Proceedings, *Windpower '95,* AWEA, Washington DC, March 27-30, 1995.

ALL WIND FARM UNCERTAINTY IS NOT THE SAME:

THE ECONOMICS OF COMMON VERSUS INDEPENDENT CAUSES*

Paul S. Veers
Sandia National Laboratories
Wind Energy Technology Department
Albuquerque, NM 87185-0708
USA

ABSTRACT

There is uncertainty in the performance of wind energy installations due to unknowns in the local wind environment, machine response to the environment, and the durability of materials. Some of the unknowns are inherently independent from machine to machine while other uncertainties are common to the entire fleet equally. The FAROW computer software for fatigue and reliability of wind turbines is used to calculate the probability of component failure due to a combination of all sources of uncertainty. Although the total probability of component failure due to all effects is sometimes interpreted as the percentage of components likely to fail, this perception is often far from correct. Different amounts of common versus independent uncertainty are reflected in economic risk due to either high probabilities that a small percentage of the fleet will experience problems or low probabilities that the entire fleet will have problems. The average, or expected cost is the same as would be calculated by combining all sources of uncertainty, but the risk to the fleet may be quite different in nature. Present values of replacement costs are compared for two examples reflecting different stages in the design and development process. Results emphasize that an engineering effort to test and evaluate the design assumptions is necessary to advance a design from the high uncertainty of the conceptual stages to the lower uncertainty of a well engineered and tested machine.

INTRODUCTION

The return on an initial capital investment in wind turbines is obtained by continuous operation of the machines over several years. The financial risk, or expected costs, must be examined and quantified before large investments can be made and large numbers of machines can be built. Certainly, investors expect some estimate of the risk they are taking with their money for comparison with the projected returns and other investment options. However, risk can be difficult to quantify with relatively new technologies or new kinds of hardware. In the case of wind turbines, the risk is driven by uncertainty, especially in the durability of the structure. A large part of the financial risk of operating wind turbines is in the replacement costs (and ancillary loss of revenue) associated with broken components.

The fatigue life of many wind turbine components is susceptible to large uncertainties for two reasons. First, the fatigue resistance of all materials has a large amount of inherently random scatter. That is, given two nominally identical pieces of material repeatedly stressed under identical conditions, the two pieces may fail at lifetimes different by factors of ten or even hundreds. Second, the nature of the fatigue process

* This work was supported by the United States Department of Energy under Contract DE-AC04-94AL85000.

is such that a small change in the loading experienced by the material will lead to a large change in the material lifetime. This sensitivity exacerbates the problem of not knowing the loadings perfectly. Small uncertainties in the loadings lead to large uncertainties in component lifetimes. The sum of these two effects is to create a wide range of possible lifetimes for fatigue-susceptible wind-turbine components.

In this paper, the economic impact of uncertainty is addressed by calculating not only component probability of failure, but by estimating the experience of a fleet of identical turbines. The number of components expected to fail in each year of operation is calculated, and the costs are assigned to those replacements. The cost in each operating year is then known, and the present value can be estimated. Thus, the nature of the risk to the fleet is quantified for use in making financial decisions.

Component fatigue life is usually calculated using the best estimates of uncertain load and resistance quantities and applying reasonable safety factors. A better measure of design adequacy is obtained by estimating the distribution of possible values for these uncertain inputs and calculating a probability of component failure at a specified target lifetime. But the probability of a component failing is not the same as the percentage of components in the fleet expected to fail. It is necessary to separate the uncertainty into two types: common, where all components share a load or strength value but that value is not known with certainty, and independent, where the value for each component varies independently of the others. The effect of these different types of uncertainty is addressed here.

SOFTWARE TO CALCULATE PROBABILITY OF FAILURE

A software tool has been developed for evaluating the probability of wind turbine components meeting a target lifetime; it is called FAROW, for Fatigue And Reliability Of Wind turbines [Veers, et al., 1994]. FAROW uses the relatively new approach of structural reliability theory to evaluate the probability of premature failure in the presence of multiple uncertain inputs with arbitrary distribution of possible values. It is specifically tailored to the wind turbine fatigue problem and does all the difficult numerical calculations internally, leaving the user to focus attention on the still formidable task of determining the distribution of possible values for all of the uncertain inputs. FAROW calculates several quantities of interest, including the median lifetime of the part, and the probability of failing before some specified target lifetime, as well as importance factors, which are estimates of how much each random variable contributes to the probability of failure. The sensitivity of the results to changes in each input quantity is also calculated.

