# Tuesday's Exercises

Please direct comments and questions to
Scott Ferson (scott@ramas.com)
Applied Biomathematics

## Interval probability

1. Explore the Risk Calc scripts BLOOD, BIAS, SAFETY, LOGIC, TURNSIGNALS. (Enter "run blood".)

   Several of the scripts use confidence intervals about frequencies to represent the sampling uncertainty of the estimates. These intervals are used in calculations based on interval analysis. The guarantee about the results of such calculations is contingent on the additional assumption that these intervals are sure to contain the true frequencies.

2. Calculate the probability of tank rupture under pumping that assumes the interval inputs and makes *no assumption* about the dependencies among the events.

   You can make the calculation in Risk Calc with the following commands.

```
t = [4.5e-6, 5.5e-6]
k2 = [2.5e-5, 3.5e-5]
s = [0.5e-4, 1.5e-4]
k1 = [2.5e-5, 3.5e-5]
r = [0.5e-4, 1.5e-4]
s1 = [2.5e-5, 3.5e-5]
t | (k2 | (s & (s1 | (k1 | r))))
   [ 2.5e-05, 0.0001905]
```

   Thus, the answer is the interval $2.5\times10^{-5}$ to $1.9\times10^{-4}$, which encloses all the other intervals computed for the tank example, but is not vacuously wide.

   You can also compute the bounds yourself, without making use of Risk Calc's & and | operators to get the probabilities of the conjunction and disjunctions. The calculation is straightforward, although rather cumbersome. Keep in mind that

   $a$ & $b$=[max(0, $a$+$b$–1), min($a$,$b$)], and
   $a$ | $b$=[max($a$,$b$), min(1, $a$+$b$)],

   where $a = [a_1, a_2]$ and $b = [b_1, b_2]$ are interval probabilities for the two events, and

   $\min(a,b) = \min([a_1,a_2], [b_1,b_2]) = [\min(a_1,b_1), \min(a_2,b_2)]$
   $\max(a,b) = \max([a_1,a_2], [b_1,b_2]) = [\max(a_1,b_1), \max(a_2,b_2)]$,
   $a+b = [a_1,a_2] + [b_1,b_2] = [a_1+b_1, a_2+b_2]$,
   $a-1 = [a_1,a_2] - [1,1] = [a_1-1, a_2-1]$, and

   you can represent zero and one as the degenerate intervals [0,0] and [1,1].

3. Derive an algorithm to compute the probability that $n$ of $k$ events occur, given intervals for the probability of each event, assuming they're independent. Derive an analogous algorithm for the Fréchet case.

   This is a very difficult task. It may be useful to obtain conservative bounds on the resulting probability that are not best possible but sure to enclose the actual probabilities. There may also be some special cases in which the calculation becomes easier. Any suggestions would be welcome.

## Robust Bayes

1. Review the Risk Calc scripts TESTPOSITIVE and BAYES.

   The TESTPOSITIVE script emphasizes the adverse consequences of repeated parameters in an expression involving uncertainty. It is fortunate that Bayes' rule is easy to simplify to remove repeated parameters. Entering the command robust as a part of the BAYES script shows an illustration of the uniform distributions in the class of posteriors computed in exercise #3 below.

2. Sketch $[\underline{P}, \overline{P}]$ for {-,B,G,B,B,G,R,G,…,(35/100)R,…(341/1000)R}.

   The sequence of samples from Walley's bag of marbles is blue, green, blue, blue, green, red, green, …, 35 red out of 100 draws, …, 341 red out of 1000 draws. To compute interval bounds on the predictive probabilities for the event "the next marble drawn will be red" before each draw, we use the formula

   $$p = [n_j / (N + s), (n_j + s)/(N + s)]$$

   where $n_j$ is the running total number of red marbles that have been drawn so far, $N$ is the total number of marbles (of any color) that have been drawn, and $s$ is a positive constant. If, as Walley recommends, we let $s = 2$, we would get

| $N$ | $n_j$ | $p$ |
|---|---|---|
| 0 | 0 | [ 0, 1] |
| 1 | 0 | [ 0, 0.66667] |
| 2 | 0 | [ 0, 0.5] |
| 3 | 0 | [ 0, 0.4] |
| 4 | 0 | [ 0, 0.33334] |
| 5 | 0 | [ 0, 0.28572] |
| 6 | 1 | [ 0.125, 0.375] |
| 7 | 1 | [ 0.11111, 0.33334] |
| ⋮ | ⋮ | ⋮ |
| 100 | 35 | [ 0.34313, 0.36275] |
| ⋮ | ⋮ | ⋮ |
| 1000 | 341 | [ 0.34031, 0.34232] |

for the observed sequence of marble draws. If, however, we let $s = 1$, we'd get

| $N$ | $n_j$ | $p$ |
|---|---|---|
| 0 | 0 | [ 0, 1] |
| 1 | 0 | [ 0, 0.5] |
| 2 | 0 | [ 0, 0.33334] |
| 3 | 0 | [ 0, 0.25] |
| 4 | 0 | [ 0, 0.2] |
| 5 | 0 | [ 0, 0.16667] |
| 6 | 1 | [ 0.14285, 0.28572] |
| 7 | 1 | [ 0.125, 0.25] |
| ⋮ | ⋮ | ⋮ |
| 100 | 35 | [ 0.34653, 0.35644] |
| ⋮ | ⋮ | ⋮ |
| 1000 | 341 | [0.34065, 0.34166] |

Both sequences are graphed below. Notice that the convergence in the left graph is slower than in the right graph.

