Flight Test *Safety* Fact

https://www.flighttestsafety.org

Carnival of Mathematics

One of the aperiodic publications I enjoy is known as the Carnival of Mathematics. Now some of you are probably running in fear from the terror clowns that your mathematics education reminds you of—somehow, I think it's appropriate to be afraid of clowns during the month of October. However, I'm also sorry that someone ruined the beauty and joy of mathematics for you. But I can assure you that there is nothing to fear in the occasional blog posts that make up the Carnival of Mathematics, but instead it has all the good stuff of nostalgic carnival memories. I recommend checking it out, because much of what you find there is the kind of stuff that will delight the novice and the expert alike: https://aperiodical.com/carnival-of-mathematics/.

In fact, the CoM might be something we can learn from. The way they present miscellany in a whimsical, enjoyable, digestible format is something I think we should consider. On the one hand, there's a low barrier of entry to "submission," which is always a bonus. That leads me to two related ideas, like a gremlin splashed with water.

The most recent SETP Cockpits magazine had letters from both the President and the Editor asking for submissions, which is a common refrain. I reached out to both and asked them: "What motivates someone to submit to Cockpits?" So I pose the same question here: What motivates someone to write? What fire is burning inside a person that overcomes all the organizational roadblocks to presenting at a conference? And how do we keep that fire burning long enough to turn a "talk" into an article or column or letter to the editor? Is it some kind of street cred? Would the motivation change if the format was different? That's the first idea.

The second is similar but focused on a different domain: What motivates someone to create a branch of a repo on Github? To contribute to an open-source software engineering project? To learn git? And do any of these questions shed any light on what a better, more organic, crowd-sourced way of building a repository of THAs might look like? And can we use Github to host this repo?

I got a question at the FTSW in May: How do we build a crowd-sourced way to share things we are actually allowed to share, like THAs? I received a similar question at the same event: How do we crowd-source the data on Flight Test Accidents? And I think they are related, and something similar to how SFTE crowd-sourced the most recent update of their FTE Handbook, creating a web-based version, seems like a fruitful line of effort.

Another thing I want to explicitly address is the format of this newsletter. It's both a throwback to an earlier time, a nod to our heritage and history, and it's a way to "thrill the reader" with something spooky for the season we are in. At least one reader will be "terrified" by the format, but such is a feature of the season. I included this format along with the following list of "orange stories" I tell during this time of year as part of a recurring theme:

"Embedded deep in the lore and legend of flight test is a uniquely colored fiber woven into the fabric of who we are and what we do. The color of that fiber is flight test orange. It appears on our patches and in our logos, but it also shows up in our aircraft, as flight test instrumentation, orange wire, and emergency equipment…

Read more by clicking any of the hyperlinks to past features:

Orange Wire from Down Under
Flight Test Orange Wire
Wear the Orange
Flight Test Orange

In other scary news ,the FAA has released its roadmap for AI:
https://www.urbanairmobilitynews.com/emerging-regulations/faas-initial-roadmap-for-ai-safety-assurance-opens-door-to-collaboration-on-autonomy/
https://www.faa.gov/aircraft/air_cert/step/roadmap_for_AI_safety_assurance

For more Carnival of Flight Test, scroll to Page 4.

The Spooky Shadow of Uncertainty – Arriving at the Signal in the Noise

What is a Markov chain Monte Carlo model, and why does it matter? What is the difference between RNAV and RNP? What is a data band? How is it different than tolerance? Why do we use them?

These are all *significant* questions, and their answer hides like a valley masked by the long shadows of towering mountains. These peaks grab our attention. They are a more familiar sight like the primary disciplines of flight test engineering—performance, flying qualities, test conduct, and systems evaluations. But in the misty fog-filled valleys below… The shadows are spooky, and things go bump in the night.

The challenge for flight test in the twenty-first century is to illuminate the data, forcing the shadows of uncertainty to flee. Then, we can transform the data into insight and information, unlocking its potential like one would unlock nuclear energy. As flight test professionals, we must adapt to the evolving nature and quantity of flight test data. We must ***continue*** to familiarize ourselves with a more diverse array of *applied tools of mathematics and statistics*, building blocks to understanding data and using it to make decisions. Expert application of these tools is critical if we will successfully navigate the shadow of uncertainty and find the signal in the noise. *Noise* – this is the most common statistical phenomena in flight test, and understanding it better is the goal of the second half of this discussion, an objective to which we will return. This discussion, though, is just part of a larger strategic discussion—one that has been conducted rigorously, among other places, in the break room and cubicles of many flight test departments, symposia, and other virtual venues. I believe that it ought to include the following points:

1. Probability is as important as airmanship;
2. The important thing is not (necessarily) the formula for standard deviation or any probability distribution but the big ideas, the fundamental principles and the way our knowledge of them guides our thinking; and
3. We need to communicate a clear and convincing explanation for flight test professionals in some format less than dissertation length.

