Managing Nuclear Spare Parts Inventories: A Data-Driven Methodology

Natalie M. Scala, Jayant Rajgopal, Kim LaScola Needy

Accepted Manuscript (AM) / Postprint

The Version of Record of this manuscript has been published and is available in *IEEE Transactions on Engineering Management.* The article was published online on October 17, 2013.

Journal Website: https://ieeexplore.ieee.org/abstract/document/6636060

Article DOI: 10.1109/TEM.2013.2283170

Please note: This is the author accepted manuscript version of the research. The final version was proofread and released for copyright. The most up-to-date version of the paper can be accessed via the DOI address above.

Title:

Managing Nuclear Spare Parts Inventories: A Data-Driven Methodology

Abstract:

This paper presents a methodology for developing a spare parts inventory management system with a focus on the nuclear power sector. Often, demand for spare parts is highly intermittent and cannot be accurately forecasted through traditional methods. Examples include nuclear power generation equipment, ground space systems, and aircraft engine parts. We take a data-driven engineering management approach and develop a four-step methodology for spare parts management in such environments. These steps comprise an influence diagram for identifying relevant factors, weighting of influences through the Analytic Hierarchy Process, grouping parts according to inventory criticality indices, and the development of base stock inventory policies for each group. This approach allows the system to be actively managed within a continuous improvement framework through employee engagement and input, and mathematical assumptions are not made in the models. To our knowledge, no such integrated, comprehensive methodology for spare parts has been developed. The techniques employed in this research can be effectively used together to holistically manage the entire spare parts process, or they may be used separately to manage portions of the process. This paper provides an overview of the methodology, and the entire approach is illustrated via a test bed nuclear power generation facility.

Managerial Relevance Statement:

This research presents a methodology for spare parts management under conditions of intermittent demand and lack of detailed statistical data on equipment failure rates. No mathematical assumptions are made, which allows for the approach to be generalized to other data sets and conditions. Personnel with or without technical backgrounds can understand and participate in the model development, which increases corporate buy-in and encourages successful implementation. The methodology balances the risk of lost revenues with costs, while not compromising on safety, which is a necessity in a deregulated electric utility generation environment. It has been applied to spare parts at a nuclear power plant, and the outcomes support efficiency and cost effective inventory management. Managers at the test bed facility estimate that application of the approach would lead to a reduction of 18% of the nuclear generation spares inventory.

Introduction and Background

This research addresses the problem of effectively managing spare parts inventory in the context of the nuclear electric utility industry. This problem has several unique characteristics that make it difficult, if not impossible, to directly apply traditional inventory control methods, and we therefore develop a systematic four-step methodology that is designed to address these characteristics. The entire approach is illustrated via a case study from a nuclear power generation facility. Examples of other classes of spare parts that also exhibit these characteristics include ground space systems (antennas, etc.) and aircraft engine parts.

Operational spare parts do not usually directly service customers, and therefore excess inventory is undesirable from a management perspective. On the other hand, a stockout of spare parts can lead to offlining a production process and/or lost sales, with significant resulting costs. In general, when the cost of consequences associated with failure is high, and one cannot accurately forecast when and in what quantities requirements will arise, large quantities of parts tend to be stockpiled in inventory. In the nuclear sector large inventories are often justified as essential for safety. However, in a deregulated market, lost revenue and costs are the inventory drivers, as safe operation of the plant is taken as a given. In fact, safety systems at modern plants are very sophisticated, either reducing power output or shutting down a plant in danger. In practice, operational spares exist in a deregulated market to ensure that the plant never shuts down for unplanned reasons, because the resulting costs and lost revenues can be very high. Given this fact, the industry can benefit from a rigorous and systematic spare parts inventory management methodology.

This paper addresses the management of spare parts inventory in environments such as the one described above. To test the approach proposed herein, we use the case of electric utilities that operate a diverse set of power generation and transmission assets in a dynamic, deregulated environment. Some form of deregulation is or has been active in twenty-two states and the District of Columbia [52]. Deregulation began in the 1990s through the Energy Policy Act and FERC Order 888 [17], [19]. Specifically, in a deregulated system the generation portion of the utility business is market-driven, but the transmission and distribution portions remain regulated by the state in which the utility is operating to ensure that power is equitably delivered to all consumers. Before deregulation, electricity prices were set after legal hearings and negotiation between the utilities and state regulatory bodies; a traditional market did not exist. The regulated electricity rate included full cost recovery plus a rate of return that could be as high as 10% [36]. We develop a four-step methodology for a deregulated environment and illustrate it through a case study at a nuclear plant in a Fortune 200 utility company's generation portfolio.

A prerequisite for a good inventory control system is an accurate characterization of demand. However, because spares are typically used intermittently, this can be a challenge. In particular, the nuclear industry experiences part demands very sporadically, and parts might be demanded only a couple of times over a multi-year period. This precludes the use of typical time-series based tools to forecast demand. There has been some research that focuses specifically on forecasting with intermittent demand, such as Croston [13], Syntetos and Boylan [47], Boylan and Syntetos [10], and Altay, Rudisill and Litteral [5]. However, if part demand is extremely sparse as in the case with nuclear spares, these methods do not work well, and a different approach is needed to address management of such parts.

One alternative might be a causal model to characterize demand as a function of some suitable set of drivers that are treated as the independent variables [31]. However, this requires accurate and detailed data on part reliabilities and specific causes for their failures; as noted in [29], such information is not commonly available. Detailed failure rates are typically not tracked accurately; rather the industry uses an aggressive preventive maintenance (PM) schedule as a way to thwart failures. While the PM schedule prescribes when to perform maintenance, it does not specify how involved the work might be. Wear and tear may vary between maintenance sessions, and parts could be in varying conditions, leading to both considerable variability in the amount of maintenance work required and uncertainty in the corresponding spare parts needed. In addition, nuclear parts are generally unique, customized, and engineered-to-order with long lead times and restrictive vendor returns. Maintenance technicians often do not know the extent of work to be done until the job commences, and as was common in the regulated era, overcompensation occurs by ordering lots of additional parts up front so that any degree of repair could be addressed.

Research on spare parts under the above conditions is limited. The seminal survey on spare parts [29] provides references to literature up to the early 2000s and identifies future research needs, including the impact of increased technology and better prediction of spare parts demand. However, that and other papers on spare parts that review the literature (e.g., [26], [32]) do not point to work that addresses the class of parts that satisfy all the conditions outlined above. Other relevant papers of note include Bachman [6], who examined the tradeoffs of wait time, order frequency, and inventory value but not risk and safety, and a series of papers by Wang and co-authors [55–58], [60], which examine nuclear spare parts and use an artificial neural network to formulate an inventory system based on economic order quantities (EOQ). However, the EOQ assumes steady demand, and when demand is sporadic (such as for the parts under consideration), it is not an appropriate option. Cavalieri, Garetti, Macchi, and Pinto [11] develop a decision making framework for spare parts that relies on forecasts of part demand, which is not an option with nuclear spares. Finally, in a very recent paper, Molenaers, Baets, Pintelon, and Waeyenbergh [30] develop a multicriteria spare parts classification model that considers equipment criticality, probability of failure, replenishment time, number of suppliers, technical specifications, and maintenance type; they apply the model to a petrochemical plant. This work has several parallels to ours in that it also uses a multi-criteria approach to criticality classification based on the Analytic Hierarchy Process (AHP), which is similar to the second and third steps of our approach. However, the criteria for a petrochemical plant are different from those in the nuclear sector. Furthermore, the authors do not explicitly consider intermittent demand and use logic decision diagrams rather than influence diagrams. More importantly, the focus is on computing criticality scores; explicit inventory control policies based on the criticality indices are not developed.

