Integrated Maintenance and Production Planning with Endogenous Uncertain Yield

Tahir Ekin

McCoy College of Business, Texas State University

March 27, 2017

Abstract

The relationships among production planning, maintenance decisions and machine yield are crucial in a number of manufacturing environments such as the semi-conductor industry. This paper presents an integrated maintenance and production decision making framework with stochastically proportional endogenous yield rate and random demand. Finding the solution for this two-stage nonlinear stochastic program with endogenous uncertainty is not straightforward, and has not been considered previously. An augmented probability simulation based method is utilized to solve for the proposed decision model. We demonstrate the use of the proposed approach by conducting a numerical study and sensitivity analysis. We discuss the trade-offs involved among the yield, and simultaneous decisions of production planning and maintenance.

Keywords : maintenance; production planning; integrated decision making; endogenous uncertainty; stochastic programs; augmented probability simulation

1 Introduction

In many production environments, diminishing machine conditions and high production rates can adversely affect the product quality. The machine yield can be jointly impacted by both the production quantity and maintenance decisions. This impact is especially important where the maintenance cycles are short compared to production cycles. Increasing wear in machine shop tool bits may cause defects over time which would result with faulty products. For instance, as part of the etch operation in semiconductor water fabrication, chemicals are used to strip materials from the surface of the silicon wafers. The contamination of the inner chambers of the etch equipment rapidly increases with higher production levels. This results in defective products and in lower machine yield. It creates unique challenges in multi-stage maintenance and production planning, and can affect the inventory and outsourcing costs. Proper understanding of such relationships has the potential to result in significant savings in operational costs and improved efficiency for the overall production system. Therefore, production and maintenance decisions as well as their impact on uncertain yield need to be considered simultaneously. The complexity of these relationships makes the integrated system-based approaches crucial. This paper introduces a novel integrated decision making framework to investigate the trade-offs among production planning and maintenance decisions under random demand and endogenous (decision dependent) yield.

Decision makers need to adapt to dynamic behavior of the production environments because of the uncertainty involved with the markets. This motivates models that allow integrated decision making under uncertainty while addressing a high number of decision alternatives. Integrated and simultaneous production and maintenance models are shown to outperform traditional sequential approaches (Sloan (2008), Batun and Maillart (2012)). Within such integrated decision making frameworks, the machine yield can be a function of maintenance and/or production. However, most of the existing models can deal with only a relatively small number of decision alternatives. The main challenge is the computational complexity; especially for high levels of decision alternatives and when uncertainty is considered. Stochastic programming methods can be employed to deal with many alternatives for the decision variables under long term uncertainty (Birge and Louveaux (2011)). Mula et al. (2006) point out the frequent use of stochastic programming models in their survey of production planning models. However, the dependence of a random quantity on previous decisions require models that consider endogenous randomness. Particularly, the probability distribution of the random variable may depend on the previous decisions. The solution of such models is not straightforward, therefore mostly models with discrete random variables and small number of scenarios are considered. Recent computational and methodogical advances in simulation based optimization algorithms such as those in Ekin et al. (2014) can be utilized to address problems with continuous uncertainty which may depend on decisions with higher number of alternatives.

This paper contributes to the literature in two main ways. First, the integrated model lets the decision maker evaluate the trade-offs involved among production quantity, maintenance, outsourcing, salvaging decisions and uncertain endogenous machine yield in a two-stage setting with random demand. We utilize a stochastically proportional yield based approach in which yield rate is modeled via a truncated Normal distibution. This lets the decision maker to explicitly model variance independently from the batch size. Our approach is general enough to accommodate any discrete or continuous probability distribution as well as any form of objective function. In addition, our stochastic program is flexible enough to accommodate more decision alternatives than the existing literature such as Sloan (2004) that can deal with relatively small number (up to 20) of maximum production quantities. Second, this is the first application of augmented probability simulation (APS) based stochastic programming solution approach of Ekin et al. (2014) to solve a production planning and maintenance problem. Finding the solution for the proposed model with endogenous uncertainty is not straightforward and had not been considered previously. We solve the proposed two-stage nonlinear stochastic program using an augmented probability simulation based optimization method. This is one of the first applications of augmented probability simulation based optimization method, the first to solve a stochastic optimization model that includes direct endogenous randomness or nonlinearity.

The paper is organized as follows. Section 2 presents the relevant literature review. Section 3 presents the modeling framework and details about modeling yield. Section 4 describes the stochastic optimization method used to solve the proposed model. Section 5 provides a numerical illustration with a large numerical study and sensitivity analysis, and presents a discussion of the results. Section 6 concludes with a summary of findings and directions for future work.

2 Literature Review

This section provides a literature review that is relevant to the proposed model and solution approach from three perspectives. First, the literature of modeling yield and integrated production models is presented. This is followed by a brief overview of endogenous stochastic models and simulation based stochastic programming approaches.

2.1 Modeling Yield and Integrated Models

In a production environment, yield is generally defined as the percentage of working products that emerge from the process. Uncertainty of yield is extensively studied in production and maintenance problems. These variable yield models can be classified depending on the uncertainty structure and whether process condition is under decision maker's control or not. Yano and Lee (1995) present a comprehensive review of production models with uncertain yield. The simplest yield model assumes that each produced unit follows a Bernoulli process, and models the total number of working products using Binomial distribution. The decision maker only needs to specify the probability of each item working properly and assume the independence of products in a batch. This can be appropriate for systems that are in statistical process control for long durations. However, such models as in Sloan (2004) do not explicitly model variance and are dependent on the batch size. In contrast, stochastically proportional yield models can specify both the mean and variance of the yield rate, independently from the batch size. Hence, they are more generally applicable especially when the variation of the batch size from production run to production run tends to be small. In most cases, the fraction of good units is between zero and one. Therefore, Truncated Normal distribution (Kalirajan (1981)) and Beta distribution (Sloan (2008)) are natural candidates when both the mean and variance are of interest.

Most of these models assume that the uncertainty of the yield is beyond the control of decision maker, and do not explicitly link the effect of the equipment condition to the yield. The limited literature that models yield as a function of previous decisions mainly focus on dependence from one perspective. A number of models consider the adverse impact of increasing production on yield, for instance see Khouja and Mehrez (1994) and Iravani and Duenyas (2002). Sana (2010a) models the total number of defective items in and out-of control environment using a function of production amount and run-time. The sole impact of maintenance on yield is generally studied with a focus on time of the repair, see the review of Wang (2002).

The decision models that consider both maintenance and production mainly focus on a single dimension, the effects of failures on production and inventory decisions (Sloan (2008)). However, increasing complexities in production systems make the system based approaches more crucial and an area of interest from researchers (Hadidi et al. (2012)). In terms of such combined models of production planning and maintenance with endogenous yield, the literature is relatively recent and limited. Sloan and Shanthikumar (2000) examine the joint determination of production and maintenance schedules in a single stage multi-product setting under deterministic demand. They model the probability of transition of machine states as a function of deterministic yield within a Markov decision process. Sloan (2008) extends this by incorporating imperfect maintenance and the uncertain time between decisions. Yield based simultaneous model is shown to outperform yield based traditional sequential approach. Batun and Maillart (2012) reassess these models and present further evidence for the advantages of simultaneous maintenance and production planning. However, none of these models consider the impact on inventory or backholding costs. Sloan (2004) presents an integrated decision model to determine the production quantity and maintenance schedule that minimizes the sum of expected production, backlog, holding costs while dealing with a discrete stochastic demand. The product yield has a binomial distribution that depends on the equipment condition and and previous decisions. Simultaneous maintenance and production decision making is shown to be more cost-effective compared to the traditional sequential approaches. These Markov decision process based models are utilized for a small number of decision alternatives. Thus, the dependence of yield is considered only for a limited number of actions.

