JEPPESEN
A Boeing Company

Global Services

Fatigue Risk Management System

A best practice for quantifying fatigue risk

Understanding and improving the situation for crew fatigue risk in an operation is greatly simplified with a well thought-out way of quantifying risk. To some extent, one can use fatigue reports, collected data and crew feedback after-the-fact. But what is the best practice metric for fatigue risk in upcoming crew pairings and rosters?

Defining Fatigue Risk

There is no formal definition of fatigue risk set by ICAO or IATA. A proven useful definition when planning crew members is: the risk of crew performing a lapse, slip, mistake or violation, negatively impacting flight safety, as an effect of low levels of alertness.

With this definition, the focus primarily lies on flight safety and human error among pilots on active flights, rather than crew comfort or sleepiness during commute, ground duties or a deadhead flight.

A Metric for One Fight

Looking at a single flight, it is clear that the potential for human error, negatively impacting flight safety, is greatly elevated during approach and landing - phases of flying most taxing on pilot capabilities. During this time, the workload is normally at its highest and there is little margin for slowing down or double checking oneself or a colleague in order to reduce risk. The consequence of a slip, lapse, mistake or violation is also potentially disastrous. A vast majority of fatiguerelated accidents in aviation are related to human error during these phases of flying.

For these reasons, it makes sense to focus a metric to estimate fatigue risk primarily using the predicted level of alertness (or sleepiness) near to the end of active flights. A commonly used point in time for collecting data is close to top of descent (TOD) - making it a good choice for predicted alertness level to represent a flight.

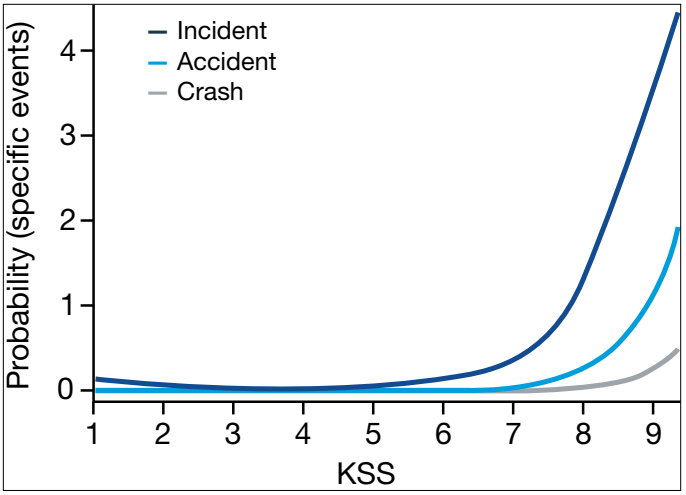

Figure 1. Probability of an accident as a function of self-assessed KSS. [1]

Risk vs. Sleepiness

The risk of a lapse, slip, mistake or violation for an individual has been shown to accelerate as sleepiness increases. Figure 1 illustrates the development in the probability of an accident in a driving simulator where an inflection point is seen just above where subjects are experiencing KSS 8.

However, when predicting future sleepiness, a fatigue model will have limited accuracy for one individual, due to a number of reasons; the models are not perfect, the models are under-informed, and there are significant inter- and intra-individual differences among crew. Figure 2 illustrates how the odds-ratio for an actual accident develops as a function of predicted sleepiness from a bio-mathematical model.

Fatigue Risk Management System

A best practice for quantifying fatigue risk

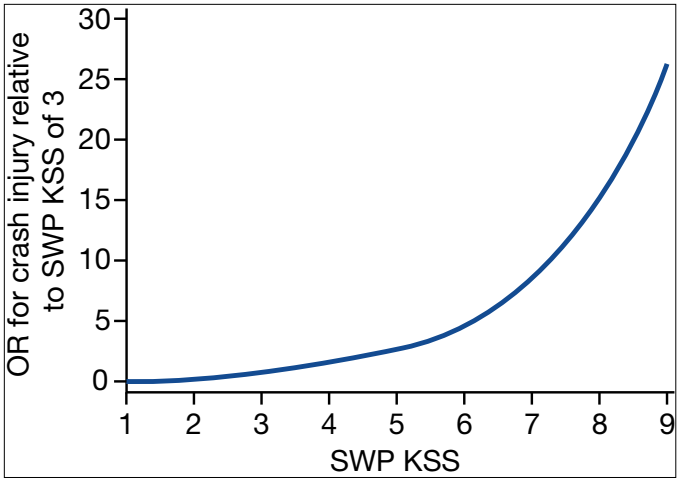

Figure 2. Odds ratio for a road crash as a function of predicted KSS [2]

The conclusion to draw is that a predictive metric capturing fatigue risk should also include a risk contribution from much lower levels of predicted sleepiness than those close to, or passing KSS 8. Human physiology, when being predicted into the future, does not have sharp thresholds separating safe from unsafe. The probability of an accident accelerates more slowly, and from lower levels, when sleepiness is predicted, compared to the risk development observed for self assessed sleepiness. Figure 3 is based on FDM data and tells a similar story.

