

Appendix A

Distributions Used in Uncertainty Analysis

The Beta distribution

The Beta distribution is one of three distributions associated with a binomial stochastic process. A binomial process is a random counting system where there are a discrete number of opportunities (trials) of some particular event happening (successes) and where each trial has the same probability of being a success. This means that each trial must be independent of every other trial.

There are many systems that closely approximate a binomial process. Random processes like the tossing of a coin are binomial, since one face of the coin can be defined as being a success and the probability of each coin being a success remains constant for all tosses. No matter how many “heads” there have been in a row, the probability of a “tails” for the next toss remains the same (e.g. 50% for a fair coin).

Random sampling from a population may also closely approximate a binomial process, where we are concerned with determining what proportion of that population has some characteristic of interest. If the population is much larger than the sample size (a rule of thumb is that the population should be at least 10 times the size of the sample) then the probability of an individual randomly sampled from the population having the characteristic of interest remains fairly constant and equivalent to the proportion of the population with that characteristic. So, for example, if we are interested in the proportion of U.S. citizens that eat meat, we can do a random survey of U.S. citizens. Providing that our sample is much smaller than the population size, the probability that each consecutive randomly selected person eats meat remains reasonably constant, though there are actually a finite number of people who eat meat and there is a finite population too, which means that the probability that the next randomly sampled person eats meat does, in fact, depend on the previous samples.

The example above for meat eaters assumes that we are sampling without replacement: in other words, we would not survey any person more than once. It also implicitly assumes that there are a fixed number of people who eat meat in the population and a fixed population size too. This would be a static system with a constant proportion of meat eaters.

The Beta distribution can be used to model the confidence one has about the probability of success of a binomial trial p where one has observed n independent trials of which s were successes (41, 66, 94, 118, 119) so that it is said that p is distributed as a Beta($s+1, n-s+1$).

This distribution is the result of applying Bayes' Theorem with a Uniform(0,1) prior distribution and a binomial likelihood function. In layman's terms, Bayes' Theorem works as follows:

1. A prior statement of the knowledge of the variable to be modelled is given. In this case, we are saying that the probability p lies somewhere between zero and one, but we would not like to say that any value within that range was anymore likely than any other value (hence the Uniform(0,1) prior distribution).
2. For each allowed value within the prior distribution's range, we calculate the probability of observing the s successes we observed from the n trials. This probability is simply the binomial probability:

$$P(s;n, p) = {}_n C_s p^s (1 - p)^{n-s}$$

3. These binomial probabilities then become the weightings given to each value of p in the prior distribution. By normalizing these weightings we arrive at a posterior (i.e. final answer) distribution.

Thus, the posterior distribution $f(p)$ for p is given by the product of the prior distribution density and the likelihood function:

$$f(p; s, n) = \frac{1 \cdot {}_n C_s p^s (1-p)^{n-s}}{1 \cdot \int_0^1 {}_n C_s t^s (1-t)^{n-s} dt} \quad 0 < p < 1$$

where the 1 in the equation is the probability density of the Uniform(0,1) prior distribution, the ${}_n C_s p^s (1-p)^{n-s}$ is the weighting given to the p value (the likelihood function) and the integral in the denominator normalizes the distribution so that the area under the curve equals unity. By omitting the 1, cancelling out the ${}_n C_s$ and using $\mathbf{a}_1 = s+1$, $\mathbf{a}_2 = n-s+1$, we arrive at the equation for the Beta(\mathbf{a}_1 , \mathbf{a}_2) distribution's probability density $f(x; \mathbf{a}_1, \mathbf{a}_2)$:

$$f(x; \mathbf{a}_1, \mathbf{a}_2) = \frac{p^{\mathbf{a}_1-1} (1-p)^{\mathbf{a}_2-1}}{\int_0^1 t^{\mathbf{a}_1-1} (1-t)^{\mathbf{a}_2-1} dt}$$

