The Identical Error paradox in aviation maintenance - more errors for safer outcomes

Kevin Hayes & Nicolò Gariglio

Airbus UK

SUMMARY

Aviation regulators require maintenance organizations to mitigate against the potential catastrophic failure of redundant systems caused by Human Error. One of these mitigating strategies, "identical error", is analysed in this paper for its efficacy at preventing Maintenance Errors. Using Rasmussen Skill-Rule-Knowledge framework (1986) in a thought experiment based scenario, and by applying a simple Human Reliability Analysis (Swain and Guttman, 1983) to the results, a potential paradox was identified; that under certain circumstances a system can intentionally produce more errors in order to result in an overall safer outcome. This finding was then considered for its real-world application to airline operators based upon the results of interview data analysed using thematic analysis (Braun & Clarke, 2006).

KEYWORDS

Aviation Maintenance Error, Aviation Safety, Maintenance, Human Error

Note: *Statements of probability in this paper are made in reference to the "Likelihood Scale" published by the International Panel on Climate Change in their Uncertainties Guidance Notes (Mastrandrea, 2010). When terms from this scale are used in the text they are denoted by being in bold italics.*

Introduction - High reliability in commercial aircraft systems

When certifying the design of an aircraft system the regulators require that a catastrophic failure of an aircraft system must be "extremely improbable" where this probability is quantified as $1x10^{-9}$ per flying hour (FAA, 2024). To achieve this, safe design principles and techniques can be employed which include design strategies such as redundancy, backup or enhanced factors of safety (FAA, 2024). However, such quantitative assessments of reliability are currently only applied to the probability of a failure for purely technical reasons and do not account for the influence of human performance during maintenance activities. For the purposes of the calculation, it is assumed that the aircraft is maintained in accordance with the prescribed instruction (EASA, 2003) or, to put it in other words, there is an assumption that all aircraft maintenance is conducted perfectly and without error by the engineers and mechanics who undertake the tasks (RAeS, 2022).

Despite the employment of various strategies to prevent a catastrophic occurrence during the design stage, there is an acknowledgment from design engineers that certain architectural choices can still be bypassed by Common Cause Failure (CCF) modes, where two systems fail in the same way for the same reason (Jones, 2012). Human Error is recognised as a possible CCF, even in a redundant system. An example of this might be if a maintenance technician is installing two identical but independent flight control computers and makes the same installation mistake on both (such as cross connecting the input feeds) due to their understanding of the system.

Both of the two largest regulators in the world, EASA and the FAA, have regulations that are aimed at aircraft operators and (rather than the manufacturers) with the intent of minimising the effect of Human Error on safety critical systems. When the same task needs to be performed at the same time on two different systems (such as in the example given previously), EASA requires maintenance organizations to try and plan for different people to perform each of the tasks. This is to try and minimise the possibility of an error being repeated on both systems, a principle known as "identical maintenance". A typical example of this applies to maintenance of aircraft that are certified for Extended Twin-Engine Operations (ETOPS). Such certification is intended to ensure that the aircraft in question is reliable enough to fly on a single engine for an extended period in the event of the other engine failing. To achieve this level of reliability, EASA also requires Continuing Airworthiness Maintenance Organizations to specify how they will prevent "identical errors" occurring during identical maintenance tasks that are taking place concurrently on redundant systems (EASA, 2010). The FAA adopts a similar approach to ETOPS but instead uses the term "dual maintenance" which not only refers to maintenance tasks occurring on the "same" ETOPs Significant Systems at the same time, but also to "substantially similar" tasks or systems (FAA, 2008).

How human error affects redundancy

What is lacking from the regulators is an understanding of what types of errors are being mitigated against through the adoption of an "identical maintenance" approach. Without an understanding of what the organization is trying to mitigate against it is not possible to analyse the efficacy of such methods and whether another approach should be employed.