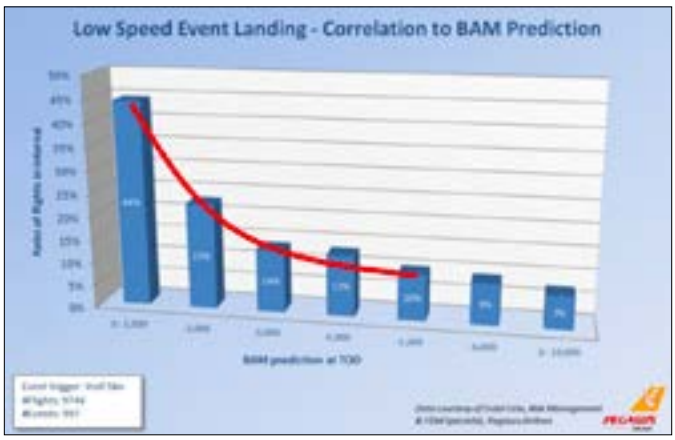

Figure 3. The same shape of acceleration in risk, or decrement in human performance, observed in FDM data (ratio of low speed events) when correlating with predicted alertness for almost 10,000 flights.

A Metric for a Set of Fights

The focus of fatigue risk management when scheduling crew should be to reduce the overall risk for the operator to suffer an incident or accident.

What is really achieved if we reduced fatigue risk on the twenty worst flight duties, if the system response from those changes is negative with the overall risk increasing? Let's take a look at an example.

Example

A reduction of maximum duty time for overnight flights may seem to be a great idea 'at first' for reducing fatigue risk. However, the flights are still present in the flight schedule and will need to be flown. The modified rule may lead to the creation of a lot more night duties, each with commute time before and after, potentially inflicting on physiologically sound timings for sleep. The change creates a need for scheduling more consecutive night duties, stacking up sleep debt for the crew.

What was perceived as an improvement when looking at one night duty in isolation, may well result in the exact opposite looking at the overall operation. This system response from changes made is far too often overlooked.

For this reason, it is crucial to have methods for quantifying, tracking and controlling overall fatigue risk, using a metric that adds up all small probabilities for the individual flight assignments, rather than working with flights in isolation. When doing so, it is logical to use a weighted sum over the set of flights with a weight that accelerates when the predicted sleepiness increases, reflecting how fatigue risk develops in individuals (Figure 2 and 3).

There is no formal standard for this, let alone a way of practically establishing the optimal shape for such a weighting function. Even so, not allowing "perfect to become the enemy of good," we can approximate a shape that mimics the acceleration of risk we do know exists. At Jeppesen, a simple quadratic shape for the risk contribution is used, accelerating from KSS 5 and assigning risk contribution to all flights above that level. Our scale is, however, the other way round as BAM is predicting alertness on the Common Alertness Scale (CAS) from 0 to 10,000, which is anchored to the KSS scale, going in the opposite direction.

The overall risk metric has been named AFR, for Absolute Fatigue Risk, and serves as a proxy for the overall probability of an incident or accident. The higher the number, the higher the risk.

$$AFR(x) = \begin{cases} (50 - x)2 & x<50 \\ 0 & otherwise \end{cases}$$

where x is the predicted alertness in CAS divided by 100

Using AFR

AFR is now our best practice metric reflecting overall fatigue risk in a set of flights. It takes both frequency of flights and severity into account and can be used in a number of helpful ways when planning crew:

- Quantify the system response. In our example, will the shortening of maximum duty time result in overall lower risk? Do we also need to limit consecutive night duties?
- Suppress risk during crew pairing and roster optimization by allowing AFR to feed into the objective function
- Distribute fatigue risk among crew, sharing the burden
- Track the risk development over time

Fatigue Risk Management System

A best practice for quantifying fatigue risk

- Direct focus to the right part of the operation
- Use as a risk profiler by dividing AFR over the number of flights; a metric named NFR (Normalized Fatigue Risk)

Figure 4 shows how the same set of flights have been planned in two different scenarios, but with a clear difference in risk. Both scenarios respect the same planning rules but we can, by just looking at the distributions, quickly confirm that scenario B is preferred as the risk is much lower. This is visible as there are fewer flights in the left tail distribution in scenario B. Our AFR and NFR metrics confirm the same but also quantify the risk has been reduced by 45%.

