

Incorporating Uncertain Costs within a Series of Sequential Probability Ratio Tests

Conor McMeel, Brett Houlding*

Discipline of Statistics, Trinity College Dublin, Dublin, Ireland

Email: mcmeelc@tcd.ie, *brett.houlding@tcd.ie

How to cite this paper: McMeel, C. and Houlding, B. (2016) Incorporating Uncertain Costs within a Series of Sequential Probability Ratio Tests. *Open Journal of Statistics*, 6, 882-897.

<http://dx.doi.org/10.4236/ojs.2016.65073>

Received: August 20, 2016

Accepted: October 18, 2016

Published: October 21, 2016

Copyright © 2016 by authors and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

<http://creativecommons.org/licenses/by/4.0/>



Open Access

Abstract

We consider an extension to Sequential Probability Ratio Tests for when we have uncertain costs, but also opportunity to learn about these in an adaptive manner. In doing so we demonstrate the effects that allowing uncertainty has on observation cost, and the costs associated with Type I and Type II error. The value of information relating to modelled uncertainties is derived and the case of statistical dependence between the parameter affecting decision outcome and the parameter affecting unknown cost is also examined. Numerical examples of the derived theory are provided, along with a simulation comparing this adaptive learning framework to the classical one.

Keywords

Adaptive Utility, Hypothesis Testing, Sequential Analysis, Sequential Probability Ratio Test

1. Introduction

Sequential Probability Ratio Tests (SPRTs) were introduced by Wald in 1945 [1] [2] as a sequential hypothesis test procedure for when data is considered in sequence rather than in entirety. They have been used in many fields of industry, for example: nuclear physics [3], medicine [4] [5], standardised testing [6] and radar detection [7], to name just a few, and even though the classical theory has now been known for some seven decades, they are still the subject of research into extensions and generalisations [8]-[10].

Generally, the objective of a SPRT is to balance the consequence of an error with the cost of acquiring further data and/or making additional observations, e.g. clinical trials, or stress tests. In this approach data is sought until the belief in the state of nature (namely the parameter controlling the decision outcome) is such that the expected cost of implementing the current optimal decision is less than that expected from seeking

additional data, updating beliefs, and then implementing the (possibly different) optimal decision.

In its simplest form a SPRT consists of the following: A choice between two decisions or courses of action (here denoted d_0 and d_A), and a state of nature w that can take one of two possible values (w_o or w_A). Depending on the decision that is selected and the true state of nature value, one of three possible losses may occur. Without loss of generality we assume a loss of $c_1 > 0$ occurs if d_A is selected when w_o is true, $c_2 > 0$ occurs if d_0 is selected but w_A is true, and a loss of 0 otherwise (Table 1).

From the above it can be seen that the objective of the Decision Maker (DM) is to choose between d_0 and d_A on the basis of their beliefs over the state of nature, seeking to match the decision to what they hope is its correct value. In general we denote such belief by the probability $\phi \in [0,1]$ with $P(w_o) = \phi = 1 - P(w_A)$. In this sense the losses c_1 and c_2 can be associated with what is commonly described as making a Type I or Type II error, and hence the connection to sequential hypothesis testing originally considered by Wald.

A graphic representation of this protocol is illustrated in Figure 1, where the x -axis varies over the possible value of ϕ and the y -axis is the resulting expected loss incurred

Table 1. Loss table applied within a SPRT.

	d_0	d_A
w_o	0	c_1
w_A	c_2	0

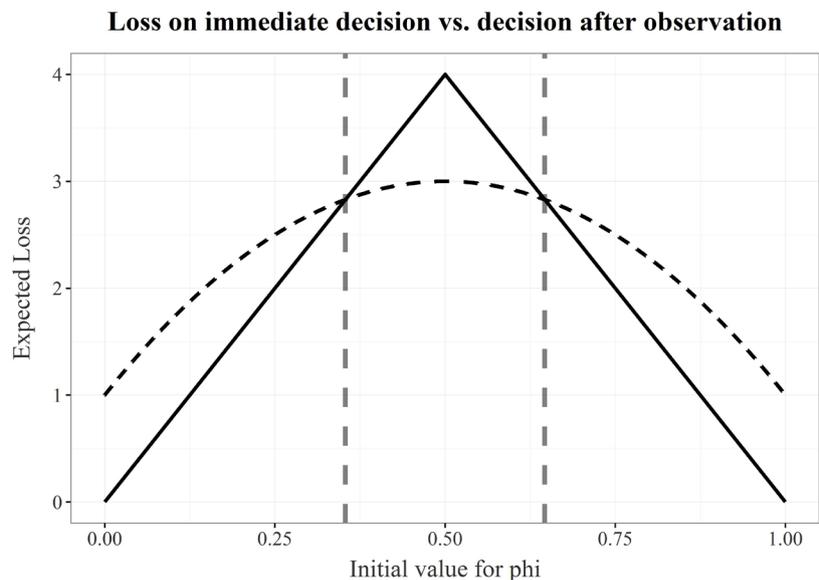


Figure 1. An illustration of the losses involved within a SPRT. The x -axis varies over $\phi = P(w_o)$, the y -axis is the associated expected loss, the solid line corresponds to making a decision immediately, while the curved dashed line corresponds to collecting further information before making a decision. Finally the vertical dashed lines indicate the bounds on ϕ within which the DM should observe further data before making a decision.

by a particular strategy. The solid line represents the expected cost of implementing a decision immediately, whilst the curved dash line corresponds to the expected loss of implementing a decision only after taking some further data (at a cost) concerning the correct value of the state of nature. It is included as a curve as it can be shown that the expected loss of deciding after data collection is a concave function of ϕ (this is because we will be taking the infimum of two further choices, namely to act once the data is collected or to again choose to sample).