Regardless of whether the term being used is "identical maintenance" or "dual maintenance" the *implied intent* of the regulators is that they are trying to reduce the probability of an Aviation Maintenance Error (AME) from occurring on critical systems that could result in a Common Cause Failure mode and thus compromise safety. This is very much in the vein of the Person Approach to human error rather than the Systems Approach (Reason, 2000). By its very nature, this strategy of having different people doing the same task on redundant systems can only be used as a mitigation against those AMEs that have causal factors relating primarily to the individual rather than those factors that relate to other parts of the Aviation Maintenance System (AMS) seen in figure 1, such as environmental, process or design related factors (Hayes, 2024).

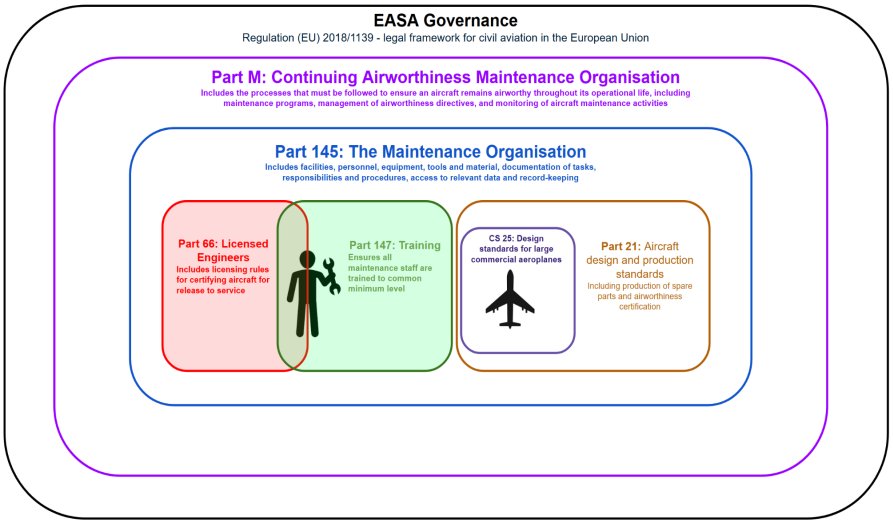

Figure 1: The Aviation Maintenance System (AMS) for an EASA approved maintenance organisation (from, EHF 2025, Maintenance Errors in Commercial Aviation: contextualising and undefined systemic problem, Hayes, et al, 2025)

Types of error addressed by these mitigations

Using Rasmussen's Skill-Rule-Knowledge framework (1986) it is possible to model which types of Human Error can be mitigated against using the "identical maintenance" principle through the application of a thought experiment involving a simple ETOPS scenario (see figure 2). In this thought experiment there is a twin engine aircraft that is having both of its engines re-installed after a period of maintenance. Because each engine acts as a redundant system for the other engine, this scenario falls into the category of "identical maintenance". The airline wishes to avoid a double engine failure in flight caused by a Maintenance Error, so the objective of the analysis in Figure 2 is to determine what types of Human Error can be mitigated against by having separate engineers installing each of the engines. This is achieved through a comparison with what is likely to happen if the same engineer installed both engines. It has been found that in aviation maintenance skill-based errors are strongly linked to Performance Shaping Factors (PSFs) that are localised by time or space, such as fatigue resulting in memory lapses or equipment deficiencies resulting in physical slips (Hobbs and Williamson, 2002). In contrast, rule-based errors and knowledge-based errors are more tightly coupled with PSFs that have a systematic influence across the cohort, regardless of when and where they are conducting the maintenance. For example, rule-based errors were found to be linked to inadequate procedures while knowledge-based errors were found to be linked to training (Hobbs and Williamson, 2002).

