Paper for "ICEAA 2024 Professional Development & Training Workshop", May 14-16, 2024, Minneapolis, MN

Deriving total project costs from WBS elements' cost probability distributions

Rainald Kasprik
Professor of Business Administration and Quantitative Marketing
Heilbronn University of Applied Sciences, Reinhold-Würth-Hochschule
22 Daimlerstrasse, Kunzelsau, 74653, Germany

Abstract

Studies on possible cost variances in major acquisition projects focus on total project costs in order to come to plausible project budgets with a confidence level of 80%. Different lognormal probability distributions have been worked out representing different states of uncertainty. However, these models cannot be applied when using risk management software for deriving the total project costs based on cost probability distributions for WBS elements. Due to a limited processing capacity, risk management software demands a division of the underlying probability distributions into intervals. A simple discretization of the models developed to date is not possible, as these models contain unrealistic values in the tails. Based on simulation studies, three lognormal probability distributions are presented that meet these challenges. Finally, some practical hints are given on the minimum number of intervals, by which the curvature of the probability distribution is still represented, and on how to interpret the joint CDF's not-defined areas.

1 Problem Definition

The studies on possible cost variances in major acquisition projects have had the focus on regarding total project costs when modelling Cost Growth Factor (CGF) probability distributions. These studies aim at assisting project cost planners when determining the Value at Risk 80 (VaR80), demanded by U.S. Weapon Systems Acquisition Reform Act of 2009. Different lognormal probability distributions had been worked out (Garvey, 2008; Lee et al. 2012; Covert, 2013; Thomas & Fitch, 2014; McClary, 2021).

When calculating the variance of total project costs, which is based on the (additive) convolution of WBS elements' cost probability distributions (e. g. Garvey, 2014, 6), a challenge arises when realizing the calculation with support of risk management software. Then, the convoluting is based on a discretization of each underlying cost probability distribution function (PDF) due to limited processing capacity (see e. g. Smith, 1993; Woodruff & Dimitrov, 2018). Each PDF is partitioned in a small number of equal width intervals (Klugman et al., 2019, 517-519; Merkhofer, 1975, 72-77; Tanaka & Toda, 2013, 447), and each interval midpoint is used as data point when convoluting the PDFs.

Three challenges result from the approach of discretizing a CGF probability distribution:

(a) Inclusion of unrealistic events. In the above-mentioned models, empirical cost variances of total project costs are approximated with lognormal probability distribution where the input value lies near zero on the left side and is not limited on the right side. Therefore, functional values (probabilities) are given for events that are implausible to occur at all.

(b) Number of data points. There is less evidence regarding the discretization of lognormal probability distributions. Further, the research published focuses on the utilization of a very small number of intervals. It is not clear whether the published results can be applied.

(c) Interpretation of the joint CDF's not-defined areas. Mathematically, the range between the data points of a discretized PDF (and, therefore, of a joint CDF based on discretized PDFs) is a not-defined area. Therefore, the VaR80, which is identical with the percentile 80 in the case of a CGF probability function, is defined by the first value which exceeds the threshold, the fractile 0.8. However, when creating a joint CDF from discretized PDFs a closer approximation can be reached.

The outline of the paper is as follows: The following section 2 presents three CGF probability distributions representing different degrees of uncertainty, which overcome the challenges given above. Section 3 gives the results of a simulation study in which the fitting with the true PDF's curvature is analyzed. Section 4 gives, based on a fictitious example, practical hints for discretizing a PDF and, further, for interpreting the joint CDF. Section 5 summarizes the experiences made.

2 Truncated lognormal Cost Growth Factor-probability density functions

The task of modelling empirical experience in major acquisition projects mainly focuses on imaging the overrun of planned cost – and have integrated in their models CGF-values which are hardly expected to occur. When regarding planned costs of milestone C on an inflation-adjusted base year rate, the following CGF values can be described as extremely improbable: a CGF below 0.5 (Covert, 2013, 24; McClary, 2022, 31) or below 0.85 (Thomas & Fitch, 2014, A60), and a CGF above 3.0 (Covert, 2013, 24; Lee et al., 2012, 41).

When carrying out probabilistic project costing or, in general, a probabilistic scenario analysis (e. g. Smith, 1993; Woodruff & Dimitrov, 2018), the task of modelling is different: The total relevant event space must be mapped, disregarding extremely improbable events for ease of calculation. This means, there is a need for modelling a CGF probability distribution without regarding extremely improbable events.

Further, to the authors mind, a modeled CGF probability distribution should take into account that the base costs are determined to the best of a cost planner's knowledge. If, within a model, base costs are not regarded as most likely costs (see Table 1, Figure 1), a disappointment on the site of the cost planners can be expected. To avoid such an irritation and to reach a higher acceptance, a CGF probability distribution should be modeled which gives the most likely case the CGF value of one (and neither 0.8 times the base costs nor 1.3 times the base costs).

