

DETC2007-35158

UPDATING UNCERTAINTY ASSESSMENTS: A COMPARISON OF STATISTICAL APPROACHES

Jason Matthew Aughenbaugh

Applied Research Laboratories
The University of Texas at Austin
Austin, TX 78759
jaughenbaugh.me03@gtalumni.org

Jeffrey W. Herrmann

Department of Mechanical Engineering
and Institute for Systems Research
University of Maryland
College Park, MD 20742
jwh2@umd.edu

ABSTRACT

The performance of a product that is being designed is affected by variations in material, manufacturing process, use, and environmental variables. As a consequence of uncertainties in these factors, some items may fail. Failure is taken very generally, but we assume that it is a random event that occurs at most once in the lifetime of an item. The designer wants the probability of failure to be less than a given threshold. In this paper, we consider three approaches for modeling the uncertainty in whether or not the failure probability meets this threshold: a classical approach, a precise Bayesian approach, and a robust Bayesian (or imprecise probability) approach. In some scenarios, the designer may have some initial beliefs about the failure probability. The designer also has the opportunity to obtain more information about product performance (e.g. from either experiments with actual items or runs of a simulation program that provides an acceptable surrogate for actual performance). The different approaches for forming and updating the designer's beliefs about the failure probability are illustrated and compared under different assumptions of available information. The goal is to gain insight into the relative strengths and weaknesses of the approaches. Examples are presented for illustrating the conclusions.

1. INTRODUCTION

Engineers make decisions based on their beliefs. These beliefs depend upon the information that the engineer has gathered, and they can change based on new information (e.g. from either experiments with actual items or runs of a simulation program that provides an acceptable surrogate for actual performance). Statistical reasoning includes a variety of approaches for

updating beliefs based on new information. Most engineers have little training on these approaches, and some approaches may be completely unfamiliar, despite their usefulness and relevance. The main goals of this article are to describe how one can apply different statistical approaches to update beliefs about the expected performance of a product and to discuss the relative strengths and weaknesses of each approach.

To meet these goals, a specific design decision-making problem is considered under a number of different scenarios. For each scenario, this paper shows how different statistical approaches can be used to update beliefs based on the results of new information (in this case, the results of additional testing).

We assume that the designer has framed the decision-making problem correctly, and that the designer may want to obtain and use more information to update his beliefs. In part, acceptance of this analysis will depend upon the reader's willingness to accept multiple theories of probability and statistics. Those who are highly committed to a specific philosophical position may find the arguments here unconvincing. This paper is intended for those who are open to trying any of these approaches but are unsure of their practical consequences.

From a practical perspective, it can be argued that the following properties are desirable for a procedure of testing and analyzing reliability (or more generally probability) data. First, it can be valuable to incorporate existing information that may be relevant into the analysis; otherwise, existing information is essentially wasted. Second, the procedure should allow for new information to be incorporated into the estimate. Third, for experimental design and planning, the procedure should help the engineers to determine if they need more information, and if so, how much.

The analysis presented here implies that designers should carefully consider how to update their beliefs. It will help designers understand (though only in a qualitative sense) the value of the information they gather. It illustrates how updating beliefs using a poor approach can lead to poor decisions.

This paper begins by stating the problem and then describes and compares the statistical approaches that will be considered. The paper introduces three scenarios that correspond to situations in which the designer has no information, substantial information, and partial information. Each statistical approach is considered for each scenario. A summary concludes the paper.

2. PROBLEM STATEMENT

The performance of a product that is being designed is unpredictable and often involves random behavior. Multiple items (i.e. instances of the product being designed) will be manufactured, sold, and used. Ideally, these items are identical. In practice, however, differences exist due to variations in materials and the manufacturing processes used to create them. Moreover, different items are used in different environments and in different ways.

We assume that the possibility of failures, though disagreeable to the designer, is unavoidable for the above reasons. Some items will fail, while others will not fail. Here, “failure” is taken very generally, but we assume it is a one-time event in the lifetime of an item. In practice, the failure could be that the item’s performance did not meet certain requirements during a test or that it catastrophically stops working at some point during its lifetime.

We assume that the designer wants the percentage of items that fail to be below a given threshold. In practice, this requirement would be based on customer expectations and the firm’s tolerance of risk. We assume that each item’s failure is a random Bernoulli process. Let θ be the failure probability, which is the parameter of interest. Let θ_{crit} be the critical threshold such that the designer wants $\theta \leq \theta_{crit}$.

Ideally, the designer would have enough data to make a precise assessment of θ , such as “ $\theta = 0.01$.” However, as discussed later in this paper, there are many practical reasons why the designer cannot or is unwilling to make a precise assessment despite holding some initial beliefs about θ .

We assume that the designer has the opportunity to obtain additional data from testing in order to update his beliefs. The designer can perform n tests in which the performance of the product—success or failure—is observed. We assume that test performance is the actual performance of interest or is an acceptable surrogate for actual performance. Let m be the number of failures observed (and consequently $n - m$

successes are observed). For convenience, we define a variable $x_i = 1$ if trial i is a failure, and $x_i = 0$ otherwise. Then $m = \sum_{i=1}^n x_i$. We assume each test is an independent Bernoulli trial, so m is a binomial random variable. The following function is the probability distribution for the number of failures in n trials, where j ranges from 0 to n .

$$P\{m = j\} = \binom{n}{j} \theta^j (1-\theta)^{n-j} \quad (1)$$

This paper considers two interrelated questions. First, given the set of results $X = \{x_i\}_{i=1}^n$, how should the designer update his beliefs about the failure probability? Second, how can the designer decide if the failure probability is acceptable? Answering these questions requires one to understand the relative strengths and weaknesses of the approaches, which may depend upon the scenario that occurs. This paper compares three approaches across three different scenarios.

3. STATISTICAL APPROACHES FOR UPDATING PARAMETER ESTIMATES

This paper compares several approaches for analyzing data and updating beliefs about the parameter θ in different scenarios. The approaches can be classified into three categories:

1. Classical sampling theory,
2. Precise Bayesian, and
3. Robust Bayesian.

These approaches are described in the following subsections.

3.1. Classical sampling theory approach

Most engineers are familiar with classical, sampling theory approaches to statistical analyses, as these are generally emphasized in introductory texts such as [1]. Standard in these approaches is the adoption of a frequentist interpretation of probabilities. Under a frequentist interpretation, a probability represents the limit of the ratio of times that one outcome occurs compared to the total number of outcomes in an endless series of identical trials. For example, the probability of *heads* for a fair coin is 0.5 because, as the number of flips tends to infinity, the fraction of flips with the result *heads* approaches 0.5.

One can estimate an unknown probability from a random sample, but, because the sample is random, the resulting conclusions are random and may be incorrect. Thus, we use confidence intervals to understand the probability of an incorrect conclusion. For example, if we repeatedly used a process to generate a 95% confidence interval, 95% of the resulting confidence intervals would contain the true value.

This approach focuses entirely on the observed data, that is the $\{x_i\}_{i=1}^n$. An unbiased point estimate of θ is the relative

frequency of failures to trials in the sample. Specifically, one can estimate θ as $\hat{\theta} = \sum_{i=1}^n x_i / n = m/n$. A commonly used approximate one-sided $100(1-\alpha)\%$ confidence interval that gives an upper bound on θ [1] is, (letting z_α be the $100(1-\alpha)$ percentile of the standard normal distribution):

$$\left[0, \frac{m}{n} + z_\alpha \sqrt{\frac{m}{n} \left(1 - \frac{m}{n}\right) \frac{1}{n}} \right] \quad (2)$$

We use the upper bound for this problem because the designer would like the failure probability to be below the threshold. If this interval includes θ_{crit} , then one cannot reliably conclude that $\theta \leq \theta_{crit}$ based on the test results. We note that other confidence intervals for estimating this parameter have been proposed. For a full discussion of them, see Brown *et al* [2].

3.2. Precise Bayesian approach

The Bayesian approach [3, 4] provides a way to combine existing knowledge and new knowledge into a single estimate. The term ‘‘Bayesian approach’’ comes from the use of Bayes’s Theorem for updating beliefs. A prior distribution $\pi(\theta)$ captures the designer’s current information about the failure probability θ . If we define Y to be additional information about θ (and let y be a particular set or realization of information), then the posterior distribution $\pi(\theta | Y = y)$ is the distribution after receiving information y . Bayes’s Theorem presents a way to update the prior using information from the experiment in order to form the posterior.

Bayes’s Theorem can be viewed as an extension of the rules of conditional probability, so its validity is generally accepted [4]. Some people use the term ‘‘Bayesian’’ to describe anything that uses Bayes’s Theorem. Others use the term to refer to any procedure that uses a subjective interpretation of probabilities (more precisely known as subjective Bayesianism or subjectivism). Under a subjective interpretation, a probability is an expression of belief based on an individual’s willingness to bet [5-8]. Bayes’s Theorem provides an objective way to update *any* initial probability with objective data. If the prior is a subjective prior, then the posterior is necessarily subjective as well. However, one can also use *objective Bayesian analyses* in which the prior distributions are defined without subjective input. In these cases, the analysis can be viewed as entirely objective. See [9] for further discussion and references.

