A BAYESIAN APPROACH TO RELIABILITY GROWTH MODELING USING THE SPACEX FALCON 9 LAUNCH VEHICLE

Chris Everett[1], Homayoon Dezfuli[2]

[1] *ISL, Rockville, MD, USA*
[2] *NASA, Washington, DC, USA*

The National Aeronautics and Space Administration (NASA) Commercial Crew Program (CCP) is working with the American aerospace industry as companies develop and operate a new generation of spacecraft and launch systems capable of carrying crews to low-Earth orbit and the International Space Station (ISS). Of critical importance to NASA is the safety of the astronauts that will be using these systems for transportation to and from the ISS. Astronaut safety is quantified in terms of the probability of loss of crew, P(LOC). One of the P(LOC) requirements that CCP participants are currently working under is 1 in 1000 for ascent.

This paper presents a method for assessing the likelihood that a system provider will be able to mature their system to the point where P(LOC) is less than or equal to 1 in 1000. The system used to illustrate the method is the SpaceX Falcon 9. The method models the growth in Falcon 9 reliability using a power law model whose coefficients are uncertain. This uncertainty is modeled using a Bayesian framework, resulting in uncertainty in forecasted Falcon 9 reliability and a corresponding uncertainty in whether or not $P(LOC)_{Ascent}$ *exceeds 1 in 1000. The analysis uses Falcon 9 flight history extant at the time of this analysis, along with prior belief concerning its initial reliability and its rate of reliability growth. The analysis indicates that under the presumed reliability growth model the probability that* $P(LOC)_{Ascent}$ *will be less than or equal to 1 in 1000 within the first forty flights is low.*

I. BACKGROUND

The National Aeronautics and Space Administration (NASA) Commercial Crew Program (CCP) is working with the American aerospace industry as companies develop and operate a new generation of spacecraft and launch systems capable of carrying crews to low-Earth orbit and the International Space Station (ISS)[1]. Of critical importance to NASA is the safety of the astronauts that will be using these vehicles for transportation to and from the ISS. Astronaut safety is quantified in terms of the probability of loss of crew, P(LOC). One of the P(LOC) requirements that CCP participants are currently working under is 1 in 1000 for ascent. Table I shows the P(LOC) requirements for ascent, entry, and a full 210-day design reference mission[2].

One of the companies participating in the CCP is SpaceX, which is developing the Falcon 9 launch vehicle and the Dragon crew vehicle[3]. This paper assesses the likelihood that SpaceX will be able to mature the Falcon 9 to the point where P(LOC) is less than or equal to 1 in 1000.[*] It models the growth in Falcon 9 reliability using a power law model whose coefficients are uncertain. This uncertainty is modeled using a Bayesian framework, accounting for prior belief as well as Falcon 9 flight history to date. This results in $P(LOC)_{Ascent}$ forecasts that are a function of flight number and are themselves uncertain.

[*] This analysis does not address the question of whether or not the $P(LOC)_{Ascent}$ requirement will be *met*. Compliance with requirements is determined based on defined verification protocols, which in the case of the $P(LOC)_{Ascent}$ requirement involves calculating $P(LOC)_{Ascent}$ via probabilistic risk analysis (PRA). However, PRAs as conventionally conducted are notoriously vulnerable to incompleteness of accident scenario identification, and a P(LOC) value calculated via PRA counts as scant evidence that the *actual* P(LOC) is equal to the calculated value.

TABLE I. CCP P(LOC) requirements.

LOC Phase of Flight	HEOMD-CSD-10001 Requirement	CCT-REQ-1130 Requirement
Ascent Phase	1 in 1000	1 in 500, for combined ascent and entry phases
Entry Phase	1 in 1000	
210 Day Mission	1 in 270	1 in 200, without operational controls implemented by ISS

II. ANALYSIS APPROACH

This analysis assumes that Falcon 9 reliability is increasing over time, i.e., from launch to launch, due to ongoing design, manufacturing, and operational enhancements that are being made as part of its development. The basis of such enhancements can be analytical (e.g., as design weaknesses are discovered as a result of failure mode and effects analysis (FMEA)) or experience-based (e.g., as a result of test or operational anomalies or mishaps). In reality, Falcon 9 reliability changes in a staggered, discrete fashion as specific enhancements are made; however, the overall growth in reliability can be characterized using a relatively simple reliability growth model that characterizes the overall reliability growth trend at a top level, rather than trying to explicitly model the effect of every enhancement made.

II.A. Core Reliability Growth Model

One of the more common reliability growth models used is the Duane model[4]. J. T. Duane derived an empirical relationship based upon observation of the improvement in mean time between failures (MTBF) of a range of items used on aircraft[5]. He observed that the cumulative MTBF, θ_c, plotted against total time on log-log paper gave a straight line, where the slope of the line, β, gave an indication of reliability growth:

$$\log \theta_c = \log \theta_o + \beta (\log T - \log T_o), \tag{1}$$

or

$$\theta_c = \theta_o (T/T_o)^{\beta} \tag{2}$$

where

T_o = initial time
θ_o = cumulative MTBF at time T_o

Equation 2 can be solved for instantaneous MTBF by differentiation, yielding:

$$\theta_i = \theta_c/(1 - \beta) = [\theta_o/(1 - \beta)] (T/T_o)^{\beta} \tag{3}$$

The Duane model is an example of a power law model, and is applicable to a population with multiple failure modes that are progressively eliminated.

