Statistical Science Research Report 9

Models for Incomplete Failure Data

Russell G. Almond

Statistical Sciences, Inc.
1700 Westlake Ave, N., Suite 500
Seattle, WA 98109
almond@statsci.com

8/5/92

This research is supported by the GRAPHICAL-BELIEF project, NASA SBIR contract NAS 9-18669.

Models for Incomplete Failure Data

Russell Almond, StatSci

ABSTRACT

This paper describes how to build probability and belief function models for situations in which there is a lot of data about the failure rate for a component, but little data about the failure modes. Simple probability and belief function models can incorporate expert judgement about the relative frequency of the failure modes while using the existing data to determine the overall probability of failures. Slightly more complex 2nd order models additionally have the capability of learning about the relative failure rates associated with the various modes given additional data.

1. Overview

The test application of the `GRAPHICAL-BELIEF` project (Almond[1992a-c]) is to model system level reliability problems. Such models are generally build by describing the state of systems in terms of the states of components. The components are usually subjected to a battery of tests, including failure testing to determine the average lifetime. This data can be used to build probabilistic or belief function models of the chance of component failure. The information in these component level models are combined to produce the chance of system failure. The purpose of this document is to describe the process of encoding component test data and engineering experience into the component models.

Although mean time between failure data is often available for components, this data may say very little about what happens when the component fails. For example, a check valve may have three different failure states: "stuck open," "stuck closed," and "other" (which is principally for a rupture, but also includes unanticipated catastrophic failure modes). There would also be a fourth state representing the properly function condition. Information about the failure state may be critical to understanding system reliability. For example, if the check valve fails in "stuck open" the system may still perform the necessary function, but there may be increased risk of failure due to damage of some component protected by the check valve.

Because typical "mean time to failure" data does not provide information about the failure state, it must be regarded as incomplete. On the one hand, we may be able to estimate the overall failure rate for the component quite accurately. On the other hand, we know next to nothing about the relative likelihood of the various failure states, and have only our engineering judgement about the failure states to rely on. The ability to model partial information of this type was one of the exciting features of belief function inference when first proposed (Dempster[1966]), although the introduction of additional assumptions (perhaps based on an engineer's initial estimates) will allow us to model such incomplete data with probabilities as well.

The following terms (Walley[1991]) are useful for comparing the difference between belief function and probability models. A model is *uncertain*, if we don't know what is going to happen, but we could predict the number of outcomes of each type if it happened repeatedly. A model is *imprecise* if we cannot predict the average behavior of the system. Thus, flipping a coin is a precise but uncertain phenomenon; although we cannot predict a given flip, if we make many flips, we can be very sure that we will see approximately as many heads as tails. On the other hand, the failure of a component of a new type may be an imprecise phenomenon, we may not be able to predict with any accuracy the average failure time of a large number of such components. We may, however, be able to determine that the failure rate is rather small, so their may not be much uncertainty about whether or not a failure occurs during a small time interval (it probably won't). The appropriate model for our state of knowledge would then be an upper bound on the failure probability rather than a precise failure probability.

Probabilities (and Bayesian probabilities) are precise uncertain models. Belief functions are imprecise uncertain models. Thus, the probabilistic model of Section 3 makes a very (perhaps unwarranted) precise guess as to how often the various failure states occur for the valve. The belief function model of Section 4 makes an imprecise guess; giving upper and lower bounds instead instead of precise probabilities. We pay a price for allowing this imprecise specification; an additional computational cost, and a weaker decision making power (we may be unable to make a decision because our critical threshold lies between the upper and lower bound). For this reason, the theory of belief functions reguard Bayesian probabilities as an ideal precise state of information, and are able to represent them as a special case.

Data about an unknown parameter (such as a failure rate) introduces uncertain information about the parameter. This can lead to a 2nd order model for the unknown parameter. An additional assumption of prior information about the parameter in the form of a precise probability distribution results in a precise Bayesian model for both the parameter and the failure of the component in question. Using the theory of belief function, an imprecise specification of prior information leads to an imprecise model for both the failure rate and the failure of the component. Because of the ability to specify imprecise models, belief functions allow a unique representation of ignorance. To produce a "non-informative" prior distribution, one must make certain assumptions about the possible values of the parameter, for most problems there are several different "non-informative" models. In the examples discussed in this paper, the belief function model with no prior information bounds the predictions made by all of the Bayesian models using "non-informative" prior information. Thus, the belief function model captures some imprecision we may have about what should be taken as the baseline "non-informative" state of information.

One of the advantages of the belief function model is the ability to simultaneously consider a number of Bayesian probability models. Again we pay the price of additional complexity for the belief function models. The Bayesian models, on the other hand, may be overly precise; they will always make an exact prediction about the failure rate, even if that does not capture our true state of ignorance.

As we gather more data about the component, our information about its failure rate becomes more precise and less uncertain. This phenomenon, called "Bayesian learning," is an important characteristic of second order probability and belief function models. In fact, it often justifies the extra effort of building second order models, allowing the information gather over the course of operation to be incorporated into the running system.