In relevant work, Kazaz and Sloan (2013) propose joint maintenance and production schedules for additional problems that do not meet the conditions of Sloan (2008). Aramon Bajestani et al. (2014) address the problem of integrated maintenance and production scheduling in a multi-machine production environment over multiple periods. Hong et al. (2014) model the dependent stochastic degradation of components and formulate the maintenance decision problem using the minimum expected cost and the stochastic dominance rules. Peng and van Houtum (2016) consider a joint optimization model to determine the production and condition based maintenance policy while using a continuous time and state degradation process. Sana (2010b) uses various product reliability parameters to account for different production and maintenance environments and initial machine conditions.

2.2 Two-stage Endogenous Stochastic Programs with Recourse

Two-stage stochastic programs with recourse allows the decision maker to postpone some decisions to the second stage after observing the realizations of the uncertain variables (Birge and Louveaux (2011)). These corrective decisions are also referred to as the recourse decisions. The objective of the decision maker is to choose a first stage decision that is feasible for all potential realizations of the uncertain variable so that the expected objective function associated with both stages is optimized. The objective function of the second stage, (a.k.a recourse function) depends on both the first stage decision and the uncertain variable. The second stage decision is determined by solving the second stage problem for the specific combination of the first stage decision and the realization of the random variable.

Most of these models assume that the random variables are independent from the previous decisions. Our focus is on the models with endogenous uncertainty where the probability distribution of the uncertain variable depends on the first stage decision. Goel and Grossmann (2006) classifies these into two groups. Firstly, the probability distribution of the random variable can be a direct function of the previous decisions (Ahmed (2000)). For instance, within a pre-disaster investment decision model for a highway network, the survival probabilities of highway links after a disaster can be modeled as a function of investment decisions (Peeta et al. (2010)). Secondly, the decision maker could resolve the uncertainty partially and can have an updated probability distribution based on previous decisions by eliminating a number of scenarios (Jonsbraten et al. (1998)). For instance, in the context of operational planning of offshore gas development, decisions of drilling locations can update the decision maker's uncertainty about the oil reserves (Goel and Grossmann (2004)). Such applications include the areas of aggregate workforce planning (Fragnière et al. (2010)), project portfolio management (Solak et al. (2010)) and reliability (Kirschenmann et al. (2014)).

To the best of our knowledge, there are not any models that consider both types of endogenous uncertainty. Our model provides a minor contribution to stochastic programming modelling literature, since it is a case of endogenous stochastic programming where both types of endogenous uncertainty are utilized.

2.3 Simulation based Stochastic Programming

Stochastic programs with recourse require the computation of the expectation function as well as its optimization over the decision variables. A variety of algorithms are proposed to solve such problems; see Birge and Louveaux (2011) for a review. When the stochastic uncertainty follows a discrete distribution with a small number of scenarios, the problem can be rewritten explicitly and solved using one of the many deterministic optimization methods such as L-shaped method (Van Slyke and Wets (1969)). However, for many real world applications, such closed form analytically available solutions cannot be obtained, therefore approximations such as simulation based methods are utilized. Simulation based stochastic programming approaches are typically based on variants of Monte Carlo sampling, see Homem-de Mello and Bayraksan (2014) and Rubinstein and Kroese (2011) for overviews. For instance, a widely used method called Sample Average Approximation (SAA), is based on estimating the expectation function using a number of scenarios and solving the resulting deterministic optimization problem for a number of replications. However, for endogenous problems, the dependence between decisions and random variables increases computational complexity since the search space increases exponentially with dimensionality (Dyer and Stougie (2006)). Therefore, many of the existing endogenous models are linear and solved with a limited number of scenarios. An exception is the model of Solak et al. (2010), which is amenable to scenario based decomposition and uses SAA as the solution approach.

The main drawback of SAA methods is potential inefficient allocation of optimization effort since Monte Carlo errors of estimating the expectation function can overwhelm the calculation of the optimal decision. APS approach of Ekin et al. (2014) avoids this by simultaneously performing the expectation and optimization. Particularly, APS involves transformation of the optimization problem into a grand simulation by treating the decision variable as random for computational purposes. In doing so, APS searches for the optimal decision in the joint space of decision and random variables. This is beneficial to address the conditional dependence between the state and the decision spaces. As another advantage, APS does not require the exact evaluation of the objective function. It is applicable for any positive objective function including nonlinear ones such as our case. Since it is a relatively recent method, its applications are limited. Aktekin and Ekin (2016) utilizes it in their call center staffing model with random arrival, service and abandonment rates. In their application, the probabilities of abandonment and delay are functions of both decision variables and random parameters, hence indirect decision dependency exists. Irie and West (2016) presents a relevant Bayesian emulation approach in their solution of a multi-step portfolio management problem.

3 Modeling Framework

The production environment for a single machine and a single product is modeled in a two stage setting. In the following, the notation for decision variables and parameters is presented. It is followed by a discussion of the the optimization model and modeling endogenous random yield.

Decision Variables

x_1: production quantity, continuous first stage

x_2: maintenance decision, integer first stage; $x_2 \in \{1, 2, 3\}$ that correspond to zero maintenance, preventive or corrective maintenance respectively

x_{2a}, x_{2b}: dummy variables used to represent x_2; $\{x_{2a}, x_{2b}\} = \{0, 0\}$, $\{x_{2a}, x_{2b}\} = \{1, 0\}$, $\{x_{2a}, x_{2b}\} = \{0, 1\}$ refer to $x_2 = \{1, 2, 3\}$ respectively

y_1: quantity of outsourced production, continuous second stage

y_2: quantity of salvaged production, continuous second stage

Parameters

c: unit product cost

m_1: unit cost of the preventive maintenance

m_2: unit cost of the corrective maintenance

PC: maximum production capacity

p: unit product sales price

s: unit scrapping price of an imperfect product

o: unit outsourcing cost

r: unit salvage price

ξ_1: random, yield rate of the machine

$d(\xi_2) = \xi_2$: random, demand for the product

α: production-yield impact parameter

$\boldsymbol{\beta}$=$[\beta_1, \beta_2, \beta_3]$: product reliability parameter vector that consists of parameters for each maintenance decision alternative

The optimization model is written as:

$$\begin{aligned} \min_{x_1, x_{2a}, x_{2b}} \quad & cx_1 + m_1 x_{2a} + m_2 x_{2b} - E[Q(x_1, x_{2a}, x_{2b}, \xi_1, \xi_2)] \\ \text{subject to} \quad & x_1 \leqslant PC, x_1 \geqslant 0, x_{2a}, x_{2b} \in \{0,1\}, x_{2a} + x_{2b} \leqslant 1 \\ \text{where } Q(x_1, x_{2a}, x_{2b}, \xi_1, \xi_2) = \max_{y_1, y_2} \quad & p(\xi_1 x_1 + y_1) + s(1-\xi_1)x_1 - oy_1 + ry_2 \\ \text{subject to} \quad & \xi_1 x_1 + y_1 - y_2 \geqslant d(\xi_2); y_1, y_2 \geqslant 0 \end{aligned} \tag{3.1}$$

Both production quantity and maintenance decisions are made simultaneously at the first stage without knowing the realization of yield and demand. Maintenance decision alternatives are chosen as zero, preventive and corrective maintenance, similar to the most of the literature of integrated models of production planning and maintenance such as those in Sloan (2004), Aramon Bajestani et al. (2014) and Peng and van Houtum (2016). It is assumed that the preventive maintenance has lower unit cost compared to corrective maintenance, $m_1 \leq m_2$. Whereas the cost of taking zero maintenance action, $x_2 = 1$, is zero. The production quantity, x_1, is nonnegative, and has the maximum capacity level, PC, as the upper bound. The objective function of the first stage is to minimize the sum of the production cost, (cx_1) and the maintenance cost, $(m_1 x_{2a} + m_2 x_{2b})$, and to maximize the expected second stage objective function $E[Q(x_1, x_{2a}, x_{2b}, \xi_1, \xi_2)]$ while satisfying the constraints. The second stage function is denoted as $Q(x_1, x_{2a}, x_{2b}, \xi_1, \xi_2)$.