For this analysis, which applies to a reusable system with a constant flight profile, the mean number of flights between failures (MFBF) was considered to be substantially equivalent to MTBF, since the flight profile and duration is roughly constant from flight to flight. A power law model was therefore considered appropriate. The power law equations used for this analysis are[6]:

$$R = 1 - \alpha T^{-\beta} \text{ for reliability growth } (\beta > 0) \tag{4}$$

$$R = (1 - \alpha) T^{\beta} \text{ for reliability degradation } (\beta < 0) \tag{5}$$

where

T = flight/trial number
α = initial failure probability ($1 - \alpha$ = initial reliability)
β = reliability growth parameter

Figure 1 shows the reliability growth model for a number of ordered pairs (α, β). In all cases the reliability of the first flight is ($1 - \alpha$). If $\beta > 0$, R approaches 1 asymptotically as T goes to infinity. If $\beta < 0$, R approaches 0 asymptotically as T goes to infinity. If $\beta = 0$, R remains constant at ($1 - \alpha$) indefinitely.

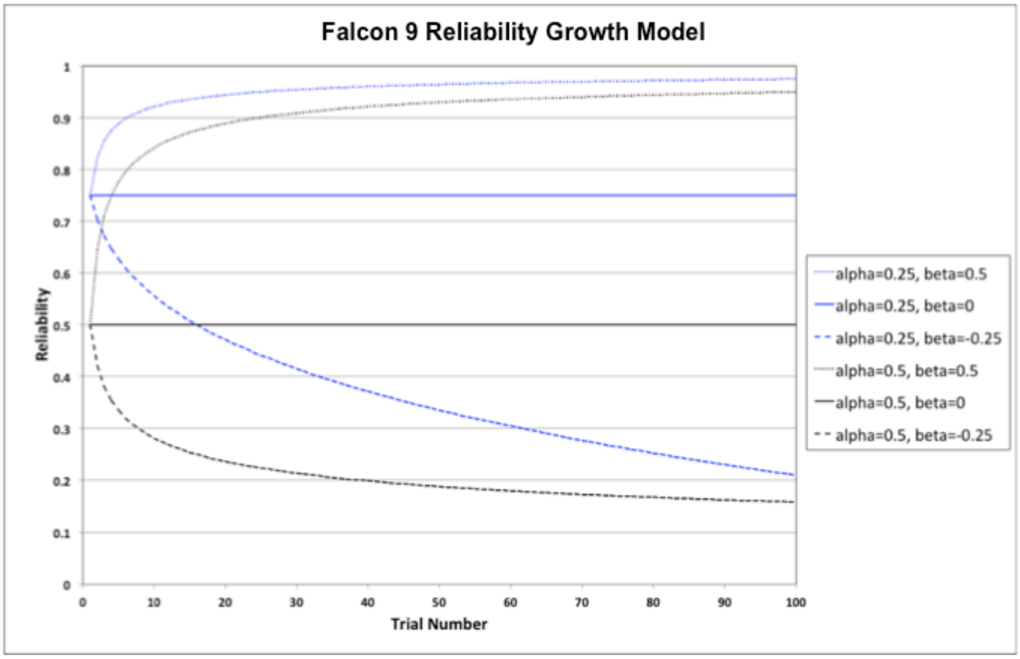

Fig. 1. Power Law Reliability Growth Model

II.B. Bayesian Uncertainty Modeling

In order to determine the values α, β of the reliability growth model, a Bayesian approach was used. Thomas Bayes developed Bayes' theorem, which is stated mathematically as:

$$P(A|B) = P(A)\, P(B|A) / P(B) \tag{6}$$

where

A and B are events
P(A) and P(B) are the probabilities of events A and B without regard to each other
P(A|B) is the probability of event A conditional on the truth of event B
P(B|A) is the probability of event B conditional on the truth of event A

As applied to this analysis, Bayes' theorem can be restated as:

$$P(\alpha, \beta \mid \text{flight experience}) \propto P(\alpha, \beta)\, P(\text{flight experience} \mid \alpha, \beta) \tag{7}$$

where the event α, β represents a particular ordered pair, (α, β), and the event, "flight experience," represents the observed flight history of the Falcon 9 as an ordered sequence of successes and failures. As such, the theorem gives the relative probabilities that particular values of (α, β) are true, given the flight experience of the Falcon 9.[†]

The factor P(α, β) in equation 7 is the *prior probability* which, when characterized over the credible range of (α, β), takes the form of a joint probability density function (pdf) over (α, β). This function is subjective in nature and represents prior belief in the relative probabilities that specific values of (α, β) are true. For this analysis, the initial reliability of the Falcon 9 was considered wholly unknown, and therefore described by a uniform pdf from a probability of zero to a probability of one, which produces a corresponding pdf over α from zero to one (α being 1 minus initial reliability). The prior belief in the reliability growth parameter β was given a triangular pdf from -0.5 to 1, with a mode at 0.45. These parameter values were chosen based on a couple of factors. First, the Engineering Statistics Handbook of the National Institute of Standards and Technology (NIST) states that "The reliability improvement slope for virtually all reliability improvement tests will be between 0.3 and 0.6. The lower end (0.3) describes a minimally effective test - perhaps the cross-functional team is inexperienced or the system has many failure mechanisms that are not well understood. The higher end (0.6) approaches the empirical state of the art for reliability improvement activities."[7]. Second, values for β were calculated for a few launch vehicles with comparatively extensive launch histories, and the results clustered around 0.45 as shown in Table II. Together, these factors suggest that the chosen prior pdf for β is reasonable and conservative, i.e., it spans the credible range of (α, β) values without representing strong belief in any particular set of values within that range, thus enabling the analytical conclusions to be responsive to the actual flight data.