The four-step approach described in this paper is general enough so that other utility plants, especially those in the nuclear sector, will be able to apply it. In addition to regulated utilities, other industries, such as aerospace and the military, could also benefit from this methodology. In general, many companies, especially utilities, need to balance new growth, modernization, obsolescence, reliability, and regulatory requirements. A model that addresses these in the context of spare parts inventory has the potential to greatly improve operations and strengthen companies' bottom lines.

Overview of the Nuclear Power Sector Environment

Although deregulation has changed the environment in which companies operate, a reengineering of the associated business processes has not been quick to occur. In particular, deregulation requires a shift in business processes and a redefinition of the management philosophy for spare parts. Under regulation, spare parts costs could readily be recovered by passing them to customers in the form of higher electricity rates negotiated with the states, but no such guarantee of recovery exists under the deregulated model, as spare parts are considered a generation expense. Furthermore, spare parts held in inventory tie up capital that could be spent on other company initiatives, such as investment in infrastructure.

On the other hand, a balance must also be maintained between reduced levels of spares and the risk of significant loss of revenue. For parts used at so-called LCO (limited condition of operation) locations, a failure compromises plant safety and must be remedied within a relatively short time (usually 72 hours), or the plant must be offlined or derated. (Once up and running, nuclear plants typically operate at 100% capacity, and a "derate" reduces the plant's output to some fraction of its capacity.) In a deregulated environment, loss of generation output translates to a loss of revenue, which impacts profitability. Thus, holding strategic inventory can definitely hedge against the possibility of significant revenue loss. While spare parts management could be addressed internally as a series of operational process initiatives, a more strategic issue is at hand; Scala [42] argues that a need exists for better understanding of what parts to order, when they should be ordered, and in what quantities.

Because of the inability to accurately quantify demand for spare parts, addressing the objective of preventing stockouts when responding to part failure or a PM request is a challenge. The fundamental tradeoff is between reducing capital investments in spare parts inventory and having to postpone maintenance or repair work due to unavailability of parts, which in turn implies a potential loss of revenue. In a competitive environment, it is crucial that any policy must mitigate the risk of stockout and consequent revenue loss, while ensuring the safe operation of the plant and minimizing parts on hand.

Given the challenges inherent in this problem, this paper develops a four-part decision analysis methodology for managing spare parts. The methodology defines and calculates the inventory criticality of a part while also considering the risk of stockout in order to improve inventory management. It has high practical value because a decision-analytic, engineering management approach is taken, which allows for personnel with non-technical backgrounds or little inventory modeling experience to actively contribute to the model development in a meaningful way. The methodology is valuable to both the body of knowledge related to spare parts management as well as the growing research on deregulated electric utilities.

Methodology

This research takes an engineering management approach to spare parts inventory management and is focused on practice in the deregulated energy sector. It incorporates all relevant factors and forces in the spare parts process, while eliminating the need for theoretical assumptions that might be required to build a traditional mathematical model. The approach develops a part scoring system that rates the importance of keeping an individual part in inventory, which is used to develop criticality groups. While this is analogous to a traditional A-B-C analysis for part classification, the usual criterion of dollar-volume is inappropriate in the environment being addressed. The literature has argued that many criteria are relevant when classifying inventory [20], [21], [33], [34], [38], [61]. There are also several inventory classification papers in the literature, including [7], [22], [49], as well as inventory management review papers that discuss classification [26], [37]. Textbooks, such as [23], present A-B-C policies in detail. Our new scoring system provides an alternative for spare parts classification. The approach then develops a modified base stock inventory stocking policy that is based on a retrospective simulation approach and balances the risk of revenue loss against capital investments in inventory. The following four-step methodology is used:

1. Development of an influence diagram of the spare parts process to identify relevant factors.
2. Use of the Analytic Hierarchy Process (AHP) to rank influences.
3. Development of a part criticality scoring system for spare parts classification.
4. Construction of inventory policies for classes of parts based on retrospective simulations.

Each step of the methodology is detailed in the following sections.

Step 1: Development of an influence diagram of the spare parts process to identify relevant factors

An influence diagram is a standard engineering management technique and is commonly used to identify and understand a problem by pictorially depicting all influences that are relevant. Square boxes depict influences that are deterministic or decisions while ovals or circles depict influences that are stochastic or uncontrollable. A directed arc between two nodes denotes that one node influences the other. Influence diagrams are easy to construct, and because they are visual, the decision maker can easily understand all the inputs to the problem, regardless of technical background. As a result, the diagrams are commonly developed as a collaborative effort between the decision maker and an analyst. For a detailed discussion of the theory of influence diagrams, see [24]. Influence diagrams have been widely used, and examples in the literature include [3], [15], [16], [27], [35], [45].

A full understanding of the problem and process is necessary for the developed model to address and incorporate all variables relevant to spare parts management. An influence diagram maps these variables and the relationships between them. It captures corporate knowledge of the process, is crucial to the engineering management approach, and serves as a basis for continuous improvement. For this research, the influence diagram depicts the current spare parts process and identifies all influences that are relevant to the process, providing a clear picture of the "as-is" process. Because spare parts are more than just a supply chain consideration, the diagram also allows for input from other plant departments, such as maintenance and planning.

The influence diagram was developed through an interactive and iterative process with the subject matter experts (SMEs) at the nuclear power generation case study facility. A list of possible, observed influences to the process was defined and presented to the SMEs who removed non-relevant influences while adding others. The influences were then consolidated into logical sets, with the SMEs providing verification and validation, so that all influences associated with a particular knowledge area were placed into a single set. Grouping influences into sets is common; for an example the reader is referred to [3]. This process of influence definition and placement into sets was repeated iteratively until all SMEs could agree with the process influence list and corresponding set placements. Discussions with the SMEs took place via conference calls and email correspondence. A total of five SMEs participated in the validation of the influence diagram; four of the SMEs held supervisory or managerial positions.

Overall, 34 influences were identified in the spare parts process and subsequently grouped into seven sets: Timeliness of Work Order, Part Failure, Vendor Availability, Part Usage in Plant, Preventive Maintenance Schedule, Outage Usage, and Cost Consequences. Each set focuses on a common theme. Details of each set including a listing of the influences placed within each one can be found in Scala [42]. Each individual influence was assigned to one set, and Figure 1 shows the diagram for the overall sets of influences. Each influence set generates a subdiagram, and an example (Set 5: Preventive Maintenance Schedule) is also shown in Figure 1. The other influence sets connect to the main diagram in a similar fashion.

The cultures at nuclear electricity generation facilities are very risk averse. Although various companies operate nuclear reactors for electricity generation, the work ethic and basic procedures are consistent across the United States nuclear fleet. Thus, an influence diagram of the nuclear spare parts process is applicable to the industry at large and can be used with other decision making processes related to inventory and maintenance. Furthermore, when generalizing the methodology, the process for the nuclear spare parts can be easily followed.

Step 2: Use of the AHP to rank influences

The second step of the methodology is to rank the influences on the diagram, which is achieved through group decision making in the AHP. The AHP is a decision analysis tool developed by Thomas Saaty [39], [40]. It is a widely used and popular method that has been extensively explored in the literature. Examples of applications include politics, technology, marketing, material handling, conflict resolution, and medicine and are summarized in [46], [53], [54], [59]. The AHP uses pairwise comparisons between criteria with respect to the goal and between alternatives with respect to each criterion. Analytically synthesizing the comparisons yields normalized prioritized alternatives; the normalized values can also be used as weights beyond a rank ordered list.

