LEARNING FROM ALL OPERATIONS

Case Study: Take-Off Rotation

Flight Safety Foundation

AUGUST 2022

1. Case Short Summary

This case study describes a safety issue regarding a specific aircraft take-off rotation. The required take-off rotation rate could not be achieved with any deflection of the controls by the flight crew — not even with inputs in accordance with the operating manuals. Because there is no rotation rate indicator in the cockpit, flight crews could not have known that the required rotation rate has never been achieved. Because this aircraft has always under-performed in this regard, flight crews' sense of normal rotation was calibrated accordingly. Hence, although the performance calculations for take-off and the flight crew actions were in accordance with procedures, safety margins were adversely impacted (increased take-off distance and a decreased obstacle clearance).

This safety issue was rooted in the aircraft certification process. It was systematic and present all the time. Because of other existing margins, at the time this case study took place, it had not manifested in an accident or serious incident. However, in situations of high airport elevation, high temperature and short runway length, safety margins were very small and close to non-existent. Take-off performance calculations should ensure safe take-off even with engine failure at V1 (defined by most civil aviation authorities as the maximum speed at which action must be taken to reject a take-off). Had an engine failure occurred during operations in these circumstances, an accident could have resulted. Even without engine failure, aircraft weight, air temperature and airport elevation at their extremes could have combined to produce a critical event. Luckily, this combination is rare.

When flight crews at the sharp end cannot mitigate a safety concern, people at the blunt end of the system must rectify the situation. In this case study, resilience was provided with the help of blunt end system analysis of all operations.

Learning From All Operations took place by looking at the distributions of key parameters. Looking at the entire range of these distributions was the only reliable way identified to understand the full relationship among the different factors affecting the safety margins. This analysis helped to characterise the latent safety risk, to identify how to affect the factors that would provide the needed margins, and to select a mitigation strategy.

The systemic solution was a change to the procedures, including an increase in the required take-off distance. In situations where the factors affecting the margins are critical, such as at high elevation airports, these changes may require a reduced payload.

2. Learning From All Operations Conceptual Framework

Before proceeding to the details of the case study, it is useful to outline the elements from the Flight Safety Foundation (FSF) Learning From All Operations conceptual framework that are relevant to this study. The whole framework is described in detail in the series of seven [Concept Notes] that can be found on the Foundation's website.

The Learning From All Operations framework is centred around the idea of using resilience capabilities to manage system pressures and to manage the resultant adaptive process. The main system pressure in this case study was an ever-present under-performing aircraft issue. This pressure was normalised by the flight crew. When combined with other pressures and environmental conditions, this pressure could have resulted in an accident, and the fact that it did not was attributed to luck.

The pressures in the FSF framework include threats, hazards and other performance-shaping factors, but are extended to include all types of demand pressures and efficiency pressures as described in Concept Note 5.

System adaptation is the reaction of a system to balancing pressures and resilience. Pressures and resilience are system performance-shaping factors —they shape the likelihood of desired and undesired outcomes. Figure 1 illustrates how resilience counteracts the pressures.

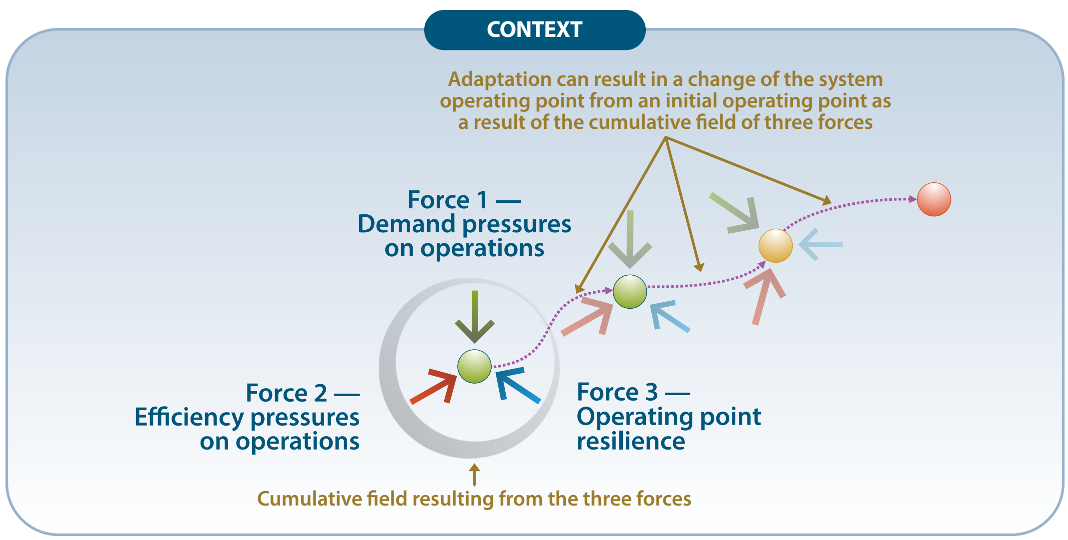

Figure 1: Three forces model of system adaptation

Learning From All Operations aims to understand how the system responds to pressures and whether, in this process, the system migrates to states of higher risk.

The performance space in which the system operates is defined by two boundaries. The first boundary is the **safety control envelope** (Figure 2, p. 3). The safety control envelope defines the actual boundaries of what is safely recoverable in operations by preventive or recovery measures (outside this boundary, safety becomes marginal to non-existent).

The safety control envelope is determined by the available capabilities to control flight safety, to enforce safety constraints, and to control the transition of the system operating point.

When the system migrates to states with higher risk, there is a tipping point at which the system becomes unstable in terms of safety (shown in Figure 2 as the boundary of the safety prevention envelope). In the middle of the envelope is the white area — the prevention space. Within this space, the system is adapting, coming closer or moving away from the critical thresholds.

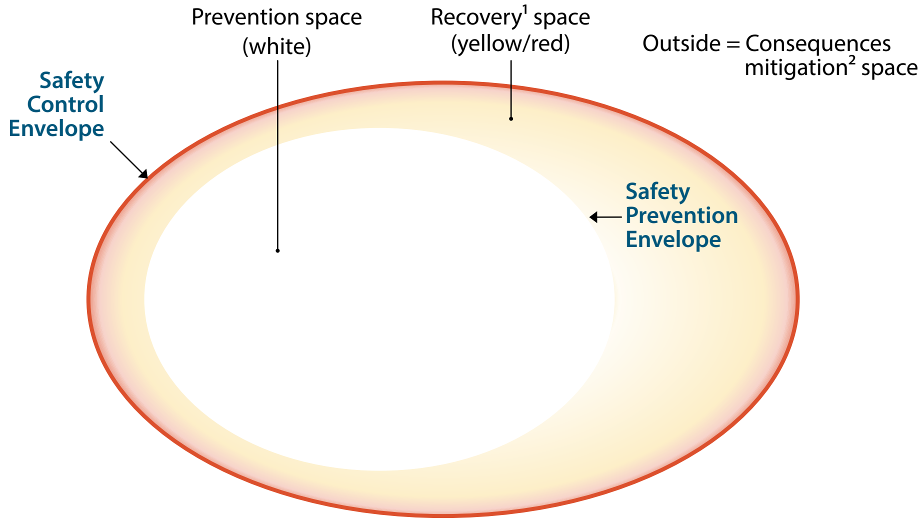

Notes

1. Recovery is possible also outside the safety control envelopve system passing beyond the safety control envelope and mitigating accident consequences (e.g., engineered materials arresting system, survivability, evacuation).
2. Consequences mitigation — mitigating consequences of the system passing beyond the safety control envelope and mitigating accident consequences (e.g., engineered materials arresting system, survivability, evacuation).