Finally, the pdfs for α and β were assumed independent, producing the joint prior (α, β) pdf shown in Figure 2.

TABLE II. β values for some established launch vehicles.

Launch Vehicle	**β Value (Posterior Mode)**[‡]
Shuttle	0.41
Atlas V	0.50
Delta IV	0.50

The factor P(flight experience | α, β) in equation 7 is the *likelihood function*, which expresses the probability that a given value of (α, β) would give rise to the observed flight data. Like the prior, the likelihood also takes the form of a joint pdf over (α, β). The result of each flight is either success or failure, which suggests that the situation is that of a Bernoulli trial, in which the probability of *s* success in *n* flights is:

$$P(s,n) = \binom{n}{s} R^{s}(1-R)^{(n-s)} \qquad (8)$$

However, the reliability R varies with flight number per equations 4 and 5, so equation 8 must be modified to account for this. Since R changes with flight number (T), every sequence of successes and failures has a unique probability of occurrence, as opposed to the situation in equation 8 where only the total number of successes is relevant. The result is:

$$\mathrm{P}(\boldsymbol{h}) = \Pi_{\mathrm{Successes}}[\mathrm{R}(\alpha, \beta, \mathrm{T})]\ \Pi_{\mathrm{Falures}}[(1-\mathrm{R}(\alpha, \beta, \mathrm{T}))] \qquad (9)$$

where

h = a flight history expressed as an ordered sequence of successes and failures
$\Pi_{\mathrm{Successes}}$ = the product of the individual probabilities of each success
Π_{Failures} = the product of the individual probabilities of each failure

[†] The term "truth" is used loosely here. This analysis *presumes* the reliability growth model of equations 4 and 5. As discussed, it is a trending model that is does not explicitly address individual enhancements. The issue of *model uncertainty* is beyond the scope of this paper.
[‡] The posterior probability density function is discussed later.

In order to calculate the likelihood over the domain of possible (α, β) values (and do other calculations), an Excel tool was developed. The likelihood pdf is shown in Figure 3 for the Falcon 9 flight history of 18 successes followed by 1 failure[§].

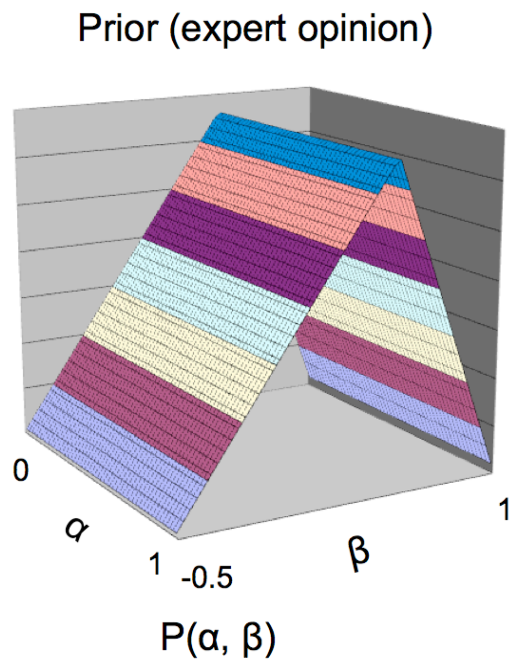

Fig. 2. Prior Probability over (α, β)

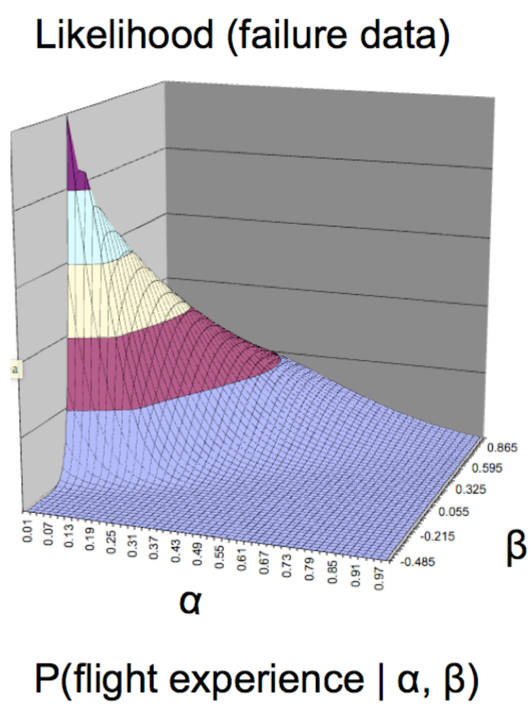

Fig. 3. Likelihood Function for (α, β) given Falcon 9 Flight History

§ All data used in this analysis is publically available.