To illustrate these points, consider the following example: learning to navigate on a completely "unscary" cross country flight.

I want you to recall an elementary idea, a cross-country flight. Do you remember back to the days when you were learning how to fly? For me, the plane was a Cessna 152, tail number four-hotel-bravo, and the place was Cook County Airport (15J) in southern Georgia. My instructor's name was Ian. He flew seaplanes somewhere in the South Pacific for many years before teaching private pilot students.

When I walked into the flight school one day, Ian told me something I would never forget. In fact, he predicted I would never forget it before he even told me: "Can ducks make vertical turns with turbulence." That mnemonic helps me remember the steps needed to plan a cross-country flight. The fact is, a pilot's head is full of crazy sayings and silly words that mean something when translated into aviation jargon. Remembering a wacky sentence about ducks is easier than remembering Compass heading, Deviation, Magnetic heading, Variance, True heading, Wind correction, and True course. Back at Cook County Airport, I opened the sectional charts and sat down to figure out where I wanted to go. Once I did, I could draw a single straight line on my chart and jot down a heading. Those two things would get me pretty close: 255° magnetic for 25 minutes. After five minutes on a 255° heading, I should cross over a major highway with an overpass to my left. Once airborne, I look outside the airplane to see where it is. It's not as far south as I thought it would be. At seven-and-a-half minutes, I should pass over the southern tip of a large pond. The pond is just north of my position. Apparently, the winds are drifting me south of my intended course. I correct my heading to 260°. At 15 minutes, I should overfly an intersection in a small town. I am just north of the intersection. A heading of 260° corrected me back to course and then a bit right. Two-five-seven is right in between. That should keep me on course.

Three steps—clock to map to ground—are the process we follow when navigating in an aircraft. There is this notion, "the pond is just north of my position." Being able to recognize that and make that judgment call is a critical element of airmanship. It's also a fundamental principle of applied tools of mathematics and statistics. Sometimes, "just north of my position," is close enough. In this case, we don't need to *quantify* what we mean by "just north of my position." Stating it qualitatively is just right.

In summary, I want to explicitly state three fundamental facts.

1. We are going to encounter uncertainty—uncertainty means we won't hit every waypoint.
2. Predict-test-evaluate is the process for navigating uncertainty. In aviation, this means plan the flight— fly the plan using the "Clock-map-ground" technique—and evaluate, using your "engineering judgment," when that's close enough.
3. Applied tools of mathematics and statistics help us evaluate when "that's close enough" (as in the case of navigation above) or when more quantitative rigor is needed.

What's next is the Spooky Part of the Story…
(continued Page 2)

At the beginning of this column, I promised to return to the subject of noise, the most common statistical phenomena in flight test. Henceforth, this is the subject of our discussion and a beta test of what I hope will be an evolution in the math section of the SFTE Reference Handbook.

Introduction to Measurement Error

In flight test as in other disciplines, measurement and calculation result in different types of errors. Some of these errors are not systematic but are random. Often we call these measurement errors noise, and it is these errors that are the focus of this section. Furthermore, it is a commonly accepted practice that the normal distribution is a suitable model for noise and measurement errors. It is imperative to emphasize that the normal distribution is just a model, a simplification of the physical world. The purpose of this next section is to demonstrate why this is a suitable model and a very relevant one.

Suppose that we are going to measure airspeed, x, with some transducer. Suppose further that at each step in the measurement process we have one of two hypothetical outcomes:

1. We measure airspeed correctly; that is, the error is zero: $\varepsilon = 0$.
2. Or there is some error in our measurement of airspeed, which we model as follows: $\varepsilon = 1$.