When the probability of a component failing in less than Y years is calculated by FAROW to be X%, one often hears the interpretation that "you can expect X out of 100 components to fail in the first Y years of operation." Unfortunately, this very simple and useful way to think is usually wrong. It would be correct if all the uncertainties in the inputs are completely *independent* from component to component. However, much of the uncertainty does not lie in the randomness of an input quantity from component to component. Rather, the quantity has some value that varies quite little from component to component, but the exact value of the quantity is simply not known. This uncertainty is *common* (perfectly correlated) between all the machines in the fleet. If *all* of the uncertainty is common between all the components, the correct interpretation of the above statement would be that either none or all of the components will fail, and the probability of all of them failing is X%. Real life is never so simple as to fit into either limiting category, but contains uncertainty of both the common and independent varieties.

SEPARATING COMMON AND INDEPENDENT CAUSES: FATIGUE PROPERTIES

Completely separating the uncertainty in component fatigue life into common and independent sources is a virtually impossible task, or at best very difficult. However, material fatigue properties have such a large and inherently independent variability that they can be used to approximate all the independent uncertainty

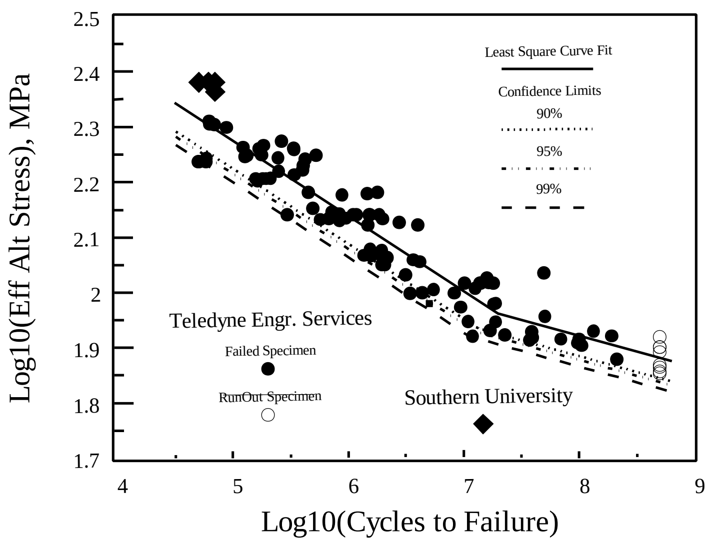

Figure 1: Typical S-N test results for identical specimens. These data are from an aluminum alloy [Van Den Avyle and Sutherland, 1989].

in the component probability of failure. Figure 1 shows fatigue test results for identical specimens, plotted as effective, alternating-stress amplitude versus number of cycles to failure, i.e., a stress-life, or S-N, plot. Notice that identical material specimens tested at the same stress level can have lifetimes that differ by a factor of ten or more. This is the norm and not the exception with fatigue properties. A typical value for the standard deviation of the cycles-to-failure is 60% of the mean value [ASCE, 1982]. It is quite possible to have common material property uncertainty due to manufacturing processes and to material lot differences. These, however, are assumed to be small relative to the inherent randomness of the material property.

The randomness in material properties is described in FAROW by using a single random variable to represent the coefficient of the S-N curve (its intercept). The S-N coefficient can also be entered in FAROW as a deterministic quantity representing a given "confidence level" or survival rate. Figure 1 shows curves for four survival rates: 50% (Least Squares Curve Fit), 90%, 95%, and 99%. FAROW then calculates the probability that the designated percentage of components (equal to the survival rate) will last for the target lifetime. Keep in mind that there are still many uncertain inputs describing the loading.

The non-material property uncertainty is dominated by common sources (i.e., values that are common to all the components, but not known with certainty). All of these inputs, although possessing some independent randomness from machine to machine, are most likely dominated by the uncertainty that is common to all components. Therefore, the material property is chosen to represent all the independent uncertainty and the rest of the inputs are assumed to be entirely common between components. This simplification is chosen as a convenience and is not necessary for the application of the procedure presented here.

The result is that FAROW can estimate the probability of achieving a fleet-wide survival rate specified by the S-N curve survival rate at any designated target lifetime. Different S-N curves are input to calculate the probability of achieving different survival rates. By applying the replacement cost to the numbers of components failing and weighting by the probability of that occurrence, the *expected*, or average, cost of fleet maintenance due to the a particular component failure and replacement is estimated. The analysis is repeated at different target lifetimes to assess the time at which replacement costs are accrued and to calculate the present value of such costs. The following examples outline the process step-by-step and illustrate some typical results.

EXAMPLES

The process of calculating the economic effect, or risk, of uncertainty from different sources, independent and common, may be illustrated with a pair of examples taken directly from the FAROW User's Manual. One case represents the situation in which extensive prototype testing has reduced the uncertainty in the machine response to the environment about as far as possible. This "low uncertainty case" has a median lifetime of 300 years, while the probability of the component failing in less than 20 years is 3%. The other

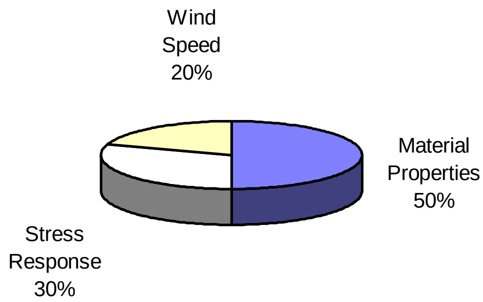

Figure 2: Relative importance of the three sources of uncertainty in the low uncertainty case.