Figure 2: Safety control envelope, safety prevention envelope, prevention space, recovery space and consequences mitigation space

The yellow (system critical states space) and red (system hazardous states space) areas represent the recovery space. The system state transition through yellow and red spaces indicates decreasing distance to the edge of the safety control envelope. If a recovery action does not bring the system back into the prevention space, the operating point of the system may pass through the safety control envelope boundary where consequence mitigation may become necessary.

The space outside the red area illustrates a situation outside the safety control envelope. Passing through the safety control envelope does not always mean an accident is certain to occur. There may be some mitigation measures to reduce that likelihood, as well as luck. But passing through the safety control envelope is associated with a significant loss of flight safety control, with only marginal possibility, if any, to recover.

The second boundary that defines the system performance space is described by the **operational limits assumptions** — the imagined boundary for operations (normative — rules, procedures, prescriptions or the subjective assumptions about where this boundary is) — illustrated by the purple arc in Figure 3 (p. 4). Limits and the assumptions by the different actors about the limits can vary, and sometimes the limits are not fully defined (missing limits).

Ideally, the operational limits will neatly protect system operations from breaching the safety control envelope. **But in reality, limits of control and limits as defined have a more complex relationship.** These and other reasons for misalignment of the operational limits and safety envelope result in some generic patterns of their relationship — as illustrated in Figure 3.

The result of the interaction of the pressures and the operational resilience can be observed in operations and understood through the lens of five patterns of operational resilience (Figure 4, p. 4):

- **Remaining within the prevention space —** Prevent, avoid or withstand pressures to stay within the safety prevention envelope. This pattern of operational resilience includes, but is not restricted to, system robustness. Remaining within the prevention space can be supported by modifying and adhering to operational limits that are more conservative.

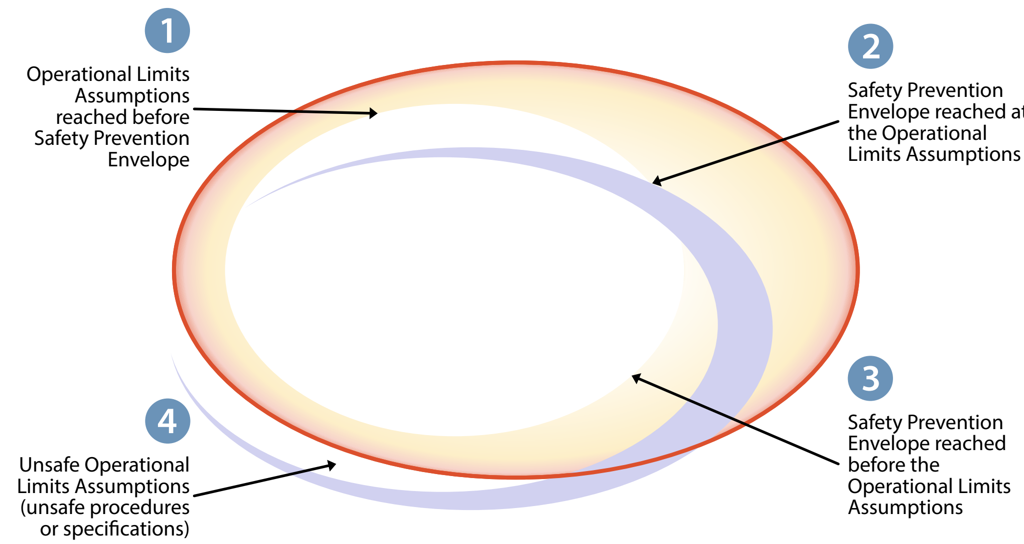

Figure 3: Operational limits and safety envelope misalignment

- **Recovering from a critical state** — This is illustrated as the system operating point transitioning through the safety prevention envelope to and back from a critical state. Here, for some amount of time, the system becomes unstable in terms of safety control, and then, the system recovers back to the prevention space.
- **Recovering from a hazardous state** — This is illustrated as the system operating point transitioning through the safety prevention envelope via critical states and to hazardous states but recovering back to the prevention space.
- **Rebounding back into the safety control envelope** — This is illustrated as the system operating point passes beyond the safety control envelope and returns back to the safety control envelope in a controlled safe manner.
- **Envelope expansion** — This is illustrated below as the safety control and/or prevention envelopes are expanded based on applying critical thinking regarding operational limits within that operational context.

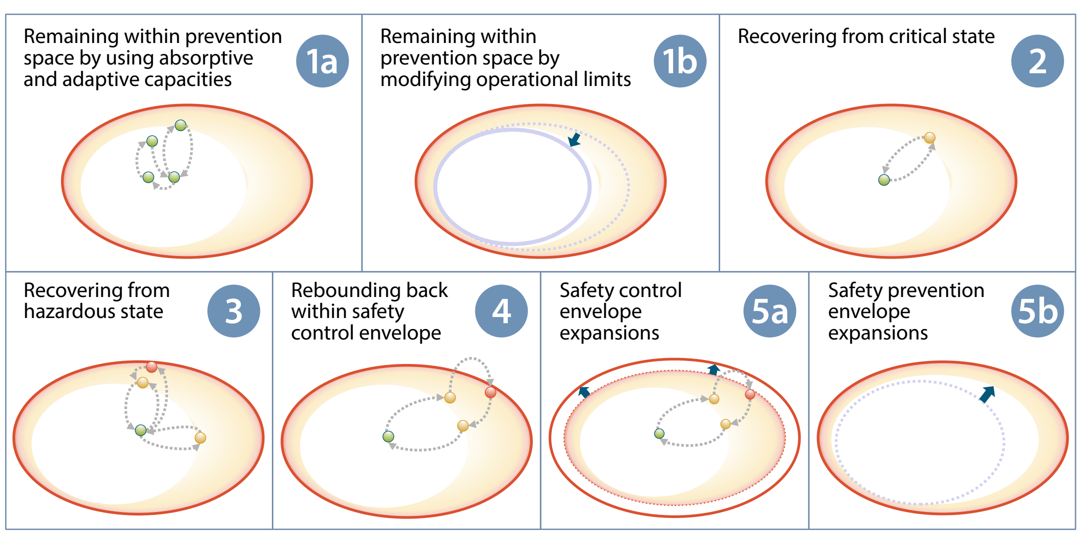

Figure 4: Five patterns of operational resilience

The described Learning From All Operations tools and concepts are used in the following sections for one specific real-life example.

3. Introduction

3.1 Purpose

This case study describes an example of learning from all operations based on flight data monitoring/flight operational quality assurance/flight data analysis programs (FDM/FOQA/FDAP), which we will characterise within Flight Safety Foundation's Learning From All Operations conceptual framework. In this way, although the case may be interesting to a wider aviation safety audience, it is predominantly targeted at an audience that is aware of FDM principles and practice. The information used for the case study is based on a real-world example at a major airline.

3.2 Background

The identified safety issue was a systematic discrepancy between the certified performance of an aircraft during take-off and the actual take-off performance. Especially during operations on performance critical runways without additional margin, this identified issue resulted in an increased risk of runway excursion or collision with obstacles during take-off, and thus, a risk of a major accident.