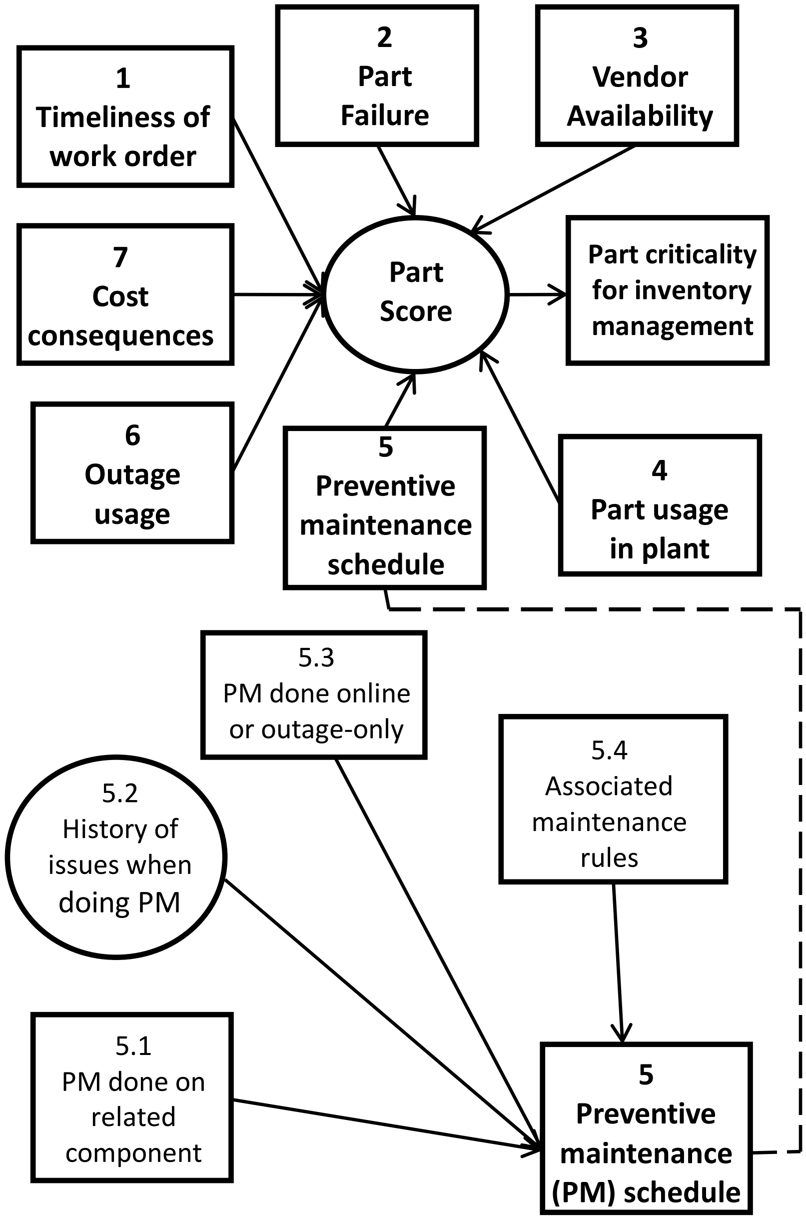

Figure 1. High level influence diagram with subdiagram detail for set 5 (Adapted from [42], [43])

The AHP is used to determine relative importance of each influence in every influence set on the diagram. The resulting priorities from synthesizing the pairwise comparisons then serve as weights for each influence. In this research, each set of influences was set up in its own two-level hierarchy, with a goal of spare parts analysis and the set influences in the second level. Figure 2 depicts the hierarchy for the "PM Schedule" set of influences.

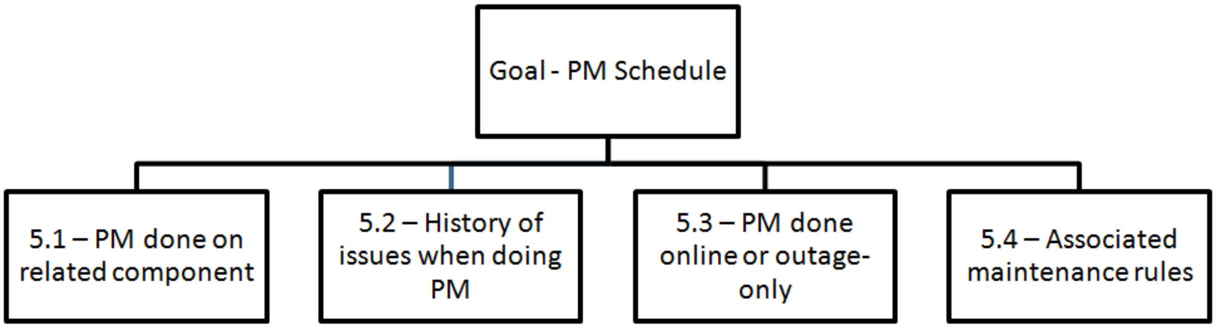

Figure 2. AHP hierarchy for the PM Schedule influence set

The pairwise comparisons were then done with respect to the goal. Thus, for PM Schedule all six possible pairwise comparisons were performed between the four influences in the set. To accurately represent the full nuclear spare parts process beyond just the storage of parts, multiple judgments of the pairwise comparisons were elicited; in particular, five unique industry SMEs performed the comparisons for the influences in each set. These individuals are experts in the influence set knowledge area to which they were assigned and were geographically dispersed at different locations within the parent company's footprint. The judgments were collected individually from each SME and each set to avoid any bias and ensure that one individual's judgments were not influenced by those of another. Furthermore, because industry SMEs from various nuclear plants (not just the test bed facility) were interviewed, the individual interviews allowed for full range of knowledge and unique perspectives to be collected. After the interview, the judgments were reviewed, and if the AHP inconsistency was too high (greater than 0.20), a follow-up interview was done with the SME. In that session, SMEs were asked to revisit and revise judgments until the inconsistency ratio fell below 0.20. Saaty [39], [40] suggests an acceptable inconsistency ratio at or below 0.10; however, because these judgments were to be combined, a higher individual inconsistency ratio was empirically allowed. During the follow-up call, although some SMEs did choose to revise judgments, others chose not to revise, and not all sets of pairwise comparisons fell below our inconsistency threshold.

In all, 25 SMEs performed comparisons, with five sets of pairwise comparisons collected for each influence set. The seven sets themselves were also compared by five SMEs. Some SMEs had knowledge of multiple areas and performed pairwise comparisons on more than one set. For an example of employee responses (original and revised) as well as a detailed interview protocol, the reader is referred to Scala [42].

Because multiple judgments were made, for each influence set the pairwise comparisons must be aggregated into one set of judgments that are representative of the five SMEs. This aggregated set of judgments can then be synthesized to weight the influences. Traditionally, a weighted or non-weighted geometric mean is used to aggregate judgments in the AHP [1], [2]. However, recent work by Saaty and Vargas [41] recommends a rigorous statistical test to ensure the aggregated group judgments do not have excessive dispersion around the geometric mean. In that case, the pairwise comparisons would be homogeneous in some responses and heterogeneous in others, which violates the axioms of the AHP. Dispersion tests were thus performed on the eight groups of pairwise comparisons—the seven influence sets as well as the overall set of influences. The test was performed for every set of pairwise comparisons; for example six separate tests were performed for the PM schedule set in Figure 1. For these data, the Saaty and Vargas [41] dispersion test failed for most influence comparisons (except for set 7), implying excessive dispersion. These comparisons therefore cannot be aggregated using a simple *non*-weighted geometric mean in this situation.