	Scenario 1: The same engineer does the task on both engines		**Scenario 2:** A different engineer does the task on each of the engines		**Conclusion**
	Engine 1 	**Engine 2** 	**Engine 1** 	**Engine 2** 	
Skill based error:	The **blue** engineer suffers a physical slip of a tool, or a cognitive lapse. This results in a Maintenance Error that remains undetected by the engineer.	It is ***exceptionally unlikely**** that the **blue** engineer will suffer the **exact same skill based error** that will result in exact same Maintenance Error **and** that the **error remains undetected.**	The **blue** engineer suffers a physical slip of a tool, or a cognitive lapse. This results in a Maintenance Error that remains undetected by the engineer.	It is ***exceptionally unlikely**** that the **red** engineer will suffer the **exact same skill based error** that will result in exact same Maintenance Error **and** that the **error will remain undetected.**	Scenario 1 and 2 are nearly as equally likely as each other. Therefore "identical errors" **are not** tightly coupled with skill based errors such as slips and lapses
Rule based error	The **blue** engineer uses the wrong learnt rule or written procedure for installing the engine. This results in a Maintenance Error that remains undetected because the engineer believes that they carried out the task correctly.	Since the installation procedure is provided to the **blue** engineer by the organization it is ***virtually certain**** that same Maintenance Error will occur again. This is because the engineer will believe that the procedure is correct and follow it as before.	The **blue** engineer uses the wrong learnt rule or written procedure for installing the engine. This results in a Maintenance Error that remains undetected because the engineer believes that they carried out the task correctly.	Since the **red** engineer will have undertaken the same Part 145 training as the blue engineer, and since the installation procedure provided to the **red** engineer will be the same procedure used by the **blue** engineer, it is ***Virtually Certain**** that the same Maintenance Error will occur again due to the wrong learnt rule or written procedure being used.	Scenario 1 and 2 are nearly as equally likely as each other. Therefore "identical errors" **are not** tightly coupled with rule based errors
Knowledge based error	The **blue** engineer needs to resort to using their own personal knowledge and experience to overcome an unforeseen issue with the engine installation that can not be resolved using a rule based solution. This results in a Maintenance Error that remains undetected because the engineer believes that they carried out the task correctly.	If the same unforeseen issue is present on the second engine (i.e. the issue is systemic in nature) then it is ***virtually certain**** that the **blue** engineer will use the same knowledge based solution and that the same Maintenance Error will occur again.	The **blue** engineer needs to resort to using their own personal knowledge and experience to overcome an unforeseen issue with the engine installation that can not be resolved using a rule based solution. This results in a Maintenance Error that remains undetected because the engineer believes that they carried out the task correctly.	The **red** engineer encounters the same unforeseen issue. However, the **red** engineer has a different set of experiences and knowledge to the **blue** engineer. Because of this it is ***very likely**** that the **red** engineer will develop a different solution that **does not** result in the exact same Maintenance Error occurring.	Depending on the nature of the issue encountered and the differing experiences of the engineers involved, it is possible to avoid the failure of both engines. Therefore "identical errors" **can be** coupled with knowledge based errors

* Statements of probability are made in reference to the "Likelihood Scale" published by the international Panel on Climate Change (IPCC) in their Uncertainties Guidance Notes (Mastrandrea et al, 2010)

Figure 2: Thought Experiment applying the SRK framework to an ETOPS scenario

From this analysis it can be seen that the concept of "identical error" avoidance is not tightly coupled with skill-based or rule-based error mechanisms. In other words, having a second engineer involved does not substantively change the pre-conditions that led to the Maintenance Error occurring. However, knowledge-based errors are more tightly coupled with the avoidance of "identical error". By having a second engineer change the second engine, it can fundamentally change the framing of the decision making process and result in a different outcome to the one reached by the first engineer. As such, the use of an "identical maintenance" approach can be an

effective mitigation against the total failure of redundant systems due to a knowledge-based Human Error CCF.

This finding, that different people with different knowledge will make different mistakes, will feel intuitively correct to the majority of people. However, when this observation is analysed further it results in a finding that will seem counter intuitive, that the "identical maintenance" approach simultaneously results in both safer outcomes but also in more Human Error overall.

Double engine failure - simplified possibility space

To illustrate this counter intuitive outcome, we need to take the thought experiment one step further and apply a very simple Human Reliability Analysis (Swain and Guttman, 1983) to only the hypothetical knowledge-based errors. The skill-based and rule-based errors can be excluded from this analysis, as it has already been demonstrated that there is no difference between the two scenarios in terms of the approximate likelihood of Human Error in an "identical maintenance" scenario.