Note that the prior distribution for p is a Uniform(0,1) distribution, shown in Figure A1. The Uniform(0,1) prior distribution is an uninformed prior, meaning that no subjective opinion or any other information has been involved in determining the prior. This is logically the most conservative approach one could take where conservatism here means expressing the maximum degree of uncertainty possible. The selection of an appropriate prior is sometimes slightly contentious. For example, Beta(0,0) is sometimes suggested as an uninformed prior, though it does not in theory exist. One criticism for using a Beta(1,1) prior is that the mean of the estimated probability is biased towards 50%, away from the observed proportion. In fact, for all applications of the Beta distribution in this analysis, the information contained in the prior distribution is generally overwhelmed by the information contained in the sample data and the results are essentially equivalent to a more traditional frequentist statistics approach which would give the confidence distribution for the probability p as: $\hat{p} = \text{Binomial}(n, s/n) / n$. The frequentist also uses the central limit theorem in situations where a large number n of samples were taken, say $n > 30$. By the central limit theorem \hat{p} has a Normal distribution $\text{Normal}(p, \sqrt{[p(1-p)/n]})$, where p is the true value of the probability of success. One then assumes (i.e. estimates) the confidence distribution for p to be :

$$\hat{p} = \text{Normal} \left(\frac{s}{n}, \sqrt{\frac{s(n-s)}{n^3}} \right)$$

Thus, from a practical viewpoint, there is little difference in using a Bayesian or frequentist approach to uncertainty estimating in this model, except that the Bayesian approach allows one to combine information from dissimilar data (66).

Perhaps the easiest way to understand the Beta distribution used in this manner is to look at a few plots of its shape for varying values of \mathbf{a}_1 and \mathbf{a}_2 . Figure A1, the Uniform(0,1) distribution, is also the Beta(1,1), i.e. the Beta distribution where we have observed $s = 0$ successes and $(n-s) = 0$ failures: this is the distribution when we have not yet done any trials and hence remains the prior distribution.




Figure A1: The Uniform(0,1) distribution, which is also the Beta(1,1) distribution

Figure A2 show the Beta(1,11) and the Beta(11,1) distributions: the former where we have observed zero successes in ten trials and the latter where we have observed ten successes in ten trials. Note that they are simply the reflection of each other since they essentially represent the same thing: one needs only to reverse the definition of a success to its opposite. Also note that, since all trials have been a success or a failure, the distributions peak at zero and one respectively. If this pattern continues with more trials, the distributions will become progressively more concentrated at zero and one.

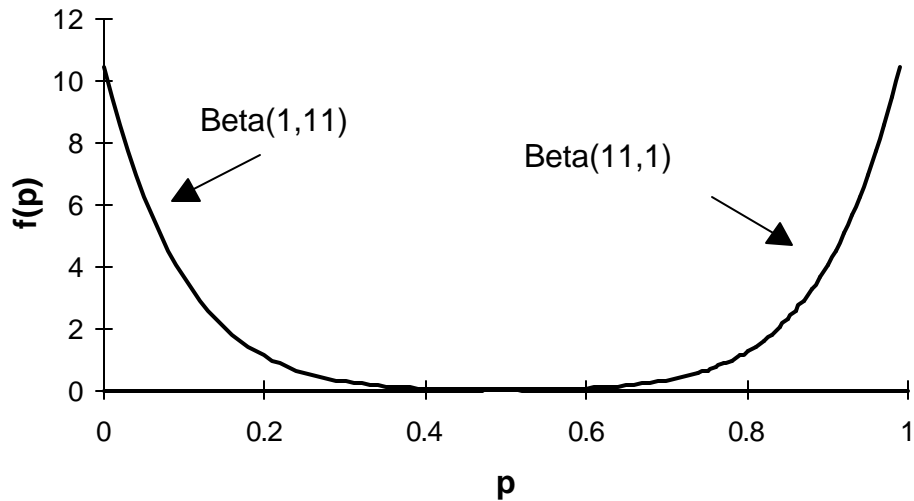


Figure A2: Examples of the Beta distribution where all trials are successes (peak at $p=1$) or all are failures (peak at $p=0$)

Figure A3 shows the Beta(3,3), Beta(11,11) and Beta(21,21) distributions representing four, 20 and 40 trials where 50% have been successes. Note that as the number of trials increases, the distribution becomes progressively narrower: in other words, one is becoming progressively more confident about what the true value of p must be.

