Understanding Assurance Cases:
An Educational Presentation in Five Parts

Module 5: Speculation

C. Michael Holloway
c.michael.holloway@nasa.gov

Senior Research Computer Engineer
Safety-Critical Avionics Systems Branch
NASA Langley Research Center, Hampton, Virginia, U.S.A.

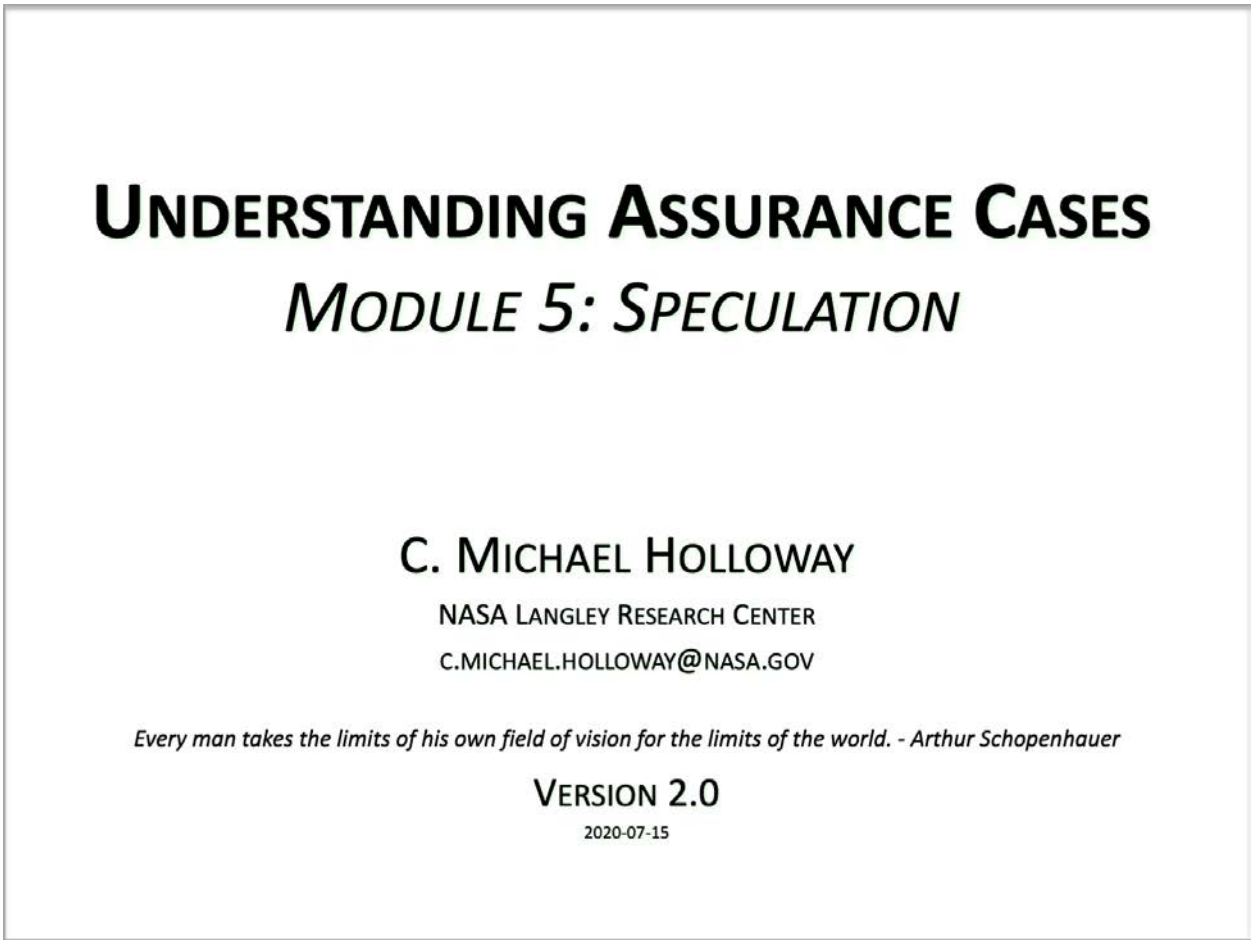

This material was created in 2015-16, as part of the Explicate '78 project. The project was supported in substantial part by the Assurance Case Applicability to Digital Systems task under the reimbursable interagency agreement with the Federal Aviation Administration for Design, Verification, and Validation of Advanced Digital Airborne Systems Technology (IAI-1073 Annex 2 for NASA; DTFACT-10-X0008, Modification 0004 for the FAA). The original presentations were delivered to a selected group of FAA civil servants and NASA Langley personnel. The audio was recorded and partial transcripts (containing only the words spoken by the presenter, Mr. Holloway) produced. The intent from the beginning was to collect the material into a form that could be made available publicly. The text adheres closely to the original transcript, with the exception of an occasional insertion of new information that arose since the original presentation dates. The full collection consists of six documents (including this one), which are available electronically through https://shemesh.larc.nasa.gov/arg/uac.html.

Hello everybody.

We've now come to the fifth, and final module in our educational series about Understanding Assurance Cases, which I've titled ***Speculation***. The topics we will discuss are a bit less concrete and a lot more speculative than the topics of the previous 4 modules.

Perhaps in talking about these topics we can provide a counter-example to Arthur Schopenhauer's assertion that "Every man takes the limits of his ***own*** field of vision for the limits of the world." [Schopenhauer, Arthur. 1951. *Studies in Pessimism: Essays from the Parerga and Paralipomena*. Translated by T. Bailey Saunders. London: Allen and Unwin.]

As always interrupt me at ***any*** point if you have a burning question. I'll either answer it or defer an answer to a more appropriate time. Also, as with the other modules, there will be a few times when I'll ask you questions, too.