In other words, we have a model that returns 0 when there is no error and 1 when there is. We can show this in tabular form as follows:

Measurement process	After a single step or factor in measurement process: $x + \varepsilon$
Possible outcomes	x
	$x + 1$

Suppose that there are two steps that affect the given measurement. For example, measurement of airspeed requires both static and dynamic pressure. The error term propagates at each step. So we have the following:

Measurement process	After a single step or factor: $x + \varepsilon$	After a second step or factor: $x + \varepsilon$
Possible outcomes	x	x
		$x + 1$
	$x + 1$	$x + 1$
		$x + 2$

After two steps, we can have any of three possible outcomes, x, $x + 1$, or $x + 2$. But the middle outcome occurred twice. Imploring upon your patience, consider outcomes after three steps.

Measurement process	After a single step: $x + \varepsilon$	After second step: $x + \varepsilon$	After third step: $x + \varepsilon$
Possible outcomes	x	x	x
			$x + 1$
		$x + 1$	$x + 1$
			$x + 2$
	$x + 1$	$x + 1$	$x + 1$
			$x + 2$
		$x + 2$	$x + 2$
			$x + 3$

Another way to see how these tables propagate is in the tree of figure 1. At each node of the tree, the value in the node represents the cumulative error, and the two branches indicate that there are two possible outcomes either $+0$ or $+1$.

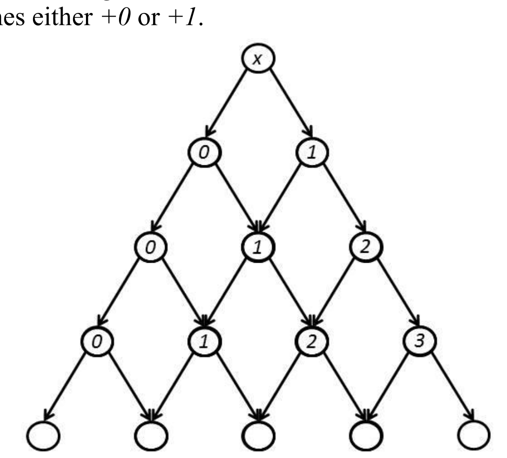

Figure 1 – Propagation of error term in a measurement

Thus after three steps or factors in the measurement process, there are four possible unique outcomes, x, $x+1$, $x+2$, $x+3$, but two of these outcomes occur more than once. In other words, there is more than one possible path through the tree to certain nodes. We can record the different outcomes and the frequency with which they occur in a rudimentary table as follows:

Number of occurrences	1	1 1 1	1 1 1	1
Outcome, $x + \varepsilon$	$x + 0$	$x + 1$	$x + 2$	$x + 3$

This table allows us to picture qualitatively, based on the height of the tallies, the relative frequency of each particular outcome. It also allows us to quantitatively compute the probabilities of a given outcome. Thus, $P(x+0)$ is the probability of no error in the measurement, and it is a ratio given by (number of times *given outcome* occurs) / (number of total possible outcomes) = 1 / 8; in other words, the number of tallies in a given column / total number of tallies.

The reader may continue this exercise for several more iterations and see two important principles:

1. Simple probabilities like the toss of a coin (a fifty-fifty chance) can quickly compound and propagate into more complex distributions.
2. There is a formula with which we can compute the number of occurrences in this discrete model, the binomial distribution.

Definition: Binomial Distribution

The *binomial distribution* is the probability of obtaining exactly n outcomes in N trials.

$$P(n) = \binom{N}{n} p^n (1-p)^{N-n}$$

From our example above, we could compute the probability of obtaining n errors in N stages or factors of the measurement process, that is, P(obtaining $x + \varepsilon$ as the outcome) for $\varepsilon = 0, 1, 2$, or 3. However, we must define a few more terms in the formula above. In this case $N = 3$ when we examine the outcome after 3 stages or factors in our measurement process.

(*See table next page.*)

Number of occurrences	1	1 1 1	1 1 1	1
Outcome, $x+\varepsilon$	$x + 0$	$x + 1$	$x + 2$	$x + 3$
$P(x + \varepsilon)$	$P(\varepsilon=0)=1/8$	$P(\varepsilon=1)=3/8$	$P(\varepsilon=2)=3/8$	$P(\varepsilon=3)=1/8$

We let $n = 0, 1, 2, or\ 3$, based on whether we want to know the probability of the outcome $x+0$, $x+1$, $x+2$, or $x+3$, respectively. Additionally, p is the probability of the error at each stage. For the purpose of our example, we can say that $\varepsilon = 0$ or 1 are equally likely, and thus we assign $p = 1/2$. We leave it to the reader to understand the formula $\binom{N}{n}$, but suffice it to say that this term allows us to compute the number of different ways a given outcome may occur. In this example there were three possible ways to arrive at $x + 1$, something we can see in both the tree and the tabular depiction above—the $\binom{N}{n}$ term allows us to compute it rigorously in any situation. Microsoft Excel or Google spreadsheets, MATLAB or python, and many other tools have functions that allow us to compute these probabilities.

