European Journal of Operational Research 234 (2014) 731–742

Contents lists available at ScienceDirect

European Journal of Operational Research

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

Stochastics and Statistics

Decision dependent stochastic processes

Thomas Kirschenmann [a], Elmira Popova [b], Paul Damien [c,*], Tim Hanson [d]

[a] Institute for Computational Engineering and Sciences, University of Texas at Austin, 1 University Station C0200, Austin, TX 78712-0027, USA
[b] Operations Research and Industrial Engineering, Department of Mechanical Engineering, University of Texas at Austin, 1 University Station C2200, Austin, TX 78712-0027, USA
[c] Information, Risk, and Operations Management, McCombs School of Business, University of Texas at Austin, 1 University Station B6500, Austin, TX 78712-0027, USA
[d] Department of Statistics, University of South Carolina, 216 LeConte College, 1523 Greene St., Columbia, SC 29208, USA

ARTICLE INFO

Article history:
Received 22 May 2012
Accepted 15 November 2013
Available online 27 November 2013

Keywords:
Bayesian inference
Decision dependence
Endogenous
uncertainty
Weibull

ABSTRACT

Managers, typically, are unaware of the significant impact their decisions could have on the random mechanism driving a data generating process. Here, a new parametric Bayesian technique is introduced that would allow managers to obtain an estimate of the impact of their decisions on the stochastic process driving the data; this, in turn, should enhance a company's overall decision-making capabilities. This general approach to modeling decision-dependency is carried out via an efficient Markov chain Monte Carlo method. A simulated example, and a real-life example, using historical maintenance and failure time data from a system at the South Texas Project Nuclear Operating Company, exemplifies the paper's theoretical contributions. Conclusive evidence of decision dependence in the failure time distribution is reported, which in turn points to an optimal maintenance policy that results in potentially large financial savings to the Texas-based company.

1. Introduction

Managers are frequently tasked with making critical decisions in the face of uncertainty without the tools to assess the impact of their choices. These decisions can influence a stochastic process in complex ways that are not intuitive and are currently difficult to quantify. Decision makers can thus benefit tremendously from estimating the dependency between decisions and a statistic of interest.

This article presents a new Bayesian inferencing technique to detect and estimate decision dependency in a stochastic process. The technique allows one to form a precise estimate of the influence a managerial decision exerts upon the process, and consequently estimate the influence the decision makes upon variables that depend on that process. There does not currently exist any similar technique that can quantify decision dependency in the generality shown in this paper. Moreover, this new approach is set in a Bayesian framework and a practitioner can easily take advantage of numerous existing Bayesian methods for assessing statistical significance of the dependency estimates, as demonstrated throughout the numerical examples. This paper applies this method to analyze the details for the case of maintenance decisions and their influence on components of a nuclear power plant.

In Section 1.1 a brief review is provided of the related literature in both reliability theory and stochastic optimization. Section 2.1 defines the set of maintenance decisions that can act upon the system and affect the failure time distribution. Section 2.1 then describes the model for decision dependent failures in a system at a nuclear power plant. The estimation procedure is implemented in Sections 2.2 and 2.3 by applying modern Markov chain Monte Carlo (MCMC) techniques to estimate the model parameters. First, a numerical example is provided in Section 2.2 that relies on artificial data which is useful to demonstrate the ability of the method in a controlled setting and to provide confidence. A second example presented in Section 2.3 uses the maintenance history and failure times from a system at the South Texas Project Nuclear Operating Company (STPNOC) located in Bay City, Texas. Section 3 describes the ensuing optimization problem to determine the optimal decision policy using the Bayesian inferencing results, resulting in a sequence of future decisions that lead to potentially large financial savings.

1.1. Related work in reliability theory and stochastic optimization with endogenous uncertainty

Barlow and Hunter (1960) published an early and important work in maintenance policy theory. They define two types of maintenance policies for performing preventive and corrective maintenance in a nuclear power plant example. The first policy requires performing preventive maintenance after a fixed amount of time of continuous operation. If the system fails before the scheduled

* Corresponding author. Tel.: +1 512 471 3322.
E-mail addresses: thk3421@ices.utexas.edu (T. Kirschenmann), elmira@mail.utexas.edu (E. Popova), paul.damien@mccombs.utexas.edu (P. Damien), hansont@stat.sc.edu (T. Hanson).

0377-2217/$ - see front matter
http://dx.doi.org/10.1016/j.ejor.2013.11.016

maintenance, maintenance is performed immediately and the next preventive maintenance is rescheduled. Maintenance of any type under this policy is assumed to restore the system to as good as new status. The second policy schedules preventive maintenance at fixed times, regardless of any failures and corresponding repairs between those fixed times. In this second policy however, the maintenance following a failure merely brings the item back to the status it was in just before failure, so that the age of the item is not reset. The second policy is exactly like the policy defined below in Section 2.1 except for two new important generalizations. The present article makes allowance for the failure rate to depend on previous types of maintenance performed, differing from that of Barlow and Hunter (1960), and this new approach allows the failure rate to change after performing the repairs following a failure event, unlike Barlow and Hunter (1960).