The second stage model allows the decision maker to make recourse decisions based on the realization of the yield and demand. The decision variables are the quantities of outsourced product, y_1 and salvaged product, y_2. The objective is to maximize the profit while meeting the demand, similar to many production models including Sloan (2004). The recourse objective function, $Q(x_1, x_{2a}, x_{2b}, \xi_1, \xi_2)$, consists of the revenue from the sold manufactured and outsourced items, $(p(\xi_1 x_1 + y_1))$ as well as the revenue from scrapped and salvaged products, $(s(1 - \xi_1)x_1 + ry_2)$ in addition to the cost of outsourced production, (oy_1). Meeting demand is guaranteed by outsourcing via $(\xi_1 x_1 + y_1 - y_2 \geqslant d(\xi_2))$. Outsourcing is assumed to be conducted at a monetary loss so that the demand is satisfied. Overall, it is assumed that $o \geqslant p \geqslant c \geqslant r \geqslant s$.

One of the important aspects of stochastic programming is non-anticipativity which refers to decisions depending only on the information available at the time of the decision. Therefore, scenarios with a common history must have the same set of decisions. Non-anticipativity can be modeled implicitly or explicitly. Explicit modeling is preferred for algorithms which benefit from that structure, including decomposition and Lagrangian based algorithms. Implicit formulation reduces the redundancies in data handling and do not report the historical information along a scenario path, see Wallace and Ziemba (2005) and Kall and Mayer (2005) for formal discussions. The proposed solution algorithm does not benefit from explicit constraints since it is not based on generating scenarios. In our setting, the first stage decisions can not depend on the realization of the random data, and non-anticipativity is guaranteed implicitly.

The demand, d is assumed to be random and exogenous from previous decisions. We assume a simple function for demand, $d(\xi_2) = \xi_2$ where $p(\xi_2)$ is Normal with parameters of (μ_d, σ_d). Next, we provide details about modeling endogenous random yield.

3.1 Modeling Endogenous Yield

This paper models the uncertain yield as a function of previous decisions. Particularly, the yield rate is assumed to follow an endogenous distribution as a function of production decision and conditional on maintenance decision as

$$p(\xi_1|x_2 = i) \sim T.Normal(1 - (\alpha \frac{x_1}{PC})^{\beta_i}, \sigma_{\xi_1}) \mathbb{I}(0 \leqslant \xi_1 \leqslant 1), i \in \{1, 2, 3\}$$

This distribution is truncated to $(0, 1)$ region which corresponds to the state space of yield rate. Normal distribution is chosen since it gives flexibility to specify the variance independently from the endogenous mean and the batch size. The variance of the yield is chosen to reflect the deviation around the endogenous mean. It is determined with respect to the confidence of the decision maker about the endogenous mean yield rate.

The mean yield rate is equal to $1 - (\alpha \frac{x_1}{PC})^{\beta_i}$ for each particular maintenance decision. In other words, the mean deterioration rate is $(\alpha \frac{x_1}{PC})^{\beta_i}$. The production-yield impact parameter, α, reflects the rate of deteroriation as a function of production. It is determined with respect to the machine condition and the overall status of the process. The yield is assumed to be nonincreasing for increasing quantities of production, and α is assumed to be bounded by 1; $0 \leqslant \alpha \leqslant 1$. A new machine which is not expected to be impacted by production, is assigned an α value of 0.01. Whereas, an α of 1 corresponds to an old machine, therefore increase in production quantity is expected to have a larger impact on the yield. The product reliability parameter vector, $\boldsymbol{\beta} = [\beta_1, \beta_2, \beta_3]$, takes the impact of maintenance on yield into account. Each product reliability parameter, β_i is specific to the maintenance decision alternative, and it represents the impact of that particular action. The decision maker has the option to conduct zero maintenance at zero cost which results with a lower expected yield. Whereas preventive maintenance improves the product at a certain rate and corrective maintenance makes the machine as good as new. In other words, the corrective maintenance has the largest positive impact on yield that is followed by preventive and zero maintenance. Therefore, product reliability parameter values are ordered in that $\beta_1 \leq \beta_2 \leq \beta_3$.

We assume that maintenance time is negligible enough to be included in the idle time and machine does not degrade in this period. Our focus is on cases where the maintenance cycle is short compared to the production cycle, so this is a reasonable assumption (Peng and van Houtum (2016)).

The proposed two stage model can be utilized within a multi-period decision framework. For instance, let's assume there is a machine that produces inputs for another process. At $t = 0$, the decisions on production quantity and maintenance are made. Then, the production is conducted, and a certain percentage of the production turn out to be defective. This yield rate is impacted by previous decisions. At $t = 1$, non-defective products are transferred to the next process for profit, and excess production can be salvaged (or stored) at a lower return. The defective items are scrapped, and the shortage in production is satisfied by outsourcing (or from inventory) at a monetary loss. The decision maker observes the machine condition and production environment to determine the values of α and $\boldsymbol{\beta}$, and makes another round of production and maintenance decisions.

4 Methodology

A closed form solution for the proposed model is not available, which requires the use of simulation based methods. SAA is computationally inefficient for endogenous problems, since the scenario trees become very large for continuous uncertainty and/or many decision alternatives. We use the APS approach of Ekin et al. (2014), which provides a natural way to deal with conditional dependence between the state and the decision spaces. In what follows, we present the implementation of APS for the proposed model.

APS is based on maximizing the expected value function via simulation from an augmented probability distribution of the decision variable and uncertain parameters. We use a number of strategies that result in proper and efficient utilization of APS. First, we restructure our formulation as a maximization problem and add a large enough constant value, M, that shifts the objective function to the positive region. This ensures proper probability densities without changing the structure of the probability density function. The decision maker can compute the miminum possible objective function value to help determine M, see Jacquier et al. (2006) for the similar solution of an asset allocation problem with negative utilities. Next, we replace the objective function by a power transformation that uses "J copies" of the random variable. This results in a more peaked surface that leads to faster convergence without changing the solution of the problem (Müller (1999)), see the discussion by Müller et al. (2004) for details.

The objective function for each j^{th} random variable is written as

$$u(\boldsymbol{x}, \boldsymbol{\xi}_j) = M - cx_1 - m_1 x_{2a} - m_2 x_{2b} + Q(\boldsymbol{x}, \boldsymbol{\xi}_j)$$

where $\boldsymbol{x} = \{x_1, x_{2a}, x_{2b}\}$ and $\boldsymbol{\xi}_j = \{\xi_{1j}, \xi_{2j}\}$. We denote "J copies" of the random variable as $\boldsymbol{\xi}_J = \{\boldsymbol{\xi_{1J}}, \boldsymbol{\xi_{2J}}\}$ where $\boldsymbol{\xi_{1J}} = \{\xi_{1,1}, \ldots, \xi_{1,j}, \ldots, \xi_{1,J}\}$ and $\boldsymbol{\xi_{2J}} = \{\xi_{2,1} \ldots, \xi_{2,j}, \ldots, \xi_{2,J}\}$ for $j = \{1, 2, \ldots, J\}$.