In the previously described scenarios, success is defined as the prevention of a CCF on both engines. It does not matter if a single engine fails, so long as the second engine remains operational. If an arbitrary value for the Human Error Probability (HEP) of a knowledge-based error occurring is applied, it is possible to see mathematically how the use of a second engineer to change the second engine can be beneficial. For a HEP value of 0.1 the following occurs:

- **Scenario 1.** There is a probability of 0.1 that the blue engineer will make a knowledge-based error when installing engine 1. If the same set of knowledge is required to install engine 2 then it is ***virtually certain*** that the blue engineer will make the same knowledge-based error again as they have already been through the decision making process and they believe, incorrectly, that they have a viable solution. Therefore, for the installation of the second engine the HEP value increases to 1.0 (the error will happen). The probability of a double engine failure in scenario 1 is therefore (0.1)*(1.0) = 0.1.
- **Scenario 2.** As before, there is a probability of 0.1 that the blue engineer will make a knowledge-based error when installing engine 1. This time the installation of engine 2 is being performed by the red engineer who has a different set of knowledge and experience. As such, if the red engineer encounters the same knowledge based problem as the blue engineer, then the red engineer also has a probability of 0.1 of making the same error. Therefore, the probability of a double engine failure in scenario 2 is (0.1)*(0.1)=0.01.

In this overly simplified model, it can be seen that the safety margin has increased tenfold by the adoption of an "identical maintenance" principle. For real world scenarios the maintenance organization would need to have accurate estimations of the HEP values but it is clear that the safety benefits can be significant since the probability of a double engine failure changes in relation to the size of the HEP factorial.

However, while the use of "identical maintenance" in this scenario has significantly reduced the chances of a complete failure of the redundant systems, it is not the entire story. While both regulators and airlines both wish to avoid the catastrophic failure of a safety critical redundant system, airlines also have other success criteria such as the operation of an efficient and profitable service.

Total success - simplified possibility space

For an airline operator, total success in any maintenance task would be when there have been zero Maintenance Errors committed. For an airline, Maintenance Errors come at a significant cost, the re-working of the incorrect task is a straight forward doubling of the maintenance cost for that task.

Depending upon when the Maintenance Error is detected, there can be significant costs from flight delays, in flight turn backs and the associated costs such as placing passengers in hotels or giving them refunds, with the annual cost to the industry being estimated to be some US$616M (Allianz, 2024).

Keeping with the arbitrary value 0.1 for the Human Error Probability (HEP) of a knowledge based error occurring, it is possible to create a simplified model that shows the consequences of using the "identical maintenance" principle on completing maintenance without inducing any Maintenance Errors. Instead of considering the thought experiment scenarios in terms of the probability of Human Error, it is necessary to consider these scenarios in terms of the probability of achieving success. If there is a probability of 0.1 of an engineer making a knowledge-based error when installing an engine, then there is a 0.9 probability that the same engineer will complete the task free of a knowledge-based error. This figure could be considered the Human Success Probability (HSP).

- **Scenario 1.** There is a HSP value of 0.9 that the blue engineer will successfully install engine 1 without any knowledge-based errors occurring. If the blue engineer successfully completes the installation of engine 1 free of any knowledge-based errors, it is ***virtually certain*** that the blue engineer will also successfully install engine 2 free of any knowledge-based errors so long as the same set of knowledge is required to install both engines. Therefore, for the installation of the second engine the HSP value increases to 1.0 (the task will be free of any knowledge-based error). The probability of total success in scenario 1 is therefore (0.9)*(1.0) =0.9
- **Scenario 2.** As before, there is a HSP value of 0.9 that the blue engineer will successfully install engine 1 without any knowledge-based errors occurring. This time the installation of engine 2 is being performed by the red engineer who has a different set of knowledge and experience. As such, the HSP for the red engineer remains at 0.9. Therefore, the probability of total success in scenario 2 is (0.9)*(0.9)=0.81.