The Phase II version of `GRAPHICAL-BELIEF` will include a "parameters" module which will allow parameters to be used in probability and belief function distributions. Such parameters will combine the simplicity of the simple belief function and probability models with the power of the 2nd order models. When only a simple calculation is required, a best guess about the uncertain parameters can be used. For more elaborate calculations, a high and a low value could be used and on rare occasions, such as final acceptance of a system, a more elaborate computation procedure can be used which fully exploits the model for component uncertainty.

Another class of models which has been proposed for solving such problems are fuzzy sets. Fuzzy sets, however, are models for imprecision (specifically the imprecision of natural language) rather than uncertainty. In reliability theory, we would like to have uncertainty, rather than imprecision. This makes probabilities the preferred model, with belief functions an alternative in the case of uncertain problem specification. Cheeseman[1986] argues that probability theory is quite adequate for many of the problems addressed by fuzzy logic, and often requires specifying fewer parameters.

We will use the valve described earlier as a running example and build first and second order probability and belief function models for its failure. Throughout we will assume that we have data available in the form of X failures where observed in T hours of testing, but through an omission in the study design, there is no information about which failure state the test component was in when it failed. There may be additional information about which component failed when, but by a sufficiency argument, that information will not become important until Section 8.

The model for this system involves two parts. The overall failure of the valve is described by a Poisson process. This model is described in Section 2. The separation of failures into the various failure states is modelled with a multinomial distribution. Section 3 describes this multinomial distribution and the parameters which must be estimated in order to use the model. Section 4 describes the belief function generalization of the multinomial model, and shows how to build a belief function model from complementary imprecise estimates.

As additional data become available, we would like to incorporate them into our models. In order to do this, we much first build 2nd order models for the various parameters of the system. Section 4 describes how to build a 2nd order model for the failure rate of the Poisson process. Section 5 describes how build a probability model for the chances of a given failure falling into a given category and Section 6 describes the corresponding belief function model.

The Poisson process model make several simplifying assumptions, one of them—the assumption of a constant failure rate—is unrealistic in many circumstance. Finally, Section 7 describes how to relax that

assumption, while still using the multinomial model of Sections 3, 4, 6 and 7 to describe the separation of failures into failure states.

2. The Poisson Process

Failures of a component can be thought of as an example of a counting process. That is we have some random function $N(t)$ which counts the number of failures which have occurred at time t (for simplicity assume that we stop the clock while we repair or replace the component). Logically, $N(0) = 0$ and $N(t_1) \geq N(t_0)$ if $t_1 > t_0$. A large proportion of the literature on reliability concentrates on building models for the stochastic process $N(t)$.

One of the simplest models for $N(t)$ is the *Poisson Process*. The Poisson process makes a number of assumptions about $N(t)$.

1. The process has "independent increments." That is knowing $N(t_1) - N(t_0)$ (the number of events in the time interval $(t_0, t_1]$) tells you nothing about $N(s_1) - N(s_0)$ (the number of events in the time interval $(s_0, s_1]$) if the two time intervals $(t_0, t_1]$ and $(s_0, s_1]$ don't overlap.
2. The failure rate is some constant α.
3. The probability of one failure in a small time interval $(t, t + h]$ is approximately αh, the probability of more than one failure is very small (smaller order than h).

For most reliability problems it is only the second assumption which seems unreasonable. Even so, it is likely to hold true for small fractions of the total component lifetime. Section 8 shows how this assumption may be relaxed. However, as a first approximation the Poisson process seems a quite reasonable model for component failure.

One of the nice features of the Poisson processes is that is parameterized by a single number, the failure rate α. It is called a Poisson process because the number of failures in a time interval of length t follows a Poisson distribution with parameter αt. Thus:

$$P\{N(t) = k\} = e^{-\alpha t} \frac{(\alpha t)^k}{k!} \qquad k = 0, 1, \ldots$$

The probability that there is no failure by time t is $e^{-\alpha t}$, and therefore the probability that the component will have failed at time t is $1 - e^{-\alpha t}$. This is the well known exponential distribution. It follows that the average lifetime of this component is $1/\alpha$, and thus the failure rate of the component is the reciprocal of the mean time between failures.

Given the particular form of our data, it is vary easy to estimate the failure probability. As X (the number of failures) and T (the total observation time for all components of this type) are the sufficient statistics, we can express our estimate in terms of them. It is simply $\hat{\alpha} = X/T$. Section 5 will look at some alternative estimates, but they will all be fairly close to X/T, especially if T (our measurement for quantity of data) is large.

3. Poisson–Multinomial Process

Consider for a moment, just the situations where the component has failed. We can divide these failures into one of M different categories. For example, the check valve had three possible failure states. We could count the number of failures in each category: $X_1, \ldots, X_M$. The total number of failures across all categories would be $N = N(t)$ the total number of failures we have observed. $X_1, \ldots, X_M$ follow the multinomial distribution (a generalization of the binomial distribution to more than two categories), with a total of N events.

