setting, demands are triggered by failures of either consumable or repairable machine components. Although we formulated our model for consumable parts, it is also applicable for repairable parts if – instead of placing orders for new parts – any failed component is immediately sent to an uncapacitated repair facility that returns the component to the stock point as a ready-for-use spare part after an i.i.d. repair lead time whose mean is scaled to 1 time unit. (Again, direct repair costs and holding costs for parts in repair would represent constant terms and can be disregarded without loss of generality.)

In this paper, we are mainly interested in the problem of fair allocation of shared costs in a spare parts inventory system operated by multiple players. We will tackle this problem for the above-described model by representing the situation as a cooperative game, allowing us to take into account what the various subgroups of players could achieve.

4. Cooperative game theory

In this section, we treat concepts from cooperative game theory that are relevant to our work. A cooperative cost game with transferable utility, which we will simply refer to as a *game*, is a pair (N, c). Here, N is the non-empty finite *set of players*, also referred to as the *grand coalition*. Any non-empty subset $M \subseteq N$, $M \neq \emptyset$ is called a *coalition*, and the set of all coalitions is denoted by 2^N_-. Given the set of players N, a game (N, c) is defined by the *characteristic cost function* c, which assigns to every coalition M its costs $c(M)$. In our spare parts pooling context, the value $c(M)$ will be interpreted as the total long-term average costs per unit time of the joint inventory system if only the players in M are involved in it.

In cooperative game theory, players are assumed to be able to draw up binding agreements, and side payments are allowed. This is also allowed in our spare parts pooling context, where furthermore cooperation by the grand coalition is socially optimal. Accordingly, a central problem is allocating $c(N)$ to the individual players in a stable way. To formalize this, we call any vector $x = (x_i)_{i \in N} \in \mathbb{R}^N$ with $\sum_{i \in N} x_i = c(N)$ an *allocation* for game (N, c). The value x_i is then interpreted as the costs assigned to player i. The game theory literature provides various allocation rules; a famous one is the Shapley value (Shapley, 1953). Loosely speaking, it is calculated as the average of marginal contributions of players to the coalitions.

An appealing property that an allocation might satisfy is stability. An allocation x for a game (N, c) is called *stable* if $\sum_{i \in M} x_i \leqslant c(M)$ for all coalitions M. Under a stable allocation, each group of players has to pay no more collectively than what they would face by acting independently as a group. The set of all stable allocations is called the *core*, introduced by Gillies (1959). A game may have an empty core, even if the costs of the grand coalition are lower than the sum of the costs over any partition of the players.

We next strengthen the notion of a stable cost allocation. In our experience, practitioners are usually interested in allocations under which each coalition becomes *strictly* better off as a result of cooperation. Indeed, an allocation x under which some subgroup of players is indifferent between cooperating or not (i.e., if $\sum_{i \in M} x_i = c(M)$ for some subcoalition $M \in 2^N_-, M \neq N$) may be hard to defend in practice because players may decide not to collaborate if they do not strictly benefit from it, out of spite. This issue is rarely addressed in the literature on cooperative game theory, with the notable exception of Zhao (2001), who introduces and characterizes the *relative interior of the core*. We will instead consider the *strict core*, which is defined as

$$\mathcal{C}(N,c)=\left\{x\in\mathbb{R}^N\middle|\sum_{i\in N}x_i=c(N)\text{ and }\sum_{i\in M}x_i<c(M)\text{ for all }M\in 2^N_-,M\neq N\right\}.$$

Given a game (N, c), we call any element of $\mathcal{C}(N,c)$ a *strictly stable* allocation. Such allocations remain stable for small perturbations of the characteristic cost function.

The last concept that we wish to introduce is a (strict) population monotonic allocation scheme (cf. Sprumont, 1990). An allocation scheme for a game (N, c) is a vector $y=(y_{i,M})_{i\in M,M\in 2^N_-}$ with $\sum_{i\in M}y_{i,M}=c(M)$ for all $M\in 2^N_-$, which specifies how to allocate the costs of every coalition to its members. This scheme is called a *population monotonic allocation scheme* (PMAS) if the amount that any player has to pay does not increase when the coalition to which he belongs grows. That is, $y_{i,\ M}\geqslant y_{i,\ L}$ for all $M,L\in 2^N_-$ with $M\subset L$ and $i\in M$. If this inequality is strict for the members of all such nested pairs of coalitions, then we call this scheme a *strictly population monotonic allocation scheme* (SPMAS). It is apparent from this definition that if a game (N, c) admits a (strict) PMAS, say y, then $(y_{i,\ N})_{i\in N}$ is a (strictly) stable allocation, which implies that (N, c) has a non-empty (strict) core.

5. Spare parts pooling games

Consider several players who may pool inventories of a common item. Each player witnesses a stationary Poisson demand process, and the demand processes of the players are assumed to be independent. The players have the same mean replenishment lead time, possibly because they use the same supplier or repair facility, and without loss of generality we rescale it to 1 time unit. The replenishment lead times of the players are mutually independent. The players' holding costs and backorder costs are the same, possibly because they operate in the same industry under similar operating conditions. For instance, the tram operators of Amsterdam, Rotterdam, and The Hague (see Section 1) typically get rather sophisticated parts from a common supplier and face approximately the same holding and tram downtime costs. To capture all relevant parameters of such a setting, we define a *spare parts situation with backordering* as a tuple $\varphi = (N, (\lambda_i)_{i\in N}, h, b)$, where N is the non-empty finite set of players, $\lambda_i > 0$ is the demand rate of player $i \in N$, $h > 0$ is the holding cost rate, and $b > 0$ is the backorder cost rate.

We assume that any coalition M can set up a single stock point from which the combined demand streams of the coalition members are fulfilled First-Come-First-Served. For instance, the Dutch tram operators might pool their spare parts in one central stock point in Leiden, a city located close to the geographical barycenter of Amsterdam, Rotterdam, and The Hague. Since the superposition of independent Poisson processes is also a Poisson process, this single stock point would face a Poisson arrival process with merged rate $\lambda_M=\sum_{i\in M}\lambda_i$. We assume throughout that players are interested in reducing their long-term average holding and backorder costs, and that other, smaller effects of setting up the pool are insignificant in comparison. For instance, if the three tram operators set up the stock point in Leiden, there might be some additional transportation costs from this stock point to the trams (which will be small compared with holding and backorder costs, given that traveling distance from Leiden to each of the three cities is only 30–40 min). On the other hand, operating a single central warehouse will be less costly than operating three separate warehouses. We disregard these minor issues in order to focus in more detail on the potentially huge savings in holding and backorder costs.

