

Truncated Normal Collocation Chasing the One-Armed Man

John Burkardt
Department of Scientific Computing
Florida State University

...

Max Gunzburger Group Seminar
07 September 2013

...

[http://people.sc.fsu.edu/~jburkardt/presentations/
truncated_normal_2013_fsu.pdf](http://people.sc.fsu.edu/~jburkardt/presentations/truncated_normal_2013_fsu.pdf)



INTRO:

In a prehistoric TV series called "The Fugitive", Dr Richard Kimble, falsely accused of murdering his wife, searched for the one-armed man who was the real killer.

The One-Armed Man



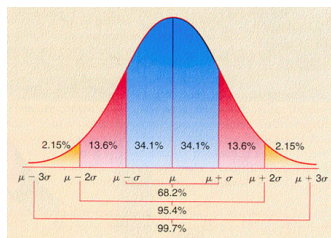
There are no murders scheduled for today, but I will be lopping off the right arm, the left arm, or both arms, of the standard normal



- **The Strange Case of the Abnormal Normal**
- Can You Describe the Suspect?
- I Have a Cunning Plan
- A Matter of Moments
- The Summing Up



ABNORMAL: The Normal Probability Distribution



I was recently invited to Ajou University, Korea, at the invitation of Professor Hyung-Chun Lee. One morning during my visit, he asked me if I could set up a collocation procedure for the **truncated normal distribution**, and I said, "*Sure, no problem!*"

The truncated normal distribution is a simple modification to our familiar friend, the normal distribution.

The normal distribution allows a natural description of how some measurable quantities (height, income, number of sick days) have



ABNORMAL: The Truncated Normal

Mathematical models are idealizations, and have their limitations. If we think the normal distribution is a good model of height distribution, then strictly speaking, we are admitting the possibility (small, but not zero!) of people who are as tall as 60 or 1000 feet - or negative 200 feet, for that matter.

This discrepancy could be a problem if we are doing a simulation, for instance. Then we treat the mathematical distribution as physical reality, we sample it, and we “believe” whatever comes out of the process. If we create a 1-inch person, then we are now stuck dealing with a physically meaningless but computationally real object.



ABNORMAL: The Truncated Normal

Sometimes these 1 inch people can actually cause the computation to crash, or to produce meaningless results.

In our research group, a commonly studied problem involves the simulation of the permeability function $a(\omega, x)$ related to groundwater flow, arising in the equation:

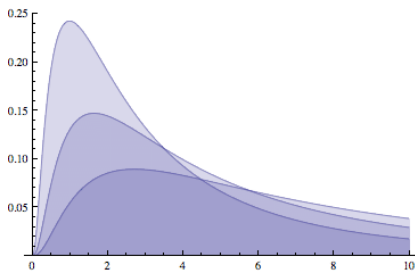
$$\nabla \cdot (a(\omega, x) \nabla u(x)) = f(x)$$

One of the requirements for existence and uniqueness of a solution requires positivity and boundedness:

$$0 < a_{min} \leq a(\omega, x) \leq a_{max} < \infty$$



ABNORMAL: The Log-Normal Probability Distribution



We can use the **log-normal** probability density function to describe $a(\omega, x)$, because if we assume that

$$\log(a()) = \alpha \sim N(\mu, \sigma)$$

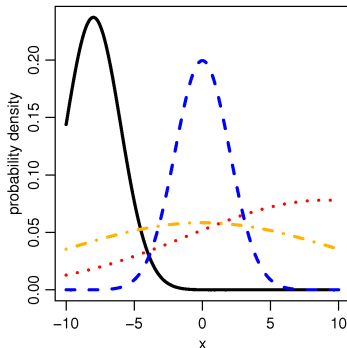
then we are guaranteed that

$$0 < e^\alpha = a()$$

so our simulation will never select a negative value for $a()$, and $a()$ is described by a mathematically tractable and plausible formula.



ABNORMAL: The Truncated Normal Distribution



Another choice uses the normal distribution, but restricts the PDF by defining an upper maximum, or a lower minimum or both.

