PK8207 - Lecture memo

Jørn Vatn
Email: jorn.vatn@ntnu.no

Updated 2020-02-05

Numerical integration of stochastic processes

Introduction

A stochastic process $\{X(t), t \in \Theta\}$ is a collection of random variables. The set Θ is called the *index set* of the process. For each index t in Θ, $X(t)$ is called the *state* of the process at time t. In this memo we only consider the situation where Θ is a continuum, that is we have a continuous-time stochastic process.

Increments

An increment of a stochastic process is the difference between two random variables of the same stochastic process. We assume that the index set (Θ) represents time, and an increment is how much the stochastic process changes over a certain time period. Let $\{X(t), t \in \Theta\}$ be a stochastic process and consider two non-negative numbers $t_1 \in [0, \infty)$ and $t_2 \in [0, \infty)$ such that $t_1 \leq t_2$, then the difference $S = X(t_2) - X(t_1)$ is defined as the increment and is a stochastic variable.

Motivation

In maintenance modelling we often describe the state of a component with a state variable. Let t denote the time axis, and $Y(t)$ the state of an item at time t. Due to many common failure mechanisms such as wear, fatigue, corrosion it is reasonable to believe that the component deteriorates over time. The deterioration in a time period from t to $t + \Delta t$ is usually considered to be associated with some randomness. It is therefore natural to consider $\{Y(t), t \in \Theta\}$ as a stochastic process where:

$$Y(t + \Delta t) = Y(t) + \Delta Y \tag{1}$$

and where ΔY is the deterioration in the interval from t to $t + \Delta t$. Typically ΔY depends on the time t, the interval length Δt and the state $Y(t)$ at time t.

Failure

In the following we assume that high deterioration corresponds to a high value of $Y(t)$. Further we assume that there is a deterioration limit L corresponding to an unacceptable deterioration level. We now define the failure time of the stochastic process by:

$$T = \inf\{h : Y(h) \geq L\} \tag{2}$$

Although not realistic, we often assume that L is fixed and known.

Positive increments

Physically there are arguments supporting that increments are positive, since a component can not improve unless it is maintained in one way or another. This calls for positive increments. If we have positive increments it follows:

$$F_T(t) = \Pr(T \leq t) = \Pr(Y(t) \geq L) \tag{3}$$

To obtain the mean time to failure we may use:

$$\mathrm{E}(T) = \int_0^\infty (1 - F_T(t))dt = \int_0^\infty \Pr(Y(t) \leq L)dt \tag{4}$$

Note that we so far only consider a system which is not maintained, and just runs to failure.

The gamma process

Background: X is said to be gamma distributed with shape parameter α, and rate parameter λ if the PDF is given by:

$$\mathrm{Ga}(x|\alpha,\lambda) = \alpha^\lambda x^{\alpha-1} e^{-\lambda x}/\Gamma(\alpha) \tag{5}$$

Let $Y(t)$ be the state of the process at time t. $Y(t)$ follows a stationary gamma process if:

- $Y(0) = 0$
- $Y(s) - Y(t) \sim \mathrm{Ga}([s-t]v, u), s > t$
- $Y(t)$ has independent increments

We observe that the increment only depends on the length of the time interval, not the actual time.

Assume that the component fails as soon as the state exceeds the value L. Let T denote the time to failure. It then follows:

$$F_T(t) = \Pr(T \leq t) = \Pr(Y(t) \geq L) = \Gamma(vt, Lu)/\Gamma(vt) \tag{6}$$

where $\Gamma(a,x)$ is the incomplete gamma function.

"Overshooting" in the gamma process

One should expect that $\mathrm{E}(T) = Lu/v$.

Exercise

Show by numerical integration that $\mathrm{E}(T) \approx Lu/v + 1/(2v)$. Also show that $\mathrm{Var}(T) \approx Lu/v^2 - 1/(12v^2)$. Hint: $\mathrm{E}(T) = \int_0^\infty R(t)dt$, and $\mathrm{E}(T^2) = \int_0^\infty 2tR(t)dt$.

This phenomenon is denoted "overshooting". Actually the gamma process can be seen as a process with an infinite countable number of small jumps. Just before the threshold we may jump slightly "above" the limit. This is called "overshooting". This means that at the time of the failure, $Y(t) > L$, hence it takes "a bit longer" to reach the limit than one should "expect".

