

COST MINIMIZATION UNDER SEQUENTIAL TESTING PROCEDURES USING A

BAYESIAN APPROACH

A THESIS

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE

MASTERS OF SCIENCE

BY

LUKAS SNYDER

DALE UMBACH

BALL STATE UNIVERSITY

MUNCIE, INDIANA

MAY 2013

## Stopping Rule

A stopping rule for sequential data is a set of criteria which, when fulfilled, will tell the observer to discontinue observation or testing. We will first consider a probability space  $(\Omega, \mathcal{A}, P)$  where  $\mathcal{A}_1 \subset \mathcal{A}_2 \subset \dots$  is a sequence of sigma-fields. Contained in  $\mathcal{A}$ . The smallest sigma-field containing all  $\mathcal{A}_n, n \geq 1$  will be denoted by  $\mathcal{A}_\infty$ . We will then call the extended random variable  $N$  a stopping time with respect to  $\mathcal{A}_n$ .

$$N = 1, 2, \dots, \infty \text{ and } \{N \leq n\} \in \mathcal{A}_n \text{ for all } n \geq 1$$

$N$  is labeled as proper with respect to  $P$  if  $P(N < \infty) = 1$ . Also, the above definition,

$\{N > n\} = \Omega - \{N \leq n\} \in \mathcal{A}_n$  must hold true for all  $n \geq 1$ . (Ghosh, Mukhopadhyay, & Sen, 1997)

One such stopping rule may tell the observer to stop after ten observations. This type of stopping rule is not dependent on the values of the data, only the number of previous observations. A rule of this type may be appropriate when the observer knows very little about what the data will look like. Another such stopping rule may be to stop when a 95% confidence interval for the parameter being estimated falls within a certain range. This type of stopping rule may be appropriate when the null hypothesis is that the true parameter does not fall within the specified range. This means that when the conditions for the stopping rule are met that the probability of a type two error will be

5%. Unlike our first stopping rule, this rule takes into account data that has been observed.

The type of stopping rule that will be explored in this paper is a rule that considers the observed data as well as cost parameters and prior information. More specifically, both the marginal cost of observation as well as a penalty cost will be considered. Often times in manufacturing processes certain characteristics of an output cannot be observed without destroying the product. In this case it is simple to think of the marginal cost of observation as the marginal cost to produce the unit. The following analysis will work sufficiently well for both destructive and nondestructive testing. The penalty cost refers to the costs associated with a type two error. In a manufacturing setting it is common that the buyer demand a certain percentage of the products to fall within some threshold. If the supplier fails to meet this percentage and ships the batch then they will have to pay the penalty cost. This cost will certainly include the cost of returning a shipment but can also include legal fees and the opportunity cost of lost business.

Many stopping rules have been proposed which incorporate the use of prior information. One of the first such rules to incorporate the marginal cost of testing was developed by Bickel and Yahav (Bickel & Yahav, 2006). The most basic form of the stopping rule instructs the observer to stop testing when

$$\Delta(k(x, c, n)) \geq 0$$

Where  $c$  is the variable cost,  $x$  is the observed data, and  $n$  is the total number of items tested. Kickel and Yahav also acknowledge that there are other stopping rules that are

asymptotically optimal (Bickel & Yahav, 2006). The stopping rule proposed in this paper will follow similar criteria but will incorporate a penalty cost associated with shipping a bad batch.

The proposed stopping rule will instruct the observer to stop observing data when it is believed that by viewing an additional data point the total cost of the testing procedure will increase. This will include the marginal cost of testing (MC), the probability of the batch not meeting the required proportion of acceptable parts (F), and the penalty cost(PC). We will stop when

$$P(F^{t+1}|data^{t+1})PC + MC * (t + 1) > P(F^t|data^t)PC + MC * t$$

Which can also be written as

$$[P(F^t|data^t) - P(F^{t+1}|data^{t+1})]PC < MC$$

Since the data at time t+1 is unknown, this probability will have to be estimated using only data at time t. This will be done by conditioning the posterior at time t+1 on the next observation which will be estimated using the posterior at time t.

The observed data can come from a variety of distributions. Using Bayesian statistics, the prior information can easily be updated with each data point. The next data point can then be estimated and used in the stopping rule. The method used to update the prior distribution will be discussed in the next section.