When this happens, the literature directs the decision maker to go back to the SMEs and ask them to revise their judgments or work together to reach consensus [1], [2], [8], [9], [41]. However, because in our case the decision makers were geographically dispersed, it was not feasible to gather all the SMEs in one location to force consensus. Each SME had a separate follow-up interview, but most of them were unwilling to revise their original judgments. Thus, the dispersion could not be reduced to an acceptable level. To aggregate the judgments, a *weighted* geometric mean was therefore used so as to account for the inherent variability associated with each judge. The corresponding weights were derived by a novel approach using principal components analysis (PCA). This method does not violate the axioms of the AHP and allows for variance around the geometric mean to be accounted for through the first principal eigenvector. In the PCA method, each decision maker is a variable or "dimension," and each judgment is an observation. For example, for the PM Schedule set, the PCA matrix will have five columns (one for each decision maker) and six rows (one for each pairwise comparison); thus we have five dimensions with six sets of observations across each. The first principal eigenvector of the corresponding 5×5 covariance matrix is then computed, with the result being a vector of five weights, one for each decision maker to be used when combining the judgments through a weighted geometric mean. The first principal eigenvector captures maximum variability by definition, and the resulting weights sum to one. Further details of this method can be found in Scala [42].

Once the judgments of the five SMEs were aggregated into one representative group judgment, the AHP priorities were found for all influences by synthesizing the aggregated judgments in each set. Ranking the influences allows for a weight to be assigned to each influence. This weight depicts the relative importance of that influence with respect to the other influences in the set and is the primary benefit of the AHP process employed in this methodology. A secondary benefit is that the relative weights are valuable to the test bed company beyond spare parts management. The nuclear workforce is very experienced, with a recent study citing eleven to fifty percent soon eligible to retire [51]. Capturing the workforce's extensive knowledge of the process is not only important but also invaluable to the industry. Thus, this approach formally captures the as-is nuclear spare parts process and how employees approach spare parts decisions as well as the corresponding risk of those decisions, which will mitigate the loss of knowledge due to employee attrition. This approach can be repeated to capture knowledge at other companies; here, the test bed company is used to illustrate the approach. This knowledge from decision makers can be applied to other company decisions and passed down to a new generation of nuclear employees to assimilate them into corporate culture.

Step 3: Development of a part criticality scoring system for spare parts classification

The third step of the methodology is the development of part criticality scores to identify the importance of keeping a part on hand in inventory. *Inventory criticality* should not be confused with *engineering criticality*; the latter is a standard nuclear industry classification of part maintenance schedule. Engineering criticality is not tied to inventory and prescribes if a part should receive preventive maintenance or if it is a run-to-failure part. Inventory criticality purely considers the need to keep parts on the shelf. Here inventory criticality scores are built from the AHP prioritization weights and historical part data. The developed scores are then used to classify parts into criticality groups for inventory management. Using the AHP is common in developing classifications and is addressed or demonstrated in [7], [20], [22], [34], and [37]. At the test bed company, part data is stored in its Enterprise Resource Planning (ERP) system and captures each part's use, maintenance schedule, etc. Each influence can be supported by various data fields, and the data in those fields are incorporated into the part criticality score. To provide consistency across the scores, the part data for each influence is converted to a unitless 1 to 5 scale. This ensures that various ranges and scales of data are considered uniformly. An example of a scale conversion is shown in Table 1 for influence 5.1: "PM done on related component."

The related part data in this instance is a count of preventive maintenance work orders on which the part was requested. A higher count implies more activity, implying increased importance or criticality for the part. An analysis of the ERP data showed a maximum of 84 preventive maintenance work orders across a sample of parts at the test bed company. Thus, the range of 0 to 85 (rounded) work orders must be mapped to the 1-5 scale, as shown in Table 1. The

"Assigned Scale Value" is the dimensionless ordinal value, with a value of 5 implying highest criticality. The "Low Value Based on Part Data" corresponding to an "Assigned Scale Value" is the lowest count of work orders assigned to that value, while the "High Value Based on Part Data" is the highest count of work orders assigned to it. The mapping of data to a scale can be done in a variety of ways; in this case, the data was mapped with feedback and validation from the test bed company SMEs.

Table 1. Assignment of part data to scale for influence 5.1

Low Value Based on Part Data	**High Value Based on Part Data**	**Assigned Scale Value**
71	85	5
41	70	4
13	40	3
1	12	2
0	0	1

In general, all scale units 1-5 do not need to be assigned to the corresponding part data for an influence. The scale is dimensionless and unitless, and part data is mapped as appropriate. The mapping is inherently subjective; however, assistance from company SMEs will help to determine the most critical data (assigned to scale value 5) and the least critical data (assigned to scale value 1). If part data does not exist for an influence, then a value of 0 is assigned for all parts, with the recommendation that the company begin tracking data relevant to that influence.

Once the data is mapped to the ordinal scale, criticality scoring equations determine the inventory importance of each part. First, a subscore is developed for each part with respect to each influence set on the influence diagram. To obtain this subscore, the part characteristic data for an influence mapped to the ordinal scale for part *j* is multiplied by a weight calculated by the AHP priority for that influence. These products are then summed across all the influences in that set. Equation (1) describes the approach:

$$Set\ k\ subscore = g_{k,j} = \sum_i \left(p_{k,i} * d_{k,i,j}\right)\ \forall j \tag{1}$$

where $p_{k,i}$ is the AHP priority for influence i within set k and $d_{k,i,j}$ is the ordinal scale characteristic data for part j corresponding to influence i within set k. For the nuclear spare parts influence diagram, a subscore is found with respect to each influence set (i.e., 7 subscores in all for each part). These subscores are then combined with the priority weight for each set of influences and aggregated to yield an overall part criticality score; this is done via equation (2):

$$Part\ score = s_j = \sum_k \left(p_k * g_{k,j}\right) \quad (2)$$

where p_k is the AHP priority for set k and $g_{k,j}$ is the subscore for set k and part j. The end result is a single criticality score for each part. For more details and an example, the reader is referred to Scala, Rajgopal, and Needy [44].

Once all parts are scored, they must be placed into groups for inventory management. In this research, parts were assigned to clusters based on the calculated criticality scores using the *k*-means algorithm. This is a standard cluster analysis technique which starts with an initial set of cluster centroids and assigns data points to the cluster with the closest centroid. The centroids are then recalculated, and data are reassigned based on the new values. This process is repeated until the clusters converge. For further details regarding the algorithm, see [18] or [48].

While other grouping techniques could certainly be used, the *k*-means method is easy to implement and straightforward to understand, while providing good results. To determine the number of groups, 200 sample parts from the test bed company were selected, and criticality scores were computed. The scores were then graphed in a histogram, and three natural breaks occurred in the data, suggesting three groups or levels of criticality. The *k*-means algorithm verified the break points between the groups in the data.

For the case study company sample data, after *k*-means analysis, 21 parts were placed in Group I, 38 parts were placed in Group II, and 141 parts were placed in Group III. Group I parts have the highest criticality scores, implying those parts are most inventory critical, while Group III parts have the lowest scores and are not of great inventory management importance. Group II parts are of moderate inventory importance. Figure 3 shows a histogram of part scores and the corresponding criticality groups.

Once the parts were grouped, an inventory management policy must be defined for every group. The last step of the methodology details this process.