One of the requirements of Bayesian analysis is a starting prior distribution. Many of the strongest arguments both for and against Bayesian analysis involve the prior. On the one hand, if a designer has existing information about θ , it is important to consider this information in the new estimate of θ . This can be true even if the available information is only subjective in nature; for example, experts are often paid handsomely for their opinions.

On the other hand, if the designer has no specific information about θ , the introduction of a particular prior may distort the posterior estimate. The objective selection of a prior distribution in the absence of relevant prior information is a topic of extensive research. The approaches proposed include the use of *non-informative priors* [10-12], maximum-entropy priors [13], and data-dependent empirical Bayes approaches [14].

In general, determining this posterior distribution can be computationally burdensome, as the output is a distribution over all possible values of θ . To support analytical solutions, the form of the prior is often restricted to *conjugate* distributions with respect to the measurement model, in which case the posterior distribution that results from the update has the same type as the prior. When using non-conjugate priors, one can employ existing codes for computing the posterior. See, for instance, the BUGS (Bayesian inference Using Gibbs Sampling) project at <http://www.mrc-bsu.cam.ac.uk/bugs/>.

For the problem considered in this paper, in which the number of failures in a given number of tests is a binomial random variable, it is convenient to model the prior as a beta distribution with parameters α and β , written $Beta(\alpha, \beta)$ and having density shown in Equation (3) where $B(\alpha, \beta)$ is the mathematical beta function.

$$f(\theta; \alpha, \beta) = \frac{1}{B(\alpha, \beta)} \theta^{\alpha-1} (1-\theta)^{\beta-1} \quad (3)$$

Using this prior and a binomial likelihood function, the posterior will also be a beta distribution. Specifically, when one starts with the prior distribution $Beta(\alpha_0, \beta_0)$ and observes m failures out of n trials, the posterior distribution is $Beta(\alpha_0 + m, \beta_0 + n - m)$. Consequently, the update can be done analytically by simple addition and subtraction, an enormous improvement in efficiency over the general case.

The problem considered here is an estimation problem: the designer wants to know the true value of the probability of failure and to compare this to the critical value. One approach is to consider a point estimate of the parameter that is based on the posterior distribution. Common point estimates are the mean and median of the posterior and the maximum likelihood estimate, i.e., the value of θ that maximizes $\pi(\theta | Y = y)$.

Alternatively, one can express probabilities about the parameter. This is a key distinction from the classical approach. For example, the designer can calculate the probability that $\theta \leq \theta_{crit}$, as given in Equation (4).

$$P[\theta \leq \theta_{crit}] = \int_0^{\theta_{crit}} \pi(\theta | y) \cdot d\theta \quad (4)$$

If this probability is large enough, the designer can be comfortable that the failure probability is acceptably small.

Alternatively, one can use intervals to estimate the parameter. Using the posterior, any *credible interval* (or more generally, any *credible set*) can be found such that the probability that the true value is an element of the set is equal to a desired value. It is natural to calculate the $100(1-\alpha)$ percentile of the posterior distribution, denoted $\theta_{1-\alpha}$ and form the $100(1-\alpha)\%$ credible set $\{\theta \mid 0 \leq \theta \leq \theta_{(1-\alpha)}\}$. The decision criteria then becomes determining whether this interval contains θ_{crit} . If it does, then the probability that $\theta \leq \theta_{crit}$ is less than $100(1-\alpha)\%$ and the design is unacceptable at that level.

3.3. Robust Bayesian approach

The third approach that we will consider uses an imprecise beta model for updating beliefs. This approach is compatible with the robust Bayesian approach and the theory of imprecise probabilities.

The core idea of robust Bayesian methods is to perform a “what-if” analysis by changing the prior [3, 15-17]. The analyst considers several reasonable prior distributions and performs the update on each to get a set of posterior distributions. If there is no significant change in the conclusion across this set of posteriors, then the conclusion is *robust* to the selection of the prior.

The theory of imprecise probabilities, formalized by Walley [18], has previously been considered in design decisions and set-based design [19, 20], but not in the context of updating parameter estimates in design settings. The theory uses the same fundamental notion of rationality as de Finetti’s work [5, 21]. However, the theory allows a range of indeterminacy—prices at which a decision-maker will not enter a gamble as either a buyer or a seller. These in turn correspond to ranges of probabilities. For the problem of updating beliefs, imprecise probability theory essentially allows for prior and posterior beliefs to be expressed as sets of density functions. Although the motivation is different from the robust Bayesian perspective, the computational approach is quite similar as it involves using Bayes’s Theorem to find a set of posterior distributions from a set of prior distributions. For the problem considered in this paper, both approaches can be implemented using the imprecise beta model.

As explained previously, the beta distribution represents a convenient prior for the example problem in this paper under the precise Bayesian approach. For the robust updating approach, it is convenient to use the imprecise beta model, described by Walley [18, 22]. To introduce the imprecise beta model, it is convenient to first re-parameterize the beta, such that the density of $beta(s, t)$ is as given in Equation (5).

$$\pi_{s,t}(\theta) \propto \theta^{st-1}(1-\theta)^{s(1-t)-1} \quad (5)$$

Compared to the standard parameterization of $Beta(\alpha, \beta)$, this means that $\alpha = s \cdot t$ and $\beta = s \cdot (1-t)$ or equivalently that

$s = \alpha + \beta$ and $t = \alpha / (\alpha + \beta)$. The convenience of this parameterization is that t is the mean of the distribution, which has an easily grasped meaning for both the prior assessment and the posterior analysis. The model is updated as follows: if the prior parameters are s_0 and t_0 , then the posterior parameters after n trials with m failures are given by $s_n = s_0 + n$ and $t_n = (s_0 t_0 + m) / (s_0 + n)$. Since $s_n = s_0 + n$, s_0 can be interpreted to be a virtual sample size of the prior information; it captures how much weight to place on the prior compared to the observed data. Selecting this parameter therefore depends on the available information and will be discussed with the different information scenarios. Following Walley [18], the parameters can be imprecise. That is, the priors are the set of beta distributions with $\alpha_0 = s_0 t_0$ and $\beta_0 = s_0 (1-t_0)$ such that $\underline{t}_0 \leq t_0 \leq \bar{t}_0$ and $\underline{s}_0 \leq s_0 \leq \bar{s}_0$. The update equations can be expressed as follows (the min and max are taken over s_0 subject to $\underline{s}_0 \leq s_0 \leq \bar{s}_0$):

$$\begin{aligned} \underline{t}_n &= \min \{ (s_0 \underline{t}_0 + m) / (s_0 + n) \} \\ \bar{t}_n &= \max \{ (s_0 \bar{t}_0 + m) / (s_0 + n) \} \\ \underline{s}_n &= \underline{s}_0 + n \\ \bar{s}_n &= \bar{s}_0 + n \end{aligned}$$

The posteriors are the set of beta distributions with $\alpha_n = s_n t_n$ and $\beta_n = s_n (1-t_n)$ such that $\underline{t}_n \leq t_n \leq \bar{t}_n$ and $\underline{s}_n \leq s_n \leq \bar{s}_n$.

The use of the imprecise beta model obviously affects the decision process in several ways, but most revealing is how it affects point estimation. Using a precise prior input, one gets a precise posterior output and a single point estimate. In the imprecise model, there are multiple posterior distributions, and consequently a range of point estimates. In particular, the interval $[\underline{t}_n, \bar{t}_n]$ is an interval estimate of θ . One approach for deciding whether or not $\theta \leq \theta_{crit}$ is to compare this interval to θ_{crit} . If $\bar{t}_n \leq \theta_{crit}$, then the designer knows that a decision that $\theta \leq \theta_{crit}$ which is based on the mean estimate of θ is robust across the range of prior assumptions. However, as with the precise case, a designer may wish to calculate the probability that $\theta \leq \theta_{crit}$ and compare this with some critical value. One can find upper and lower probabilities for this based on the set of posterior distributions. Although this varies across the set of posterior distributions, the minimal and maximal values occur at the extreme points of $[\underline{t}_n, \bar{t}_n] \times [\underline{s}_n, \bar{s}_n]$. Finally, one can determine a set of credible intervals from these posterior distributions.

4. DISCUSSION OF APPROACHES

The debate between Bayesians and non-Bayesians has been long, contentious, and without universal resolution [23-25]. However, there are certain properties of Bayesian analysis that are often attractive in practice. First, Bayesian analyses obey the *likelihood principle*, which basically states that all relevant information is contained in the likelihood function, and that

two likelihoods contain the same information if they are proportional to each other (for example, see [3]). This implies that inferences should be made based on the data that was actually observed, not based on hypothetical data that may have been observed but was not. Lindley and Phillips [26] provide an example.

Second, a Bayesian analysis can take advantage of existing information. Thus, in some cases the Bayesian approach can incorporate more information than a classical approach. Of course, the impact of prior beliefs decreases as more data is obtained. Third, Bayesian approaches enable analysts to make direct probability statements about hypotheses and model parameters, statements that cannot be made in classical approaches.

However, the classical and Bayesian approaches are compatible under certain scenarios. For example, Hogg and Tanis [1] point out that if the Bayesian prior is a constant non-informative one and the analyst uses the mode of the posterior distribution as the point estimate, then this is exactly the same as the maximum likelihood estimate in the classical approach.