The multinomial distribution is parameterized by a vector of probabilities $(p_1, \ldots, p_M)$, where p_1 respresent the probability that any given observation will fall into the first category, p_2 the second and so forth. As the probability that the distribution will fall into any category is always 1, the equation $\sum_{i=1}^{M} p_i = 1$ must hold. Thus there are only $M - 1$ free parameters. Enforcing the condition that the probabilities sum to 1 is known as normalization and can be done by simply summing over all the relative probabilities and dividing by the sum. Thus the statement "The probability of being stuck open and stuck closed are equally likely, and both together are ten times more likely than other failures" can be represented by a vector of relative probability judgements $(5, 5, 1)$ which is normalized by dividing all entries by 11.

The multinomial and Poisson processes have very convenient interaction properties. Suppose that failures for a certain type of component follow a Poisson process with rate α, and given that the component has failed, the probability that it is in failure state i is p_i. We can think of $X_i(t)$ as being a stochastic process counting the number of failure of type i by time t. This will be a Poisson process with rate αp_i! Thus from the parameters of the Poisson and multinomial model we can easily produce prediction equations for specific failure states.

The relation works the other way around as well. Suppose we have a series of independent Poisson processes $X_i(t)$ with failure rates λ_i. Then their sum $N(t) = \sum_{i=1}^{M} X_i(t)$ is a process which counts the number of times in which any of the events occurs. $N(t)$ is also a Poisson process, with failure rate $\alpha = \sum_{i=1}^{M} \lambda_i$. Furthermore, the conditional distribution of $(X_1(t), \ldots, X_M(t))$ (the number of events in each failure category) given $N(t)$ (the total number of events) is multinomial with $p_i = \lambda_i / \alpha$.

This combination of multinomial and Poisson processes actually allows us to separate the component failure model into two pieces. The first piece (about which we have data) is the overall failure rate of the component. The second piece (about which we have no data) is the relative frequency of the various failure modes. The set of parameters describing this should be much easier to elicit from engineers than the overall failure rate.

This model is adequate for first pass computations of system failure, however, it essentially assumes that the relative rates for the various failure states is known. Often this is the information which is specifically missing from the failure data, therefore we need ways of relaxing this assumption. The next section describes a belief function model for the same process which allows us to express our incomplete information about the relative rate of the failure states. As we collect information about the relative likelihood of the various failure states we would like to be able to update our state of knowledge about the parameters $(p_1, \ldots, p_M)$ of the multinomial distribution. Sections 6 describes how to this can be done.

4. Belief Function Poisson–Multinomial models

Consider once again the check valve which has three possible failure states: stuck *open*, stuck *closed* or some *other* state (such as a rupture). For simplicity, we will use the word in italics to denote the state. If we know that the valve is failed, that knowledge could correspond to one of seven possible sets of failure states: $\{open\}$, $\{closed\}$, $\{other\}$, $\{open, closed\}$, $\{open, other\}$, $\{closed, other\}$, or any of the three. A belief function would allow us to ascribe a probability mass or *mass* to each of those sets of states. From these we can easily assign upper and lower bounds to the probability of being in each failure state. For example the lower probability or *belief* associated with failure *open* is the probability mass assigned to the set $\{open\}$ and the upper probability or *plausibility* is the sum of the probability masses associated with the sets $\{open\}$, $\{open, closed\}$, $\{open, other\}$, and $\{open, closed, other\}$.

It might also be possible to go from upper and lower probability judgements to belief functions, but you must be somewhat careful here, because not all upper and lower bound judgements can be mapped to the probability distributions over sets necessary for belief functions. It is usually easier to work directly from the states of information which generate the mass functions rather than trying to specify the probability distribution directly.

The power of belief functions comes from their flexibility in modelling incomplete states of information such as might be associated with the failure states of the valve. An important example is the *vacuous belief function* created by placing probability mass 1.0 on the set $\{open, closed, other\}$ (any state). This conveys the complete lack of information in the failure data about the relative frequencies of the various failure states. This model will result in conservative estimates, that is the upper bound on the *other* failure rate is likely to be too high and the lower bound too low.

As conceptually appealing as the vacuous model is, it is not very realistic. In the case of our valve, the probability of an *other* failure (specifically, a rupture) known to be smaller than the probability of either of the other two failure states. For example, from our experience we may decide that one of the two stuck states is between 9 and 99 times more likely than the *other* state. This results in the following mass function:

$$m(\{open, closed\}) = .9 \qquad m(\{open, closed, other\}) = .09 \qquad m(\{other) = .01 \tag{1}$$

This translates into an upper bound of 10% and a lower bound of 1% for the *other* state. Both the *open* and *closed* states have a belief and plausibility of 0 and .99 respectively, however, the set of states $\{open, closed\}$ has belief .9 and plausibility .99. Thus this belief function tells us that the failure is very likely to be a stuck *open* or *closed*, but not which one.

We might further know that stuck *open* and *closed* failures were roughly equally likely. We need to quantify what we mean by this; it might translate into a judgement of the type that the probability of seeing *open* is no more than three times that of seeing *closed* and similarly the probability of seeing *closed* is no more than three times that of seeing *open*. Because this judgement says nothing about the state *other* we can express it like this:

$$m(\{open, other\}) = .25 \qquad m(\{open, closed, other\}) = .5 \qquad m(\{closed, other\}) = .25 \tag{2}$$

This gives a belief of .25 and plausibility of .75 for the both the set $\{open, other\}$ and the set $\{closed, other\}$.