This gives us a great deal of flexibility, but if we take this approach we must be able to produce the same kind of mathematical information that is already well known for standard PDF's.



ABNORMAL: Parameters

A standard PDF is described by parameters. To describe a truncated PDF, we start by describing the normal PDF that is the “parent”, that is, we must supply the values of $\bar{\mu}$ and $\bar{\sigma}$.

Then we must list the values that define the truncation interval $[a, b]$, and we should allow for all four possibilities:

- non-truncated normal, $(-\infty, +\infty)$
- lower truncated normal, $[a, +\infty)$
- upper truncated normal, $(-\infty, b]$
- doubly truncated normal, $[a, b]$

So a truncated normal is described by $\bar{\mu}, \bar{\sigma}, a, b$, and note that $\bar{\mu}$ is **not** the mean μ of the truncated normal, and $\bar{\sigma}^2$ is **not** its variance σ^2 !



- The Strange Case of the Abnormal Normal
- **Can You Describe the Suspect?**
- I Have a Cunning Plan
- A Matter of Moments
- The Summing Up



Suppose we are given the vital statistics of a truncated normal distribution, namely, the values of $\bar{\mu}, \bar{\sigma}, a, b$.

There are six standard chores we should be able to do:

- 1 evaluate $\text{pdf}(\bar{\mu}, \bar{\sigma}, a, b; x)$;
- 2 evaluate $\text{cdf}(\bar{\mu}, \bar{\sigma}, a, b; x) = \int_a^x \text{pdf}(\bar{\mu}, \bar{\sigma}, a, b; \xi) d\xi$;
- 3 solve $C = \text{cdf}(\bar{\mu}, \bar{\sigma}, a, b; x)$ for the value of x ;
- 4 uniformly sample $\text{pdf}(\bar{\mu}, \bar{\sigma}, a, b; x)$.
- 5 evaluate $\mu = \int_a^b x \text{pdf}(\bar{\mu}, \bar{\sigma}, a, b; x) dx$;
- 6 evaluate $\sigma^2 = \int_a^b (x - \mu)^2 \text{pdf}(\bar{\mu}, \bar{\sigma}, a, b; x) dx$;



DESCRIBE: PDF

Task 1: evaluate $\psi(x) = \text{pdf}(\bar{\mu}, \bar{\sigma}, a, b; x)$;

Denote by $\phi(\xi)$ and $\Phi(\xi)$ the PDF and CDF for the standard normal distribution with mean 0 and variance 1.

To adjust for the effects of the nonstandard mean and variance, define:

$$\xi(x) = \frac{x - \bar{\mu}}{\bar{\sigma}}$$

Then we can normalize the PDF over the nontruncated range:

$$\psi(x) = \begin{cases} 0 & \text{if } x < a \\ \frac{\phi(\xi(x))}{\Phi(\xi(b)) - \Phi(\xi(a))} & \text{if } a \leq x \leq b \\ 0 & \text{if } b < x \end{cases}$$

The quantity

$$S = \Phi(\xi(b)) - \Phi(\xi(a))$$

is a scale factor which we will need in order to normalize our integrals involving the truncated normal PDF.



Task 2: evaluate

$$\Psi(x) = \text{cdf}(\bar{\mu}, \bar{\sigma}, a, b; x) = \int_a^x \text{pdf}(\bar{\mu}, \bar{\sigma}, a, b; \xi) d\xi.$$

If we think about it, the CDF has to be 0 at a and 1 at b , and in between it's simply integrating the scaled PDF. So the formula has to be:

$$\Psi(x) = \begin{cases} 0 & \text{if } x < a \\ \frac{\Phi(\xi(x)) - \Phi(\xi(a))}{s} & \text{if } a \leq x \leq b \\ 1 & \text{if } b < x \end{cases}$$



Task 3: solve $C = \text{cdf}(\bar{\mu}, \bar{\sigma}, a, b; x)$ for the value of x ;