II.C. Falcon 9 Failure Description

The lone Falcon 9 failure, which occurred on the nineteenth flight (which is also the most recent flight), is believed to have been initiated by a flawed strut inside a second stage liquid oxygen (LOX) propellant tank whose purpose was to secure a high-pressure helium bottle[9,10]. (The helium, which was pressurized to 5,500 psi, is used to pressurize the propellant tank in order to keep the tank structurally sound). With the failure of the strut, the helium system integrity was breached and the high-pressure helium was released into the propellant tank, which ruptured due to overpressurization, as shown in Figure 4.

Fig. 4. Failure of the Falcon 9 on its Nineteenth Launch

The strut, which was designed and certified to withstand 10,000 lbs of force, was made mostly of steel and was approximately two feet long and an inch thick. Figure 5 shows an example of such a strut (Figure 5 is not a SpaceX graphic). Close-out stage construction photos that were reviewed during the accident investigation showed no visible flaws or damage. However, at the time of the accident 138 seconds into ascent, telemetry data showed a momentary drop in helium pressure, then a rise back to the system's starting pressure, according to SpaceX chief Elon Musk[11]. One explanation of this is that the tank broke free, introducing helium into the LOX tank until a kink in a feed line stopped the leak, allowing the pressure in the helium system to rise again.

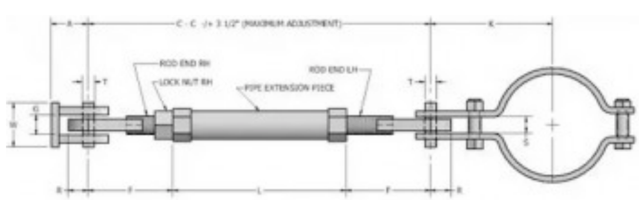

Fig. 5. Example Support Strut

A manufacturing defect is believed to have been responsible for the strut failure, which SpaceX engineers have concluded likely fractured near a bolt attach point. At the time of the accident, the strut was under a nominal 2,000 lbs of stress – five times less than its certified capacity. However, since the accident, tests of thousands of similar struts have produced several failures well below the 10,000-pound certification load, and material analysis found problems with the grain structure of the steel. "It hadn't been formed correctly, so we think that was the problem — a bad bolt that snuck through, that looked good, but wasn't actually good in the inside," said Musk.

Despite having successfully used the struts on all previous Falcon 9 flights, SpaceX has discontinued their use going forward. Additionally, SpaceX is implementing additional hardware quality audits throughout the vehicle to further assure that all parts, such as the failed strut, that are received from outside vendors will perform as expected per their certification documentation. Pressurization system support struts will be individually pull-tested before being installed on future launch vehicles. This change in the design and management of the Falcon 9 is an example of the kind of specific enhancement that cumulatively produces the reliability growth that is modeled at a high level by equations 4 and 5. In the words of SpaceX, "While the CRS-7[**] loss is regrettable, this review process invariably will, in the end, yield a safer and more reliable launch vehicle for all of our customers, including NASA, the United States Air Force, and commercial purchasers of launch services. Critically, the vehicle will be even safer as we begin to carry U.S. astronauts to the International Space Station in 2017"[9].

III. RESULTS

III.A. Falcon 9 Reliability

The term P(α, β | flight experience) in equation 7 is the *posterior probability*, which is interpreted as the degree of belief in a particular set of values (α, β) given the flight experience. The posterior pdf, presented in Figure 6, is a joint pdf over (α, β) accounting for both prior belief and Falcon 9 flight experience. Each ordered pair (α, β) on the domain represents a distinct reliability growth curve for the Falcon 9. The relative heights of the posterior pdf above different values of (α, β) represent the different degrees of belief in the truth of those values. Collectively, cross sections can be taken through this probabilistically weighted family of reliability growth curves, producing reliability forecasts in the form of pdfs that vary as a function of flight number, as shown in Figure 7.

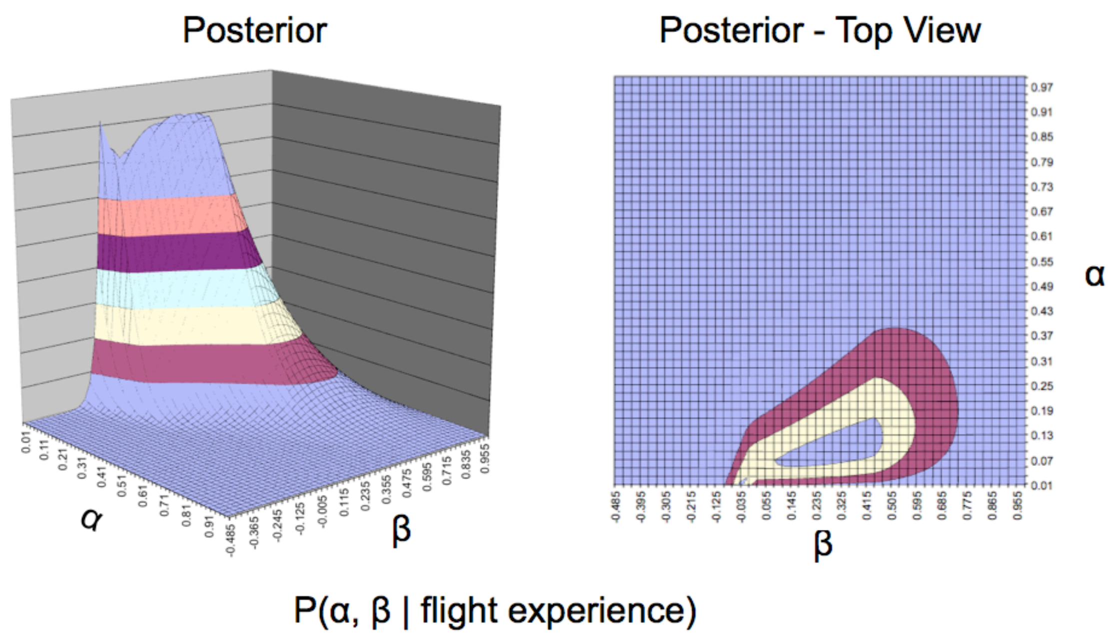

Figure 6. Posterior Probability over (α, β)