## Bayesian updating

Although the observed data can be viewed as following any number of continuous or discrete distributions, we will consider data generated by a series of Bernoulli trials. If an observation falls within tolerance then we will consider it a success and if it falls out of tolerance it will be considered a failure. The parameter that is of interest will therefore be the probability of success for an individual observation, which can also be thought of as the overall proportion of parts falling within tolerance. In Bayesian statistics it is common to use the beta distribution to model parameters ranging between 0 and 1. A second reason that this distribution is useful is because when a beta is updated, the posterior is also a beta distribution which is known as conjugacy. The beta distribution has two parameters  $\alpha$  and  $\beta$  which both must be greater than 0. The density for P is

$$f(p|\alpha, \beta) = f(p) = \begin{cases} \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1} (1-p)^{\beta-1} & \text{if } 0 \leq p \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

If during our next N Bernoulli trials we view k successes the conditional probability for K=k given P=p is

$$P(K = k|P = p) = \binom{N}{k} p^k (1-p)^{N-k}$$

The general form of Bayes' theorem tells us that the new density function for P given that K=k is

$$f'(p|k) = \frac{P(K = k|P = p)f(p)}{P(K = k)}$$

Solving for  $f'$  in terms of  $N$ ,  $k$ , and  $p$  we get

$$f'(p|k) = \frac{\binom{N}{k} p^k (1-p)^{N-k} \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1} (1-p)^{\beta-1}}{\int_0^1 \binom{N}{k} p^k (1-p)^{N-k} \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1} (1-p)^{\beta-1} dp}$$

We can quickly see that the binomial coefficients and gamma functions cancel out. We are now left with

$$f'(p|k) = \frac{p^{\alpha+k-1} (1-p)^{\beta+(N-k)-1}}{\int_0^1 p^{\alpha+k-1} (1-p)^{\beta+(N-k)-1} dp}$$

Finally, integrating over the range from (0,1) we obtain

$$f(p|\alpha, \beta) = f(p) = \begin{cases} \frac{\Gamma(\alpha + \beta + N)}{\Gamma(\alpha + k)\Gamma(\beta + N - k)} p^{\alpha+k-1} (1-p)^{\beta+N-k-1} & \text{if } 0 \leq p \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

If the data is being sampled sequentially then the previous equation can be through of as

$$f(p|\alpha_t, \beta_t) = \begin{cases} \frac{\Gamma(\alpha_t + \beta_t + 1)}{\Gamma(\alpha_t + i)\Gamma(\beta_t + 1 - i)} p^{\alpha+i-1} (1-p)^{\beta-i} & \text{if } 0 \leq p \leq 1 \\ 0 & \text{otherwise} \end{cases}$$

Where  $i$  is 1 if a success is viewed and 0 if a failure is viewed. If the current distribution at time  $t$  is  $\text{Beta}(\alpha, \beta)$  and at time  $t+1$  the observer tests a part which is deemed acceptable then the posterior distribution at time  $t+1$  is  $\text{Beta}(\alpha+1, \beta)$ . If the tested part is unacceptable then the posterior distribution at time  $t+1$  is  $\text{Beta}(\alpha, \beta+1)$ . This convenient result will

simplify the simulation process but is only one such pair of distributions which can be utilized for the stopping rule. The possibility of using other distributions will be discussed in the conclusion.

Mukhopadhyay Ghosh outlines the criteria for a stopping rule which uses Bayesian techniques in the text *Sequential Estimation* (Ghosh, 1997). Let  $\delta_j^0(X_1, X_2, \dots, X_j)$  be a bayes rule for the statistical decision problem  $(\theta, a, L)$  based on a fixed sample with respect to a prior  $\pi \in \theta^*, j \geq 0$ . Then for every fixed stopping rule  $\psi$ ,  $r(\pi, (\psi, \delta))$  is minimized by  $\delta^0 = (\delta_0^0, \delta_1^0, \delta_2^0 \dots)$ .

The authors go on to say that the optimal stopping time is the following. If at time  $t$  the conditional risk, given the data up to time  $t$  is less than the conditional risk of taking one more observation and then stop immediately (Ghosh, 1997). This means that we should stop if we feel that the conditional risk will increase after the next observation. The proposed stopping rule will use the same idea except that the conditional risk will factor in both the marginal cost of testing as well as the penalty cost of making an incorrect decision.

When predicting the probability of batch failure in time  $t+1$  we will look at the conditional probability given that we see either an acceptable result or an unacceptable result. The mean of the beta distribution will serve as our best guess for the proportion of parts within tolerance at time  $t+1$  and will be used to weight our conditional probability. The expected probability of batch failure in time  $t+1$  will be

$$E[Failure^{t+1}] = \mu^t * P[(p < Req)|Beta^t(\alpha + 1, \beta)] + (1 - \mu^t) * P[(p < Req)|Beta^t(\alpha, \beta + 1)]$$

Where  $\mu^t$  is the mean of the current beta distribution and  $Req$  is the required proportion of good products in a batch. Intuitively, we can think of this process as narrowing the beta distribution. Both of the parameters will increase by a small amount. If the mean of the beta distribution is in the acceptable region, then this expectation will generally be smaller than the current value.