Selecting the prior distribution is an important part of the Bayesian approach. Although a particular conjugate prior may be useful for computational convenience, it may not represent the prior information perfectly. When there is little prior information available, additional challenges occur. Using a single precise prior in this case seems to conceal the lack of information, as shown later in this paper.

There are at least two arguments as to why a designer should not consider a single prior distribution—one practical, and one philosophical [27]. First, because eliciting and assessing an individual's beliefs is a resource intensive process, it will often be impractical to fully characterize them [18, 28-30]. Consequently, only a partial (imprecise) characterization will be available. This is the view that advocates of the robust Bayesian or Bayesian sensitivity analysis approaches hold [3, 15-17].

The second argument claims that true beliefs need not be precise. A person may be unwilling to commit to either side of a gamble when he has no information about it. It is possible that one would be satisfied with some price, but this is not a condition for rationality. This view is the one that advocates of imprecise probabilities hold [18, 22, 27, 31].

The use of upper and lower probabilities has a number of advantages in this situation [22, 31]. Most relevant to the decision being discussed here, the results reflect the amount of information on which the conclusions are drawn; one can distinguish between probabilities based on a large amount of data and those based on ignorance. Moreover, the approach is able to model prior ignorance in a very intuitive way, and it includes common Bayesian models for prior ignorance as special cases.

The robust Bayesian approach can require more effort than the precise Bayesian approach, especially when the imprecise Beta (or imprecise Dirichlet) models do not apply. Also, as multiple sources of uncertainty are propagated, it may be that the resulting intervals are quite large, leading to considerable regions of indecision, as we will see later in this paper.

5. INFORMATION SCENARIO DESCRIPTIONS

The feasibility and desirability of a given approach for updating one's beliefs about an uncertain quantity depend upon the initial conditions (the amount of data currently available and the beliefs that are currently held).

In some applications, the designer may have relevant information about a product before he plans specific testing. For example, if a new product is a derivative of a previous product, certain aspects of its performance may be expected to be similar. In this case, data from experiments with the predecessor product give a rough idea of the performance of the new product, and testing is pursued to confirm and refine these assumptions. In other applications, the new product may be so novel or employed in such a different environment that the designer believes that existing data have little value in predicting the performance of the new product. In this paper, three scenarios are considered:

1. No relevant prior information is available.
2. Substantial, highly-relevant data histories are available.
3. Partially relevant information or a small amount of relevant information is available.

The placement of an analysis task into one of these categories is a subjective judgment of the designer, and his assessment certainly could be incorrect; in this case, it may be valuable for subsequent testing and analysis to indicate this error.

5.1. Scenario 1: no prior information

In this scenario, it is assumed that the designer has concluded that there is no available information that is relevant to the desired reliability analysis. Essentially, the designer needs to construct an estimate of the reliability from scratch. All inferences about the product's reliability will be made using only the data samples received from the planned experiments, specifically $X = \{x_i\}_{i=1}^n$.

5.1.1. Classical sampling approach

From one perspective, a classical approach is perfectly suited to a scenario in which no prior information exists. A classical approach deals exactly with analyzing the probability of the data given a particular truth, and therefore focuses entirely on the observed data, that is, $X = \{x_i\}_{i=1}^n$. Prior information plays no role in the analysis. Consequently, the classical approach is identical in all information scenarios.

5.1.2. Precise Bayesian approach

A key element of the precise Bayesian approach is the choice of a prior distribution for θ . In the case of no prior information, the use of a *non-informative* prior is generally advocated [10-13], as they favor no value of θ over any others. In this example, a natural choice is the uniform [0, 1] distribution (which in this case is also the maximum entropy prior). Conveniently, the uniform [0, 1] distribution is a special case of the beta distribution, expressed by letting $\alpha_0 = 1$ and $\beta_0 = 1$, so the conjugate prior can be used..

5.1.3. Robust Bayesian approach

The robust Bayesian approach begins with a set of priors. When there is a complete lack of prior information, then the appropriate starting point is a *vacuous prior* (so-called because it contains the least information). In the imprecise beta model, this means setting $\underline{t}_0 = 0$ and $\bar{t}_0 = 1$; all that is said is that the probability is somewhere between 0 and 1, which is the least specific statement that is possible. We must also choose a “learning parameter” s . Because it reflects how much “importance” to assign to the prior data and there is no prior data, one could select a small learning parameter. However, too small a parameter may cause the posterior to react too quickly to the data. As described by Walley [18], $s_0 = 2$ has a number of good properties. One can also allow for a range of learning parameters, such as $\underline{s}_0 = 0$ and $\bar{s}_0 = 2$.

Note that $t_0 = 0.5$ and $s_0 = 2$ corresponds to the beta distribution with $\alpha_0 = 1$ and $\beta_0 = 1$. Thus, this set of priors includes the single uniform prior considered in the precise Bayesian approach with no information. This highlights the difference between a uniform (precise) prior and a vacuous prior. A vacuous prior is a set of prior distributions which happens to contain the uniform prior; the vacuous prior is less specific (i.e. more general and generic) than a uniform prior.

5.1.4. Comparing behavior of estimates

The choice of approach naturally will affect the conclusions that are drawn from the test results. The classical approach will

give a confidence interval on the failure probability. The precise Bayesian approach will give a single distribution for θ . The imprecise beta model yields a set of distributions for θ .

To illustrate the differences, consider the following example: the designer wants θ to be less than $\theta_{crit} = 0.05$; the classical approach uses a 95% one-sided confidence interval, the Bayesian prior is the uniform ($\alpha_0 = 1, \beta_0 = 1$), and the robust Bayesian prior is the vacuous prior ($\underline{t}_0 = 0, \bar{t}_0 = 1$) with a learning parameter of $\underline{s}_0 = \bar{s}_0 = 2$.

If the designer conducts $n = 100$ trials, Table 1 describes the conclusions for some of the possible outcomes. As the number of observed failures increases, the confidence interval on θ increases, and the posterior beta distributions shift, which increases the 95% credible intervals for both the precise Bayesian approach and the imprecise beta model. In this example, the 95% credible intervals for the precise Bayesian approach are significantly larger than the confidence intervals. The smallest credible intervals from the imprecise beta model are nearly the same size as the confidence intervals.

Alternatively, consider how the conclusions would change based on the number of trials, for a fixed relative frequency of observed failures to trials. For instance, if $m/n = 0.04$, then the conclusions change based on the number of trials as shown in Tables 2 and 3. The confidence intervals decrease in width as the number of trials increases because the large number of trials reduces the sample variance. Similarly, the credible intervals decrease in width because the variance of the posterior distribution decreases. This also increases the probability that $\theta \leq \theta_{crit}$. The probability from the precise Bayesian approach is in the middle of the range of upper and lower probabilities that result from the imprecise beta model. Note also that this range significantly decreases as the number of trials increases.

Tables 1 and 2 illustrate how the results of the robust Bayesian approach reflect the amount of information available, while the precise Bayesian approach conceals it. If we compare the results for 6 failures out of 100 trials from Table 1 and the results for 2 failures out of 50 trials from Table 2, we see that the credible intervals for the precise Bayesian approach are nearly the same, despite the difference in observed rate (6%

Table 1. Conclusions based on 100 trials

# of failures	Classical approach	Precise Bayesian	Robust Bayesian
m	95% confidence interval	95% credible interval	Range of 95% credible intervals
0	[0, 0.000]	[0, 0.029]	[0, 0] to [0, 0.046]
1	[0, 0.026]	[0, 0.046]	[0, 0.029] to [0, 0.061]
2	[0, 0.043]	[0, 0.061]	[0, 0.046] to [0, 0.075]
3	[0, 0.058]	[0, 0.075]	[0, 0.061] to [0, 0.088]
4	[0, 0.072]	[0, 0.088]	[0, 0.075] to [0, 0.101]
5	[0, 0.086]	[0, 0.101]	[0, 0.088] to [0, 0.114]
6	[0, 0.099]	[0, 0.114]	[0, 0.101] to [0, 0.126]
7	[0, 0.112]	[0, 0.126]	[0, 0.114] to [0, 0.138]
8	[0, 0.125]	[0, 0.138]	[0, 0.126] to [0, 0.150]
9	[0, 0.137]	[0, 0.150]	[0, 0.138] to [0, 0.162]
10	[0, 0.149]	[0, 0.162]	[0, 0.150] to [0, 0.174]

Table 2. Conclusions with observed failures at 4%.

Number of trials	Classical approach	Precise Bayesian	Robust Bayesian
n	95% confidence interval	95% credible interval	Range of 95% credible intervals
25	[0, 0.104]	[0, 0.170]	[0, 0.109] to [0, 0.223]
50	[0, 0.086]	[0, 0.118]	[0, 0.090] to [0, 0.145]
100	[0, 0.072]	[0, 0.088]	[0, 0.075] to [0, 0.101]
200	[0, 0.063]	[0, 0.071]	[0, 0.064] to [0, 0.077]
500	[0, 0.054]	[0, 0.057]	[0, 0.055] to [0, 0.060]
1000	[0, 0.050]	[0, 0.052]	[0, 0.051] to [0, 0.053]
2000	[0, 0.047]	[0, 0.048]	[0, 0.047] to [0, 0.048]

Table 3. Conclusions with observed failures at 4%.