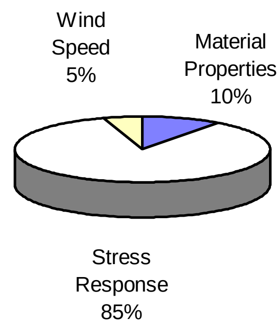

Figure 3: Relative importance of the three sources of uncertainty in the high uncertainty case.

case reflects a situation earlier in the design and development process before there has been much testing. Structural response levels may have been calculated, but have not been test validated. There would therefore be a high uncertainty on stress levels. This "high uncertainty case" has an 11% probability of component failure in less than 10 years although the median lifetime is 600 years.

Details on the entire description of the input quantities reflecting the appropriate degree of randomness and uncertainty for these examples can be found in the FAROW User's Manual. The exact inputs are not important to the topic of this study. Rather, they can be summarized using the importance factors calculated by FAROW. Importance factors reflect the contribution to the probability of failure due to each of the random inputs. Figures 2 and 3 show the importance factors for the two cases lumped into three areas: wind speed, stress response and material property inputs.

Material properties include the inherent randomness in fatigue properties, and represent all of the independent uncertainty in these examples. The wind speed category describes the annual wind speed distribution. The stress response category includes such quantities as stress concentration factors, nominal stress levels as a function of wind speed and cyclic stress amplitude distribution parameters. As stated above, all the latter two categories are designated as common sources of uncertainty.

EXAMPLE WITH LOW UNCERTAINTY

If all median properties are used to calculate the fatigue life in this example, the life of this component (a blade joint) is estimated at about 300 years. Of course, no designer worth his or her salt would ever design with median properties. Some substantial factors of safety would be applied. Here, we calculate the probability that a component will last for a predetermined period of time using the FAROW software. As stated above, this probability of failure in a 20-year lifetime for an individual component has been calculated at 3%. The probability of failure for lifetimes less than the 20-year

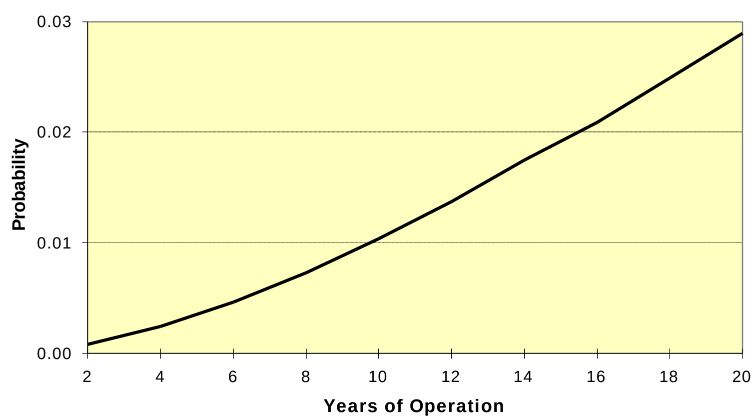

Figure 4: The probability of component failure grows with time as the fatigue damage accumulates (low uncertainty case).

target is also easily estimated and is plotted in Figure 4. The question remains: What do all these numbers mean? We would like to know how many machines are likely to fail and at what time, rather than the probability that any individual component will fail.

The number of components likely to fail is assessed by first setting the material property to a specified percentage-survival level. A target lifetime is then selected. FAROW is then used to take all the remaining uncertain quantities and calculates the probability that the survival level will be achieved at the target lifetime. The results of this analysis for several survival levels and a 20-year target lifetime are shown in Figure 5. While it is practically a sure thing that this component will exceed 50% or 60% survival rates, it has only a slim chance of achieving a 99.9% survival percentage. The chances of achieving survival percentages between these extremes are shown in the figure. For example, the chance of achieving a 98% or higher survival level at 20 years is 0.71, and exceeding the 99% level has only a 0.57 probability.