For given values of $(\boldsymbol{x}, \boldsymbol{\xi}_j)$, we can solve our second stage linear program to optimality. The optimal second stage decisions are written as functions of first stage decisions and realizations of random variables as in

$$y^*_{1j}(x_1, \boldsymbol{\xi}_j) = max(0, \xi_{2j} - \xi_{1j} x_1)$$

and

$$y^*_{2j}(x_1, \boldsymbol{\xi}_j) = max(0, \xi_{1j} x_1 - \xi_{2j}).$$

For parsimony, we abbreviate these as y^*_{1j} and y^*_{2j}. Finally, an auxiliary distribution is constructed on the augmented space of first stage decision variables and "J" copies of the random variables, such that

$$\pi(\boldsymbol{x}, \boldsymbol{\xi}_J) \propto \prod_{j=1}^{J} u(\boldsymbol{x}, \boldsymbol{\xi}_j) p(\xi_{2j}) p(\xi_{1j}|\boldsymbol{x}) p(\boldsymbol{x}) \mathbb{I}(x_1 \leqslant PC, x_1 \geqslant 0, x_{2a}, x_{2b} \in \{0, 1\}, x_{2a} + x_{2b} \leqslant 1).$$

The prior density of decision variables, $p(\boldsymbol{x})$ is assumed to follow a Uniform distribution bounded by the search space and the objective function for the j^{th} copy, $u(\boldsymbol{x}, \boldsymbol{\xi}_j)$, is explicitly written as $(M - cx_1 - m_1 x_{2a} - m_2 x_{2b} + p(\xi_{1j} x_1 + y^*_{1j}) + s(1 - \xi_{1,j}) x_1 - o y^*_{1j} + r y^*_{2j})$. The optimal first stage decision $\boldsymbol{x}^*$ is given by the multivariate mode of the marginal distribution of decision variables;

$$\pi(\boldsymbol{x}) \propto [\int u(\boldsymbol{x}, \boldsymbol{\xi}_j) p(\xi_{2j}) p(\xi_{1j}|\boldsymbol{x}) d\boldsymbol{\xi}_j]^J \mathbb{I}(x_1 \leqslant PC, x_1 \geqslant 0, x_{2a}, x_{2b} \in \{0,1\}, x_{2a} + x_{2b} \leqslant 1)$$

The major challenge of APS is to sample from these augmented and marginal distributions. We utilize Markov chain Monte Carlo (MCMC) methods, particularly Metropolis Hastings within Gibbs sampling algorithm. A Gibbs sampler is constructed using conditional distributions $\pi(\boldsymbol{x}|\boldsymbol{\xi}_J)$ and $\pi(\boldsymbol{\xi}_j|\boldsymbol{x})$ for every j^{th} copy, which are

$$\pi(x|\boldsymbol{\xi}_J) \propto \prod_{j=1}^{J} u(\boldsymbol{x}, \boldsymbol{\xi}_j) p(\boldsymbol{x}) \mathbb{I}(x_1 \leqslant PC, x_1 \geqslant 0, x_{2a}, x_{2b} \in \{0,1\}, x_{2a} + x_{2b} \leqslant 1)$$

and

$$\pi(\boldsymbol{\xi}_j|\boldsymbol{x}) \propto u(\boldsymbol{x}, \boldsymbol{\xi}_j) p(\xi_{2j}) p(\xi_{1j}|\boldsymbol{x}), \ \forall\, j = 1, 2, .., J.$$

The draws from these full conditional distributions provide samples from the joint distribution under certain conditions, see Casella and George (1992) for an overview of Gibbs sampling. However, these are nonstandard distributions which are not possible to directly sample from. Hence, we use the Metropolis-Hastings algorithm, which is based on sampling from a proposal distribution which results with an irreducible and aperiodic Markov chain (Chib and Greenberg (1995)). The details to sample from $\pi(\boldsymbol{\xi}_j|\boldsymbol{x})$ follows. We use multivariate Normal distribution as the proposal density, $h(\boldsymbol{\xi}_j)$. The choice of proposal density does not affect the theoretical convergence properties, but it impacts the convergence rate, see Bielza et al. (1999) for a detailed discussion. For a MCMC algorithm with G iterations, let's summarize the procedure for the $(g)^{th}$ draw. We draw $\boldsymbol{\xi}_{*j}^{(g)}$ from the proposal density $h(\boldsymbol{\xi}_j)$. The proposal density has the current draw, $\boldsymbol{\xi}_j^{(g-1)}$ as mean and a pre-specified variance covariance matrix that can result in reasonable acceptance rate such as $30-40\%$. This is repeated J times, and each draw is referred to as the j^{th} candidate draw in the g^{th} iteration. Next, we compare this candidate draw with the existing $(g-1)^{th}$ draw, $\boldsymbol{\xi}_j^{(g-1)}$. The acceptance probability is

$$\alpha^{(g)}(\boldsymbol{\xi}_j, \boldsymbol{\xi}_{*j}) = min\left[\frac{(u(\boldsymbol{x}, \boldsymbol{\xi}_{*j}^{(g)})h(\xi_{*j}^{(g)})}{(u(\boldsymbol{x}, \boldsymbol{\xi}_j^{(g-1)})h(\xi_j^{(g-1)})}, 1\right].$$

If the Uniform draw $U(0,1)$ is smaller than this acceptance probability, the candidate draw is accepted and the value of $\boldsymbol{\xi}_j^{(g)}$ is set as $\boldsymbol{\xi}_{*j}^{(g)}$; otherwise the value of $\boldsymbol{\xi}_j^{(g)}$ is kept as $\boldsymbol{\xi}_j^{(g-1)}$. This is repeated for G iterations or until assessment of practical convergence of the Markov chain. A similar procedure is used to sample from $\pi(\boldsymbol{x}|\boldsymbol{\xi}_J)$, but not described in detail to preserve space. Once the MCMC convergence is attained, the marginal samples from $p(\boldsymbol{x})$, can be retrieved, and their mode would provide the optimal decision.

APS results in a number of computational advantages. First, one does not have to generate scenarios for each decision alternative which grows exponentially for endogenous problems. The draws of the random variable are tilted away from the probability density $p(\boldsymbol{\xi}_j|\boldsymbol{x})$ toward $u(\boldsymbol{x}, \boldsymbol{\xi}_j)p(\xi_{1j}|\boldsymbol{x})p(\xi_{2j})$. The algorithm samples "smart" values of ξ_j where the importance function is the objective function that tightens around the optimal decision with the convergence. The proposed algorithm draws samples of the decision variable from the regions of the decision space with high objective function values. This results in reduced sample variance and smaller Monte Carlo error (Ekin et al. (2014)).

The choice of J affects the convergence rate. For practical purposes, instead of a formal implementation, we use a diagnostic approach that increases $J^{(g)}$ until MCMC draws stabilize (Gelman and Rubin (1992)). Then, it is sufficient to sample from the joint density and to estimate the mode along the chain for a value of J on a given schedule, which depends on the flatness of the objective function surface (Jacquier et al. (2007)).

The asympotic properties are discussed in detail and the proof of the convergence of the APS based stochastic optimization method is provided by Bielza et al. (1999) and Ekin et al. (2014). The convergence of MCMC, particularly of the Markov chain as a function of the length of chain, G, is well studied in literature, see the overview of Gamerman and Lopes (2006). It is a standard result that $(x, \boldsymbol{\xi}_J^g)_{g=1}^G$ become draws from the joint distribution, $\pi_J(x, \boldsymbol{\xi}_J)$ as $G \to \infty$. For a fixed value of J, Pincus (1970) and Tierney (1994) show the MCMC convergence in number of iterations, G. In a relevant work, Müller et al. (2004) show that a time inhomogeneous Markov chain, whose draws converge to $\pi_\infty(x)$ with a Dirac measure on the desired optimum, can be constructed by increasing $J^{(g)}$ at a logarithmic rate. We utilize Brooks-Gelman-Rubin (BGR) statistics (Brooks and Roberts (1998)) to assess practical convergence. BGR values that are smaller than 1.10 are assumed to be enough to judge that the respective Markov Chains have practically converged (Gill (2014)).

5 Numerical Illustration

This section presents the computational results for various production environments. We investigate the trade-offs among the uncertain yield, optimal production and maintenance decisions and objective function. First, we report and analyze the results for $7,776$ test problems that consider different specifications of the parameters. Then, we use particular cases to investigate the sensitivity of the results with respect to changes in some parameters and discuss the potential use of other distributions. We also present a brief illustration of the solution method.

5.1 Overview of Test Problems

For the comprehensive numerical study, each parameter is tested for at least two levels which resulted in a total of $7,776$ test problems. Table 1 presents the summary of the parameters and their levels. The parameter levels are chosen to provide a fair coverage of different production settings. We assume $o \geqslant p \geqslant c \geqslant r \geqslant s$. The values of o, p, c, r and s are chosen to reflect cases where they are close or very different from each other. Our focus has been on production environments where the maintenance cycle is short compared to production. Therefore, the values of unit maintenance costs, m_1 and m_2, are chosen so that the overall maintenance cost is comparable to the overall production cost.