While the aircraft acceleration was usually in accordance with the expected performance, the rotation was not. It was found that take-off rotation in daily operation was significantly different from the rotation model used for certification. Since the actual take-off rotation was slower than calculated, the pitch which was required for lift-off was reached at a later point in time and thus, the overall take-off distance was significantly increased (Figure 5).

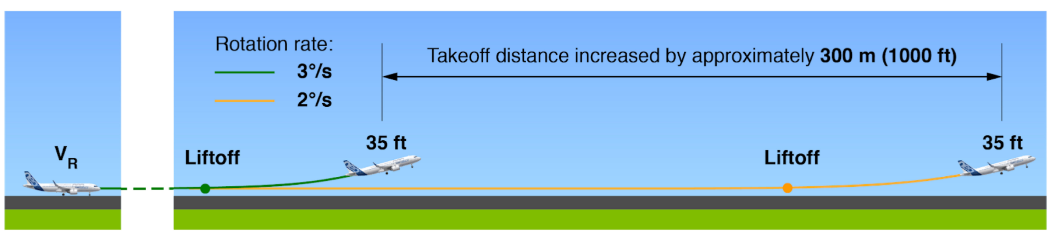

Source: https://safetyfirst.airbus.com/a-focus-on-the-takeoff-rotation/

Figure 5: Influence of lower take-off rotation rate on take-off distance

The issue was initially revealed by an air safety report in which the flight crew reported that they had crossed the departure end of the runway at a very low height. A subsequent FDM analysis identified a systematic problem. This analysis was mainly focused on take-offs at El Dorado International Airport (BOG), Bogotá, Colombia, which is generally performance-critical due to high-density altitudes combined with heavy take-off weights.

For the analysis, a box plot (also called a box and whisker plot) graph was used. A box plot is a standardised way to display relevant statistical data instead of showing the distribution itself (for further explanation, see Figure 6, p. 6). In this graph, the "box" represents one standard deviation above and below the mean value (i.e., approximately 68 percent of all values).

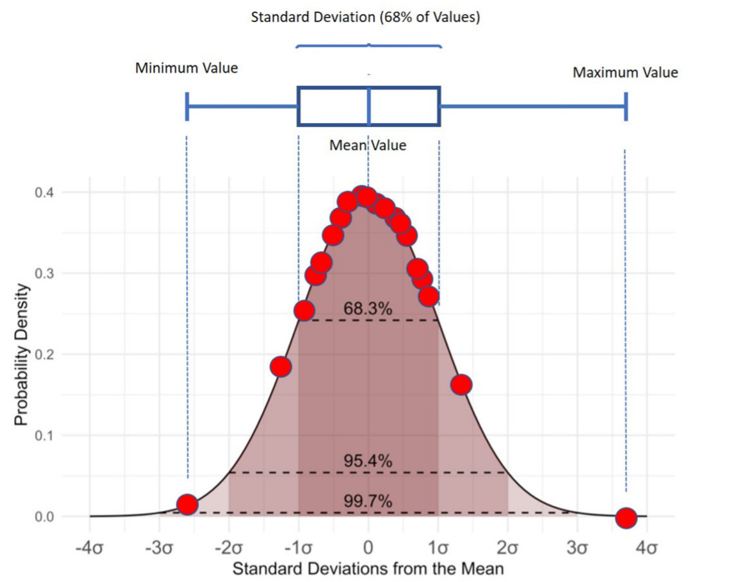

Figure 6: Schematic of a box plot graph

Figure 7 shows a box plot of the threshold crossing heights for a selected time period. Each box plot represents an underlying distribution of crossing heights for the respective month. The box depicts the mean value plus and minus one standard deviation, including approximately 68 percent of all values. The whiskers represent the minimum and maximum values of the respective month.

The red vertical line represents a crossing height which is considered to be safe. All depicted take-offs are with all engines operating, and thus, on average, well above the red line. However, one must always take into account that a certain margin is required if, for example, an engine fails.

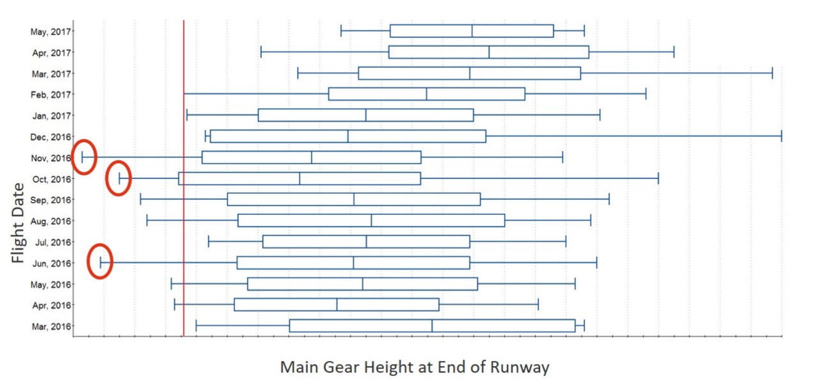

Figure 7: Systematic analysis of runway threshold crossing heights by means of FDM

Particularly large deviations (marked by a red circle) were identified from June 2016 on, shortly after the operation with a new aircraft type started; earlier, another aircraft type with better take-off performance was used at this airport. These initial results suggested a focus on a specific type of aircraft. Such outliers indicate a high risk of runway excursion or collision with obstacles during take-off, especially if an engine fails.

Following the introduction of the particular aircraft type into the BOG operation, the distributions of the crossing heights shifted towards lower values and as a result, values below the red line were observed more frequently. At first, the reason for this shift could not be explained.

Suspicions of erroneous loading of the aircraft or incorrect meteorological data used for the take-off performance calculation were resolved after an initial investigation. Hence, the focus turned to other influencing factors during the take-off roll.

It was concluded that the aircraft rotated more slowly than described in the flight crew operating manuals. However, the details were only revealed through flight data analysis because there is no cockpit instrument indicating rotation rate which would allow a crosscheck of the actual rate with the written procedures. Thus, for flight crews as well as training personnel, it was hard to recognise significant deviations from prescribed rotation rates. Instead, the sense of "normal rotation" for flight crews and training personnel was calibrated accordingly and, therefore, flight crews were trained on rotation rates that did not meet the specifications of the flight crew operating manual.

As long as sufficient margins existed during take-off, the problem was not evident. Because the majority of take-offs fall into this category, the problem isn't evident in daily operations for many operators. Only in the case of performance-critical take-offs, such as those at BOG, did the problem become evident.

Beginning in December 2016, the improvement in crossing heights, seen in Figure 7, resulted from the implementation of mitigation measures, such as artificially reducing the runway length available for take-off performance calculations, once the risk was identified and analysed in depth. This analysis process is described in the following sections.

4. The Studied System

The studied system was the take-off phase of different flights of a certain aircraft type which departs from performance critical airports. The rotation rate issue becomes a problem if the physical limits of the aircraft performance are reached.

When the aircraft performance is at its limit, the flight crew very carefully considers the circumstances so as not to give up the available margins. If no margins are left, flight crews strive to create margins, for instance by reducing the payload to enable the take-off. As long as the calculated performance stays within the certified limits, the flight crew might feel comfortable, even though in reality, there might be no margins remaining.