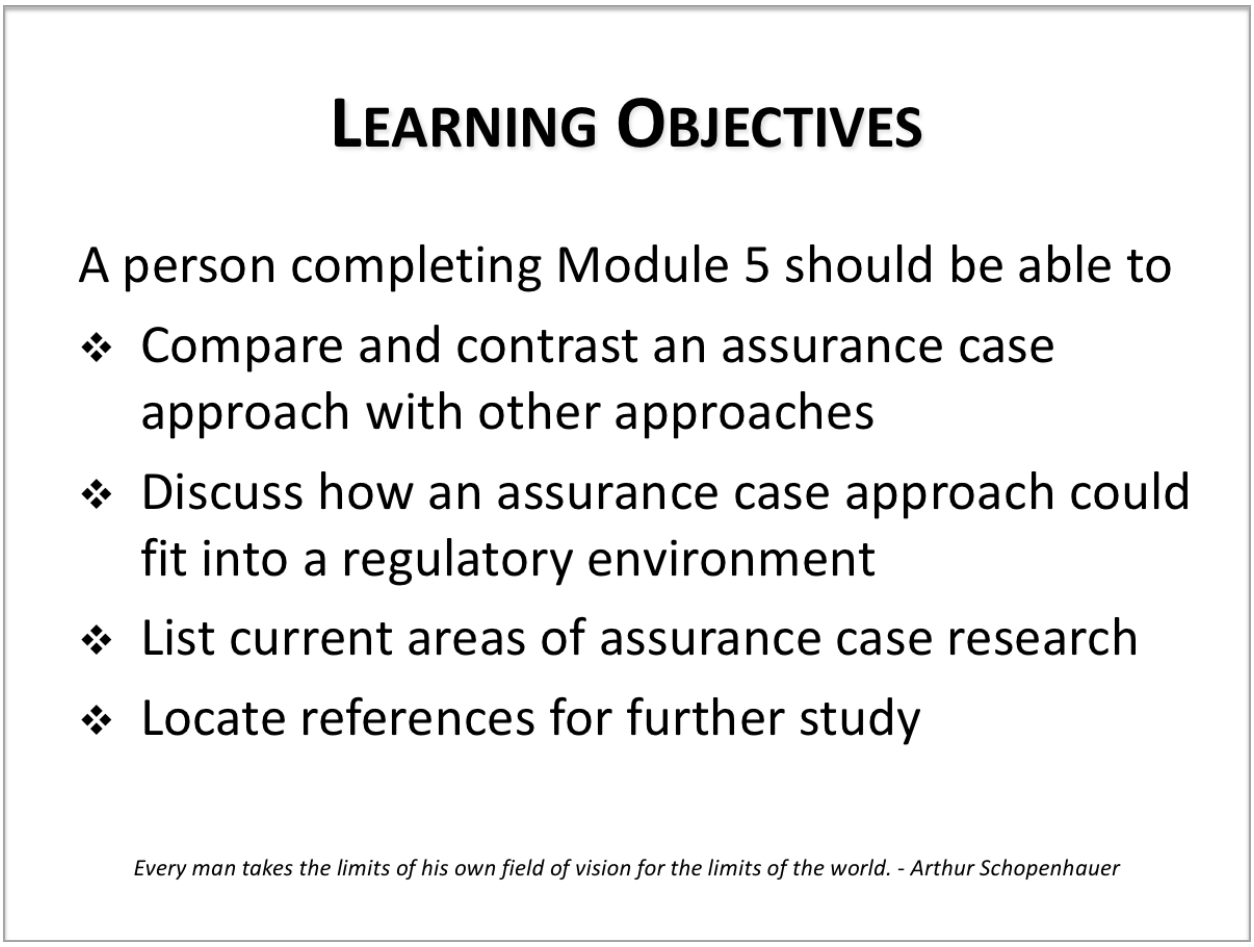

Here are our learning objectives for this module:

- Compare and contrast an assurance case approach with other approaches
- Discuss how an assurance case approach could fit into a regulatory environment
- List current areas of assurance case research
- Locate references for further study

In Module 4, we left Jon and Mike with Jon thanking his Dad for agreeing to show him how to create an assurance case concerning Tim driving Jon to the game.

We still don't know exactly who is playing in the game, or even what sport it is, which I find a bit disconcerting. But today we learn that Jon's case convinced his dad to let him go to the game.

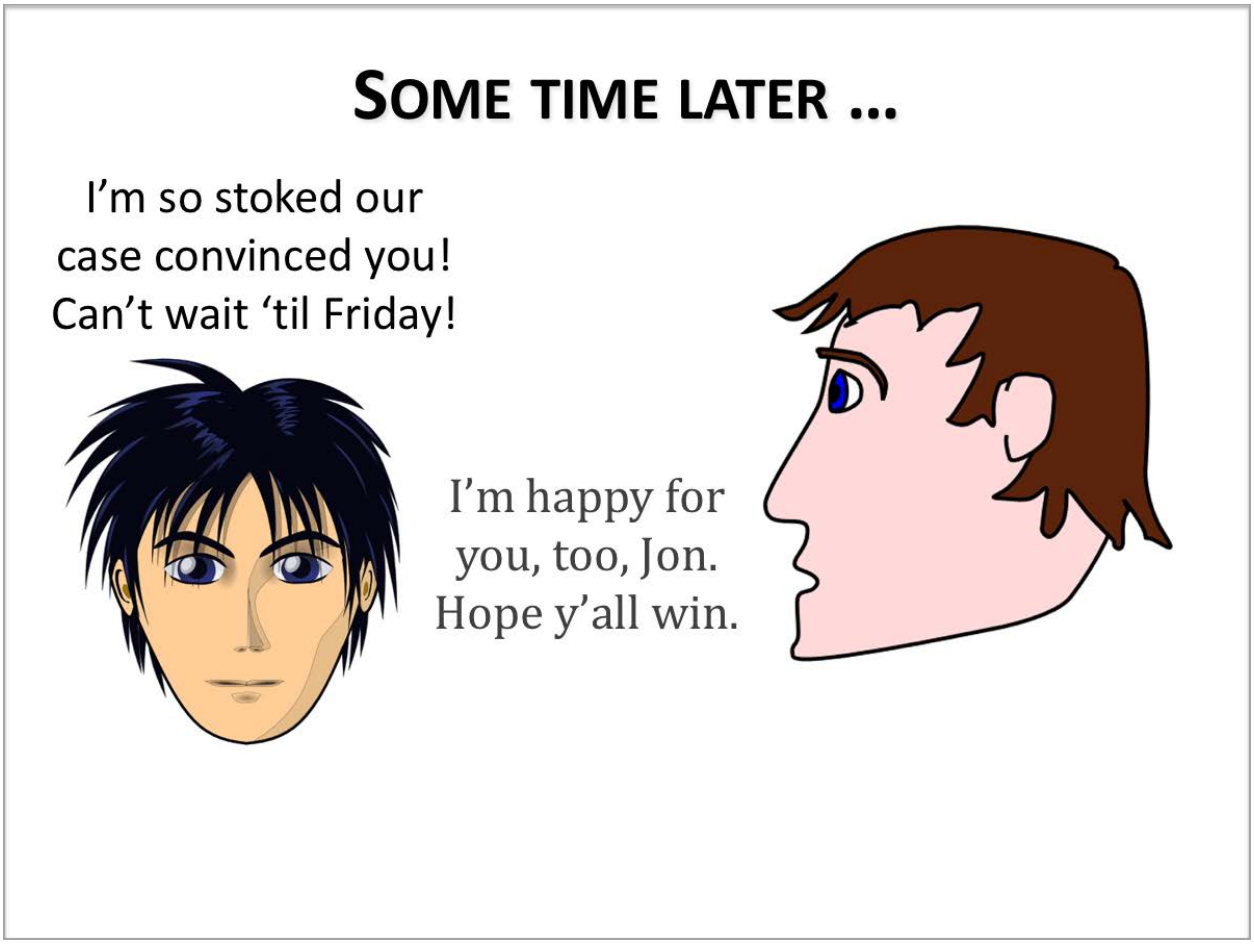

"I'm so stoked our case convinced you! Can't wait 'til Friday!"

(So we now know when the game is taking place.)

"I'm happy for you, too, Jon. Hope y'all win."

