1-2016

Representation of Confidence in Assurance Cases Using the Beta Distribution

Lian Duan

Sanjai Rayadurgam

Mats Heimdahl

Oleg Sokolsky

University of Pennsylvania, sokolsky@cis.upenn.edu

Insup Lee

University of Pennsylvania, lee@cis.upenn.edu

Follow this and additional works at: https://repository.upenn.edu/cis_papers

Part of the [Computer Engineering Commons](#), and the [Computer Sciences Commons](#)

Recommended Citation

Lian Duan, Sanjai Rayadurgam, Mats Heimdahl, Oleg Sokolsky, and Insup Lee, "Representation of Confidence in Assurance Cases Using the Beta Distribution", *2016 IEEE 17th International Symposium on High Assurance Systems Engineering (HASE)*, 86-93. January 2016. <http://dx.doi.org/10.1109/HASE.2016.52>

IEEE High Assurance Systems Engineering Symposium ([HASE 2016](#)), Orlando, Florida, USA, January 7 - 9, 2016.

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/cis_papers/812
For more information, please contact repository@pobox.upenn.edu.

Representation of Confidence in Assurance Cases Using the Beta Distribution

Abstract

Assurance cases are used to document an argument that a system—such as a critical software system—satisfies some desirable property (e.g., safety, security, or reliability). Demonstrating high confidence that the claims made based on an assurance case can be trusted is crucial to the success of the case. Researchers have proposed quantification of confidence as a Baconian probability ratio of eliminated concerns about the assurance case to the total number of identified concerns. In this paper, we extend their work by mapping this discrete ratio to a continuous probability distribution—a beta distribution—enabling different visualizations of the confidence in a claim. Further, the beta distribution allows us to quantify and visualize the uncertainty associated with the expressed confidence. Additionally, by transforming the assurance case into a reasoning structure, we show how confidence calculations can be performed using beta distributions.

Disciplines

Computer Engineering | Computer Sciences

Comments

IEEE High Assurance Systems Engineering Symposium ([HASE 2016](#)), Orlando, Florida, USA, January 7 - 9, 2016.

Representation of Confidence in Assurance Cases using the Beta Distribution

Lian Duan*, Sanjai Rayadurgam*, Mats Heimdahl*, Oleg Sokolsky[†], Insup Lee[†]

*Department of Computer Science and Engineering

University of Minnesota

{lduan, rsanjai, heimdahl}@cs.umn.edu

[†]Department of Computer and Information Science

University of Pennsylvania

{sokolsky, lee}@cis.upenn.edu

Abstract—

Assurance cases are used to document an argument that a system—such as a critical software system—satisfies some desirable property (e.g., safety, security, or reliability). Demonstrating high confidence that the claims made based on an assurance case can be trusted is crucial to the success of the case. Researchers have proposed quantification of confidence as a Baconian probability ratio of eliminated concerns about the assurance case to the total number of identified concerns. In this paper, we extend their work by mapping this discrete ratio to a continuous probability distribution—a beta distribution—enabling different visualizations of the confidence in a claim. Further, the beta distribution allows us to quantify and visualize the uncertainty associated with the expressed confidence. Additionally, by transforming the assurance case into a reasoning structure, we show how confidence calculations can be performed using beta distributions.

I. INTRODUCTION

Assurance cases are structured logical arguments that are used to document how a claim is supported by evidence. Assurance cases are becoming increasingly popular, especially in safety-critical areas such as medical devices and civil aviation. A well-argued assurance case can be used to show that a system satisfies desirable properties such as safety (thus becoming a safety case), reliability, or security. The UK Ministry of Defence describes a safety case as:

“A structured argument, supported by a body of evidence that provides a compelling, comprehensible and valid case that a system is safe for a given application in a given operating environment.” [1].

In the United States, manufacturers of certain medical devices that seek FDA approval must show that they are safe through, among other things, the use of an assurance case. The preparation of the assurance case demonstrates that the manufacturer has considered how various pieces of evidence such as testing results, pre-clinical trials, user studies, and documentation support a claim of safety. The reviewer must now decide, given all supporting information, if he or she agrees with the manufacturer. An important factor in this

process is the amount of confidence the reviewer has in the evidence given and on how the assurance case has been prepared and structured. Being able to evaluate this confidence in some systematic way is crucial to the evaluation of assurance cases. While the process has a certain subjectiveness due to differences in reviewers, it is our hope that such a systematic approach will bring a consistency to the results.