Step 4: Construction of inventory policies for classes of parts based on retrospective simulations

The fourth step of the methodology develops a modified base stock inventory policy for each part group via retrospective simulation using historical data. Results from the simulation yield a set of policies that can balance the cost of base stock inventory against delayed work days per part per month. The final policy to manage inventory can then be selected by the decision makers based on their risk tolerance profiles. Inventory control with intermittent demand has been discussed in the literature; review papers include [26], [29] and [32]. Recent papers for intermittent demand include Teunter, Syntetos, and Babai [50], who develop a policy based on a compound binomial process for demand; Chang, Chou, and Huang [12], who develop a stochastic continually reviewed constant reorder policy based on criticality; and Dekker, Kleijn, and de Rooij [14], who develop a stocking policy reserving stock for critical demand but assume critical and non-critical demand follows a Poisson process. However, as we have seen, for this class of parts, a demand process cannot be determined, and the intermittent nature of parts prohibits accurate calculation of a consistent reorder point. As a result, we develop a base stock policy to manage nuclear spare parts inventory.

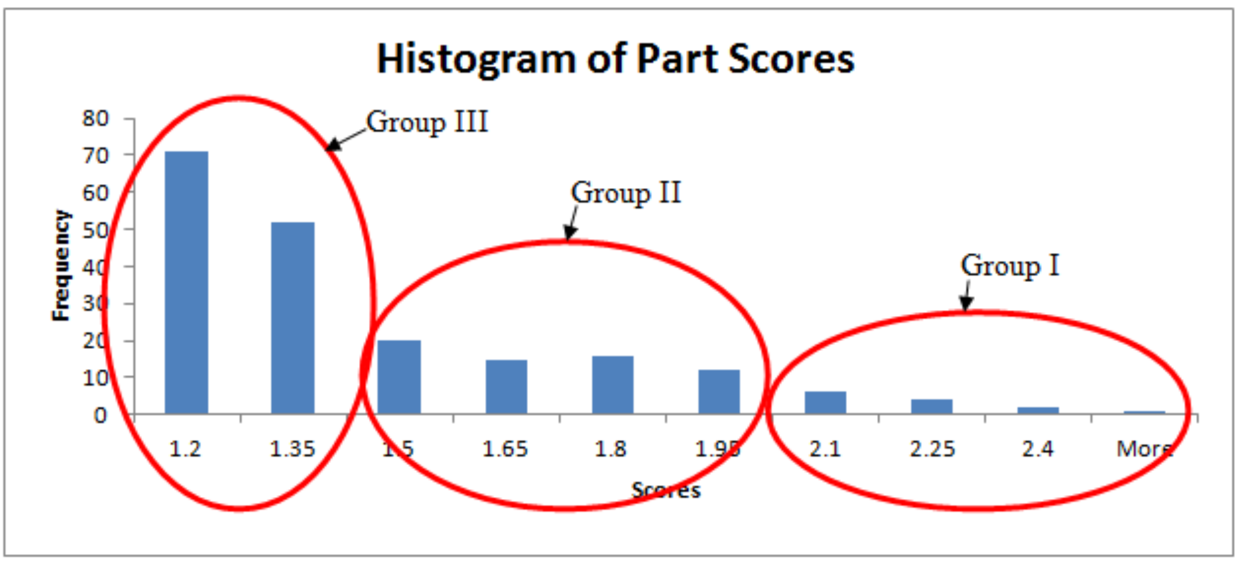

Figure 3. Histogram of part scores and criticality groups (Adapted from [42], [44])

Base stock inventory policies are continuous review policies that are predicated on always maintaining an inventory position equal to some "base stock" level. This level is typically set so that enough inventory is held to cover expected demand during the lead time (plus some safety stock that is proportional to the standard deviation of lead time demand). Inventory levels are depleted only when a demand occurs, and replacement orders are placed after each demand occasion for the amount demanded. When demand occurs in unit quantities and the lead time equals zero, a base stock policy is equivalent to a continuous review (s,S) policy, with $s = S-1$. Discussions on base stock inventory policies may be found in texts such as [25] and [28]; the literature has also shown [4] that (*s,S*) policies are appropriate for intermittent demand items. These policies are inherently conservative and suited to situations where items are expensive and demand is intermittent as in the nuclear power generation sector.

Because base stock policies are based on holding enough inventory to cover demand during lead time, a distribution for such demand is necessary to implement the policy. However, this is virtually impossible to determine for nuclear parts due to the extremely intermittent nature of demand, the lack of consistent demand patterns across parts in a group, and the absence of statistical data on failure rates in this sector. As a result, any base stock policy for nuclear spares must be modified using something other than a probability distribution of lead time demand. Our approach was to look at demand *occasions* over time and estimate how much was demanded on each occasion because these values are independent of time and relatively easy to measure.

To do this we developed a retrospective numerical simulation that examines 2,618 days (over 7 years) of historical demand occurrences for each part in our test bed sample set with different multiples of average demand per occurrence (i.e., per plant request) considered as potential base stock inventory positions. The average demand per plant request is defined as the average number requested by the plant across all historical work orders for each part. This demand calculation must also account for part returns of unused parts to the warehouse in order to estimate the actual demand. Details on our approach may be found in Scala [42]. Essentially, through analysis of the historical data, a value z_j was calculated for each part j in the sample set, where z_j denotes the typical quantity used by the plant for each work order or request for parts.

The historical numerical simulation outputs a series of potential base stock policies, with each policy identified by the number of parts to keep on the shelf, the cost of the corresponding inventory, and the associated delays in work (if any) that would have been historically experienced had that policy been in effect. Recall that the simulation is historical, so we are evaluating potential policies against a previous demand pattern, which we are assuming can also be expected in the future. The selected base stock policy is the simulated policy that calls for a base stock level that is some multiple of the average demand per request and that would have resulted in the minimum number of work days delayed due to unavailability of parts over the period considered. The simulation tests various multiples of average demand; specifically, multiples from 0 up to 50 in increments of 0.5. We denote this multiplier value of average demand by b. The upper bound of 50 (5,000% of average demand per request) is of course unrealistic in practice but was selected so that the point at which a 100% service rate is obtained could be determined for any part. The simulation tracks missed days at both LCO and non-LCO locations for every value of b as well as the corresponding inventory investment cost. Recall that failure of parts installed at LCO locations can cause the plant to shut down or derate if the situation is not remedied in a specified time window. Table 2 illustrates simulation results for one of the part groups (Group II). For every part in group 2, $b * z_j$ would be the quantity to store on the shelf. Similar results were obtained for the other groups as well; details of which may be found in Scala [42]. Figures 4, 5, and 6 show the relationships between *b,* cost, average non-LCO missed days, and average LCO missed days.

Table 2 indicates that in order to completely eliminate LCO missed days along with any possibility of a derate or shutdown due to lack of parts, the test bed company should have stored 11 times the average demand per plant request for the LCO parts in the group. While this would also greatly reduce delaying work days at non-LCO locations, it would not have completely eliminated them. To do so for the non-LCO parts in the group, the company would have had to store 22.5 times the average demand per plant request. In all cases, the simulation also yields the exact values of these inventory levels for all parts as well as dollar investments for both types of parts in the group (although Table 2 only shows totals across both types for a fixed multiple).