Mike's response suggests that perhaps Jon's school is one of the teams in the game.

After a brief pause, Jon says, "One more question … then I've gotta do homework."

"OK, Let's hear it," replies Mike.

"This case stuff really made me think. Tim said so, too. It seems like such a great idea. Why don't more people use it?" Jon asks.

His dad replies, "Wow. … that's a hard question with lots of different parts to the answer."

“Like what, for instance? ” inquires Jon.

“Lack of understanding … thinking it is harder than it is … or thinking it adds nothing to what’s usually done … or even just believing rants from some experts.”

“Sounds confusing. When’s it gonna end?”

“Don’t know if it ever will,” says Mike frowning, “unless someone figures out a good way to compare how well different methods work.”

Jon says nothing for a second or two, then exclaims, “Better go do my homework now.”

“Watcha gotta do,” asks Mike.

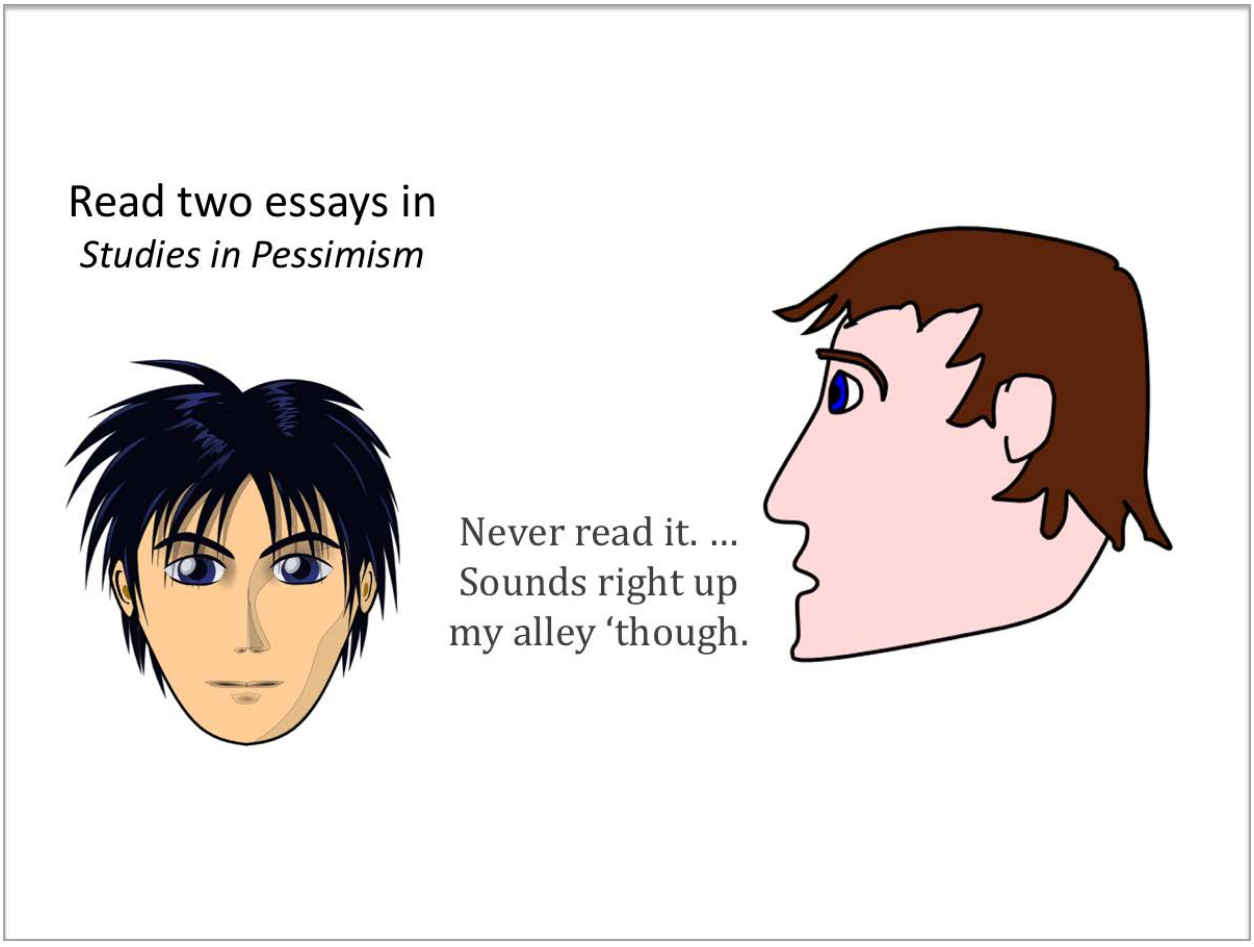

“Read two essays in *Studies in Pessimism.”*

“Never read it,” says Mike, “Sounds right up my alley ‘though.”

Studies in Pessimism is by Schopenhauer by the way, and is the source of the quote for Module 5[1].

[1] Well, to be more precise, *Studies in Pessimism* is the English title of a translation of a compilation of some of Schopenhauer’s writings

To continue along the lines Mike mentioned, let's now spend a little bit of time comparing and contrasting an assurance case approach to other approaches.

We'll begin with similarities.