Hawkins et al. [9] introduced a qualitative approach to establishing confidence in an assurance case. They proposed the concept of an assured safety argument—an assurance case and its corresponding confidence case. The confidence case is based on arguing the sufficiency of the implications in an assurance case written in a graphical notation like GSN [12]. The crux of their approach is to find assurance deficits—anything that could reduce one’s confidence in the assurance case—and argue why these are acceptable.

Goodenough et al. [5] extend Hawkins et al.’s work by quantifying confidence as a Baconian probability. This ratio is derived from inductive reasoning—increasing confidence through increased knowledge. This Baconian probability ratio is a pair of integers frequently written in a fractional notation for convenience. The number of assurance deficits (“doubts” or “defeaters”) that have been eliminated or mitigated is the numerator-like value while the total number of doubts identified is the denominator-like value. The Baconian probabilities are then summed up the assurance case to result in a confidence probability value for the entire assurance case.

We extend the work of Hawkins et al. and Goodenough et al. by quantifying and visualizing the Baconian probability with the beta distribution. The beta distribution can take on a variety of shapes and usually has the range $0 - 1$, making it ideal for probabilities. It has two parameters, α and β , that affect the shape and scale of the distribution. As α and β increase, the variance of the graph decreases. Increasing α shifts the mode of the distribution to the right (towards 1) while increasing β shifts the mode of the distribution to the left (towards 0). When α and β are equal, the distribution is symmetric. The number of doubts eliminated can be mapped directly onto the α parameter and the number of doubts remaining can be mapped directly onto the β parameter. Intu-

This work has been partially supported by NSF grants CNS-0931931 and CNS-1035715.

itively, this makes sense—as the number of doubts eliminated grows, so does our confidence. We then can use properties of the beta distribution to calculate the uncertainty. Instead of simply summing up the Baconian probability values, however, we propose a weighting scheme with the beta distribution parameters and a logical restructuring of the assurance case to make a more intuitive argument. The use of the beta distribution also allows us to use Jøsang’s opinion triangle [10] as an additional visualization tool. Jøsang introduces a direct mapping from the beta distribution to the opinion triangle and back via subjective logic. The opinion triangle visualizes an opinion on an intuitive three-dimensional scale of belief, disbelief, and uncertainty. According to Habib et al. [8], the opinion triangle is “very well suited for analysis done by experts,” which fits what we are proposing. Although they ultimately conclude that a different trust analysis system is more intuitive for the lay person, the opinion triangle will be apt in our case.

II. BACKGROUND AND RELATED WORK

We will go into a bit of background related to assurance cases, relevant approaches to evaluating confidence in assurance cases, and the beta distribution. Additionally, we will highlight relevant related work.

Figure 1 is a sample assurance case which seeks to show that radiation over the standard amount (henceforth, overradiation) will not occur with an x-ray backscattering machine, similar to what one might see at an airport. For illustrative purposes, this assurance case is very small. Our sample assurance case follows the GSN notation [12], a popular approach to writing assurance cases that has been adopted by a variety of companies in Europe [6].

In our sample assurance case, we have a top-level claim (G1) that “All causes of overradiation have been eliminated” in the system. This claim is supported by the argument strategy, S1, “Argument over all identified causes.” Our assumption, that all causes of overradiation have been identified, serves as the context (C1) for our argument. Our argument is supported by two sub-claims, G2, “Software operating as intended,” and G3, “Timer interlock operating correctly.” From this figure, we can see that G2 is supported by two pieces of evidence (solutions), Sn1 (Formal verification results) and Sn2 (Testing results), and G3 is supported by two pieces of evidence, Sn3 (Fault Tree Analysis) and Sn4 (Testing results). The question that we must then answer is, given the four pieces of evidence, and this argument structure claiming that the evidence supports the top-level claim, how much confidence can we have in this assurance case?