However, the slow take-off rotation rate reduced the margins to below the certified limits without the flight crew having any way to be aware of it. Flight crews were complying with procedures and did not know about the problem, resulting in limited options to counteract a possible accident risk at the sharp end of the system. Thus, a major contribution to the safety of the system had to be generated by the blunt end, with a significant time lag between the process at the system's blunt end and the operation at the system's sharp end.

4.1 The system's sharp end

The results from the initial FDM analysis as shown in Figure 7 revealed a critical combination of a specific aircraft type and conditions at BOG. The combination of a high-density altitude airport and heavy gross weights at take-off resulted in a critical performance with regard to the take-off distance. This critical combination might be relevant at other airports with similar characteristics and was therefore included in further investigation. Furthermore, Learning From All Operations also means the application of the learning to all operations. Thus, not only other airports had to be analysed, but also other aircraft types.

Regarding the affected aircraft type, a more in-depth analysis had to be conducted to better understand the problem. Since the inquiry already focused on the rotation rate as the main contributing factor, as a first step, the rotation rates of different aircraft types of the same manufacturer were compared, as shown in Figure 8 (p. 8).

The nominal rotation rate is a rate which has to be achieved in order to fulfil the certification requirements of the performance calculation. This nominal rotation rate is described in the Flight Crew Operating Manual (FCOM).

For all three aircraft types studied, this nominal rotation rate is 3 degrees per second, shown in the graph as a black vertical line. Hence, we would expect the mean value of the distributions to be at 3 degrees/second. Flight data analysis, however, shows a significant difference between this nominal rate and the observed values. This analysis nicely illustrates how identifying an issue in one place in the operation can lead to Learning From All Operations.

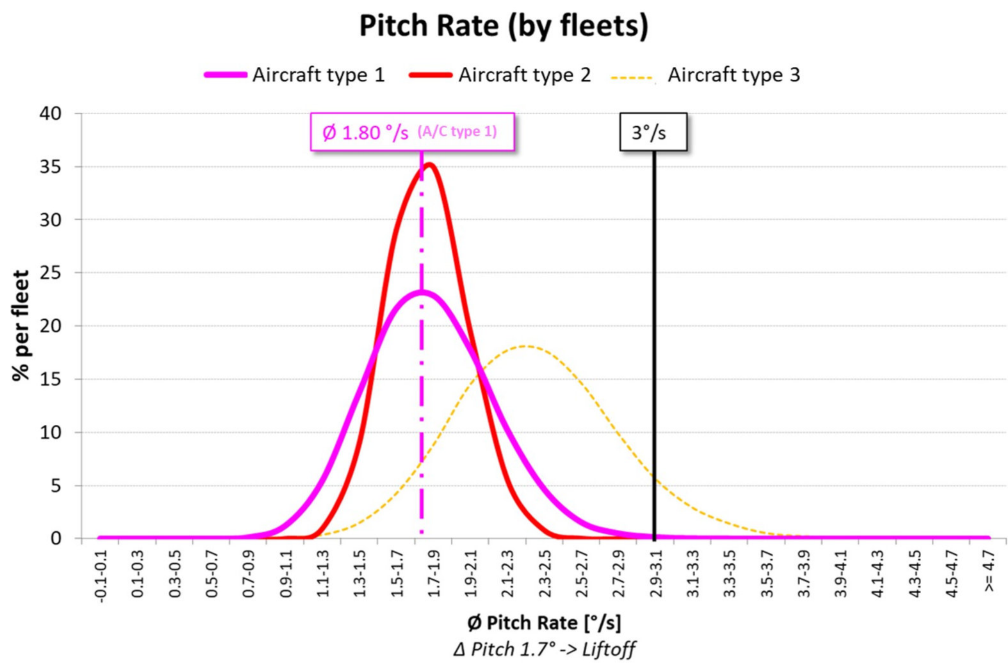

Figure 8: Distributions of observed pitch rates for the three aircraft types, aggregated for all take-offs of the respective aircraft type

Interestingly, an in-depth analysis of the rotation model used for the certified take-off performance calculation showed a significant difference between aircraft type 1 and aircraft type 2. The rotation model of aircraft type 1 uses a value of 1.0 degrees per second for the first 1.7 seconds after start of rotation, and 3.1 degrees per second thereafter. In contrast, the rotation model of aircraft type 2 uses a time delay of 0.7 seconds at rotation speed to allow for a delayed pilot reaction; then there is a rate of 1.3 degrees per second for 2.5 seconds, and finally a rate of 2.3 degrees per second. This model results in a slower rotation for aircraft type 2 than for type 1. This different rotation model results in a better fit of the actual performance versus the certified take-off performance, and hence, leads to no loss in safety margin for aircraft type 2 as long as the take-off is conducted within the certified limits. Even though the FCOM calls for a 3 degrees per second rotation rate, the take-off performance calculations are based on a different model than just straight rate throughout, resulting in performance predictions that are very close to the actual performance of most flight crews.

The reason for these different rotation models is a change in certification standards which came into effect after the certification of aircraft type 1 and before the certification of aircraft type 2.

The safety issue described in this case study is focused on aircraft type 1. Further investigation had to be conducted to reveal whether flight crews could be trained to perform the certified nominal rotation rate and thus, if the risk could be reduced just with training. However, such training would have required a considerable amount of time, and immediate action was required to reduce the risk. Regardless, training of flight crews had to be conducted, an action for the system's blunt end as it required a coordinated decision by the airline management.

The system's blunt end

Because there were limited options for the flight crews to counteract a possible risk at the sharp end of the system, a major contribution to the safety of the system had to come from the blunt end of the system. The system's blunt end in this case included the airline as a whole, the aircraft manufacturer, and the aircraft certification authority.

Two aspects had to be investigated. First, has the training been emphasising the assumptions underlying the performance calculations and the expected rotation rate? If not, possible effects of re-training on the improvement of rotation rate had to be evaluated. This investigation involved the airline's training management.

Second, it had to be investigated whether the manufacturer could modify the take-off performance calculation module according to the parameters observed in actual flight operations. This investigation also required the involvement of the aircraft certification authorities.

5. Selecting the Learning Approach

5.1 Learning direction

The selected approach for the case study learning direction is top-down (starting from the undesired aircraft states) because the issue was not known in detail before the first safety-critical event occurred.

This safety issue had been identified through the standardised hazard identification process of the implemented safety management system (SMS) by means of an air safety report (ASR).

Once the safety issue had been identified, the first step of the safety risk management (SRM) process was to build a qualitative model of the issue. This process typically starts with the definition and description of a hazard, which is an undesired aircraft state. In our example, the hazard is defined as "low height at departure end of runway overfly". Besides the hazard, our qualitative risk model also contains pressures/threats, which lead to the hazard, as well as possible outcomes.

The main threat leading to the hazard was identified as the slower-than-calculated rotation rate. Other threats are a critical take-off performance, because during non-critical take-offs, the hazard would not arise.

The worst foreseeable outcome would be a hull loss because of a lift-off beyond the physical runway. In the next step of the risk management process, the quantification of the risk of this worst foreseeable outcome had to be determined. However, quantification required data. Since airlines are required to implement FDM, this tool was used for the subsequent analysis.

5.2 Learning scope

The SRM process as described above requires another step beyond quantification. The overall goal is the management of the risk, which requires adequate and efficient mitigation measures. To develop and implement such mitigation measures, the risk has to be understood in depth. This requires not only the knowledge about quantity of the risk, but also the pathway of how the risk evolves. This makes necessary the use of insights into routine operations (i.e., Learning From All Operations).