Figure 7 shows statistics from distinct reliability pdfs at each flight number, with a clear positive reliability growth trend not only in the mean but also at each of the three percentiles shown. The pdfs for flights 1 – 19 are retrospective, showing current belief about past Falcon 9 reliability. The pdf for flight 20, the statistics from which are shown in Table 3, represents current belief about the reliability of the next flight. The pdfs for flights 20 and above represent reliability forecasts for the future given the flight history to date. For example, Figure 7 indicates that Falcon 9 reliability is expected to increase to 95% by the 40th flight, with a 5% chance that it will increase to 99.4% or above.

[**] CRS-7 was the 7th SpaceX cargo resupply mission to the ISS under its Commercial Resupply Services (CRS) contract[12].

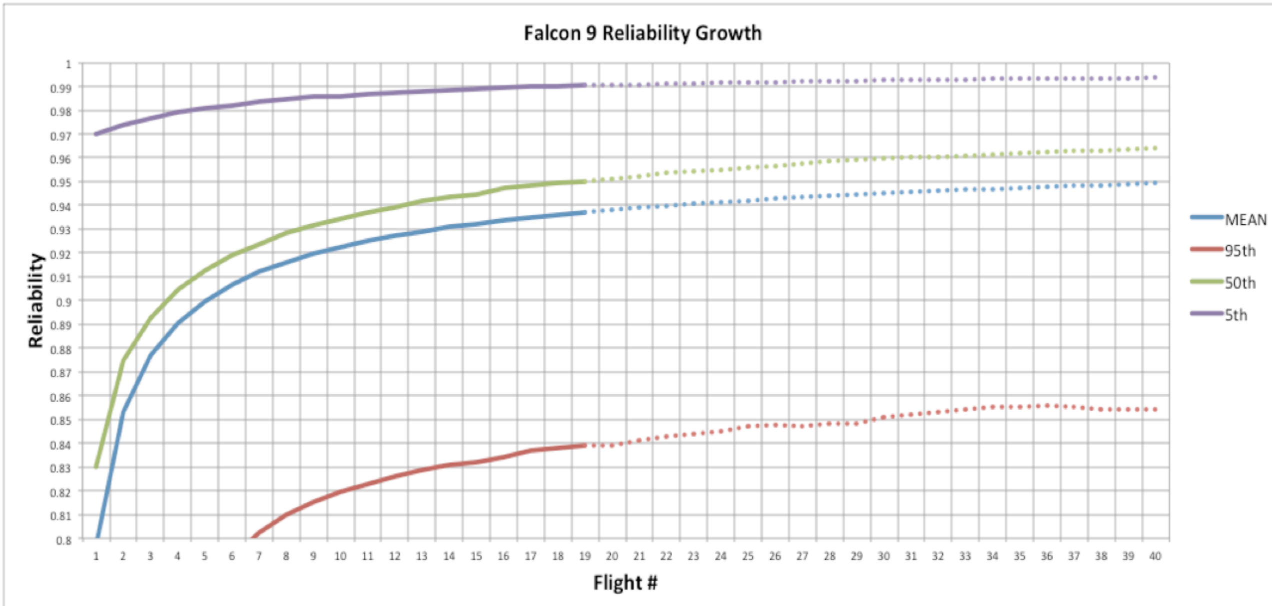

Fig. 7. Forecasted Falcon 9 Reliability

TABLE III. Falcon 9 reliability forecast for the next (20th) flight.

Falcon 9 Reliability	
Mean	93.8%
95th Percentile	83.9%
50th Percentile	95.1%
5th Percentile	99.1%

It is notable that the reliability growth trend is positive despite the fact that the one Falcon 9 failure that occurred did so on its most recent flight. Based on this evidence alone one might conclude that the vehicle's reliability growth is trending negative. In fact, the likelihood function, which is affected only by the data and not by prior belief, has its mode at $(\alpha, \beta) = (0.01, -0.035)$, such that a maximum likelihood approach to reliability growth estimation would show negative reliability growth. However, when the prior is factored in, the region around (0.01, -0.035), and the region $\beta < 0$ generally, turns out to be an outlier relative to the overall mass of the posterior, as shown in Figure 6. This can be interpreted as saying that given the still relatively sparse data set, the fact that the one Falcon 9 failure occurred late can be chalked up to random chance rather than to an actual negative trend. Although a negative reliability growth trend is possible, past experience from other systems strongly suggests that positive reliability growth is the norm, and one late failure in an otherwise consistent record of success is not enough to disabuse one of the belief that the same is true of the Falcon 9.