Number of trials	Precise Bayesian	Robust Bayesian
n	$P\{\theta \leq \theta_{crit}\}$	$[P\{\theta \leq \theta_{crit}\}, \bar{P}\{\theta \leq \theta_{crit}\}]$
25	0.376	[0.139, 0.736]
50	0.473	[0.251, 0.731]
100	0.573	[0.393, 0.749]
200	0.679	[0.552, 0.791]
500	0.824	[0.762, 0.875]
1000	0.920	[0.895, 0.941]
2000	0.980	[0.974, 0.985]

versus 4%) and the different amounts of data. However, the ranges of credible intervals for the robust Bayesian approach are quite different and reflect the underlying amount of information. In the first case (where there is more information), the difference of the upper bounds is 0.025. In the second case (where there is less information), the difference of the upper bounds is 0.055.

5.2. Scenario 2: substantial prior information

In this scenario, there exists substantial information that the designer believes is relevant to the desired reliability analysis. This information may be based on previous tests, the performance of similar products, or other sources. The designer is considering testing as a way of verifying this information.

5.2.1. Classical approach

The classical approach does not include any prior information. As mentioned before, it focuses entirely on the observed data. Thus, the approach does not change for this scenario.

5.2.2. Precise Bayesian approach

A key to the performance of this approach is again the choice of a prior distribution for θ . The designer should choose a prior that represents the substantial, relevant information. It will be convenient once again to use a beta distribution. Larger values of α_0 and β_0 yield a prior distribution with less variance.

Consider two different priors for the following scenario: the designer wants θ to be less than $\theta_{crit} = 0.05$. Suppose that the true $\theta = 0.04$. The first (or “good”) prior is $\alpha_0 = 4$ and $\beta_0 = 96$. The second (or “bad”) prior is $\alpha_0 = 6$ and $\beta_0 = 94$ (it is a “bad” prior because it gives a mean-point estimate for θ that is too high, which means that the prior information is actually wrong or irrelevant). Suppose the number of observed failures is $m = 0.04n$. Then Figure 1 shows how the probability that $\theta \leq \theta_{crit}$ increases as more trials are conducted. As expected, the use of a “bad” precise prior can adversely affect the analysis of the acceptability of θ . Additional cases are considered subsequently in the paper and summarized in Tables 4 and 5.

5.2.3. Robust Bayesian approach

The robust Bayesian approach begins with a choice of priors. Given substantial prior information, the designer can choose \bar{t}_0 and \underline{t}_0 to be upper and lower bounds on θ . If the designer is unsure of how much “importance” to assign to the prior data, then the designer could specify upper and lower learning parameters \bar{s}_0 and \underline{s}_0 . In the case of substantial information, we would expect \bar{t}_0 and \underline{t}_0 to be close to each other and \bar{s}_0 and \underline{s}_0 to be large.

Naturally, the quality of the prior information affects the conclusions. To illustrate the differences, consider again the previous example: the designer wants θ to be less than $\theta_{crit} = 0.05$. Given that the true $\theta = 0.04$, we will consider two different sets of prior information: the first set has the following ranges of parameters: $[\underline{t}_0, \bar{t}_0] = [0.035, 0.045]$ (which includes the truth and is considered “good”) and $[\underline{s}_0, \bar{s}_0] = [80, 120]$. The second set has the following ranges of parameters: $[\underline{t}_0, \bar{t}_0] = [0.055, 0.065]$ (which overestimates the truth and is considered “bad”) and $[\underline{s}_0, \bar{s}_0] = [80, 120]$.

Each set of posterior distributions (the one based on good priors and the one based on incorrect priors) gives a range of values for the probability that θ is acceptable. Figure 2 shows these ranges [where the spread across the priors of each type are shown using the error bars] for each set as the number of trials increases. In all cases, the number of observed failures is the expected value; that is, $m = 0.04 n$. As the number of trials increases, the posteriors based on the good set of priors converge, and their variance decreases, significantly increasing the probability that θ is acceptable.

The posteriors based on the bad set of priors converge more slowly, which keeps the probability that θ is acceptable too low. Interestingly, if we consider the results from Table 3, we see that the posteriors based on the vacuous priors lead to higher probabilities that θ is acceptable than the posteriors based on the bad set of priors. In the vacuous case, the new information needs to overcome a lack of knowledge; in the other case, the

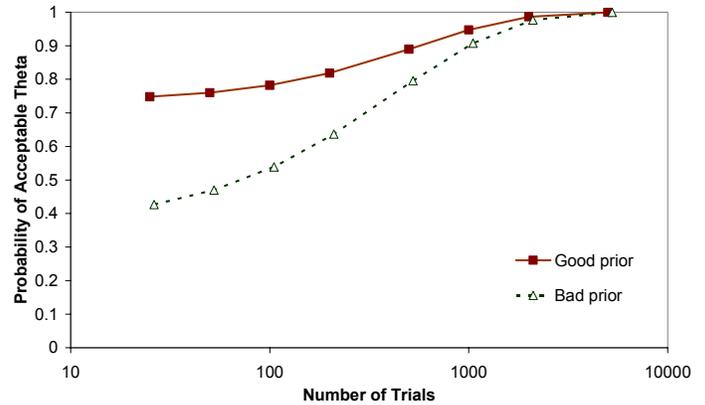


Figure 1. The probability that $\theta \leq \theta_{crit}$ for a good prior and a bad prior using the precise Bayesian approach.

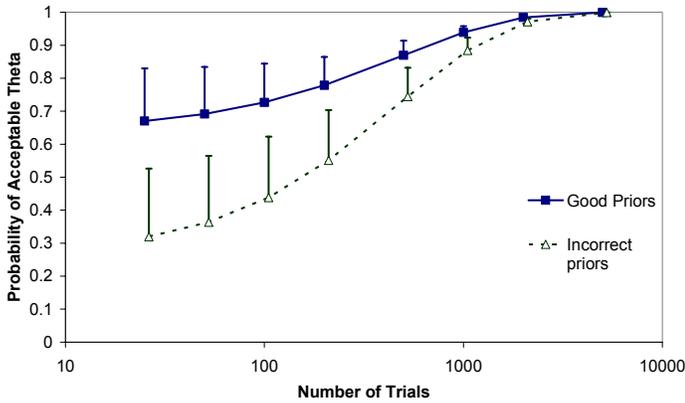


Figure 2. Range of probability of acceptable θ for different sets of priors in the robust Bayesian approach.

new information needs to overcome a large amount of incorrect knowledge. Thus it is important to not overstate the available information.

5.3. Scenario 3: Partial prior information

In this scenario, there exists some information that the designer believes is relevant to the desired reliability analysis. This partial information may be based on some previous tests, the performance of similar products, or other sources. The designer is considering testing in order to augment this partial information.

5.3.1. Classical approach

The classical approach does not include any prior information, so the approach is unchanged from the previous scenarios.

5.3.2. Precise Bayesian approach

It will be convenient once again to use a beta distribution. It is natural to choose the parameters α_0 and β_0 such that the mean $\alpha_0 / (\alpha_0 + \beta_0)$ is close to where the designer believes that the true value of θ may be. In this scenario, the designer should choose small absolute values to reflect the inadequate amount of information. Choosing $\alpha_0 \geq 1$ ensures that the density function has a bell-shaped curve. Lindley and Phillips [26] discuss this process and give some relevant (though non-engineering) examples.

For illustrating these approaches, we choose the priors shown in Table 4 for the case when the true $\theta = 0.06$ and the priors shown in Table 5 for the case when $\theta = 0.04$. We chose these priors so that they have a mean either near the true θ (for the good prior case) or farther away from the true θ (for the bad case). However, compared to the priors

chosen for the significant information scenario, these priors have smaller parameter values and larger variance. Because of the smaller initial parameter values, the posteriors react more quickly to new data, thus de-emphasizing the importance of the prior.

5.3.3. Robust Bayesian approach

As discussed in Scenario 2, this approach starts with a set of prior distributions for θ and creates a set of posterior distributions for θ . The chosen priors are summarized in Table 4 and Table 5. Since these priors reflect less information than in Scenario 2, the range of probabilities is wider, and the values of s_0 are smaller, indicating a smaller “pseudo” sample size for the prior. In general, the range for t_0 will be larger since there is more uncertainty in the estimates when there is less information.

