FLIGHT SAFETY FOUNDATION

LEARNING FROM ALL OPERATIONS

Concept Note 3

Safety Envelopes and Operational Limits Assumptions

Flight Safety Foundation

MARCH 2022

FLIGHT SAFETY FOUNDATION

A note to the reader:

The goal of these Concept Notes is to provide a common framework and common language for talking about aviation safety. Such a new framework and language are needed because the existing language of safety is built around learning from failures and cannot easily express learning from success. Similarly, the existing frameworks of safety data collection and analysis are designed for incidents and accidents, and we want to learn from all operations.

As we expand our understanding of what constitutes a safety-relevant occurrence—an expansion that encompasses learning from all operations—we need a shared means of articulating what we are already learning that also allows us to discuss new ways of learning. Positing a separate framework for describing safety successes, however, can create challenges for relating what can be learned from success to what has been learned from failure. Therefore, the goal is to describe a unitary framework for safety based on learning from all that happens, rather than separate frameworks for different "kinds" of safety. To achieve this goal, each of these Concept Notes establishes part of the necessary foundation, which is then integrated and translated into practical implications and applications in Concept Note 7.

1. Introduction

This note introduces an important learning dimension of Learning From All Operations. It concerns monitoring aviation operations and **learning about the difference between what is assumed to be safe and what is actually safe.** These are two important concepts:

- **Safety control envelope** —the actual boundaries of what is safely recoverable in operations by preventive or recovery measures (outside this boundary, the safety control becomes marginal to non-existent). It is to be noted that sometimes where the actual safety envelope is may not be known in foresight but only in hindsight; and,
- **Operational limits assumptions** —the imagined boundaries for operations (normative—rules, procedures, prescriptions or the subjective assumptions about where these boundaries are).

To describe these two concepts and their relationship, this note uses the system and system states terminology introduced in Concept Note 2.

2. Safety Envelope

2.1 What is a safety control envelope?

Concept Note 2 described how a system dynamically changes its states in the performance space defined by some relevant system parameters (Figure 1). We showed, as an example, how the system operating point transitions for a flight from take-off to landing in a space defined by the values of aircraft altitude and speed. It is well known in aviation that in this specific example, the possible positions and transitions of the operating point are constrained by the airplane flight envelope. The airplane flight envelope is confined by the stall, maximum altitude, and top speed boundaries as illustrated with the red line in Figure 1.

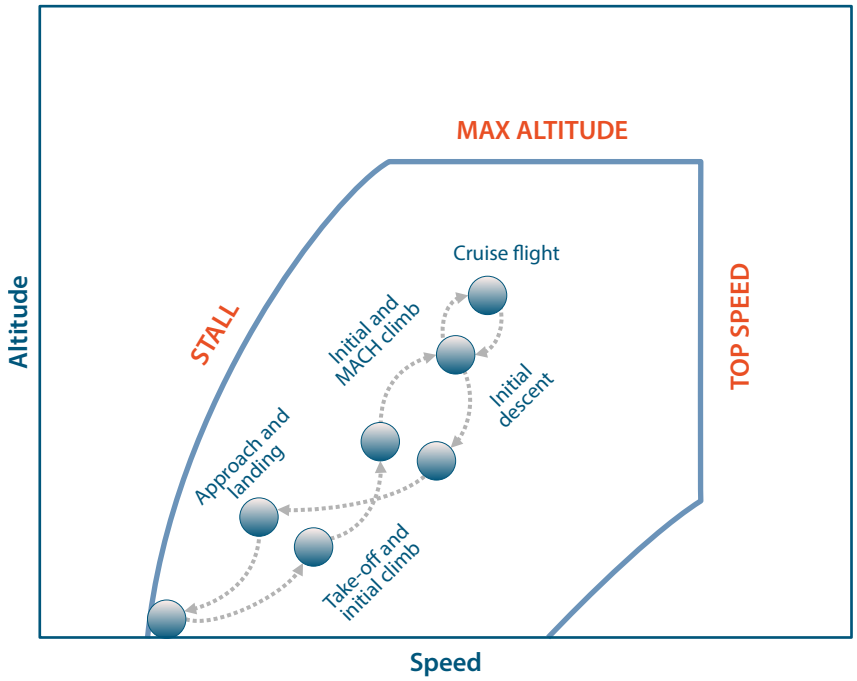

Figure 1: Flight Envelope Example

The envelope is the region in which aircraft can operate safely. It is objective—it exists independently of our understanding of it. In the example of the flight envelope, the critical boundaries are determined by the objective laws of aerodynamics. If the aircraft state transitions outside the envelope, then the aircraft will be in a hazardous state (stall or structural damage) with a limited possibility of preventing an accident.