## Methods

The method used to analyze this stopping rule will be a computer simulation using R. Data points following a Bernoulli distribution will be generated for a certain batch size. This batch can be of any size but for the following analysis we will fix the batch size to 500 units. The other parameters used for generating data will be the true probability of success, which will be varied and compared in the results section.

Once the data are generated a prior distribution will be selected. As discussed, the prior distribution will follow a Beta distribution. The values of  $\alpha$  and  $\beta$  will be varied. The analysis of the prior will include an uninformed prior, which will be discussed at greater length later in the paper, as well as uninformative or flat priors. Informed priors will be generated using an R function called `beta.select`. `Beta.select` outputs a pair of  $\alpha$  and  $\beta$  parameters for corresponding quantile inputs. The informed prior information will be varied and discussed in the results section but will be restricted to values close to both the generated data and the batch requirement amount.

After the data and the prior are generated, all of the posterior distributions will be calculated. In practice the posterior distribution will be calculated after each observation and batch decision, but for simplicity of calculation, all of the posteriors will be determined at the beginning but the decision rule at time  $t$  will only be based on information available at time  $t$ .

Recall that the criteria for the stopping rule instructs the observer to stop testing when

$$[P(F^t|data^t) - P(F^{t+1}|data^{t+1})]PC < MC$$

The probability of batch failure will be directly calculated using the simulated data up to time t. The probability of batch failure in time t+1 will be estimated using

$$E[Failure^{t+1}] = \mu^t * P[p < Req|Beta^t(\alpha + 1, \beta)] + (1 - \mu^t) * P[p < Req|Beta^t(\alpha, \beta + 1)]$$

Once the stopping rule criteria are met the observation point and the expected cost at that point will be recorded.

In order to analyze the usefulness of this stopping rule the results will be compared to the optimal stopping point as well as the results of a 5% stopping rule. The simulation will be repeated for 100 trials. The results of these trials will be summarized and compared by both mean and standard deviation of average expected cost at the stopping point.

### Example Results

The analysis of this process will begin by examining how the parameters will affect one batch trial and then be extended to include more trials and summary results. For the first example we will consider the following parameter setup. The supplier has an agreement with a buyer that it will provide a batch of 500 products where 90% fall within a certain threshold. The quality tester will test products one at a time at a marginal cost of \$120. If the supplier ships the batch and fewer than 90% of the products are not within specification then the supplier will be charged a penalty cost of \$900,000. The supplier believes that there is a 90% probability that the batch being produced has a proportion of acceptable parts below 95%. The supplier also believes that there is a 10% probability that the batch being produced has a proportion of acceptable parts below 90%.

The data for the above example will be simulated using Bernoulli trials with probability of success .95. The quantile information tells us that the company is in fact making a conservative estimate regarding their process. A Beta(163.67,13.07) will satisfy the quantile information and serve as our prior guess for what the probability of batch success actually looks like. As stated above the data will now be simulated and a new prior will now be formed for each sequential data point. From this information, along with our cost parameters we will generate the expected total process cost using

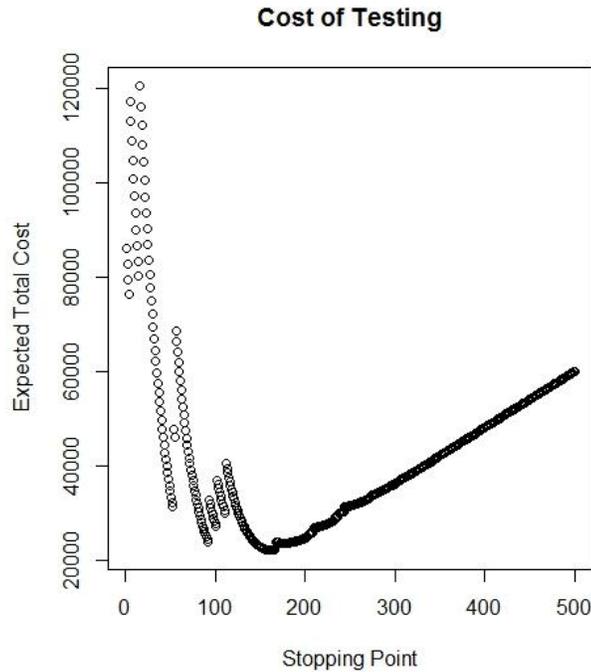
$$P(F^t | data^t)PC + MC * t$$

Where the probability of batch failure is found by calculated the area under the current posterior distribution at time t below the required proportion.