5.4. Summary of scenarios

The priors selected for the precise Bayesian and the robust Bayesian approaches are summarized in Table 4 for the case when $\theta = 0.06$ and in Table 5 for the case when $\theta = 0.04$. The classical approach does not use any prior information, so it is unnecessary to describe it. It is important to note that although some priors are labeled “good” or “bad”, not all good (or all bad) priors are created equal. The term “good” merely refers to a prior with an estimated mean for θ that is relatively “close” to the truth compared to the “bad” prior. It is impossible to define an equivalent “good” for a substantial information case and a partial information case, because the quality of the prior

Table 4. Bayesian priors for $\theta = 0.06$ case

Approach	No prior info.	Partial prior info. (good)	Partial prior info. (bad)	Significant prior info. (good)	Significant prior info. (bad)
Precise Bayesian	$\alpha = 1$ $\beta = 1$	$\alpha = 1.7$ $\beta = 23$	$\alpha = 1.0$ $\beta = 26$	$\alpha = 6$ $\beta = 94$	$\alpha = 4$ $\beta = 96$
Robust Bayesian	$t_0 = 0$ $\bar{t}_0 = 1$ $\underline{s}_0 = 0$ $\bar{s}_0 = 2$	$t_0 = 0.05$ $\bar{t}_0 = 0.10$ $\underline{s}_0 = 20$ $\bar{s}_0 = 30$	$t_0 = 0.00$ $\bar{t}_0 = 0.05$ $\underline{s}_0 = 20$ $\bar{s}_0 = 30$	$t_0 = 0.055$ $\bar{t}_0 = 0.065$ $\underline{s}_0 = 80$ $\bar{s}_0 = 120$	$t_0 = 0.035$ $\bar{t}_0 = 0.045$ $\underline{s}_0 = 80$ $\bar{s}_0 = 120$

Table 5. Bayesian priors for $\theta = 0.04$ case

Approach	No prior info.	Partial prior info. (good)	Partial prior info. (bad)	Significant prior info. (good)	Significant prior info. (bad)
Precise Bayesian	$\alpha = 1$ $\beta = 1$	$\alpha = 1.1$ $\beta = 25$	$\alpha = 1.8$ $\beta = 22$	$\alpha = 4$ $\beta = 96$	$\alpha = 6$ $\beta = 94$
Robust Bayesian	$t_0 = 0$ $\bar{t}_0 = 1$ $\underline{s}_0 = 0$ $\bar{s}_0 = 2$	$t_0 = 0.00$ $\bar{t}_0 = 0.05$ $\underline{s}_0 = 20$ $\bar{s}_0 = 30$	$t_0 = 0.05$ $\bar{t}_0 = 0.10$ $\underline{s}_0 = 20$ $\bar{s}_0 = 30$	$t_0 = 0.035$ $\bar{t}_0 = 0.045$ $\underline{s}_0 = 80$ $\bar{s}_0 = 120$	$t_0 = 0.055$ $\bar{t}_0 = 0.065$ $\underline{s}_0 = 80$ $\bar{s}_0 = 120$

is a combination of the accuracy of the estimate (where accuracy refers to how on target an estimate is) and the amount of information on which the estimate is based.

6. COMPARISONS ACROSS SCENARIOS

In this section we will compare the approaches in a variety of ways, focusing first on the precise and robust Bayesian approaches. We will consider both the quality of the prior distribution (is it close to θ or not) and the magnitude of the prior (does it represent no information, partial information, or substantial information).

We will first consider the case in which the number of observed failures equals the expected number. We will then consider the randomness of the outcomes. Specifically, given policies for deciding when θ is acceptably small, we explore how the probability that the designer will make a correct decision depends upon the quality and magnitude of the priors and the number of trials conducted.

6.1. Probability of an acceptable design

The probability that $\theta \leq \theta_{crit}$ is a key output of the precise Bayesian approach. For simplicity, we first assume that the number of observed failures equals the expected number; that is, $m = \theta n$, with $\theta = 0.04$. Figure 3 shows the change in this probability for five priors: the non-informative prior (from Scenario 1), a correct prior based on substantial information (from Scenario 2), an incorrect prior based on substantial information (from Scenario 2), a correct prior based on partial information (from Scenario 3), and an incorrect prior based on partial information (from Scenario 3). The details of the priors are listed in Table 5.

The plot reveals that, as expected, more good information leads to a better assessment. However, it also reveals that having a small amount of bad prior information is useful (compared to the non-informative prior), but a lot of bad information can be

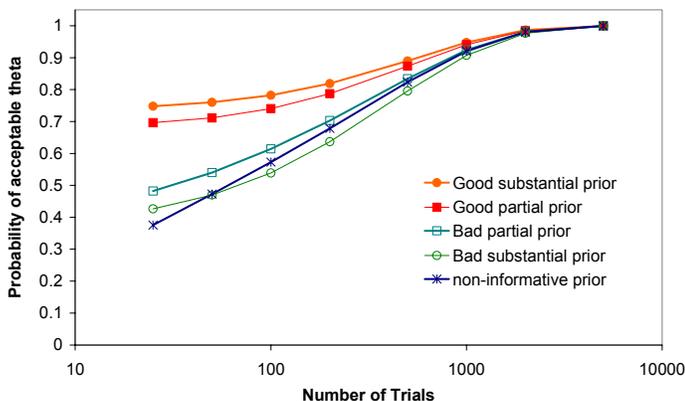


Figure 3. For the precise Bayesian approach, the posterior probability of an acceptable θ depends upon the quantity and quality of the prior information.

detrimental. A “bad” estimate that is near the truth but has little weight essentially provides a good starting point for fast updates based on the observed data, especially compared to the uniform prior, which essentially estimates that θ is just as likely to be near 0.96 (or any other value) as near the true value of 0.04. A bad estimate that is given a lot of weight will not allow the posterior adapt quickly enough to the observed data.

6.2. Point Estimates for θ

The imprecise beta model can be used to determine upper and lower bounds on θ . When $\theta = 0.04$, Figure 4 shows these bounds over a range of trial sizes (assuming that the number of observed failures equals the expected number; that is, $m = \theta n$) for five priors: the non-informative prior (from Scenario 1). The priors are: a correct prior based on substantial information (from Scenario 2), an incorrect prior based on substantial information (from Scenario 2), a correct prior based on partial information (from Scenario 3), and an incorrect prior based on partial information (from Scenario 3). To improve the readability of the chart somewhat, the lower bounds are shown with lines, the upper bounds are shown using the error bars, and the abscissa for each point for an informative prior is shifted to the right or left of the point for the non-informative prior. Details of the priors are located in Table 5.

In all cases, the lower and upper bounds converge as the number of trials increases. Additionally, it should be noted that initially the bad substantial prior result is a subset of the bad partial prior result, which in turn is a subset of the vacuous prior result. As n increases, the vacuous prior and partial prior results adapt faster to the observed data and converge more quickly to the truth than the substantial prior case does.

6.3. Accepting the Design

The above discussion is based upon the expected results, in which the number of observed failures is the expected number for any number of trials (that is, $m = \theta n$). However, the observed results are random. Thus, if the designer uses a fixed policy for judging that θ is acceptably small (and therefore accepting the design), the outcome of the designer’s decision is random. For a given policy, we can look at the range of possible outcomes and calculate the probability of the resulting decision. For convenience, let “accept” mean that the designer decides that the design is acceptable ($\theta \leq \theta_{crit}$). Let “reject” mean that the designer decides that the design is not acceptable ($\theta > \theta_{crit}$). Let “no decision” mean that the designer does not decide, but rather concludes that additional information is needed.

6.3.1. Classical Approach Results

In the classical sampling approach, suppose that the designer will accept the design if the upper bound of the 95% one-sided confidence interval on θ is less than or equal to θ_{crit} .

Otherwise, he rejects the design. Figure 5 shows the decision probabilities for each value of n in two cases: in the first, $\theta = 0.06$ (and is thus larger than θ_{crit}); in the second, $\theta = 0.04$ (and is thus smaller than θ_{crit}).

There are several interesting behaviors apparent in Figure 5. First, note that the classical method performs better when $\theta > \theta_{crit}$. This occurs because the classical approach, using a 95% one-sided confidence interval, is (in a sense) biased towards concluding $\theta > \theta_{crit}$. In a hypothesis testing framework, this is equivalent to formulating $\theta > \theta_{crit}$ as the null hypothesis and then requiring a substantial amount of information to reject it in favor of the alternative hypothesis that $\theta \leq \theta_{crit}$.

A second point worth noting is the nature of the curves in Figure 5. They are not smooth, as one might expect, but are instead almost saw-tooth in nature. This leads to non-intuitive results. For example, in the case with $\theta = 0.04$, the probability of a correct decision is larger when $n = 54$ than when $n = 278$. How can more information lead to worse performance? The answer is due to the discrete nature of the binomial and the approximation made in Equation (2). The actual coverage probability of the approximate 95% interval fluctuates as n varies, an effect that is accentuated for small θ [2]. Similar patterns are seen for the other approaches and metrics, too, again due to the discrete nature of the underlying problem.

6.3.2. Precise Bayesian Results

If the designer uses the precise Bayesian approach, the designer will accept the design (meaning he or she concludes that $\theta \leq \theta_{crit}$) if the posterior probability that $\theta \leq \theta_{crit}$ is at least 0.8. Otherwise, he will reject the design (concluding that $\theta \geq \theta_{crit}$). Figure 6 shows the decision probabilities for each value of n in two cases: in the first, $\theta = 0.06$ (and is thus larger than θ_{crit}); in the second, $\theta = 0.04$ (and is thus smaller than θ_{crit}).

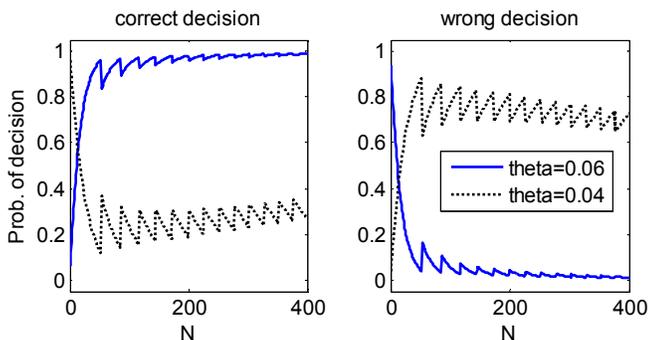


Figure 5. Decision probabilities for the classical approach.