We can almost solve the previous CDF equation:

$$\Phi(\xi(x)) = S * C + \Phi(\xi(a))$$

Presumably, we can invert the CDF of the normal distribution:

$$\xi(x) = \Phi^{-1}(S * C + \Phi(\xi(a)))$$

and so the corresponding value of x is simply:

$$x = \bar{\mu} + \bar{\sigma} \xi$$



DESCRIBE: SAMPLE

Task 4: uniformly sample pdf($\bar{\mu}, \bar{\sigma}, a, b; x$);

Luckily, solving Task 3 makes this task trivial. To sample from the distribution, simply generate a uniform random value $C \in [0, 1]$. Regard C as the value of the CDF at a point x , and compute x .

Values chosen in this way are uniformly distributed with respect to the truncated normal distribution.

By the way, here's an alternative way that's correct, but generally bad because it can reject lots of data:

Sample a value from the normal distribution with mean $\bar{\mu}$ and variance $\bar{\sigma}^2$, but if the value is less than a or greater than b , try again.



DESCRIBE: MEAN

Task 5: evaluate $\mu = \int_a^b x \text{pdf}(\bar{\mu}, \bar{\sigma}, a, b; x) dx$;

A formula is available for this task:

$$\mu = \bar{\mu} + \frac{\phi(\alpha) - \phi(\beta)}{S} \bar{\sigma}$$

where

$$\alpha = \frac{a - \bar{\mu}}{\bar{\sigma}}; \quad \beta = \frac{b - \bar{\mu}}{\bar{\sigma}};$$

For the lower truncated normal, we have $b = \infty$ so $\phi(\beta) = 0$; we can handle the upper truncated normal similarly, and we see that the formula will also be correct for the normal distribution as well, simply returning $\mu = \bar{\mu}$.



DESCRIBE: VARIANCE

Task 6: evaluate $\sigma^2 = \int_a^b (x - \mu)^2 \text{pdf}(\bar{\mu}, \bar{\sigma}, a, b; x) dx$;

A formula is also available for this task:

$$\sigma^2 = \bar{\sigma}^2 \left(1 + \frac{\alpha\phi(\alpha) - \beta\phi(\beta)}{S} - \left(\frac{\phi(\alpha) - \phi(\beta)}{S} \right)^2 \right)$$

Again, the formula is written for the doubly-truncated case, but can easily be used for the lower, upper, and non-truncated distributions as well.



DESCRIBE: Checking the Facts

I set up a library called **truncated_normal** which included code for the six tasks, for all four truncation possibilities,

Now I figured I needed some confidence in my formulas before moving on, so I constructed a set of tests.

One simple test is to start with a value of X , compute its CDF, then compute `invCDF` and see if we get back to X .

The second test was to do a simulation. That is, use the `SAMPLE` function to compute, say, 10,000 sample values of a distribution, compute the sample mean and variance, and compare them to the `MEAN` and `VARIANCE` functions.

After banging on the code, the tests behaved, and I felt much better.

http://people.sc.fsu.edu/~jburkardt/m_src/truncated_normal/truncated_normal.html



DESCRIBE: You're Not Done!

I kept busy all morning long working out these details about the truncated normal distribution and trying to program, document, and test them.

That afternoon, Professor Lee came back into the office and asked if I had been able to complete the collocation task.

"Well...", I said hesitantly, *"I can PDF, CDF, invCDF, SAMPLE, MEAN and VARIANCE."*

"What about the collocation?" he asked.

"Actually," I said, *"I might need one more day..."*

As it turned out, I ended up working on this problem for another month.



Truncated Normal Collocation

- The Strange Case of the Abnormal Normal
- Can You Describe the Suspect?
- **I Have a Cunning Plan**
- A Matter of Moments
- The Summing Up



PLAN: What Do We Need?

We weren't pursuing the truncated normal distribution for its own sake - what we were really after was the ability to do collocation.

We wanted to estimate the expected value of “quantities of interest” associated with a system of partial differential equations that included stochastic input terms.