III.B. Probability of Loss of Crew on Ascent

The motivating question of this analysis is the desire to forecast the likelihood that a commercial provider will be able to mature its system to the point where its reliability will be high enough to support a $P(LOC)_{Ascent}$ that is less than or equal to 1 in 1000. System reliability is one of two major factors affecting $P(LOC)_{Ascent}$. The other is abort effectiveness. Loss of crew is the product of launch vehicle failure and abort system failure:

$$P(LOC)_{Ascent} = (1 - \text{system reliability})\,(1 - \text{abort effectiveness}) \qquad (10)$$

This paper does not delve into abort effectiveness. Available analyses indicate that reasonable estimates of abort effectiveness are in the 80% - 90% range[13,14]. When combined with the assessed Falcon 9 reliability, this results in $P(LOC)_{Ascent}$ values in the range of 1 in 81 to 1 in 160 for the next flight, which is considerably higher than the NASA requirement of 1 in 1000, and higher even than the total P(LOC) requirement of 1 in 200 for the full 210-day mission. Assuming an 85% effective abort capability, a Falcon 9 reliability of 99.3% is needed in order to produce a $P(LOC)_{Ascent}$ of 1 in 1000. However, 99.3% reliability is only at the 3% confidence level of the forecast for the next flight, and only at the 6% confidence level of the forecast for the 40th flight. The bottom line is that the $P(LOC)_{Ascent}$ of the SpaceX architecture (Falcon 9 and Dragon) seems highly unlikely to meet or fall below the current requirement value of 1 in 1000 within a reasonable number of flights, absent an unusually aggressive program of reliability growth or the development of a surprisingly effective abort capability.

Additionally, the question can be asked, "Under the best case scenario of no further failures (other than the one just experienced), how many flights would it take to bring the mean assessed Falcon 9 reliability up to 99.3%. To answer this question, the Excel tool was run with dummy successes from flight 20 until a next-flight mean reliability of 99.3% or greater was hit. The result was that even with continual successes from here on in, it wouldn't be until flight 129 that the mean reliability forecast would hit 99.3%. The (uncertain) reliability growth curve for this analysis is shown in Figure 8.

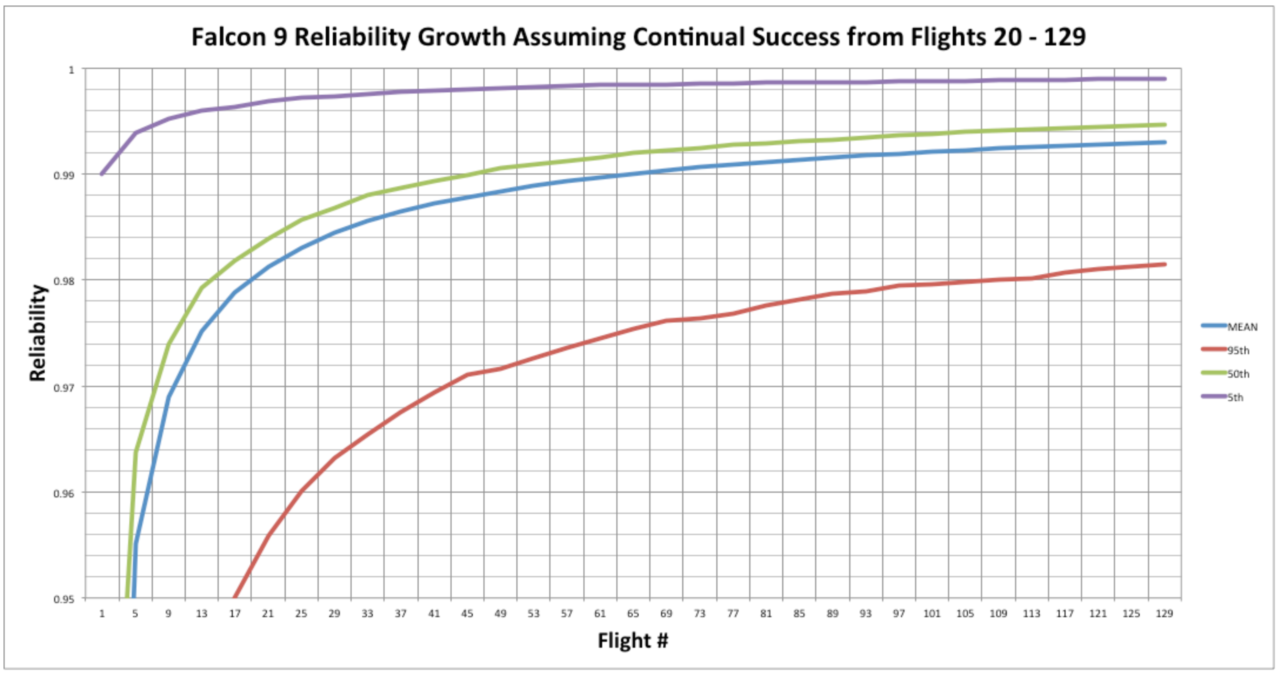

Fig. 8. Falcon 9 Reliability Growth Assuming only Successes Going Forward

It is interesting to compare Figures 7 and 8 to see how additional information, in this case good news, modifies belief. For example, the current best-estimate of next-flight reliability is 93.8%, but if the unlikely string of successes assumed in Figure 8 comes to pass, it will be seen as having been 98.1%.