Non-stationary gamma process

The gamma process could be extended to a non-stationary process by letting the shape parameter be a function of time, i.e., $v(t)$ is the shape function, and we have:

- $Y(0) = 0$
- $Y(s) - Y(t) \sim \mathrm{Ga}(v(s) - v(t), u), s > t$
- $Y(t)$ has independent increments

The CDF now reads

$$F_T(t) = \Pr(T \le t) = \Pr(Y(t) \ge L) = \Gamma(v(t), Lu)/\Gamma(v(t)) \tag{7}$$

Note that the expected increment (deterioration) in the period from time t to s is $[v(s) - v(t)]/u$. In maintenance modelling it is common to use a form for the shape function like $v(t) = at^b$, where $b > 1$. If this is the case, the expected increment in a fixed length interval increases with time. This means that the deterioration goes faster and faster as the component is ageing. It should be mentioned that it is the age that determines the deterioration, not the state in terms of the deterioration level.

Positive and negative increments

Although we from physical arguments would believe that increments are always non-negative, we often observe negative increments. This points to the fact that what we measure is not necessary the real underlying deterioration level. What we could do would be to establish two stochastic processes, $\{X(t), t \in \Theta\}$ and $\{Y(t), t \in \Theta\}$ where $X(t)$ is the actual state at time t and $Y(t)$ is what we measure or calculate at time t. This idea is not pursued here, instead we will for the time being focus on only $\{Y(t), t \in \Theta\}$, and then allow increments also to be negative.

Brownian motion and the Wiener process

Brownian motion is the random motion of particles in a fluid where collisions between particles result in a rather chaotic behaviour. The phenomenon was first described by Robert Brown in 1827. In mathematics, Brownian motion is described by a continuous-time stochastic process named the Wiener process.

The Wiener process $\{W(t), t \in \Theta\}$ is characterized by:

- $W(0) = 0$
- $\{W(t)\}$ is almost surely continuous
- $\{W(t)\}$ has independent increments
- $W(s) - W(t) \sim \mathcal{N}(0, s-t)$ (for $0 \leq s \leq t$)

where $\mathcal{N}(\mu, \sigma^2)$ denotes the normal distribution with expected value μ and variance σ^2.

The Wiener process defined above is fluctuating around zero. A related stochastic process is defined by:

$$Y(t) = \mu t + \sigma W(t) \tag{8}$$

This process is called a Wiener process with drift μ and infinitesimal variance σ^2.

For the Wiener process with drift we have:

$$Y(t) - Y(s) \sim \mathcal{N}(\mu(s-t), \sigma^2(t-s)) \tag{9}$$

Note that since the increments could be negative, the relation $F_T(t) = \Pr(T \leq t) = \Pr(Y(t) \geq L)$ *does not hold* for the Wiener process. However, it may be shown that the time to failure, i.e., the time to the first passage of the limit L is inverse-Gauss distributed. That is:

$$F_T(t) = \Phi\left(\frac{\sqrt{\lambda}}{\nu}\sqrt{t} - \sqrt{\lambda}\frac{1}{\sqrt{t}}\right) + \Phi\left(-\frac{\sqrt{\lambda}}{\nu}\sqrt{t} - \sqrt{\lambda}\frac{1}{\sqrt{t}}\right)e^{2\lambda/\nu} \tag{10}$$

and

$$\mathrm{E}(T) = \nu \tag{11}$$

$$\mathrm{Var}(T) = \nu^3/\lambda \tag{12}$$

where $\nu = L/\mu$ and $\lambda = L^2/\sigma^2$. In terms of the original parameters this means that we have $\mathrm{E}(T) = L/\mu$ and $\mathrm{Var}(T) = L\sigma^2/\mu^3$.

Now, assume that we at time t observe $Y(t) = y_t$. Let $\tilde{T}_t$ be the time from t until a failure occurs. It then follows that $\tilde{T}_t$ is inverse-Gauss distributed with parameters $\tilde{\nu}_t = (L - y_t)/\mu$ and $\tilde{\lambda}_t = (L - y_t)^2/\sigma^2$.

Later on we will use the notation remaining useful lifetime (RUL) rather than $\tilde{T}_t$.