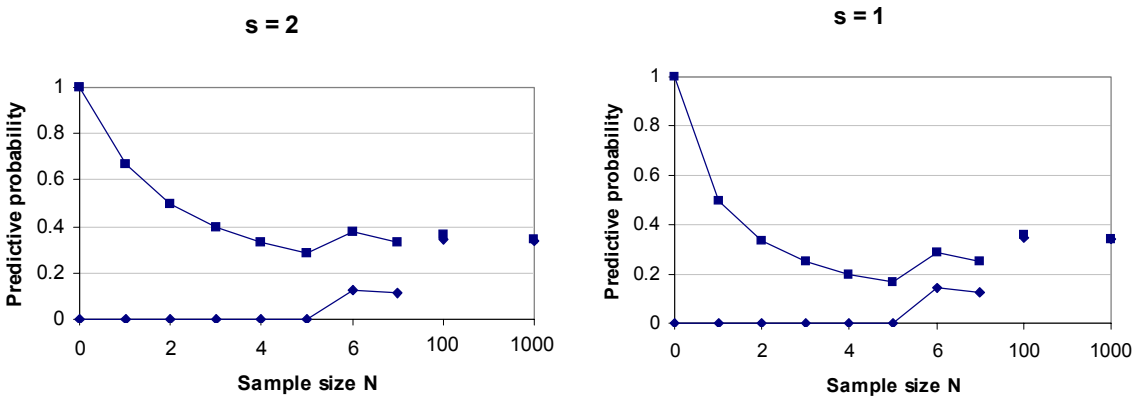


3. What can be said about the posterior if the prior is uniform($a$,$b$) and the likelihood function is an a non-zero constant for values of θ between $c$ and $d$ and zero elsewhere? Sketch the answer for $a = 2$, $b = 14$, $c = 5$, $d = 21$. What is the posterior if $a \in [1,3]$, $b \in [11,17]$, $c \in [4,6]$, $d \in [20,22]$?

The posterior distribution would be uniform(max($a$,$c$), min($b$,$d$)), that is, a uniform ranging from the greater of $a$ and c to the lesser of $b$ and $d$. In the first case, the posterior would be a uniform (flat density, straight line cumulative) between 5 and 14. The graphs below illustrate this case, in a density plot on the left and a cumulative probability plot on the right. In both, the posterior is shown in gray.

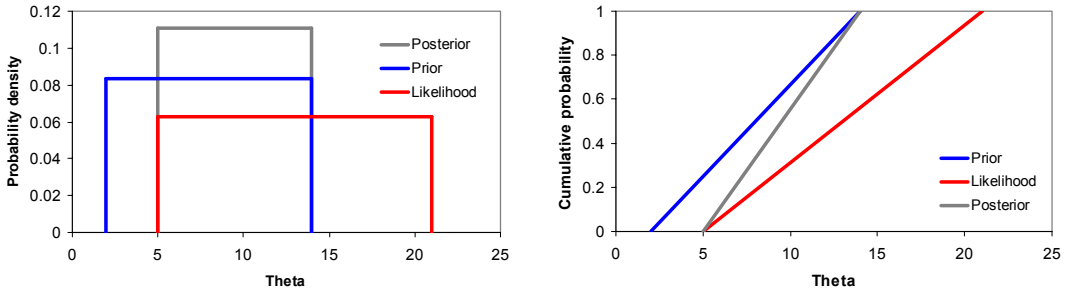


In the second case, the posterior would be the *class* of uniforms whose minima are within the interval [4,6] and whose maxima are within the interval [11,17]. The class of priors includes all distribution functions that are straight lines inside the blue region depicted in the graph below. The class of likelihood functions (normalized and cumulated for display on the same scale) consists of those functions corresponding to any straight line inside the red region in the graph, that is, any straight line that intersects theta axis between 4 and 6 and the top edge between 20 and 22. Any straight line that can be drawn inside the gray region in the graph representing a uniform distribution is a possible posterior from this analysis.

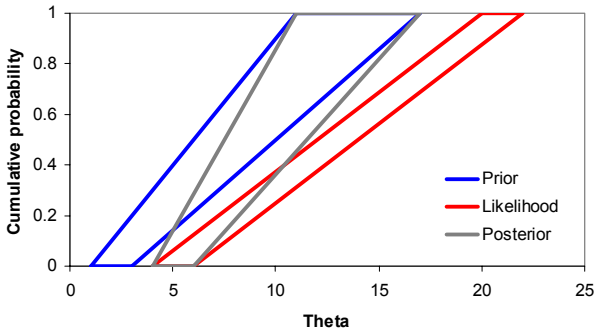


Notice that the formula uniform(max(*a*,*c*), min(*b*,*d*)) works even when *a*, *b*, *c* and *d* are intervals. You can make the calculations and get the (cumulative) graphs in Risk Calc by entering the following commands.

```
a = 2
b = 14
c = 5
d = 21
prior = uniform(a, b)
likelihood = uniform(c, d)
posterior = uniform(max(a,c), min(b,d))
posterior
```

**~uniform(range=[5,14], mean=9.5, var=6.75)**

```
a = [1,3]
b = [11,17]
c = [4,6]
d = [20,22]
```

prior = uniform(a, b)
likelihood = uniform(c, d)
posterior = uniform(max(a,c), min(b,d))
posterior
**~uniform(range=[4,17], mean=[7.5,11.5], var=[4,10.1])**

4. How should the interval bounds on a posterior arising from interval bounds on the prior *density* distribution and interval bounds on the likelihood function be normalized?

Suppose the bounds on the prior distribution are depicted in blue on the graph below and that the bounds on the likelihood function are depicted in red on the same graph. (You should check that the area below the lower blue bound is strictly less than unity and the area below the upper blue bound is strictly greater than unity, and, in this case, the area under the lower bound is 0.8 and under the upper bound is 2.5.)