Remark 5.1. In practice, penalty costs may sometimes be hard to quantify, whereas service constraints are more readily adoptable. However, as discussed by van Houtum and Zijm, 2000, cost models and service models are equivalent, i.e., there exists a one-to-one relationship between them. We chose to concentrate on a cost model in the present paper because it permits adequate comparisons between coalitions. To illustrate, suppose that we would instead have chosen to study a service model where each coalition would minimize the integer base stock level subject to, say, a 95% fill rate constraint. Then, improvements in service beyond that 95% would unjustly appear worthless: In such a service model, a player would prefer a coalition in which he would face 1000 \$/month in holding costs under a 95.0% fill rate over a coalition costing 1001 \$/month for 99.9%. Such unnatural outcomes are avoided by considering backlogging costs explicitly.

The first natural question to be asked is whether or not pooling the inventory and demand streams of several players is always beneficial. We start by giving a positive answer for a situation where the base stock level of the joint stock point is set at the sum of the initial stocking levels of all players. This could happen if players already possess repairable spare parts which cannot be sold or produced anymore due to their specificity (in which case re-optimization of base stock levels would not be possible and all inventories are taken over in full). For this situation, the following proposition states that pooling leads to a strict reduction in expected backorders (Eq. (3)), expected on-hand inventory (Eq. (4)), and expected costs (as defined in Eq. (5)). The intuition behind this is that pooling allows one player's backorder to cancel against another player's on-hand part, and the proof we provide may help understand this. (We remark that a cost decrease would not be guaranteed in case the players would have vastly *different* backorder cost rates, as then a player with low backlogging costs may often take a part that would have better been saved to guard against the higher backorder costs of another player. See the appendix for an example illustrating this effect.)

Proposition 5.1. *Consider a set of players N where any player $i\in N$ faces demand rate λ_i and owns $S_i\in\mathbb{N}_0$ parts. Suppose that $\sum_{i\in N}S_i>0$ and that $|N|\geqslant 2$.*

(i) $\sum_{i\in N}\mathbb{E}B(S_i,\lambda_i)>\mathbb{E}B\left(\sum_{i\in N}S_i,\sum_{i\in N}\lambda_i\right)$.
(ii) $\sum_{i\in N}\mathbb{E}I(S_i,\lambda_i)>\mathbb{E}I\left(\sum_{i\in N}S_i,\sum_{i\in N}\lambda_i\right)$.
(iii) $\sum_{i\in N}K(S_i,\lambda_i)>K\left(\sum_{i\in N}S_i,\sum_{i\in N}\lambda_i\right)$.

Proof. Consider, at an arbitrary point in time, the pipeline stocks $(X_i)_{i\in N}$ belonging to the various players. The pipeline stock of any player $i\in N$ is drawn from a Poisson distribution with mean λ_i, i.e., unaffected by any pooling arrangement. Consider any realization of pipeline stocks $x=(x_i)_{i\in N}$. Given x, the total number of backorders under no pooling, $\sum_{i\in N}\max\{x_i-S_i,0\}$, can never be less than the total number of backorders under pooling, $\max\{\sum_{i\in N}x_i-\sum_{i\in N}S_i,0\}$. The former is even *strictly* larger than the latter if $x_i-S_i<0$ while $x_j-S_j>0$ for two distinct players i, $j\in N$. Such a realization x exists with positive probability because, by assumption, there is a player $i\in N$ with $S_i>0$. Hence, we conclude that $\sum_{i\in N}\mathbb{E}B(S_i,\lambda_i)>\mathbb{E}B\left(\sum_{i\in N}S_i,\sum_{i\in N}\lambda_i\right)$, i.e., Part (i) holds. Parts (ii) and (iii) follow trivially from Part (i) by Eqs. (4) and (5). □

So, it is beneficial to share players' inventories while maintaining the aggregate of their base stock levels, but we may reduce costs even further by allowing any coalition to re-optimize their joint base stock level. After all, due to the risk pooling effect, lower base stock levels may suffice to jointly serve all demand streams in a cost-effective way. (Although such a lower aggregate base stock level might result in an *increase* in the expected backorders compared to the no-pooling situation, this would be counterweighted by the reduction in the expected on-hand inventory.) Since stocking levels in most real collaborations are adjustable rather than fixed, the base stock level of any coalition's stock point will be a decision variable in the remainder of this paper.

For any particular choice of the base stock level $S \in \mathbb{N}_0$, the behavior of the stock point would correspond to the model described in Section 3, and thus the expected relevant costs per time unit faced by coalition M would be equal to $K(S, \lambda_M)$. Assuming that any coalition picks an optimal base stock level, the game (N, c^{φ}) corresponding to our spare parts situation with backordering φ is defined by

$$c^{\varphi}(M) = \min_{S \in \mathbb{N}_0} K(S, \lambda_M) \tag{6}$$

for all coalitions M. We call this game the associated *spare parts pooling game*.

As mentioned, we are interested in methods to fairly distribute the collective expected costs of the grand coalition over the cooperating players in any spare parts situation with backordering $\varphi = (N, (\lambda_i)_{i \in N}, h, b)$. A simple rule is to divide these costs proportional to the demand rate of each player. Formally, we define this rule $\mathcal{P}$ by $\mathcal{P}_i(\varphi) = c^{\varphi}(N) \cdot \lambda_i / \lambda_N$ for each $i \in N$ in situation φ. Extending this idea to every coalition, we define the proportional allocation scheme rule $\mathcal{P}$ by

$$\mathcal{P}_{i,M}(\varphi) = c^{\varphi}(M) \cdot \lambda_i / \lambda_M \tag{7}$$

for each coalition M and $i \in M$ in situation φ. The following example illustrates this proportional rule numerically and simultaneously shows that the Shapley value is not necessarily in the core. In the appendix we give another example showing that another seemingly natural allocation, in which the difference between the costs of the grand coalition and the sum of the costs for the single-player coalitions is allocated proportional to the demand rates, is also not necessarily in the core.

Example 5.1. Consider three companies in the capital goods industry that wish to pool common parts. One company expects to face 0.1 demands per month on average. The monthly demand rates are 0.8005 and ln 2 for the other two companies. The part in question is very expensive; a single part on hand costs 10,000 dollars per month, and a machine that is down will cost 10,000 dollars per month as well.[1] This can be modeled as a spare parts situation with backordering $\varphi = (N, (\lambda_i)_{i \in N}, h, b)$ with player set $N = \{1, 2, 3\}$, $\lambda_1 = 0.1$, $\lambda_2 = 0.8005$, $\lambda_3 = \ln 2 (\approx 0.6931)$, and $h = b = 1$. To illustrate the determination of an optimal base stock level and associated costs, consider the singleton coalition {3}. By Eq. (1), $\mathbb{P}[X(\lambda_3) = 0] = 0.5$. Combining this with Eqs. (3)–(5), we obtain for the case were player 3 would decide to stock zero parts that $\mathbb{E}B(0, \lambda_3) = \ln 2$, $\mathbb{E}I(0, \lambda_3) = 0$, and $K(0, \lambda_3) = \ln 2$. If player 3 would decide to use a base stock level of one instead, then $\mathbb{E}B(1, \lambda_3) = \ln 2 - 1 + 0.5$, $\mathbb{E}I(1, \lambda_3) = 0.5$, and $K(1, \lambda_3) = \ln 2$. As $K(0, \lambda_3) = K(1, \lambda_3)$ and this cost function is strictly convex in the base stock level (which we shall prove later in Lemma 6.3), minimal costs are achieved with a base stock level of either 0 or 1, and $c^{\varphi}(\{3\}) = \ln 2$. In the remainder of this example, we will round values to four decimals for notational convenience.