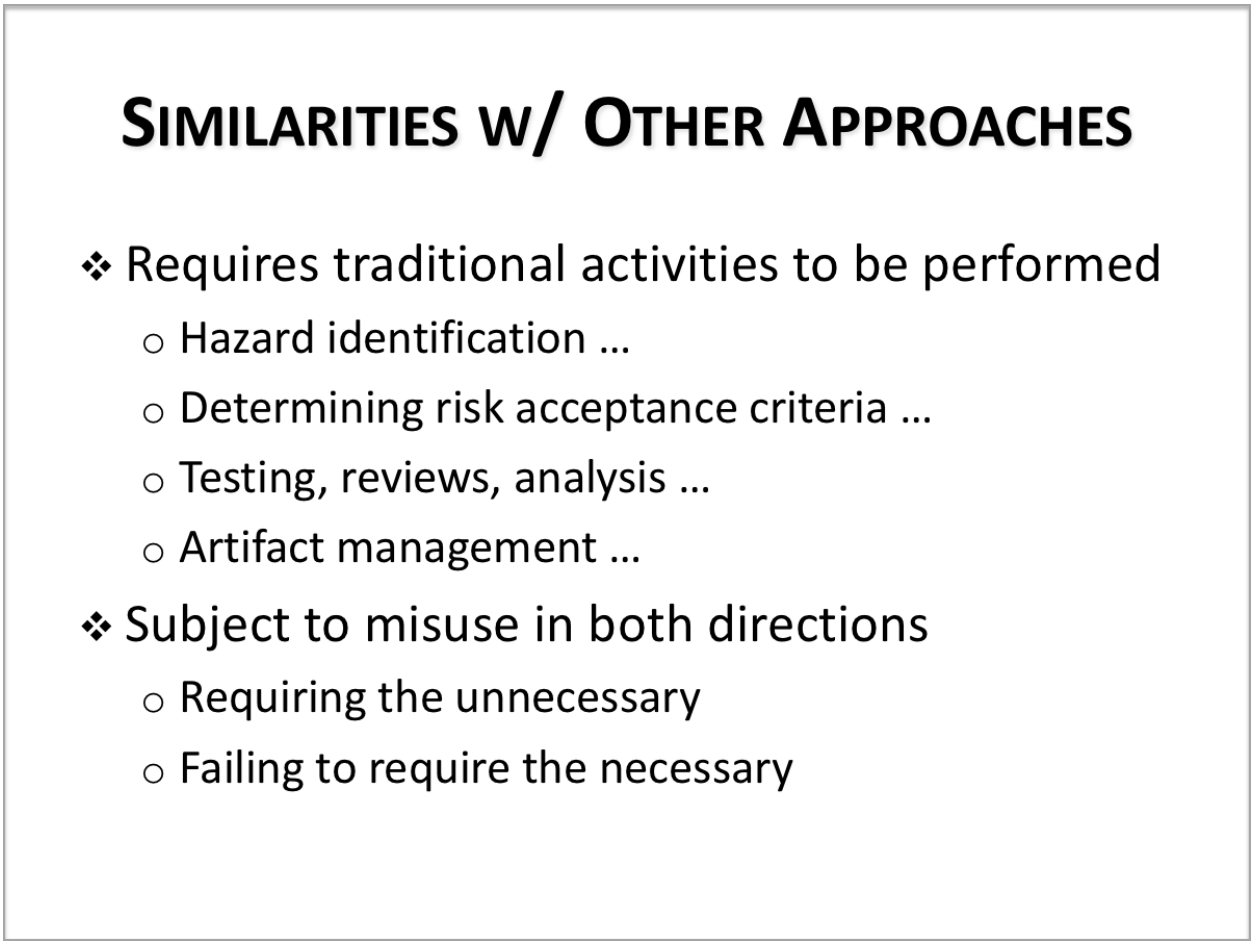

One of the most important similarities between an assurance case approach and traditional approaches to safety assurance is that writing an assurance case still requires traditional activities to be performed. We've talked about this several times in previous modules, but I want to emphasis it again here. Assurance cases are not a way to get out of doing necessary technical work. They may provide a way to get out of doing unnecessary "administrative"-type work, but *they're not a shortcut to safety*.

In particular, writing an assurance case doesn't absolve you of the need to identify hazards or to determine risk acceptance criteria or to do testing, reviews, and analysis or to manage all your artifacts well. All of those things still need to be done, perhaps in slightly different ways, using different notations or techniques, but they still have to be done.

Another similarity between assurance case techniques and other techniques is that they do not provide a guarantee against misuse. Someone may adopt an insurance case approach and still require things that are unnecessary. Or fail to require things that are necessary. *Assurances cases are not a silver bullet*.

To emphasize again, they're neither a shortcut nor a silver bullet. Unsafe, incorrect, bad systems can be mistakenly assured as safe, correct, good under an assurance case regime just as they can be under some other regime. Perhaps you recall from Module 2 the Nimod accident, before which a bad system was proclaimed safe through an abysmal safety case.[2]

[Question to participants: Who has questions about similarities? Or perhaps suggestions for other similarities?]

Assurance cases are not identical to other approaches, 'though, so let's now talk now about some differences.

I'll show two slides about differences. Here is the first.

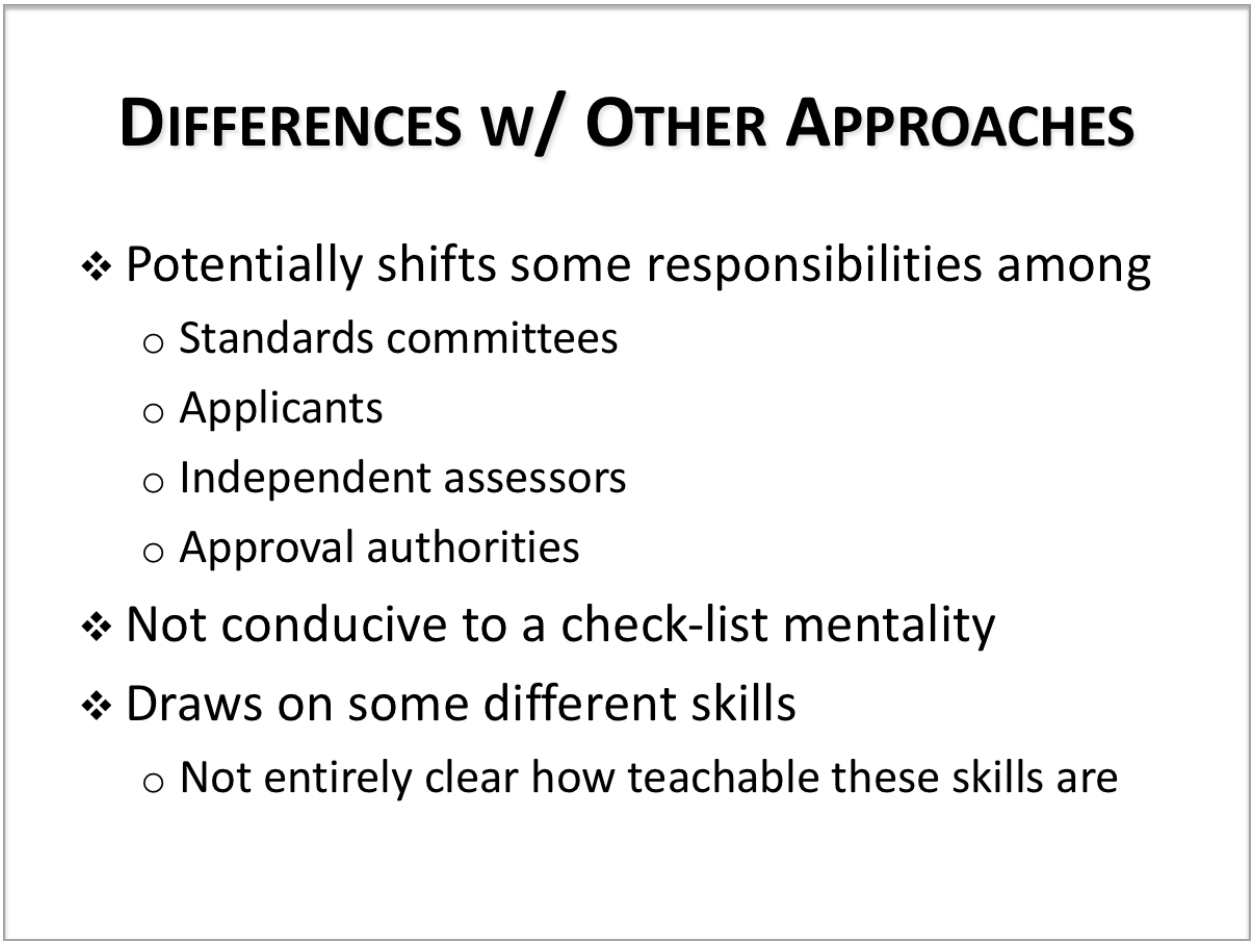

One of the most important differences concerns the potential for shifting responsibilities among various entities such as standards committees, applicants (in FAA terminology), independent assessors, and approval authorities. Exactly how these responsibilities may shift depends on the particular approaches that are taken to employing assurance cases.