Modeling Error with the Normal Distribution

Up to this point, our example has highlighted use of the binomial distribution, a model capable of handling discrete cases. In other words, we can compute the probability for any $N = 1, 2, 3\ldots$ including any whole number value. However, this model cannot accept continuous or fractional values, and it is cumbersome for even nominally large values of N—it becomes an unnecessary burden on memory and computational resources. There is a natural relationship between the binomial distribution and the normal distribution, and the bell curve is an excellent model for continuous and fractional measurement errors.

Consider the bar chart of probabilities of the binomial distribution with $N = 30$ and $p = ½$. The chart shows us all the possible values of $P(n)$, the probability of n errors in $N = 30$ trials. For example, if we compute $P(n=15)$ using the formula for the binomial distribution, we would find that $P(15) = 0.14446$. We also see that the height of the bar at $n=15$ is 0.14446. The x-axis depicts n, and we can see that it ranges from 0 to 30, and the y-axis is the probability $P(n)$, a number between 0 and 1. Recall that in our example of propagating errors, we were counting the number of errors (ε) after a certain number of steps or factors. Thus n is equivalent to the number of errors, and N is equivalent to the number of steps or factors. As we can see, the shape of this discrete distribution begins to resemble the familiar bell curve shape of the normal distribution.

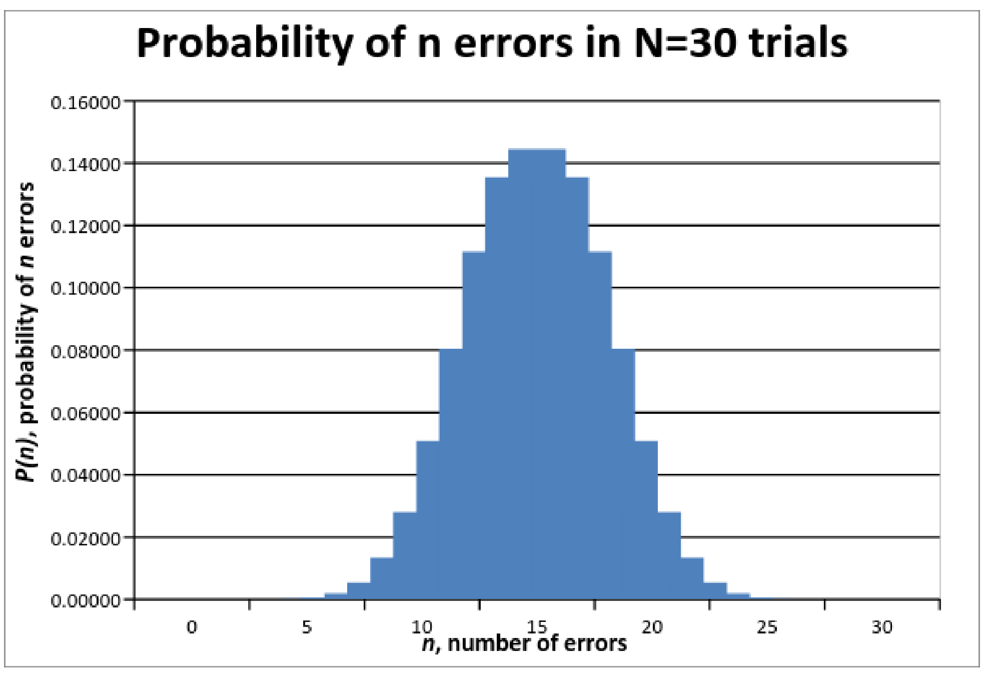

Consider now a continuous model of measurement error. Suppose again that we are going to measure airspeed, x, with some transducer. Suppose further that at each step in the measurement process we can have fractional errors. In other words, we measure 202.5 when the truth is $x = 200$ knots. Here we have as the error term $\varepsilon = 2.5$, as an example. This is more like the physical reality than the binomial example above. To describe it adequately, we need the following additional definition.