It can therefore be seen that the probability of completing the task error free has decreased directly as a result of the organisation adopting the "identical maintenance" principle as a mitigation strategy that helps ensure the overall safety of the aircraft. This is the "identical error paradox", that under certain circumstances a system can intentionally produce more errors in order to result in an overall safer outcome.

Limitations of the thought experiment

Of course this is just a thought experiment, empirical data is required to prove if this is a veridical paradox (a counterintuitive truth) or a falsidical paradox (appears true but based on a flawed logical error). Additionally, the thought experiment is a vast oversimplification of a complicated and safety critical maintenance task, one that is normally conducted by a team of aeronautical engineers. The Skill-Rule-Knowledge framework is often more a spectrum rather than ineffably distinct modes of cognitive operation (Proctor, 2008) and as such the team of engineers will be constantly switching between skill-based, rule-based and knowledge-based behaviours depending on the immediate context and requirements of the activity. This places additional caveats on how well the thought experiment represents the reality of aviation maintenance. Finally, the HEP figures presented in the thought experiment were chosen for the sake of simplicity and are unlikely to be representative of actual HEP figures in real aviation maintenance tasks. As such the practical effect of the "identical error paradox" may be very small, assuming that the hypothesis proves valid at all.

The challenge of calculating HEP for aviation maintenance tasks

In order to validate the hypothesis of the "identical error paradox", there is a requirement to calculate the real-world HEP value for a maintenance task that is sometimes conducted using the

"identical maintenance" safety strategy and sometimes without this strategy. The two HEP values can then be compared to determine if the error rate changes. There are, however, practical issues with trying to determine HEP for aviation maintenance tasks and they are not routinely calculated by maintenance organisations in commercial aviation. To calculate the error rate two pieces of data are required, the number of errors that have taken place and the number of opportunities that were available for the errors to occur (how often the task was conducted). These two data sets are normally stored on two separate IT systems that are unconnected and which have been designed for different purposes. How often a task has been conducted can be found on the organizations MIS (maintenance information system) which is used as a work recording tool to ensure compliance with the organizations continuing airworthiness responsibilities. These work recording systems do not normally record if unscheduled maintenance is being conducted as a result of a technical failure or as a result of human error that needs re-work. Meanwhile, instances of Maintenance Errors will be recorded by the organizations Safety Management System (SMS) using some sort of reporting tool (normally electronic). These reporting tools do not normally allow the safety team to dashboard and trend Maintenance Errors in a manner that shows what types of error are routinely occurring in a given maintenance task. Additionally, maintenance data is not widely shared, due to commercial sensitivity reasons and safety data is shared on a limited basis, with normally only the most severe occurrences being shared between organisations. This currently makes a global assessment of error rates impractical for regulatory bodies and equipment manufacturers alike. Airline operators are marginally better placed to calculate error rates for their own operation but are hampered by the tools available to them to do so and the cost of building their own bespoke tools.

Searching for the "identical error paradox" effect

As part of this study an exploratory attempt was made at trying to determine if a difference could be detected between the HEP rate for a single engine change versus the HEP rate for a double engine change. Hypothetically, these two figures should be identical if "identical maintenance" has no effect on the error rate. Should "identical maintenance" have an effect on the error rate then we should expect to see a higher value for the HEP on a double engine change.