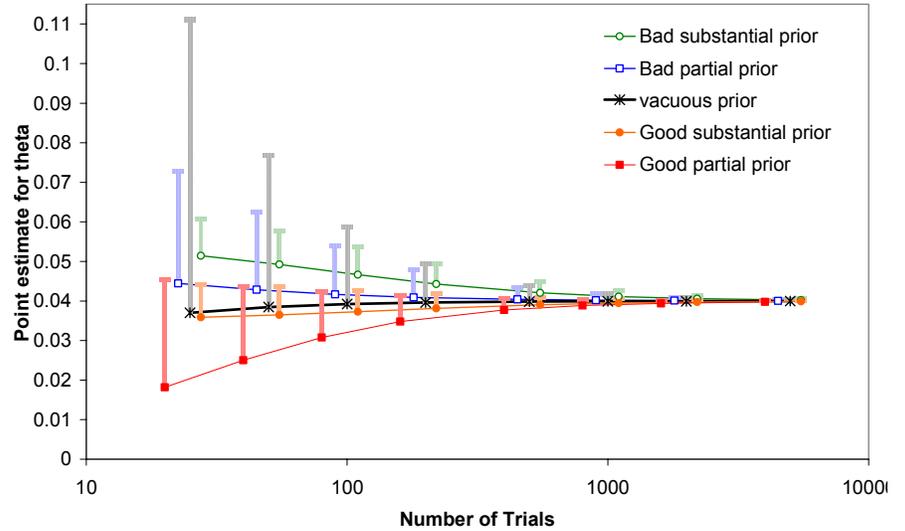


Figure 4. For the robust Bayesian approach, the range of point estimates for θ for different quantities and qualities of prior information.

We begin by examining the upper left plot, which shows the probability of making a correct decision when $\theta = 0.06$. This probability is highest when the designer starts with a good prior that is based on substantial information. Starting with a good prior that is based on partial information and starting with non-informative prior yield similar results. For small n , the probability of a correct decision is lowest when the designer starts with a bad prior that is based on partial information. For larger n , the probability of a correct decision is lowest when the designer starts with a bad prior that is based on substantial information. When n is small, the partial, bad prior is more easily influenced and therefore requires less information to push the probabilities over the decision threshold of 0.80. As n increases, this responsiveness causes the posterior to move

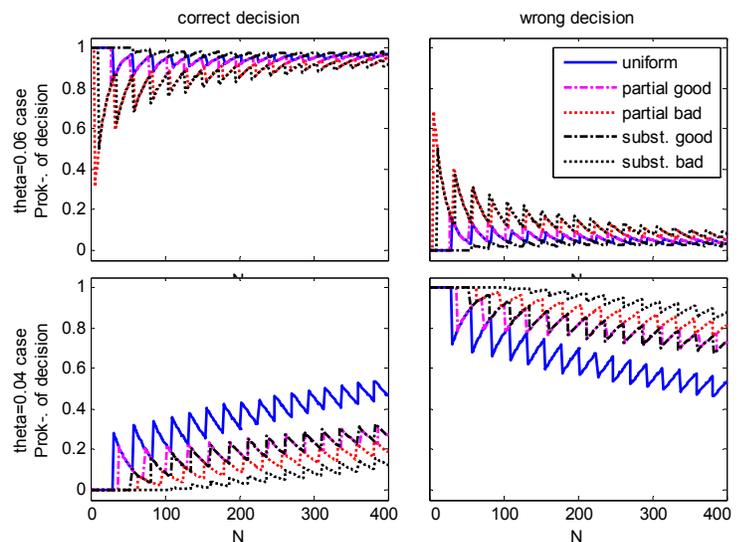


Figure 6. Decision probabilities for the precise Bayesian approach.

away from the bad prior faster.

The lower left plot shows the probability of making a correct decision when $\theta = 0.04$. In general, the results are as expected, since having substantial, good prior information leads to the best performance, followed by having partial good information. Having substantial, bad prior information leads to the worst performance. Interestingly, the partial, bad prior scenario performs slightly better than the no information, uniform $[0, 1]$ prior. As defined, the partial, bad prior was “bad” in that the mean of the prior was 0.076, which is much greater than the true 0.040. Despite the error in the mean, this prior is still in the vicinity of the truth, especially compared to the mean of the uniform, which is at 0.500.

6.3.3. Robust Bayesian Results

In the robust Bayesian (imprecise beta) approach, we assumed that the designer will accept the design if (over the set of posteriors) the minimum probability that $\theta \leq \theta_{crit}$ is at least 0.8. The designer will reject the design if (over the set of posteriors) the maximum probability that $\theta \leq \theta_{crit}$ is less than 0.8. Otherwise, he makes no decision, essentially concluding that there is still not sufficient information to warrant a firm decision. Figure 7 shows the decision probabilities for each value of n in two cases: in the first, $\theta = 0.06$ (and is thus larger than θ_{crit} and unacceptable); in the second, $\theta = 0.04$ (and is thus smaller than θ_{crit} and acceptable).

We begin by considering the bottom row of charts, which are for the $\theta = 0.04$ scenario. The probability of making a correct decision is lowest when the designer starts with a bad prior that is based on substantial information. Because there is less prior information to overcome, the probability of making a correct decision is higher when starting with a bad prior that is based on partial information. The probability of making a correct decision is highest when the designer starts with a good prior that is based on substantial information or a good prior that is based on partial information.

The partial, bad prior and the vacuous prior lead to similar performance. Starting with a small amount of bad information and starting with no information and are similar, in that in both cases it requires a lot of additional information in order to conclude that the design is acceptable. With the partial, bad prior, one needs enough additional information to counter the small amount of bad prior information that, while bad, is still in the vicinity of the truth. In the no information scenario, one requires enough additional information to lower the upper

bound of the point estimate from 1.0 to something less than 0.04.

The performance of the partial, good prior actually exceeds that of the substantial, good prior for small n , with the performance shifting to favor the substantial, good prior as n approaches 400. This is related to the quality and responsiveness of the priors. With the priors listed Table 5, the initial bounds on the probability that $\theta \leq \theta_{crit}$ are $[0.60, 1.00]$ for the partial, good prior case and $[0.64, 0.83]$ for the substantial, good prior case, so both actually lead to no decision.

Suppose $n = 20$ trials are performed and $m = 0$ failures are observed. This outcome has a probability of 0.44 when $\theta = 0.04$. For the partial, good prior, the posterior bounds on the probability that $\theta \leq \theta_{crit}$ are $[0.82, 1.00]$, so the conclusion is that the design is acceptable. For the substantial, good prior, the posterior bounds are $[0.76, 0.93]$ and no conclusion can be reached. Thus, it is evident that as expected, the partial prior is more sensitive to small sample sizes. Although this is advantageous in the $\theta = 0.04$ case, it is detrimental in the $\theta = 0.06$ case, as discussed next.

The probability of making a correct decision when $\theta = 0.06$ is higher in general than it is when $\theta = 0.04$. The probability is highest when the designer starts with a good prior that is based on substantial information. Starting with a good prior that is based on partial information and starting with a vacuous prior perform similarly to each other. The probability of a correct decision is lowest when the designer starts with a bad prior that is based on substantial information or a bad prior that is based on partial information.

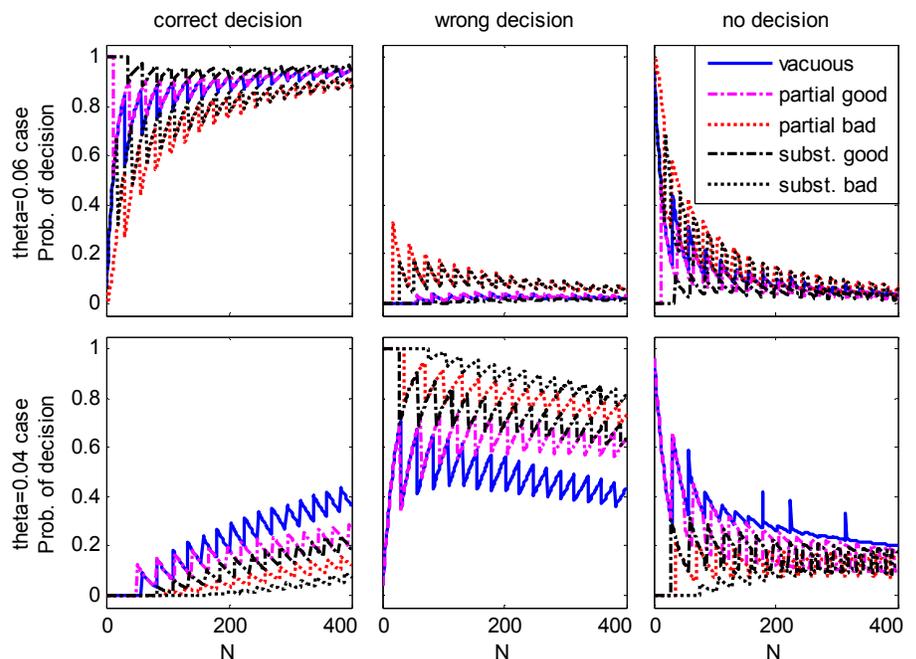


Figure 7. Decision probabilities for the robust Bayesian approach.