A reviewer looking at this assurance case would need to establish a confidence level for the case, whether implicitly or explicitly. Duan et al. provide a general survey that summarizes various approaches to evaluating confidence in assurance cases [3].

Fig. 1: An example of an assurance case

A. Assured Safety Argument

Hawkins et al. [9] introduces the idea of an assured safety argument, an assurance case accompanied by a corresponding confidence case. The confidence case seeks to address all uncertainties that may exist in the assurance case. This purely qualitative approach claims that one can arrive at consistent and reasonable conclusions given a systematic approach.

The assurance case and the confidence case are connected through three types of *assurance claim points*, or ACP. An ACP exists between a claim or argument and a context, a claim and an argument, and a claim and a solution (evidence). Each ACP has a corresponding confidence case which is meant to strengthen and argue why the connection is valid and necessary.

The general approach to creating confidence cases is to identify assurance deficits in the argument—anything that can reduce our confidence in the assurance case. Then, we try to eliminate those. Any that have not been eliminated readily become a “residual assurance deficit.” These must be argued further, by showing that they are acceptable either because they will not severely impact that top-level claim or that significant counter-evidence against them does not exist.

B. Baconian Probabilities

Goodenough et al. [5] proposed the idea of using inductive reasoning, as introduced by Francis Bacon, to quantify the confidence one has in an assurance case. The general premise is that when evaluating an assurance case, one comes up with a list of “defeaters” or “doubts”—anything that might cause one to decrease confidence in the assurance case. This is similar to the assurance deficit idea that Hawkins et al. used. Example sources of doubt could be credibility of the testing, validity of the results, trustworthiness of the testers, or quality of documentation. For example, suppose we have identified all assurance deficits for evidence node Sn1—there are 10 of them—and thus have 10 sources of doubt. Then, we go

through each doubt and eliminate it by resolving or mitigating it—for example, 8 doubts have been eliminated—or decide it cannot be eliminated—2 doubts remain unresolved. These unresolved doubts are similar to the residual assurance deficits mentioned by Hawkins et al. The confidence one has in an assurance case is then represented as a Baconian probability—a ratio of the number of doubts eliminated to the total number of doubts. This ratio is irreducible, as $\frac{4}{5}$ would represent an entirely different confidence value than $\frac{8}{10}$ —the latter shows a higher confidence value and indicates that more doubts had been found and ultimately mitigated or eliminated.

Figure 2 shows our sample assurance case with example Baconian probability values and how they would be combined, by summing up the assurance case, to arrive at a final confidence value for the entire case.

Fig. 2: Assurance case with sample Baconian probabilities

C. Dempster-Shafer Theory

Dempster-Shafer theory seeks to reason about uncertainty by quantifying confidence (or belief) as a mass value. This mass value, which is between 0 and 1, grows as our confidence grows. The confidence value can exist on a range bounded by belief on the lower end and plausibility on the upper end. Existing information is used to find the boundaries of this possible range.

What makes Dempster-Shafer theory unique is that it separates out uncertainty from belief and disbelief. Uncertainty is treated explicitly as the quantity left over after belief and disbelief have been accounted for. Jøsang visualizes this concept as an opinion triangle, as seen in Figure 3 (a). An opinion, ω , is represented by three parameters—belief (b), disbelief (d), and uncertainty (u), on this bounded, triangular plane.

Cyra and Gorski [2] use Dempster-Shafer theory, the opinion triangle, and Toulmin’s argumentation theory [13] in their calculations and visualizations of confidence in assurance cases. Figure 4 shows our sample assurance case as visualized

Fig. 3: Jøsang’s Opinion Triangle

by the tool created by Cyra and Gorski. The left side shows the text-style assurance case as espoused by Cyra and Gorski. Similar to GSN, it has a top level claim, a context for that claim, an argument for the entire assurance case, subclaims, and pieces of evidence (solutions). The biggest change from GSN is the requirement of a justification for the argument. On the right, we see how users can choose to set confidence and decision ratings for different components of the assurance case.