Flight data is the most precise and efficient way to gain insight into a mechanically caused safety issue. To gain more insight into all operations, the distribution of flight data is a good way to describe certain relationships between causes and outcomes.

The learning scope of our case study was therefore mainly based on flight data. Even though in the first step of the hazard identification process, only exceedances were searched for to get an idea about the extent of the safety issue, understanding of the evolution of the issue required examining distributions across all operations.

5.3 Key learning parameters

Key learning parameters were selected for the purpose of the case study specific analysis. FDM exceedance events, which indicate a possible safety event, may be different than measurements used for the further analysis of the risk. In this case study, an event indicating a crossing height below a certain identified threshold was used. The correlation between a low crossing height with an increased risk indicates a situation that might lead to an accident.

To understand the whole safety issue, a more in-depth investigation of rotation rates was necessary. Moreover, the interaction between input (flight crew input on controls) and output (achieved aircraft rotation rate) is relevant to verify possible discrepancies between actions of flight crew and the outcome (i.e., the observed rotation rate). This enables possible improvements by means of training.

6. Operational Limits and Safety Envelopes

6.1 Operational limits

Performance-limited take-offs are rather unusual. Hence, in daily operation, large margins typically exist between the system's operating point and the boundary of the safety control envelope, represented by position 1 in Figure 9, below.

A normal take-off has multiple sources of safety margins, such as:

- Use of reduced take-off thrust with assumed temperature (FLEX T/O, TASS) results in a stopping margin due to a lower true airspeed than anticipated for the thrust calculation;
- Conservative use of environmental parameters (e.g., no incorporation of head wind, use of small margins in outside temperature and atmospheric pressure); and,
- Stopping margin due to longer runway than required, even if maximum thrust reduction has been reached.

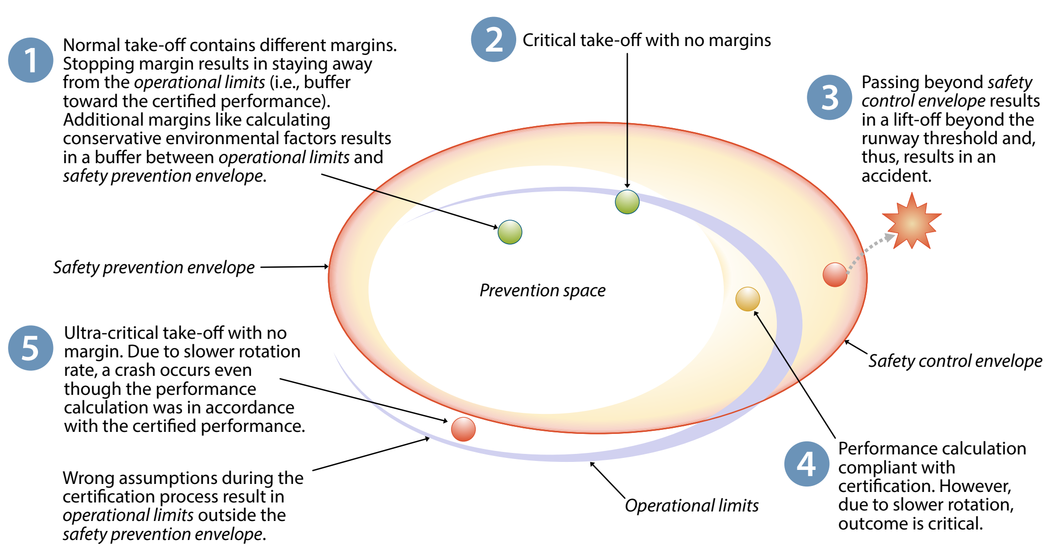

Figure 9: Examples of system states of take-offs in different environments

The operational limits are primarily defined by the aircraft manufacturer during flight certification in concordance with the certification authority. These limits are directly derived from the physical performance limits of the aircraft, usually incorporating certain margins. These operational limits are integrated in the performance calculation software used by the flight crews before each take-off. They are considered valid operational limits, which still include a certain margin before reaching the boundary of the safety control envelope.

Especially for normal take-offs, the operational limits arc in Figure 9 is located away from the yellow safety area (i.e., there is sufficient margin between the system operating point 1 of the particular take-off and the operational limits arc itself, as well as between the operational limits arc and the boundary of the safety control envelope).

However, sometimes the performance calculation results in reduced margins or even no margins at all, especially if the environmental factors are critical and the runway is relatively short. In this case, the operational limits arc could touch the boundary of the safety prevention envelope.

If a discrepancy develops between the assumed location of the system operating point based on the certified performance calculations and the actual location in daily operations, the operational limit arc may be outside the safety prevention envelope, or even outside the safety control envelope. In this case, an accident may occur even if every participant acts in full compliance with the certification of the aircraft.

6.2 Safety envelopes

The safety prevention envelope describes the space within which a safe take-off can be achieved even if critical systems fail. A salient example of such a failure is an engine failure, which is covered by the take-off performance calculation as required by the certification authority. As long as no critical system fails, a take-off can still be conducted without damage, even if the system operating point shifts beyond the safety prevention envelope but remains within the safety control envelope.

As soon as a critical system fails, and the system operating point is located beyond the safety prevention envelope, the system operating point immediately shifts further towards or even beyond the safety control envelope boundary and may become an accident.

If a take-off is ultra-critical regarding the performance, the system operating point could also shift beyond the safety control envelope without any failure of critical systems.

7. Information Collection and Analysis

Once the issue has been identified and assessed, the SMS requires adequate risk mitigations to control the risk. Because the operational limits arc penetrates the boundary of the safety prevention envelope and even the boundary of the safety control envelope, in certain cases, the main objective of this effort to learn from all operations was to shift the operational limits arc so it is contained within the boundary of the safety prevention envelope. Only this shift would guarantee that the system operating point will always stay within the safe nominal operations space (prevention space), as long as the flight crews remain compliant.

7.1 Would training be a solution? Rotation rate vs. controls deflection

The most desirable mitigation would be the elimination of the risk. That would be the case only if the actual rotation rate could be increased to match the certification rotation rate. Since the certification rotation rate could not be observed in daily operation, it had to be determined whether such a rate could be achieved at all.

It also had to be determined if increasing the rotation rate could lead to undesirable consequences such as an increased probability of a tail-strike. A few years prior to this study, the airline experienced several tail-strikes due to higher-than-normal rotation rates and as a consequence adapted conservative rotation techniques in its training.

To determine whether the operational rotation rate could be shifted to match the certification rotation rate, the relationship between the deflection of the flight controls by the pilot during rotation and the corresponding rotation rate had to be analysed. Figure 10 shows that relationship.