A single simulation of a batch yields the following results

Number of acceptable parts	Number of parts tested	Expected total cost of testing	Optimal Number of parts tested	Expected total cost of testing for optimal number
475	165	\$ 22115.70	165	\$22115.70

In this example we see that the proposed stopping rule aligned perfectly with what the optimal stopping point would be given the parameters and prior information. Note that the proposed stopping rule uses only information at time t while the optimal testing results use data over the entire range of batch data. If the expected total cost of testing is plotted for each possible stopping point, the following graph is obtained.



From the graph we can see that the expected total cost of testing increases sharply at certain points near the beginning of the process. These points result from viewing a negative observation, which, when not much information is known about the actual proportion of acceptable units in the batch, will cause a sharp increase in expected total cost. When the observer is fairly certain that the proportion of acceptable units is above the required amount, the probability that the batch is not within specification converges to zero. Near the end of the process the total cost is increasing linearly. The slope of this line will equal the marginal cost because the penalty cost becomes such a small portion of the total expected cost. For further analysis of the stopping rule and how it responds to different combinations of parameters, this example model will be replicated and parameters will be altered one at a time. Batches will be simulated 100 times and the mean and standard deviation will be returned.

## Summary Results

The following summary provides a look at variations in the process described above. The parameters will be varied one at a time and the summary of 100 batches will be compared among the proposed stopping rule, the optimal stopping amounts, and a 5% stopping rule.

### Prior

The prior information used by the observer will be varied between an uninformed prior and the relatively highly informed prior used in the example. Sun and Berger explain how to obtain an uninformed prior for the binomial distribution (Berger & Sun, 2008).

For a sequence of i.i.d. random variables with known density function and known proper stopping time for the sequential experiment, we find that

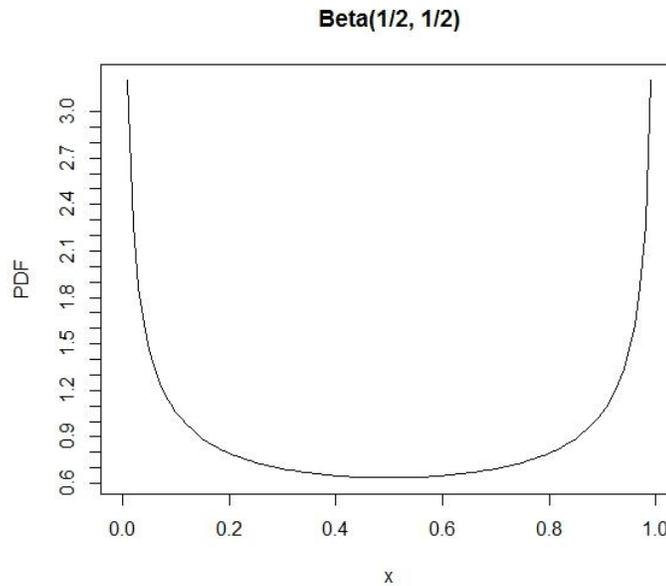
$$I^* = E_{\theta}(N)I(\theta).$$

Where N is the proper stopping time and  $I(\theta)$  is the Fisher information matrix. For sequential testing, an uninformed prior known as the Jeffery's prior is

$$\pi_j^*(\theta) \propto \{E_{\theta}(N)\}^{\frac{q}{2}} |I(\theta)|^{\frac{1}{2}} \propto \{E_{\theta}(N)\}^{\frac{q}{2}} \pi_j(\theta).$$

Applying the Jeffrey's prior to the binomial distribution will give us the following prior for the probability of success

$$\pi \sim \text{Beta}\left(\frac{1}{2}, \frac{1}{2}\right)$$



This result is well known and easily proved (Berger & Sun, 2008). This prior will be used to show our results for the unbiased case. In general, an informed prior will be used. This is due to the vast amount of information on past batches which manufacturers have at their disposal.

The following table summarizes the results of the simulations where the prior distribution is varied.