[2] Critics of assurance/safety case approaches sometimes cite examples of poorly constructed and inadequately evaluated cases as conclusive evidence of inherent flaws in the approach. Supporters, on the other hand, point out that examples of improper use do not mean that proper use is impossible.

If, for example, a wide-open, un-fettered use of assurance cases is permitted, then it could well be the case that standards committees would be irrelevant, except perhaps if there is a standard for the particular notation used. In such an environment, the role of independent assessors, who would perhaps evaluate the assurance case arguments, may be greatly expanded.

As another example, in a somewhat more structured assurance case environment, perhaps the approval authorities would have a catalog (that's probably not the right word, but I think it may convey the general idea) of acceptable assurance case structures, and applicants would generally be expected to create their arguments using those structures.

Another difference between assurance case approaches and some traditional approaches is that assurance cases are generally not conducive to a checklist mentality. That is, one cannot easily create a checklist against which an argument can be evaluated unthinkingly. You wouldn't for example simply have a checklist that says look for at least three conclusions, six premises, and two reasons. As we saw in Module three evaluating assurance case arguments is not trivial.

Which brings us to the third difference, namely, that the use of assurance cases seems to draw on some different skills than perhaps are usually possessed by engineering organizations and regulators. It is not entirely certain that different skills are essential, but it seems intuitively to be so. At present the jury is still out on how teachable these skills may be. Looking into this issue seems like a fruitful, but difficult, area of research.

We'll talk about research in just a few minutes, but let's continue with the differences.

[Question to participants: Any questions on this slide before I go to the next one?]

Another difference between assurance cases and traditional techniques is that, within the US at least, assurance case methods are less well understood.

We've talked in previous modules about the sorts of mistakes that novices can make; I won't reiterate those here, unless someone wants me to do so.

Perhaps more dangerously, the general lack of understanding means that recognizing actual experts can be hard. Because assurance cases have become a somewhat trendy topic within academic circles people have jumped on the bandwagon without necessarily having the knowledge to contribute anything useful or to even recognize they are unable to do so. (The temptation to go into a rant at this point is great, but I shall resist it.)

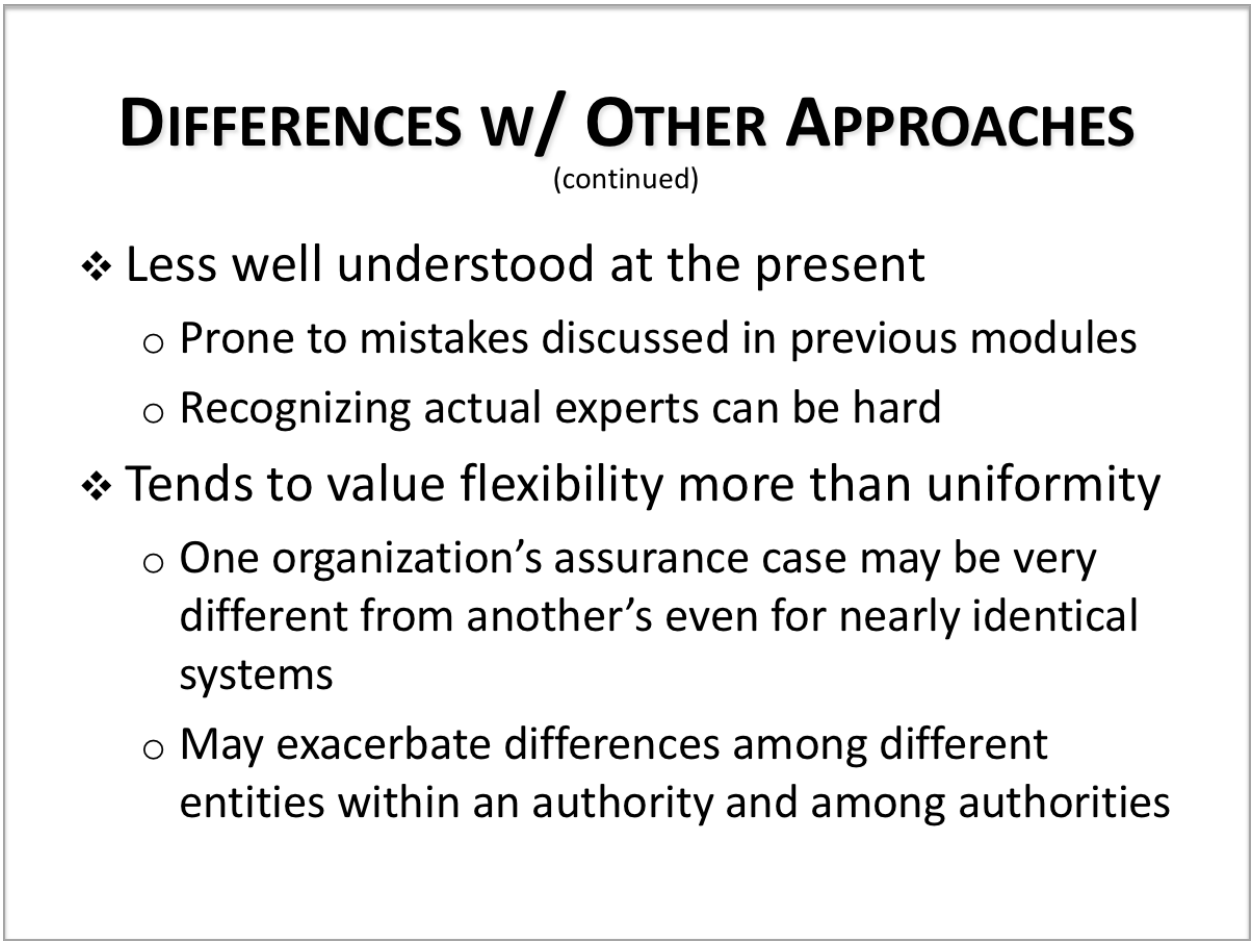

Another difference, which in some ways may be the biggest one, particularly as far as the use in regulatory environments may be concerned, is that the assurance case approach tends to value flexibility more than uniformity.

In a general assurance case regime, one organization's assurance case may look very different from another's even for nearly identical products or systems or subsystems. The case might be structured differently, it might use different notations, it might take it different approach to specifying reasoning, it might take a different approach to addressing defeaters, and so on.

These differences could very well serve to exacerbate already existing differences among entities within a single approval authority and among different approval authorities. A case that is accepted in one region may be rejected in another, for example.

This difference leads directly into our next subject, which is talking a bit about questions that need to be considered concerning using some form of assurance case approach within an FAA regulatory environment[3].

[Question to participants: But before we do that, are there any questions?]