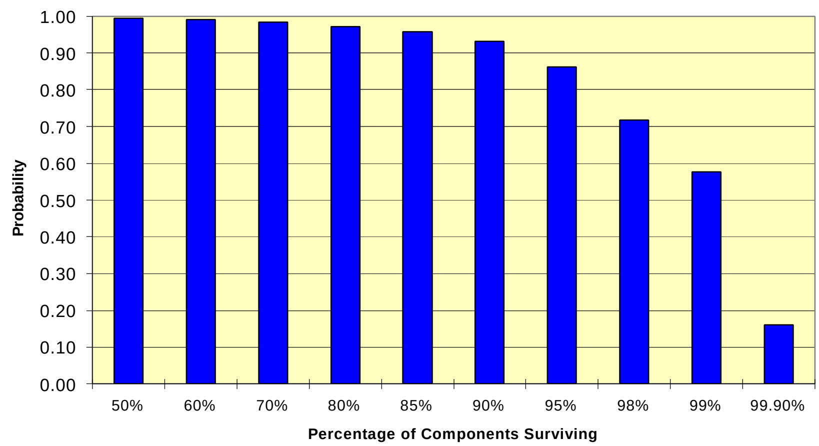

Figure 5: Probability of achieving various component survival percentages after 20 years of operation.

The expected cost of replacement is calculated by first determining the probability that different percentages of components will fail, then assessing a cost to that number of replacements, and finally adding up the costs over all possible percentages of failures. The difference between the probabilities at each level in Figure 5 is the probability that the percentage of components failing will be in the range between those levels. For example the probability that the number of failures after 20 years will be between 1% and 2% is 0.71 - 0.57 = 0.14. The cost associated with between 1% and 2% of all components failing can be lumped at 1.5% of the total fleet replacement cost. Let the fleet replacement cost be 100 units for both this example and the next. The expected cost of a fleet failure rate exactly between 1% and 2% is therefore 1.5% times the cost of replacement times the probability that the failure rate will be in the specified range: 0.015 x 100 x 0.14 = 0.21. Similar cost estimates can be made for all the other ranges of survival percentages and for all the other years of operation. Keep in mind that the calculated costs are cumulative over time.

This same calculation can be done at earlier target lifetimes to fill in a complete description of the number of components that are likely to fail and after how many years. Figure 6 shows a compilation of these results from 2 to 20 years for this example.

Figure 7 shows the cumulative cost breakdown by number of components expected to have failed at 2 and at 20 years. The greatest cost associated with any particular percentage point is in the <1% bin, but because range widths are not uniform in Fig. 7, the costs are larger in bins associated with the wider ranges. Also, the greatest probability of occurrence is for the <1% bin. Because of the greater costs associated with greater number of component failures, the costs do not drop off as fast as the probabilities. This example illustrates a case for which about half the expected cumulative costs are due to the risk that small numbers of components (less than 10% of the total installed) will have failed in 20 years of operation.

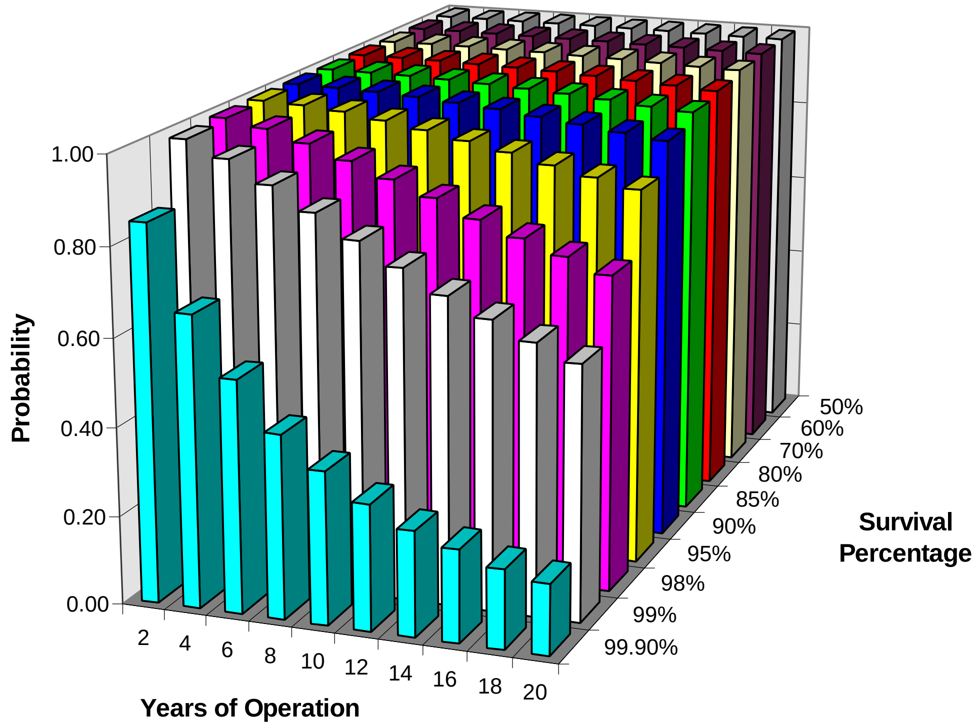

Figure 6: Probability of achieving different component survival percentages from 2 to 20 years for the low uncertainty case.

It is perhaps more interesting that *half the financial risk* comes from the chance that *more that 10%* of the components will fail. This outcome may not have been apparent from the combined sources calculation, which produced a 3% probability that any particular component will fail in the 20-year target lifetime.