For values of $\phi \in \{0,1\}$ the DM assumes for sure that they know what the state of nature will be, and hence will make a decision in the belief that they will receive a cost of 0. As ϕ varies away from these extremes however, the DM will not presume to be certain in their knowledge of the state of nature, and hence expects a risk of making either a Type I or Type II error and incurring the associated cost. This risk can be shown to increase and then decrease linearly between the extreme values of ϕ (the change from increase to decrease occurring at $\phi = 0.5$ when $c_1 = c_2$). The dashed curve line, corresponding to making a decision only after further data collection, does not have an expected loss of 0 at $\phi \in \{0,1\}$ because of the additional cost of collecting data. Depending on what this particular data collection cost is, the DM should either always collect further information (when the cost is 0), never collect additional information (when the cost of doing so is prohibitive compared to the cost of actually making a Type I or Type II error), or as is the case in **Figure 1**, either choose to collect additional data or not to depending on the value of ϕ that they assign to $P(w_0)$. The vertical dashed lines of

Figure 1 indicates, for the particular numerical example displayed, the range of values for ϕ within which the DM expects it is better to collect further data before making a decision.

Whilst the approach described above outlines the classical way of performing a SPRT, it fails to take into account that in practice, many of the costs involved will not be known for certain. For example, in the case of an observation cost, the cost associated with undertaking clinical trials prior to deciding to market a drug may not be known for sure, or in the case of a Type I or Type II error, the reputational or financial effect of implementing a poor decision may be unknown, *e.g.*, releasing poorly coded software when there was opportunity to have more testing to determine unknown bugs. In such instances it is then natural for us to model our beliefs and uncertainties about relevant costs according to some parameter, say θ , to which we only specify a prior distribution. The question then arises as to the effect this has on how we perform a SPRT, given that we may now learn between successive SPRTs, or in the case of unknown observation cost, between successive observations.

The concept of unknown utility (utility defined to be negative loss), but which may instead be learned through experience, is the topic of adaptive utility theory first considered by Cyert and De Groot [11]. Here not only do we have uncertainty concerning decision outcome (as modelled through the unknown state of nature), but also in the preferences over those outcomes (or equivalently attitudes to risk) [12]-[14].

In the case of only performing a solitary SPRT, and where the uncertainty relates to only the consequence of a Type I or Type II error, the appropriate procedure is equivalent to the classical one with the cost assigned to its expected value, as there is no possibility to learn about the relevant costs before implementing a decision. However, if the DM has opportunity to purchase information about such costs, e.g. by performing some market survey or enlisting the assistance of a knowledgeable expert, then the value that such information is worth may be calculated as the expected difference between the expected loss without the information, and the expected loss with it. Determining this value will be our primary interest [14] [15]:

$$\max_{d \in D} E[L(d)] - E_{i \in I} \left[\max_{d \in D} E[L(d|i)] \right]. \quad (1)$$

Here I is the set of information statements we could receive, i is an actual information statement, D represents the set of available decisions, d a particular decision, and $L(d)$ the expected loss for implementing decision d .

The remainder of this paper is as follows: In Section 2 we consider SPRTs with uncertain Type I or Type II error cost followed by uncertain observation cost in Section 3. In the former we consider the value of perfect information and that of noisy information, along with providing numerical examples. The details of a simulation carried out in the case of perfect information are also given. Finally we conclude in Section 4.

2. Unknown Consequence of Error

Suppose our uncertainty does not concern the cost of taking further observations, but rather the cost of a Type I or Type II error, or both (possibly with different distributions describing these). Without loss of generality assume the uncertainty is with respect to the cost of a Type I error only. In this case our loss table is as in **Table 2**, where θ represents the uncertain cost of a Type I error.

There are three steps to perform to generate the expected value of information relating to these uncertain costs. In the case of information being perfect then these are the following:

- 1) Consider a SPRT when no information is learned.
- 2) Obtain expected loss following learning of the uncertain parameter(s).
- 3) Subtract to obtain the expected value of information, which can then be subtracted from the unknown loss consequence(s).

In performing step 1 we utilize the expected costs using the loss in **Table 3**. Hence the expected loss on making an immediate decision, as a function of ϕ , is:

$$\rho(\phi) = \min(E[\theta]\phi, c_2(1-\phi)) = \begin{cases} E[\theta] & \text{if } \phi \leq \frac{c_2}{E[\theta] + c_2} \\ c_2(1-\phi) & \text{otherwise} \end{cases}. \quad (2)$$

If a Type I or Type II error is made, we learn the value of θ . The process of performing an SPRT is repeated, but now with the exact value of θ rather than its prior expectation $E[\theta]$, resulting in a change in the expected risk profile. An *expected*

Table 2. Loss table with uncertain cost of Type I error.

	d_0	d_A
w_o	0	θ
w_A	c_2	0

Table 3. Loss table assuming no learning about uncertain cost θ .

	d_0	d_A
w_o	0	$E[\theta]$
w_A	c_2	0

Table 4. Loss table for numerical example in Section 2.1.

	d_0	d_A
w_o	0	2
w_A	$E[\theta]$	0

expected risk profile concerning how this may look depending on what is learned, based on the prior distribution of θ , is now determined. This is then subtracted from the original risk profile (using $E[\theta]$) to obtain the expected value of the cost information.