[3] The discussion here is in the context of the FAA environment, but the general questions should be similar, or have analogs, in just about any regulatory situation.

IN FAA REGULATORY ENVIRONMENT?

- Questions to consider include …
  - What's broke that needs fixing?
  - Are people & resources available to facilitate a cultural change?
  - Is it possible to conduct 'clinical trials'?
  - Could two separate but equal approval tracks be established?
  - Might the UAS domain be appropriate as a 'testbed'?

In thinking about assurance cases and the FAA environment, I'm going to suggest five general questions that I think need to be carefully considered. Let's read them all together, then discuss each one a bit more.

One very important question is "What's broke that needs fixing?"

Question two is "Are people and resources available to facilitate a cultural change?"

The third question is, "Is it possible to conduct 'clinical trials'?"

Question four: "Could two separate but equal approval tracks be established?"

And the final question that I propose is, "Might the UAS domain be appropriate as a 'testbed'?

The "what's broke" question is critically important, because its answer may go a long ways towards helping to decide whether some form of assurance case approach is likely to help fix the perceived problems.

As we just discussed, in general assurance case approaches tend to promote flexibility at the expense of uniformity.

If the biggest problems that are currently facing the FAA in terms of the regulatory environment is that the environment is too rigid, that it tends to discourage or even prevent useful innovation, then moving towards an assurance case regime may well help

address those sorts of problems. If, on the other hand, the biggest problems involve inconsistency among different approvers within the FAA or between the FAA and EASA, then moving towards an assurance case regime may not help at all. (I'm not saying it wouldn't be possible to create a specialized assurance case regime that could help with such a situation, just that it may be difficult).

Concerning resources being available to facilitate a cultural change, I think what we've discussed in these series of lessons have made clear that some cultural changes would be necessary. As we just discussed a few minutes ago, there may need to be some skill-set changes, too. Unless resources will be available to make these things happen, moving towards an assurance case approach is not likely to succeed.

The last three questions on the slide all are about the same general theme, "How can you go about establishing that an assurance case approach 'works' well for the FAA?"

I mention the idea of a 'clinical trial' approach because it may provide a fairly inexpensive initial assessment of feasibility. The idea of 'separate but equal' approval tracks is meant to suggest the possibility of allowing organizations to continue what they're doing now if they like, or to try going down an assurance case based track instead[4]. If it turns out that the assurance case based track doesn't work, then all that would be necessary is to remove that track; no changes would otherwise be necessary. Finally, suggesting that the UAS domain might be appropriate as a 'testbed' stems simply from my perception that this area seems to be in a bit of turmoil right now, and trying out assurance case approaches to regulation there might (or might not) help resolve some of the turmoil.

[Question to participants: That's all I've planned to say on this topic, does anyone have some questions?]

We'll move now to talking a little bit about the research that is currently going on in the assurance case / safety case arena. This discussion will be necessarily quite subjective, and you will easily be able to find people who have very different opinions from my own. Please remember that I giving you only my personal thoughts, none of which should be construed to represent an official NASA position.

Shortly I will show you two different lists of current research topics. First, you will see a list ordered by my subjective evaluation of current popularity. The ordering is entirely subjective, but I did ask some other people within the community for their opinions, and they generally agreed with my ordering, with only an occasional exception. Second, you will see a list ordered by my opinion of the priority that ought to be given to the various topics[5].

---

[4] The Overarching Properties work which arose after, and was partially motivated by, the Explicate '78 project is based on applying this principle.

[5] Although these orderings were developed in 2016, I do not think any significant changes (to either side) have happened in the last four years.

The first three areas that you see here — quantifying confidence, formalizing arguments, and generating cases automatically — are almost certainly the currently most popular research areas.

Each of these research areas has some first glance appeal.

If it is possible to place useful numbers on the degree of justified confidence that one should have in an assurance case argument …

Before continuing that sentence, let me explain what I mean by useful numbers.

I mean numbers that can be compared and manipulated, so that, for example, a confidence score of 995 would be known to always be better than a score of 850, and that if a minimum threshold of say 990 was required, we could be sure that a score of 993 indicated sufficient justified confidence.

So, repeating the sentence I started ….

If it is possible to place useful numbers on the degree of justified confidence that one should have in an assurance case argument, then having such numbers would seem to be clearly a good thing.

Similarly, if it is possible to formalize assurance arguments, particularly to make them purely deductive (if you don't remember from Modules 1 & 3 what purely deductive

means then ask, or look it up if you're reading the material), then formalizing them seems like a good thing, too. Much of the evaluation of formal arguments could be done automatically. And, if it is possible to generate arguments automatically, based purely on things that engineers are already doing, then this, too, seems to be a great thing to do.

The next three items that you see on the slide are also being research fairly actively.

Exploring modularity & composition refers to efforts aimed at creating arguments that can be reused directly in other contexts and to developing ways to compose existing arguments into higher-level arguments without having to change or reevaluate the individual original arguments.

Creating & extending notations is pretty self-explanatory.

Developing argument patterns is a bit similar to the modularity and composition idea, but on a different scale. Rather than trying to create completely reusable arguments, pattern research seeks to create general frameworks for certain types of arguments, which then may be instantiated with system specifics as necessary.

The final item, you see here, assessing efficacy, refers to efforts to determine whether, and if so, how, assurance cases truly provide the benefits that proponents claim. Think back to Mike's comment to Jon. To date, all of the efforts in this area have tended to involve case studies, retrospective evaluations, or non-public proprietary studies.

[Question to participants: Any questions about what I mean by any of these areas?]

I will now show you a different ordering and slightly different set of research areas that corresponds to what I personally think ought to be going on.

Once again, please remember that I am showing you *only* my opinion. Plenty of smart people within the assurance case research community disagree with me.