The expected cumulative costs are shown in Figure 8. Notice the similarity with Figure 4. The total expected cost could also have been calculated by taking the probability of individual component failure (from all sources of uncertainty, common and independent) and multiplying by the cost of replacement. The expected cost is the same whether the sources of the uncertainty are broken out or not. It would not be the same if the cost associated with larger failure percentages were higher than for smaller percentages. Such might be the case if loss of production due to

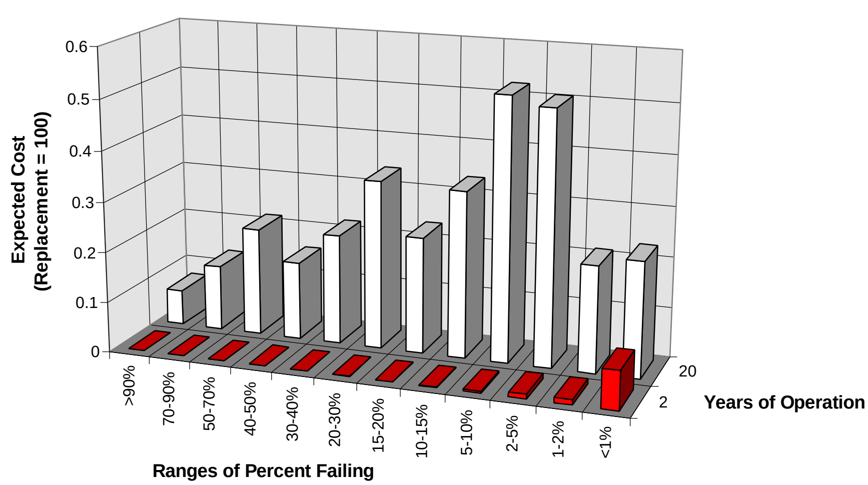

Figure 7: Source of replacement cost risk divided into ranges of percent of components failing. Notice that the range widths are not uniform.

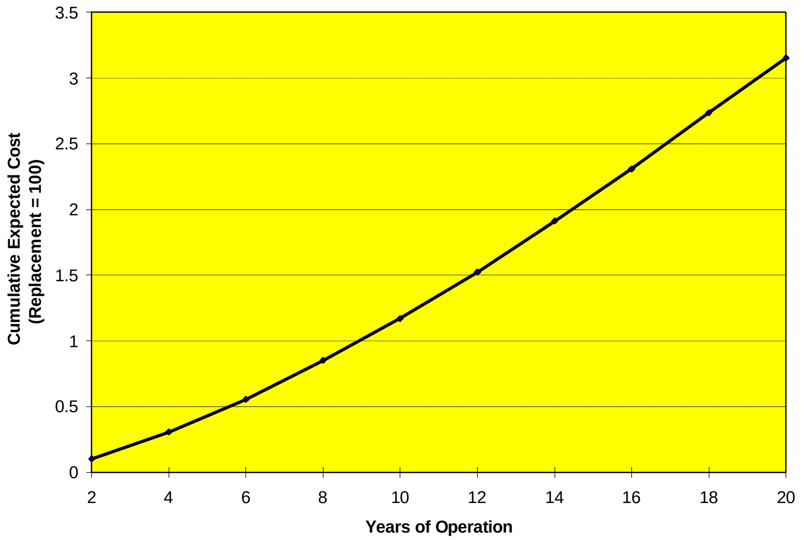

Figure 8: Cumulative expected cost of replacement from summing the expected cost of replacing the percentages of failures in each range of Figure 7 in each 2-year interval from 2 to 20 years.

maintenance down time were included in the replacement cost. High rates of component failure could lead to an inability to fix the machines in a timely manner, resulting in low availability. These effects, and other more detailed cost effects such as replacement of the replaced parts, are not included in these examples.

Incremental costs in each 2-year interval are calculated by taking the difference in cumulative costs between consecutive 2-year increments. The present value of these incremental costs is shown in Figure 9 for discount rates of zero, 5% and 9%. The present values for 20 years of operation come out to be 3.1, 1.7 and 0.8, respectively. Notice that because the incremental costs initially increase and then level out, the present values are dominated by the costs in the later years, unless the discount rate is high.