Using Monte Carlo simulation to obtain the distribution of T

Since we know that T is inverse-Gauss distributed it seems meaningless to find the distribution by Monte Carlo simulation (MCS). However, we will later see that we do not have an easy expression for T in many cases, and MCS might be an easy way out. A pseudo code for the simulation is given below:

```
nSim = 1000
dt = some small time step, e.g., 1/1000 of E(T)
mu = some value
sigma = some value
L = some value
For n = 1 To nSim
  Y = 0
  T = 0
  Do While Y < L
    Y = Y + rndNormal(mu*dt,sigma*sqr(dt))
    T = T + dt
  Loop
  Tvector(n) = T
Next n
makeHistogram(Tvector)
...we are done
```

where `rndNormal()` returns a pseudo normal distributed variable with the given mean and standard deviation.

Integrating the Wiener process to find the first passage CDF

The use of MCS is usually not recommended if we are able to apply numerical approaches for obtaining e.g., the PDF or CDF of a quantity of interest.

Let $Y(t)$ denote the state (deterioration level) of an item at time t, $t \geq 0$, and let $f(y|t)$ be the probability density function (PDF) of the deterioration level at time t. In a small time interval from t to $t+dt$ it is assumed that the item deteriorates, where the deterioration is a random quantity, say $S_{y,t,dt}$, i.e.,

$$Y(t+dt) = Y(t) + S_{y,t,dt} \tag{13}$$

Now let $g(s|y,t,dt)$ be the PDF of $S_{y,t,dt}$ in a small time period from t to $t+dt$, given that the degradation level at time t is y.

In the Wiener process the increments are independent both on time and state, so therefore we will here simplify and use the notation $g(s)$, where we also assume that we have determined the integration step length dt.

Now, if the PDF of the degradation level at time t is known, i.e., $f(y|t)$, the law of total probability gives the PDF of the deterioration level at time $t+dt$:

$$f(y|t+dt) = \int_{-\infty}^{\infty} f(y-s|t)g(s)ds \tag{14}$$

Since we assume a failure occurs the first time the deterioration level exceeds the threshold L, we need a modified expression to use in the numerical integration:

$$f(y|t+dt) = \int_{y-L}^{\infty} f(y-s|t)g(s)ds \tag{15}$$

If we are able to integrate Equation (15) numerically, we can also obtain at any point of time:

$$F_T(t) = \Pr(T \le t) = 1 - \int_{-\infty}^{L} f(y|t)dy \tag{16}$$

In a numerical implementation it is convenient to store $f(y|t)$ in an array which is denoted `f`. Indexing starts from a negative number so low that we never go below that value. The upper value could be e.g., $n = 1000$. The "points" are now defined such that `f(i)` $= f(y_i|t)$, where $y_i = idy$ and $dy = L/n$ is the interval length. To find the numerical solution for each step Equation (15) is now rewritten:

$$f(y_i|t+dt) = \sum_j \int_{(j-1)dy}^{jdy} f(y_i-s|t)g(s)ds \tag{17}$$

where we in the sum only include "valid" values.

Note that when s runs from $(j-1)dy$ to jdy, the argument in $f(y_i-s)$ runs from $(i-j+1)dy$ down to $(i-j)dy$. If $f(y)$ is approximated by the linear function $f(y) = a'y + b'$ in this interval, this yields $f(y_i - s) = -a's + (b' + a'y_i) = as + b$. The parameters a' and b' forming the linear interpolation are calculated from the values `f(i-j)` and `f(i-j+1)` obtained from the previous iteration of Equation (15).

This means that we approximate $f(y_i - s|t)$ with a linear function on the form $as + b$ and $g(s)$ is the probability density function of the normal distribution.

In many situations we are interested in calculating the "truncated expectation" $\int_{-\infty}^{\beta} xf(x)dx$. This is also the case here where we need $\int_{\alpha}^{\beta} sg(s)ds$. For the normal distribution the following result may be used:

If X is normally distributed with parameters μ and σ then:

$$\int_{-\infty}^{\beta} xf(x)dx = \mu\Phi\left(\frac{\beta-\mu}{\sigma}\right) - \sigma\phi\left(\frac{\beta-\mu}{\sigma}\right) \tag{18}$$

where $\Phi()$ and $\phi()$ are the CDF and PDF for the standard normal distribution respectively. Since efficient numerical routines exist for $\Phi()$ and $\phi()$ it is should be rather easy to implement Equation (17).

Exercise

Assume a Wiener process with $L = 100$, $\mu = 1$ and $\sigma = 0.3$. Implement the numerical approach indicated above to find CDF for the first passage time (failure time). Compare with the inverse-Gauss analytical result.