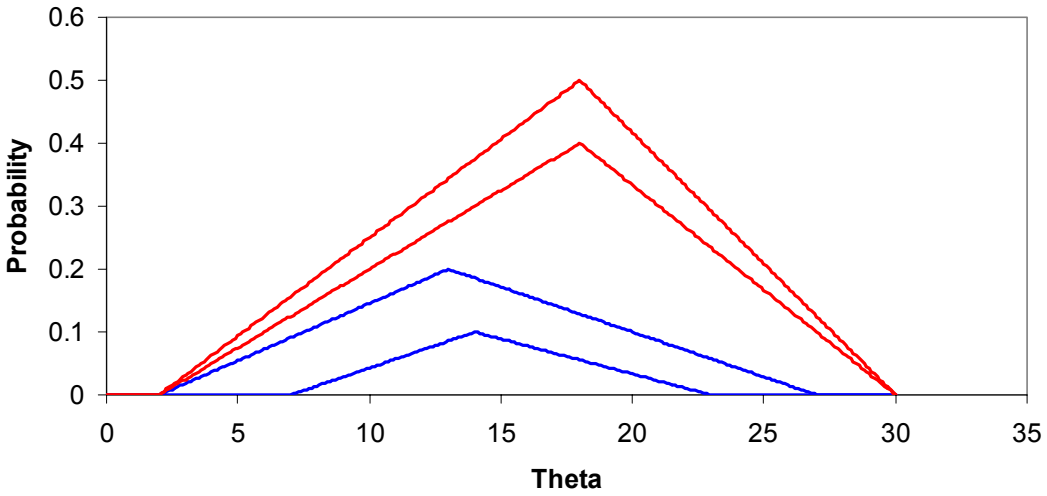


The first step toward computing the implied bounds on the posterior is to multiply the bounds on the prior with the bounds on the likelihood for every value of theta. For instance, at the value theta=16, the prior bounds [0.08, 0.16] and the likelihood bounds [0.35, 0.44] yield the product [0.028, 0.07] by interval multiplication. The bounds resulting from this process are depicted in gray in the graph below.

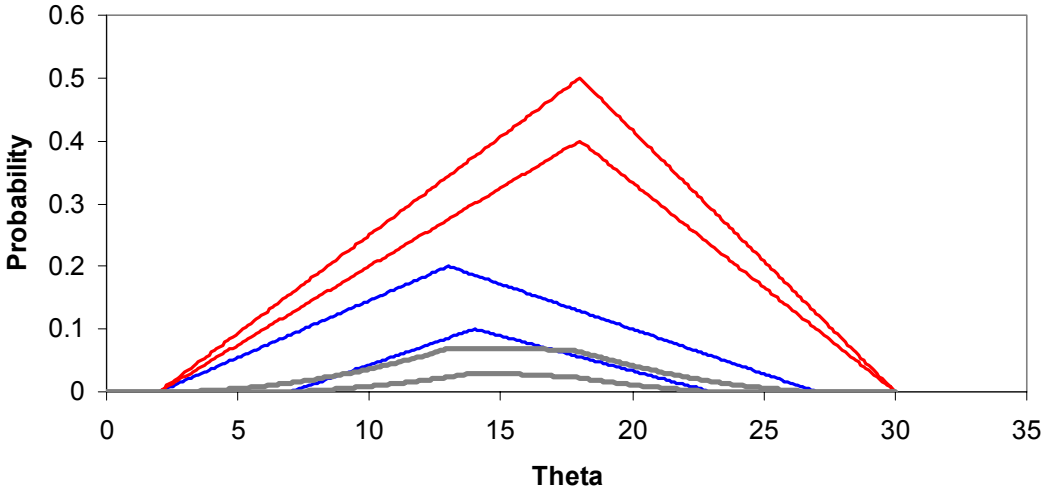


Any function that lies between these gray bounds is an un-normalized posterior distribution. The normalization that will make the function into a posterior is independent of that needed for other functions within the bounds that become posteriors. What are the bounds on all *normalized* posteriors? The upper bound is just the upper gray bound divided by the area under the lower gray bound. The lower bound is the lower gray bound divided by the area under the upper gray bound. The resulting bounds on the normalized priors are depicted in black in the graph below. These black bounds contain all posteriors that could result from this analysis, and they are the best possible bounds that do so, although it is not the case that any function lying within the bounds that has unit area will necessarily be one of the posteriors. To see why the black bounds in the graph below are the best possible bounds on the density of posterior distributions, consider extreme cases of un-normalized posteriors selected from within the gray bounds in the graph above. In particular, consider the function that traces the lower gray bound at every value of theta except the value 15, at which the function jumps discontinuously to the upper gray bound. Clearly, this function is a possible un-normalized posterior, given the bounds on the prior and the likelihood. How will it be normalized? Because its area is obviously just the same as the area of the lower gray bound, we divide by that area to obtain the normalized posterior. The value of this normalized posterior at the value of theta=15 is obviously the largest possible value of any posterior because it corresponds to the largest numerator and the smallest denominator possible in any normalization. Similar upper bounds on the normalized posterior can be obtained for any other value of theta in the same way. If the area beneath the lower gray bound is zero, then the densities of the posterior are unbounded above. A reciprocal argument explains why the lower bound on the posteriors is computed as the lower gray bound divided by the area of the upper gray bound.