D. Beta Distribution

The beta distribution is part of “a most flexible family of distributions” [7], a continuous version of the binomial distribution on the range from 0 to 1, making it very appropriate for modeling probabilities. It is popular for use in Bayesian analysis [7]. The beta function is defined by:

$$B(\alpha, \beta) = \int_0^1 z^{\alpha-1} (1-z)^{\beta-1} dz \quad (1)$$

Its probability density function (pdf) is described by two parameters, α and β . The pdf is described by:

$$f(x; \alpha, \beta) = \frac{1}{B(\alpha, \beta)} x^{\alpha-1} (1-x)^{\beta-1}, 0 < x < 1 \quad (2)$$

where $\alpha > 0$, $\beta > 0$, and $B(\alpha, \beta)$ is the beta function.

When α and β are both 1, $beta(1, 1)$ is the uniform distribution. When α and β are equal, the beta distribution models an approximate Gaussian shape. When α and β are both less than 1, the beta distribution takes a “U” shape, with the special case of $beta(0.5, 0.5)$ being the arc-sine distribution [7]. When β is 1, the distribution takes on a power shape. When α and β are both greater than 1 and not equal, the beta distribution takes on a skewed Gaussian shape, where the mean and mode of the curve are not equal. Figure 5 shows some of these shapes.

As α increases while β stays the same, the peak of the curve shifts to the right. This increase can be viewed as more positive information incoming, such as with reviews. Each review increases our α value, and increases our trust while also reducing the uncertainty due to the fact that we have more information. As β increases while α stays the same, the peak of the curve shifts to the left. This movement can be viewed as more negative information incoming, such as negative reviews that reduce our trust, but also reduce uncertainty. As α or

Fig. 4: Assurance case and assessment example using Argevide Nor-Sta

Fig. 5: Sample beta distribution shapes

β increases, the maximum value of the peak increases while the uncertainty, or variance of the curve, decreases. This can be viewed as knowledge being increased, thereby reducing uncertainty or ignorance.

Josang [11] provides a mapping between his opinion triangle and the beta distribution, as seen in Figure 3. This mapping depends on the values of belief, disbelief, uncertainty, and the prior belief value, a known quantity with no uncertainty. In Figure 3, we have an opinion $\{\text{belief}, \text{disbelief}, \text{uncertainty}\}$ of $\{0.7, 0.1, 0.2\}$, which translates to a beta distribution with parameters $\text{beta}(8, 2)$. We can see in the opinion triangle that this opinion has fairly high belief and low disbelief, but also has a fair amount of uncertainty. These facts are reflected in the beta distribution with a peak around 0.8, but quite a high variance.

Duan et al. [4] first proposed the use of the beta distribution to represent confidence in assurance cases. This paper extends that work with an application of the beta distribution with Baconian probabilities, the use of a weighting scale, and a more formal approach to restructuring the assurance case into

a logical argument.

III. TECHNICAL APPROACH

A. Visualization of Confidence

The beta distribution can be used to visualize the Baconian probability ($\frac{n}{d}$) representing confidence. One can view the α parameter of the beta distribution as representing the number of doubts that have been eliminated—increasing our confidence and decreasing our uncertainty. The β parameter would represent the number of doubts that still remain (the numerator subtracted from the denominator)—items that would reduce our confidence while also decreasing our uncertainty.

When all the doubts have been eliminated in Baconian probability, the numerator and denominator components are equal and we have full confidence. As visualized by a beta distribution, the β parameter is 0, and we would have a discontinuity (spike) at 1. For a Baconian probability value of $\frac{16}{30}$, we would get the corresponding α and β values of $\text{beta}(16, 14)$.

Assuming n is the Baconian numerator and d is the Baconian denominator, the mapping is thus:

$$\alpha = n \quad (3)$$

$$\beta = d - n \quad (4)$$

Figure 6 shows the beta distributions as mapped from Baconian probability values for each of the evidence nodes, Sn1 ($\text{beta}(5, 1)$), Sn2 ($\text{beta}(6, 1)$), Sn3 ($\text{beta}(2, 6)$), and Sn4 ($\text{beta}(3, 6)$). The distributions for each node are given in different colors and line styles. We see that for Sn1 and Sn2, which only have one defeater each, the graph is a low-sloped J-shape. Sn2 has a higher slope than Sn1, due to the extra information that exists for it. Sn3 and Sn4 have their modes more in the left side of the graph, reflecting our low confidence (or higher disbelief) in them. Figure 7 shows the distributions for the evidence nodes Sn1 ($\text{beta}(5, 1)$) and Sn2 ($\text{beta}(6, 1)$), and the combination of the two, the confidence for the software claim node G2 ($\text{beta}(11, 2)$). We see that the additional defeater minimizes the ability for our confidence