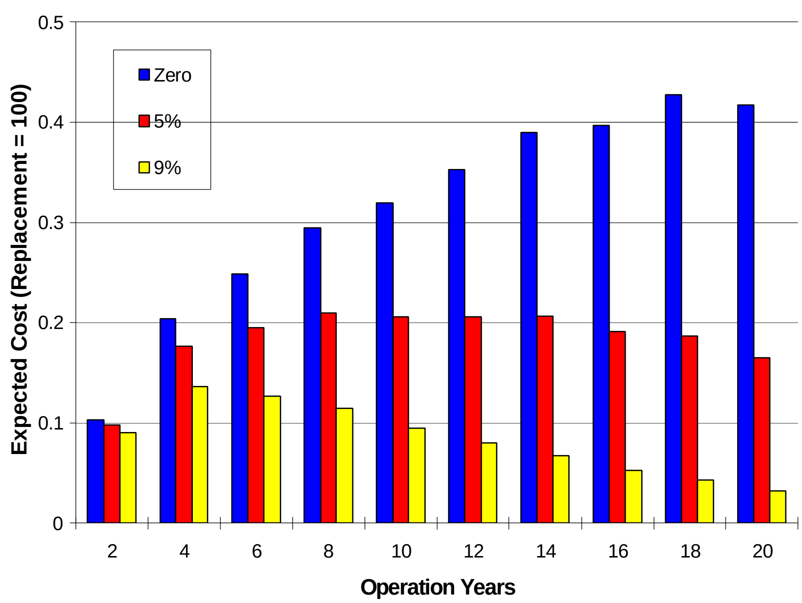

Figure 9: Present value of incremental costs (in two year increments) for the low uncertainty case using different discount rates.

EXAMPLE WITH HIGH UNCERTAINTY

Now let's examine a very different situation, one in which there is substantially higher uncertainty in the stress response[1] of the turbine. The wind speed and material property uncertainty are assumed to be about the same as in the low uncertainty example. Here, however, the stress response uncertainty is so much higher it accounts for 85% of the probability of failure, as shown in Figure 3. There is sufficient randomness to produce an 11% probability of individual component failure in less than 10 years, even

[1] It is not essential to this discussion to know exactly where all the uncertainty comes from. It is usually a combination of imperfect knowledge of the overall level of response, distribution of stress cycle amplitudes, mean stress, and stress concentration factors.

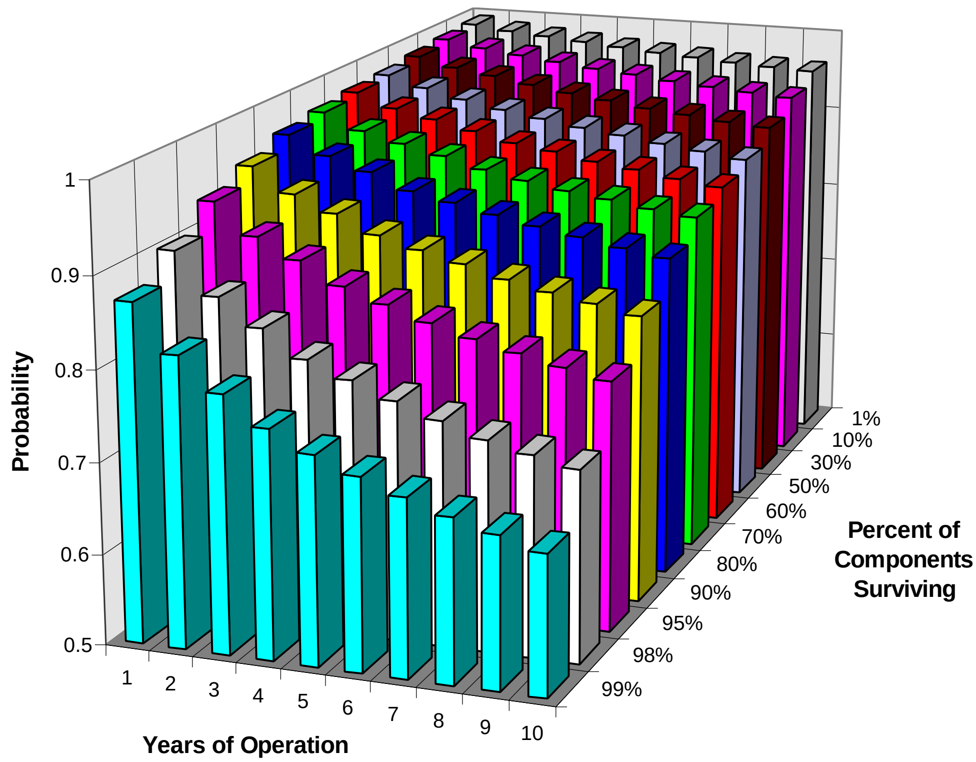

Figure 10: Probabilities of achieving component survival levels for various years of operation in the high uncertainty case.

when a 600-year lifetime is calculated using median values of all the inputs. Should this be interpreted as 11 out of 100 components having failed after 10 years? No. Let's see why.