It is interesting to note that having partial, bad prior information is actually worse than having substantial, bad prior information over the region of n shown. To some extent, this is an artifact of “bad” being defined ambiguously. For example, the partial, bad prior gives the probability that $\theta \leq \theta_{\text{crit}}$ as the interval $[0.18, 0.62]$, whereas the substantial, bad prior gives the interval $[0.27, 0.48]$. The upper bound for the partial information case (0.62), is much closer to the threshold (0.80) for making a bad decision than the upper bound for the substantial information case (0.48). Since the partial information prior is also more sensitive to the data, more of the possible random results of an experiment will lead to accepting the design, which is the wrong decision. A more substantial, yet also incorrect, prior will be more robust to the variation in small samples.

A major difference between these results and those for the precise Bayesian approach is the option of making no decision. In the $\theta = 0.06$ scenario, the probability of not reaching a conclusion is highest for the partial, bad prior and the substantial, bad prior. This is desirable; when the prior and data conflict, something is amiss and more data are needed to resolve the conflict.

Recall that the decision policy is to accept the design only when the minimum probability that $\theta \leq \theta_{\text{crit}}$ is at least 0.80. For the partial, bad prior scenario, the set of prior distributions actually leads to no conclusion, and hence the initially high probability of no conclusion. As evidence is gathered that $\theta > \theta_{\text{crit}}$, eventually the prior belief is countered by the evidence, and both the probability of reaching no conclusion and the probability of reaching the wrong decision (i.e. accepted the design when $\theta > \theta_{\text{crit}}$) both decrease (and naturally the probability of the correct decision increases).

When there is substantial, good prior information, there is a much smaller probability of reaching no conclusion, because on average the prior and data agree with each other. In particular, for the substantial, good prior case, there is initially no chance of reaching no conclusion, because the prior always would lead to a correct decision, and the probability of getting of getting a small data set that is so “bad” that it cancels out the influence of the substantial, good prior is very small.

In the $\theta = 0.04$ scenario, the probability of not reaching a conclusion is highest for the vacuous prior. This reflects that when there is very little information, acceptance of the design requires that this information be strongly in favor of concluding that the design is acceptable, and this has a very low probability of occurring. It may appear that there is an inconsistency compared to the $\theta = 0.06$ scenario in that the “bad” priors now have the lowest probabilities of reaching no decision. Additionally, the substantial, bad prior actually has a generally increasing probability of no decision for n between 0 and 200. This results from the complicated tradeoffs between prior

information, observed data, and the “bias” in the decision rule, as follows.

Recall that a design is accepted only if the minimum probability that $\theta \leq \theta_{\text{crit}}$ is at least 0.8. Thus, when there is substantial prior information that the decision is not acceptable, the designer will tend to reject the design. As one acquires data that starts to suggest that the design may be acceptable, the designer moves from concluding it is unacceptable to reaching no conclusion. As the data increasingly suggests that $\theta \leq \theta_{\text{crit}}$, the prior becomes less relevant, and eventually the designer will conclude that the design is acceptable.

More generally, the probability of reaching no conclusion reflects the overall weight of evidence available. As the number of samples increases, the probability of no decision will eventually tend to zero, regardless of the prior. However, the probability of no decision may increase initially, a situation that reflects conflict between the prior and the data. The ability of the robust Bayesian approach to reflect both the overall quality and consistency of the combined prior and data evidence appears to be an advantage over the other methods, at least qualitatively.

6.3.4. Results across methods for the no prior information scenario

All of the results about the decision probabilities depend upon the policies used to make decisions. Therefore, it is also important to consider two different sets of policies and compare their performance in a specific scenario (we will use Scenario 1, the no information scenario). A 95% confidence interval is used with the classical approach as a benchmark for both sets.

The first set uses “generous” policies for the Bayesian approaches. The designer accepts the design if the probability that $\theta \leq \theta_{\text{crit}}$ is at least 0.80. (In the robust Bayesian, the minimum probability must be at least 0.8 to accept, and the maximum probability must be less than 0.8 to reject.) The results of these policies are shown in Figure 8. The second set uses “strict” policies that accept the design if the probability that $\theta \leq \theta_{\text{crit}}$ is at least 0.95. (In the robust Bayesian, the minimum probability must be at least 0.95 to accept, and the maximum probability must be less than 0.95 to reject.) The results of these policies are shown in Figure 9. Both Figure 8 and Figure 9 show how the decision probabilities change for two possible truth scenarios: in the first, $\theta = 0.06$ (and is thus larger than θ_{crit}); in the second, $\theta = 0.04$ (and is thus smaller than θ_{crit}).

Under either set of policies, because it has the “no decision” option, the robust Bayesian approach leads to a lower probability of correct decision and a lower probability of a wrong decision. When $\theta = 0.06$, the generous policies lead to a lower probability of a correct decision (compared to the strict policies) because they are more likely to accept the design (which is wrong in this case). When $\theta = 0.04$, the generous

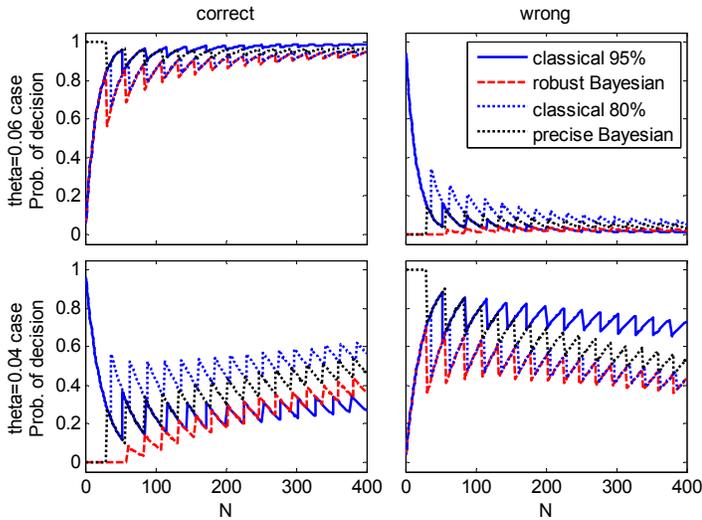


Figure 8. Decision probabilities for the no information scenario using the “generous” policies (0.80 probability test).

policies lead to a higher probability of a correct decision (compared to the strict policies), because accepting the design is the correct decision.

The results for the $\theta = 0.06$ scenario in Figure 9 all indicate that performance is very good. Using the strict decision rule, the decision is biased towards concluding that the design is not acceptable. In the classical approach, the 95% confidence level denotes that, when using the approach many times, the true θ does not fall in the interval at most 5% of the time. However, if the upper bound of the confidence interval is greater than θ_{crit} and the true θ is less than θ_{crit} , then θ is in the confidence interval but the decision will be incorrect (it rejects an acceptable design).

For the Bayesian approaches, the strict policy requires the probability that $\theta \leq \theta_{crit}$ to be greater than 0.95 in order to accept the design. Even if the true θ is less than θ_{crit} , there is a significant probability that, due to the randomness of the outcomes, the posterior probability that $\theta \leq \theta_{crit}$ will be less than 0.95, especially when the number of trials is small. This policy has a stricter acceptance criterion than the 0.80 policy of Figure 8, so the results are even more biased towards concluding that $\theta > \theta_{crit}$. When this is the true conclusion (e.g. when $\theta = 0.06$), then the performance of the method is very good. When this is the false conclusion (e.g. when $\theta = 0.04$), this method performs badly. This tradeoff shares some aspects of the tradeoff between type I and type II errors in traditional hypothesis testing.

In all cases, the robust Bayesian method yields a lower probability of making a wrong decision compared to the precise Bayesian method. It also yields a lower probability of making a wrong decision than the classical approach, except for the $\theta = 0.04$ scenario using the strict decision policy (Figure 9), for which the results are similar. The cost of

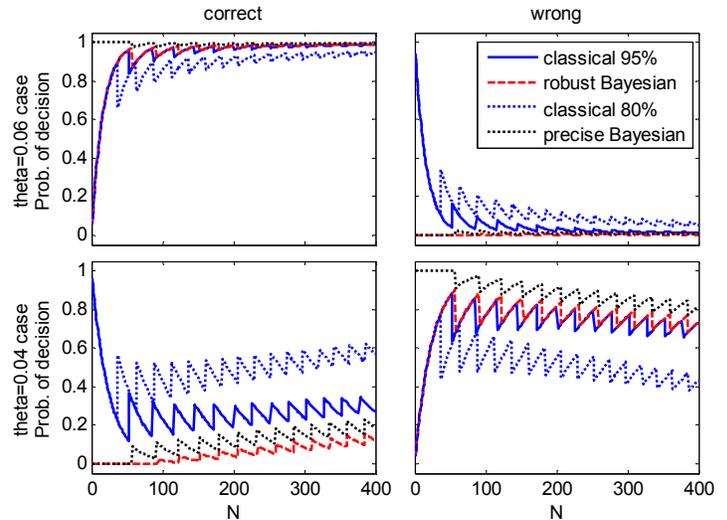


Figure 9. Decision probabilities for the no information scenario using the “strict” policies (0.95 probability test).

avoiding a wrong decision is the inability to make a correct decision for small n . Unlike the precise Bayesian approach, the robust Bayesian approach acknowledges the inability to reach a conclusion when information is scarce or conflicting.