to be at 1, but the disbelief is also reduced slightly. Figure 8 shows the distributions for the evidence nodes Sn3 ($\text{beta}(2, 6)$) and Sn4 ($\text{beta}(3, 6)$), and the combination of the two, the confidence for the hardware claim node G3 ($\text{beta}(5, 12)$). The combination graph has a higher peak and lower uncertainty than the two component graphs, and its mode is situated between the two, as one would expect. Lastly, Figure 9 shows the distributions for the software node ($\text{beta}(11, 2)$) and the hardware node ($\text{beta}(5, 12)$), and the combination of the two, the confidence for the overradiation node, G1, and the entire assurance case ($\text{beta}(16, 14)$). We see that when combining two very disparate graphs, we arrive at a distribution almost evenly between the two. These examples show the added richness of understanding that is provided when visualizing the Baconian probability ratio as a beta distribution.

Fig. 6: Beta distributions for confidence in the evidence nodes Sn1 - Sn4

Fig. 7: Beta distributions for confidence in the software node G2

By using the beta distribution, we explicitly add a third dimension of information that did not previously exist in the Baconian approach—the separation out of uncertainty. Uncertainty in the beta distribution can be calculated as the difference between the two inflection points of the curve. This

Fig. 8: Beta distributions for confidence in the hardware node G3

Fig. 9: Beta distributions for confidence in the overradiation node G1

gives a specific, quantifiable value to uncertainty that can aid in assessment. The calculation for the inflection points of the beta distribution is:

$$\frac{\alpha - 1}{\alpha + \beta - 2} \pm \frac{1}{\alpha + \beta - 2} \sqrt{\frac{(\alpha - 1)(\beta - 1)}{\alpha + \beta - 3}} \quad (5)$$

Note that when α or β is equal to one, the inflection points are both either 0 or 1, respectively, so Eq. 5 cannot be used to calculate an uncertainty value. In such a situation, we propose an alternative method for evaluating uncertainty. Two times the distance from 0 (for α) or 1 (for β) of the expected value of the beta distribution,

$$\frac{\alpha}{\alpha + \beta} \quad (6)$$

should be used. The inflection points of the beta distribution are approximately located on either side of expected value. But

with a J-shaped graph like we see when α or β is 1, we don't have inflection points—the slope is constantly increasing as it gets closer and closer to 0 or 1, respectively. The expected value, which can be viewed as the center of the distribution, would then give us a measure of how far from either extreme most of the data is centered. We then multiply this value by 2 to account for the two inflection points that exist otherwise. Table I shows the uncertainty values for the evidence nodes.

An additional visualization tool is available with the use of the beta distribution. Introduced by Josang [10], it is called an opinion triangle, a triangle whose vertices represent belief (right bottom corner), disbelief (left bottom corner), and uncertainty (top corner). Josang developed a mapping between the beta distribution and the opinion triangle, via the use of subjective logic [11]. Figure 10 shows our Baconian probability value of $\frac{16}{30}$ (represented by a beta distribution of $\text{beta}(16, 14)$) as plotted on Josang's opinion triangle. We can see that we have fairly low uncertainty, due to the quantity of information, but our information is a bit conflicting, so we are almost at the middle between disbelief and belief. A reviewer examining this assurance case with a confidence of a Baconian probability ratio of $\frac{16}{30}$ will then have two options for visualizing this confidence value beyond the numerical values given in the table—either as the central beta distribution from Figure 9 or as the dot on the opinion triangle in Figure 10. Even though the distribution seemed to have a high uncertainty, from the opinion triangle we can see that it is actually not so high as to cause alarm. We can also clearly see, from the triangle, how evenly divided we are in the opinion.