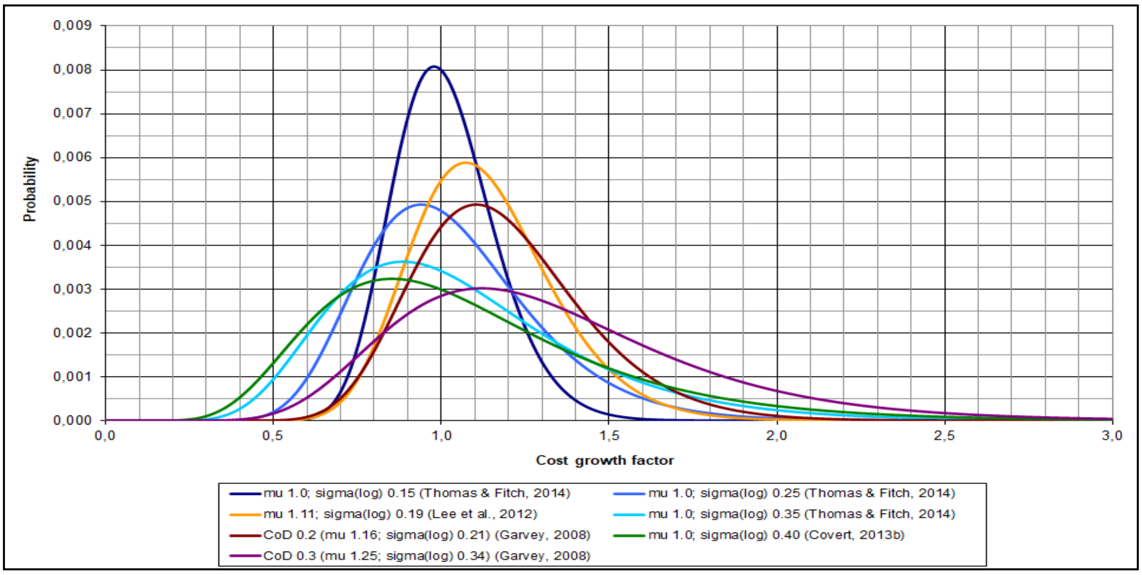

Figure 1: Modeled Cost Growth Factor probability distributions in research

Table 1: Key figures of Cost Growth Factor probability distributions

	Description	CGF "1" at Perc.:	Me-dian	Mode	Mean	P70	P80	P90	CV (natural scale)
mu 1.0; sigma(log) 0.15 (Thomas & Fitch, 2014)	low	50.5	1.001	0.977	1.011	1.082	1.136	1.211	0.151
mu 1.0; sigma(log) 0.25 (Thomas & Fitch, 2014)	medium	50.3	1.001	0.941	1.032	1.142	1.235	1.379	0.254
mu 1.11; sigma(log) 0.19 (Lee et al., 2012)	not defined	29.2	1.109	1.073	1.129	1.226	1.298	1.409	0.188
mu 1.0; sigma(log) 0.35 (Thomas & Fitch, 2014)	high	50.2	1.001	0.884	1.061	1.202	1.343	1.565	0.357
CoD 0.2 (mu 1.16; sigma(log) 0.21) (Garvey, 2008)	low	25.3	1.157	1.103	1.183	1.295	1.385	1.523	0.217
mu 1.0; sigma(log) 0.40 (Covert, 2013b)	high	50.2	1.001	0.854	1.073	1.235	1.400	1.670	0.403
CoD 0.3 (mu 1.25; sigma(log) 0.34) (Garvey, 2008)	medium	25.2	1.253	1.121	1.311	1.496	1.664	1.928	0.333
mu 1.4; sigma(log) 0.27 (McClary, 2021)	not defined	8.6	1.451	1.349	1.501	1.673	1.823	2.048	0.269
CoD 0.4 (mu 1.37; sigma(log) 0.47) (Garvey, 2008)	high	25.1	1.370	1.103	1.356	1.748	2.027	2.489	0.425

The first goal can be achieved by modeling a truncated lognormal PDF that ignores extremely improbable events. Setting of the location parameter mu of a lognormal CGF probability distribution above the value one, and so disregarding the questionable assumption that the baseline costs would always represent the median of a lognormal CGF probability distribution, can reach the second goal. This is founded by the observations of Coleman et al. (2009): They have shown that the planned costs (CGF=1) mainly lay below the median and that a setting of the location parameter mu above one is therefore indicated.

To achieve these goals, the author developed various lognormal probability distributions. Methodologically, a simulation model was created in which the input variable of a standard lognormal distribution was set between the values of 0.75 and 3.0, and the parameter mu above the value 1.0. The input values were discretized in 1000 equal width intervals and the probability distribution was derived for these intervals deploying Microsoft Excel 2016. Different simulations were run by varying the parameters mu and sigma through that an iteration to the numerical values for VaR80 and VaR90, given in Thomas & Fitch (2014, 32) for low, medium, and high uncertainty, were achieved. The decision criterion for the goodness of the solution was the minimum deviation of the sum of the absolute differences from the difference between the approximated and the given values of VaR80 and, further, VaR90.[1]

Table 2 gives the results achieved. Figure 2 illustrates the result by displaying the CGF in comparison to the CGFs of the distributions given in Table 1.

The author proposes a review of these results based on the published empirical SAR data, assuming that the CDF is derived for relevant bandwidth between 0.75 and 3.0.

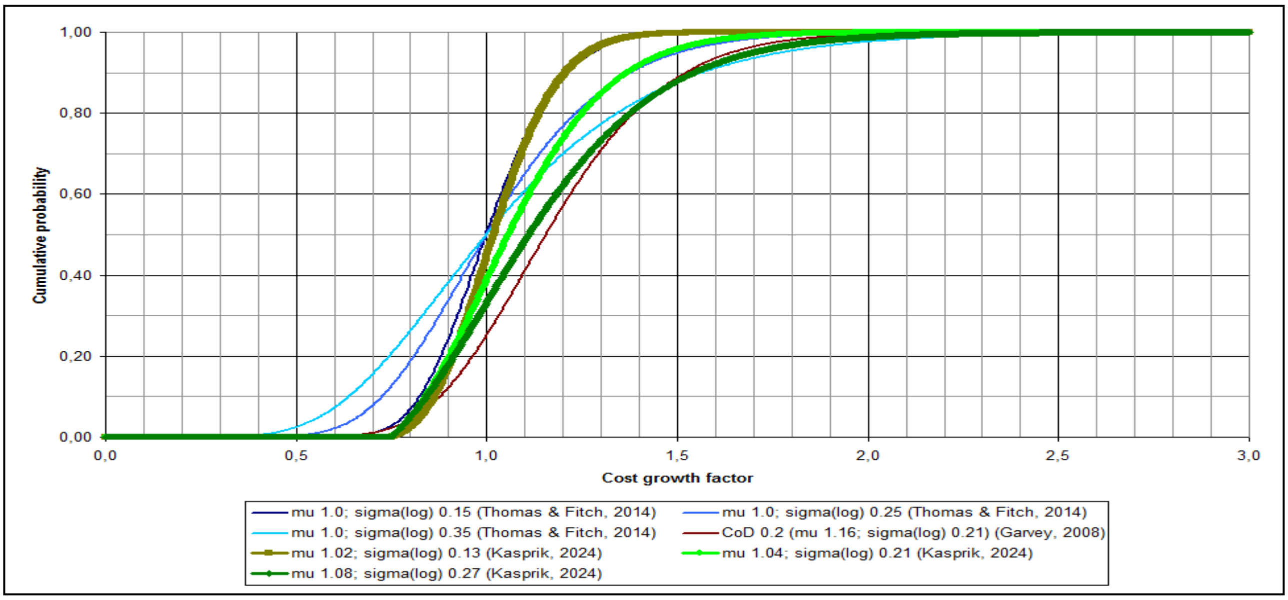

Figure 2: Cumulative distribution functions of the Cost Growth Factor probability distributions developed

Table 3: Key figures of Cost Growth Factor probability distributions developed

	Description	CGF "1" at Perc.:	Median	Mode	Mean	P70	P80	P90	CV (natural scale)
mu 1.02; sigma(log) 0.14 (Kasprik, 2024; truncated data)	low	44.8	1.019	1.001	1.030	1.094	1.139	1.208	0.130
mu 1.04; sigma(log) 0.21 (Kasprik, 2024; truncated data)	medium	38.8	1.061	1.001	1.089	1.175	1.253	1.370	0.192
mu 1.08; sigma(log) 0.27 (Kasprik, 2024; truncated data)	high	33.4	1.112	1.001	1.162	1.268	1.376	1.544	0.245

[1] This follows the "Method-of-Moments-Matching" approach (Shore, 1995; Au-Yeung, 2003).