A new machine which is not expected to be impacted by production much, is assigned an α value of 0.01. Whereas, an α of 1 corresponds to an old machine, therefore an increase in production quantity is expected to have a larger impact on the yield. We test three levels of product reliability parameter vector to consider the impact of maintenance. For instance, when $\boldsymbol{\beta} = [1, 2, 3]$, the impact of maintenance on yield is small. Whereas for the third level,

Parameters	Description	Factor values
p	unit product sale price	25, 50, 100
c	unit product cost	10, 15, 20
o	unit outsourcing cost	100, 200
s	unit scrapping price	1, 3
r	unit salvage price	3, 9
m_1	cost of the preventive maintenance	1,000, 10,000
m_2	cost of the corrective maintenance	20,000, 50,000
α	production-yield impact parameter	0.01, 0.1, 1
$\boldsymbol{\beta}$	product reliability parameter vector	[1,2,3], [1,5,10], [1,10,20]
μ_d	mean demand	300, 500, 1000

Table 1: Summary of parameters and their levels for test problems

$\boldsymbol{\beta} = [1, 10, 20]$, even preventive maintenance has a huge impact on yield, and makes the machine almost as good as new. The levels of mean demand are chosen to reflect the cases where it is lower, equal to and higher than the production capacity, which is assumed to be 500 for all cases. The standard deviation of yield, σ_{ξ_1}, is assumed to be relatively low, 0.01, whereas the standard deviation of demand, σ_d, is set as 10. These let us to demonstrate the interplay among the first and second stage decisions and yield. However, later in this section, we analyze the sensitivity of the solutions to the changes in these parameters as well.

5.2 Analysis of Results

Table 2 presents the average function values over all test problems for various levels. For instance, for the parameter p at level $p = 25$, there are $2,592$ test problems. We compute the relevant functions for all these cases and compute their averages. For each parameter level, we report the average values of the optimal decisions as well as the expected objective function and yield values given those optimal decisions. This is followed by the values of the optimal second stage decisions computed for the combination of the optimal first stage decisions and the expected values of random variables.

The proposed model with the endogenous structure reveals a number of insights about the trade-offs within a production environment. Figure 1 presents the histogram of all optimal objective function values that represents the wide spectrum of production settings. A few high negative optimal objective function values explain the negative averages for some

Parameter	Level	x_1^*	x_2^*	$E[u(\boldsymbol{x}^*, \boldsymbol{\xi})]$	$E[\xi_1 \| \boldsymbol{x}^*]$	$y_1^*(\boldsymbol{x}^*, E[\boldsymbol{\xi}])$	$y_2^*(\boldsymbol{x}^*, E[\boldsymbol{\xi}])$
	All	409.13	1.43	-8,363	0.88	244.91	4.74
p	25	409.08	1.43	-28,365	0.88	244.93	4.74
	50	409.08	1.43	-13,362	0.88	244.87	4.71
	100	409.22	1.43	16,637	0.88	244.94	4.76
c	10	413.59	1.43	-6,312	0.88	244.64	7.21
	15	408.68	1.43	-8,374	0.88	244.87	4.48
	20	405.12	1.43	-10,402	0.89	245.23	2.55
o	100	406.41	1.39	3,491	0.88	248.15	4.18
	200	411.85	1.47	-20,218	0.89	241.68	5.30
s	1	409.37	1.43	-8,416	0.88	244.96	4.93
	3	408.89	1.43	-8,310	0.88	244.87	4.54
r	3	406.21	1.43	-8,369	0.88	244.88	2.66
	9	412.05	1.43	-8,357	0.88	244.95	6.82
m_1	1,000	417.82	1.58	-5,625	0.90	232.64	4.87
	10,000	400.44	1.30	-11,101	0.87	257.19	4.60
m_2	20,000	409.67	1.44	-8,365	0.88	243.30	4.75
	50,000	409.60	1.41	-8,360	0.88	246.52	4.72
α	0.01	438.74	1.00	2,608	0.99	172.47	5.80
	0.1	423.89	1.34	-2,072	0.95	203.17	6.73
	1	364.76	1.95	-25,626	0.71	359.10	1.67
$\boldsymbol{\beta}$	[1,2,3]	391.20	1.34	-11,936	0.82	273.23	4.15
	[1,5,10]	416.85	1.47	-7,722	0.89	239.13	4.19
	[1,10,20]	419.34	1.47	-5,430	0.94	222.39	5.87
μ_d	300	334.94	1.30	8,713	0.89	20.88	14.22
	500	446.09	1.49	6,009	0.88	107.01	0.00
	1000	446.35	1.49	-39,813	0.88	606.85	0.00

Table 2: The average function values over all test problems for various levels

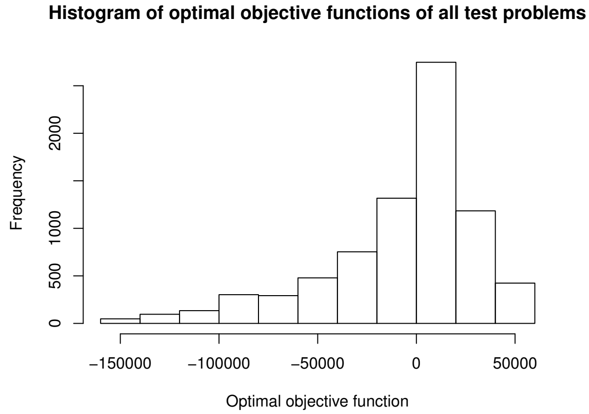

Figure 1: Histogram of optimal objective functions, $E[u(\boldsymbol{x}^*, \xi)]$, of all test problems

levels. It makes sense that average optimal objective function values become higher for increasing levels of p, s and r and decreasing levels of o and c. It may initially found to be counterintuitive that the average optimal objective function is lower for higher unit corrective maintenance cost level. However, this can be explained by optimal maintenance decision not being the corrective maintenance since it is very expensive. The decrease in average optimal objective function values for the level of $\alpha = 1$ can also be recognized and emphasizes the importance of machine condition. Another large decrease in average optimal objective function occurs for mean demand values that are greater than the maximum production quantity. Large demand and lack of production opportunities result in higher outsourcing quantities in order to satisfy the demand, which increase the expected cost.

Overall, the average values for different levels of p, c, o, s, and r are very similar to the overall averages. The unit preventive maintenance cost has a relatively large impact on the average results. For instance, when the unit preventive maintenance cost, m_1, is $1,000$, the optimal maintenance decision in all $3,888$ test problems are either preventive or zero maintenance, $2,166$ and $1,722$ times respectively. Whereas, when the unit preventive maintenance cost becomes more expensive ($m_1 = 10,000$), the frequencies of optimal maintenance decisions are $2,870$, 874 and 144 for zero, preventive and corrective maintenance respectively. Similarly, when the unit corrective maintenance cost becomes expensive $(50,000)$, none of the test problems results with the optimal decision of corrective maintenance.

In the test problems with smaller α values, the impact of production on yield is small, therefore the need for any maintenance action is not imminent. For instance, the optimal maintenance decision is zero maintenance for all $2,592$ test problems where α is 0.01. These test problems also result in high expected yield, 0.99. On the other hand, test problems with large α values require some level of maintenance to prevent yield from being very small. For instance, when $\alpha = 1$, preventive maintenance is the optimal decision $2,170$ times within $2,592$ test problems. This large impact of production on yield also results with smaller levels of production quantity, and higher amount of outsourced products. On the other hand, higher product reliability parameters result in outcomes with small expected deterioration rates and high expected yield. Particularly, the increases in β_3 result in higher number of corrective maintenance decision alternative as the optimal, higher yield, higher production, and lower cost values.

As expected, changes in the mean demand level compared to the maximum production quantity affect the average results. When the mean demand is 300, less than the maximum production quantity of 500, the optimal decision becomes zero maintenance $1,816$ times within $2,592$ test problems. For these cases, the level of outsourcing is lower compared to the average. The increase in mean demand is satisfied mainly by production with higher yield due to preventive maintenance and outsourcing. When the outsourcing is expensive, the decision maker conducts corrective maintenance instead of outsourcing products, which can improve the mean yield rate.