Fig. 10: Baconian probability example on an opinion triangle

B. Weighting Scheme

In Goodenough et. al.'s approach, any uncertainty that exists is accounted for in the doubts that have accumulated. These doubts only exist at the leaf level of the assurance case—the evidence nodes—and to find a confidence value for the entire case, one simply sums up the tree structure (as seen in Figure 2).

We, however, believe that simple addition will not be enough for calculating a final confidence value. This is because there are factors that have not been accounted for—for example, the assurance case structure might not be correct, or one evidence should have more importance (weight) than another one. As an example for the latter, suppose two of the sources of doubts we have for Sn2: Testing Results are (1) whether or not proper documentation was employed, and (2) all the tests have passed to a statement coverage criteria, instead of a higher standard such as MC/DC. It is unrealistic to put equal weights to these two sources of doubts, as (2) is more important than (1).

We propose a *weighting scheme* for such a situation, to be applied to the parameters of the beta distribution for each of the evidence nodes. The weighting scheme will have five options: very low impact (scale by 0.5), low impact (scale by 0.75), medium (default) impact (scale by 1.0), high impact (scale by 1.5), and very high impact (scale by 2.0). Each evidence node should have a corresponding weight, or importance, which will figure into ultimate calculations. Each sub-claim node should also have a corresponding weight. Figure 11 shows sample weights for our Baconian probability assurance case. We have deemed the software branch of the assurance case to have more weight than the hardware branch, and have weighted the various nodes according. As such, even though we have low confidence in the hardware node, the fact that we have high confidence in the software node should be more important. Figure 12 shows the beta distribution for the top-level claim for both the original Baconian probability values (red, dotted line) and the weighted Baconian probability values (blue, solid line). We can see that, as expected, the weighted distribution has a peak at a higher confidence value, reflecting the higher confidence we have in the assurance case due to the weighting system. As a result of our weighting system, we no longer have integers for the confidence values. This is a fairly trivial concern—the beta distribution has no such limit on whether its parameters are integers or reals, and our use of the Baconian probability only extends to the evidence nodes.

C. Combination of Evidence and Uncertainty

When analyzing the way confidence is combined in an assurance case, we can switch to a fault tree analysis-like structure to aid us. We are thus no longer limited to the straightforward addition approach as used by Goodenough et al. We call this a *reasoning structure*. Depending on the assurance case, evidence can be combined with an AND or OR operator (and in the future, others as needed). An OR operator would not be accurately represented by simple addition, which would disproportionately favor the weaker argument.

For an assurance case, we want to make an argument over all causes, so each branch needs to be evaluated carefully, because we seek to make the strongest argument possible. When evaluating the confidence, however, the argument structure could be changed, because, depending on the reasoning structure, confidence in only one branch would be enough.

TABLE I: Uncertainty values calculated from the beta distribution

Node Name	α	β	Expected Value	Inflection Point 1	Inflection Point 2	Uncertainty
Sn1 Formal verification results	5	1	0.833	1	1	0.333
Sn2 Testing results	6	1	0.857	1	1	0.286
Sn3 Fault Tree Analysis	2	6	0.25	0.333	0	0.333
Sn4 Testing results	3	6	0.333	0.470	0.101	0.369
G2 Software node	11	2	0.846	1	0.818	0.182
G3 Hardware node	5	12	0.294	0.385	0.148	0.236
G1 Overradiation node	16	14	0.533	0.632	0.440	0.192

Fig. 11: Assurance case with Baconian probabilities and weights

Fig. 12: Beta distributions for weighted evidence nodes

As an example, consider our assurance case with Baconian probability values assigned to the nodes as seen in Figure 2. The timer interlock (hardware) branch has a lot of doubts that have not been eliminated, so it has a very low confidence value of $\frac{5}{17}$. If we continued with the Baconian approach, we would have a moderate confidence value of $\frac{16}{30}$ for the entire assurance case. Such a (comparedly) low confidence value would not give one much confidence in the system being assured. Figure 9 shows the Baconian probability values for the software, hardware, and overradiation nodes as converted to beta distribution parameters and then plotted. We see that one branch has high confidence (mode of graph is shifted to the right) while the other branch has low confidence (mode of graph is shifted to the left), resulting in a medium confidence for the assurance case.