Nguyen and Murthy (1981) develop optimal maintenance policies while considering the case of a failure rate that increases with the number of repairs. Their work, like most of the work on these topics, assumes an infinite time horizon for constructing an optimal policy. It is also possible to improve upon the notion that preventive maintenance perfectly restores an item. Nakagawa (1988) introduced the idea of an improvement factor, such as maintenance that can not only reduce the hazard rate of a system, but also reduce the age of a system without perfectly restoring it. Nakagawa (1988) forms optimal policies of sequences under both of these scenarios using the usual assumption of a Weibull shaped failure time distribution. Nguyen and Murthy (1981) and Nakagawa (1988) are important contributions for constructing a modern approach to modeling failure times, however neither project addresses the effect of decision dependency.

Singh (2011) further developed models for periodic preventive maintenance using virtual-age based age-reduction factors as well as factors describing the change in the failure rate of the item. The parameters in these models are estimated in a Bayesian framework using a Markov chain Monte Carlo method and are then input into a two-stage stochastic optimization program for determining the optimal interval to perform periodic imperfect preventive maintenance.

The problem of choosing the optimal time interval for scheduled maintenance is also explored in the work of Damien, Galenko, Popova, and Hanson (2007). This work demonstrates a semiparametric Bayesian model to determine the optimal time to schedule preventive maintenance while allowing for corrective maintenance in order to minimize the expected cost of operating a single item in the context of a nuclear power plant. Their work is the first to report that the problem of scheduling preventive maintenance and performing corrective maintenance after failures has a decision dependent influence on the expected lifetime and total cost of operation. This finding is just one of many decision dependent scenarios and serves as the ideal motivation for the present article to quantify decision dependent uncertainty.

The related literature on stochastic programming can be divided into two classes, problems with exogenous uncertainty and problems with endogenous uncertainty. Exogenous uncertainty in a stochastic process is uncertainty that is not influenced by optimization decisions. Endogenous uncertainty in a stochastic process is uncertainty that depends on the optimization decisions, either explicitly or implicitly. Goel and Grossmann (2004) make this distinction in their literature review and note that most previous work focuses on exogenous uncertainty, where the optimization decisions cannot influence the stochastic process. Goel and Grossman also note that Pflug (1990) was the first work towards solving problems with endogenous uncertainty. Goel and Grossmann (2004) further classify endogenous uncertainty by two types of effects produced. The first type of decision dependent uncertainty is where a decision-maker may change the probability distribution of the process by making one outcome more likely than another. The second type of decision dependent uncertainty is that of increased information available to the decision-maker by partially resolving a particular uncertainty. Goel and Grossmann (2004) further clarify the distinction by noting that in the first case the decision-maker can force one possibility to become more probable. In the second case the decision-maker can only become more sure as to which possibility may occur in the future. The decision dependency estimated in the present article falls into the first category of endogenous uncertainty, wherein the decision alters the underlying probability distribution.

Goel and Grossmann (2004, 2006) present a planning problem for gas fields where the decisions influence the available information by further resolving uncertainty about the potential yield from that area. A costly investment can be made to build infrastructure in a gas field that gives information about the region and reduces uncertainty. The decision of whether to build the infrastructure in a particular region influences when the uncertainty is resolved. In Goel and Grossmann (2006) a branch and bound method is demonstrated, and in Goel and Grossmann (2004) a scenario tree is used to describe the evolution of the stochastic process and also the uncertainty that is resolved after each decision. The solution for this class of problems uses a hybrid mixed-integer disjunctive program that includes non-anticipativity constraints. It is clear from this and other important works that each instance of decision dependency must be solved with the context of the problem in mind, greatly increasing the difficulty of generalizing solutions to decision dependent problems.

An interesting example of decision dependency in an optimization method itself is presented in Pflug (1990). Pflug (1990) introduces a method for determining the minimization of an objective function involving a Markovian process with a recursive estimation procedure. His method simulates several steps of the Markov process, each under different control values, then computes a stochastic gradient. Depending on the results, the procedure is iteratively repeated with different control values. This method is similar to a gradient descent method in deterministic systems, however the decision of how to adapt the control variables changes the ultimate results of the optimization routine, making the result inherently decision dependent. The presentation of decision dependency in an optimization method itself, not just in a system to be modeled, demonstrates the ubiquitous nature of decision dependency and also underscores the importance for further understanding.

Peeta, Salman, Gunnec, and Viswanath (2010) solve a two-stage stochastic program where investment decisions must be made for strengthening a highway network before a disaster. The first stage decisions influence the probability distributions for the subsequent damage to the network links. The first stage variables are restricted to integers and the distributions of the random parameters depend on these variables. The problem thus falls under the umbrella of stochastic integer programming, although it has a decision dependent structure that adds tremendous complexity. Peeta et al. (2010) state that models with decision dependent probabilities are typically known to be quite difficult to solve and suggests trying the sample average approximation (SAA) method originated in Ahmed and Shapiro (2002) and Kleywegt, Shapiro, and Homem-de Mello (2002). The present article takes advantage of some empirical structure found in the objective function values and uses a genetic algorithm to quickly converge to an optimal solution.

2. Decision dependent Bayesian estimation

In this section the definition and estimation of a decision dependent stochastic process in a Bayesian setting are introduced. The stochastic process is described in the context of maintenance