If player 1 would join to form coalition {1, 3}, then it can be verified that a base stock level of one for their combined stock point is optimal; thus, $c^{\varphi}(\{1, 3\}) = K(1, \lambda_1 + \lambda_3) \approx 0.6980$. Under the proportional allocation scheme rule $\mathcal{P}$, player 3 would have to pay $\mathcal{P}_{3,\{1,3\}}(\varphi) \approx 0.6100$ in coalition {1, 3}, which is lower than $\mathcal{P}_{3,\{3\}}(\varphi) \approx 0.6931$ (see Table 1). This strict population monotonicity can be verified for the members of all other nested pairs of coalitions as well, implying that $\mathcal{P}(\varphi)$ is strictly population monotonic. Accordingly, the spare parts pooling game (N, c^{φ}) has a non-empty strict core containing $\mathcal{P}(\varphi)$.

However, the game's Shapley value $\Phi(N, c^{\varphi})$, which assigns $\Phi_1(N, c^{\varphi}) \approx 0.0556$, $\Phi_2(N, c^{\varphi}) \approx 0.4774$, and $\Phi_3(N, c^{\varphi}) \approx 0.4670$, is not in the core of this game because $\Phi_2(N, c^{\varphi}) + \Phi_3(N, c^{\varphi}) > c^{\varphi}(\{2, 3\})$. Accordingly, the characteristic cost function is not submodular; indeed, $c^{\varphi}(\{1, 3\}) - c^{\psi}(\{3\}) < c^{\phi}(\{1, 2, 3\}) - c^{\phi}(\{2, 3\})$. In other words, player 1's marginal cost contribution may increase if he joins a larger coalition.

In this example, cost allocation could be carried out in a stable and population monotonic way via the proportional rules. To show that this is not a coincidence, we will exploit various new analytical properties of our inventory model's cost function, which are derived in the next section.

6. Analysis of the underlying inventory model

In this section, we first provide a characterization of the optimal base stock levels and subsequently derive partial derivatives of the cost function K with respect to the demand rate. These intermediate results ultimately enable us to analyze how the cost performance under optimal base stock levels behaves as the demand rate varies on $\mathbb{R}_{++} = (0, \infty)$, which will enable us to prove stability and population monotonicity of the proportional allocations in Section 7. The holding and backorder cost rates, h and b, will remain fixed in the ensuing analysis.

We show in Lemma 6.3 that the optimal base stock levels are intricately related to the steady-state probability of having no backorders, $\mathbb{P}[B(S, \lambda) = 0]$. We therefore start by stating several properties of $\mathbb{P}[B(S, \lambda) = 0]$ in Lemmas 6.1 and 6.2. Although these properties are rather straightforward, we were unable to find a proof in the literature, and therefore we provide a proof in the appendix.

Lemma 6.1. *Let the demand rate $\lambda > 0$ be fixed.*

(i) $\mathbb{P}[B(S, \lambda) = 0]$ *is strictly increasing as a function of S (for S on $\mathbb{N}_0$).*
(ii) $\lim_{S \to \infty} \mathbb{P}[B(S, \lambda) = 0] = 1$.

Lemma 6.2. *Let the base stock level $S \in \mathbb{N}_0$ be fixed.*

(i) $\mathbb{P}[B(S, \lambda) = 0]$ *is differentiable as a function of λ (for λ on $\mathbb{R}_{++}$).*
(ii) $\mathbb{P}[B(S, \lambda) = 0]$ *is strictly decreasing as a function of λ (for λ on $\mathbb{R}_{++}$).*
(iii) $\lim_{\lambda \downarrow 0} \mathbb{P}[B(S, \lambda) = 0] = 1$ *and* $\lim_{\lambda \to \infty} \mathbb{P}[B(S, \lambda) = 0] = 0$.

The following lemma states that the cost function in our model is strictly convex in the base stock level and provides a standard characterization of the cost-minimizing base stock level (s) in terms of a newsvendor fractile. This characterization is illustrated in Fig. 1. Although convexity is relatively well-known for the inventory model under consideration (see, e.g., Zipkin, 2000, p. 215) we show *strict* convexity and address the uniqueness and multiplicity of optimal base stock levels more formally, which will facilitate our analysis. The proof of this and subsequent lemmas are given in the appendix.

[1] We are aware that these downtime costs are rather low relative to the holding costs. We chose these parameter values because they simultaneously yield a computationally convenient illustration of the costs and allocations involved, in addition to a game whose Shapley value lies outside the core.

Table 1
The spare parts pooling game and proportional allocation scheme of Example 5.1.

Coalition M	Optimal base stock levels	$c^{\varphi}(M)$	$\mathcal{P}_{1,M}(\varphi)$	$\mathcal{P}_{2,M}(\varphi)$	$\mathcal{P}_{3,M}(\varphi)$
{1}	0	0.1000	0.1000	*	*
{2}	1	0.6987	*	0.6987	*
{3}	0 and 1	0.6931	*	*	0.6931
{1, 2}	1	0.7132	0.0792	0.6340	*
{1, 3}	1	0.6980	0.0880	*	0.6100
{2, 3}	1	0.9428	*	0.5053	0.4375
N	1	1.0000	0.0628	0.5023	0.4349

Lemma 6.3. *Let the demand rate $\lambda > 0$ be fixed.*

(i) *$K(S, \lambda)$ is strictly convex as a function of S (for S on $\mathbb{N}_0$).*
(ii) *There is at least one $S \in \mathbb{N}_0$ for which $\mathbb{P}[B(S,\lambda) = 0] \geqslant b/(b+h)$; let S^* denote the smallest such S. Then, S^* is the unique optimal base stock level unless $\mathbb{P}[B(S^*,\lambda) = 0] = b/(b+h)$; in that case, both S^* and $S^* + 1$ (and no other) are optimal.*

We now introduce some additional notation: for any $n \in \mathbb{N}$ (note that $0 \notin \mathbb{N}$), we define ℓ_n to be the unique positive real number that satisfies

$$\mathbb{P}[B(n-1,\ell_n) = 0] = b/(b+h),$$

which is well-defined due to Lemma 6.2. Notice that, by Part (*ii*) of Lemma 6.3, for any $n \in \mathbb{N}$ it holds that ℓ_n is the demand rate for which both base stock levels $n - 1$ and n are optimal. The following lemma formally establishes several additional properties, which are illustrated in Fig. 1.