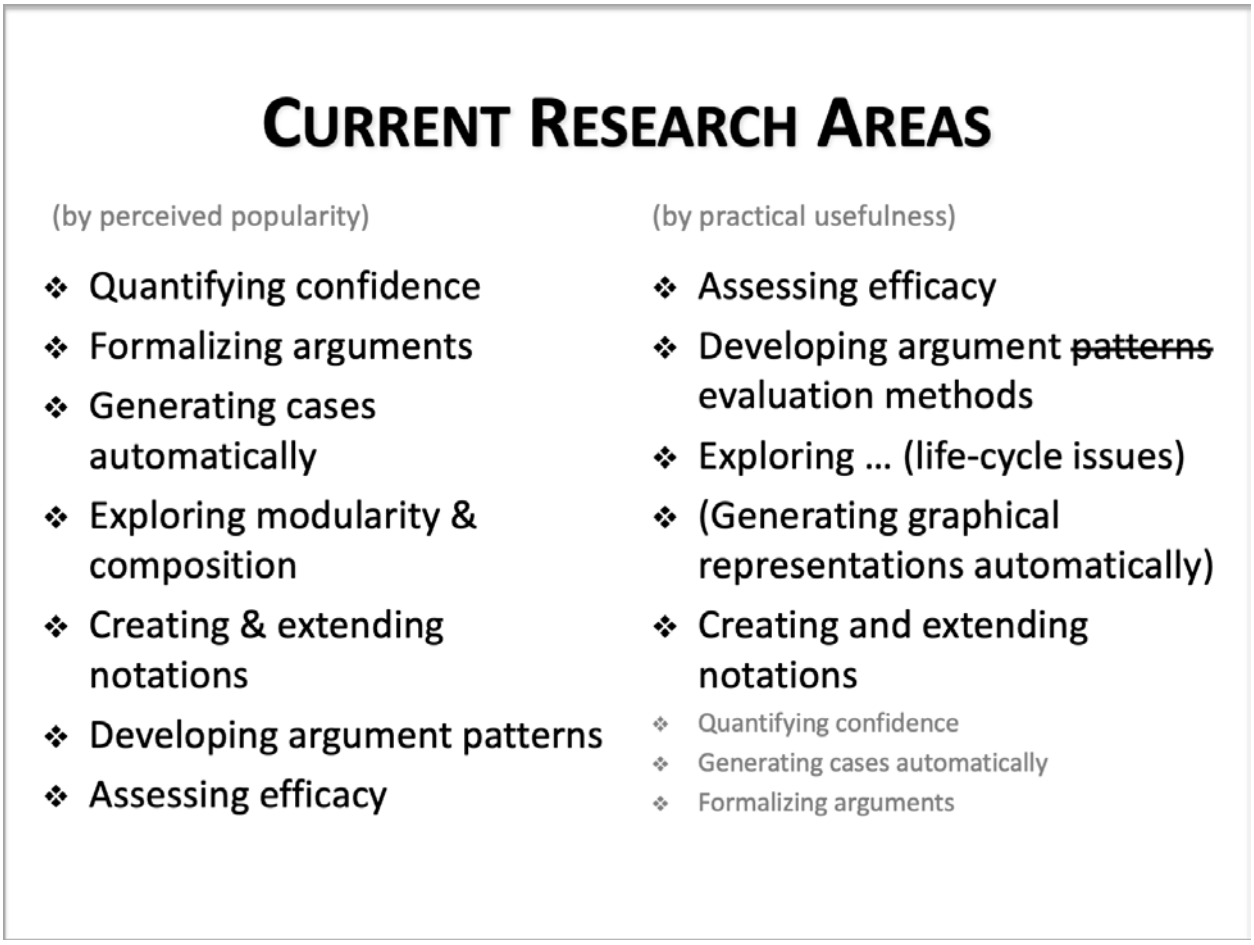

Perhaps the first thing you'll notice is that the order is almost directly inverted from the current popularity order. Or perhaps the first thing you'll notice is that I have written the three areas that are currently the most popular in grey and a small font. I did this because, despite the first-glance intuitive appeal of these areas, in practice, given the current state of the art and state of our knowledge, the "if" clauses for all three are practically false.

It is *not* possible to generate useful numbers. Well, to be more precise, none of the proposals thus far to do so can withstand scrutiny[6].

It is *not* possible to formalize important parts of assurance case arguments. The concepts with which these arguments are concerned are often not formal concepts themselves, but rather emergent, non-deductive properties that can't be described precisely in any existing logical formalism.

And finally, it is *not* feasible to generate very many useful assurance case arguments automatically, partially because automatic generation assumes some sort of formalization.

---

[6] See the journal article [Graydon, P.J., Holloway, C.M. 2017. An Investigation of Proposed Techniques for Quantifying Confidence in Assurance Arguments. *Safety Science,* vol. 92, pp. 53-65.] and the significantly longer and more detailed technical report [Graydon, P.J., Holloway C.M. 2016. An Investigation of Proposed Techniques for Quantifying Confidence in Assurance Arguments. NASA/TM-2016-219195.] for the results of applying scrutiny to existing quantification techniques.

I think that the current top three most popular areas of research should be mostly abandoned, or left entirely in the hands of academic departments who have no interest in practicality.

The area that is currently the least studied, should be, in my opinion, the most studied, namely assessing efficacy: determining whether assurance cases truly provide the benefits proponents claim they do. Such research is not easy, nor is it cheap, nor is the sort of thing that seems to be currently in vogue with funding agencies at the moment, but I think it is critical. We at NASA have started a bit of work in this area, and are hoping to be able to expand the work further. Perhaps you will be reading about the results of the work one day.

I won't go into anything more about these other areas unless someone has a question.

Let's now talk a bit about what you can do to further your study of assurance cases.

(At this point in the original presentation, I showed three slides. On each slide was a list of five references for further study. Since the original presentation in 2016, I have revised my recommendations slightly, and created three different priority orderings. Rather than replicate the original slides, I will show the new material at the end of this document.)

Following the practice we established in Module 1, we will review the learning objectives, formulated as questions.

REVIEW OF LEARNING OBJECTIVES

Are you able to

- Compare and contrast an assurance case approach with other approaches?
- Discuss how an assurance case approach could fit into a regulatory environment?
- List current areas of assurance case research?
- Locate references for further study?

Every man takes the limits of his own field of vision for the limits of the world. - Arthur Schopenhauer

Think to yourself how you'd answer these questions.

After you've thought about the questions for a little bit, I'll end with the superb quotation from the Nimrod report that I used back in Module 2.

> "At all stages of the safety pilgrimage it is vital to ask questions such as 'What if?', 'Why?', 'Can you explain?', 'Can you show me?', 'Can you prove it?'. Questions are the antidote to assumptions, which so often incubate mistakes."
>
> "A Questioning Culture is the key to a true Safety Culture. In my view, people and organisations need constant reminding of the importance of asking questions rather than making assumptions, of probing and testing rather than assuming safety based on past success, of independent challenge of conventional wisdom …, of the exercise of judgment rather than retreat behind the assignment of arbitrary quantitative values."
>
> "Questioning is a catalyst for thinking. As Professor McDermid told me, if he could replace all of the regulations with one word it would be: 'THINK'".
>
> *Haddon-Cave, C. (2009) The Nimrod Review. London: The Stationary Office. p. 574.*
> *www.official-documents.gov.uk/document/hc0809/hc10/1025/1025.pdf*

If you remember nothing else from these five modules about Understanding Assurance Cases, please remember those words. If you cannot remember all of these words, then at least remember Professor McDermid's single word: **think**.

Thank you for your attention, and I'll be happy to field any remaining questions or comments about this module in particular, or the whole series in particular.

Thus ended the educational presentations.

If you have questions or comments about this material, contact its author at `c.michael.holloway@nasa.gov`.

Recommendations for additional reading.

These suggested references are intended to provide a broad overview of philosophy, principles, and practices associated with the assurance / safety case approach to obtaining confidence in the safety and efficacy of systems and services. Reading all of the suggested references will not tell you everything you need to know, but it should provide you with the knowledge that is needed to understand most everything else that you will encounter. The length of the material various considerably, from a low of 6 pages to a high of nearly 600 pages.

No single one of the references is complete in itself. Also, some of the references take points of view that are different from others. Inclusion on the list does not imply endorsement of the content.