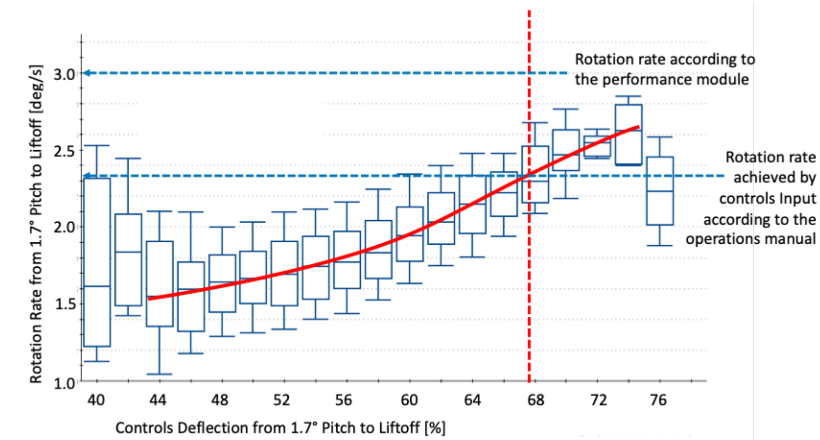

Figure 10: Rotation rate as a function of control deflection

Each box plot in Figure 10 (as well as in Figure 11, p. 13) shows the distribution of rotation rates in degrees per second for a given percent control deflection. For instance, the left-most box shows that crews who pulled back on the controls (yoke/stick/sidestick) to deflect 40 percent of the control device movement range achieved on average a 1.62 degrees per second rotation rate. Approximately 68 percent of these crews achieved a rotation rate between 1.23 and 2.31 degrees per second (the length of the box), with a low extreme at 1.12 and a high extreme at 2.53 degrees per second (the tips of the whiskers). The solid red line in the figures is a function of the average values of all the distributions.

The FCOM calls for rotation using ⅔ of full control deflection, which corresponds to about 67 percent, as shown by the dashed red vertical line in the graph. The intersection of the dashed red line with the solid red line shows that on average, using 67 percent control deflection would lead to a rotation rate of 2.32 degrees per second (as noted by the dashed blue line). This result shows that the controls deflection required by the FCOM would not be sufficient to generate the required certification rotation rate.

Moreover, this analysis shows that the certification rotation rate could not be achieved by any of the control deflections seen in operation, which is a finding worthy of further investigation.

Another learning opportunity provided by such an analysis is the possibility to predict outcomes (i.e., the resulting rotation rate, by varying controls deflection). For example, if the controls deflection input of a given pilot would be 55 percent, the resulting rotation rate could be expected to be around 1.8 degrees per second, as can be seen in Figure 11. Thus, proper training goals could be set using such analyses.

This analysis could also be used in one more way. Recall from Figure 8 that the average rotation rate for aircraft type 1, the focus of this case study, was 1.8 degrees per second. Given Figure 11, it can be determined that the average crew applies a 55 percent control deflection during rotation. Again, this is an important piece of feedback to the flight crews that can be used during training, and as part of an educational campaign.

7.2 A different solution? Virtual shortening of the runway

Instead of shifting the system operating point into the safe operations space (prevention space) by means of crew training, the organisation could ensure that the operational limits themselves always stay within this safe space to ensure no unintentional outliers beyond the safety prevention envelope.

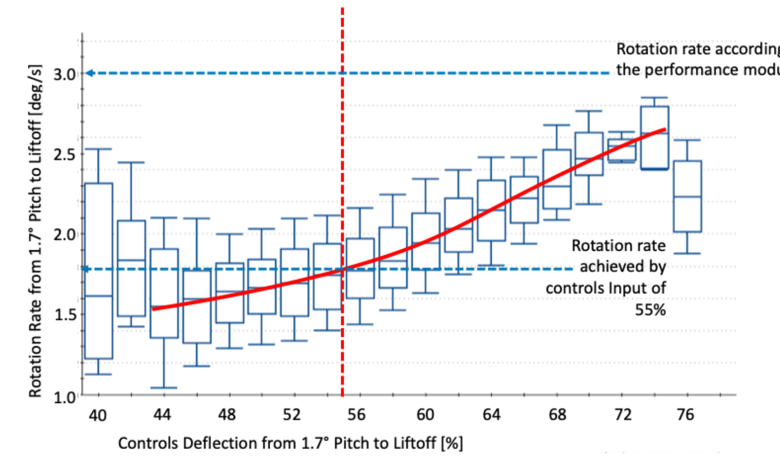

Figure 11: Learning from the graph. How would a certain change in training influence the overall outcome?

If the risk of a slow rotation rate is not mitigated by the manufacturer and the certification authority by changing the performance calculation to more practical conditions, the organisation (airline) has to mitigate the risk by implementing adequate margins. A likely mitigation is a payload reduction that results in an economic impact.

To minimise this economic impact, a new balance between protection and production had to be found. The amount of payload reduction, resulting in lower earnings, had to be balanced against the safety increase by gaining safety margin in crossing height at the departure end of the runway. A thorough analysis of FDM data was required to determine this balance.

Slower-than-normal rotation rates require additional take-off distance. Flight data were used to determine the typical additional distance. An example of how this additional distance can be measured is shown in Figure 12. In this graph, the rotation rate from 1.7 degrees to lift-off is plotted against the distance used from the first controls input until lift-off.

Beside the scatter plot (a graph displaying the values of the two variables), a box plot is shown for airspeeds of 155 knots +/–3 knots. The regression line connects all the median[1] values of the

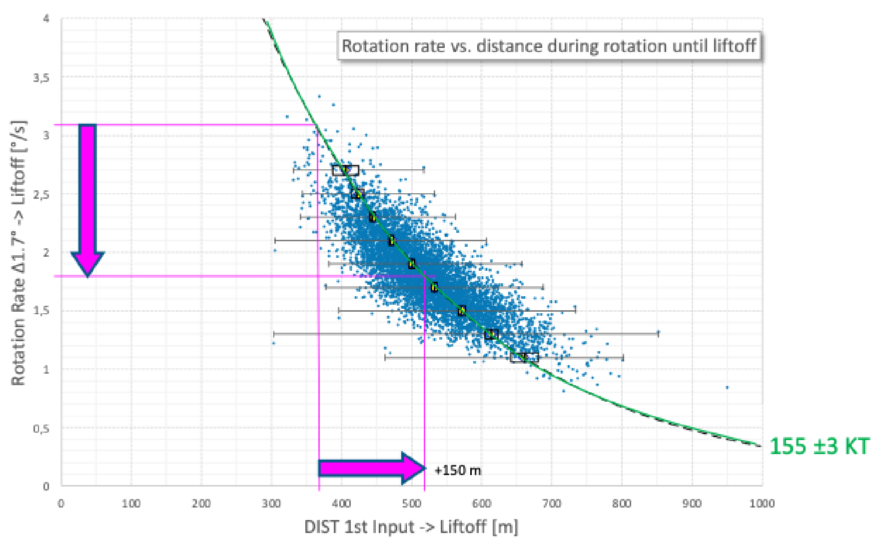

Figure 12: Scatter plot of the rotation rate and the resulting distance between first control input and lift-off

[1]Median is the point from which 50 percent of all measured values are higher, and 50 percent of all measured values are lower.

respective boxes. This regression line enables the evaluation of the additional distance required for slower-than-nominal rotation rates. The graph shows the increase of required take-off distance due to the slower than calculated rotation rate. For airspeeds of around 155 knots this results in an increase of take-off distance of 150 meters.

The distance resulting from the rotation rate that is required for the certified performance can be compared with the distance which results from the typically observed rotation rate in all operations. To cover the majority of flights, the standard deviation can be added to the mean value of the observed rotation rate. Provided the observed rotation rates are symmetrically distributed, this would cover approximately 84 percent of all take-offs (Figure 13).