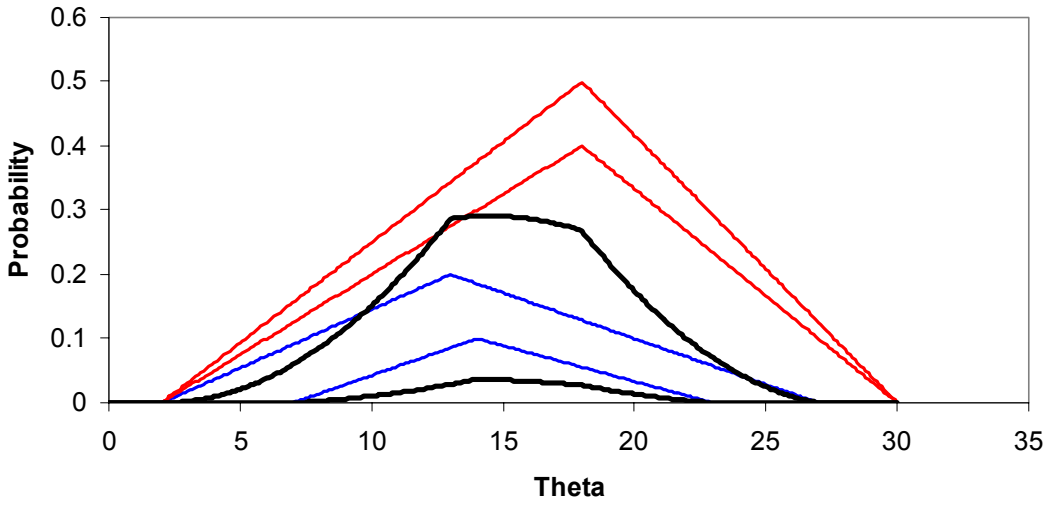


## Dempster-Shafer theory

1. Explore the script DIKE. Create Dempster-Shafer structures like "mixture( 0.2, [1,3], 0.5, [2,5], 0.3, [4,8])". Compute with them in the Listener window.

   The mixture function in Risk Calc does not require that the masses you specify for the focal elements add up to unity. If they don't, it will scale them so that they do. So, for example, the function "mixture(2, [1,3], 5, [2,5], 3, [4,8])" produces the same Dempster-Shafer structure as "mixture( 0.2, [1,3], 0.5, [2,5], 0.3, [4,8])".

   The script DIKE actually makes use of probability distributions rather than Dempster-Shafer structures for the significant wave height and offshore peak wave steepness. To replace those definitions so that the script matches the example given in the slides, enter the assignments

   ```
   Hs = mixture(1, [1,1.5],  2, [1.1,1.5],  2, [1.3,1.6],  2, [1.3,1.4],  1, [1.5,1.7]) * units('m')
   s_0p = mixture(1,[0.03,0.036],  9,[0.034,0.042],  9,[0.039, 0.04],  1,[0.045, 0.048])
   ```

   All you need to do is enter the above two lines, and then cursor up to the line defining *Z*, and press carriage return 9 times to execute the subsequent lines.

2. Compute the DS structure for *B*–*A*, and sketch the cumulative belief and plausibility functions. Concoct a numerical example of a logical operation such as AND or OR with DS structures.

   Yager's Cartesian product for the subtraction is below. Notice that the interval subtractions are anti-elementwise, so, for example, [2,8] – [1,3] is [–1,7], and not [1,5].

| $B - A$ independence | $A \in [1,3]$ $p_1 = 1/3$ | $A \in [2,4]$ $p_2 = 1/3$ | $A \in [3,5]$ $p_3 = 1/3$ |
|---|---|---|---|
| $B \in [2,8]$ $q_1 = 1/3$ | $B-A \in [-1,7]$ prob=1/9 | $B-A \in [-2,6]$ prob=1/9 | $B-A \in [-3,5]$ prob=1/9 |
| $B \in [6,10]$ $q_2 = 1/3$ | $B-A \in [3,9]$ prob=1/9 | $B-A \in [2,8]$ prob=1/9 | $B-A \in [1,7]$ prob=1/9 |
| $B \in [8,12]$ $q_3 = 1/3$ | $B-A \in [5,11]$ prob=1/9 | $B-A \in [4,10]$ prob=1/9 | $B-A \in [3,9]$ prob=1/9 |

Stacking up the intervals from this Cartesian product yields the CPF and CBF depicted in the graph below.

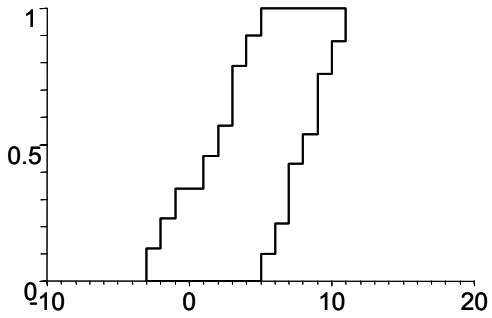


A numerical example involving logical operations on Dempster-Shafer structures requires that the operands be logical quantities. This means they must be dimensionless and range only between zero and one. This would be satisfied by the Dempster-Shafer structures

$C$ = {([0.1, 0.3], 1/3), ([0.2, 0.4], 1/3), ([0.3, 0.5], 1/3)} and
$D$ = {([0.1, 0.4], 1/3), ([0.3, 0.5], 1/3), ([0.4, 0.6], 1/3)},

which are the previous structures $A$ and $B$ scaled by division by 10 and 20 respectively. The Cartesian product for conjunction (&) with these operands is displayed below. Note that conjunction under independence corresponds to interval multiplication, so the table is easy to complete.

| $C \,\&\, D$ independence | $C \in [0.1, 0.3]$ $p_1 = 1/3$ | $C \in [0.2, 0.4]$ $p_2 = 1/3$ | $C \in [0.3, 0.5]$ $p_3 = 1/3$ |
|---|---|---|---|
| $D \in [0.1, 0.4]$ $q_1 = 1/3$ | $C\&D \in [0.01,0.12]$ prob=1/9 | $C\&D \in [0.02,0.16]$ prob=1/9 | $C\&D \in [0.03,0.2]$ prob=1/9 |
| $D \in [0.3, 0.5]$ $q_2 = 1/3$ | $C\&D \in [0.03,0.15]$ prob=1/9 | $C\&D \in [0.06,0.2]$ prob=1/9 | $C\&D \in [0.09,0.25]$ prob=1/9 |
| $D \in [0.4, 0.6]$ $q_3 = 1/3$ | $C\&D \in [0.04,0.18]$ prob=1/9 | $C\&D \in [0.08,0.24]$ prob=1/9 | $C\&D \in [0.12,0.3]$ prob=1/9 |

Stacking up the intervals from this Cartesian product yields the CPF and CBF depicted in the graph below.