Given the limitations discussed in the previous section, this analysis only looked at maintenance records on a very large MIS from which the number of engine changes could be extracted. The MIS data set covered a period from 01 Jan 2010 to 01 Jan 2026 and had records of over 10,000 engine changes for different types of commercial aircraft operating different engine types. Through this dataset it was not possible to determine which double engine changes were conducted under an "identical maintenance" regime and so there is an assumption that the second engine change is being conducted by a separate team of people. Also, the study did not have access to a comparably large dataset of safety records for the same set of aircraft that could be used to calculate a true HEP. Instead, a proxy error rate was calculated by looking at how many unscheduled maintenance events occurred on the engine within 150 flight cycles of the engine being changed. Unscheduled maintenance are those tasks performed in response to unforeseen defects or malfunctions (EASA, 2024). These defects occur for ultimately for one of two reasons 1) a technical failure of a system that originates from the reliability of the system and how it is used, such as a crack appearing in landing gear from usage 2) a Maintenance Error induced into the system through human activity, such as a component failing because it was incorrectly assembled. While it was not possible to tell from the MIS data whether any of the unscheduled maintenance was conducted due to a technical failure or a Maintenance Error, it is reasonable to assume that on average the technical reliability of a like-for-like engine will be the same and therefore any differences between the single engine and double engine datasets can be attributable to a difference in the error rates between the two activities. A periodicity of 150 flight cycles was chosen as it is approximately analogous to a 30 day operating period for a typical commercial Single Aisle aircraft and approximately 30% of the period

until the next A check (where possible A checks will be conducted at the same time as an engine change to ensure a longer continuous operating period). The MIS showed a spike of unscheduled maintenance in the first 5 cycles after maintenance and another around the 30 cycle point. Deeper analysis would be required to explain these spikes but it is likely, at least in part, to be a function of how and when the symptoms of the faults become detectable by the organization. A summary of the data used to calculate the proxy HEP can be seen in Table1.

Table 1: A comparison of proxy HEP rates for single and double engine changes

Parameter	Single engine change	Double engine change
Number of unscheduled maintenance events within 150 flight cycles the engine change [Proxy for the number of Errors (*e*)]	2,492	1,023
Number of engine changes [Opportunities for error (*n*)]	7,237	2,849
Proxy Error Rate (*p^*)	0.344	0.359

Analysis of the results

It was found that the double engine scenario had a proxy error rate 1.47% higher than that of the single engine change scenario. While this shows some evidence in support of the "identical error paradox" the difference between the two figures is not statistically significant with a p-value of 0.0811 and so the null hypothesis cannot be rejected. Given the extreme caveats with the dataset it is not surprising that the results are inconclusive but given the safety critical nature of aviation maintenance the potential effect should not be completely dismissed at this stage. Future work should look to create a valid data set that does not use proxy information and which can link MIS data to safety report data to generate a true HEP. Such a study should also be conducted on a safety critical task that is less complicated than an engine change and which has fewer confounding variables (such as the number of engineers required for the task). Finally, it should be confirmed that an "identical maintenance" regime was being enforced when two tasks were being conducted simultaneously on the same aircraft.

Practical application of the findings

The important takeaway from this paper is less around whether the "identical error paradox" is a real effect or not, and more around how the aviation industry thinks about Maintenance Error, why these errors happen and what mitigations should be put in place to reduce the chances of such an error happening again. Regulators have mandated certain strategies to mitigate against human error in maintenance but there is little research on the efficacy of these techniques and so many questions remain unanswered such as; Are we (as an industry) adequately addressing all of the factors that lead to Maintenance Errors? Are our mitigations targeting the correct PSF's? Are we over emphasising or under emphasising the right strategies? A strategy like "identical maintenance" comes at a cost to maintenance organizations, at best it results in a simple doubling of the manpower cost. At worst it could be increasing the error rate and the amount of rework required. Organisations need to understand all of these issues if the industry is to make informed decisions about continuing to keep people safe, whilst delivering effective operations at an acceptable cost.