Table 2. Historical simulation results for group II [42]

b	**Cost**	**Avg. Non-LCO Missed**	**Avg. LCO Missed**	***b***	**Cost**	**Avg. Non-LCO Missed**	**Avg. LCO Missed**
0	$ -	2.673	0.08	11.5	$768,612.93	0.022	0
0.5	$42,013.52	1.309	0.013	12	$786,552.71	0.022	0
1	$73,660.13	0.930	0.013	12.5	$827,751.42	0.022	0
1.5	$110,914.38	0.833	0.013	13	$859,378.28	0.022	0
2	$133,346.48	0.707	0.013	13.5	$896,871.17	0.018	0
2.5	$174,322.45	0.631	0.013	14	$919,084.38	0.018	0
3	$202,114.73	0.334	0.013	14.5	$959,933.49	0.018	0
3.5	$243,222.84	0.186	0.013	15	$987,815.27	0.018	0
4	$265,655.41	0.186	0.013	15.5	$1,028,606.41	0.018	0
4.5	$302,650.19	0.180	0.013	16	$1,050,910.16	0.018	0

5	$334,403.91	0.172	0.013	16.5	$1,087,855.19	0.018	0
5.5	$375,478.33	0.168	0.013	17	$1,120,369.64	0.01	0
6	$394,140.01	0.168	0.013	17.5	$1,161,198.62	0.01	0
6.5	$434,939.37	0.168	0.013	18	$1,179,058.11	0.01	0
7	$467,095.62	0.168	0.013	18.5	$1,220,194.12	0.01	0
7.5	$503,621.32	0.163	0.013	19	$1,252,539.73	0.002	0
8	$525,935.79	0.163	0.013	19.5	$1,289,499.51	0.002	0
8.5	$566,998.61	0.163	0.013	20	$1,311,367.04	0.002	0
9	$594,516.71	0.138	0.009	20.5	$1,352,797.05	0.002	0
9.5	$636,384.29	0.110	0.009	21	$1,380,719.61	0.002	0
10	$658,171.53	0.106	0.009	21.5	$1,422,248.38	0.002	0
10.5	$695,209.03	0.098	0.009	22	$1,444,069.47	0.002	0
11	$727,089.10	0.055	0	22.5	$1,480,581.88	0	0

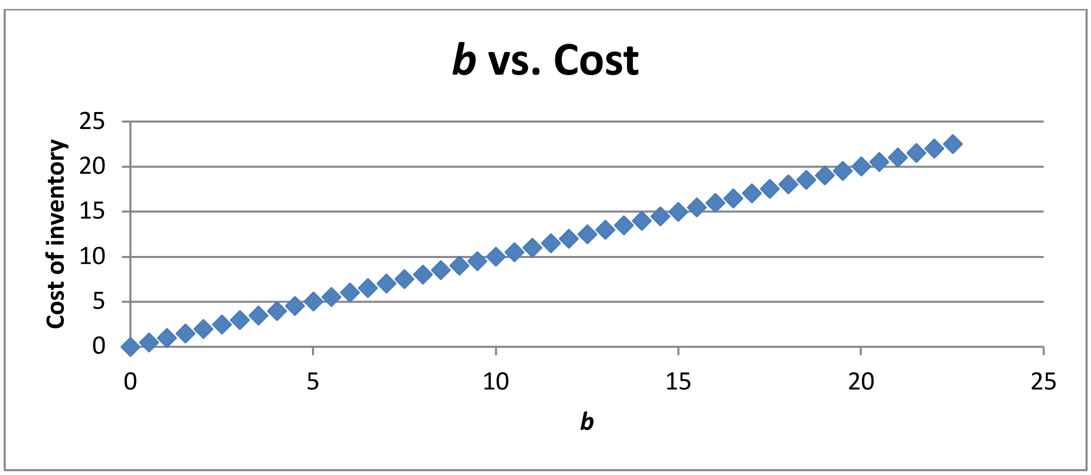

Figure 4. *b* vs. cost

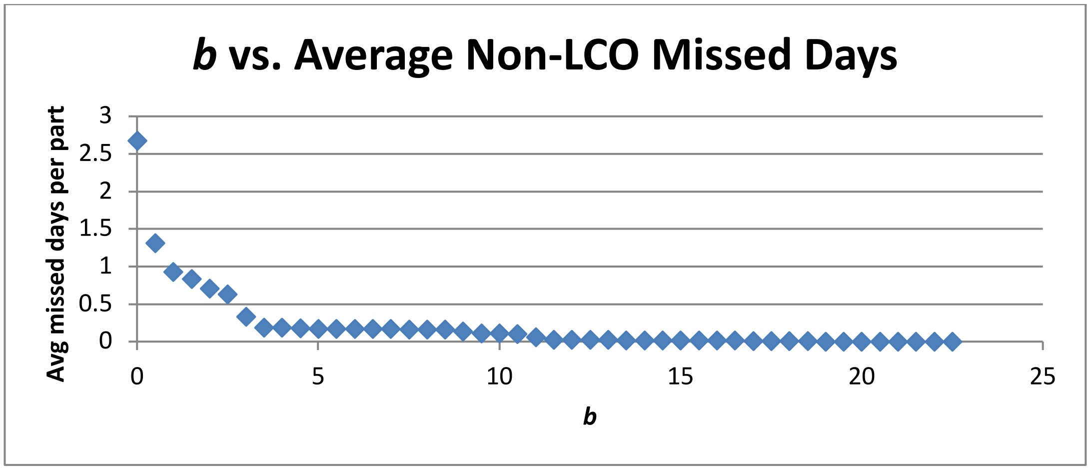

Figure 5. *b* vs. average non-LCO missed days

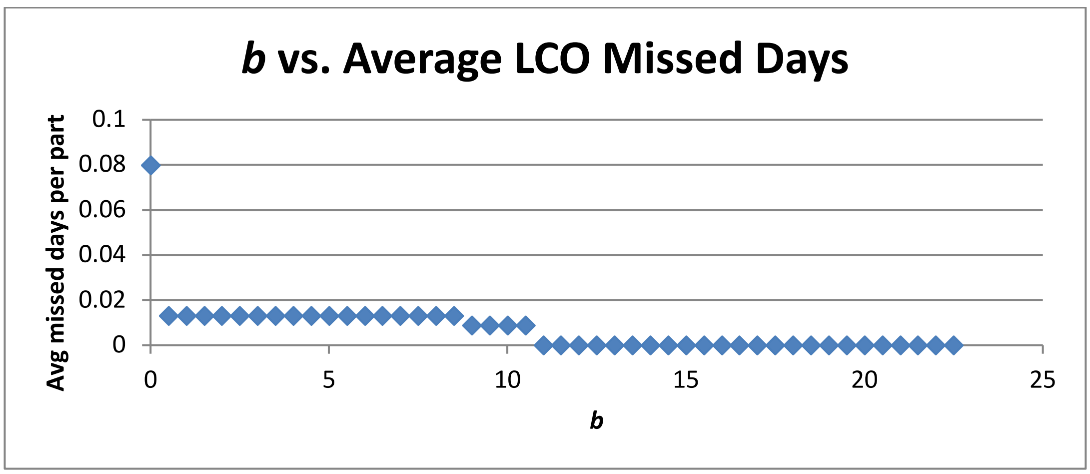

Figure 6. *b* vs. average LCO missed days

Costs of delaying work at LCO locations can be substantial (approximately $500,000 to $1,000,000 per day), depending on the size of plant and price of power, and failures at LCO locations can lead to plant shutdowns or derates, which cause a loss of revenue in a deregulated market [42]. Determining costs at non-LCO locations is trickier, as penalties associated with delaying work at these locations are not currently tracked by the test bed company. As a result, reducing the non-LCO days to zero may not be necessary, and the test bed company might be comfortable with permitting some work days delayed at non-LCO locations in exchange for inventory cost savings. The final decision rests with the decision maker at the test bed facility and is associated with the decision maker's tolerance for risk.