Because the probability of a correct decision with $\theta = 0.06$ is generally high with either set of policies, the designer may seek to increase the probability of a correct decision with $\theta = 0.04$, which would lead to recommending the generous policies.

Finally, as part of the classical approach, one could use a different confidence interval for the decision policy. Increasing the confidence interval (e.g., using a 99% confidence interval) will make accepting the design more difficult. Decreasing the confidence interval (e.g., using a 80% confidence interval) will make accepting the design easier.

7. SUMMARY

A variety of statistical approaches are available for updating beliefs about the expected performance of a product based on test results. This paper has compared the classical sampling theory approach, the precise Bayesian approach, and the robust Bayesian approach. This last approach uses the imprecise beta model, which makes it compatible with the imprecise probabilities approach. The emphasis is on exploring how each method behaves as more information is collected. Direct quantitative comparisons and judgments of absolute superiority are not possible due to the inherently different assumptions and structures underlying each method.

The classical approach does not take into account any information that might be available before testing. This feature has the advantage of being insensitive to the quality of the

information. If one ignores the information, it doesn't matter if it is wrong.

The Bayesian approaches, on the other hand, can use prior information to help the designer make a correct decision with less additional information. However, they are greatly affected by the quality and amount of existing information. In general, it appears that using a prior based on partial information is less "risky": if it is "bad," new information corrects it quickly and good decisions can be made. If it is good, it leads to decisions nearly as good as those obtained using a good prior based on significant information. Interestingly, if it is not too "bad," using a prior based on partial, incorrect information leads to better results than using the non-informative prior because while not perfect, the "bad" information is still more closely related to the truth than a non-informative prior is.

Using a prior based on significant information can lead to poor results if it is a "bad" prior. Overcoming the incorrect information will take a great amount of new information. Of course, if it is "good," then good decisions can be made quickly. Therefore, using partial information may be a good compromise. A prior (or set of priors) based on such information provides an effective head start compared to the non-informative prior (or vacuous priors) with minimal investment in information gathering, but, if the prior information is incorrect, new information can overcome the priors more easily than it would if the prior was based on significant information.

The robust Bayesian approach allows the designer to use ranges to express his beliefs about the value of the parameter (via a range of means) and the importance of the prior data (via a range of learning parameters). Such caution about the prior information reduces the probability of making a bad decision. However, the probability of making a good decision also decreases because outcomes for which the available information is not conclusive explicitly result in no decision being made. Thus, the robust Bayesian method reflects both the quality and quantity of information available, including indirectly any conflict between the prior information and the new data.

Given a failure probability threshold that should not be exceeded, accepting a design when it is truly acceptable is a challenge because the threshold is close to zero and there is a significant probability of test results that lead to a wrong decision. This indicates that the designer should use generous policies for making the acceptance decision. The probability of incorrectly accepting a design whose failure probability is too high remains very low.

These conclusions have been illustrated using a specific example decision. More generally, a designer can use the methods presented here to get insight into the tradeoffs that exist in a specific domain.

ACKNOWLEDGEMENTS

This work was supported in part by the Applied Research Laboratories at the University of Texas at Austin IR&D grant 07-09.

REFERENCES

- [1] Hogg, R. V., and Tanis, E. A., 2001, *Probability and Statistical Inference*, Sixth Ed., Prentice Hall, Upper Saddle River, NJ.
- [2] Brown, L. D., Cai, T. T., and DasGupta, A., 2001, "Interval Estimation for a Binomial Proportion," *Statistical Science*, 16.2, pp. 101-133.
- [3] Berger, J. O., 1985, *Statistical Decision Theory and Bayesian Analysis*, 2nd Ed., Springer.
- [4] Box, G. E. P., and Tiao, G. C., 1973, *Bayesian Inference in Statistical Analysis*, Addison-Wesley, Reading, Massachusetts.
- [5] de Finetti, B., 1974, *Theory of Probability Volume 1: A Critical Introductory Treatment*, Wiley, New York.
- [6] Lindley, D. V., 1982, "Subjectivist View of Decision Making," *European Journal of Operational Research*, 9.3, pp. 213.
- [7] Winkler, R. L., 1996, "Uncertainty in Probabilistic Risk Assessment," *Reliability Engineering & System Safety*, 54.2-3, pp. 127-132.
- [8] Savage, L. J., 1972, *The Foundation of Statistics*, Dover Publications, New York.
- [9] Berger, J. O., 2000, "Bayesian Analysis: A Look at Today and Thoughts of Tomorrow," *Journal of the American Statistical Association*, 95.452, pp. 1269-1276.
- [10] Jeffreys, H., 1961, *Theory of Probability*, 3rd Ed., Oxford University Press, London.
- [11] Zellner, A., 1977, "Maximal Data Information Prior Distributions," *New Methods in the Applications of Bayesian Methods*, Aykac, A. and Brumat, C. eds., North-Holland, Amsterdam.
- [12] Berger, J. O., and Bernardo, J. M., 1992, "On the Development of the Reference Prior Method," *Bayesian Statistics 4*, Bernardo, J. M., Berger, J. O., Dawid, A. P., and Smith, A. F. M. eds., Oxford University Press, Oxford.
- [13] Fougere, P., ed. 1990, *Maximum Entropy and Bayesian Methods*, Kluwer Academic Publishers, Dordrecht.
- [14] Maritz, J. S., and Lewin, T., 1989, *Empirical Bayes Methods*, 2nd Ed., Chapman and Hall, London.
- [15] Berger, J. O., 1993, "An Overview of Robust Bayesian Analysis," Technical Report #93-53c, Purdue University, West Lafayette, IN.
- [16] Insua, D. R., and Ruggeri, F., 2000, *Robust Bayesian Analysis*, Springer, New York.

- [17] Berger, J. O., 1984, "The Robust Bayesian Viewpoint," *Robustness of Bayesian Analysis*, Kadane, J. B. ed., North-Holland, New York, pp. 63-124.
- [18] Walley, P., 1991, *Statistical Reasoning with Imprecise Probabilities*, Chapman and Hall, New York.
- [19] Aughenbaugh, J. M., and Paredis, C. J. J., 2006, "The Value of Using Imprecise Probabilities in Engineering Design," *Journal of Mechanical Design*, 128.4, pp. 969-979.
- [20] Rekuc, S. J., Aughenbaugh, J. M., Bruns, M., and Paredis, C. J. J., 2006, "Eliminating Design Alternatives Based on Imprecise Information," *Society of Automotive Engineering World Congress 2006-01-0272*.
- [21] de Finetti, B., 1980, "Foresight. Its Logical Laws, Its Subjective Sources (Translated)," *Studies in Subjective Probability*, Kyburg, H. E. and Smokler, H. E. eds., E. Krieger Publishing Company.
- [22] Walley, P., Gurrin, L., and Burton, P., 1996, "Analysis of Clinical Data Using Imprecise Prior Probabilities," *The Statistician*, 45.4, pp. 457-485.
- [23] Rothenberg, T. J., 1974, "The Bayesian Approach and Alternatives in Econometrics-II," *Studies in Bayesian Econometrics and Statistics: In Honor of Leonard J. Savage*, Fienberg, S. E. and Zellner, A. eds., North-Holland Publishing Company and American Elsevier Publishing Company, New York, pp. 55-67.
- [24] Zellner, A., 1974, "The Bayesian Approach and Alternatives in Econometrics-I," *Studies in Bayesian Econometrics and Statistics: In Honor of Leonard J. Savage*, Fienberg, S. E. and Zellner, A. eds., North-Holland Publishing Company and American Elsevier Publishing Company, New York, pp. 39-54.
- [25] Clarotti, C. A., 1993, "Decision Making: Drawbacks of the Classical Statistics Approach," *Reliability and Decision Making*, Barlow, R. E., Clarotti, C. A., and Spizzichino, F. eds., Chapman & Hall, New York, pp. 1-12.
- [26] Lindley, D. V., and Phillips, L. D., 1976, "Inference for a Bernoulli Process (a Bayesian View)," *American Statistician*, 30.112-119.
- [27] Pericchi, L. R., and Walley, P., 1991, "Robust Bayesian Credible Intervals and Prior Ignorance," *International Statistical Review / Revue Internationale de Statistique*, 59.1, pp. 1-23.
- [28] Groen, F. J., and Mosleh, A., 2005, "Foundations of Probabilistic Inference with Uncertain Evidence," *International Journal of Approximate Reasoning*, 39.1, pp. 49-83.
- [29] Weber, M., 1987, "Decision Making with Incomplete Information," *European Journal of Operational Research*, 28.1, pp. 44.
- [30] Savage, L. J., 1971, "Elicitation of Personal Probabilities and Expectations," *Journal of the American Statistical Association*, 66, pp. 783-801.
- [31] Walley, P., 1996, "Inferences from Multinomial Data: Learning About a Bag of Marbles," *Journal of the Royal Statistical Society. Series B (Methodological)*, 58.1, pp. 3-57.