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
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