2.1. Perfect Information Numerical Example

As a toy example illustrating this situation consider testing if a sequence of coins are fair (w_o) or biased (w_A) meaning that $P(H) = 0.8$. A Type I error corresponds to throwing away a fair coin and we suppose this has known loss of 2 units, namely the coin's value. A Type II error would correspond to accepting a biased coin, of which we have little experience. This could be very bad resulting in a loss of $\theta = 4$ units, or not so bad resulting in a loss of $\theta = 1$ unit. Further suppose the prior on θ is such that $P(\theta = 1) = 2/3$ and let ϕ represent the probability that the coin is fair.

The expected loss table is given in **Table 4**. From the description we see $E[\theta] = 2$ so that if $\phi = 0.5$ then *a priori* we are indifferent between saving the coin or throwing it away. Then the expected risk of an immediate decision is:

$$\rho(\phi) = \min(2\phi, 2(1-\phi)) = \begin{cases} 2\phi & \text{if } \phi \leq 1/2 \\ 2(1-\phi) & \text{otherwise} \end{cases} \tag{3}$$

We may also consider sampling data by flipping a coin which is assumed to cost 0.1 units. We now calculate the range of ϕ where it is beneficial to flip the coin. To do so we determine posteriors on ϕ after observing the possible results of a coin flip:

$$P(w = w_o | H) = \phi | H = \frac{0.5\phi}{0.5\phi + 0.8(1-\phi)} = \frac{5\phi}{8-3\phi} \tag{4}$$

$$P(w = w_o | T) = \phi | H = \frac{0.5\phi}{0.5\phi + 0.2(1-\phi)} = \frac{5\phi}{2+3\phi} \tag{5}$$

The predictive probability of observing heads or tails, at any point, is:

$$P(H) = 0.5\phi + 0.8(1-\phi) \quad (6)$$

$$P(T) = 0.5\phi + 0.2(1-\phi). \quad (7)$$

Letting ϕ' denote our posterior probability of the coin being biased after being flipped, the risk profiles of the decision following an observation is:

$$\rho(\phi') = \begin{cases} 2\phi' & \text{if } \phi' \leq 1/2 \\ 2(1-\phi') & \text{otherwise} \end{cases}. \quad (8)$$

Now we can relate the bounds on ϕ' to bounds on ϕ :

$$\phi | H \leq \frac{1}{2} \Rightarrow \frac{5\phi}{8-3\phi} \leq \frac{1}{2} \Rightarrow \phi \leq \frac{8}{13} \quad (9)$$

$$\phi | T \leq \frac{1}{2} \Rightarrow \frac{5\phi}{2+3\phi} \leq \frac{1}{2} \Rightarrow \phi \leq \frac{2}{7}. \quad (10)$$

Suppose $\phi \in [2/7, 8/13]$, then the expected loss is:

$$\begin{aligned} & 2 \times \frac{5\phi}{8-3\phi} (0.5\phi + 0.8(1-\phi)) + 2 \times \left(1 - \frac{5\phi}{2+3\phi}\right) (0.5\phi + 0.2(1-\phi)) \\ & = 0.4 + 0.6\phi \end{aligned} \quad (11)$$

We also need to include the cost of flipping (0.1 units) resulting with an expected loss for observing once then deciding of $0.5 + 0.6\phi$. To see when this risk is preferable to deciding immediately we solve the following inequalities for ϕ :

$$0.5 + 0.6\phi \leq 2\phi \Rightarrow \phi \geq \frac{5}{14} \quad (12)$$

$$0.5 + 0.6\phi \leq 2(1-\phi) \Rightarrow \phi \leq \frac{15}{26}. \quad (13)$$

Observations should continue to be taken until ϕ leaves this interval, at which a point a decision should be made. Hence the expected risk profile is:

$$\rho(\phi) = \begin{cases} 2\phi & \text{if } \phi \leq 5/14 \\ 0.5 + 0.6\phi & \text{if } 5/14 < \phi \leq 15/26 \\ 2(1-\phi) & \text{if } 15/26 < \phi \end{cases}. \quad (14)$$

With this risk profile we now compute the expected loss assuming we know the parameter θ . There are two cases: when $\theta = 1$ or $\theta = 4$. In each case the expected loss as a function of ϕ is computed. The process is identical to the above so we just report them:

For $\theta = 1$:

$$\rho(\phi) = \begin{cases} 2\phi & \text{if } \phi \leq 1/4 \\ 0.3 + 0.8\phi & \text{if } 1/4 < \phi \leq 7/18 \\ (1-\phi) & \text{if } 7/18 < \phi \end{cases}. \quad (15)$$

For $\theta = 4$:

$$\rho(\phi) = \begin{cases} 2\phi & \text{if } \phi \leq 1/2 \\ 0.9 + 0.2\phi & \text{if } 1/2 < \phi \leq 31/42 \\ 4(1-\phi) & \text{if } 31/42 < \phi \end{cases}. \quad (16)$$

Recall the prior on θ was such that $P(\theta = 1) = 2/3$. Hence, the expected risk after learning θ is:

$$\rho(\phi) = \begin{cases} 2\phi & \text{if } \phi \leq 1/4 \\ 1/5 + 6\phi/5 & \text{if } 1/4 < \phi \leq 7/18 \\ 2/3 & \text{if } 7/18 < \phi \leq 1/2 \\ 29/30 - 3\phi/5 & \text{if } 1/2 < \phi \leq 31/42 \\ 2(1-\phi) & \text{if } 31/42 < \phi \end{cases} \quad (17)$$