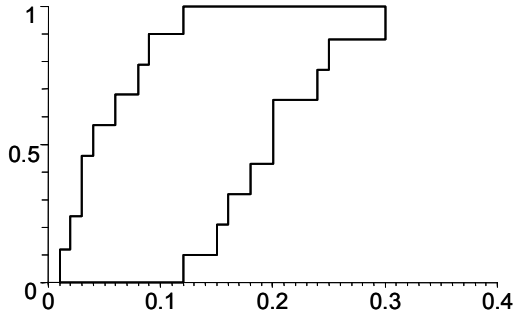


Had we chosen an example involving disjunction, the calculation would have required computing $C|D = [c_1,c_2] \mid [d_1,d_2] = 1 - (1 - [c_1,c_2])(1 - [d_1,d_2])$ for each matrix element.

The operations could have been done in Risk Calc with the following commands.

```
A = mixture(1/3, [1,3],   1/3, [2,4],   1/3, [3,5])
B = mixture(1/3, [2,8],   1/3, [6,10],   1/3, [8,12])
B |-| A
   ~(range=[-3,11],  mean=[1.3,8],  var=[0,22])
C = A / 10
D = B / 20
C |&| D
   ~(range=[0.01,0.3],  mean=[0.053,0.2],  var=[0,0.0131])
```

where the vertical bars around the minus sign and the ampersand indicate that the operation is to assume independence. The symbol for disjunction under independence is ||| (a vertical bar enclosed in vertical bars). The CPF and CBF would be graphed automatically.

3. How would you compute the Yager convolution assuming perfect dependence if the operands had unequal numbers of focal elements? What if the focal elements are not nicely ordered?

The calculation assuming perfect dependence presumes the Cartesian product creates a square matrix whose diagonal absorbs all the mass. If the Dempster-Shafer structure for $A$ has three focal elements and that for $B$ had only two, the matrix wouldn't be square. It can be made square by repartitioning the mass into the least common multiple of focal elements. For example, if

$A = \{([1,3], 1/3),\ ([2,4], 1/3),\ ([3,5], 1/3)\}$ and
$B = \{([2,8], 1/2),\ ([6,10], 1/2)\}$,

then the equivalent structures

$A = \{([1,3],1/6),([1,3],1/6),\ ([2,4],1/6),([2,4],1/6),\ ([3,5],1/6),([3,5],1/6)\}$

$$B = \{([2,8],1/6), ([2,8],1/6),([2,8],1/6), ([6,10],1/6),([6,10],1/6),([6,10],1/6)\}.$$

would allow the matrix to be square. The structures are equivalent to the ones above because they make precisely the same claims about where along the real line the mass resides. If one Dempster-Shafer structure had two focal elements and the other had four, then only the first would need to be repartitioned.

The calculation assuming perfect dependence also assumes that the focal elements of the Dempster-Shafer structures can be ordered from smallest to largest. The elements of each Dempster-Shafer structure can be reordered to facilitate this so long as some least-to-greatest ordering of the elements by their left bounds corresponds to a least-to-greatest ordering of the elements by their right bounds. This will not be possible, however, if some of the focal elements within a Dempster-Shafer structure are nested. If this is the case, then applying this algorithm for computing the convolution under perfect dependence would require that the Dempster-Shafer structure first be transformed into a related Dempster-Shafer structure that can be ordered. For instance, the Dempster-Shafer structure {([2,10], 1/2), ([6,8], 1/2)} would be transformed to {([2,8], 1/2), ([6,10], 1/2)}. These structures have the same CPF and CBF, but they are obviously not the same Dempster-Shafer structures. Numerical experiments suggest that this is a reasonable strategy, but it is an open question whether such transformations are legitimate, and, if not, how convolutions assuming perfect dependence for such cases could be computed. Your thoughts on this issue would be welcome.

## Probability bounds analysis

1. Explore the scripts ADDP, DEPEND, CONVOLVE, BOXES, MOMENTS, and CORRS.

   The ADDP and DEPEND scripts show how an interval and a probability distribution can be added together and how the convolution of distributions is influenced by the assumption that is made about their dependence. Note that an interval is entirely different from a uniform distribution and the two behave completely differently in calculations. Consider, for example, the following convolutions under independence, in which the expression U($a$,$b$) denotes a uniform probability distribution over the range of values between $a$ and $b$, and the expression [$a$, $b$] denotes an interval over the range of values between $a$ and $b$. Note that, although the range is always [3,6], the moments vary noticeably as do the shapes of the p-boxes.

U(1,2) |+| U(2,4)

**~trapezoidal(range=[3,6], mean=4.5, var=0.416667)**

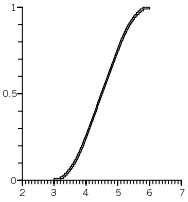

[1,2] |+| U(2,4)

**~(range=[3,6], mean=[4,5], var=[0.333,0.584])**

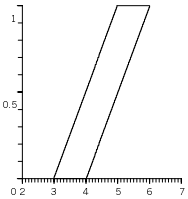

U(1,2) |+| [2,4]

**~(range=[3,6], mean=[3.5,5.5], var=[0.08,1.09])**

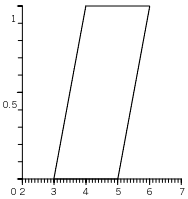

[1,2] |+| [2,4]

**[ 3, 6]**

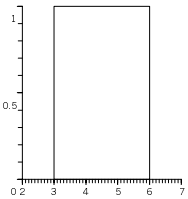

Similar calculations can be made with without any assumption about dependence between the operands (the Fréchet case).