To better understand the issues faced by airlines when conducting analysis of in-service events that involve Maintenance Error, Airbus engaged with stakeholders and safety investigation professionals from six different major airlines that had a combined fleet of over 800 passenger planes. As part of this research 24 semi-structured interviews were conducted by videoconference and lasted between 1.30hr and 2hrs each with the results being analysed using Thematic Analysis (Braun & Clarke, 2006). One of the major findings was that airlines often struggled to mitigate against Maintenance Error reoccurring. The reasons for this often involved the misidentification of what the causal factors were that led to the Maintenance Error occurring and/or the inability to identify an appropriate mitigation that would both systemically address the primary causal factors and that was also within the ability of the organization to implement. It was found that airlines get frustrated by similar types of Maintenance Errors appearing again and again in their data analysis, despite the implementation of mitigating measures. Because concepts like “dual maintenance” and “identical maintenance” were seen as mitigations there was often an expectation by airlines that tasks undertaken using such concepts should always be “error free”, so when error continued to happen during these tasks, it was often perceived as a failing of the individual to adhere to the mitigations that the organization had already implemented. This shows a disconnect between the types of error being observed, the selection of an appropriate mitigation strategy and the expectations of the organization. The chasing of “zero Maintenance Errors” is very much in the vein of Safety-I thinking (Hollnagel, 2014) but when an error occurs during an “identical maintenance” scenario then it is the perfect natural experiment (Petticrew et al, 2005) for understanding how to apply a Safety-II philosophy in the aviation maintenance domain. Such a situation represents a direct like-for-like comparison in as near to identical conditions as possible. If an investigator can identify why two individuals with access to the same tooling, infrastructure and procedures can have two different outcomes then, it will allow the organization to identify why things normally go right, as well as why errors sometimes occur. Additionally, such an analysis can show where the error sits on the Skill-Rule-Knowledge spectrum and whether the failure mode is tightly coupled with “identical maintenance”. This can enable the most appropriate means of mitigation to be selected for the error type (Hobbs and Williamson, 2002) and reduces the chance of a similar error happening again.

Through having a deeper understanding of the reality of aircraft maintenance and why Maintenance Errors keep happening then it can help organizations to foster a better Just Culture in the organization. Understanding that the correct mitigation is needed to address a given error type can change how leadership teams view Maintenance Error and the role of the individual. Should the "identical error paradox” ever prove to be an observable phenomenon then it will have implications for how safety KPIs are measured, an organization cannot chase zero Maintenance Errors if part of the safety strategy actively generates more errors.

Conclusion

In theory, success in aviation maintenance should not be contingent upon the unique knowledge and experience of the individual performing the task. By regulatory intent, any trained and qualified maintainer should have all the resources and information required for them to successfully complete the task or know when the task must be stopped. Practically speaking, the industry knows this to not be the case, which is why regulation mandates mitigation strategies such as “dual maintenance” and “identical maintenance”. When a knowledge-based error provoking situation arises, it is essentially symptomatic of issues within the Aviation Maintenance System (Hayes, 2024) as it shows that there has been a failure with one of the systemic barriers to an error occurring. This is exactly the type of error which the regulated space is supposed to mitigate against. The thought experiment presented in this paper not only shows how mitigations need to be aligned with appropriate error modes but also how such an alignment can better support analysis through a Safety II approach and also influence and improve the Just Culture within airline operators and maintenance organizations.

The "identical error paradox" remains an unproven hypothesis, but one that is useful as a lens for questioning the efficacy of mitigation strategies and whether they are truly addressing the mechanisms that are causing the error. Such a questioning philosophy is important to regulators, aircraft operators and Maintenance Repair Organizations alike. It can shape future regulation, foster a Just Culture and reduce costs all while maintaining and even improving safety in aviation maintenance.

References

Allianz. (2024). Aviation Risk Report 2024. Available for download at: https://commercial.allianz.com/news-and-insights/news/aviation-trends-2024.html

Braun, V., & Clarke, V. (2006). Using thematic analysis in psychology. Qualitative Research in Psychology, 3(2), 77–101. Available at: https://doi.org/10.1191/1478088706qp063oa

European Union (EU). (2014). Regulation (EU) No 376/2014 Occurrence Reporting. Available at: https://eur-lex.europa.eu/legal-content/EN/TXT/?qid=1576245532595&uri=CELEX%3A32014R0376

European Union Aviation Safety Agency (EASA). (2003). AMC and GM to Part 21, Gm21.A.3b(b). Available at: https://www.easa.europa.eu/et/downloads/1860/en

European Union Aviation Safety Agency (EASA). (2010). AMC 20-6 Annex II Extended Range Operation with Two-Engine Aeroplanes ETOPS Certification and Operation. Available at: https://www.easa.europa.eu/en/downloads/1696/en