3 The appropriate number of equal width interval when discretizing a probability density function

Previous research has focused on how to discretize a probability distribution into three to five data points without the properties of the PDF deviating from the true PDF (Barbiero & Hitaj, 2021; Hadlock, 2017; Hammond & Bickel, 2015; Tanaka & Toda, 2013, 447; Woodruff & Dimitrov, 2018). It is not clear whether the published results can be applied in particular in the case of a lognormal distribution, as is the case when working with CGF distributions.

To recognize the general conditions, a pre-study was conducted. The aim was to analyze the effects of discretizing a probability distribution into five equally wide intervals covering the entire bandwidth. The effects were studied on the lognormal distribution presented in section 2 modelling a medium uncertainty (mu [natural scale] = 1.044; sigma [log scale] = 0.207; input values between 0.75 and 3.0) as well as for a normal distribution also representing a medium uncertainty. A normal distribution representing a medium uncertainty can be described as possessing a Coefficient of Variation (CV) of 0.27 (Thomas & Fitch, 2014, 32). A normal distribution within the bandwidth of plus / minus 3.66 z reaches that CV-value if the input-values are transformed to represent a CGF between zero and two through which the CGF-value of one (the baseline costs) comes to lay at the mode of the distribution. This transformation is done by following the methodology of normalizing. For performing the calculations, the lognormal distribution is discretized in 1000 (equal width) intervals, the normal distribution in 999 (equal width) intervals. The calculations are performed with Microsoft Excel 2016.

A visual inspection has shown (and has further been supported by the goodness of fit measures Goodness of Variance Fit,[2] Kolmogorov Distance and Total Variation Distance) that the Five-interval-solution cannot be regarded as appropriate image of the true probability density function.[3] This is because the most relevant part of the distribution, i. e. the range in which the strongest change in the rate of direction occurs (here the marked range between red columns), is covered by just three to four intervals leading to data points that dramatically equalize the curvature, especially in case of the lognormal distribution (Figure 3 and 4).

Inspired by this result, the effect of excluding the outer input-values by truncating the PDF at certain percentiles is examined to explore whether an appropriate Five-interval solution could be found.

Both types of probability distributions from the preliminary study were selected, and two further degrees of uncertainty were modeled. (a) A low degree of uncertainty: a lognormal distribution with mu [natural scale] = 1.017, sigma [log scale] = 0.133 (with input values between 0.75 and 3.0), a normal distribution with z-values (plus / minus 5.14) leading to a CV of 0.19. (b) A high degree of uncertainty: a lognormal distribution with mu [natural scale] = 1.078, sigma [log scale] = 0.27 (with input values between 0.75 and 3.0), a normal distribution with z-values (plus / minus 2.59) leading to a CV of 0.37.

---

[2] The Goodness of Variance Fit (GVF) measure is designed to measure the classification's goodness when regarding the total bandwidth via the internal (within-class) homogeneity (Armstrong et al., 2003; Schiewe, 2018; Smith, 1986)

[3] Gained values for Lognormal PDF: Goodness of Variance Fit: 0.692; Kolmogorov Distance (midpoints): 0.40; normalized TVD (data points): 0.139. Gained values for Standard Normal PDF: Goodness of Variance Fit: 0.848; Kolmogorov Distance (midpoints): 0.266; normalized TVD (data points): 0.133.

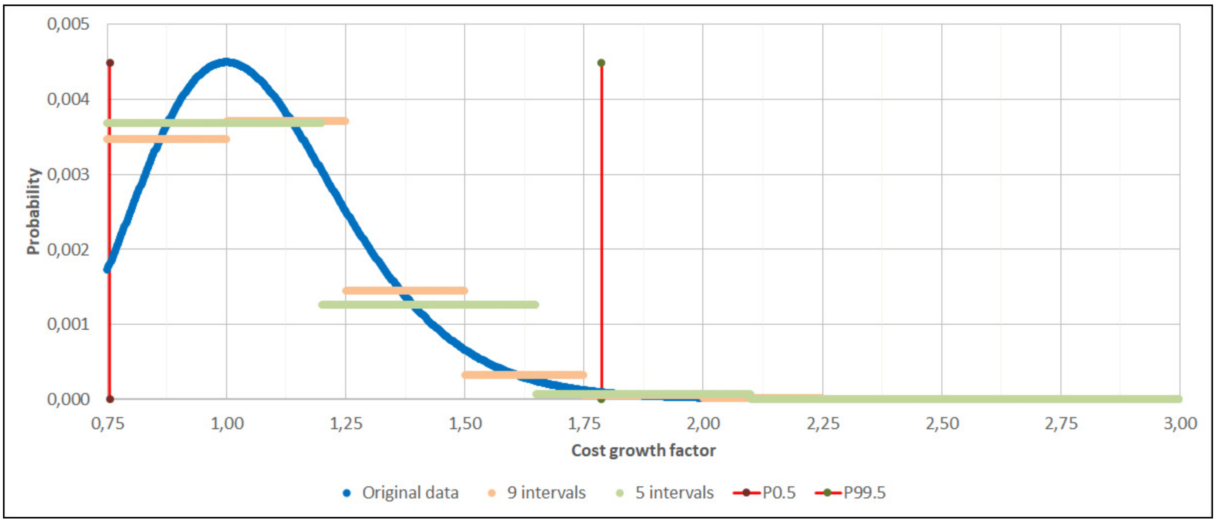

Figure 3: Nine- and Five-interval discretization - over the whole bandwidth (lognormal distribution)

Figure 4: Nine- and Five-interval discretization - over the whole bandwidth (normal distribution)

The heuristic, which was found in the pre-study (truncating the probability distribution at fractile 0.005 and fractile 0.995), is cross-checked by truncating the probability distribution at fractile 0.025 and fractile 0.975. When truncating the probability distribution at fractile 0.025 and fractile 0.975, the positive effect of an increased number of intervals within the range, in which higher rates of change of direction predominantly occur, could be counteracted by an inacceptable increase of the interval probability. The reason for this is that the number of events in the event space is further reduced, which leads to a higher probability per event (data point).

The discretization's goodness of fit is measured by calculating the (normalized) „Total Variation Distance" (TVD) (Gibbs & Su, 2002, 424; Markatou & Sofikitou, 2019, 8). For each data point, the estimated functional value is subtracted from the true value, is then set in absolute numbers and finally summed. The sum of all differences gives TVD. TVD measures the total (and unweighted) deviation derived from each data point's functional value. When evaluating a probability function, the normalization of the TVD allows to determine the extent to which the estimated probability of all data points deviates from the true probability function in relation to the average probability of occurrence. The normalized TVD value of one represents an average deviation of 1 time the average probability per approximated data point, the value of zero for an average deviation of zero times the average probability (a perfect fit). In order to draw this conclusion, the same number of data points for the approximation is required as is given in the true probability distribution. This is achieved by following the assumption of uniformity that each data point within an interval has the same probability of occurrence (Anderson et al., 2007, 136-138; Klugman et al., 2019, 517-518). A data point's approximated probability is derived by averaging the interval probability over the number of data points in each interval. The author chooses 0.1 as critical value for the normalized TVD – a deviation of 10% the average probability.