U(1,2) + U(2,4)

**~(range=[3,6], mean=4.5, var=[0.08,0.75])**

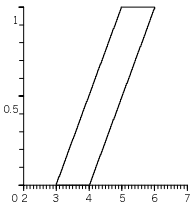

[1,2] + U(2,4)

**~(range=[3,6], mean=[4,5], var=[0.03,1.11])**

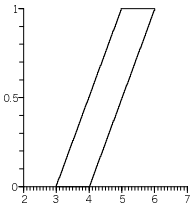


U(1,2) + [2,4]

**~(range=[3,6], mean=[3.5,5.5], var=[0,1.6])**

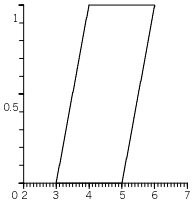


[1,2] + [2,4]

**[ 3, 6]**

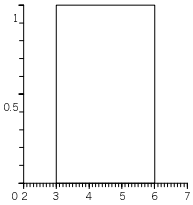


2. Type in the inputs and model for the mercury and PCB examples into Risk Calc. Do you get the same answers? Did you enter the units too?

The Risk Calc commands you would enter and the graphs they would produce are shown below.

```
// PCBs in ducks eaten by hunters
EF = minmaxmeanstddev(1, 52, 5.4, 10) * units('meal per year')
```

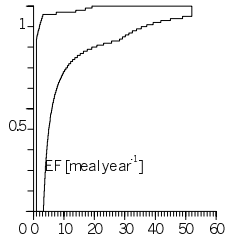


```
IR = minmaxmeanstddev (1.5, 675, 188, 113) * units('gram per meal')
```

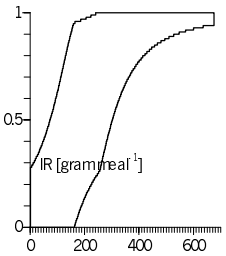


C = [7.1, 9.73] mg per kg

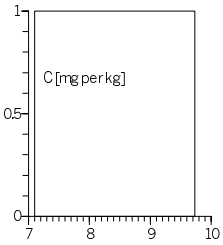


LOSS = 0

AT = 365.25 days per year

BWmale = lognormal(171 pounds,  30 pounds)

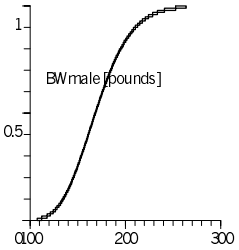


BWfemale = lognormal(145 pounds,  30 pounds)

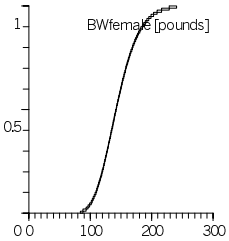


BW = mixture(BWfemale, BWmale)

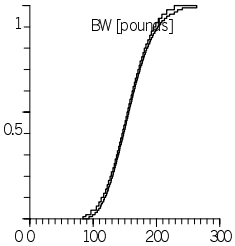


RfD = 0.00002 mg per kg per day

HQ = (EF |*| IR |*| C |*| (1 |-| LOSS)) |/| (AT |*| BW |*| RfD) + 0  //adding 0 cancels units

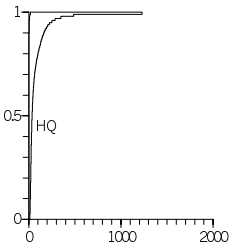


mean(HQ) //Mean hazard quotient

**[ 14, 21]**

sd(HQ) //Standard deviation

**[ 32, 47]**

median(HQ) //Median (50th percentile)

**[ 1, 36]**

cut(HQ,95%) //95th upper percentile

**[ 5.6, 2.3e+02]**

mean(HQ) //Range of possible values

**[ 14, 21]**

////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////////

//mercury in wild mink

FMR = lognormal ([90,120] Kcal per kg per day, [22,30] Kcal per kg per day)

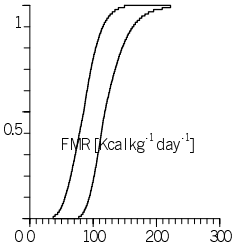


BW = normal( 608 gram, 66.9 gram)

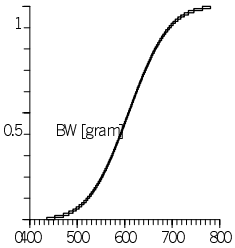


AEfish = minmaxmean(0.77, 0.98, 0.91)

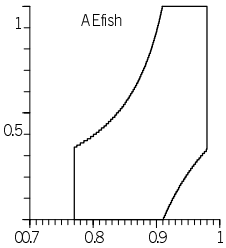


AEinverts = minmaxmean(0.72, 0.96, 0.87)

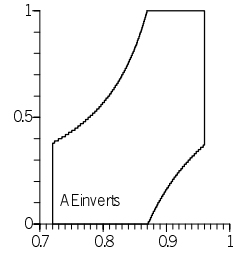


GEfish = normal(1200 Kcal per kg, 240 Kcal per kg)

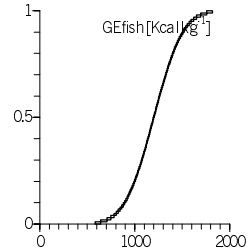


GEinverts = normal(1050 Kcal per kg, 225 Kcal per kg)