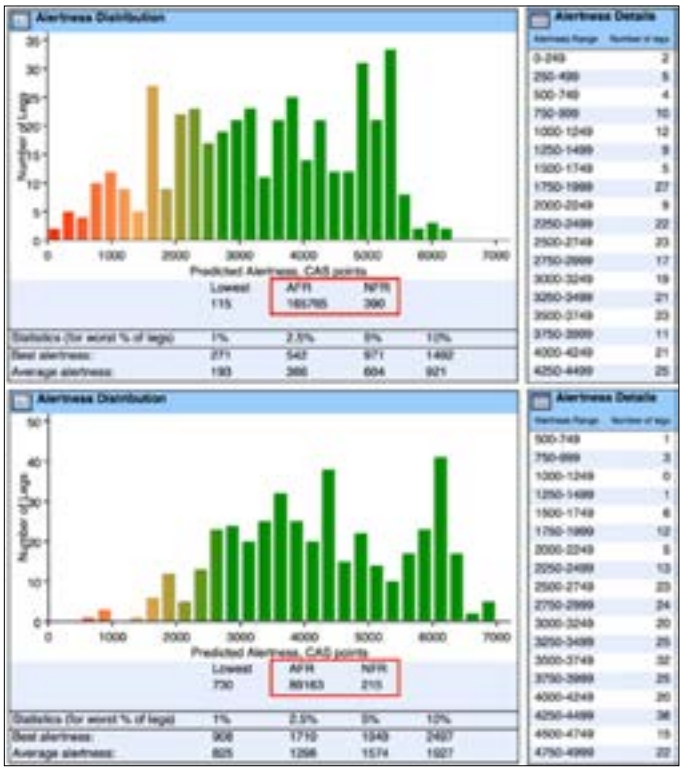

Figure 4. Scenario A (top) and scenario B (bottom) with the same set of flights planned in two different ways into crew rosters.

The AFR/NFR approach is used as an established best practice by a large number of Jeppesen customers today. It is used to control and reduce overall fatigue risk, allowing for greater risk management.

The NFR metric turns out can be very useful when comparing the risk profile between different fleets, bases, ranks, destinations and even airlines. It can also be beneficial for tracking risk development over time. Table 1 lists some typical NFR ranges for various types of operations.

Whereas AFR varies with the volume of the operation, NFR is ideal for benchmarking and has overtaken the position PA5 held, a few years ago [3], as the main risk profile metric. NFR may be less intuitive, but has the advantage of working well on small sets of flights, like a pairing or one roster. It is also capable of reflecting fatigue risk development in human physiology in a more detailed way.

NFR	Type of operation/business model/risk appetite
>400	Long-haul cargo operations with multiple crew bases in different time zones, with sub-standard FRM practices or rules dealing with TZ transitions.
250-400	Challenging long-haul operations such as charter operations with two pilots with multiple TZ crossings and short out-station layovers. Also some 24/7 short haul PAX operations in South and Mid America, the gulf region and Russia, as well as some domestic night cargo operations.
100-250	Challenging short-haul operations and mid-haul with a fair amount of back-of-the-clock operations requiring stacking up consecutive late (or night) operation.
50-100	Normal short-haul PAX operation, performed by many European and North American flag carriers.
<50	Predominantly day-time PAX operation with good FRM practices

Table 1. Typical NFR (risk profile) values for different operations. [4]

Further reading:

- Aligning rules with human physiology
- Are your processes in control?
- Jeppesen Concert
- Assignment-centric performance Indicators
- BAM safety performance indicators
- The secret behind pro-active risk reduction
- Your opinion is interesting, but optimization matters

[1] Ingre M, Åkerstedt T, Peters B, Anund A, Kecklund G, et al. (2006) Subjective sleepiness and accident risk avoiding the ecological fallacy. J Sleep Res 15: 142–148.

[2] Åkerstedt, T., Connor, J., Grey A., & Kecklund, G. (2008). Predicting road crashes from a mathematical model of alertness regulation—The Sleep/Wake Predictor. Accident Analysis & Prevention, 40(4), 1480-1485.

[3] Olbert Å, Hellerström D, Klemets T (2011) A Comprehensive Investigation of Flight and Duty Time Limitations and their Ability to Control Crew Fatigue. Jeppesen White Paper, link.

[4] Approximate ranges only. An individual operator may deviate from ranges stated.

Learn more about Fatigue Risk Management System at jeppesen.com/frm