Secondly, the Kolmogorov distance (KD, also "Kolmogorov-Smirnov Distance") is chosen. It compares a true cumulative probability of a data point with its approximated cumulative probability by subtracting the approximated cumulative probability from its true cumulative probability (Barbiero & Hitaj, 2021, 10; Cohen et al., 2018; Drakakis & Radulovic, 2018; Rohatgi & Saleh, 2015, 585; Wilcox, 2012, 25-26). The maximum difference in absolute numbers gives KD. KD reaches the value one in the (hypothetical) case when a data point's cumulative probability is one and the corresponding value reaches zero (Wilcox, 2012, 26). KD captures not only the differences between the comparison points but also the impact of consecutive deviations of the same sign with regard to the ordered attributes of the variable. When comparing a discretized probability distribution with the true probability distribution, the cumulative probability of the interval midpoint is compared with the cumulative probability of the corresponding data point. KD values between 0.1 and 0.3 are reported in discretization studies (Keefer & Bodily, 1983, 603; Smith, 1993, 345). The author chooses 0.15 as the lowest acceptable deviation.

For informational purposes, the rate of divergence between the first moments of the distributions (here: the expected value) is calculated: the difference between the probability-weighted interval midpoints to the true expected value, which is divided through true expected value ("Percentage Error Mean" [PEM]; Miller & Rice, 1983; Shore, 1995; Hadlock, 2017). This is done regarding the possible situation that the expected value is used as a decision criterion. PEM reaches the value of one in case of identical means. The author takes a critical view of the applicability of this criterion when assessing the goodness of a discretization because, when using this measure, a divergence in PEM is not only influenced by a divergence in the probability of the data point, but also by the numerical value of this data point. An identical difference between the approximated and the true probability of a data point leads to a higher divergence of the measure PEM in the case that the difference occurs apart from the true expected value. In the author's opinion, however, the goodness of a discretization should be measured independently of the position of the discrepancy on the bandwidth.

The results of discretizing these six variants in 5, 7, and 9 equal width intervals within the range between fractile 0.005 and fractile 0.995 and, further, fractile 0.025 and fractile 0.975 are given in appendix 1. They show that under both conditions, following a truncation at the fractile 0.005 and 0.995 resp. 0.025 and 0.975, seven equal width intervals can be regarded as appropriate minimum number of intervals. The critical values of TVD and KD are not exceeded.

Figures 5 and 6 show the 5- and 7-interval solutions for the situation of discretizing the lognormal probability distribution, which range shows the strongest change in the rate of direction in a small partition of the bandwidth (the situation of low uncertainty).

However, it should be noted that this study focuses on the curvature of the distribution and that some other criteria may require a higher number of intervals.

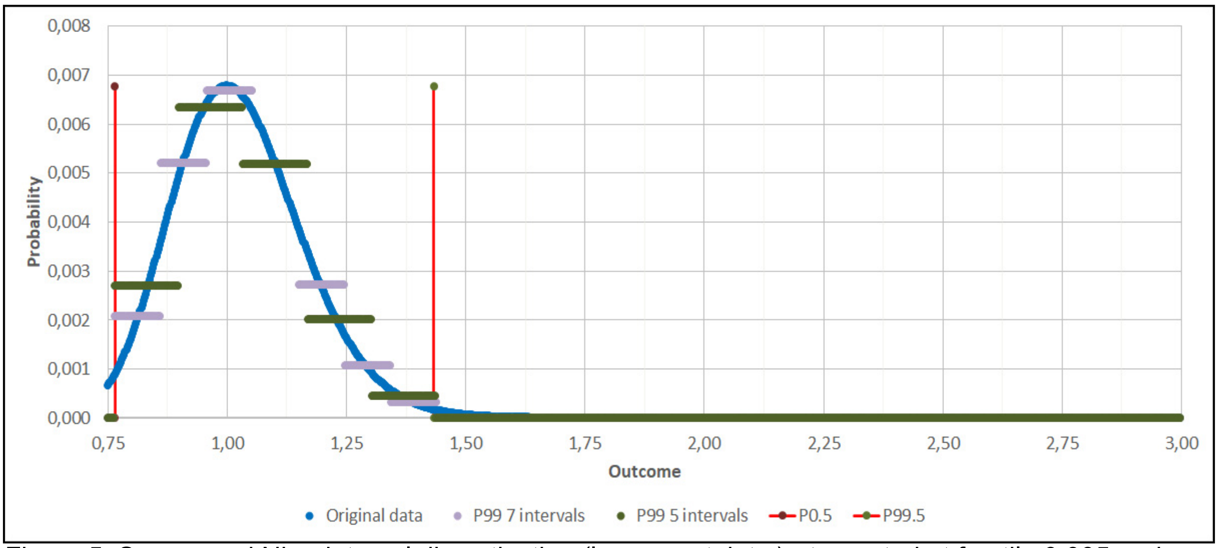

Figure 5: Seven- and Nine-interval discretization (low uncertainty) - truncated at fractile 0.005 and 0.995

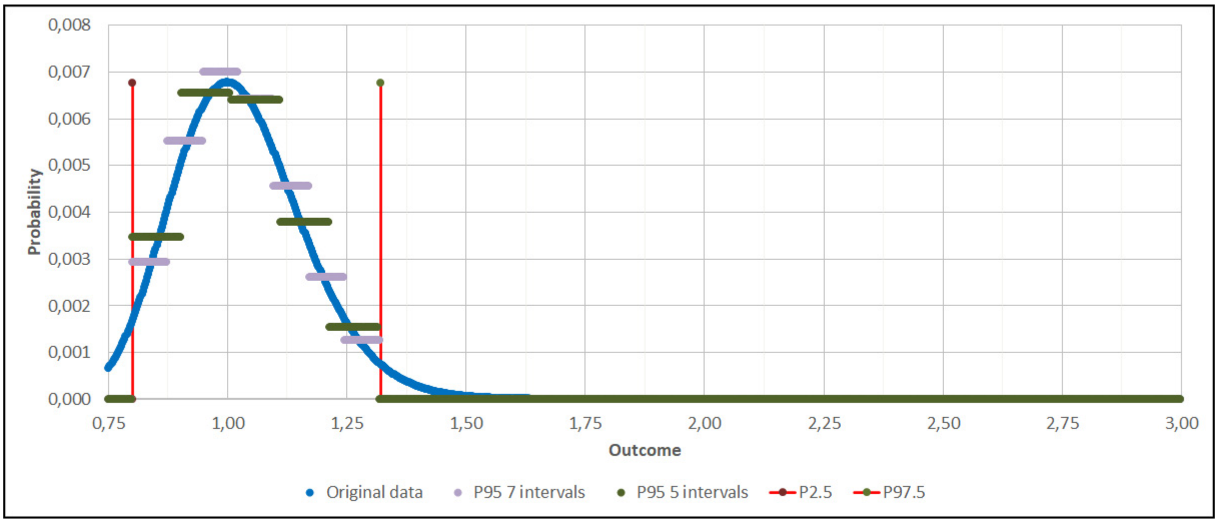

Figure 6: Seven- and Nine-interval discretization (low uncertainty) - truncated at fractile 0.025 and 0.975