We can also speak of a safety envelope in the performance space that is defined by any other safety-relevant parameter. The **safety control envelope** boundaries will be determined by those critical values of the parameters beyond which the system will critically lose control over flight safety.

2.2 Stability of the systems states and safety envelope

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** (Leveson, 2004). If we introduce more capabilities that increase our level of control, then the safety control envelope will expand. For example, in a surveillance air traffic control (ATC) environment, the critical distance between two aircraft in the air will be much less than the critical distance in a non-surveillance environment.

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 (Rasmussen, 1997). The system response to pressures can lead to a system state that is either stable or unstable. In this way, the concept of controlling flight safety is connected to the concept of system stability.

A useful way to visualize the concept of system stability and system control is with a ball-in-cup analogy (Holling et al., 1995; Holling, 1996; Reason, 2008; Figure 2). In this representation, the state of the system is represented as the position of a ball rolling on a surface.

The illustration is for a system state defined by one system parameter—for example, the distance between an airplane and terrain. All possible states of the system can be represented by a specific position of the ball. As the ball (the system state) moves under some pressures, the distance to terrain is changing. When the movement is away from the central, stable operating point, then the distance to terrain is decreasing.

There is a tipping point when the system becomes unstable, represented in Figure 2 by the colour of the ball becoming yellow. **This is a system critical state.** In our example, it means that the trajectory of the aircraft consistently turns towards terrain, and from this moment, the distance

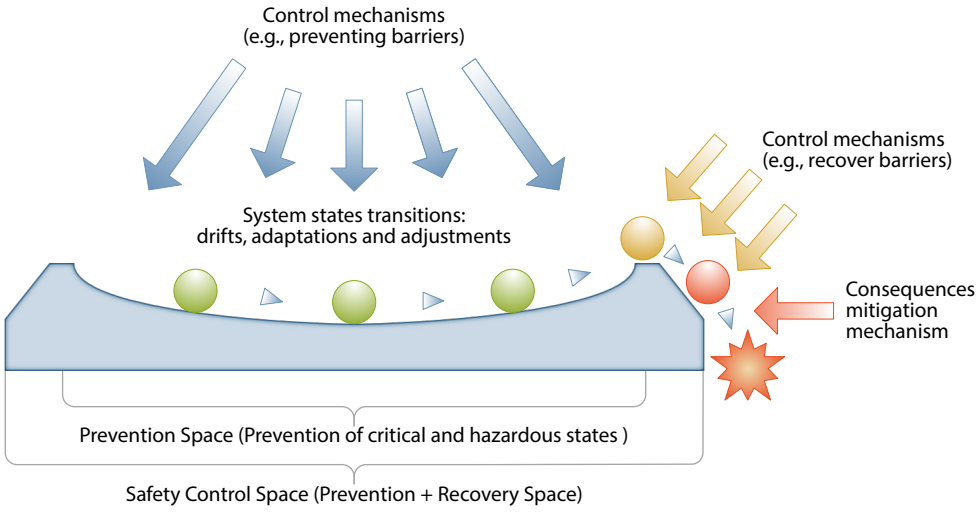

Figure 2: Visualising System Stability and System Control

to terrain is continuously decreasing, and there is a need for a recovery action. In the absence of such recovery action, the ball will roll downhill—in other words, the system will transition to a **hazardous state with marginal control on safety and a prompt need for recovery action to keep the system within the safety control envelope,** represented in Figure 2 by the colour of the ball becoming red. In our example, the hazardous state would be a critical distance between the aircraft and terrain, indicating a prompt need for recovery actions to prevent an accident.