The probability of achieving various component survival levels has been calculated in the same manner as for the first example. The results are quite different in this case. While in the low uncertainty example there was virtual certainty that not all the components would experience problems, here there is a substantially higher chance of fleet-wide problems. Figure 10 shows calculations of the probability of achieving different percentage survival levels at each of the years from one to ten. Notice that the scales are different than in Figure 6, especially on the probability axis, which runs from 0.5 to 1.0 in Fig. 10 and 0 to 1 in Fig. 6. There is no certainty that even very low survival rates will be achieved. Conversely, there is almost a 2 in 3 chance that there will be a 99% survival rate or better. Because the stress response uncertainty is high, there is a slight chance that stresses will be too high throughout the fleet, but also a good chance that the component has been over designed and will have needlessly low stresses in all applications.

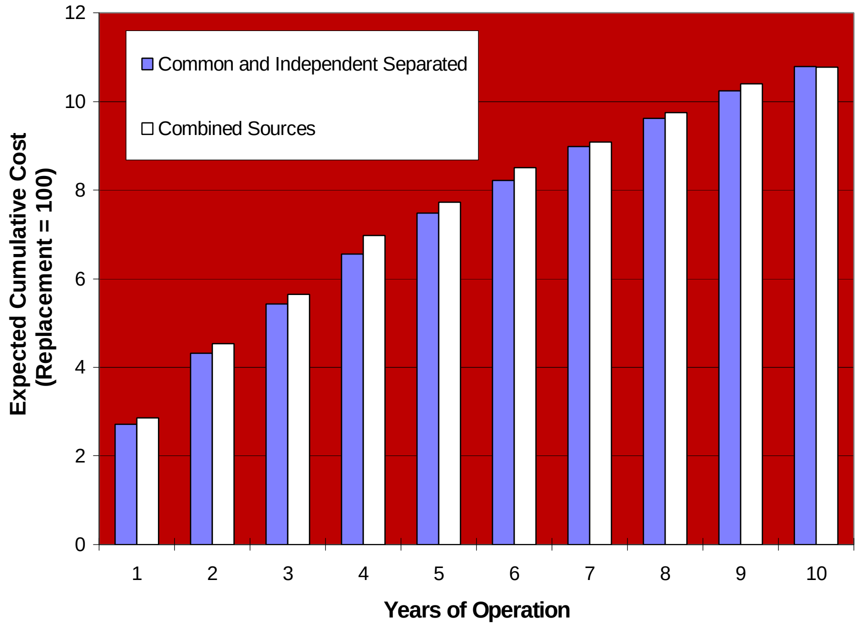

Figure 12: Cumulative costs for replacement of failed components estimated by combining all sources of uncertainty (from FAROW, see Figure 4) and by separating the sources, calculating the percentages of components likely to fail at any time, and summing all the costs.

The source of the costs is also calculated as before with dramatically different results. Figure 11 shows the source of replacement costs in terms of the percentage of components failing. The risk is shown to be dominated by the cost associated with the finite probability that all the components will fail. The total expected cost (with replacement cost set to 100 as in the first example) is 11. The expected cost of more than 99% of the components failing (1% survival level) is about 4, which is more than a third of the total cost of replacements after 10 years. Even in the first year, the expected cost of full-fleet replacement is a large part of the total. On the other hand, there is so little chance that only a small part of the fleet will be experiencing problems that it is not a significant part of the total. The expected costs associated with less than half the components failing is about 40% of the total in the first year and drops to about 25% of the total risk after 10 years.

The total expected cost of 11 matches the probability of failure calculated from combined common and independent sources of uncertainty (11%) times the replacement cost. The cumulative costs match the combined sources probability of failure in every year, as seen in Figure 12. The separated-sources result comes from the sum of the costs across all of the percentages of components failing (in Figure 10) for each target lifetime. The combined result is obtained by using the probability of failure calculated by FAROW, which combines all sources of uncertainty. The results are nearly identical and, in fact, would be exactly the same except for numerical errors accumulated by

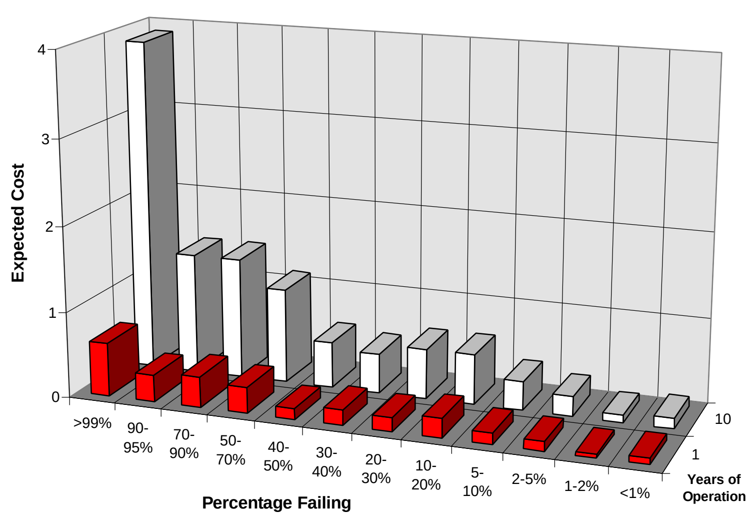

Figure 11: Source of the expected replacement-cost risk by ranges of percentages of components failing after one and ten years.

using the wide bin widths of Figures 10 and 11.