Lemma 6.4. *Let $n \in \mathbb{N}$. Then, the following statements hold.*

(i) *$\ell_n < \ell_{n+1}$.*
(ii) *For any λ in the interval (ℓ_n, ℓ_{n+1}), the unique optimal base stock level for the inventory system with demand rate λ is n.*
(iii) *$\lim_{n\to\infty} \ell_n = \infty$.*

The following lemma considers, for a fixed base stock level greater than zero, the expected steady-state costs per unit time as a function of the demand rate: the lemma states a simple expression for its derivative and shows that this cost function is strictly convex. This convexity is illustrated in Fig. 2.

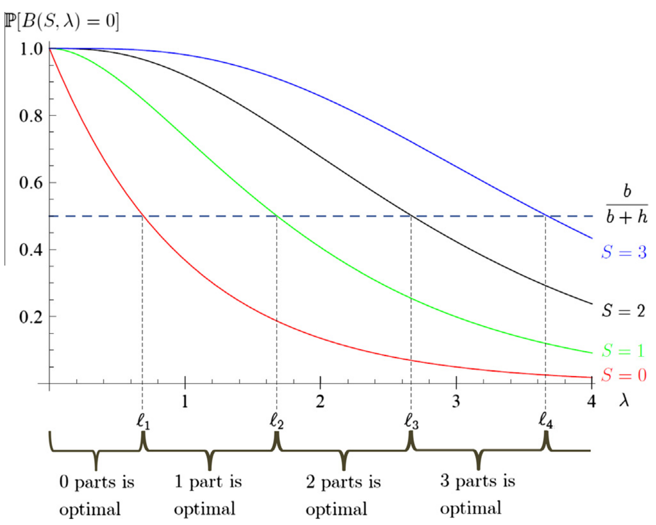

Fig. 1. The probability of having no backorders, $\mathbb{P}[B(S,\lambda) = 0]$, as a function of the demand rate λ for various base stock levels S. ($h = b = 1$.)

Lemma 6.5. *Let the base stock level $S \in \mathbb{N}$ be fixed.*

(i) $\frac{\partial}{\partial\lambda} K(S,\lambda) = b - (b+h) \cdot \mathbb{P}[B(S-1,\lambda) = 0]$.
(ii) *$K(S, \lambda)$ is twice differentiable and strictly convex as a function of λ (for λ on $\mathbb{R}_{++}$).*

The following lemma provides insightful expressions for the partial derivatives of the cost function with respect to the demand rate, evaluated at any ℓ_n (the demand rate at which both base stock levels n and $n - 1$ are optimal). In particular, when this cost function is considered as a function of λ, the tangent line to $K(n, \lambda)$ at $\lambda = \ell_n$ is flat and the tangent line to $K(n - 1, \lambda)$ at $\lambda = \ell_n$ goes through the origin, as illustrated in Fig. 2.

Lemma 6.6. *Let $n \in \mathbb{N}$ be an arbitrary positive integer.*

(i) $\frac{\partial}{\partial\lambda} K(n,\lambda)\big|_{\lambda=\ell_n} = 0$.
(ii) $\frac{\partial}{\partial\lambda} K(n-1,\lambda)\big|_{\lambda=\ell_n} = \frac{K(n-1,\ell_n)}{\ell_n}$.

We finally consider how the cost of an inventory system with optimal base stock levels behaves as the demand rate varies. To this end, we define the optimal cost function $\widetilde{K} : \mathbb{R}_{++} \to \mathbb{R}_{++}$ by

$$\widetilde{K}(\lambda) = \min_{S\in\mathbb{N}_0} K(S,\lambda). \quad (8)$$

This function is well-defined due to Part (*ii*) of Lemma 6.3.

Now, we say that a function $f : \mathbb{R}_{++} \to \mathbb{R}_{++}$ is *elastic* if $f(x_1)/x_1 \geqslant f(x_2)/x_2$ for all $x_1, x_2 \in \mathbb{R}_{++}$ with $x_1 \leqslant x_2$. Intuitively, if $f(x)$ expresses the cost of, say, serving demand level x, then elasticity of f says that the per-demand cost is non-increasing in the total demand served, i.e., f exhibits economies of scale. The use of the term "elasticity" is based on the economics literature, as motivated by Özen et al. (2011, p. 386). It is easy to prove that concavity of f implies elasticity (but not vice versa, as can be seen in Fig. 2). The following theorem states that the optimal cost function in our spare parts inventory model, $\widetilde{K}$, is elastic, as illustrated in Fig. 3.

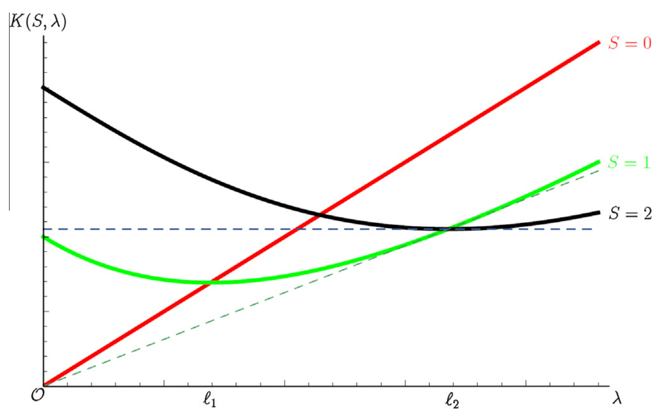

Fig. 2. The costs, $K(S, \lambda)$, as a function of the demand rate λ for various base stock levels S. Also shown (dashed) are the tangent lines to $K(1, \lambda)$ and $K(2, \lambda)$ at $\lambda = \ell_2$.

Theorem 6.7. *The function $\widetilde{K}$ is elastic. In particular, $\widetilde{K}(\lambda)/\lambda$ is constant for λ on $(0, \ell_1]$ and strictly decreasing for λ on $[\ell_1, \infty)$.*

Proof. First, for any $\lambda \in (0, \ell_1]$, a base stock level of zero is optimal, by the combination of Part (*ii*) of Lemma 6.2, Part (*ii*) of Lemma 6.3, and the definition of ℓ_1. Hence, for λ on $(0, \ell_1]$, $\widetilde{K}(\lambda) = K(0, \lambda) = b\lambda$, and thus $\widetilde{K}(\lambda)/\lambda$ is constant.