Increments depending on time and/or deterioration level

In the Wiener process with constant drift the increments are independent on both current time and current state. Some mechanisms may develop with time. For example we could imagine that the corrosion rate increases with time because the steel has been exposed to more and more corrosive medium. Further Paris' law (also known as the Paris-Erdogan equation) is a crack growth equation stating that crack growth during one load cycle depends on the crack length. Therefore the assumptions behind the Wiener process is not matching our physical understanding of the failure mechanism.

In the modelling it is therefore required to investigate relevant failure mechanism and propose a model for the increments (deterioration) based on our physical understanding where focus on time and deterioration level are important to consider. The geometric Brownian motion is a stochastic process taking the deterioration level into account when the next time interval is considered.

Geometric Brownian motion

A stochastic process $\{S(t)\}$ is said to follow a GBM if it satisfies the following stochastic differential equation (SDE):

$$dS(t) = \mu S(t)dt + \sigma S(t)dW(t) \tag{19}$$

where $W(t)$ is a normalized Wiener process and the notation d is denoting the increment in a small time interval of lengty dt. μ is "the percentage drift" and σ is "the percentage volatility". Note the slightly different notation here. We use $dS(t)$ rather than $S_{t,dt}$ as used in Equation (13) for the increment. The term $\mu S(t)dt$ is used to model deterministic trends, while the term $\sigma S(t)dW(t)$ is used to model the stochastic behaviour. The model has gain popularity in mathematical finance to model stock prices. The fact that deterioration in this model is proportional to the deterioration level also makes the model appropriate for many physical failure mechanisms.

Note that if we start with a perfect system, i.e., $S(0) = 0$ there will be no development of this process! This corresponds to a situation where a perfect piece of material will not suffer fatigue cracks! In order to use the GMB we need to assume that the initial state is positive. Further, if the volatility is large compared to the drift, the process might go "backwards", and being absorbed at zero. In fact, if $\mu > 0.5\sigma^2$ it seems like the process never reaches zero. This has to be investigated further! Find some references!

There is no known easy way to find the CDF for the first passage time. But we may rewrite Equation (15):

$$f(y|t+dt) = \int_{y-L}^{\infty} f(y-s|t)g(s|y-s)ds \tag{20}$$

where $g(s|y-s)$ is the PDF of the increment in a small time interval of length dt when the sate of the process at time t is $y-s$. From $dS(t) = \mu S(t)dt + \sigma S(t)dW(t)$ it follows that the PDF is the normal density function with mean value $\mu(y-s)dt$ and standard deviation $\sigma(y-s)\sqrt{dt}$.

Exercise

Write a computer program to calculate the CDF of the first passage time. Compare the result with what you get from MCS.

RUL = Remaining Useful Lifetime

There exist many definitions of RUL, framing something like RUL is *the useful life left on an asset at a particular time of operation* or RUL of an asset or system is *the length from the current time to the end of the useful life*. For definitions it is required to define what is meant by *useful*. A system may still function, but it is considered not useful any more because the risk of complete breakdown is considered high, the performance is low etc. To be formal, we introduce S_L as a set of unacceptable states of an item or a system. In many cases we may define an explicit threshold, say L describing some failure limit, or "stopping rule" for the operation of the item. Further we define $Y(t)$ as the condition, or health indicator of the system at time t. RUL is now defined as a stochastic variable:

$$\text{RUL}(t_j) = \inf\{h : Y(t_j+h) \in S_L | y(s) \notin S_L, 0 \le s \le t_j\} \tag{21}$$

Often we consider S_L as a fixed set corresponding to an explicit failure limit for the health indicator $Y(t)$. In other cases S_L could be seen as a "moving" set of states. For example we might use the vibration level measured in decibel as a stopping rule for operation of a rotating machinery. On the other hand, the vibration level would not necessarily be the best predictive health indicator. For example the kurtosis of the signal is often considered as a very good predictor. Hence, there will be no fixed set of unacceptable state for the kurtosis which corresponds to the decibel limit. In such a case S_L needs to be considered as a stochastic set of states. In a simplified situation with only a limit, say L, we may model this L as a stochastic variable.