Definition: Normal Distribution

The bell curve is formally known as the *normal distribution* and is the continuous probability distribution given by the probability density function below, where μ, the mean, and σ, standard deviation, are given.

$$f(x) = \frac{1}{\sigma\sqrt{2\pi}} e^{-\frac{(x-\mu)^2}{2\sigma^2}}$$

Practically speaking, the parameters μ and σ, help us define the shape of the bell curve, whether it is tall and skinny or short and fat, for example.[1] There are two ways to observe our measurement error. We could plot a bell curve centered at 200 knots, the truth value, or we could subtract our measured value from the truth value and obtain an error term, $\varepsilon = 202.5 – 200 = 2.5$. In this second case, our curve would be centered at 0.

When we plot flight test data, we normally plot the raw values, so the former may occur more naturally. Additionally, plotting error terms is not always analytically tractable, and thus it is advantageous to plot the raw data. However, strictly speaking, it is the error term that we model with a normal distribution, not the airspeed term. Therefore, one must apply great care to avoid the mistakes in reason caused by a misunderstanding of what data are actually normally distributed.

This has been a very quick introduction to noise and the normal distribution, a vital tool that we use to model noise in flight test data, and we will close with one example that uses this tool: RNAV and RNP.
Flight test is the place where validation of the accuracy of navigation systems occurs by comparison to a truth position source, and understanding the normal distribution is essential to applying the definitions of RNAV and RNP.

Figure 1 is from an FAA advisory circular and illustrates the definition of RNAV. Ninety-five percent is a common value used in many confidence intervals, based on approximately two standard deviations from the mean, and it is this concept used in the definition of RNAV. We won't take the time to discuss the subtleties of confidence intervals here, however.

In this second illustration, also from an FAA advisory circular, we see the definition of RNP. Again, confidence intervals are brought to bear on the navigation position, but this time, we see a much higher level of confidence, one exceeding three standard deviations from the mean.

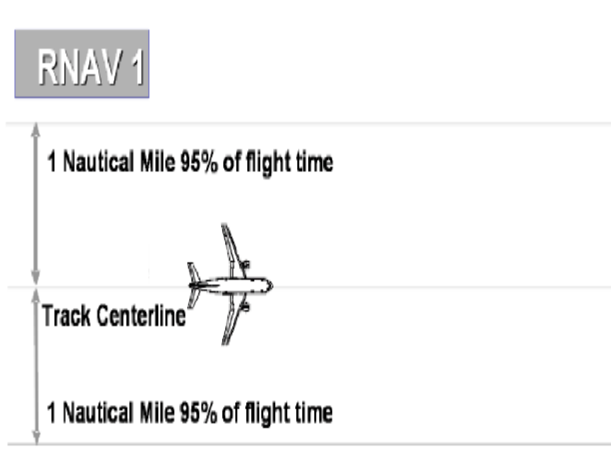

Figure 1 – RNAV lateral position error

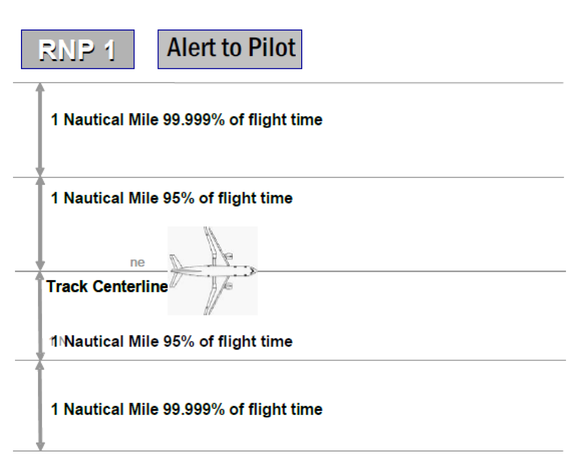

Figure 2 – RNP lateral position error

[1] See http://en.wikipedia.org/wiki/Normal_distribution for examples of normal distributions with different shapes.

Other Signs of the Season

September 28, 2024 marks one hundred years since the first successful global circumnavigation. How did I miss that? https://pioneersofflight.si.edu/content/first-flight-around-world.

There is a three-episode recap of the event on the Fighter Pilot Podcast, which has golden hues, which fits our fall colors theme. https://podcasts.apple.com/us/podcast/world-flight-centennial-1-of-3-build-up/id1330534712?i=1000669704016

Also if you have an interest in Flight Control Systems from the perspective of a Test Pilot but written for ordinary aircrew, check out Episode 147 .