Next, let $n \in \mathbb{N}$. By Part (*ii*) of Lemma 6.4, for any $\lambda \in [\ell_n, \ell_{n+1})$, it holds that $\widetilde{K}(\lambda) = K(n, \lambda)$. We now fix $\tilde{\lambda} \in [\ell_n, \ell_{n+1})$ arbitrarily. Note that, by Part (*i*) of Lemma 6.4, $\tilde{\lambda} < \ell_{n+1}$. Using this, we find that

$$K(n, \tilde{\lambda}) > K(n, \ell_{n+1}) + (\tilde{\lambda} - \ell_{n+1}) \cdot \frac{\partial}{\partial \lambda} K(n, \lambda)\bigg|_{\lambda = \ell_{n+1}}$$
$$= K(n, \ell_{n+1}) + (\tilde{\lambda} - \ell_{n+1}) \cdot \frac{K(n, \ell_{n+1})}{\ell_{n+1}} = K(n, \ell_{n+1}) \cdot \frac{\tilde{\lambda}}{\ell_{n+1}}. \quad (9)$$

The inequality holds because any strictly convex, twice differentiable function – properties satisfied by $K(n, \lambda)$ as a function of λ for λ on $\mathbb{R}_{++}$, cf. Part (*ii*) of Lemma 6.5 – lies strictly above any of its tangent lines (except, of course, at the point where the tangent line touches the function's curve, but that does not pose a problem since $\tilde{\lambda} < \ell_{n+1}$). The first equality holds by Part (*ii*) of Lemma 6.6.

Using this, we obtain

$$\begin{aligned}
\frac{\partial}{\partial \lambda}\left(\frac{K(n,\lambda)}{\lambda}\right)\bigg|_{\lambda=\tilde{\lambda}} &= \left(\tilde{\lambda} \cdot \frac{\partial}{\partial \lambda} K(n,\lambda)\bigg|_{\lambda=\tilde{\lambda}} - K(n,\tilde{\lambda})\right) \Big/ \tilde{\lambda}^2 \\
&< \left(\tilde{\lambda} \cdot \frac{\partial}{\partial \lambda} K(n,\lambda)\bigg|_{\lambda=\ell_{n+1}} - K(n,\tilde{\lambda})\right) \Big/ \tilde{\lambda}^2 \\
&= \left(\tilde{\lambda} \cdot \frac{K(n,\ell_{n+1})}{\ell_{n+1}} - K(n,\tilde{\lambda})\right) \Big/ \tilde{\lambda}^2 \\
&< \left(\tilde{\lambda} \cdot \frac{K(n,\tilde{\lambda})}{\tilde{\lambda}} - K(n,\tilde{\lambda})\right) \Big/ \tilde{\lambda}^2 = 0.
\end{aligned}$$

The first inequality holds by Part (*ii*) of Lemma 6.5. The subsequent equality holds by Part (*ii*) of Lemma 6.6. The second inequality holds by Inequality (9). We conclude that $\widetilde{K}(\lambda)/\lambda$ is strictly decreasing for λ on $[\ell_n, \ell_{n+1})$.

As, by Part (*ii*) of Lemma 6.3, both n and $n+1$ are optimal base stock levels for demand rate ℓ_{n+1}, it holds that $\widetilde{K}(\ell_{n+1}) = K(n, \ell_{n+1}) = K(n+1, \ell_{n+1})$. Furthermore, it follows from Part (*ii*) of Lemma 6.5 that both $K(n, \lambda)$ and $K(n+1, \lambda)$ as functions of λ are continuous at $\lambda = \ell_{n+1}$. We conclude that $\widetilde{K}$ is continuous at ℓ_{n+1}.

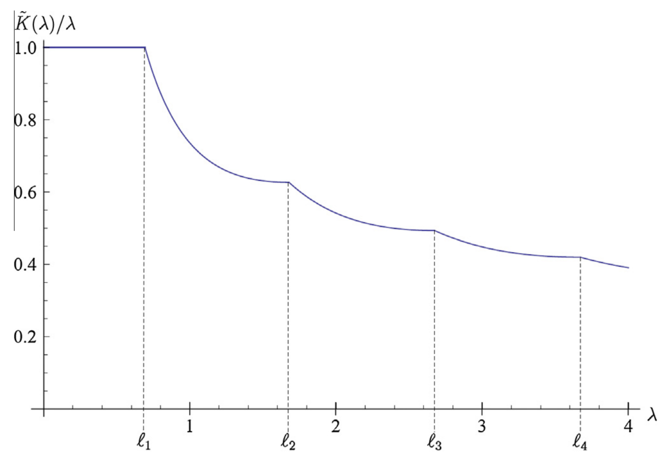

Fig. 3. The optimal per-demand costs, $\widetilde{K}(\lambda)/\lambda$. ($h = b = 1$).

To summarize, we have established that $\widetilde{K}(\lambda)/\lambda$ is non-increasing for λ on $(0, \ell_1]$ and that, for arbitrary positive integer n, this function is strictly decreasing on $[\ell_n, \ell_{n+1})$ and continuous at ℓ_{n+1}. Now, since by Part (*iii*) of Lemma 6.4 it holds that $\bigcup_{n\in\mathbb{N}}[\ell_n, \ell_{n+1}) = [\ell_1, \infty)$, this implies that $\widetilde{K}(\lambda)/\lambda$ is strictly decreasing for λ on $[\ell_1, \infty)$. Elasticity of $\widetilde{K}$ follows, and the proof is complete. □

7. Fair allocations of expected costs

We first give our main results for spare parts pooling games in Section 7.1. Subsequently, we study who reaps the benefits of cooperation in Section 7.2. Finally, we discuss connections with so-called single-attribute games and newsvendor games in Section 7.3.

7.1. Stability and population monotonicity

Using Theorem 6.7, we can now show that spare parts pooling games have a non-empty core and that the proportional allocation scheme rule $\mathcal{P}$ (see Eq. (7)) accomplishes a PMAS. As stated in the following theorem, the population monotonicity is strict if demand rates are sufficiently high for each coalition with two or more players to have an optimal base stock level greater than zero. To formally state this and later results, we let

$$S^*(\lambda) = \min\{S \in \mathbb{N}_0 : K(S, \lambda) = \widetilde{K}(\lambda)\}$$

denote the (smallest) optimal base stock level for any demand rate $\lambda > 0$.

Theorem 7.1. *Let $\varphi = (N, (\lambda_i)_{i\in N}, h, b)$ be a spare parts situation with backordering.*

(i) *The associated spare parts pooling game (N, c^φ) has a non-empty core containing $\mathcal{P}(\varphi)$, and $\mathcal{P}(\varphi)$ is a PMAS.*
(ii) *If $S^*(\lambda_L) > 0$ for each $L \in 2^N_-$ with $|L| \geqslant 2$, then $\mathcal{P}(\varphi)$ is an SPMAS and (N, c^φ) has a non-empty strict core containing $\mathcal{P}(\varphi)$.*

Proof. Part (*i*). We use a straightforward implication of $\widetilde{K}$'s elasticity, thereby specializing a known implication (see, e.g., Hamlen, Hamlen, & Tschirhart, 1977, p. 621, or Özen et al., 2011, Theorem 1) for general cost sharing problems to our spare parts pooling games. Let $M, L \in 2^N_-$ with $M \subset L$, and let $i \in M$. Then $\mathcal{P}(\varphi)$ is a PMAS because, by Theorem 6.7,

$$\mathcal{P}_{i,L}(\varphi) = c^\varphi(L)\frac{\lambda_i}{\lambda_L} = \widetilde{K}(\lambda_L)\frac{\lambda_i}{\lambda_L} \leqslant \widetilde{K}(\lambda_M)\frac{\lambda_i}{\lambda_M} = c^\varphi(M)\frac{\lambda_i}{\lambda_M}$$
$$= \mathcal{P}_{i,M}(\varphi). \quad (10)$$

Core inclusion of $\mathcal{P}(\varphi)$ follows from the closing sentence of Section 4.