Note that $\text{RUL}(t_j)$ is a stochastic variable. The terms 'prediction' and 'estimation' is therefore not straight forward. In statistics we usually use the term 'estimation' in relation to model *parameters*, hence we should avoid

the term 'RUL estimation'. Also 'RUL prediction' is a challenging word since there is no unambiguous interpretation. To be precise we rather define:

$$\Pr(\text{RUL}(t_j) \leq \tau) = F_{\text{RUL}(t_j)}(\tau) \tag{22}$$

If we use the term 'RUL prediction' we mean $\text{E}(\text{RUL}(t_j))$, i.e., the expected remaining useful lifetime of the item.

When assessing $\Pr(\text{RUL}(t_j) \leq \tau)$ or 'predicting RUL' we need to treat the following:

- Information regarding the operational condition up to time t_j
- The historical values of the health indicator, $y(s), 0 \leq s \leq t_j$
- Future operational conditions.

In most cases we assume some Markovian properties, i.e., it is only the current value of the health indicator, $y(t_j)$, that matters. Historical and future operational conditions will be ignored. This will be the case for the Wiener process and the geometric Brownian motion.

We will, however, make the following comment: The uncertainty in the remaining useful life is caused by *aleatory* uncertainty, *epistemic* uncertainty and *model* uncertainty. Aleatory uncertainty is the variability, e.g., described by σ in the above mentioned models. Epistemic uncertainty relates to the lack of knowledge regarding the value of the model parameters, e.g., μ, σ and S_L. Model uncertainty relates to the fact that the applied model, e.g., a Wiener process might be rather inappropriate.

In theoretical work and in exercises we often only consider aleatory uncertainty and take for granted that we know the model parameters and the model. In reality, when observing the health indicator of an item, we might consider to update our knowledge regarding the model parameters. This calls for *adaptive* methods. To be discussed later.

Bayesian approaches

In adaptive procedures we often update the model parameters, here μ and σ^2 as we get evidence for the running process. To interpret this, assume that we have some prior believe about the true underlying μ and σ^2 before we start observing a new system for which we would like to make decision based on the condition. This prior believe could be established from analysis of similar systems. When we get evidence from the actual system under consideration we update this prior believe into what is denoted a posterior believe. In the following we present some basic results from the Bayesian approach that applies for the Gaussian distribution, see, e.g., [1] for proofs and further explanations. For mathematical convenience it is common to work with the

precision, $\eta = 1/\sigma^2$ rather than variance. We thus assume that future deviations, say X_i's are $\sim \mathcal{N}(\mu, \eta^{-1})$. To specify our prior distribution we use a conjugated prior denoted the normal-Gamma distribution:

$$\mathrm{NG}(\mu, \eta | \mu_0, \kappa_0, \alpha_0, \beta_0) \stackrel{\mathrm{Def}}{=} \mathcal{N}\left(\mu | \mu_0, (\kappa_0 \eta)^{-1}\right) \mathrm{Ga}(\eta | \alpha_0, \beta_0) \tag{23}$$

This unusual way to describe a distribution is explained as follows. The parameter μ is unknown, and our prior knowledge regarding this parameter is described by a normal distribution with mean value μ_0 and precision given as the product $\kappa_0\eta$. μ_0 and κ_0 are numerical values, whereas η is the unknown precision of the observations. Then next, the unknown precision η is described by a gamma distribution with the parameters α_0 and η_0. The interpretation of κ_0 is then a precision factor relative to the precision in the X_i's. Note that μ is normally distributed conditional on the value of η. The marginal distribution of μ is however not the normal distribution but the student-t distribution.

To assess the numerical values for $\mu_0, \kappa_0, \alpha_0$ and β_0 we my proceed as follows. Assume that we have observed deterioration from say m systems. These systems are consider to be a representative sample of actual systems we will observe in the future. For each system, i, we estimate μ_i and $\eta_i = {\sigma_i}^{-2}$ by standard average and square sum formulas. We then let $\mu_0 = \frac{1}{m}\sum \mu_i$ be the average of the mean values. A preliminary precision figure is calculated by $p_0 = m/\sum(\mu_i - \mu_0)^2$. From the η_i's we use standard MLE to fit a gamma distribution with parameters α_0 and β_0. Finally we set $\kappa_0 = p_0\beta_0/(\alpha_0 - 1)$. The argument for $\alpha_0 - 1$ rather than the intuitive α_0 follows from the marginal distribution for μ which is a student-t distribution.