I heard an interesting (open) question during the podcast which prompted me to post the question on Aviation Stack Exchange. The website values questions with answers that have definitive references, not opinions, and it too is a research project into the efficacy of organic, high-engagement audiences, something we should consider.

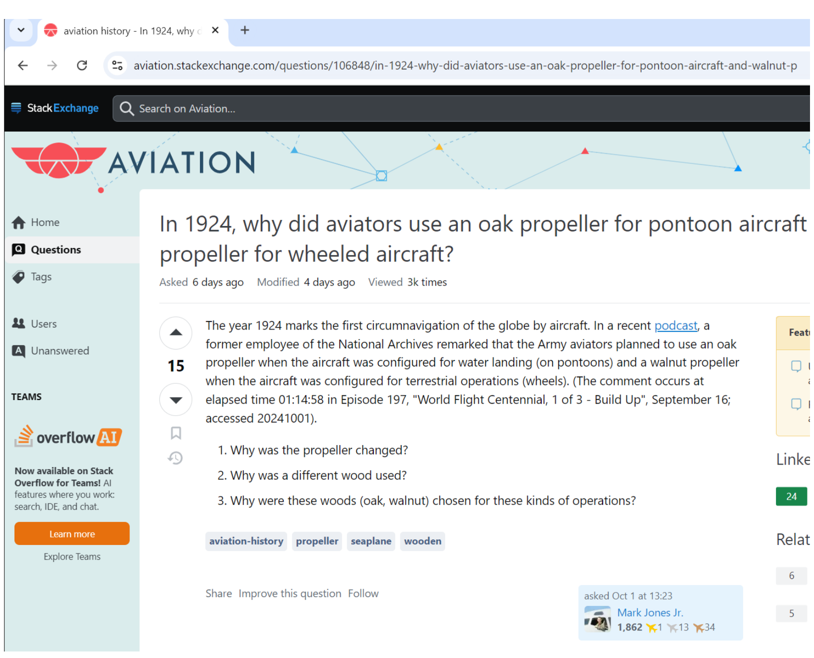

Contact Flight Test *Safety* Fact

Stuart "Chia" Rogerson, Chairman chairman@flighttestsafety.org

Susan Bennett, FTSC Administrator susan@setp.org

Society of Flight Test Engineers edir@sfte.org

Society of Experimental Test Pilots setp@setp.org

AIAA Flight Test Group derek.spear@gmail.com

Connect with us by joining the LinkedIn Group: "Flight Test Safety Committee."

Website: flighttestsafety.org

Chia Chat

Fall is upon us and with it, kids, college students and teachers headed back to school. To ensure you didn't miss out on that academic opportunity, I hope you enjoyed Mark's refresher on normal distributions and how it applies to flight test. It's okay if you need to be excused because your brain is full! My certainly was. However, when it comes to safety, I lean more to non-Gaussian distribution functions. If we go back to 2015 and the 59th SETP Annual Symposium and Banquet, Daniel "Animal" Javorsek II presented a fantastic paper entitled: "Modernizing the Safety Process: Complexity Theory for Flight Test Professionals." There is plenty more math in that presentation if your brain isn't quite full, but I really like how Animal relates probability to Thanksgiving.

Here is a short excerpt from that paper which are words of wisdom we should strive to live by as flight test professionals:

> One of the fallacies that continuously affect our decisions is how we interpret past data. In a world governed by Gaussian distributions, confidence in extrapolation into the future is bolstered by successes due to each outlier's low probability and ultimately a matter of incomplete information. For example, consider the life of a Thanksgiving turkey. Initially, he is concerned about the farmer and his intentions. However, after three years of being fed every morning, safety is apparently confirmed, and on Thanksgiving morning he boldly, yet naively predicts the same positive outcome. If the turkey instead lived in a world governed by non-Gaussian distributions, he might have a more shrewd outlook and view each feeding as confirmation of the persistent danger of becoming dinner.
>
> Since highly unexpected events have occurred in flight test when past data suggested that a maneuver or technique worked means that the shrewd outlook of the Turkey may be more appropriate. **In short, high risk test points accomplished uneventfully in the past provide no guarantee of future risk reduction and, to the contrary, may mean that we are simply primed for the conditions to change just enough that a mishap is imminent.** The intent of this discussion is not to make us paranoid but rather to point out our own limitations and biases as well as to honestly acknowledge and identify uncertainties. To use the words of Nassim Taleb, "It is much more sound to take risks you can measure than to measure the risks you are taking."