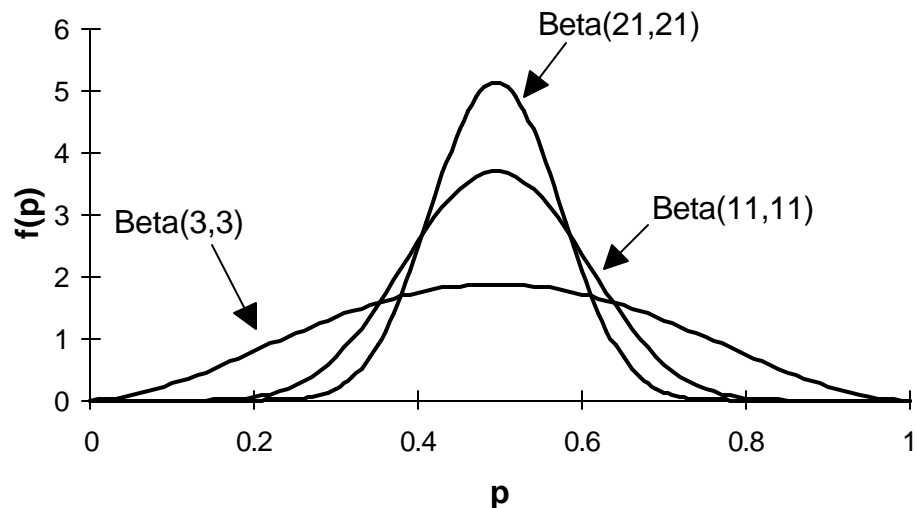


Figure A3: Examples of the Beta distribution where there are equal successes and failures (i.e. $\alpha=\beta$)

The figures below show that except for the case with very small n , the Binomial distributions or the Normal approximations to them that a frequentist would use to model the confidence distributions modeled as Beta posterior distributions above are practically indistinguishable from the Beta distributions. The cumulative probability distributions for the Binomial(4,0.5)/4, the Normal(0.5, 0.25), and the Beta(3,3) which are the ones for the case $n=4$, are shown in Figure A4a. One can see that for the proportion 0.5 the deviations between the Binomial and the Normal or Beta are fairly evenly arranged around the center of the graph. For a proportion of success less than 0.5, the deviations between the Binomial and the Normal or Beta are larger to the left half of the graph and very small toward the right side of the graph. Conversely, for a proportion of success greater than 0.5, the deviations between the Binomial and the Normal or Beta are larger to the right half of the graph and very small toward the left side of the graph. The major source of the deviation is the relative size of the steps in the Binomial when n is small. The cumulative distribution for the Binomial is a step function. As sample size increases the deviations between the discrete Binomial cumulative probability distributions and those of the corresponding continuous Normal or Beta distributions diminish. Figure A4b displays the cumulative probability distributions for the case $n=20$ and Figure A4c, for the case $n=40$. Because the Beta distribution is able to model success probability using non-integer numbers of successes it is extremely useful for risk assessment modeling where inputs are weighted survey results. Bayesian modeling is more conducive to combining various types of probability distributions than frequentist modeling is.




Figure A4a. Cumulative probability distributions for p based on the Beta, Binomial and Normal for the case $n=4$ and the number of observed successes is 2.