Firstly, the interpretability of the joint distribution data points. When creating a joint PDF from discretized PDFs one should note that each data point is the midpoint of an (equal width) interval. When additively convoluting PDFs, each joint distribution data point represents the midpoint of an interval bounded by the sum of all lower respective of all upper bounds of the PDF intervals integrated (Nguyen, 2012, 180). Then it could be that the resulting interval width of the joint PDF becomes too wide so that the value of the information obtained is unacceptably diminished.

Secondly, the position of the true mode within the mode interval. If the true mode (the baseline costs) does not come to lay at the interval midpoint of the mode interval, a remarkable distortion of the result must be expected. In that case, a flattening effect occurs leading to an increase in probability of the neighbor intervals. Even if that aspect showed to be acceptable with regard to the study above, it might, nevertheless, not be acceptable when taking into account the consequences on the joint distribution. To meet this challenge, the author suggests either an increase of the number of intervals or a shift of the relevant bandwidth upwards or downwards the probability distribution to let the mode become the midpoint of the mode interval. When selecting a bandwidth between fractile 0.005 and 0.995 as relevant bandwidth, then the shifted new relevant bandwidth will still not exceed the critical lower fractile of 0.025 or undercut the critical upper fractile 0.995. The input- and output-values for the shifted solution are given in appendix 2 to 4 for the three CGF probability distributions developed in section 2.

4 How to interpret the joint CDF's not-defined areas?

Assume that for three WBS elements the costs have been estimated, representing the most likely costs. Due to uncertain future circumstances, the cost variance of each WBS element has been assessed to follow a different lognormal probability distribution. The situation for WBS element 1 is that of a minimum uncertainty (representing e. g. a known technology and a low complexity), for WBS element 2 that of a medium uncertainty, and for WBS element 3 that of a maximum uncertainty.

The lognormal PDFs presented in section 2 are applied. In order to come to a sufficiently accurate estimate, the practical approach of discretizing the PDFs and of calculating the joint distribution with the help of a risk management software is chosen. As planned costs were determined for WBS element 1: Currency Units [CU] 2000, for element 2: CU 500, for element 3: CU 1000. The relevant bandwidth is seen as laying between the minimum (0.75 the planned costs) and fractile 0.999 for each CGF probability distribution.

The decider aims at an identical interval width per probability distribution for ease of interpretation: then, the cost share of each WBS element within the joint distribution interval is identical.[4] He or she chooses an interval width of 90.48 CU in order to reach more than five intervals in the PDF with lowest bandwidth (WBS element 2) between minimum and fractile 0.999 and, further, to come to an interpretable joint distribution interval width after an additive convolution (271.43 CU).

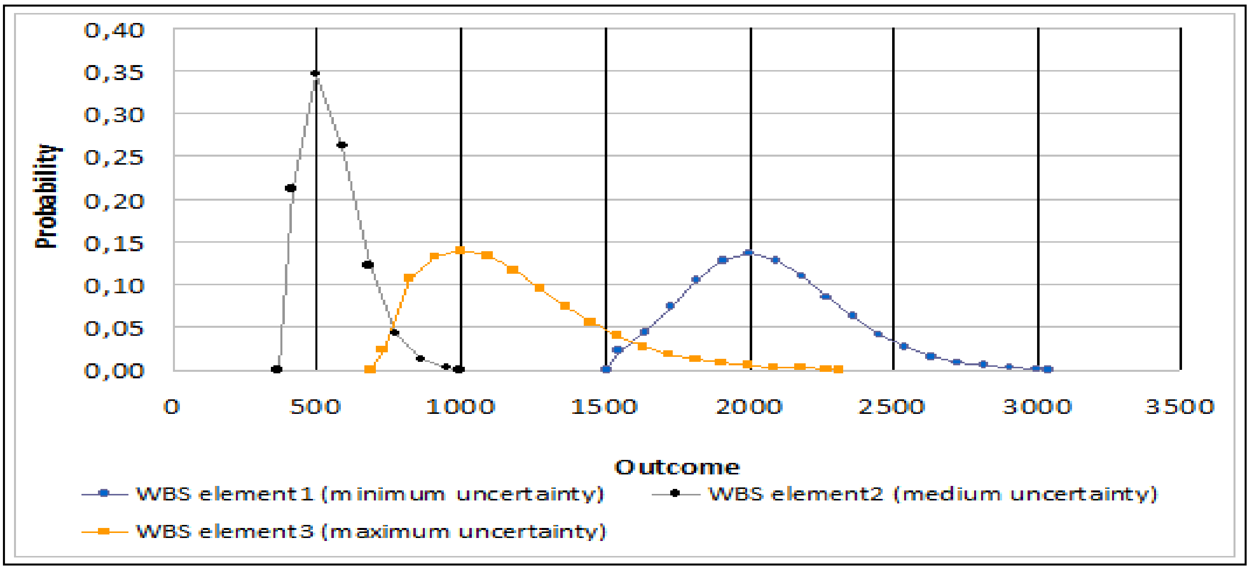

Figure 7: Cost probability distribution of the three elements regarded

[4] It should be mentioned that an identical interval width across all underlying PDFs drastically reduces the number of joint distribution data points, as then one data point results from a higher number of combinations. A second supporting condition is that the input values are calculated with eight decimal places. Both ways allow a full enumeration even if the calculated number combinations would theoretically exceed the processing capacity. In the example presented, the number combinations is (17*7*18) 2142. However, the joint distribution data points are 40.

Next, the chosen interval width is cross-checked for the resulting number of intervals for the other two WBS elements (17 for WBS element 1 and 18 for WBS element 3) with regard to the data processing capacity of the software deployed.[5]

A challenge arises for WBS element 3: the true mode (CU 1000) comes to lay slightly apart the midpoint of the mode interval (CU 974.66) is resolved by shifting the relevant bandwidth to the left, which leads to a negligible reduction of the upper limit (the PDF is truncated at fractile 0.997). Figure 7 shows the PDFs of the WBS elements' cost variances.