Fortuitously, the next paper presented immediately after Animal's was "Lessons Learned and Murphy's Corollary" by Michael Meier, Wills Wing, Inc and it was a perfect practical example of the previous paper from the world of hang gliders, where continuous past successes led to complacency and then a mishap occurred. I recommend you check out both papers sometime for a good refresher. For me, this all circles back to my non-paranoid, but pessimistic review of Test Hazard Analysis and test team assumptions when planning flight test. Never assume it "should work" in your risk mitigation and just because the system or the test worked in the past doesn't mean it will in the future. Gaining knowledge and reducing our uncertainty is a big risk mitigator, but do really know how much margin your successful outcome had left before disaster would have struck?

On the topic of papers from SETP's Annual get together, I just returned from the 68th Symposium and Banquet where there was a fantastic line up of presentations and great networking opportunities. When the videocasts and papers are published make sure you check them out. I won't list all my favorites, but here are a few that have some great lessons learned with a safety perspective.

First, was the paper "Everyone Who Knows That is Dead: A B-1 Story." For those involved in what we at Textron Aviation call sustaining programs, this had so many relatable stories. The expertise on the system under test (in this case the B-1) had long since moved on or even died, the type of planned testing had not been conducted in a long time, and the data available to the test team was less than they are used to on newer programs. Flight test at Edwards relies heavily on TM support, but in this case this standard practice of defaulting to the TM room to keep you safe was less than ideal because of those unrecognized gaps in their normal capability. I hope to use this videocast at a future safety meeting at Textron Aviation because of all the relatable lessons.

Immediately after the B-1 story, the next paper was “Building and Flight Testing a Wing on a Formula One Race Airplane.” The speaker, JP O’Dell is just beginning his career in flight test and the entire audience was super impressed with his presentation and candor. He talked about one particular test point he was having trouble completing and he suspected it was to do with the fatiguing environmental conditions he was under, and it was impacting his own personal performance. He decided to try it one more time and nailed it! Congratulations, right? However, in post flight reflection he realized that, although he was successful, it was not the right call. He should have knocked it off and headed home. His question to the audience was do you define how many attempts a test point should be conducted before you call it off? Does that limit change if you are getting close? Again, great paper and I will be very interested to see where this speaker ends up in 10-15 years!

On Friday morning, Eviation presented their paper “Smoke, Smoke, Smoke.” How often do we hear folks say it is just a ground test? This particular ground test resulted in a loss of a valuable test asset, but because the team implemented some basic safety rigor ensuring an effective ground egress could be conducted saved lives. However, it was a close call and this presentation will give you pause next time you are planning “just a ground test.”

Finally, we should always make the time to learn from those accidents that involved the loss of life in memory of those that paid the ultimate sacrifice. Leonardo presented their paper “AW609 AC2 N609AG Fatal Accident, Evidence, Recommendations and Lessons Learned” to the symposium. The causes are certainly skew heavily to the technical aspects of tilt rotocraft, but there are good lessons for all flight test professionals in that paper so please check it out or talk to someone who was present at SETP this year and get more details. Please ensure their legacy lives on.

In other news, we just recently finalized our Board of Directors. I would like to thank all of those that volunteered. We had a huge amount of interest with insufficient openings to use all of them, which bodes well for the future of our flight test community. I look forward to the contributions that Tom Fields, John Rudzis, Shawn Kern, Paul Smith and Bob Stoney will make to the FTSC. For those that are still interested in helping with the FTSC, we have plenty of sub-committees that don’t require board membership to be involved with, including a new sub-committee focused on bringing AI and large language models to our FTSC treasure trove of data. If you are interested, feel free to reach out to me or Susan at SETP.

Finally, mark your calendars for 2025. On May 6th and 7th, we will be holding the North American FTSW in Greensboro, NC. As well, we just signed a contract with the hotel in Trieste, Italy for the European FTSW November 4th and 5th. It will be great to get our European folks together again to network and learn from each other. Hope you can make it to at least one of those events.

Fly safe and be pessimistic with your non-Gaussian distributions!

Stuart “Chia” Rogerson

Editor’s note: Congratulations to Chia. He earned the distinction of SETP Fellow, pictured here in Anaheim just a few days ago.