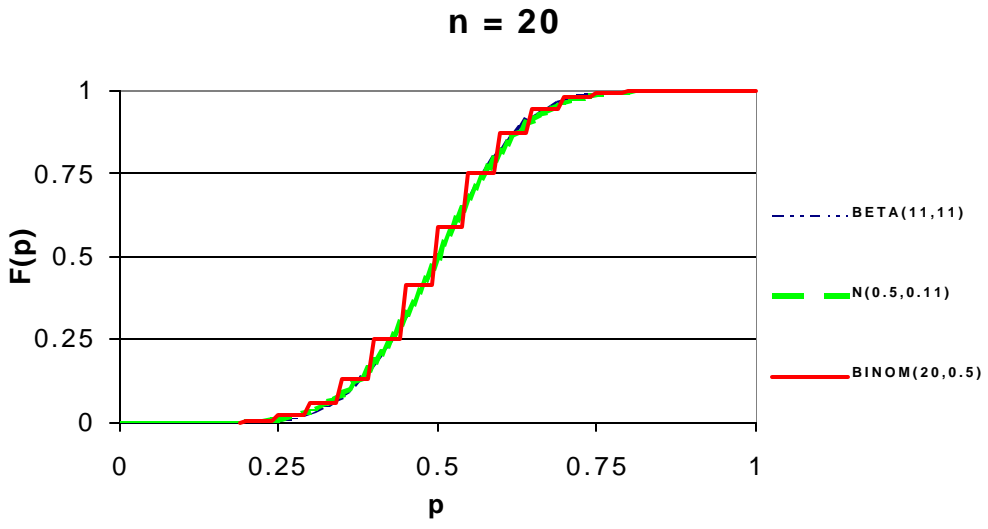


Figure A4b. Cumulative probability distributions for p based on the Beta, Binomial and Normal for the case $n=20$ and the number of observed successes is 10.




Figure A4c. Cumulative probability distributions for p based on the Beta, Binomial and Normal for the case $n=40$ and the number of observed successes is 20.

Use of the Gamma distribution to describe the uncertainty about a Poisson mean

Like the binomial probability p , the mean events per period \mathbf{I} is a fundamental property of the stochastic system in question. It can never be observed and it can never be exactly known. However, we can become progressively more certain about its value as more data are collected. Bayesian inference again provides us with a means of quantifying the state of our knowledge as we accumulate data.

Let us assume an uninformed prior $p(\mathbf{I}) = 1/\mathbf{I}$. The Poisson likelihood function for observing X events in period t is given by:

$$l(X | \mathbf{I}, t) = \frac{e^{-\mathbf{I}t} (\mathbf{I}t)^X}{X!} \propto e^{-\mathbf{I}t} (\mathbf{I})^X$$

The posterior distribution density is the product of the prior density and the likelihood function. We can disregard terms that not involving \mathbf{I} , recognizing that the distribution will be normalized eventually, and we then get the posterior distribution:

$$p(\mathbf{I} | X) \propto e^{-\mathbf{I}t} \mathbf{I}^{X-1} \quad \mathbf{I} > 0$$

which has the functional form of, and therefore is, a Gamma($X, 1/t$) distribution where \mathbf{I} is the variable.

The shape of the posterior Gamma distribution becomes progressively less sensitive to the prior distribution as data is collected. In essence, the sensitivity of the Gamma distribution to the prior amounts to whether $(X-1)$ is approximately the same as X . So, if X was 100, the difference would be roughly 1% influenced by the prior and 99% influenced by the data. In this model, the information contained in the quantity of data available *always* overpowers the prior.

Gamma distributions used in the risk assessment to model the rate of invasive infection, λ_i which has a Gamma(43,1) distribution and the rate of enteric infection, λ_e which has a Gamma(3985,1) distribution. Those distributions are shown in the following graphs.