Thus the expected value of perfect information is the difference between Equation (14) (without perfect information) and Equation (17) (with perfect information):

$$g(\phi) = \begin{cases} 0 & \text{if } \phi \leq 1/4 \\ -1/5 + 4\phi/5 & \text{if } 1/4 < \phi \leq 5/14 \\ 3/10 - 3\phi/5 & \text{if } 5/14 < \phi \leq 7/18 \\ -1/6 + 3\phi/5 & \text{if } 7/18 < \phi \leq 1/2 \\ -7/15 + 6\phi/5 & \text{if } 1/2 < \phi \leq 15/26 \\ 31/30 - 7\phi/5 & \text{if } 15/26 < \phi \leq 31/42 \\ 0 & \text{if } 31/42 < \phi \end{cases} \quad (18)$$

This represents the maximum amount of units we should be prepared to forsake in order to be informed the true value of the cost parameter θ prior to commencing the SPRT. From this we obtain a new function $\theta(\phi) = E[\theta] - g(\phi)$ that represents the loss resulting from the occurrence of a Type II error:

$$\theta(\phi) = \begin{cases} 2 & \text{if } \phi \leq 1/4 \\ 11/5 - 4\phi/5 & \text{if } 1/4 < \phi \leq 5/14 \\ 17/10 + 3\phi/5 & \text{if } 5/14 < \phi \leq 7/18 \\ 13/6 - 3\phi/5 & \text{if } 7/18 < \phi \leq 1/2 \\ 37/15 - 6\phi/5 & \text{if } 1/2 < \phi \leq 15/26 \\ 29/30 + 7\phi/5 & \text{if } 15/26 < \phi \leq 31/42 \\ 2 & \text{if } 31/42 < \phi \end{cases} \quad (19)$$

Equation (19) represents the expected value of the loss of making a Type II error, but discounted by the fact that we obtain information which allows more informed decisions to be made in subsequent SPRTs. A plot of (19) is given in **Figure 2** where it can be observed that local minima in the expected loss occur at boundaries of indifference between choices in the initial SPRT, and that plateaus in the expected loss coincide with values of ϕ where it is never beneficial to take an observation for any value of θ .

2.2. Noisy Information

Now assume we only receive noisy observations concerning θ meaning that following observation we are not certain of its value. The procedure is similar to that for perfect information and again we first perform an SPRT without considering the value of the information.

Denoting the true value of θ as θ_T and our observation as θ_{ob} then this setting

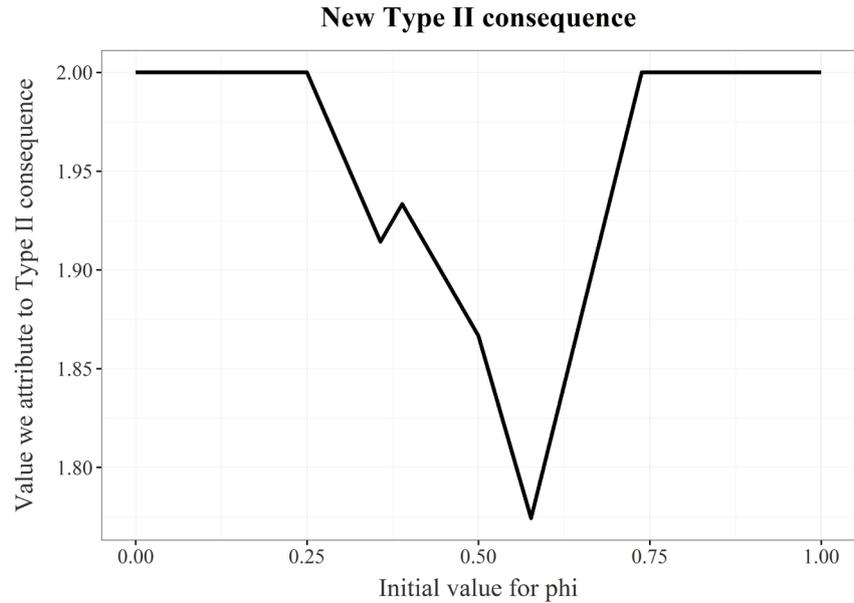


Figure 2. A plot of the expected loss incurred from committing a Type II error for Example 2.1 generated from Equation (19). The x-axis varies over the prior probability for the state of nature w , whilst the y-axis indicates the resulting expected loss.

means $P(\theta_T = \theta_{ob}) = p < 1$. We also allow the distribution over what is observed to depend on the true value θ_T and hence generate a likelihood $P(\theta_{ob} | \theta_T)$, from which a marginal distribution $P(\theta_{ob})$, expected value $E[\theta_{ob}]$, and posterior distribution $P(\theta_T | \theta_{ob})$ may be calculated in the usual way.

For each potential θ_{ob} a new expected value of θ , $E[\theta | \theta_{ob}]$, is generated. Denoting $E[L_\theta]$ as the expected loss before observation and $E[L_\theta | \theta_{ob}]$ as the expected loss after observing θ_{ob} , the expected value of the noisy observation is calculated as:

$$E[L_\theta] - \sum_{\theta_{ob}} E[L_\theta | \theta_{ob}] P(\theta_{ob}). \tag{20}$$

Once this has been generated the consequence of the error will be reduced in the risk table just as was the case with perfect information, allowing a classical SPRT to be performed.

2.3. Noisy Information Numerical Example

We return to the setting of Example 2.1, but now assume that the probability that the true value is observed is only 0.8, i.e., $P(\theta_T = \theta_{ob}) = 0.8$. This results in $P(\theta_{ob} = 1) = 3/5$ and $P(\theta_{ob} = 4) = 2/5$.