European Union Aviation Safety Agency (EASA). (2024). Proposed CM-ICA-002 SORA OSO#03 Medium robustness Airworthiness requirements. Available at: https://www.easa.europa.eu/en/downloads/139680/en

European Union Aviation Safety Agency (EASA). (2025). Easy Access Rules for Continuing Airworthiness (Regulation (EU) No 1321/2014), AMC1 145.A.48(c)(3) performance of maintenance. Available at: https://www.easa.europa.eu/en/downloads/95788/en

Federal Aviation Administration (FAA). (2008). AC 120-42B Extended Operations (ETOPS and Polar Operations). Available at: https://www.faa.gov/documentLibrary/media/Advisory_Circular/120-42B.pdf

Federal Aviation Administration (FAA). (2024). AC 25.1309-1B System Design and Analysis. Available at: https://www.faa.gov/documentLibrary/media/Advisory_Circular/AC_25.1309-1B.pdf

Hayes, K., Marquee, L., & Dann, L. (2024). Maintenance errors in commercial aviation: contextualising an undefined systemic problem. Contemporary Ergonomics and Human Factors 2025. Available at: https://publications.ergonomics.org.uk/uploads/Maintenance-errors-in-commercial-aviation-contextualising-an-undefined-systemic-problem.pdf

Hobbs, A., Williamson, A. (2002). "Skills, Rules and Knowledge in aircraft maintenance: errors in context", Ergonomics 45 (4), 2002, pp.290-308. Available at: https://doi.org/10.1080/00140130110116100

Hollnagel, E. (2014). Safety-I and Safety-II. The Past and Future of Safety Management. CRC Press. Boca Raton. FL.

Jones, H. (2012). Common Cause Failures and Ultra Reliability. 42nd International Conference on Environmental Systems, 2012, ICES. Available at: https://doi.org/10.2514/6.2012-3602

Kinnison, H. A., & Siddiqui, T. (2013). Aviation maintenance management (2nd ed). McGraw-Hill. https://accessengineeringlibrary.com/browse/aviation-maintenance-management-second-edition

Mastrandrea, M., Filed, C., Stocker, T., Edenhofer, O., Ebi, K., Frame, D., Held, H., Kriegler, E., Mach, K., Matschoss, P., Plattner, G., Yohe, G., Zwiers, F. (2010). Guidance Note for Lead Authors of the IPCC Fifth Assessment Report on Consistent Treatment of Uncertainties. Jasper Ridge, CA: WMO. Available at: https://www.ipcc.ch/site/assets/uploads/2017/08/AR5_Uncertainty_Guidance_Note.pdf

Petticrew, M., Cummins, S., Ferrell, C., Findlay, A., Higgins, C., Hoy, C., & Kearns, A. (2005). Natural experiments: an unused tool for public health?. Public Health, Volume 119, Issue 9, pages 751-757. Available at: https://doi.org/10.1016/j.puhe.2004.11.008.

Proctor, T., Van Zandt, T. (2008) Human Factors in Simple and Complex Systems (2nd ed). CRC Press. Boca Raton. FL.

Rasmussen, J. (1986). The role of hierarchical knowledge representation in decision making and system management. IEEE Transactions on Systems, Man and Cybernetics, SMC-13, 257-266. Available at: 10.1109/TSMC.1985.6313353

Reason, J. (2000). Human Error: Models and management. British Medical Journal 18;320(7237): 768-70. Available at: https://doi.org/10.1136/bmj.320.7237.768

Royal Aeronautical Society (RAeS) Human Factors: Engineering Sub-Group Project. (2022). Development of a strategy to enhance human-centred design for maintenance. Available at: https://www.aerosociety.com/media/19526/human-factors-engineering-sub-group-report-oct-2022.pdf

Swain, A.D., & Guttman, H.E. (1983). Handbook of Human Reliability Analysis with Emphasis on Nuclear Power Plant Applications. Albuquerque, NM: Sandia National Laboratories. Available at: https://www.nrc.gov/docs/ML0712/ML071210299.pdf