Part (*ii*). For arbitrary $M, L \in 2^N_-$ with $M \subset L$, assume that $S^*(\lambda_L) > 0$. This implies that λ_L, the collective demand rate of coalition L, is strictly larger than ℓ_1, the demand rate for which both base stock levels 0 and 1 are optimal. Therefore, the inequality in (10) is strict by Theorem 6.7. Accordingly, $\mathcal{P}(\varphi)$ is an SPMAS, and $\mathcal{P}(\varphi)$ is strictly stable. □

Theorem 7.1 states an important result, because a proportional allocation rule is easy to understand and computationally attractive. Moreover, it satisfies the appealing property of immunity to manipulations of the players via artificial splitting and merging. This means that, if collective costs are divided proportionally according to the rule $\mathcal{P}$, no group of players will have an incentive to artificially represent themselves together as a single player, or vice versa. Indeed, their total cost allocation will remain the same because splitting or merging does not affect their sum of demand rates. So in a collaboration between, e.g., company A with a single business unit and company B with two business units, the costs assigned to company A by rule $\mathcal{P}$ are unaffected by whether the business units comprising company B claim they should be treated as one player together or two players separately. Finally, in line with a result in Karsten, Slikker, and van Houtum (2011a) for games in Erlang loss queues, $\mathcal{P}$ can be axiomatically characterized as the *unique* continuous rule satisfying this non-manipulability property, which strengthens its fairness as a allocation rule.

7.2. Who reaps the benefits?

One might wonder who actually reaps most of the benefits of the collaboration. With benefits we mean the absolute difference between the costs incurred by a player when acting alone and the cost assigned to this player under rule $\mathcal{P}$. Thus, the benefits (i.e., gains) for a player with demand rate $\lambda > 0$ when participating in a spare parts pool with aggregate demand rate $\Lambda \geqslant \lambda$ are given by

$$G(\lambda, \Lambda) = \widetilde{K}(\lambda) - \lambda \cdot \widetilde{K}(\Lambda)/\Lambda.$$

Now, will the smallest player (i.e., the player with the lowest demand rate) always gain the most, or will the largest player always take the lion's share? The following example shows that it could actually be neither of them. Thus, the proportional rule appealingly does not categorically favor smaller or larger players.

Example 7.1. Reconsider the spare parts situation with backordering φ as described in Example 5.1. For this situation, the players' benefits (rounded to four decimals) are $G_1 = G(\lambda_1, \lambda_N) = 0.0372$, $G_2 = G(\lambda_2, \lambda_N) = 0.1964$, $G_3 = G(\lambda_3, \lambda_N) = 0.2582$. Since $\lambda_3 < \lambda_2$ and $G_3 > G_2$, a small player might reap more benefits than a large player. Yet, since $\lambda_3 > \lambda_1$ and $G_3 > G_1$, a large player might reap more benefits than a small player as well. This is represented graphically in Fig. 4. We remark that if we would consider benefits in relative rather than absolute terms, then there are two players tied for highest per-demand benefit, but player 3 is still one of them: $G_1/\lambda_1 = 0.3725$, $G_2/\lambda_2 = 0.2453$, and $G_3/\lambda_3 = 0.3725$.

This example suggests that the largest benefits are typically reaped by a player with demand rate equal to ℓ_n for some $n \in \mathbb{N}$, i.e., a demand rate for which two base stock levels are optimal. In the following theorem, we show that this holds in general, provided that it is not optimal for the grand coalition to stock zero parts (in which case there would clearly be no pooling benefits at all). The proof is deferred to the appendix.

Theorem 7.2. *Consider a spare parts pool with total demand rate $\Lambda > \ell_1$. Then, there exists an $n \in \{1, \ldots, S^*(\Lambda)\}$ such that $G(\ell_n, \Lambda) > G(\lambda, \Lambda)$ for all $\lambda \in (0, \Lambda]$ with $\lambda \neq \ell_{n'}$ for some $n' \in \mathbb{N}$.*

From this theorem, we immediately obtain the following corollary, which concerns situations where the grand coalition optimally stocks a single part. This is quite common for low-demand, expensive spare parts.

Corollary 7.3. *Let $\varphi = (N, (\lambda_i)_{i\in N}, h, b)$ be a spare parts situation such that $S^*(\lambda_N) = 1$ and $\lambda_i = \ell_1$ for some $i \in N$. Then, $G(\lambda_i, \lambda_N) > G(\lambda_j, \lambda_N)$ for all $j \in N$ with $\lambda_j \neq \lambda_i$.*

7.3. An alternative proof approach

The fact that spare parts pooling games have a non-empty core with a proportional PMAS (Part (*i*) of Theorem 7.1) can be proven in a different way, via a connection with so-called single-attribute

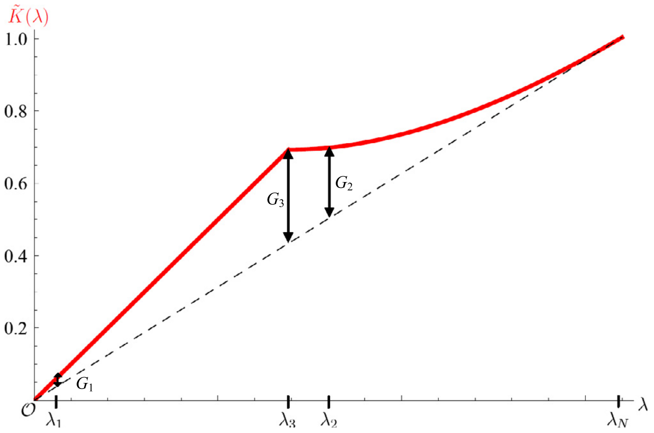

Fig. 4. The optimal costs $\widetilde{K}(\lambda)$ as a solid line, a dotted line through the origin with slope $\widetilde{K}(\lambda_N)/\lambda_N$, and the benefits or gains G_1, G_2, and G_3 for the players in Example 7.1.

games and newsvendor games, together with the contraposition of a recent result in Özen et al. (2011). This alternative proof approach, which is detailed in the appendix, employs a result that applies more generally: a certain class of newsvendor games admits a PMAS in which costs are assigned proportional to player's mean demands. However, this alternative proof approach does not provide insights into the structure of the problem as our analysis in Section 6 did. Moreover, our structural analysis in Section 6 allowed us to identify a *strictly* stable allocation (Part (*ii*) of Theorem 7.1) and the player who benefits most (Section 7.2); these additional results do not follow from the alternative proof approach.