As described above, **it is critical to determine our capabilities to control such situations**. If we have high controllability of terrain clearance (and no additional disturbances), then we can allow ourselves to fly close to the terrain. If we have lower capability to control the terrain clearance—for example, flying over mountains with possible mountain wave effects—then the critical distance to terrain will have to be greater.

The system state transitions can be anticipated from the shape of the surface and from the dynamics of the pressures and counter-pressures applied to the system at any given moment. In aviation, we usually refer to these counter-pressures as preventing mechanisms (i.e., preventing the system state from transitioning to a critical state), recovery mechanisms (i.e., helping the system recover from a critical or hazardous state), or consequence-mitigating mechanisms (i.e., mitigating the consequences of the system passing beyond the safety control envelope, by, for example, lessening accident consequences).

For stable systems, the ball is in the "valley" (i.e., the low point in the cup) and the pressures and counter-pressures keep the ball away from the tipping point of critical system states. Unstable systems may have sudden large operating point shifts that can take them close to and through the safety control envelope boundaries (Cook, 2005). The slopes and the size of the valley represent one way, but not the only way, to characterise some aspects of system resilience.

2.3 Safety prevention envelope and safety recovery envelope

As previously described, the objective boundaries of safety-controlled operations are confined within the safety control envelope (Cook, 2005). Figure 3 provides a view from above the ball-in-cup metaphorical diagram in Figure 2.

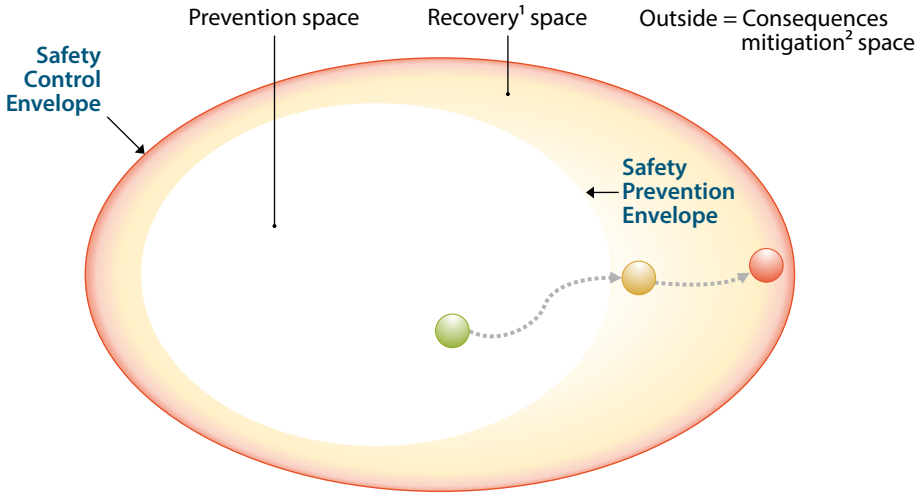

Notes:

1. Recovery is possible also outside the safety control envelope, but it is mostly by chance.
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 3: Prevention Space and Recovery Space

In the middle of the envelope is the white area—the safety prevention envelope. Within this space, the system is adapting, coming closer or moving away from the critical thresholds (Leveson, 2011). Prevention here is used to denote prevention of system transitioning to critical or hazardous states.

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 space indicates growing closeness to the safety control envelope. Here the system is unstable: At least one of the safety critical parameters is consistently changing in the direction of its critical threshold, and there is a need for control action so that the system can recover back to the prevention space.

If the 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.

The space outside the red area illustrates a situation outside the safety control envelope, (e.g., a distance and rate of closure between two aircraft that does not provide sufficient time for collision avoidance). 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, including luck. But passing through the safety control envelope is associated with a significant loss of control over flight safety, with only marginal safety control, if any, available.

The illustration of safety envelopes in Figure 3 defines the spaces where the system operating point can transition. The system adaptation process can result in system operating point transitions within the prevention space as well as the recovery and mitigation space. During these transitions, the system operating point can enter and cross the red area of system hazardous states. This conceptual framework brings together the concepts and the language of the existing risk and hazard terminology with the concept of system adaptation and system space for adaptation.