After observing a value for θ we update its expected value to the following:

$$E[\theta | \theta_{ob} = 1] = \frac{8}{9} \times 1 \times \frac{1}{9} \times 4 = \frac{4}{3} \tag{21}$$

$$E[\theta | \theta_{ob} = 4] = \frac{1}{3} \times 1 \times \frac{2}{3} \times 4 = 3. \tag{22}$$

Letting $E[L_x]$ represent the expected loss when $E[\theta] = x$ (so, for example, $E[L_2]$ is the loss from step 1 where future trials are not considered), then the value of

information from our noisy observation θ_{ob} for each ϕ is calculated as:

$$E[L_2] - \frac{3}{5}E[L_{4/3}] - \frac{2}{5}E[L_3]. \tag{23}$$

Note that each term in the above implicitly depends on the initial value assigned to ϕ . We then simply proceed as in Example 2.1 to obtain the final decision rule. The resulting loss tables are provided below for the three quantities listed in Equation (23):

$$E[L_2] = \begin{cases} 2\phi & \text{if } \phi \leq 5/14 \\ 0.5 + 0.6\phi & \text{if } 5/14 < \phi \leq 15/26 \\ 2(1-\phi) & \text{if } 15/26 < \phi \end{cases} \tag{24}$$

$$E[L_{4/3}] = \begin{cases} 2\phi & \text{if } \phi \leq 11/38 \\ 11/30 + 11\phi/15 & \text{if } 11/38 < \phi \leq 29/62 \\ 4(1-\phi)/3 & \text{if } 29/62 < \phi \end{cases} \tag{25}$$

$$E[L_3] = \begin{cases} 2\phi & \text{if } \phi \leq 7/16 \\ 0.7 + 0.4\phi & \text{if } 7/16 < \phi \leq 23/34 \\ 3(1-\phi) & \text{if } 23/34 < \phi \end{cases} \tag{26}$$

As a result the expected value of noisy information $g(\phi)$ is calculated as:

$$g(\phi) = \begin{cases} 0 & \text{if } \phi \leq 11/38 \\ -11/50 + 19\phi/25 & \text{if } 11/38 < \phi \leq 5/14 \\ 8/25 - 16\phi/25 & \text{if } 5/14 < \phi \leq 7/16 \\ 0 & \text{if } 7/16 < \phi \leq 29/62 \\ -29/50 + 31\phi/25 & \text{if } 29/62 < \phi \leq 15/26 \\ 23/25 - 34\phi/25 & \text{if } 15/26 < \phi \leq 23/34 \\ 0 & \text{if } 23/34 < \phi \end{cases} \tag{27}$$

Now the new expected cost of a Type II error $\theta(\phi)$ for the noisy information example can be determined as in see Equation (26). A plot of this function is given in **Figure 3** which can be contrasted with the perfect information case given in **Figure 2**. Note that as before the minima occur at boundaries of indifference and that plateaus occur where we would always (or never) take an observation no matter the value of θ . Also note that in comparison to **Figure 2**, the result for noisy information results in a larger expected cost of Type II error when the true value of θ does play a role in the decision making. This is to be expected due to the weaker and less useful noisy information in comparison to what we learn from perfect information.

$$\theta(\phi) = \begin{cases} 2 & \text{if } \phi \leq 11/38 \\ 111/50 - 19\phi/25 & \text{if } 11/38 < \phi \leq 5/14 \\ 42/25 + 16\phi/25 & \text{if } 5/14 < \phi \leq 7/16 \\ 2 & \text{if } 7/16 < \phi \leq 29/62 \\ 129/50 - 31\phi/25 & \text{if } 29/62 < \phi \leq 15/26 \\ 27/25 + 34\phi/25 & \text{if } 15/26 < \phi \leq 23/34 \\ 2 & \text{if } 23/34 < \phi \end{cases} \tag{28}$$

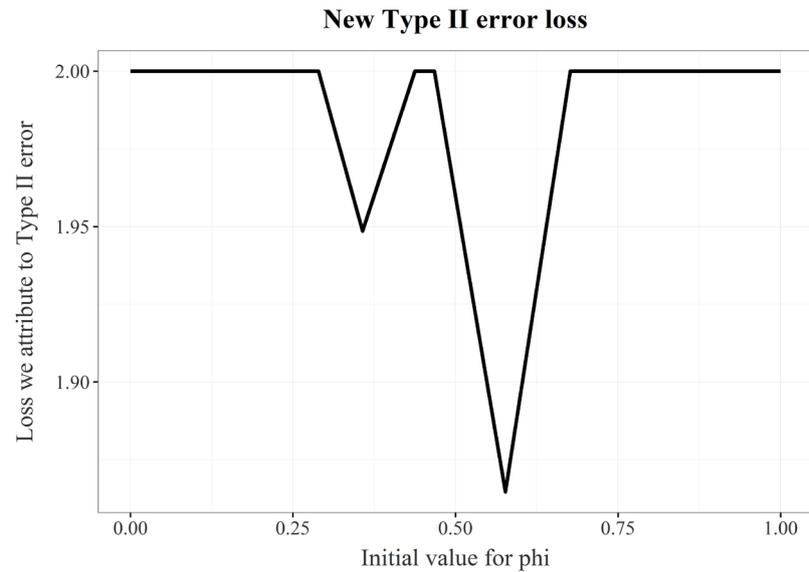


Figure 3. A plot of the expected loss incurred from committing a Type II error for Example 2.3 generated from Equation (28). The x -axis varies over the prior probability for the state of nature w , whilst the y -axis indicates the resulting expected loss.

2.4. Numerical Simulation

Details of a numerical simulation are now provided. The scenario detailed in Example 2.1 was tested in R [16] by considering the outcome of 3 million trials of both the classic and adaptive framework.

Each classical trial consisted of:

- 1) A SPRT with consequence of Type I/II error of 2 and cost of observation 0.1 run repeatedly until a Type II error is made. The bounds used are those in Equation (14), namely $[5/14, 15/26]$, before value of information is considered.