Prior	Proposed stopping rule		Minimum cost		5% stopping rule	
	Mean	SD	Mean	SD	Mean	SD
Uninformed	\$28,791	\$17,418	\$19,701	\$11,496	\$51,174	\$7,781
P(x<.95)=.6 P(x<.9)=.4	\$26,772	\$14,672	\$18,179	\$8,735	\$49,055	\$6,636
P(x<.95)=.7 P(x<.9)=.3	\$30,093	\$18,056	\$22,124	\$10,371	\$50,956	\$7,435
P(x<.95)=.8 P(x<.9)=.2	\$26,958	\$12,135	\$20,997	\$7,774	\$51,348	\$5,316
P(x<.95)=.9 P(x<.9)=.1	\$26,670	\$11,424	\$20,947	\$7,444	\$49,234	\$5,256

The prior information refers to how certain the observer is that the data will be within a given range. For example,  $P(x<.95)=.6$  and  $P(x<.9)=.4$  means that the observer is 60% certain that the proportion of acceptable parts in the batch is below 95% and is 60% certain that the proportion is greater than 90%. As we can see, variation in the prior information has little effect on which stopping rule is preferred in terms of total cost. The proposed stopping rule yields an average expected total cost which is lower than that of the 5% stopping rule. There also does not seem to be any linear relationship between the strength of the prior and the average expected total costs.

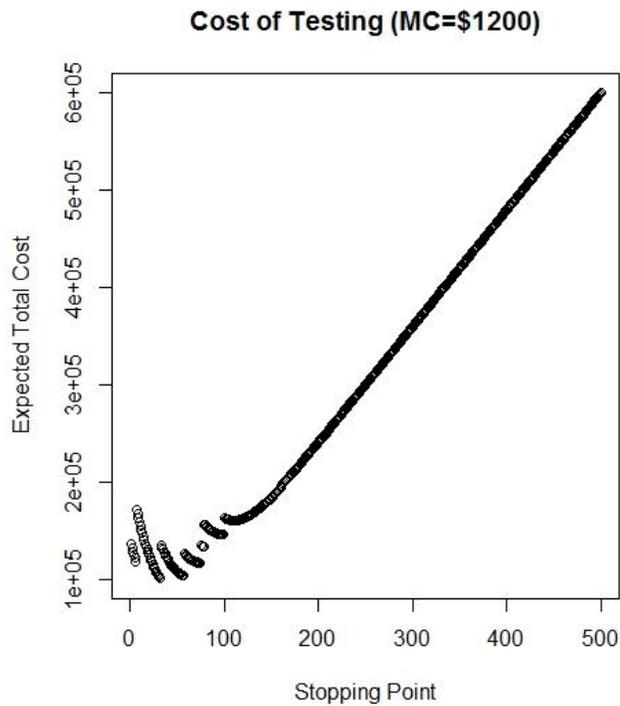
Another notable difference is that the 5% stopping rule yields a significantly lower standard deviation. This trend will be seen throughout most of the analysis. This low variability is due to the ignorance of the penalty cost for the 5% stopping rule. It is likely that the observer will be 95% certain that the batch falls within specification at relatively the same time in the observation process which will yield relatively similar expected total cost values. In our example, the 5% stopping rule, in general, tends to look at far fewer data points than the proposed stopping rule.

## Marginal Cost

The next variable to be examined is the marginal cost of testing. As stated, the marginal cost of testing can be either simply the costs associated with paying a worker to observe and record an observation, or, in the case of destructive testing, can be the actual marginal cost of the product. Like the other simulations the marginal cost will be simulated 100 times and summarized for each set of parameters.

Marginal Cost	Proposed stopping rule		Minimum cost		5% stopping rule	
	Mean	SD	Mean	SD	Mean	SD
\$12.00	\$4,290	\$2,851	\$3,557	\$1,344	\$45,048	\$487
\$60.00	\$15,483	\$6,171	\$12,774	\$4,624	\$46,860	\$2,605
\$120.00	\$25,391	\$10,391	\$19,867	\$6,516	\$48,558	\$3,450
\$240.00	\$46,153	\$20,747	\$32,613	\$9,538	\$54,238	\$10,151
\$1,200.00	\$107,490	\$55,924	\$68,460	\$10,299	\$86,457	\$37,080

The first trend that we see is the steep increase in the average expected cost for the proposed stopping rule as the marginal cost increases. We also notice that the average expected cost for the 5% rule also increases, but at a much slower rate. For a marginal cost of \$240 and \$1200 seen above we see that the 5% stopping rule actually leads to better results than the proposed stopping rule. To analyze this problem, we will examine one batch trial and see how the expected cost varies at different stopping times throughout the process.