3. Limits as Assumed

3.1 What are the limits?

In addition to the objective safety control envelope, there are subjective limits, often describing threshold values of certain parameters based on our understanding of "how far can we go?" For example, speed limit road signs define the operational limits in terms of maximum allowed speed for a given stretch of road. Our understanding of the speed limit is the operational limit assumption, whereas the objective safety control envelope is the speed at which we will lose control of our car at a particular curve in the road, given the specific outside conditions and driver abilities.

Similarly, maximum aircraft taxi speed when vacating a runway can be limited by procedures, for example, to 30 kt. A flight crew may assume that this limit is 25 kt or 35 kt, and this will be their understanding of the limits—their operational limit assumption. The objective safety control envelope when an aircraft will lose directional control, depending on the operational conditions may be 20 kt in icing conditions or 45 kt if the runway is vacated via a rapid exit taxiway, the runway and the taxiway are dry, and the rudder pedals are used for directional control. As seen from these examples, **the subjective operational limits assumptions and the objective safety control envelope are two different things and their relationship is part of the more generic distinction between operations as they are imagined and operations as they actually exist** (Dekker, 2006).

Rules and regulations are a major source for operational limits. The specifications of technological systems and the associated manufacturer-suggested procedures also define what is assumed to be acceptable system behaviour. Using the rules, regulations, manufacturer guidance and their own experience and knowledge, companies develop standard operating procedures (SOPs) that define the operational limits to be used by professionals.

3.2 Missing operations limits

Sometimes the limits are not fully defined. In Figure 4, the operational limits are illustrated by the purple arc (1). The arc is confining the system operating space as defined by the operational limits, and the gap in the arc (2) illustrates missing operational limits.

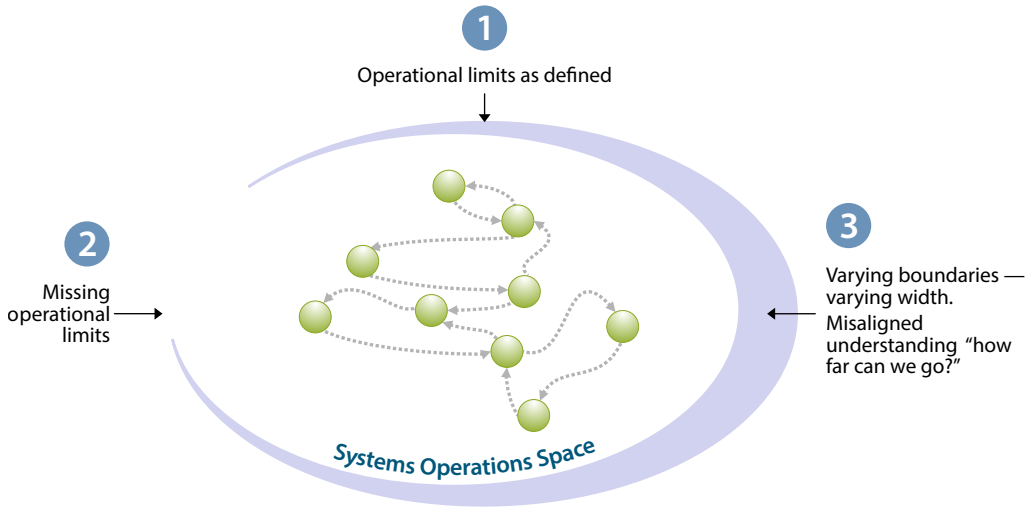

Figure 4: Operational Limits

Here are some examples of missing operational limits as reported by members of the aviation community:

- What action should an airport and air navigation service providers (ANSPs) take when a bird strike is reported? Shall the runway be closed for inspection, or shall the action be only to inform the next aircraft on approach?
- How should it be decided what constitutes heavy or moderate turbulence?
- How should the system react to reports of space weather hazards (e.g., solar storms)?

3.3 Varying operational limits

Limits can vary (Cook, 2005)—as indicated by the varying width of the purple arc (3) in Figure 4. Different actors could assume different limits, for example:

- ATC uses a minimum approach speed that may be different from that included in the flight crew stabilised approach criteria; or,
- Separation at the border of controlled and uncontrolled airspace may be different for different air traffic services (ATS) units.