The stochastic input terms were going to be modeled by truncated normal distributions, so that they behaved like normal variables, but over a truncated range.

In order to estimate the quantities of interest, a collocation procedure selects many test values of the input terms, weighted by their probabilities, and computes an average that is really a multidimensional integral.



PLAN: What Do We Need?

If every input term is controlled by a truncated normal distribution, then the crucial tool we need is a sequence of quadrature rules, of increasing accuracy, for that distribution.

A quadrature rule for the truncated normal distribution is a set of n points x_i and weights w_i for which we make the estimate

$$\frac{1}{S\sqrt{2\pi\sigma^2}} \int_a^b f(x) e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx \approx \sum_{i=1}^n w_i f(x_i)$$

We are usually looking for a quadrature rule of Gaussian type, so that the n -point rule will integrate precisely any function $f(x)$ which is a polynomial of degree $2n - 1$ or less.

Analytic formulas are known for special weight functions; otherwise, a famous paper by Golub and Welsch shows how to construct a matrix whose eigendecomposition will produce the desired rule.



PLAN: What Do We Need?

The most common algorithm described by Golub and Welsch assumes that the user knows a family of polynomials $\phi_i(x)$, $i = 0, \dots$ which are orthogonal with respect to the PDF.

For the normal distribution, this family is known as the Hermite polynomials, whose first elements are:

$$H_0(x) = 1$$

$$H_1(x) = x$$

$$H_2(x) = x^2 - 1$$

$$H_3(x) = x^3 - 3x$$

Such families of orthogonal polynomials always satisfy a three term recurrence relationship, of the form:

$$\phi_{i+1}(x) = \alpha_i x \phi_i(x) + \beta_i \phi_{i-1}(x)$$

and for the Hermite polynomials, $\alpha_i = 1$ and $\beta_i = -i$.



PLAN: What Do We Need?

In cases where we can determine the recurrence coefficients, the Golub-Welsch procedure forms what is known as the **Jacobi matrix**:

$$J = \begin{pmatrix} \alpha_0 & \sqrt{\beta_1} & 0 & \dots & 0 \\ \sqrt{\beta_1} & \alpha_1 & \sqrt{\beta_2} & \dots & 0 \\ 0 & \sqrt{\beta_2} & \alpha_2 & \dots & 0 \\ \dots & \dots & \dots & \dots & \dots \\ 0 & 0 & 0 & \dots & \alpha_{n-1} \end{pmatrix}$$

The eigenvalues give us the quadrature points, and the weights are computed from the first components of the normalized eigenvectors.

Because this matrix is symmetric and tridiagonal, the computation is not difficult - except that I can't figure out an orthogonal polynomial family for the truncated normal distribution!

Fortunately, Golub and Welsch described an alternative procedure, based on the moments of the PDF, and this, I think, I can do!



PLAN: The Moment Matrix Approach

The k -th moment of a PDF is simply the value of the integral

$$m_k = \int_a^b x^k \text{pdf}(x) dx$$

Then, to compute an n -point Gauss quadrature rule for the PDF, you construct the $(n+1) \times (n+1)$ **moment matrix**:

$$M = \begin{pmatrix} m_0 & m_1 & m_2 & \dots & m_n \\ m_1 & m_2 & m_3 & \dots & m_{n+1} \\ m_2 & m_3 & m_4 & \dots & m_{n+2} \\ \dots & \dots & \dots & \dots & \dots \\ m_n & m_{n+1} & m_{n+2} & \dots & m_{2n} \end{pmatrix}$$

We compute the upper Cholesky factorization $M = R'R$. From the entries of R it is possible to compute the vectors α and β needed to construct the Jacobi matrix J and hence to determine the points and weights of the quadrature rule.



PLAN: So Now What's Missing?

Now it seems like a quadrature rule is within our grasp, if only we can compute these moments. So we are asking for the values of integrals like

$$m_k = \frac{1}{S\sqrt{2\pi\sigma^2}} \int_a^b x^k e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx$$

I suppose I could compute one of these integrals with Mathematica, but does this solve my problem? I can hope that Mathematica can return a single formula as an answer, with a and b left as variables, but can I also get away with leaving k as a variable? Don't I have to separately treat the cases where $a = -\infty$ or $b = +\infty$?