All of the listed references except for the Toulmin book are available for free in electronic form. The lists below include URLs that worked as of 14 July 2020.

Three different suggested reading orders are provided: one for students, researchers, and the simply curious; one for practicing engineers and approval authorities; and one for managers, which contains only five suggestions.

Recommended order for students, researchers, and the curious

1. The Uses of Argument (Updated edition) Toulmin, S. E. (2003, 1958). This book must be purchased. One place to get it is www.amazon.com/Uses-Argument-Stephen-E-Toulmin/dp/0521534836/

2. The Safety Argumentation Schools of Thought. Graydon, P. J. (2017). hdl.handle.net/2060/20180000378

3. A Taxonomy of Fallacies in System Safety Arguments. Greenwell, W. S., et al (2006). hdl.handle.net/2060/20060027794

4. Current Practices in Constructing and Evaluating Assurance Cases With Applications to Aviation. Rinehart, D. J., Knight, J. C., & Rowanhill, J. (2015). hdl.handle.net/2060/20150002819

5. The Purpose, Scope, and Content of Safety Cases. Office for Nuclear Regulation (2013). www.onr.org.uk/operational/tech_asst_guides/ns-tast-gd-051.pdf

6. Arguing Safety - A Systematic Approach to Managing Safety Cases. Kelly, T. P. (1998). www-users.cs.york.ac.uk/tpk/tpkthesis.pdf

7. Reviewing Assurance Arguments: A Step-By-Step Approach. Kelly, T. P. (2007). www-users.cs.york.ac.uk/~tpk/dsnworkshop07.pdf

8. A New Approach to Creating Clear Safety Arguments. Hawkins, R., Kelly, T., Knight, J., & Graydon, P. (2011). www.cs.virginia.edu/~jck/publications/SSS.2011.safety.cases.pdf

9. The Nimrod Review: An independent review into the broader issues surrounding the loss of the RAF Nimrod MR2 Aircraft XV230 in Afghanistan in 2006. Haddon-Cave, C. (2009). www.official-documents.gov.uk/document/hc0809/hc10/1025/1025.pdf

10. The Friendly Argument Notation (FAN). Holloway, C. Michael. (2020). shemesh.larc.nasa.gov/arg/fantm.pdf

11. Regulatory Report: Chevron Richmond Refinery Pipe Rupture and Fire. U. S. Chemical Safety and Hazard Investigation Board (2014). www.csb.gov/assets/1/20/chevron_regulatory_report_06272014.pdf

12. Certification and Safety Cases. Graydon, P., Knight, J., & Green, M. (2010). www.cs.virginia.edu/~jck/publications/ISSC.2010.pdf

13. Assurance cases and prescriptive software safety certification: A comparative study. Hawkins, R., Habli, I., Kelly, T. P., & McDermid, J. (2013). www.sciencedirect.com/science/article/pii/S0925753513001021

14. Explicate '78: Uncovering the Implicit Assurance Case in DO-178C. Holloway, C. M. (2015). hdl.handle.net/2060/20150009473

15. An Investigation of Proposed Techniques for Quantifying Confidence in Assurance Arguments. Graydon, P. J., Holloway, C. M. (2016). hdl.handle.net/2060/20160006526

Recommended order for practicing engineers & approval authorities

1. The Purpose, Scope, and Content of Safety Cases. Office for Nuclear Regulation (2013). www.onr.org.uk/operational/tech_asst_guides/ns-tast-gd-051.pdf

2. Regulatory Report: Chevron Richmond Refinery Pipe Rupture and Fire. U. S. Chemical Safety and Hazard Investigation Board (2014). www.csb.gov/assets/1/20/chevron_regulatory_report_06272014.pdf

3. Current Practices in Constructing and Evaluating Assurance Cases With Applications to Aviation. Rinehart, D. J., Knight, J. C., & Rowanhill, J. (2015). hdl.handle.net/2060/20150002819

4. The Safety Argumentation Schools of Thought. Graydon, P. J. (2017). hdl.handle.net/2060/20180000378

5. Arguing Safety - A Systematic Approach to Managing Safety Cases. Kelly, T. P. (1998). www-users.cs.york.ac.uk/tpk/tpkthesis.pdf

6. Reviewing Assurance Arguments: A Step-By-Step Approach. Kelly, T. P. (2007). www-users.cs.york.ac.uk/~tpk/dsnworkshop07.pdf

7. The Nimrod Review: An independent review into the broader issues surrounding the loss of the RAF Nimrod MR2 Aircraft XV230 in Afghanistan in 2006. Haddon-Cave, C. (2009). www.official-documents.gov.uk/document/hc0809/hc10/1025/1025.pdf

8. Assurance cases and prescriptive software safety certification: A comparative study. Hawkins, R., Habli, I., Kelly, T. P., & McDermid, J. (2013). www.sciencedirect.com/science/article/pii/S0925753513001021

9. Certification and Safety Cases. Graydon, P., Knight, J., & Green, M. (2010). www.cs.virginia.edu/~jck/publications/ISSC.2010.pdf

10. Explicate '78: Uncovering the Implicit Assurance Case in DO-178C. Holloway, C. M. (2015). hdl.handle.net/2060/20150009473

11. A Taxonomy of Fallacies in System Safety Arguments. Greenwell, W. S., et al (2006). hdl.handle.net/2060/20060027794

12. A New Approach to Creating Clear Safety Arguments. Hawkins, R., Kelly, T., Knight, J., & Graydon, P. (2011). www.cs.virginia.edu/~jck/publications/SSS.2011.safety.cases.pdf

13. The Friendly Argument Notation (FAN). Holloway, C. Michael. (2020). shemesh.larc.nasa.gov/arg/fantm.pdf

14. An Investigation of Proposed Techniques for Quantifying Confidence in Assurance Arguments. Graydon, P. J., Holloway, C. M. (2016). hdl.handle.net/2060/20160006526

15. The Uses of Argument (Updated edition) Toulmin, S. E. (2003, 1958). This book must be purchased. One place to get it is www.amazon.com/Uses-Argument-Stephen-E-Toulmin/dp/0521534836/

Recommended order for managers

1. The Safety Argumentation Schools of Thought. Graydon, P. J. (2017). hdl.handle.net/2060/20170007188

2. The Purpose, Scope, and Content of Safety Cases. Office for Nuclear Regulation (2013). www.onr.org.uk/operational/tech_asst_guides/ns-tast-gd-051.pdf

3. Regulatory Report: Chevron Richmond Refinery Pipe Rupture and Fire. U. S. Chemical Safety and Hazard Investigation Board (2014). www.csb.gov/assets/1/20/chevron_regulatory_report_06272014.pdf

4. Current Practices in Constructing and Evaluating Assurance Cases With Applications to Aviation. Rinehart, D. J., Knight, J. C., & Rowanhill, J. (2015). hdl.handle.net/2060/20150002819

5. An Investigation of Proposed Techniques for Quantifying Confidence in Assurance Arguments. Graydon, P. J., Holloway, C. M. (2016). hdl.handle.net/2060/20160006526