In sum, the methodology and corresponding policies presented are not prescriptive but rather a guideline for decision makers to support trading-off the appropriate risk of delaying work versus the capital cost of inventory, based on their own risk profile. Determining the important parts helps to buffer the risk of revenue loss associated with plant derate or shutdown, which results from an emergent situation that compromises the plant with parts not immediately available to remedy the issue. Through this methodology, the company may focus on and provide resources to the most important parts for inventory management.

The approach outlined in step 4 can be extended to any situation where lead time demand is small, intermittent, and cannot be characterized by a probability distribution. It should also be noted that we chose to examine multiples of the mean value corresponding to each demand occasion; other measures such as the mode or the median could also be used. Such situations would typically arise in spare parts problems, such as nuclear generation, aircraft engines, and ground space systems, but may also arise in new product development or a job shop environment where little but highly unique inventory is held.

Summary and Future Research

This research details a four-step methodology for managing spare parts in the nuclear electricity generation industry, utilizing actual data from a test bed company for illustration. The models developed in this research address the new competitive business environment in which many utilities are now operating and were verified and validated through real data. They can be readily adapted to other utilities to build decision support tools that can allow management decision makers to perform what-if analyses to assess the tradeoffs of their decisions to balance risk, safety, and costs. The methodology could also be generalized to other industries, such as aerospace and the military. It is possible for a company to use the entire methodology or just individual portions of it in order to develop its own influence diagram, AHP prioritized weights, criticality scores, and/or corresponding inventory policies. To our knowledge, no such integrated methodology for spare parts has been previously developed nor has one used all the techniques outlined here together in an integrated fashion.

Overall, the methodology is easy-to-use and implementable, without mathematical assumptions, allowing operational employees to work with and update the models. Employees who understand their work and can contribute to decision making will inevitably take ownership, leading to a higher probability of successful, sustained implementation and deeper employee engagement.

A natural extension of this work is the application of the methodology to nuclear spare parts suppliers. Each nuclear plant in the United States is uniquely designed, and suppliers provide customized support and parts to the plants they build. If suppliers have a better understanding of when parts will be ordered by plants, then intermittent demands at the supplier could be better managed, leading to reduced part lead time. Better management of demands and inventory at the supplier end would enable better response and improved servicing of nuclear plants, with both parties (suppliers and plants) operating under the same philosophy and similar methodologies for spare parts management.

A second extension would be to explore the development of utility functions for both average missed days per part per month and the cost of base stock inventory by interfacing with the subject matter experts. These functions would provide further analysis and depth to the inventory policy simulation results and support quantitative identification of the optimal base stock policy.

In conclusion, the conditions of intermittent demand, lack of failure rates, and limited data present a challenge for managing spare parts. In future research we plan to explore other opportunities to improve spare parts inventory management, especially at nuclear facilities, under these conditions.

References

[1] J. Aczél and C. Alsina, "Synthesizing judgements: A functional equations approach," *Math. Modelling,* vol. 9, pp. 311-320, 1987.

[2] J. Aczél and T. L. Saaty, "Procedures for synthesizing ratio judgements," *J. of Math. Psychology,* vol. 27*,* pp. 93-102, 1983.

[3] A. M. Agogino, O. Nour-Omid, W. Imaino, and S. S. Wang*,* "Decision-analytic methodology for cost-benefit evaluation of diagnostic testers," *IIE Trans.,* vol. 24, pp. 39-54, 1992.

[4] N. Altay, L. A. Litteral, and F. Rudisill, "Effects of correlation on intermittent demand forecasting and stock control," *Int. J. of Prod. Econ.,* vol. 135, pp. 275-283, 2012.

[5] N. Altay, F. Rudisill, and L. A. Litteral, "Adapting Wright's modification of Holt's method to forecasting intermittent demand," *Int. J. of Prod. Econ.,* vol. 111, pp. 389-408, 2008.

[6] T. C. Bachman, "Reducing aircraft down for lack of parts with sporadic demand," *Military Operations Research,* vol. 12, pp. 39-53, 2007.

[7] M. Braglia, A. Grassi, and R. Montanari, "Multi-attribute classification method for spare parts inventory management," *J. of Quality in Maintenance Eng.,* vol. 10, pp. 55-65, 2004.

[8] I. Basak, "When to combine group judgments and when not to in the Analytic Hierarchy Process: A new method," *Math. and Comput. Modelling*, vol. 10, pp. 395-404, 1988.

[9] I. Basak and T. Saaty, "Group decision making using the Analytic Hierarchy Process," *Math. and Comput. Modelling*, vol. 17, pp. 101-109, 1993.

[10] J. E. Boylan and A. A. Syntetos, "The accuracy of a modified Croston procedure," *Int. J. of Prod. Econ.,* vol. 107, pp. 511-517, 2007.

[11] S. Cavalieri, M. Garetti, M. Macchi, and R. Pinto, "A decision-making framework for managing maintenance spare parts," *Prod. Planning & Control,* vol. 19, pp. 379-396, 2008.

[12] P. L. Chang, Y. C. Chou, and M. G. Huang, "A *(r, r, Q)* inventory model for spare parts involving equipment criticality," *Int. J. of Prod. Econ.,* vol. 97, pp. 66-74, 2005.

[13] J. D. Croston, “Forecasting and stock control for intermittent demands,” *Operational Research Quart.,* vol. 23, pp. 289-303, 1972.

[14] R. Dekker, M. J. Kleijn, and P. J. de Rooij, “A spare parts stocking policy based on equipment criticality,” *Int. J. of Prod. Econ.,* vol. 56-57, pp. 69-77, 1998.

[15] R. J. Duintjer-Tebbens, M. A. Pallansch, O. M. Kew, R. W. Sutter, R. B. Aylward, M. Watkins, H. Gary, J. Alexander, H. Jafari, S. L. Cochi, and K. M. Thompson, “Uncertainty and sensitivity analyses of a decision analytic model for posteradication polio risk management,” *Risk Anal.,* vol. 28, pp. 855-876, 2008.

[16] A. G. Eleye-Datubo, A. Wall, A. Saajedi, and J. Wang, “Enabling a powerful marine and offshore decision-support solution through Bayesian network technique,” *Risk Anal.,* vol. 26, pp. 695-721, 2006.

[17] *Energy Policy Act of 1992,* 7 U.S.C. § 711, 1992.

[18] B. S. Everitt, S. Landau, and M. Leese, *Cluster Analysis,* 4th ed. London, England: Arnold, 2001.

[19] *Federal Energy Regulatory Commission*. 2010. “Order No. 888.” Last modified June 28. http://www.ferc.gov/legal/maj-ord-reg/land-docs/order888.asp.

[20] B. E. Flores, D. L. Olson, and V. K. Dorai, “Management of multicriteria inventory classification,” *Math. and Comput. Modelling,* vol. 16, pp. 71-82, 1992.

[21] B. E. Flores and D. C. Whybark, “Implementing multiple criteria ABC analysis,” *J. of Operations Manage.,* vol. 7, pp. 79-85, 1987.

[22] P. P. Gajpal, L. S. Ganesh, and C. Rajendran, “Criticality analysis of spare parts using the Analytic Hierarchy Process,” *Int. J. of Prod. Econ.,* vol. 35, pp. 293-297, 1994.

[23] J. Heizer and B. Render, *Operations Management,* 10th ed. Boston, MA: Prentice Hall, 2011.

[24] R. A. Howard and J. E. Matheson, “Influence diagrams,” *Decision Anal.,* vol. 2, pp. 127-143, 2005.

[25] W. J. Hopp and M. L. Spearman, *Factory Physics: Foundations of Manufacturing Management,* 2nd ed. Boston, MA: Irwin McGraw-Hill, 2001.