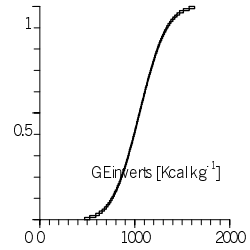


Cfish = [0.1,0.3] mg per kg

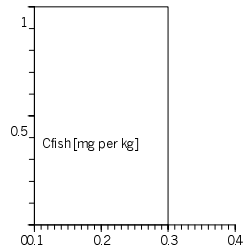


Cinverts = [0.02, 0.06] mg per kg

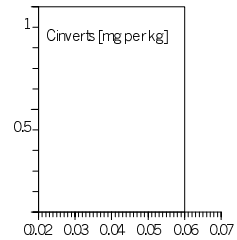


Pfish = 0.90

Pinverts = 0.10

TDI = FMR |*| ((Cfish |*| Pfish) |/| (AEfish |*| GEfish) |+| (Cinverts |*| Pinverts) |/| (AEinverts |*| GEinverts))

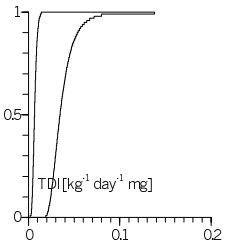


mean(TDI) // Total Daily Intake (units of mg of mercury per kilogram of mink tissue per day)

**[ 0.0074365, 0.034531] $kg^{-1}$ $day^{-1}$ mg**

median(TDI)

**[ 0.006601, 0.034378] $kg^{-1}$ $day^{-1}$ mg**

cut(TDI, 95%) // 95th percentile

**[ 0.011073, 0.061117] $kg^{-1}$ $day^{-1}$ mg**

sd(TDI)

**[ 0.0021782, 0.018363] $kg^{-1}$ $day^{-1}$ mg**

3. The "sample rule" says that *n* independent samples of a random variable divide the real line into *n* + 1 segments of equal probability. Is this reasonable? What would the p-box look like?

Suppose that the values 57, 15, 76, 37, 55, 11, and 23 were observed in seven random samples of some random variable. The assertion that there is equal probability within each of the resulting eight segments of the real line could be represented by the probability distribution shown in the graph below. This distribution assumes equiprobability within as well as among the segments of the real line. The tails, in principle, go out to infinity in both directions. This is the maximum entropy distribution corresponding to this sample of seven data points.

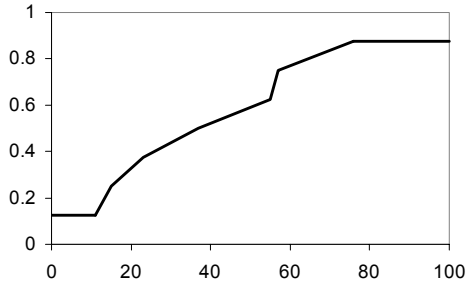


A p-box that represents a somewhat weaker statement is shown in the graph below. This p-box makes no claim about where within each of the segments the probability (1/8) lies.

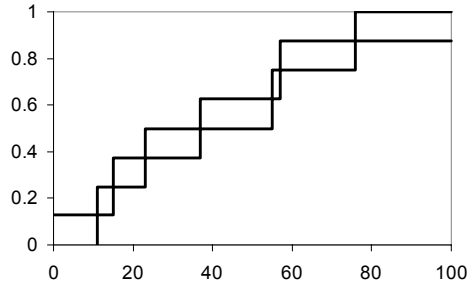


A p-box that makes a still weaker statement is shown below. This p-box represents the upper and lower 95% Kolmogrov-Smirnov confidence bounds on the distribution from which the seven data values were sampled. The breadth of these bounds acknowledges the small sample size of the data on which they are based. As the number of samples increases, the bounds tend to become closer together. The calculation of the Kolmogorov-Smirnov confidence bounds assumes that the data are random (that is, independent and identically distributed). The result is distribution-free except that it assumes continuity. Like other bounds based on confidence procedures, these bounds are not rigorous. Nevertheless, they do seem considerably more reasonable than either of the previous representations of the data in view of the smallness of the sample size.

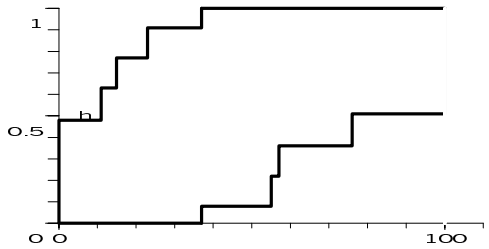


The Risk Calc commands below can be used to construct the three objects depicted in the graphs above.

```
a=[0,100]   // this should be the range of the random variable or perhaps [-infinity, infinity]
mixture(1,U(left(a),11),1,U(11,15),1,U(15,23),1,U(23,37),1,U(37,55),1,U(55,57),1,U(57,76),1,U(76,right(a)))
   ~(range=[0,100],  mean=40.5006,  var=736.829)
mixture(1,[left(a),11], 1,[11,15], 1,[15,23], 1,[23,37], 1,[37,55], 1,[55,57], 1,[57,76], 1,[76,right(a)])
   ~(range=[0,100],  mean=[34.2,46.8],  var=[434,1122])
histogram(a,a,11,15,23,37,55,57,76)
   ~(range=[0,100],  mean=[10,79],  var=[0,2470])
```

The U's denote uniform distributions. The right and left functions return the upper and lower endpoints of an interval. The first object is simply a stochastic mixture of uniform distributions. The second object is likewise a mixture of intervals. Note that the variable *a* is ascribed the range [0,100], which might be perfectly appropriate if the values are percentages. The choice of the range will strongly affect the estimates of the mean and variance.

4. Consider the "box-interval" conjecture (see BOXINT).

   The calculations in the script suggest that p-boxes are not merely bounds on probability distribution functions. A p-box carries a little bit more information than just these bounds. It also carries (interval bounds on) the mean and variance of the unknown distribution. This additional information can be used to improve the specificity of the results of calculations. For instance, moment projection formulas tell us that the means of sums, products and differences of independent random numbers are the sums, products and differences of the means. If the quantities are not independent, it is still possible to bound the moments of the functions. Distribution shape is also carried along with a p-box, such as when it is known to be normal, lognormal, Weibull, etc. Information about the shape of the distribution sometimes allows calculations to be improved. For example, additive convolutions of normal distributions under independence always yield normal distributions; multiplicative convolutions of lognormals yield lognormals; etc. There are a variety of theorems from mathematical statistics that can be employed to improve the results of calculations involving p-boxes. Finally, a p-box also carries its units, which is actually a very important piece of information that is commonly overlooked. Knowing the units allows the software to detect a large class of error conditions, such as adding quantities with incompatible units ("apples and oranges") or raising a quantity to the "2 inch" power.