There are 144 cases where the corrective maintenance is the optimal decision. These all have close levels of unit preventive and corrective maintenance costs; $m_1 = 10,000$ and $m_2 = 20,000$ respectively. The average demand is either 500 or 1000, $\boldsymbol{\beta}$ is either at second or third level, and α is 1. These are compatible with general intuition that corrective maintenance is the better decision when it can improve the operation significantly at a reasonable cost difference.

Another observation can be made with regards to the cases where the unit cost of preventive maintenance is very cheap and the corrective maintenance is expensive ($m_1 = 1,000$ and $m_2 = 50,000$). A real world example would be the manufacturing environments in that relatively minor tool changes can increase the mean yield rate at a cheap rate compared to a

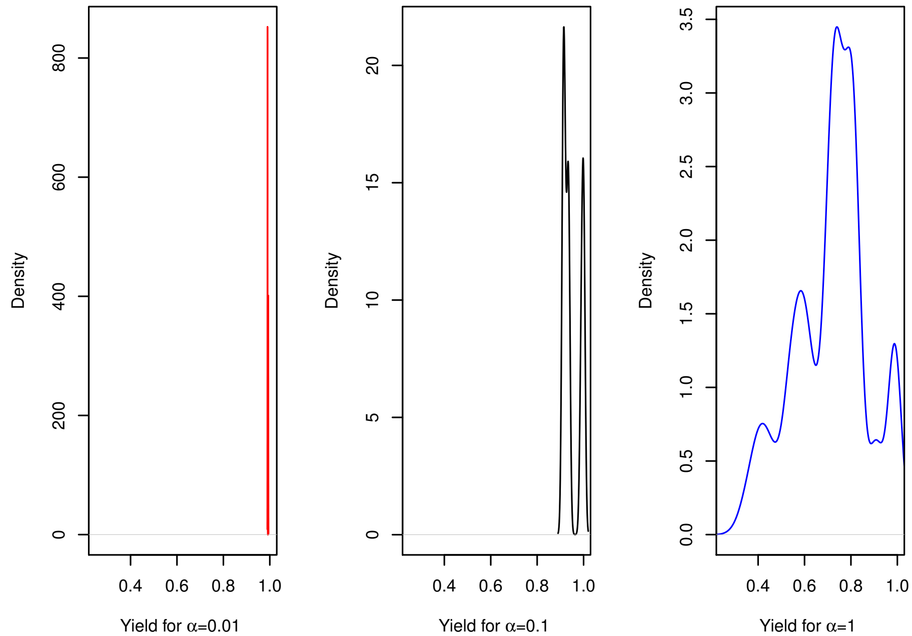

Figure 2: The density plots of yield of test problems for different levels of α”

replacement. Our model results with preventive maintenance as the optimal decision $1,062$ times in $1,944$ such test problems.

The variability of expected yield given the optimal decision for a given parameter level is also of interest. Figure 2 presents the kernel density plots of expected yield given the optimal decision for different levels of α. The changing levels of expected yield can be explained by the levels of α and the respective optimal decisions. When $\alpha = 0.01$, the expected yield is very high. Whereas a medium level of production-yield impact parameter ($\alpha = 0.1$) results in expected yield with a range of 0.91 and 1. In contrast, the range of values is large for cases with $\alpha = 1$ where the minimum expected yield is 0.28.

We retrieve the random variable realizations used in our MCMC algorithm, and compare them with the draws from the endogenous distribution given optimal decisions via mean absolute error function. The average estimation error is computed as 2.3% over all test problems.

5.3 Sensitivity Analysis

We analyze the sensitivity of the results with respect to the choice of standard deviation of yield and demand as well as the maximum production capacity. For brevity and ease of demonstration, we conduct the sensitivity analysis using a base case where $p = 25$, $c = 10$, $o = 100$, $s = 1$, $r = 3$, $\alpha = 1$, $\boldsymbol{\beta} = [1, 2, 3]$, $m_1 = 1,000$, $m_2 = 20,000$, $\mu_d = 300$. This case has optimal decisions of $x_1^* = 318$ and $x_2^* = 1$.

First, we test the model using higher standard deviations of yield, 0.05 and 0.10. Table 3 presents the optimal function values for this case. The optimal decision for production quantity increases when the standard deviation of yield increases. The optimal maintenance decision for all cases is zero maintenance, and the expected yield is very high, at 0.99. For increasing standard deviation, the expected objective function value decreases, mainly because of the increase in outsourced amounts.

Parameter	Level	x_1^*	x_2^*	$E[u(\boldsymbol{x}^*, \boldsymbol{\xi})]$	$E[\xi_1 \vert \boldsymbol{x}^*]$	$y_1^*(\boldsymbol{x}^*, E[\boldsymbol{\xi}])$	$y_2^*(\boldsymbol{x}^*, E[\boldsymbol{\xi}])$
σ_ξ	0.01	318	1	4,367	0.99	0.00	14.45
	0.05	335	1	4,227	0.99	0.00	20.78
	0.10	365	1	3,985	0.99	0.00	34.88

Table 3: The optimal function values for different values of σ_ξ

Second, we repeat our analysis with different standard deviation values for demand, particularly 5 and 30. Table 4 presents the optimal function values for this case. It can be argued that less variability results in higher expected profit. The production quantity for the higher demand variance of 30 is the highest to address potentially high demand. Hence, this results in excess production and preventive maintenance which increase the costs.

Parameter	Level	x_1^*	x_2^*	$E[u(\boldsymbol{x}^*, \boldsymbol{\xi})]$	$E[\xi_1 \vert \boldsymbol{x}^*]$	$y_1^*(\boldsymbol{x}^*, E[\boldsymbol{\xi}])$	$y_2^*(\boldsymbol{x}^*, E[\boldsymbol{\xi}])$
σ_d	10	318	1	4,367	0.99	0.00	14.45
	5	311	1	4,416	0.99	0.00	7.68
	30	334	2	-9,414	0.55	127.23	0.00

Table 4: The optimal function values for different values of σ_d

Then, we test for the sensitivity of decisions with respect to changes in PC, maximum production capacity. We have tested for PC values of 100 and $1,000$, in addition to the base value of 500. Table 5 presents the optimal function values for this case. Lower production capacity results with the production quantity becoming maximum possible, 100 at an expected yield of 0.99. The outsourcing is done to satisfy the demand, which increases the expected costs.

Parameter	Level	x_1^*	x_2^*	$E[u(\boldsymbol{x}^*, \boldsymbol{\xi})]$	$E[\xi_1 \vert \boldsymbol{x}^*]$	$y_1^*(\boldsymbol{x}^*, E[\boldsymbol{\xi}])$	$y_2^*(\boldsymbol{x}^*, E[\boldsymbol{\xi}])$
PC	500	318	1	4,367	0.99	0.00	14.45
	100	100	1	-13,627	0.99	201.28	0.00
	1,000	318	1 4,367	0.99	0.00	15.07	0.00

Table 5: The optimal function values for different values of PC

Lastly, we briefly present the implementation details of the proposed method for the base case. We set M as $50,000$ to shift the objective function to a positive region. The standard deviation of the proposal density for demand and yield are set as 0.01 and 3. They are assumed to be independent, so covariance values of the multivariate Normal proposal density are set as 0. Practical assessment of MCMC convergence can also be judged from the trace plots, see Figure 3 that presents the trace plots of x_1 for different values of J ($J = 5, 10, 25, 50$). The variance of the draws decreases and the draws stabilize for higher values of J. The left panel of Figure 4 presents the density plot of the production quantity, x_1. The vertical dashed line corresponds to the mode of the marginal draws which is the optimal decision. The right panel of Figure 4 presents the bar plot of the maintenance decision, x_2. It can easily be judged that zero maintenance, $x_2 = 1$ is the optimal decision.