III.C. Sensitivity Studies

In order to develop confidence that the results of this analysis are robust against assumptions concerning the prior probability of (α, β), several sensitivity runs were conducted. Figure 9 shows the baseline results (upper left table) as compared to several runs that vary the range of the uniform pdf over α, and the minimum, maximum, and mode of the triangular pdf over β. The results are shown in terms of probability of failure, which is 1 minus reliability. The runs show variation, but the overall conclusions of this analysis remain valid over the range of variation.

III.D. Evolving Belief vs. Reliability Growth

The analysis above used the entire Falcon 9 flight history to date. In particular, the reliability results shown in Figure 7 for flights 1 to 19 represent current belief about past reliability, given what we now know. It is interesting, however, to apply the analysis methodology to flights 1 through 19 using only the flight history that was available at that time.

P(Falcon 9 Failure)

MEAN	6.19E-02	1 in 16.17
95th Percentile	1.61E-01	1 in 6.21
50th Percentile	4.87E-02	1 in 20.53
5th Percentile	9.39E-03	1 in 106.50

Alpha = [0.9, 1.0]

MEAN	1.06E-01	1 in 9.42
95th Percentile	1.93E-01	1 in 5.19
50th Percentile	9.60E-02	1 in 10.41
5th Percentile	5.66E-02	1 in 17.67

Alpha = [0.0, 0.5]

MEAN	6.00E-02	1 in 16.66
95th Percentile	1.56E-01	1 in 6.40
50th Percentile	4.66E-02	1 in 21.46
5th Percentile	9.44E-03	1 in 105.90

Beta = [0.3, 0.45, 0.6]

MEAN	5.52E-02	1 in 18.13
95th Percentile	1.24E-01	1 in 8.03
50th Percentile	4.79E-02	1 in 20.89
5th Percentile	1.02E-02	1 in 98.15

Alpha = [0.0, 0.1]

MEAN	3.92E-02	1 in 25.48
95th Percentile	1.33E-01	1 in 7.53
50th Percentile	2.44E-02	1 in 41.07
5th Percentile	4.97E-03	1 in 201.28

Beta = [-0.01, 0.0, 0.01]

MEAN	9.57E-02	1 in 10.45
95th Percentile	2.13E-01	1 in 4.69
50th Percentile	8.60E-02	1 in 11.63
5th Percentile	1.65E-02	1 in 60.59

Fig. 9. Some Sensitivity Runs on Falcon 9 Next-Flight Reliability

Figure 10 shows what this analysis method would have forecasted for Falcon 9 next-flight reliability for flights 1 to 20 (labeled "evolving belief"), compared to current belief about what the reliability was (labeled "reliability growth"). The results are presented in terms of MFBF and displayed logarithmically for clearer comparison. The figure shows that the forecast for the first flight would have been overly pessimistic, driven exclusively by the conservative prior pdf over α (β is immaterial for the first flight). Then, as SpaceX logged a string of consecutive successes, by the 8th flight the assessed reliability would have increased beyond what the current flight history now justifies. Eventually, with the failure on the 19th flight, the overly high reliabilities that had been reached would have been tempered and the assessment would have converged on the current forecast for the 20th flight.

IV. CONCLUSIONS

The primary conclusion of this analysis is that the reliability of the Falcon 9 is unlikely to be high enough to support the level of safety that might seem implied by NASA's $P(LOC)_{Ascent}$ requirement of 1 in 1000. Assuming an 85% abort effectiveness for the integrated Falcon 9 / Dragon architecture, a Falcon 9 reliability of 99.3% is needed. The expected reliability for the next flight is 93.8%, with only a 3% chance of exceeding 99.3%. Even at the 40th flight the forecasted expected reliability is 95%, with only a 6% chance of exceeding 99.3%. Even under the assumption of continual successes going forward, the forecasted expected reliability would not reach 99.3% until the 129th flight.

The analysis indicates positive reliability growth for the Falcon 9. This is notable because the Falcon 9 flight history consists of 18 successes followed by 1 failure. Based on flight history alone one might conclude that if a reliability trend exists, it is negative, not positive – i.e., reliability is getting worse over time. The counterintuitive result of positive reliability can be traced to the prior joint pdf over (α, β), and in particular on the prior belief that β is likely to be greater than 0. This belief is rooted in the flight history of other launch vehicles, as well as on NIST institutional knowledge about reliability growth efforts generally. It is also supported, at least qualitatively, by SpaceX's response to the CRS-7 failure, which by all indications suggests that flights going forward will be more reliable than they have been up to now. The conclusion of positive reliability growth is reconciled with the flight history by attributing the timing of the one Falcon 9 failure to chance, rather than to a manifesting trend.

Bayesian analysis of reliability, specifically of the uncertain coefficients of an assumed reliability model, is a powerful tool, especially early on when flight history is limited. It provides a means of incorporating subjective belief into the analysis as appropriate, with a mathematically rigorous way of transitioning the driver of the analysis results from subjective, prior belief to objective observation as flight history accumulates.