Because Judgement 1 tells us about the relationship of *other* to the *open* and *closed* states, and Judgement 2 tells us about the relationship between the *open* and *closed* states, they are independent and we can combine them using Dempster's product-intersection rule. This results in the belief function represented by the following mass function:

$$\begin{array}{ccc} m(\{open\}) = .225 & m(\{closed\}) = .225 & m(\{open, closed\}) = .45 \\ m(\{open, other\}) = .0225 & m(\{closed, other\}) = .0225 & m(\{open, closed, other\}) = .045 \\ & m(\{other\}) = .01 & \end{array}$$

This model is more complex the one of the previous section; it also more closely models our imprecision state of information about the relative rates of the various failure categories. Thus we trade complexity of modelling and representation for a more specific notation for ignorance. By building a second order model, we can also learn from our experience. This is described in Section 7.

5. Gamma–Poisson Model

The preceding discussion has assumed that certain parameters of the models, such as the failure rate of the component where known with certainty. In practice, this is rarely the case, there may be some uncertainty about these parameter values; however, additional data about the component will decrease that uncertainty. This phenomenon is known as Bayesian learning, and is important enough that we will want to capture it in `GRAPHICAL-BELIEF`.

The key idea of Bayesian statistics is that such uncertainty about parameters can be represented by probability distributions. Such probability distributions can be easily updated in the presence of new data by a simple application of Bayes theorem. Thus the Bayesian model "learns" from experience. In general the model for the parameter *posterior* to obtaining the data with be less uncertain (as measured by the variance of the parameter) than the best model for the parameter *prior* to obtaining the data. The terms *prior* and *posterior* are usually used for the model or distribution of the parameter prior to and after obtaining the data.

Building a model for the overall failure rate of the component requires knowledge about the likely form of the data. The simplest way to gather information about the failure rate is to take a sample of N components and see at what time they fail. A sufficiency argument states that for the Poisson process we can add the failure times of all N components together to get the total failure time T. (In fact, they individual failure times are useful for checking the model, they should be roughly exponentially distributed. If N is reasonably large, this should be checked with simple graphical techniques.) Because the waiting time for one component to fail follows an exponential distribution, the waiting time T for all X components to fail must follow a gamma distribution with shape parameter X and Scale parameter $1/\alpha$. A little bit of algebra results in a $Gamma(X+1, 1/T)$ distribution as being the best estimate of our state of knowledge about the unknown failure rate α.

It is very easy to update a gamma model for the failure rate α of a Poisson process. Suppose that our current state of information for α follows a gamma distribution with shape parameter a and scale parameter b; that is, it has a probability density function:

$$f(\alpha) = \frac{1}{\Gamma(a)b^a}\alpha^{a-1}e^{-\alpha/b} \qquad Gamma(a,b)$$

The average estimate for α is ab, the variance of our information about α is ab^2, and the most likely value of α (mode) is $(a-1)b$. If we later observe (independent) data of the form X failures in T hours of operation, we can update our knowledge about α to a gamma distribution with shape parameter $a+X$ and scale parameter $[1/b+T]^{-1}$. This property is known as conjugacy, and hence the family of models for the process and parameter is known as the *Gamma–Poisson* family.

Conjugate distributions are easy to use as models of prior information. In order to specify the gamma distribution, we must specify two parameters a and b. The mean of the distribution is easily obtained as an engineer's best guess for the failure rate α. The second parameter can be obtained by equating the engineer's experience to so many hours of operation of the process.

The argument above which lead to the $Gamma(X+1, 1/T)$ distribution for the unknown parameter α is actually not a Bayesain one, but a fiducial one (Fisher[1973]). However, Fisher's model is equivalent to a Bayesian model in which we assume that *a prior* to seeing any data, all values of α are equally likely (a $Gamma(1,0)$ distribution). Unfortunately, this choice of a "non-informative" prior is not unique. An information theoretic argument due to Jefferys[1961] gives a non-informative prior of $Gamma(1/2, 0)$ and a scale invariance argument of Jaynes [1968] gives a prior of $Gamma(0,0)$. Note that when the amount of data (in terms of total observation time T), the three posterior distributions resulting from the three non-informative prior distributions are remarkably similar, yielding average estimates of $(X+1)/T$ $(X+1/2)/T$ and X/T respectively. It is only when the amount of information (data plus expert experience) is small, that the choice of baseline for the model is critical.

A variation on the Fisher[1973] argument and the Dempster[1966] application of it to the binomial problem, produces the following belief function model for the failure rate (Almond[1991a,1991b]). Instead of producing a random estimate for the unknown failure rate, the belief function model produces a random

interval $[\underline{\alpha}, \overline{\alpha}]$ which contains it:

$$m\left([\underline{\alpha}, \overline{\alpha}]\right) = \frac{s^{X+1}}{(X-1)!}\underline{a}^{X-1}e^{-\overline{a}T}$$

This model yields the following upper and lower expected values for the rate:

$$\overline{\mathrm{E}}[\lambda] = \frac{X+1}{s} \qquad \underline{\mathrm{E}}[\lambda] = \frac{X}{s}$$

Thus the belief function model is roughly equivalent to looking at all three of the "non-informative" priors. This model not only captures the uncertainty about the unknown failure rate due to the data, but the imprecision is specifying a base-line non-informative prior.