Cost variances are seen as independent of each other. The additive convolution is done with the risk management software DPL 9 of Syncopation Software. Figure 8 shows the software output of the cumulative density function (CDF). The value of CU 4042.92 represents VaR80 and an interval (the sum of all lower respective of all upper bounds of the PDF intervals integrated) between CU 3907.21 and CU 4178.64. The (point) estimates are for expected value: CU 3759.85, for VaR50: CU 3771.50, and for the most likely costs: CU 3681.02.

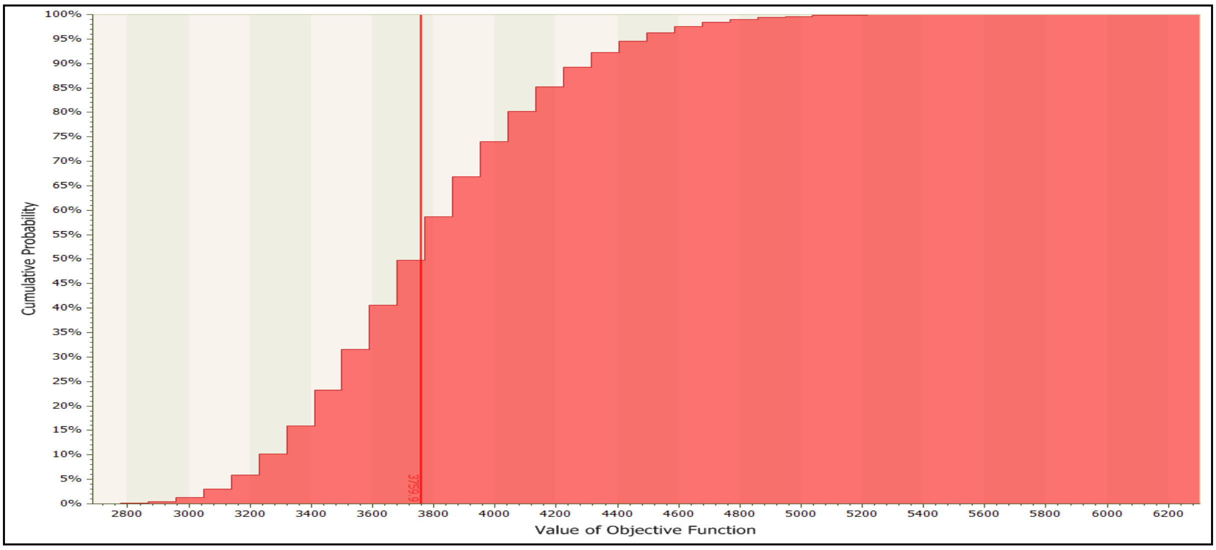

Figure 8: Joint cumulative density function

A challenge arises out of the fact that, mathematically, the distance between the joint distribution data points is a not-defined area. If a decider wishes to come to a closer solution, he or she can take into account that in the case of a single discretized PDF the cumulative probability of an interval midpoint represents the sum of the probabilities of the class members. This means, the right bound of each convoluted PDF interval gives the true cumulative probability and not the interval midpoint.

Then, given continuous PDFs which are discretized and, further, given an additive convolution, the assumption can be made that this holds for the joint distribution interval as well because of its equal interval width in cases of an additive convolution. Therefore, it seems appropriate to interpolate the right bounds of the joint intervals. In the example presented here, the width of each joint distribution interval is CU 90.48. The value of CU 4088.16 then represents the VaR80 (Figure 9). It is right to the midpoint (CU 4042,92). Its interval (based on the sum of all lower respective of all upper bounds of the PDF intervals integrated) covers the values between CU 3907.21 and CU 4178.64.

---

[5] The number of computable combinations is bounded in the case of a full enumeration. The number depends on the resulting number of the joint distribution data points (between 750 and 1000).

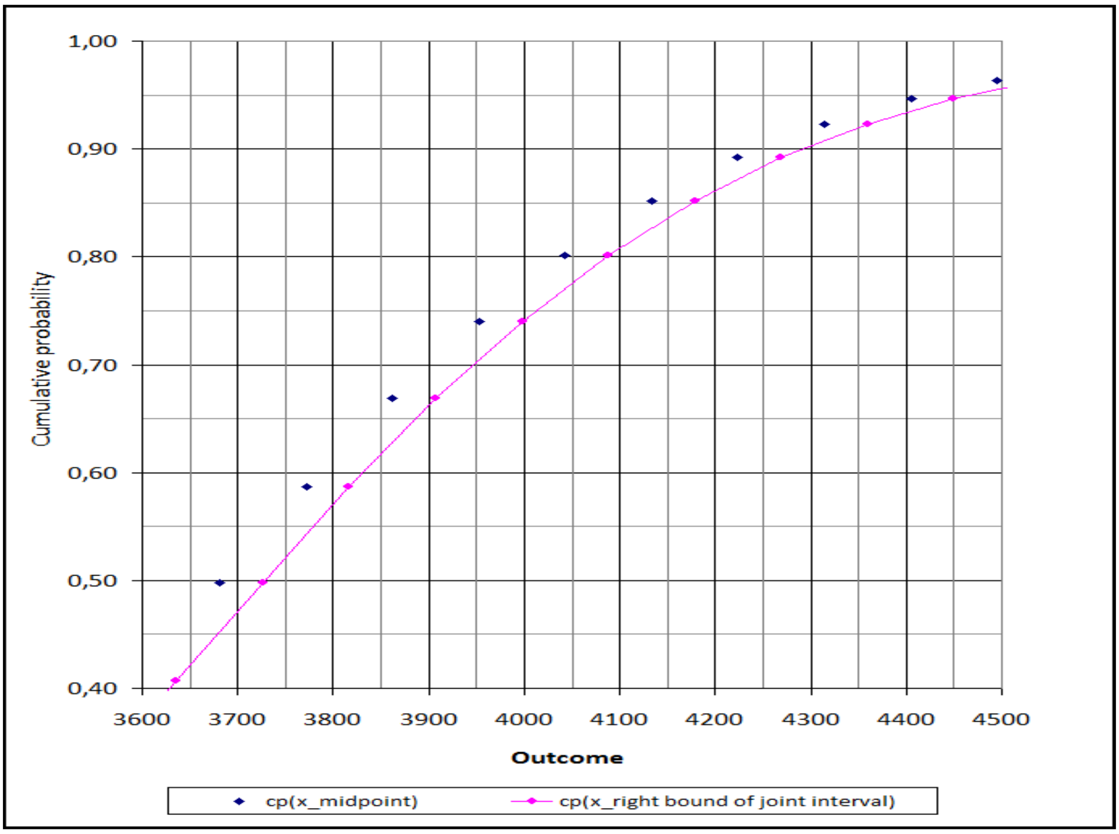

Figure 9: Cumulative probability of interval midpoints and of the joint interval right bounds