In Figure 13, the mean observed rotation rate including one standard deviation towards the lower rotation rates result in an additional take-off distance of approximately 200 meters. As a possible risk reduction measure, the runway could be virtually shortened by this amount of additional distance for the calculation of the take-off performance. This distance, which can be understood as an additional safety margin to cover the deficiencies of the slower-than-nominal rotation rate, can be incorporated into the performance calculation module as a standard for all take-off calculations by the operator.

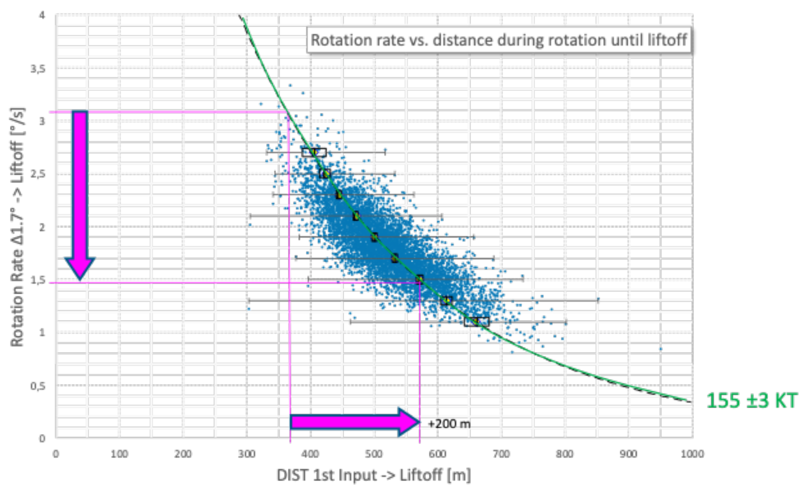

Figure 13: Increase of required take-off distance due to the slower than calculated rotation rate. Including one standard deviation of all observed rotation rates, this results in an increase of take-off distance of 200 meters.

7.3 Three forces and resilience mechanism

To fulfil the operational system purpose, the demand pressures include flight operations at different types of airports (including some that may be performance-critical), the use of aircraft types that are available in the airline fleet that are best suited for the specific operation, and the management of other threats like adverse weather, engine failure or flight crew personal pressures (Figure 14, p. 15).

The demand pressures are influenced by the aircraft type, which is critical for this type of airport. However, change of aircraft types could lead to other issues (e.g., an aircraft type that is less critical with respect to take-off rotation, but is much more limited with respect to landing performance).

The efficiency pressures are mainly economic factors. The most important goal is to carry as much payload as possible. At first glance, the aircraft could be operated safely as long as the operation is compliant with its certification.

Another efficiency pressure is time. Even though an ultra-critical take-off can be planned carefully, conditions can change quickly. Both efficiency pressures are affecting the system's sharp end and require resilience by the flight crew to counteract those pressures.

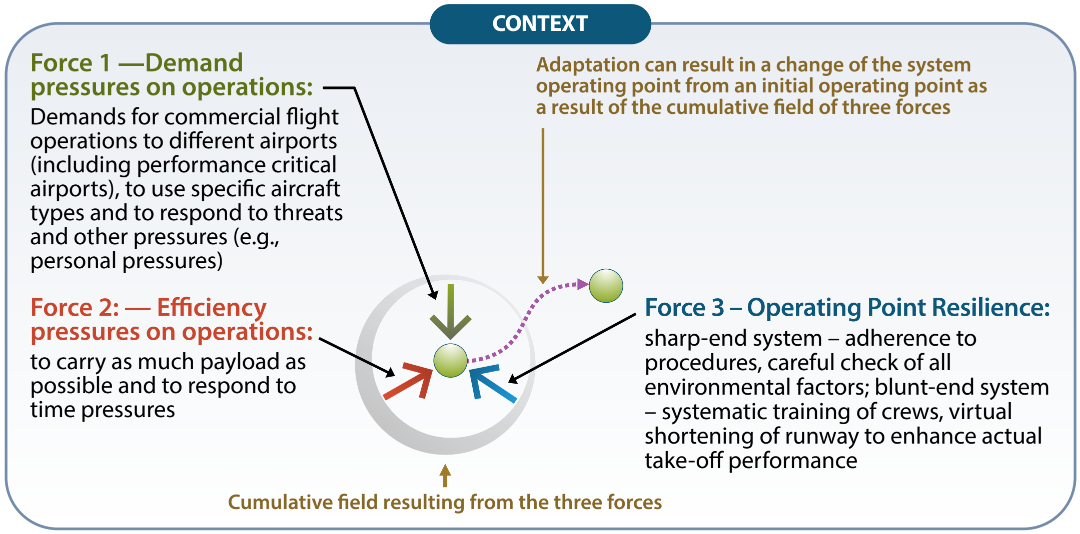

Figure 14: Three forces of the case study which influence the system state

However, the resilience potential in this specific case study is mainly achieved by the system's blunt end. System resilience is achieved by implementing risk reduction measures, which must be implemented before the actual operation.

7.4 Manifestation of resilience

In the context of the described safety issue, resilience to counteract a negative shift of the system's operating point towards the safety prevention envelope does not come from the actions of the sharp end due to their limited possibilities. Resilience by the flight crew has limited effect in this case because the aircraft is operating at the boundary of its physical limits.

The proactive building of skills and capabilities lies mainly with the blunt end and was achieved by a detailed analysis of flight data and the resulting distributions (Figure 15). It required learning from all operations and is part of the learning process of the organisation. It's this learning that forms the foundation of the system's resilience.

Following the implementation of risk reduction measures, it is vitally important to monitor the implemented measures closely and adapt them if necessary.

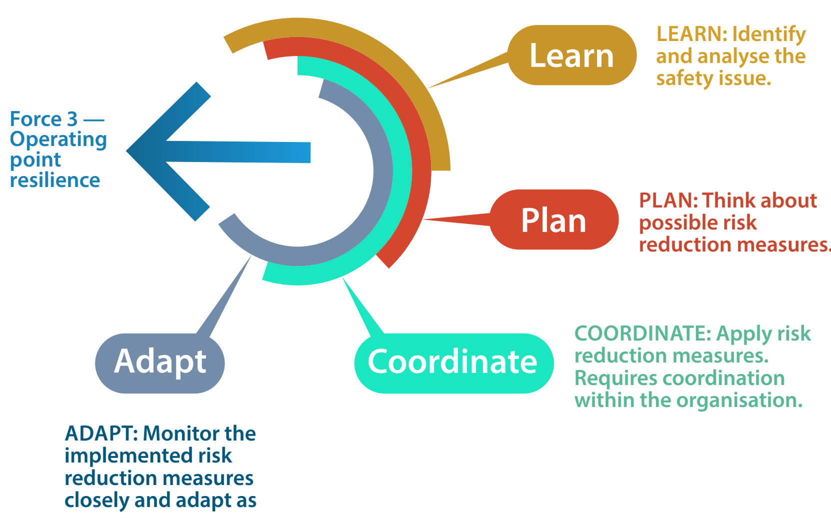

Figure 15: Overview of the generic resilience capabilities at the system's blunt end

8. Conclusion

Two different approaches were considered to counteract the identified risk of not achieving required take-off rotation that leads to increased take-off distance and a decreased obstacle clearance. The first approach was the systematic training of flight crews to apply higher rotation rates, to shift the system operating point towards the prevention space. In Figure 16, this effect is indicated by the transparent arrow. This approach was found to be inefficient in this case.