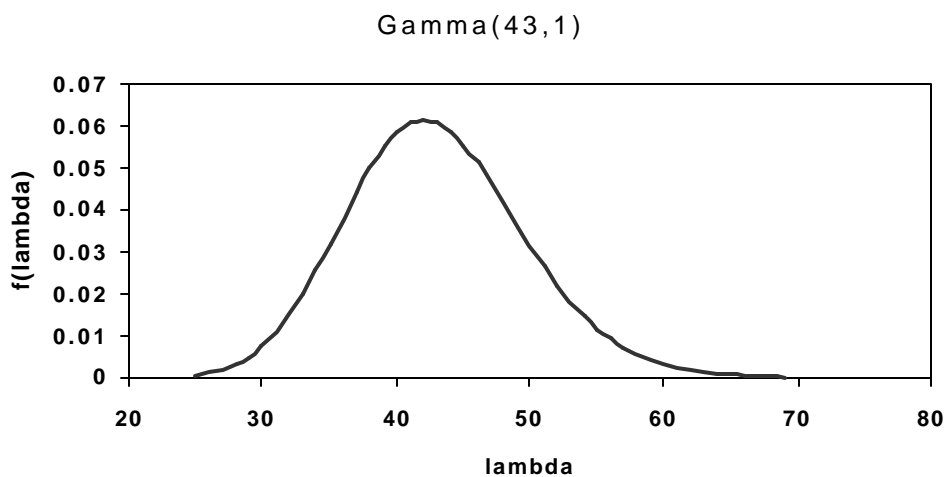


Figure A5a. The Gamma(43,1) distribution used to model the rate of invasive disease, λ_i .

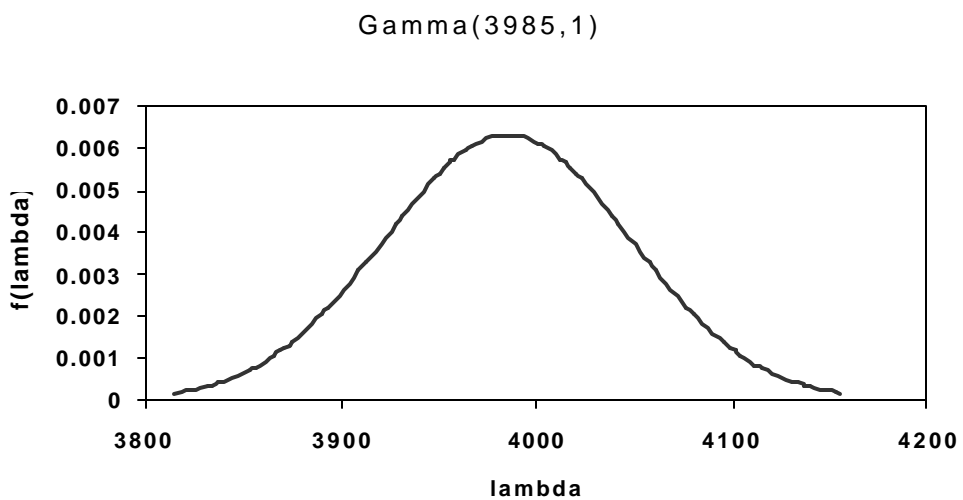


Figure A5b. The Gamma(3985,1) distribution used to model the rate of enteric disease, λ_e .

As was the case for the Beta distribution, there is a distributional counterpart to the gamma that the frequentist would apply to this estimation problem. When X , the number of events in one unit of time ($t=1$), is distributed $Poisson(\lambda)$, the maximum likelihood estimator for λ , $\hat{\lambda}$ is the observed value of X . Then the $Poisson(X)$ distribution is used as the uncertainty distribution for λ . The Poisson distribution is the limiting distribution for a Binomial random variable when n is large and p is very small. Because of this relationship, the central limit theorem applies for the Poisson under the same limiting conditions. By the central limit theorem $\hat{\lambda}$ is normal with mean λ and variance λ , where λ is the true value of the parameter. This means that the confidence distribution for λ is estimated as:

$$Normal(X, \sqrt{X})$$

The three possible choices for the confidence distribution for λ when 43 cases were observed during the year are shown to demonstrate how similar they are.




Figure A6. The cumulative distribution functions for Gamma(43,1), Poisson(43), and Normal(43,6.56)