It should be noted that practical convergence is assessed after $10,000$ iterations with $J = 50$ for all test problems. As expected, the variability of the draws of the decision variable increases and convergence takes longer when the standard deviation of the uncertain variables are higher.

The proposed model and methodology are general enough to accommodate any probability distribution. Beta distribution can be argued to be a viable alternative since it is defined in the space of $(0, 1)$. However, it is difficult for the decision maker to specify the variance independently from the mean within an endogenous model. If the decision maker has available information about the production environment, another alternative would be to use triangular distribution by specifying the minimum, maximum and mode of the yield.

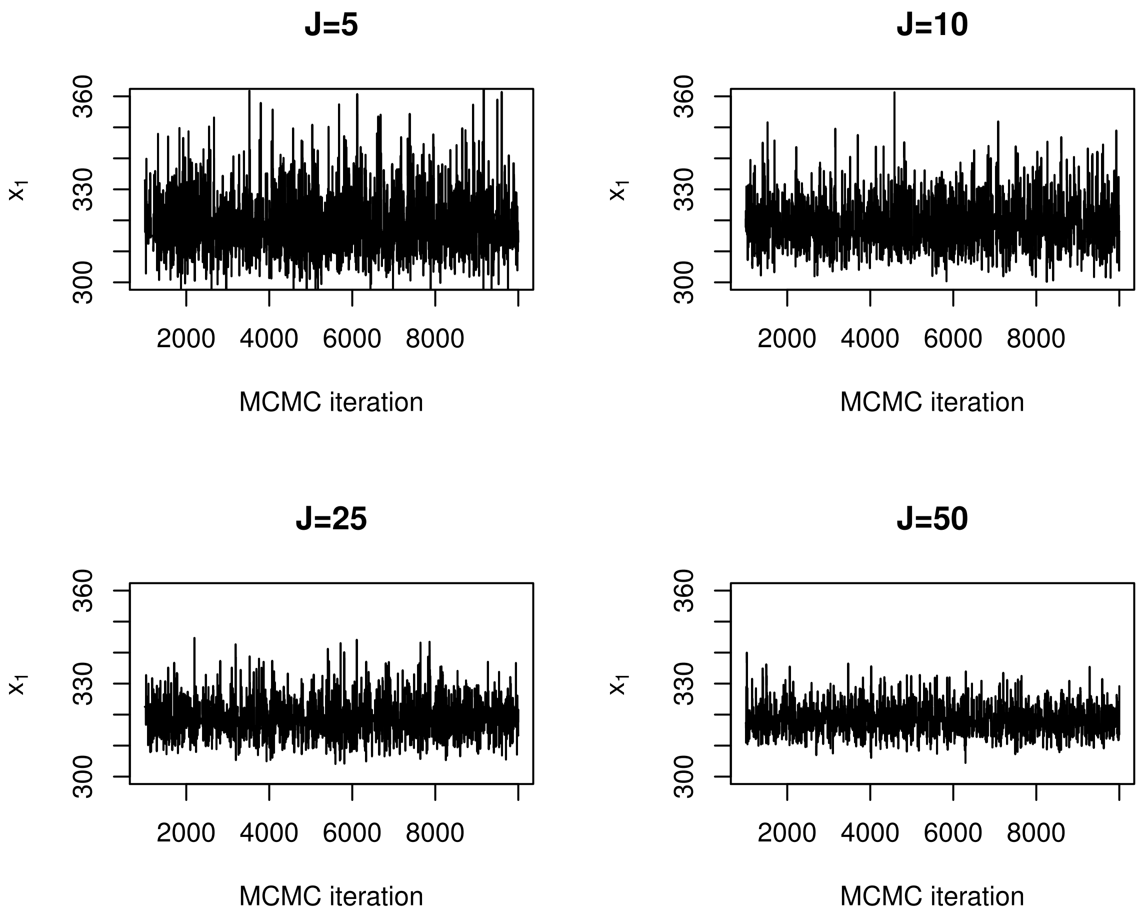

Figure 3: Trace plot of the decision variable for different values of J

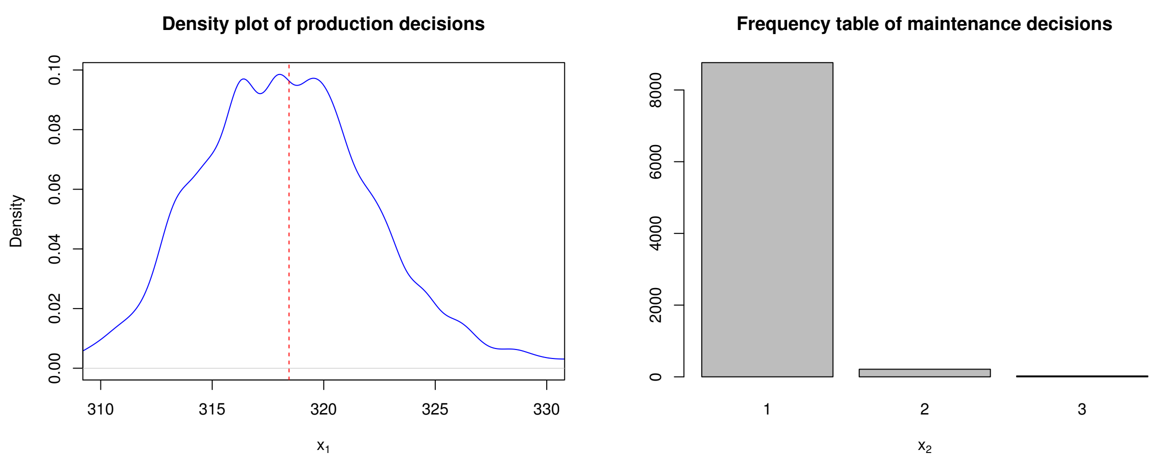

Figure 4: Density plot (left) and bar plot (right) of the decision variables

However, endogeneity also makes the choice of these parameters complicated. The decision maker can set the minimum and maximum values as 0 and 1, and model the mode using an endogenous function. These distributions relax the symmetry assumption of Normal distribution at the expense of further complexity in modeling endogenous mean and independent variance.

6 Conclusion and Directions for Future Work

In this study, we present a novel integrated decision making framework for production planning and maintenance under endogenous uncertain yield. This is the first such approach that treats the uncertain yield using a stochastically proportional model with endogenous mean. This lets the decision maker measure the effects of previous decisions on production environment and machine condition. In addition, the variance of the yield rate can be specified independently from the endogenous mean and the batch size. Our two-stage nonlinear stochastic program formulation allows for random demand and recourse decisions for outsourcing and salvaging. Finding the solution for this nonlinear stochastic recourse model with endogenous uncertainty is not straightforward and had not been considered previously. This model is solve by APS, which is general enough to accommodate any objective function, probability distribution and number of decision alternatives. This is one of the first applications of APS based optimization method, the first to solve a stochastic recourse model that includes direct endogenous randomness or nonlinearity. We demonstrate the use of the proposed approach by conducting a comprehensive numerical study and sensitivity analysis. We discuss the trade-offs involved among the yield, and simultaneous decisions of production planning and maintenance.

One of the limitations of the proposed model is the assumption of negligible maintenance time. Treating the equipment repair time as a random variable may make the model more general at the expense of further complexity. Another improvement would be to extend the model to production environments that include multiple product types, each with a different yield distribution. Special forms of the proposed model can be considered such as the models with only maintenance or only production dependent mean yield rates. For instance, models with production dependent yield rate can be applicable for systems with equipments that have long life time, and can result in high expected yield without any need for maintenance for a long period. In a number of production settings where the production is done in batches and specifying the yield variance is not important, the decision maker can use a discrete yield distribution within an optimization model that assumes integer production quantities. For instance, the decision maker can utilize a binomial yield function to model the number of defective items by assuming that defectiveness of each item is independent and identically Bernoulli distributed. Then, the probability of a defective item can be modeled as a function of previous decisions.

Lastly, the inference of the yield distribution is of further interest. The issue of learning from the data is complicated since one cannot observe the random variables or generate scenarios directly but only can observe the function value dependent on the previous decisions. This is similar to multi armed bandit problems, but needs to be investigated further for stochastic endogenous problems.