[26] J. Huiskonen, “Maintenance spare parts logistics: Special characteristics and strategic choices,” *Int. J. of Prod. Econ.,* vol. 71, pp. 125-133, 2001.

[27] D. L. Keefer, “Practice abstracts,” *Interfaces,* vol. 34, pp. 206-207, 2004.

[28] L. J. Krajewski and L. P. Ritzman, *Operations Management: Processes and Value Chains,* 7th ed. Upper Saddle River, NJ: Pearson Prentice Hall, 2005.

[29] W. J. Kennedy, J. W. Patterson, and L. D. Fredendall, “An overview of recent literature on spare parts inventories,” *Int. J. of Prod. Econ.*, vol. 76, pp. 201-215, 2002.

[30] A. Molenaers, H. Baets, L. Pintelon, and G. Waeyenbergh, “Criticality classification of spare parts: A case study,” *Int. J. of Prod. Econ.,* vol. 140, pp. 570-578, 2012.

[31] S. Nahmias, *Production and Operations Analysis,* 6th ed. New York, NY: McGraw-Hill Irwin, 2009.

[32] G. Nenes, S. Panagiotidou, and G. Tagaras, “Inventory management of multiple items with irregular demand: A case study,” *European J. of Operational Research,* vol. 205, pp. 313-324, 2010.

[33] W. L. Ng, “A simple classifier for multiple criteria ABC analysis,” *European J. of Operational Research,* vol. 177, pp. 344-353, 2007.

[34] F. Y. Partovi and J. Burton, “Using the Analytic Hierarchy Process for ABC analysis,” *Int. J. of Operations and Prod. Manage.,* vol. 13, pp. 29-44, 1993.

[35] P. E. Pfeifer, S. E. Bodily, R. L. Carraway, D. R. Clyman, and S. C. Frey, “Preparing our students to be newsvendors,” *Interfaces,* vol. 31, pp. 112-122, 2001.

[36] L. Philipson and H. L. Willis, *Understanding Electric Utilities and De-regulation,* 2nd ed. Boca Raton, FL: Taylor and Francis, 2006.

[37] J. Rezaei and S. Dowlatshahi, “A rule-based multi-criteria approach to inventory Classification,” *Int. J. of Prod. Research,* vol. 48, pp. 7107-7126, 2010.

[38] R. Ramanathan, “ABC inventory classification with multiple-criteria using weighted linear optimization,” *Comput. & Operations Research,* vol. 33, pp. 695-700, 2006.

[39] T. L. Saaty, *The Analytic Hierarchy Process: Planning, Priority Setting, Resource Allocation*. New York, NY: McGraw-Hill, 1980.

[40] T. L. Saaty, “How to make a decision: The Analytic Hierarchy Process,” *European J. of Operational Research,* vol. 48, pp. 9-26, 1990.

[41] T. L. Saaty and L. G. Vargas, “Dispersion of group judgments,” *Math. and Comput. Modelling,* vol. 46, pp. 918-925, 2007.

[42] N. M. Scala, “Spare parts management for nuclear power generation facilities,” Ph.D. dissertation, Dept. Ind. Eng., Univ. of Pittsburgh, Pittsburgh, PA, 2011.

[43] N. M. Scala, J. Rajgopal, and K. L. Needy, “Influence diagram modeling of nuclear spare parts process,” in *Proc. of Industrial Engineering Research Conf.,* Cancun, Mexico, 2010.

[44] N. M. Scala, J. Rajgopal, and K. L. Needy, “An inventory criticality classification method for nuclear spare parts: A case study,” in *Decision Making in Service Industries: A Practical Approach,* J. Faulin, A. A. Juan, S. E. Grasman, and M. J. Fry, Eds. Boca Raton, FL: CRC Press, 2012, pp. 365-392.

[45] S. Stafira, G. S. Parnell, and J. T. Moore, “A methodology for evaluating military systems in a counterproliferation role,” *Manage. Sci.,* vol. 43, pp. 1420-1430, 1997.

[46] N. Subramanian and R. Ramanathan, “A review of applications of Analytic Hierarchy Process in operations management,” *Int. J. of Prod. Econ.,* vol. 138, pp. 215-241, 2012.

[47] A. A. Syntetos and J. E. Boylan, “On the stock control performance of intermittent demand estimators,” *Int. J. of Prod. Econ.*, vol. 103, pp. 36-47, 2006.

[48] P. Tan, M. Steinbach, and V. Kumar, *Introduction to Data Mining.* Boston, MA: Pearson Addison Wesley, 2006.

[49] R. H. Teunter, M. Z. Babai, and A. A. Syntetos, “ABC classification: Service levels and inventory costs,” *Prod. and Operations Manage.,* vol. 19, pp. 343-352, 2010.

[50] R. H. Teunter, A. A. Syntetos, and M. Z. Babai, “Determining order-up-to levels under periodic for compound binomial (intermittent) demand,” *European J. of Operational Research,* vol. 203, pp. 619-624, 2010.

[51] United States Department of Energy. (2006). Workforce trends in the electric utility industry: A report to the United States Congress pursuant to Section 1101 of the Energy Policy Act of 2005. [Online]. Available: http://www.oe.energy.gov/DocumentsandMedia/Workforce_Trends_Report_090706_FINAL.pdf.

[52] United States Energy Information Administration. (2010). Status of electricity restructuring by State. [Online]. Available: http://www.eia.gov/cneaf/electricity/page/restructuring/restructure_elect.html.

[53] O. S. Viadya and S. Kumar, “Analytic Hierarchy Process: An overview of applications,” *European J. of Operational Research*, vol. 169, pp. 1-29, 2006.

[54] L. G. Vargas, “An overview of the Analytic Hierarchy Process and its applications,” *European J. of Operational Research,* vol. 48, pp. 2-8, 1990.

[55] L. Wang, Y. Zeng, C. Gui, and H. Wang, “Application of artificial neural network supported by BP and particle swarm optimization algorithm for evaluating the criticality class of spare parts,” in *Third Int. Conf. on Natural Computation.,* Haikou, 2007, pp. 528-532.

[56] L. Wang, Y. Zeng, Y. Li, and H. Wang, “An intelligent decision support system for spare parts joint replenishment,” in *Int. Conf. on Hybrid Information Technology,* Cheju Island, 2006, pp. 514-519.

[57] L. Wang, Y. Zeng, and J. Zhang, “A case study on joint replenishment problem for slow moving spare parts in a nuclear power plant,” in *Int. Conf. on Wireless Communications, Networking and Mobile Computing,* Shanghai, China, 2007, pp. 6597-6600.

[58] L. Wang, Y. Zeng, J. Zhang, W. Huang and Y. Bao, “The criticality of spare parts evaluating model using artificial neural network approach,” in *Computational Sci. - ICCS 2006,* V. N. Alexandrov, G. D. van Albada, P. M. A. Sloot, and J. Dongarra, Eds. Heidelberg, Germany: Springer-Verlag, 2006, pp. 728-735.

[59] F. Zahedi, “The Analytic Hierarchy Process—A survey of the method and its applications,” *Interfaces,* vol. 16, pp. 96-108, 1986.

[60] Y. Zheng, L. Wang, and J. Zhang, “A web-based fuzzy decision support system for spare parts inventory control,” in *Fuzzy Information and Engineering,* B. Y. Cao, Ed. Berlin, Germany: Springer-Verlag, 2007, pp. 601-609.

[61] G. W. Zimmerman, “The ABC’s of Vilfredo Pareto,” *Prod. and Inventory Manage.,* vol. 16, pp. 1-9, 1975.