Now assume we start observing the process, i.e., we get values for the day to day deterioration, say $x_1, x_2, \ldots, x_n$. The posterior distribution is now also normal-Gamma distributed with parameters $\mu_n, \kappa_n, \alpha_n$ and β_n where

$$\begin{aligned} \mu_n &= \frac{\kappa_0\mu_0 + n\bar{x}}{\kappa_0 + n} \\ \kappa_n &= \kappa_0 + n \\ \alpha_n &= \alpha_0 + n/2 \\ \beta_n &= \beta_0 + \tfrac{1}{2}\sum_{i=1}^{n}(x_i - \bar{x})^2 + \frac{\kappa_0 n(\bar{x} - \mu_0)^2}{2(\kappa_0 + n)} \end{aligned} \tag{24}$$

This means that we get better and better knowledge regarding the true underlying parameters μ and σ. Given μ and σ we can find the PDF or CDF of the RUL. Then, by integrating over all possible values μ and σ with the posterior normal-Gamma distribution we may obtain the unconditional PDF or CDF as we go along, i.e., get more and more data.

As an example, consider the Wiener process with drift. Assume that the deterioration level at time t is y_t. Further assume that we have updated the posterior hyper parameters $\mu_n, \kappa_n, \alpha_n$ and β_n according to Equation (24).

$$\Pr(\mathrm{RUL}(t|y_t,\mu_n,\kappa_n,\alpha_n,\beta_n) \le \tau) =$$
$$\int_0^\infty \int_{-\infty}^\infty F_{\mathrm{IG}}(\tau|\nu,\lambda) f_{\mathrm{N}}\left(\mu|\mu_n,(\kappa_n\eta)^{-1}\right) f_{\mathrm{Ga}}(\eta|\alpha_n,\beta_n) d\mu, d\eta \tag{25}$$

where $F_{\mathrm{IG}}(\tau|\nu,\lambda)$ is the CDF of the inverse Gauss distribution given by Equation (10), where $\nu = (L-y_t)/\mu$ and $\lambda = (L-y_t)^2\eta$. $f_{\mathrm{N}}\left(\mu|\mu_n,(\kappa_n\eta)^{-1}\right)$ and $f_{\mathrm{Ga}}(\eta|\alpha_n,\beta_n)$ are the PDFs of the normal and gamma distributions respectively.

If we also consider the threshold L to be uncertain, we could integrate the solution in Equation (25) over the corresponding PDF, $f_L(l)$.

How to perform discretization?

Equation (25) involves several integrals where we integrate over the probability distributions. The integrals could be replaced by sums if we perform discretization of the distribution.

[2] propose a standard approach for approximating a continuous distribution by a discrete distribution where (i) the outcome region is divided into intervals, then (ii) for each interval a representing point is chosen, and (iii) a probability, p_i is assigned to each point. Usually the intervals are found by dividing the outcome region into k equally probable intervals, where the representative point is the mean of the corresponding interval, and the assigned probability is $1/k$. When generating a limited number of discrete outcomes, some statistical properties should be specified. It is common to include the first four central moments as properties, i.e., the moments for the discrete stochastic variable should be as close as possible to the moments of the variable we are approximating.

[3] have used such an approach if the stochastic variable is normally distributed with mean value μ and standard deviation σ. Odd numbers of points are used, and Table 1 presents standardized distances from the mid point. The midpoint is always given by μ. If a k point approximation is required, the two nearest points to the mid point are given by the (left and right) distance $d_{k,1}\sigma$ from the midpoint at μ, the second nearest points to the mid point are given by the distance $d_{k,2}\sigma$ from the midpoint at μ and so forth.

Table 1: Standardized distances from mid point

k = # of scenarios	$d_{k,1}$	$d_{k,2}$	$d_{k,3}$	$d_{k,4}$	$d_{k,5}$
3	1.22474				
5	0.87889	1.31436			
7	0.16787	0.49042	1.79758		
9	0.21902	0.60872	0.67431	1.90442	
11	0.26459	0.43883	0.70498	0.89051	1.98681

References

[1] DeGroot, M. (1970) Optimal Statistical Decisions. McGraw-Hill.

[2] Høyland, K. and Wallace, SW. (2001) Generating Scenario Trees for Multistage Decision Problems. *Management Science*, 47(2), pp 295 - 307

[3] Bjerke, MM. and Bykvist, HR. (2018) *Managing Supply Uncertainty in the Operational Production Planning in a Whitefish Value Chain - A Stochastic Programming Approach*. Master thesis. NTNU - Norwegian University of Science and Technology.