Note that all three Bayesian non-informative priors, and even the vacuous belief function model are somewhat unrealistic. They are actually improper prior distributions because the excessive mass they put on large values of the failure rate prevents them from being normalized (they cannot be integrated from 0 to infinity to calculate the normalization constant). Although this problem goes away as soon as data are introduced, it is somewhat disquieting, especially as near infinite failure rates are never seen in practice.

It is always almost possible to do better than the "non-informative" priors as a model for for ignorance, even about a brand new component. For example, an engineer could easily put a conservative upper bound on the failure rate (fewer than 1 failure per hour, for example). This could be used to select an appropriate non-zero scale parameter for the appropriate gamma distribution. In the Belief function case, the same gamma distribution is used for the upper bound on the interval (the lower bound of zero is usually acceptable).

The preceding discussion of prior and posterior models seems to add needless complication to the straightforward problem of calculating system failure rates, especially as there may be a great deal of data available about the component in question. In fact, in many situations, where only a nominal estimate of failure probability is needed (for example, diagnosis and first pass design sensitivity studies) it gets in the way of the main issues. In `GRAPHICAL-BELIEF`, the uncertainty about model parameters (such as failure rates) is hidden by the *parameter* mechanism.

In this case, the failure rate of a given component is a parameter. If similar components appear many places in the system, that component may appear many times. If only a nominal calculation is required, then a nominal estimate (such as the mean of the distribution) can be substituted wherever the parameter appears, and the calculations proceed according to the normal mechanism. For more elaborate calculations, a high and a low estimate (for example a 95% and 5% quantile of the distribution) can be used to get a "approximate best" and "approximate worst" case estimate. A complete and correct estimate can be given by the Monte Carlo procedure described in Almond[1990]. This is more time consuming, but would be needed only rarely, (for example, for final sign-off on a design) and hence can be safely assigned to overnight calculation or supercomputing.

The important idea is that the parameters hide the distribution from the casual user of the system. Once in place in the library, the parameter can be used to form nominal, "approximate best case" and "approximate worst case" estimates, or simulated draws from the distribution without the user needing to know the details of the distribution. Furthermore, data gathered in the course of using the system to diagnose real failures can be recorded by the parameters for updating the distributions (probability or belief function) reducing the amount of uncertainty about the parameter.

Although the indirection of building two stage models is not crucial in the case of the overall failure rate, where we have a good deal of data, it become much more important when the data is much more scarce, especially about the failure state information. These models are the topics of the next two sections.

6. Dirchlet–Multinomial Model

Section 3 discussed the idea of using the multinomial distribution to model the separation of failures into distinct failure states. The critical parameter was a vector of probabilities $\mathbf{p} = (p_1, \ldots, p_M)$ describing the probability that a given failure will fall into each of the M possible states. In Section 2 we derived a simple "back of the envelope" estimate for those failure probabilities. Using a fully Bayesian model, we can update our first guess given additional data, such as experience with the system in the field.

The natural conjugate family for the multinomial distribution is the Dirchlet distribution. This is characterized by a vector of hyperparameters $(\eta_1, \ldots, \eta_M)$ corresponding to the vector $\mathbf{p}$. For convenience, let $\eta_0 = \sum_{i=1}^{M} \eta_i$. The exact form of the distribution is:

$$f(\mathbf{p}) = \frac{\Gamma(\eta_0)}{\prod_{i=1}^{M} \Gamma(\eta_i)} \prod_{i=1}^{M} p_i^{\eta_i - 1} \qquad Dirchlet(\eta_1, \ldots, \eta_M)$$

Note that $\sum_{i=1}^{M} p_i = 1$ so that this is really a $M - 1$ dimensional distribution; there is also a small amount of negative dependence among the p_i's. If we have a $Dirchlet(\eta_i, \ldots, \eta_M)$ prior distribution over $\mathbf{p}$ and we observe some number of failures in each category $(X_1, \ldots, X_M)$, then the posterior distribution about $\mathbf{p}$ is $Dirchlet(\eta_1 + X_1, \ldots, \eta_M + X_M)$. Note that this means that η_0 can be interpreted as the number of cases observed.

The mean of the Dirchlet distribution is $(\eta_1/\eta_0, \ldots, \eta_M/\eta_0)$. This gives us a useful method for assessing Dirchlet priors: first assess the "best guess" for the vector of probabilities, and then assign our guess a relative worth in amount of data. Take for example, the check valve with the three possible failure states described in Section 2. Because the chance of *open* and *closed* are judged to be equal $\eta_1 = \eta_2$. Because the chance of *open* and *closed* together are 10 times more likely than *other* $\eta_1 + \eta_2 = 10\eta_3$. Finally, if we judge that our data is worth approximately 50 cases, we get $\eta_0 = 50$ Solving these three equations, gives us the approximate solution $(22.5, 22.5, 4.5)$ (with a $\eta_0 = 49.5$ instead of 50. Note that we have "fractional cases" in the solution to this problem. This presents no difficulty in the theory or mathematics, and such fractional cases are often useful in assessing prior information.