From this simulation we see that we do not have the same checkmark shape present in the example problem. The 5% stopping rule dictates that for this data, only 40 observations will be taken. The proposed stopping rule takes 76 observations and will, as stated, yield a higher cost. This result means that if the Marginal cost is relatively high compared to the weighted penalty cost that the proposed stopping rule will not perform as well as other fixed percentage stopping rules. In order for the proposed stopping rule to be effective, the data must have an expected minimum point that is not at the very beginning of the process. When the expected total cost displays this strictly linear trend, it is likely that the proposed stopping rule will not yield the minimum expected cost. Next, we will explore whether the same results are found by varying the penalty cost rather than the marginal cost.

## Penalty Cost

The penalty cost represents additional costs for the producer associated with shipping a batch of products which is outside of the buyer specification. If the only consequence of shipping a bad batch is the expense of a return shipment then the penalty cost will be relatively low. The penalty cost may also include factors such as legal fees, explicit costs stated in the contract, and even the opportunity cost of lost business. For this reason it is very important to look at penalty costs over a large range of values.

Penalty Cost	Proposed stopping rule		Minimum cost		5% stopping rule	
	Mean	SD	Mean	SD	Mean	SD
\$90,000.00	\$10,969	\$7,259	\$6,925	\$1,106	\$9,258	\$4,318
\$450,000.00	\$23,488	\$10,143	\$17,528	\$6,528	\$27,944	\$5,785
\$900,000.00	\$28,436	\$13,605	\$22,237	\$8,172	\$49,159	\$4,977
\$1,800,000.00	\$34,147	\$19,372	\$24,532	\$9,259	\$94,071	\$3,936
\$9,000,000.00	\$43,806	\$17,240	\$35,967	\$12,188	\$450,868	\$5,512

After 100 batch trials the average expected cost for the proposed stopping seems to follow the minimum cost criteria closely. In fact, as the penalty cost increases by two orders of magnitude, the average expected cost for the proposed stopping rule increases by less than a multiple of four. The 5% stopping rule does outperform the proposed stopping rule at very low penalty cost levels, but increases sharply as penalty costs increase. This result supports the hypothesis proposed in the previous section. It does seem that as penalty cost and marginal cost become relatively closer together that the 5% stopping rule can outperform the proposed stopping rule with other factors held constant. The exact threshold that dictates when the 5% stopping rule becomes better is beyond the scope of this simulation.

## Requirement Amount

The contract between the buyer and producer will specify what proportion of the shipped products must be within some specification. It is possible that the required proportion of acceptable products will be 1. In this instance, the proposed stopping rule as well as the 5% stopping rule break down. This is because regardless of the posterior distribution, the area between 0 and 1 will always be 1. Also, at very low requirement amounts, holding the prior distribution constant, the decision can be to stop after the first observation. This means that the most interesting range of requirement proportions ranges between .85 and .95.

Required proportion	Proposed stopping rule		Minimum cost		5% stopping rule	
	Mean	SD	Mean	SD	Mean	SD
0.85	\$806	\$203	\$797	\$186	\$798	\$188
0.88	\$10,401	\$5,235	\$7,965	\$2,780	\$16,704	\$3,484
0.9	\$25,692	\$11,448	\$20,767	\$7,595	\$48,940	\$3,788
0.92	\$78,182	\$67,499	\$54,599	\$31,817	\$86,434	\$60,531
0.95	\$688,155	\$201,367	\$516,564	\$179,502	\$686,726	\$202,645

As predicted, when the required proportion is very low compared to the prior, both stopping rules will normally only take one observation and stop. Also, when the required proportion becomes closer to both the data generation as well as our prior information, the average expected costs increase sharply. When the required proportion of parts is very close to the actual data generation parameter, both stopping rules will very often test the entire batch. Since this simulation and analysis does not include adjustments for a known population size, when the entire batch is tested, regardless of

whether the batch is within or out of the requirement proportion, the probability that the true population parameter will be below .95 will be close to .5. In the future, a finite population correction may be added to the model for more definite results as the requirement proportion nears the data generation parameter.

### Data Generation

The data generation parameter refers to how the data will actually be simulated regardless of prior distribution. Ideally, the prior distribution should represent how the data generation process varies between batches but other values will also be explored. As hypothesized in the last section, it is believed that as the requirement amount and the data generation parameter near the same value that the process will begin to break down due to the lack of a finite population correction. Also, it is reasonable to believe that as the data generation parameter increases further above the requirement proportion that the proposed stopping rule will yield significantly better results than the 5% stopping rule.