5 Lessons learned when working with risk management software

1. Test the maximum number of combinations that a risk management software can calculate if you plan to deploy a full enumeration of all combinations (preferable).
2. Choose equal interval widths per PDF. Otherwise, the probability of each data point is not comparable: The probability of an interval depends on the number of the class members (data points), which are condensed into one interval.
3. Start the decision process on the number of intervals by looking at the resulting interval width of the joint distribution data points: Probabilistic project costing represents a situation of an additive convolution. Therefore, a joint distribution data point represents the midpoint of an interval that is bounded by the sum of all lower respective of all upper bounds of the PDF intervals integrated.
4. Choose an identical interval width for each PDF: Then (a) the cost variance for each WBS element is of identical importance within the joint distribution interval and (b) the number of joint data points decreases drastically.
5. Make the midpoint of the discretized mode interval get close to the true mode of the PDF by shifting the relevant bandwidth: The midpoint of the mode interval influences the resulting joint distribution more than other data points.
6. In the case of the lognormal CGF distributions presented: the minimum number of equal width intervals is seven if truncated between fractile 0.005 and 0.995 or fractile 0.025 and fractile 0.975.

7. Interpret the joint CDF of additively convoluted PDFs as follows: Assess the distance of the joint distribution data points and shift the CDF half the distance to the right. The CDF data point given at each interval's right bound indicates the true CDF, and a linear interpolation between these points approximates the true curvature of the joint CDF.

Acknowledgements

This paper was partially funded by the Foundation for the Promotion of the Reinhold-Würth-Hochschule, Künzelsau (grant number 17/1/1.1).

Conflicts of interest

The author declares no conflict of interest. The funders had no role in the design of the study, in the collection, analysis or interpretation of data, in the writing of the manuscript, or in the decision to publish the results.

Literature

Anderson, David R., Sweeney, Dennis J., and Williams, Thomas A. (2007): Essentials of Modern Business Statistics. Mason, OH: Thomson.

Armstrong, Marc P., Xiao, Ningchuan, and Bennett, David A. (2003): Using genetic algorithms to create multicriteria class intervals for choropleth maps. Annals, Association of American Geographers, Vol. 93, No. 3, pp. 595–623.

Au-Yeung, Susanna W. M. (2003): Finding Probability Distributions From Moments. Master Thesis (supervisor: William J. Knottenbelt). London, UK: University of London.

Barbiero, Alessandro, and Hitaj, Asmerilda (2021): Approximation of continuous random variables for the evaluation of the reliability parameter of complex stress–strength models. Annals of Operations Research. Open access; https://link.springer.com/article/10.1007/s10479-021-04010-6.

Cohen, Liat, Fried, Dror, and Weiss, Gera (2018): An optimal approximation of discrete random variables with respect to the Kolmogorov distance. eprint arXiv:1805.07535. https://arxiv.org/abs/1805.07535.

Coleman, Richard L., Druker, Eric R., Braxton, Peter J., Cullis, Bethia L., and Kanick, Christina M. (2009): What Percentile Are We At Now (And Where Are We Going?). International Cost Estimating & Analysis Association, IPSA/SCEA Joint Annual Conference, 2-5 June 2009, St. Louis, MO.

Covert, Raymond W. (2013): Analytic Method for Cost and Schedule Risk Analysis. Report for National Aeronautics and Space Administration (NASA). Washington, DC: NASA Cost Analysis Division.

Drakakis, Konstantinos, and Radulovic, Dragan (2018). Results on the multiresolution structure of Internet traffic traces. Scalability and Traffic Control in IP Networks II, Vol. 4868, pp. 234-244. Revised version 15 October 2018.

Garvey, Paul R. (2008): A scenario-based method for cost risk analysis. Journal of Cost Analysis and Parametrics. Vol. 1, Issue 1, pp. 65-76.

Garvey, Paul R. (2014): Elements of Cost Risk Analysis. Report No. MP130729. Bedford, MA: The MITRE Corporation.

Gibbs, Alison L., and Su, Francis Edward (2002): On choosing and bounding probability metrics. International Statistical Review. Vol. 70, No.3, pp. 419–435.

Hadlock, Christopher C. (2017): Quantile-Parameterized Methods for Quantifying Uncertainty in Decision Analysis. Dissertation. Austin, TX: University of Texas at Austin.

Hammond, Robert K., and Bickel, J. Eric (2015): Discretization methods for continous probability distributions. In: CochranJames J. (ed.): Wiley Encyclopedia of Operations Research and Management Science, online- edition. Chichester, UK: John Wiley & Sons Ltd, pp. 1-13.

Klugman, Stuart A., Panjer, Harry H., and Willmot, Gordon E. (2019): Loss models. Fifth edition. Hoboken, NJ: John Wiley & Sons.

Lee, Richard C., Braxton, Peter J., Cincotta, Kevin M., and Flynn, Brian J. (2012): SAR Data Analysis, CV Benchmarks, and the Updated NCCA S-Curve Tool. Paper prepared for SCEA/ISPA Conference, 26-29 June 2012 in Orlando, FL. Annandale, VA: International Cost Estimating & Analysis Association.

Markatou, Marianthi, and Sofikitou, Elisavet M. (2019): Statistical Distances and the Construction of Evidence Functions for Model Adequacy. Frontiers in Ecology and Evolution, Section Environmental Informatics and Remote Sensing. Vol. 7, Article 44.

McClary, Matthew A. (2021): Methodology for modeling cost and schedule risk associated with resource decision involving the U.S. army's modernization efforts for 2023. Bachelor Thesis (thesis advisor: Gregory K. Mislick). Monterey, CA: Naval Postgraduate School.

Merkhofer, Miley W. (1975): Flexibility and decision analysis. Dissertation. Stanford, CA: Stanford University.

Miller, Allen C., and Rice, Thomas R. (1983): Discrete Approximations of Probability-Distributions. Management Science. Vol. 29, pp. 352-362.

Nguyen, Hong Diep (2012). Efficient Implementation of Interval Matrix Multiplication. In: Jónasson, Kristján (ed.): Applied Parallel and Scientific Computing. 10th International Conference, PARA 2010, Reykjavík, Iceland, June 6-9, 2010. Berlin and Heidelberg, DE: Springer, pp. 179–188.

Randolph, KaDonna C., and Seaver, Bill (2007): An Alternative to Traditional Goodness-of-Fit Tests for Discretely Measured Continuous Data. Forest Science, Vol. 53, No. 5, pp. 590-599.