PLAN: Mathematica Tries to Help

Before we try the truncated normal distribution, let's just see if Mathematica can tell us, once for all, a formula for the moments of the normal distribution itself. Even for the normal distribution, it's not easy to find information beyond the fourth moment. So here goes:

```
In[1] = 1/(S Sqrt[2 Pi s^2]
  Integrate [ x^k Exp[-(x - m)^2/(2 s^2)],
  {x, -Infinity, +Infinity}]
```

```
Out[1] = ConditionalExpression[(1/(S Sqrt[\[Pi]] s))
  2^(-(1/2) + 1/2 (-1 + k)) E^(-(m^2/(2 s^2))) (1/s^2)^(
  1/2 (-1 - k)) (-Sqrt[2] (-1 + (-1)^k) m Sqrt[1/s^2]
  Gamma[1 + k/2] Hypergeometric1F1[1 + k/2, 3/2, m^2/(
  2 s^2)] + (1 + (-1)^k) Gamma[(1 + k)/2] Hypergeometric1F1[
  1 + k)/2, 1/2, m^2/(2 s^2)]), Re[1/s^2] > 0 && Re[k] > -1
```

which was the kind of answer I was afraid I would get.



PLAN: Maybe the Internet isn't entirely bad

Since our library was closed for six weeks, in order to throw out most of the books, and hide the others, I decided to try to search on the internet for a useful formula for the moments of the normal distribution. Since the words **moment** and **normal** occur in many contexts, this was not so easy.

But I came across a wonderful page by John D. Cook, titled *General formula for normal moments*:

$$\begin{aligned} m_k = E\{x^k\} &= E\{(\sigma\xi + \mu)^k\} = \sum_{i=0}^k \binom{k}{i} E(\xi^i) \sigma^i \mu^{k-i} \\ &= \sum_{j=0}^{\lfloor k/2 \rfloor} \binom{k}{2j} (2j-1)!! \sigma^{2j} \mu^{k-2j} \end{aligned}$$

This is straightforward algebra and easy to program. So, if we can get the moment-based quadrature scheme written, we can first input the moments of the normal distribution as a test.



Truncated Normal Collocation

- The Strange Case of the Abnormal Normal
- Can You Describe the Suspect?
- I Have a Cunning Plan
- **A Matter of Moments**
- The Summing Up



MOMENTS: The Mysterious Phoebus J Dhrymes

Returning to the Internet, I ran across a reference in a paper to an unpublished note (not a paper) by Phoebus J Dhrymes, titled *Moments of Truncated (Normal) Distributions*. I was able to find a web site for this combination statistician and psychiatrist, and indeed, there was a 3-page discussion that included the formula:

$$E(x^k | x \leq b) = \sum_{i=0}^k \binom{k}{i} \sigma^i \mu^{k-i} l_i$$

Recall that $\phi(x)$ and $\Phi(x)$ are the PDF and CDF, respectively, of the standard normal distribution and define:

$$\beta = \frac{b - \bar{\mu}}{\bar{\sigma}}$$

Then the quantity l_i satisfies the recursion:

$$l_0 = 1$$

$$l_1 = -\frac{\phi(\beta)}{\Phi(\beta)}$$



MOMENTS: Upper Truncated Normal Moments!

So now I had a formula for the moments of the upper truncated normal distribution! I programmed it, and as a test, I asked Mathematica to compute the corresponding integral for specific choices of the exponent, the normal parameters μ and σ , and the upper truncation limit b ;

For example, for $\bar{\mu} = 5$, $\bar{\sigma} = 1$, $b = 10$:

Order	Moment	Mathematica
0	1	1
1	5	5
2	26	26
3	140	140
4	777.997	777.997
5	4449.97	4449.97
6	26139.69	26139.67
7	157396.75	157396.71
8	969946.73	969946.45



MOMENTS: Lower Truncated Normal Moments!