- 2) Upon making a Type II error, the cost from that particular SPRT is stored. The value of θ is then learned and another SPRT is run using the true value for the consequence of Type II error. The two costs are added to provide the total value for that trial.

In accordance with our prior on θ , two-thirds of the trials were performed with the true value of $\theta = 1$, while the others had $\theta = 4$.

A further 3 million trials were then run using the adaptive framework under the same procedure but with the bounds in step 1 being different. This is due to the different values used for consequence of Type II error seen in Equation (19). Using initial values of $\phi = 1/2$ corresponded to using $\theta = 28/15$, resulting in bounds of approximately $[0.34, 0.56]$. The second step remains the same as the classical trial.

The average costs are given in Table 5. As can be seen, this indicated a substantial improvement (21% with the numerical scenario here) in using the adaptive framework and formally taking such uncertainty into account.

2.5. Statistical Dependence

To conclude we give a brief discussion on the effect of their being statistical dependence

between the state of nature w and cost parameter θ . Without loss of generality, consider a joint distribution as taking on the values (and associated probability) given in **Table 6**. This implies conditional probabilities as given in **Table 7**. Note that this specification ensures that w and θ are not independent.

Now consider the implementation of the SPRT. The initial loss table when w and θ were independent is given in **Table 8**. However, note that we can only incur losses governed by θ when the state of nature is w_A . So any loss that occurs in the joint distribution when w_0 is true should not be considered here. Also note that an equivalent scenario will occur if the uncertainties were in both Type I and Type II errors. Thus, **Table 8** should be corrected to that given in **Table 9**, where, as can be seen, $E[\theta | w = w_A]$ remains constant at $(1/6 \times 2) + (5/6 \times 1.5) = 19/12$ independently of the value of ϕ , and hence the value of $1 - \phi = P(w_A)$. This means the SPRT will have constant losses that do not change between observations, and so we simply proceed as before.

Table 5. Average costs from the simulation described in Section 2.4.

	Average cost
Classical	3.66
Adaptive	2.88

Table 6. Assumed joint distribution between w and θ .

	θ_1	θ_2
w_0	1 (0.2)	1 (0.2)
w_A	2 (0.1)	1.5 (0.5)

Table 7. Implied conditional probabilities.

	θ_1	θ_2
$P(\theta w = w_0)$	0.5	0.5
$P(\theta w = w_A)$	1/6	5/6

Table 8. Initial loss table in the case of independence.

	d_0	d_A
w_0	0	2
w_A	$E[\theta]$	0

Table 9. Loss table in the case of statistical dependence.

	d_0	d_A
w_0	0	2
w_A	$E[\theta w = w_A]$	0

3. Unknown Observation Cost

Now suppose that the costs of making a Type I (c_1) or Type II (c_2) error are known. This means that if we were to implement an immediate decision the expected loss will be unchanged from the classical setting. However, we assume the observation cost θ is uncertain but subject to some prior distribution $f(\theta)$ and some specified data likelihood, in which case the expected loss of making a decision after observation will have to take into account not only the uncertainty concerning the information we may receive in relation to the true state of nature, but also the uncertainty in the additional cost of having taken a further observation.

If we take the expected value of θ , $E[\theta]$ as the observation cost, then we can determine bounds $[\phi_{L_{mit}}, \phi_{U_{mit}}]$ on values of ϕ within which we should seek additional data before implementing a decision. The expected risk profile $\rho(\phi)$ (expected loss), as a function of ϕ would then be:

$$\rho(\phi) = \begin{cases} c_1\phi & \text{if } \phi \leq \phi_{L_{mit}} \\ c(\phi) + E[\theta] & \text{if } \phi_{L_{mit}} < \phi \leq \phi_{U_{mit}} \\ c_2(1-\phi) & \text{if } \phi_{U_{mit}} < \phi \end{cases} \quad (29)$$

where $c(\phi)$ is a concave (or linear) function of ϕ determined by the data generating mechanism. Then, for each possible information statement i we may receive (where here i contains both the information concerning the true state of nature and any information we gain concerning θ the cost of sampling), we can determine a posterior distribution on θ and updated expected value $E[\theta|i]$. With this we continue the SPRT leading to updated intervals $[\phi_{L|i}, \phi_{U|i}]$ which if $\phi' = \phi|i = P(w_0|i)$ does not fall within, would result in our now taking an immediate decision. The updated risk table would now have form:

$$\rho(\phi') = \begin{cases} c_1\phi' & \text{if } \phi' \leq \phi_{L|i} \\ c'(\phi') + E[\theta|i] & \text{if } \phi_{L|i} < \phi' \leq \phi_{U|i} \\ c_2(1-\phi') & \text{if } \phi_{U|i} < \phi' \end{cases} \quad (30)$$

Here c' is another concave (or linear) function in ϕ' .

As the information i we may receive is currently unknown, we take the expectation of Equation (30). Subtracting this from Equation (29) (the expected risk without learning information) we obtain the expected value of that information, which can be thought of as the most we would be willing to pay for it in advance of seeing it. This should now be subtracted from $E[\theta]$, the original expected observation cost, to obtain what we would use as the adaptive information cost for the adaptive SPRT. Note that this value will be a function of ϕ . A classical SPRT is then performed with this adaptive observation cost until the true cost has been learned, at which point the test continues with the cost uncertainty removed, *i.e.*, in the classical way.

Remark. The expected value of information is zero for any $\phi \notin [\phi_L, \phi_U]$ (the bounds on ϕ for which we would take further samples) and also for any ϕ that is always contained in $[\phi_L, \phi_U]$.