Once again it there is some controversy over the choice of "non-informative" prior distribution for the Dirchlet. The $Dirchlet(1, \ldots, 1)$ distribution is uniform over all allowable values of $\mathbf{p}$. Jeffreys' argument leads to $Dirchlet(1/2, \ldots, 1/2)$ as the non-informative prior and Jaynes' leads to the $Dirchlet(0, \ldots, 0)$ (which is again an improper distribution). Again, the choice of baseline is not critical when we have a large amount of data; however when the amount of data is small, the choice of "non-informative" prior could be the basis of the decision.

One way to approach this difficulty is to due a sensitivity analysis to the choice of prior. If the choice of prior (or non-informative prior) model makes a large difference in the final conclusions, then an engineer should spend a good deal of time developing a prior model for the failure rate for the component. It might also be a good idea to send the component out for more testing, thus reducing the amount uncertainty about that component. Another approach is to look at the belief function models described in the next section.

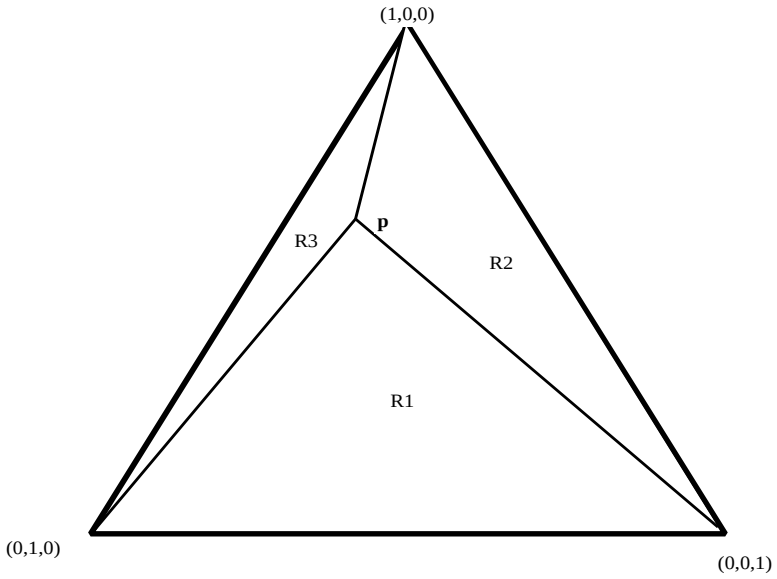

Figure 1. Probability for Valve and regions corresponding to various states

7. Dempster's Multinomial Models

Dempster[1968], often cited as the origin of belief functions, contains a specific model for the multinomial distribution. It starts with a slightly different way of viewing sampling from a multinomial distribution. Consider the trinomial distribution which describes which failure state the valve is in given that it has failed. The vector of probabilities $\mathbf{p} = (p_1, p_2, p_3)$, where $p_1 + p_2 + p_3 = 1$ describes the state of the result. Because of this we can plot $\mathbf{p}$ in a triangle using barycentric co-ordinates as shown in Figure 1.

There are three regions marked on the figure. We can think of a random vector $\mathbf{u}$ which is uniform on the triangle. If $\mathbf{u}$ falls into Region R1, then the failure will be of Mode 1 (stuck open) and so forth. The probability of failure is proportional to the area of the regions, and a simple geometric argument shows that the area of Region Ri is p_i. Thus this triangle is a model for the trinomial distribution (it becomes a simplex instead of a triangle for $M > 3$).

If we had absolutely no information about $\mathbf{p}$ except that it is a trinomial probability, that would result in it lying anywhere in the triangle. Specific information about $\mathbf{p}$, such as a belief function model like the one in Section 3 could be a specific region containing $\mathbf{p}$. Probabilitistic information about $\mathbf{p}$ can be represented by a distribution over this triangle, such as a Dirchlet (whose space conveniently follows the triangle).

Dempster[1968] notes that starting with vacuous and after observing some multinomial data, the information about the multinomial parameter $\mathbf{p}$ can be described by a random region in this triangle, as shown in Figure 2. Let (x_1, x_2, x_3) be the vector of observations and $X = x_1 + x_2 + x_3$. Then the endpoints of the figure will have a sort of joint Dirchlet distribution (described in Dempster[1968]); the marginal distributions of the endpoints are given as follows:

W_1: $Dirchlet(x_1 + 1, x_2, x_3)$	Z_1: $Dirchlet(x_1, x_2 + 1, x_3 + 1)$
W_2: $Dirchlet(x_1, x_2 + 1, x_3)$	Z_2: $Dirchlet(x_1 + 1, x_2, x_3 + 1)$
W_3: $Dirchlet(x_1, x_2, x_3 + 1)$	Z_2: $Dirchlet(x_1 + 1, x_2 + 1, x_3)$

From this we can see that the random region will always contain the random point defined by the Dirchlet distribution starting with any of the "non-informative" priors of the previous session. Thus, Dempster's distribution contains not only the uncertainty about the unknown parameter, but also the the imprecision in specification of a baseline model.