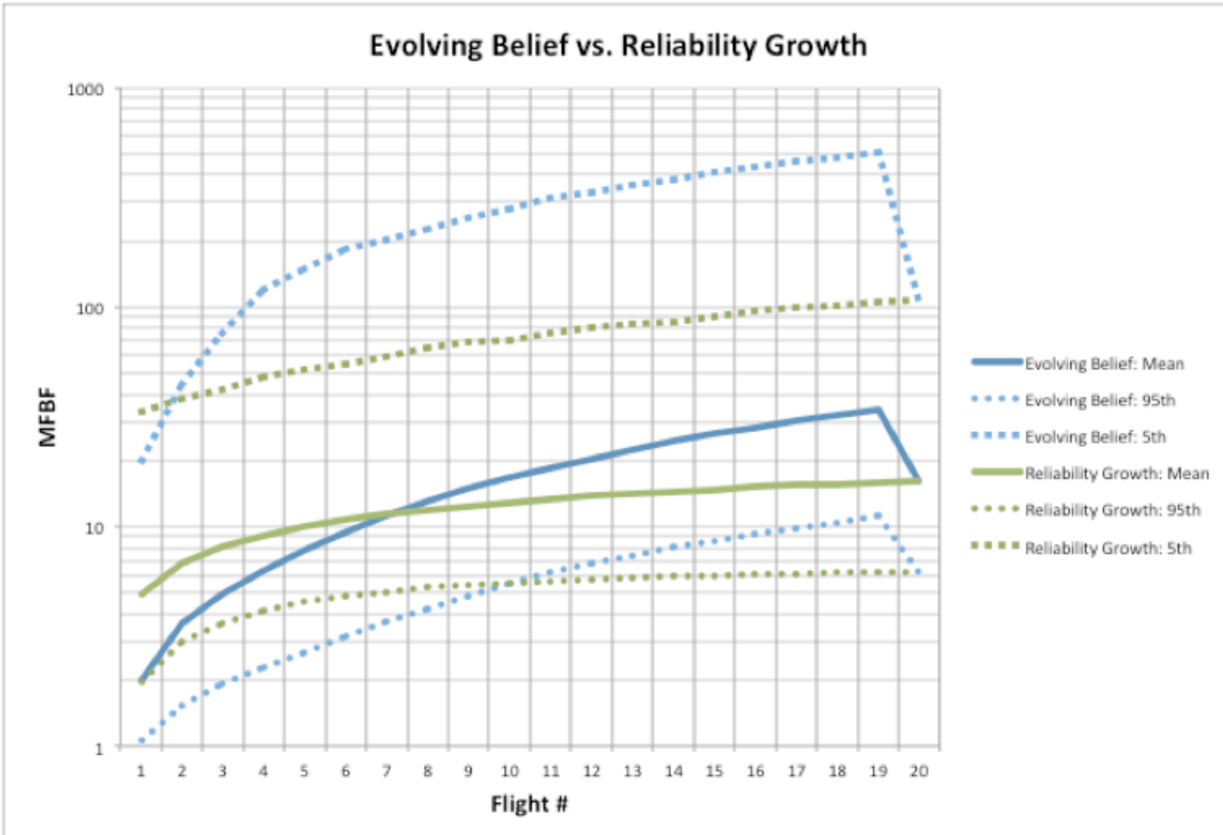

Fig. 10. Evolving Belief vs. Reliability Growth for the Falcon 9

DISCLAIMER

Nothing in this paper should be viewed as an official NASA position on the reliability of the Falcon 9 or any other system, nor should the contents of this paper be viewed as an endorsement or non-endorsement of any particular system. The use of Falcon 9 flight history as it existed at the time the analysis was conducted was strictly for the purpose of illustrating the method in the context of a developmental system with a limited flight history.

REFERENCES

1. NASA, Commercial Crew Program website, https://www.nasa.gov/content/commercial-crew-overview.
2. NASA, “Commercial Crew Transportation System Certification Requirements for NASA Low Earth Orbit Missions,” HEOMD-CSD-10001, Revision A.
3. SpaceX, Falcon 9 website, http://www.spacex.com/falcon9.
4. J. T. DUANE, "Learning Curve Approach To Reliability Monitoring," IEEE Transactions on Aerospace, Vol. 2, pp. 563-566, 1964.
5. P. D. T. O’CONNOR, and A. KLEYNER, *Practical Reliability Engineering*, Section 1.13.1, 2012.
6. C. EVERETT, et al, “Bayesian Reliability Analysis Tool (BRAT),” PowerPoint presentation, SAIC, May 2006.
7. NIST, *NIST/SEMATECH Engineering Statistics Handbook*, Section 8.1.9.2, http://www.itl.nist.gov/div898/handbook/apr/section1/apr192.htm.
8. Wikipedia, “List of Falcon 9 and Falcon Heavy launches,” https://en.wikipedia.org/wiki/List_of_Falcon_9_and_Falcon_Heavy_launches
9. SpaceX, CRS-7 Investigation Update, http://www.spacex.com/news/2015/07/20/crs-7-investigation-update, July 2015.
10. Space News, “SpaceX Wrapping Up Falcon 9 Failure Investigation,” October 9, 2015.

11. Spaceflight Now, "Support strut probable cause of Falcon 9 failure," July 20, 2015.
12. Wikipedia, "SpaceX CRS-7," https://en.wikipedia.org/wiki/SpaceX_CRS-7.
13. C. EVERETT, et al, "Dynamic Abort Risk Evaluator (DARE) Analysis Process Document," SAIC, February 2007.
14. D. MATHIAS, et al, "Simulation Assisted Risk Assessment Applied to Launch Vehicle Conceptual Design," NASA Ames Research Center, January 2008.