Rohatgi, Vijay K., and Saleh, A. K. Md. Ehsanes (2015): An Introduction to Probability and Statistics. Third edition. Hoboken, NJ: John Wiley & Sons.

Schiewe, Jochen (2018): Development and comparison of uncertainty measures in the framework of a data classification. The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences, Vol. 52, No. 4, pp. 551-558.

Shore, Haim (1995): Fitting a distribution by the first two moments (partial and complete). Computational Statistics & Data Analysis. Vol. 19, pp. 563-577.

Smith, Richard M. (1986): Comparing traditional methods for selecting class intervals on choropleth maps. Professional Geographer, Vol. 38, No. 1, pp. 62-67.

Smith, James E. (1993): Moment Methods for Decision Analysis. In Decision Science, Vol. 39, No. 3, pp. 340-358.

Tanaka, Ken'ichiro, and Toda, Alexis Akira (2013): Discrete approximations of continuous distributions by maximum entropy. Economics Letters, Vol. 118, pp. 445–450.

Thomas, Duncan, and Fitch, John (2014): Joint Agency Cost Schedule Risk and Uncertainty Handbook. Washington, DC: Naval Center for Cost Analysis.

Wilcox, Rand (2012): Introduction to Robust Estimation and Hypothesis Testing. Third edition. Cambridge, MA: Elsevier.

Woodruff, Joshua, and Dimitrov, Nedialko B. (2018): Optimal discretization for decision analysis. Operations Research Perspectives. Vol. 5, pp. 288-305.

Appendix 1: Goodness of fit

Low uncertainty, truncated between fractile 0.005 and fractile 0.995	**TVD (normalized, per data point)**	**KD (interval midpoint)**	**PEM (interval midpoint)**
P99 9 intervals	0.0568	0.1044	-0.0010
P99 7 intervals	0.0708	0.1390	-0.0006
P99 5 intervals	0.0992	0.2015	-0.0005
Low uncertainty, truncated between fractile 0.025 and fractile 0.975	**TVD (normalized, per data point)**	**KD (interval midpoint)**	**PEM (interval midpoint)**
P95 9 intervals	0.0666	0.0887	-0.0024
P95 7 intervals	0.0767	0.1092	-0.0030
P95 5 intervals	0.0924	0.1523	-0.0031
Medium uncertainty, truncated between fractile 0.005 and fractile 0.995	**TVD (normalized, per data point)**	**KD (interval midpoint)**	**PEM (interval midpoint)**
P99 9 intervals	0.0500	0.1087	-0.0020
P99 7 intervals	0.0620	0.1476	-0.0013
P99 5 intervals	0.0879	0.1830	-0.0002
Medium uncertainty, truncated between fractile 0.025 and fractile 0.975	**TVD (normalized, per data point)**	**KD (interval midpoint)**	**PEM (interval midpoint)**
P95 9 intervals	0.0634	0.0847	-0.0079
P95 7 intervals	0.0703	0.1108	-0.0078
P95 5 intervals	0.0859	0.1570	-0.0070

High uncertainty, truncated between fractile 0.005 and fractile 0.995	**TVD (normalized, per data point)**	**KD (interval midpoint)**	**PEM (interval midpoint)**
P99 9 intervals	0.0471	0.1220	-0.0023
P99 7 intervals	0.0604	0.1480	-0.0013
P99 5 intervals	0.0844	0.2165	0.0013
High uncertainty, truncated between fractile 0.025 and fractile 0.975	**TVD (normalized, per data point)**	**KD (interval midpoint)**	**PEM (interval midpoint)**
P95 9 intervals	0.0613	0.0882	-0.0107
P95 7 intervals	0.0689	0.1168	-0.0103
P95 5 intervals	0.0808	0.1539	-0.0081

Appendix 2: Low uncertainty – CGF data points

	CGF (interval midpoint)	**p(interval)**			**CGF (interval midpoint)**	**p(interval)**
Lower limit	*0.753303571429*			*Lower limit*	*0.810625000000*	
Interval1	0.809071428571	0.083292		**Interval1**	0.777250000000	0.037401
Interval2	0.904535714286	0.217604		**Interval2**	0.851500000000	0.109998
Interval3	1.000000000000	0.281188		**Interval3**	0.925750000000	0.189059
Interval4	1.095464285714	0.230165		**Interval4**	1.000000000000	0.222305
Interval5	1.190928571429	0.122039		**Interval5**	1.074250000000	0.192510
Interval6	1.286392857143	0.050211		**Interval6**	1.148500000000	0.130033
Interval7	1.381857142857	0.015501		**Interval7**	1.222750000000	0.071682
Upper limit	*1.429589285714*			**Interval8**	1.297000000000	0.033432
				Interval9	1.371250000000	0.013580
				Upper limit	*1.408375000000*	

Appendix 3: Medium uncertainty – CGF data points

	CGF (interval midpoint)	p(interval)			CGF (interval midpoint)	p(interval)
Lower limit	0.776285715286			*Lower limit*	0.827125000000	
Interval1	0.852464285714	0.219510		**Interval1**	0.775000000000	0.085396
Interval2	1.000000000000	0.298819		**Interval2**	0.887500000000	0.188992
Interval3	1.147535714286	0.239881		**Interval3**	1.000000000000	0.223481
Interval4	1.295071428571	0.141960		**Interval4**	1.112500000000	0.197148
Interval5	1.442607142857	0.065352		**Interval5**	1.225000000000	0.140418
Interval6	1.590142857143	0.025373		**Interval6**	1.337500000000	0.085343
Interval7	1.737678571429	0.009105		**Interval7**	1.450000000000	0.046052
Upper limit	1.811446428571			**Interval8**	1.562500000000	0.022711
				Interval9	1.675000000000	0.010459
				Upper limit	*1.731250000000*	

Appendix 4: High uncertainty – CGF data points

	CGF (interval midpoint)	p(interval)			CGF (interval midpoint)	p(interval)
Lower limit	0.898428572429			*Lower limit*	0.763000001000	
Interval1	0.799750000000	0.181309		**Interval1**	0.844750000000	0.200893
Interval2	1.000000000000	0.308178		**Interval2**	1.000000000000	0.244716
Interval3	1.200250000000	0.249875		**Interval3**	1.155250000000	0.214116
Interval4	1.400500000000	0.146877		**Interval4**	1.310500000000	0.151043
Interval5	1.600750000000	0.071080		**Interval5**	1.465750000000	0.092272
Interval6	1.801000000000	0.030515		**Interval6**	1.621000000000	0.051115
Interval7	2.001250000000	0.012166		**Interval7**	1.776250000000	0.026476
Upper limit	2.101375000000			**Interval8**	1.931500000000	0.013094
				Interval9	2.086750000000	0.006275
				Upper limit	*2.164375000000*	