We can get an average case model which looks like the simple belief function model of Section 4, by taking the average values of the endpoints. The seven regions of Figure 2 each correspond to one of the focal elements of the belief function, we can calculate the belief mass associated with each one by calculating the

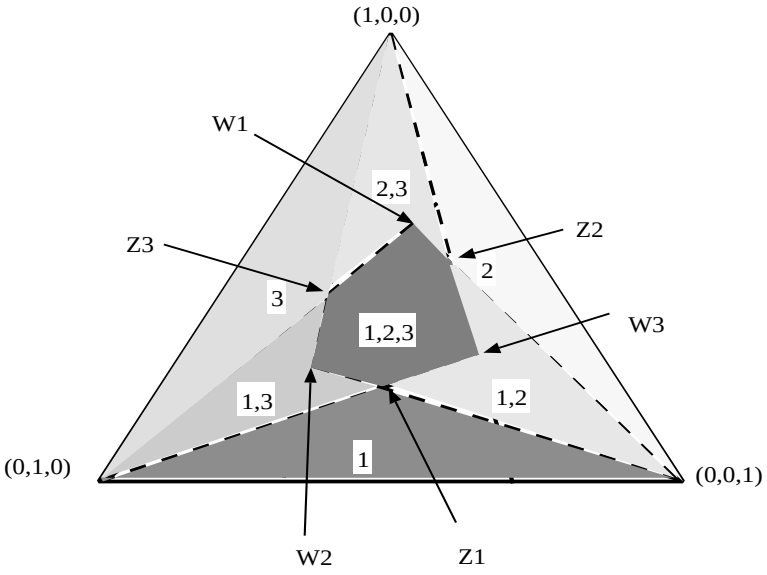

Figure 2. Dempster's Random Region distribution

area. This results in the following belief function (over the possible failure states $1, 2, 3$):

$$\begin{array}{ccc} m(\{1\}) = \frac{X_1}{X+2} & m(\{2\}) = \frac{X_2}{X+2} & m(\{3\}) = \frac{X_3}{X+2} \\ m(\{1,2\}) = \frac{X_1+X_2}{(X+2)(X+1)} & m(\{1,3\}) = \frac{X_1+X_3}{(X+2)(X+1)} & m(\{2,3\}) = \frac{X_2+X_3}{(X+2)(X+1)} \\ & m(\{1,2,3\}) = \frac{2}{(X+2)(X+1)} & \end{array}$$

This results in the following upper and lower bounds for the probability of each state:

$$\begin{array}{ll} \mathrm{BEL}(\{1\}) = \frac{X_1}{X+2} & \mathrm{PL}(\{1\}) = \frac{X_1+1}{X+1} \\ \mathrm{BEL}(\{2\}) = \frac{X_2}{X+2} & \mathrm{PL}(\{2\}) = \frac{X_2+1}{X+1} \\ \mathrm{BEL}(\{3\}) = \frac{X_3}{X+2} & \mathrm{PL}(\{3\}) = \frac{X_3+1}{X+1} \end{array}$$

These bounds cover the estimate of the failure rates for each state which would by the three non-informative priors (respectively X_i/X, $(X_i + .5)/(X + 1.5)$ and $(X_i + 1)/(X + 3)$.

Note that although this distribution is quite complex, it is specified simply by three number, the number of observed failures of each type. Thus it is quite simple to specify a distribution of this type, all of the complexity is hidden within the model specification. It is also straightforward to combine data of this type with bounds on the failure probability of the kind described in Section 4.

8. Nonhomogeneous Poisson Processes

Section 2 introduced the Poisson process and the three assumptions that went into this model. The second of those assumption, the constant failure rate was somewhat suspect. In fact, it is probably much more common for components to have a high failure rate early in their lifetime (during an initial "break-in" period), a low failure rate during the first few years of service and then a slowly rising failure rate as the component nears the end of its service lifetime.

Such a situation can be modelled with a *nonhomogeneous Poisson process*. Instead of the constant failure rate α, the failure rate is a function of time $\alpha(t)$. Now the number of failures during any time interval $(t_0, t_1]$ follows a Poisson distribution with average number of failures $\lambda(t_0, t_1) = \int_{t_0}^{t_1} \alpha(t)dt$. Thus it is quite possible to estimate the failure rate of the system for any desired time interval.

The difficulty lies in estimating the function $\alpha(t)$. Actually, the literature on this process is quite extensive including both parametric (mostly developed by researchers in reliability) and nonparametric (often developed by medical researchers for patient survival) estimates. The individual failure times of the tested components are used to fit models for $\alpha(t)$, so such models can often be fitted from existing pools of data.

For the purposes of graphical belief modelling of systems, we will most often be interested in specific chunks of time. For example, we may want to build a model for the first month, the rest of the first year, the second year, and so forth. This means we only need a step function approximation to $\alpha(t)$ and not the full model.

Finally, we can still combine the more complex nonhomogeneous Poisson process model with the multinomial models for the various failure states.

9. Comparison of Proposed Models

The models presented in the paper have been in roughly increasing order of complexity. The probability models have been simpler and required fewer parameters than each of their belief function counterparts. It is not explored here, but there is an additional computational cost associated with the belief function models as well. On the other hand, the class of belief functions is richer and more expressive than the class of probability models. The question still remains as to whether or not that extra expressive power is necessary.