Data Generation	Proposed stopping rule		Minimum cost		5% stopping rule	
	Mean	SD	Mean	SD	Mean	SD
0.9	\$390,610	\$218,104	\$52,642	\$22,278	\$261,386	\$262,201
0.91	\$204,235	\$173,919	\$45,699	\$23,346	\$141,244	\$173,084
0.92	\$103,313	\$94,770	\$38,391	\$164	\$71,805	\$78,213
0.95	\$25,543	\$11,002	\$20,059	\$6,995	\$49,257	\$4,393
0.98	\$14,403	\$3,065	\$13,739	\$2,668	\$47,157	\$1,345

As hypothesized, the process does indeed break down as the data generation parameter is very close to the requirement proportion. The average expected costs are

very high when data generation is .9 and .91. The simulations also show that as the data generation parameter increases, the proposed stopping rule becomes closer to the minimum cost. In fact, it is intuitive and simple to prove that if the data is generated with 100% of observations falling within specification, that the proposed stopping rule will converge exactly to the minimum expected cost and trials will have absolutely no variation when other parameters are held constant.

Comparison to Other Fixed Percent Stopping Rules

The final parameter to be varied is the percentage that the comparison stopping rule will use for its stopping criteria. Since it has been established that the 5% stopping rule, in general, will not observe as many observations as the proposed stopping rule or the optimal amount, perhaps another stopping rule will perform better. It would seem that a stricter stopping rule based on probability of batch failure may work better than the 5% rule.

X	Proposed stopping rule		Minimum cost		X% stopping rule	
	Mean	SD	Mean	SD	Mean	SD
0.001	\$26,634	\$10,202	\$20,239	\$6,713	\$31,665	\$12,286
0.025	\$27,361	\$10,557	\$20,909	\$6,863	\$31,501	\$6,210
0.05	\$25,603	\$11,510	\$20,305	\$7,102	\$49,023	\$4,121
0.075	\$27,014	\$12,462	\$20,175	\$6,300	\$68,529	\$2,307
0.1	\$26,681	\$13,086	\$20,244	\$7,348	\$86,382	\$882

All of the parameters in the simulations are the same as those used in the first example problem. The only change being made is the probability of batch failure criteria used in the comparison stopping rule. Because of this, the average expected cost should be relatively the same for the proposed stopping rule across the simulations. As expected,

the average expected cost is closer to the optimal amount for stricter stopping rules. The simulations do not, however, show that any of the considered fixed percent stopping rules yields a lower average expected cost than the proposed stopping rule. Also, as the fixed percent stopping rule becomes less strict, the average expected cost increases. This is due to the smaller samples and greater probability of incurring the penalty cost under these conditions.

## Conclusion

After analysis of the simulations, a number of relationships between the variable may be discussed. The first interesting result is that when the marginal cost is relatively high compared to the penalty cost, then the proposed stopping rule will rarely find the minimum expected cost. This is due to the strict linear relationship that results from this type of simulation. The proposed stopping rule performs better under conditions where the minimum expected cost does not occur at the beginning of the testing process. The simulations merely give the observer intuition about how the cost parameters will affect the expected cost under the proposed stopping rule. In the future, this relationship between the cost variables and the prior distribution may be analyzed theoretically. This analysis could result in exact threshold which the proposed stopping rule will approximate the true minimum expected cost.

A second implication of the simulations is that the process begins to break down as the required proportion of acceptable observations and the data generation parameter become closer. This result is partially due to the lack of a finite population correction term. This term would adjust the probability that the batch is within specification to account for the fixed batch size. Future research involving this stopping rule would benefit from including this term when analyzing how the rule reacts when the data generation parameter is close to the requirement proportion.

One final interesting result is that the proposed stopping rule converges to the optimal stopping rule as data generation nears 1. This means that if the data is being generated perfectly within specification that the stopping rule will perfectly align with what the minimum cost of testing truly is. From this result it seems likely that very strict processes would benefit highly from the proposed stopping rule.

In the future the proposed stopping rule can be applied to different sampling distributions. For example, if the test is for measuring tensile strength and it is believed that tensile strength is normally distributed then instead of Bernoulli trials, the data would be generated from a normal distribution. This method would be useful because observations which are severely outside of the acceptance region will affect the posterior more than observations marginally outside of the acceptance region. This modification would alter the way that the expected cost is generated but will not change the criteria for the stopping rule.

### Works Cited

Bickel, P. J., & Yahav, J. A. (1968). Asymptotically optimal bayes and minimax procedures in sequential estimation. *The Annals of Mathematical Statistics*, 39(2), 442-456

Ghosh, M., Mukhopadhyay, N., & Sen, P. K. (1997). *Sequential estimation*. New York: John Wiley & Sons, Inc.