We can, however, change the assurance case into our reasoning structure (Figure 13). The conversion requires one to analyze the specific situation and see what kind of reasoning structure should best be used. In our case, for overradiation to occur, we would need a failure in *both* the software and the hardware components. However, since our claim is that “No overradiation will occur,” and we know that there is an existing redundancy system in place where both components have to fail for overradiation to occur, we just need to have enough confidence in one branch not failing to have enough confidence in the whole assurance case—thus, we use an OR argument. No overradiation will occur if the software does not fail or the hardware does not fail (or neither fails). The new beta distributions for the logical argumentation structure is shown in Figure 14. We can see that now we have fairly high confidence in the assurance case, reflecting our confidence in the redundancy system that exists in the x-ray machine. In normal probability, a logical-OR construct has a higher probability of occurrence than either of its components. By choosing the branch in which we have the highest confidence, we are actually taking a conservative approach to ensure that we have at least accounted for a minimum possible confidence value.

IV. CONCLUSION

Assurance cases have become increasingly popular in areas of safety-critical systems. Being able to capture the confidence and uncertainty one has in an assurance case in a quantifiable way is extremely helpful. We have introduced a way of visualizing the Baconian probability method of representing confidence in an assurance case with the use of the beta distribution and the opinion triangle and provided an example

Fig. 13: Logical Argument for Example Assurance Case

Fig. 14: Beta distribution representation of Baconian probabilities in the logical argument structure

of its use. Additionally, we have established a method to calculate the uncertainty associated with an evidence node of an assurance case based on the inflection points and expected values of the beta distribution. These approaches will give a sound, mathematical basis to calculating uncertainty. We also introduced a weighting scheme to make calculations of confidences more realistic. Lastly, we introduced a reasoning structure to better combine confidences for different nodes of an assurance case. It is our hope that these approaches will provide a clearer, more consistent method of quantifying confidence and uncertainty for an assurance case. Our future work will seek to further clarify and codify the set of “rules” for our approach, as well as provide a more objective evaluation of the claims.

REFERENCES

- [1] Safety management requirements for defence systems. Defence Standard 00-56 4, Ministry of Defense, June 2007.
- [2] L. Cyra and J. Górski. Supporting expert assessment of argument structures in trust cases. In *9th International Probability Safety Assessment and Management Conference PSAM*, 2008.
- [3] L. Duan, S. Rayadurgam, M. Heimdahl, A. Ayoub, O. Sokolsky, and I. Lee. Reasoning about confidence and uncertainty in assurance cases: A survey. In *Software Engineering in Health Care*, 2014.
- [4] L. Duan, S. Rayadurgam, M. Heimdahl, O. Sokolsky, and I. Lee. Representation of confidence in assurance case evidence. In *ASSURE 2015*, 2015.
- [5] J. B. Goodenough, C. B. Weinstock, and A. Z. Klein. Toward a theory of assurance case confidence. Technical report, Carnegie Mellon, 2012.
- [6] GSN. <http://www.goalstructuringnotation.info/archives/29>.
- [7] A. K. Gupta and S. Nadarajah, editors. *Handbook of Beta Distribution and Its Applications*. Marcel Dekker, Inc., 2004.
- [8] S. M. Habib, F. Volk, S. Hauke, and M. Mhlhuser. Chapter 21 - computational trust methods for security quantification in the cloud ecosystem. In R. K.-K. R. Choo, editor, *The Cloud Security Ecosystem*, pages 463 – 493. Syngress, Boston, 2015.
- [9] R. Hawkins, T. Kelly, J. Knight, and P. Graydon. A new approach to creating clear safety arguments. In *Advances in Systems Safety*, 2011.
- [10] A. Jøsang. Artificial reasoning with subjective logic. In *Proceedings of the Second Australian Workshop on Commonsense Reasoning*, 1997.
- [11] A. Jøsang. A logic for uncertain probabilities. 1999.
- [12] T. Kelly and R. Weaver. The goal structuring notation - a safety argument notation. In *Dependable Systems and Netowkrs 2004 Workshop on Assurance Cases*, 2004.
- [13] S. Toulmin. *The Uses of Argument*. Cambridge University Press, 1958.