A more detailed illustration of varying operational limits is provided in Figure 5 (p. 6). The illustration is a zoomed-in perspective of part of the operational limits that shows different understandings by the various actors that are involved in operations—illustrated with the yellow lines, which represent different assumptions about "how far can we go?" What is also to be noted is the possible difference between the assumptions and the operational limits, as actually defined by specifications, rules or procedures—illustrated with the blue line.

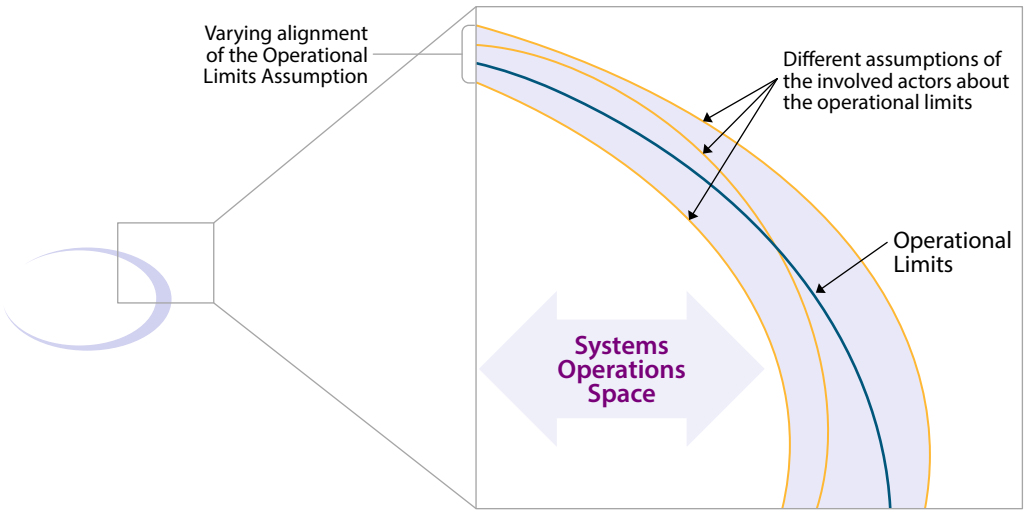

Figure 5: Varying Operational Limits

4. How do operational limits relate to safety boundaries?

Ideally, the operational limits will neatly protect system operations from breaching the safety control envelope. This is shown in Figure 6. Some of the yellow space is not protected by the operational limits. This may well be a deliberate choice. As a reminder, when the system is in the yellow space, it means that at least one safety-relevant parameter is consistently reducing towards its critical threshold. For example, when trajectories of two aircraft become conflicting, then the system operating point will be in the yellow space. Airspace design and procedures try to strategically deconflict the aircraft trajectories, but some conflicts are left to ATC to tactically resolve. In this example, the purple line of operational limits would be the ATC separation minima.

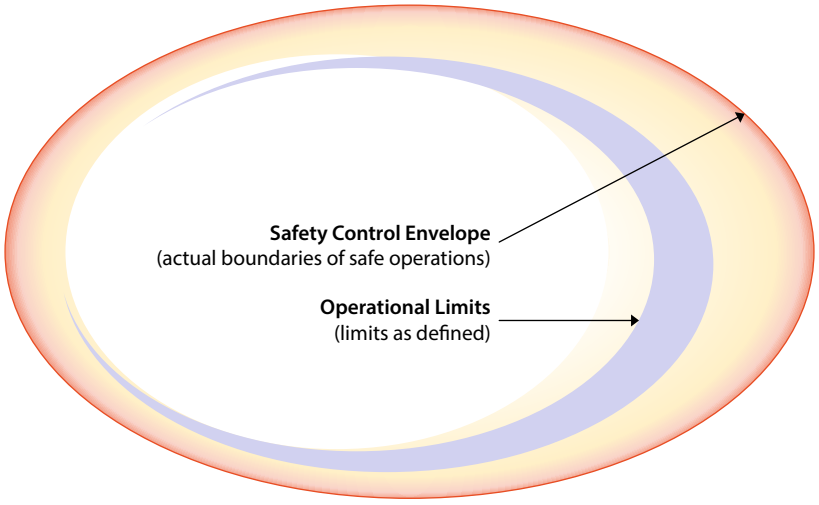

Figure 6: Operational Limits Protect Against Breaching the Safety Envelope

But in reality, limits of control and limits as defined have a more complex relationship. For example, system states can be within operational limits (assumptions) but beyond the safety envelope (e.g., flight duty time limitations do not always guarantee a particular level of pilot alertness).