The incremental costs for this example are shown in Figure 13. Notice that the largest expected cost is in the first year. This indicates that if the entire fleet is to have problems, it is sure to be manifest early in the fleet operations. The early years of operation are the equivalent of a testing program providing the knowledge that is lacking in the absence of adequate prototype testing. The present values are not sensitive to the discount rate because of this heavy weighting on the early years of operation.

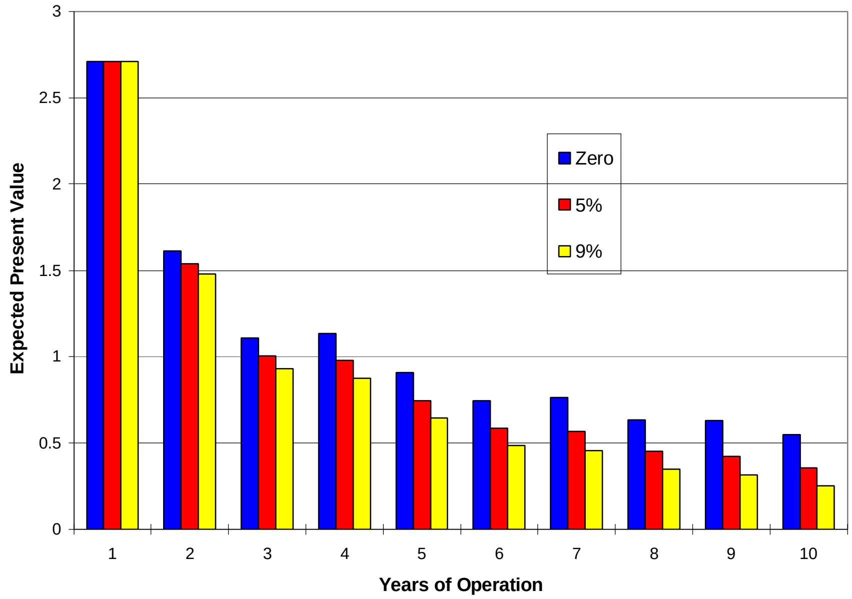

Figure 13: Present value of incremental costs (in two year increments) for the high uncertainty case using discount rates of zero, 5% and 9%. Total present values after 10 years of operation are 11, 9.4 and 8.5, respectively.

SUMMARY

The expected cost of replacing failed components during the operating life of a wind farm can be calculated in either of two ways. The easiest is to estimate the probability of individual component failure and multiply by the total replacement cost. This should not be confused with an equivalent percentage of the components failing. The percentage of components expected to fail is the probability of individual component failure only if the source of uncertainty responsible for the probability of failure is completely independent for each component. Common sources of uncertainty, shared by all turbines in the fleet, are also prevalent. A second approach is to separate (as well as possible) the sources of uncertainty into common and independent sources. The probability of different percentages of components failing is then estimated and the expected (average) cost of replacement is calculated. These two approaches result in the same expected total replacement cost if the individual replacement costs are independent of the number of components failing. More sophisticated cost models, reflecting a change in replacement costs when different fractions of the fleet have problems, will result in different expected costs from the two approaches. The simple first method would not be able to handle this variable cost case. The incremental costs in each year of operation estimated by either method are used to calculate the present value of replacement costs.

It is quite apparent that in the high uncertainty example the nature of the risk, components possibly failing early and in large numbers, is very different from the low uncertainty case of small numbers of components almost certainly failing gradually over time. In the low uncertainty case the risk is easily managed, while in the high uncertainty case the risk amounts to gambling with the short-term viability of the enterprise. It should be clear that an engineering effort to test and evaluate the design assumptions is necessary to advance a design from the high uncertainty of the conceptual stages to the lower uncertainty of a well engineered and tested machine.

REFERENCES

ASCE, 1982, "Fatigue Reliability: Introduction," by The Committee on Fatigue and Fracture Reliability of the Committee on Structural Safety and Reliability of the Structural Division, *Journal of the Structural Division*, American Society of Civil Engineers, Vol. 108, No. ST1.

Van Den Avyle, J. A., and H. J. Sutherland, 1989, "Fatigue Characterization of a VAWT Blade Material," *Eighth ASME Wind Energy Symposium*, SED Vol. 7, Berg and Klimas, eds., ASME, New York.

Veers, P. S., S. R. Winterstein, C. H. Lange and T. A. Wilson, 1994, *User's Manual for FAROW: Fatigue and Reliability of Wind Turbine Components*, SAND94-2460, Sandia National Laboratories, Albuquerque, New Mexico.