Flight Test Safety Committee - Calendar

Flight Test Safety Workshop 2025 Announced

When: 6-7 May 2025
The Koury Convention Center in Greensboro is the likely location for next year’s event.

European FTSW
When: 4-5 November 2025
Trieste, Italy

https://www.flighttestsafety.org/workshops

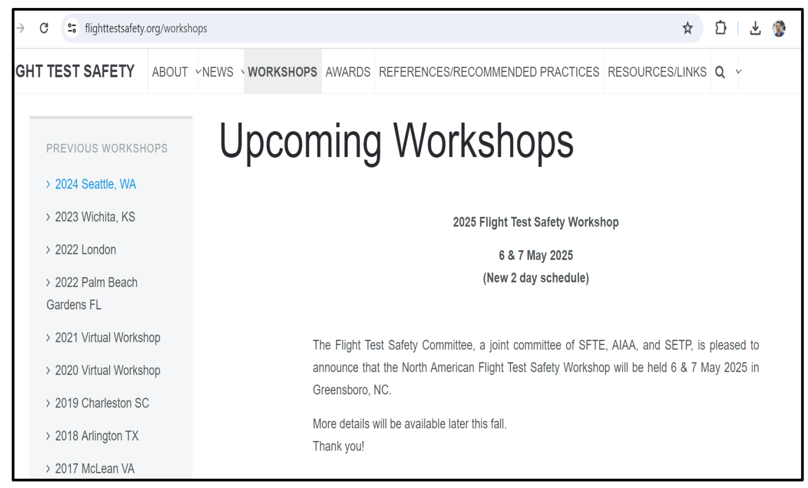

flighttestsafety.org/workshops

GHT TEST SAFETY ABOUT NEWS **WORKSHOPS** AWARDS REFERENCES/RECOMMENDED PRACTICES RESOURCES/LINKS

PREVIOUS WORKSHOPS
- 2024 Seattle, WA
- 2023 Wichita, KS
- 2022 London
- 2022 Palm Beach Gardens FL
- 2021 Virtual Workshop
- 2020 Virtual Workshop
- 2019 Charleston SC
- 2018 Arlington TX
- 2017 McLean VA

Upcoming Workshops

2025 Flight Test Safety Workshop

6 & 7 May 2025
(New 2 day schedule)

The Flight Test Safety Committee, a joint committee of SFTE, AIAA, and SETP, is pleased to announce that the North American Flight Test Safety Workshop will be held 6 & 7 May 2025 in Greensboro, NC.

More details will be available later this fall.
Thank you!

Editor’s note: The spookiest thing that happened in this edition of the FTSF was when Chia implied I was Gaussian. (I say implied, because he implicitly compared our positions when he said he “leaned more non-Guassian.”) That was a mean thing to do, but the average person won’t know I’m just joking. I have no more hard feelings than I would if Chia jumped out from behind the stage in a Star Wars costume to frighten me. The normal distribution is my least favorite thing, but I do like the binomial distribution, which strictly speaking, falls in the domain of probability, not statistics. I’ve covered many of the nuances of my position many times in these pages, and as a bonus, share some of them here.

FTSF covered three of Animal’s papers in FTSF 20-08: https://flighttestfact.com/flight-test-safety-fact-20-08/.

FTSF covered big data—a strictly non-Guassian conversation—in FTSF 20-04: https://flighttestfact.com/flight-test-safety-fact-20-04/.

Big Data is related to probability and statistics and predictions: https://flighttestfact.com/whats-the-big-deal-about-big-data/

“Counting is Hard” is another probability and stats and data science rabbit hole that first appeared in an FTSF, and three column supplement starts here: https://flighttestfact.com/counting-is-hard/.

If you aren’t frightened by math, read up on computing the probability of something that is really, really unlikely, but that happens quite often: https://flighttestfact.com/expressing-probability-qualitatively/.

Three Heuristics for Communicating Uncertainty in Flight Test: https://flighttestfact.com/three-heuristics-for-communicating-uncertainty-in-flight-test/
I don’t even mention the word Gauss, normal, distribution, or statistics, not even once.

In FTSF 23-09, I experimented with the “Carnival of Mathematics” format, and—BONUS—I also talked about probability: https://safe.menlosecurity.com/https://flighttestfact.com/flight-test-safety-fact-23-09/.

I’m ALWAYS open to long talks about probability or statistics. Send me a email mark@flighttestfact.com. Call me. Challenge me to a duel. Whatever gets us talking about numbers, I’m in.