   This strategy of compiling ad hoc propagation methods together may continue to be useful until convenient software to handle calculations involving credal sets or other representations of imprecise probabilities becomes available. In principle, arbitrary additional information could be carried along with a p-box and propagated through calculations so long as propagation formulas are known. This approach does not in general yield best possible results, of course, but may give results that are good enough for practical problems. What additional information could be carried along with a p-box that could be used to improve the calculations? For instance, would it be possible and useful to carry along with a p-box the fact that the distribution it represents is integer-valued? Would it be possible to keep track of impossible values or distributions or even p-boxes of disallowed distributions inside of a p-box? How could knowing such information be used to obtain improved numerical results? Suggestions are most welcome.

## Imprecise probability

1. Of the p-boxes shown in BOXES, which have the property that every distribution they contain is a member of the specified class?

   Running the BOXES script and selecting Windows/Tile from the main menu produces the display below. Only the boxes in the graphs labeled 'normal', 'minmaxmedian', and 'minmax' have the property that every distribution within the box has the specified property (of being normal, having the specified minimum, maximum and

median, and having the specified minimum and maximum, respectively). For example, every distribution inside the p-box in the bottom, right-hand graph has the specified minimum and maximum. See the script BOXES.UC itself for details about how the p-boxes are specified. Note that the p-box in the 'normal' graph has the property only because the parameters are precise that the p-box is degenerate so that it contains a single distribution.

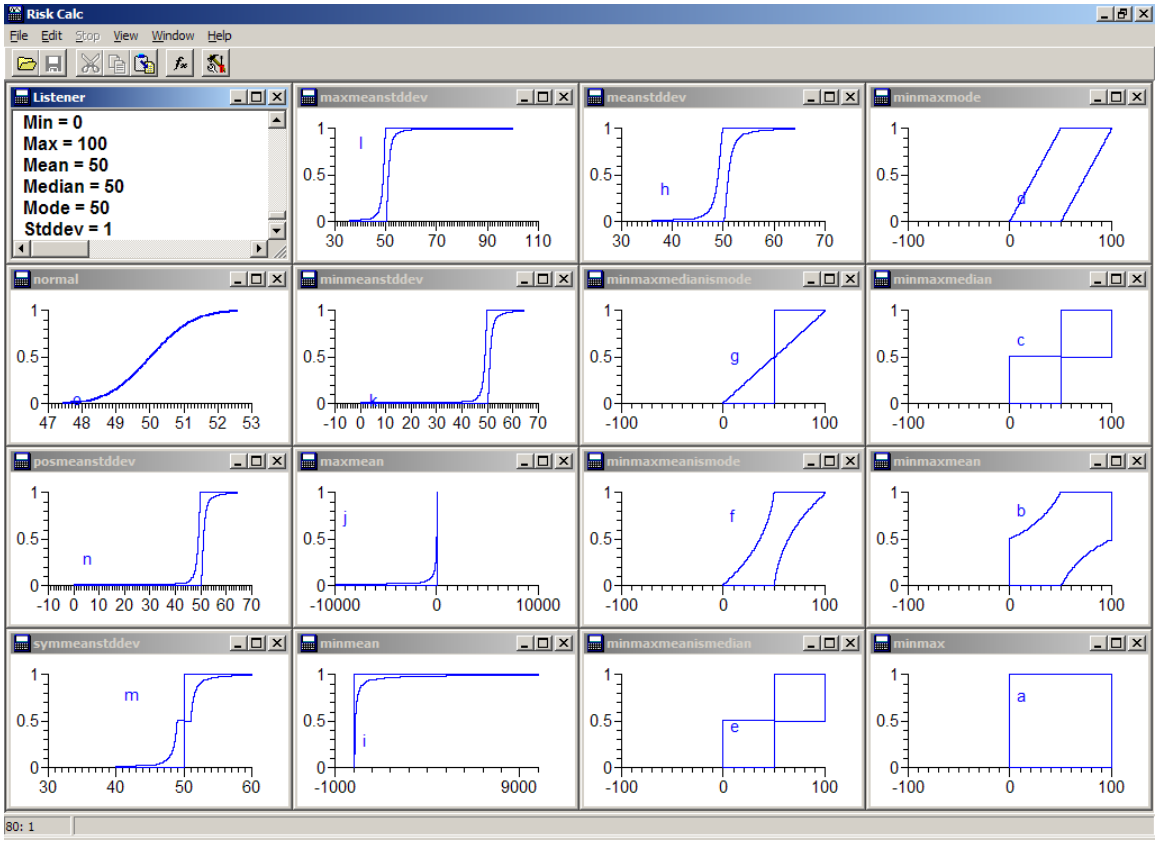


2. Would be the expectations of *TDI* and *HQ* in the mercury and PCB examples be tighter if computed with imprecise probabilities than the interval means from the probability bounds analysis? Would the estimates improve if the assumption of random-set independence were replaced by an assumption of strong independence?

The interval expectations for *TDI* and *HQ* might be tighter than those we obtained in exercise #2 in the section on probability bounds analysis. The intervals might also get tighter still if we assumed strong independence. However, computing the actual intervals would require mathematical programming, and there is currently no convenient software to make these calculations. Knowing how much improvement there could be in these practical cases would be very interesting and important. Suggestions, and especially sample calculations, would be extremely welcome.