8. A truth-inducing allocation process for cost realizations

The previous section showed that the proportional allocation rule satisfies the first two requirements posed in the introduction (Section 1). In this section, we treat the remaining two requirements. First, we show that the proportional allocation rule can be easily implemented in practice via a simple cost division per realization. Although our games have been formulated in expected terms to investigate a priori attractiveness of pooling, fair assignments of *realized* costs in any finite time period will be required to sustain cooperation in practice. To propose a method for assigning realized costs, we make the natural assumption of a First-In-First-Out (FIFO) stock discipline: whenever more than one part is available in the on-hand stock when a demand is placed, the demand is fulfilled by the oldest part in the on-hand stock. We now propose the following method to allocate costs as they materialize in an inventory system with any base stock level $S \in \mathbb{N}_0$ operated by the grand coalition in any spare parts situation with backordering $(N, (\lambda_i)_{i\in N}, h, b)$.

Process 8.1. Realized costs for the grand coalition are assigned as follows:

- Each player, upon placing a demand when the on-hand stock is positive, pays all holding costs incurred for the part taken (according to FIFO). That is, if the taken part was delivered at the stock point at a time τ and the player's demand occurs at time t, then this player pays $h(t - \tau)$ upon placing his demand.
- Each player, upon placing a demand that is backordered, pays all backorder costs incurred for this backorder. That is, if the demand occurs at time t and the associated backorder is later fulfilled via delivery of a new part at time τ, then this player will have to pay $b(\tau - t)$ over the duration of his backorder.

This process assigns the holding costs incurred for some part in an intuitively fair way to the player who directly benefits from the part. No holding costs are assigned to a player who does not benefit from on-hand inventory (as a result of not facing any demands over a certain time period or due to unfortunate stock-outs at his demand epochs). Additionally, since players fully incur all costs for their own backorders, the process eliminates the need for transfer payments of backorder costs, thereby avoiding disputes about their exact magnitude – this is an important property for the capital goods context wherein backorder costs typically comprise the downtime costs of a player's machine due to unavailability of a spare part. Moreover, as stated in the following lemma, the process can indeed be used to *implement*[2] the proportional allocation.

Lemma 8.1. *Under Process 8.1, the share of long-term average costs (of an inventory system operated by the grand coalition with base stock level $S \in \mathbb{N}_0$) borne by player $i \in N$ is λ_i/λ_N. In particular, if the grand coalition optimally stocks $S^*(\lambda_N)$ parts, the long-term average costs assigned to each player under Process 8.1 coincide with the assignment of expected costs under the proportional allocation rule $\mathcal{P}$.*

Proof. Follows directly from the well-known property that Poisson arrivals see time averages (Wolff, 1982). □

A final appealing property of Process 8.1 is that it removes any incentive for players to lie about their demand rates a priori. Although thus far we have adhered to the assumption of full and open information (a standard assumption in cooperative game theory), in reality a player's demand rate may be private information. During initial negotiations, when all cooperating players have to state their demand rate for the purpose of joint base stock level optimization, a player might lie if the collaboration would be implemented via an inappropriate cost realization assignment method. For example, under a method that charges each player a yearly fee based on his stated demand rates independent of that player's realized demand volume in that year, a player might have a reason to understate his actual demand rate a priori. However, under Process 8.1, truth telling is a Nash equilibrium (formally defined in Nash, 1951) in the non-cooperative information disclosure game in which each player in N has to state any demand rate in $\mathbb{R}_{++}$. In this game, the payoff to each player for any strategy profile $(\hat{\lambda}_i)_{i\in N}$, containing each player's stated demand rates, is equal to the long-term average costs assigned under Process 8.1 in the inventory system with base stock level $S^*\left(\sum_{i\in N}\hat{\lambda}_i\right)$.

Theorem 8.2. *The strategy profile $(\lambda_i)_{i\in N}$, in which each player $i \in N$ states his true demand rate λ_i, is a Nash equilibrium in this non-cooperative game.*

Proof. Consider a player $i \in N$ and suppose that all other players $j \in N\backslash\{i\}$ announce their true demand rate λ_j. By lying, i.e., stating any demand rate $\mathcal{L} \in \mathbb{R}_{++}$ other than λ_i, player i can only effect a possibly suboptimal base stock level since $K(S^*(\lambda_N), \lambda_N) \leqslant K(S^*\left(\sum_{j\in N\backslash\{i\}}\lambda_j + \mathcal{L}\right), \lambda_N)$ by definition of S^*. Yet, the fraction of long-term average costs assigned to player i under Process 8.1 is equal to λ_i/λ_N by Lemma 8.1, i.e., is independent of the demand rate that he states. Thus, player i minimizes his costs by stating λ_i. This completes the proof. □

In similar fashion, we can derive that Process 8.1 gives each player an incentive to immediately disclose any change in his expected demand rate, which is relevant for collaborations in a dynamic world where a player's number of installed machines may change over time or if forecasts improve.

9. Conclusion

We have studied the cost allocation problem in a spare parts inventory model with backordering. We have derived new structural properties of the resulting cost function, in particular concerning its behavior for varying demand rates, which may be relevant beyond the context of our games. Using these properties, we were able to show that the associated cooperative games have non-empty cores. We have further shown that the allocation of total expected costs proportional to each player's demand rate is stable, and that this cost allocation has appealing properties that enable easy implementation in practice. Indeed, it satisfies the four requirements posed in Section 1.

Our results have important managerial implications for companies facing high spare parts holding and/or downtime costs: For the model we considered, pooling is not only beneficial from the whole system's point of view, but can also be supported by a stable cost allocation. This means that inventory pooling, which is already commonly exploited in the case of a single player who owns all the parts and all the demand streams, can also be achieved if there are multiple, self-interested players. By using our easy-to-implement proportional allocation rule, these players can be assured that everyone will get strictly better off. This should pave the way for sustainable collaborations in practice.

One limitation to the practical implications of our findings is our assumption of symmetric backlogging costs. In practice, due to differences in downtime costs or service level agreements with customers, backlogging costs may be asymmetric. In those cases, discussed Appendix A.1, full pooling may actually be detrimental. For further research, it would be interesting to investigate when the assumption of symmetric backlogging costs is justified and how the FIFO stock discipline might be adjusted for asymmetric backlogging costs.

Another avenue for future research is to extend the model to two echelon levels: a one warehouse, multiple retailers setting. The spare parts at the central warehouse may be owned by a coalition of retailers, or by a third party. If this third party is the original equipment manufacturer, then it may also be interesting to allow this party to exert additional design effort to improve component reliability. This may be beneficial from the whole system's point of view, but it also raises the question of what share of the benefits the manufacturer is entitled to. Cooperative game theory may provide the tools to determine (existence of) fair allocations of collective costs.