References

Ahmed, S. (2000). *Strategic planning under uncertainty: Stochastic integer programming approaches*. PhD thesis, University of Illinois Urbana Champaign.

Aktekin, T. and Ekin, T. (2016). Stochastic call center staffing with uncertain arrival, service and abandonment rates: A Bayesian perspective. *Naval Research Logistics*, 63(6):460–478.

Aramon Bajestani, M., Banjevic, D., and Beck, J. C. (2014). Integrated maintenance planning and production scheduling with markovian deteriorating machine conditions. *International Journal of Production Research*, 52(24):7377–7400.

Batun, S. and Maillart, L. M. (2012). Reassessing tradeoffs inherent to simultaneous maintenance and production planning. *Production and Operations Management*, 21(2):396–403.

Bielza, C., Müller, P., and Ríos Insua, D. (1999). Decision analysis by augmented probability simulation. *Management Science*, 45(7):995–1007.

Birge, J. and Louveaux, F. (2011). *Introduction to Stochastic Programming. Second edition.* Springer-Verlag, Berlin.

Brooks, S. P. and Roberts, G. O. (1998). Assessing convergence of Markov chain Monte Carlo algorithms. *Statistics and Computing*, 8(4):319–335.

Casella, G. and George, E. (1992). Explaining the Gibbs sampler. *The American Statistician*, 46(3):167–174.

Chib, S. and Greenberg, E. (1995). Understanding the Metropolis-Hastings algorithm. *The American Statistician*, 49(4):327–335.

Dyer, M. and Stougie, L. (2006). Computational complexity of stochastic programming problems. *Mathematical Programming*, 106(3):423–432.

Ekin, T., Polson, N. G., and Soyer, R. (2014). Augmented Markov chain Monte Carlo simulation for two-stage stochastic programs with recourse. *Decision Analysis*, 11(4):250–264.

Fragnière, E., Gondzio, J., and Yang, X. (2010). Operations risk management by optimally planning the qualified workforce capacity. *European Journal of Operational Research*, 202(2):518–527.

Gamerman, D. and Lopes, H. (2006). *Markov chain Monte Carlo: Stochastic simulation for Bayesian inference. Second edition.* Chapman & Hall.London, United Kingdom.

Gelman, A. and Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4):457–472.

Gill, J. (2014). *Bayesian methods: A social and behavioral sciences approach*, volume 20. CRC press. Boca Raton, Florida.

Goel, V. and Grossmann, I. (2004). A stochastic programming approach to planning of offshore gas field developments under uncertainty in reserves. *Computers & Chemical Engineering*, 28(8):1409–1429.

Goel, V. and Grossmann, I. (2006). A class of stochastic programs with decision dependent uncertainty. *Mathematical Programming*, 108(2-3):355–394.

Hadidi, L., Al-Turki, U., and Rahim, A. (2012). Integrated models in production planning and scheduling, maintenance and quality: a review. *International Journal of Industrial and Systems Engineering*, 10(1):21–50.

Homem-de Mello, T. and Bayraksan, G. (2014). Monte Carlo sampling-based methods for stochastic optimization. *Surveys in Operations Research and Management Science*, 19(1):56–85.

Hong, H., Zhou, W., Zhang, S., and Ye, W. (2014). Optimal condition-based maintenance decisions for systems with dependent stochastic degradation of components. *Reliability Engineering & System Safety*, 121:276–288.

Iravani, S. and Duenyas, I. (2002). Integrated maintenance and production control of a deteriorating production system. *IIE Transactions*, 34(5):423–435.

Irie, K. and West, M. (2016). Bayesian emulation for optimization in multi-step portfolio decisions. *arXiv preprint arXiv:1607.01631*.

Jacquier, E., Johannes, M., and Polson, N. (2006). Maximum expected utility via MCMC. Available at http://faculty.chicagobooth.edu/nicholas.polson/research/papers/mcmceu.pdf.

Jacquier, E., Johannes, M., and Polson, N. (2007). MCMC maximum likelihood for latent state models. *Journal of Econometrics*, 137(2):615–640.

Jonsbraten, T. W., Wets, R. J.-B., and Woodruff, D. L. (1998). A class of stochastic programs with decision dependent random elements. *Annals of Operations Research*, 82:83–106.

Kalirajan, K. (1981). An econometric analysis of yield variability in paddy production. *Canadian Journal of Agricultural Economics/Revue canadienne d'agroeconomie*, 29(3):283–294.

Kall, P. and Mayer, J. (2005). Stochastic linear programming: models, theory and computation. International series in operations research and management science.

Kazaz, B. and Sloan, T. W. (2013). The impact of process deterioration on production and maintenance policies. *European Journal of Operational Research*, 227(1):88–100.

Khouja, M. and Mehrez, A. (1994). Economic production lot size model with variable production rate and imperfect quality. *Journal of the Operational Research Society*, 45(12):1405–1417.

Kirschenmann, T., Popova, E., Damien, P., and Hanson, T. (2014). Decision dependent stochastic processes. *European Journal of Operational Research*, 234(3):731–742.

Mula, J., Poler, R., Garcia-Sabater, J., and Lario, F. (2006). Models for production planning under uncertainty: A review. *International Journal of Production Economics*, 103(1):271–285.

Müller, P. (1999). Simulation-based optimal design. In Bernardo, J. and Berger J.O., A.P.Dawid, A., editors, *Bayesian Statistics 6*, pages 459–474. Oxford University Press.

Müller, P., Sansó, B., and De Iorio, M. (2004). Optimal Bayesian design by inhomogeneous Markov chain simulation. *Journal of the American Statistical Association*, 99(467):788–798.

Peeta, S., Salman, F. S., Gunnec, D., and Viswanath, K. (2010). Pre-disaster investment decisions for strengthening a highway network. *Computers & Operations Research*, 37(10):1708–1719.

Peng, H. and van Houtum, G.-J. (2016). Joint optimization of condition-based maintenance and production lot-sizing. *European Journal of Operational Research*, 253(1):94–107.

Pincus, M. (1970). A Monte Carlo method for the approximate solution of certain types of constrained optimization problems. *Operations Research*, 18(6):1225–1228.

Rubinstein, R. Y. and Kroese, D. P. (2011). *Simulation and the Monte Carlo method*, volume 707. John Wiley & Sons. New York.

Sana, S. (2010a). An economic production lot size model in an imperfect production system. *European Journal of Operational Research*, 201(1):158–170.

Sana, S. S. (2010b). A production–inventory model in an imperfect production process. *European Journal of Operational Research*, 200(2):451–464.

Sloan, T. (2004). A periodic review production and maintenance model with random demand, deteriorating equipment, and binomial yield. *Journal of the Operational Research Society*, 55(6):647–656.

Sloan, T. (2008). Simultaneous determination of production and maintenance schedules using in-line equipment condition and yield information. *Naval Research Logistics*, 55(2):116–129.

Sloan, T. and Shanthikumar, J. (2000). Combined production and maintenance scheduling for a multiple-product, single-machine production system. *Production and Operations Management*, 9(4):379–399.

Solak, S., Clarke, J., Johnson, E., and Barnes, E. (2010). Optimization of R & D project portfolios under endogenous uncertainty. *European Journal of Operational Research*, 207(1):420–433.

Tierney, L. (1994). Markov chains for exploring posterior distributions(with discussion). *The Annals of Statistics*, 22(4):1701–1762.

Van Slyke, R. and Wets, R. (1969). L-shaped linear programming with application to optimal control and stochastic programming. *SIAM Journal on Applied Mathematics*, 17(4):638–663.

Wallace, S. W. and Ziemba, W. T. (2005). *Applications of stochastic programming*, volume 5. Siam.

Wang, H. (2002). A survey of maintenance policies of deteriorating systems. *European Journal of Operational Research*, 139(3):469–489.

Yano, C. and Lee, H. (1995). Lot sizing with random yields: A review. *Operations Research*, 43(2):311–334.