Here are some reasons why operational limits and the safety envelopes may not be properly aligned:

- Operational limits designed for some average performance may leave some part of the normal operational performance distribution outside the limits.
- Operational limits may be based on inaccurate assumptions.
- Operational limits may not accurately define the boundary conditions and may miss a significant parameter.

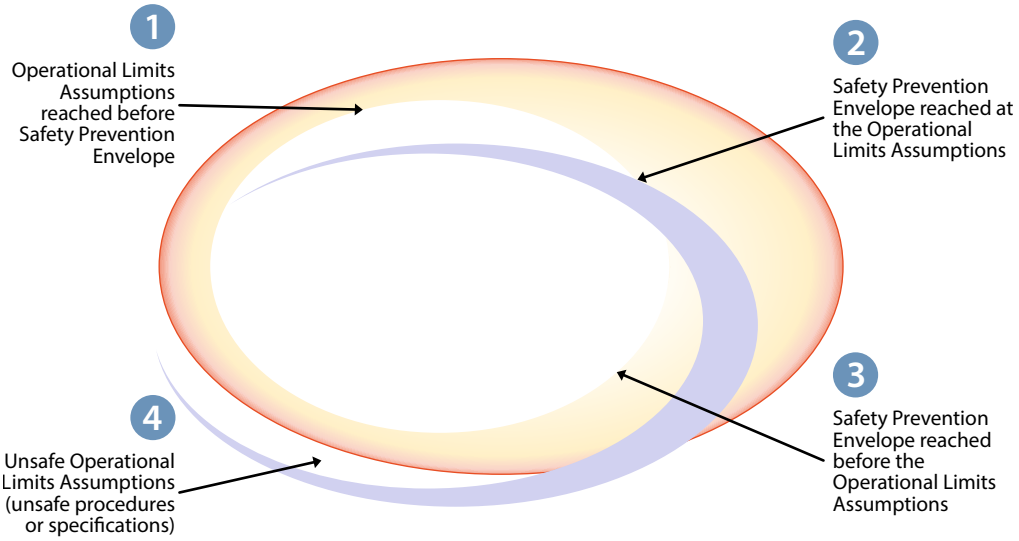

Figure 7: Operational Limits and Safety Envelope Misalignment

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 7. Some operational examples illustrating the patterns are described in Table 1 (p. 8).

The illustration and study of the relative position and the patterns of misalignment of the safety envelope and operational assumptions are important elements of the Learning From All Operations concept. The concepts of the safety control envelope and operational limits assumptions are part of the larger picture of charting of the distance between operations as they actually exist and operations as they are imagined in the minds of managers or rule-makers (Dekker, 2006). This distance is a critical component in understanding an organisation's resilience and tells something about the models of risk currently applied, and how well calibrated they are. Moreover, shared, accurate, and precise knowledge of the dynamics and location of the current operating point and boundary locations is a necessary component of a safety culture (Cook & Rasmussen, 2005).