Numerical Example

As a toy example to aid in clarification of the above, suppose we are testing the efficacy of a drug and are certain of the costs incurred in making a Type I or Type II error (say 2 and 4 units respectively). Assume, however, that we have little experience in running clinical trials (our observation costs) and are not sure if it will be easy and cheap to organise ($\theta = 0.1$) or relatively expensive ($\theta = 0.25$). Prior beliefs are that it is more likely to be cheap so that $P(\theta = 0.1) = 0.6$. Also suppose that the probability a bad drug passes the clinical trial is 0.5 whilst the probability that a drug that works passes is 0.8.

As we begin testing of the first drug we determine how to modify the SPRT procedure to take into account this uncertainty. Interest lies in the expected value of information of the observation cost, and we assume that the information will be of a perfect nature (namely remove all uncertainties). Noting that $E[\theta] = 0.16$, the risk profile, without information, is:

$$\rho(\phi) = \begin{cases} 2\phi & \text{if } \phi \leq 8/15 \\ 0.96 + 0.2\phi & \text{if } 8/15 < \phi \leq 76/105. \\ 4(1 - \phi) & \text{if } 76/105 < \phi \end{cases} \tag{31}$$

So if $\phi \in [8/15, 76/105]$ we take a further observation and hence also determine the true value of θ . This leads to two possible further risk profiles depending on if we learn $\theta = 0.1$ or $\theta = 0.25$.

For $\theta = 0.1$:

$$\rho(\phi) = \begin{cases} 2\phi & \text{if } \phi \leq 1/2 \\ 0.9 + 0.2\phi & \text{if } 1/2 < \phi \leq 31/42. \\ 4(1 - \phi) & \text{if } 31/42 < \phi \end{cases} \tag{32}$$

For $\theta = 0.25$:

$$\rho(\phi) = \begin{cases} 2\phi & \text{if } \phi \leq 7/12 \\ 1.05 + 0.2\phi & \text{if } 7/12 < \phi \leq 59/84. \\ 4(1 - \phi) & \text{if } 59/84 < \phi \end{cases} \tag{33}$$

Recalling the prior on θ is such that $P(\theta = 0.1) = 0.6$ leads to an expected risk after learning information of:

$$\rho(\phi) = \begin{cases} 2\phi & \text{if } \phi \leq 1/2 \\ 0.54 + 0.92\phi & \text{if } 1/2 < \phi \leq 7/12 \\ 0.96 + 0.2\phi & \text{if } 7/12 < \phi \leq 59/84. \\ 2.14 - 1.48\phi & \text{if } 59/84 < \phi \leq 31/42 \\ 4(1 - \phi) & \text{if } 31/42 < \phi \end{cases} \tag{34}$$

Subtracting Equation (34) (expected risk with knowledge of θ) from Equation (31) provides the expected value of perfect information for the observation cost:

$$g(\phi) = \begin{cases} 0 & \text{if } \phi \leq 1/2 \\ 1.08\phi - 0.54 & \text{if } 1/2 < \phi \leq 8/15 \\ 0.42 - 0.72\phi & \text{if } 8/15 < \phi \leq 7/12 \\ 0 & \text{if } 7/12 < \phi \leq 59/84 \\ 1.68\phi - 1.18 & \text{if } 59/84 < \phi \leq 76/105 \\ 1.86 - 2.52\phi & \text{if } 76/105 < \phi \leq 31/42 \\ 0 & \text{if } 31/42 < \phi \end{cases} \quad (35)$$

A plot of Equation (35) is provided in **Figure 4**. Note that the areas where the expected value of information is zero are where the decision rule is the same regardless of the information concerning the cost of sampling, agreeing with our earlier remark, and that the expected value of sampling information increases to be maximal where we are currently indifferent between making an immediate decision or taking further samples. With this to hand, we would continue by performing the SPRT as if we had an observation cost of $E[\theta] - g(\phi) = 0.16 - g(\phi)$, and if we do take an observation we learn the true value of θ and continue the SPRT with this knowledge.

4. Conclusions

In this paper we have considered the generalisation of SPRTs from a classical to adaptive utility setting where preferences or associated costs are not assumed fully known but are instead learned through experience or by funding additional information through survey or trial etc. Both unknown cost of Type I/II error was examined before subsequently considering the effect of uncertain observation cost.

Both perfect and noisy information were discussed, where we demonstrated the methods of quantifying the value for such information and numerical examples were

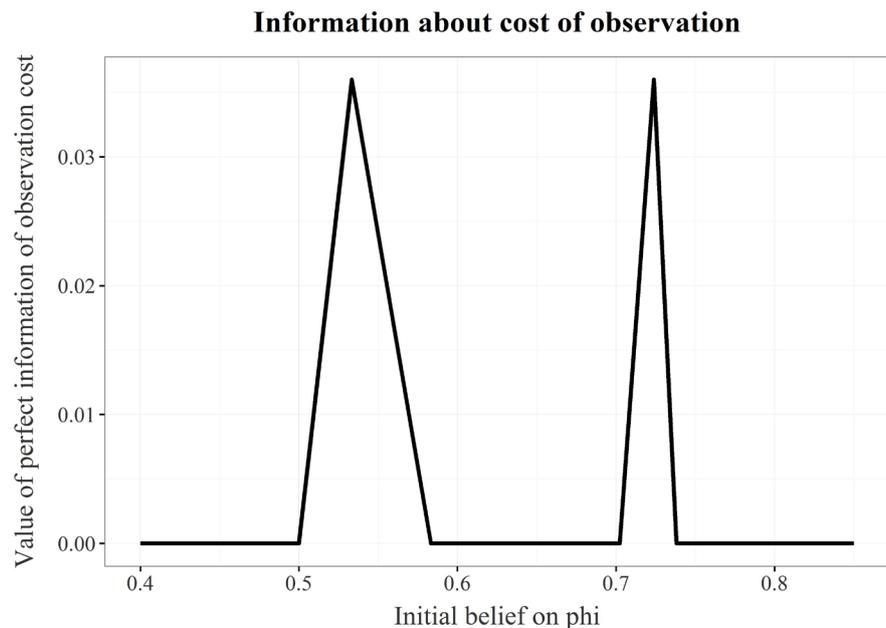


Figure 4. A plot of the expected value of information in Example 3.1 given by Equation (35).