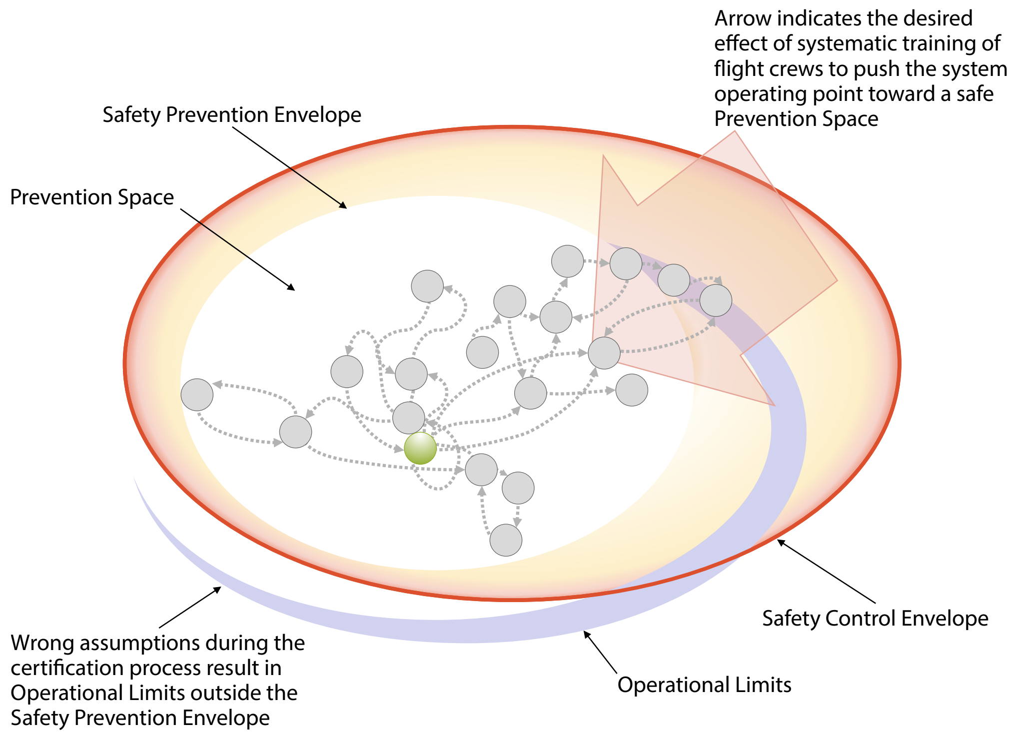

Figure 16: Effect of systematic training of flight crews with regard to higher rotation rates

The second approach, which was adopted by the airline, was the virtual shortening of the runway in order to provide additional safety margin. This measure places the operational limits inside the prevention space and counteracts the wrong assumptions that had been made during the certification process. This effect can be seen in Figure 17 (p. 17).

In summary, this case study provides an operational example that illustrates critical aspects of Learning From All Operations:

- **Learning from everyday work in addition to learning from exceptions.** Distributions of routine operational data were used to create a clearer picture of overall operations, which provided context for recognising and understanding exceedances.
- **Leveraging and evolving existing processes, practices and tools to expand opportunities to collect, analyse and act upon safety critical insights.** Learning From All Operations does not require wholesale replacement of existing safety approaches. In this case study, different sources of safety data successfully worked together — the initial trigger was provided through flight-crew safety reporting, and subsequently, flight data monitoring was used to look at the entire range of key parameters to understand the relationships among the different factors affecting the safety margins.

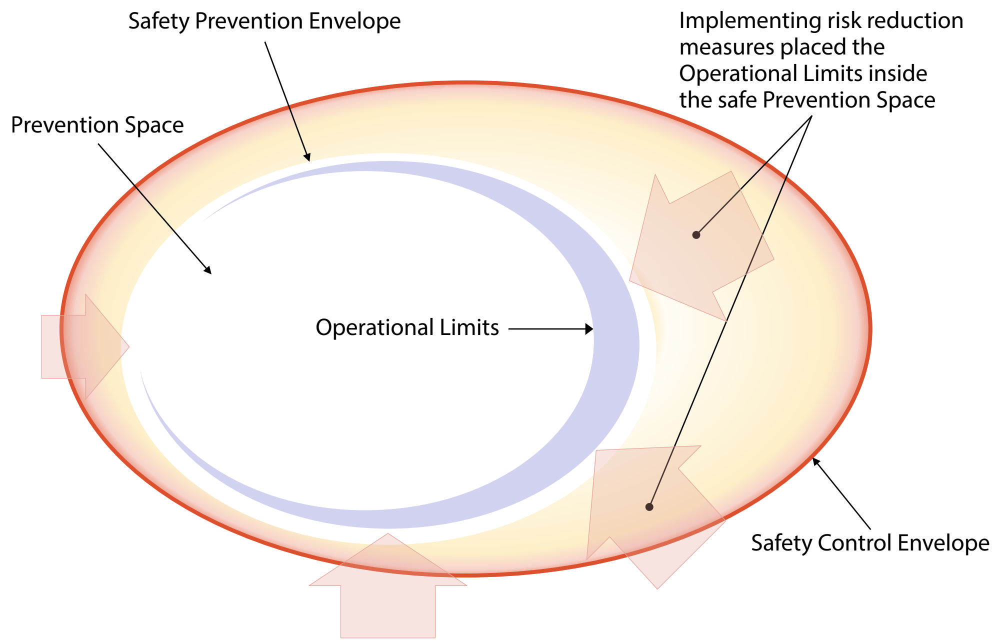

Figure 17: Effect of virtual shortening of the take-off runway

- **Involving everyone.** Operators at the system sharp end are frequently the focus of safety management, due to their physical and temporal proximity to hazards, and thus their frequent potential to influence safety outcomes. This case study, however, illustrated a situation in which a systemic risk was difficult to detect and likely impossible to resolve without the help of the system blunt end. Analysis of a broader range of safety data enabled the system blunt end to implement changes that could be transferred to all operations.

While this case study was focused specifically on how one organisation learned from all operations to create safer take-off rotation procedures, the lessons have far broader implications. As illustrated in this case study, Learning From All Operations can help organisations to understand the adaptations that personnel make to keep systems operating, recognise slow or unobtrusive changes and respond before unwanted events occur. While many ultrasafe organisations may be able to identify specific examples of learning from all operations within their institution, the full benefits of learning from all operations are realised when an organisation is able to do so systematically and consistently — when Learning From All Operations becomes part of a routine safety mindset that expands our understanding of what constitutes a safety-relevant occurrence and improves our ability to consistently learn from what happens.

9. Acknowledgements

This work was funded under grant and cooperative agreement 80NSSC21M0187 from the U.S. National Aeronautics and Space Administration's (NASA) System-Wide Safety Project, part of the Aeronautics Research Mission Directorate's Aviation Operations and Safety Program.

This case study was drafted by Dr. Max Butter, Lufthansa captain, Tzvetomir Blajev from Flight Safety Foundation, Dr. Immanuel Barshi from NASA and Dr. Jon Holbrook from NASA. Thank you to the members of Flight Safety Foundation's Learning From All Operations Working Group, who contributed to the ideas and clarity of this case study: Valerie Stait, Cathay Pacific Airways; Capt. Hideaki Miyachi, Japan Airlines; Capt. Tom Becker, TUI Fly; and Capt. Bertrand de Courville.

Suggested citation: Flight Safety Foundation. (2022). Learning From All Operations Concept Case Study: Take-off Rotation.