Table 1: Operational Limits — Safety Envelope Patterns

Operational Limits — Safety Envelope Pattern	Example	Notes
Operational limits assumptions reached before the safety prevention envelope (1)	Design of "stabilization on final" criteria: An approach to a very long dry runway with a lot of head wind may result in a safe landing, even with a major exceedance of approach speed. This would be an example of non-compliance but still within the safety prevention envelope, because beyond the procedure limits, there is still a buffer before the safety prevention envelope is breached.	Regarding the interaction between safety and compliance, this setup should be normal in aviation: The procedures are designed to still leave a safety buffer towards the safety prevention envelope. This kind of procedure design can be called a safe procedure/rule.
Safety prevention envelope reached at the operational limits assumptions (2)	Rejected take-off at V_1, if the take-off is stop-critical (that is, if the remaining stop margin is low). Because take-off performance is calculated beforehand in a very precise way, the calculated parameters are very predictable. If stop margin is 300 m, for example, the aircraft would come to a stop 300 m before the physical runway end. In other words, there is still a non-critical safety control between the stop and the runway excursion scenario. In this case, the results of the take-off performance calculation represent the operational limits assumptions.	This kind of procedure design can be called efficient, because the procedure limit corresponds with the safety prevention envelope. Thus, compliance meets safety requirements with no additional buffers or margins.
Safety prevention envelope reached before the operational limits assumptions (3)	An approach towards a short, wet runway with a tail wind (which is still in operational limits) with an overspeed within procedural limits may result in a risk which is not tolerable by the organisation. This may apply even if performance limits published by the manufacturer are still adhered to. However, in this case, the margin may be reduced in a manner which is not in accordance with the operator's safety targets.	In this example, compliance with the procedure is already unsafe. The safety prevention envelope is penetrated.
Unsafe assumptions (unsafe procedures or specifications): Safety prevention envelope and safety recovery envelope reached before the operational limits assumptions (4)	Certification assumptions for aircraft take-off rotation provides safety buffers at Airport A, but those assumptions are violated for Airport B when all environmental circumstances are at the limit. In this circumstance, at Airport B, adherence to the operational limits assumptions would result in an exceedance of the safety envelope.	In routine operations, this risk may not be evident, because most runways under most environmental conditions still meet the certification assumptions. If this risk is recognised, operators may choose to modify other operational limits assumptions to compensate (e.g., reduce allowable take-off weight).

5. Acknowledgements

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

This note was drafted by Tzvetomir Blajev from Flight Safety Foundation and by Dr. Jon Holbrook from NASA. Thank you to Dr. Immanuel Barshi from NASA and to the members of Flight Safety Foundation's Learning From All Operations Working Group, who contributed to the ideas and clarity of this report: Valerie Stait, Cathay Pacific Airways; Capt. Tom Becker, TUI Fly; Capt. Nick Peterson, American Airlines; Capt. Max Butter, Lufthansa; Sonnie Bates, WYVERN Ltd.; and Capt. Bertrand de Courville.

Suggested citation: Flight Safety Foundation. (2022). Learning From All Operations Concept Note 3: Safety Envelopes and Operational Limits Assumptions.

References

Cook, R., & Rasmussen, J. (2005). "Going Solid": A Model of System Dynamics and Consequences for Patient Safety. BMJ Quality & Safety, 14(2), 130–134. doi:10.1136/qshc.2003.009530

Dekker, S. (2006). Resilience Engineering: Chronicling the Emergence of Confused Consensus. In E. Hollnagel, D. D. Woods, & N. G. Leveson, Resilience Engineering: Concepts and Precepts (pp. 77–78, 81, 85). Aldershot, U.K.: Ashgate.

Holling, C.S., Schindler, D.W., Walker, B.W. and Roughgarden, J. (1995) Biodiversity in the Functioning of Ecosystems: An Ecological Synthesis. In: Perrings, C., Maler, L.G., Folke, C., Holling, C.S. and Jansson, B.O., Eds., Biodiversity and Loss: Economic and Ecological Issues, Cambridge University Press, Cambridge, 44–83. https://dx.doi.org/10.1017/cbo9781139174329.005

Holling, C.S. (1996) Engineering Resilience versus Ecological Resilience. In Schulze, P.E., Ed., Engineering within Ecological Constraints, National Academy.

Leveson, N. (2004, April). A New Accident Model for Engineering Safer Systems. Safety Science, 42(4), 237–270.

Leveson, N.G. (2011). Engineering a Safer World: Systems Thinking Applied to Safety. Cambridge, Massachusetts, U.S.: The MIT Press. doi:10.7551/mitpress/8179.001.0001

Rasmussen, J. (1987, July). Mental Models and the Control of Actions in Complex Environments. Risø-M-2656. Roskilde, Denmark: Risø National Laboratory.

Rasmussen, J. (1997). Risk management in a dynamic society: a modelling problem. Safety Science, 27(2–3), 183–213. doi:10.1016/S0925-7535(97)00052-0

Reason, J., (2008) The Human Contribution: Unsafe Acts, Accidents and Heroic Recoveries. U.K.: Ashgate.