Consider once again the valve about which we known the failure rate, but not the relative frequency of the failure modes. We can build a model for this valve using two parameters, α the failure rate and $\mathbf{p}$ the vector of probabilities for each failure state given a failure. We may have quite exact information about α but relatively little information about $\mathbf{p}$. Here the ability of belief functions to model imprecise states of information becomes quite exciting, maybe even worth the additional computational expense.

All of the probability models (Sections 2, 3, 5 and 6) require precise assumptions about unknown parameter values. While this is not uncomfortable in the models of Sections 2 and 5 for which we have a great deal of data, it is uncomfortable in the case the models for Sections 3 and 6 which depend on human judgement of unknown accuracy. There are two solutions to this problem: to expend a lot of manpower and expertise on the problem to make the judgements as accurate as possible, or to perform a *sensitivity analysis* to determine which of the judgements are critical in the final estimates or decision making process. Those judgements which are critical are candidates for further refinement either by further data gathering or the application of more expertise.

The belief function models, which are able to capture imprecise states of information, can describe states of information in terms of upper and lower bounds. On the other hand, a belief function model may not be decisive, consider what would happen if the upper and lower bound of the final estimate for the system failure probability straddled the accepted limit. The model would need further refinement (or the collection of more data) before the system could be accepted. This could act as a partial sensitivity analysis, although a full sensitivity analysis to find and test critical assumptions is important in the belief function models as well.

An important characteristic of both Bayesian probability models and belief function models is their ability to learn. As actual experience and tests with the component will tell us more about the relative frequency of the failure modes, we should be able to combine that information with the prior expert judgement to make an updated model. Fortunately, using the models of Sections 5, 6 and 7 that is quite

simple. For that reason the second order models are preferable to the first order ones of the earlier sections. Graphical models, combined with this learning ability using Bayesian and belief function models, creates a model which almost looks like a neural-net, but has the ability to explain itself. Research in this area is just beginning, but we hope that GRAPHICAL-BELIEF will be a platform to support such future developments.

One of the big open questions in this kind of work, is whether or not 2nd order probability models are sufficient to capture the knowledge state about an uncertain event. It may be quite possible to model imprecision about the unknown parameter with a Bayesian probability statement describing uncertainty about it. As long as an honest assessment of the uncertainty in system failure rates due to uncertainty about the parameter is made, we may have sufficient information with which to make a decision, and to estimate our ignorance. Calculating the uncertainty of the final estimate (perhaps as a variance or standard deviation) may provide a sufficient guide to the amount of information in the model. This is the argument of many of the critics of belief functions (see some of the discussion of Dempster[1968]). Dempster himself argues that exact Bayesian information is preferable when obtainable. By allowing GRAPHICAL-BELIEF to support both kinds of models we hope to be able to better understand the trade-offs between belief function and probabilistic modelling.

10. References

Almond, Russell G. [1990]. *Fusion and Propagation in Graphical Belief Models: An Implementation and an Example.* Ph.D. dissertation and Harvard University, Department of Statistics Technical Report S-130. To be published as a monograph from Van Nostrand Reinhold.

Almond, Russell G. [1991a]. "Fiducial Inference and Belief Functions." Personal Technical Report A-4. Technical Report 206, University of Washington, Department of Statistics.

Almond, Russell G. [1991b]. "Belief Function Models for Simple Series and Parallel Systems." Personal Technical Report A-6. Technical Report 207, University of Washington, Department of Statistics.

Almond, Russell G.[1992a]. "GRAPHICAL-BELIEF: Design Overview and Review" Statistical Science Research Report 14.

Almond, Russell G.[1992b]. "GRAPHICAL-BELIEF: Users commands and Rule Bases." Statistical Science Research Report 15.

Almond, Russell G.[1992c]. "GRAPHICAL-BELIEF: Phase I Project Report." Statistical Science Research Report 16.

Cheeseman, Peter[1986]. "Probabilistic versus Fuzzy Reasoning." *Uncertainty in Artificial Intelligence.* Kanal, L.N. and Lemmer J.F. (ed.) Elsevier Science Publishers B.V.

Dempster, Arthur P. [1966]. "New Methods for Reasoning Towards Posterior Distributions Based on Sample Data." *Annals of Mathematical Statistics,* **37**, 355-374.

Dempster, Arthur P. [1968]. "A Generalization of Bayesian Inference (with discussion)." *Journal of the Royal Statistical Society, Series B,* **30**, 205-247.

Fisher, R.A. [1973]. *Statistical Methods and Scientific Inference*. Hafner Press. (This edition was revised posthumously to include some notations Fisher had made in his copy for later revision). Earlier editions: 1956, 1959.

Jaynes, E.T. [1968]. "Prior probabilities." *IEEE Transactions on Systems Science and Cybernetics,* **SSC-4**, 227–241.

Jeffreys, H. [1961]. *Theory of Probability* (3rd edn.). Oxford University Press.

Walley, Peter[1991]. *Statistical Reasoning with Imprecise Probabilities.* Chapman and Hall.