I soon realized that the formula for upper truncated moments also gave me the formula for lower truncated moments, since, by a change of variable $y = -x$, we can get:

$$\begin{aligned}m_k &= \frac{1}{S\sqrt{2\pi\bar{\sigma}^2}} \int_a^\infty x^k e^{-\frac{(x-\bar{\mu})^2}{2\bar{\sigma}^2}} dx \\&= \frac{1}{S\sqrt{2\pi\bar{\sigma}^2}} \int_a^\infty (-y)^k e^{-\frac{(-y-\bar{\mu})^2}{2\bar{\sigma}^2}} dy \\&= \frac{(-1)^k}{S\sqrt{2\pi\bar{\sigma}^2}} \int_{-\infty}^{-a} y^k e^{-\frac{(y-\bar{\mu})^2}{2\bar{\sigma}^2}} dy\end{aligned}$$

or ± 1 times the k -th upper truncated normal moment for $-\bar{\mu}$ and $\bar{\sigma}$, with $-a$ as the upper limit.

So I had moment formulas for the normal, lower truncated, and upper truncated distributions, but nothing on the doubly truncated distribution!



MOMENTS: Double Truncated Normal Moments!

At last, I found a paper online that referenced Phoebus J Dhrymes, and stated that Dhrymes's result also implied a simple formula for moments of the doubly truncated normal distribution.

Define

$$\alpha = \frac{a - \bar{\mu}}{\bar{\sigma}}; \quad \beta = \frac{b - \bar{\mu}}{\bar{\sigma}}$$

Then (as before) we have

$$m_k = \sum_{i=0}^k \binom{k}{i} \bar{\sigma}^i \bar{\mu}^{k-i} l_i$$

where the quantity l_i satisfies the recursion:

$$l_0 = 1$$

$$l_1 = - \frac{\phi(\beta) - \phi(\alpha)}{\Phi(\beta) - \Phi(\alpha)}$$

$$l_i = - \frac{\beta^{i-1} \phi(\beta) - \alpha^{i-1} \phi(\alpha)}{\Phi(\beta) - \Phi(\alpha)} + (i-1) l_{i-2}$$



MOMENTS: Double Truncated Normal Moments!

Now that I had usable moment formulas for all four cases, I added a seventh “task”, to compute the k -th moment, to the **truncated_normal** library.

And I was now ready to consider the next step, which was to implement the moment formulation of the Golub-Welsch algorithm for computing quadrature rules.

http://people.sc.fsu.edu/~jburkardt/m_src/truncated_normal/truncated_normal.html



Truncated Normal Collocation

- The Strange Case of the Abnormal Normal
- Can You Describe the Suspect?
- I Have a Cunning Plan
- A Matter of Moments
- **The Summing Up**



SUMUP: One Step Beyond

We are now ready to try to compute an n -point quadrature rule using the moment-based version of the Golub Welsch algorithm - that is, once we write a program to implement the Golub Welsch algorithm!

The first step of the Golub Welsch algorithm requires us to form the order $n + 1$ moment matrix M , filling it with the values of moments m_0 through m_{2n} , which we just figured out how to compute.

The second step requires us to compute the upper Cholesky factor R such that $M = R'R$ - but it's not difficult to put together the code for this calculation either.

Golub and Welsch then supply formulas for extracting the vectors α and β from the information in R , and with these, we can construct the $n \times n$ symmetric tridiagonal Jacobi matrix J .

Now comes the tricky step - compute eigenvalues and eigenvectors



SUMUP: Jacobi will get me the eigenvalues

I want my code to be accessible in several languages. So I really want to write down an eigenvalue routine explicitly. Luckily, our symmetric matrix J is ideal for the Jacobi eigenvalue algorithm.

The basic idea of the Jacobi algorithm is to pre- and post-multiply the matrix by Jacobi rotation matrices that zero out the largest off-diagonal element. As this process is repeated, the matrix rapidly approximates a diagonal matrix, from which the eigenvalues can be read off.