Acknowledgements

The authors thank Marco Slikker, Geert-Jan van Houtum, Ulaş Özen, and two anonymous referees for valuable suggestions and comments. The second author gratefully acknowledges the support of the Lloyds Register Foundation (LRF). LRF supports the advancement of engineering-related education, and funds research and development that enhances safety of life at sea, on land and in the air.

Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at http://dx.doi.org/10.1016/j.ejor.2013.08.029.

[2] Here, we assume that any player takes a part from the pool if *and only if* he faces a component failure. Opportunistic behavior (not demanding a part upon a failure, or demanding a part before a failure occurs) may be prohibited by requiring a failed part in exchange for any part taken from the inventory and/or by hiding the pool's actual on-hand inventory level from the players.

References

Anily, S., & Haviv, M. (2007). The cost allocation problem for the first order interaction joint replenishment model. *Operations Research, 55*(2), 292–302.

Chen, X., & Zhang, J. (2009). A stochastic programming duality approach to inventory centralization games. *Operations Research, 57*(4), 840–851.

Cohen, M. A., Agrawal, N., & Agrawal, V. (2006). Winning in the aftermarket. *Harvard Business Review, 84*(5), 129–138.

Dror, M., Guardiola, L. A., Meca, A., & Puerto, J. (2008). Dynamic realization games in newsvendor inventory centralization. *International Journal of Game Theory, 37*, 139–153.

Dror, M., & Hartman, B. C. (2011). Survey of cooperative inventory games and extensions. *Journal of the Operational Research Society, 62*(4), 565–580.

Feeney, G. J., & Sherbrooke, C. C. (1966). The $(s-1, s)$ inventory policy under compound Poisson demand. *Management Science, 12*(5), 391–411.

Fiestras-Janeiro, M. G., García-Jurado, I., Meca, A., & Mosquera, M. A. (2011). Cooperative game theory and inventory management. *European Journal of Operational Research, 210*(3), 459–466.

Flint, P. (2006). Plenty of Life in the Cycle. <http://atwonline.com/airline-financedata/article/plenty-life-cycle-0309> Accessed 24.06.13.

Gillies, D. B. (1959). Solutions to general non-zero-sum games. In A. Tucker & R. Luce (Eds.), *Contribution to the theory of games IV. Annals of mathematics studies* (Vol. 40, pp. 47–85). Princeton University Press.

Guajardo, M., Ronnqvist, M., Halvorsenb, A. M., & Kallevik, S. I. (2012). Inventory management of spare parts in an energy company. Working paper, Norwegian School of Economics.

Hale, J. (2011). Dutch call on EU to pool, share capabilities. <http://www.defensenews.com/story.php?i=6951686> Accessed 24.06.13.

Hamlen, S. S., Hamlen, W. A., Jr., & Tschirhart, J. T. (1977). The use of core theory in evaluating joint cost allocation schemes. *The Accounting Review, 52*(3), 616–627.

Hartman, B. C., & Dror, M. (1996). Cost allocation in continuous-review inventory models. *Naval Research Logistics, 43*(4), 549–561.

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

Karsten, F., Slikker, M., & van Houtum, G. J. (2011a). Analysis of resource pooling games via a new extension of the Erlang loss function. BETA working paper 344, Eindhoven University of Technology.

Karsten, F., Slikker, M., & van Houtum, G. J. (2011b). Resource pooling and cost allocation among independent service providers. BETA working paper 352, Eindhoven University of Technology.

Karsten, F., Slikker, M., & van Houtum, G. J. (2012). Inventory pooling games for expensive, low-demand spare parts. *Naval Research Logistics, 59*(5), 311–395.

Kemahlıoğlu-Ziya, E., & Bartholdi, J. J. III, (2011). Centralizing inventory in supply chains by using Shapley value to allocate the profits. *Manufacturing & Service Operations Management, 13*(2), 146–162.

Kranenburg, A. A. (2006). Spare parts inventory control under system availability constraints. PhD thesis, Eindhoven University of Technology.

Kukreja, A., Schmidt, C. P., & Miller, D. M. (2001). Stocking decisions for low-usage items in a multilocation inventory system. *Management Science, 47*(10), 1371–1383.

Muckstadt, J. A. (2005). *Analysis and algorithms for service parts supply chains*. New York: Springer.

Müller, A., Scarsini, M., & Shaked, M. (2002). The newsvendor game has a nonempty core. *Games and Economic Behavior, 38*(1), 118–126.

Nash, J. (1951). Non-cooperative games. *Annals of Mathematics, 54*(2), 286–295.

Norde, H., Özen, U., & Slikker, M. (2011). Setting the right incentives for global planning and operations. Working paper.

Özen, U., Fransoo, J., Norde, H., & Slikker, M. (2008). Cooperation between multiple newsvendors with warehouses. *Manufacturing & Service Operations Management, 10*(2), 311–324.

Özen, U., Reiman, M. I., & Wang, Q. (2011). On the core of cooperative queueing games. *Operations Research Letters, 39*(5), 385–389.

Palm, C. (1938). Analysis of the Erlang traffic formulae for busy-signal arrangements. *Ericsson Technics, 5*, 39–58.

Shapley, L. S. (1953). A value for n-person games. In H. Kuhn & A. Tucker (Eds.), *Contribution to the theory of games II. Annals of mathematics studies* (Vol. 28, pp. 307–317). Princeton University Press.

Sherbrooke, C. C. (2004). *Optimal inventory modeling of systems: multi-echelon techniques* (second ed.). Dordrecht: Kluwer Academic.

Slikker, M., Fransoo, J., & Wouters, M. (2001). Joint ordering in multiple news-vendor situations: A game-theoretical approach. BETA working paper 64, Eindhoven University of Technology.

Sprumont, Y. (1990). Population monotonic allocation schemes for cooperative games with transferable utility. *Games and Economic Behavior, 2*(4), 378–394.

van Houtum, G. J., & Zijm, W. H. M. (2000). On the relationship between cost and service models for general inventory systems. *Statistica Neerlandica, 54*(2), 127–147.

Wolff, R. W. (1982). Poisson arrivals see time averages. *Operations Research, 30*(2), 223–231.

Wong, H., Cattrysse, D., & van Oudheusden, D. (2005). Stocking decisions for repairable spare parts pooling in a multi-hub system. *International Journal of Production Economics, 93*, 309–317.

Wong, H., van Houtum, G. J., Cattrysse, D., & van Oudheusden, D. (2006). Multi-item spare parts systems with lateral transshipments and waiting time constraints. *European Journal of Operational Research, 171*(3), 1071–1093.

Wong, H., van Oudheusden, D., & Cattrysse, D. (2007). Cost allocation in spare parts inventory pooling. *Transportation Research Part E, 43*(4), 370–386.

Zhao, J. (2001). The relative interior of the base polyhedron and the core. *Economic Theory, 18*(3), 635–648.

Zipkin, P. H. (2000). *Foundations of inventory management*. McGraw-Hill.