Sun, D., & Berger, J. O. (2008). Objective bayesian analysis under sequential experimentation. *IMS Collections*, 3, 19-32.

## R Code

```
#####  
#Testing costs of dichotomous data with informed prior  
#####  
###  
#Before running program be sure to load the LearnBayes pack onto the computer  
###  
library(LearnBayes)  
  
#next, enter the required inputs  
batch=500    #batch size of goods produced  
req=.9  
#Cost Parameters  
MC=1200      #Marginal cost  
C2=900000    #penalty cost  
trials=100   #Number of batch testing procedures  
  
###This loop applies to one batch testing process. The results of this loop will be output  
and compared  
Mypoint<-rep(0,trials)  
Truepoint<-rep(0,trials)  
Estmin<-rep(0,trials)  
Truemin<-rep(0,trials)  
Fivepoint<-rep(0,trials)  
Fivecost<-rep(0,trials)  
for (t in 1:trials){  
  
    ###generate prior distribution by entering quantile information  
    quantile2=list(p=.9,x=.95)  
    quantile1=list(p=.1,x=.90)  
    BPrior<-beta.select(quantile1,quantile2)  
    #BPrior<-c(.5,.5)  
    posterior_a<-matrix(0,batch,1)  
    posterior_b<-matrix(0,batch,1)  
    #posterior_a[1]=BPrior[1] #Prior for shape parameter 1  
    #posterior_b[1]=BPrior[2] #Prior for shape parameter 2
```

```

###Sampling Procedure
Data<-matrix(0,batch,1)
Data=rbinom(size=1,n=batch,p=.95)
Sample<-matrix(0,batch,batch)
for(i in 1:batch){
  Sample[i,1:i]=Data[1:i]
}
###This loop will generate the posterior distribution after every sample from the batch
for(i in 1:batch){
  posterior_a[i,1] =sum(Sample[i,1:i]) + BPrior[1]
  posterior_b[i,1] =i - sum(Sample[i,1:i]) + BPrior[2]
}

###This loop will calculate the probability of each beta distribution will not meet the
batch requirement
pfail<-matrix(0,batch,1)
for(k in 1:batch){
  pfail[k]=pbeta(req,posterior_a[k],posterior_b[k])
}

efail<-matrix(0,batch-1,1)
for (k in 1:batch){
  efail[k-1]=((posterior_a[k-1]/(posterior_a[k-1]+posterior_b[k-1]))*(pbeta(req,posterior_a[k-1]+1,posterior_b[k-1])))+((1-(posterior_a[k-1]/(posterior_a[k-1]+posterior_b[k-1])))*(pbeta(req,posterior_a[k-1],posterior_b[k-1]+1)))
}

totalcost<-matrix(0,batch,1)
for(l in 1:batch){
  totalcost[l]=l*MC+pfail[l]*C2
}
excost<-matrix(0,batch-1,1)
for(l in 1:batch-1){
  excost[l]=l*MC+efail[l]*C2
}

chcost<-matrix(0,batch-1,1)
index<-rep(0,batch-1)

for (k in 1:batch-1){
  chcost[k]=C2*(efail[k+1]-pfail[k])
}

```

```

    index[k]=k
  }

  ### This code will retrieve my stopping point for this batch trial
  d<-matrix(0,batch-1,1)
  e<-matrix(0,batch-1,1)
  d=ifelse (abs(chcost) < MC,1,0)
  e=d*index
  e=ifelse(e==0,499,e)
  Mypoint[t]=min(e[1:498])
  Estmin[t]=excost[Mypoint[t]]

  ### This code will retrieve the 5% stopping point
  d<-matrix(0,batch-1,1)
  e<-matrix(0,batch-1,1)
  d=ifelse (efail<.05,1,0)
  e=d*index
  e=ifelse(e==0,500,e)
  Fivepoint[t]=min(e)
  Fivecost[t]=totalcost[Fivepoint[t]]

  for (k in 1:batch){
    index[k]=k
  }

  ### This code will retrieve the optimal stopping point
  d<-matrix(0,batch,1)
  e<-matrix(0,batch,1)
  d=ifelse (min(totalcost)==totalcost,1,0)
  e=d*index
  e=ifelse(e==0,500,e)
  Truepoint[t]=min(e[1:500])
  Truemin[t]=totalcost[Truepoint[t]]
}

sum(Estmin)/trials
sqrt(var(Estmin))
sum(Truemin)/trials
sqrt(var(Truemin))
sum(Fivecost)/trials
sqrt(var(Fivecost))

```