provided to demonstrate the theory. Statistical dependence between the parameter and the state of nature was also considered and shown to not influence results. The numerical simulation indicated the enhanced performance by formally treating uncertainties and opportunities to learn within a SPRT in comparison to the somewhat easier modeling assumption of equating uncertainties in costs to their expected values.

References

- [1] Wald, A. (1945) Sequential Tests of Statistical Hypotheses. *Annals of Mathematical Statistics*, **16**, 117-186. <http://dx.doi.org/10.1214/aoms/1177731118>
- [2] Wald, A. and Wolfowitz, J. (1948) Optimum Character of the Sequential Probability Ratio Test. *Annals of Mathematical Statistics*, **19**, 326-339. <http://dx.doi.org/10.1214/aoms/1177730197>
- [3] Gross, K. and Humenik, K. (1991) Sequential Probability Ratio Test for Nuclear Plant Component Surveillance. *Nuclear Technology*, **93**, 131-137.
- [4] Kulldorf, M., Davis, R.L., Kolczak, M., Lewis, E., Lieu, T. and Platt, R. (2011) A Maximised Sequential Probability Ratio Test for Drug and Vaccine Safety Surveillance. *Sequential Analysis*, **30**, 58-70. <http://dx.doi.org/10.1080/07474946.2011.539924>
- [5] Spiegelhalter, D., Grigg, O., Kinsman, R. and Treasure, T. (2003) Risk-Adjusted Sequential Probability Ratio Tests: Applications to Bristol, Shipman and Adult Cardiac Surgery. *International Journal for Quality in Health Care*, **15**, 7-13. <http://dx.doi.org/10.1093/intqhc/15.1.7>
- [6] Spray, J.A. and Reckase, M.D. (1996) Comparison of SPRT and Sequential Bayes Procedures for Classifying Examinees into two Categories using a Computerized Test. *Journal of Educational and Behavioral Statistics*, **21**, 405-414. <http://dx.doi.org/10.3102/10769986021004405>
- [7] Corsini, G., DalleMese, E., Marchetti, G. and Verrazzani, L. (1985) Design of the SPRT for Radar Target Detection. *IEE Proceedings F Communications, Radar and Signal Processing*, **132**, 139-148. <http://dx.doi.org/10.1049/ip-f-1.1985.0035>
- [8] Fay, M., Kim, H. and Hachey, M. (2007) On Using Truncated Sequential Probability Ratio Test Boundaries for Monte Carlo Implementation of Hypothesis Tests. *Journal of Computational and Graphical Statistics*, **16**, 946-967. <http://dx.doi.org/10.1198/106186007X257025>
- [9] Malladi, D. and Speyer, J. (1999) A Generalized Shirayev Sequential Probability Ratio Test for Change Detection and Isolation. *IEEE Transactions on Automatic Control*, **44**, 1522-1534. <http://dx.doi.org/10.1109/9.780416>
- [10] Matas, J. and Chum, O. (2005) Randomized Ransac with Sequential Probability Ratio Test. *Tenth IEEE International Conference on Computer Vision (ICCV05)*, **1**, 1727-1732. <http://dx.doi.org/10.1109/ICCV.2005.198>
- [11] Cyert, R. and DeGroot, M. (1975) Adaptive Economic Models. Academic Press.
- [12] Boutilier, C. (2003) On the Foundations of Expected Utility. *Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence (IJCAI-03)*, Acapulco, 9-15 August 2003, 285-290.
- [13] Chajewska, U., Koller, D. and Parr, R. (2000) Making Rational Decisions Using Adaptive Utility Elicitation. *Proceedings of the Seventeenth National Conference on Artificial Intelligence and Twelfth Conference on Innovative Applications of Artificial Intelligence*, Austin, Texas, 30 July-3 August 2000, 363-369.
- [14] Houlding, B. and Coolen, F.P.A. (2011) Adaptive Utility and Trial Aversion. *Journal of Statistical Planning and Inference*, **141**, 734-747. <http://dx.doi.org/10.1016/j.jspi.2010.07.023>

- [15] DeGroot, M. (1984) Changes in Utility as Information. *Theory and Decision*, **17**, 287-303.
<http://dx.doi.org/10.1007/BF00132613>
- [16] R Core Team (2014) R: A Language and Environment for Statistical Computing. R Foundation for Statistical Computing, Vienna.



Scientific Research Publishing

Submit or recommend next manuscript to SCIRP and we will provide best service for you:

Accepting pre-submission inquiries through Email, Facebook, LinkedIn, Twitter, etc.

A wide selection of journals (inclusive of 9 subjects, more than 200 journals)

Providing 24-hour high-quality service

User-friendly online submission system

Fair and swift peer-review system

Efficient typesetting and proofreading procedure

Display of the result of downloads and visits, as well as the number of cited articles

Maximum dissemination of your research work

Submit your manuscript at: <http://papersubmission.scirp.org/>

Or contact ojs@scirp.org