Moreover, I had forgotten that this algorithm can also return the corresponding eigenvectors, which we need to have for the quadrature weights.

So my next task was to write a library called **jacobi_eigenvalue** which would allow me to test my implementation on some standard problems.



SUMUP: The Golub-Welsch Moment Algorithm Executes

Once the Jacobi eigenvalue algorithm was working, I had all the pieces needed to carry out a Golub-Welsch momentum algorithm.

Since I had never done this before, I first set up momentum calculations for Legendre, Laguerre, and Normal distributions, because I knew what the associated quadrature rules should be in these cases, and I caught some small errors this way.

Then I added in the momentum calculations for the truncated normal distribution, and was finally able to compute some example rules.

I called this hunk of software **quadmom** for *quadrature by moments*.

http://people.sc.fsu.edu/~jburkardt/m_src/quadmom/quadmom.html



SUMUP: An Interactive Rule Calculator

Using **quadmom** requires writing and compiling a calling program.

A more convenient approach is an executable program, **truncated_normal_rule**, that only asks for input:

- *option* 0/1/2/3 for none, lower, upper, double truncation;
- *n* the number of points in the rule;
- *mu*, the mean of the original normal distribution;
- *sigma* the standard deviation of the original normal distribution;
- *a* the left endpoint (for options 1 or 3);
- *b* the right endpoint (for options 2 or 3);
- *filename*, the root name of the output files.

http://people.sc.fsu.edu/~jburkardt/m_src/truncated_normal_rule/truncated_normal_rule.html



SUMUP: An Interactive Rule Calculator

To compute a doubly truncated quadrature rule of 10 points, with $\bar{\mu} = 0$ and $\bar{\sigma} = 1$, over the interval $[-3.0, +3.0]$, we write:

```
truncated_normal_rule 3 10 0.0 1.0 -3.0 +3.0 double10
```

For a lower truncated rule over $[-3.0, \infty)$, write:

```
truncated_normal_rule 1 10 0.0 1.0 -3.0 lower10
```

Dropping both limits, we get a non-truncated normal rule:

```
truncated_normal_rule 0 10 0.0 1.0 normal10
```



SUMUP: A Sample Rule Computation

The program writes the rule to three output files, containing the points, the weights, and the integration limits. Thus, the lower truncated normal rule we requested above would be stored in the following files:

lower10_x.txt	lower10_w.txt	lower10_r.txt
-2.83915	0.00279	-3.00000
-2.28008	0.02023	1.0E+30 <-- ' 'Inf'
-1.53530	0.09714	
-0.71994	0.25745	
0.12958	0.34188	
1.00659	0.21473	
1.91770	0.05925	
2.88043	0.00629	
3.93128	0.00019	
5.16662	0.8E-06	

(We actually write more digits in the real files)



SUMUP: Testing the Rule

Let's stick with the lower truncated normal distribution, with parameters $\bar{\mu} = 0$ and $\bar{\sigma} = 1$, over the interval $[-3.0, \infty)$, and estimate the following integral using an n -point rule.

$$Q = \frac{1}{S\sqrt{2\pi\bar{\sigma}^2}} \int_a^{+\infty} \sin(x) e^{-\frac{(x-\bar{\mu})^2}{2\bar{\sigma}^2}} dx$$

N	Estimated Q
1	0.004437820
2	-0.002956940
3	0.000399622
4	-0.000236540
5	-0.000173932
6	-0.000177684
7	-0.000177529
8	-0.000177534
9	-0.000177534
Mathematica	-0.000177531



CONCLUSION:

I hope I've given you an idea of the kinds of problems I look at, and how I go about trying to solve them,

The main moral I can give is that, in a scientific computing project, you are almost always “an infinite distance away” from your final, happy working code, and so you have to just pay very close attention to what you are doing right now, making sure you have understood what you need to do, how it fits into the big picture, and why you can demonstrate that it is correct.

CASE CLOSED

